

THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

---

# Development of high-coherence superconducting devices for quantum computing

*Fabrication process development, materials analysis, and device characterization*

JANKA BIZNÁROVÁ

Department of Microtechnology and Nanoscience  
Quantum Technology Laboratory  
Chalmers University of Technology  
Gothenburg, Sweden, 2024

Development of high-coherence superconducting devices for quantum computing  
Fabrication process development, materials analysis, and device characterization  
JANKA BIZNÁROVÁ  
ISBN 978-91-8103-096-9

Copyright © 2024 JANKA BIZNÁROVÁ  
All rights reserved.

Doktorsavhandlingar vid Chalmers tekniska högskola  
Ny serie nr. 5554  
ISSN 0346-718X

Quantum Technology Laboratory  
Department of Microtechnology and Nanoscience  
Chalmers University of Technology  
SE-412 96 Gothenburg  
Sweden  
Phone: +46 (0)31 772 1000  
[www.chalmers.se](http://www.chalmers.se)

Printed by Chalmers Reproservice  
Gothenburg, Sweden, August 2024

## Abstract

Superconducting quantum circuits are a promising platform for the experimental realization of quantum computers. One of the main challenges in quantum computing hardware is the limited time over which we can sustain the information encoded in a quantum state: the state is easily perturbed by its environment in a process known as decoherence, leading to a loss of the information stored within.

A major source of decoherence in superconducting circuits are parasitic two-level systems (TLSs), which can couple to the device and act as a source of dielectric loss. In this thesis, I study the impact of device design, fabrication procedure, and materials properties on the TLS loss of our devices through materials analysis, cryogenic microwave measurements, and simulations. The devices I study include coplanar waveguide (CPW) resonators, 3D cavity resonators, and aluminium-on-silicon transmon qubits.

In our quantum processor architecture, the CPW resonators used for qubit readout face similar loss sources as the qubits; however, resonators have a faster fabrication and characterization turnaround. Therefore, we use resonators as proxies to investigate the decoherence mechanisms in quantum circuits, and we show how the extracted information can be applied to improve coherence in transmon qubits.

Having identified a dominant source of TLS loss at the substrate-metal interface of our devices, we find that by increasing the grain size of the superconducting films through increasing the film thickness, we can decrease the contribution of intergranular oxide to the TLS loss at this interface. We show that this approach can yield time-averaged energy relaxation times  $T_1 > 200 \mu\text{s}$ , with the best qubit reaching an average  $T_1 = 270 \mu\text{s}$  and a highest observed  $T_1 = 501 \mu\text{s}$ , improving on our previously standard  $T_1 \sim 100 \mu\text{s}$ .

Studying the performance of our 3D cavity resonators as a function of design, material quality, and fabrication procedure, our improved cavities can reproducibly yield resonance quality factors above  $80 \times 10^6$ .

We also integrate our quantum processors into a flip-chip architecture to improve the scalability of our devices. We find that our approach does not measurably degrade the performance of the integrated qubits.

Additionally, we study the nonlinearity of the dielectric susceptibility of TLSs through the intermodulation products generated in a resonator driven by two detuned tones. Our analysis method can reconstruct the standard TLS model parameters from a single spectrum measured at relatively high drive powers corresponding to  $\sim 10^3$  photons.

**Keywords:** Superconducting qubits, superconducting circuits, transmon, quantum computing, coherence, two-level system loss, TLS



## List of Publications

This thesis is based on the following publications:

- [A] J. Biznárová, A. Osman, E. Rehnman, L. Chayanun, C. Križan, P. Malmberg, M. Rommel, C. Warren, P. Delsing, A. Yurgens, J. Bylander, A. Fadavi Roudsari. *Mitigation of interfacial dielectric loss in aluminum-on-silicon superconducting qubits*. npj Quantum Inf., 10, 78, 2024.
- [B] L. Chayanun\*, J. Biznárová\*, L. Zeng, P. Malmberg, A. Nylander, A. Osman, M. Rommel, P. L. Tam, E. Olsson, P. Delsing, A. Yurgens, J. Bylander, and A. Fadavi Roudsari. *Characterization of process-related interfacial dielectric loss in aluminum-on-silicon by resonator microwave measurements, materials analysis, and imaging*. APL Quantum, 1, 026115, 2024.
- [C] J. Biznárová, J. C. R. Hernández, D. Forchheimer, J. Bylander, D. B. Haviland, and G. Andersson. *Intermodulation spectroscopy and the nonlinear response of two-level systems in superconducting coplanar-waveguide resonators*. Phys. Rev. Applied, 22, 014063, 2024.
- [D] M. Kudra, J. Biznárová, A. Fadavi Roudsari, J. J. Burnett, D. Niepce, S. Gasparinetti, B. Wickman, and P. Delsing. *High quality three-dimensional aluminum microwave cavities*. Appl. Phys. Lett., 117(7), 2020.
- [E] S. Kosen, H.-X. Li, M. Rommel, D. Shiri, C. Warren, L. Grönberg, J. Salonen, T. Abad, J. Biznárová, M. Caputo, L. Chen, K. Grigoras, G. Johansson, A. F. Kockum, C. Križan, D. P. Lozano, G. J. Norris, A. Osman, J. Fernández-Pendás, A. Ronzani, A. F. Roudsari, S. Simbierowicz, G. Tancredi, A. Wallraff, C. Eichler, J. Govenius, and J. Bylander. *Building blocks of a flip-chip integrated superconducting quantum processor*. Quantum Sci. Technol., 7(3):035018, 2022.

## Publications outside the scope of this thesis:

- [F] K. Grigoras, N. Yurttagül, J.-P. Kaikkonen, E. T. Mannila, P. Eskelinen, D. P. Lozano, H.-X. Li, M. Rommel, D. Shiri, N. Tiencken, S. Simbierowicz, A. Ronzani, J. Hättinen, D. Datta, V. Vesterinen, L. Grönberg, J. Biznárová, A. F. Roudsari, S. Kosen, A. Osman, M. Prunnila, J. Hassel, J. Bylander, and J. Govenius. *Qubit-Compatible Substrates With Superconducting Through-Silicon Vias*. IEEE Transactions on Quantum Engineering, 3, 5100310, 2022.
- [G] L. Chen, H.-X. Li, Y. Lu, C. W. Warren, C. J. Križan, S. Kosen, M. Rommel, S. Ahmed, A. Osman, J. Biznárová, A. Fadavi Roudsari, B. Lienhard, M. Caputo, K. Grigoras, L. Grönberg, J. Govenius, A. F. Kockum, P. Delsing, J. Bylander, and G. Tancredi. *Transmon qubit readout fidelity at the threshold for quantum error correction without a quantum-limited amplifier*. npj Quantum Inf., 9(26):1–7, 2023.
- [H] C. W. Warren, J. Fernández-Pendás, S. Ahmed, T. Abad, A. Bengtsson, J. Biznárová, K. Debnath, X. Gu, C. Križan, A. Osman, A. Fadavi Roudsari, P. Delsing, G. Johansson, A. Frisk Kockum, G. Tancredi, and J. Bylander. *Extensive characterization and implementation of a family of three-qubit gates at the coherence limit*. npj Quantum Inf., 9(44):1–9, 2023.
- [I] S. Kosen, H.-X. Li, M. Rommel, R. Rehammar, M. Caputo, L. Grönberg, J. Fernández-Pendás, A. F. Kockum, J. Biznárová, L. Chen, C. Križan, A. Nylander, A. Osman, A. F. Roudsari, D. Shiri, G. Tancredi, J. Govenius, and J. Bylander. *Signal crosstalk in a flip-chip quantum processor*. arXiv: 2403.00285
- [J] M. A. Aamir, C. C. Moreno, S. Sundelin, J. Biznárová, M. Scigliuzzo, K. E. Patel, A. Osman, D. P. Lozano, I. Strandberg, and S. Gasparinetti. *Engineering Symmetry-Selective Couplings of a Superconducting Artificial Molecule to Microwave Waveguides*. Phys. Rev. Lett., 129(12):123604, 2022.
- [K] P. Lolur, M. Skogh, W. Dobrautz, C. Warren, J. Biznárová, A. Osman, G. Tancredi, G. Wendin, J. Bylander, and M. Rahm. *Reference-State Error Mitigation: A Strategy for High Accuracy Quantum Computation of Chemistry*. J. Chem. Theory Comput., 19(3):783–789, 2023.
- [L] M. Skogh, W. Dobrautz, P. Lolur, C. Warren, J. Biznárová, A. Osman, G. Tancredi, J. Bylander, and M. Rahm. *The electron density: a fidelity witness for quantum computation*. Chem. Sci., 15(6):2257–2265, 2024.

## Acknowledgments

First and foremost, I would like to extend my gratitude to Jonas for entrusting me with the opportunity to attempt this project under his supervision. Even when not all experiments went *quite* as we envisioned, I knew I could always rely on your support in getting the project back on track.

I am also thankful to Anita for her unwavering dedication to this project and for always taking the time to proofread write-ups or give feedback on measurement results, no matter how busy you were.

My gratitude belongs also to August for his sharp yet kind honesty – your occasional “I don’t believe this!” played a crucial role in prompting us to reexamine our conclusions more thoroughly and to support them more convincingly.

My journey at the Quantum Technology Laboratory at Chalmers began with a master’s thesis project. After taking the course in superconductivity and low-temperature physics – still one of my favourite courses I have ever taken – I approached Per to ask if there was a possibility to do a thesis project in his group. Despite my clear warning about my educational background being mostly in chemical engineering and nanotechnology rather than quantum physics, Per actually ended up having the perfect thesis project in mind. So I am very grateful to Per for being the first one to trust my ability to pull this off, and for giving me the opportunity to try.

I then completed my master’s thesis under the skilful and enthusiastic supervision of Marina. Marina – your patient guidance helped transform me from someone who found quantum technology fascinating yet intimidating into someone confident enough to apply for a PhD in my dream topic. Good that we ended up being more proficient at experimental quantum physics than at canoeing, too!

I started my PhD work in late 2019, which was the beginning of some “unprecedented times”. While Chalmers facilities remained largely open, the presence of staff on campus was reduced, and getting training on lab equipment or procedures was a real challenge. I remain forever grateful to Marcus for taking the time to help get me on track, and teaching me most of my fabrication skills.

In an emergent multidisciplinary field like this, an aspiring experimentalist must quickly acquire a diverse collection of knowledge and skills that typically do not come conveniently pre-packaged in any undergraduate format. We must know the relevant aspects of physics, chemistry, microwave engineering, nanofabrication, and many practical experimental skills that are really developed on the go. I was fortunate to be blessed with incredibly knowledgeable and helpful people to turn to, without whose help my learning curve would have been much steeper. A huge thank you to Chris for teaching me how

to negotiate with qubits and measurement instrument drivers. Your patience with silly questions is an inspiration. Also, my eternal gratitude goes to David for teaching me all the necessary basics of microwave engineering, resonators, and TLSs.

The papers appended to this thesis were the result of many fruitful collaborations. Besides the people already acknowledged in this text, I am grateful to Gustav for envisioning the intermodulation project, and to JC, David and Daniel for pushing through on the modeling and simulations side. It took just a tiny bit longer than we all thought – but the result was pretty satisfactory! The work focused on materials analysis would also not be possible without the skilful contributions of Per Malmberg, Lunjie Zeng, Stefan Gustafsson, Erik Tam, and Eva Olsson.

The work of an experimentalist can sometimes become frustrating. For those cases, it is invaluable to have supportive colleagues to lift one’s spirits, and my colleagues from QTL have often achieved exactly that. From the original crew that drew me in – Marina, Marco, Gustav, Andreas B, Martí, Philip – to you all who joined with me during “the wave”, or later during my time at QTL – Amr, Christian, Chris, Hang-Xi, Liangyu, Lert, Sandoko, Anuj, Miroslav, Tom, Olga, Anna, Andreas N, Hampus, Irshad – you all have been fantastic colleagues and a pleasure to work with – not to mention one of the funniest bunch of people I know.

I also got to supervise Emil’s master’s thesis work, which was a true privilege – I look forward to reading your PhD thesis in the future, Emil. I am certain it’s going to be great. Good luck!

None of my devices presented in the thesis would have been anywhere near as cutting-edge without access to the world-class research cleanroom facility at Myfab Chalmers, and the extremely knowledgeable and helpful staff of the Nanofabrication Laboratory. The smooth running of such an operation is surely a team effort from all involved, so my gratitude belongs to the entire NFL team under the lead of Peter Modh. I would also like to especially to thank the people who have maintained the equipment I used in such pristine conditions, and helped me develop my fabrication processes: Marcus, Niclas, Henrik, Mats, Mattias, Johan, Ruggero, Grigory, Mahdad, and Petra.

Lastly, I would like to acknowledge the support of my family: my parents, who gave me the possibility to reach the point in my education where I could begin a PhD at such a prestigious lab, and to David who was a source of constant support throughout this work – and for putting up with the months leading up to the publication of this thesis.



---

# Contents

---

|   |            |
|---|------------|
| <b>Abstract</b>   | <b>i</b>   |
| <b>List of Papers</b>   | <b>iii</b> |
| <b>Acknowledgements</b>   | <b>v</b>   |
| <br>  |            |
| <b>I Thesis</b>   | <b>1</b>   |
| <br>  |            |
| <b>1 Introduction</b>   | <b>3</b>   |
| 1.1 Quantum computing . . . . .                                 | 5          |
| 1.2 Superconductivity . . . . .                                 | 7          |
| 1.3 Superconducting quantum devices . . . . .                   | 10         |
| <br>  |            |
| <b>2 Loss and decoherence</b>                                   | <b>15</b>  |
| 2.1 Quality factors and coherence times . . . . .               | 15         |
| 2.2 TLS loss . . . . .  | 18         |
| 2.2.1 TLS loss as a function of power and temperature . . . . . | 18         |
| 2.2.2 Modeling TLS loss . . . . .                               | 21         |
| 2.3 Quasiparticle loss . . . . .                                | 22         |
| 2.4 Magnetic noise . . . . .                                    | 24         |

|          |   |           |
|----------|---|-----------|
| 2.5      | Radiation . . . . .                               | 25        |
| 2.5.1    | Purcell decay . . . . .                           | 25        |
| <b>3</b> | <b>Device design</b>                              | <b>27</b> |
| 3.1      | Coplanar waveguide resonator design . . . . .     | 28        |
| 3.2      | Qubit design . . . . .                            | 31        |
| <b>4</b> | <b>Device fabrication</b>                         | <b>35</b> |
| 4.1      | Process overview . . . . .                        | 36        |
| 4.2      | Substrate preparation . . . . .                   | 38        |
| 4.3      | Thin film deposition . . . . .                    | 40        |
| 4.3.1    | Evaporation . . . . .                             | 41        |
| 4.3.2    | Sputtering . . . . .                              | 42        |
| 4.4      | In-situ passivation . . . . .                     | 42        |
| 4.5      | Resist patterning . . . . .                       | 43        |
| 4.5.1    | Optical lithography . . . . .                     | 43        |
| 4.5.2    | Electron beam lithography . . . . .               | 44        |
| 4.5.3    | Ashing . . . . .                                  | 45        |
| 4.6      | Etching . . . . .                                 | 46        |
| 4.6.1    | Wet etching . . . . .                             | 46        |
| 4.6.2    | Dry etching . . . . .                             | 47        |
| 4.7      | Resist removal . . . . .                          | 47        |
| 4.8      | Josephson junctions . . . . .                     | 48        |
| 4.9      | Dicing and packaging . . . . .                    | 50        |
| 4.10     | Choices . . . . .                                 | 51        |
| 4.10.1   | Substrate . . . . .                               | 51        |
| 4.10.2   | Superconductor . . . . .                          | 52        |
| 4.10.3   | Junctions . . . . .                               | 53        |
| 4.10.4   | Deposition type . . . . .                         | 53        |
| 4.10.5   | Etching type . . . . .                            | 54        |
| 4.11     | Common fabrication challenges . . . . .           | 54        |
| 4.11.1   | Finding the correct exposure parameters . . . . . | 54        |
| 4.11.2   | Photolithography on Al . . . . .                  | 55        |
| 4.11.3   | JJ resistance targeting . . . . .                 | 56        |

|          |   |           |
|----------|---|-----------|
| <b>5</b> | <b>Microwave measurements</b>   | <b>59</b> |
| 5.1      | Cryogenic microwave measurement setup . . . . .                                 | 59        |
| 5.1.1    | Dilution refrigeration . . . . .  | 60        |
| 5.1.2    | Microwave measurement set-up . . . . .  | 60        |
| 5.1.3    | Measurement instrumentation . . . . .   | 63        |
| 5.2      | Resonator measurements . . . . .  | 64        |
| 5.2.1    | Resonance circle fit analysis . . . . .   | 65        |
| 5.2.2    | Circulating photon number . . . . .   | 66        |
| 5.3      | Qubit measurements . . . . .  | 68        |
| 5.3.1    | Dispersive readout of a qubit state . . . . .                                   | 68        |
| 5.3.2    | Qubit calibration . . . . .   | 69        |
| 5.3.3    | Anharmonicity . . . . .   | 72        |
| 5.3.4    | Coherence measurements . . . . .  | 72        |
| 5.3.5    | Purcell decay rate . . . . .  | 75        |
| 5.4      | Input line attenuation calibration . . . . .                                    | 75        |
| <b>6</b> | <b>Mitigating loss - film quality and the substrate-metal interface</b>         | <b>77</b> |
| 6.1      | Participation ratio simulations . . . . .                                       | 78        |
| 6.2      | Pre-deposition substrate treatment in evaporated films . . . . .                | 80        |
| 6.3      | Pre-deposition treatment and deposition conditions in sputtered films . . . . . | 84        |
| 6.4      | Film thickness and grain size . . . . .   | 87        |
| 6.4.1    | Materials analysis . . . . .  | 87        |
| 6.4.2    | TLS loss in resonators . . . . .  | 89        |
| 6.4.3    | Qubit coherence . . . . .   | 89        |
| 6.4.4    | High-power loss . . . . .   | 92        |
| 6.4.5    | Discussion . . . . .  | 94        |
| <b>7</b> | <b>Mitigating loss - metal-air and substrate-air interfaces</b>                 | <b>97</b> |
| 7.1      | Passivation of metal films . . . . .  | 98        |
| 7.1.1    | In-situ oxidation . . . . .   | 98        |
| 7.1.2    | In-situ capping . . . . .   | 100       |
| 7.2      | Post-fabrication treatment with HF and HMDS . . . . .                           | 103       |
| 7.2.1    | Al on Si . . . . .  | 104       |
| 7.2.2    | Nb on sapphire . . . . .  | 107       |
| 7.3      | 3D cavities . . . . .   | 108       |
| 7.4      | Flip-chip . . . . .   | 110       |

|           |   |            |
|-----------|---|------------|
| <b>8</b>  | <b>Intermodulation spectroscopy of TLSs</b> | <b>113</b> |
| 8.1       | Intermodulation spectroscopy . . . . .      | 114        |
| 8.2       | Power dependence . . . . .                  | 116        |
| 8.3       | Temporal stability . . . . .                | 116        |
| 8.4       | Modeling and reconstruction . . . . .       | 118        |
| <b>9</b>  | <b>Summary and Concluding Remarks</b>       | <b>121</b> |
|           | <b>References</b>                           | <b>127</b> |
| <b>II</b> | <b>Papers</b>                               | <b>145</b> |
| <b>A</b>  | <b>Paper A</b>                              | <b>A1</b>  |
| <b>B</b>  | <b>Paper B</b>                              | <b>B1</b>  |
| <b>C</b>  | <b>Paper C</b>                              | <b>C1</b>  |
| <b>D</b>  | <b>Paper D</b>                              | <b>D1</b>  |
| <b>E</b>  | <b>Paper E</b>                              | <b>E1</b>  |

**Part I**

**Thesis**



# CHAPTER 1

---

## Introduction

---

This thesis is based on five appended scientific journal articles, in which we study quantum state decoherence in superconducting quantum circuits developed for application in quantum information processing.

In quantum information science, we encode and manipulate information represented in a quantum state. Promising platforms for quantum computing include trapped Rydberg ions, the spin states of nuclei in magnetic fields, electron spin states inside nitrogen-vacancy centers, or states of circuit variables in superconducting circuits.

One of the main challenges in quantum computing across all the aforementioned platforms is the limited timescale over which we can sustain the information encoded in the quantum state: the quantum state is easily perturbed by its environment in a phenomenon known as *decoherence*, resulting in a loss of the information encoded within. The timescale of this decoherence (also referred to as the coherence time) limits the number of meaningful logical operations we can perform on the quantum state before the information is lost.

In superconducting quantum devices, which are the quantum computing platform investigated in this thesis, a major source of decoherence is the presence of so-called parasitic two-level system (TLS) defects in the devices. These TLSs can couple to the device and cause decoherence by dielectric loss and parameter fluctuations.

In the work contained in this thesis, we investigate the origins of decoherence in various architectures of superconducting devices through cryogenic microwave measurements, materials analysis, imaging, modeling, and simulations. Based on our findings, we identify the dominant microscopic sources of decoherence in our devices. We develop and test strategies to mitigate the immediately dominant decoherence mechanisms, improving the resulting device coherence.

In Paper **A**, we find a dominant source of TLS loss at the substrate-metal interface of our qubits and resonators, and mitigate this loss through increasing the grain size in the metal films, thereby reducing the contribution of intergranular oxide to interfacial loss. In Paper **B**, we perform a systematic study of the impact of fabrication process variations on loss in coplanar waveguide resonators. In Paper **C**, we demonstrate and characterize the frequency mixing products generated by the nonlinearity of parasitic TLSs in a superconducting resonator. In Paper **D**, we investigate loss in 3D cavity resonators, focusing on the impact of metal purity, geometry, and fabrication treatments. Finally, in Paper **E**, we improve the scalability of our devices by integrating the qubits into a flip-chip architecture. We show that our approach does not degrade the qubits' performance compared to our standard planar architecture. In addition, this thesis contains previously unpublished results of coherence investigations, which are outside the scope of the appended papers.

In the present chapter, I will contextualize the work presented in the thesis and introduce the underlying background concepts. The following chapter will cover the standard device quality measures used to characterize the performance of the qubits and resonators we study. I will also discuss the theory of the various loss channels known to contribute to decoherence in superconducting devices.

The subsequent chapters will contain the theoretical concepts behind our chosen device design, as well as behind the fabrication and measurement procedures, interspersed with the relevant experimental details.

Afterwards, I will present the main results of my work on coherence in superconducting devices, connecting the results published across the different articles together with results of previously unpublished experiments. However, these results overview chapters serve only as a summary of the highlights of the papers appended to this thesis, and not all results contained in these papers will be mentioned. Therefore, the interested reader is encouraged to find more details in the appended papers.

The final chapter concludes the thesis with a summary of the main results, with suggestions of how the findings presented in this thesis could be applied in future works on extending the coherent lifetimes of quantum states in superconducting devices.




## 1.1 Quantum computing

Over the last one hundred years, the meaning of the word “computer” has shifted markedly from originally signifying a person whose task it was to manually compute numerical calculations. A big shift came with the advent of the first electromechanical computers – behemoths that could barely fit through the door, possessing the ability to perform numerical calculations at the speed of hundreds of operations per second [1]. After decades of considerable research and development efforts in the field of computer science, computers became omnipresent in the world. Today, computers range from pocket versions for personal use with the ability to access a large catalogue of mankind’s collective information and perform complex calculations at trillion operations per second, all the way to industrial supercomputers able to perform as many as  $10^{18}$  operations per second [2].

However, despite these astounding advancements, these so-called *classical* computers we are familiar with are bound to struggle with certain classes of problems, such as high-complexity optimization problems, prime factorization of large integers, or the quantum simulation of molecules [3, 4]. The aim of quantum computing is to exploit quantum mechanical phenomena such as the superposition of degrees of freedom in the physical devices that represent quantum information, as well as quantum state entanglement, to provide up to an exponential speed-up on certain classes of problems [5].

On the flip side, quantum computers are not likely to be particularly efficient in many standard algorithms performed by classical computers today. Therefore, the intent behind building a quantum computer is not to replace classical computers in the everyday tasks they already perform to a more than satisfactory degree; rather, quantum computers are intended to complement classical computers in tasks that are not practically feasible on classical platforms. In many cases the relationship between classical and quantum computers is likely to be symbiotic – such as today’s quantum computer prototypes being controlled and programmed by classical computers.

Both classical and quantum computers encode information in bits within a basis conventionally denoted as 0 and 1. A classical bit can only exist in either 0 or 1; however, a quantum bit (or *qubit*) can exist in an arbitrary *quantum superposition* of the two. A superposition of the ground state  $|0\rangle$  and the first excited state  $|1\rangle$  in a qubit can be expressed as a linear combination of the two states:  $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$ , where  $\alpha$  and  $\beta$  are the probability amplitudes of finding the qubit in the associated state. All possible qubit state vectors  $|\psi\rangle$  fulfilling the coherent superposition condition  $|\alpha|^2 + |\beta|^2 = 1$  can be represented as a vector on the surface of the so-called Bloch sphere (see Fig. 1.1).



**Figure 1.1:** The Bloch vector representing a coherent quantum state  $\psi$ , which is a linear combination of the ground state  $|0\rangle$  and the excited state  $|1\rangle$  on the surface of the Bloch sphere. In case of decoherence by energy relaxation and the loss of phase coherence, the qubit decays into the so-called mixed state and the state vector retracts within the sphere (see Chapter 2).

Additionally, the quantum mechanical nature of qubits enables the *quantum entanglement* of multiple qubits, which gives rise to *quantum parallelism*. Quantum parallelism means that we have parallel access to all  $2^n$  states that  $n$  entangled qubits can represent, which is at the heart of the theoretical exponential speed-up compared to classical algorithms that we strive to achieve. Consequently, with just 270 qubits, we can represent  $2^{270}$  states – more than the estimated number of particles in the known Universe [6]. Achieving such a large-scale superposition experimentally in a real system is, however, significantly more difficult than what it might sound like. It is not sufficient to build 270 independent qubits on a chip, but one needs to solve a number of device integration and quality challenges such as wiring scalability, device parameter targeting, individual qubit control, and coherence, to name only a few. This thesis focuses on the effort of increasing the coherent lifetimes of quantum states encoded in superconducting devices.

The first prototypes of superconducting qubits supported coherent lifetimes on the order of nanoseconds [7], the mere observation of which was an admirable feat. Steady progress during the subsequent years in understanding the coherence limitations of qubits and making improvements in qubit design [8–11], reducing the noise introduced by the qubit’s environment [11, 12], improving fabrication methods [13, 14], and performing materials studies [10, 15, 16], have pushed the standard coherence values of qubits with the so-called transmon architecture (see Sec. 1.3) to around  $100\ \mu\text{s}$  [17, 18]. This value of  $\sim 100\ \mu\text{s}$  remained the standard for several years regardless of the materials platform on which these transmons were fabricated, until finally surpassed by qubits made on tantalum films which were the first to demonstrate energy relaxation times


of 300 – 480  $\mu\text{s}$  [19, 20], closely followed by transmons made on titanium nitride and niobium with average relaxation times up to 291  $\mu\text{s}$  and 210  $\mu\text{s}$ , respectively [21, 22]. In this thesis, I demonstrate the ability of aluminium-based transmons to reach energy relaxation times well above 200  $\mu\text{s}$ , with the best qubit reaching the time-averaged value of 270  $\mu\text{s}$  and the best observed value of 501  $\mu\text{s}$ .

An obvious question presents itself: how much do we need for practical applications in quantum computing? How long of a coherence time is enough? The answer is complex. It is not only the qubit coherence that is the deciding factor, but also the gate speed and number of operations that enter the equation [23–25]. Faster gates and more efficient algorithms reduce the time during which the device needs to retain the coherence of the encoded quantum state, and longer coherence times enable more complex calculations. While the current state-of-the-art coherence values of  $\sim 100 \mu\text{s}$  and gate times below 100 ns are sufficient to enable demonstrations of quantum supremacy [26, 27] – calculations performed on a quantum processor that are intractable on a classical computer within a reasonable time – these values are insufficient in order to implement truly useful algorithms on a quantum computer [28]. To help bridge the gap between the available and desirable coherence times, efforts in quantum error correction seek to encode a “logical qubit” into a set of multiple physical qubits with different functions, including error detection. Such logical qubits can yield lower error rates – however, at the cost of an increased overhead in the number of physical qubits needed on the device [24].

On the other hand, there is no such thing as too much coherence – if we find ourselves in such a fortunate position that a qubit’s coherence time far exceeds our operation time, we have the option to actively reset our qubits to the ground state before attempting a new calculation [29].

## 1.2 Superconductivity

Superconductivity is a property granted to select materials, where below a certain critical temperature  $T_c$ , electrical resistance abruptly drops to zero and direct current is transported losslessly [30]. I show this phenomenon experimentally for a superconducting Al transmission line in Fig. 1.2. The superconducting critical temperatures are typically very low for pure elemental superconductors – a few kelvins or below [31] – which is why it took until Heike Kammerlingh Onnes’s experiments on mercury with liquid helium as a refrigerant in 1911 for the phenomenon to be observed for the first time [32]. Since then, superconductivity has been discovered in several other chemical elements, inter-



**Figure 1.2:** DC-transport measurement of an aluminium transmission line. The temperature-dependent resistance above  $T_c$  is consistent with the expected trend for a normal metal. Below  $T_c$ , the measured resistance drops to 0. The inset shows a closer look at the region of the main figure close to  $T_c$ .

metallic compounds, organo-metallic compounds, and ceramics [33]. Many of these more recently discovered superconductors fall into the category of so-called high-temperature superconductors, with  $T_c$  values above the boiling point of nitrogen (77 K). As liquid nitrogen is a readily available low-cost coolant, this makes high-temperature superconductors more suitable for practical applications in e.g. magnetic resonance imaging, nuclear magnetic resonance, particle accelerators, or fusion reactors [34–36].

Superconductors are incredibly valuable in quantum processors: by providing a dissipation-free environment they, quite predictably, help extend the coherent lifetimes of the fragile quantum information encoded in single quanta of energy. Superconducting wires also help transport the extremely weak device output signals up the measurement output line until these signals can be amplified.

However, the perfect conductivity statement is only truly valid for direct currents – alternating currents will experience frequency-dependent dissipation, which we need to be mindful of when choosing the frequencies of our devices (typically in the GHz regime). Additionally, even for DC there is a maximum critical current  $I_c$ , above which the superconductor will act as a normal metal. This  $I_c$  depends on material properties, as well as the conductor’s geometry.


The first explanation of superconductivity was offered by Gorter and Casimir in the the so-called *two-fluid model* of superconductivity [37]. They proposed the existence of two different types of charge carriers – “normal” carriers that lose energy in scattering events involving the ion lattice of the metal, and so-called “supercarriers” that somehow

avoid this type of loss.

In the equivalent circuit of this model, shown in Fig. 1.3a, all DC current is passed through the path of least resistance - the supercarriers. At AC, however, some current will flow through the resistive path, and we observe dissipation. The relative density of the normal carriers and supercarriers is temperature dependent (see Fig. 1.3b). In the original two-fluid model, the density of supercarriers  $n_s$ , normalized by the total carrier density  $n$ , scales as  $n_s/n = 1 - (T/T_c)^4$ . Subsequent models, such as the one discussed in Section 2.3, expand on this model and refine its temperature dependence. In either model, just below  $T_c$ , most carriers are still normal, while at  $T = 0$ , all carriers have turned superconducting. At finite temperatures, there will always be normal electrons present (also referred to as *quasiparticles* in the context of superconductivity), which can cause dielectric losses – particularly at AC (see Sec. 2.3).

The two-fluid model is fully classical and makes no assumptions as to the nature of these mysterious “supercarriers”, yet it describes the observed behaviour of superconductors remarkably well. Only some 20 years later, Leon Cooper realized that in certain metallic lattice structures, the negative electron passing through the lattice can ever-so-slightly distort its structure, leaving a local maximum of positive charge in its wake. This positive charge can then attract another electron, challenging the the age-old stipulation that electrons repel each other. Electron pairs bonded in this way are referred to as Cooper pairs in honour of their discoverer [38]. The conditions for Cooper pair formation are met at  $T_c$  for select materials. Unlike their fermionic single electron counterparts, the newly formed Cooper pairs, consisting of electrons of opposite spin, are in fact bosons. Since bosons are not subject to Pauli’s exclusion principle, the Cooper pairs are all free to condense to the ground state. This delocalized state can be described by a single wave function  $\psi$ , and allows for current to be transported losslessly within the bounds of  $\psi$ . Thus, superconductivity is at its core a quantum mechanical phenomenon. The Bardeen-Cooper-Schrieffer (BCS) theory of superconductivity describes the microscopic model of superconductivity based on the Cooper pair notion in more detail [39, 40].

The second defining feature of superconductors, apart from perfect conductivity, is their perfect diamagnetism - the expulsion of magnetic fields from the superconductor’s volume, also called the *Meissner effect* [30]. Based on their diamagnetic behaviour, we distinguish two types of superconductors – Type-I and Type-II [41]. Type-I superconductors will expel all magnetic fields below a critical value  $H_c$ . In Type-II superconductors, we recognize two different critical fields. Below the lower value  $H_{c1}$ , we achieve the Meissner state of perfect diamagnetism where all field is expelled. Between  $H_{c1}$



**Figure 1.3:** Two-fluid model of superconductivity. **a** The equivalent circuit for the two-fluid model, where a current  $J_{tot}$  can flow losslessly through the inductive path as a supercurrent  $J_s$  carried by the supercarriers  $n_s$ , or through the resistive path as a normal current  $J_n$  carried by the normal carriers  $n_n$ . **b** Temperature-dependent density of the supercarriers  $n_s$  and the normal carriers  $n_n$ .

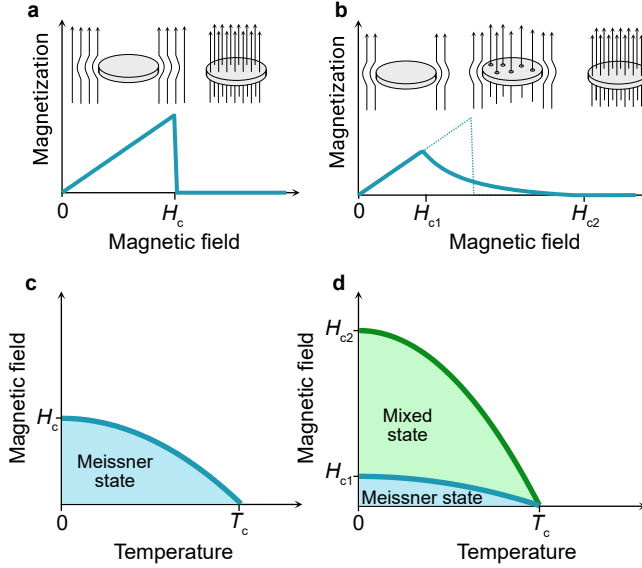
and  $H_{c2}$ , while magnetic field is still being expelled, some field is allowed to enter the superconductor, where it is pinned in so-called Abrikosov vortices containing a single quantum of flux each [42]. For an illustration of the diamagnetic properties of both superconductor types, see Fig. 1.4.

It should be noted that this diamagnetism is only perfect for bulk superconductors. The London brothers found that the magnetic field decays on the scale of the so-called London penetration depth  $\lambda_L$  [43]. As a result, thin films with thicknesses on the scale of  $\lambda_L$ , even when made from Type-I films, can also exhibit Type-II behaviours and pin flux vortices inside the superconductor, particularly at material defect sites (see Sec. 2.4). In a related effect, the Londons found that the current carried by the superconductor is located within this penetration depth  $\lambda_L$ .

### 1.3 Superconducting quantum devices

A strong advantage of superconducting quantum computing lies in the freedom of choice over the key parameters of our system, such as the frequencies of the different elements, and their mutual couplings.

We can in essence engineer “artificial atoms” - macroscopic devices that despite consisting of quadrillions of atoms, can be made to absorb or emit single photons at anharmonic spectra, just like regular atoms. In the context of this work, this is achieved through circuit quantum electrodynamics (circuit QED), which enables us to engineer light-matter interactions using basic circuit elements.

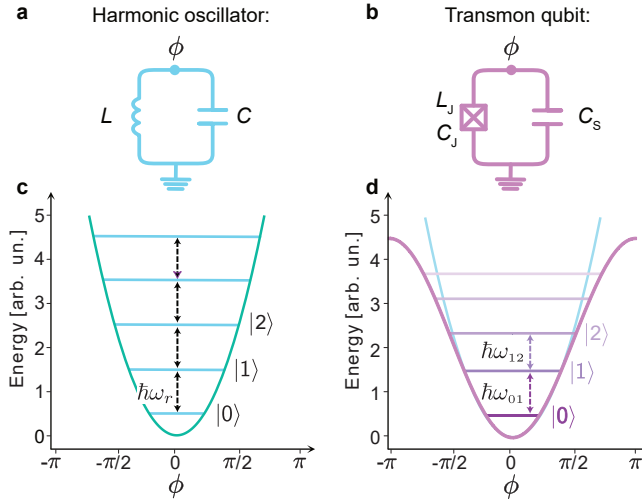


**Figure 1.4:** Diamagnetic properties of bulk Type-I and Type-II superconductors. **a** Below  $T_c$ , the Type-I superconductor expels magnetic fields lower than the critical field  $H_c$  from its volume. **b** The Type-II superconductor likewise expels magnetic fields below the first critical value  $H_{c1}$ . When a field between the two critical values  $H_{c1}$  and  $H_{c2}$  are applied, single magnetic flux quanta are pinned in vortices inside the superconductor. **c–d** Phase diagrams of the Type-I and Type-II superconductors, respectively, showing the critical field values as a function of temperature.

In contrast, platforms relying on actual atoms, such as trapped ion-based quantum computing – another promising platform for quantum computing – are tied to a few specific atoms with fixed properties.

Perhaps the most basic element in circuit QED is the resonator, which can be modeled as an  $LC$  oscillator (see Fig. 1.5a). The energy spectrum of these resonators is that of a quantum harmonic oscillator – a parabolic potential with discrete energy levels separated by  $\hbar\omega_r$ , where  $\omega_r = 1/\sqrt{LC}$  is the resonance frequency. Some examples of the physical realizations of such quantum harmonic oscillators are lumped-element resonators, microstrip-line resonators, planar resonators, coplanar-waveguide resonators, or coaxial-waveguide resonators. The work contained in this thesis utilizes the latter two of these architectures.

Unfortunately, since all transition frequencies are equal in a harmonic oscillator, we




**Figure 1.5:** Lumped-element models and energy spectra for a harmonic oscillator and a transmon qubit. **a** Resonator modeled as an  $LC$  oscillator.  $\phi$  denotes the superconducting phase on the island, with the ground referenced as zero. **b** The transmon qubit consists of a JJ element with a capacitance  $C_J$  and a nonlinear inductance  $L_J$ , shunted by a large capacitor  $C_S$ . **c** Energy diagram of a quantum harmonic oscillator with equidistant energy levels, represented here by the photon number states  $|0\rangle$ ,  $|1\rangle$ , etc., separated by  $\hbar\omega_r$ . **d** The harmonic oscillator potential (blue line) of a transmon is distorted by the cosine term brought in by the JJ element. In the resulting anharmonic potential, the energy level separation between the ground state  $|0\rangle$  and the first excited state  $|1\rangle$  is  $\hbar\omega_{01}$ , which is no longer equal to the transition energy from  $|1\rangle \rightarrow |2\rangle$  characterized by  $\hbar\omega_{12}$ .

cannot easily distinguish between the different states, or address the desired levels at will. Therefore, while resonators remain a crucial element in quantum processors, they alone cannot constitute a qubit.

What we need to do is to turn the energy spectrum anharmonic – meaning, that the transitions between the different energy levels have different frequencies, and we can selectively drive the transitions of our choice by selecting the corresponding frequency. This can be done by introducing a nonlinear element into the circuit, such as the nonlinear inductance of a Josephson junction (JJ). The sinusoidal term introduced to the Hamiltonian by the JJ distorts the quadratic potential of the linear  $LC$  terms, resulting in an anharmonic energy spectrum. For an example of this effect, demonstrated on a transmon qubit consisting of a JJ element shunted by a large capacitance [8], see





**Figure 1.6:** Superconducting wave functions inside an SIS JJ. The wave functions  $\psi_1$  and  $\psi_2$  for the two electrodes decay exponentially outside the superconducting regions, but due to their overlap inside the weak link, supercurrent can flow across this insulating barrier.

Fig. 1.5b.

A Josephson junction consists of two superconductors separated by a weak link. This weak link could be an insulator, a normal metal, or a constriction. In our devices, we use superconductor-insulator-superconductor (SIS) junctions with Al as the superconductor, and a thin layer of Al oxide as the insulator.

If the weak link between two superconducting electrodes is small enough, the wave functions of the two superconducting electrodes can overlap [44, 45], see Fig. 1.6. As a result, Cooper pairs can tunnel losslessly through the insulating layer from one electrode to the other. The resulting current flowing through the JJ is described in the DC Josephson effect as a function of the phase difference across the JJ, as well as the critical current  $I_c$  above which the JJ becomes resistive:

$$I = I_c \sin(\varphi). \quad (1.1)$$

The voltage across the JJ is given in the AC Josephson effect as:

$$V = \left( \frac{\hbar}{2e} \right) \frac{d\varphi}{dt}, \quad (1.2)$$

where  $(\hbar/2e)$  can also be represented by  $(\Phi_0/2\pi)$ , where  $\Phi_0$  is the flux quantum. Combining equations 1.1 and 1.2, we find:

$$\frac{dI}{dt} = I_c \cos(\varphi) \frac{d\varphi}{dt}, \quad (1.3)$$

together with

$$V = \frac{\hbar}{2eI_c \cos(\varphi)} \frac{dI}{dt} = L_J \frac{dI}{dt}, \quad (1.4)$$

where  $L_J = \hbar/2eI_c \cos(\varphi)$  is the nonlinear Josephson inductance of the junction.

The theory presented above gives the framework for the Cooper pair box - a trail-

blazer for superconducting qubits [7]. In the years after its demonstration, the Cooper pair box design was improved in the transmon qubit, where an additional shunt capacitance isolates the qubit from charge noise. Adding two JJs in parallel allows for the resulting inductance to be tunable by an applied magnetic flux, paving way for the frequency-tunable transmon [46]. I will discuss the various design elements of the transmon qubit in more detail in Chapter 3.

---

## Loss and decoherence

---

Quantum information encoded in a single-photon resonant mode is fragile. Over time, the quantum state is affected by its environment, resulting in dephasing and energy relaxation, which causes loss of quantum coherence and thereby also loss of the information stored within. The timescale of this decoherence limits the number of meaningful logical operations we can consecutively perform on the qubit. In this chapter, I define the various decoherence timescales and device quality figures of merit we use to benchmark quantum devices. I will also discuss the various loss mechanisms leading to decoherence.

### 2.1 Quality factors and coherence times

The figure of merit for resonators is the quality factor  $Q$ , which quantifies the energy stored vs. the power dissipated during a single oscillation:

$$Q = 2\pi \frac{\text{stored energy}}{\text{energy dissipated per cycle}} = 2\pi f_r \frac{\text{stored energy}}{\text{dissipated power}} = \omega_r T_1, \quad (2.1)$$

where  $f_r$  and  $\omega_r$  are the natural and angular resonance frequencies, respectively, and  $T_1$  the time constant of the resonator's exponential energy decay.


When integrated in a circuit, the total  $Q$  of the observed resonance can be broken down into the internal quality factor  $Q_i$  which quantifies the loss intrinsic to the resonator, and the external quality factor  $Q_c$  quantifying losses due to coupling to the surrounding circuitry, such as the input/output transmission line. As  $Q_c$  is a design parameter which can be modified arbitrarily, our interest lies in extracting the  $Q_i$ . We discuss the details of this in section 5.2.

In qubits, we recognize two paths of decoherence - *energy relaxation*, and the loss of phase coherence (*dephasing*) [46]. The two decoherence mechanisms, as well as the consequence of their combination, are depicted on the Bloch sphere in Fig. 2.1.

Measuring the energy relaxation time constant ( $T_1$ ) is relatively straightforward, as we can observe energy loss by monitoring the decay from one pole of the sphere ( $\Psi = |1\rangle$ ) to the opposite pole ( $\Psi = |0\rangle$ ). Phase is undefined at the poles, and so observing the decay from  $|1\rangle$  to  $|0\rangle$  gives us direct information about the longitudinal decay along the  $z$  axis.

The pure dephasing time constant ( $T_\varphi$ ), on the other hand, cannot be measured directly, as at any point on the Bloch sphere where dephasing is observable, the system also suffers from energy relaxation. Thus, a combination of both mechanisms is observed.

The effects of dephasing are most impactful at the equator of the sphere. Therefore, we define the decoherence time  $T_2$  of a qubit as the decay of the equal superposition  $\Psi = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$  where the qubit, subject to both energy relaxation and dephasing, decays into the mixed state. In the mixed state, the qubit state vector is no longer a coherent superposition represented by the surface of the Bloch sphere, but retracts



**Figure 2.1:** **a** Energy relaxation, also known as the longitudinal decay rate, brings the qubit to  $|0\rangle$ . **b** Pure dephasing results in loss of phase coherence. **c** Decoherence as a combination of energy relaxation and pure dephasing brings the qubit into the mixed state, and the Bloch vector retracts within the sphere.

within the sphere and is instead a statistical mixture of the possible states.

Provided we have information about the qubit's characteristic  $T_1$ , we can extract  $T_\varphi$  from the measured  $T_2$ :

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\varphi}. \quad (2.2)$$

The above relation also sets the theoretical limit of  $T_2 = 2T_1$  in the absence of pure dephasing.

While considering  $T_1$  and  $T_2$  is important from the point of view of logical gates and algorithms, it is a rather inconvenient measure for comparing the quality of different devices, as  $T_1$  depends directly on frequency (see Eq. 2.1). Therefore, when assessing the impact of process development or design changes on qubit performance, comparing the qubit's frequency-independent  $Q$  values gives a more accurate insight.

The inverse of the  $Q$  factor quantifies the loss of the device, which is the sum of the various contributing loss mechanisms:

$$\frac{1}{Q} = \frac{1}{Q_{\text{TLS}}} + \frac{1}{Q_{\text{qp}}} + \frac{1}{Q_{\text{mag}}} + \frac{1}{Q_{\text{p}}} + \frac{1}{Q_{\text{rad}}} + \dots, \quad (2.3)$$

where the partial  $Q$  values describe the theoretical quality factors for a qubit limited solely by dielectric loss due to parasitic two-level systems (TLSs) in  $Q_{\text{TLS}}$ , by quasiparticle loss in  $Q_{\text{qp}}$ , by magnetic flux noise in  $Q_{\text{mag}}$ , by Purcell decay in  $Q_{\text{p}}$ , and by other radiation losses unrelated to Purcell decay in  $Q_{\text{rad}}$ . The associated frequency-dependent loss rates for these mechanisms are conventionally denoted as  $\Gamma_i$ .

As these losses are cumulative, it is challenging to quantify their relative contributions in order to resolve which loss path is limiting the performance of the device and should therefore be targeted with the highest priority. Some mechanisms are uniquely dependent on a certain measurement condition, such as the drive-power dependent TLS loss. However, variations in device temperature can affect TLSs as well as quasiparticles, and a change in geometry may impact radiation loss, TLS loss, and even quasiparticle loss.

As  $T_1$  is generally the coherence-limiting factor of the devices presented in this work, I focus on mitigating energy loss as opposed to dephasing. Dephasing, associated with qubit frequency instabilities, can share some of its culprits with relaxation, such as the coupling to interacting TLSs [12, 47]. In frequency-tunable qubits based on the flux-tunable superconducting quantum interference devices (SQUIDs), magnetic flux fluctuations play a significant role in qubit frequency instability [48]; however, this issue is suppressed in our fixed-frequency qubits based on single Josephson junctions. Addi-

tionally, qubit frequencies can fluctuate due to ac-Stark shifts caused by varying photon number populations in readout resonators [49].

## 2.2 TLS loss

At cryogenic temperatures and single-photon drive powers, the quality of superconducting devices is often limited by dielectric loss due to parasitic two-level system (TLS) defects [12, 50, 51]. The term TLS here refers to any microscopic system that can absorb energy to switch between two predominant states. The device can couple to near-resonant TLSs, resulting in energy dissipation. Additionally, the coupling between the device and the TLS fluctuates over time, leading to device frequency fluctuations [47, 52].

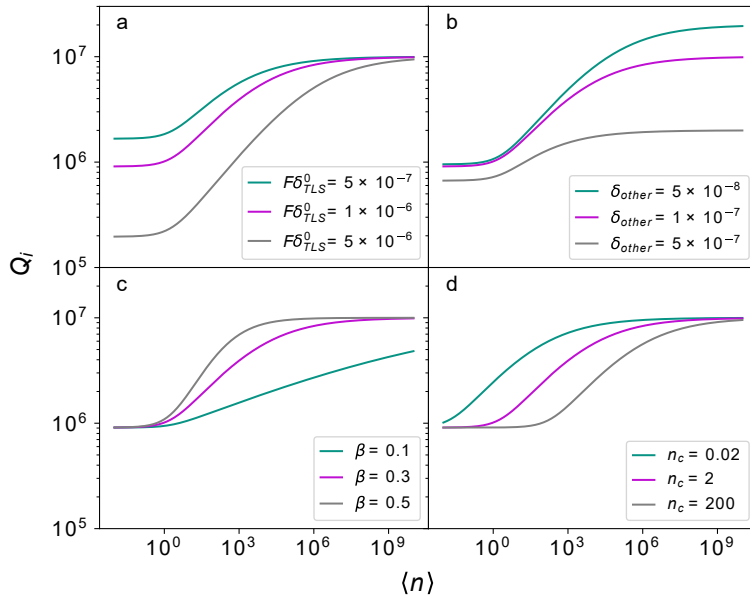
The physical nature of these TLSs is a topic of ongoing active research, but they are often found to reside in amorphous materials at interfaces between the different materials in the device (see the coplanar waveguide cross-section in Fig. 2.3), such as surface oxides, fabrication residues and adsorbates [53, 54]. Microscopically, TLSs have been traced to the tunneling of atoms and electrons promoted by positional defects in amorphous materials, as well as the presence of dangling bonds, atomic hydrogen, and molecular oxygen at interfaces [55–58]. In qubits, TLS losses can also occur inside the Josephson junction barrier [59].

The coupling strength between a parasitic TLS and the device depends on the physical location of the TLS - for instance, a defect near the minimum of the electric field distribution of a coplanar mode couples to the device less strongly than a defect inside the junction oxide.

In paper C, we demonstrate the inherent nonlinearity of the dielectric susceptibility due to TLSs through measuring the mixing products generated when driving a TLS-limited resonator by two slightly detuned drive tones.

### 2.2.1 TLS loss as a function of power and temperature

In superconducting resonators, TLS loss manifests itself in the characteristic dependence of  $Q_i$  on the number of photons  $\langle n \rangle$  circulating in the resonant mode. At low drive powers, with only a few photons circulating in the device, the competition between the TLSs and the device over these photons is significant, and the resulting  $Q_i$  can be greatly diminished. However, as a single TLS can only absorb a single energy quantum at a time, increasing the drive power can effectively “saturate” the TLS bath, resulting in a



**Figure 2.2:** Simulation of the TLS model for different loss parameters. The curve in purple is identical in all four panels, showing loss parameters typical for my devices, with  $F\delta_{TLS}^0 = 1 \times 10^{-6}$ ,  $\delta_{other} = 1 \times 10^{-7}$ ,  $\beta = 0.3$  and  $n_c = 2$ . In each of the panels **a–d**, one of the model parameters is decreased (green) and increased (grey) at a time. **a** A higher TLS loss quantified by  $F\delta_{TLS}^0$  suppresses  $Q_i$  at low power. **b** Other, power-independent losses show at high power when the TLS bath is saturated. When  $\delta_{other} \ll F\delta_{TLS}^0$ , the  $Q_i$  at low  $\langle n \rangle$  is largely determined by TLS loss (green and purple). However, due to the cumulative nature of the various independent loss channels that contribute to the total loss, the  $Q_i$  becomes even more suppressed at low  $\langle n \rangle$  as  $\delta_{other}$  and  $F\delta_{TLS}^0$  become comparable (grey). **c** A weaker slope of the power-dependence  $\beta$  increases the power at which the TLS bath is saturated. For low  $\beta$  such as 0.1 for the cavities in paper **D**, it becomes challenging to saturate TLS with power and capture this dependence. **d** A lower  $n_c$  decreases the input power necessary to start saturating individual TLSs.

higher apparent  $Q_i$ . TLSs can be similarly saturated by thermal photons. The power and temperature dependence of the resonator  $Q_i$  is described by the modified standard tunneling model for TLSs [60, 61]:

$$\frac{1}{Q_i} = F\delta_{TLS}^0 \frac{\tanh(\hbar\omega_r/2k_B T)}{(1 + \langle n \rangle/n_c)^\beta} + \delta_{other}. \quad (2.4)$$

I demonstrate the impact of the various parameters of this model on the achieved  $Q_i$  in Fig. 2.2.

In this model,  $F$  quantifies the total TLS filling factor (the ratio between the electric field stored in the lossy regions vs. the total energy of the mode), and  $\delta_{TLS}^0$  represents the effective dielectric loss tangent of the TLS ensemble.

The temperature dependent term  $\tanh(\hbar\omega_r/2k_B T) \approx 1$  at 10 mK and GHz frequencies. Raising the sample temperature may lead to an increase of resonator  $Q_i$  (11% for a 6 GHz resonator at 100 mK); however, this may come at the expense of an increased quasiparticle loss negating or even countering this improvement.

The power dependence of TLS loss is governed by two parameters:  $n_c$  and  $\beta$ . The critical photon number  $n_c$  gives the power necessary to saturate a single TLS.  $\beta$  tells how sharply the TLS bath saturates with applied drive power. In the traditional standard tunneling TLS model for non-interacting TLSs,  $\beta = 0.5$ , i.e. TLS loss scales with  $1/\sqrt{1 + \langle n \rangle/n_c}$  [62, 63]. However, this often fits high-quality superconducting resonators poorly, as they tend to have a weaker power-dependence, which we see in the appended papers **A–D**.

To accommodate this discrepancy, the modeled power-dependence acquires a phenomenological fitting parameter – the aforementioned  $\beta$  – either by directly replacing the square root with a power of  $\beta$  as is done in Eq. 2.4, or by placing the fitting parameter on the  $(\langle n \rangle/n_c)$  term underneath the square root, i.e.  $1/\sqrt{1 + (\langle n \rangle/n_c)^\beta}$  [55, 60]. The two model alterations produce similar results, with a deviation around  $n_c$  for small  $\beta$ . We find that we can fit our data to both models with good accuracy, but fitting to Eq. 2.4 can give smaller error bars. The reasons for the deviation of  $\beta$  from the standard tunneling model are not fully known, but include the TLS-TLS interactions neglected in the standard model which only considers interactions between a TLS and the device, as well as the presence of various species of TLS with different power-saturation behaviours (this can especially be the case in devices with different species at the various materials interfaces), non-resonant absorption, and the geometric variation of the field density along a distributed structure (such as is the case in our quarter-wavelength resonators).

Finally,  $\delta_{other}$  represents other, non-TLS related losses. These can be due to loss channels such as radiation, magnetic flux noise, or quasiparticle loss. The  $\delta_{other}$  term is presumed to be power-independent and indeed, in our measurements, we observe that the  $Q_i$  vs.  $\langle n \rangle$  curve flattens out once the TLS bath is saturated, and the power-independent  $Q_i$  at high  $\langle n \rangle$  ( $\gtrsim 10^7$ ) shows the limitations imposed by the next-dominant loss mechanism.

Thus, we can distinguish TLS loss from other mechanisms in superconducting res-



onators by measuring  $Q_i$  as a function of applied drive power. Since resonators and qubits are typically built from identical materials and undergo similar fabrication procedures (barring the addition of Josephson junctions in qubits), they also share many of the same loss channels. Due to the relative ease of resonator fabrication and characterization, resonators become powerful tools for investigating TLS loss in superconducting qubits.

### 2.2.2 Modeling TLS loss

The electromagnetic mode of a coplanar waveguide (CPW) structure sees various materials, which can have vastly different dielectric properties. These materials and their interfaces, which can all contribute to dielectric loss, are illustrated in the CPW cross-section in Fig. 2.3 and include the substrate, air, and the various interfaces: substrate-metal (SM), substrate-air (SA), and metal-air (MA). As the substrate and the metal can grow distinct oxides that contribute to dielectric loss, a corner region at the junction of the MA and SA interfaces can be defined.

The TLS loss parameter  $F\delta_{\text{TLS}}^0$  can be modeled as a weighted sum of the individual contributions of the various lossy regions, which each contribute to the total loss with a dielectric loss tangent  $\delta_i^0$ , weighted by the corresponding participation ratio  $p_i$  [51, 64]:


$$\frac{1}{Q_{\text{TLS}}} = F\delta_{\text{TLS}}^0 = \sum_i p_i \delta_i^0. \quad (2.5)$$

The energy loss participation ratio  $p_i$  is defined as the energy stored in the lossy region, divided by the total energy of the resonant mode:

$$p_i = \frac{W_i}{W_{\text{total}}} = \frac{\int_{V_i} \epsilon_i \vec{E}^2(\vec{r}) d\vec{r}}{\int_V \epsilon_{\text{tot}} \vec{E}^2(\vec{r}) d\vec{r}}. \quad (2.6)$$

The dielectric loss tangent  $\delta_i^0$  of a given material is strongly temperature and frequency dependent [65]. Acquiring precise values for the  $\delta_i^0$  of a specific material, especially for an interface, at millikelvin temperatures is nontrivial and can require elaborate cryogenic microwave measurement strategies such as in Ref. [66]. Additionally, the loss tangents extracted from these experiments may vary between devices fabricated in different fabrication rounds, especially if the fabrication process is modified, which we explore in detail in papers **A** and **B**.

Therefore, predicting the  $F\delta_{\text{TLS}}^0$  for a given device design is challenging, as inaccuracies in the assumed  $\delta_i^0$  values, as well as in the effective thicknesses of all regions



**Figure 2.3:** Micrograph of a CPW resonator coupled to an input/output transmission line, with an illustration of the CPW cross-section, showing the various materials and their interfaces (not to scale).

contributing to dielectric loss, will distort the result.

## 2.3 Quasiparticle loss

In a superconductor, electrons paired into superconducting Cooper pairs, and unpaired “normal” electrons (referred to as quasiparticles), coexist in an equilibrium dictated by temperature in the ideal case. Just below  $T_c$ , only a few of the electrons form Cooper pairs, while at  $T = 0$ , all electrons have paired up (unless one lone electron remains in case of an odd total number of free electrons). This results in a temperature-dependent complex conductivity  $\sigma$  in superconductors, as described by the Mattis-Bardeen theory [67]:

$$\sigma(\omega, T) = \sigma_1(\omega, T) - i\sigma_2(\omega, T). \quad (2.7)$$

At DC, current flows through the path of least resistance - the superconducting Cooper pairs - and the current transport is lossless ( $\sigma_1(\omega = 0, T) = 0$ ). However, at AC, some current will be carried by quasiparticles, which is a cause of resistive loss ( $\sigma_1(\omega, T) \neq 0$ ). This temperature dependent  $\sigma$  causes a temperature dependent dissipation rate which scales with the real part of the complex conductivity  $\sigma_1$ , as well as a temperature dependent inductance (ergo device frequency), which scales with the

imaginary part  $\sigma_2$  [68]:

$$\frac{1}{Q_i(T)} - \frac{1}{Q_{i,0}} = \alpha \frac{\sigma_1(T) - \sigma_{1,0}}{\sigma_{1,0}}, \quad (2.8)$$

$$\frac{f_r(T) - f_{r,0}}{f_{r,0}} = -\frac{\alpha}{2} \frac{\sigma_2(T) - \sigma_{2,0}}{\sigma_{2,0}}, \quad (2.9)$$

where  $Q_{i,0}$ ,  $f_{r,0}$ ,  $\sigma_{1,0}$  and  $\sigma_{2,0}$  are values at  $T = 0$ , and  $\alpha = L_k/L$  is the ratio of kinetic inductance  $L_k$  to total inductance  $L$ . This ratio depends on the superconductor properties, as well as the geometry of the superconducting structure ( $\alpha$  increases for narrower structures). I show the effects of the temperature-dependent  $\sigma$  on the device  $Q_i$  and  $f_r$  in Fig. 2.4.

The loss rate due to quasiparticles can be expressed as a function of the thermal equilibrium ratio  $x_{qp} = n_n/n_s$  of the density of normal quasiparticles  $n_n$  and the density of the superconducting Cooper pairs  $n_s$  as [69, 70]:


$$\begin{aligned} \Gamma_{qp} &= \frac{\omega_r}{\pi} \alpha \sqrt{\frac{2\Delta_s}{\hbar\omega_r}} x_{qp} = \frac{\omega_r}{\pi} \frac{L_k}{L} \sqrt{\frac{2\Delta_s}{\hbar\omega_r}} \frac{\sqrt{2\pi\Delta_s k_B T}}{\Delta_s} e^{-\Delta_s/k_B T} \\ &= \frac{\omega_r}{\pi} \alpha \sqrt{4\pi \frac{k_B T}{\hbar\omega_r}} e^{-\Delta_s/k_B T}, \end{aligned} \quad (2.10)$$

where  $\Delta_s$  is the superconducting energy gap of the material,  $k_B$  the Boltzmann constant, and  $T$  temperature.

It is common to find higher quasiparticle densities in devices than what is expected from the equilibrium distribution associated with the cryostat temperature [71, 72]. The origin of these excess nonequilibrium quasiparticles is not yet fully understood, but some known sources are Cooper pair-breaking events such as the incidence of high-energy photons with frequencies larger than  $\Delta$ , or even cosmic rays. In addition to causing resistive loss, high-energy incidences may give the generated quasiparticles large enough momentum to stimulate emission in the device.

While one may expect nonequilibrium quasiparticles to accumulate with consecutive pair-breaking events and over time fully take over all the electrons present in the device, fortunately, that is not true. Quasiparticles may recombine into Cooper pairs again, and so the density of nonequilibrium quasiparticles will depend on the quasiparticle generation rate, as well as their recombination rate [69].

To reduce the number of equilibrium quasiparticles, we can use materials and geome-



**Figure 2.4:** Experimental data showing a temperature sweep of a CPW resonator. The drive power corresponds to  $\sim 200$  photons at low temperature.

tries with a higher  $\Delta_s$ , and reduce the device temperature to  $T \ll T_c$ . The generation of nonequilibrium quasiparticles by incident high-energy particles may be prevented with good radiation shielding of the device, as well as high frequency filtering on the input/output lines. Strategies such as normal-metal quasiparticle traps [73, 74] and active cooling [75] have also seen considerable success.

In addition to the resistive loss due to quasiparticles, qubits suffer additional decoherence from quasiparticles tunneling through the Josephson junction. This can be combated through engineering the junction electrode gaps [76].

## 2.4 Magnetic noise

The Meissner effect of magnetic field expulsion is a defining feature of superconductors. While bulk type-I superconductors expel magnetic fields fully when cooled below  $T_c$ , type-II superconductors pin single quanta of magnetic flux into so-called Abrikosov vortices [42]. These vortices also occur in thin films of type-I superconductors, with defects such as grain boundaries and normal inclusions serving as favored vortex-pinning sites [77]. Another source of magnetic flux noise is the stochastic flipping of spins in the device materials.

While suitably positioned normal-core vortices may be beneficial since they can act as quasiparticle traps, the positional instability of these vortices as the trapped flux moves between different defects acting as pinning centers can cause dissipation. This

effect can be mitigated through the inclusion of flux-trapping holes in the device design [77].

In flux-tunable qubits, magnetic-flux fluctuations cause additional issues in the form of frequency noise [48].

## 2.5 Radiation

Loss due to radiation can take different forms, the simplest of which is spontaneous emission into free space and antenna radiation. This radiation rate is determined by geometry and frequency. While the exact geometric dependence of radiation for qubits and coplanar structures is non-trivial to compute [78], a simplified semiclassical argument considering the spontaneous emission of an electric dipole predicts the frequency dependence of spontaneous emission as  $\Gamma_{rad} \propto \omega^3$  [8].

Besides spontaneous emission into free space, the resonant mode can couple to its surrounding elements and decay into potential slotline modes on the chip, the cavity mode of the sample box, or back into the transmission line. The qubit, which might otherwise be hard-pressed to decay into the far-away transmission line, might do so through the coupling to its readout resonator in a phenomenon dubbed the Purcell effect.

### 2.5.1 Purcell decay

The rate of spontaneous emission of a quantum system can be enhanced by its environment in the so-called Purcell effect [79]. In our case, the qubit decays into the resonant mode of the readout resonator, which in turn decays into the transmission line.

The Purcell decay rate is governed by:

$$\Gamma_P = \kappa \frac{g^2}{\Delta^2}, \quad (2.11)$$

where  $\kappa = \omega_r/Q$  is the readout resonator decay rate,  $\Delta = |\omega_q - \omega_r|$  is the frequency difference (detuning) between the qubit and the resonator, and  $g$  is the coupling between the qubit and the resonator. We take a close look at how to extract the Purcell decay rate in chapter 5.

The Purcell decay rate can be decreased through decreasing  $g$  and  $\kappa$ , and increasing  $\Delta$ . This may however complicate readout, as it results in smaller dispersive shift  $\chi$  (for more details on dispersive readout, see chapter 5.3.1). A smaller  $\chi$  can compromise the readout contrast, which necessitates more averaging, slowing down readout.

An alternative to compromising the qubit and resonator design parameters is implementing a so-called Purcell filter between the resonator and the transmission line, or on the transmission line itself [80]. This may, however, make qubit control more cumbersome by necessitating a separate qubit drive line.

# CHAPTER 3

---


## Device design

---

In this chapter, I discuss some basic design considerations for the devices containing qubits and resonators discussed throughout this thesis. The primary purpose of these devices is the investigation of the decoherence mechanisms in superconducting quantum circuits. Therefore, the device design is kept as simple as possible in order to reduce the number of parameters that need to be taken into account when disentangling the various sources of decoherence the device is experiencing. For this reason, we do not add the typical features of our multiqubit quantum processors such as dedicated qubit drive lines, Purcell filters, or frequency-tunable elements to our test devices, as these elements turn the system more complex and can introduce additional loss sources [8, 81, 82].

Our typical qubit test device, which I show in Fig. 3.1, consists of a transmission line coupled to a coplanar waveguide (CPW) resonator, which is in turn coupled to a transmon qubit. The resonator serves for both qubit readout and control. The qubit itself comprises two electrodes separated by a Josephson junction (JJ), shunted by a large cross-type capacitor. This qubit type is referred to as the cross-type transmon, or Xmon [83].

In our decoherence studies, we often measure a CPW resonator coupled only to an input/output transmission line, omitting the qubit altogether. This is because the qubits



**Figure 3.1:** Circuit diagram (a) and false-colour micrograph (b) of a transmon qubit, coupled to a readout resonator, which is in turn coupled to an input/output transmission line. The inset in b is a scanning electron micrograph of the qubit’s JJ element, consisting of two superconducting electrodes separated by a thin layer of insulator. The fringed patches connect the JJ electrodes to the remaining circuitry (the qubit capacitor on one end, and the ground plane on the other).

and resonators are fabricated using the same materials and fabrication procedures, and thereby share many of the same sources of loss. However, resonators are faster to fabricate and characterize, and the characteristic dependence of their performance as a function of temperature, applied drive power, or even magnetic field, is more straightforward to interpret when assigning the measured loss rates to the contributing loss sources.

In addition, we surround our qubits and resonators with an array of flux-trapping holes to suppress loss due to magnetic vortex migration [77, 84]. The flux-trapping hole array consists of squares with a  $2\ \mu\text{m}$  long side and a center-to-center distance of  $10\ \mu\text{m}$ , etched into the metal film.

### 3.1 Coplanar waveguide resonator design

A coplanar waveguide (CPW) resonator [85] consists of a center conductor separated from a ground plane by a gap. Both the center conductor and the ground plane are atop a dielectric substrate. In Fig. 3.2, I show a micrograph of such a CPW resonator, together with an illustration of the CPW cross-section.

The electric field distribution depends on the boundary conditions of the CPW. In



our devices, we use so-called quarter-wavelength ( $\lambda/4$ ) resonators, where one end of the center conductor forms an open circuit, and the other end is shorted to the ground plane (see Fig. 3.2). As the name implies, the length of this resonator is one quarter of its resonance wavelength. The resonator is meandered to make its footprint more compact, as it is quite substantial in length ( $\sim 1$  cm). The bending radius is kept at least five times the width of the center conductor.

The resonance frequency is determined by the resonator inductance  $L$  and capacitance  $C$  as:

$$\omega_r = \frac{1}{\sqrt{LC}}. \quad (3.1)$$


We determine the capacitance and inductance per length of the center conductor via conformal mapping [86]. These values are governed by the center conductor width  $w$ , the gap to ground  $g$ , the dielectric constants of the substrate and the gap material (vacuum in our case), as well as the proximity to other conductors (in our case, this could be the lid of the sample box, which we keep at least 5 mm away from the substrate so as to not interfere with the resonant mode).

In our standard design where  $w = 2g = 20 \mu\text{m}$ , we keep the impedance of our resonators close to  $50 \Omega$ , which is the typical impedance in microwave equipment. Keeping the impedance of a resonator coupled to a  $50 \Omega$  transmission line in a notch configuration to the matching value is, however, not strictly necessary: while an impedance mismatch will affect the coupling between the two elements, this should not introduce unwanted signal reflections. In our experiments published in Ref. [87], we find that a wider gap to ground may improve the performance of our resonators at low drive powers, despite the impedance mismatch this induces. On the other hand, we keep the impedance of the input/output transmission line strictly to  $50 \Omega$  to reduce unwanted reflections when the device is placed in series with the  $50 \Omega$ -matched measurement lines.

The coupling strength between the resonator and the transmission line is determined by the length of the center conductor section that is facing the transmission line, as well as by its proximity to the transmission line [88].

We decide the target coupling based on application. For a qubit readout resonator, a strong coupling is beneficial for increasing the readout speed. On the other hand, a strongly coupled resonator increases the Purcell decay rate of the qubit into the transmission line through the resonator (See Sections 2.11 and 5.3.5). For the readout resonators on our qubit test chips, we keep the coupling quality factors  $Q_c$  [89] in the range of  $10 - 20 \times 10^3$ .

For bare resonators intended for loss studies, we aim for the so-called *critical* coupling




**Figure 3.2:**  $\lambda/4$  CPW resonator, coupled capacitively to an input/output transmission line via the open-circuit end. The resonator is surrounded by an array of flux-trapping holes. The cartoon shows a CPW cross-section with its constituting materials and their interfaces.

where  $Q_i \sim Q_c$ . This is because the accuracy of the data fitting routine used to extract  $Q_i$  is most accurate when  $Q_i \sim Q_c$  [90]. I discuss this in more detail in Section 5.2.

The choice of the widths of  $w$  and  $g$  is based on the geometric scaling of the contributing loss sources. Smaller geometries tend to have larger energy participation ratios at lossy interfaces, which makes these geometries useful for probing TLS loss, but less likely to achieve high coherence. Larger geometries, on the other hand, can become more susceptible to other losses, such as radiation.

In Ref. [87], we explore the geometric scaling of losses in our devices in more detail. An excerpt from this study in Fig. 3.3, performed with resonators designed with the  $g : w = 1 : 2$  ratio kept constant, indicates that our resonators reach optimal low-power performance at the  $g/w/g = 15/30/15 \mu\text{m}$  geometry. Smaller geometries are more strongly limited by TLS loss, and the  $Q_i$  vs.  $\langle n \rangle$  curve is steep. For larger geometries, the curve flattens out, and the resonators are less limited by TLS loss. However, another loss source becomes dominant at the  $g/w/g = 25/50/25 \mu\text{m}$  size and higher, which suppresses the  $Q_i$  across the entire measured power span. The high-power loss of the narrower resonators ( $w \leq 20 \mu\text{m}$ ) has a frequency dependence that can be fitted with  $Q_{\text{rad}} \propto f^{-\alpha}$ , with  $\alpha$  ranging between 1.8 and 2.3, which is close to the  $\alpha = 2$  predicted



**Figure 3.3:** Resonator loss as a function of applied drive power for resonators with different CPW geometries. The resonance frequency is  $\sim 5$  GHz.

by the theory in Ref. [91]. Our wider resonators do not fit this model, and can therefore be limited by other types of radiations with a different frequency dependence, or by loss due to quasiparticles or magnetic flux.

The approach of diluting the electric field into a larger volume to minimize interfacial participation ratios is taken into a successful extreme in 3D cavity resonators. I discuss the superior performance of these 3D cavities with regards to the more compact 2D geometry in Section 7.3.

## 3.2 Qubit design

When adding a transmon qubit to the circuit described above, there are several things to consider, including: the qubit frequency  $\omega_q$ , the qubit capacitance  $C_\Sigma$  and the related anharmonicity  $\alpha$ , the coupling strength between the qubit and the resonator  $g$ , and their mutual detuning  $\Delta$ . The coupling between the resonator and the transmission line might also need to be adjusted with the addition of the qubit, to optimize readout or minimize Purcell decay, as discussed above.

The transmon qubit frequency is determined by the Josephson energy  $E_J$ , and the

charging energy  $E_C$  [92]:

$$\omega_q = \frac{\sqrt{8E_J E_C} - E_C}{\hbar}. \quad (3.2)$$

The charging energy  $E_C$  is a function of the sum of all capacitances in the system  $C_\Sigma$ , which includes the Josephson junction capacitance, as well that of the qubit capacitor:

$$E_C = \frac{e^2}{2C_\Sigma}. \quad (3.3)$$

$E_C$  also dictates the qubit's anharmonicity  $\alpha$ , which relates the transition frequency  $\omega_{01}$  between the states  $|0\rangle$  and  $|1\rangle$ , and the transition frequency  $\omega_{12}$  between  $|1\rangle$  and  $|2\rangle$  as:

$$\alpha = \omega_{12} - \omega_{01} = -E_C/\hbar. \quad (3.4)$$

In order to be in the charge-insensitive transmon regime, the qubit needs to meet the condition  $E_J/E_C \gg 1$ . In this regime, the qubit is isolated from charge noise, although at the expense of a reduced anharmonicity compared to the Cooper pair box from which the transmon is originally derived [8]. In our devices, we typically target  $\alpha \sim -2\pi \times 200$  MHz.

The Josephson energy, which is the potential energy accumulated in the junction as the supercurrent flows through it, is a function of the junction inductance, and thereby its critical current  $I_c$ :

$$E_J = \frac{\Phi_0 I_c}{2\pi}, \quad (3.5)$$

where  $\Phi_0$  is the magnetic flux quantum. We can calculate the critical current of the JJ from its normal-state resistance using the Ambegaokar-Baratoff relation, which at near-zero temperatures can be approximated by:

$$I_c R_N = \frac{\pi \Delta_s}{2e}, \quad (3.6)$$

where  $\Delta_s$  is the superconducting gap. Fortunately, this relation enables us to calculate the qubit frequency from the room-temperature properties of its JJ. Since the junction area is small, variations can influence the frequency significantly – I discuss this in more detail in Section 4.11.3.

Choosing the right qubit frequency is an optimization problem. Lower-frequency qubits of similar quality will automatically yield longer  $T_1$  (see Eq. 2.1). Lower frequencies also have lower rates of spontaneous decay [8], and are less susceptible to quasiparticle loss (see Eq. 2.10). However, lower frequencies also increase the chances

of spontaneous thermal excitations as  $\hbar\omega_q$  approaches  $k_B T$ . We need to also consider the frequency of the resonator, and the mutual qubit-resonator detuning  $\Delta = |\omega_q - \omega_r|$ . The two elements need to be close enough in frequency to interact – however, small detunings increase the Purcell decay rate, which can be a limiting factor in high-quality qubits (see Eq. 2.11). As it is the qubit performance that is of interest to us, more so than the performance of the resonator, we set  $\omega_q < \omega_r$  due to the aforementioned reasons. We keep  $\Delta$  at a minimum of 1 GHz as a standard – in order to suppress Purcell decay in our longer-lived qubits, we increase  $\Delta$  to 2–3 GHz.

We keep the coupling strength between the qubit and the resonator to  $g \sim 2\pi \times 50$  MHz. A stronger coupling facilitates faster interactions between the resonator and the qubit, enabling faster readout and control of the qubit through its resonator. On the other hand, predictably, this stronger coupling enhances the Purcell decay rate (see Eq. 2.11).



# CHAPTER 4

---

## Device fabrication

---


This chapter covers the qubit fabrication process, starting with an overview of the full fabrication procedure, which is intended to act as a big picture reference for the subsequent deep dive into the specifics of the individual techniques. The ensuing section contains a discussion on the choice of materials and fabrication processes, followed by a section presenting selected challenges one may encounter in the outlined fabrication processes, with some suggestions on how to overcome them.

While the microwave frequency regime in which we operate our resonators sets their length on the order of  $\sim 1$  cm, their cross-sectional dimensions are significantly smaller on the order of  $\sim 10\ \mu\text{m}$ , and the qubits' JJ dimensions are  $\sim 100$  nm, landing our device fabrication process in the domain of micro- and nanofabrication. Fabrication procedures requiring nanoscale precision are sensitive to particle contamination and vibrations, as well as environmental factors that affect the rate of chemical reactions such as temperature, humidity, and pressure. This necessitates a laboratory with good control over such parameters, such as a cleanroom laboratory. The qubits presented in this thesis were fabricated in the class 100 cleanroom of Myfab Chalmers, which permits up to 100 particles with a diameter  $> 0.5\ \mu\text{m}$  per cubic foot of air.

## 4.1 Process overview

The qubit fabrication process starts with a bare substrate wafer that is 2"-4" in diameter. This wafer size is typical in research applications - for an industrial-scale production, the wafer size would typically be larger.

Our standard qubits, shown in Fig.4.1, are fabricated using aluminium (Al) as the superconductor and silicon (Si) as the substrate. The junctions are Al/AlOx/Al, where AlOx is native Al oxide.



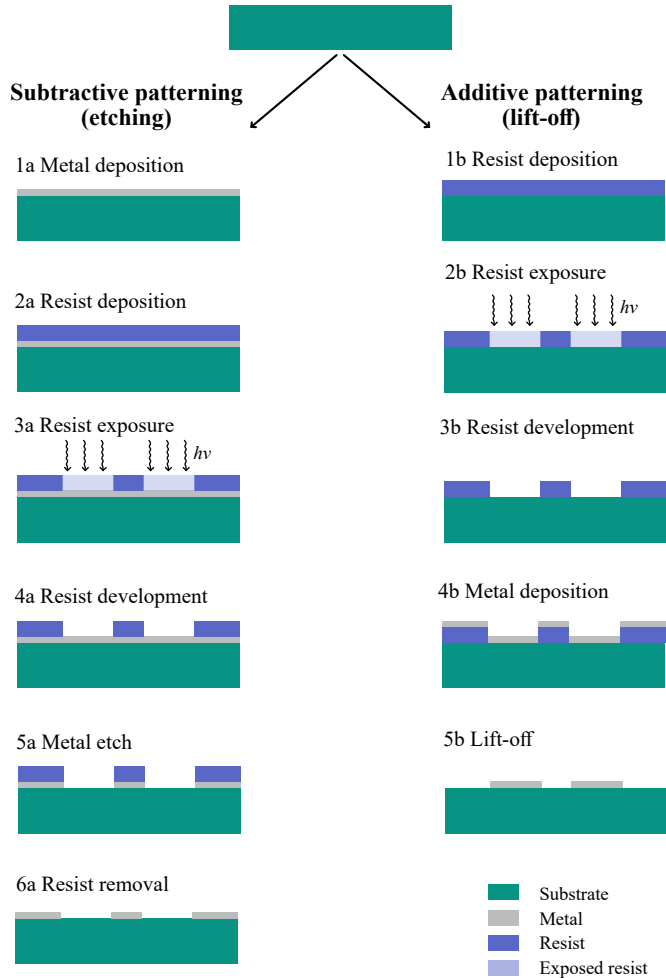
**Figure 4.1:** Micrograph of a transmon qubit device. The optical image depicts the circuitry defined in an Al film by photolithography and wet etching. The scanning electron micrograph inset depicts the Josephson junction, defined by electron beam lithography and metal lift-off.

Our fabrication process utilizes the two conventional patterning techniques – subtractive patterning via etching, and additive patterning via lift-off – illustrated in Fig. 4.2.

In **subtractive** patterning, a continuous layer of metal is deposited on the substrate, and is then covered by a layer of polymer (a positive resist). The resist is then patterned by exposing selected areas to a laser or electron beam, which chemically changes the resist in those areas. The exposed resist is then dissolved in a developer solution. In the subsequent etching step, the metal in areas unprotected by resist is etched. Afterwards, the resist can be removed, leaving behind a substrate with a patterned metal on top.

In **additive** patterning, the resist is deposited directly on the substrate, then a pattern is exposed and developed in the resist. A metal is deposited on the wafer, making direct contact with the substrate in the exposed areas, while being spaced by the resist where unexposed. Once the resist is removed, only the metal in direct contact with the substrate remains on the wafer, while the dissolving resist lifts off the unwanted metal.





**Figure 4.2:** Additive and subtractive patterning processes with positive resists. In **subtractive** patterning, the metal is deposited first and then covered with a resist. The resist is then patterned through exposure and development. The areas of the metal that remain unprotected by the patterned resist are then etched away. In **additive** patterning, the resist is deposited directly onto the substrate. The resist is then patterned, after which a metal film is deposited on the wafer. The wafer areas unprotected by resist will make direct contact with the metal, while the metal separated by the resist layer will be lifted off during resist removal.

The qubit fabrication process consists of the following steps:


- Substrate cleaning
- Metal deposition + in-situ oxidation
- Patterning
  - Resist deposition + soft bake
  - Resist patterning (exposure, development, ashing)
  - Pattern transfer into the metal film (etching)
  - Resist removal
- JJ fabrication
  - Resist stack deposition + soft bake
  - Resist patterning (exposure, development, ashing)
  - JJ and patch deposition
  - Lift-off
- Dicing
- Chip cleaning and packaging



## 4.2 Substrate preparation

The quality of the substrate-metal interface is of utmost importance to the performance of our devices. We strive to achieve an atomically clean interface between the Si and Al, avoiding the presence of any contaminants or oxides at this interface.

The three most established industrial processes for Si wafer cleaning are the standard clean 1 (SC1) which removes particles, organics, and some metals, standard clean 2 (SC2) which removes metallic contaminants, and the hydrofluoric acid (HF) dip which removes silicon oxide.


SC1 involves submerging the wafer into a heated mixture of ammonium hydroxide ( $\text{NH}_4\text{OH}$ ), hydrogen peroxide ( $\text{H}_2\text{O}_2$ ), and deionized water. In our process we expose the wafer to a 1:1:5 solution of  $\text{NH}_4\text{OH}:\text{H}_2\text{O}_2:\text{H}_2\text{O}$  heated to  $80^\circ\text{C}$  for 10 min. The process works by a repeating tandem of Si surface oxidation (Eq. 4.1) and a subsequent oxide removal (Eq. 4.2). This can lead to surface roughening, and leaves the Si surface oxidized and hydrophilic [93], as illustrated in Fig. 4.3.





**Figure 4.3:** Oxidized Si surface, with a water molecule attached through hydrogen bonds.


The resulting oxide layer, or the native  $\text{SiO}_2$  layer of an untreated Si wafer, can be removed by dipping the wafer in dilute HF. The HF is selective towards  $\text{SiO}_2$  and does not further etch the Si substrate. In the first stage of the etching process, the HF attacks the  $\text{SiO}_2$ , leaving behind an F-terminated surface, as shown in Fig 4.4.



**Figure 4.4:** Silicon oxide removal by HF. **a** Lone pair binds a  $\text{H}^+$  cation from the HF solution. **b** The resulting positive charge on the O attracts the electrons in the Si-O bond, leaving a positive charge on the Si (**c**). **c** The positively charged Si attracts a  $\text{F}^-$  from the solution. Another lone pair on O binds  $\text{H}^+$ , the resulting positive charge on O attracts the Si-O pair, releasing a water molecule in **d**. **d-g** The process repeats until the surface is F-terminated, where the strong polarity of the Si-F bond results in the partial charges  $\delta$ .

Initially it was believed that the Si remains F-terminated, as the Si-F bond is one

of the strongest single bonds known in chemistry, with a bond energy of  $\sim 6\text{ eV}$  [94]. However, the highly polarized nature of this bond makes it vulnerable to further HF attack, releasing  $\text{SiF}_4$  and leaving the surface H-terminated [95] (Fig. 4.5).



**Figure 4.5:** The transformation of a F-terminated Si surface to H-terminated, releasing  $\text{SiF}_4$ . The arrows indicate the occurring reactions, and  $\delta$  the partial charges resulting from the strong polarity of the Si-F bonds.

The H-passivated surface is conveniently hydrophobic, and will resist reoxidation for some time. In low-oxygen environments ( $<0.1$  ppm) the surface can survive for weeks without forming an oxide layer, but in normal cleanroom air, a 0.5 nm layer of oxide will form within a few hours [93].

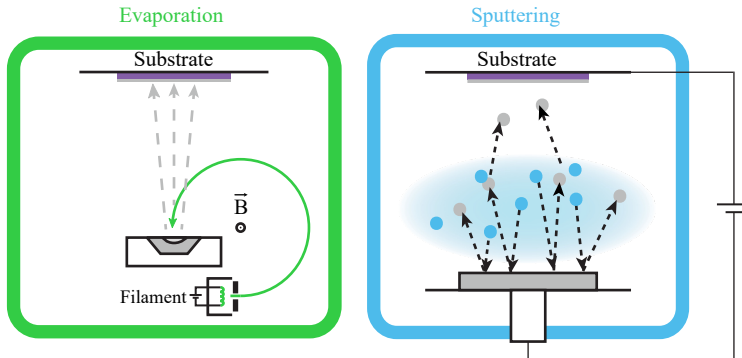
SC2 is an ionic clean where a heated solution of hydrochloric acid (HCl) and  $\text{H}_2\text{O}_2$  is used to remove metallic contaminants such as alkali residues, residual trace metals, and metal hydroxides. The solution does not etch Si or  $\text{SiO}_2$  and has no further beneficial surfactant activity beyond the removal of metallic contaminants. As we do not suspect metallic contamination in our wafers prior to processing, we do not perform this cleaning procedure on our wafers.

For cleaning non-silicon substrates which do not respond to the standard Si-cleaning solutions, the so-called piranha solution can be used. This solution is a mixture of sulfuric acid ( $\text{H}_2\text{SO}_4$ ) and  $\text{H}_2\text{O}_2$ , which vigorously decomposes organic matter. This method can be used on Si, as well as other substrates, but extreme caution should be applied during handling due to the hazardous nature of this process.

### 4.3 Thin film deposition

There are various methods for depositing metallic thin films, most commonly physical vapor deposition (sputtering, evaporation), chemical vapor deposition, atomic layer deposition, and electrochemical deposition [93].

In this work we keep to physical vapor deposition, where the substrate and a source



**Figure 4.6:** PVD by evaporation and sputtering. In **evaporation**, the tungsten filament driven by high voltage emits an electron beam, which is deflected by a magnetic field towards the target. This melts the metal, which evaporates and makes its way towards the substrate. In **sputtering**, the applied voltage ignites a plasma, which knocks out atoms from the target.

of the metal (also called a target or crucible) are placed in a high vacuum environment. Then, material is ejected from the target surface in the direction of the substrate, where it precipitates and forms a thin film.

The deposition chamber is kept under good vacuum ( $\lesssim 1 \times 10^{-7}$  mbar) during deposition in order to provide a mean-free path for the ejected material on its journey towards the substrate, to minimize the presence of atmospheric reagents, to avoid interference with the crystallization process, and to avoid contamination and defects.

The crystalline structure of the deposited film is affected by substrate properties (lattice matching, roughness), the temperature and pressure during deposition, the deposition rate, and film thickness.

### 4.3.1 Evaporation

Evaporation is a physical vapor deposition (PVD) technique where a solid target is heated past its melting point. As hot metals have high vapour pressures, the metal vapour will disperse in the high vacuum of the deposition chamber, some of it condensing on the substrate facing the target.

The metal can be evaporated by resistive or electron beam-assisted heating. We evaporate our Al using electron-beam (e-beam) deposition, where a tungsten filament is biased by a high voltage ( $\simeq 10$  kV), emitting electrons which are guided towards the target by a magnetic field (see Fig. 4.6).

In our standard qubit fabrication process, we always evaporate the JJs, as well as the patches connecting the the JJ electrodes to the rest of the circuitry. The remaining circuitry is often evaporated as well.

### 4.3.2 Sputtering

In sputtering, a heavy gas is introduced in the chamber, and a voltage applied across the substrate-target space ignites the gas into a plasma. This plasma attack knocks out material from the target, which then makes its way towards the substrate (see Fig. 4.6).

In nonreactive ion sputtering, such as when depositing Al, we want the plasma ions to be heavy enough for optimal battering effect, but also inert such as to not influence the ejected metal atoms chemically. A popular candidate for such purposes would be  $\text{Ar}^+$ . For depositing compounds such as oxides or nitrides, the sputtering gas can be complemented by another, reactive gas.

In our qubit fabrication process, the wiring layer for the transmission line, resonators, and capacitors, and other supporting elements can be sputtered.

## 4.4 In-situ passivation

Most metals inevitably oxidize in contact with ambient air, albeit on varying timescales. Some metals, like Al, self-passivate, meaning that after a certain oxide thickness has been reached, this oxide will protect the underlying metal and further corrosion is imperceptible [96]. Other metals can oxidize all the way through, even if the kinetics of this oxidation changes with the oxide thickness.

Accepting the inevitability of an oxide layer afflicting the metal-air interface, one can mitigate the consequences by controlling how the oxide is grown. Oxidation in ambient air can introduce variations in the oxide thickness and structure based on the atmospheric conditions controlling the kinetics of oxidation (temperature, humidity, pressure), and contamination. These variations can affect device performance and etch reproducibility, as well as introduce defects into the film (see Sec.7.1.1). Oxidation reproducibility can be increased by oxidizing the film in pure oxygen inside the PVD chamber directly after deposition. In our standard process, we oxidize every newly deposited Al layer in-situ.

One can also attempt to avoid this dielectrically disadvantaged fate by covering an oxidation-prone metal with another material in-situ directly after deposition. Oxidation-resistant metals are practical for this purpose. Due to the superconducting proximity

effect, it is not necessary for the capping metal to be superconducting, as long as the thickness of the capping film is well below the coherence length of the superconductor (see Sec. 7.1.2). However, the capping film cannot be arbitrarily thin, due to practical concerns such as the bottom metal roughness and the formation of pinholes in very thin films. Additionally, intermetallization between the two metals can have an anomalous effect on the  $T_c$  of the proximitized stack, which limits the choice of capping metals for metals which intermetallize readily with a wide variety of metals (such as Al).

## 4.5 Resist patterning

The process to create a resist-based etching mask consists of resist deposition (Fig. 4.2 **1b** and **2a**), exposure (Fig. 4.2 **2b** and **3a**), development (Fig. 4.2 **3b** and **4a**), and often a mild ashing.

The resist is typically deposited on the substrate or the metal by spin-coating, although other options such as spray-coating or lamination are available. In spin-coating, a dollop of resist suspended in a solvent is poured on the center of the wafer, then the wafer is spun at a high rotation speed to achieve uniform coverage across the wafer area. Afterwards, the solvent is evaporated during a “soft-bake” on a hot plate or in an oven. The resist thickness affects the minimum achievable feature size, as well as the maximum thickness of a lift-off metal.

A resist is sensitive to exposure to a narrow spectrum of optical wavelengths, or to electron currents. Exposing select areas of resist changes the resist chemically, affecting the solubility of these areas of resist in a solution referred to as the developer. Depending on whether the solubility of the exposed resist is increased or decreased by exposure, we refer to the resist type as positive or negative, respectively.

### 4.5.1 Optical lithography

Optical lithography typically involves exposures to light with wavelengths in the ultra-violet (UV) range [93]. The development of photoresist is typically an acid-base reaction process, often containing tetramethylammonium hydroxide (TMAH). The development is stopped by immersion in water.

Traditionally, this is done by shining a UV lamp on the wafer, while a “hard mask” with a predefined pattern is placed in between the wafer and the source, blocking the light where necessary. This mask, placed as close to the wafer as possible or in direct contact with the wafer to minimize the effects of diffraction, is made from an optically

transparent material such as quartz, and the pattern is defined on the mask with a thin layer of chromium.

An alternative to masked flood exposure is direct-write lithography (DWL), where a laser write head programmed with a digital pattern scans the wafer, selectively exposing the target areas.

The advantage of DWL is the utmost flexibility with which a pattern can be modified between exposures, which is highly practical for research applications. Additionally, DWL does not require hard masks, which take additional resources to produce and degrade with use (particularly during contact lithography). The disadvantage is that it is significantly slower than flood exposure, which can take mere seconds per wafer, as opposed to the 1 – 1000 mm<sup>2</sup>/min range available for DWL, depending on feature size [97]. Due to this speedup, masked exposure remains popular in industrial applications.

The theoretical minimal feature size of an optically defined pattern is given by [98]

$$d_{min} = \frac{3}{2} \sqrt{\lambda \left( g + \frac{1}{2} h \right)}, \quad (4.3)$$

where  $\lambda$  is the optical wavelength,  $g$  is the gap between the photoresist and the mask or the laser write head, and  $h$  is the thickness of the resist.

The resolution of the lithography is limited by wavelength, which becomes impractical past the deep UV limit. To achieve a higher precision, we turn to electron beam lithography.

In practice, we achieve a resolution of 0.8 – 1  $\mu\text{m}$  in our DWL system, which offers sufficient precision for patterning transmission lines, resonators, capacitors, and flux-trapping holes.

## 4.5.2 Electron beam lithography

Electron beam (e-beam) lithography is another type of direct-write lithography, which uses a collimated electron beam to expose patterns.

The electron wavelength can be calculated from the de Broglie relation as [99]


$$\lambda_e = h/p_e = \frac{h}{\sqrt{2m_e\varepsilon_{kin}}}, \quad (4.4)$$

where  $p_e = \sqrt{2m_e\varepsilon_{kin}}$  is the electron momentum,  $m_e$  the electron mass and  $\varepsilon_{kin}$  the kinetic energy. Our e-beam system supplies an acceleration voltage of 100 kV, which translates to an electron kinetic energy  $\varepsilon_{kin} = 100$  keV. This makes the electron wave-



length  $\lambda_e \sim 4$  pm. In such a case, the lithography resolution is no longer limited by the electron wavelength, but rather by practical issues such as beam collimator precision, stability, and resist thickness and uniformity.

A popular e-beam resist is poly(methyl methacrylate) (PMMA), which is a polymerized chain of methyl methacrylate (MMA) monomers (Fig. 4.7).




**Figure 4.7:** The chemical formula of PMMA, which is a chain of  $n$  MMA monomers.

E-beam exposure causes scission of the polymer chains, which increases their solubility of the exposed areas in the appropriate developer. E-beam resists often develop in solvents - for PMMA, this would typically be a mixture of methyl isobutyl ketone (MIBK) and isopropanol (IPA), or a mixture of IPA and water. The development is stopped by submersion in pure IPA.

### 4.5.3 Ashing

After development, there is a possibility of resist residues remaining on the exposed film, or around the resist profile edges, which can impact an etching process by masking the areas underneath, or get trapped between the substrate and the metal in a lift-off process.

To combat this effect, we can use resist descumming, or ashing, where a mild oxygen plasma effectively burns off the resist. This can oxidize the material underneath and introduce extra roughness. In an etching process this is of no concern, as the affected metal will be etched away in the next step. However, in a lift-off process, this could affect device performance. It should also be noted that this process thins down the resist mask, which can be a concern for very thin resists.



**Figure 4.8:** Isotropic (a) and anisotropic (b) etching. **a** In an isotropic profile, etching happens in all directions, and the topmost areas which are exposed to the etchant for the longest time have a wider opening. **b** An anisotropic profile is directional and results in a straight edge, defined by the opening in the resist.

## 4.6 Etching


A pattern defined in a resist on top of a metal (Fig. 4.2 **4a**) can be transferred into the metallic layer (Fig. 4.2 **5a**) by wet or dry etching.

### 4.6.1 Wet etching

In wet etching, the wafer is submerged in a liquid solution (typically an acid or a base), that chemically reacts with the metal, producing soluble compounds. The rate of the etching reaction is determined mainly by etchant concentration and temperature. This process can be highly selective towards a target metal, while the surrounded material remains unaffected.

A typical downside of wet etching is that the process is generally isotropic, leading to potentially undesirable etch profiles (see Fig. 4.8).

A popular solution for etching Al is a mixture of phosphoric, nitric and acetic acids. In a process reminiscent of the SC1 method, this etching process relies on a tandem of oxidation and oxide removal. The phosphoric acid ( $\text{H}_3\text{PO}_4$ ) etches the native  $\text{AlOx}$  covering the Al, which is then reoxidized by nitric acid ( $\text{HNO}_3$ ), only to be etched by  $\text{H}_3\text{PO}_4$  as the cycle continues [100]:



### 4.6.2 Dry etching

Dry etching is done under vacuum, and can take the form of reactive ion etching (RIE), or ion milling. Both process types are driven by a voltage gradient, making them highly directional (anisotropic), which leads to straight etch profiles (see Fig. 4.8). On the downside, as the process relies on high kinetic energy processes such as ion bombardment or plasma, it can introduce excessive roughness and defects in the underlying substrate.

Ion milling is a physical etching technique where a heavy inert gas (typically Ar) is accelerated towards the sample, where it sputters material away from the surface. An advantage is that as a physical process, ion milling does not leave behind unwanted chemical reaction products. The disadvantage is that it is not a selective process, and will damage the resist mask during etching, as well as the substrate underneath. In our qubit fabrication process, a gentle ion mill is used to remove the native AlO<sub>x</sub> before connecting the JJ electrodes to other circuitry, avoiding the formation of parasitic JJs between the JJ electrode and the remaining circuitry [101].

In RIE, an AC voltage at radio frequencies is used to ignite a chemically reactive plasma. The voltage both ionizes the reactive gas, and accelerates the positively charged particles towards the sample, where they react on collision. The volatile reaction products are pumped away from the chamber.

The etch rate of RIE can be increased by generating an inductively coupled plasma (ICP) via a coil wound around the etching chamber - however, at the expense of diminished anisotropy.

## 4.7 Resist removal

After etching, the resist mask can be removed. In our standard process, we dissolve the resist in a sequence of three solvents: an N-methylpyrrolidone (NMP)-based remover, acetone, and IPA.

NMP-based removers are very effective at resist removal, particularly when heated. However, as NMP is on the candidate list of substances of very high concern by the European Chemicals Agency due to its toxicity for reproduction [102], it should be avoided where possible, and replaced with other solvents. NMP should always be handled with utmost care and in well ventilated premises.

Acetone works well to dissolve many resists and to clean remover residues from the wafer. However, it does not dry well, leaving behind residues on the wafer surface. IPA is a gentler solvent that removes acetone residues, and leaves behind a clean surface when blow-dried by a flow of nitrogen gas.

In addition to heating, sonication helps in guiding dissolved resist away from the substrate, and keep a fresh supply of solvent circulating within the beaker. Once JJs are present on the wafer, the sonication should be kept to mild conditions to avoid JJ damage.

## 4.8 Josephson junctions


The theoretical concept behind the fabrication of Al/AlO<sub>x</sub>/Al JJs is straightforward - given that Al readily oxidizes upon contact with oxygen, all one needs to do to create an Al/AlO<sub>x</sub>/Al sandwich is to deposit a bottom Al layer, then introduce a source of oxygen to grow the oxide (preferably in a controlled manner in-situ), and then deposit the second electrode on top. However, achieving all this in a single lithography process with the two electrodes connected to different parts of the circuit becomes a geometrically challenging endeavor.

JJs can be fabricated via etching or lift-off, and the geometric difficulties can be resolved through double-angle evaporation using the Dolan bridge technique [103], the bridge-free technique [104], or the cross-type technique. Our process involves the cross-type technique and lift-off.

An example of a cross-type design, with the JJ at the cross where the two electrodes overlap, is shown in Fig. 4.9 **a**. The challenge of this design is to ensure that the electrodes are evaporated one at a time, in a single process without breaking vacuum and performing additional lithography. This can be achieved with double-angle evaporation.

The idea behind double-angle evaporation is that when evaporating either electrode at a tilt with an angle  $\theta$  relative to the metal target, the other electrode which is normal relative to the currently deposited electrode will be shadowed by the resist layer, and vice versa (see Fig. 4.9 **b**). This technique is suitable for electrodes narrower than  $s/\tan(\theta)$ , where  $s$  is the resist thickness.


After the JJs are fabricated, they need to be connected to the qubit capacitor, and in our design also the ground plane. This can be done by an additional lithography process consisting of patterning, evaporation, and lift-off, which is time-consuming but versatile (Fig. 4.11**a**). Alternatively, a carefully designed angular patch can be deposited in the same lithography layer as the JJs, using the patch-integrated cross-type (PICT) technique [101]. Here, the patch is deposited at the same tilt  $\theta$  as the JJ electrodes, with the wafer rotated by a relative angle of  $45^\circ$  (Fig. 4.11**b**).



**Figure 4.9:** Cross-type JJs. **a** The two electrodes (E1 and E2) forming a JJ at their overlap. E1 can be evaporated from whichever end, but E2 should be evaporated from the direction from which it will connect to the patch, as E1 may shadow on the opposite end and form open circuit on that side. **b** Double-angle evaporation. Both electrodes are deposited from two different directions normal to each other, at an angle  $\theta$ . When E1 is evaporated from the north, the resist stack north of E2 shadows the opening for E2. Afterwards, when E2 is evaporated from the east, the resist east of E1 protects it from deposition. The resist mask is a bilayer with an undercut to facilitate lift-off, see Fig. 4.10



**Figure 4.10:** Lift-off with two different resist profiles. **a** A too thin resist without an undercut results in a continuous layer of metal on the sidewall, preventing lift-off. **b** A thicker resist stack with an undercut profile in the bottom resist facilitates liftoff. In our stack we use PMMA as the imaging layer, with MMA underneath. As MMA is more soluble in developer than PMMA, it will develop faster than the exposed PMMA, creating an undercut.



**Figure 4.11:** Two types of patch for connecting the JJ electrodes to their surrounding circuitry. **a** A patch deposited in a separate lithography step without angular evaporation. **b** An angular patch fabricated in the same lithography step as the JJs, using the PICT technique.

## 4.9 Dicing and packaging

After the wafer-scale fabrication has been completed, it is time to separate the wafer into individual chips. The wafer is mounted on a sticky surface (such as a tape) from the backside to keep the diced chips together, then a thin blade spinning at several thousand rotations per minute slices the wafer into chips.

Some materials dice more easily than others, Si being a good example of a material that can be diced at high cutting speed with a thin ( $\sim 30\ \mu\text{m}$  wide) diamond blade. Hard materials such as sapphire do not yield to dicing easily, and may need to be diced significantly slower, sometimes in multiple passes partially cutting deeper and deeper into the substrate to avoid damaging the blade (and by extension the sample) during cutting.

It is crucial to have good contact between the wafer backside and the mounting tape, as small chips may come loose under the relentless flow of blade cooling water. If one has forgiving packaging and loose requirements for cleanness of cut, then an option is to not dice all the way through the wafer, but leave the bottom-most  $\sim 100\ \mu\text{m}$  undiced, which can easily be broken off manually after removing the wafer from the dicing tool.

The dicing process can be quite harsh for the devices, as they will be exposed to debris, cooling water, and mechanical manipulation. Therefore, the front side of the wafer is typically coated with a protective layer of resist. A thicker resist may help, but I have found  $0.5\ \mu\text{m}$  of resist to be sufficient. The surviving diced chips can then be stripped of the protective resist layer as usual.

The chip is then mounted on a sample box, which in our case is an oxygen-free high thermal conductivity (OFHC) copper box, equipped with SMA connectors. The chips are attached to the Cu box by a polyvinyl butyral-based medical glue (BF6), which

is surprisingly resilient to cryogenic temperatures, and dissolves in acetone or ethanol. The electrical connection between the chip's ground plane and the box, and between the transmission line and the connector centre pin, is done through Al wire bonds.

## 4.10 Choices

In planar transmon qubits, we have three opportunities for choosing the optimal material platform - the substrate, the JJ stack, and the superconductor forming the resonators, capacitors, and the supporting circuitry.

From the processing side, we can choose the lithography approach (etching vs. lift-off), the type of etching and the etching chemistry, and the type of deposition.

### 4.10.1 Substrate

From a device performance perspective, we consider the substrate's dielectric constant  $\epsilon$ , its dielectric loss tangent  $\tan(\delta)$ , and the properties of any oxides or other compounds growing on the substrate surface. The substrate's  $\epsilon$  affects the field distribution - a higher  $\epsilon$  helpfully pushes more field into the lossless vacuum. The  $\tan(\delta)$  of the substrate should be as low as possible, such as the  $1 \times 10^{-6} - 1 \times 10^{-8}$  for Si and sapphire [66, 105]. The substrate should preferably not oxidize, or if oxidizing, it should grow a thin, easily removable oxide with a low  $\tan(\delta)$ . A related advantage is the ability to thoroughly clean the substrate surface, such as the SC1, SC2 and HF processes for Si.

Regarding the substrate-metal combination, we consider the lattice matching between the substrate and metal, as well as differences in thermal expansion, the chemistry of the materials in direct contact, and the possibility of inter-diffusion.

Practical concerns may include chemical safety of cleaning procedures, as well as how cumbersome a material is to work with - brittle materials may break during fabrication, while hard materials may resist dicing.

Two popular choices of substrate for qubit fabrication are silicon and sapphire. Both materials have low dielectric loss tangents. Sapphire is crystalline aluminium oxide, whose advantage is that it does not grow amorphous oxide on the surface. Si does grow an oxide, which is easily removed with HF, although it will regrow when exposed to atmosphere. Si processing technology has been honed by decades of research in the semiconductor industry, which facilitates superconducting qubit process development. Si can be cleaned easily, has good mechanical properties, is well matched with Al. Si can even be deep-etched relatively easily with a high aspect ratio, which moves the substrate-

air interface further away from the regions of highest electric field concentrations in coplanar waveguide structures.

We use high resistivity, intrinsic Si for our devices. While Si is a semiconductor at room temperature, the charge carriers are frozen out at millikelvin temperatures, and therefore we can assume Si to act as a dielectric at qubit operating conditions.

### 4.10.2 Superconductor

The chief qualifying property of a superconductor is a reasonably high  $T_c$ . This is not merely so that operating in the superconducting regime is practical, but also to minimize resistive loss due to quasiparticles. The equilibrium quasiparticle distribution depends on the superconducting gap, favouring higher  $T_c$  values as the number of thermally excited quasiparticles gets exponentially suppressed for operating temperatures further away from  $T_c$ .

The type of oxide that grows on the surface of the metal is of tremendous importance - its thickness, potential for self-passivation, oxide type variety (some metals like Nb grow a whole family of oxides with different properties [106]), magnetic properties, and  $\tan(\delta)$ .

Last but not least come the practical concerns - ease, variability, and reproducibility of deposition, price, availability, chemical compatibility, etc. Sensitivity to magnetic flux can also influence device coherence.

The superconductors currently facing the biggest popularity in the field of superconducting qubits are Al, Nb, Ta, and TiN [10, 19–22, 107]

The advantage of Al is that it is straightforward to reproducibly deposit high-quality Al films using a wide variety of deposition techniques and parameters. In addition, Al has a self-limiting oxide, matches well with Si as a substrate, and has a reasonably high  $T_c$  of 1.2 K [31]. As a resource it is relatively cheap and available. A disadvantage is the rich chemistry of Al, which constrains processing options as soon as the metal is present on the wafer, notably the harsh but effective acid-based cleaning procedures or standard photoresist developers.

Ta-based transmon qubits were the first to exhibit  $T_1$  values well above 100  $\mu$ s in a planar geometry [19]. An extraordinary advantage of Ta is its resilience to the buffered oxide etch (BOE) as well as the piranha solution, meaning that the surfaces can be thoroughly cleaned of fabrication residues or other impurities in-between fabrication steps until the point where JJs are deposited, resulting in pristine interfaces. A downside of Ta is that it can grow in two different phases,  $\alpha$ -Ta and  $\beta$ -Ta, with different  $T_c$  values (4.5 K and 0.6 – 1 K, respectively) [108, 109]. Only the  $\alpha$ -Ta phase has been observed to yield good coherence [19].



Nb is likewise resilient to BOE, which enables the stripping of both Nb-oxides and Si-oxides in one step, while leaving Nb intact. The  $T_c$  of Nb is also relatively high at 9.3 K [31]. The disadvantage is that Nb oxide is not self-limiting, and will keep growing in thickness. Instead of a single oxide type, Nb has a whole family of oxides, one of which is even magnetic [110].

Superconducting compounds such as TiN and NbTiN have also yielded high-quality devices, and can be prepared in relatively corrosion-resistant configurations [21, 111]. However, these materials can suffer from low wafer-to-wafer reproducibility in nitrogen content as well as poor uniformity across the wafer, which impedes scalability.

### 4.10.3 Junctions

In principle, it is possible to fabricate superconductor-insulator-superconductor (SIS) JJs out of any oxidizing superconductor, as well as a combination thereof (e.g. the Nb/AlO<sub>x</sub>/Nb trilayer JJ). However, qubits with  $T_1 \geq 100 \mu\text{s}$  are yet to be demonstrated with any other JJ platform than Al/AlO<sub>x</sub>/Al.

Al/AlO<sub>x</sub>/Al JJs can be prepared via etching or lift-off. Lift-off junctions have a long history of yielding highly coherent qubits; however, the process is not well scalable due to the loss of deposition rate uniformity in angular evaporation, as well as the chance of resist residues being trapped underneath the metal. An industrial-compatible etch-based JJ fabrication process yielding  $T_1 > 100 \mu\text{s}$  has been demonstrated [112].

### 4.10.4 Deposition type

The choice of deposition technique depends on the metallic material, the substrate, other materials present on the wafer, and the process itself.

The target metal properties determine which techniques are viable from the start. One such property is the melting temperature of the metal, where for example Al with its melting temperature of 660 °C lends itself to a wide variety of deposition methods and conditions, while it is challenging to evaporate Nb with its melting point of 2477 °C or tantalum at 3020 °C at any reasonable deposition rates, in which case sputtering would be preferable.

The substrate and the presence of sensitive materials (for example resist) may also influence suitability. Depositing a lift-off metal on a resist mask may not be suitable for the high energy sputtering process, where the heat evolution or high energy particle bombardment may reflow the resist, or burn it to a degree where it may become difficult to remove.

The desirability of step coverage is also an influence. Evaporation is more directional than sputtering and is thus more suitable for lift-off, but where good step coverage is desired, sputtering might be the better choice.

### 4.10.5 Etching type

In this section I discuss the advantages and disadvantages of wet and dry etching for materials where both options are available.

Wet etching is convenient, as it can be done in a beaker, although this requires manual handling of corrosive chemicals, which is a safety concern. It can also be highly selective towards the target metal. A disadvantage is the poor anisotropy of the process, which results in curved or slanted straight profiles, as well as in features smaller than the openings in the resist mask (referred to as “overetch”). In a stable process, an overetch can be compensated for by biasing the pattern by the expected overetch depth. Additionally, etch rates can be affected by changes in ambient conditions.

Dry etching requires a more advanced set-up than wet etching but can be highly reproducible and anisotropic. The associated ion or plasma bombardment, together with a heat evolution at the substrate, can damage the substrate and other materials on the wafer, such as the resist mask. A mild burnt “skin” on top of a resist can be ashed off before dissolution, but a badly damaged resist might resist removal, leaving behind significant residues.


## 4.11 Common fabrication challenges

### 4.11.1 Finding the correct exposure parameters

A crucial step for the success of lithography is finding the correct exposure and development parameters.

The difference between the exposed and unexposed resist is not a “yes or no” question from the point of view of dissolution in the developer, but rather a question of the *rate* at which the resist is developed away. For positive resists, the exposed resist will develop significantly faster than the unexposed regions, but given enough time, the developer consumes resist across the entire wafer.

The parameters affecting the development is the resist thickness, time and temperature of the soft-bake, the exposure dose (focus, intensity and transmission for laser lithography, current and beam dwell time in e-beam lithography), the developer concentration and time of development, agitation during development and, as usual, ambient



**Figure 4.12:** Dose evaluation patterns. The darker shade is unexposed resist, lighter areas are where resist is developed away. **a** Positive and negative lines with widths ranging  $0.5 - 1 \mu\text{m}$ . As the resolution limit for our DWL tool is  $\sim 0.8 \mu\text{m}$ , narrower lines should not develop well in the positive pattern on top, and should be developed away in the negative pattern at the bottom. **b** A slightly overdeveloped grid of  $2 \times 2 \mu\text{m}^2$  squares, with the exposed squares larger than unexposed ones. **c** The grid pattern developed at the correct dose.

conditions such as temperature, humidity and pressure.


The correct dose can be found by fixing the resist spinning, baking, and development conditions, and performing a succession of exposures of the same pattern, sweeping the relevant exposure parameters. An example of such a test is shown in Fig. 4.12, where we include a test based on the DWL tool resolution, as well as clearance of our pattern's smallest dimension ( $2 \mu\text{m}$ ).

### 4.11.2 Photolithography on Al

For better or for worse, Al lends itself readily to a wide range of acid-base reactions, including with the prime ingredient of photodevelopers - TMAH. This means that once the TMAH has made its way through the exposed photoresist during development, it will start attacking the Al underneath. This may not necessarily be a problem, and can even be exploited to develop and etch in the same step, provided one takes the resulting overdevelopment of the pattern into account.

However, this is not ideal if one would like to verify the patterning prior to committing to etching into the Al, as by the time the development is through, the pattern will be irrevocably imprinted in the top layers of the metal. The use of dilute “Al-friendly” photodevelopers with low TMAH contents mitigates this problem, but does not eradicate it.

In our process we deposit a thin protective layer of e-beam resist between the Al and



**Figure 4.13:** Process for patterning Al by photolithography with a protective e-beam resist layer. **a** shows the stack before patterning. The e-beam resist protects the Al from corrosion during development (**b**). After the photoresist has been exposed and developed, the patterning can be inspected. In case of successful patterning, the thin e-beam resist can be ashed away in oxygen plasma (**c**). The resulting patterned resist stack (**d**) can then be used as an etching mask as usual.


the photoresist. As e-beam resists develop in solvents rather than bases, the e-beam resist layer will be unaffected by the developer, shielding the Al underneath. The patterned photoresist can then be inspected after development. In case issues are discovered in the patterning at this stage, the resist stack can be dissolved as usual, and the process can be reattempted on the undamaged Al. Once the patterning of the photoresist is satisfactory, the underlying e-beam resist can be ashed away in oxygen plasma. Here it is important that the e-beam resist is significantly thinner than the photoresist, such that a sufficiently thick resist mask is left behind after the ashing, which will affect both resists indiscriminately, thinning down the stack. The process is illustrated in Fig. 4.13. Afterwards, the pattern can be transferred into the Al as usual.

### 4.11.3 JJ resistance targeting

The qubit frequency is determined by the charging energy  $E_C$  of the capacitive elements and the Josephson energy  $E_J$ , which depends on the JJ resistance.  $E_C$  is dominated by the qubit capacitor, which is relatively large, and thus forgiving in terms of lithography variation and reproducibility. The JJ, on the other hand, is a small and sensitive element, which can introduce significant variations to the qubit frequency.

There are two factors that determine the resistance of a JJ - the junction area, and the oxide barrier properties.

The oxide properties can be determined by the granular structure of the bottom electrode and by the oxidation parameters, including the oxidation mode (static/dynamic), pressure, temperature, and time. In our high-vacuum deposition system, we can assume having a good control over these parameters.



**Figure 4.14:** The so-called "resistance curve" for JJs - the resistance as a function of the junction area.

Figure 4.14 shows the measured room-temperature resistance  $R$  of 10 JJs on the same chip, designed with different square areas  $A$ . For  $R$  in the range of 10–20 k $\Omega$ , typical for qubit designs, variations in size impact the obtained  $R$  significantly. For larger junction areas, these variations affect  $R$  to a lower degree. It is possible to modify the oxidation parameters to get higher  $R$  for large junctions [113], although the probability of trapping a TLS in the JJ oxide may increase.

For JJs on Si, the substrate can capacitively shunt the JJ. We observe a sensitivity of this shunting effect to the ambient light conditions in the laboratory, changing the measured resistance. Therefore, in our set-up, we measure JJ resistance in the dark using an automatic probe station.



---

## Microwave measurements

---

In this chapter, I introduce the cryogenic microwave measurement set-up used to characterize both qubits and resonators. I also describe the measurement and analysis procedures for extracting the relevant device characteristics, as well as the procedures for device and set-up calibration.

### 5.1 Cryogenic microwave measurement setup

All of the microwave measurements discussed in this work are performed with the device at cryogenic temperatures  $\sim 10$  mK inside a dilution refrigerator. There are two main reasons that necessitate this rather cumbersome choice of set-up – achieving superconductivity, and suppressing noise due to thermal fluctuations.

The superconducting transition temperature  $T_c$  for Al, our usual superconductor of choice, is 1.2 K [31]. While other popular superconductors for CPWs and capacitors in quantum processors (e.g. Ta and Nb) can achieve higher  $T_c$  than Al, the state-of-the-art highly coherent devices still rely on Al electrodes for their Josephson junctions – in which case, reaching  $\lesssim 1$  K would be sufficient to fulfil the condition for superconductivity.

Suppressing thermal fluctuations puts a more stringent bound on the device temperature. To avoid thermal excitations in the qubit, it is necessary that the qubit transition

energy  $\hbar\omega_{01} \gg k_B T$ . For a frequency of 5 GHz, this corresponds to 240 mK, planting our experiments firmly in the mK-regime. Additionally, even for a higher frequency device,  $T \ll T_c$  is preferred in order to limit the presence of quasiparticles - a cause of resistive loss.

### 5.1.1 Dilution refrigeration

There are currently two leading technologies capable of reaching and sustaining millikelvin temperatures on practical timescales - adiabatic demagnetization refrigeration (ADR) and dilution refrigeration (DR). In our experiments we use the latter.

Dilution refrigeration relies on a mixture of two helium isotopes,  $^3\text{He}$  and  $^4\text{He}$ .  $^4\text{He}$  is a boson, which allows for multiple particles to occupy the same energy state. Below 2.17 K,  $^4\text{He}$  particles condense into the lowest energy state, forming a so-called Bose-Einstein condensate, marking the onset of superfluidity in  $^4\text{He}$  - a condition where a fluid loses all viscosity.  $^3\text{He}$ , however, being a fermion, is only allowed up to two particles (of opposite spin, no less) to occupy the same energy state. As fermions thus cannot form a Bose-Einstein condensate, achieving superfluidity is challenging for  $^3\text{He}$ , and will not occur under normal DR operating conditions.


A mixture of these two  $^3\text{He}/^4\text{He}$  isotopes separates into two phases when cooled below 870 mK: a  $^3\text{He}$ -rich phase, which floats on top of the heavier,  $^3\text{He}$ -poor phase. The equilibrium between the two phases is disrupted by pumping away some of the  $^3\text{He}$ -rich phase, after which a new equilibrium needs to be restored. This equilibrium restoration process is endothermic, which results in cooling.

The cooldown from room temperature to 10 mK begins with pumping out the gas contents of the cryostat. Then, a pulse tube with He as an exchange gas brings the DR temperature to  $\sim 4\text{K}$ . Afterwards, the  $^3\text{He}/^4\text{He}$  gas is condensed inside the gas-handling system, resulting in a rather rapid further cooling, which the gas-handling system aids further by facilitating evaporative cooling via gas circulation. Below 870 mK, the phase separation of the  $^3\text{He}/^4\text{He}$  occurs, and the dilution refrigeration process itself is responsible for further cooling down to the cryostat's base temperature.

### 5.1.2 Microwave measurement set-up

Inside the DR, there are multiple stages with descending temperatures, shielded from one another in a nesting doll-style architecture. The bottom-most (and simultaneously inner-most) stage to which the mixing chamber of the dilution unit is attached reaches a temperature of 10 mK, and it is to this stage that we attach a Cu plate onto which






**Figure 5.1:** The first stages of chip packaging and mounting to the cryostat. **a** The chip is glued to a Cu box on the backside. The transmission line is connected to the SMA connectors via wirebonds, and wirebonds along the chip periphery ground the chip through the box. **b** The sample box is closed with a light-tight lid and mounted on a Cu plate attached to the mixing chamber stage of the cryostat. **c** The cold plate with samples is enclosed in a Cu can. This can will be enclosed in a Cryoperm shield afterwards.

we mount our devices. The Cu plate with the device is protected by a light-tight Cu shield, which is itself enclosed in a Cryoperm magnetic shield. The magnetic shield attached to the mixing chamber plate is also coated with SiC mixed with Stycast to absorb radiation. The temperature stages, as well as the various components of the input-output measurement line used to probe the devices, are depicted in Fig 5.2.

The input line is attenuated at every temperature stage to mitigate the associated thermal noise. While the total attenuation is significant, this is not a problem, as we need to deliver very weak powers to the device in order to induce single-photon transitions. A low-pass filter and an Eccosorb-based filter are also present to filter out high-frequency noise.

Any power dissipation affecting the often extremely weak device output signal becomes undesirable, therefore the output line cabling is superconducting until it reaches the cryogenic low-noise amplifier on the 4K-stage. The amplified signal with a more favourable signal-to-noise ratio (SNR) can then be brought up to room temperature, where it is further amplified and measured. To minimize the incidence of unattenuated thermal noise from the upper stages, signal reflections, or noise leakage from the amplifiers, the output signal is passed through microwave isolators and band-pass filters. In addition, the poor heat conductivity of the superconducting wiring helps isolate the lower-temperature stages from thermal noise.

For the sake of device measurement throughput per cooldown, microwave switches that break out to connect multiple samples are installed on both the input and output lines.



**Figure 5.2:** Cryogenic microwave measurement setup for a transmission measurement. The device-under-test (DUT) is attached to the mixing chamber plate at  $\sim 10$  mK.

### 5.1.3 Measurement instrumentation

For continuous wave spectroscopy, such as in resonator characterization, we use a vector network analyzer (VNA). A VNA can send continuous signals at a specified frequency and power, and measure the complex response of the device.

For qubit readout and control, we need to generate signal pulses instead of continuous waves. Direct generation of pulses with short and precise envelopes on the order of ns with a high-frequency content of a few GHz is an ambitious and expensive affair - an alternative is to generate the high frequency content and the envelope separately, and then combine the two with the help of a signal mixer. The envelope carrying a low-frequency signal at an intermediate frequency  $\omega_{IF}$  can be generated using an arbitrary waveform generator (AWG), which can then be mixed with a local oscillator carrier frequency  $\omega_{LO}$  produced by a signal generator (SG), see Fig. 5.3a. Such a frequency upconversion can be described in the Fourier space as:


$$S = \cos(\omega_{LO}t) \cos(\omega_{IF}t) = \frac{1}{2} \cos[(\omega_{LO} + \omega_{IF})t] + \frac{1}{2} \cos[(\omega_{LO} - \omega_{IF})t]. \quad (5.1)$$

The result of this frequency mixing are signals at two frequencies, or two “sidebands”, which coincide with the sum and the difference of the two original frequencies. The unwelcome sideband can be suppressed with the use of a slightly more complex set-up involving the separate generation of an in-phase (I) and quadrature (Q) signal, and IQ mixers (Fig. 5.3b). By generating the I and Q components of the envelope separately in the AWG, and by inserting a hybrid coupler that shifts the LO phase by  $\pi/2$  we get the following spectrum:

$$\begin{aligned} S &= \cos(\omega_{LO}t) \cos(\omega_{IF,I}t) + \sin(\omega_{LO}t) \cos(\omega_{IF,Q}t) \\ &= \frac{1}{2} \{ \cos[(\omega_{LO} + \omega_{IF,I})t] + \cos[(\omega_{LO} - \omega_{IF,I})t] \\ &\quad + \sin[(\omega_{LO} + \omega_{IF,Q})t] + \sin[(\omega_{LO} - \omega_{IF,Q})t] \}. \end{aligned} \quad (5.2)$$

If the I and Q components of the IF signal generated in the AWG are in quadrature,  $\cos(\omega_{IF,Q}t) = \sin(\omega_{IF,I}t)$ , resulting in a single sideband:

$$\begin{aligned} S &= \cos(\omega_{LO}t) \cos(\omega_{IF,I}t) + \sin(\omega_{LO}t) \sin(\omega_{IF,I}t) \\ &= \frac{1}{2} \cos[(\omega_{LO} - \omega_{IF,I})t]. \end{aligned} \quad (5.3)$$



**Figure 5.3:** High-frequency short pulse generation via frequency upconversion. **a** A mixer upconverts a low-frequency pulse with a precise envelope and a continuous high-frequency signal into a short high-frequency pulse. **b** The desired sideband is filtered out with the use of an IQ mixer consisting of a 90 deg hybrid coupler, two mixers driven in quadrature, and a power combiner.

On the output of the cryostat, the device output signal is downconverted and digitized by a digitizer.

An alternative to this set-up is the use of an FPGA-based multi-frequency lock-in amplifier platform, which synthesises the signal directly on the FPGA with the help of digital mixing. The device output signal is likewise digitized on the FPGA board. We have used both these set-ups to measure qubit coherence interchangeably with indistinguishable results. The lock-in amplifier platform is also used in our intermodulation spectroscopy measurements.

## 5.2 Resonator measurements

A practical definition of the resonator quality factor  $Q$  quantifies its value as a ratio of the resonance frequency  $f_r$  and the resonance linewidth  $\Delta f$  (i.e. the full width at half maximum) [114]:

$$Q = \frac{f_r}{\Delta f}. \quad (5.4)$$

This information can be extracted from a simple spectroscopy where we measure the magnitude of the transmitted or reflected signal as a function of frequency. However, a more precise result and a great deal of additional information can be obtained by measuring the complex amplitude of the device response with a VNA instead of just the magnitude.

In a two-port measurement, the relation between an incident wave  $(a_1, a_2)$  and the reflected wave  $(b_1, b_2)$ , where the indices 1, 2 stand for the VNA output/input ports, can be expressed by the following scattering matrix:

$$\begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = \begin{pmatrix} S_{11} & S_{12} \\ S_{21} & S_{22} \end{pmatrix} \begin{pmatrix} a_1 \\ a_2 \end{pmatrix}, \quad (5.5)$$

where the scattering parameters  $S_{11}$  and  $S_{22}$  quantify the reflected signal output and measured by the same port, while  $S_{12}$  and  $S_{21}$  quantify the transmitted signal ( $S_{12}$  for signal output on port 2 and measured on port 1, and  $S_{21}$  for signal output on port 1 and measured on port 2).


In a resonator coupled to an input/output transmission line in a notch configuration (see Fig. 2.3), with the VNA outputting the probe signal from port 1 and receiving the transmitted signal at port 2, the relevant scattering parameter becomes  $S_{21}$ . In such a configuration,  $S_{21}$  shows a dip in the frequency spectrum that coincides with the resonance frequency, as the resonator absorbs near-resonant signals and rejects frequencies outside of the resonance linewidth (see Fig. 5.4 **a** and **d**).

### 5.2.1 Resonance circle fit analysis

A resonator integrated in a circuit will lose energy to its internal loss mechanisms (internal quality factor  $Q_i$ ), as well as to the coupling to the surrounding microwave circuitry such as the input/output transmission line (real part of the coupling quality factor  $Q_c$ , whose complex nature I will address soon). The total (or loaded) quality factor  $Q_l$  becomes:

$$\frac{1}{Q_l} = \frac{1}{Q_i} + \frac{1}{\Re(Q_c)}. \quad (5.6)$$

The loaded resonator quality factor  $Q_l$  can be calculated from the resonance linewidth and frequency as described above; to separate it into the two constituents requires information about the amplitude and phase across the resonance, see Fig. 5.4. The transmission of an ideal resonator plots a circle with a diameter  $d = Q_l/|Q_c|$  on the complex plane (Fig. 5.4 **c** and **f**). In case of an impedance mismatch  $\phi$ , this circle



**Figure 5.4:** Circle fit of resonator data collected at two different drive powers: the single-photon level in **a–c** and at  $\sim 10^6$  photons in **d–f**. At high probe power the SNR is high, while at the single-photon level the noisy data increases the fit error.

becomes skewed, which is handled by introducing a complex  $|Q_c| \cdot e^{-i\phi}$  [115].

For a resonator in a notch-type configuration,  $S_{21}$  as a function of the probe frequency  $f$  has the following form [116]:

$$S_{21}(f) = ae^{i\alpha} e^{-2\pi i f \tau} \left[ 1 - \frac{(Q_l/|Q_c|)e^{i\phi}}{1 + 2iQ_l(f/f_r - 1)} \right], \quad (5.7)$$

where the terms in black describe the transmission of an ideal resonator, and the terms in red adjust for the effects on an imperfect microwave environment: a background transmission amplitude  $a$ , a phase shift  $\alpha$ , and an electronic delay  $\tau$ . In our analysis, we use the circle fit routine adapted directly from [116].

The circle fit is at its most accurate in extracting both  $Q_i$  and  $Q_c$  when the resonator is critically coupled, i.e.  $Q_i \approx Q_c$ , and the accuracy falls when  $Q_i$  and  $Q_c$  differ by orders of magnitude [90].

## 5.2.2 Circulating photon number

The quality of superconducting resonators is power dependent. For this reason among others, we want to calculate the power circulating inside the resonator for a given power

$P_{in}$  delivered to the device input. We start by assuming that any signal delivered to the device that is not transmitted further along to the output line ( $P_{trans}$ ) or reflected back to the transmission line ( $P_{refl}$ ) is absorbed by the resonator:

$$P_{abs} = P_{in} - P_{trans} - P_{refl}. \quad (5.8)$$

On resonance,  $S_{21} = S_{21}^{min} = Q_c / (Q_c + Q_i)$ , and we can express  $P_{trans}$  and  $P_{refl}$  in the following way, adjusting for a potential mismatch between the resonator impedance  $Z_r$  and the transmission line impedance  $Z_0$  [111, 117]:

$$P_{trans} = \frac{Z_0}{Z_r} P_{in} |S_{21}|^2 = \frac{Z_0}{Z_r} P_{in} \left| \frac{Q_c}{Q_c + Q_i} \right|^2, \quad (5.9)$$

$$P_{refl} = \frac{Z_0}{Z_r} P_{in} |S_{11}|^2 = \frac{Z_0}{Z_r} P_{in} |S_{21} - 1|^2 = \frac{Z_0}{Z_r} P_{in} \left| \frac{-Q_i}{Q_c + Q_i} \right|^2. \quad (5.10)$$

Inserting equations 5.10 and 5.9 into 5.8, we obtain the power absorbed by the resonator:

$$P_{abs} = \frac{Z_0}{Z_r} \frac{2Q_l^2}{Q_i Q_c} P_{in}. \quad (5.11)$$

The power circulating in the resonator then becomes:

$$P_{circ} = P_{abs} Q_i = \frac{Z_0}{Z_r} \frac{2Q_l^2}{Q_c} P_{in}. \quad (5.12)$$

It is customary to express the circulating power in the equivalent average number of photons  $\langle n \rangle$ , which we obtain through the average photon energy  $\langle E \rangle = \langle n \rangle \hbar \omega_r$  and its relation to power  $\langle E \rangle = P / \omega_r$ :

$$\langle n \rangle = \frac{\langle E_{circ} \rangle}{\hbar \omega_r} = \frac{P_{circ} / \omega_r}{\hbar \omega_r} = 2 \frac{Z_0}{Z_r} \frac{Q_l^2}{Q_c} \frac{P_{in}}{\hbar \omega_r^2}. \quad (5.13)$$

I will discuss the calibration of the power level at the input of the device in the *Input line calibration* section.

## 5.3 Qubit measurements

The main purpose of the qubit measurements outlined in this section is to extract relevant device characteristics (frequencies, couplings) and the qubit coherence times. I will start by introducing the dispersive qubit readout scheme, then I will describe the qubit calibration, the coherence measurement procedures, and finally the calculation of the qubit's Purcell decay rate.

### 5.3.1 Dispersive readout of a qubit state

Reading out the state of a qubit without altering its state is complicated by the impressionable nature of quantum states, which are easily perturbed by measurement efforts.

One example of a quantum non-demolition measurement protocol that achieves this goal is the dispersive readout [118], where the state of the qubit is inferred from the frequency shift of a readout resonator coupled to the qubit in the dispersive regime. This regime occurs when the frequency detuning between the qubit and the resonator far exceeds their mutual coupling strength ( $\Delta \gg g$ ). Under this condition, the field inside the resonator and the qubit cannot induce any real mutual transitions, but the coupling will renormalize the energies of the system such that the frequency of each component is affected by the state of the other.

The Hamiltonian operator describing the dynamics between the qubit coupled to a resonator, derived from the Jaynes-Cummings model while taking the rotating wave approximation, takes the following form [92]:

$$\hat{\mathcal{H}}_{JC} = \hbar\omega_r \left( \hat{a}^\dagger \hat{a} + \frac{1}{2} \right) + \frac{\hbar\omega_q}{2} \hat{\sigma}_z + \hbar g (\hat{\sigma}_+ \hat{a} + \hat{\sigma}_- \hat{a}^\dagger). \quad (5.14)$$

The first term of the Jaynes-Cummings Hamiltonian is that of a harmonic oscillator, with  $\hat{a}^\dagger$  and  $\hat{a}$  as the creation and annihilation operators for a photon in the readout resonator with a frequency  $\omega_r$ . Similarly, the second term treats the excitation and relaxation of the qubit with a frequency  $\omega_q$ , where  $\hat{\sigma}_z = |1\rangle\langle 1| - |0\rangle\langle 0|$ . The last term describes the mutual interaction between the resonator and the qubit, resulting in a coherent exchange of a single quantum between the two at dynamics dictated by  $g$ . In this term,  $\hat{\sigma}_+ = |1\rangle\langle 0|$  denotes the qubit excitation and  $\hat{\sigma}_- = |0\rangle\langle 1|$  its relaxation.

In the dispersive regime where  $\Delta \gg g$ , the Jaynes-Cummings Hamiltonian can be approximated to the following form using second order perturbation theory:



$$\hat{\mathcal{H}}_{disp} \approx \hbar \left( \omega_r + \frac{g^2}{\Delta} \hat{\sigma}_z \right) \left( \hat{a}^\dagger \hat{a} + \frac{1}{2} \right) + \frac{\hbar}{2} \left( \omega_q + \frac{g^2}{\Delta} \right) \hat{\sigma}_z. \quad (5.15)$$

In this dispersive Hamiltonian, the frequencies of the two components are modified by the dispersive shift  $\chi = g^2/\Delta$ . It should be noted that the dispersive regime is only valid when the system is driven by a power below a certain critical photon number,  $n_c = \Delta^2/(4g^2)$ . Above  $n_c$ , the dispersive condition no longer holds, the qubit becomes saturated and we observe the bare frequency of the resonator, unaffected by the presence of the qubit.

The equations above hold for a true two-level system, or for a highly anharmonic spectrum such as the Cooper-pair box. For the transmon qubit, which sacrifices high anharmonicity  $\alpha$  to suppress charge noise [8], the attainable higher levels affect the dispersive shift of the resonator frequency when the qubit is brought from  $|0\rangle$  to  $|1\rangle$ , which becomes:

$$\chi_{01} = \frac{g_{01}^2}{\Delta} \frac{\alpha}{\Delta + \alpha}. \quad (5.16)$$


### 5.3.2 Qubit calibration

The qubit calibration sequence begins by locating the bare and dressed frequencies of the readout resonator, and their corresponding drive power thresholds. Then, we find the frequency of the qubit, and through a Rabi oscillation measurement, we can find the parameters of the  $\pi$  pulse that brings the qubit from  $|0\rangle$  to  $|1\rangle$ . Then, with a targeted use of  $\pi/2$  pulses, we can hone in on the qubit frequency in a Ramsey sequence.

#### Readout resonator spectroscopy

The first step in calibrating a qubit measurement is locating the “bare” and “dressed” resonator frequencies and the corresponding powers at which they occur. We do this by a resonator spectroscopy at different drive power levels. At high powers well above  $n_c$ , we find the resonator in its bare state with a frequency unshifted by the presence of the qubit, see Fig. 5.5. Below  $n_c$ , we find the resonator at a frequency “dressed” by the presence of the qubit, and thus shifted by  $\chi$  in the dispersive regime.

When picking a readout power, we prefer to pick a relatively high power while still staying firmly in the dispersive regime. Under the dispersive condition, driving the resonator with multiple photons does not induce unwanted transitions in the qubit, and stronger signal powers give a higher SNR, enabling faster readout with lower averaging.



**Figure 5.5:** Readout resonator spectroscopy. **a** Resonator spectroscopy at different power levels show the dressed resonator frequency in the dispersive regime at  $\langle n \rangle \ll n_c$ , and the bare resonator frequency at high drive powers. **b** The bare and dressed frequencies, shown at the two representative power levels indicated by the dashed line in **a**.

Additionally, stronger drives yield higher resonator  $Q_i$  values due to the power-dependent TLS loss.

### Qubit spectroscopy

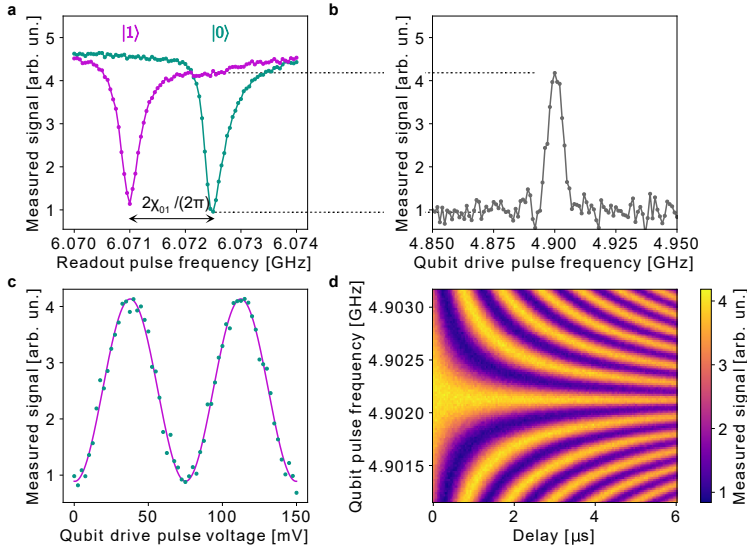
Now that we have established the necessary parameters of the readout pulse (power and frequency), it is time to start calibrating the qubit excitation pulse parameters, starting with its frequency. We do this through the so-called two-tone spectroscopy, where we send two pulses in immediate succession. The first of these two pulses is intended to excite the qubit, while the second pulse checks the outcome by probing the resonator.

In the dispersive regime, a successful qubit excitation shifts the readout resonator frequency by  $2\chi_{01}$ , as I show in Fig. 5.6a. Therefore, if we repeatedly send a pulse to the resonator at the dressed resonance frequency  $f_{dressed}$  which we established in the previous measurement with the qubit in  $|0\rangle$ , we read out the transmission amplitude at the resonance dip. Sweeping the frequency of the qubit pulse we send preceding the readout pulse, if this pulse fails to excite the qubit, we continue to read out the amplitude of the resonance dip. If the qubit pulse excites the qubit, the resonance dip moves away, and we sample the background at a higher amplitude. This results in a peak in the spectrum coinciding with the qubit frequency, as illustrated in Fig. 5.6b.

## Rabi oscillations

Having found the qubit frequency, it is time to find out what pulse parameters induce a rotation on the Bloch sphere by a single  $\pi$  from  $|0\rangle$  to  $|1\rangle$ , such that we can later produce arbitrary rotations in this basis.

We do this through measuring the Rabi oscillations of the qubit. In this step, we fix one of the two remaining free pulse parameters governing the pulse power - amplitude or length - and sweep the other. As gate speed is of essence in quantum algorithms, we choose to fix the pulse length and sweep the amplitude. The result of such a measurement is shown in Fig. 5.6c, where increasing the pulse amplitude from 0 to  $\sim 37$  mV brings the qubit all the way from  $|0\rangle$  to  $|1\rangle$ , and overshooting this amplitude sends the vector further along the surface of the sphere, until it makes its way back to  $|0\rangle$  at  $\sim 74$  mV. Increasing the amplitude further repeats this rotation. Similarly, to bring the qubit to an equal superposition between  $|0\rangle$  and  $|1\rangle$  at the equator of the Bloch sphere in a  $\pi/2$  rotation, one need only cut the amplitude of the newly established  $\pi$ -pulse in half.



**Figure 5.6:** Qubit calibration. **a** Resonator spectroscopy with the qubit in  $|0\rangle$  and  $|1\rangle$ . **b** Qubit spectroscopy. **c** Rabi oscillations. **d** Ramsey fringes.

## Ramsey fringes

The precision of the qubit frequency, which has been assigned from a qubit spectroscopy, can be further improved through the Ramsey sequence. In the rotating wave approximation behind the dispersive Hamiltonian, we assume that the qubit vector is precessing along the Bloch sphere with the qubit transition frequency  $f_{01}$ , and we move to the rotating frame so that we can treat the qubit vector as stationary [46]. However, in the case of a near-resonant drive detuned from  $f_{01}$ , the qubit vector left to freely evolve will appear to precess around the Bloch sphere with a frequency equal to the drive detuning.

In the Ramsey sequence, the qubit vector is brought to the equator of the Bloch sphere, where this effect is most pronounced, by a  $\pi/2$ -pulse. Then, the qubit is left to freely evolve for a time  $\Delta t$ , before we send another  $\pi/2$ -pulse and read out the state of the qubit immediately (for a pulse schematic, see Fig. 5.7b). Varying the time  $\Delta t$ ,  $\Delta t = 0$  the two  $\pi/2$  pulses simply add up to a  $\pi$ . Allowing the vector to precess around the equator to the opposite side of the sphere makes the second  $\pi/2$  pulse bring the vector back to  $|0\rangle$  (disregarding decoherence for the moment). Sweeping  $\Delta t$  results in oscillations between these two extremes with the frequency of the pulse detuning.

Therefore, a two-dimensional parameter sweep varying the delay time between the two pulses on one axis, and the qubit pulse frequency on the other – see Fig. 5.6d – results in the so-called Ramsey fringes, and we find the true  $f_{01}$  when the oscillations disappear.

### 5.3.3 Anharmonicity

If we are interested in the qubit anharmonicity  $\alpha = 2\pi(f_{12} - f_{01})$ , we can perform a spectroscopy starting with the qubit in  $|1\rangle$ . This sequence uses three pulses - the first excites the qubit, the second searches for  $f_{12}$ , and the third pulse reads out the result through the resonator.

Sweeping the frequency of the second pulse, when its frequency coincides with  $f_{12}$ , the readout resonator frequency shifts once again. This we observe through the change in the readout signal amplitude, just as we did in the two-tone spectroscopy in search for  $f_{01}$ .

### 5.3.4 Coherence measurements

In this section I will show how to measure the coherence times defined in Chapter 2.

### Energy relaxation time

The energy relaxation time  $T_1$  can be measured relatively easily by bringing the qubit to  $|1\rangle$  with a  $\pi$ -pulse and waiting some time before measuring the qubit state. Increasing this readout delay  $\Delta t$ , we observe an exponential decay to  $|0\rangle$  with a time constant of  $T_1$ , see Fig. 5.7. To ensure that we are starting each measurement of the decay curve in  $|0\rangle$ , we can either wait a multiple of the expected  $T_1$  before measuring the next point, or for the more ambitious, perform an active reset of the qubit between each measurement [29].

### Ramsey decay time

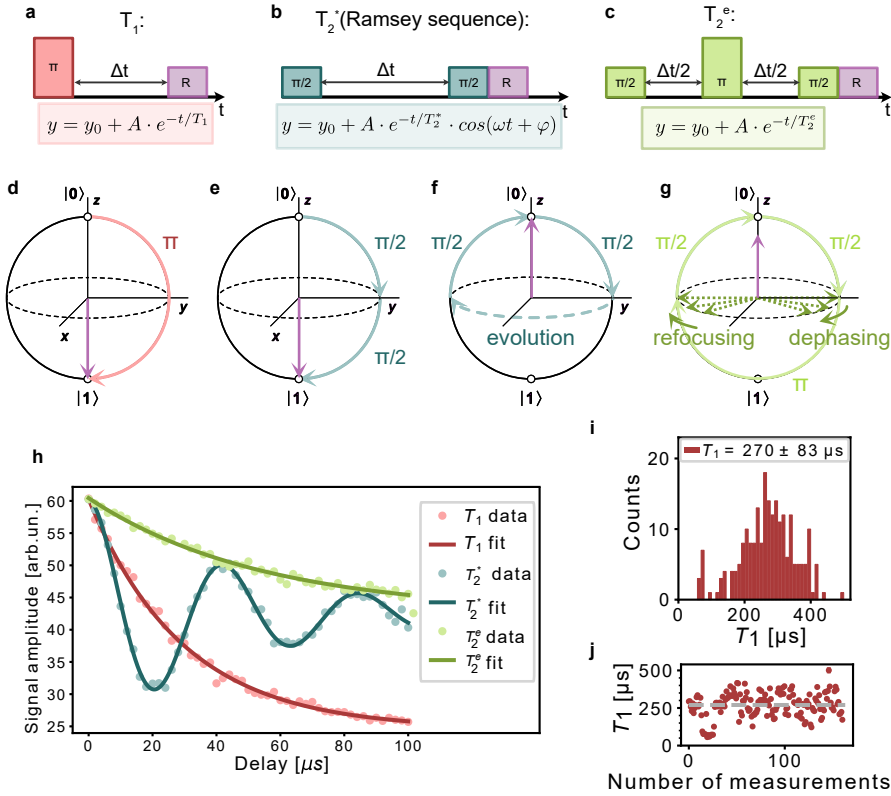
There are multiple strategies for obtaining the decoherence time  $T_2$ , the simplest of which is the Ramsey sequence described in the *Ramsey fringes* section. When the sequence is applied perfectly on resonance, the qubit vector brought to the equator of the Bloch sphere experiences both energy relaxation and dephasing, and decays into the mixed state. The time constant obtained by increasing  $\Delta t$  between the two  $\pi/2$  pulses is commonly denoted as  $T_2^*$ .

In case of a detuning between the pulse frequency and the true qubit  $f_{01}$ , we get a decaying oscillation with the frequency of the drive detuning and the decay time constant of  $T_2^*$ . With a near-resonant drive, this oscillation on top of the exponential decay can be barely perceptible and can decrease the accuracy of the exponential fit. This issue can be exacerbated by TLS-induced qubit frequency fluctuations when tracking  $T_2^*$  over time. Therefore, I prefer to intentionally detune the  $\pi/2$ -pulse by a couple kHz to induce an oscillation I can reliably fit, such as is the case in Fig. 5.7h.

### Spin echo decay time

Another popular method for quantifying qubit coherence is the spin echo sequence, where a so-called refocusing  $\pi$  pulse is added between the two  $\pi/2$  pulses. The idea is that if a systematic environmental noise source is dephasing the qubit in one direction, then flipping the qubit to the opposite pole of the Bloch sphere will counteract the effect of this dephasing. If the qubit spends an equal amount of time  $\Delta t/2$  on both sides of the sphere, the effect of constant systematic noise could cancel out.

This method can be expanded with an arbitrary number of  $\pi$ -pulses in the Carr-Purcell-Meiboom-Gill (CPMG) sequence, where in the event of facing systematic noise, an increasing number of  $\pi$ -pulses brings  $T_2^{\text{CPMG}}$  closer to the theoretical maximum of  $2T_1$  [46]. Such an experiment reveals how much of the qubit dephasing is caused by



**Figure 5.7:** Qubit coherence measurements. **a–c** The pulse sequences for measuring  $T_1$ ,  $T_2^*$  and  $T_2^e$ , respectively, where **R** is the readout pulse. In the accompanying fitting equations for the time constant extraction,  $y_0$  is the  $y$ -axis offset and  $A$  the amplitude.  $\omega$  is the Ramsey oscillation frequency. **d–g** The pulse sequences in **a–c** represented on the Bloch sphere, with the lilac arrow showing the qubit state at the end of the sequence. No decoherence is assumed in the cases of **d–f**. **e–f** The two extremes of applying two slightly detuned  $\pi/2$ -pulses in the Ramsey sequence, bringing the qubit to  $|1\rangle$  or  $|0\rangle$ . **g** The spin echo sequence, where the  $\pi$  pulse refocuses the dephasing incurred on the opposite side of the sphere. Since we are considering decoherence now, we find the qubit in the mixed state inside the Bloch sphere at the end of the sequence, as opposed to the surface. **h** Data showing an example of  $T_1$ ,  $T_2^*$  and  $T_2^e$  decays, fitted to the equations in **a–c**. For visual comparison purposes,  $T_2^e$  has been flipped by another  $\pi$  pulse immediately prior to readout, such that the maximum at  $\Delta t = 0$  corresponds to  $|1\rangle$ . In the  $T_1$  procedure the qubit decays into  $|0\rangle$ , while the two  $T_2$  sequences see the qubit decay into the mixed state. **i** shows the histogram of the  $T_1$  values of a single qubit collected over 48 h. The individual  $T_1$  values over time are displayed in **j**.

systematic vs. stochastic sources.

### Coherence fluctuations

It is typical to observe fluctuations in qubit coherence over time, due to TLSs as well as other other parameter fluctuations [12]. We show this on a repeated measurement of over 48 h in Fig. 5.7i–j. Therefore, when quoting device coherence values, it is beneficial to collect a larger statistical sample and quote the obtained mean values.

#### 5.3.5 Purcell decay rate

The Purcell decay time  $T_p$  gives the timescale on which the qubit decays to the transmission line through the resonator. It is the inverse of the Purcell decay rate  $\gamma = \kappa g^2 / \Delta^2$ , where the readout resonator decay rate is given by  $\kappa = 2\pi f_r / Q_l$ . The coupling  $g$  between the qubit and the resonator can be calculated from the  $\chi$ -shift between the bare and dressed resonator frequencies and the qubit-resonator detuning  $\Delta$  as  $g = \sqrt{\chi \Delta}$ .

As we are interested in the Purcell decay rate at the qubit operating conditions, we extract the relevant resonator-related parameters  $\Delta$  and  $\kappa$  in the dressed state at the drive power corresponding to the qubit readout power.


Therefore, we can explicitly calculate  $T_p$  from our measured value as:

$$T_p = \kappa^{-1} \frac{\Delta^2}{g^2} = \kappa^{-1} \frac{\chi}{\Delta} = \frac{Q_l}{2\pi f_r} \frac{2\pi(f_r - f_{01})}{2\pi(f_{dressed} - f_{bare})}. \quad (5.17)$$

## 5.4 Input line attenuation calibration

To convert the measurement instrument output power to the power delivered to the sample, we need to know the input line attenuation at the measurement conditions. Measuring the line attenuation at room temperature before connecting the sample is only marginally helpful, as this changes once the system is cold, at which stage we no longer have the possibility to directly measure the input line resistance. Therefore, our best bet is to infer the input power from measuring the device properties.

Qubits are the ideal candidates for this, as due to the anharmonicity of their energy spectrum, we can tell how much power is needed to populate the qubit with a single photon. We have discussed how in the dispersive regime, the qubit and the resonator shift each other's frequencies by  $2\chi_{01}$  for every additional photon in the respective element. For a qubit shifted by a driven resonator, this frequency shift is called the ac-Stark shift [49]. We can use this shift to calculate the population of the resonator, and subsequently



**Figure 5.8:** Input line calibration. **a** AC-Stark shift of a qubit as a function of the VNA output power driving the readout resonator. **b** Resonator population as a function of the VNA drive power.

by considering coupling between the resonator and the input/output transmission line, we can extract the power delivered at the device input, which gives us the precise input line attenuation.

In this procedure, we measure the the qubit frequency as a function of resonator drive power. This can be done by pulsed qubit spectroscopy while continuously driving (and reading out) the resonator at different powers using a VNA. The linear fit of this relation gives the unshifted frequency  $f_{01}^0$  of a qubit in the absence of a drive, see Fig. 5.8a.

To convert this into the photon number population of the resonator, we divide the ac-Stark shift at a given power by the dispersive shift  $2\chi_{01}$ :  $\langle n \rangle = (f_{01} - f_{01}^0)/(2\chi_{01})$ . The dispersive shift is extracted by resonator spectroscopy with the qubit in  $|0\rangle$  and  $|1\rangle$ , see Fig. 5.6a. A linear fit of  $\langle n \rangle$  as a function of VNA output power gives the resonator population at an arbitrary drive power, see Fig. 5.8b.

Extracting the  $Q$  values from a circle fit of the dressed resonance, we can calculate the power at the device input  $P$  at an arbitrary  $\langle n \rangle$  from the modified Eq. 5.13:

$$P = \frac{Q_c \hbar \omega_r^2}{2Q_l^2 \langle n \rangle} \quad (5.18)$$

The difference between the VNA drive power and the  $P$  at the device input extracted above then gives the input line attenuation at the resonator frequency.



---

## Mitigating loss - film quality and the substrate-metal interface

---

In the following chapters, I will summarize the results of our investigations into the loss mechanisms limiting the coherence of our devices, starting with losses incurred at the substrate-metal (SM) interface, in conjunction with losses associated with the film quality.

Our goal is to understand the origin of the loss currently limiting device performance, and to find ways to mitigate this loss through optimizing the device fabrication process or device design.

The impact of various changes in the fabrication process, design layout and choice of materials on the device quality can be inferred through coherence measurements - however, the procedure of full qubit fabrication, chip packaging, cooldown to millikelvin temperatures and appropriately averaged microwave characterization is time-consuming and resource-heavy. Therefore, together with studying the impact of the introduced changes on device coherence, we are also looking for indicators of device performance which could be collected at an early stage of the fabrication process, ideally at room temperature. These could be materials characteristics such as surface roughness, elemental profile, or the presence of various defects or contaminants – potentially in conjunction with reliable modeling and simulation methods. Mapping the aforementioned parameters to coherence data not only broadens our understanding of the loss mechanisms

at play, but also paves the way to finding early-stage device performance indicators for quantum processor fabrication.

Dielectric loss due to parasitic TLSs is typically a major limiting factor in superconducting qubits [12, 59, 119], therefore our investigations are focused around this loss channel. CPW resonators make convenient proxies for studying TLS loss in qubits thanks to the power-dependent resonator  $Q$  caused by TLS loss, and the fact that qubits and resonators share many of the same loss channels due to their shared fabrication procedures and material platforms.

In TLS loss studies, the focus is typically on the interfaces between the different materials [51, 120]. We have three primary materials in our CPW structures – the substrate, the metal, and vacuum. We assume the cryogenic vacuum to be lossless, and there to be zero field inside the superconducting metal. We also choose a substrate with a very low bulk dielectric loss tangent (such as Si or sapphire with loss tangents in the range of  $10^{-8} - 10^{-7}$  [51, 66]).

Assuming zero field inside a uniform superconductor, it is tempting to also consider the film quality (in terms of granular structure, defects or contaminants) inconsequential to the device loss. However, the implications of different film structures for the morphology of the lossy interfaces play a role in dielectric loss, which we explore in papers **A** and **B**.

Additionally, in paper **D** we see that the quality of 3D cavities improves with the purity of the metal – the cavities fabricated from a 6081 Al alloy (96%-98% Al with traces of mostly Mg, Si and Mn) show a significantly lower  $Q_i$  than cavities fabricated from 99.99% (4N) pure Al, which are in turn outperformed by Al with a 99.999% (5N) purity.

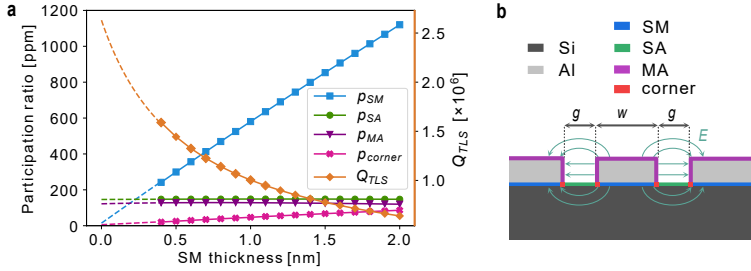
## 6.1 Participation ratio simulations

As discussed in section 2.2.2, energy participation ratio simulations are a useful tool for shedding light on the relative contributions of different materials and their interfaces to the total loss of a device; however, the accuracy of these simulations relies on assuming correct values for the effective thicknesses of the contributing lossy regions and their dielectric loss tangents. In paper **B**, we show that fabricating CPW resonators of identical design, while modifying one step of the fabrication procedure affecting a single interface at a time, has a noticeable impact on the total TLS loss  $F\delta_{\text{TLS}}^0$  in ways that are nontrivial to predict.

In our participation ratio investigations, we simulate a CPW cross-section with our

| Interface | $\tan \delta$ | $\epsilon_r$ | Thickness         |
|-----------|---------------|--------------|-------------------|
| Al        | -             | -            | 150 nm            |
| Air       | 0             | 1.0          | 2 mm              |
| Si        | $10^{-7}$     | 11.7         | 280 $\mu\text{m}$ |
| MA        | $10^{-3}$     | 7.0          | 5 nm              |
| SA        | $10^{-3}$     | 4.0          | 2 nm              |
| SM        | $10^{-3}$     | 4.0          | 0.4 - 2 nm        |
| corner    | $10^{-3}$     | 4.0          |                   |

**Table 6.1:** Materials parameters used in the simulations shown in Fig. 6.1.  $\tan \delta = 10^{-3}$  stands at the lower bound of the limits reported by [51] .



**Figure 6.1:** **a** Participation ratio simulation results as a function of the thickness of a dielectric at the SM interface. **b** Cartoon of the simulated CPW cross-section, its materials and their interfaces. The electric field lines of the resonant mode are sketched as the turquoise lines.

standard resonator geometry (center conductor width  $w = 20 \mu\text{m}$  and the gap to ground  $g = 10 \mu\text{m}$ ), using the materials parameters listed in Table 6.1. Studying the impact of variations in the simulated parameters, we find that as soon as we allow for almost any dielectric to exist at the substrate-metal interface, the loss at this interface quickly becomes dominant (see Fig. 6.1). Below a certain effective thickness of the SM interface, it is the losses at the MA and SA interfaces that dominate. However, due to the cumulative nature of losses at all participating interfaces, reductions of loss at the MA and SA interfaces can still result in measurable coherence improvements in cases where the SM loss has a dominant status.

It is tempting to assume negligible loss at the SM interface, as this is the first interface we form during the fabrication process, and thus we have good control over its formation with a relatively large freedom in choosing the optimal surface cleaning procedures. At

later stages of the fabrication process, we have to consider the durability of the structures present on the wafer, and we become more limited in the harshness of cleaning treatments permitted. During fabrication, we accumulate adsorbates and fabrication residues on surfaces [121], and can even promote excessive oxidation beyond the ambient native oxidation through thermal or acid-base treatments [122]. Surficial oxide is not only a prominent source of dielectric loss by itself, but can also encapsulate the aforementioned contaminants, exacerbating the associated loss. Some materials, such as Nb and Ta, are resilient towards harsh chemical treatments that can remove surface contaminants and even oxides while leaving the metal intact, such as the buffered oxide etch or the piranha clean [19, 106, 123]. This is very helpful for cleaning the SA and MA interfaces in-between the various processing steps. However, once the Al/AlOx-based JJs are on, the processing limitations become more stringent.

## 6.2 Pre-deposition substrate treatment in evaporated films

In our standard fabrication process, we form the SM interface of the CPW structures by first cleaning the Si surface with SC1, which removes organic contaminants (see Section 4.2). Afterwards, we remove the Si oxide (SiOx) by dipping the wafer in 2% HF for 60 s, which should leave behind a hydrophobic, oxide-free Si surface, passivated by a Si-H monolayer (see section 4.2). As this Si-H passivation is temporary, the native SiOx will regrow with time. Therefore, we immediately load the HF-treated wafer into a vacuum-pumped loadlock of a thin-film deposition tool, which in the case of our standard process is a Plassys evaporator. Inside, the wafer is heated to 300 °C for 10 min to desorb moisture and other volatile adsorbates, such as atmospheric organic matter [124–126]. Afterwards, the wafer is left to passively cool to room temperature (RT) while pumping to reach a lower chamber pressure, and to promote the desorption of matter from the surfaces. If we pump for 24 h, which is the starting point of the following experiments, the pressure in the chamber reaches  $\sim 4 \times 10^{-8}$  mbar. Then, a 150 nm thick Al film is deposited at  $1 \text{ nm s}^{-1}$ , after which the film is oxidized *in situ* in pure oxygen.

In the experiment published in Paper **B**, we fabricate five wafers, where we systematically vary one step of the pre-deposition Si treatment at a time. Sample #1 serves as a control sample, with all the aforementioned standard steps (SC1, HF, preheat, 24h pumpdown) present in the fabrication process. In samples #2–4, we skip the HF dip, SC1, and preheat, respectively. In sample #5, we significantly cut down the pre-deposition vacuum pumping time from the standard 24h to 4h, after which the pressure

|                                   | Sample 1<br>(#1)       | Sample 2<br>(#2)      | Sample 3<br>(#3)      | Sample 4<br>(#4)      | Sample 5<br>(#5)       |
|-----------------------------------|------------------------|-----------------------|-----------------------|-----------------------|------------------------|
| Process variation                 | Standard               | No HF                 | No SC1                | No preheat            | Short pumping          |
| No. of resonators                 | 8                      | 6                     | 6                     | 7                     | 6                      |
| $Q_i$ ( $\langle n \rangle = 1$ ) | $1.1 \times 10^6$      | $0.24 \times 10^6$    | $0.95 \times 10^6$    | $0.8 \times 10^6$     | $1.1 \times 10^6$      |
| $F\delta_{\text{TLS}}^0/10^{-6}$  | 0.87<br>( $\pm 0.29$ ) | 4.4<br>( $\pm 0.79$ ) | 1.0<br>( $\pm 0.23$ ) | 1.4<br>( $\pm 0.70$ ) | 0.87<br>( $\pm 0.20$ ) |
| $\delta_{\text{other}}/10^{-7}$   | 2.3<br>( $\pm 0.53$ )  | 2.3<br>( $\pm 0.30$ ) | 2.0<br>( $\pm 0.72$ ) | 2.1<br>( $\pm 0.43$ ) | 2.4<br>( $\pm 0.76$ )  |

**Table 6.2:** Summary of samples #1-5 with different pre-evaporation substrate treatments. The table includes the number of resonators on each sample, the average  $Q_i$  at  $\langle n \rangle = 1$  across all the resonators on the sample, and the average TLS model fit parameters  $F\delta_{\text{TLS}}^0$  and  $\delta_{\text{other}}$ , including their standard deviations.

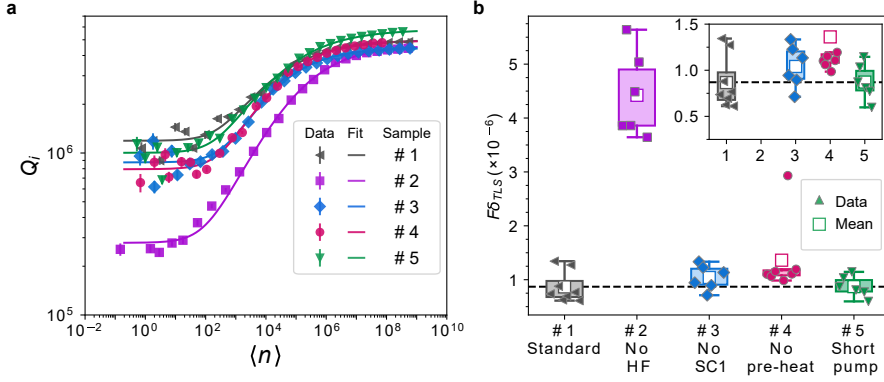
in the deposition chamber reaches  $\sim 6 \times 10^{-8}$  mbar and the temperature is  $50^\circ\text{C}$ .

We then fabricate chips with CPW resonators with our standard  $g : w : g = 10 : 20 : 10 \mu\text{m}$  geometry from wafers #1–5. We measure the  $Q_i$  vs  $\langle n \rangle$  dependence of these resonators, which we fit to the TLS model in Eq. 2.4. The extracted parameters are shown in Fig. 6.2 and summarized in Table 6.2.

Despite keeping the resonator geometry and the post-deposition fabrication procedure identical between the samples, we observe measurable differences in the extracted  $F\delta_{\text{TLS}}^0$ , which we attribute to differences in loss at the SM interface caused by the pre-deposition procedure deviations.

In sample #2, we study the effects of failing to remove the native SiOx from the substrate surface prior to deposition, which increases the average  $F\delta_{\text{TLS}}^0$  of this sample 5-fold compared to the standard sample #1. This is unsurprising, as the transmission electron micrograph (TEM) of sample #2 in Fig. 6.3b shows a 1.5 nm thick interfacial oxide layer, which is a prominent source of TLS loss. In contrast, the TEM of sample #1 in Fig. 6.3a shows that our standard procedure can result in a pristine interface between the Si and Al, with no continuous oxide layer in-between.

We also perform elemental depth profiling of the different samples by time-of-flight secondary-ion-mass-spectrometry (ToF-SIMS) – a highly accurate method capable of simultaneous analysis of species with different masses [127]. In our ToF-SIMS set-up, an auxiliary  $\text{Cs}^+$  beam sputters through the material stack, starting from the MA interface all the way through until the bulk of the Si wafer. Concurrently, the elemental com-




**Figure 6.2:** Energy loss in CPW resonators fabricated with different pre-deposition treatments. **a**  $Q_i$  vs. circulating power for one resonator each from samples #1–5 with a frequency  $\sim 4.4$  GHz, fitted to the TLS model in Eq. 2.4. **b** Box plot of the TLS loss  $F\delta_{\text{TLS}}^0$  extracted from the TLS fit of each resonator on samples #1–5. The inset shows a close-up of the area of panel **b** where most data points for samples #1 and #3–5 are located.

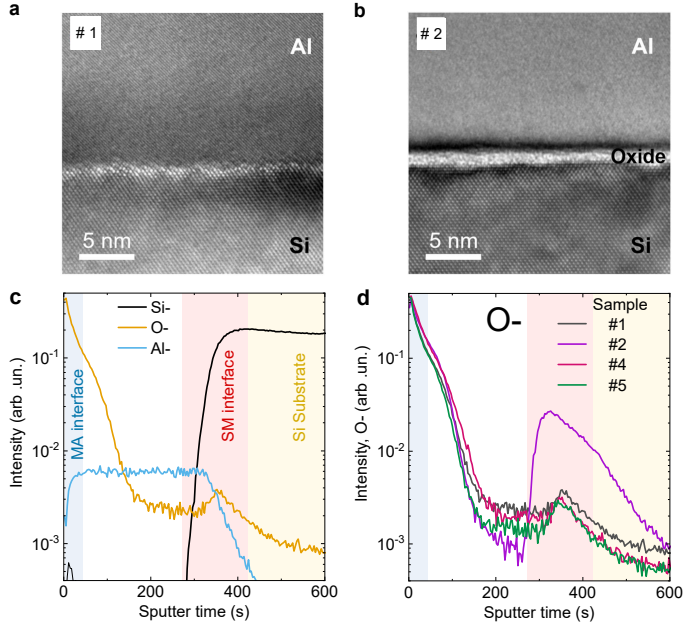
position of the stack is analyzed as a function of sputtered depth – a primary  $\text{Bi}^{3+}$  ion beam generates secondary ions from the exposed matter, which are accelerated towards the mass spectrometer, where their time-of-flight and mass/charge ratios are analyzed. An example of this analysis for sample #1 is shown in Fig. 6.3c, and the intensity of the measured oxygen content for the various samples of interest is shown in Fig. 6.3d.

The measured oxygen intensity for sample #2 shows a distinct peak at the SM interface, which is consistent with the oxide layer observed in the TEM of this sample. Surprisingly, this SM peak is also present in the remaining samples, although significantly weaker than in sample #2. The peak is also present in sample #1, where no oxide layer was observed at the SM interface in the TEM in Fig. 6.3a. We study the origin of this peak in more detail in Paper **A**, as well as section 6.4 of this chapter.

In Al films grown on SiOx, it is thermodynamically favourable for the SiOx to be reduced into AlOx [128, 129]:



We confirm that this is the case in our samples through ToF-SIMS measurements in the Supplementary materials to Paper **A**, which show that the SM peaks consist of predominantly AlOx species, as opposed to SiOx.



**Figure 6.3:** TEM and ToF-SIMS investigation of the SM interface of samples in Fig. 6.2. **a** TEM of sample #1 with a pre-deposition oxide strip in HF shows a direct interface between Si and Al with no continuous intermediate oxide layer. The bright contrast at the interface is due to Fresnel fringe contrast. **b** TEM of sample #2 with no HF treatment, showing a 1.5 nm thick oxide layer at the SM interface. **c** Depth profile of the standard-process sample (#1) characterized by ToF-SIMS. The depth is proportional to the time during which the sample was etched by the  $\text{Cs}^+$  sputter beam. The MA and SM interfaces are highlighted with blue and red, respectively, while the Si substrate is in yellow. **d**) The depth profile of  $\text{O}^-$  from sample #1 (standard process), #2 (no HF), #4 (no preheat), and #5 (short pumping).

In sample #3, even though the  $\text{SiO}_x$  was stripped by HF, the omission of SC1 results in a 15% increase of  $F\delta_{\text{TLS}}^0$  compared to the standard #1, which can be attributed to residual contamination of the substrate surface.

Sample #4, which was treated by both SC1 and HF but did not receive the pre-heat treatment to  $300^\circ\text{C}$  inside the vacuum chamber, shows a 61% increase in  $F\delta_{\text{TLS}}^0$  compared with sample #1 when considering all resonators on the sample, and a 26% increase when the one outlier resonator is discarded, in which case the average  $F\delta_{\text{TLS}}^0$  of sample #4 falls to  $1.1 \pm 0.07 \times 10^{-6}$ . We observe no discernible difference in the oxygen concentration at the SM interface of sample #4 compared to the standard sample #1 in

the ToF-SIMS measurement shown in Fig. 6.3d. It is possible that change in the oxygen level is below the observable limit, or that the factor behind the increased performance of the devices on sample #1 when the preheat step is included is not the desorption of oxidizing species.

On the other hand, reducing the pre-deposition pumping time with the wafer inside the vacuum chamber from our previously standard 24 h to 4 h showed no adverse effect on the TLS loss of resonators on sample #5.

Some of the observed deviations in the extracted  $F\delta_{TLS}^0$  of samples #1–5 are to be expected - particularly in sample #2, where we obtain a relatively thick layer of lossy dielectric at the dominant SM interface in the absence of the HF oxide strip step. However, variations of  $F\delta_{TLS}^0$  when skipping the preheat step or SC1 are harder to predict, and therefore become difficult to account for in device modeling and simulations. This highlights an issue in adapting materials parameters values from literature, which can differ measurably between different laboratory environments and fabrication procedures.

### 6.3 Pre-deposition treatment and deposition conditions in sputtered films

In a related experiment, also published in Paper **B**, we study the loss of resonators made from sputtered Al films. In samples #6–11, we vary the pre-deposition substrate treatments, as well as the deposition conditions of the sputtered films.

The post-deposition fabrication of the CPW resonators made from wafers #6–11 is identical to those of #1–5, as is the device design. As for pre-deposition treatments, SC1 was omitted for samples #6–11 unless otherwise specified, but the Si wafers did receive the standard 2% HF treatment, after which they were loaded into a heated loadlock (80 °C) of the sputter tool for 40 min until a pressure below  $5 \times 10^{-7}$  mbar was reached. Afterwards, the substrate was heated to 300 °C for 10 min and pumped for 16 h, at which point the base pressure of the deposition chamber reached  $2.2 \times 10^{-8}$  mbar. Al was then deposited at a rate of  $1 \text{ nm s}^{-1}$  by direct-current magnetron sputtering in Ar plasma, followed by in-situ oxidation of the film surface.

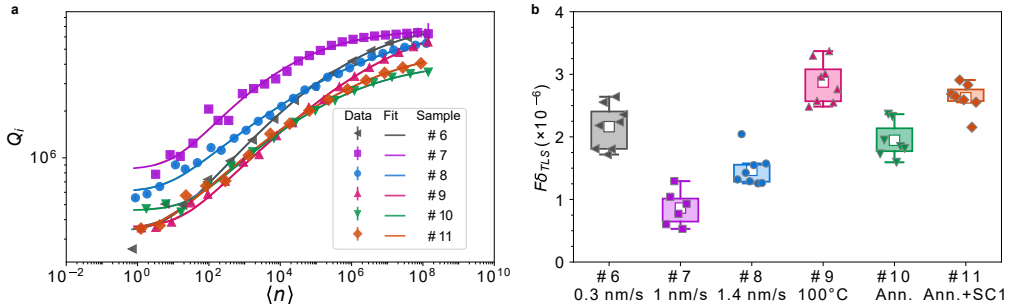
In samples #6–8, we investigate the impact of sputtering rate ( $0.3 \text{ nm s}^{-1}$ ,  $1 \text{ nm s}^{-1}$ , and  $1.4 \text{ nm s}^{-1}$ ), which we vary by changing the forward power delivered to the DC magnetron that ignites the Ar plasma. In sample #9, we heat the wafer to 100 °C immediately prior to the Al deposition with the intent to promote larger grain sizes [130]. Finally, the Si substrates in samples #10–11 are annealed at 700 °C for 25 min instead of the standard 300 °C for 10 min in order to promote Si surface reconstruction



and achieve a better atomic transition between the metal and substrate [131, 132]. In addition, sample #11 received the SC1 treatment. The varied fabrication parameters for samples #6–11 are summarized in Table 6.3, together with the key parameters extracted from CPW resonators fabricated from these samples. The resonator results are also shown in Fig. 6.4. Additionally, we measure the surface roughness of Al films prepared using the above conditions by scanning a  $5 \times 5 \mu\text{m}$  area of the Al surface using atomic force microscopy. The obtained values for the quadratic mean roughness  $R_q$ , also known as the root mean square (RMS) roughness, are summarized in Tab. 6.4.

Samples #6–9 show the impact of deposition conditions, which can influence the crystallographic parameters of the film, its purity, and the properties of both the SM and MA interfaces. Sample #7 is the closest in deposition conditions to the standard evaporated sample #1 (barring the omitted SC1 and the sputtered deposition at an equivalent deposition rate). Despite these differences and a doubled surface roughness of sample #7 compared to #1, the two samples show comparable TLS loss. Therefore, we can use sputtered films for the wiring layer of our devices.

The resonator data from samples #6–8 indicates that there is an optimal intermediate sputtering rate for decreasing TLS loss, as both the higher and lower deposition rates yield higher TLS loss than the intermediate  $1 \text{ nm s}^{-1}$ . This optimum may be governed by grain size, or perhaps by the variations in the system conditions that govern the sputtering rate (plasma power and the associated heat evolution, or alternatively the target-substrate distance and substrate position relative to the plasma profile). Gener-



**Figure 6.4:** Energy loss in CPW resonators fabricated on sputtered films with different pre-deposition treatments and deposition conditions. **a**  $Q_i$  vs. circulating power for one resonator each from samples #6–11 with a resonance frequency  $\sim 4.8 \text{ GHz}$ , fitted to the TLS model in Eq. 2.4. **b** Box plot of the TLS loss  $F\delta_{\text{TLS}}^0$  extracted from the TLS fit of each resonator on samples #6–11.

|                                  | Sample 6<br>(#6)       | Sample 7<br>(#7)     | Sample 8<br>(#8)       |
|----------------------------------|------------------------|----------------------|------------------------|
| Variation                        | 0.3 nm s <sup>-1</sup> | 1 nm s <sup>-1</sup> | 1.4 nm s <sup>-1</sup> |
| No. of resonators                | 8                      | 6                    | 8                      |
| $Q_i (\langle n \rangle \sim 1)$ | $0.40 \times 10^6$     | $1.0 \times 10^6$    | $0.58 \times 10^6$     |
| $F\delta_{\text{TLS}}^0/10^{-6}$ | 2.2                    | 0.86                 | 1.5                    |
|                                  | (±0.35)                | (±0.29)              | (±0.26)                |
| $\delta_{\text{other}}/10^{-7}$  | 1.6                    | 1.5                  | 1.3                    |
|                                  | (±0.55)                | (±0.24)              | (±0.38)                |

|                                  | Sample 9<br>(#9)   | Sample 10<br>(#10) | Sample 11<br>(#11)   |
|----------------------------------|--------------------|--------------------|----------------------|
| Variation                        | 100 °C dep.        | 700 °C preheat     | 700 °C preheat + SC1 |
| No. of resonators                | 8                  | 7                  | 7                    |
| $Q_i (\langle n \rangle \sim 1)$ | $0.32 \times 10^6$ | $0.48 \times 10^6$ | $0.38 \times 10^6$   |
| $F\delta_{\text{TLS}}^0/10^{-6}$ | 2.9                | 2.0                | 2.6                  |
|                                  | (±0.34)            | (±0.29)            | (±0.24)              |
| $\delta_{\text{other}}/10^{-7}$  | 1.9                | 2.0                | 1.8                  |
|                                  | (±0.98)            | (±0.66)            | (±1.7)               |

**Table 6.3:** Summary of samples #6-11 with sputtered Al. The table includes the number of resonator on each sample, the average  $Q_i$  at  $\langle n \rangle \sim 1$  across all the resonators on the sample, as well as the average TLS model fit parameters  $F\delta_{\text{TLS}}^0$  and  $\delta_{\text{other}}$ , including their standard deviations.

| Substrate treatment | Deposition method  | Equivalent sample # | $R_q$   |
|---------------------|--------------------|---------------------|---------|
| 300 °C, 10 min      | Evaporation, RT    | #1                  | 1.4 nm  |
| 300 °C, 10 min      | Sputtering, RT     | #7                  | 3.2 nm  |
| 300 °C, 10 min      | Sputtering, 100 °C | #9                  | 58.9 nm |
| 700 °C, 25 min      | Sputtering, RT     | #10                 | 11.3 nm |

**Table 6.4:** Surface roughness values of different Al films, alongside the corresponding deposition parameters. The equivalent sample # denotes the resonator sample prepared with a film fabricated using identical deposition parameters.

ally, increased deposition rates are desired to reduce the probability of trapping impurities in the film during growth. In addition, the higher kinetic energy of atoms sputtered at high power gives the atoms increased mobility, which affects the grain size. However, this increased kinetic energy can also enable implantation of the target material in the substrate, degrading the quality of the SM interface.

In sample #9, raising the substrate temperature to 100 °C prior to deposition yields a significantly higher surface roughness than the standard - so much so that the film looks white to the naked eye, as opposed to the typical mirror-like finish of Al thin films. It is plausible that such an increase in roughness would exacerbate loss at the MA interface.

In both samples #10-11, the 700 °C annealing step degraded the low-power  $Q_i$  of the resonators. This may have been caused by the deterioration of the residual hydrogen monolayer on the Si substrate after the HF treatment, and a subsequent re-oxidation of the Si surface. Additionally, the surface roughness increased in comparison with the standard sample, which may have caused additional losses at the MA interface.

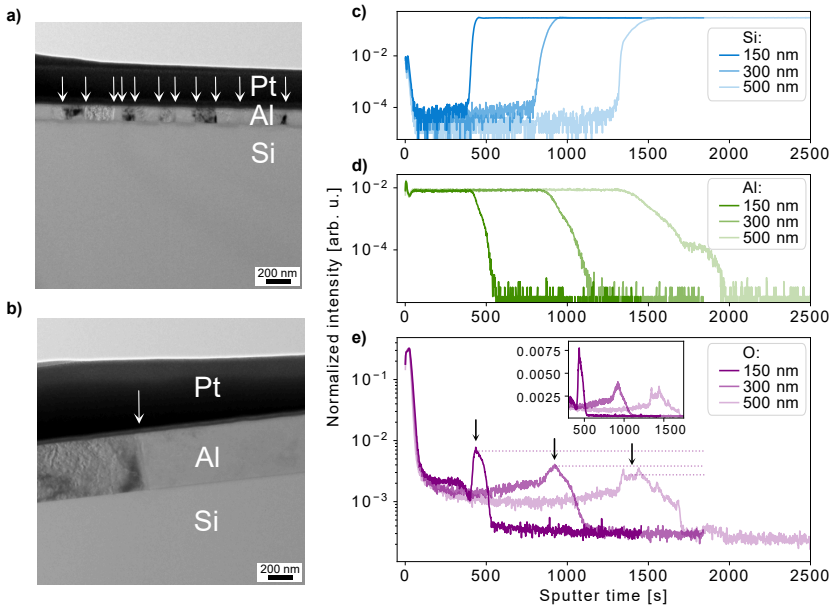
We also observe a decrease in the other, non-TLS related losses in the sputtered films in comparison to evaporated films. Further studies into the reason behind this decrease are needed; however, a decrease in quasiparticle loss may be partially responsible, as the  $T_c$  of the sputtered sample #7 was higher than of the evaporated sample #1 (1.2 K and 0.91 K, respectively).

## 6.4 Film thickness and grain size

As shown in Fig. 6.3, the TEM of the SM interface of our standard sample (prepared by evaporating Al on HF-treated Si) shows a direct interface between the Si and Al, which is seemingly at odds with the peak in the oxygen concentration at this interface inferred from the ToF-SIMS of the same sample. In the following experiments, published in Paper A, we aim to investigate the origin of this oxygen peak, devise a strategy to mitigate it, and study the impact of this mitigation of oxygen content at the SM interface on the coherence of our devices.

### 6.4.1 Materials analysis

While a TEM of the type shown in Fig. 6.3a would detect a continuous layer of oxide such as the one in Fig. 6.3b, localized and discontinuous oxide species may be overlooked – however, these could be picked up by the exceptionally sensitive ToF-SIMS. Additionally, a key difference between the two investigations presented in Fig. 6.3 is that the ToF-



**Figure 6.5:** TEM and ToF-SIMS of Al films with different thicknesses. **a–b** TEM of a 150 nm film in **a**, and a 500 nm film in **b** at identical magnification. Each arrow points at an identified grain boundary. **c–d** Elemental depth profiles of three different film thicknesses obtained through ToF-SIMS, shown here for Si, Al and O, respectively. The  $x$ -axis shows the sputtering time, which is proportional to the sputtered depth. The  $y$ -axis shows the intensity of the measured signal identifying each element, normalized by the total beam current. The sharp increase in the intensity of Si in **c**, and the arrows in **e**, mark the SM interface of each sample. The inset in **e** shows the intensity of the measured O signal at the SM peak on the linear scale.

SIMS measurement samples a significantly larger area of the sample ( $100 \times 100 \mu\text{m}^2$ ). This makes the ToF-SIMS more likely to also sample areas with local defects, as well as one very important part of thin film structure – the grain boundaries.

As shown in another, lower-magnification TEM in Fig. 6.5**a**, under high-vacuum conditions, Al can grow in columnar structures of parallel grain boundaries [133]. These grain boundaries can support oxide growth in-between grains, and can also act as diffusion channels for contaminants, including oxidizing agents. The size of the parallel grains depends on the film thickness [134, 135], which we verify on our system by imaging the grains of Al films of two different thicknesses – 150 nm and 500 nm – in Fig. 6.3**a–b**. This method of changing grain size by film thickness and keeping all other growth parameters

constant is convenient for isolating the impact of grain size on device performance, as this method does not require a change of system parameters influencing the deposition rate (temperature, electron beam current, target position, etc.). Such changes may by themselves jeopardize the quality of the interfaces or the film purity – as observed in the sputter rate experiments of Section 6.3.

Having verified that growing thicker films gives rise to larger grains and therefore fewer grain boundaries, we study the effects on the elemental depth profiles of Si/Al stacks with Al films of three different thicknesses - 150 nm, 300 nm and 500 nm - prepared using the updated standard process from Section 6.2 with a 4 h pre-deposition pumping time. In Fig. 6.5e, we find that the oxygen level inside the film decreases in the thicker films. Additionally, the SM oxygen peak intensity, indicated by the dashed lines, falls with a ratio of 5:3:2 from the thinnest to the thickest film.


While it appears in Fig. 6.5e as if the Si/Al interface were broadening with increasing film thickness, this is unlikely to be a real feature of the samples. Instead, this broadening is a known artifact caused by surface roughening and atomic mixing during the sputter-assisted depth profiling [136, 137]. Although disregarding the width of the peaks in similar scenarios is not unusual, we compare the peak intensity and the integrated intensity for O across the entire Al film, including the MA and SM interfaces. The integrated intensity divided by film thickness scales with a ratio of 6:3:2.

### 6.4.2 TLS loss in resonators

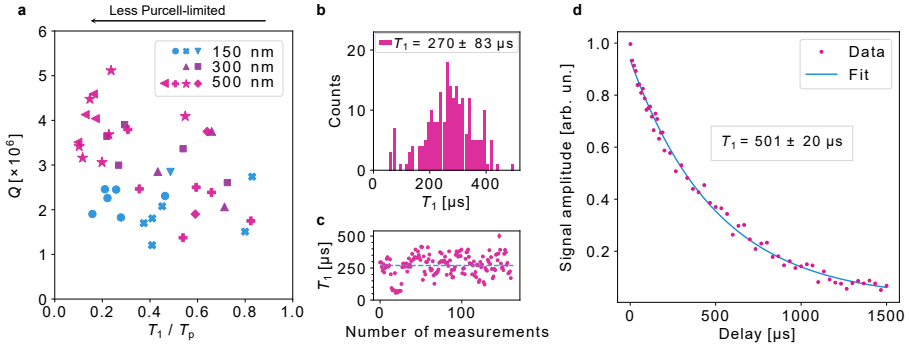
Having established that growing thicker Al films gives rise to larger grain sizes and consequently fewer grain boundaries that support oxide formation, which is a common suspect in TLS loss, we set out to study the effect of Al film thickness on the TLS loss of CPW resonators. In Fig. 6.6a, we show the  $F\delta_{TLS}^0$  of 50 resonators, fabricated from Al films of the three thicknesses studied in the SIMS analysis. This data shows that the average  $F\delta_{TLS}^0$  of resonators made from 150 nm thick films is  $1 \times 10^{-6}$ , while the  $F\delta_{TLS}^0$  of 500 nm films is half that at  $5 \times 10^{-7}$  on average. The 300 nm films resulted in resonators with average  $F\delta_{TLS}^0$  of  $8 \times 10^{-7}$ .

### 6.4.3 Qubit coherence

Comparing the quality of different qubits fairly can be challenging due to the frequency-dependence of the decoherence time constants we measure ( $T_1$ ,  $T_2$ ), as well as due to the design-dependent Purcell decay into the readout resonator or other elements coupled to the qubit.



**Figure 6.6:** TLS loss in bare CPW resonators with Al films of different thicknesses. **a**  $Q_i$  as a function of the average number of photons  $\langle n \rangle$  fitted to the TLS model in Eq. 2.4, for two resonators with a comparable  $\delta_{other}$  and resonant frequency of 4.45 GHz. **b**  $F\delta_{TLS}$  extracted from the TLS model fits, plotted as a function of frequency for resonators with the three different film thicknesses of 150 nm, 300 nm and 500 nm. Data indicated by the same marker belong to resonators on a single chip. The dashed lines indicate the average value of  $F\delta_{TLS}$  for a given film thickness. For best visual comparability, the resonators showcased in **a** are chosen for their proximity in high-power  $Q_i$ , and are indicated in **b** with a dashed-line square.



**Figure 6.7:** Energy relaxation of qubits with Al films of different thicknesses. **a** Time-averaged  $Q$  of 40 qubits fabricated with 150 nm, 300 nm and 500 nm Al films, plotted as a function of the measured relaxation time relative to the calculated Purcell decay time. Qubits with  $T_1 / T_p \rightarrow 1$  are limited by Purcell decay, while the quality factor of those qubits with  $T_1 \ll T_p$  is closer to the limit set by TLS loss. Qubits represented by identical markers are fabricated on the same wafer. **b** Histogram of the  $T_1$  values of the best qubit from **a** measured over 48 h, reaching an average  $T_1 = 270 \mu s$  plus/minus one standard deviation of  $83 \mu s$ . **c** The individual  $T_1$  measurements of the qubit in **b**, measured over a span of 48 h. Fit error bars are smaller than markers where not visible. **d** Demonstration of the exponential fit to the longest  $T_1$  measured.

We solve the former issue by converting the energy relaxation times  $T_1$  into the frequency-independent quality factors  $Q = \omega_q T_1$ , where  $\omega_q$  is the angular frequency of the qubit. As established in Section 2.1, we expect the total qubit loss  $1/Q$  to be a linear combination of several loss mechanisms. Analogously to the common practice in resonator characterization, we separate these loss contributions into the internal qubit quality factor  $Q_i$  with contributions from all losses native to the qubit, and the external (Purcell) factor  $Q_p$ :

$$\frac{1}{Q} = \frac{1}{Q_{\text{TLS}}} + \frac{1}{Q_{\text{qp}}} + \frac{1}{Q_{\text{mag}}} + \frac{1}{Q_{\text{rad}}} + \frac{1}{Q_p} \dots = \frac{1}{Q_i} + \frac{1}{Q_p}. \quad (6.2)$$

To compute the theoretical  $Q_i$  of a qubit in the absence of Purcell decay, we could calculate  $Q_p$  and subtract it from the measured qubit  $Q$ . However, the theoretical Purcell decay rate, as introduced in Eq. 2.11, might not be accurate for some cases (such as large detunings or asymmetric resonances [138]), and require correction factors. Instead, we can transform the equation  $1/Q = 1/Q_i + 1/Q_p$  into:

$$\frac{1}{Q_i} = \frac{1}{Q} \left( 1 - \frac{T_1}{T_p} \right). \quad (6.3)$$

In this form, we find that the qubit  $Q$  approaches its limit set by its internal loss mechanisms (such as by  $Q_{\text{TLS}}$ , as suggested by the resonator measurements), when  $T_1/T_p \ll 1$ . On the other hand, as  $T_1/T_p$  approaches 1, the qubit becomes close to fully limited by Purcell decay. Therefore, we plot the measured  $Q$  of 40 qubits, fabricated from films of the three familiar thicknesses, as a function of  $T_1/T_p$  in Fig. 6.7.

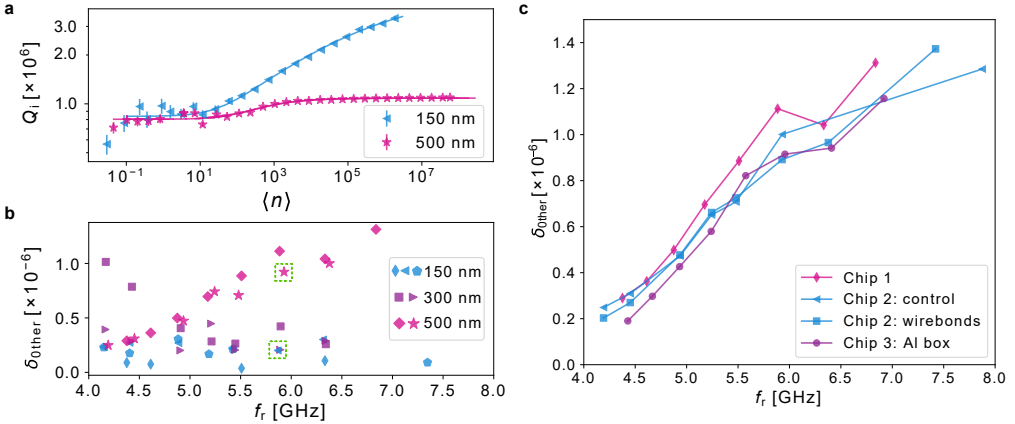
The average qubit  $Q$  for the 150 nm thick films is  $2.1 \times 10^6 \pm 23\%$ , which increases to  $3.3 \times 10^6 \pm 26\%$  for the “thicker” films (300 nm and 500 nm). The differences between the two “thicker” films are not sufficiently pronounced as to consider them separately. The separation between the “thin” and “thick” datasets scales with the Purcell decay. In the Purcell-limited regime of  $T_1/T_p \rightarrow 1$ , the data is scattered and the performance of different thicknesses is indistinguishable. For  $T_1 \leq 0.5 T_p$ , where the qubits are not significantly limited by Purcell decay, the data for the thicker and thinner films are clearly separated, with the qubits fabricated from the 150 nm thick Al showing a comparatively lower  $Q$ . The average quality factor in this range is  $2.1 \times 10^6 \pm 22\%$  for the qubits on the thin film, while it increases by 76% to  $3.7 \times 10^6 \pm 18\%$  for those on the thicker films. For a more stringent case of  $T_1 \leq 0.25 T_p$ , the average quality factor becomes  $3.9 \times 10^6 \pm 16\%$  for the thicker films.

### 6.4.4 High-power loss

In the resonator studies, we find a frequency-dependent increase in the other, non-TLS-related losses for the 500 nm thick films, see Fig. 6.8a–b. As the  $F\delta_{\text{TLS}}^0$  values of resonators made from 500 nm thick film are now so small they become comparable to  $\delta_{\text{other}}$ , this other loss warrants further investigation.

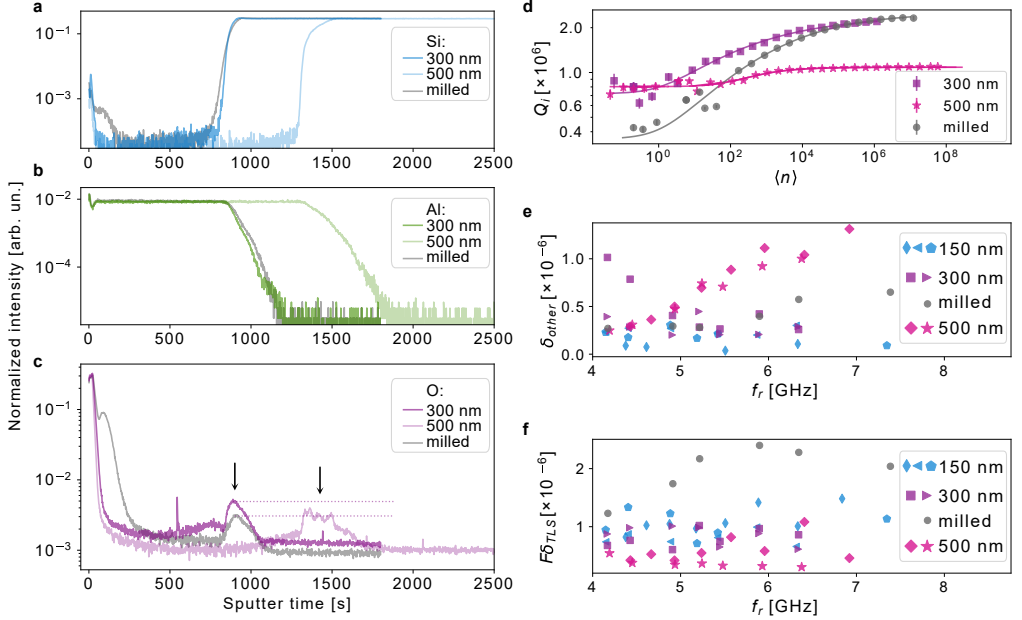
Depending on the loss mechanism, qubits can be affected in a similar way to resonators, or barely at all. For example, radiation loss will depend on geometry, and the relatively small qubit and the distributed resonator are impacted differently. On the other hand, quasiparticles or magnetic flux noise may affect both structures in similar ways.

We hypothesize that the radiation loss of the resonators could be exacerbated by decay into the resonant mode of the copper sample box, or into parasitic slotline modes.



**Figure 6.8:** Non-TLS related loss in bare CPW resonators with Al films of different thicknesses. **a**  $Q_i$  as a function of average number of photons  $\langle n \rangle$  fitted to the TLS model in Equation 2.4, for two resonators. The resonance frequency is 5.9 GHz, which is high enough to observe the exacerbated  $\delta_{\text{other}}$  for high-frequency resonators made from the 500 nm films. **b** The loss due to other sources  $\delta_{\text{other}}$  extracted from the TLS model fits, as a function of frequency for resonators of three different film thicknesses. Data indicated by the same marker belong to resonators on a single chip. For best visual comparability, the resonators showcased in **a** are chosen for their proximity in frequency and low-power  $Q_i$  and are indicated in **b** with a dashed-line square. **c**  $\delta_{\text{other}}$  of resonators made on 500 nm thick films. *Chip 1* and *Chip 2* are placed (each) in a copper sample box, while *Chip 3* is in an aluminum box. In *Chip 2*, we add wirebonds across the main transmission line and the resonators.





**Figure 6.9:** Elemental analysis of the Al/Si material stack by ToF-SIMS and resonator loss from corresponding films. Depth profiles of **a)** Si, **b)** Al, and **c)** O. The sample in grey shows an Al film milled from 500 nm to 300 nm, the other two samples are Al films of the two respective thicknesses. The sharp increase in the intensity of Si in **(a)**, and the arrows in **(c)**, mark the SM interface of each sample. **d)**  $Q_i$  as a function of average number of photons  $\langle n \rangle$  fitted to the TLS model in Equation 2.4, for three resonators with a resonant frequency of 5.9 GHz. **e–f)**  $\delta_{other}$  and  $F\delta_{TLS}$  of resonators made from films of three different thicknesses, as well the film milled from 500 nm to 300 nm show that this milling reduces the  $\delta_{other}$  exacerbated in thick films at the expense of increasing TLS loss.

To test this, we measure three identical chips coming from the same wafer (see Fig. 6.8c), packaged in different ways. We place one chip in an Al box instead of Cu, which we expect will change the  $Q$  of the sample box mode, and thus the resonators' decay rate into this mode. We measure another chip (*Chip 2*) in a Cu box as usual, and then we add wirebonds across the transmission line and the resonator to suppress slotline modes, after which we re-measure this chip. We do not observe a significant decrease in  $\delta_{other}$  in either of these experiments.

To investigate to which extent this increased  $\delta_{other}$  has to do with the film thick-

ness and how much with its structure, we fabricate another 500 nm film and use  $\text{Ar}^+$  ion milling to thin the film down to 300 nm. We perform elemental depth profiling of the thinned-down film using ToF-SIMS. We also fabricate CPW resonators, which we compare to regular films of the corresponding two thicknesses (300 nm and 500 nm), see Fig. 6.9. We find that the SM oxygen peak of the 500 nm film milled down to 300 nm remains unchanged after the milling, and the SM interfacial O content corresponds to the comparatively low value of an unmilled 500 nm film. On the other hand,  $\delta_{other}$  drops to lower values consistent with the 300 nm films, and loses its frequency dependence. This is encouraging; however, the significant damage to the SM interface increases TLS loss to such an extent, that the resulting low-power  $Q_i$  is significantly diminished. This milling-induced TLS loss is also observed in [139].

### 6.4.5 Discussion

Throughout this chapter, I have identified the interface of our CPW geometry at which excess dielectric loss has the largest impact - the SM interface. I have identified a potential source of loss at this interface - a peak in oxygen concentration measured by ToF-SIMS, which I have correlated with the grain size associated with a given film thickness. I have shown that increasing the Al film thickness while keeping all deposition conditions constant increases the grain size, which reduces the number of grain boundaries, which in turn lowers the peak in oxygen concentration. Fabricating CPW resonators from films of different thicknesses showed that dielectric loss due to TLSs is reduced in thicker films. Finally, I showed that qubits fabricated from thicker films ( $\geq 300$  nm) outperform our previously standard 150 nm thick films.

Now, I explore alternative explanations for the observed decrease in qubit and resonator loss when increasing the film thickness from 150 nm to 300 nm or higher.

The dilution of electric field associated with an increased film thickness decreases the participation ratios of the interfaces. Simulating this effect on our standard CPW cross-section, we expect  $Q_{TLS}$  to increase by about 5% from 150 nm to 500 nm, which does not explain the magnitude of improvement we observe in the measured data.

Another TLS-related hypothesis would concern the impact of change in granular structure on the Al surface roughness, and associated losses at the MA interface. However, atomic force microscopy reveals negligible differences in the surface roughness of the 150 nm and 500 nm thick films. Additionally, the ToF-SIMS data from the MA interfaces of the various film thicknesses are indistinguishable - any improvement from reduced grain boundary oxidation is negated by the  $\sim 5$  nm of native  $\text{AlOx}$  we expect on our films (see Chapter 7).

Additionally, thicker films could result in higher  $T_c$  values, which could suppress equilibrium quasiparticle loss. However, in our dc-transport measurements, we observe no such trend in  $T_c$  between the three measured film thicknesses (0.91 K, 0.85 K, and 0.93 K for 150 nm, 300 nm and 500 nm, respectively). On the other hand, the measured residual-resistance ratio increases with film thickness (8, 17, and 28 from the thinnest of the three to the thickest), confirming the higher crystallographic quality of the thicker films.



---

## Mitigating loss - metal-air and substrate-air interfaces

---

In Section 6.1, I have discussed how participation ratio simulations of our standard CPW geometry indicate the dominant status of loss due to TLS at the substrate-metal (SM) interface.

Once loss at the SM interface is sufficiently mitigated, the losses at the metal-air (MA) and substrate-air (SA) interfaces are expected to dominate – provided that the device performance is still limited by TLS loss after the SM loss is mitigated. However, as the losses at the various interfaces are additive, we may observe the impact of a change in the non-dominant loss as well. In this chapter, I present our investigations into losses at the MA and SA interfaces in CPW resonators. The data collected during the investigations shown in this chapter includes results published in Paper **B**, as well as previously unpublished data.

Additionally, I will discuss the results published in Paper **D**, where we investigate loss in 3D superconducting cavities. We find evidence of TLS-related loss in these cavities, which, in the absence of a substrate and any associated substrate-related interfaces, we attribute to the MA interface.

I will also discuss the results of Paper **E**, where we integrate our resonators and qubits into a more scalable flip-chip architecture without measurably degrading qubit performance compared to our standard planar architecture.

## 7.1 Passivation of metal films

Most metals, including the popular choices for superconductors in quantum processors, grow amorphous oxide when exposed to ambient air. Throughout the previous chapters, I have discussed how such interfacial oxide amounts to dielectric loss in quantum circuits. In this section, I explore strategies for controlling or hindering the growth of surfacial oxide in thin metal films.

### 7.1.1 In-situ oxidation

Despite how counter-intuitive this step might sound after the discussion concerning the lossy nature of amorphous oxides preceding this chapter, in our standard qubit and resonator fabrication process, we conclude each Al deposition by oxidizing the film in-situ in pure oxygen at 10 mbar for 10 min.


An important reason for controlling the oxidation parameters in materials that would otherwise inevitably grow native oxide upon exposure to the atmosphere is reproducibility. When exposing a pure metal film to ambient air, conditions such as temperature, pressure, moisture and the presence of contaminants will affect the kinetics of oxidation, and may also influence the oxide thickness and composition. Such discrepancies may influence not only device performance, but also processing parameters such as the etch time during patterning, as metals and metal oxides can have different etch rates [140].

In paper **B**, we compare the materials properties of films oxidized in-situ with films oxidized in ambient air. We also compare the quality of CPW resonators produced from two different wafers - one prepared using the standard process including in-situ oxidation (sample #1 in Section 6.2), and a sample that differs from the standard sample by only one fabrication step - skipping the post-deposition in-situ oxidation of the Al film.

We find that skipping the oxidation step increases the average TLS loss of the 8 resonators on the sample only marginally: the  $F\delta_{TLS}^0$  of the sample oxidized in atmospheric conditions is  $9.7 \pm 2.7 \times 10^{-7}$ , while the value for the standard sample with in-situ oxidation is  $8.7 \pm 2.9 \times 10^{-7}$ .

However, we do find differences in the microscopic properties of the two films. The surface oxide thickness is  $\sim 5$  nm for both films; however, a scanning electron micrograph (SEM) in Fig. 7.1a reveals unexpected features in the film oxidized in the ambient air, as well as on the surface of the underlying Si substrate after the patterned Al has been etched away.

We find a high presence of voids in the Al film oxidized in ambient air. The voids show as dark spots in the Al-covered areas in Fig. 7.1a, as confirmed by transmission electron



**Figure 7.1:** Materials analysis via SEM, AFM, TEM and EDS of a sample where the Al film is oxidized in ambient air instead of our standard in-situ oxidation in pure oxygen. **a** The SEM shows voids in the Al film (dark spots in the light gray area), and residues on the underlying Si where the Al has been etched away (light spots in the dark gray area). **b** AFM showing the residues in an area of Si where the Al has been etched away. **c** TEM cross-section of the Al/Si material stack, showing a void in the Al film. The Pt film was deposited prior to the TEM imaging in order to protect the structure during sample preparation. **d** EDS composition map of a cross-section of an area where a residue flake remains after the etching of Al on Si. The signals of Al, O, and Si are plotted in yellow, blue, and cyan, respectively. **e** The EDS cross-sectional profiles extracted from the area indicated by the dashed-line square in **d**, plotted in the direction shown by the arrow.

microscopy (TEM) in panel **c**. To investigate the origin of these voids, we prepared a second film where we omitted the in-situ oxidation. At first, we did not see any defects in the film. We heated this sample to 160 °C for 5 min to create similar conditions to that of baking steps during resonator patterning, which did not produce voids immediately; however, monitoring this sample over time, we noticed the appearance of the voids after about two days.

After patterning the Al film with CPW structures, we found residues on the exposed Si surface where the Al is etched away (Fig. 7.1a), even for films that were never heated after deposition. These residues appeared insoluble in Al etchant, even if exposed to the

etchant for extensive periods of time. The atomic force micrograph (AFM) in Fig. 7.1b places the lateral size of these features in the range of few tens of nanometers up to 200 nm, and the thickness to about 2 nm. The compositional distribution of one of the residual features obtained via energy-dispersive x-ray spectroscopy (EDS) in Fig. 7.1d–e shows intermixing of Al, Si, and O in the features, which can explain the insolubility in Al etchant.

Due to the thermal expansion mismatch between Al and Si, changes in temperature can induce stress in the Al film ( $\alpha_{\text{Al}} = 23.1 \times 10^{-6} \text{ }^\circ\text{C}^{-1}$  and  $\alpha_{\text{Si}} = 2.6 \times 10^{-6} \text{ }^\circ\text{C}^{-1}$ ), which can promote recrystallization in the polycrystalline film. Recrystallization under stress is known to lead to the formation of hillocks and voids in pure Al films grown on Si [130]. We suggest that the structural differences in the Al oxide formed under ambient conditions and in pure oxygen introduce different levels of stress into the underlying Al film, which can promote the formation of defects in the film oxidized in ambient air.

As the TLS loss of the resonators fabricated from the film without in-situ oxidation is only marginally higher than of the standard, one might conclude that the in-situ oxidation is not a necessary step for device performance. However, the defects formed in the absence of the in-situ oxidation steps may result in random and unwanted open circuits in the device, decreasing the fabrication yield.

### 7.1.2 In-situ capping

While controlling the oxide growth and mitigating its negative impact on the device quality is helpful, the ideal strategy would be to prevent the oxide growth altogether. In metals that inevitably oxidize in ambient air, a solution can be to encapsulate (or *cap*) the surfaces that will be exposed to air by a material that is less prone to oxidation. This can be done in-situ inside the deposition chamber without breaking the vacuum by depositing the superconducting film, and then immediately depositing the capping material layer before the superconductor has a chance to oxidize. Ideally, the capping material would resist oxidation entirely, but improved coherence may be achieved by capping the superconductor with a metal that grows a thinner oxide, or an oxide with a lower dielectric loss tangent [141].

It should be noted that while the in-situ passivation of a film prior to patterning is an interesting tool in studying dielectric loss in our circuits, it is only a partial solution to the MA interface loss of the CPWs fabricated using our traditional approach. In an etch-based patterning process, if the film is deposited first, then immediately passivated, and only then a pattern is etched into the film, only the *top* surface of the CPW structures will be protected from oxidation – the sidewalls of the CPW will be oxidized as usual.



Given that the majority the electric field is concentrated in the gaps of the CPW, the top-surface passivation only partially protects the critical structures from oxidation.

Due to the proximity effect in superconductors, the capping material does not necessarily need to be a superconductor in order to avoid resistive losses at the capped interface [142]. If a thin layer of normal metal (thinner than the coherence length of the superconducting film) is placed in proximity to a superconductor, the superconducting wave function will extend to the normal metal, facilitating lossless DC transport. This, however, comes at the expense of a reduced  $T_c$ .

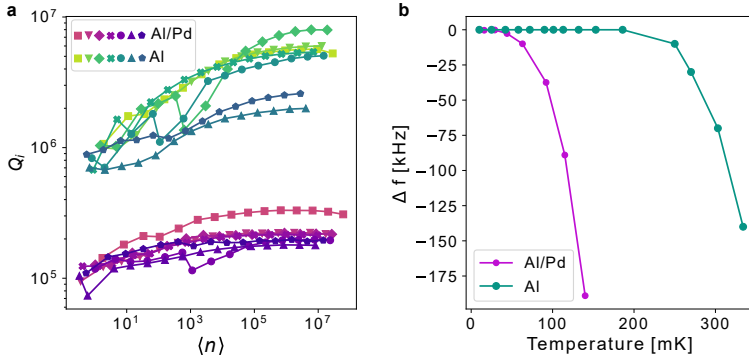
In practical applications, we need to consider not only the coherence lengths of the materials, but also the surface roughness of the bottom film – the capping layer has to be thick enough to form a continuous film without forming pinholes [143]. Thus, we find a maximum capping film thickness governed by the coherence length of the superconductor, and a minimum thickness governed by the superconductor’s roughness.

The superconducting coherence length of Al is relatively long at 1600 nm[144]. However, the capping material possibilities are compromised by the tendency of Al to form intermetallic compounds with many other metals. The formation of these compounds can cause anomalous effects in  $T_c$  [145].

### Palladium capping

For example, while it is possible to cap Nb with Pd [146], Al capped with Pd will form an intermetallic alloy [147]. In an experiment shown in Fig 7.2, I set out to see whether capping Al with a more oxidation-resistant Pd would lower the TLS loss at the MA interface. I compare resonators of our standard CPW design fabricated on two films: a 150 nm thick evaporated Al film, and another 150 nm thick Al film evaporated the same way and immediately followed by a 10 nm Pd deposition.

The Pd layer thickness was chosen to be the thinnest found continuous layer. I determined the continuity of the capping layer by dipping the wafer in the standard Al etchant (Sec. 4.6.1), which etches Al but not Pd. If the film corroded (as it did at 5 nm of Pd), it was a sign that the film was not continuous. At 10 nm Pd thickness, there was no corrosion. During resonator patterning, I milled the Pd away with oxygen plasma, and then etched the Al as usual. In Fig 7.2a, we find that any potential benefits of decreasing TLS loss this way are outweighed by an increase in other, non-TLS-related losses, which have a profound negative effect on the resonator  $Q_i$  over the entirety of the measured input power span. Measuring the resonator properties as a function of temperature, we see that the frequency of the Al/Pd resonators begins shifting at a significantly lower temperature than the uncapped Al (see Fig 7.2b). The frequency



**Figure 7.2:** Properties of resonators with frequencies 4–8 GHz fabricated on a 150 nm thick film of evaporated Al (shades of green), and a comparable Al film capped in-situ by 10 nm of Pd (shades of purple). **a** Resonator  $Q_i$  as a function of drive power for 7 resonators each present on one chip with an Al film, and one chip with a Al/Pd film, respectively. **b** The frequency shift of one resonator out of each of the two Al and Al/Pd samples plotted as a function of the temperature of the cryostat stage where the sample is mounted. The frequency shift is calculated with respect to the frequency at 10 mK

shifts are associated with an increase of quasiparticle density, and their occurrence at lower temperatures indicates a lower  $T_c$  (see Sec. 2.3). Therefore, the deterioration of the Al/Pd resonator performance could be explained by the anomalous effects caused by the Al/Pd intermetallization.

### Aluminium fluoride capping

In another experiment, we explore the potential of passivating Al with aluminium fluoride ( $\text{AlF}_3$ ), which forms on Al when exposed to fluorine-based plasma in dry-etching processes [148]. This formation of  $\text{AlF}_3$  makes Al resistant to fluorine-based gas etching [149]. Unlike efforts in capping the superconductor with oxidation-resistant metals, we do expect  $\text{AlF}_3$  to contribute to dielectric loss. We aim to compare whether the differences between the thicknesses and dielectric loss tangents of the passivating  $\text{AlF}_3$  and the native  $\text{AlO}_x$  amount to differences in TLS loss of CPW resonators.

In the fabrication process for this test, we first cap the Al in-situ with a Nb film. Nb is chosen based on availability in our deposition system, as well as on findings in literature demonstrating how Nb and Al are a suitable material combination for quantum circuits in Nb-based trilayer junctions [150]. In our process, the Nb serves as a temporary

protective layer, and it will oxidize – however, it will prevent the underlying Al from oxidation.

After patterning the resonators into an Al/Nb material stack, we etch the protective Nb away using  $\text{NF}_3$  plasma, intending to leave behind an  $\text{AlF}_3$ -passivated Al.

At the time of writing, the results of these tests are inconclusive. The resonators fabricated from a wafer using the process outlined above show a comparable performance to the standard samples. We do detect traces of  $\text{AlF}_3$  in the passivated sample in ToF-SIMS; however, we also find a native oxide layer in a TEM cross-section.

Further tests will be needed to ascertain whether the  $\text{AlF}_3$  passivation was incomplete and thus insufficient in preventing oxidation, or whether this passivation is temporary, leading to an inevitable native oxide growth over time.

## 7.2 Post-fabrication treatment with HF and HMDS

As I discussed in the previous sections, after the conclusion of the qubit and resonator fabrication processes, the device surfaces are covered in amorphous native oxide together with potential organic residues from the various fabrication steps, which is a source of dielectric loss in these devices. Other contamination may also occur during fabrication.

In this section, I explore the impact on the quality of our CPW resonators of stripping surface oxides with an HF dip of the chips as the last step of the fabrication process. The potential benefits are two-fold: removal of the surface oxide, and together with it, any contaminants accumulated on top or encased within this oxide. On Si substrates, the SA interface can also be temporarily passivated by an Si–H monolayer after the HF treatment (see Sec. 4.2).

While the removal of contaminants may have a long-lasting effect, the native oxide will regrow over time. Therefore, in our investigations of the effect of native oxide removal, we begin the chip packaging procedure immediately after the HF treatment, after which we cool the chips down. In practice, this translates to about  $\sim 2.5$  h between when the chip is treated in HF, and the start of the vacuum pumpdown of the cryostat.

A similar approach has been successfully demonstrated to improve the performance of resonators fabricated using Nb-on-Si and Ta-on-Si [106, 123, 151]. However, this process is less straightforward for Al-based devices, as the exposed Al itself reacts with HF after its oxide has been consumed. This can result in damages in the metal film, and a degradation of the device quality [152]. Therefore, we optimize the process to avoid Al film degradation.

To extend the passivation at the SA interface, we test following the HF dip by a

vapour deposition of hexamethyldisilazane (HMDS) on a hot plate heated to 100°C. HMDS is typically used to enhance the adhesion of photoresist to Si substrates. It can be deposited by spin-coating, in which case it forms an amorphous coat. However, when deposited as a vapour, HMDS can bind to SiO<sub>x</sub> and create a hydrophobic monolayer, which could be expected to protect the Si from re-oxidizing [111, 153].

We experiment with this post-fabrication cleaning process on CPW resonators fabricated on two material platforms: Al on Si, and Nb on sapphire. Some of the data presented in this section is published in Paper **B**.

### 7.2.1 Al on Si

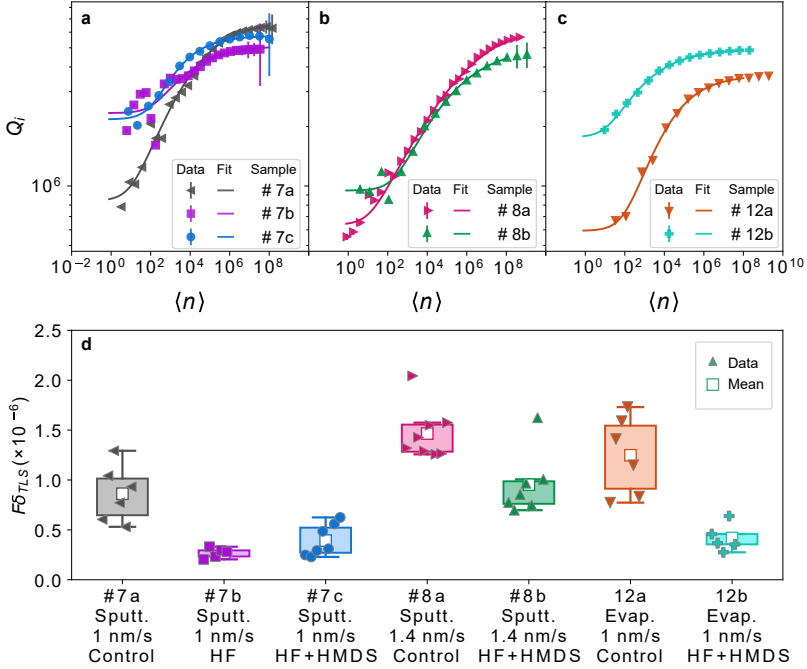
In the Al-on-Si material platform, both the Al and Si oxides (AlO<sub>x</sub>, SiO<sub>x</sub>) can be removed by HF. Removing SiO<sub>x</sub> without damaging the underlying Si is relatively straightforward: SiO<sub>x</sub> is removed in an HF solution, but the underlying Si becomes terminated by an Si-H monolayer and does not degrade further (see Sec. 4.2). In our standard qubit and resonator fabrication process, we expose the bare Si wafer to a 2% HF solution for 60s prior to the Al deposition, although 30s has also been shown to be sufficient in literature [123].

Removing AlO<sub>x</sub> is less straightforward, as HF will start etching Al once the surficial AlO<sub>x</sub> has been removed. We found that on our previously standard 150nm thick films, a 30s HF dip was already sufficient to damage the Al film. Therefore, we submerge the chips for 15s only before rinsing the HF off in deionized water and proceeding to chip packaging.

We measure CPW resonator fabricated on films from three different wafers. Wafers #7 and #8 have Al films sputtered at 1 nm s<sup>-1</sup> and 1.4 nm s<sup>-1</sup>, as discussed in Sec. 6.2. The Al film on wafer #12 was evaporated at the standard rate of 1 nm s<sup>-1</sup>. From each of these wafers, we measure resonators on one chip without any post-fabrication treatment (denoted with the index “a” following the respective wafer number), and at least one chip with either post-fabrication treatment (HF or HF + HMDS, denoted by “b” and “c”). The results of the resonator analysis are summarized in Table 7.1 and Fig. 7.3.

Apart from the deposition conditions outlined above, the films are fabricated using the standard process (see Sec. 6.2) with one exception – the SC1 step is omitted.

Comparing the HF-treated chip #7b to the control #7a, we find that the average low-power  $Q_i$  of resonators with the HF treatment doubles. The subsequent HMDS deposition degrades the  $Q_i$  of resonators on chip #7c slightly, but the dielectric loss of this sample is still well below half that of the control sample. This degradation is presumably at least partially due to the heated treatment which can accelerate reox-



**Figure 7.3:** Resonator loss as fitted to the TLS model (Eq. 2.4) for the samples with process variations listed in Table 7.1. **a–c**  $Q_i$  vs.  $\langle n \rangle$  for one resonator each from samples fabricated in a sputtered film deposited at a standard rate of  $1 \text{ nm s}^{-1}$ , a film sputtered at  $1.4 \text{ nm s}^{-1}$ , and a film evaporated at the standard rate of  $1 \text{ nm s}^{-1}$ , respectively. The resonance frequency is  $\sim 4.5 \text{ GHz}$ . **d** Box plot of TLS loss from the fits to the TLS loss model for all of the resonators on each sample. Each data point represents the loss of an individual resonator.

idation of Al and Si. Further studies into the longevity of the HF+HMDS passivation treatment are needed, together with material characterization to ascertain its nature. HMDS is expected to bond as a monolayer to SiOx, and therefore a study into the microscopic morphology of the Si surface after the short HF dip, as well as at the start of the passivation step, is important in optimizing this step in terms of timing and deposition temperature.

Samples from wafers #8 and #12 likewise see an improvement in low-power  $Q_i$  and TLS after the HF+HMDS treatment, compared to the untreated samples. The otherwise-standard sample #12 shows a slightly subpar resonator performance compared to our standard devices (see Sec. 6.2). I attribute this to persistent lithographic residues

|                                   | <b>Sample 7a</b><br>(#7a) | <b>Sample 7b</b><br>(#7b) | <b>Sample 7c</b><br>(#7c) |  |
|-----------------------------------|---------------------------|---------------------------|---------------------------|--|
| Process variation                 | Sputter                   | Sputter                   | Sputter                   |  |
|                                   | $1 \text{ nm s}^{-1}$     | $1 \text{ nm s}^{-1}$     | $1 \text{ nm s}^{-1}$     |  |
|                                   | Reference                 | HF                        | HF + HMDS                 |  |
| No. of resonators                 | 6                         | 5                         | 7                         |  |
| $Q_i (\langle n \rangle \sim 10)$ | $1.0 \times 10^6$         | $2.2 \times 10^6$         | $2.1 \times 10^6$         |  |
| $F\delta_{\text{TLS}}^0/10^{-6}$  | 0.86                      | 0.27                      | 0.39                      |  |
|                                   | ( $\pm 0.29$ )            | ( $\pm 0.05$ )            | ( $\pm 0.16$ )            |  |
| $\delta_{\text{other}}/10^{-7}$   | 1.5                       | 2.0                       | 1.7                       |  |
|                                   | ( $\pm 0.24$ )            | ( $\pm 0.37$ )            | ( $\pm 0.49$ )            |  |

|                                   | <b>Sample 8a</b><br>(#8a) | <b>Sample 8b</b><br>(#8b) | <b>Sample 12a</b><br>(#12a) | <b>Sample 12b</b><br>(#12b) |
|-----------------------------------|---------------------------|---------------------------|-----------------------------|-----------------------------|
| Process variation                 | Sputtered                 | Sputtered                 | Evaporated                  | Evaporated                  |
|                                   | $1.4 \text{ nm s}^{-1}$   | $1.4 \text{ nm s}^{-1}$   | $1 \text{ nm s}^{-1}$       | $1 \text{ nm s}^{-1}$       |
|                                   | Reference                 | HF + HMDS                 | Reference                   | HF + HMDS                   |
| No. of resonators                 | 8                         | 7                         | 6                           | 5                           |
| $Q_i (\langle n \rangle \sim 10)$ | $0.58 \times 10^6$        | $0.91 \times 10^6$        | $0.72 \times 10^6$          | $1.5 \times 10^6$           |
| $F\delta_{\text{TLS}}^0/10^{-6}$  | 1.5                       | 0.95                      | 1.3                         | 0.42                        |
|                                   | ( $\pm 0.26$ )            | ( $\pm 0.32$ )            | ( $\pm 0.40$ )              | ( $\pm 0.14$ )              |
| $\delta_{\text{other}}/10^{-7}$   | 1.3                       | 2.2                       | 6.2                         | 3.5                         |
|                                   | ( $\pm 0.38$ )            | ( $\pm 1.2$ )             | ( $\pm 5.4$ )               | ( $\pm 1.5$ )               |

**Table 7.1:** List of samples in the post-fabrication HF dip and HMDS deposition study. Samples with the same numerical index are from the same wafer. Samples with the numerical index **a** are reference samples with no post-fabrication HF or HMDS treatment. Samples with indices **b–c** have received the HF or HF + HMDS treatment indicated in the table. The table includes the number of resonators on each sample, the average  $Q_i$  at  $\langle n \rangle \sim 10$  across all the resonators on the sample, as well as the average  $F\delta_{\text{TLS}}^0$  and  $\delta_{\text{other}}$  for the TLS model fit. Standard deviations are provided for the parameters of the TLS model.


that appeared during processing and were not removed by our gentle cleaning steps. Likewise, the performance of chips from this wafer improved with the discussed post-fabrication treatment.

## 7.2.2 Nb on sapphire

When treating sapphire substrates with HF, we do not expect the same passivation benefits as with Si, but we can remove contaminants, as well as the metal oxide.

In Fig. 7.4, I show resonators fabricated from high crystallographic quality Nb with a thickness of 60 nm, grown via molecular beam epitaxy on sapphire. The residual resistance ratio of this film is 37, and  $T_c = 9.2$  K. The films are etched with an  $\text{NF}_3$  plasma, and the resonator design is identical to those on our standard Al films. We investigate the effect of different HF treatment times, as well as the effect of an HMDS treatment. The samples are measured in two different cooldowns – the resonators shown in panel Fig. 7.4b were allowed to age inside a dry box with a nitrogen atmosphere for about two months longer than the resonators in panel Fig. 7.4a.

The control sample in Fig. 7.4b, which has not undergone any HF or HMDS steps, shows a single-photon  $Q_i$  between  $3\text{--}4 \times 10^5$ , which is relatively low for the state of the art quality that can be achieved in the same geometry using Nb-on-Si [106]. A post-fabrication dip in a 2% aqueous solution of HF for 2 min increases the total quality of the resonator, with the low-power  $Q_i$  between  $1\text{--}2 \times 10^6$ , which is slightly lower than our HF-treated Al resonators. The HMDS-treated resonators deteriorate – this is because we do not expect any benefits from the HMDS–sapphire interaction as opposed to the



**Figure 7.4:** Resonator quality of CPW resonators with  $f_r \sim 5.1$  GHz, made from Nb on sapphire. Each data set belongs to a different chip, with the post-fabrication treatment noted in the legend. All chips are from the same wafer. The resonators shown in panel b were aged in a nitrogen atmosphere for two months after the resonator patterning and chip dicing was completed.

interaction with Si; on the other hand, the heated HMDS vapor process may increase the kinetics of the Nb reoxidation.

When aged for two months, the  $Q_i$  of the control resonators degrades slightly. A treatment in 2% HF for 10 min does not increase the quality of these resonators further, which may indicate that such prolonged treatment incurs damage to either the substrate or the metal surface on the Nb-on-sapphire material platform, negating the cleaning and oxide-stripping benefits of this processing step.

### 7.3 3D cavities

In Paper **D**, we investigate the performance of 3D aluminium cavities as a function of metal purity, cavity design, resonance frequency, and surface treatment. Due to the significantly reduced dielectric participation ratios in 3D cavities as opposed to 2D waveguides, we also observed significantly higher  $Q_i$  compared to our CPW resonators.

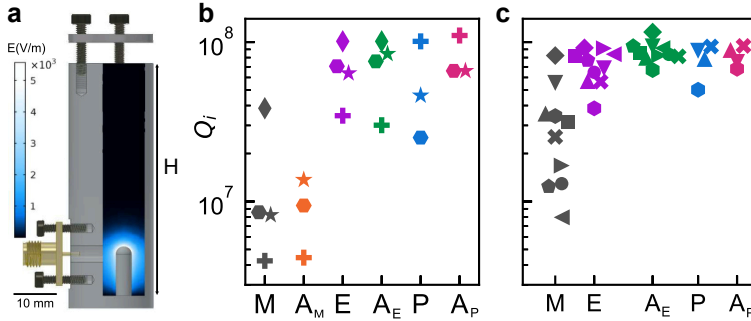
Studying the impact of frequency and design, we observe only a slight trend of higher  $Q_i$  at lower frequencies, which may be attributed to a lower density of TLS at these frequencies. We also find that cavities with a taller cylindrical waveguide section above the coaxial pin 35  $\mu\text{m}$  (see Fig. 7.5a) perform better than the shorter cavities. This we attribute to a reduction of seam loss at the top of the cavity – the intensity of the electric field falls exponentially from the maximum at the top of the coaxial post, and cavities with a top seam sufficiently far away from the post are not limited by seam loss.

We find that the purity of the Al metal has an impact on the cavity  $Q_i$ : cavities made from the 6081 alloy (96%-98% Al with traces of mostly Mg, Si and Mn) achieve  $Q_i$  in the range of  $10^6 - 10^7$ , which is relatively low for our 3D architecture. Cavities made from the 4N Al (99.99% purity) perform better than the 6081 alloy, with  $Q_i$  in the  $10^7 - 10^8$  range. On average, 5N Al (99.999% purity) cavities perform even better than 4N.

We measure these cavities as-machined, and after various material treatments – etching, annealing, and electrochemical polishing. The results are shown in Fig. 7.5b–c.

It has been observed that etching 100–200  $\mu\text{m}$  of material is instrumental for achieving millisecond coherence times in Nb and Al cavities [154, 155]. In our etching process, we immerse the cavities in a beaker of the standard aluminium etchant (see Section 4.6.1), heated to 50 °C, for 2 h. We repeat this process twice (totaling 4 h of etch time), refreshing the etchant between the two etching rounds to maintain a constant etch rate without saturating the etchant. The etch rate at 50 °C is about 10  $\text{nm s}^{-1}$ , therefore we expect to etch about 150  $\mu\text{m}$  of Al during this process [156].





**Figure 7.5:** Design and  $Q_i$  of cavities that have undergone different treatments, measured at the single-photon power level. **a** The stub-geometry cavity design and the simulated electric field amplitude. **b–c** The cavities are measured after each step of a sequence consisting of machining (M), post-machining annealing ( $A_M$ ), etching (E), post-etch annealing ( $A_E$ ), electropolishing (P), and post-polishing annealing ( $A_P$ ). Data plotted with identical markers belong to the same cavity. The cavities in **b** were machined from Al with a 4N purity, while in **c** the material purity was 5N.

Due to the relatively low melting point of Al at  $660^\circ\text{C}$ , it is possible to anneal this metal at  $500^\circ\text{C}$  to promote recrystallization, which stimulates the migration of atoms in the crystal lattice and a subsequent reduction of defects [157]. To discourage excessive surface oxidation at high temperature, we anneal the cavities in a nitrogen atmosphere. The annealing furnace ramps up from room temperature to  $500^\circ\text{C}$  at a rate of  $5^\circ\text{C min}^{-1}$ , which is limited by the furnace specifications. Then, we anneal the cavity for 3 h. Afterwards, the cooldown from  $500^\circ\text{C}$  to room temperature takes about 4–5 h, which favours the relaxation of defects in the Al lattice [158].

We also investigate the potential of electrochemical polishing (or *electropolishing*) of cavities in reducing the surface roughness of the cavities, and therefore TLS loss at the surface. The polishing process, developed in [159], is done in a solution of phosphoric and sulfuric acids with a 60:40 ratio using graphite as a counter-electrode. The polishing cycle begins by a voltage sweep into the diffusion-limited polishing regime, followed by a constant voltage step lasting 20 min. This cycle is repeated three times for each polished cavity.

We find that electrochemical polishing can markedly improve the  $Q_i$  of short cavities in which we can achieve a better flow of reagents. However, in taller cavities with worse flow,  $Q_i$  is unaffected or can even degrade.

Our 5N cavities, just after machining and solvent-cleaning the cavities in acetone and

IPA, achieve  $Q_i$  between  $10^7 - 10^8$ . A subsequent etching treatment reduces the spread of the scattered  $Q_i$  values, increasing the average value. A post-etch annealing reduces the spread further, with the mean  $Q_i$  tending towards the higher values. Subsequent treatment by electropolishing and annealing does not further improve the performance of these cavities.

We find that by machining cavities from 5N Al, followed by etching and annealing, we can reproducibly achieve  $Q_i$  over  $80 \times 10^6$ .

## 7.4 Flip-chip

To reduce the physical footprint as well as to facilitate the routing of the input/output signals of our quantum processors in order to increase their scalability, we integrate our devices into a flip-chip architecture, which we present in Paper **E**.

In our chosen approach, we separate the quantum processor into two modules – a qubit chip and a control chip. The qubit chip contains all qubits and couplers, while the control chip contains the input/output transmission line, the readout resonators, and the control lines for the qubits and couplers. The qubit and control chips are fabricated on separate wafers and then bonded facing each other through so-called bump bonds at a set separation designed to achieve the desired coupling between the elements on the two respective chips.

We use superconducting indium (In) for the bump bonds. In is soft at room temperature, and therefore it is possible to bond chips by compressing the In bumps on the respective chips together.

As I discussed in Section 7.1, Al has a tendency to intermetallize with many other metals, which includes indium. Therefore, we deposit an under-bump metallization layer from superconducting NbN between the Al ground plane and the indium bumps to act as a diffusion barrier.

In the fabrication process flow, we first deposit the Al ground plane, after which we deposit the under-bump metallization via lift-off. Afterwards, we pattern the Al, and deposit Josephson junctions on the qubit chip. Finally,  $8\mu\text{m}$  tall indium bumps are fabricated on each wafer via evaporation and lift-off, after which the wafers are diced into chips. The control and qubit chips are then bonded by compression at room temperature with the target chip separation of  $8\mu\text{m}$ .

While the flip-chip modules are designed with great care to reduce potential losses introduced by the proximity of the added elements to critical device structures compared to our standard planar chips, it is still possible for dielectric losses to be exacerbated

during the added processing steps, e.g. lithographic residues or excessive oxidation during heating steps. These potential losses would likely affect primarily the MA or SA interfaces, although heating and added stress in the film could also influence the film structure and the SM interface. In Paper **E**, we find the performance of the qubits integrated into the flip-chip modules to be comparable to our planar devices – therefore, we conclude that the added processing steps and elements do not degrade the device performance when integrated into our flip-chip architecture.



---

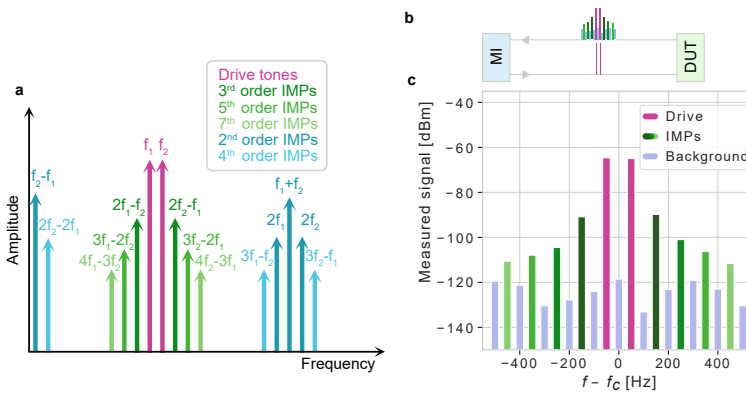
## Intermodulation spectroscopy of TLSs

---

In the previous chapters, I have identified parasitic TLSs as a major source of decoherence in our devices, with the focus on investigating their microscopic origin and developing strategies to mitigate them through device design and fabrication. I have also discussed how the presence of these TLSs in our devices gives rise to nonlinearly power-dependent internal quality factors  $Q_i$  with a TLS-limited performance at qubit-operating conditions. In the following experiments, published in Paper C, we demonstrate and investigate the inherent nonlinear nature of TLSs through intermodulation spectroscopy of a CPW resonator.

A hallmark feature of nonlinear elements is that when driven by multiple frequencies, the nonlinearity gives rise to frequency mixing (also referred to as *intermodulation*). Driven nonlinear media thus generate so-called intermodulation products (IMPs) – signals at frequencies not contained in the original drive tones, but rather at linear combinations of the incident drive frequencies ( $f_{\text{IMP}} = \sum_i k_i f_i$ , where  $f_i$  are drive frequencies and  $k_i$  are integers). I illustrate the strongest IMPs generated from a two-tone drive in Fig. 8.1a.

All experiments in this chapter are performed using a CPW resonator with a frequency of 4.11 GHz. Conventionally, CPW resonators are treated as fully linear elements and modeled as lumped-element  $LC$  harmonic oscillators. The resonator used in



**Figure 8.1:** Intermodulation in nonlinear media. **a** Illustration of the strongest IMPs generated by frequency mixing of two drive tones  $f_1$  and  $f_2$ , which are depicted in pink. Odd-order IMPs (i.e. when  $|k_1| + |k_2|$  is an odd number) are shown in shades of green, and even-order IMPs in shades of blue. **b** Measurement schematic: The measurement instrument (MI) outputs two drive tones and measures the response of the device under test (DUT). **c** IMPs measured when two slightly detuned drive tones are applied at resonance. The drive power corresponds to  $\langle n \rangle = 10^3$ . The  $x$ -axis shows the detection comb frequencies  $f$  relative to the centre of the detection comb  $f_c = (f_1 + f_2)/2$ . IMPs are detected at frequencies fulfilling the  $f_{\text{IMP}} = |k_1|f_1 + |k_2|f_2$  condition, which falls on every other tone in this detection comb. The remaining comb tones sample the microwave background.

these experiments is coupled only to its input-output transmission line and there are no engineered nonlinear elements present on the chip; therefore, nonlinear behaviour is not expected from this resonator.

## 8.1 Intermodulation spectroscopy

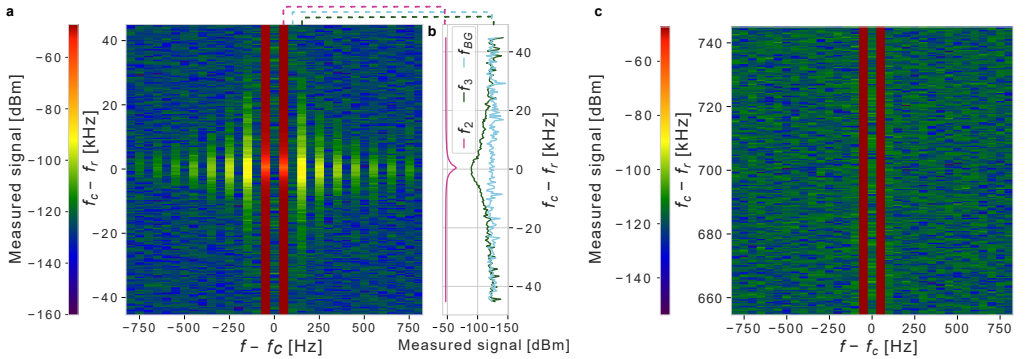
For the intermodulation spectroscopy experiments, we use a multi-frequency lock-in amplifier to generate two drive tones separated by a comparatively small spacing  $\Delta = 100$  Hz with respect to the resonance linewidth. The power-dependent  $Q_i$  (see Fig.8.3a) translates to a power-dependent linewidth, which is in the range of 8-11 kHz for this resonator. Setting the tone spacing  $\Delta$  much smaller than the linewidth ensures that both drives, as well as numerous lowest-order IMPs, can fall well within the resonance linewidth.

We measure the response of the device to the two drive tones using a detection

frequency comb demodulating the in-phase and quadrature components. We define the detection frequency comb such that every other frequency of the comb coincides with a frequency  $f_{IMP}$  where an odd-order IMP is expected, and the remaining comb tones sample the microwave background, see Fig. 8.1b. This way we not only collect information about the background, but also confirm that the signal observed at  $f_{IMP}$  is indeed the result of intermodulation. We are not able to observe any even-order IMPs, as their frequencies would fall far outside the resonance linewidth.

The generation of the drive tones and the digitization of the response are performed in the second Nyquist zone, without the use of analog mixers for frequency conversion.

We perform IMP spectroscopy by sweeping the frequency comb with the two drives in the center, defined in Fig. 8.1c, across the resonance (see Fig. 8.2). We detect no intermodulation when the two drive tones are far detuned from the resonance frequency  $f_r$ , outside the resonance linewidth. Approaching  $f_r$ , IMPs start to appear at the comb tones corresponding to  $f_{IMP}$ , with a peak when  $f_c = f_r$ . The observation of IMPs only when the drive tones are near-resonant with the device – verified in another, far-



**Figure 8.2:** Intermodulation spectroscopy. The  $x$ -axis in **a** and **c** shows the detection comb tones relative to the comb center frequency  $f_c$ , set up in an identical fashion to that shown in Fig. 8.1c – every other tone of the comb corresponds to an  $f_{IMP}$  with the two drive tones in the center of the comb. **a** Sweeping  $f_c$  across the resonance on the  $y$ -axis, we observe IMPs at the expected comb frequencies as we approach  $f_c = f_r$ . **b** Vertical line-cuts of the data shown in **a**: the device resonance demonstrates as a dip in the transmission of drive tone  $f_2$ , the third IMP  $f_3$  presents as a peak at  $2f_2 - f_1$ , and the background noise is shown at a frequency  $f_{BG}$  that does not correspond to any IMP. **c** The drive tones are off-resonant by  $\sim 700$  kHz – so far outside the resonance linewidth that the drives are rejected by the resonator, and no IMPs are generated.

detuned spectroscopy in Fig. 8.2c – confirms that the observed nonlinearity is native to the device, and is not an artefact of nonlinearities present elsewhere in the measurement set-up (e.g. the amplification chain).

## 8.2 Power dependence

We perform intermodulation spectroscopy as a function of applied drive power and compare the acquired data with the standard TLS analysis performed with a VNA. In Fig. 8.3c, we plot the measured power of the lowest-order IMPs when  $f_c = f_r$  as a function of applied drive power. The drive powers used in this measurement exceed the critical photon number  $n_c$  (see Fig. 8.3a). We detect IMPs well above the noise floor in the entirety of the measurement power span and identify three regions with differing slopes of this power dependence on the logarithmic scale.


In the low-power regime where  $\langle n \rangle \lesssim 200$ , TLSs dominate the total loss of the resonator, and the dependence of the IMPs on power is strong with a variable slope. In the intermediate drive power regime where  $\langle n \rangle \gg n_c$ , the four strongest IMPs show a power-law dependence on applied drive power. We fit the power dependence in this regime to  $\langle n \rangle^m$  and summarize the fit parameters  $m$  for the different orders of IMPs in Table 8.1. It is notable that the IMP power increases more subtly than the drive power (i.e.  $m < 1$  for our device, while  $m = 1$  would be typical for intermodulation distortion in electronic devices [160]). This weaker power dependence can be attributed to the saturable character of the TLS bath.

At drive powers exceeding  $5 \times 10^5$  photons, the slope of the nonlinearity changes abruptly, which relates to the onset of kinetic inductance effects at high input power due to current-induced Cooper pair breaking [161, 162]. This increase in kinetic inductance also demonstrates in  $f_r$  shifts and  $Q_i$  degradation in the high-power regime, see Fig. 8.3a–b. This high-power behaviour is not specific to the resonator used in our IMP experiments, but is common for all CPW resonators studied in the previous chapters. As this kinetic inductance-induced behaviour is not a subject of the studies discussed in chapters 6–7, this high-power data is typically disregarded in our analyses.

## 8.3 Temporal stability

The coupling between a quantum circuit and the TLS bath is known to fluctuate on slow time scales, which results in fluctuations of device quality and frequency over time [12, 47, 55]. Tracing these device parameter fluctuations to their underlying mechanisms





**Figure 8.3:** TLS characterization of the same device using different methods. **a**  $Q_i$  as a function of  $\langle n \rangle$  extracted from a VNA measurement, fitted to the standard TLS model (Eq. 2.4). The data plotted in a lighter shade of green was disregarded in the fit. **b** Resonator frequency  $f_r$  as a function of applied power, likewise collected by a VNA. At low power the frequency is scattered due to the spectrally unstable TLSs. At high powers, coinciding with the degradation of  $Q_i$  in **a**,  $f_r$  drops due to increased kinetic inductance. **c** The power of the four strongest IMPs, as well as the background noise at a non-IMP frequency, plotted as a function of applied drive power when the comb center  $f_c$  is at resonance. The slope of the power dependence of the data plotted on a logarithmic scale is fitted for  $\langle n \rangle$  in the range of  $200\text{--}2 \times 10^5$  photons. **d** IMPs measured repeatedly for 25h in 5-minute intervals, at a drive power level of  $10^3$  photons.

such as the interaction with TLSs, or environmental effects like cryostat temperature or magnetic field fluctuations, or even cosmic ray incidence can require advanced analysis.

In Fig. 8.3d, we study the temporal stability of the measured IMPs by repeating an identical measurement over a span of 25 h. The measurement is performed with the frequency comb center on resonance, and the applied drive power corresponds to

$\langle n \rangle = 10^3$ . This power is chosen from the region of Fig 8.3a where the power is low enough that the TLS bath is not saturated, but high enough to achieve a favourable signal-to-noise ratio for several orders of IMPs. Each of the measurements plotted in Fig. 8.3d is averaged for 5 min.

Over the measurement time span, we observe fluctuations of varying amplitudes in the measured IMPs. Low-amplitude fluctuations are present over the entire time span, and at some time stamps, such as the 7–8 h mark, the amplitude jumps are significant. These amplitude jumps are correlated between the different IMP orders, but do not extend to the measured noise floor or the drive tones, which are stable in comparison. From this we conclude that the fluctuations are native to the nonlinearity in the device, and do not originate in the measurement setup from sources such as drifts in room-temperature electronics. As the IMPs are unique to the nonlinear mechanism producing them, this method enables us to study TLS fluctuations over time.

In the IMP power sweep in Fig. 8.3c, we also observed fluctuations of the measured IMPs. The temporal stability measurement in Fig. 8.3d suggests that the fluctuations in Fig. 8.3c may not necessarily have been a function of applied power (such as the saturation of a prominent TLS), but were likely at least partially caused by fluctuations in time.


## 8.4 Modeling and reconstruction

In this section, we model the IMPs of a driven harmonic oscillator with nonlinear damping modeled according to the standard tunneling model for TLSs. Simulating the IMP amplitude as a function of drive power, we achieve a qualitative match with the experimental data. We also reconstruct parameters of the standard TLS model by fitting a single intermodulation spectrum using harmonic balance analysis [163, 164].

We first solve the forward problem by numerically integrating the following equation of motion for a harmonic oscillator with nonlinear damping due TLS loss to obtain the time-dependent intracavity field:

$$\dot{\hat{a}} = -i\omega_r \hat{a} - \frac{\kappa_0 + \kappa_{ext}}{2} \hat{a} - \frac{\kappa_{TLS}(|a|)}{2} \hat{a} - \sqrt{\kappa_{ext}} a_{in}, \quad (8.1)$$

where  $\hat{a}$  and  $a_{in}$  are the time-dependent intracavity modes and the input drives, respectively,  $\omega_r$  is the resonant frequency,  $\kappa_0$  the internal linear loss rate,  $\kappa_{ext}$  the external coupling to the transmission line, and  $\kappa_{TLS}(|a|)$  the power-dependent loss due to TLSs.



**Figure 8.4:** Simulation of the power of the generated IMPs as a function of applied drive power. The parameters of the fit shown with dashed lines are shown in Table 8.1.

| IMP order | Experimental fit $m$ | Model fit $m$ |
|-----------|----------------------|---------------|
| 3         | $0.41 \pm 0.05$      | 0.42          |
| 5         | $0.38 \pm 0.04$      | 0.43          |
| 7         | $0.42 \pm 0.04$      | 0.44          |
| 9         | $0.48 \pm 0.05$      | 0.45          |

**Table 8.1:** Power dependence of the measured IMPs in Fig. 8.3c, and of the modeled IMPs in Fig. 8.4, fitted to  $\langle n \rangle^m \cdot 10^{y_0}$ , where  $y_0$  is the  $y$ -axis offset.

We consider a power-dependent damping of the standard form of Eq. (2.4), giving

$$\kappa_{TLS}(|a|)a = \frac{\kappa_{TLS}}{\left[1 + (|a|/a_c)^2\right]^\beta}, \quad (8.2)$$

where  $\kappa_{TLS} = f_r F \delta_{TLS}^0$  represents the TLS loss rate, and  $a_c^2 = n_c$  the critical photon number.

The simulation results, plotted in Fig 8.4, qualitatively match the experimentally observed behaviour in Fig. 8.3c. For strong drives  $a \gg a_c$ , the modelled IMPs scale with power as  $\langle n \rangle^m$ , where  $m$  matches well with the values obtained by fitting experimental data – for a side-by-side comparison of  $m$  obtained by fitting the experimental and modelled data, see Table 8.1.

Using an harmonic balance analysis, we can reconstruct the standard TLS model parameters  $F \delta_{TLS}^0$  and  $n_c$  from a single IMP spectrum taken at a relatively high drive power of  $\sim 10^3$  photons, such as the one in Fig. 8.1b. The reconstructed parameters are in good agreement with the parameters of the standard TLS model fitted to the  $Q_i$  vs.  $\langle n \rangle$  relation, see Table 8.2. Differences may be due to variable conditions across cooling cycles as the standard resonance sweeps and IMP data were obtained in separate cooldowns [165].

However, this method of analyzing a single IMP spectrum in its current state does not

| Parameters             | Standard model | Harmonic balance |
|------------------------|----------------|------------------|
| $f_r$ [GHz]            | 4.11           | 4.11             |
| $\kappa_{ext}$ [kHz]   | $6.8 \pm 0.1$  | $6.6 \pm 0.2$    |
| $\kappa_{TLS}$ [kHz]   | $4.0 \pm 0.3$  | $5.2 \pm 0.2$    |
| $a_c$ [ $\sqrt{n_c}$ ] | $1.4 \pm 0.4$  | $1.6 \pm 0.4$    |

**Table 8.2:** Reconstructed resonator parameters via fitting the standard tunnelling model Eq. (8.1) and the intermodulation spectral method with harmonic balance.

yield consistent  $\beta$  with the standard model. In the reconstruction in Table 8.2,  $\beta = 0.3$  is fixed to the value obtained from the standard analysis in Fig. 8.3a. Nonetheless, while a VNA power sweep is necessary to obtain  $\beta$ , we find that it is sufficient to perform that power sweep at high powers, which is fast due to a favourable signal-to-noise ratio. Since it is the measurements at single-photon power levels requiring heavy averaging that are necessary to extract  $F\delta_{TLS}^0$  in the standard method that are time consuming, our method allows for an efficient characterisation of loss in superconducting resonators compared to the traditional approach.

---

## Summary and Concluding Remarks

---

In these chapter, I will summarize the main results of the thesis, and provide concluding remarks on the significance of these results, as well as an outlook on what I would consider interesting future work that could be done to continue the line of research established throughout this thesis.

In the preceding chapters, I have summarized my work on studying the mechanisms of decoherence in superconducting quantum devices of various architectures – transmon qubits, coplanar waveguide resonators, and three-dimensional cavity resonators. The main focus of this work was the investigation of the microscopic nature of the typical performance-limiting factor of these devices – dielectric loss due to parasitic TLSs residing at the various materials interfaces – and strategies towards mitigating this decoherence channel.

Our energy participation ratio simulations show that our coplanar geometry is particularly sensitive to dielectric loss at the substrate-metal (SM) interface. Thus, typical sources of TLS loss such as amorphous interfacial oxides, dangling bonds, molecular oxygen, hydrogen, or organic adsorbates tend to exacerbate the loss rate of the device if present at this SM interface. Once the SM dielectric loss is mitigated, loss at the remaining materials interfaces (metal-air and substrate-air) is presumed to contribute to loss almost equally in our standard geometry – provided that TLS loss is still the

dominant loss mechanisms once the SM loss has been mitigated.

We see the effect of the sensitivity of the SM interface to variations in dielectric loss in our experiment with systematic variations of the fabrication procedure of CPW resonators. When we focus on varying one step of the SM interface formation at a time (substrate cleaning with SC1 and HF, or substrate annealing in vacuum), we find that subtle changes in this procedure can measurably affect the quality of the resulting devices in ways that are nontrivial to predict.

Unsurprisingly, we find that failing to strip the native oxide of the substrate wafer prior to the metal deposition has major repercussions on the device quality. However, TEM shows that following our standard fabrication procedure outlined in Section 6.2, including a Si native oxide removal in 2% HF prior to deposition, can result in a pristine Si/Al interface with no continuous interfacial oxide observed. Therefore, the subsequent ToF-SIMS observation of a distinct peak in the oxygen concentration at the SM interface of a Si/Al sample, fabricated using the standard process, came as a surprise. Due to the broad scanning area of the ToF-SIMS analysis, we hypothesized that this was due to oxide growth along the grain boundaries of the Al film. We do not expect such intergranular oxide to contribute to any additional losses within the metal film as the field is zero inside a superconductor; however, this oxide could contribute to dielectric loss in the vicinity of the materials interfaces.

To test this hypothesis, we want to find a way to increase the grain size of Al films without affecting the systemic parameters of the deposition procedure, as we know from the experiments in Section 6.3 that such variations can have complex consequences for device quality. Al films deposited in high vacuum tend to grow in columnar structures of parallel grain boundaries, and the grain size depends on film thickness, which we confirm is the case for our films via TEM. Therefore, increasing the film thickness also increases the grain size, decreasing the number of grain boundaries that are known diffusion channels and hosts of oxide.

We fabricate samples with Al films of three different thicknesses: 150 nm (our previous standard), 300 nm, and 500 nm. We confirm through TEM that the grain size scales with film thickness, and ToF-SIMS analysis shows that the oxygen content of the Al films, as well as the oxygen concentration of the SM interface, decreases in the thicker films. This is consistent with our hypothesis that the oxygen detected in ToF-SIMS originates at the grain boundaries.

Having found that growing thicker films can mitigate the presence of oxygen – a known suspect in TLS loss – at the SM interface in our films, we fabricate and measure CPW resonators with the three aforementioned film thicknesses. Analyzing the

---

power-dependence of the quality of these resonators, we find a systematic decrease in the characteristic TLS loss parameter  $F\delta_{TLS}^0$ , which halves when increasing the film thickness from 150 nm to 500 nm.

Expanding this study to a set of 42 qubits fabricated with Al films of the three thicknesses, we find that the average qubit quality factor increases from  $2.1 \times 10^6$  to  $3.3 \times 10^6$  across all devices when increasing the film thickness from 150 nm to  $\geq 300$  nm. The qubits in this study are all impacted differently by Purcell decay, which is a device design parameter that does not relate to the internal quality of the qubit; however, this Purcell decay can limit the observed  $T_1$  significantly. Therefore, we introduce the  $T_1/T_p$  parameter, which shows how closely a qubit's performance is limited by Purcell decay, when comparing different qubits. We find that if we constrain our data set to the less Purcell-limited qubits with  $T_1/T_p \leq 0.5$ , the average qubit  $Q$  of the 150 nm films remains constant; however, the average value for the thicker films rises to  $3.7 \times 10^6$ . Constraining the data set further to  $T_1/T_p \leq 0.25$  where the qubits should not be significantly limited by Purcell decay, the average  $Q$  for the thicker-film qubits becomes  $3.9 \times 10^6$ . For a qubit frequency of 3 GHz, an average  $Q$  of  $3.9 \times 10^6$  corresponds to a  $T_1$  of 200  $\mu$ s.

Therefore, we conclude that the presence of oxygen at the SM interface, where the energy participation ratio is typically high, is a source of TLS loss in our devices. We mitigate this loss by increasing the Al film thickness to at least 300 nm, which increases the grain size, thereby reducing the contribution of intergranular oxide to interfacial dielectric loss.

We consider several alternative hypotheses, such as changes to the metal-air interface loss as a result of the modified film morphology, an increase of film  $T_c$  in the thicker films suppressing quasiparticle-related losses, or a decrease of the participation ratios of the lossy interfaces associated with a dilution of the electric field in a thicker CPW; however, none of these hypotheses explain the observed improvement in device quality sufficiently.

Unfortunately, the CPW resonators fabricated on the thickest, 500 nm-thick films, show a frequency-dependent increase in non-TLS related loss. This loss shows at high drive powers where the TLS bath is saturated; however, it impacts the low-power  $Q$  as well due to the cumulative nature of the various loss mechanisms. We find that at the lower end of our frequency range  $\sim 4$  GHz, the non-TLS related loss of the thicker films is comparable to their thinner counterparts. However, towards 8 GHz, this other loss becomes significant. Further research is much needed to establish the nature of this non-TLS-related loss mechanism, and how it relates to transmon qubits. For example, if the mechanisms is related to radiation, qubits may be less affected due to their smaller

footprint. Other losses, however, may affect qubits similarly to resonators. Therefore, until this question is resolved, 300 nm thick films may be preferable.

Having decreased the TLS-loss of our devices to such an extent that other losses become comparable, further studies are needed into what other, non-TLS-related losses our devices are facing – even if the 300 nm film with lower other losses is selected. Additionally, now that the SM loss is mitigated further, the remaining TLS loss may originate predominantly at the substrate-air and metal-air interfaces. We investigate some strategies towards mitigating loss at these interfaces in Chapter 7; however, considerable further research is needed on that topic, particularly for Al-based devices.

Following the success of tantalum- and niobium-based qubits in breaking through the  $100\ \mu\text{s}$   $T_1$  threshold, we have now demonstrated that Al-based platforms are also capable of achieving comparable coherence times. This suggests that the important factor in achieving high coherence in superconducting qubits might not necessarily be the metal itself, but more importantly, the surface oxide, the fabrication procedure and the resulting materials interfaces, as well as the device design. Al has long been a popular choice of superconductor for quantum circuits for many reasons: availability, ease and variability of deposition, reproducibility, self-passivating native oxide, and the relatively uncomplicated Al/ $\text{AlO}_x$ /Al Josephson junction platform, among others. Al has proved itself a reliable material; however, the question now becomes whether it is the material that will yield the highest-performing devices. Future works in side-by-side comparisons of state-of-the-art qubits fabricated from the different popular superconductors, using identical designs and matching fabrication procedures, would be of special interest. However, this is challenging to execute: differences will occur during etching, as some metals perform better when etched using wet chemistry, while others should be etched by a plasma, and they rarely share a similar etching chemistry.

The fact that Al reacts with many acids and bases is convenient in certain applications – on the other hand, it severely limits our options for efficient but harsh chemical treatments designed to clean the substrate-air or metal-air interfaces, as the metal film incurs damage during these treatments. The choice of metals suitable for in-situ capping is also limited for Al due to its tendency to intermetallize with many materials. Finding a reliable in-situ passivation method for Al that lowers the metal-air loss, as well as protects the underlying metal so that harsh surface treatments could be executed without damaging the metal, might remove the negative implications of using Al while keeping the advantages.

Additionally to our efforts in TLS loss mitigation, we have investigated the nonlinear behaviour inherent to TLSs through intermodulation spectroscopy of a CPW resonator.



---

This resonator with no engineered nonlinear elements produced frequency mixing products when driven by two slightly detuned drive tones – a hallmark feature of nonlinear media. We show that the behaviour of this measured nonlinearity is consistent with the standard tunneling model for interacting TLSs. This method of analysis is particularly interesting in distinguishing various loss sources in a device, as it picks up only on the mechanisms with a nonlinear fingerprint. Furthermore, we find potential in applying this measurement method continuously to track the spectrally unstable coupling between the TLSs and the device over time. As the TLS model parameters can be reconstructed from a single spectrum taken at an intermediate power level, the method also provides a framework for a loss characterization speedup.

We also developed an improved fabrication recipe and design for three-dimensional coaxial cavity resonators. This recipe has since been implemented in architectures integrating a qubit inside the cavity, and used to produce bosonic states in a step towards continuous variable quantum computing [166].

The flip-chip architecture demonstrated in the appended Paper **E** has also been developed further, producing a 25-qubit chip used to investigate cross-talk between different qubits on a chip, which affects quantum gate fidelities [167]. Future work for the extensibility of this architecture could plausibly include adding through-silicon-vias to aid with a more compact wiring.



---

## References

---

- [1] H. H. Goldstine and A. Goldstine, “The electronic numerical integrator and computer (ENIAC),” *Math. Comput.*, vol. 2, no. 15, pp. 97–110, 1946, ISSN: 0025-5718. DOI: 10.1090/S0025-5718-1946-0018977-0.
- [2] *Frontier supercomputer debuts as world’s fastest, breaking exascale barrier | ORNL*, [Online; accessed 4. Aug. 2024], May 2022. [Online]. Available: <https://www.ornl.gov/news/frontier-supercomputer-debuts-worlds-fastest-breaking-exascale-barrier>.
- [3] P. W. Shor, “Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer,” *SIAM J. Comput.*, vol. 26, no. 5, 1997. [Online]. Available: <https://epubs.siam.org/doi/10.1137/S0097539795293172>.
- [4] D. S. Abrams and S. Lloyd, “Simulation of Many-Body Fermi Systems on a Universal Quantum Computer,” *Phys. Rev. Lett.*, vol. 79, no. 13, pp. 2586–2589, Sep. 1997, ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.79.2586.
- [5] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge, England, UK: Cambridge University Press, Dec. 2010, ISBN: 978-0-51197666-7. DOI: 10.1017/CB09780511976667.
- [6] M. M. Vopson, “Estimation of the information contained in the visible matter of the universe,” *AIP Adv.*, vol. 11, no. 10, Oct. 2021. DOI: 10.1063/5.0064475.
- [7] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai, “Coherent control of macroscopic quantum states in a single-Cooper-pair box,” *Nature*, vol. 398, pp. 786–788, Apr. 1999, ISSN: 1476-4687. DOI: 10.1038/19718.

- [8] J. Koch *et al.*, “Charge-insensitive qubit design derived from the Cooper pair box,” *Phys. Rev. A*, vol. 76, no. 4, p. 042319, Oct. 2007, ISSN: 2469-9934. DOI: 10.1103/PhysRevA.76.042319.
- [9] J. M. Sage, V. Bolkhovskiy, W. D. Oliver, B. Turek, and P. B. Welander, “Study of loss in superconducting coplanar waveguide resonators,” *Journal of Applied Physics*, vol. 109, no. 6, p. 063915, Mar. 2011, ISSN: 0021-8979. DOI: 10.1063/1.3552890. [Online]. Available: <https://doi.org/10.1063/1.3552890>.
- [10] A. D. *et al.*, “Characterization and reduction of capacitive loss induced by sub-micron josephson junction fabrication in superconducting qubits,” *Appl. Phys. Lett.*, vol. 111, no. 2, p. 022601, 2017. DOI: 10.1063/1.4993577.
- [11] R. T. Gordon *et al.*, “Environmental radiation impact on lifetimes and quasiparticle tunneling rates of fixed-frequency transmon qubits,” *Applied Physics Letters*, vol. 120, no. 7, p. 074002, Feb. 2022, ISSN: 0003-6951. DOI: 10.1063/5.0078785.
- [12] J. J. Burnett, A. Bengtsson, M. Scigliuzzo, D. Niepce, M. Kudra, P. Delsing, and J. Bylander, “Decoherence benchmarking of superconducting qubits,” *npj Quantum Inf.*, vol. 5, no. 54, pp. 1–8, Jun. 2019, ISSN: 2056-6387. DOI: 10.1038/s41534-019-0168-5.
- [13] A. Megrant *et al.*, “Planar superconducting resonators with internal quality factors above one million,” *Applied Physics Letters*, vol. 100, no. 11, p. 113510, Mar. 2012, ISSN: 0003-6951. DOI: 10.1063/1.3693409.
- [14] A. Nersisyan *et al.*, “Manufacturing low dissipation superconducting quantum processors,” *2019 IEEE International Electron Devices Meeting (IEDM)*, pp. 31.1.1–31.1.4, 2019. DOI: 10.1109/IEDM19573.2019.8993458.
- [15] J. B. Chang *et al.*, “Improved superconducting qubit coherence using titanium nitride,” *Applied Physics Letters*, vol. 103, no. 1, p. 012602, Jul. 2013, ISSN: 0003-6951. DOI: 10.1063/1.4813269.
- [16] C. E. Murray, “Material matters in superconducting qubits,” *Materials Science and Engineering: R: Reports*, vol. 146, p. 100646, 2021, ISSN: 0927-796X. DOI: <https://doi.org/10.1016/j.mser.2021.100646>.
- [17] P. A. Spring *et al.*, “High coherence and low cross-talk in a tileable 3D integrated superconducting circuit architecture,” *Sci. Adv.*, vol. 8, no. 16, Apr. 2022, ISSN: 2375-2548. DOI: 10.1126/sciadv.ab16698.

- 
- [18] K. *et al.*, “Superconducting Qubits: Current State of Play,” *Annu. Rev. Condens. Matter Phys.*, vol. 11, no. 1, pp. 369–395, Mar. 2020, ISSN: 1947-5454. DOI: 10.1146/annurev-conmatphys-031119-050605.
- [19] A. P. M. Place *et al.*, “New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds,” *Nat. Commun.*, vol. 12, no. 1779, pp. 1–6, Mar. 2021, ISSN: 2041-1723. DOI: 10.1038/s41467-021-22030-5.
- [20] C. Wang *et al.*, “Towards practical quantum computers: transmon qubit with a lifetime approaching 0.5 milliseconds,” *npj Quantum Inf.*, vol. 8, no. 3, pp. 1–6, Jan. 2022, ISSN: 2056-6387. DOI: 10.1038/s41534-021-00510-2.
- [21] H. Deng *et al.*, “Titanium nitride film on sapphire substrate with low dielectric loss for superconducting qubits,” *Phys. Rev. Appl.*, vol. 19, p. 024013, 2 Feb. 2023. DOI: 10.1103/PhysRevApplied.19.024013.
- [22] S. Kono, J. Pan, M. Chegnizadeh, X. Wang, A. Youssefi, M. Scigliuzzo, and T. J. Kippenberg, “Mechanically Induced Correlated Errors on Superconducting Qubits with Relaxation Times Exceeding 0.4 Milliseconds,” *arXiv*, May 2023. eprint: 2305.02591.
- [23] R. Barends *et al.*, “Superconducting quantum circuits at the surface code threshold for fault tolerance,” *Nature*, vol. 508, no. 7497, pp. 500–503, Apr. 2014, ISSN: 1476-4687. DOI: 10.1038/nature13171. eprint: 24759412.
- [24] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, “Surface codes: Towards practical large-scale quantum computation,” *Phys. Rev. A*, vol. 86, no. 3, p. 032324, Sep. 2012, ISSN: 2469-9934. DOI: 10.1103/PhysRevA.86.032324.
- [25] T. Abad, J. Fernández-Pendás, A. Frisk Kockum, and G. Johansson, “Universal Fidelity Reduction of Quantum Operations from Weak Dissipation,” *Phys. Rev. Lett.*, vol. 129, no. 15, p. 150504, Oct. 2022, ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.129.150504.
- [26] F. Arute *et al.*, “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, pp. 505–510, Oct. 2019, ISSN: 1476-4687. DOI: 10.1038/s41586-019-1666-5.
- [27] Y. Wu *et al.*, “Strong Quantum Computational Advantage Using a Superconducting Quantum Processor,” *Phys. Rev. Lett.*, vol. 127, no. 18, p. 180501, Oct. 2021, ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.127.180501.

- [28] H. Linn, I. Brundin, L. García-Álvarez, and G. Johansson, “Resource analysis of quantum algorithms for coarse-grained protein folding models,” *arXiv*, Nov. 2023. DOI: 10.48550/arXiv.2311.04186. eprint: 2311.04186.
- [29] L.-C. Han *et al.*, “Active reset of superconducting qubits using the electronics based on RF switches,” *AIP Adv.*, vol. 13, no. 9, Sep. 2023. DOI: 10.1063/5.0166535.
- [30] M. Tinkham, *Introduction to Superconductivity*. Mineola, NY, USA: Dover Publications, 2004, ISBN: 978-0-48643503-9.
- [31] J. W. Rohlf, *Modern Physics from alpha to Z0*. Chichester, England, UK: Wiley, Apr. 1994.
- [32] K. H. Onnes, “The Superconductivity of Mercury,” *Comm. Phys. Lab. Univ. Leiden*, vol. 251, no. 122, 1911.
- [33] F. E. Wang, *Bonding Theory for Metals and Alloys*. Waltham, MA, USA: Elsevier Science, 2005, ISBN: 978-0-444-51978-8. DOI: 10.1016/B978-0-444-51978-8.X5000-4.
- [34] D. C. Hawkworth, “Development of Superconducting Magnet Systems for MRI,” in *Advances in Cryogenic Engineering*, Boston, MA, USA: Springer, Boston, MA, 1990, pp. 529–538, ISBN: 978-1-4613-0639-9. DOI: 10.1007/978-1-4613-0639-9\_62.
- [35] T. Nakamura, D. Tamada, Y. Yanagi, Y. Itoh, T. Nemoto, H. Utumi, and K. Kose, “Development of a superconducting bulk magnet for NMR and MRI,” *J. Magn. Reson.*, vol. 259, pp. 68–75, Oct. 2015, ISSN: 1090-7807. DOI: 10.1016/j.jmr.2015.07.012.
- [36] N. Yanagi *et al.*, “Design and development of high-temperature superconducting magnet system with joint-winding for the helical fusion reactor,” *Nucl. Fusion*, vol. 55, no. 5, p. 053021, Apr. 2015, ISSN: 0029-5515. DOI: 10.1088/0029-5515/55/5/053021.
- [37] C. J. Gorter and H. Casimir, “On supraconductivity I,” *Physica*, vol. 1, no. 1, pp. 306–320, Jan. 1934, ISSN: 0031-8914. DOI: 10.1016/S0031-8914(34)90037-9.
- [38] L. N. Cooper, “Bound Electron Pairs in a Degenerate Fermi Gas,” *Phys. Rev.*, vol. 104, no. 4, pp. 1189–1190, Nov. 1956, ISSN: 1536-6065. DOI: 10.1103/PhysRev.104.1189.

- 
- [39] J. Bardeen, L. N. Cooper, and J. R. Schrieffer, “Microscopic Theory of Superconductivity,” *Phys. Rev.*, vol. 106, no. 1, pp. 162–164, Apr. 1957, ISSN: 1536-6065. DOI: 10.1103/PhysRev.106.162.
- [40] J. Bardeen, L. N. Cooper, and J. R. Schrieffer, “Theory of Superconductivity,” *Phys. Rev.*, vol. 108, no. 5, pp. 1175–1204, Dec. 1957, ISSN: 1536-6065. DOI: 10.1103/PhysRev.108.1175.
- [41] V. L. Ginzburg and L. D. Landau, “On the Theory of superconductivity,” *Zh. Eksp. Teor. Fiz.*, vol. 20, pp. 1064–1082, 1950.
- [42] A. A. Abrikosov, “The magnetic properties of superconducting alloys,” *J. Phys. Chem. Solids*, vol. 2, no. 3, pp. 199–208, Jan. 1957, ISSN: 0022-3697. DOI: 10.1016/0022-3697(57)90083-5.
- [43] F. London and H. London, “The electromagnetic equations of the superconductor,” *Proc. Royal Soc. Lond. A*, vol. 149, no. 866, pp. 71–88, Mar. 1935. DOI: 10.1098/rspa.1935.0048.
- [44] B. D. Josephson, “Possible new effects in superconductive tunnelling,” *Physics Letters*, vol. 1, no. 7, pp. 251–253, Jul. 1962, ISSN: 0031-9163. DOI: 10.1016/0031-9163(62)91369-0.
- [45] B. D. Josephson, “The discovery of tunnelling supercurrents,” *Rev. Mod. Phys.*, vol. 46, no. 2, pp. 251–254, Apr. 1974, ISSN: 1539-0756. DOI: 10.1103/RevModPhys.46.251.
- [46] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, “A quantum engineer’s guide to superconducting qubits,” *Appl. Phys. Rev.*, vol. 6, no. 2, Jun. 2019. DOI: 10.1063/1.5089550.
- [47] D. Niepce, J. J. Burnett, M. Kudra, J. H. Cole, and J. Bylander, “Stability of superconducting resonators: Motional narrowing and the role of Landau-Zener driving of two-level defects,” *Sci. Adv.*, vol. 7, no. 39, Sep. 2021, ISSN: 2375-2548. DOI: 10.1126/sciadv.abh0462.
- [48] J. Bylander *et al.*, “Noise spectroscopy through dynamical decoupling with a superconducting flux qubit,” *Nat. Phys.*, vol. 7, pp. 565–570, Jul. 2011, ISSN: 1745-2481. DOI: 10.1038/nphys1994.
- [49] D. I. Schuster, A. Wallraff, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. M. Girvin, and R. J. Schoelkopf, “ac Stark Shift and Dephasing of a Superconducting Qubit Strongly Coupled to a Cavity Field,” *Phys. Rev. Lett.*, vol. 94, no. 12, p. 123602, Mar. 2005, ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.94.123602.

- [50] A. Megrant *et al.*, “Planar superconducting resonators with internal quality factors above one million,” *Appl. Phys. Lett.*, vol. 100, no. 11, Mar. 2012, ISSN: 0003-6951. DOI: 10.1063/1.3693409.
- [51] C. Wang, C. Axline, Y. Y. Gao, T. Brecht, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf, “Surface participation and dielectric loss in superconducting qubits,” *Appl. Phys. Lett.*, vol. 107, no. 16, Oct. 2015, ISSN: 0003-6951. DOI: 10.1063/1.4934486.
- [52] J. Burnett *et al.*, “Evidence for interacting two-level systems from the  $1/f$  noise of a superconducting resonator,” *Nat. Commun.*, vol. 5, no. 4119, pp. 1–6, Jun. 2014, ISSN: 2041-1723. DOI: 10.1038/ncomms5119.
- [53] A. Bilmes, A. Megrant, P. Klimov, G. Weiss, J. M. Martinis, A. V. Ustinov, and J. Lisenfeld, “Resolving the positions of defects in superconducting quantum bits,” *Sci. Rep.*, vol. 10, no. 3090, pp. 1–6, Feb. 2020, ISSN: 2045-2322. DOI: 10.1038/s41598-020-59749-y.
- [54] D. Niepce, J. J. Burnett, M. G. Latorre, and J. Bylander, “Geometric scaling of two-level-system loss in superconducting resonators,” *Supercond. Sci. Technol.*, vol. 33, no. 2, p. 025 013, Jan. 2020, ISSN: 0953-2048. DOI: 10.1088/1361-6668/ab6179.
- [55] C. Müller, J. H. Cole, and J. Lisenfeld, “Towards understanding two-level-systems in amorphous solids: insights from quantum circuits,” *Rep. Prog. Phys.*, vol. 82, no. 12, p. 124 501, Oct. 2019, ISSN: 0034-4885. DOI: 10.1088/1361-6633/ab3a7e.
- [56] S. E. de Graaf, A. A. Adamyan, T. Lindström, D. Erts, S. E. Kubatkin, A. Y. Tzalenchuk, and A. V. Danilov, “Direct Identification of Dilute Surface Spins on  $\text{Al}_2\text{O}_3$ : Origin of Flux Noise in Quantum Circuits,” *Phys. Rev. Lett.*, vol. 118, no. 5, p. 057 703, Jan. 2017, ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.118.057703.
- [57] P. Kumar *et al.*, “Origin and Reduction of  $1/f$  Magnetic Flux Noise in Superconducting Devices,” *Phys. Rev. Appl.*, vol. 6, no. 4, p. 041 001, Oct. 2016, ISSN: 2331-7019. DOI: 10.1103/PhysRevApplied.6.041001.
- [58] A. Jayaraman, A. V. Danilov, J. Bylander, and S. E. Kubatkin, “Loss and decoherence in superconducting circuits on silicon: Insights from electron spin resonance,” *Phys. Rev. Appl.*, vol. 22, no. 1, p. 014 030, Jul. 2024, ISSN: 2331-7019. DOI: 10.1103/PhysRevApplied.22.014030.



- 
- [59] J. M. Martinis *et al.*, “Decoherence in Josephson Qubits from Dielectric Loss,” *Phys. Rev. Lett.*, vol. 95, no. 21, p. 210503, Nov. 2005, ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.95.210503.
- [60] J. Burnett, J. Sagar, O. W. Kennedy, P. A. Warburton, and J. C. Fenton, “Low-Loss Superconducting Nanowire Circuits Using a Neon Focused Ion Beam,” *Phys. Rev. Appl.*, vol. 8, no. 1, p. 014039, Jul. 2017, ISSN: 2331-7019. DOI: 10.1103/PhysRevApplied.8.014039.
- [61] P. Macha, S. H. W. van der Ploeg, G. Oelsner, E. Il’ichev, H.-G. Meyer, S. Wünsch, and M. Siegel, “Losses in coplanar waveguide resonators at millikelvin temperatures,” *Appl. Phys. Lett.*, vol. 96, no. 6, Feb. 2010, ISSN: 0003-6951. DOI: 10.1063/1.3309754.
- [62] D. P. Pappas, M. R. Vissers, D. S. Wisbey, J. S. Kline, and J. Gao, “Two Level System Loss in Superconducting Microwave Resonators,” *IEEE Trans. Appl. Supercond.*, vol. 21, no. 3, pp. 871–874, Jan. 2011. DOI: 10.1109/TASC.2010.2097578.
- [63] M. R. Vissers, J. Gao, D. S. Wisbey, D. A. Hite, C. C. Tsuei, A. D. Corcoles, M. Steffen, and D. P. Pappas, “Low loss superconducting titanium nitride coplanar waveguide resonators,” *Appl. Phys. Lett.*, vol. 97, no. 23, Dec. 2010, ISSN: 0003-6951. DOI: 10.1063/1.3517252.
- [64] J. Wenner *et al.*, “Surface loss simulations of superconducting coplanar waveguide resonators,” *Appl. Phys. Lett.*, vol. 99, no. 11, p. 113513, Sep. 2011, ISSN: 0003-6951. DOI: 10.1063/1.3637047.
- [65] F. Argall and A. K. Jonscher, “Dielectric properties of thin films of aluminium oxide and silicon oxide,” *Thin Solid Films*, vol. 2, no. 3, pp. 185–210, Sep. 1968, ISSN: 0040-6090. DOI: 10.1016/0040-6090(68)90002-3.
- [66] A. P. Read, B. J. Chapman, C. U. Lei, J. C. Curtis, S. Ganjam, L. Krayzman, L. Frunzio, and R. J. Schoelkopf, “Precision Measurement of the Microwave Dielectric Loss of Sapphire in the Quantum Regime with Parts-per-Billion Sensitivity,” *Phys. Rev. Appl.*, vol. 19, no. 3, p. 034064, Mar. 2023, ISSN: 2331-7019. DOI: 10.1103/PhysRevApplied.19.034064.
- [67] D. C. Mattis and J. Bardeen, “Theory of the Anomalous Skin Effect in Normal and Superconducting Metals,” *Phys. Rev.*, vol. 111, no. 2, pp. 412–417, Jul. 1958, ISSN: 1536-6065. DOI: 10.1103/PhysRev.111.412.

- [68] J. Gao, “The physics of superconducting microwave resonators,” Ph.D. dissertation, California Institute of Technology, 2008.
- [69] R. Barends *et al.*, “Minimizing quasiparticle generation from stray infrared light in superconducting quantum circuits,” *Appl. Phys. Lett.*, vol. 99, no. 11, Sep. 2011, ISSN: 0003-6951. DOI: 10.1063/1.3638063.
- [70] G. Catelani, R. J. Schoelkopf, M. H. Devoret, and L. I. Glazman, “Relaxation and frequency shifts induced by quasiparticles in superconducting qubits,” *Phys. Rev. B*, vol. 84, no. 6, p. 064517, Aug. 2011, ISSN: 2469-9969. DOI: 10.1103/PhysRevB.84.064517.
- [71] K. Serniak *et al.*, “Hot Nonequilibrium Quasiparticles in Transmon Qubits,” *Phys. Rev. Lett.*, vol. 121, no. 15, p. 157701, Oct. 2018, ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.121.157701.
- [72] X. Pan *et al.*, “Engineering superconducting qubits to reduce quasiparticles and charge noise,” *Nat. Commun.*, vol. 13, no. 7196, pp. 1–7, Nov. 2022, ISSN: 2041-1723. DOI: 10.1038/s41467-022-34727-2.
- [73] R.-P. Riwar, A. Hosseinkhani, L. D. Burkhardt, Y. Y. Gao, R. J. Schoelkopf, L. I. Glazman, and G. Catelani, “Normal-metal quasiparticle traps for superconducting qubits,” *Phys. Rev. B*, vol. 94, no. 10, p. 104516, Sep. 2016, ISSN: 2469-9969. DOI: 10.1103/PhysRevB.94.104516.
- [74] R.-P. Riwar and G. Catelani, “Efficient quasiparticle traps with low dissipation through gap engineering,” *Phys. Rev. B*, vol. 100, no. 14, p. 144514, Oct. 2019, ISSN: 2469-9969. DOI: 10.1103/PhysRevB.100.144514.
- [75] S. Gustavsson *et al.*, “Suppressing relaxation in superconducting qubits by quasiparticle pumping,” *Science*, vol. 354, no. 6319, pp. 1573–1577, Dec. 2016, ISSN: 0036-8075. DOI: 10.1126/science.aah5844.
- [76] M. McEwen *et al.*, “Resisting high-energy impact events through gap engineering in superconducting qubit arrays,” *arXiv*, Feb. 2024. DOI: 10.48550/arXiv.2402.15644. eprint: 2402.15644.
- [77] B. Chiaro *et al.*, “Dielectric surface loss in superconducting resonators with flux-trapping holes,” *Supercond. Sci. Technol.*, vol. 29, no. 10, p. 104006, Aug. 2016, ISSN: 0953-2048. DOI: 10.1088/0953-2048/29/10/104006.

- 
- [78] S. Hähnle, N. v. Marrewijk, A. Endo, K. Karatsu, D. J. Thoen, V. Murugesan, and J. J. A. Baselmans, “Suppression of radiation loss in high kinetic inductance superconducting co-planar waveguides,” *Appl. Phys. Lett.*, vol. 116, no. 18, May 2020, ISSN: 0003-6951. DOI: 10.1063/5.0005047.
- [79] K. J. Vahala, “Optical microcavities,” *Nature*, vol. 424, pp. 839–846, Aug. 2003, ISSN: 1476-4687. DOI: 10.1038/nature01939.
- [80] L. Chen *et al.*, “Transmon qubit readout fidelity at the threshold for quantum error correction without a quantum-limited amplifier,” *npj Quantum Inf.*, vol. 9, no. 26, pp. 1–7, Mar. 2023, ISSN: 2056-6387. DOI: 10.1038/s41534-023-00689-6.
- [81] E. A. Sete, J. M. Gambetta, and A. N. Korotkov, “Purcell effect with microwave drive: Suppression of qubit relaxation rate,” *Phys. Rev. B*, vol. 89, no. 10, p. 104516, Mar. 2014, ISSN: 2469-9969. DOI: 10.1103/PhysRevB.89.104516.
- [82] C. W. Warren *et al.*, “Extensive characterization and implementation of a family of three-qubit gates at the coherence limit,” *npj Quantum Inf.*, vol. 9, no. 44, pp. 1–9, May 2023, ISSN: 2056-6387. DOI: 10.1038/s41534-023-00711-x.
- [83] R. Barends *et al.*, “Coherent Josephson Qubit Suitable for Scalable Quantum Integrated Circuits,” *Phys. Rev. Lett.*, vol. 111, no. 8, p. 080502, Aug. 2013, ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.111.080502.
- [84] G. Stan, S. B. Field, and J. M. Martinis, “Critical Field for Complete Vortex Expulsion from Narrow Superconducting Strips,” *Phys. Rev. Lett.*, vol. 92, no. 9, p. 097003, Mar. 2004, ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.92.097003.
- [85] C. P. Wen, “Coplanar Waveguide: A Surface Strip Transmission Line Suitable for Nonreciprocal Gyromagnetic Device Applications,” *IEEE Trans. Microwave Theory Tech.*, vol. 17, no. 12, pp. 1087–1090, Dec. 1969. DOI: 10.1109/TMTT.1969.1127105.
- [86] R. N. Simons, *Coplanar Waveguide Circuits, Components, and Systems*. Mar. 2001, ISBN: 978-0-47116121-9. DOI: 10.1002/0471224758.
- [87] E. Rehnman, “Geometric scaling of competing loss mechanisms in superconducting coplanar waveguide resonators,” M.S. thesis, Chalmers University of Technology, 2023.
- [88] J. M. Martinis, R. Barends, and A. N. Korotkov, “Calculation of Coupling Capacitance in Planar Electrodes,” *arXiv*, Oct. 2014. DOI: 10.48550/arXiv.1410.3458. eprint: 1410.3458.

- [89] M. Göppl *et al.*, “Coplanar waveguide resonators for circuit quantum electrodynamics,” *J. Appl. Phys.*, vol. 104, no. 11, p. 113 904, Dec. 2008, ISSN: 0021-8979. DOI: 10.1063/1.3010859.
- [90] P. G. Baity, C. Maclean, V. Seferai, J. Bronstein, Y. Shu, T. Hemakumara, and M. Weides, “Circle fit optimization for resonator quality factor measurements: Point redistribution for maximal accuracy,” *Phys. Rev. Res.*, vol. 6, no. 1, p. 013 329, Mar. 2024, ISSN: 2643-1564. DOI: 10.1103/PhysRevResearch.6.013329.
- [91] B. A. Mazin, “Microwave Kinetic Inductance Detectors,” Ph.D. dissertation, California Institute of Technology, 2005. [Online]. Available: <https://thesis.library.caltech.edu/3910>.
- [92] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, “Circuit quantum electrodynamics,” *Rev. Mod. Phys.*, vol. 93, no. 2, p. 025 005, May 2021, ISSN: 1539-0756. DOI: 10.1103/RevModPhys.93.025005.
- [93] S. Franssila, *Introduction to microfabrication*. Wiley, 2010.
- [94] K. A. Reinhardt and W. Kern, *Handbook of Silicon Wafer Cleaning Technology*. William Andrew, 2018, ISBN: 978-0-323-51084-4. DOI: 10.1016/C2016-0-01001-X.
- [95] H. Ubara, T. Imura, and A. Hiraki, “Formation of  $x_n$ -SiH- $x_{25}q$  bonds on the surface of microcrystalline silicon covered with SiO<sub>x</sub> by HF treatment,” *Solid State Commun.*, vol. 50, no. 7, pp. 673–675, May 1984, ISSN: 0038-1098. DOI: 10.1016/0038-1098(84)90156-X.
- [96] R. I. Revilla, H. Terryn, and I. De Graeve, “On the use of SKPFM for in situ studies of the repassivation of the native oxide film on aluminium in air,” *Electrochem. Commun.*, vol. 93, pp. 162–165, Aug. 2018, ISSN: 1388-2481. DOI: 10.1016/j.elecom.2018.07.010.
- [97] *DWL 66+ System Specifications*, [Online; accessed 25. Aug. 2024], Apr. 2022. [Online]. Available: <https://heidelberg-instruments.com/wp-content/uploads/2021/02/Fact-Sheet-DWL-66-v220408.pdf>.
- [98] C. Mack, *Fundamental Principles of Optical Lithography*. Nov. 2007, ISBN: 978-047001893-4. DOI: 10.1002/9780470723876.
- [99] L. de Broglie, “Recherches sur la théorie des quanta,” Ph.D. dissertation, Ann. Phys. (Paris), 1924.
- [100] D. J. Elliott, *Integrated Circuit Fabrication Technology*. Maidenhead, England, UK: McGraw-Hill, 1982, ISBN: 978-0-07019238-6.

- 
- [101] A. Osman *et al.*, “Simplified Josephson-junction fabrication process for reproducibly high-performance superconducting qubits,” *Appl. Phys. Lett.*, vol. 118, no. 6, Feb. 2021, ISSN: 0003-6951. DOI: 10.1063/5.0037093.
- [102] *Candidate List of substances of very high concern for Authorisation - ECHA*, [Online; accessed 12. Apr. 2024], Apr. 2024. [Online]. Available: <https://echa.europa.eu/candidate-list-table/-/dislist/details/0b0236e1807da281>.
- [103] G. J. Dolan, “Offset masks for lift-off photoprocessing,” *Appl. Phys. Lett.*, vol. 31, no. 5, pp. 337–339, Sep. 1977, ISSN: 0003-6951. DOI: 10.1063/1.89690.
- [104] F. Lecocq *et al.*, “Junction fabrication by shadow evaporation without a suspended bridge,” *Nanotechnology*, vol. 22, no. 31, p. 315302, Jul. 2011, ISSN: 0957-4484. DOI: 10.1088/0957-4484/22/31/315302.
- [105] M. Checchin, D. Frolov, A. Lunin, A. Grassellino, and A. Romanenko, “Measurement of the Low-Temperature Loss Tangent of High-Resistivity Silicon Using a High- $Q$  Superconducting Resonator,” *Phys. Rev. Appl.*, vol. 18, no. 3, p. 034013, Sep. 2022, ISSN: 2331-7019. DOI: 10.1103/PhysRevApplied.18.034013.
- [106] M. V. P. Altoé *et al.*, “Localization and Mitigation of Loss in Niobium Superconducting Circuits,” *PRX Quantum*, vol. 3, no. 2, p. 020312, Apr. 2022, ISSN: 2691-3399. DOI: 10.1103/PRXQuantum.3.020312.
- [107] K. Li *et al.*, “Long-lived transmons with different electrode layouts,” *MRS Adv.*, vol. 7, no. 13, pp. 273–277, May 2022, ISSN: 2059-8521. DOI: 10.1557/s43580-022-00265-8.
- [108] J. G. C. Milne, “Superconducting Transition Temperature of High-Purity Tantalum Metal,” *Phys. Rev.*, vol. 122, no. 2, pp. 387–388, Apr. 1961, ISSN: 1536-6065. DOI: 10.1103/PhysRev.122.387.
- [109] K. Kouwenhoven *et al.*, “Resolving Power of Visible-To-Near-Infrared Hybrid  $\beta$ -Ta/Nb-Ti-N Kinetic Inductance Detectors,” *Phys. Rev. Appl.*, vol. 19, no. 3, p. 034007, Mar. 2023, ISSN: 2331-7019. DOI: 10.1103/PhysRevApplied.19.034007.
- [110] Y. Krasnikova *et al.*, “Magnetic fluctuations in niobium pentoxide,” *arXiv*, Dec. 2023. DOI: 10.48550/arXiv.2312.10697. eprint: 2312.10697.
- [111] A. Bruno, G. de Lange, S. Asaad, K. L. van der Enden, N. K. Langford, and L. DiCarlo, “Reducing intrinsic loss in superconducting resonators by surface treatment and deep etching of silicon substrates,” *Appl. Phys. Lett.*, vol. 106, no. 18, May 2015, ISSN: 0003-6951. DOI: 10.1063/1.4919761.

- [112] J. Van Damme *et al.*, “High-coherence superconducting qubits made using industry-standard, advanced semiconductor manufacturing,” *arXiv*, Mar. 2024. DOI: 10.48550/arXiv.2403.01312. eprint: 2403.01312.
- [113] A. Osman *et al.*, “Mitigation of frequency collisions in superconducting quantum processors,” *Phys. Rev. Res.*, vol. 5, no. 4, p. 043001, Oct. 2023, ISSN: 2643-1564. DOI: 10.1103/PhysRevResearch.5.043001.
- [114] D. M. Pozar, *Microwave Engineering, 4th Edition*. Hoboken, NJ, USA: Wiley, Apr. 2012, ISBN: 978-1-118-21363-6.
- [115] M. S. Khalil, M. J. A. Stoutimore, F. C. Wellstood, and K. D. Osborn, “An analysis method for asymmetric resonator transmission applied to superconducting devices,” *J. Appl. Phys.*, vol. 111, no. 5, Mar. 2012, ISSN: 0021-8979. DOI: 10.1063/1.3692073.
- [116] S. Probst, F. B. Song, P. A. Bushev, A. V. Ustinov, and M. Weides, “Efficient and robust analysis of complex scattering data under noise in microwave resonators,” *Rev. Sci. Instrum.*, vol. 86, no. 2, Feb. 2015, ISSN: 0034-6748. DOI: 10.1063/1.4907935.
- [117] A. Adamyan, “Slow propagation line-based superconducting devices for quantum technology,” Ph.D. dissertation, Chalmers University of Technology, Sweden, 2016, ISBN: 978-91-7597-413-2. eprint: 240894.
- [118] K. N. Nesterov and I. V. Pechenezhskiy, “Measurement-induced state transitions in dispersive qubit readout schemes,” *arXiv*, Feb. 2024. DOI: 10.48550/arXiv.2402.07360. eprint: 2402.07360.
- [119] N. S. Smirnov, E. A. Krivko, A. A. Solovyova, A. I. Ivanov, and I. A. Rodionov, “Wiring surface loss of a superconducting transmon qubit,” *Sci. Rep.*, vol. 14, no. 7326, pp. 1–7, Mar. 2024, ISSN: 2045-2322. DOI: 10.1038/s41598-024-57248-y.
- [120] A. Melville *et al.*, “Comparison of dielectric loss in titanium nitride and aluminum superconducting resonators,” *Appl. Phys. Lett.*, vol. 117, no. 12, Sep. 2020, ISSN: 0003-6951. DOI: 10.1063/5.0021950.
- [121] I. Tsioutsios *et al.*, “Free-standing silicon shadow masks for transmon qubit fabrication,” *AIP Adv.*, vol. 10, no. 6, Jun. 2020. DOI: 10.1063/1.5138953.
- [122] R. C. Jaeger, *Introduction to Microelectronic Fabrication, Volym 5*. Upper Saddle River, NJ, USA: Prentice Hall, 2002, ISBN: 978-0-20144494-0.

- 
- [123] D. P. Lozano *et al.*, “Low-loss  $\alpha$ -tantalum coplanar waveguide resonators on silicon wafers: fabrication, characterization and surface modification,” *Mater. Quantum Technol.*, vol. 4, no. 2, p. 025 801, May 2024, ISSN: 2633-4356. DOI: 10.1088/2633-4356/ad4b8c.
- [124] S. Fritz, A. Seiler, L. Radtke, R. Schneider, M. Weides, G. Weiß, and D. Gerthsen, “Correlating the nanostructure of Al-oxide with deposition conditions and dielectric contributions of two-level systems in perspective of superconducting quantum circuits,” *Scientific Reports*, vol. 8, no. 1, p. 7956, May 2018.
- [125] P. K. Naik, *Vacuum: Science, technology and applications*. CRC PRESS, 2018.
- [126] M.-C. G. Chalbot and I. G. Kavouras, “Nuclear magnetic resonance spectroscopy for determining the functional content of organic aerosols: A review,” *Environ. Pollut.*, vol. 191, pp. 232–249, Aug. 2014, ISSN: 0269-7491. DOI: 10.1016/j.envpol.2014.04.034.
- [127] A. Benninghoven, “Chemical Analysis of Inorganic and Organic Surfaces and Thin Films by Static Time-of-Flight Secondary Ion Mass Spectrometry (TOF-SIMS),” *Angew. Chem., Int. Ed. Engl.*, vol. 33, no. 10, pp. 1023–1043, Jun. 1994, ISSN: 0570-0833. DOI: 10.1002/anie.199410231.
- [128] R. S. Bauer, R. Z. Bachrach, and L. J. Brillson, “Au and Al interface reactions with SiO<sub>2</sub>,” *Applied Physics Letters*, vol. 37, no. 11, pp. 1006–1008, Jul. 1980, ISSN: 0003-6951. DOI: 10.1063/1.91720.
- [129] L. J. Zeng, T. Greibe, S. Nik, C. M. Wilson, P. Delsing, and E. Olsson, “Nanoscale interaction layer at the interface between Al films and SiO<sub>2</sub> substrates of Al/AlO<sub>x</sub>/Al Josephson tunnel junctions,” *Journal of Applied Physics*, vol. 113, no. 14, p. 143 905, Apr. 2013, ISSN: 0021-8979. DOI: 10.1063/1.4801798.
- [130] J. D. Plummer, M. Deal, and P. B. Griffin, *Silicon VLSI Technology: Fundamentals, Practice and Modeling*. Prentice Hall, 2000, 696–698, and 565–566, ISBN: 0130850373.
- [131] S. Fritz, L. Radtke, R. Schneider, M. Weides, and D. Gerthsen, “Optimization of Al/AlO<sub>x</sub>/Al-layer systems for Josephson junctions from a microstructure point of view,” *Journal of Applied Physics*, vol. 125, no. 16, p. 165 301, Apr. 2019, ISSN: 0021-8979. DOI: 10.1063/1.5089871.

- [132] B. M. McSkimming, A. Alexander, M. H. Samuels, B. Arey, I. Arslan, and C. J. K. Richardson, "Metamorphic growth of relaxed single crystalline aluminum on silicon (111)," *Journal of Vacuum Science & Technology A*, vol. 35, no. 2, p. 021 401, Dec. 2016, ISSN: 0734-2101. DOI: 10.1116/1.4971200.
- [133] M. Adamik, P. B. Barna, and I. Tomov, "Columnar structures in polycrystalline thin films developed by competitive growth," *Thin Solid Films*, vol. 317, no. 1, pp. 64–68, 1998, ISSN: 0040-6090. DOI: [https://doi.org/10.1016/S0040-6090\(97\)00661-5](https://doi.org/10.1016/S0040-6090(97)00661-5).
- [134] D. Chaverri, A. Saenz, and V. Castano, "Grain size and electrical resistivity measurements on aluminum polycrystalline thin films," *Materials Letters*, vol. 12, no. 5, pp. 344–348, 1991, ISSN: 0167-577X. DOI: [https://doi.org/10.1016/0167-577X\(91\)90114-L](https://doi.org/10.1016/0167-577X(91)90114-L).
- [135] A. E. Lita and J. E. Sanchez, "Effects of grain growth on dynamic surface scaling during the deposition of Al polycrystalline thin films," *Phys. Rev. B*, vol. 61, pp. 7692–7699, 11 Mar. 2000. DOI: 10.1103/PhysRevB.61.7692.
- [136] X. L. Yan, M. M. Duvenhage, J. Y. Wang, H. C. Swart, and J. J. Terblans, "Evaluation of sputtering induced surface roughness development of Ni/Cu multilayers thin films by Time-of-Flight Secondary Ion Mass Spectrometry depth profiling with different energies O<sub>2</sub><sup>+</sup> ion bombardment," *Thin Solid Films*, vol. 669, pp. 188–197, Jan. 2019, ISSN: 0040-6090. DOI: 10.1016/j.tsf.2018.10.049.
- [137] V. Gorbenko *et al.*, "SIMS depth profiling and topography studies of repetitive III-V trenches under low energy oxygen ion beam sputtering," *J. Vac. Sci. Technol., B*, vol. 34, no. 3, 03H131, May 2016, ISSN: 2166-2746. DOI: 10.1116/1.4944632.
- [138] E. Jeffrey *et al.*, "Fast Accurate State Measurement with Superconducting Qubits," *Phys. Rev. Lett.*, vol. 112, no. 19, p. 190 504, May 2014, ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.112.190504.
- [139] J. Van Damme *et al.*, "Argon-Milling-Induced Decoherence Mechanisms in Superconducting Quantum Circuits," *Phys. Rev. Appl.*, vol. 20, no. 1, p. 014 034, Jul. 2023, ISSN: 2331-7019. DOI: 10.1103/PhysRevApplied.20.014034.
- [140] K. R. Williams and R. S. Muller, "Etch rates for micromachining processing," *J. Microelectromech. Syst.*, vol. 5, no. 4, pp. 256–269, Dec. 1996. DOI: 10.1109/84.546406.



- 
- [141] M. Bal *et al.*, “Systematic improvements in transmon qubit coherence enabled by niobium surface encapsulation,” *npj Quantum Inf.*, vol. 10, no. 43, pp. 1–8, Apr. 2024, ISSN: 2056-6387. DOI: 10.1038/s41534-024-00840-x.
- [142] H. Meissner, “Superconductivity of Contacts with Interposed Barriers,” *Phys. Rev.*, vol. 117, no. 3, pp. 672–680, Feb. 1960, ISSN: 1536-6065. DOI: 10.1103/PhysRev.117.672.
- [143] H. Takatsuji and T. Arai, “Pinholes in Al thin films: their effects on TFT characteristics and a taguchi method analysis of their origins,” *Vacuum*, vol. 59, no. 2, pp. 606–613, Nov. 2000, ISSN: 0042-207X. DOI: 10.1016/S0042-207X(00)00323-7.
- [144] C. Kittel, *Introduction to Solid State Physics*. Chichester, England, UK: Wiley, 1996, ISBN: 978-0-47111181-8.
- [145] J. Clarke, “The proximity effect between superconducting and normal thin films in zero field,” *Journal de Physique Colloques*, no. 29, pp. C2-3-C2-16, 1968.
- [146] C. Cirillo, S. L. Prischepa, A. Romano, M. Salvato, and C. Attanasio, “Proximity effect in superconductor/highly paramagnetic Nb/Pd systems,” *Physica C*, vol. 404, no. 1, pp. 95–98, May 2004, ISSN: 0921-4534. DOI: 10.1016/j.physc.2003.10.036.
- [147] A. J. McAlister, “The Al-Pd (Aluminum-Palladium) system,” *Bulletin of Alloy Phase Diagrams*, vol. 7, no. 4, pp. 368–374, Aug. 1986. DOI: 10.1007/BF02873025.
- [148] A. Asaduzzaman, “Atomic-Scale Etching Mechanism of Aluminum with Fluorine-Based Plasma,” *J. Phys. Chem. C*, vol. 126, no. 33, pp. 14 180–14 186, Aug. 2022, ISSN: 1932-7447. DOI: 10.1021/acs.jpcc.2c03963.
- [149] A. Bagolini, P. Scauso, S. Sanguinetti, and P. Bellutti, “Silicon Deep Reactive Ion Etching with aluminum hard mask,” *Mater. Res. Express*, vol. 6, no. 8, p. 085 913, May 2019, ISSN: 2053-1591. DOI: 10.1088/2053-1591/ab2423.
- [150] A. Anferov, K.-H. Lee, F. Zhao, J. Simon, and D. I. Schuster, “Improved coherence in optically defined niobium trilayer-junction qubits,” *Phys. Rev. Appl.*, vol. 21, no. 2, p. 024 047, Feb. 2024, ISSN: 2331-7019. DOI: 10.1103/PhysRevApplied.21.024047.
- [151] J. Verjauw *et al.*, “Investigation of Microwave Loss Induced by Oxide Regrowth in High-Q Niobium Resonators,” *Phys. Rev. Appl.*, vol. 16, no. 1, p. 014 018, Jul. 2021, ISSN: 2331-7019. DOI: 10.1103/PhysRevApplied.16.014018.

- [152] S. Vallés Sanclemente, “Surface treatments for increased coherence of superconducting quantum processors,” M.S. thesis, Delft University of Technology, 2021.
- [153] N. Li and C.-M. Ho, “Photolithographic patterning of organosilane monolayer for generating large area two-dimensional B lymphocyte arrays,” *Lab Chip*, vol. 8, no. 12, p. 2105, Dec. 2008. DOI: 10.1039/b810329a.
- [154] H. Padamsee, *RF Superconductivity: Science, Technology, and Applications*. Weinheim, Germany: Wiley, Apr. 2009, ISBN: 978-3-52740572-5.
- [155] M. Reagor *et al.*, “Reaching 10 ms single photon lifetimes for superconducting aluminum cavities,” *Appl. Phys. Lett.*, vol. 102, no. 19, May 2013, ISSN: 0003-6951. DOI: 10.1063/1.4807015.
- [156] M. J. Reagor, “Superconducting cavities for circuit quantum electrodynamics,” Ph.D. dissertation, Yale University, 2015.
- [157] R. M. Brick, A. W. Pense, and R. B. Gordon, *Structure and Properties of Engineering Materials (McGraw-Hill Series in Materials Science and Engineering)*. McGraw-Hill College, Jan. 1977, ISBN: 978-0-07007721-8.
- [158] D. L. Creedon, M. Goryachev, N. Kostylev, T. B. Sercombe, and M. E. Tobar, “A 3D printed superconducting aluminium microwave cavity,” *Appl. Phys. Lett.*, vol. 109, no. 3, Jul. 2016, ISSN: 0003-6951. DOI: 10.1063/1.4958684.
- [159] J. Biznárová, “Electropolishing of superconducting cavities for quantum memory applications,” M.S. thesis, Chalmers University of Technology, 2019.
- [160] H. Zumbahlen and t. e. s. Of Analog Devices, *Linear Circuit Design Handbook*. Newnes, 2008, ISBN: 978-0-7506-8703-4. DOI: 10.1016/B978-0-7506-8703-4.X0001-6.
- [161] L. J. Swenson, P. K. Day, B. H. Eom, H. G. Leduc, N. Llombart, C. M. McKenney, O. Noroozian, and J. Zmuidzinas, “Operation of a titanium nitride superconducting microresonator detector in the nonlinear regime,” *J. Appl. Phys.*, vol. 113, no. 10, Mar. 2013, ISSN: 0021-8979. DOI: 10.1063/1.4794808.
- [162] J. Zmuidzinas, “Superconducting Microresonators: Physics and Applications,” *Annu. Rev. Condens. Matter Phys.*, vol. 3, no. 1, pp. 169–214, Feb. 2012, ISSN: 1947-5454. DOI: 10.1146/annurev-conmatphys-020911-125022.
- [163] K. Yasuda, S. Kawamura, and K. Watanabe, *Identification of Nonlinear Multi-Degree-of-Freedom Systems : Presentation of an Identification Technique*, Mar. 1988. DOI: 10.1299/JSMEC1988.31.8.

- [164] C. Hutter, D. Platz, E. A. Tholén, T. H. Hansson, and D. B. Haviland, “Reconstructing Nonlinearities with Intermodulation Spectroscopy,” *Phys. Rev. Lett.*, vol. 104, no. 5, p. 050 801, Feb. 2010, ISSN: 1079-7114. DOI: 10.1103/PhysRevLett.104.050801.
- [165] C. R. H. McRae, G. M. Stiehl, H. Wang, S.-X. Lin, S. A. Caldwell, D. P. Pappas, J. Mutus, and J. Combes, “Reproducible coherence characterization of superconducting quantum devices,” *Appl. Phys. Lett.*, vol. 119, no. 10, Sep. 2021, ISSN: 0003-6951. DOI: 10.1063/5.0060370.
- [166] M. Kudra *et al.*, “Robust Preparation of Wigner-Negative States with Optimized SNAP-Displacement Sequences,” *PRX Quantum*, vol. 3, no. 3, p. 030 301, Jul. 2022, ISSN: 2691-3399. DOI: 10.1103/PRXQuantum.3.030301.
- [167] S. Kosen *et al.*, “Signal crosstalk in a flip-chip quantum processor,” *arXiv*, Mar. 2024. DOI: 10.48550/arXiv.2403.00285. eprint: 2403.00285.

