

Advances in Superconducting Circuit Quantum Electrodynamics

by

Jérémy H. Béjanin

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Physics (Quantum Information)

Waterloo, Ontario, Canada, 2022

© Jérémy H. Béjanin 2022

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner:	John P. Davis Professor, University of Alberta
Supervisor:	Matteo Mariantoni Associate Professor
Co-Supervisor:	James Martin Associate Professor
Internal Member:	Roger Melko Professor
Internal Member:	Johnathan Baugh Associate Professor
Internal/External Examiner:	Christopher Wilson Professor

Author's Declaration

This thesis consists of material all of which I authored or co-authored: see Statement of Contributions included in the thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Statement of Contributions

Jérémy Béjanin was the sole author for Chapter 1, 2, 3 and 7. The material presented in those chapters was researched under the supervision of Dr. Matteo Mariantoni, and stems in part from personal research and discussions with members of the DQM Lab group and the larger IQC community. A significant part of the material in Chapter 2 was elaborated over many discussions with Matteo Mariantoni. The fabrication recipe shown in Appendix A.1 was developed by C. T. Earnest and C. R. H. McRae.

Chapters 4, 5, and 6, and the corresponding appendices A, B, and C, consist of material written for publication and were authored by multiple people. The contributions for those chapters are detailed below.

Chapter 4

The work of Chapter 4 was initially submitted in May 2016 and was published as: J. H. Béjanin, T. G. McConkey, J. R. Rinehart, C. T. Earnest, C. R. H. McRae, D. Shiri, J. D. Bateman, Y. Rohanizadegan, B. Penava, P. Breul, S. Royak, M. Zapatka, A. G. Fowler, and M. Mariantoni, Three-Dimensional Wiring for Extensible Quantum Computing: The Quantum Socket, [Physical Review Applied](#) **6**, 044010 (2016).

J. H. Béjanin and T. G. McConkey contributed equally to this work.

- J. H. Béjanin: microwave and DC measurements, package holder design, computer platform design and implementation, data analysis (inc. resonator fitting), magnetic measurements, editing
- T. G. McConkey: initial concept and design, microwave and DC measurements, microwave package design, spring characterization, resonator fitting, fabrication, circuit design, simulations (both circuitry and package), editing
- J. R. Rinehart: TDR measurements, computer platform design and implementation, editing
- C. T. Earnest: sample fabrication
- C. R. H. McRae: sample fabrication, resonator fitting, circuit design
- D. Shiri: microwave package design, spring characterization, resonator fitting
- J. D. Bateman: simulations, spring characterization
- Y. Rohanizadegan: laboratory setup
- B. Penava, P. Breul, S. Royak, and M. Zapatka (INGUN): three-dimensional wire design and manufacturing, in-house microwave measurements (not shown)
- A. G. Fowler: initial concept and theory

- M. Mariantoni: initial concept, supervision, project development and management, data measurements and analysis, editing

Chapter 5

The work of Chapter 5 was initially submitted in August 2020 and was published as: J. H. Béjanin, C. T. Earnest, Y. R. Sanders, and M. Mariantoni, Resonant Coupling Parameter Estimation with Superconducting Qubits, [PRX Quantum 2, 040343 \(2021\)](#).

- J. H. Béjanin: measurements, simulations, experimental design and implementation, data analysis, editing
- C. T. Earnest: sample fabrication, measurements, data analysis
- Y. R. Sanders: initial concept and theory
- M. Mariantoni: editing

Chapter 6

The work of Chapter 6 was initially submitted in June 2021 and was published as: J. H. Béjanin, C. T. Earnest, A. S. Sharafeldin, and M. Mariantoni, Interacting Defects Generate Stochastic Fluctuations in Superconducting Qubits, [Physical Review B 104, 094106 \(2021\)](#).

- J. H. Béjanin: measurements, simulations, theory, experimental design and implementation, data analysis, editing
- C. T. Earnest: sample fabrication, measurements
- A. S. Sharafeldin: measurements
- M. Mariantoni: initial concept, theory, editing

Abstract

The topics of this thesis are based on *circuit quantum electrodynamics* (cQED), a theoretical and experimental platform allowing the study of light-matter interaction. This platform is rich both in observable physical phenomena and future practical applications. A “circuit” in cQED may comprise various elements, with the two main types being electromagnetic quantum harmonic oscillators, or resonators, and superconducting Josephson quantum bits, qubits. Because of the relative ease to fabricate and control quantum circuits—especially when compared to the more traditional cavity quantum electrodynamics—cQED has quickly grown in popularity in research labs across the world and is regarded as one of the major contenders for quantum computing.

The advances referred to in the title of this thesis address three significant challenges to practical applications of cQED; they are relevant not only to quantum computing, but also to other applications, such as simulations of physical systems. The first advance is related to *control scalability*. Practical applications require large circuits, and the current approaches used to send control signals to those circuits will not scale indefinitely. A solution to this challenge, the quantum socket, is presented and evaluated in depth. The second advance concerns *calibration*. Any application of cQED requires knowing the precise parameters defining the interactions between the various components of a circuit. Two cutting edge methods for the calibration of interaction parameters are explained and benchmarked; they show a remarkable improvement over existing, inefficient, methods. The third advance involves the physics of *dielectric defects* in the samples on which circuits are fabricated. These unwanted defects are modeled as two-level systems (TLS) that interact with circuit elements such as qubits. Experimental measurements and novel simulations conclusively demonstrate that interactions between TLS are responsible for the stochastic relaxation-time fluctuations observed in superconducting qubits.

Résumé

Les sujets de cette thèse portent sur *l'électrodynamique quantique des circuits* (EDQc), une plateforme théorique et expérimentale permettant l'étude de l'interaction entre la lumière et la matière. Cette plateforme présente de nombreux phénomènes physiques observables et aussi un fort potentiel pour des applications futures. Un "circuit" en EDQc peut contenir divers éléments, les deux principaux étant l'oscillateur harmonique quantique électromagnétique et le bit quantique supraconducteur de Josephson. Les circuits quantiques sont relativement faciles à fabriquer et à contrôler, surtout lorsque comparés à l'électrodynamique quantique en cavité traditionnelle. C'est pourquoi l'EDQc a rapidement gagné en popularité dans les laboratoires du monde entier et est maintenant considérée comme une voie prometteuse vers la fabrication d'un ordinateur quantique.

Les "avancées" mentionnées dans le titre de cette thèse abordent trois difficultés qui entravent les applications pratiques de l'EDQc. Ces avancées sont pertinentes non seulement pour le calcul quantique, mais aussi pour d'autres applications, telles que la simulation de systèmes physiques quantiques, difficile à réaliser avec un ordinateur classique. La première avancée est liée à *l'extensibilité du contrôle*. Les applications pratiques nécessitent des circuits de grande taille et les approches actuelles utilisées pour envoyer des signaux de contrôle à ces circuits ne pourront pas être étendues indéfiniment. Une solution à ce problème, le *connecteur quantique*, est présentée et évaluée de manière approfondie. La deuxième avancée concerne *l'étalonnage*. Toute application de l'EDQc nécessite de connaître les paramètres précis définissant les interactions entre les différentes composantes d'un circuit. Deux techniques de pointe pour le calibrage des paramètres d'interaction sont expliquées et évaluées. Ces techniques montrent une amélioration notable par rapport aux méthodes existantes inefficaces. La troisième avancée concerne la physique des *défauts diélectriques* présents dans les puces fabriquées pour les expériences. Ces défauts indésirables sont modélisés comme des systèmes à deux niveaux (SDN) qui interagissent avec les éléments du circuit tels que les qubits. Des mesures expérimentales et des simulations novatrices démontrent de manière concluante que les interactions entre les SDN sont responsables des fluctuations stochastiques du temps de relaxation, omniprésentes dans les qubits supraconducteurs.

Acknowledgements

When I moved to Waterloo for my graduate studies in the summer of 2013, I certainly did not realize that I would spend so much of my life there! Initially, I had signed up just for a Master’s program, which would have normally taken about two years. After three years of undergrad at McGill, that did not sound like a long time, and I figured I would soon be back with my friends and family in Montreal. Little did I know that setting up a full low-temperature quantum lab would take so long! Fortunately, the people I met, worked with, and became friends with, made my time in Waterloo an amazing and unforgettable experience.

This thesis is the product of the research made over the course of approximately 8 years by myself and the other group members of the Digital Quantum Matter lab. It would not have been possible without the guidance, help, care, support and love of the many people in my life.

I would like to thank Matteo Mariantoni, my supervisor. He has taught me most of what I know about superconducting qubits. Though the innumerable meetings, discussions, blackboard and whiteboard derivations, and lab work sessions that we had—either one-on-one or with the group—Matteo managed to slowly teach me about the intricate art of low-temperature measurements of a circuit QED device.

I would like to thank all the members of the DQM Lab group over the years. When I started the program, the group consisted of one postdoc, Daryoush Shiri, four PhD students, Carolyn Earnest, Corey Rae McRae, Thomas McConkey, and John Rinehart, a research assistant, James Bateman, and me. At this level, experimental research is necessarily a group effort. I remember poring over the many decisions we had to make when designing the fridge wiring with Daryoush and Thomas or listening in over nanofabrication meetings with Corey Rae and Carolyn. None of the research in this thesis would have been a reality without them. I will always remember the moment—immortalized in Fig. 3.12 (a)—when Carolyn and I finally managed to measure a qubit!

I am also thankful for all the help given by Jean-Luc Orgiazzi. When he came back to Waterloo in 2018, we were just started with qubit measurements. Despite the fact he had to graduate ASAP, Jean-Luc still managed to spend many hours with me in the lab explaining the basic qubit measurement concepts. We were also roommates for a short while, and we cooked up a number of delicious meals together!

I would like to thank staff members at the Institute for Quantum Computing and the University of Waterloo: Roberto Romero, from IQC’s technical services, Steve Weiss from the IQC’s IT group, Zhenwen Wang from the Science Technical Services Electronics Shop, and Andrew Dubé and Harmen Vander Heide and the Science Technical Services Machine Shop. The lab setup would not have been completed without their help and work. I would also like to thank Nathan Nelson-Fitzpatrick, Vito Logiudice, and the staff at the Quantum-Nano Fabrication and Characterization Facility (QNFCF) at

the University of Waterloo for their support with device fabrication. All the devices presented in this thesis were fabricated at the QNFCF.

Je tiens à remercier ma famille. Mes parents, Sandrine et Jean-Baptiste, m'ont toujours soutenu dans les choix que j'ai faits tout au long de ma vie. Ils m'ont appris à toujours faire de mon mieux, à me dépasser, et à être ouvert d'esprit. C'est grâce à eux et à leur affection que j'ai eu la chance de pouvoir me concentrer sur mes études, depuis mon entrée à l'Université McGill en 2010. Sans eux, je n'aurais jamais pu terminer cette thèse. Avant de déménager, je vivais en famille à Montréal et j'étais proche de mes deux frères, Julien et Anatole, que je côtoyais tous les jours. Mes frères ont toujours été importants pour moi, ce fut tout un choc de les quitter.

Finally, I must acknowledge that my life in Waterloo would not have been bearable in the slightest without my friends. Whether during travels (to British Columbia, Quebec, India, or Spain), camping trips (to Bruce Peninsula, Algonquin, and French River), road trips (to Jasper and Banff in Alberta), karaoke nights, music shows and festivals, museum outings, pot lucks, dinners, or parties. My friends made the time I spent outside of research worth it. It would have been difficult for me to spend time working without the off-time spent with them, and I want to give them all my gratitude. I want to thank, in historical-ish order: Jean-Philippe MacLean, John Donohue, Mike Mazurek, Corey Rae McRae, Carolyn Earnest, Beenie Benhamron, Dennis Clyne, John Rinehart, Piers Lillystone, Sumit Sijher, Anirudh Krishna, Hammam Quassim, Hina Bandukwala, Thomas Alexander, Jean-Luc Orgiazzi, Vadiraj AM, Jeremy Flannery, Cameron Vickers, Supratik Sarkar, Christopher Warren, Anand Lopez, Nachiket Sherlekar, Manjit Grewal, Rahul Deshpande, Júlia Amorós Binefa, Aditya Jain, Irene Melgarejo Lermas, Nikhil Kotibhaskar, Evan Peters, Laura Edmonds, Kimia Mohammadi, and Claire Phelps, along with the Chainsaw (RIP) and Duke crew, the Spin Ice hockey crew, the Hadamards FC soccer crew and the volleyball crew.

I would like to express special thanks to Claire, who has supported me as I was finishing the writing of this thesis, a time during which I certainly had trouble concentrating on anything else than my thesis.

Along with my co-authors, I would to acknowledge certain people for the research presented in Chapters 4, 5, and 6. For Chapter 4, we acknowledge our fruitful discussions with John M. Martinis and the Martinis Group, Eric Bogatin, and William D. Oliver. We thank the group of Adrian Lupaşcu as well as Patrick Smutek of Plassys-Bestek SAS for their assistance in the deposition of Al films. We thank Dario Mariani and Alexandra V. Bardysheva for their help with a few images and figures. For Chapter 5, we thank Christopher Warren for useful discussions. For Chapter 6, we acknowledge M. I. Dykman for his involvement in the early stages of this work.

The research in this thesis was undertaken in part thanks to funding from the Alfred P. Sloan Foundation, the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada First Research Excellence Fund (CFREF), and the Ministry of Research and Innovation (MRI) of Ontario. I also acknowledge funding from Provost's Entrance Award given by the University of Waterloo. We would like to acknowledge the Canadian Microelectronics Corporation (CMC) Microsystems for the provision of products and services that facilitated this research, including CAD tools and design methodology.

Dedication

*Pour mes grands-parents,
Jacques et Odile Béjanin
Roland Sicsic et Claudine Karsenty-Sicsic*

Table of Contents

List of Figures	xv
List of Tables	xvii
List of Publications	xviii
1 Introduction	1
1.1 Cavity Quantum Electrodynamics	2
1.2 Superconductivity and the Josephson Effect	3
1.3 Quantum Computing and Quantum Information	5
1.4 Thesis Overview	8
2 Circuit Quantum Electrodynamics	9
2.1 Resonators	10
2.1.1 Circuit Equations	10
2.1.2 Resonator Hamiltonian	15
2.2 Superconducting Qubits	17
2.2.1 The Josephson Equations	17
2.2.2 Qubit Hamiltonian	19
2.3 Quantum Circuits	23
2.4 Quantum Gates	28
2.5 Dynamics and Decoherence	31
2.5.1 The Rotating Frame	32
2.5.2 The Driven Qubit Hamiltonian	33
2.5.3 Decoherence	35
3 How to Measure a Qubit	38
3.1 Device Design & Fabrication	38
3.2 Dilution Refrigerator Wiring	42
3.2.1 Thermal Engineering	42
3.2.2 Amplification & Attenuation Chains	46
3.2.3 Eccosorb Tee Filter	49
3.3 Control & Measurement	50
3.3.1 Continuous Wave Measurements with a Network Analyzer	51

3.3.2	Time-Domain Measurements	53
3.3.3	Qubit Readout	55
3.4	Initial Characterization Experiments	60
3.5	Randomized Benchmarking	66
4	The Quantum Socket	71
4.1	Introduction	72
4.2	Design	73
4.2.1	Three-dimensional wires	74
4.2.2	Microwave package	77
4.2.3	Package holder	78
4.2.4	Microwave simulations	79
4.3	Implementation	81
4.3.1	Magnetic properties	82
4.3.2	Thermal properties	83
4.3.3	Spring characterization	84
4.3.4	Alignment	84
4.4	Characterization	88
4.4.1	Four-point measurements	89
4.4.2	Two-port scattering parameters	90
4.4.3	Time-domain reflectometry	97
4.4.4	Signal crosstalk	100
4.5	Applications to Superconducting Resonators	102
4.6	Conclusions	107
5	Resonant Coupling Parameter Estimation	111
5.1	Introduction	112
5.2	Qubit Calibration in Frequency-Tunable Architectures	113
5.3	Traditional Swap Spectroscopy	115
5.4	Offline Octave Sampling	117
5.4.1	Theoretical Method	117
5.4.2	Experimental Results	121
5.5	Online Bayesian Learning Algorithm	124
5.5.1	Theoretical Method	124
5.5.2	Experimental Results	127
5.6	Discussion	130
5.7	Conclusions	132
6	Time Fluctuations	133
6.1	Introduction	133

6.2	Theory	135
6.2.1	Physical Models of TLSs	135
6.2.2	Qubit-Q-TLS Interaction	137
6.2.3	Q-TLS-T-TLS Interaction	139
6.2.4	Qubit Stochastic Fluctuations	140
6.3	Methods	140
6.3.1	Experiments	140
6.3.2	Simulations	141
6.4	Results	144
6.5	Discussion	147
6.5.1	Physical Characteristics of a T-TLS	147
6.5.2	Density of TLSs	149
6.5.3	On the Interpretation of the Allan Deviation and Power Spectral Density	150
6.6	Conclusions	152
7	Conclusion	154
	Bibliography	156
A	Appendix for Chapter 3	184
A.1	Qubit Chip Fabrication Recipe	184
B	Appendix for Chapter 4	188
B.1	Wire Compression	188
B.2	Magnetism	189
B.3	Thermal Conductance	191
B.4	Thermo-Mechanical Tests	193
B.5	Alignment Errors	195
B.6	Sample Fabrication	197
C	Appendix for Chapter 5	198
C.1	Qubit Flux Pulse	198
C.2	Details on Octave Analysis	199
C.3	Particle Distribution Generation	200
C.4	Bayesian Resampling Procedure	202
D	Appendix for Chapter 6	203
D.1	Experimental Details	203
D.2	Long-Time Stability	204
D.3	Qubit Electric Field	205

List of Figures

1.1	Schematic representation of a cavity QED experiment	2
1.2	Schematic representation of a Josephson junction	5
2.1	Circuit diagram of a parallel RLC resonator	10
2.2	S_{21} parameter for a capacitively coupled resonator	14
2.3	Circuit diagram of a tunable transmon qubit	18
2.4	First four eigenfunctions of the transmon Hamiltonian for $E_J/E_C = 100$	22
2.5	Chip design with two Xmon qubits connected by a resonator	23
2.6	Circuit diagram and corresponding graph for two qubits connected by a resonator	24
2.7	The Bloch Sphere	29
3.1	Complete mask design for a two-qubit device	39
3.2	Images of a fabricated sample	41
3.3	Picture of the open dilution refrigerator	43
3.4	Steady-state occupation probability of the first excited state of a qubit	44
3.5	Dilution refrigerator wiring schematics	47
3.6	Fabrication of the Eccosorb tee filter	50
3.7	Qubit spectroscopy of the same sample before and after installation of infrared filtering	51
3.8	Readout transmitter and receiver setup	56
3.9	Qubit readout characterization	57
3.10	Qubit single shot measurement IQ calibration	59
3.11	Punchout measurement	61
3.12	Rabi and Ramsey experiments	62
3.13	T_1 and T_2 experiments	64
3.14	Full Hamiltonian calibration spectroscopy	65
3.15	Single-qubit randomized benchmarking	68
3.16	Randomized benchmarking over time	69
4.1	Computer-aided designs of the three-dimensional wire, microwave package, and package holder	75
4.2	Two-dimensional cut view of the three-dimensional wire, contact pad, and screw-in micro connector	76

4.3	Numerical simulations of the electric field distribution	80
4.4	Images of the quantum socket as implemented	82
4.5	Micro images used to evaluate the alignment procedure of the three-dimensional wires	85
4.6	I-V characteristic curve for R_{ig}	89
4.7	S-parameter measurements and simulations for a three-dimensional wire at room temperature	91
4.8	Microwave characterization setup	92
4.9	S-parameter measurements for the Au sample	93
4.10	S-parameter measurements for the Ag sample	94
4.11	Quantum socket microwave parameters	96
4.12	TDR measurements for the Au sample at room temperature and 77 K	98
4.13	TDR measurements for the Ag sample at room temperature	99
4.14	Quantum socket signal crosstalk	101
4.15	Macro photograph of an Al chip on Si substrate mounted in a sample holder with grounding washer	103
4.16	Measurements of Al on Si resonators	104
4.17	Measurement of Q_i as a function of $\langle n_{ph} \rangle$ for one of the Al on sapphire resonators	105
4.18	Extensible quantum computing architecture	108
5.1	Pulse sequence for a swap spectroscopy experiment	115
5.2	Swap spectra for two frequency ranges	116
5.3	Offline octave sampling	119
5.4	Illustration of a simulated iteration of the online particle filter algorithm	125
5.5	Performance of the online estimation algorithm over 1000 runs	129
6.1	Experimental and simulated spectrotemporal charts of T_1 vs. f_q and t	145
6.2	Three spectral-diffusion patterns	146
6.3	Comparison between the statistical analyses of two simulated times series	151
B.1	Side and top view images of a three-dimensional wire after 200000 compression cycles.	194
B.2	Micro images showing three-dimensional wire alignment errors	195
C.1	Peaks found in the highest octave of Fig. 5.3 (b)	200
C.2	Initial particle distribution generated from the octave data in Fig. 5.3 (b)	201
D.1	Scatter plot of T_1 vs. f_q and t	203
D.2	Time series of T_1 vs. t showing the relative time between measurements	205
D.3	Qubit electric field $\ \vec{E}_q\ $ for $\phi_0 = 1$ V vs. width x	206

List of Tables

2.1	Classification of the integer order solutions to Mathieu's equation	21
3.1	Qubit-resonator device parameters	66
3.2	Single-qubit Cliffords	67
4.1	Simulation results for the first three box modes of the lower cavity inside the assembled microwave package	79
4.2	DC resistance tests	87
4.3	Resonator parameters	106
5.1	Resonance modes detected after analyzing the octave sampling data . . .	123
5.2	Estimated parameters for the three resonance modes detected after running the online Bayesian learning algorithm	128
6.1	Experimental parameters for the three datasets	141
6.2	T-TLS and Q-TLS parameters used in the simulations of Fig. 6.2.	148
6.3	Time-series simulation parameters used in Fig. 6.3	150
B.1	Chemical composition of the two main materials used in the three-dimensional wires	189
B.2	ZGC calibration	190
B.3	Magnetic field measurements of the materials used for the main compo- nents of the quantum socket	191
B.4	Parameters used in the estimate of the heat transfer rate	192
B.5	Thermo-mechanical tests on hardened BeCu springs	193
C.1	Frequency and coupling parameters statistics of the particle distributions generated for RM1, RM2 and RM3	201

List of Publications

1. J. H. Béjanin, T. G. McConkey, J. R. Rinehart, C. T. Earnest, C. R. H. McRae, D. Shiri, J. D. Bateman, Y. Rohanizadegan, B. Penava, P. Breul, S. Royak, M. Zapatka, A. G. Fowler, and M. Mariantoni, Three-Dimensional Wiring for Extensible Quantum Computing: The Quantum Socket, [Physical Review Applied](#) **6**, 044010 (2016).
2. C. R. H. McRae, J. H. Béjanin, Z. Pagel, A. O. Abdallah, T. G. McConkey, C. T. Earnest, J. R. Rinehart, and M. Mariantoni, Thermocompression Bonding Technology for Multilayer Superconducting Quantum Circuits, [Applied Physics Letters](#) **111**, 123501 (2017).
3. C. T. Earnest, J. H. Béjanin, T. G. McConkey, E. A. Peters, A. Korinek, H. Yuan, and M. Mariantoni, Substrate Surface Engineering for High-Quality Silicon/Aluminum Superconducting Resonators, [Superconductor Science and Technology](#) **31**, 125013 (2018).
4. T. G. McConkey, J. H. Béjanin, C. T. Earnest, C. R. H. McRae, Z. Pagel, J. R. Rinehart, and M. Mariantoni, Mitigating Leakage Errors Due to Cavity Modes in a Superconducting Quantum Computer, [Quantum Science and Technology](#) **3**, 034004 (2018).
5. C. R. H. McRae, J. H. Béjanin, C. T. Earnest, T. G. McConkey, J. R. Rinehart, C. Deimert, J. P. Thomas, Z. R. Wasilewski, and M. Mariantoni, Thin Film Metrology and Microwave Loss Characterization of Indium and Aluminum/Indium Superconducting Planar Resonators, [Journal of Applied Physics](#) **123**, 205304 (2018).
6. M. S. Moeed, C. T. Earnest, J. H. Béjanin, A. S. Sharafeldin, and M. Mariantoni, Improving the Time Stability of Superconducting Planar Resonators, [MRS Advances](#) **4**, 2201 (2019).
7. J. H. Béjanin, C. T. Earnest, Y. R. Sanders, and M. Mariantoni, Resonant Coupling Parameter Estimation with Superconducting Qubits, [PRX Quantum](#) **2**, 040343 (2021).
8. J. H. Béjanin, C. T. Earnest, A. S. Sharafeldin, and M. Mariantoni, Interacting Defects Generate Stochastic Fluctuations in Superconducting Qubits, [Physical Review B](#) **104**, 094106 (2021).



The author of this thesis in October 2013, below the dilution refrigerator that would occupy much of his time for the next few years...

Introduction

As a discipline, the study and knowledge of the natural world has its origins in Archaic Greece, between 650 and 480 BCE. The word “physics” itself traces its roots back to the Ancient Greek word for “nature”: φύσις (phýsis). It was there that some of the first natural philosophers started seeking explanations for natural phenomena in nature itself, rather than in supernatural beliefs. One of those early philosophers was Democritus of Abdera, probably born between 470 and 460 BCE. Democritus was one of the first known proponents of *atomism*, the idea that matter was constituted of indivisible discrete pieces, or *atoms*, from the word ἄτομος (átomos) meaning “indivisible”. In-between the atoms was simply void that one could cut through in order to separate chunks of matter. The argument of Democritus was surprisingly simple: If matter was infinitely divisible, then we could carry out this division (which would of course take infinite time), and we would be left with infinitely many parts. However, those parts of matter should also be divisible, thereby contradicting the hypothesis that the initial chunk was divided infinitely. Therefore, Democritus concluded, there must be a point at which we must have to stop dividing matter: the atom [1].

As we now know, Democritus ended up right about the atom (or nearly right—let’s not forget quarks!) In some ways, the indivisibility of the atom is an idea precursor to quantum mechanics¹. In quantum mechanics, particles are discrete, certainly², but many other, less obvious, physical quantities—like light and energy—come in discrete chunks as well.

This thesis is concerned with the physics resulting from the interaction of light and matter in small numbers of discrete quanta of energy. Reaching this low-energy *regime* can be accomplished in a few ways, with the traditional approach being cavity quantum electrodynamics experiments, which deal with natural atoms and their absorption of

¹The word “quantum” comes from the Latin interrogative adjective meaning “how much”.

²And there are many of them! See the standard model.




Figure 1.1: Schematic representation of a cavity QED experiment. An atom (red, center) is sent in an electromagnetic cavity formed by two mirrors (gray, on each side). The cavity is populated with a photon field (blue). The atom interacts with the cavity (yellow) and can be excited externally by a laser (purple). Both atom and cavity also slowly decay into the environment, at different rates γ and κ (green).

light quanta. A second approach is to use superconductivity to create *artificial* atoms: quantum states that are *condensates* of many particles, but, just like Democritus' atoms, behave as a single, indivisible, one.

1.1 Cavity Quantum Electrodynamics

Building interesting quantum systems that we can manipulate is no small endeavor. Historically, most quantum experiments were implemented within the context of cavity quantum electrodynamics, or cavity QED. Such a platform let physicists explore light-matter interactions by constraining light in a cavity (a region of space enclosed by mirrors), and observing its interactions with matter, usually individual atoms. The cavity, in particular, turned out to be a very useful experimental apparatus. Indeed, free space is a very difficult environment to work in, electromagnetically-speaking, for the simple reason that it allows *all* light propagation. In the classical world, this is useful, permitting light emitted from the Sun to bounce around freely, lighting up our world and allowing us to see. But for quantum mechanics, this freedom restricts the amount of phenomena that would otherwise occur. Atoms “excited” to a particular energy level can usually release some of that energy as a photon at a particular frequency $\omega_{nm} = (E_m - E_n)/\hbar$, which corresponds to the energy difference between the start and end levels m and n . If an excited atom is sitting around in free space, a photon can thus easily get emitted *spontaneously* and fly away. This is problematic if, for example, you are trying to entangle two atoms in a Bell state $|00\rangle + |11\rangle$; the state will not be stable for very long, since the atoms will quickly lose their energy and fall back to the ground state $|00\rangle$.

Enter cavities: A cavity, which can generally be fabricated with two parallel mirrors,

is different from free space in that it only allows photons with a very restricted set of frequencies to exist within it. If an atom is excited by a quantum of energy corresponding to a frequency that is *not* supported by the cavity, there is nowhere for a hypothetical photon to go if emitted. Fascinatingly, this restriction on the photon actually prevents the atom's excitation from being converted into light in the first place!

The theoretical ground for such effects is historically attributed to Purcell. In 1946, he was the first to calculate that a single-electron atom coupled to a near-resonance³ lossy cavity mode would have its decay rate enhanced relative to an atom in free space [2]. Conversely, if the atom and cavity are far off-resonance, the decay rate can be reduced. The first experiments on suppressed spontaneous emission were performed by Drexhage, Kuhn and Schäfer with fluorescing organic dye [3].

A cavity also enhances the strength of the interaction with the atom by increasing the amplitude of the electromagnetic field. Placing atoms within cavities allowed researchers to start experimenting with strong atom–light interactions, with long atom lifetime. Early experiments in this direction were started in the 1980s. 20 years later, it became possible to observe individual excitations swap back and forth between the electromagnetic field and the atom [4] and create entangled quantum states [5]. For pioneering this kind of work, Serge Haroche (along with David Wineland, who worked on similar light–matter effects) won the 2012 Nobel Prize in Physics [6].

1.2 Superconductivity and the Josephson Effect

The phenomenon of superconductivity was first discovered in 1911 by Dutch physicist Heike Kamerlingh Onnes. That discovery itself relied on a second discovery, the production of liquid Helium, which Onnes had achieved three years earlier, in 1908. Thanks to the low temperature of liquid Helium, Onnes was able to cool down metals to very low temperature, which let him observe the transition of Mercury to a state of near-zero resistance at 4.2 K [7]. A few years later, experiments with persistent currents showed that the resistivity of superconductors was so small that currents could not be observed to decay over hours. Onnes received the 1913 Nobel Prize in Physics “for his investigations on the properties of matter at low temperatures” [8]. Of course, perfect conductivity is only one of the features of superconductors, the second is the Meissner effect, discovered in 1933 by the German scientists Walther Meissner and Robert Ochsenfeld [9]. The effect describes the expulsion of magnetic fields from a superconductor as it transitions into the superconducting state. This effect does not result simply from the perfect conductivity. As such, the Meissner effect was a clear sign that superconductivity was a distinct phase of matter that required the development of new theories for a proper explanation.

³This means that the atom and cavity have nearly the same frequency

By most accounts, that theory is the BCS theory, named after John Bardeen, Leon Cooper, and John Robert Schrieffer, the researchers who proposed it in 1957 [10, 11]. The BCS theory relied on the idea of bound electron states, where two electrons—which are fermions—combine to form a new particle: a *Cooper pair*—a boson [12]. The BCS theory is therefore a microscopic theory of superconductivity, and though it is not a universal one, meaning that certain superconductors do not behave according to BCS theory, it still enabled many predictions, in particular, the Josephson effect. For their theory, Bardeen, Cooper and Schrieffer received the 1972 Nobel Prize in Physics⁴.

Just a few years after the BCS theory was introduced, in 1962, Brian Josephson predicted what is now called the Josephson effect [14]. The Josephson effect refers to the quantum tunneling of superconducting electrons—Cooper pairs—across a thin non-superconducting barrier. While it was known at the time that normal electrons could tunnel across a barrier, this had not been conclusively observed with Cooper pairs. In all likelihood, the electrical noise in high-resistance samples created currents in excess of the critical current⁵, thereby preventing observation of supercurrents [15]. Once this was understood, Anderson and Rowell quickly conducted experiments with low-resistance samples that demonstrated probable tunneling supercurrent [16]. Further experiments soon confirmed the discovery by observing the magnetic field and AC dependence of this tunneling supercurrent [17, 18]. The significance of Josephson’s prediction was quickly recognized: In 1973, he received the Nobel Prize in Physics, shared with Leo Esaki and Ivar Giaever [19].

Almost immediately after the Josephson effect was measured, researchers realized that creating a loop comprising two Josephson junctions would result in a highly sensitive magnetometer [20]. The DC SQUID, which stands for “superconducting quantum interference device,” exploits the Josephson effect in combination with superconducting flux quantization to create a flux-to-voltage converter, thereby allowing measurement of minuscule changes of the magnetic flux threading the loop [21].

A second notable feature of the Josephson junction concerns its dynamics. While an isolated superconductor in the ground state has a fully defined condensate wave function, a Josephson junction possesses a macroscopic degree of freedom: the wave function phase difference across the barrier $\varphi = \phi_R - \phi_L$, where ϕ_R and ϕ_L are the order parameter phases⁶. This variable is macroscopic because it determines the relative phase between every Cooper pair in the left and right superconductors, i.e. a “macroscopic” number of Cooper pairs.

⁴This was actually Bardeen’s *second* Nobel Prize! He had already received one in 1956 for the invention of the transistor [13].

⁵Superconductors and Josephson junctions are only superconducting if the current flowing through them is less than the critical current.

⁶The order parameter is a concept used to solve superconductivity phase transition problems in the Ginzburg-Landau theory. See Chapter 4 of Ref. [22] for more details.




Figure 1.2: Schematic representation of a Josephson junction. A Josephson junction consists of a superconductor (gray) interrupted by an insulating barrier (blue). If the barrier is thin enough, Cooper pairs may tunnel across, resulting in a superconducting tunneling current. The behavior of the tunneling current is determined by a single *macroscopic* variable: the wave function phase difference $\varphi = \phi_R - \phi_L$.

The work more directly relevant to this thesis was started by demonstrating various macroscopic quantum effects with Josephson junctions. In 1981, Voss and Webb demonstrated macroscopic quantum tunneling (MQT) [23], which is unrelated to the tunneling of Cooper pairs through the junction and instead refers to the tunneling of a φ “particle” out of its potential well⁷. The existence of discrete energy levels within the potential well was confirmed in 1985 by Martinis, Devoret and Clarke [24]. These results ignited a flurry of experiments meant to show the viability of using Josephson junctions to build an “artificial atom”. Such a device would be in many ways equivalent to a real atom, as those used for cavity QED experiments. Unlike traditional atoms, however, an artificial one was expected to provide many more experimental “knobs,” for example, controllable energy and couplings to electromagnetic modes.

1.3 Quantum Computing and Quantum Information

A significant reason for the current popularity of superconducting circuit QED is, of course, its potential as a platform for quantum information and quantum computing. Building a *universal quantum computer* [25, 26, 27, 28, 29, 30] will make it possible to execute quantum algorithms [31], which would have profound implications on science and society. As a field, quantum computing began in the early 1980s, when researchers started thinking about performing classical computation using quantum mechanics. The first articles about quantum computation showed that a quantum-mechanical Turing machine would in fact not perform worse than a classical computer, and that classical computation performed with such a “quantum computer” was logically reversible [32, 33, 34].

It was in 1982 that Feynman introduced the idea that a quantum computer could actually have a computational advantage over classical computers at performing certain

⁷A φ particle is represented by a junction wavefunction that is well-defined in the φ coordinate.

tasks. The example he gave was simple: A quantum computer could be used to simulate a quantum mechanical system [35]. Feynman's idea spurred research in this direction, with David Deutsch first formulating the workings of a universal quantum computer in 1985, followed by further work on the now-ubiquitous circuit model of quantum computation [36, 37, 38, 39]. Deutsch was also the first to demonstrate a quantum algorithm providing a provable speedup over a classical computer, the Deutsch-Jozsa algorithm [40].

To this day, however, Shor's algorithm for integer factorization remains the most important quantum algorithm [41, 42]. Though the Deutsch-Jozsa algorithm provided a speedup, its potential for real world application was very limited. Shor's algorithm, on the other hand, could break public key cryptography based on the RSA system, which relies on the hardness of factorization and had been recently developed [43]. At the same time, Shor introduced an algorithm to solve the discrete logarithm problem, which was used for Diffie-Hellman key exchange [44]. By exploiting the properties of the quantum Fourier transform [45], both of Shor's algorithms provide an exponential speedup over the best classical algorithms. Although no proof exists, it is widely believed that there are no efficient classical solutions to integer factorization or the discrete log problem [26]. For that reason, Shor's result attracted considerable interest and is still one of the main sources of motivation for the realization of a quantum computer.

Soon after, a second important result was published, and is now known as Grover's search algorithm [46, 47]. Unlike Shor's algorithm, the search algorithm was proven to be faster than any classical algorithm. The speedup, however, is only quadratic. While a classical algorithm will require $\mathcal{O}(N)$ steps to search through a list containing N items, a quantum computer can search through the list in only $\mathcal{O}(\sqrt{N})$ steps. The relatively smaller advantage of Grover's algorithm is counterbalanced by the wide applicability it has to many problems [48].

The language of quantum computing and quantum algorithms is provided by the field of quantum information, which deals with the processing, storage, and transmission of quantum bits, or qubits. The material in this thesis does not depend explicitly on those concepts, and we instead refer readers to literature dedicated to those topics, such as Refs. [26, 28, 49, 50]. Within quantum information, the topic most relevant to this thesis is probably error correction. Indeed, it was quickly realized, once the first physical qubits started appearing in the 1990s, that many sources of errors would prevent the algorithms discussed above from being implemented. A similar problem affects classical information—bits—and there the solution is to encode bits with some amount of redundancy, a concept called error correction. This classical solution, however, could not be applied directly to quantum information since qubits are fundamentally analogue (i.e., continuous) quantities, and, even more critical, qubit states cannot be copied [51]. Fortunately, after a few years it was discovered that quantum

error correction was possible [52, 53, 54, 55]. Today, the most promising avenue for scalable quantum error correction is provided by *surface codes*, a family of topological error correcting codes based on the stabilizer formalism [56, 57, 58, 59, 60]. Although quantum computers will be lacking *fault tolerance* for the near future, a variety of potential applications have been devised for such systems. Examples include solving optimization problems [61, 62], machine learning [63, 64, 65, 66], materials science and chemistry [67, 68, 69].

Quantum computing architectures based on different types of physical qubits have been investigated since the late 1990s [29]. This thesis is focused on superconducting devices, and will not cover the many other approaches to experimental quantum information. In brief, those other architectures may be implemented using photons [70, 71], trapped ions [72], and spins in molecules [73], quantum dots [74, 75, 76, 77, 78], and silicon [77, 79]. For an overview of the realization of various quantum algorithms on various types of experimental platforms, see Ref. [31] and the references within.

The very first superconducting qubit was the *charge qubit*, which was made out of a voltage-biased Josephson junction. It was developed almost at the same time by the Devoret group, in France, and the Nakamura group, in Japan [80, 81, 82]. The quantum mechanical description of the charge qubit is given by the Cooper-pair box Hamiltonian. The charge qubit managed to show that it was possible to create a quantum two-level system and control it. Unfortunately, immature experimental techniques and pervasive charge noise reduced the coherence time to only tens of nanoseconds. Soon after, the Devoret group, now at Yale University, improved the charge qubit by replacing the single junction with two junctions in a loop, allowing for tunability of the qubit transition frequency and operation at the “sweet spot,” improving coherence times [83, 84]. The second improvement was the development of an efficient readout scheme, though it was destructive to the qubit state. Two other qubit types were developed around the same time: the *flux qubit* [85, 86, 87, 88] and the *phase qubit* [89]; together with the charge qubit, they form the original superconducting qubit trifecta [90, 91].

Today, better qubit circuits, setup designs and fabrication improvements have pushed the coherence times of superconducting qubits to 100 μs and above [92, 93, 94], and improved quantum-limited amplifier designs allow for high-fidelity nondemolition readout [95]. Large companies have also joined the race to increase the number of qubits, with the goal of improving the computational power of quantum devices, and eventually creating a fault-tolerant qubit. The latest major milestone reached demonstrated a quantum advantage over classical computers, though with computational problems specifically designed for that purpose [96, 97]. As a result of these advances, quantum computing has transitioned into the so-called noisy intermediate scale quantum (NISQ) era [98]. The progress in the last 20 years has therefore been substantial, and we suggest a few recent reviews to get a sense of it [91, 99, 100, 101].

1.4 Thesis Overview

This thesis presents three advances, or results, in the field of circuit quantum electrodynamics (circuit QED, or cQED, if you really want to save space). For this purpose, we first explain in Chapter 2 the theory foundational to circuit QED. This includes the two basic circuit elements, the superconducting resonator and the superconducting qubit, and the interaction between them. We also briefly cover quantum gates, which is how qubits are controlled in experiments, and basic dynamics and decoherence concepts useful for the later chapters.

Chapter 3 switches the tone to explain in depth how to realize cQED experiments. No more theory (or very little)! The first part covers maybe the most important aspect of cQED experiments: the setup. Often overlooked in articles, the setup, more than anything else, will determine the quality of the results. We attempt to touch as many topics as possible: sample design, thermalization, wiring, filtering, control, readout, and instrumentation. The rest of the chapter demonstrates the “bring-up” of a device, including characterization experiments to perform after first cooling down, and evaluation of the primary circuit parameters. The end of the chapter covers more advanced experiments, including a demonstration of randomized benchmarking.

The following chapters, Chapter 4, 5, and 6 also present experiments that can be done with a circuit QED setup. Chapter 4, which presents the *quantum socket* in conjunction with superconducting resonator measurements, is in many ways an extension of the setup sections. The quantum socket is a fully vertical wiring system for qubit devices introduced in 2016. Such an approach is necessary for scalable superconducting quantum computing. Even today, we are one of very few groups with a true 3D wiring setup in the lab.

Chapter 5 addresses important issues in cQED involving *parameter estimation*. As the size of circuit increases, it becomes important to properly and efficiently calibrate the various system parameters. The two algorithms presented are able to efficiently estimate coupling parameters, e.g., between two qubits, for gate calibration, or between a qubit and a noisy defect, for coherence time improvements.

Chapter 6 explores the basic physics of a pervasive problem in superconducting devices: the *time-fluctuations* of qubit relaxation times. Because these fluctuations can cause T_1 to vary by an order of magnitude over days, it is critical to understand their origin and mechanism. The material of this chapter presents measurements and simulations that strongly support the mechanism of these fluctuations to be due to interacting two-level systems.

Finally, Chapter 7 concludes by summarizing the advances in this thesis and placing them in the broader context of the current challenges.

Circuit Quantum Electrodynamics

The field of circuit quantum electrodynamics was born shortly after the demonstration of the first superconducting qubits, then called artificial atoms, or artificial two-level systems [80, 81, 87, 89, 83]. Indeed, it was quickly realized that the relative ease of fabrication of superconducting devices made them convenient to embed in complex circuits comprising many components. In particular, the first step was to demonstrate a coupling between a qubit and a superconducting cavity, or resonator, thereby replicating the setup of cavity QED experiments entirely on a chip [102, 103]. The field quickly expanded, driven by the interest in quantum computing. Coherence times improved [104, 105, 92], and the number of qubits and resonators on a chip increased [106, 107, 108, 109]. Today's state of the art quantum circuits are made with dozens of qubits, and they pose incredible challenges to fabricate, measure, calibrate and characterize [109, 97].

In this chapter, we discuss the most important elements of circuit QED experiments, namely the superconducting resonators and qubits. The circuit equations of parallel RLC resonators are derived in Sec. 2.1. We also cover the transmission line resonators, which are used in experiments and may be modelled similarly. The quantum Hamiltonian of the resonator is attained by a method relying on canonical quantization of conjugate variables. In Section 2.2 we then cover the Josephson equations, which we use to obtain the qubit Hamiltonian. Section 2.3 then explains how multiple elements can be wired together to form a proper quantum circuit. The derivation of the full circuit Hamiltonian hinges on a systematic formulation of the circuit equations as a directed graph, in addition to the quantization concepts used previously for the resonator and qubit. Finally, Sections 2.4 and 2.5 cover basic concepts underlying gate-based operations and the dynamics of quantum circuits.

The material of Sec. 2.3 was developed in collaboration with M. Mariani. In particular, the idea of using the classical instantaneous energy of the circuit via Tellegen's




Figure 2.1: Circuit diagram of a parallel RLC resonator.

theorem to derive the quantum Hamiltonian was initially conceived by M. Mariani.

2.1 Resonators

Superconducting resonators are the circuit QED equivalent of optical cavities [110]. Instead of trapping photons, they trap electrons (though, of course, both result in confined electromagnetic waves). Just like cavities, resonators in quantum mechanics are represented by quantum harmonic oscillators. Resonators are extremely important in circuit QED: they are used for measuring qubits. Ideally, they must be of very high quality to avoid losing information to dissipation.

Note that this section is a simple overview focused on the applications of resonators to the experiments in this thesis. This field of research is vast and there are many kinds of circuits, coupling methods and applications for superconducting resonators. For more details, see Refs. [111, 112, 113] and the references therein.

This section also does not cover the fabrication of resonators. For all the experiments covered in this thesis, the resonators are $\lambda/4$ aluminum coplanar waveguides (CPW). More details can be found in Earnest *et al.* [114].

2.1.1 Circuit Equations

The circuit of Fig. 2.1 represents a lumped-element parallel RLC circuit, connected to a voltage source via a coupling capacitance C_c . In superconducting experiments, distributed-element—rather than lumped-element—resonators are generally used because of their ease of fabrication and quality. Distributed resonators are made from a finite section of transmission line, the length of which determines the resonance frequency. The correspondence between the two kinds of circuits is given later in this section.

We initially focus on the RLC section of the circuit. We can immediately calculate the resonance frequency ω_i , for which the input impedance $Z_r(\omega)$ is purely real. The impedance is [115]

$$Z_r(\omega) = \left(\frac{1}{R} + \frac{1}{i\omega L} + i\omega C \right)^{-1}. \quad (2.1)$$

Since the resistive part of the impedance is real, the condition for resonance occurs when the inductive and capacitive parts sum to zero:

$$\frac{1}{i\omega_i L} + i\omega_i C = 0 \quad (2.2)$$

$$\omega_i^2 = \frac{1}{LC} \quad (2.3)$$

and therefore $\omega_i = 1/\sqrt{LC}$.

A second quantity we may calculate is the internal quality factor. The quality factor of a resonant circuit is defined as 2π times the number of cycles it can sustain after being energized. It is calculated by dividing the total energy stored by the amount of power dissipated, and multiplying by the angular resonance frequency to convert the resulting time to a number. This quantity is therefore calculated at resonance.

$$Q_i = \omega_i \frac{E_{\text{total}}}{P_{\text{dissipated}}} \quad (2.4)$$

On average, the energy in an RLC circuit is stored in equal amounts in the capacitor and the inductor. The resistor only dissipates energy as heat. For a root mean square (rms) voltage V across the resonator, and on resonance, the *total* energy is

$$E_{\text{total}} = CV^2 = LI_L^2 \quad (2.5)$$

where I_L is the rms current flowing through the inductor.

Given the rms voltage, the average power dissipated by the resistor is found by using Ohm's law to solve for I_R , the rms current flowing through the resistor. Then,

$$P_{\text{dissipated}} = I_R V = \frac{V^2}{R}. \quad (2.6)$$

Therefore, the internal quality factor is

$$Q_i = \omega_i \frac{R}{V^2} CV^2 = \omega_i RC. \quad (2.7)$$

For a parallel RLC circuit, increasing R improves the quality factor.

As mentioned above, distributed element resonators are often made of a section of transmission line with $50\ \Omega$ characteristic impedance. For example, a section of transmission line with length ℓ open at one end and grounded at the other resonates at the fundamental frequency ω_i corresponding to a wavelength $\lambda_i = 4\ell$. The circuit diagram for such a “quarter-wave” resonator is very similar to Fig. 2.1, with the boxed RLC network simply replaced by a piece of transmission line. The impedance of such a line is

$$Z_r(\omega) = Z_0 \tanh(\alpha\ell + i\beta\ell) \quad (2.8)$$

where Z_0 is the characteristic impedance, α is the attenuation constant, $\beta = \omega/v_p$ is the phase constant and $v_p = c/\sqrt{\epsilon_r}$ is the phase velocity which must be determined according to the effective relative permittivity ϵ_r of the transmission line¹.

Parallel RLC resonators behave very similar to $\lambda/4$ resonators. In fact, if we approximate either Eq. (2.1) or Eq. (2.8) near ω_i , i.e., for small $\Delta\omega_i = \omega - \omega_i$, we obtain an expression of the form [115]

$$Z_r(\omega) \approx Z_0 (\pi/4Q_i + i\pi\Delta\omega_i/2\omega_i)^{-1}. \quad (2.9)$$

Comparing terms gives the correspondence between the RLC and $\lambda/4$ resonator:

$$R = Z_0/(\alpha\ell) \quad (2.10)$$

$$C = \pi/(4\omega_i Z_0) \quad (2.11)$$

$$L = 1/(\omega_i^2 C) \quad (2.12)$$

$$\omega_i = 2\pi v_p/\lambda_i \quad (2.13)$$

$$Q_i = \beta/(2\alpha\ell). \quad (2.14)$$

It is thus clear that for a transmission line resonator, decreasing the attenuation constant improves the quality factor. One notable difference between the RLC and $\lambda/4$ resonator is that the latter supports infinitely many resonance modes at the frequencies $\omega_n = (2n - 1)\omega_i$ for $n \in 1, 2, 3, \dots$. When deriving the Hamiltonian (Subsection 2.1.2) we usually ignore the higher modes under the assumption they do not interact with other circuit elements. This assumption is not always valid, however, and it is important to keep the other modes in mind.

In an experiment, a resonator cannot be isolated as we have assumed above. It must be connected to a voltage source, which provides the energy to drive the resonance, and a load, across which a voltage may be measured. Figure 2.1 shows a simple scheme making use of a coupling capacitor C_c to connect the resonator in parallel with a load Z_0 . We also include the source-side impedance Z_0 . For a distributed element circuit, Z_0 is the characteristic impedance.

¹For a CPW transmission line, the effective relative permittivity can be calculated analytically given dimensions and substrate type. See, e.g., Chapter 2 in Simons [116].

The use of a coupling circuit to measure a resonator means that ω_i and Q_i cannot be directly observed. This is because the coupling circuit *loads* the resonator, modifying the resonance frequency and the quality factor. The input impedance of the resonator in series with the capacitor is

$$Z_{\text{in}}(\omega) = \frac{1}{i\omega C_c} + Z_r(\omega). \quad (2.15)$$

Finding the coupled resonance frequency ω_0 once again requires finding the frequency at which the imaginary part of the impedance is zero, i.e., the solution to $\text{Im}[Z_{\text{in}}(\omega_0)] = 0$. This calculation is done most conveniently via a root-finding computer program.

The total quality factor Q may be expressed as a combination of Q_i and the coupling quality factor Q_c as

$$Q = \left(\frac{1}{Q_i} + \frac{1}{Q_c} \right)^{-1} \quad (2.16)$$

where $Q_c \approx 2\pi/(2\omega_0 Z_0 C_c)^2$ for $C_c \ll C$.

When measuring a resonator such as the one of Fig. 2.1, the quantity of interest is the S_{21} S-parameter:

$$S_{21} = \frac{2V_1}{V_s}. \quad (2.17)$$

The voltage on the load impedance is given by current division:

$$V_1 = I_1 Z_0 \quad (2.18)$$

$$= I_s \left(\frac{1}{Z_0} + \frac{1}{Z_{\text{in}}} \right)^{-1} \quad (2.19)$$

the total source voltage is

$$V_s = I_s \left(Z_0 + \left(\frac{1}{Z_0} + \frac{1}{Z_{\text{in}}} \right)^{-1} \right) \quad (2.20)$$

and therefore

$$S_{21} = \frac{2V_1}{V_s} \quad (2.21)$$

$$= 2 \left(Z_0 \left(\frac{1}{Z_0} + \frac{1}{Z_{\text{in}}} \right) + 1 \right)^{-1} \quad (2.22)$$

$$= \left(1 + \frac{Z_0}{2Z_{\text{in}}} \right)^{-1}. \quad (2.23)$$




Figure 2.2: S_{21} parameter for a capacitively coupled RLC or $\lambda/4$ resonator with $R = 63.66 \text{ M}\Omega$, $L = 2.0264 \text{ nH}$, $C = 500 \text{ fF}$, $C_c = 3.5 \text{ fF}$, and $Z_0 = 50 \Omega$.

We plot the amplitude in dB and phase of S_{21} as a function of ω on Fig. 2.2, for $R = 63.66 \text{ M}\Omega$, $L = 2.0264 \text{ nH}$, $C = 500 \text{ fF}$, $C_c = 3.5 \text{ fF}$, and $Z_0 = 50 \Omega$. The resonator thus has $\omega_i/2\pi = 5 \text{ GHz}$ and $Q_i = 1000000$. However, we can see that resonance occurs at a lower frequency $\omega_0/2\pi = 4.98256 \text{ GHz}$. The effect of the coupling capacitor is thus to decrease the resonance frequency. For a transmission line resonator this effect can be intuitively understood by picturing the coupling capacitor elongating the wavelength past the open end of the line.

If we approximate Eq. (2.23) given Z_r from Eq. (2.9) near ω_0 , i.e., for small $\Delta\omega = \omega - \omega_0$, we obtain the following expression [111]:

$$S_{21}(\omega) \approx \left(1 + \frac{Q_i}{Q_c} \frac{1}{1 + 2iQ_i \frac{\Delta\omega}{\omega_0}} \right)^{-1}. \quad (2.24)$$

This expression can be adjusted to account for imperfect characteristic impedance on the source or load side [105]. Thus, given a measurement of S_{21} , we are able to fit the data and extract Q_i and Q_c , resulting in a useful characterization of the resonator [114].

2.1.2 Resonator Hamiltonian

Until now, the quantities we have seen could be calculated in a completely classical manner. In order to calculate the quantum Hamiltonian, we start from the classical circuit equations, although we assume the loss of the resonator to be zero, i.e., $1/R = 0$.

For any closed circuit the total power absorbed is 0. This fact is known as Tellegen's theorem. Because the power is 0, the energy is constant. Integrating the power therefore results in the classical Hamiltonian. Given the instantaneous voltages and currents v_k and i_k on each branch of a circuit, we may write

$$\sum_k v_k i_k = 0 \quad \rightarrow \quad H = \int_{-\infty}^t dt' \sum_k v_k i_k. \quad (2.25)$$

Equation (2.25) is general and applies to any kind of circuit, we will make use of it in Sec. 2.2 and 2.3. In the case of a simple LC resonator, there are only two branches, one each for the inductor and capacitor, thus we may substitute the constituent equations

$$v_L = L \frac{di_L}{dt} \quad (2.26)$$

$$i_C = C \frac{dv_C}{dt} \quad (2.27)$$

and obtain the classical Hamiltonian

$$H = \int_{-\infty}^t dt' (v_L i_L + v_C i_C) \quad (2.28)$$

$$= \int_{-\infty}^t dt' \left(L \frac{di_L}{dt} i_L + C v_C \frac{dv_C}{dt} \right) \quad (2.29)$$

$$= \int_{-\infty}^t dt' \frac{d}{dt'} \frac{1}{2} (L i_L^2 + C v_C^2) \quad (2.30)$$

$$= \frac{1}{2} (L i_L^2 + C v_C^2). \quad (2.31)$$

$$(2.32)$$

Note here that for this simple circuit, $v_C = v_L$ and $i_C = i_L$. While the current and voltage are perfectly valid quantities to analyze, it is common to change variables to flux and charge. In general, the flux ϕ is defined as the time integral of the voltage across an element, while the charge q is the integral of the current flowing through an

element [117]

$$\phi(t) = \int_{-\infty}^t v dt' \quad (2.33)$$

$$q(t) = \int_{-\infty}^t i dt'. \quad (2.34)$$

For the LC circuit then, we obtain

$$i_L = \frac{\phi_L}{L} \quad (2.35)$$

$$v_C = \frac{q_C}{C} \quad (2.36)$$

$$H = \frac{\phi_L^2}{2L} + \frac{q_C^2}{2C}. \quad (2.37)$$

It is easy to see that ϕ and q are canonically conjugate variables, meaning they have a Poisson bracket of one. To “quantize” the Hamiltonian, we use the method of classical analogy, which is generally attributed to Dirac [118]. According to this method, we may use the Poisson bracket relation between canonically conjugate variables to define the *commutation relation* of analogous quantum operators. Therefore, we *promote* ϕ and q to $\hat{\phi}$ and \hat{q} as such:

$$\{\phi, q\} = 1 \quad \rightarrow \quad [\hat{\phi}, \hat{q}] = i\hbar. \quad (2.38)$$

The quantum Hamiltonian is then

$$\hat{H} = \frac{\hat{\phi}_L^2}{2L} + \frac{\hat{q}_C^2}{2C}. \quad (2.39)$$

Because $\hat{\phi}$ and \hat{q} obey the canonical commutation relation, like the quantum harmonic oscillator, \hat{H} may be diagonalized with ladder operators. By analogy, we define

$$\hat{\phi} = \sqrt{\frac{L\hbar\omega_i}{2}}(\hat{a}^\dagger + \hat{a}) \quad (2.40)$$

$$\hat{q} = i\sqrt{\frac{C\hbar\omega_i}{2}}(\hat{a}^\dagger - \hat{a}) \quad (2.41)$$

where $\omega_i = 1/\sqrt{LC}$, as before. After substituting the ladder operators, Eq. (2.39) finally results in the well-known Hamiltonian

$$\hat{H} = \hbar\omega_i(\hat{a}^\dagger\hat{a} + 1/2). \quad (2.42)$$

2.2 Superconducting Qubits

Qubits are the “artificial atoms” of circuit QED. Unlike the harmonic oscillator of the preceding section, qubits are nonlinear devices in the sense that their energy levels are not equally-spaced, making them *anharmonic* oscillators. This fact makes them an essential component in circuit QED. A resonator has all of its energy levels equally spaced by an energy $\hbar\omega_0$. When driven by a microwave source, a resonator is thus excited into a *coherent state*, a quantum superposition over many states that is experimentally indistinguishable from the state of a classical oscillator. While superconducting qubits are generally not true two-level systems, their uneven level spacing makes it possible to address particular transitions individually. This paves the way to experiments where we create and manipulate arbitrary quantum states, thereby revealing the quantum nature of our world.

There exist many types of superconducting Josephson qubits, some were mentioned in Section 1.3. We focus here on a particular type initially developed by the Schoelkopf group at Yale: the transmon [104, 119]. The transmon was simplified and made more modular by the Martinis group in Santa Barbara, where they called it the Xmon due to its shape [92]. This is the design we use for the experiments of this thesis, though we will use the two names interchangeably.

The transmon resembles the original Cooper-pair box charge qubit [82]. It is a superconducting island connected to another island (or directly to ground plane for the Xmon) via a pair of parallel Josephson junctions. The major difference when compared to the original Cooper-pair box is the large capacitance C_Q provided by the island(s), in parallel to the junctions. The larger capacitance exponentially reduces the sensitivity to charge noise, at a trade-off of smaller anharmonicity [104]. This modification, along with fabrication improvements, significantly improved coherence times and propelled the transmon to wide usage across the world, both in academia and industry.

Because the two Josephson junctions are arranged in a loop, they form a DC SQUID. However, instead of using the interference effect to measure a magnetic field, qubits use it as a way to *tune* the transition frequency via a controlled external flux threading the loop [83]. Figure 2.3 shows the circuit diagram of a qubit with Josephson energy E_J and capacitance C_Q . Note that this circuit is very similar to that of an LC resonator, in fact, the Josephson element (either a single junction or two in a loop) is often called a “nonlinear inductor”.

2.2.1 The Josephson Equations

The Josephson equations were derived by Josephson in 1962 [14]. There are two equations: the first describes the supercurrent flow as a function of the order parameter phase difference φ (the current-phase relation) while the second relates the voltage




Figure 2.3: Circuit diagram of a tunable transmon qubit with Josephson energy E_J and capacitance C_Q . The two parallel Josephson junctions form a superconducting loop, allowing the critical current I_0 —or equivalently E_J —to be tuned by an external magnetic field.

across the junction to the time-derivative of φ (the phase evolution equation)

$$i_J = I_0 \sin(\varphi) \quad (2.43)$$

$$v_J = \frac{\hbar}{2e} \frac{d\varphi}{dt} \quad (2.44)$$

where I_0 is the junction critical current and $e = 1.602176634 \times 10^{-19}$ C is the elementary charge. If the current flowing through the junction exceeds I_0 , a voltage will be produced resulting in phase evolution.

Recall from Eq. (2.33) that the flux is the integral of the voltage, therefore

$$\phi_J = \frac{\hbar}{2e} \varphi = \frac{\Phi_0}{2\pi} \varphi \quad (2.45)$$

where $\Phi_0 = h/2e$ is the magnetic flux quantum.

If we substitute φ for ϕ in the first Josephson equation, we obtain an equation similar to Eq. (2.35). In fact, if we define the Josephson inductance to be $L_J = \hbar/(2eI_0)$, we see that when the phase difference is small we recover Eq. (2.35) exactly:

$$i_J = I_0 \sin\left(\frac{2e}{\hbar} \phi\right) \quad (2.46)$$

$$= \frac{\hbar}{2eL_J} \sin\left(\frac{2e}{\hbar} \phi\right) \quad (2.47)$$

$$\approx \frac{1}{L_J} \phi \quad \text{for small } \phi. \quad (2.48)$$

This is why a Josephson junction is a nonlinear inductor. Note however, that L_J is not a magnetic inductance; it is a *kinetic* inductance, generated by the kinetic energy of the Cooper pairs carrying the supercurrent. It is this nonlinear inductance that makes a

qubit an *anharmonic* oscillator rather than a harmonic oscillator, and which therefore enables quantum operations.

Before moving to the derivation of the Hamiltonian, we introduce one last equation which is very useful when fabricating junctions: the Ambegaokar-Baratoff relation [120], which relates the junction critical current to its normal state resistance R_n :

$$I_0 = \frac{\pi\Delta}{2eR_n} \quad (2.49)$$

where Δ is the superconducting gap, assuming that both superconductors are made of the same material. For aluminum, $\Delta \sim 200 \mu\text{eV}$ [121]. A typical single-junction resistance for the devices used in this thesis is $R_n \sim 10 \text{ k}\Omega$, resulting in $I_0 \sim 30 \text{ nA}$.

2.2.2 Qubit Hamiltonian

We proceed in a way similar to the last section: we start by writing the classical energy and quantize the conjugate variables. Given the Josephson equations and Tellegen's theorem [Eq. (2.25)], the classical energy is

$$H = \int_{-\infty}^t dt' (i_j v_j + i_c v_c) \quad (2.50)$$

$$= \int_{-\infty}^t dt' \left(\frac{\hbar I_0}{2e} \sin(\varphi) \frac{d\varphi}{dt} + C_Q v_c \frac{dv_c}{dt} \right) \quad (2.51)$$

$$= \int_{-\infty}^t dt' \frac{d}{dt'} \left(-\frac{\hbar I_0}{2e} \cos(\varphi) + \frac{1}{2} C_Q v_c^2 \right) \quad (2.52)$$

$$= -E_J \cos(\varphi) + \frac{1}{2} C_Q v_c^2 \quad (2.53)$$

where $E_J = \hbar I_0 / 2e$ is the Josephson energy. Note that we have assumed that there was a single Josephson element, with critical current I_0 and energy E_J . Nonetheless, taking into account two junctions in parallel actually leads to the same Hamiltonian. This is because quantization of the magnetic flux flowing through the SQUID loop constrains the phase differences of the two junctions, resulting in a single degree of freedom (rather than two). The Hamiltonian of the SQUID is

$$H_{\text{SQUID}} = -E_{J_1} \cos(\varphi_1) - E_{J_2} \cos(\varphi_2) \quad (2.54)$$

Given the flux quantization condition², we may rewrite the SQUID Hamiltonian to the form of a single junction with a rescaling of E_J [104, 122, 123]:

$$E_J \rightarrow E_{J_\Sigma} \left| \cos\left(\frac{\pi\Phi}{\Phi_0}\right) \right| \sqrt{1 + d^2 \tan^2\left(\frac{\pi\Phi}{\Phi_0}\right)^2} \quad (2.55)$$

where $E_{J_\Sigma} = E_{J_1} + E_{J_2}$, $d = \frac{E_{J_2} - E_{J_1}}{E_{J_1} + E_{J_2}}$ is the asymmetry, and Φ is an externally applied magnetic flux. The external magnetic flux, in particular, is how the effective Josephson energy may be controlled, therefore allowing for tunability of the qubit frequency. Tuning is essential for many experiments, in particular those of Chapters 5 and 6.

Instead of changing variables to (ϕ, q) , we leave φ as is and change the voltage variable to *charge number* $n = q_C/2e$, which is simply equivalent to the number of Cooper pairs on the qubit island, $v_C = 2en/C_Q$. Those two variables (φ, n) therefore have the same commutation relation as (ϕ, q) (up to a factor of \hbar) and we may quantize according to Eq. (2.38) to obtain the quantum Hamiltonian

$$\hat{H} = -E_J \cos(\hat{\varphi}) + 4E_C \hat{n}^2 \quad (2.56)$$

where $E_C = e^2/2C_Q$ is the single-electron charging energy.

Equation (2.56) is known as the Cooper-pair box Hamiltonian. As mentioned at the beginning of this section, however, the added parallel capacitance changes the regime of the system. Instead of \hat{n} being the well-defined variable, as in the Cooper-pair box, the large E_J/E_C ratio of the transmon causes $\hat{\varphi}$ to become the “localized” coordinate. Counterintuitively, increasing the capacitance C_Q *decreases* the charging energy.

It is interesting to compare Eq. (2.56) and Eq. (2.39). The qubit Hamiltonian looks different at first glance, since $\hat{\varphi}$ is enclosed in a cosine, rather than being squared. We can notice, however, that for a large value of E_J , and thus a large potential well, the phase coordinate will be confined to small values. It is therefore possible to expand the cosine in the small angle approximation, yielding, to fourth order

$$-\cos(\hat{\varphi}) \approx -1 + \frac{\hat{\varphi}^2}{2} - \frac{\hat{\varphi}^4}{24}. \quad (2.57)$$

Now the two Hamiltonians look very similar: they both contain the same harmonic potential $\hat{\varphi}^2$. The qubit Hamiltonian of course contains more terms which provide the required anharmonicity. This small angle approximation is called the transmon

²The flux through a superconducting loop is quantized because the superconducting order parameter is single-valued in the phase; i.e., the phase must vary by an integer multiple of 2π when going around the loop such that the order parameter is equal when returning to the same point [22].

Table 2.1: Classification of the integer order solutions to Mathieu’s equation. The solutions are categorized according to parity and periodicity, resulting in four classes of characteristic values and corresponding Mathieu functions. The index of the characteristic value is called the *order* and is defined for $m \in 0, 1, 2, 3, \dots$. The function names “ce” and “se” stand for *cosine-* and *sine-elliptic*. We are using the same notation as Ref. [126].

Characteristic Value	Mathieu Function	Parity	Periodicity
$a_{2m}(q)$	$ce_{2m}(x, q)$	even	π -periodic
$a_{2m+1}(q)$	$ce_{2m+1}(x, q)$	even	π -antiperiodic
$b_{2m+1}(q)$	$se_{2m+1}(x, q)$	odd	π -antiperiodic
$b_{2m+2}(q)$	$se_{2m+2}(x, q)$	odd	π -periodic

approximation [124], and it may be used to diagonalize the Hamiltonian. We choose not to use it here, and instead provide an exact method.

It is clear that we cannot use the same harmonic oscillator ladder operators to diagonalize the transmon Hamiltonian. Instead, we express \hat{n} in the conjugate momentum derivative form

$$\hat{n} = -i \frac{d}{d\hat{\phi}} \quad (2.58)$$

resulting in the time-independent Schrödinger equation in the *phase basis*:

$$-4E_C \frac{d^2 \psi_n}{d\hat{\phi}^2} - E_J \cos(\hat{\phi}) \psi_n = E_n \psi_n. \quad (2.59)$$

This equation has the same form as Mathieu’s differential equation [125, 126] with

$$\varphi = 2x \quad (2.60)$$

$$y(x) = \psi_n(2x) \quad (2.61)$$

$$q = -E_J/2E_C \quad (2.62)$$

$$a = E_n/E_C. \quad (2.63)$$

For real q , there is an infinite number of solutions with a particular *characteristic value* and a corresponding Mathieu function that may be categorized according to parity and periodicity. These solutions, the *integer order* Mathieu functions, are tabulated in Table 2.1.

It is clear from the form of the Schrödinger equation that the wave function must be periodic in φ :

$$\psi_n(\varphi) = \psi_n(\varphi + 2\pi). \quad (2.64)$$




Figure 2.4: First four eigenfunctions of the transmon Hamiltonian for $E_J/E_C = 100$.

Thus, we need $y(x)$ to be π -periodic for $\psi(\varphi)$ to be 2π -periodic. We may now define the eigenfunctions and eigenvalues of Eq. (2.59) for levels $n \in 0, 1, 2, 3, \dots$: If n is even, $\psi_n = ce_n$ with $E_n = E_C a_n$, and if n is odd, $\psi_n = se_{n+1}$ with $E_n = E_C b_{n+1}$. We thus select Mathieu functions of even order, alternating between cosine-elliptic and sine-elliptic. It may also be convenient to offset the eigenenergies such that $E_0 = 0$.

We usually label the transmon basis state vectors according to their corresponding energy level number as $|0\rangle, |1\rangle, |2\rangle, \dots$. When using the transmon as a qubit, it is possible to only consider the two or three lowest levels, which are then denoted $|g\rangle, |e\rangle, |f\rangle$, where the label “g” and “e” stand for ground and excited. In the qubit approximation, the Hamiltonian is simply

$$\hat{H} = -\frac{1}{2}(E_1 - E_0)\hat{\sigma}_z = -\frac{1}{2}\hbar\omega_{01}\hat{\sigma}_z \quad (2.65)$$

where $\hbar\omega_{01} = E_1 - E_0$, and we negate the Pauli Z matrix to obtain the correct level ordering.

The first four eigenfunctions of Eq. (2.59) are plotted in Fig. 2.4 with the help of the `Mathieu.jl` computer package [127]. Unsurprisingly, they look quite similar to the eigenfunctions of the quantum harmonic oscillator.

The anharmonicity of the transmon is the difference between the transition fre-




Figure 2.5: Chip design with two Xmon qubits connected by a $\lambda/2$ resonator in the middle. The $\lambda/4$ readout resonators are connected to the top branch of each qubit and are capacitively coupled to the readout transmission line at the top. The transmission lines on the left and right sides are for flux biasing; they are inductively coupled to each qubit's SQUID. The lines at the bottom of the cross are microwave lines used for qubit drive pulses; they are capacitively coupled to the cross-shaped island.

quencies ω_{01} and ω_{12}

$$\alpha = \hbar\omega_{12} - \hbar\omega_{01} = E_2 - 2E_1 + E_0. \quad (2.66)$$

Because the cosine potential of the transmon is “softer” than a true harmonic potential, $\omega_{12} < \omega_{01}$, and α is therefore negative. In the transmon approximation, $\hbar\omega_{01} \approx \sqrt{8E_J E_C} - E_C$, and $\alpha \approx -E_C$. More accurate closed-form analytical expressions can be derived by expanding the cosine to higher order [128].

2.3 Quantum Circuits

In the last two sections, we introduced the two main elements of circuit QED individually. We now explain how to analyze a circuit comprising multiple connected elements in a general and systematic way, assuming that the couplings between elements are made with capacitors. This method is, however, easily extensible to inductive couplings; this extension is left to the reader.

We choose as an example a realistic circuit comprising two transmon qubits and one resonator connecting them; it is pictured in Fig. 2.5 (we ignore the readout resonators

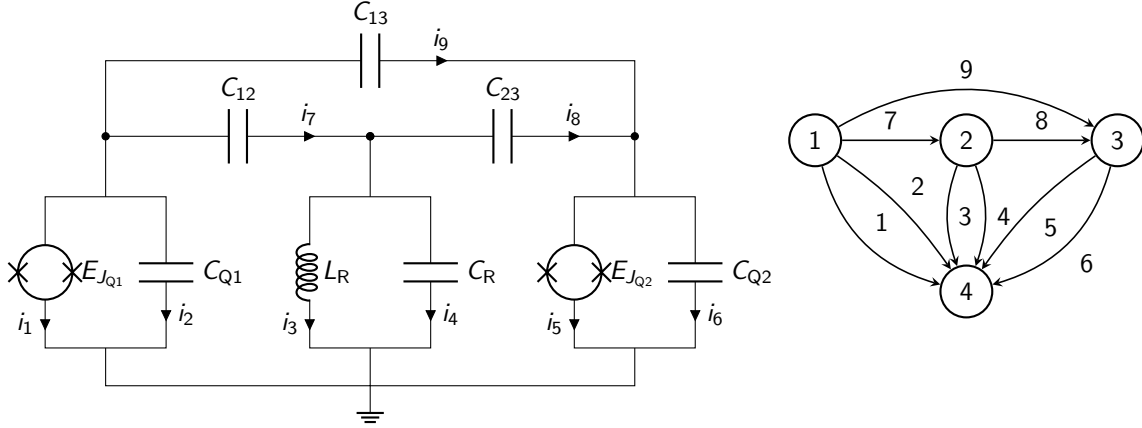


Figure 2.6: Circuit diagram (left) and corresponding directed graph (right) for two qubits connected by a resonator. Though no direct capacitive coupling was designed between the two qubits, we still include a *parasitic capacitance* C_{13} in the model. The four circuit nodes are situated at the three dots and the ground.

above each qubit). The circuit diagram for such a design is shown in Fig. 2.6, along with the directed graph corresponding to the circuit. We use this graph as a tool to identify the degrees of freedoms in the circuit and subsequently write the circuit equations.

To create the directed graph, we must first locate the branches and nodes in the circuit. Branches correspond to individual currents flowing through one or multiple elements in series. Nodes are located wherever currents from different branches split or join. We also choose a reference or *datum* node. The datum may be any node, but it is generally chosen to be the circuit ground.

The next step is to write the *incidence matrix*, which is simply a different representation of the graph. For a circuit with n nodes and b branches, the incidence matrix \mathbf{A} is an $n \times b$ matrix where each entry \mathbf{A}_{ij} is associated to a node/branch combination. If branch j is directed out of (into) node i , $\mathbf{A}_{ij} = 1(-1)$. If branch j is not connected to node i , $\mathbf{A}_{ij} = 0$. For large circuits the matrix will thus likely be sparse.

If the circuit is connected, meaning that there is a path from any node to any other (this will be true for any relevant circuit), we can remove the row corresponding to the datum node from the incidence matrix. This is because \mathbf{A} is a representation of Kirchhoff's circuit laws which, for n nodes, determine $n - 1$ linearly independent equations [129]. The reduced incidence matrix of the circuit, with the datum row

crossed out, is then

$$\mathbf{A} = \begin{pmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 1 & 0 & 0 & -1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & -1 & -1 \\ \hline -1 & -1 & -1 & -1 & -1 & -1 & 0 & 0 & 0 \end{pmatrix}. \quad (2.67)$$

If we denote the branch currents with the vector $\vec{i} = (i_1, i_2, \dots, i_b)^T$, the branch voltages with the vector $\vec{v} = (v_1, v_2, \dots, v_b)^T$, and the node-to-datum voltages with the vector $\vec{e} = (e_1, e_2, \dots, e_{n-1})^T$, Kirchhoff's current and voltage laws can be succinctly written as

$$\mathbf{A}\vec{i} = 0 \quad \text{and} \quad \mathbf{A}^T\vec{e} = \vec{v}. \quad (2.68)$$

With those two equations, we can choose $n - 1$ currents and voltages as our independent degrees of freedom. For this circuit, we have, e.g., the voltages

$$v_7 = e_1 - e_2 = v_1 - v_3 \quad (2.69)$$

$$v_8 = e_2 - e_3 = v_3 - v_5 \quad (2.70)$$

$$v_9 = e_1 - e_3 = v_1 - v_5 \quad (2.71)$$

therefore leaving 3 independent voltages, v_1 , v_3 and v_5 . We use only those independent voltages and currents to write the classical energy with Tellegen's theorem [Eq. (2.25)] and the constitutive equations covered in the last sections for the inductor, the capacitor, and the Josephson junction:

$$\begin{aligned} H = \int_{-\infty}^t & \left[E_{J_{Q1}} \dot{\phi}_{Q1} \sin \varphi_{Q1} + v_1 C_{Q1} \dot{v}_1 + L_R i_3 \dot{i}_3 + v_3 C_R \dot{v}_3 + E_{J_{Q2}} \dot{\phi}_{Q2} \sin \varphi_{Q2} + v_5 C_{Q2} \dot{v}_5 \right. \\ & \left. + (v_1 - v_3) C_{12} (\dot{v}_1 - \dot{v}_3) + (v_3 - v_5) C_{23} (\dot{v}_3 - \dot{v}_5) + (v_1 - v_5) C_{13} (\dot{v}_1 - \dot{v}_5) \right] dt' \end{aligned} \quad (2.72)$$

$$\begin{aligned} = \int_{-\infty}^t & \frac{d}{dt'} \left[-E_{J_{Q1}} \cos \varphi_{Q1} + \frac{1}{2} L_R i_3^2 - E_{J_{Q2}} \cos \varphi_{Q2} + \frac{1}{2} (C_{Q1} v_1^2 + C_R v_3^2 + C_{Q2} v_5^2 \right. \\ & \left. + C_{12} (v_1 - v_3)^2 + C_{23} (v_3 - v_5)^2 + C_{13} (v_1 - v_5)^2 \right] dt' \end{aligned} \quad (2.73)$$

$$= -E_{J_{Q1}} \cos \varphi_{Q1} + \frac{1}{2} L_R i_3^2 - E_{J_{Q2}} \cos \varphi_{Q2} + \frac{1}{2} \vec{V}^T \mathbf{C} \vec{V} \quad (2.74)$$

where \vec{V} is the voltage vector and \mathbf{C} is the capacitance matrix:

$$\vec{V} = \begin{pmatrix} v_1 \\ v_3 \\ v_5 \end{pmatrix} \quad \mathbf{C} = \begin{pmatrix} C_{Q1} + C_{12} + C_{13} & -C_{12} & -C_{13} \\ -C_{12} & C_R + C_{12} + C_{23} & -C_{23} \\ -C_{13} & -C_{23} & C_{Q2} + C_{13} + C_{23} \end{pmatrix}. \quad (2.75)$$

If the circuit had more inductive elements, we would have also defined the appropriate current vector \vec{I} and inductance matrix \mathbf{L} , which would have led to the inductive part of the energy $\frac{1}{2}\vec{I}^T\mathbf{L}\vec{I}$. We can now substitute \vec{V} and \vec{I} for $\vec{\phi}$ and \vec{q} by inverting \mathbf{C} and \mathbf{L} :

$$\vec{V} = \mathbf{C}^{-1}\vec{q} \quad \text{and} \quad \vec{I} = \mathbf{L}^{-1}\vec{\phi} \quad (2.76)$$

where \mathbf{C}^{-1} and \mathbf{L}^{-1} are symmetric. In those coordinates, the Hamiltonian is

$$H = -E_{J_{Q1}} \cos \varphi_{Q1} + \frac{\phi_R^2}{2} L_R^{-1} - E_{J_{Q2}} \cos \varphi_{Q2} + \frac{1}{2} \vec{q}^T \mathbf{C}^{-1} \vec{q}. \quad (2.77)$$

Finally, we quantize according to Eq. (2.38), with a slight addition: Because there are multiple sets of conjugate variables, the quantized coordinates will correspond to separate Hilbert spaces. The qubit and resonator parts therefore live in their own Hilbert space. The charge cross-terms, however, are tensor products of two distinct charge coordinates. They represent the coupling energies between the circuit elements. For example, the coupling energy between the first qubit and the resonator is

$$\hat{H}_{Q1,R} = [\mathbf{C}^{-1}]_{(1,2)} \hat{q}_{Q1} \otimes \hat{q}_R. \quad (2.78)$$

On the resonator side, we already know the form of \hat{q}_R in the diagonal basis, which is given by Eq. (2.41). To obtain an operator representation for \hat{q}_{Q1} , we switch to the conjugate momentum form and multiply on the left and right by the identity operator in the diagonal transmon basis $\mathbf{I} = \sum_l |\lambda\rangle\langle\lambda|$. We then insert the identity in the $\hat{\phi}$ basis $\mathbf{I}_{\hat{\phi}} = \int_{-\pi}^{\pi} d\hat{\phi} |\hat{\phi}\rangle\langle\hat{\phi}|$.

$$\hat{q} = \frac{2e}{i} \frac{d}{d\hat{\phi}} \quad \text{according to Eq. (2.58)} \quad (2.79)$$

$$= \frac{2e}{i} \sum_{l,m} |\lambda\rangle\langle\lambda| \frac{d}{d\hat{\phi}} |m\rangle\langle m| \quad \text{multiplicating by } \mathbf{I} \quad (2.80)$$

$$= \frac{2e}{i} \sum_{l,m} \int_{-\pi}^{\pi} d\hat{\phi} |\lambda\rangle\langle\lambda| |\hat{\phi}\rangle\langle\hat{\phi}| \frac{d}{d\hat{\phi}} |m\rangle\langle m| \quad \text{inserting } \mathbf{I}_{\hat{\phi}} \quad (2.81)$$

$$= \frac{2e}{i} \sum_{l,m} |\lambda\rangle\langle\lambda| \int_{-\pi}^{\pi} d\hat{\phi} \psi_l^* \frac{d\psi_m}{d\hat{\phi}} \quad \text{because } \langle\hat{\phi}|m\rangle = \psi_m(\hat{\phi}). \quad (2.82)$$

It is therefore easy to express \hat{q} as a matrix in the transmon basis, with each entry given by an integral over the product of Mathieu functions and derivatives.

One notable point concerns the inverse capacitance matrix in Eq. (2.77): In general, \mathbf{C}^{-1} will be completely dense. Thus, even if we had not included the parasitic capacitance C_{13} , there would still be a term proportional to $\hat{q}_1 \otimes \hat{q}_3$ in the Hamiltonian. This

second order coupling [130] would be fairly small, however, and could be neglected depending on the level of physical accuracy desired. A second consequence is that, unlike in Eqs. (2.39) and (2.56), the capacitive part of the resonator and qubit energy does not depend only on the “local” capacitor (i.e., C_R and C_Q). Instead, there is a dependence on the entire capacitive network of the circuit, as can be seen, e.g., in the charging energy of the first qubit

$$E_{C_{Q1}} = \frac{e^2}{2} [\mathbf{C}^{-1}]_{(1,1)}. \quad (2.83)$$

This effect must be properly accounted for when calculating the level energies.

In this section, we have demonstrated a formalism for determining the Hamiltonian of arbitrary quantum circuits, based on the topology of the network expressed as a graph. This formalism is convenient, especially when compared to the method used, e.g., in Ref. [131], which depends on guessing a Lagrangian from the circuit equations.

We finish this section by mentioning that if we perform a two-level approximation and a rotating-wave approximation (see Sec. 2.5) to a qubit–resonator Hamiltonian as obtained by the method above, we obtain the well-known Jaynes-Cummings Hamiltonian

$$\hat{H}_{\text{JC}} = \hbar\omega_R \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) - \frac{1}{2} \hbar\omega_Q \hat{\sigma}_z + g(\hat{a}^\dagger \hat{\sigma}_- + \hat{a} \hat{\sigma}_+) \quad (2.84)$$

where $\omega_Q = \omega_{01}$, $\hat{\sigma}_{(+,-)}$ are respectively the qubit raising and lowering operators, and g is the coupling coefficient.

An interesting transformation to Eq. (2.84) provides a hint as to how qubit readout might be achieved in circuit QED, the Schrieffer-Wolf transformation [132, 133]. For the Jaynes-Cummings Hamiltonian, this unitary transformation is [102]

$$\hat{U} = e^{\frac{g}{\Delta}(\hat{a}\hat{\sigma}_+ - \hat{a}^\dagger\hat{\sigma}_-)} \quad (2.85)$$

where $\Delta = \omega_Q - \omega_R$. The Schrieffer-Wolf transformation *decouples* the Hamiltonian, eliminating the last term of Eq. (2.84), and therefore diagonalizing it. The resulting Hamiltonian is an approximation however, and is only valid in the *dispersive* regime, for $\Delta \gg g$,

$$\hat{U} \hat{H} \hat{U}^\dagger \approx \hbar(\omega_R - \chi \hat{\sigma}_z) \hat{a}^\dagger \hat{a} - \frac{1}{2} \hbar(\omega_Q + \chi) \hat{\sigma}_z \quad (2.86)$$

where $\chi = g^2/\Delta$. The first term of the above equation shows that the resonator’s frequency is *state-dependent*. It changes by an amount 2χ depending on whether the qubit is in $|g\rangle$ or $|e\rangle$. This frequency shift is often called the dispersive shift, or the AC Stark shift, from its quantum optics origins.

For qubit readout, we simply measure the resonator. This can be done in a variety of ways. A simple one would be to measure the transmission amplitude at $\omega_R - \chi$. If

the qubit is in $|g\rangle$, the resonator will be there, and we will see the corresponding S_{21} transmission. If the qubit is in $|e\rangle$, S_{21} will show the resonator is not there.

One important caveat to the above derivation is that the transformation was applied in the two-level approximation of the qubit. For a transmon, the χ shift is only accurate when the qubit is in $|g\rangle$. In $|e\rangle$, level repulsion caused by the $|f\rangle$ state renders the approximation invalid. Instead, the excited-qubit-state frequency of the resonator is

$$\omega_{R|e} \approx \omega_R - \chi \frac{(\Delta + E_C)}{(\Delta - E_C)} \quad (2.87)$$

as can be observed in Fig. 3.9 (b).

The Jaynes-Cummings Hamiltonian is a staple of quantum optics, and by extension of circuit QED. It is covered extensively in the literature and, thus, we do not explain its many properties and applications. Instead, see, e.g., Refs. [134, 135, 136, 137].

2.4 Quantum Gates

A qubit is, by definition, a quantum two-level system, and, just like its classical counterpart a variety of operations or “gates” are used to control its state. Note that other approaches to quantum computation exist, e.g., measurement based computing [138, 139] or even quantum annealing [140, 141, 142], but gate-based quantum computing, through its generality and simple correspondence with the way algorithms themselves are designed, remains the principal contender.

It is convenient to adopt the picture given by the Bloch sphere when describing qubit operations. The Bloch sphere represents the state space of a two-level system in 3D, and was originally formulated for nuclear and atomic spins [143, 144]. In that representation, an arbitrary qubit pure state

$$|\psi\rangle = \alpha |g\rangle + \beta |e\rangle \quad (2.88)$$

$$= \cos(\theta/2) |g\rangle + e^{i\varphi} \sin(\theta/2) |e\rangle \quad (2.89)$$

is represented as a unit vector in spherical coordinates $(1, \theta, \varphi)$, or, in \mathbb{R}^3

$$\vec{a} = (\sin \theta \cos \varphi, \sin \theta \sin \varphi, \cos \theta) \quad (2.90)$$

with the angles given by $\theta = 2 \cos^{-1}(\alpha)$ and $\varphi = \tan^{-1}(\text{Im}[\beta]/\text{Re}[\beta])$, assuming that $\alpha \in \mathbb{R}$ [26]. We can always make α real because we have the freedom to set the *global phase* of a state (i.e. we have the freedom to multiply $|\psi\rangle$ by an arbitrary factor $e^{i\phi}$). The global phase of a state can never be observed in a physical measurement. In contrast, φ is a relative phase between $|g\rangle$ and $|e\rangle$, and it can be observed. It is often just called “the” phase of a state, since the global phase is ignored.




Figure 2.7: The Bloch Sphere. The Bloch sphere represents the state space of a qubit. An arbitrary qubit state may be expressed as a Bloch vector \vec{a} via its length $\|\vec{a}\|$, azimuthal angle θ , and polar angle φ in spherical coordinates. The Euclidean axes of the Bloch sphere correspond to basic qubit states, with the x , y , and z -axes corresponding to the $|+\rangle$, $|+i\rangle$, and $|g\rangle$ states.

Figure 2.7 illustrates the Bloch sphere, including labels for the six states corresponding to the axes. The $|g\rangle$ and $|e\rangle$ states are located at the north and south poles, respectively. The four states on the equator correspond to equal superpositions. The $|\pm\rangle = (|g\rangle \pm |e\rangle) / \sqrt{2}$ states are aligned with the x -axis, and the $|\pm i\rangle = (|g\rangle \pm i |e\rangle) / \sqrt{2}$ states with the y -axis. The orange vector arrow shows an example state $|\psi\rangle$ with $\theta \sim \pi/2$ and $\varphi \sim \pi/25$.

In that representation, qubit gates can be described as *rotations* around the x , y , and z axes, which is why that language is often used in the literature (and in this thesis). If a qubit is initialized in $|g\rangle$, a π rotation about the x -axis will set it to $|e\rangle$. Such a gate is often denoted X or X_π . Similarly, a $\pi/2$ rotation about the y -axis is denoted $Y/2$ or $Y_{\pi/2}$. Note that gates that rotate the state by a negative angle differ from their positive complement, e.g., a $\pi/2$ rotation is not identical to a $-3\pi/2$ rotation. This is because rotations are not made instantaneously and other physical processes—notably decoherence—may occur during the trajectory. For example, imagine that $1/f$ noise temporarily changes the qubit frequency by an amount $\delta\omega$. This detuning causes φ to increase at a constant rate during the gate, or, equivalently, causes the axis of the rotation to change. The state after a $-3\pi/2$ rotation will therefore differ from the one where a $\pi/2$ rotation was made. The two trajectories in Fig. 2.7 illustrate this difference,

and more details on the dynamics are explained in Sec. 2.5.

The state of a qubit after decoherence is no longer described by a pure state vector, as in Eq. (2.88), but instead by a density matrix ρ . The Bloch sphere can again be used to represent such a state by allowing the length of the Bloch vector \vec{a} to be less than 1. As a result, the vector will point to the interior of the sphere. Given \vec{a} , an arbitrary single-qubit state can be described by a density matrix

$$\rho = \frac{1}{2} (\mathbf{I} + \vec{a} \cdot \vec{\sigma}) \quad (2.91)$$

where $\|\vec{a}\| \leq 1$, and $\vec{\sigma} = (\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z)$ is the Pauli vector.

The main drawback of the Bloch sphere representation is that it can only be used for single-qubit states. Multi-qubit states require exponentially more dimensions, and therefore cannot fit in this 3D picture. Nonetheless, because single-qubit gates are part of most quantum algorithms, the language of qubit “rotations” remains useful even in multi-qubit contexts.

We finish this section by briefly discussing *gate sets*. With superconducting qubits, it is physically possible to rotate the state about any axis in the $x - y$ plane with a microwave pulse. Such a pulse is made by multiplying an envelope, e.g., Gaussian, with a carrier wave (see Section 3.3.2). The relative *phase* of the carrier between pulses determines the relative angle between rotation axes. A mathematical explanation for this effect is provided in Section 2.5.2, where we analyze the dynamics of the driven qubit. Rotations about the z -axis are done very differently. They are equivalent to changing the frequency of the qubit by an amount $\delta\omega$ for a time Δt , resulting in a rotation by $\Delta\varphi = \delta\omega\Delta t$ about the z -axis. Such rotations can be generated by actually changing the qubit frequency with a flux pulse, or they may be done *virtually*, by tracking rotations of the reference frame and modifying the axis of subsequent X and Y rotations [145].

Despite this freedom, quantum algorithms are generally constructed with a finite gate set comprising a number of single- and two-qubit gates (and sometimes three-qubit gates). To implement algorithms then, it is necessary to convert the theoretical description to an experimental pulse sequence, ideally minimizing the number of pulses necessary [146], or even changing gates to minimize coherent errors [147]. The pulses forming the pulse sequence are part of the experimental gate set. The gates in this set should be optimized to produce high-quality operations, and include, in particular, two-qubit gates specific to the architecture. For tunable qubits, this gate is often a controlled-phase (CPHASE) implemented with a flux pulse [148, 149, 150, 151]. The calibration of such a gate requires knowing the coupling strength between two qubits; this can be achieved with the methods described in Chapter 5.

Finally, it is important to characterize the quality of the gates in the experimental gate set. The gold standard is *randomized benchmarking*, along with its many exten-

sions [152, 153, 154, 155, 156, 157, 158, 159, 160], which results in the average gate fidelity [161, 162]

$$\bar{F}_g = \int d\psi \langle \psi | U^\dagger \mathcal{E}_U(\psi) U | \psi \rangle \quad (2.92)$$

where U is the ideal gate to be implemented, $\mathcal{E}_U(\psi)$ is the result of its implementation on a pure state $|\psi\rangle$ and is, in general, a density operator, and the integration is performed uniformly over all pure states (the Haar measure) and is normalized such that $\int d\psi = \mathbf{I}$. The average gate fidelity is thus a metric that characterizes how well the physical implementation of a gate matches its ideal target U on average when the gate is applied to any possible input state. In Sec. 3.5 we perform single-qubit randomized benchmarking, purity benchmarking, and randomized benchmarking over time.

2.5 Dynamics and Decoherence

The dynamics of all coherent quantum systems are governed by the *Schrödinger equation*, a linear partial differential equation that relates the time derivative of the wave function to the Hamiltonian [163]

$$i\hbar \frac{\partial}{\partial t} |\psi\rangle = \hat{H} |\psi\rangle. \quad (2.93)$$

Despite its simple appearance, the Schrödinger equation does not allow many analytical closed-form solutions. Some well-known solvable problems include, of course, the harmonic oscillator, the particle in a box, and the hydrogen atom. More complicated many-body Hamiltonians, such as those used for the circuits in this thesis, are generally not solvable analytically. It is therefore important to learn and make use of both analytical approximations and numerical simulations when studying the dynamics of circuit QED systems.

Two important approximations used in circuit QED are the two-level approximation and the rotating-wave approximation. Both of these approximations are used to obtain the Jaynes-Cummings Hamiltonian [Eq. (2.84)] when starting from the exact circuit Hamiltonian of a qubit–resonator system.

The two-level approximation is self-explanatory: Instead of including the many energy levels of a particular system, we *truncate* the Hilbert space to keep only the ground state and the first excited state. The validity of this approximation depends on many factors. For a transmon qubit, the approximation is partially justified by the anharmonicity α . The anharmonicity makes it possible to drive the transition between $|0\rangle$ and $|1\rangle$ without exciting other energy levels (in contrast, many energy levels are populated when driving a resonator). The approximation becomes more accurate with a larger anharmonicity and a smaller pulse bandwidth. Short drive pulses have a larger

frequency bandwidth and therefore tend to cause leakage of the qubit state out of the two-level subspace. Pulse shapes that minimize this leakage have been designed, leading to higher gate fidelity [164, 165, 166, 167]. Other problems with the two-level approximation of transmon includes energy level crossings with the qubit in the $|1\rangle$ state, which are generally influenced by the second excited state $|2\rangle$. The two-level approximation can also be used with other types of systems if a single excitation is involved in the dynamics. In that case, the Hamiltonian can be reduced to an effective two-level subspace. The methods in Chapter 5 exploit this property to simplify the dynamics between a qubit and, e.g., a resonator. Although the two-level approximation is conceptually useful, its narrow range of correct applications means that it should not be used for accurate calculations, and numerical simulations should instead be preferred.

The rotating-wave approximation is often applied in the context of a *rotating frame*, and since that is a useful concept in its own right, we explain both the rotating-wave approximation and the rotating frame, as applied to the dynamics of a qubit driven by a microwave pulse³.

2.5.1 The Rotating Frame

The rotating frame is a special case of the standard interaction picture in quantum mechanics. What makes it special is that it is designed to eliminate the time-independent part of a Hamiltonian while simplifying a time-dependent part (in combination with the rotating-wave approximation). Given a Hamiltonian

$$\hat{H} = \hat{H}_0 + \hat{H}_1(t) \quad (2.94)$$

where H_0 is time-independent and $H_1(t)$ is time-dependent, we define the state in the rotating frame

$$|\psi'(t)\rangle = \hat{R} |\psi(t)\rangle \quad (2.95)$$

$$= e^{\frac{i}{\hbar}\hat{H}_0 t} |\psi(t)\rangle \quad (2.96)$$

³Thereby making this example a rare “three birds, one stone” event...

where \hat{R} is a rotation operator and is usually chosen to be $e^{\frac{i}{\hbar}\hat{H}_0 t}$. We insert $|\psi'(t)\rangle$ on the left side of Eq. (2.93) to determine the “rotated dynamics”:

$$i\hbar \frac{\partial}{\partial t} |\psi'\rangle = i\hbar \frac{\partial}{\partial t} (\hat{R} |\psi\rangle) \quad (2.97)$$

$$= i\hbar \frac{\partial \hat{R}}{\partial t} \hat{R}^\dagger \hat{R} |\psi\rangle + \hat{R} i\hbar \frac{\partial}{\partial t} |\psi\rangle \quad \text{product rule and inserting } \hat{R}^\dagger \hat{R} \quad (2.98)$$

$$= i\hbar \frac{\partial \hat{R}}{\partial t} \hat{R}^\dagger |\psi'\rangle + \hat{R} \hat{H} |\psi\rangle \quad \text{substituting Eqs. (2.95) and (2.93)} \quad (2.99)$$

$$= i\hbar \frac{\partial \hat{R}}{\partial t} \hat{R}^\dagger |\psi'\rangle + \hat{R} \hat{H} \hat{R}^\dagger |\psi'\rangle \quad \text{inserting Eq. (2.95)} \quad (2.100)$$

$$= \left(i\hbar \frac{\partial \hat{R}}{\partial t} \hat{R}^\dagger + \hat{R} \hat{H} \hat{R}^\dagger \right) |\psi'\rangle. \quad (2.101)$$

It is clear that the time evolution of this new rotating state is also described by the Schrödinger equation, but with a modified Hamiltonian

$$\hat{H}' = i\hbar \frac{\partial \hat{R}}{\partial t} \hat{R}^\dagger + \hat{R} \hat{H} \hat{R}^\dagger. \quad (2.102)$$

2.5.2 The Driven Qubit Hamiltonian

The Hamiltonian of a driven qubit in the two-level approximation is

$$\hat{H} = H_0 + H_1 = -\frac{1}{2}\hbar\omega_Q \hat{\sigma}_z + i\hbar\varepsilon(t) \cos(\omega_d t + \phi_d) (\hat{\sigma}_+ - \hat{\sigma}_-). \quad (2.103)$$

For reference, note that the drive Hamiltonian H_1 is derived from the application of a classical field—an oscillating voltage—capacitively coupled to a qubit. The form preferred for numerical simulations is thus given by $H_1 = \varepsilon(t) \cos(\omega_d t + \phi_d) \hat{q}$, where \hat{q} is given by Eq. (2.82) and $\varepsilon(t)$ is the scaled amplitude of the drive in units of angular frequency [101].

We could now use a rotating frame with $\hat{R} = e^{\frac{i}{\hbar}\hat{H}_0 t}$, which would entirely eliminate H_0 , but for slightly more generality, we instead enter a frame rotating at the drive

frequency $\hat{R} = e^{-\frac{i}{2}\omega_d t \hat{\sigma}_z}$. Following Eq. (2.102)

$$\hat{H}' = i\hbar \frac{\partial \hat{R}}{\partial t} \hat{R}^\dagger + \hat{R} \hat{H} \hat{R}^\dagger \quad (2.104)$$

$$= -\frac{i^2 \hbar}{2} \omega_d \hat{\sigma}_z R R^\dagger + \hat{R} \hat{H} \hat{R}^\dagger \quad \text{taking the derivative of } \hat{R} \quad (2.105)$$

$$= \frac{\hbar}{2} \omega_d \hat{\sigma}_z + \hat{H}_0 + \hat{R} \hat{H}_1 \hat{R}^\dagger \quad \text{since } \hat{H}_0 \text{ and } \hat{R} \text{ commute} \quad (2.106)$$

$$= -\frac{\hbar}{2} (\omega_Q - \omega_d) \hat{\sigma}_z + \hbar \varepsilon(t) \cos(\omega_d t + \phi_d) \hat{R} i (\hat{\sigma}_+ - \hat{\sigma}_-) \hat{R}^\dagger \quad (2.107)$$

$$= -\frac{\hbar}{2} (\omega_Q - \omega_d) \hat{\sigma}_z + \hbar \varepsilon(t) \cos(\omega_d t + \phi_d) (\cos(\omega_d t) \hat{\sigma}_y + \sin(\omega_d t) \hat{\sigma}_x) \quad (2.108)$$

where we simplified $\hat{R} i (\hat{\sigma}_+ - \hat{\sigma}_-) \hat{R}^\dagger = \hat{R} \hat{\sigma}_y \hat{R}^\dagger$ with a useful corollary of the Baker-Campbell-Hausdorff formula

$$\hat{R} \hat{\sigma}_y \hat{R}^\dagger = e^{-\frac{i}{2}\omega_d t \hat{\sigma}_z} \hat{\sigma}_y e^{\frac{i}{2}\omega_d t \hat{\sigma}_z} \quad (2.109)$$

$$= \hat{\sigma}_y + \left(\frac{i\omega_d t}{2}\right) \frac{[\hat{\sigma}_z, \hat{\sigma}_y]}{1!} + \left(\frac{i\omega_d t}{2}\right)^2 \frac{[\hat{\sigma}_z, [\hat{\sigma}_z, \hat{\sigma}_y]]}{2!} + \left(\frac{i\omega_d t}{2}\right)^3 \frac{[\hat{\sigma}_z, [\hat{\sigma}_z, [\hat{\sigma}_z, \hat{\sigma}_y]]]}{3!} + \dots \quad (2.110)$$

$$= \hat{\sigma}_y - \left(\frac{i\omega_d t}{2}\right) \frac{2i\hat{\sigma}_x}{1!} + \left(\frac{i\omega_d t}{2}\right)^2 \frac{4\hat{\sigma}_y}{2!} - \left(\frac{i\omega_d t}{2}\right)^3 \frac{8i\hat{\sigma}_x}{3!} + \left(\frac{i\omega_d t}{2}\right)^4 \frac{16\hat{\sigma}_y}{4!} + \dots \quad (2.111)$$

$$= \sum_{k=0}^{\infty} \frac{(-1)^k (\omega_d t)^{2k}}{(2k)!} \hat{\sigma}_y + \sum_{k=0}^{\infty} \frac{(-1)^k (\omega_d t)^{2k+1}}{(2k+1)!} \hat{\sigma}_x \quad (2.112)$$

$$= \cos(\omega_d t) \hat{\sigma}_y + \sin(\omega_d t) \hat{\sigma}_x. \quad (2.113)$$

We continue where we left off, focusing on the second term of Eq. (2.108) to obtain

$$\hbar \varepsilon(t) \cos(\omega_d t + \phi_d) (\cos(\omega_d t) \hat{\sigma}_y + \sin(\omega_d t) \hat{\sigma}_x) \quad (2.114)$$

$$= \frac{\hbar \varepsilon(t)}{2} ([\cos(2\omega_d t + \phi_d) + \cos(\phi_d)] \hat{\sigma}_y + [\sin(2\omega_d t + \phi_d) - \sin(\phi_d)] \hat{\sigma}_x) \quad (2.115)$$

$$\approx \frac{\hbar \varepsilon(t)}{2} (\cos(\phi_d) \hat{\sigma}_y - \sin(\phi_d) \hat{\sigma}_x) \quad (2.116)$$

where the rotating-wave approximation was used in the last line to eliminate the fast terms oscillating at a frequency $2\omega_d$. The idea behind the rotating-wave approximation is that the fast terms average out to zero over timescales longer than $2\pi/2\omega_d$. For $\omega_d/2\pi \sim 5$ GHz, this is 0.1 ns. Since most qubit operations are in the range of 10-100 ns, this approximation is generally valid. Large-amplitude pulses that change the qubit state very fast are an exception, and should be modelled without the approximation [168]. Also notice that, had we chosen a frame rotating at a different frequency, e.g., at ω_Q , the two leftover terms would be oscillating at a frequency $\omega_Q - \omega_d$.

The final approximated rotating frame Hamiltonian is

$$\hat{H}' = \frac{\hbar}{2}\Delta\omega\hat{\sigma}_z + \frac{\hbar\varepsilon(t)}{2}(\cos(\phi_d)\hat{\sigma}_y - \sin(\phi_d)\hat{\sigma}_x) \quad (2.117)$$

where $\Delta\omega = \omega_Q - \omega_d$ is the drive detuning. This Hamiltonian shows that the axis of rotation can be controlled through the phase ϕ of the microwave drive. Solving for the time-evolution dynamics of \hat{H}' leads to Rabi oscillations, where the qubit rotates between the ground and excited state. This is how the Bloch sphere trajectories in Fig. 2.7 were calculated, with $\Delta\omega/2\pi = 5$ MHz and a constant-amplitude 20 ns drive pulse. If the drive frequency is swept around ω_Q , we would notice that the rotation frequency is higher away from $\Delta\omega = 0$, creating a “chevron” pattern. These kinds of dynamics also occur for two coupled systems, and are studied in Chapter 5.

2.5.3 Decoherence

Experimental quantum systems always interact with the environment, whether through couplings designed in the circuit for control and measurement, or through unwanted interactions. This second category includes a wide variety of effects. The most prominent problem today is relaxation caused by TLSs, which are likely found in the disordered oxide layers at the interfaces between the device and the outside air. Other important issues include couplings to free electric charges, causing charge noise [169, 104, 119]; couplings to stray magnetic fields, causing flux noise [170, 171]; couplings to free space modes or package box modes, leading to relaxation [104, 172]; thermal radiation or cosmic rays which break down Cooper pairs into quasiparticles, leading to both relaxation and decoherence [173, 174, 175, 176]. Designed couplings also lead to decoherence by directly carrying heat or electrical noise from room-temperature electronics, which is why it is critical to design the experimental setup very carefully (see Chapter 3). Even couplings between circuit elements can lead indirectly to decoherence, e.g., through the Purcell effect [177, 178].

In principle, these interactions could be handled by including them directly in the Hamiltonian of Eq. (2.93). However, the extremely large number of degree of freedoms (i.e., coordinates) associated with the various couplings make this approach completely impractical, even numerically. Instead, we want a way to describe the various decoherence channels while including only the Hilbert space of the systems of interest in calculations. Fortunately, such methods exist, although they come with certain concessions in the form of assumptions. One of these approaches is given by the Lindblad equation, which resembles the Schrödinger equation, but instead describes the time evolution of density operators [179, 180]. We write it in diagonal form below:

$$\frac{d\rho}{dt} = -\frac{i}{\hbar}[\hat{H}, \rho] + \sum_k \left(A_k \rho A_k^\dagger - \frac{1}{2} (A_k^\dagger A_k \rho + \rho A_k^\dagger A_k) \right) \quad (2.118)$$

where ρ and \hat{H} are the density matrix and Hamiltonian of the system, and the operators A_k are Lindblad operators and are sometime also called “collapse” operators.

The Lindblad equation is based on the ideas sketched above, whereas we want to describe the dynamics of a system coupled to an environment. To make it tractable, we make the *Born-Markov* approximation, which states that the system and environment are uncorrelated—i.e., described by a product-state—and that the interactions between the two are weak enough that they *stay* uncorrelated, even after some time evolution. This assumption has a few consequences, the most important for us being that we may now ignore the coordinates of the environment and consider the resulting effects on the system ρ only. These effects are included into “collapse” operators A_k , which describe how the environment causes relaxation and dephasing of the system. A second consequence is that the dynamics of the system are time-local, or *Markovian*. The description of a state at a single point in time by a density operator $\rho(t)$ entirely determines the future dynamics. While this may be convenient, it prohibits application of the Lindblad equation to any non-Markovian dynamics; those for which the state of the system would depend on what happened to the environment in the past. In other words, the environment cannot have a “memory” of past interactions with the system. Dynamics between a qubit and a coherent TLS are highly non-Markovian, since near resonance the two systems swap energy back and forth. In this case, a good compromise would be to model coherent TLSs within the system Hamiltonian, and include the effect of weakly coupled ones as Lindblad operators. This is the approach used in Chapter 5, where we determine the coupling parameters of a coherent TLS (or another resonant mode), but still include “background” incoherent relaxation in the model.

The two most commonly encountered Lindblad operators in the context of qubits are the amplitude and phase damping operators. They can be expressed with Pauli operators:

$$A_1 = \sqrt{\Gamma_1} \hat{\sigma}_- \quad (2.119)$$

$$A_2 = \sqrt{\frac{\Gamma_\varphi}{2}} \hat{\sigma}_z. \quad (2.120)$$

The amplitude dampening operator A_1 models energy relaxation, with a rate Γ_1 and the phase damping operator A_2 models dephasing with a rate Γ_φ . If we solve Eq. (2.118) for the state of a qubit in the rotating frame (i.e. $\hat{H} = 0$), we obtain

$$\rho(t) = \begin{pmatrix} c_{gg} + c_{ee}(1 - e^{-\Gamma_1 t}) & c_{ge} e^{-(\Gamma_1/2 + \Gamma_\varphi)t} \\ c_{eg} e^{-(\Gamma_1/2 + \Gamma_\varphi)t} & c_{ee} e^{-\Gamma_1 t} \end{pmatrix} \text{ where } \rho(0) = \begin{pmatrix} c_{gg} & c_{ge} \\ c_{eg} & c_{ee} \end{pmatrix} \quad (2.121)$$

is the initial density matrix. It is clear that A_1 leads to both exponential relaxation and dephasing and that A_2 results only in dephasing. Note here that dephasing simply

means a decay of the off-diagonal elements. We define the T_1 and T_2 relaxation times, which are the ubiquitous qubit quality metrics:

$$T_1 = \frac{1}{\Gamma_1} \tag{2.122}$$

$$T_2 = \frac{1}{\Gamma_1/2 + \Gamma_\phi}. \tag{2.123}$$

Notice that if there is no pure dephasing ($\Gamma_\phi = 0$), $T_2 = 2T_1$; in this case T_2 is said to be T_1 -limited. In frequency tunable qubits, the additional sensitivity to frequency noise usually means that T_2 will be smaller than T_1 . Today's fixed frequency qubits, on the other hand, have very little pure dephasing [181].

The above equations are used to fit qubit experimental data, as is done in Sec. 3.4. Chapter 6 expands on the above concepts to describe the relaxation caused by an ensemble of TLSs for a qubit operated at different frequencies. Section 6.2.2 in Chapter 6 contains a derivation of the qubit relaxation $\Gamma_1^{\text{Q-TLS}}$ caused by a TLS with a particular coupling strength and frequency.

How to Measure a Qubit

In this chapter, we tackle the experimental art of controlling and measuring qubits and resonators in circuit QED. Though the word “art” is used tongue-in-cheek, it is true that the experimental techniques used in circuit QED are somewhat arcane and, to some degree, obscure. These methods have been developed and carried through generations of experimentalists, generally taught to students in the lab by their professor. As a result, that knowledge is not always readily available in the literature and only the most careful reading of the many theses written by students over the years will yield morsels of information. This chapter is not meant to fully solve this problem, but it will attempt to discuss the important details¹ that enter the preparation and execution of superconducting circuit QED experiments.

This chapter begins with a short discussion on device design and fabrication. Then, assuming that we have a chip in-hand, we proceed in the fashion of a *tutorial*, starting from the experimental setup needed, including dilution refrigerator wiring (see Chapter 4 as well) and electronic instruments used for control and readout, and continuing with the basic techniques used to characterize resonators and qubits as if we had just cooled down a device. The last section covers randomized benchmarking, a *de facto* standard used in quantum computing.

3.1 Device Design & Fabrication

We first discuss device design which, for our purpose, essentially consists of all the tasks involved in the creation of a *mask*, such as the one of Fig. 3.1 which corresponds to the design in Fig. 2.5. The mask is created by arranging the various structures needed for the quantum circuit. Those include, for example, transmission lines for measurement

¹Though many will unavoidably be omitted or simply forgotten...




Figure 3.1: Complete mask design for a two-qubit device. The shaded areas get exposed to ultraviolet light; they correspond to gap areas. The rest of the mask corresponds to metal. There are test resonators on the measurement line and test SQUIDs on the sides. The normal state resistance R_n of test SQUIDs is measured to estimate the qubit frequency prior to a cooldown.

and control pulses, transmission lines for resonators, cross-shaped islands for qubits—including a loop for the tunable SQUID element—and pads to connect the device to external electronics. Then, we cover fabrication and present the parameters of the qubit used in this thesis.

The very first step in creating a design is to choose the “high-level” parameters of the circuit according to the experiments that we want to run. These high-level parameters include the layout of the circuit, the number of qubits and resonators and their frequencies, the coupling strength between the elements, etc.

Once these parameters have been chosen, the “low-level” parameters must be determined. That includes the parameters of the individual circuit elements: transmission line impedance, capacitors, inductors, junctions. These parameters are generally fixed by the geometry of the metal traces and gaps on the chip. Thus, the various elements should be modelled with CAD tools, and the resulting parameters calculated via an appropriate physical simulation. For example, the coupling capacitance between a qubit and a resonator may be determined via electromagnetic simulations with, e.g., Ansys Q3D [182] or FastCap [183]. The mutual inductance between the flux tuning line and the SQUID can be simulated with FastHenry [184]. The resonance frequency of distributed element resonators should also be simulated, e.g., with HFSS [185]. If possible, those circuit parameters should then be validated by running quantum simulations based on the Hamiltonian derived according to the method of Chapter 2. These simulations can confirm that the frequencies and coupling strengths selected in the first step are correct for the experiment.

The final step in the design is to assemble the various geometries of the components into a full chip design. The various control and readout transmission lines must be connected to pads, and other test structures may be added. Figure 3.1 shows an example of a completed *mask*. The structure arrays on the left and right are test SQUIDs. The normal resistance of those SQUIDs can be measured after fabrication, giving an estimate of the Josephson energy [via Eq. (2.49)]. If the energy is not what was designed, the junction evaporation step will need to be adjusted.

The fabrication of a qubit chip comprises many steps, with the two main ones being the optical lithography, used to create the “larger” features, and the electron-beam lithography, where the Josephson junctions are formed. The full recipe used for the sample of this thesis is presented in Appendix A.1. The qubit is fabricated by depositing and patterning thin-film Al on a thoroughly-cleaned silicon wafer; we use the same cleaning process as in our work of Ref. [114]. The mask is used during optical lithography to expose a resist layer resting on top of a metal layer (which itself was evaporated on the surface of a wafer). When the wafer is exposed, ultraviolet light passes through the open areas of the mask and changes the chemical properties of the resist, allowing it to be dissolved away with a solvent and revealing the metal layer below. This metal itself is then etched away, creating gaps in the layer. More details on optical lithography fabrication can be found in Ref. [114]. The Josephson tunnel junctions are fabricated with electron-beam lithography in a later step using a standard double-angle Niemeyer–Dolan technique [186, 187]. This technique requires two additional metal evaporations along with an intermediate oxidation step to create the insulating barrier. See Ref. [108] for more details on electron beam lithography. Photographs of a fabricated qubit—including the SQUID and junctions—can be seen in Fig. 3.2. A device with resonators only solely requires optical lithography; a fabricated resonator chip can be seen in Fig. 4.15.




Figure 3.2: Images of a fabricated sample. (a) Optical image of an Xmon transmon qubit, with the drive line on the left, measurement resonator coupler above, and flux bias line below. The cross-shaped island constitutes a capacitor to ground while the SQUID at the end of the lower arm acts as a tunable nonlinear inductor. (b) Scanning electron microscope (SEM) image of the SQUID. The brighter layer is the aluminum deposited during the junction fabrication. (c) SEM image of a Josephson junction made with a Dolan bridge.

A superconducting qubit identical to the one used in Chapters 5 and 6 is pictured in Fig. 3.2. The qubit consists of an island in parallel with a SQUID. The island forms a capacitor that is composed of two intersecting CPW segments in the shape of a Greek cross, where each segment has length $L = 376 \mu\text{m}$. One segment is formed by a center conductor, or strip, of width $S = 24 \mu\text{m}$ and is separated by a distance $W = 24 \mu\text{m}$ from a ground plane on each side of the strip. The capacitance of the island is $C_q \approx 100 \text{ fF}$ (corresponding to a single-electron charge energy $E_c/h \approx 188.6 \text{ MHz}$). The qubit capacitor is connected in parallel with the SQUID, which is made of a loop interrupted by two parallel Josephson tunnel junctions with critical current $I_{c0} \approx 17.4 \text{ nA}$ (corresponding to a Josephson energy $E_J/h \approx 8.6 \text{ GHz}$) for each junction. The SQUID forms the inductive element of the qubit. The qubit state can be measured by means of a capacitively coupled readout resonator, with capacitance $C_M \approx 3.4 \text{ fF}$. The qubit can be energized by means of X or Y microwave pulses (see Section 2.5.2), which are applied through a capacitive network with coupling capacitor of capacitance $C_{XY} \approx 100 \text{ aF}$.

We are able to tune the frequency of the qubit in situ during an experiment by threading the SQUID loop with a flux $\phi_Z = M_Z i_Z$, where $M_Z \sim 3 \text{ pH}$ is the mutual inductance between the loop and an external circuit with current i_Z . A quasi-static flux bias ϕ_Z^{qs} allows us to set the qubit frequency $f_q(\phi_Z^{\text{qs}})$, i.e., the qubit bias point. The

qubit parameters given above result in a zero-bias $f_q(\phi_Z^{\text{qs}} = 0) \approx 4.8$ GHz.

3.2 Dilution Refrigerator Wiring

You have a chip; now what? As may be expected, if you want to run experiments you will need a well-equipped cryogenics laboratory. The centerpiece of such a lab is the *dilution refrigerator* (DR), a voluminous apparatus that will reach the very low temperature of approximately 10 mK. Figure 3.3 shows the DR used for the experiments of this thesis. A DR consists of multiple horizontal metal plates in a vertical arrangement, with each plate corresponding to a different temperature stage. The very top outer stage is at room temperature, and lower plates get progressively colder. The first stage is maintained at ~ 50 K, the second stage is cooled to ~ 3 K, the third stage is called the still, and is at ~ 700 mK, the fourth stage, the cold plate, is at ~ 100 mK and the fifth stage—last and coldest—the mixing chamber, is cooled to ~ 10 mK. Dilution refrigerators depend on two separate cooling systems. The first system is used to cool the 3 K stage... to 3 K. For that stage, DRs use either a liquid helium heat exchanger or a pulse tube cryocooler (if a pulse tube is used, the DR is a “dry” fridge). The second stage indirectly cools the first stage, which is why its temperature is in-between room temperature and 3 K. The second cooling system, which gives dilution refrigerators their name, is a heat exchanger exploiting the properties of helium-3 and 4 dilution. The top part of the dilution unit must be maintained at 3 K by the second stage while the bottom part directly cools the mixing chamber to ~ 10 mK.

The wiring and electrical equipment inside the fridge are used to relay signals from room-temperature instruments to the chip mounted at the mixing chamber stage. The many instruments that generate and read out those signals are located outside the fridge, in the lab. A complete diagram of the experimental setup can be seen in Fig. 3.5. This section will cover various aspects of the DR wiring with the goal of providing a clear picture of the requirements that must be met to perform high-quality experiments. Additional wiring topics are covered in Chapter 4, including the package and chip 3D wires, the package holder, and magnetism.

3.2.1 Thermal Engineering

Superconducting quantum devices need a low operating temperature for many reasons. The obvious one is that the metal must become superconducting for the Josephson effect to work. For aluminum, the transition temperature is $T_c \sim 1.2$ K. The temperature however, is also constrained by the frequencies of the devices. Superconducting qubits and resonators are designed with a frequency in the 4 – 10 GHz range. To a good approximation, the thermal occupation of the first excited state $|e\rangle$ is described by the




Figure 3.3: Picture of the open dilution refrigerator showing the wiring on the mixing chamber stage (bottom most plate), the cold plate, the still stage, and the 3K stage (aluminum, at the top). The quantum socket is visible at the very bottom, just below the mu-metal can lid. The back of the two instrument racks is seen in the background.




Figure 3.4: Steady-state occupation probability of the first excited state of a qubit. At higher temperatures the qubit has a significant probability to be found in the first excited state; this is generally undesirable.

Boltzmann distribution:

$$P_e = \frac{e^{-\hbar\omega_q/k_B T}}{1 + e^{-\hbar\omega_q/k_B T}} \quad (3.1)$$

where k_B is the Boltzmann constant. We plot the probability of occupation as a function of temperature for a qubit with frequency $\omega_q/2\pi = 5$ GHz in Fig. 3.4. Below 30 mK the occupation probability is $< 0.1\%$; above 40 mK however, the probability starts rising very rapidly. A high excitation probability is generally undesirable because it makes the steady-state of the qubit unpredictable, among other issues [188]. Experiments generally require knowledge of the start state, thus if the qubit is thermally excited an additional initialization step involving nondemolition measurement will be required [189, 190, 191].

One additional important problem with operating temperatures between T_c and 40 mK concerns *quasiparticles* [22]. For any BCS superconductor, thermal noise caused by temperatures above absolute zero will tend to excite Cooper pairs and break them into quasiparticles, thereby forcing the superconductor out of the ground state and causing dissipation [192, 173, 104]. This effect also depends on the device frequency, with lower frequencies commanding lower temperatures.

It is clear then, that the chip should be cooled well below T_c for optimal performance. To achieve low device temperatures, we apply two directives in the thermal design: First, we maximize the cooling of the chip by thermally anchoring it as well as possible to the mixing chamber stage of the DR. Second, we minimize the heating of the chip caused by electrical connections to instruments outside the DR. In most setups, thermal anchoring is accomplished by creating a strong mechanical connection between the package where

the chip is housed and the mixing chamber with a high thermal conductivity metal and through the thermal conductivity of the outer conductor of the connectors to the ground plane of the chip (e.g., with wire bonds). In our setup, we use the quantum socket to mount the device to the DR. The quantum socket—which comprises the package, package holder, and 3D wires—is described extensively in Ch. 4.

At this time, the control and readout signals that we send to the device come from electronic instruments located outside the fridge, at room temperature². Because the cables and wires that carry signals to the device are made of metal, they form a direct thermal path. The heating caused by this path is managed primarily by careful selection of the wiring materials. For example, we use low thermal conductivity stainless steel coaxial cables to connect the different stages of the DR. For DC wires, we choose phosphor bronze, which also has a low thermal conductivity. For the readout line, which must carry very weak signals from the sample back to room temperatures, we cannot use stainless steel due to its poor electrical conductivity. Instead, we use a superconducting niobium cable which, obviously, has excellent conductivity but also poor thermal conductivity. The electrical and thermal conductivity of metals usually go hand in hand; superconductors, when in the superconducting state, are the exception.

At each DR stage, we thermally anchor the wiring such that its temperature matches the temperature of the stage. The cables are therefore progressively cooled to lower and lower temperatures. For coaxial cables, thermally anchoring the outer conductor is easy: because we use the metal structure of the fridge as the electrical ground, we may simply attach the outer conductor to each plate. This is done with coaxial feedthroughs bolted to the stage to which cables connect using an SMA connector. The inner conductor is more tricky to thermalize since it is electrically isolated from the outer conductor by an insulator—usually polytetrafluoroethylene (PTFE)—with low thermal conductivity. Fortunately, attenuators, which are discussed in Sec. 3.2.2, are used at every stage and provide a second thermal pathway. We ensure that the attenuators themselves are well thermalized by placing their body in contact with a custom copper fixture mounted to the stage. The fixtures can be seen in-between the attenuators on the mixing chamber stage in Fig. 3.3. To anchor DC wires, we wind them around custom copper bobbins that are mounted to the stage plates. These bobbins can be seen on each stage in Fig. 3.3, on the right, in front of the stainless steel posts.

We finish this section by mentioning the effect of thermal radiation. Indeed, although the stages might be thermally isolated from each other thanks to the use of low conductivity material, blackbody radiation emitted by the higher temperature stages could find its way to the mixing chamber, heating up the stage and exciting quasiparticles. To prevent this, the DR stages are isolated from each other by metal

²In the future it might become possible to move certain instruments inside the fridge; for example, signal generators based on single flux quantum electronics [193, 194].

cans that shield the lower stages from the higher stages. In addition, we coat the inside of the mixing chamber can with a black epoxy mixture. This helps ensure that stray radiation is absorbed by the stage rather than by the sample [195].

For more details on the various experimental techniques used for low-temperature measurements, see Ref. [196].

3.2.2 Amplification & Attenuation Chains

The signals sent to the chip should be free of noise and other distortions, in other words, they should have a large signal-to-noise ratio (SNR). However, the energy contained in the signal itself should be relatively small: a single 5 GHz photon has an energy of 3.3×10^{-24} J, or about 2×10^{-5} eV. Assuming that this photon is sent as a pulse of length 20 ns, as might be the case for a qubit π -pulse, the average signal power is

$$S = \frac{\hbar\omega}{\Delta t} = 1.7 \times 10^{-16} \text{ W} = -128 \text{ dBm} \quad (3.2)$$

where dBm is the *decibel-milliwatts* unit. We can calculate the corresponding room-temperature (~ 300 K) Johnson–Nyquist noise power as [197, 198]

$$N = k_{\text{B}}TB = 2.1 \times 10^{-13} \text{ W} = -97 \text{ dBm} \quad (3.3)$$

where $B \sim 1/\Delta t$ is the bandwidth corresponding to the signal. The thermal noise is therefore three orders of magnitude larger than the signal. The energy of a single photon is tiny, so this result is not surprising. Instead of sending a single photon, we thus need to emit a much higher signal power at room temperature and *attenuate* this signal as it travels down the fridge until it reaches the desired power at the device. As a comparison, the noise power at 10 mK is given by Planck’s radiation law (including the zero-point energy):

$$N = \hbar\omega \left(\frac{1}{2} + \frac{1}{e^{\frac{\hbar\omega}{k_{\text{B}}T}} - 1} \right) B = 8.3 \times 10^{-17} \text{ W} = -131 \text{ dBm}. \quad (3.4)$$

The guiding equation in choosing the specifics of the attenuation chain is Friis’ formula [199], which calculates the output SNR given an input SNR and the noise added by each stage in the chain. For a single stage, the equation is

$$\text{SNR}_o = \frac{S_o}{N_o} = \frac{G_s S_i}{G_s N_i + N_s} \quad (3.5)$$

where the subscripts i and o refer to the input and output of the stage, N_s is the noise power added by the stage [given by Eq. (3.4) for the temperature of the stage], and G_s

CHAPTER 3. HOW TO MEASURE A QUBIT

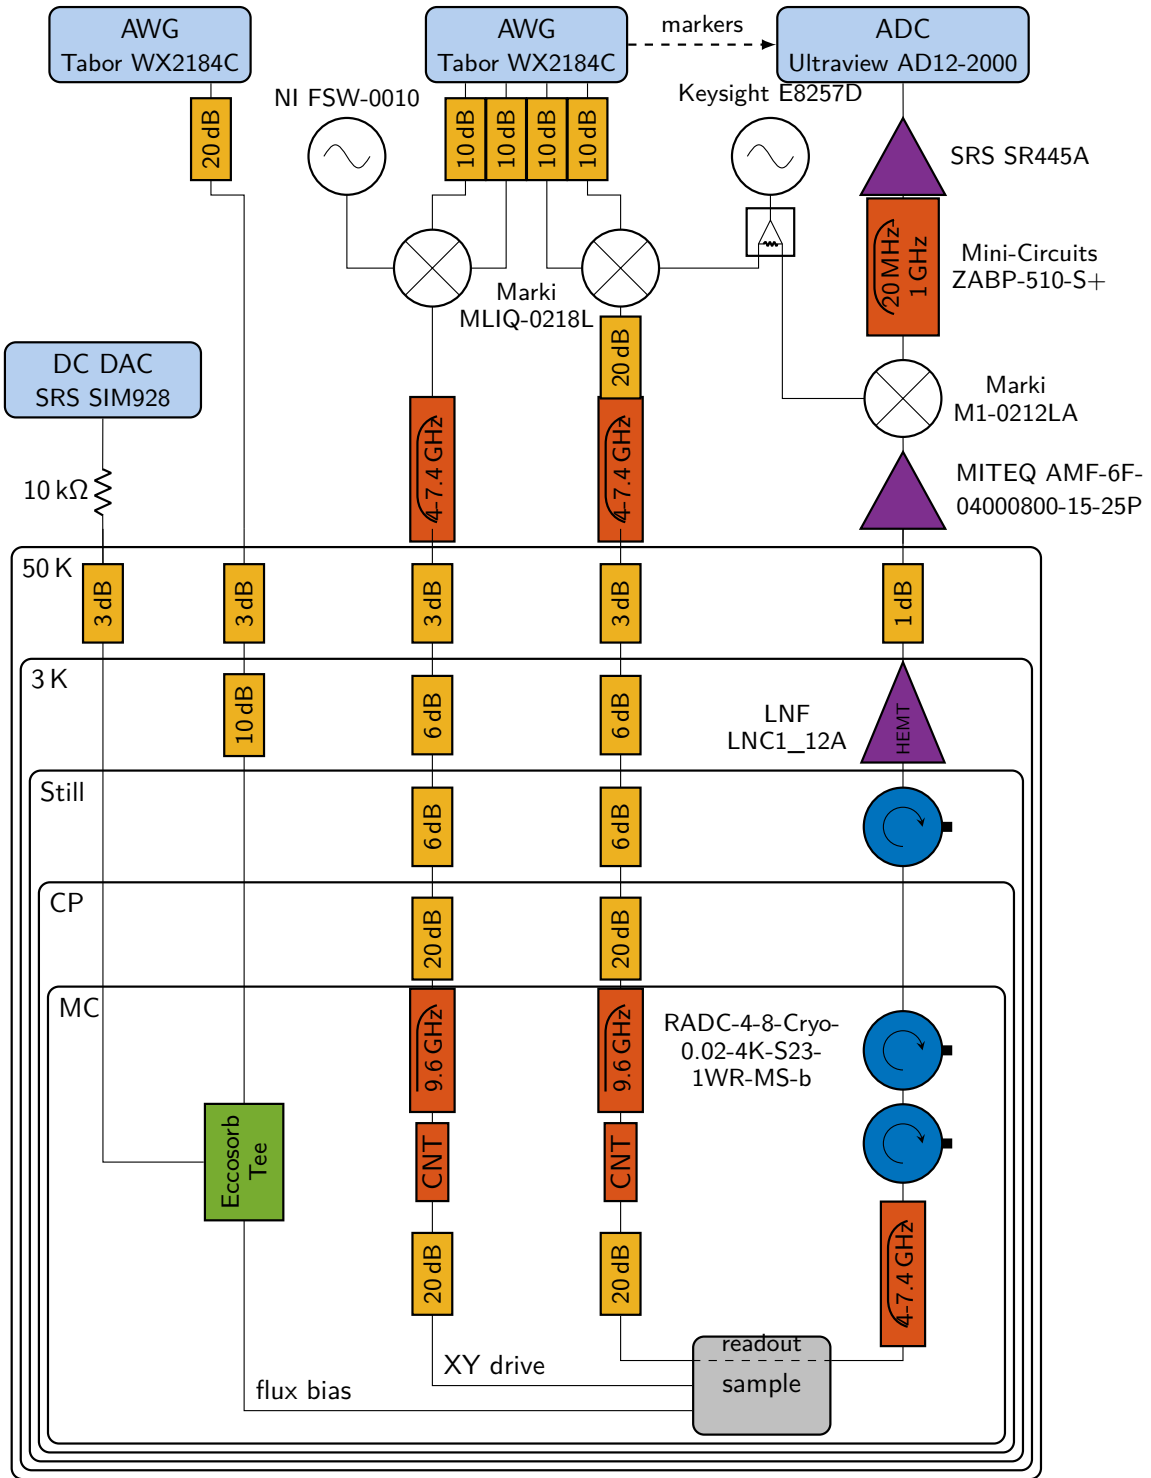


Figure 3.5: Dilution refrigerator wiring schematics of instruments and control lines.

is the gain (< 1 for attenuation). At each stage, the signal and noise at the input are scaled by the gain and extra thermal noise is added. This equation may be cascaded for multiple stages by feeding the signal and noise at the output of a stage to the input of the next. Since the temperature of the stages drops, progressively smaller amounts of noise are added thereby maintaining a high SNR.

Ideally, we would attenuate only at the mixing chamber stage to obtain the highest SNR. This is not possible, however, because the mixing chamber has only a very small amount of cooling power available and would not be able to handle the large amount of energy that this would dissipate. In addition, attenuators have a second purpose: They help thermalize the inner conductors of coaxial cables, which they mechanically connect to the DR plates via a resistor. Taking those two points into account, the best scheme is thus to gradually attenuate the signal with an attenuator at each stage, preferring lower attenuation at the higher temperature stages and higher attenuation at the lower temperature stages. The attenuation values chosen for the experiments of this thesis can be seen in Fig. 3.5, where two attenuation chains are depicted: one for qubit control (the XY line), and one for readout. We use cryogenic SMA attenuators made by XMA Corporation.

The pulses going out of the device and back to the lab for readout encounter a similar problem: they must be amplified. If we didn't amplify them, they would quickly be overpowered by thermal noise as they travel to the warmer stages. Friis' formula also applies to amplification chains, except that this time, the gain is > 1 and the added noise is determined by the noise of the amplifier itself³. Amplifying at the lowest temperature stage would be the most advantageous, but, once again, thermal aspects must be considered, as amplifiers dissipate a lot of heat. At the mixing chamber stage, only an extremely low power amplifier may be used, such as a Josephson traveling-wave parametric amplifier [95]. At higher stages, more power-hungry high-electron mobility transistors (HEMT) amplifiers may be used (a HEMT will dissipate 10 mW). In our setup we only use a Low-Noise Factory cryogenic HEMT at the 3 K stage, in addition to room-temperature low-noise amplifiers (LNA). To minimize losses, we use superconducting coaxial cables from the mixing chamber to the HEMT. Because a HEMT still emits a large amount of thermal noise, we install three *circulators* on the amplification line. Circulators are ferrite-based nonreciprocal three-port devices; signals entering one port only exit at one of the two others. In this way, signals from the device enter port one, exit out of port two and continue on up the line, whereas thermal noise coming from the HEMT up the chain enter port two and exit out of port *three*, which is simply terminated with a 50Ω load. Two circulators can be seen on the mixing chamber stage in Fig. 3.3 encased in a gray Cryoperm package for magnetic

³Amazingly, the “noise temperature” of an amplifier—a measure of the noise added—can be lower than the temperature of the amplifier itself!

shielding.

Both attenuation and amplification chains are heavily filtered to prevent as much noise as possible from reaching the sample. The XY and readout lines are filtered with microwave low-pass and bandpass filters, in addition to custom carbon-nanotube filters (CNT) which block frequencies in the infrared range and higher [200]. The flux bias line is filtered with a custom magnetically loaded Eccosorb filter.

3.2.3 Eccosorb Tee Filter

The Eccosorb tee discussed in this section is a novel cryogenic filter that was designed and fabricated for qubit flux biasing. As explained in Sec. 2.2, superconducting qubits can be made frequency-tunable by designing them with a loop interrupted by two Josephson junctions (a SQUID). The qubit frequency is then controlled by applying a magnetic flux through that loop with a current line. Such a line is visible coming from the bottom in Fig. 3.2 (a). Because that line is inductively coupled to the loop, any current noise will therefore manifest as flux noise, changing the qubit frequency and causing dephasing. It is thus critical to minimize noise as much as possible.

The Eccosorb filter is a *tee*, meaning that it combines DC and AC signals. One of the ways in which we minimize noise is by setting the idle frequency of the qubit with a pure DC bias voltage coming from a battery. As a result, this voltage is extremely low-noise, and, in particular, it is free of 60 Hz harmonics that would otherwise come from mains electricity powered instruments (wall power). AC signals are required to send pulses that set the instantaneous frequency of the qubit. They are emitted by an arbitrary waveform generator at high power and can be more strongly attenuated. The purpose of the idle and pulsed bias is explained in Chapter 5.

There are three stages of filtering: the first stage, which applies only to the DC signal, is a lumped-element low-pass LC filter made with surface-mount devices (SMD). Its 3 dB cutoff frequency is $f_c = 600$ kHz. The second stage is a distributed low-pass stepped-impedance filter, with $f_c = 800$ MHz. While these filters are very effective below and near f_c , their attenuation diminishes at high-frequency (IR frequencies), and they eventually become completely transparent. To block high-frequency radiation, we cast Eccosorb CRS-117 [201] in the package. Eccosorb is an EM absorber, it consists of a silicone rubber base loaded with ferromagnetic particles, and therefore has a high magnetic permeability [202]. The Eccosorb filter stage has $f_c = 3$ GHz (determined by the package length), and attenuation keeps increasing for higher frequencies. The Eccosorb absorber is also characterized by a high thermal conductivity and low out-gassing, making it particularly suited to cryogenic applications. To ensure maximum absorption, we use a thin polyimide film as substrate, DuPont Pyralux AP 8515R [203], which is 25.4 μm thick. The PCB with the copper traces (18 μm thick) was manufac-




Figure 3.6: Fabrication of the Eccosorb tee filter. Left: Picture showing the package before the top Eccosorb layer is poured in. We can see the lumped-element LC filter on the DC side (left) and the distributed stepped-impedance filter on the AC side (right). Note that the DC signal goes through the distributed filter. A broadband conical inductor is used to prevent the AC pulses from traveling back out of the DC port. The PCB is translucent, letting us see the Eccosorb-filled cavity below. Right: Picture showing the package with the top Eccosorb layer poured-in, just before curing.

tured by Printech Circuit Laboratories in the United Kingdom. Figure 3.6 shows images of the filter during fabrication.

Figure 3.7 shows qubit spectroscopy before and after installation of the CNT and Eccosorb tee filters. We can see impressive improvements in the resonator linewidth. This suggests that thermal noise was strongly reduced. For more details on spectroscopy experiments, see the next sections.

3.3 Control & Measurement

The dilution refrigerator is now fully wired and the care we have taken to guarantee that noise and heating effects are well-managed should lead to excellent measurement results. The last thing we must discuss is the instrumentation needed to perform those measurements. To keep this section manageable, we will discuss two kinds of measurements: first, *continuous-wave* (CW) measurements, which can be done entirely with a network analyzer, and second, *time-domain* measurements, which require an arbitrary waveform generator (AWG) to shape pulses.

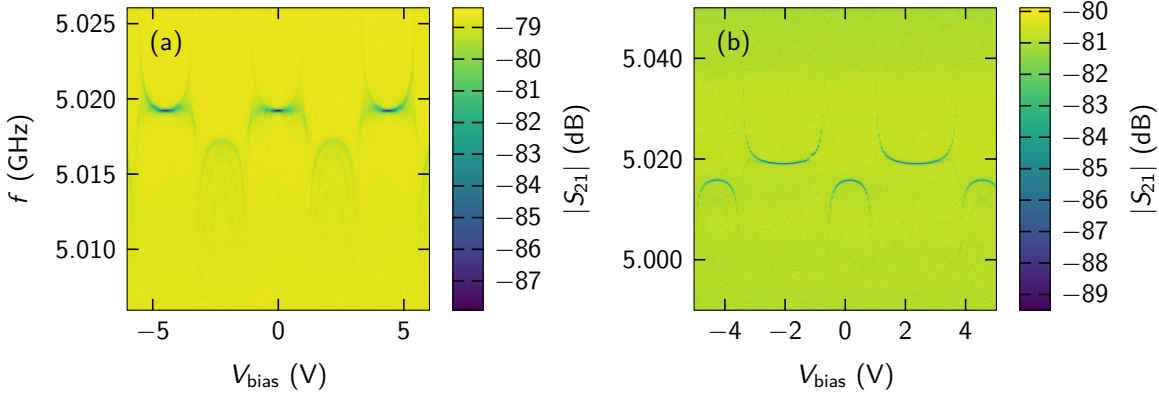


Figure 3.7: Qubit spectroscopy of the same sample (a) before and (b) after installation of infrared filtering. For this experiment, we measure the resonator with a VNA at low power while varying the bias voltage of the qubit. The spectroscopic lines of the resonator are much sharper and better defined in (b), corresponding to a lower level of noise.

3.3.1 Continuous Wave Measurements with a Network Analyzer

A vector network analyzer (VNA) is a convenient all-in-one instrument: It takes care of both sending and receiving signals. Because the received signals can be filtered heavily, VNAs boast excellent sensitivity and extremely low noise floors. Network analyzers generally comprise at least two ports, each capable of both transmitting and receiving simultaneously. They are therefore able to characterize the S-parameters of a network [115]. S-parameters are ratios of the voltage received at a port to the voltage emitted at a port, and are measured as a function of frequency. $S_{i,j} = V_i/V_j$ thus represents the amplitude of the voltage wave measured at port i divided by the amplitude of the voltage wave emitted at port j . For example, S_{11} measures the reflectivity of the network at port 1. S_{21} measures the transmission from port 1 to port 2. A vector network analyzer is capable of measuring the amplitude and phase of signals, the S-parameters are therefore complex quantities.

One important aspect to keep in mind when making VNA measurements is that, though an S-parameter is a relative measurement, i.e., a ratio, and is therefore unitless, it still may depend on the absolute voltage that was sent by the transmitter. In other words, S-parameters measured with a particular output power may not be representative of other powers. In fact, this is exactly what we see when measuring resonators, which have a power-dependent transmission (because of TLS, see e.g., Fig. 4.17).

Internally, a VNA contains a plethora of components: microwave sources to send signals, mixers for frequency up- and down-conversion, filters for noise reduction, and

attenuators for power control. The transmitter sends a CW signal at the selected frequency. Generally, we will *sweep* through multiple frequencies in order to characterize a network over a particular bandwidth. The time spent at each frequency depends on the settings of the receiver. The receiver is generally a super-heterodyne design. When the incoming signal enters the receiver (after exiting the network), it is mixed with a local oscillator (LO), which is simply another CW signal generated by a source. This mixing results in waves at the sum and difference of the input signal and LO frequencies. Only the signal at the difference frequency is kept; this is a lower frequency, called the intermediate frequency (IF). For example, if the signal frequency is 6.1 GHz and the LO frequency is 6 GHz, the resulting IF is 100 MHz. It is at this stage that the IF filter is applied. Because the tone emitted by the transmitter is almost a pure sine wave, it is extremely narrowband (<1 Hz). It is therefore possible to filter the received signal with a narrow bandpass filter without losing useful information. The only disadvantage is that a low IF bandwidth (IFBW) will cause the measurement to take longer due to the signal needing more time to pass through the filter. After the IF filter, the signal is digitized by a detector and its amplitude and phase is determined.

A VNA is indispensable for many types of microwave measurements. We have used a VNA to characterize almost all the microwave components we purchased or made: cables, attenuators, power dividers, circulators, filters, etc. We used a VNA extensively to characterize the quantum socket, as is detailed in Chapter 4. Finally, superconducting resonators are also measured with a VNA.

VNA measurements may be augmented by adding a second microwave source (also a CW tone). As an example, such a tone could be used to drive a qubit while the measurement resonator is measured with a VNA. This type of experiment is called two-tone spectroscopy. For a two-tone experiment, we set the VNA to measure at a single frequency, that of the resonator. When the drive tone matches the qubit frequency, the resonator will move due to the dispersive shift (see Sec. 2.3); this is reflected by S_{21} . Note that because both tones are CW, they are always “on”. As a result, the qubit cannot be excited to $|e\rangle$; instead, the state is continuously rotating between $|g\rangle$ and $|e\rangle$ until decoherence results in a statistical 50-50 mixture. On the side of the resonator, the continuous tone creates a coherent state which also devolves into a statistical mixture described by a density matrix, with multiple levels incoherently occupied [204]. The result is that, on average, the resonator displacement is dampened, and therefore the contrast of the measured signal is lower. Despite this disadvantage, the simplicity of VNA measurements make them very useful, especially in the initial stages of a qubit sample characterization. This is discussed in Section 3.4.

3.3.2 Time-Domain Measurements

Time-domain measurements are meant to rectify the principal problem of continuous wave measurements: they allow the qubit and the resonator to reach their ground states and to be controlled independently over time. Instead of applying continuous tones that necessarily overlap, we want to send *pulses*, one after the other. The second advantage is that we are now able to generate pure states and measure them before they decohere. These two points are, of course, requirements for gate-based quantum computation.

The other aspects of the pulsed readout setup are otherwise similar to the VNA: We still want the ability to measure the amplitude and phase of the voltage transmitted through the DR. In effect, we simply want to build a “pulsed VNA”. The main part of this setup is the arbitrary waveform generator. An AWG is a digital-to-analog converter (DAC): we give it a list of voltage points, and it outputs them. The rate at which an AWG can emit voltage points is the clock rate. On the measurement side, we use an analog-to-digital converter (ADC): the ADC measures the incoming voltage and digitizes it, giving us a list of voltage points. The rate at which the ADC can digitize points is the sampling rate.

At the time that the setup was built (circa 2015), the fastest available AWGs and ADCs had rates of a few Gsps (giga samples per second)⁴. These rates are insufficient to directly generate and measure signals at ~ 5 GHz. Instead, we use microwave mixers on the transmitter side to upconvert the signal generated by the AWG before sending it to the device in the DR. We do the opposite on the receiver side: we use a mixer to downconvert the signal before it is digitized. Many types of transmitter and receiver designs exist, we now describe the one we have built.

The transmitter uses an IQ (in-phase and quadrature) mixer with single-sideband (SSB) modulation. This means that the AWG emits two identical signals that are 90° out of phase at an intermediate frequency f_{IF} between 100 and 300 MHz. The IQ mixer multiplies the two signals with a microwave carrier, or local oscillator (LO), $f_{LO} \sim 5$ GHz. This multiplication results in each signal being duplicated in two *sidebands* at the sum and difference frequencies $f_{LO} \pm f_{IF}$. This is not ideal because the qubit will then be driven at two different frequencies. With SSB modulation, we can eliminate one of the two sidebands by destructive interference. This is possible because the signal sent to the Q port of the mixer is transmitted with an additional phase difference of 90° relative to the I port. By choosing the sign of the relative phase emitted by the AWG, we can suppress either the left or right sideband. Note that IQ mixers are not perfect, the exact phase difference between the I and Q signals must be properly calibrated to obtain good sideband suppression. In addition, the LO signal itself can also leak through the

⁴There are very fast (>20 GHz) AWGs and ADCs, but they are extremely expensive.

mixers. This can be minimized by calibrating the DC level of the signal going in the I and Q ports of the mixer.

The transmitter design described above can send high-quality pulses to the DR. In our setup, we use it both to drive qubit rotations and for resonator readout pulses. For qubit pulses, control of the rotation axis is achieved by digitally setting the phase of the IF cosine generated by the AWG. This is substantially more accurate than varying the amplitude of the I and Q signals. The schematics of Fig. 3.5 show the instruments used for the transmitter. The IF pulses are generated with a Tabor 2.3 Gsps four-channel 14-bit AWG (Tabor WX2184C). We use a National Instruments source for the qubit drive LO (NI FSW-0010) and a Keysight source for measurement LO (Keysight E8257D). The same model Marki IQ mixer is used for upconversion (Marki IQ MLIQ-0218L).

The receiver is built almost exactly like the inverse of a transmitter. We use the same LO signal to downconvert the incoming pulse back to f_{IF} , although a normal (non IQ) mixer is used. Once again, two “sidebands” are generated in the multiplication, however, the sum frequency is high, near $2f_{\text{LO}}$. It is therefore very simple to filter out the high-frequency image. We use a bandpass filter to additionally suppress any DC and near-DC frequency components before we digitize. In order to obtain the amplitude and phase of the voltage samples—or, equivalently, the I and Q samples—we use a custom-programmed digital downconversion (DDC) scheme.

The first step, directly after digitization, is to decimate the signal in order to obtain a sample rate equal to $4f_{\text{IF}}$. The ADC clock is supplied by a 1 GHz reference signal coming from the readout LO and is internally doubled. The sample rate of the digitized data is thus 2 Gsps. If $f_{\text{IF}} = 100$ MHz, we decimate by a factor of five. The decimation is executed with a finite impulse response (FIR) digital filter which effectively averages every five samples together. The next step is the downconversion, which is very simple to perform. Because the sample rate is $4f_{\text{IF}}$, multiplying with a sine and cosine at f_{IF} is equivalent to multiplying by ± 1 and 0. Thus, we transform every group of 4 samples identically, with samples $(s_i, s_{i+1}, s_{i+2}, s_{i+3})$ being mapped to I and Q as

$$(I_{i/2}, I_{i/2+1}) = (s_i, -s_{i+2}) \quad (3.6)$$

$$(Q_{i/2}, Q_{i/2+1}) = (s_{i+1}, -s_{i+3}), \quad (3.7)$$

where i is the group index. Notice that the downconversion results in another reduction of the sample rate by a factor of two simply because we are creating a single complex sample out of two real ones. A critical element that should not be forgotten is that the digitizer must be synchronized to the AWG for the phase measurements made over multiple pulse repetitions to be consistent. This is accomplished by triggering the ADC digitization on a marker signal positioned at the same point in time for each measurement pulse. This signal is sent via a separate cable connecting the AWG marker output to the ADC trigger input. As a result, the ADC always measures the first point

of the incoming waveform “at the same time,” allowing relative phase changes to be detected.

The output of the DDC step is the baseband (DC) complex envelope of the pulse. For the purpose of qubit or resonator measurement we average the envelope over the length of the pulse, clipping out the leading and trailing edges. The result is thus a single complex IQ point representing the amplitude and phase of the transmitted signal in millivolts. To obtain S_{21} , we divide by the amplitude of the pulse emitted by the AWG. Note that while we choose the envelope of the emitted readout pulse to be a perfect rectangular shape with Gaussian leading and trailing edges, the envelope of the measured pulse will look different. The exact portion of the received pulse over which to average might therefore need to be adjusted, e.g., depending on the resonator Q_c .

Schematics of the receiver are shown in the top right of Fig. 3.5. Downconversion is done with a non-IQ Marki mixer (Marki M1-0212LA). The IF signal is filtered by a Mini-Circuits 20 MHz-1 GHz bandpass filter (Mini-Circuits ZABP-510-S+). The IF signal is further amplified by an SRS IF amplifier (SRS SR445A) and then digitized by an Ultraview 2Gsp/s 12-bit ADC (Ultraview AD12-2000).

Figure 3.8 shows a picture of the transmitter and receiver setup. The large blue power divider is used to split the LO signal and drive the transmitter and receiver mixers. We installed directional couplers at the output of the upconversion mixers (only one pictured here) to monitor the mixed signals in real time on an oscilloscope.

With time-domain experiments, we usually must repeat the measurement many times, either to achieve lower noise by averaging, or to collect qubit readout statistics (see the next section). In the experiments of this thesis, we usually acquire between 500 and 1000 readout pulses per “point”. Because the qubit and resonator must fully decay back to their ground state before a new measurement can be made, the repetition rate of the readout pulse is limited by relaxation times. We use repetition rates between 2 and 4 kHz, resulting in acquisition times of less than 0.5 s per point.

3.3.3 Qubit Readout

As was briefly explained at the end of Sec. 2.3, readout in circuit QED is usually done by taking advantage of the qubit state dependent frequency of a coupled resonator. There are in fact two different methods that make use of this effect in very different power regimes. The first method is the low-power dispersive readout that was initially envisioned [102, 205]. The low-power readout is a nondemolition process [206, 207]. During the readout, the qubit state is projected onto the measurement basis, and, if measurement is repeated, it should yield the same state. This is obviously advantageous for any kind of experiment where the qubit needs to be operated on after being measured, e.g., for state initialization [191] or quantum error correction [59, 108, 109]. The major disadvantage of the low-power readout is the low signal power, which needs




Figure 3.8: Readout transmitter and receiver setup. The IQ mixer on the left multiplies the LO with the AWG signals at the I and Q ports. The mixer on the right downconverts the signal coming from the DR with the same LO. An IF bandpass filter is used for image rejection. Two room-temperature amplifiers and a microwave switch can be seen in the background.

to be on the single-photon level. Consequently, we must either forgo single-shot readout and instead average the measurement signal over many shots, or we need to install a near-quantum limited amplifier at the mixing chamber stage.

For the low power readout, we send a pulse at one of the displaced resonator frequencies and measure the transmitted signal. For example, we can choose to measure at the eigenfrequency that the resonator has when the qubit is in the excited state. In that case, when the qubit is in $|g\rangle$, the resonator will not be at the readout frequency, leading to a large S_{21} value. When the qubit is in $|e\rangle$, the resonator will absorb the signal, leading to a small S_{21} value. The different transmission amplitudes are thus correlated with each qubit state. Figures 3.9 (a) and (b) show the magnitude of the transmitted power $|S_{21}|$ averaged over many shots as a function of the readout pulse frequency and power for the qubit initialized in the ground (a) and excited (b) state. At low power, the length of the pulse is limited by qubit decoherence, which is itself increased during the measurement [208]. Notice that, unlike the prediction of Eq. (2.86), the resonator




Figure 3.9: Qubit readout characterization. The top row shows time-domain (pulsed) measurements of S_{21} transmission for the qubit initialized in $|g\rangle$ (a) and $|e\rangle$ (b). The S_{21} samples are averaged together and we plot the magnitude. The state-dependent dispersive shift is visible at low power. In (c), the same S_{21} samples are used to calculate the single-shot readout visibility.

is *not* evenly shifted by $\pm\chi$ around ω_R . The correct resonator frequency at low power in Fig. 3.9 (b) is given by Eq. (2.87).

High power readout, which we choose to perform for all the experiments of this thesis, solves many of these issues. Instead of relying directly on measuring the frequency-shifted resonator, the high-power scheme exploits the state-dependent sensitivity of the onset of the “bright” state of the resonator, at the bare resonator frequency [209]. We can see on Fig. 3.9 (a) and (b) that, at very high power ($P_{\text{readout}} = 10$ dBm), the frequency shift caused by the qubit is completely eliminated, and the “bright” resonator

is visible at its bare frequency, ω_R (see also the “punchout” measurement in Sec. 3.4). At a slightly lower—but still high—power, the appearance of this bright state depends on the qubit state. Indeed, if we observe closely Fig. 3.9, we notice that the power at which the resonator appears is slightly lower when the qubit is prepared in $|e\rangle$. We thus configure readout by setting the power *just below* the power required to trigger the bright state in $|g\rangle$.

The power used for this kind of readout ensures that the SNR of the pulse measured by the ADC is high. This allows us to determine the qubit state in a “single-shot,” without averaging over many pulses. The second advantage is that the readout is not sensitive to qubit decoherence. If the resonator switches to the bright state, it will stay in that state for as long as power is maintained, we can thus increase SNR by averaging over a longer pulse length (if necessary). The main disadvantage is that the high-power readout is a destructive readout. After the pulse, both the qubit and the resonator will need time to relax back to the ground state.

It is possible to calibrate the measurement such that we assign the result of each shot to a binary result, i.e., the qubit state, instead of returning the raw IQ sample. This makes it possible to plot the results of experiments using the average population P_e instead of an arbitrary voltage scale. Figure 3.10 illustrates the calibration process. We first collect raw IQ samples for the qubit initialized in the $|g\rangle$ (pale blue) and $|e\rangle$ (pale orange) state. We may confirm that this is indeed the correct state correspondence by looking at the magnitude of the voltage: It is higher in the ground state since the resonator is not present at the measurement frequency (high transmission). Note, however, that there are outlier samples, i.e., samples for which the resonator was measured in the bright state even though the qubit was initialized in $|g\rangle$, and vice-versa. These errors stem either from incorrect state preparation or measurement (so-called SPAM errors).

These errors make it impossible to use a simple average to identify the IQ values corresponding to a particular state. Instead, we combine both datasets and run a k -means clustering analysis using the means of the $|g\rangle$ and $|e\rangle$ datasets as starting points. If successful, the clustering analysis returns the position of the centroid of each cluster, c_g and c_e . We draw a partition line at a point C along the axis defined by the centroids. Each sample is thus classified according to the side of the partition line it falls on. The measurement visibility is defined as

$$V = P_{e|e} + P_{g|g} - 1 \quad (3.8)$$

where $P_{e|e}$ is the proportion of measurement shots classified as $|e\rangle$ when the qubit was prepared in $|e\rangle$ and equivalently for $P_{g|g}$. The point C itself is determined by maximizing the visibility (or, if the two clusters are well separated, we may simply set it at the midpoint between c_g and c_e).




Figure 3.10: Qubit single shot measurement IQ calibration. Two single-shot measurements are made, one each for a qubit prepared in $|g\rangle$ (pale blue) and $|e\rangle$. Each measurement consists of 786 individual “shots” or points. The position of the two clusters c_g and c_e is determined via k -means clustering. After finding an optimal midpoint C , we calculate the transformation taking the IQ samples to the I axis. Subsequent measurements can then be classified as $|e\rangle$ if they are on the positive side of the axis, and $|g\rangle$ otherwise.

Numerically, it is convenient to perform the classification by calculating the *complex* plane transformation—a translation by $-C$ followed by a rotation by $\theta = \arg(c_e^* - C)$ —that places c_g and c_e on the real axis (corresponding to I). A raw complex IQ sample z and transformed as $z' = e^{i\theta}(z - C)$ is thus classified as an excited state if $\text{Re}[z'] > 0$, and as a ground state otherwise. The bright blue and orange points in Fig. 3.10 correspond to those transformed samples. For these two measurements, $P_{g|g} = 0.990$ and $P_{e|e} = 0.965$, resulting in a visibility $V = 0.955$.

For each experiment, we calibrate the readout by measuring preparations of $|g\rangle$ and $|e\rangle$ and calculating the transformation. Any subsequent measurement is then classified according to its real part after the transformation. The average population, for an array of complex IQ samples $\{z_i\}$ is calculated by taking the mean

$$P_e = \frac{\sum_{i=1}^N \text{Re}[e^{i\theta}(z_i - C)] > 0}{N} \quad (3.9)$$

where N is the number of shots and the comparison results in a Boolean $\{0, 1\}$.

Figure 3.9 (c) shows the calculated single-shot visibility as a function of measurement power and frequency. The best visibility is actually obtained at a frequency slightly higher than ω_r . The visibility drops at low power because the SNR is insufficient for a single-shot readout.

3.4 Initial Characterization Experiments

Finally! The fridge is cold, instruments are connected, and we are ready to begin measurements. The first step is to take a full high-power VNA frequency sweep. With the filters, circulators and amplifiers, the measurement frequency range is from 4 to 8 GHz. Such a sweep is visible in Fig. 4.16 (c). This measurement will reveal the resonance frequencies of the devices that are coupled to the transmission line. For us, that means resonators. At this point, you can compare the designed resonance frequencies to the measured ones; if all went well, they should match.

We can now make a more detailed measurement of each resonator in order to fit the transmission and extract Q_i and Q_c [see Fig. 4.16 (d)]. This can be done as we vary the power, to see how much the resonators are affected by TLS [see Fig. 4.17]. If we are measuring a resonator that is coupled to a qubit (e.g., a readout resonator), a power sweep will also immediately indicate whether the qubit is working or not. This is because the resonator gradually transitions into the quantum regime as power is lowered [210] and during this transition—if there is a qubit—the resonator often disappears entirely!

If this effect is observed during resonator measurements, a “punchout” experiment should be executed next [211]. For this measurement, we sweep a larger bandwidth around the resonator and over a range of powers. The goal is to observe the frequency of the resonator change as a function of power. An example of a punchout measurement can be seen in Fig. 3.11, where we plot the amplitude and the phase of S_{21} as a function of VNA power P_{readout} and frequency f_{readout} . At high power, there is so much energy going into the resonator that the system becomes classical—the qubit is “punched out”—and we see the resonator at its “bare” frequency. As the power is lowered below 3 dBm, the resonator briefly fades out and subsequently reappears as its frequency shifts. At low power, the system is well-described by the Jaynes-Cummings Hamiltonian [Eq. (2.84)]. The difference between the high-power and low-power frequency is given by χ in Eq. 2.86, the AC Stark shift [102]. If the shift is positive (as in Fig. 3.11), $\omega_R > \omega_Q$. A punchout measurement therefore gives a very simple way to calculate the qubit frequency given an estimate of g .

If it appears that we have working qubits, the next step (before moving to time-domain measurement) is to perform CW spectroscopy, where we tune the idle qubit frequency via the DC voltage bias and look at the frequency of the resonator change,




Figure 3.11: Punchout measurement. We measure S_{21} with a VNA over a range of frequencies around the resonator, varying the emitted power. At low power, we observe the resonator at a frequency shifted by χ . At high power, the qubit is “punched out,” and the resonator reverts to its “bare” frequency.

again, because of the Stark shift. This measurement must be made at low power: the resonator should be well in the quantum regime. We can use the punchout results to determine the appropriate power. Figure 3.7 shows two examples of “single-tone” spectroscopy. We are able to observe anti-crossings between the qubit and resonator, and, though we cannot observe the qubit directly, we can imagine (and calculate) where it is. If we wanted, we could also perform “two-tone” spectroscopy to pinpoint the qubit frequency directly. If we have a good estimate for g , though, it is not necessary.

We are done with VNA measurements, time for... time-domain! We disconnect the resonators, and instead connect the homodyne readout setup to the measurement input and outputs of the DR. We also connect the pulsed drive equipment to the qubit XY line. The setup now looks exactly as the one of Fig. 3.5. The first thing we will need to do is to roughly calibrate the measurement power. Because we are using a high-power readout scheme, this is relatively simple: we simply start at very high power, where we see the resonator at its bare frequency, and lower the power until we see the resonator disappear (in effect, we replicate the punchout measurement). At this stage, the measurement does not need to be perfect. We then drive the qubit at its estimated frequency with a pulse, measuring right after. We repeat this pulse sequence many times to average the measured voltage, sweeping either the drive pulse length or amplitude. Assuming that readout was set up properly, we will see the average voltage *oscillate*; this is a Rabi measurement. A Rabi measurement where we swept the pulse length t_{pulse} is pictured in Fig. 3.12 (a). We can easily determine the length of a π -pulse: it is half the oscillation period; here, approximately $0.3 \mu\text{s}$. Note that Fig. 3.12 (a) was




Figure 3.12: Rabi and Ramsey experiments. Left: Rabi experiment. The amplitude of the measured voltage oscillates as the length of the pulse applied to the qubit increases. The period of a full rotation is approximately $0.6 \mu\text{s}$. The data pictured here is from the author’s first Rabi experiment, which was done with Carolyn Earnest on Oct. 19th 2018. Right: Ramsey experiment. The oscillation is caused by the difference between the drive and qubit frequencies. Solid lines are fits.

the first Rabi experiment ever done in our lab, it used a very basic rectangular pulse. The Rabi calibrations done for the experiments in this thesis use fixed length Gaussian or cosine pulses and we instead vary the amplitude. Also note that once a basic π -pulse is working, we can change the measurement scale from the voltage amplitude measured by the digitizer to the average excited population P_e , as explained in the last section.

While we can measure the qubit frequency via spectroscopy, the best way to accurately pinpoint ω_Q is a Ramsey experiment. The pulse sequence of a Ramsey experiment consists of a first $\pi/2$ -pulse⁵, followed by a delay time Δt , and then a second $\pi/2$ -pulse. Since we do not know the exact qubit frequency, the pulses are made at a frequency ω_{pulse} that we think is relatively close to ω_Q . If we are not too far off, the off-resonance pulses will successfully rotate the qubit. The frequency difference, however, causes the superposition state $|g\rangle + |e\rangle$ obtained after the first pulse to *precess*, with the phase of $|e\rangle$ increasing at a rate $f_R = (\omega_{\text{pulse}} - \omega_Q)/2\pi$ during the delay Δt . This increasing phase is what causes the oscillations visible in Fig. 3.12 (b). If we fit these oscillations, we can obtain f_R , thereby allowing us to determine ω_Q . In

⁵ $\pi/2$ -pulses are simply half-amplitude π -pulses

practice, we perform multiple Ramsey runs at different pulse frequencies in order to get a very precise fix on ω_Q . In Fig. 3.12 (b), we fit four Ramsey frequencies obtained at $\omega_{\text{pulse}}/2\pi \in \{4.7237, 4.7277, 4.7597, 4.7637\}$ GHz. We can thus calculate the qubit frequency $\omega_Q = 4.74383(1)$ GHz.

With an accurate qubit frequency, we can repeat the Rabi experiment to better calibrate the π -pulse amplitude. We can then improve the measurement visibility by fine-tuning the amplitude and frequency of the readout pulse (see Fig. 3.9). This simple calibration is generally sufficient for the experiments described in this thesis. To properly optimize gate fidelity, however, more sensitive experiments are needed. For example, instead of calibrating the π -pulse amplitude by sending a single pulse, we can send multiple. The error is thus amplified and can be detected more easily. For examples of these more sophisticated pulse calibration schemes, see, e.g., Refs. [212, 213, 214].

We are now in a good shape: We know the qubit frequency accurately, and we calibrated a high-quality π -pulse and readout. This allows us to properly characterize the decoherence of the qubit with T_1 and T_2 experiments. The T_1 experiment is simple: we pulse the qubit to the excited state, wait a delay time Δt , and measure. Because the qubit is coupled to the environment (and in particular TLS, see Chapters 5 and 6), its excitation leaks into the environment at a rate $\Gamma_1 = 1/T_1$. As was explained in Sec. 2.5, as Δt is increased the average population decreases exponentially. Figure 3.13 (a) shows a T_1 experiment; we fit the decay, giving $T_1 = 19.75(30)$ μs .

From the point of view of gate-based quantum computing, T_2 is usually the more critical number because it encompasses both the relaxation rate Γ_1 and the pure dephasing rate Γ_ϕ and, for tunable qubits, is often significantly smaller than T_1 . A T_2 experiment is simply an on-resonance Ramsey experiment. As can be seen in Fig. 3.13 (b), because $f_R = 0$, there are no oscillations, only a decay. Note that due to $1/f$ noise [215], T_2 decays are generally not exponential ($\propto e^{-\Delta t/T_2}$), but closer to Gaussian ($\propto e^{-(\Delta t/T_2)^k}$, with $1 \lesssim k \lesssim 2$). Therefore, we fit the decay with a model containing an extra exponent parameter k . The T_2 time obtained with this method can be considered to be a worst-case number because a pure exponential with the same time constant would actually decay faster⁶.

It is possible to make a T_2 experiment where we filter out the sensitivity to low-frequency noise [84]. By inserting a π -pulse exactly in the middle of the delay time, noise that causes a qubit frequency change—and therefore a precession—that is *constant* during the length of the pulse sequence can be eliminated. This is because the direction of the precession is inverted after the π -pulse, thereby bringing the qubit back to its original state when the measurement is made. Higher-frequency noise

⁶A pure exponential would decay faster in the time period $0 < \Delta t < T_2$. After reaching $\Delta t > T_2$, a Gaussian decays faster. We consider the initial period to be more important since that is the time during which operations will be performed reliably.

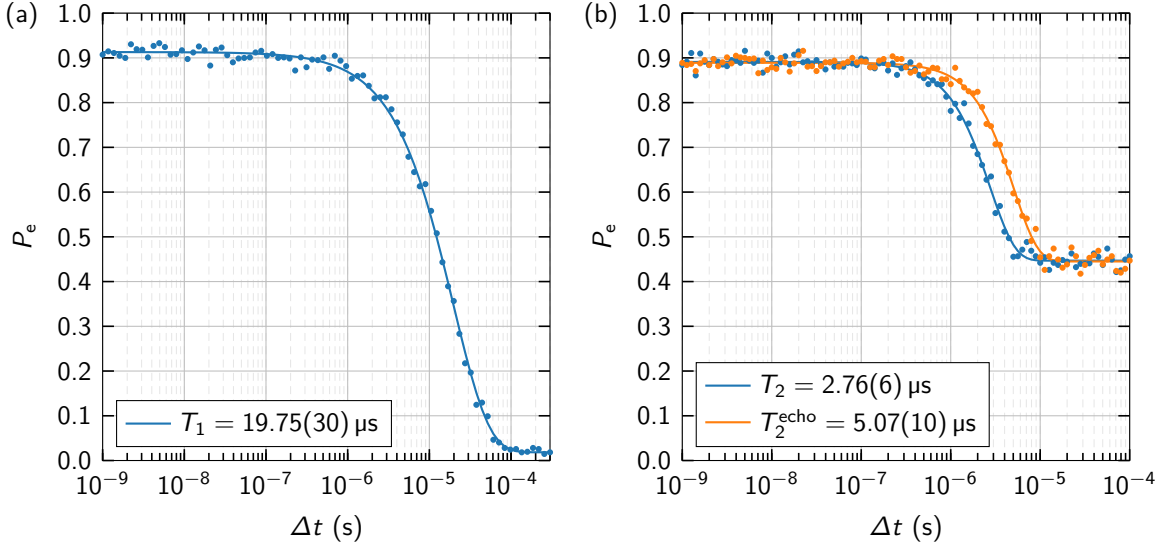


Figure 3.13: T_1 and T_2 experiments. Solid lines are fits. Left: T_1 experiment. Right: T_2 experiments.

causes the precession to increase at a randomly fluctuating rate, rendering the π -pulse ineffective. This type of T_2 experiment is called an “echo” experiment because, in old nuclear magnetic resonance setups, the middle π -pulse caused a resurgence of the measurement signal [216]. If multiple π -pulses are inserted, the shape of the noise filtering function can be tuned, making possible to probe noise in different regions of the spectrum [217, 181]. A T_2^{echo} decay is shown in Fig. 3.13; as would be expected, reducing low-frequency noise increases the time constant. Because the source of low-frequency noise is often $1/f$ in nature, eliminating it often renders the decay more exponential.

To finish the initial characterization of the device, we perform spectroscopy of the qubit and resonator over a large range of V_{bias} , including spectroscopy of the ω_{02} transition, which is accessible at high drive power via a two-photon process at $\omega_{02}/2$. If desired, we can also measure T_1 and T_2 during that same experiment. The relaxation rate tends to vary as a function of $\Delta = \omega_R - \omega_Q$ due to the Purcell effect [2, 177, 178]. At smaller detunings, the qubit relaxation rate is enhanced by the resonator relaxation rate. The dephasing rate directly depends on the flux biasing of the qubit. At zero net flux, the qubit is at the highest frequency, and the frequency derivative is zero $\frac{d\omega_Q}{d\varphi} = 0$. As a result, to first order, flux noise in the SQUID loop will not change the frequency of the qubit. This is the “sweet spot”. Conversely, as flux bias is increased and the qubit frequency decreases, the derivative becomes larger and larger, increasing sensitivity to




Figure 3.14: Full Hamiltonian calibration spectroscopy. To obtain this data, we first perform low-power resonator spectroscopy with the VNA. We then switch to the time-domain measurement setup and perform qubit spectroscopy of both $\omega_Q = \omega_{01}$ and ω_{02} . The second excited state is accessible at high drive power via a two-photon process at $\omega_{02}/2$. We then simultaneously fit all three frequencies as a function of V_{bias} using a full circuit model. The extracted parameters are shown in 3.1. Though this is not shown, we also measure T_1 and T_2 during the experiment.

flux noise, and reducing T_2 .

Figure 3.14 shows the result of such spectroscopy, along with lines fit according to a physical model derived with the method of Chapter 2. Such a fit allows us to determine all the relevant circuit parameters; they are tabulated in Table 3.1. We can then use those circuit parameters to check whether design was done correctly. One particularly useful number is E_J , which, in combination with the room-temperature resistance of the SQUID loop and Eq. (2.49), allows us to calculate the value of the superconducting gap Δ for the fabricated film. This value may then be used for future designs. With this sample we measured $R_n = 6.3(1)$ k Ω , corresponding to a gap $\Delta \sim 160$ μeV .

Another useful experiment for characterization and calibration of a qubit over a large range of frequencies is *swap spectroscopy*, and its optimized extension, *octave sampling*. This experiment uses flux pulses to instantaneously set the qubit frequency rather than setting a DC flux with V_{bias} . Swap spectroscopy and octave sampling are covered in Chapter 5.

Table 3.1: Qubit–resonator device parameters calculated from the fit of Fig. 3.14. This device was used in some experiments of this thesis, in particular those of Chapter 5. C_c is the coupling capacitance between the qubit and its readout resonator. Note that the resonator frequency is that of the “bare” resonator. If we include the effect of the coupling capacitance, the resonator frequency is 5.032 GHz. The Hamiltonian parameters are determined via a fit of the resonator and qubit frequency, and the qubit anharmonicity. The relative error on each parameter is less than 1%.

Parameter	Value
E_J (GHz)	19.614(5)
E_C (MHz)	188.92(5)
f_R (GHz)	5.04844(3)
C_c (fF)	3.371(5)
T_1 (μ s)	10 to 25
T_2 (ns)	400 to 1000

3.5 Randomized Benchmarking

We have shown how to characterize a qubit–resonator system. More advanced experiments will now involve more parameter dimensions (e.g., time), multiple circuit elements, or longer pulse sequences. The first two points are covered by the experiments done in Chapters 5 and 6, respectively. A great example of an experiment involving long pulse sequences is randomized benchmarking (RB), the topic of this section.

As mentioned in Sec. 2.4, randomized benchmarking is used to characterize the *fidelity* of quantum gates. The classic RB protocol [152] uses progressively longer sequences of random gates G_1, G_2, \dots, G_m designed to modify the state of a qubit (or of multiple qubits) such that, on average, the qubit covers the entire state space uniformly during the sequence. A final gate G_{m+1} applied at the end of each sequence returns the qubit to the ground state, after which the qubit is measured. With a perfect qubit and perfect gates, that would result in a constant measurement of $|g\rangle$, no matter the sequence length. For a physical qubit and imperfect gates, any kind of noise or errors—energy relaxation, dephasing, or even unitary gate errors—transforms the sequence to a simple depolarizing channel. As a result, the qubit is not always returned to the ground state, and on average P_g decays exponentially as a function of the sequence length m . The decay is fit to obtain the average gate fidelity \bar{F}_g .

The gates used in the sequence are drawn randomly from a set called the *Cliffords*. This gate set has two advantages: First, it is a discrete set, though its size grows extremely quickly with the number of qubits, and, second, the final inversion gate G_{m+1}

Table 3.2: Single-qubit Cliffords. The single-qubit Clifford group contains 24 elements. The notation $R_{\vec{u}}(\theta)$ represents a Clifford element as a Bloch sphere rotation by an angle θ about the axis represented by the Euclidean vector $\vec{u} = (x, y, z)$. The right column gives an implementation of the Clifford in terms of basic X and Z pulses. Note that a single Clifford may have multiple possible implementations. The notation $X/2$ represents a $\pi/2$ rotation around the x -axis.

Clifford Element	Implementation	Clifford Element	Implementation
Identity	I	$R_{(1,0,0)}(\pi/2)$	$X/2$
$R_{(1,0,0)}(\pi)$	X	$R_{(-1,0,0)}(\pi/2)$	$-X/2$
$R_{(0,1,0)}(\pi)$	X, Z	$R_{(0,1,0)}(\pi/2)$	$X/2, Z/2, -X/2$
$R_{(0,0,1)}(\pi)$	Z	$R_{(0,-1,0)}(\pi/2)$	$X/2, -Z/2, -X/2$
$R_{(1,1,1)}(2\pi/3)$	$X/2, Z/2$	$R_{(0,0,1)}(\pi/2)$	$Z/2$
$R_{(1,1,-1)}(2\pi/3)$	$-Z/2, X/2$	$R_{(0,0,-1)}(\pi/2)$	$-Z/2$
$R_{(1,-1,1)}(2\pi/3)$	$Z/2, X/2$	$R_{(1,0,1)}(\pi)$	$X/2, Z/2, X/2$
$R_{(1,-1,-1)}(2\pi/3)$	$X/2, -Z/2$	$R_{(1,0,-1)}(\pi)$	$X/2, -Z/2, X/2$
$R_{(-1,1,1)}(2\pi/3)$	$Z/2, -X/2$	$R_{(0,1,-1)}(\pi)$	$Z, X/2$
$R_{(-1,1,-1)}(2\pi/3)$	$-X/2, -Z/2$	$R_{(0,1,+1)}(\pi)$	$Z, -X/2$
$R_{(-1,-1,1)}(2\pi/3)$	$-X/2, Z/2$	$R_{(1,1,0)}(\pi)$	$X, Z/2$
$R_{(-1,-1,-1)}(2\pi/3)$	$-Z/2, -X/2$	$R_{(-1,1,0)}(\pi)$	$X, -Z/2$

can be computed efficiently. A consequence of the first point is that the qubit will not cover the entire Bloch sphere; it is confined to a finite set of states. Fortunately, this still results in exponential decay [154]. The second point is particularly important because it means that the Cliffords do not constitute a *universal* gate set. In other words, quantum computation is not realizable by using gates solely from the Clifford set. Table 3.2 lists the single-qubit Clifford gates. The two-qubit Clifford set can be built out of tensor products of single-qubit Cliffords, plus a two-qubit gate, e.g., the CNOT. Classically efficient methods exist to randomly generate Cliffords for any number of qubits, making randomized benchmarking *scalable*. We have implemented the algorithm described in Ref. [218] to generate Cliffords.

We run a randomized benchmarking experiment by executing pulse sequences that realize random Clifford sequences. Because the Clifford set contains rotations about “diagonal” axes, e.g., the Hadamard gate, they cannot all be realized with a single microwave pulse (for XY rotations), or flux pulse (for Z rotations). Each Clifford must instead be decomposed into a set of achievable pulses. Since generating Z rotations with flux pulses generally yields lower fidelity gates, it is advantageous to use “all-microwave” gates instead. This is possible because X and Y rotations can be combined to create a Z rotation. The disadvantage is that each Clifford now requires more pulses.




Figure 3.15: Single-qubit randomized benchmarking (a) and purity benchmarking (b). For each sequence length m , 50 random sequences are generated. We choose 28 lengths logarithmically spaced between 1 and 5000 for RB, and 1 and 3000 for PB. Rotations about the z-axis are done virtually by tracking the rotation of the qubit reference frame. Each microwave pulse is 20 ns long and has a DRAG-optimized cosine shape. We calibrate the pulse amplitude, DRAG parameter, and qubit frequency with the “AllXY” protocol from Ref. [212] prior to running the Clifford sequences.

Another approach is to replace physical Z rotations with *virtual* rotations, where we instead track rotations of the qubit reference frame, and adjust the rotation axis of subsequent gates [145]. This method is not only applicable to randomized benchmarking, but to any algorithm, and is compatible with two-qubit gates.

Figure 3.15 (a) shows the result of a randomized benchmarking experiment with each Clifford element decomposed in terms of microwave X rotations and virtual Z rotations. On average, this decomposition results in $23/24 \approx 0.9583$ physical pulses per Clifford. We generate 50 random sequences for each sequence length m and record the ground population P_g after each sequence. We then fit the average P_g over each of the 50 sequences to an exponential model

$$P_g(m) = Ap^m + B \quad (3.10)$$

where A , B and p are fitting parameters. The average Clifford gate fidelity is then

$$\bar{F}_g = \frac{1}{d} + p \left(1 - \frac{1}{d} \right) \quad (3.11)$$




Figure 3.16: Randomized benchmarking over time. We measure \bar{F}_g over time by continuously running sequences of increasing length m . We choose 19 lengths logarithmically spaced between 1 and 3000. A run of a single sequence of each length takes 30 s. Each point in the plot is the result of a fit performed over a 25 min-window moving average, with each P_g value in the exponential decay therefore averaged over 50 sequences. We measure a total of 4800 sequences of each length for a total experimental time of 40 h. The pulse parameters are calibrated in the same way as the experiment of Fig. 3.15. The \bar{F}_g fluctuations are larger than the fit standard error, which is 0.00002 on average.

where d is the space dimension. For a single qubit, $d = 2$.

Figure 3.15 (b) shows the result of a “purity” benchmarking (PB) experiment [158, 219]. Purity benchmarking is very similar to RB. The same random Clifford sequences are run, but instead of measuring only P_g —equivalent to a measurement of $\langle \hat{\sigma}_z \rangle$ —at the end of each sequence, we also measure $\langle \hat{\sigma}_x \rangle$ and $\langle \hat{\sigma}_y \rangle$. This means that each individual random sequence must be repeated three times, and a rotation pulse must be added to measure $\langle \hat{\sigma}_x \rangle$ and $\langle \hat{\sigma}_y \rangle$. These three expectation values are then used to obtain an estimate of the purity $P = \langle \sigma_x \rangle^2 + \langle \sigma_y \rangle^2 + \langle \sigma_z \rangle^2$. The purpose of a purity benchmarking experiment is to evaluate the *incoherent* error per gate ϵ_{in} . As might be expected, incoherent errors are caused by decoherence. The incoherent error is useful because, in combination with the total error $\epsilon = 1 - \bar{F}_g$, it allows us to estimate the amount of *coherent* error, $\epsilon_{\text{coh}} = \epsilon - \epsilon_{\text{in}}$. In principle, the coherent error may be reduced via better calibration or pulse quality, and it is therefore a metric of how well the system has been optimized.

It is now well-known that taking individual measurements of T_1 , T_2 , or \bar{F}_g is not sufficient to properly characterize a quantum device. Two-level systems cause fluctua-

tions of those parameters over time, sometimes decreasing T_1 by an order of magnitude. A proper characterization therefore requires measurement over long time periods, sometime even multiple days. Chapter 6 details the origin of those fluctuations, and presents measurement and simulations of T_1 over long time periods and over a range of qubit frequencies. Similar experiments may be run for other metrics, such as \bar{F}_g .

Figure 3.16 shows the result of a “monitoring” RB experiment, where we continually run RB sequences in order to monitor \bar{F}_g over time. The protocol for this experiment is slightly different from a standard RB experiment. Instead of directly generating 50 random sequences for each length, we run a *single* sequence for each length m , and start over after running the last sequence. We insert a short variable delay time after the last sequence to ensure that a full run takes exactly 30 s. The sequences are repeated for as long as is desired. We calculate the average P_g value over 50 sequences for every time point with a centered moving window. At the endpoints, the window is truncated, e.g., for the first and last point the average is only performed over 25 sequences. The average P_g decay is fit to obtain an estimate of \bar{F}_g at each time point.

The Quantum Socket

Quantum computing architectures are on the verge of scalability, a key requirement for the implementation of a universal quantum computer. The next stage in this quest is the realization of quantum error correction codes, which will mitigate the impact of faulty quantum information on a quantum computer. Architectures with ten or more quantum bits (qubits) have been realized using trapped ions and superconducting circuits. While these implementations are potentially scalable, true scalability will require systems engineering to combine quantum and classical hardware. One technology demanding imminent efforts is the realization of a suitable wiring method for the control and measurement of a large number of qubits. In this chapter¹, we introduce an interconnect solution for solid-state qubits: *The quantum socket*. The quantum socket fully exploits the third dimension to connect classical electronics to qubits with higher density and better performance than two-dimensional methods based on wire bonding. The quantum socket is based on spring-mounted microwires – *the three-dimensional wires* – that push directly on a micro-fabricated chip, making electrical contact. A small wire cross-section ($\sim 1 \mu\text{m}$), nearly non-magnetic components, and functionality at low temperatures make the quantum socket ideal to operate solid-state qubits. The wires have a coaxial geometry and operate over a frequency range from DC to 8 GHz, with a contact resistance of $\sim 150 \text{ m}\Omega$, an impedance mismatch of $\sim 10 \Omega$, and minimal crosstalk. As a proof of principle, we fabricated and used a quantum socket to measure high-quality superconducting resonators at a temperature of $\sim 10 \text{ mK}$. Quantum error correction codes such as the surface code will largely benefit from the quantum socket, which will make it possible to address qubits located on a two-dimensional lattice. The present implementation of the socket can be readily extended to accommodate a quantum processor with a 10×10 qubit lattice, which would allow the realization of a

¹This chapter was largely adapted from Ref. [220]. The list of author contributions can be found in the Statement of Contributions within the front matter of this thesis. © 2016 American Physical Society

simple quantum memory.

4.1 Introduction

At present, one of the main objectives in the quantum computing community is to build and prototype practical hardware technology for scalable architectures that may lead to the realization of a universal quantum computer [36]. In this chapter, we undertake the task of implementing an extensible wiring method for the operation of a quantum processor based on solid-state devices, e.g., superconducting qubits [90, 91, 221]. Possible experimental solutions based on wafer bonding techniques [222, 223, 224, 225, 226] or coaxial through-silicon vias [227] as well as theoretical proposals [228, 229] have recently addressed the wiring issue, highlighting it as a priority for quantum computing.

Despite the recent experimental accomplishments mentioned in Chapter 1, a truly scalable qubit architecture has yet to be demonstrated. Wiring is one of the most basic unsolved scalability issues common to most solid-state qubit implementations, where qubit arrays are fabricated on a chip. The conventional wiring method based on wire bonding suffers from fundamental scaling limitations as well as mechanical and electrical restrictions. Wire bonding relies on bonding pads located at the edges of the chip. Given a two-dimensional lattice of $N \times N$ physical qubits on a square chip, the number of wire bonds that can be placed scales approximately as $4N$ (N bonds for each chip side). Wire bonding will thus never be able to reach the required N^2 law according to which physical qubits scale on a two-dimensional lattice. Furthermore, for large N , wire bonding precludes the possibility of accessing physical qubits in the center region of the chip, which is unacceptable for a physical implementation of the surface code. In the case of superconducting qubits, for example, qubit control and measurement are typically realized by means of microwave pulses or, in general, pulses requiring large frequency bandwidths. By their nature, these pulses cannot be reliably transmitted through long sections of quasi-filiform wire bonds. In fact, stray capacitances and inductances associated with wire bonds as well as the self-inductance of the bond itself limit the available frequency bandwidth, thus compromising the integrity of the control and measurement signals [230].

In this chapter, we set out to solve the wiring bottleneck common to almost all solid-state qubit implementations. Our solution is based on suitably packaged *three-dimensional microwires* that can reach any area on a given chip from above. We define this wiring system as the *quantum socket*. The wires are coaxial structures consisting of a spring-loaded inner and outer conductor with diameters of 380 μm and 1290 μm , respectively, at the smallest point and with a maximum outer diameter of 2.5 mm. The movable section of the wire is characterized by a maximum stroke of approxi-

mately 2.5 mm, allowing for a wide range of on-chip mechanical compression. All wire components are nearly non-magnetic, thereby minimizing any interference with the qubits. The three-dimensional wires work both at room temperature and at cryogenic temperatures as low as ~ 10 mK. The wires' test-retest reliability (repeatability) is excellent, with marginal variability over hundreds of measurements. Their electrical performance is good from DC to at least 8 GHz, with a contact resistance smaller than $150\text{ m}\Omega$ and an instantaneous impedance mismatch of approximately $10\ \Omega$. Notably, the coaxial design of the wires strongly reduces unwanted crosstalk, which we measured to be at most -45 dB for a realistic quantum computing application.

This chapter is organized as follows. In Sec. 4.2, we introduce the quantum socket design and microwave simulations. In Sec. 4.3, we show the socket physical implementation with emphasis on materials and alignment procedures. In Sec. 4.4, we present a comprehensive DC and microwave characterization of the quantum socket operation at room and cryogenic temperatures. In Sec. 4.5, we show an application of the quantum socket relevant to superconducting quantum computing, where the socket is used to measure aluminum (Al) superconducting resonators at a temperature of approximately 10 mK. Finally, in Sec. 4.6, we envision an extensible quantum computing architecture where a quantum socket is used to connect to a 10×10 lattice of superconducting qubits.

4.2 Design

The development of the quantum socket required a stage of meticulous micromechanical and microwave design and simulations. Based on signal integrity and quantum device quality criteria, it was determined that a spring-loaded interconnect – the three-dimensional wire – was the optimal method to electrically access devices lithographically fabricated on a chip² and operated in a cryogenic environment. An on-chip contact pad geometrically and electrically matched to the bottom interface of the wire can be placed easily at any desired location on the chip as part of the fabrication process, thus making it possible to reach any point on a two-dimensional lattice of qubits. The coaxial design of the wire provides a wide operating frequency bandwidth, while the springs allow for mechanical stress relief during the cooling process. The three-dimensional wires used in this chapter take advantage of the knowledge in the existing field of microwave circuit testing [231]. However, reducing the wire dimensions to a few hundred micrometers and using it to connect to quantum-mechanical micro-fabricated circuits at low temperatures resulted in a significant extension of existing implementations and applications.

²A typical chip comprises a dielectric substrate (e.g., silicon or sapphire) and a metallic surface.

4.2.1 Three-dimensional wires

Figure 4.1 shows the design of the quantum socket components. Figure 4.1 (a) displays a model of a three-dimensional wire. The coaxial design of the wire is visible from the image, which features a wire 30.5 mm long when uncompressed. The wire is characterized by an inner cylindrical pin of diameter 380 μm and an outer cylindrical body (the electrical ground) of diameter 1290 μm at its narrowest region; this region is the bottommost section of the wire and, hereafter, will be referred to as the wire *contact head* (cf. the inset of Fig. 4.1 (a), as well as the dashed box on the left of Fig. 4.2 (a)). The contact head terminates at the wire *bottom interface*; this interface is designed to mate with a pad on a chip (cf. Fig. 4.2 (b) and (c)). The outer body includes a rectangular aperture, the *tunnel*, to prevent shorting the inner conductor of an on-chip coplanar waveguide (CPW) transmission line [232, 115]; the transmission line can connect the pad with any other structure on the chip. Two different tunnel dimensions were designed, with the largest one reducing potential alignment errors³. The tunnel height was 300 μm in both cases, with a width of 500 μm or 650 μm . The internal spring mechanisms of the wire allow the contact head to be compressed; the maximum stroke was designed to be 2.5 mm, corresponding to a working stroke of 2.0 mm.

The outer body of the three-dimensional wire is an M2.5 male thread used to fix the wire to the lid of the microwave package (cf. Fig. 4.1 (b) and (d)). The thread is split into two segments of length 3.75 mm and 11.75 mm that are separated by a constriction with outer diameter 1.90 mm. The constriction is necessary to assemble and maintain in place the inner components of the three-dimensional wire. A laser-printed marker is engraved into the top of the outer body. The marker is aligned with the center of the tunnel, making it possible to mate the wire bottom interface with a pad on the underlying chip with a high degree of angular precision.

Figure 4.2 (a) shows a lateral two-dimensional cut view of the three-dimensional wire. Two of the main wire components are the inner and outer barrel, which compose part of the inner and outer conductor. The inner conductor barrel is a hollow cylinder with outer and inner diameters of 380 μm and 290 μm (indicated as part iv in Fig. 4.2 (a)), respectively. This barrel encapsulates the inner conductor spring. The outer conductor barrel is a hollow cylinder as well, in this case with an inner diameter of 870 μm (parts ii and vii). Three polytetrafluoroethylene (PTFE) disks serve as spacers between the inner and outer conductor; such disks contribute marginally to the wire dielectric volume, the majority of which is air or vacuum. The outer spring is housed within the outer barrel towards its back end, just before the last PTFE disk on the right-hand side of the wire. The *back end* of the wire is a region comprising a female thread on the outer conductor and an inner conductor barrel (cf. dashed box on the right-hand side of Fig. 4.2 (a)).

³These errors can result in undesired short-circuit connections to ground.




Figure 4.1: Computer-aided designs of the three-dimensional wire, microwave package, and package holder. (a) A wire of length $\ell = 30.5$ mm along with a detail of the contact head (inset). (b) Assembled microwave package including six three-dimensional wires, washer, washer springs, and chip (shown in green). The arrow indicates the screw-in micro connector mated to the back end of the wire. Forward hatching indicates the washer cutaway, whereas backward hatching indicates both lid and sample holder cutaways. (c) cross-section of the microwave package showing the height of the upper cavity, which coincides with the minimum compression distance ℓ_c of the three-dimensional wires (cf. Appendix B.1). (d) Microwave package mounted to the package holder, connected, in turn, to the mounting plate of a DR with SMP connectors. A channel with a cross-sectional area of $800 \mu\text{m} \times 800 \mu\text{m}$ connects the inner cavities of the package to the outside, thus making it possible to evacuate the inner compartments of the package. This channel meanders to prevent external electromagnetic radiation from interfering with the sample.




Figure 4.2: Two-dimensional cut view of the three-dimensional wire, contact pad, and screw-in micro connector. (a) Side view of the wire cross-section. The wire components are: i, spring-loaded center conductor of the contact head; ii, spring-loaded outer conductor of the contact head; iii, vi, and ix, dielectric spacers; iv, center conductor barrel; v, center conductor spring; vii, outer conductor barrel; viii, outer conductor spring; x, center conductor tail; xi, outer conductor tail; xii, threaded outer body. The dashed box on the left indicates the contact head, whereas that on the right indicates the female threads included for use with the screw-in micro connector. (b) Front view of the wire. The blue surface indicates the wire bottom interface. (c) On-chip contact pad. Here, the blue surface indicates the pad dielectric gap, whereas the white surfaces refer to conductors (thin metallic films deposited on a dielectric substrate). (d) Screw-in micro connector. The left end of the micro connector mates with the back end of the three-dimensional wire; the right end is soldered to a coaxial cable, the inner conductor of which serves as the inner conductor of the micro connector (slotting into x).

The inner conductor tip is characterized by a conical geometry with an opening angle of 30° . Such a sharp design was chosen to ensure that the tip would pierce through any possible oxide layer forming on the contact pad metallic surface, thus allowing for a good electrical contact.

Figure 4.2 (c) shows the design of a typical on-chip pad used to make contact with the bottom interface of a three-dimensional wire. The pad comprises an inner and outer conductor, with the outer conductor being grounded. The pad in the figure was designed for a silver (Ag) film of thickness $3\ \mu\text{m}$. A variety of similar pads were designed for gold (Au) and Al films with thickness ranging between approximately 100 nm and

200 nm. The pad inner conductor is a circle with diameter 320 μm that narrows to a linear trace (i.e., the inner conductor of a CPW transmission line) by means of a raised-cosine taper. The raised cosine makes it possible to maximize the pad area, while minimizing impedance mismatch. As designed, the wire and pad allow for lateral and rotational misalignment of $\mp 140 \mu\text{m}$ and $\mp 28^\circ$, respectively. The substrate underneath the pad is assumed to be silicon (Si) with a relative electric permittivity $\epsilon_r \simeq 11$. The dielectric gap between the inner and outer conductor is 180 μm in the circular region of the pad; the outer edge of the dielectric gap then follows a similar raised-cosine taper as the inner conductor. The pad characteristic impedance is designed to be $Z_0 = 50 \Omega$.

4.2.2 Microwave package

The microwave package comprises three main parts: The lid; the sample holder; the grounding washer. The package is a parallelepiped with a height of 30 mm and with a square base of side length 50 mm. The chip is housed inside the sample holder. All these components mate as shown in Fig. 4.1 (b) and (c).

In order to connect a three-dimensional wire to a device on a chip, the wire is screwed into an M2.5 female thread that is tapped into the lid of the microwave package, as depicted in Fig. 4.1 (b). The pressure applied by the wire to the chip is set by the depth of the wire in the package. The wire stroke, package dimensions, thread pitch, and alignment constraints impose discrete pressure settings (cf. Appendix B.1). In the present implementation of the quantum socket, the lid is designed to hold a set of six three-dimensional wires, which are arranged in two parallel rows. In each row, the wires are spaced by 5.75 mm from center to center, with the two rows being separated by a distance of 11.5 mm.

A square chip of lateral dimensions 15 mm \times 15 mm is mounted in the sample holder in a similar fashion as in Ref. [233]. The outer edges of the chip rest on four protruding lips, which are 1 mm wide. Hereafter, those lips will be referred to as the *chip recess*. For design purposes, a chip thickness of 550 μm is assumed. Correspondingly, the chip recess is designed so that the top of the chip protrudes by 100 μm with respect to the adjacent surface of the chip holder, i.e., the depth of the recess is 450 μm (cf. Fig. 4.1 (c)). The outer edges of the chip are pushed on by a spring-loaded grounding washer. The 100 μm chip protrusion ensures a good electrical connection between chip and washer, as shown in Fig. 4.1 (c).

The grounding washer was designed to substitute the large number of lateral bonding wires that would otherwise be required to provide a good ground to the chip (as shown, for example, in Fig. 6 of Ref. [233]). The washer springs are visible in Fig. 4.1 (b), which also shows a cut view of the washer. The washer itself is electrically grounded by means of the springs as well as through galvanic connection to the surface of the lid. The four feet of the washer, which can be seen in the cut view of Fig. 4.1 (b), can be

designed to be shorter or longer. This makes it possible to choose different pressure settings for the washer.

After assembling the package, there exist two electrical cavities (cf. Fig. 4.1 (c)): One above the chip, formed by the lid, washer, and metallic surface of the chip (*upper cavity*), and one below the chip, formed by the sample holder and metallic surface of the chip (*lower cavity*). The hollow cavity above the chip surface has dimensions $14\text{ mm} \times 14\text{ mm} \times 3.05\text{ mm}$. The dimensions of the cavity below the chip surface are $13\text{ mm} \times 13\text{ mm} \times 2\text{ mm}$. The lower cavity helps mitigate any parasitic capacitance between the chip and the box (ground). Additionally, it serves to lower the effective electric permittivity in the region below the chip surface, increasing the frequency of the substrate modes (cf. Sec. 4.2.4).

A pillar of square cross-section with side length of 1 mm is placed right below the chip at its center; the pillar touches the bottom of the chip, thus providing mechanical support⁴. The impact of such a pillar on the microwave performance of the package will be described in Sec. 4.2.4.

4.2.3 Package holder

The three-dimensional wires, which are screwed into the microwave package, must be connected to the qubit control and measurement electronics. In addition, for cryogenic applications, the package must be thermally anchored to a refrigeration system in order to be cooled to the desired temperature. Figure 4.1 (d) shows the mechanical module we designed to perform both electrical and thermal connections. In this design, each three-dimensional wire is connected to a *screw-in micro connector*, which is indicated by an arrow in Fig. 4.1 (b) and is shown in detail in Fig. 4.2 (d). One end of the micro connector comprises a male thread and an inner conductor pin that mate with the back end of the three-dimensional wire. The other end of the micro connector is soldered to a coaxial cable⁵.

The end of each coaxial cable opposite to the three-dimensional wire is soldered to a sub-miniature push-on (SMP) connector. The SMP connectors are bolted to a horizontal plate attached to the microwave package by means of two vertical fixtures, as shown in Fig. 4.1 (d). The vertical fixtures and the horizontal plate constitute the package holder. The package holder and microwave package form an independent assembly. A

⁴The pillar was included in the design as there was concern over potential damage to the large $15\text{ mm} \times 15\text{ mm}$ substrates (particularly the Si ones) from mechanical strain due to the three-dimensional wires pushing on the top of the chip.

⁵The micro connector is necessary because the high temperatures generated by soldering a coaxial cable directly to the wire back end would damage some of the inner wire components. The breaking temperature of those components is even lower than the melting temperature of available eutectic solders.

horizontal mounting plate, designed to interface with the package holder, houses a set of matching SMP connectors. The mounting plate is mechanically and, thus, thermally anchored to the mixing chamber (MC) stage of a dilution refrigerator (DR).

4.2.4 Microwave simulations

The three-dimensional wires, the 90° transition between the wire and the on-chip pad as well as the inner cavities of a fully-assembled microwave package were extensively simulated numerically at microwave frequencies⁶. The results for the electromagnetic field distribution at a frequency of approximately 6 GHz, which is a typical operation frequency for superconducting qubits, are shown in Fig. 4.3. Figure 4.3 (a) shows the field behavior for a bare three-dimensional wire. The field distribution resembles that of a coaxial transmission line except for noticeable perturbations at the dielectric PTFE spacers. Figure 4.3 (b) shows the 90° transition region. This is a critical region for signal integrity since abrupt changes in physical geometry cause electrical reflections [116, 231]. In order to minimize such reflections, an impedance-matched pad was designed. However, this leads to a large electromagnetic volume in proximity of the pad, as seen in Fig. 4.3 (b), possibly resulting in parasitic capacitance and crosstalk.

In addition to considering the wire and the transition region, the electrical behavior of the inner cavities of the package was studied analytically and simulated numerically.

⁶The simulation software used was the high frequency three-dimensional full-wave electromagnetic field simulation software (also known as HFSS) by Ansys, Inc.

Table 4.1: Simulation results for the first three box modes of the lower cavity inside the assembled microwave package shown in Fig. 4.1 (b). The dielectric used for these simulations was Si at room temperature with relative electric permittivity $\epsilon_r = 11.68$. “Vacuum” indicates that no Si is present in the simulation. “with pillar” indicates that the $1.0 \text{ mm} \times 1.0 \text{ mm} \times 2.0 \text{ mm}$ support pillar is present. TE_{xyz} indicates the number of half-wavelengths spanned by the electric field in the x , y , and z directions, respectively (cf. Fig 4.3 (c)). Note that the frequency of the first mode of the upper cavity is $\sim 17.2 \text{ GHz}$.

	TE_{110} (GHz)	TE_{120} (GHz)	TE_{210} (GHz)
Vacuum	15.7	24.2	24.2
Vacuum with pillar	13.1	23.6	23.6
Si	13.5	16.8	16.8
Si with pillar	6.3	16.2	16.9




Figure 4.3: Numerical simulations of the electric field distribution. (a) Field for a three-dimensional wire at 6 GHz. (b) Field in proximity of the 90° transition region also at 6 GHz. (c) Field for the first box mode at 6.3 GHz. Color bar scales are indicated in their respective panels. The x , y , and z directions of a Cartesian coordinate system are also indicated. In (b), the cross-section of the transition region is shown. Note the large volume occupied by the electric field beneath the contact pad. In (c), an offset cross-section of the first box mode is shown. The field confinement due to the pillar is clearly visible. Additionally, the simulation shows a slight field confinement in the region surrounding the chip recess. A time-domain animation of the simulated electric field distributions can be found in the Supplemental Material of Ref. [220].

As described in Sec. 4.2.2, the metallic surface of the chip effectively divides the cavity of the sample holder into two regions: A vacuum cavity above the metal surface and a cavity partially filled with dielectric below the metal surface. The latter is of greatest concern as the dielectric acts as a perturbation to the cavity vacuum⁷, thus lowering the box modes. For a simple rectangular cavity, the frequency f of the first mode due to this perturbation can be found as [115],

$$f = f_0 - \frac{f_0 (\epsilon_r - 1) d_s}{2b} , \quad (4.1)$$

where f_0 is the frequency of the unperturbed mode, ϵ_r the relative electric permittivity of the dielectric, d_s the substrate thickness, and b the cavity height. From Eq. (4.1), we estimated this box mode to be 12.8 GHz. However, considering the presence of the pillar,

⁷Provided the vacuum still constitutes the majority of the volume of the cavity.

the three-dimensional wires, etc., we had to use numerical simulations to obtain a more accurate estimate of the lowest box modes. The results for the first three modes are reported in Table 4.1. Discounting the pillar, the analytical and simulated values are in good agreement with each other. The addition of the support pillar significantly lowers the frequency of the modes. In fact, it increases the relative filling factor of the cavity by confining more of the electromagnetic field to the dielectric than to vacuum. Given the dimensions of this design, the pillar leads to a first mode which could interfere with typical qubit frequencies. In spite of this, the pillar was included in the design in order to provide a degree of mechanical support. Note that the pillar can alternatively be realized as a dielectric material, e.g., PTFE; a dielectric pillar would no longer cause field confinement between the top surface of the pillar and the metallic surface of the chip.

4.3 Implementation

The physical implementation of the main components of the quantum socket is displayed in Fig. 4.4. In particular, Fig. 4.4 (a) shows a macro photograph of a three-dimensional wire. The inset shows a scanning electron microscope (SEM) image of the wire contact head, featuring the 500 μm version of the tunnel. This wire was cycled approximately ten times; as a consequence, the center conductor of the contact head, which had a conical, sharp shape originally, flattened at the top. The metallic components of the wire were made from bronze and brass (cf. Sec. 4.3.1), and all springs from hardened beryllium copper (BeCu). Except for the springs, all components were gold plated without any nickel (Ni) adhesion underlayer.

Figure 4.4 (b) displays the entire microwave package in the process of locking the package lid and sample holder together, with a chip and grounding washer already installed. As shown in the figure, two rows of three-dimensional wires, for a total number of six wires, are screwed into the lid with pressure settings as described in Appendix B.1; each wire is associated with one on-chip CPW pad. The four springs that mate with the grounding washer feet are embedded in corresponding recesses in the lid; the springs are glued in these recesses by way of a medium-strength thread locker that was tested at low temperatures. Figure 4.4 (c) shows a picture of the assembled microwave package attached to the package holder; the entire structure is attached to the MC stage of a DR. More details about materials and microwave components can be found in the Supplemental Material of Ref. [220].




Figure 4.4: Images of the quantum socket as implemented. (a) Macro photograph of a three-dimensional wire; inset: SEM image of the contact head. Note that the tip of the inner conductor retained small metallic flakes that were scraped off the on chip pads. (b) Microwave package lid with six three-dimensional wires and four washer springs, washer, and sample holder with chip installed. (c) Package holder with attached microwave package mounted to the MC stage of a DR. The lid of a custom-made magnetic shield can be seen at the top of the panel.

4.3.1 Magnetic properties

An important stage in the physical implementation of the quantum socket was the choice of materials to be used for the three-dimensional wires. In fact, it has been shown that non-magnetic components in proximity of superconducting qubits are critical to preserve long qubit coherence [234, 235, 236, 105]. The three-dimensional

wires are the closest devices to the qubits. For this reason, all their components should be made using non-magnetic materials. Due to machining constraints, however, alloys containing ferromagnetic impurities (iron (Fe), cobalt (Co), and Ni) had to be used. For the outer conductor components we used brass, which is easy to thread; the chosen type was CW724R⁸. For the inner conductor components, brass CW724R did not meet the machining requirements. Consequently, we decided to use copper alloy (phosphor bronze) CW453K⁹. The chemical composition for these two materials is reported in Table B.1 of Appendix B.2. The dielectric spacers were made from PTFE and the rest of the components from hardened BeCu; both materials are non-magnetic. The weight percentage of ferromagnetic materials is non-negligible for both CW453K and CW724R. Thus, we performed a series of tests using a zero Gauss chamber (ZGC) in order to ensure both materials were sufficiently non-magnetic. The results are given in Appendix B.2 and show that the magnetic impurities should be small enough not to disturb the operation of superconducting quantum devices.

The microwave package and grounding washer were made from high-purity Al alloy 5N5 (99.9995 % purity). The very low level of impurities in this alloy assures minimal stray magnetic fields generated by the package and washer, as confirmed by the magnetic tests discussed in Appendix B.2.

4.3.2 Thermal properties

The thermal conductance of the three-dimensional wires is a critical parameter to be analyzed for the interconnection with devices at cryogenic temperatures. Low thermal conductivity would result in poor cooling of the devices, which, in the case of qubits, may lead to an incoherent thermal mixture of the qubit ground state $|g\rangle$ and excited state $|e\rangle$ [237]. Even a slightly mixed state would significantly deteriorate the fidelity of the operations required for QEC [238]. It has been estimated that some of the qubits in the experiment of Ref. [108], which relies solely on Al wire bonds as a means of thermalization, were characterized by an excited state population $P_e \approx 0.04$. Among other possible factors, it is believed that this population was due to the poor thermal conductance of the Al wire bonds. In fact, these bonds become superconductive at the desired qubit operation temperature of ~ 10 mK, preventing the qubits from thermalizing and, thus, from being initialized in $|g\rangle$ with high fidelity.

In order to compare the thermal performance of an Al wire bond with that of a three-dimensional wire, we estimated the heat transfer rate per kelvin of a wire, Π_t , using a simplified coaxial geometry. At a temperature of 25 mK, we calculated $\Pi_t \approx 6 \times 10^{-7} \text{ W K}^{-1}$. At the same temperature, the heat transfer rate per kelvin of a typical

⁸Alloy 430, grade ISO CuZn21Si3P, UNS C69300.

⁹Grade DIN 2.1030 - CuSn8, UNS C52100.

Al wire bond was estimated to be $\Pi_b \simeq 4 \times 10^{-12} \text{ W K}^{-1}$ (cf. Appendix B.3 for more details). A very large number of Al wire bonds would thus be required to obtain a thermal performance comparable to that of a single three-dimensional wire.

4.3.3 Spring characterization

Another critical step in the physical implementation of the quantum socket was to select springs that work at cryogenic temperatures. In fact, the force that a wire applies to a chip depends on these springs. This force, in turn, determines the wire-chip contact resistance, which impacts the socket's DC and, possibly, microwave performance. Among various options, we chose custom springs made from hardened BeCu.

It is noteworthy to mention that the mean number of cycles before mechanical failure for the three-dimensional wires is larger than 200000 at room temperature (cf. Appendix B.4 for details); at 10 mK, we were able to use the same wire more than ten times without any mechanical or electrical damage.

To characterize the springs, their compression was assessed at room temperature, in liquid nitrogen (i.e., at a temperature $T \simeq 77 \text{ K}$), and in liquid helium ($T \simeq 4.2 \text{ K}$). Note that a spring working at 4.2 K is expected to perform similarly at a temperature of 10 mK. A summary of the thermo-mechanical tests is reported in Appendix B.4. The main conclusion of the tests is that the springs do not break (even after numerous temperature cycles) and have similar spring constants at all measured temperatures.

4.3.4 Alignment

In order to implement a quantum socket with excellent interconnectivity properties, it was imperative to minimize machining errors and mitigate the effects of any residual errors. These errors are mainly due to: Dicing tolerances; tapping tolerances of the M2.5-threaded holes of the lid; tolerances of the mating parts for the inner cavities of the lid and sample holder; tolerances of the chip recess. These errors can cause both lateral and rotational misalignment and become likely worse when cooling the quantum socket to low temperatures. More details on alignment errors can be found in Appendix B.5.

The procedure to obtain an ideal and repeatable alignment comprises three main steps: Optimization of the contact pad and tunnel geometry; accurate and precise chip dicing; accurate and precise package machining. For the quantum socket described in this chapter, the optimal tunnel width was found to be $650 \mu\text{m}$. This maintained reasonable impedance matching, while allowing greater CPW contact pad and tapering dimensions. The contact pad width W_p and taper length T_p were chosen to be $W_p = 320 \mu\text{m}$ and $T_p = 360 \mu\text{m}$. These are the maximum dimensions allowable that accommodate




Figure 4.5: Micro images used to evaluate the alignment procedure of the three-dimensional wires. (a)-(b) Ag pads. The magenta arrows indicate the first (1) and second (2) mating instance. The lengths W_p and T_p are indicated in (a) by means of magenta bars. (c)-(d) Al pad before and after a cooling cycle to ~ 10 mK. Center conductor dragging due to cooling is indicated by a green bar. The magenta dashed line in (a) indicates tunnel (i.e., rotational) alignment for the Ag pad. Note that the geometries for the pads in panels (a) and (b) are optimized for a $3\ \mu\text{m}$ Ag film and, thus, are slightly different than those for the pads in panels (c) and (d), which are designed for a $120\ \text{nm}$ Al film.

the geometry of the wire bottom interface for a nominal lateral and rotational misalignment of $\mp 140\ \mu\text{m}$ and $\mp 28^\circ$, respectively. In order to select the given pad dimensions, we had to resort to a $50\ \Omega$ matched raised-cosine tapering.

The majority of the chips used in the experiments presented here was diced with a dicing saw from the DISCO Corporation, model DAD3240. To obtain a desired die length,

both the precision of the saw stage movement and the blade's kerf had to be considered. For the DAD3240 saw, the former is $\sim 4 \mu\text{m}$, whereas the latter changes with usage and materials. For the highest accuracy cut, we measured the kerf on the same type of wafer just prior to cutting the actual die. In order to achieve maximum benefit from the saw, rotational and lateral alignment dicing markers were incorporated on the wafer. Such a meticulous chip dicing procedure is only effective in conjunction with a correspondingly high level of machining accuracy and precision. We used standard computer numerical control (CNC) machining with a tolerance of 1 thou ($25.4 \mu\text{m}$), although electrical discharge machining can be pursued if more stringent tolerances ($\lesssim 10 \mu\text{m}$) are required.

Following the aforementioned procedures we were able to achieve the desired wire-pad matching accuracy and precision, which resulted in a repeatability of 100 % over 94 instances. These figures of merit were tested in two steps: First, by micro imaging several on-chip pads that were mated to a three-dimensional wire; second, by means of DC resistance tests.

On-chip pad micro imaging

Micro imaging was performed on a variety of different samples, four of which are exemplified in Fig. 4.5. The figure shows a set of micro images for Ag and Al pads (details regarding the fabrication of these samples are available in Appendix B.6). Figure 4.5 (a) and (b) show two Ag pads that were mated with the three-dimensional wires at room temperature. Panel (a) shows a mating instance where the wire bottom interface perfectly matched the on chip pad. Panel (b) shows two mating instances that, even though not perfectly matched, remained within the designed tolerances. Notably, simulations of imperfect mating instances revealed that an off-centered wire does not significantly affect the microwave performance of the quantum socket. Finally, panels (c) and (d) display two Al pads which were both mated with a wire one time. While the pad in (c) was operated only at room temperature, the pad in (d) was part of an assembly that was cooled to $\sim 10 \text{ mK}$ for approximately three months. The image was taken after the assembly was cycled back to room temperature and shows dragging of the wire by a few tens of micrometers. Such a displacement can likely be attributed to the difference in the thermal expansion of Si and Al (cf. Appendix B.5).

As a diagnostic tool, micro images of a sample already mounted in the sample holder after a mating cycle can be obtained readily by means of a handheld digital microscope.

DC resistance tests

In contrast to the micro imaging tests, which require the removal of the microwave package's lid, DC resistance tests can be performed *in situ* at room temperature after the

Table 4.2: DC resistance tests. Multiple Au samples were measured. For all samples the length from the center of one pad to that of the opposite pad of the CPW center conductor is $L_{pp} = 11.5$ mm. In the table are reported: The width W of each CPW transmission line; the thickness d of the metal; the metal volume resistivity ρ at room temperature or at 77 K; the input and output wire pressure settings ℓ_p^i and ℓ_p^o , respectively; the operating temperature T ; the number of measurements N ; the estimated trace resistance R^t (for the Au samples, the very large parallel resistance ~ 46 k Ω at room temperature due to the titanium (Ti) adhesion layer was neglected); the measured resistances R_{io} , R_{ig} , and R_{og} . For a given chip, each resistance was measured independently N times under similar measurement conditions. The mean values and standard deviations of R_{io} are provided; the minimum values of R_{ig} and R_{og} are given. Note that because $R^c + R^{wc} \ll R^t$, we expect $R_{io} \approx R^t$. The discrepancy between the estimated and measured values (R^t and R_{io}) for the Au and Al samples is mainly due to uncertainties associated with the metal thickness d . The inaccuracies are smaller for thicker films, as in the case of the 3 μm Ag samples.

Metal (-)	W (μm)	d (nm)	ρ (Ωm)	ℓ_p^i (mm)	ℓ_p^o (mm)	T (K)	N (-)	R^t (Ω)	R_{io} (Ω)	R_{ig} (M Ω)	R_{og} (M Ω)
Au	10	100	22	4.52	4.44	300	30	253	218(3)	31	31
Au	10	100	22	4.97	4.89	300	2	253	223(0)	38	38
Au	10	100	22	4.18	4.11	300	2	253	217(0)	39	39
Au	10	100	22	4.57	4.45	300	2	253	229(0)	28.8	28.6
Au	10	200	22	4.60	4.70	300	10	126.5	98.0(7)	50	50
Au	10	200	4.55	4.60 ^a	4.70 ^a	77	6	26.16	36.02(2)	77.3	81.8
Ag	30	3000	16	4.60	4.70	300	6	2.04	2.71(4)	0.0043	0.0043
Al	15	120	26	4.25	4.07	300	24	166.1	171(1)	0.0042	0.0042

^a At 300 K.

package and package holder have been fully assembled. These tests were performed on all devices presented in this chapter, including Au, Ag, and Al samples.

The typical test setup comprises a microwave package with two three-dimensional wires each mating with an on-chip pad. The two pads are connected by means of a CPW transmission line with series resistance R^t . The back end of the wires is connected to a coaxial cable ending in a microwave connector, similar to the setup in Figs. 4.1 (d) and 4.4 (c). The DC equivalent circuit of this setup can be represented by way of a four-terminal Pi network. The circuit comprises an input “i” and output “o” terminal, two terminals connected to a common ground “g,” an input-output resistor with resistance R_{io} , and two resistors to ground with resistance R_{ig} and R_{og} . The i and o terminals correspond to the inner conductor of the two microwave connectors. The

outer conductor of both connectors is grounded.

The resistance R_{i_0} is that of the center conductor of the CPW transmission line, including the contact resistance R^c for each wire-pad interface and the series resistance R^{wc} of the wire's and coaxial cable's inner conductor, $R_{i_0} = R^t + 2(R^c + R^{wc})$. The resistances R_{i_g} and R_{o_g} are those of the path between each center conductor and ground and include the resistance of the inner and outer conductor of the various coaxial cables and wires as well as any wire-pad contact resistance. Ideally, these ground resistances should be open circuits. In reality, they are expected to have a finite but large value because of the intrinsic resistance of the Si wafers used as a substrate.

The design parameters, electrical properties, measurement conditions as well as the measured values of R_{i_0} , R_{i_g} , and R_{o_g} for various Au, Ag, and Al samples are reported in Table 4.2. Measuring resistances significantly different from the expected values meant that either a lateral or rotational misalignment occurred. The resistances for some Au samples were also measured at 77 K to verify whether a good room temperature alignment persisted in cryogenic conditions. The cold measurements were realized by dunking the package holder into liquid nitrogen. Note that one chip with sapphire substrate and Al conductors was also measured; in this case, both R_{i_g} and R_{o_g} were larger than 500 M Ω . Notably, we found a 100 % correlation between a successful DC test at room temperature and a microwave measurement at 10 mK.

The measured value of R_{i_0} for the Ag samples is larger than the estimated trace resistance by ~ 650 m Ω . This simple result makes it possible to find an upper bound value for the contact resistance, $R^c \lesssim 325$ m Ω . A more accurate estimate of the contact resistance based on four-point measurements will be described in Sec. 4.4.1.

The DC resistance testing procedure presented here will be useful in integrated-circuit quantum information processing, where, for example, CPW transmission lines can serve as qubit readout lines [107, 239, 108]. These tests can be expanded to encompass different circuit structures such as the qubit control lines utilized in Ref. [108].

4.4 Characterization

The three-dimensional wires are multipurpose interconnects that can be used to transmit signals over a wide frequency range, from DC to 10 GHz. These signals can be: The current bias used to tune the transition frequency of a superconducting qubit; the Gaussian-modulated sinusoidal or the rectangular pulses that, respectively, make it possible to perform XY and Z control on a qubit; the continuous monochromatic microwave tones used to read out a qubit state or to populate and measure a superconducting resonator [90, 107, 108, 105]. In general, the wires can be used to transmit any baseband modulated carrier signal within the specified frequency spectrum, at room and cryogenic temperatures. In this section, we report experimental results for a




Figure 4.6: I-V characteristic curve for R_{ig} . The sweeps were conducted by both increasing (red) and decreasing (blue) the applied current between -7 mA and $+7$ mA. The shaded region indicates two standard deviations. The dashed black lines indicate the region (± 1.5 mA) for which the resistance value was found using linear regression. The origin of the hysteresis is explained in the text. The inset shows the circuit diagram of the device under test, including all resistors measured by means of the four-point measurement. The position of the pad is indicated by an arrow.

series of measurements aiming at a complete electrical characterization of the quantum socket at room temperature and at approximately 77 K (i.e., in liquid nitrogen).

4.4.1 Four-point measurements

The wire-pad contact resistance R^c is an important property of the quantum socket. In fact, a large R^c would result in significant heating when applying DC bias signals and rectangular pulses, thus deteriorating qubit performance.

In order to assess R^c for the inner and outer conductor of a three-dimensional wire, we performed four-point measurements using the setup shown in the inset of Fig. 4.6. Using this setup, we were able to measure both the series resistance of the wire R^w and the contact resistance R^c .

The setup comprises a microwave package with a chip entirely coated with a 120 nm thick Al film; no grounding washer was used. The package featured three three-dimensional wires, of which two were actually measured; the third wire was included to provide mechanical stability. The package was attached to the MC stage of a DR and

measured at room temperature by means of a precision source-measure unit (SMU); more details in the Supplemental Material of Ref. [220].

We measured the resistance between the inner conductor of a wire and ground, R_{ig} . This resistance comprises the inner conductor wire resistance R_i^w in series with the inner conductor contact resistance R_i^c and any resistance to ground, R_g . Note that, at the operation temperature of the experiment (~ 10 mK), Al is superconducting and, thus, the metal resistance can be neglected.

Figure 4.6 shows the current-voltage (I-V) characteristic curve for R_{ig} . With increasing bias currents, the contact resistance results in hot-spot generation leading to a local breakdown of superconductivity. For sufficiently high bias currents, superconductivity breaks down completely. At such currents, the observed hysteretic behavior indicates the thermal limitations of our setup. Note, however, that these currents are at least one order of magnitude larger than the largest bias current required in typical superconducting qubit experiments [92].

In order to estimate R_{ig} from the I-V characteristic curve, we selected the bias current region from -1.5 mA to $+1.5$ mA and fitted the corresponding slope. We obtained $R_{ig} \approx 148$ m Ω . This value, which represents an upper bound for the wire resistance and the wire-pad contact resistance, ($R_i^w + R_i^c$), is significantly larger than that associated with Al wire bonds [240] or indium bump bonds [226]. In future versions of the three-dimensional wires, we will attempt to reduce the wire-pad contact resistance by rounding the tip of the center conductor, stiffening the wire springs, using a thicker metal film for the pads, depositing Au or titanium nitride (TiN) on the pads, and plating the wires with TiN. We note, however, that even a large value of the wire and/or wire-pad contact resistance will not significantly impair the quantum socket microwave performance; for example, the quantum architecture in Ref. [228] would be mostly unaffected by the contact resistance of our three-dimensional wires.

4.4.2 Two-port scattering parameters

The two-port scattering parameter (S-parameter) measurements of a bare three-dimensional wire were realized by means of the setup shown in the inset of Fig. 4.7 (a) and described in detail in the Supplemental Material of Ref. [220]. The device under test (DUT) comprises a cable assembly attached to a three-dimensional wire by means of a screw-in micro connector. The bottom interface of the wire is connected to a 2.92 mm end launch connector, which is characterized by a flush coaxial back plane; this plane mates with the wire bottom interface well enough to allow for S-parameter measurements up to 10 GHz. In order to measure the S-parameters of the DUT, we used a vector network analyzer (VNA) and performed a two-tier calibration (cf. supplement), which made it possible to set the measurement planes to the ports of the DUT.




Figure 4.7: S-parameter measurements and simulations for a three-dimensional wire at room temperature. (a) Magnitude of the measured S-parameters $|S_{mn}|$, with $m, n = \{1, 2\}$. Inset: Image of the measurement setup. From left to right: Segment of flexible coaxial cable (gray); Sub-miniature type A (SMA) female connector (red); after plane i, SMA male connector (orange); segment of semi-rigid coaxial cable EZ 47 cable (gray; cf. supplement); screw-in micro connector (green); three-dimensional wire (purple); after plane ii, 2.92 mm end launch connector (white and black); SMA female connector (red); segment of flexible coaxial cable (gray). (b) S-parameter simulations. The lower attenuation is due to idealized material properties and connections.

The magnitudes of the measured reflection and transmission S-parameters are displayed in Fig. 4.7 (a). We performed microwave simulations of a three-dimensional wire for the same S-parameters (cf. Sec. 4.2.4 for the electric field distribution), the results of which are plotted in Fig. 4.7 (b). The S-parameters were measured and simulated between 10 MHz and 10 GHz. The S-parameters $|S_{21}|$ and $|S_{12}|$ show a featureless microwave response, similar to that of a coaxial transmission line. The attenuation at 6 GHz is $|S_{21}| \approx -0.58$ dB and the magnitude of the reflection coefficients at the same frequency is $|S_{11}| \approx -13.8$ dB and $|S_{22}| \approx -14.0$ dB. The phase of the various S-parameters (not shown) behaves as expected for a coaxial transmission line. All measurements were performed at room temperature.

The S-parameter measurements of a three-dimensional wire indicate a very good microwave performance. However, these measurements alone are insufficient to fully characterize the quantum socket operation. A critical feature that deserves special




Figure 4.8: Microwave characterization setup. The vertical black dashed lines indicate main reflection planes. The yellow terminations correspond to SMA male connectors at the end of each cable. The input(output) flexible cable corresponds to the region in between planes i and ii(xii and xiii), in gray; the blue blocks correspond to SMA female bulkhead adapters; the plane ii(xii) corresponds to the input(output) port of the DUT; the orange block corresponds to an SMA male to SMA female adapter; the EZ 47 input(output) cable corresponds to the region in between planes iv and v(x and xi), in gray; the plane v(x) corresponds to the solder connection on the three-dimensional wire; the plane vi(ix) is associated with the screw-in micro connector; the plane vii(viii) corresponds to the 90° interface connecting each three-dimensional wire to the input(output) of the CPW transmission line (pale blue). The three-dimensional wires are indicated in purple.

attention is the 90° transition region between the wire bottom interface and the on-chip CPW pad. It is well-known that 90° transitions can cause significant impedance mismatch and, thus, signal reflection [116]. In quantum computing applications, these reflections could degrade both the qubit control and readout fidelity.

Figure 4.8 shows a typical setup for the characterization of a wiring configuration analogous to that used for qubit operations (cf. Supplemental Material of Ref. [220] for details). The setup comprises a DUT with ports 1 and 2 connected to a VNA; the DUT incorporates a microwave package with a pair of three-dimensional wires, which address one CPW transmission line on an Au or Ag chip. The microwave package was attached to the package holder, as described in Sec. 4.2.3 and Sec. 4.3 (cf. also Figs. 4.1 (d) and 4.4 (c)). The transmission line geometrical dimensions and wire pressure settings are reported in Table 4.2; only the 200 nm Au samples and the Ag samples were characterized at microwave frequencies. The back end of each three-dimensional wire is connected to one end of an EZ 47 cable by means of the screw-in micro connector described in Sec. 4.2.3; the other end of the EZ 47 cable is soldered to an SMA male connector. A calibration was performed for all measurements.

We performed a two-port S-parameter measurement of the DUT from 10 MHz to 10 GHz. The measurement results at room temperature for the Au and Ag samples are shown in Figs. 4.9 (a) and 4.10 (a), respectively. The results for the Au sample at 77 K are shown in Fig. 4.9 (b).

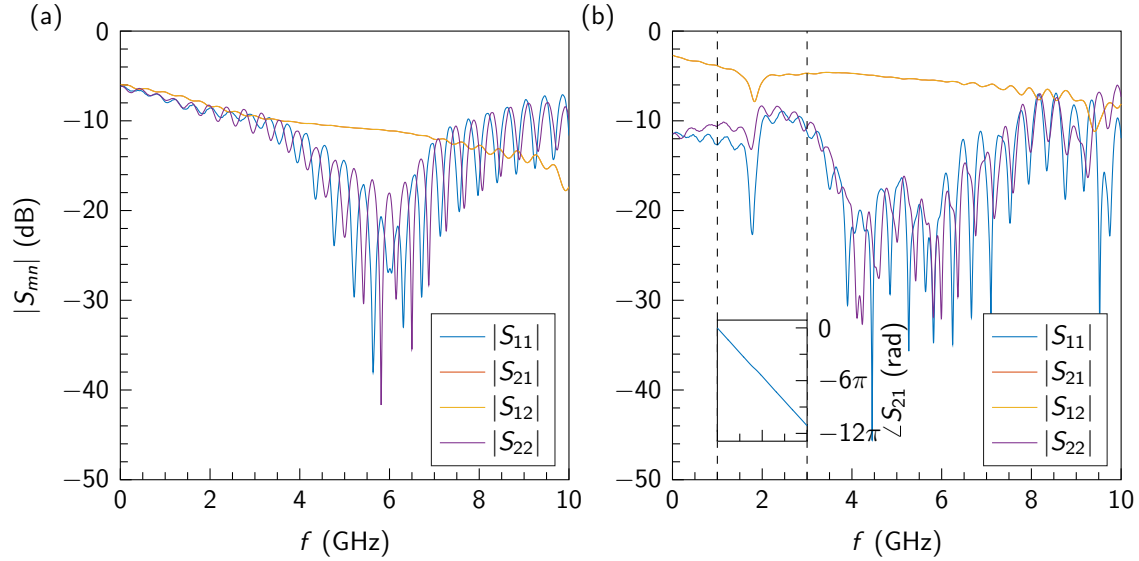


Figure 4.9: S-parameter measurements for the Au sample. (a) $|S_{mn}|$ at room temperature. (b) $|S_{mn}|$ at 77 K. The transmission coefficients show that the DUT is a reciprocal device (i.e., $S_{21} \simeq S_{12}$), as expected for a passive structure. The inset in (b) shows the unwrapped phase angle $\angle S_{21}$; the black dashed lines delimit the frequency region between 1 GHz and 3 GHz. Note that the reflection coefficients S_{11} and S_{22} are relatively large at very low frequency. This is expected for a very lossy transmission line. In fact, the center conductor for the Au sample is characterized by a series resistance $R_{i_0} \simeq 98 \Omega$ at room temperature (cf. Table 4.2), which corresponds to $S_{11} \sim S_{22} \simeq -6$ dB at 10 MHz, and $R_{i_0} \simeq 36 \Omega$ at 77 K, which corresponds to $S_{11} \sim S_{22} \simeq -12$ dB at 10 MHz. These findings are consistent with the time-domain results to be shown in Fig. 4.12, where the large impedance steps are also due to the large series resistance (cf. Sec. 4.4.3). The low-loss Ag sample shows much lower reflection coefficients at low frequency (cf. Fig. 4.10), whereas the lossy Al sample shows high reflections at low frequency and room temperature (cf. Fig. 4.16 (a) in Sec. 4.5).

The S-parameter measurements of the Au sample show that the quantum socket functions well at microwave frequencies, both at room temperature and at 77 K. Since most of the mechanical shifts have already occurred when cooling to 77 K [241], this measurement allows us to deduce that the socket will continue functioning even at lower temperatures, e.g., ~ 10 mK. The Au sample, however, is characterized by a large value of R_{i_0} , which may conceal unwanted features both in the transmission and reflection measurements. Therefore, we prepared an Ag sample that exhibits a much lower resistance even at room temperature. The behavior of the Ag S-parameters




Figure 4.10: S-parameter measurements for the Ag sample. (a) $|S_{mn}|$ measurement at room temperature. (b) $|S_{mn}|$ microwave simulation. Note that the attenuation for the measured data is larger than that in the simulation because the latter does not include the EZ 47 cables used in the DUT of Fig. 4.8.

is similar to that of a transmission line or coaxial connector. For example, $|S_{11}|$ is approximately -15 dB; as a reference, for a high-precision SMA connector at the same frequency $|S_{11}| \approx -30$ dB.

The presence of the screw-in micro connector can occasionally deteriorate the microwave performance of the quantum socket. In fact, if the micro connector is not firmly tightened, a dip in the microwave transmission is observed. At room temperature, it is straightforward to remove the dip by simply re-tightening the connector when required. On the contrary, for the measurements at 77 K and for any other application in a cryogenic environment assuring that the micro connector is properly torqued at all times can be challenging. Figure 4.9 (b), for example, shows the S-parameters for an Au sample measured at 77 K. A microwave dip appeared at approximately 1.8 GHz, with a 3 dB bandwidth of approximately 200 MHz. The inset in Fig. 4.9 (b) displays the phase angle of S_{21} between 1 GHz and 3 GHz, showing that the dip is unlikely a Lorentzian-type resonance (more details in the Supplemental Material of Ref. [220]). Note that the dip is far from the typical operation frequencies for superconducting qubits. Additionally, as briefly described in Sec. 4.6, we will remove the screw-in micro connector from future generations of the three-dimensional wires.

Figure 4.10 (b) shows a simulation of the S-parameters for the Ag sample, for the

same frequency range as the actual measurements. While there are visible discrepancies between the measured and simulated S-parameters, the latter capture well some of the characteristic features of the microwave response of the DUT. In particular, the measured and simulated reflection coefficients display a similar frequency dependence. It is worth mentioning that we also simulated the case where the wire bottom interface is not perfectly aligned with the on-chip pad (results not shown). We considered lateral misalignments of 100 μm and rotational misalignments of $\sim 20^\circ$. This allowed us to study more realistic scenarios, such as those shown in Fig. 4.5. We found that the departure between the misaligned and the perfectly aligned simulations was marginal. For example, the transmission S-parameters varied only by approximately ∓ 0.5 dB.

Figure 4.11 presents a set of microwave parameters that help further analyze the performance of the quantum socket. These parameters were calculated from the measured S-parameter data of Figs. 4.9 and 4.10 (a) for the Au sample at room temperature and at 77 K and for the Ag sample at room temperature. They make it possible to characterize the input and output impedance as well as the dispersion properties of the quantum socket.

The complex input impedance can be obtained from the frequency dependent impedance matrix $\mathbf{Z} = [Z_{mn}]$ as [199]

$$Z_{\text{in}} = Z_{11} - \frac{Z_{12}Z_{21}}{Z_{22} - Z_L} \quad , \quad (4.2)$$

where $Z_L = Z_c = 50 \Omega$ is the load impedance. The impedance matrix was obtained using the measured complex S-parameter matrix $\mathbf{S} = [S_{mn}]$ from

$$\mathbf{z} = \sqrt{Z_c} \left(\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + \mathbf{S} \right) \left(\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} - \mathbf{S} \right)^{-1} \sqrt{Z_c} \quad . \quad (4.3)$$

The magnitude of Z_{in} is shown in Fig. 4.11 (a).

The input voltage standing wave ratio (VSWR) was obtained from [115]

$$\text{VSWR}_{\text{in}} = \frac{1 + |S_{11}|}{1 - |S_{11}|} \quad (4.4)$$

and is displayed in Fig. 4.11 (b).

The phase delay was calculated as [199]

$$\tau_\phi = -\frac{1}{2\pi} \frac{\angle S_{21}}{f} \quad (4.5)$$

and is displayed Fig. 4.11 (c).

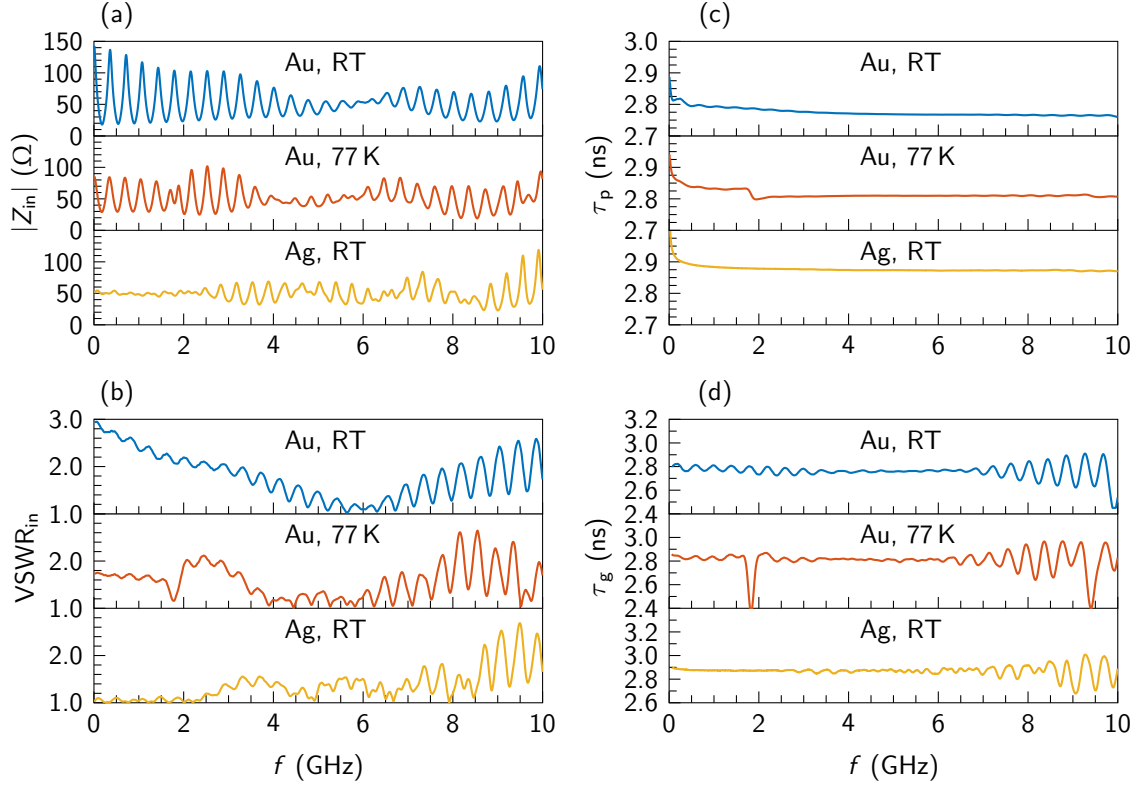


Figure 4.11: Quantum socket microwave parameters. (a) Input impedance magnitude $|Z_{in}|$. (b) Input VSWR, $VSWR_{in}$. (c) Phase delay τ_p . (d) Group delay τ_g . Blue corresponds to the Au sample at room temperature (RT), red to the Au sample at 77 K, and orange to the Ag sample at room temperature (RT).

Finally, the group delay was obtained from [115]

$$\tau_g = -\frac{1}{2\pi} \frac{\partial}{\partial f} (\angle S_{21}) \quad (4.6)$$

and is displayed in Fig. 4.11 (d). The derivative in Eq. (4.6) was evaluated numerically by means of central finite differences with 6th order accuracy. The data in Fig. 4.11 (d) were post-processed using 1% smoothing. Note that the output impedance and VSWR were also evaluated and resembled the corresponding input parameters.

The input and output impedances as well as the VSWRs indicate a good impedance matching up to approximately 8 GHz. The phase and group delays, which are directly related to the frequency dispersion associated with the quantum socket, indicate minimal dispersion. This is expected for a combination of coaxial structures (the three-dimensional wires) and a CPW transmission line. Thus, we expect wideband control

pulses to be transmitted without significant distortion in applications with superconducting qubits (cf. Supplemental Material of Ref. [220] for further details about microwave pulse transmission).

4.4.3 Time-domain reflectometry

In time-domain reflectometry (TDR) measurements, a rectangular pulse with fast rise time and fixed length is applied to a DUT; the reflections (and all re-reflections) due to all reflection planes in the system (i.e., connectors, geometrical changes, etc.) are then measured by way of a fast electrical sampling module. The reflections are, in turn, related to the impedances of all of the system components. Thus, TDR makes it possible to estimate any impedance mismatch and its approximate spatial location in the system.

TDR measurements were performed on the DUT shown in Fig. 4.8, with the same Au or Ag sample as for the measurements in Sec. 4.4.2. As always, the Au sample was measured both at room temperature and at 77 K, whereas the Ag sample was measured only at room temperature. The TDR setup is analogous to that used for the S-parameter measurements, with the following differences: The DUT input and output reference planes were extended from planes ii and xii to planes i and xiii of Fig. 4.8; when testing the DUT input port, the output port was terminated in a load with impedance $Z_L = Z_c$ and vice versa when testing the DUT output port. The TDR measurements were realized by means of a sampling oscilloscope with key features reported in the Supplemental Material of Ref. [220]. The voltage reflected by the DUT, V^- , is acquired as a function of time t by means of the oscilloscope. The time t is the round-trip interval necessary for the voltage pulse to reach a DUT reflection plane and return back to the oscilloscope. The measured quantity is given by¹⁰

$$V_{\text{meas}}(t) = V^+(t) + V^-(t) \quad , \quad (4.7)$$

where V^+ is the amplitude of the incident voltage square wave. From Eq. (4.7), we can obtain the first-order instantaneous impedance as [242]

$$Z(t) = Z_c \frac{1 + \xi(t)}{1 - \xi(t)} \quad , \quad (4.8)$$

where $\xi(t) = (V_{\text{meas}}(t) - V^+) / V^+$.

¹⁰Note that the DUT is a piecewise transmission line inhomogeneously filled with dielectric materials. Transforming the time t into distance is only possible with detailed knowledge of geometries and materials for all regions of the DUT. Since this information is not known to a high degree of accuracy, we prefer to express all measured quantities as a function of t .




Figure 4.12: TDR measurements for the Au sample at room temperature (blue) and 77 K (red). The inset shows the room temperature data associated with part of the EZ 47 cable, the input three-dimensional wire, and part of the CPW transmission line.

Figure 4.12 shows $Z(t)$ for the DUT with the Au sample at room temperature and at 77 K; the measurement refers to the input port of the DUT, i.e., plane i in Fig. 4.8. The figure inset shows the room temperature data for a shorter time interval. This corresponds to a space interval beginning at a point between planes iv and v and ending at a point between planes vii and viii in Fig. 4.8.

Figure 4.13 (a) shows $Z(t)$ for the Ag sample at room temperature. Figure 4.13 (b) displays the data in (a) for a time interval corresponding to a space interval beginning at a point between planes iv and v and ending at a point between planes x and xi in Fig. 4.8; as a reference, the Au data are overlaid with the Ag data.

For the Au sample, the first main reflection plane (plane ii) is encountered at $t \approx 18$ ns. The second main reflection plane (plane v) appears after ~ 2.5 ns relative to the first plane, at $t \approx 20.5$ ns. From that time instant and for a span of approximately 250 ps, the TDR measurement corresponds to $Z(t)$ of the three-dimensional wire itself. The maximum impedance mismatch between the EZ 47 coaxial cable and the three-dimensional wire is approximately 10Ω . The third main reflection plane (plane vii) corresponds to the 90° transition region; for the Au sample, it is impossible to identify features beyond this plane owing to the large series resistance of the on-chip CPW transmission line. From empirical evidence, the impedance $Z(t)$ of a lossy line with series resistivity ρ increases linearly with the length of the line L as $\rho L/(Wd)$. In fact, for the Au sample we measured an impedance step across the CPW transmission




Figure 4.13: TDR measurements for the Ag sample at room temperature. (a) Measurement of port 1 of the setup in Fig. 4.8. (b) Zoomin of (a) addressing the three-dimensional wire and the 90° transition region between the wire and the CPW transmission line (blue). The room temperature Au data (red) are also displayed as a reference.

line of approximately 100Ω at room temperature and 40Ω at 77 K. These steps are approximately the R_{i0} values reported in Table 4.2.

In order to obtain a detailed measurement of the impedance mismatch beyond the 90° transition region, we resorted to the TDR measurements of the DUT with the much less resistive Ag sample. First, we confirmed that $Z(t)$ of the input three-dimensional wire for the Ag sample is consistent with the TDR measurements of the Au sample; this is readily verified by inspecting Fig. 4.13 (b). The three-dimensional wire is the structure ending at the onset of the large impedance step shown by the Au overlaid data. The structure spanning the time interval from $t \approx 20.75$ ns to $t \approx 21$ ns is associated with the input transition region, the CPW transmission line, and the output transition region. The output three-dimensional wire starts at $t \approx 21$ ns, followed by the EZ 47 coaxial cable, which finally ends at the SMA bulkhead adapter at $t \approx 23.5$ ns. The maximum impedance mismatch associated with the transition regions and the CPW transmission line is $\sim 5 \Omega$. Notably, this mismatch is smaller than the mismatch between the three-dimensional wire and the coaxial cable. This is an important result. In fact, while it would be hard to diminish the impedance mismatch due to the transition region, it is feasible to further minimize the wire mismatch by creating accurate lumped-element models of the wire and use them to minimize stray

capacitances and/or inductances ¹¹.

It is worth comparing $Z(t)$ of the quantum socket with that of a standard package for superconducting qubits, where wire bonds are used to make interconnections between a printed circuit board and the control and measurement lines of a qubit on a chip. A detailed study of the impedance mismatch associated with wire bonds is found in Ref. [230], where the authors have shown that a long wire bond (of length between ~ 1 mm and 1.5 mm; typical length in most applications) can lead to an impedance mismatch larger than 15Ω (cf. Fig. S3 in the supplementary information of Ref. [230]); on the contrary, a short wire bond (between ~ 0.3 mm and 0.5 mm; less typical and more challenging to realize) results in a much smaller mismatch, approximately 2Ω . In terms of impedance mismatch the current implementation of the quantum socket, which is limited by the mismatch of the three-dimensional wires, lies in between these two extreme scenarios.

4.4.4 Signal crosstalk

Crosstalk is a phenomenon where a signal being transmitted through a channel generates an undesired signal in a different channel. Inter-channel isolation is the figure of merit that quantifies signal crosstalk and that has to be maximized to improve signal integrity. Crosstalk can be particularly large in systems operating at microwave frequencies, where, if not properly designed, physically adjacent channels can be significantly affected by coupling capacitances and/or inductances. In quantum computing implementations based on superconducting quantum circuits, signal crosstalk due to wire bonds has been identified to be an important source of errors and methods to mitigate it have been developed [233, 243, 244]. However, crosstalk remains an open challenge and isolations (opposite of crosstalk) lower than 20 dB are routinely observed when using wire bonds ¹². The coaxial design of the three-dimensional wires represents an advantage over wire bonds. The latter, being open structures, radiate more electromagnetic energy that is transferred to adjacent circuits. The former, being enclosed by the outer conductor, limit crosstalk due to electromagnetic radiation.

In realistic applications of the quantum socket, the three-dimensional wires must land in close proximity of several on-chip transmission lines. In order to study inter-channel isolation in such scenarios, we designed a special device comprising a pair of CPW transmission lines, as shown in the inset of Fig. 4.14 (a). One transmission line connects two three-dimensional wires (ports 1 and 2), exactly as for the devices studied in Subsecs. 4.4.2 and 4.4.3; the other line, which also connects two three-dimensional wires (ports 3 and 4), circumvents the wire at port 1 by means of a CPW semicircle.

¹¹Work in progress.

¹²Daniel T. Sank (private communication).

The distance between the semicircle and the wire outer conductor is designed to be as short as possible, $\sim 100 \mu\text{m}$.




Figure 4.14: Signal crosstalk. (a) Transmission and crosstalk coefficients for the Ag sample shown in the inset. The numbers adjacent to the pads in the inset correspond to the device ports. Reciprocal and reflection S-parameters are not shown. (b) Microwave simulation of the same device. The origin of the peaks at approximately 7 GHz is explained in the text.

The chip employed for the crosstalk tests is similar to the Ag sample used for the quantum socket microwave characterization and was part of a DUT analogous to that shown in Fig. 4.8. The DC resistances of the center trace of the 1 – 2 and 3 – 4 transmission lines were measured and found to be $\sim 2.8 \Omega$ and $\sim 4.5 \Omega$, respectively (note that the 3 – 4 transmission line is $\sim 18.0 \text{ mm}$ long, hence, the larger resistance). All DC resistances to ground and between the two transmission lines were found to be on the order of a few kilohms, demonstrating the absence of undesired short circuit paths. A four-port calibration and measurement of the DUT were conducted by means of a VNA (cf. Supplemental Material of Ref. [220] for details). Among the 16 S-parameters, Fig. 4.14 (a) shows the magnitude of the transmission coefficients S_{21} and S_{43} , along with the magnitude of the crosstalk coefficients S_{31} , S_{41} , S_{32} , and S_{42} .

The results show that the isolation in the typical qubit operation bandwidth, between 4 GHz and 8 GHz, is larger than $\sim 45 \text{ dB}$. Note that the crosstalk coefficients shown in Fig. 4.14 (a) include attenuation owing to the series resistance of the Ag

transmission lines. The actual isolation, due only to spurious coupling, would thus be smaller by a few decibels.

Figure 4.14 (b) shows the microwave simulations of the crosstalk coefficients, which agree reasonably well with the experimental results. These simulations are based on the models explained in Sec. 4.2.4. From simulations, we believe the isolation is limited by the crosstalk between the CPW transmission lines, instead of the three-dimensional wires. Note that the peaks at approximately 7 GHz correspond to an enhanced crosstalk due to a box mode in the microwave package. The peaks appear in the simulations, which are made for a highly conductive package, and may appear in measurements performed below ~ 1 K, when the Al package becomes superconductive. For the room temperature measurements shown in Fig. 4.14 (a), these peaks are smeared out due to the highly lossy Al package.

4.5 Applications to Superconducting Resonators

Thus far, we have shown a detailed characterization of the quantum socket in DC and at microwave frequencies, both at room temperature and at 77 K. In order to demonstrate the quantum socket operation in a realistic quantum computing scenario, we used a socket to wire a set of superconducting CPW resonators cooled to approximately 10 mK in a DR. We were able to show an excellent performance in the frequency range from 4 GHz to 8 GHz, which is the bandwidth of our measurement apparatus. Multiple chips were measured over multiple cycles using the same quantum socket; this demonstrates the high level of repeatability of our wiring method. We measured five Al on Si samples, as well as one Al on gallium-arsenide (GaAs) sample [245] (data not shown) and one Al on sapphire sample. The Al on sapphire device, in particular, featured a few resonators with quality factor comparable to the state-of-the-art in the literature [105], both at high and low excitation power.

The experimental setup is described in the Supplemental Material of Ref. [220] and shown in Fig. S1. Figure 4.15 shows a macro photograph of a 15 mm \times 15 mm chip housed in the sample holder; the chip is the Al on Si sample described in Sec. 4.3.4, with geometrical and DC electrical parameters reported in Table 4.2. The sample comprises a set of three CPW transmission lines, each connecting a pair of three-dimensional wire pads; multiple shunted CPW resonators are coupled to each transmission line. In this section, we will focus only on transmission line three and its five resonators. The transmission line has a center conductor width of 15 μm and gap width of 9 μm , resulting in a characteristic impedance of approximately 50 Ω . The resonators are $\lambda/4$ -wave resonators, each characterized by a center conductor of width W and a dielectric gap of width G . The open end of the resonators runs parallel to the transmission line for a length ℓ_{κ} , providing a capacitive coupling; a 5 μm ground section separates the

gaps of the transmission line and resonators (cf. Fig. S3 in the Supplemental Material of Ref. [220]). The nominal resonance frequency \tilde{f}_0 as well as all the other resonator parameters are reported in Table 4.3.

A typical DR experiment employing the quantum socket consists of the following steps. First, the chip is mounted in the microwave package, which has already been attached to the package holder (cf. Sec. 4.2.3 and Sec. 4.3). Second, a series of DC tests is performed at room temperature. The results for a few Al on Si samples are reported in Table 4.2. Third, the package holder assembly is characterized at room temperature




Figure 4.15: Macro photograph of an Al chip on Si substrate mounted in a sample holder with grounding washer. The image shows three CPW transmission lines each coupled to a set of $\lambda/4$ -wave resonators. The grounding washer, with its four protruding feet, is placed above the chip covering the chip edges. The marks imprinted by the bottom interface of the three-dimensional wires on the Al pads are noticeable. More detailed images of these marks are shown in Fig. 4.5. This chip and similar other chips with analogous microwave structures and geometries, including one Al on GaAs sample as well as one Al on sapphire sample, were used in the measurements at ~ 10 mK in the DR.

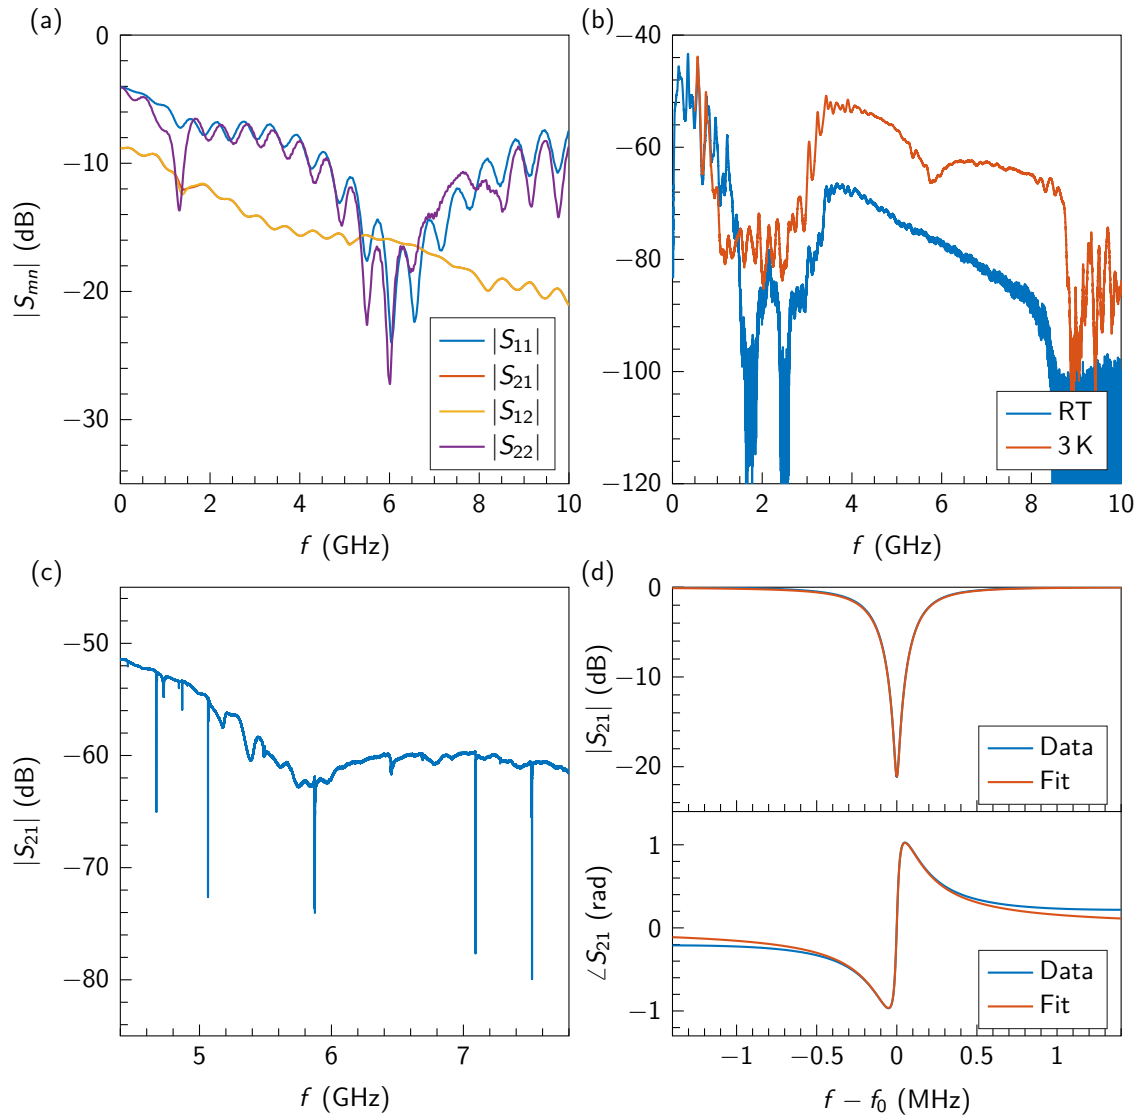


Figure 4.16: Measurements of Al on Si resonators. (a) Benchtop measurement of the S-parameters of the CPW transmission line three conducted at room temperature. (b) $|S_{21}|$ measurement of the same line with the chip mounted on the MC stage of the DR at room temperature (blue) and ~ 3 K (red). (c) $|S_{21}|$ measurement of the sample at approximately 10 mK. The five dips correspond to $\lambda/4$ -wave resonators. (d) Magnitude and phase of S_{21} for resonator 2.

by measuring its S-parameters. The results of such a measurement are shown in Fig. 4.16 (a). Fourth, the package holder is mounted by means of the SMP connectors to




Figure 4.17: Measurement of Q_i as a function of $\langle n_{\text{ph}} \rangle$ for one of the Al on sapphire resonators. The experimental settings used in the measurements are reported in the Supplemental Material of Ref. [220]; the confidence intervals were evaluated from the standard errors of the fitting parameters of the normalized inverse transmission coefficient \tilde{S}_{21}^{-1} and are smaller than the blue diamond symbols. All measurements were performed at ~ 10 mK. The typical quality factor “S-curve” is observed; the plateaus on the leftmost and rightmost regions of the curve indicate the reaching of low values of $\langle n_{\text{ph}} \rangle$ and saturation of two-level systems, respectively [105].

the MC stage of the DR and an S_{21} measurement is performed. The results (magnitude only) are shown in Fig. 4.16 (b) in the frequency range between 10 MHz and 10 GHz. Fifth, the various magnetic and radiation shields of the DR are closed and the DR is cooled down. Sixth, during cooldown the S_{21} measurement is repeated first at ~ 3 K and, then, at the DR base temperature of approximately 10 mK. The results are shown in Fig. 4.16 (b) and (c), respectively. At ~ 3 K we note the appearance of a shallow dip at approximately 5.7 GHz, probably due to a screw-in micro connector becoming slightly loose while cooling (cf. Sec. 4.4.2). It is important to mention that in the next generation of three-dimensional wires we will eliminate the screw-in micro connector, since we believe we found a technique to overcome the soldering issues detailed in Sec. 4.2.3 (cf. Sec. 4.6 for a brief description). At the base temperature, all five resonators are clearly distinguishable as sharp dips on the relatively flat microwave background of the measurement network. We then select a narrower frequency range around each resonator and make a finer S_{21} measurement. For example, Fig. 4.16 (d) shows the magnitude and phase of the resonance dip associated with resonator number 2.

The normalized inverse transmission coefficient \tilde{S}_{21}^{-1} was fitted as in Ref. [105]. This procedure makes it possible to accurately estimate both the internal Q_i and the rescaled coupling Q_c^* quality factors of a resonator. The fit results are shown in Table 4.3. The plot of the fits for the magnitude and phase of S_{21} for resonator 2 are overlaid with the measured data in Fig. 4.16 (d). The real and imaginary parts of \tilde{S}_{21}^{-1} for the same resonator, as well as the associated fit, are shown in Fig. S4 in the Supplemental Material of Ref. [220].

The resonator mean photon number $\langle n_{\text{ph}} \rangle$ can be estimated from the room temperature power at the input channel P_{in} and the knowledge of the total input channel attenuation α (cf. Supplemental Material of Ref. [220]). From basic circuit theory and Ref. [113], we obtain

$$\langle n_{\text{ph}} \rangle = \frac{2}{h\pi^2} \frac{Q_1^2}{Q_c^*} \frac{P'_{\text{in}}}{\tilde{f}_0^2}, \quad (4.9)$$

where h is the Planck constant, $1/Q_1 = 1/Q_i + 1/Q_c^*$ is the inverse loaded quality factor of the resonator, and $P'_{\text{in}} = P_{\text{in}}/\alpha$ is the power at the resonator input. For example, $\langle n_{\text{ph}} \rangle \simeq 4.1 \times 10^7$ for resonator 2.

The fabrication process of the resonators described in Table 4.3 was not optimized for high values of Q_i , which, however, is an important figure of merit for applications to quantum computing. In order to verify the compatibility of the quantum socket with resonators of higher quality, we decided to fabricate a sample featuring an Al thin film deposited by means of a ultra-high vacuum electron beam physical vapor deposition (EBPVD) system; the substrate of choice was, in this case, sapphire. The sample design is similar to that shown in Fig. 4.15 and the sample preparation analogous

Table 4.3: Resonator parameters. The measured resonance frequency is f_0 . The rescaled coupling and internal quality factors Q_c^* and Q_i , respectively, are obtained from the fits of the measured transmission coefficients (cf. text for details). These quality factors were measured at a high resonator excitation power, corresponding to $\langle n_{\text{ph}} \rangle > 10^5$.

i (-)	\tilde{f}_0 (MHz)	f_0 (MHz)	W (μm)	G (μm)	ℓ_κ (μm)	Q_c^* (-)	Q_i (-)
1	4600.0	4673.2	8	5	400	5012	21243
2	5000.0	5064.5	15	9	300	16002	165790
3	5800.0	5872.9	25	15	400	10269	47165
4	7000.0	7091.7	15	9	300	6230	54894
5	7400.0	7520.1	8	5	400	4173	28353
6	4700.0	4717.6	15	9	45	244960	1977551

to that in Ref. [105]; fabrication details are in Appendix B.6. We were able to measure a few resonators with $Q_i > 10^6$ for large values of $\langle n_{\text{ph}} \rangle$. For one of these resonators, we measured Q_i as a function of $\langle n_{\text{ph}} \rangle$, as shown in Fig. 4.17. As expected from literature measurements [105], Q_i decreased by approximately one order of magnitude when the resonator mean photon number was reduced from $\langle n_{\text{ph}} \rangle \simeq 10^6$ to $\simeq 10^{-2}$. For the lowest mean photon number, $Q_i \simeq 2.8 \times 10^5$; such a quality factor is a good indication that the quantum socket will likely preserve quantum coherence sufficiently well when utilized for the manipulation of superconducting qubits.

4.6 Conclusions

In a recent work [91], seven sequential stages necessary to the development of a quantum computer were introduced. At this time, the next stage to be reached is the implementation of a single logical qubit characterized by an error rate that is at least one order of magnitude lower than that of the underlying physical qubits. In order to achieve this task, a two-dimensional lattice of 10×10 physical qubits with an error rate of at most 10^{-3} is required [59].

Figure 4.18 shows an extensible quantum computing architecture where a two-dimensional square lattice of superconducting qubits is wired by means of a quantum socket analogous to that introduced in this chapter. The architecture comprises three main layers: The quantum hardware; the shielding interlayer; the three-dimensional wiring mesh.

As shown in Fig. 4.18 (a), the quantum hardware is realized as a two-dimensional lattice of superconducting qubits with nearest neighbor interactions. The qubits are a modified version of the Xmon presented in Ref. [92]. Each qubit is characterized by seven arms for coupling to one XY and one Z control line as well as one measurement resonator and four inter-qubit coupling resonators. We name this type of qubit the *heptaton*. The inter-qubit coupling is mediated by means of superconducting CPW resonators that allow the implementation of control Z (CZ) gates between two neighboring qubits [148, 151]. A set of four heptatons can be readout by way of a single CPW transmission line connected to four CPW resonators, each with a different resonant frequency. Figure 4.18 also shows the on-chip pads associated with each three-dimensional wire. In the Supplemental Material of Ref. [220], we propose an extended architecture where each qubit can be measured by means of two different resonators, one with frequency above and the other with frequency below all coupling resonator frequencies.

Assuming a pitch between two adjacent three-dimensional wires of 1 mm, the lateral dimension of one square cell having four heptatons at its edges is 8 mm. The three distances A , B , and C between wire pads and resonators leading to this quantity

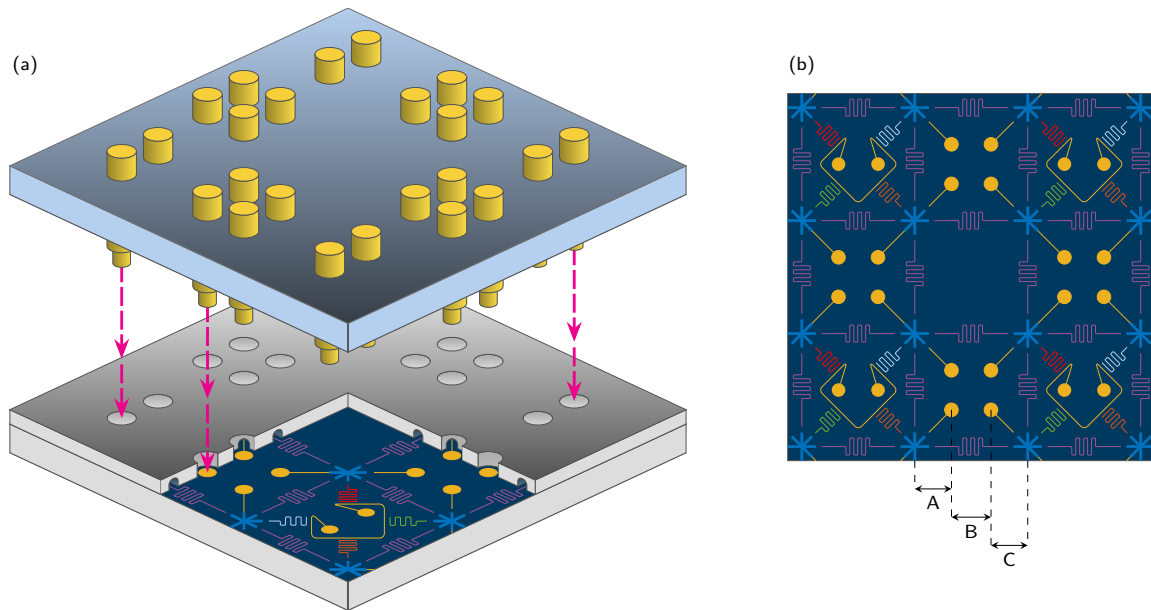


Figure 4.18: Extensible quantum computing architecture. (a) Main three layers of the architecture: The quantum hardware (bottom); the shielding interlayer (middle); the three-dimensional wiring mesh (top), with wires indicated in yellow. The vertical magenta dashed lines with double arrows show the mounting procedure to be used to prepare the assembly. The middle layer (thinner) will be metalized on the bottom and wafer bonded to the quantum layer beneath; both the thru-holes that accommodate the three-dimensional wires and the tunnels above the qubits and manipulation lines are shown. The back end of the wires (top) will be connected to SMPS connectors (not shown). (b) Two-dimensional view of the quantum hardware. The substrate is indicated in dark blue, the heptatons in light blue, the coupling resonators in magenta, the four readout resonators in red, green, orange, and cyan, and, finally, the wire pads and associated lines in yellow. The distances between coupling resonators and wire pads are $A = C = 2.25$ mm and $B = 3.5$ mm. Note that the Z control lines are represented as galvanically connected to the heptatons, similar to Ref. [92]. The measurement can be multiplexed so that four qubits are readout by one line only.

are indicated in Fig. 4.18 (b). It is thus possible to construct a two-dimensional lattice of 10×10 heptatons on a square chip with lateral dimension 9×8 mm = 72 mm. A 72 mm \times 72 mm square chip is the largest chip that can be diced from a standard 4 inch wafer. This will allow the implementation of a logical qubit based on the surface code, with at least distance five [59]. In this architecture, the coupling resonators act as a *coherent spacer* between pairs of qubits, i.e., they allow a sufficient separation to

accommodate the three-dimensional wires, while maintaining quantum coherence during the CZ gates. Additionally, these resonators will help mitigate qubit crosstalk compared to architectures based on direct capacitive coupling between adjacent qubits (cf. Ref. [108]). In fact, they will suppress qubit-mediated coupling between neighboring control lines (similar coupling mechanism as in Ref. [130]). It is worth noting that adjacent coupling resonators can be suitably designed to be at different frequencies, thus further diminishing qubit-mediated crosstalk.

Implementing a large qubit chip with a lateral dimension of 72 mm presents significant challenges to the qubit operation at microwave frequencies. A large chip must be housed in a large microwave package, causing the appearance of box modes that can interfere with the qubit control and measurement sequences [233]. Moreover, a large chip will inevitably lead to floating ground planes that can generate unwanted slotline modes [233]. All these parasitic effects can be suppressed by means of the shielding interlayer, as shown in Fig. 4.18 (a). This layer can be wafer bonded [222, 223, 224, 225] to the quantum layer. Thru-holes and cavities on the bottom part of the layer can be readily fabricated using standard Si etching techniques. The holes will house the three-dimensional wires, whereas the cavities will accommodate the underlying qubit and resonator structures. Large substrates also generate chip modes that, however, can be mitigated using buried metal layers and/or metalized through-silicon vias [243].

The three-dimensional wires to be used for the 10×10 qubit architecture will be an upgraded version of the wires used in this chapter. In particular, the M2.5 thread will be removed and the wires will be inserted in a dedicated substrate (cf. Fig. 4.18 (a)); additionally, the screw-in micro connector will be substituted by a direct connection to a subminiature push-on sub-micro (SMPS) connector (not shown in the figure).

In future applications of the quantum socket, we envision an architecture where the three-dimensional wires will be used as interconnect between the quantum layer and a classical control/measurement layer. The classical layer could be realized using RSFQ digital circuitry [246, 247]. For example, high-sensitivity digital down-converters (DDCs) have been fabricated based on RSFQ electronics [248]. Such circuitry is operated at very low temperatures and can substitute the room temperature electronics used for qubit readout. Note that cryogenic DDC chips with dimensions of less than $5 \text{ mm} \times 5 \text{ mm}$ can perform the same operations presently carried out by room temperature microwave equipment with an overall footprint of $\sim 50 \text{ cm} \times 50 \text{ cm}$. Recent interest in reducing dissipation in RSFQ electronics¹³ will possibly enable the operation of the classical electronics in close proximity to the quantum hardware. We also believe it is feasible to further miniaturize the three-dimensional wires so that the wire outer diameter would be on the order of $500 \text{ }\mu\text{m}$. Assuming a wire-wire pitch also of $500 \text{ }\mu\text{m}$, it will therefore be possible to realize a lattice of 250000 wires connecting to $\sim 10^5$ qubits arranged

¹³Confer <https://www.iarpa.gov/index.php/research-programs/c3>.

on a 315×315 two-dimensional qubit grid with dimensions of $1 \text{ m} \times 1 \text{ m}$. This will allow the implementation of simple fault-tolerant operations between a few tens of logical qubits [59].

Resonant Coupling Parameter Estimation

Today's quantum computers are composed of tens of qubits interacting with each other and the environment in increasingly complex networks. In order to achieve the best possible performance when operating such systems, it is necessary to have accurate knowledge of all parameters in the quantum computer Hamiltonian. In this chapter¹, we demonstrate theoretically and experimentally a method to efficiently learn the parameters of resonant interactions for quantum computers consisting of frequency-tunable superconducting qubits. Such interactions include, for example, those to other qubits, resonators, two-level systems, or other wanted or unwanted modes. Our method is based on a significantly improved swap spectroscopy calibration and consists of an *offline* data collection algorithm, followed by an *online* Bayesian learning algorithm. The purpose of the offline algorithm is to detect and coarsely estimate resonant interactions from a state of zero knowledge. It produces a square-root reduction in the scaling of the number of measurements. The online algorithm subsequently refines the estimate of the parameters to comparable accuracy as traditional swap spectroscopy calibration, but in constant time. We perform an experiment implementing our technique with a superconducting qubit. By combining both algorithms, we observe a reduction of the calibration time by one order of magnitude. Our method will improve present medium-scale superconducting quantum computers and will also scale up to larger systems. Finally, the two algorithms presented here can be readily adopted by communities working on different physical implementations of quantum computing architectures.

¹This chapter was largely adapted from Ref. [249]. The list of author contributions can be found in the Statement of Contributions within the front matter of this thesis. Published by the American Physical Society under the terms of the [Creative Commons Attribution 4.0 International](https://creativecommons.org/licenses/by/4.0/) license.

5.1 Introduction

As quantum computing systems become larger and more complex, the high-fidelity operation of those devices is accompanied by the daunting task of calibrating numerous physical qubits. In particular, calibrating tunable qubits requires the estimation of resonant interaction parameters, i.e., the frequency and coupling strength, of both wanted and unwanted resonances. Wanted resonances include those between pairs of interacting qubits [149, 150], qubits and resonators [151], and qubits and tunable couplers [250, 251]. In this case, knowledge of the interaction parameters allows for the implementation of two-qubit gates and readout. Unwanted resonance are mainly those between qubits and two-level systems (TLSs) [252, 114, 253] as well as substrate and box modes [172]. Microwave control crosstalk may also behave as an unwanted resonance, causing Rabi oscillations at the frequency of a nearby qubit. Such unwanted resonances must be avoided through proper frequency spacing. TLSs, especially, are a pervasive source of errors in superconducting architectures that must be remediated [254]. In summary, learning all the interaction parameters allows for a comprehensive calibration and, thereby, minimization of coherent and incoherent errors.

In this chapter, we study theoretically and demonstrate experimentally a data-efficient and automated method for identifying and estimating the parameters of resonant interactions based on *swap spectroscopy* [255, 151]. We realize swap spectroscopy by performing energy relaxation time T_1 measurements of a frequency-tunable Xmon transmon qubit [92] at different qubit frequencies. The identification and estimation method is divided into two parts: an *offline* data collection algorithm [256] and an *online* Bayesian learning algorithm [257, 258]. Both algorithms are based on the dynamics of interacting quantum systems. The former is used from a state of zero knowledge about a particular frequency range to coarsely identify resonance parameters within that range. The latter focuses on improving the estimate of those parameters. In this context, the term “online” means that measurements taken during the execution of the algorithm inform the subsequent ones. For the “offline” method, the execution of the entire algorithm is predetermined. It is worth noting that the offline algorithm is applicable to any of the resonance types listed above. The online algorithm can be used only for coherent resonances, either wanted or unwanted.

By means of our parameter-estimation method, we can shorten the calibration time of an Xmon transmon qubit significantly. The offline data collection algorithm makes it possible to reduce the scaling of the number of measurements by a square-root when compared to a traditional swap spectroscopy calibration. In our experiment, this algorithm takes ≈ 30 min to detect resonances in a 1 GHz bandwidth: one order of magnitude less time than with traditional methods. The online Bayesian learning algorithm runs in ≈ 25 s per resonance, bringing the estimation accuracy to the same

level as high-resolution traditional swap spectroscopy.

In order to test our method and compare it against traditional swap spectroscopy, we characterize the resonances within a 1 GHz bandwidth of an Xmon transmon qubit. We additionally synthesize two resonance modes that emulate the interaction with another qubit, resonator, or TLS. These modes are created by applying a coherent drive with a microwave source to the qubit under test. The synthesized resonance mode is a convenient and flexible tool to test our method since we can arbitrarily change its *resonance frequency* by tuning the source frequency as well as its *coupling strength* by changing the emitted source power.

Our method is not confined to the realm of superconducting quantum computing. In fact, it can easily be adopted by practitioners working on different physical implementations of quantum computing architectures such as trapped ions and semiconductor qubits [29].

This chapter is organized as follows: In Sec. 5.2, we explain qubit calibration in frequency-tunable architectures. In Sec. 5.3, we summarize the working principle of traditional swap spectroscopy, explaining why it is inefficient for the purpose of detecting resonances. In Sec. 5.4, we introduce the offline octave sampling algorithm (Sec. 5.4.1) and demonstrate its experimental implementation and associated data analysis procedure to detect interactions between an Xmon transmon qubit and four resonance modes (RMs), including an incoherent one (Sec. 5.4.2). In Sec. 5.5, we explain the online Bayesian learning algorithm (Sec. 5.5.1) and demonstrate its performance at finding an accurate estimate for the parameters of a resonance (Sec. 5.5.2). In Sec. 5.6, we discuss additional concerns with the algorithms and the relevance of our methods for quantum computing. Finally, in Sec. 5.7, we provide an outlook and conclusions.

5.2 Qubit Calibration in Frequency-Tunable Architectures

A fundamental requirement to the operation of a quantum computer is the proper calibration of the physical qubits in the system. This calibration includes many specific operations. One of the most basic tasks, for example, is to run a Rabi experiment on each qubit. This allows the determination of some experimental parameters needed to set up, e.g., a π -pulse and perform a measurement. Once this first task is realized, further experiments can refine the knowledge of the pulse amplitude, rotation axis, measurement parameters, etc. Finally, a full calibration requires knowing the precise parameters of the system Hamiltonian and the interaction to the environment, allowing for the systematic optimization of the fidelity of one- and two-qubit gates as well as measurement.

In a frequency-tunable superconducting qubit architecture such as the Google architecture [96] or the one used in this work, an additional degree of freedom must be considered during calibration: the qubit frequency f_q . Xmon transmon qubits are one example of tunable qubits [92]. In this design, an on-chip capacitive island made from aluminum (Al) is coupled in parallel to a superconducting quantum interference device (SQUID) comprised of two Josephson tunnel junctions in parallel, forming a superconducting loop [259]. An Xmon transmon qubit is a quantum *anharmonic* oscillator, characterized by a non-equally spaced ladder of quantum states. The frequency (i.e., energy) difference f_q between the ground state $|g\rangle$ and first excited state $|e\rangle$ differs from that between $|e\rangle$ and the second excited state $|f\rangle$ by the so-called qubit anharmonicity α [104]. The qubit transition frequency f_q is controlled *in situ* by applying a local external magnetic flux that threads the SQUID, tuning the Josephson energy E_J and therefore the level separation.

Frequency tunability leads to a few distinct advantages to the operation of a quantum computer. For instance, it allows for adjustable qubit-qubit interactions because the effective coupling strength between two qubits depends on the frequency difference between them. This enables the implementation of several types of two-qubit gates such as the controlled-phase (CPHASE) gate that takes advantage of state $|f\rangle$ as an auxiliary state [148, 149, 150, 151], as well as the \sqrt{i} SWAP and i SWAP gates [151]. In addition, setting the frequency of spatially neighboring qubits away from each other helps avoid control crosstalk and frequency crowding issues, the latter being endemic in fixed-frequency systems [100].

Another advantage inherent to frequency-tunable architectures is related to energy relaxation. On-chip superconducting qubits interact with a distribution of TLSs, which are present in the various amorphous dielectric materials surrounding the qubit metallic structures (e.g., Si and Al oxides). While the microscopic origin of TLSs is still under debate [252], their effect on the qubit leads to either a T_1 reduction or T_1 and f_q time fluctuations (see Chapter 6). In particular, the T_1 reduction is caused by the coherent or incoherent exchange of energy between a qubit and semi-resonant TLSs. The ability to set the frequency of a qubit away from that of TLSs is therefore desirable and realizable only with tunability.

Calibrating qubits to implement two-qubit gates or to avoid TLSs is a *parameter estimation* problem. We need to determine the Hamiltonian parameters that define the resonant interactions between a qubit and another system. In all the aforementioned cases, two parameters must be found: the resonance frequency and coupling strength of the interaction.

Historically, swap spectroscopy has been a prominent method to perform this kind of calibration. Unfortunately, traditional swap spectroscopy is inefficient in the amount of data it requires and therefore slow. This is inconvenient for multiple reasons. First, as the number of qubits in a system grows, so does the number of calibrations that must

be performed. This is particularly relevant to qubit-qubit coupling calibration, which cannot be performed in parallel on all qubits. Second, TLSs in the environment are known to fluctuate over time [254, 260, 261, 253, 262]. Similarly, f_q itself can shift in time. The identification of resonant interactions must therefore be repeated at regular intervals. We thus require a robust, accurate, and time-efficient method to identify the parameters associated with resonant interactions.

5.3 Traditional Swap Spectroscopy

Swap spectroscopy is an experimental method that allows exploring the environment of a qubit at various frequencies by using the qubit itself as a probe. Traditionally, swap spectroscopy has been used to select the operating frequency of qubits, making it possible to avoid TLSs or regions of low T_1 . Additionally, it has been used to explore resonant interactions, such as those with other qubits [150] or resonators [255]. Performing swap spectroscopy requires a minimally calibrated qubit and, thus, is suitable as a tune-up experiment.

In a swap spectroscopy experiment the qubit is initialized at the so-called *idle frequency*. A π -pulse is then applied to the qubit, energizing it from $|g\rangle$ to $|e\rangle$. At the end of the π -pulse, a flux pulse is applied to the SQUID in order to tune the qubit to a different frequency, the *probe frequency* $f_p(A)$, where A is the pulse amplitude. This procedure requires knowledge of the correspondence between f_q and A , which can be calibrated via regular pulse spectroscopy (see Appendix C.1). After a time t , the flux pulse is turned off and the qubit is measured back at the idle frequency. This pulse sequence is illustrated in Fig. 5.1. Note that using a flux pulse to set f_p presents advantages over quasi-statically changing the idle qubit frequency by means of a DC current to the SQUID. Namely, it avoids recalibrating the π -pulse and measurement pulse at each qubit frequency.

In a traditional experiment, t and f_p are swept linearly over a desired range and the




Figure 5.1: Pulse sequence for a swap spectroscopy experiment. The initial π -pulse (red) excites the qubit, which is initialized at the so-called idle frequency. The flux pulse (blue) changes the qubit transition frequency f_q to the probe frequency $f_p(A)$ for a duration t . The qubit is then measured (green) after being set back to its idle frequency.




Figure 5.2: Swap spectra for two frequency ranges. The x -axis shows the probe frequency of the qubit, which is set by the amplitude of the flux pulse applied to the SQUID. The y -axis indicates the length of the flux pulse before measurement and, therefore, corresponds to the interaction time with potential resonance modes. (a) Distinct features are visible in the full spectrum, including two chevron patterns around 4.8 GHz and one at 5.1 GHz. The resonance at 5.1 GHz looks aliased because of the low-resolution sampling of the time axis. A low T_1 streak is also visible at 4.35 GHz, likely caused by an incoherent TLS. (b) Zoom in the region with the slower coherent chevron patterns caused by synthesized resonance modes. For this experiment, the synthesizers were set at 4.8100 and 4.8314 GHz. A wide frequency scan is needed to see if and where there are resonance modes, as in (a), but a more detailed experiment, as in (b), is needed to properly estimate the resonance parameters.

qubit is measured at each point, recording how many measurement shots correspond to an excited or ground state, n_e or n_g , respectively. For the experiments of this work, each pulse sequence is repeated to make 786 high-power single-shot measurements (see Section 3.3.3). We estimate a readout visibility $\gtrsim 90\%$. As a result of measuring the qubit in the energy basis, swap spectroscopy is insensitive to dephasing.

Figure 5.2 (a) shows the result of a typical swap spectroscopy experiment, a *swap spectrum*, with data taken between 4.146 and 5.170 GHz and for times up to 500 ns. Resonant couplings appear as oscillations, or chevron patterns, of the measured average population $P_e = n_e / (n_e + n_g)$ in time. For example, on the far right of the spectrum it is possible to observe very fast oscillations, corresponding to a strong coupling

of $g \approx 40$ MHz between the qubit and the measurement resonator (see Chapter 3 for details on the sample layout and experimental setup). To the left of the resonator we observe a slower oscillation corresponding to a weaker interaction between the qubit and synthesized resonance mode. Finally, at an even lower frequency, around 4.35 GHz, we observe a “streaky” structure. In this region, the qubit excitation is lost faster than elsewhere, and we cannot observe any oscillation. This structure is caused by an incoherent resonance, most likely a TLS.

The features observed in Fig. 5.2 (a) demonstrate a selection of possible resonant interactions: strong interactions, where $g \gg 1/T_1$ resulting in multiple coherent oscillation cycles and weak interactions appearing as regions of lower T_1 . Neither of them is ideal for the operation of a qubit. In the case shown in Fig. 5.2, the best choice for the qubit idle frequency is around 4.6 GHz; far away from any unwanted interactions.

The data in Fig. 5.2 (a) gives us a rough idea about the parameters of any possible resonance modes coupled to the qubit within the measured spectrum. It is hard to tell, however, that there are in fact *two* resonance modes at 4.8 GHz, or what the frequency of the oscillation for the resonator is. A more detailed scan, such as the one in Fig. 5.2 (b), might be necessary to estimate the parameters with sufficient accuracy. Traditional swap spectroscopy, with data taken in a linear grid, is a possible method to detect and estimate resonance modes. We show in Sec. 5.4, however, that it is inefficient, and that there exists a much better way to perform this task: octave sampling.

5.4 Offline Octave Sampling

The offline octave sampling algorithm has a similar objective as swap spectroscopy, i.e., to determine if there are any systems interacting resonantly with the qubit and provide an estimate for their coupling parameters. However, we want to achieve this purpose in a more efficient fashion by acquiring less data, therefore saving valuable experimental time. Note that the pulse sequence employed to perform octave sampling is the same as for swap spectroscopy (see Fig. 5.1). The difference lies in how the spectrum is sampled. Whereas traditional swap spectroscopy samples the frequency-time space in a regular grid, octave sampling takes advantage of resonant dynamics to acquire as few data as possible.

5.4.1 Theoretical Method

In order to explain the data collection strategy, we analyze the time dynamics of the systems at play. Since we are searching for resonant interactions with a qubit, we work in a single-excitation manifold ($|g\rangle \leftrightarrow |e\rangle$). Thus, even if a system is characterized by

more than two energy levels (e.g., a resonator), we can still treat it as a two-level system because higher levels are never occupied. This is our working assumption throughout the rest of the chapter.

Note that we can probe the environment of an anharmonic oscillator, e.g., a transmon, within a different single-excitation manifold. In that case, we can populate the second excited state and look for systems coupled to the $|e\rangle \leftrightarrow |f\rangle$ transition. This allows for the calibration of certain two-qubit CPHASE gates [149, 150, 151, 106]. In either case, because we consider the exchange of a single excitation, the effective Hamiltonian remains unchanged.

After a rotating wave approximation, the Hamiltonian of a qubit at the probe frequency f_p interacting with a resonance mode at a frequency f_{RM} reads

$$\hat{H} = \frac{hf_p}{2}\hat{\sigma}_{z,1} + \frac{hf_{RM}}{2}\hat{\sigma}_{z,2} + hg(\hat{\sigma}_1^+\hat{\sigma}_2^- + \hat{\sigma}_1^-\hat{\sigma}_2^+), \quad (5.1)$$

where g is the coupling strength of the qubit–resonance mode interaction, $\hat{\sigma}_{z,1(2)}$ are Pauli matrices for the qubit (1) and resonance mode (2), and $\hat{\sigma}_{1(2)}^+$ and $\hat{\sigma}_{1(2)}^-$ are raising and lowering operators for the qubit and resonance mode. We solve for the time evolution of the qubit when it is initialized in state $|e\rangle$ and with the resonance mode starting in $|g\rangle$. The theoretical probability of finding the qubit in the excited state after a time t is then given by

$$\tilde{P}_e(t) = 1 - \left(\frac{2g}{\Omega}\right)^2 \sin^2(2\pi\Omega t/2), \quad (5.2)$$

where $\Omega^2 = \delta f^2 + 4g^2$, with $\delta f = f_p - f_{RM}$. The probability \tilde{P}_e of Eq. (5.2) is plotted in Fig. 5.3 (a) as contours. Close to resonance, the excitation swaps between the qubit and the resonance mode with frequency Ω increasing at larger δf , resulting in the familiar chevron pattern. Both the width of the pattern, which we quantify by the full-width half maximum of the amplitude, $4g$, and Ω depend on g . Crucially, the width is proportional to g , while the period of the oscillation and the position in time of the first minimum is proportional to $1/g$. The goal is to detect a resonance mode by finding the *first minimum* of an oscillation, where $\tilde{P}_e \sim 0$ because the excitation has swapped into the resonance mode.

With these observations in mind, we choose to divide the frequency-time space into rows of *bins* within which we take a constant number n_s of swap spectroscopy measurements. Instead of naively sampling the spectrum in a uniform grid, we adapt the measurement based on the value of g that we are trying to detect. The coupling strength determines the *time* t at which we measure and the *bin size*. On the one hand, a resonance mode with large coupling strength g has a large width and a short period. For short time t , then, we choose the bins to be wide and short [see the bottom rows of

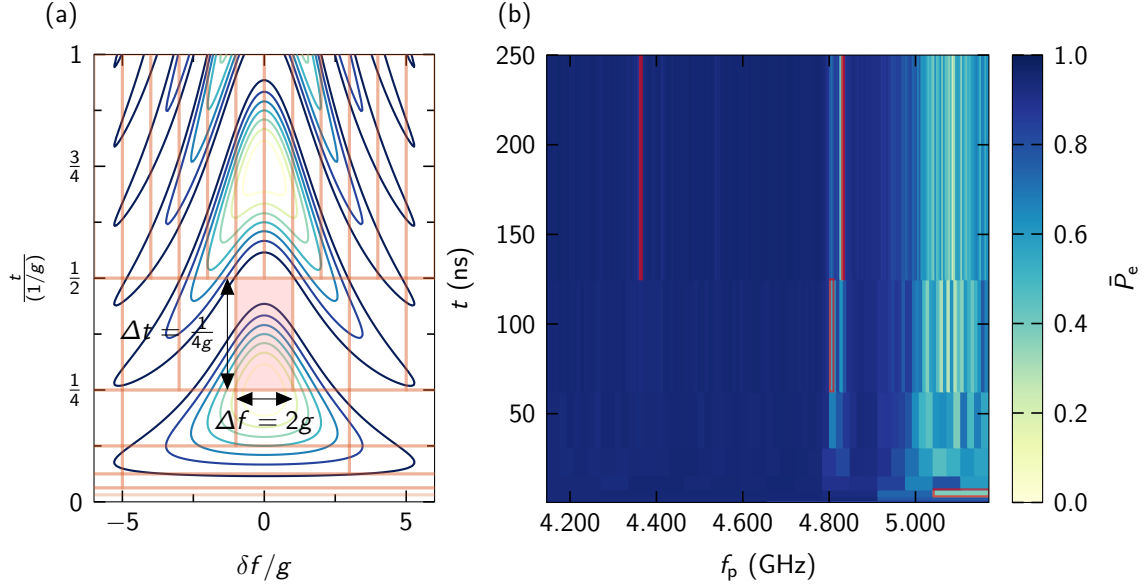


Figure 5.3: Offline octave sampling. (a) Contour plot of the probability of finding the qubit in $|e\rangle$ [see Eq. (5.2)] as a function of time t and frequency detuning Δf ; both axes are normalized by the coupling strength g . The highlighted box in the center indicates which portion of the chevron pattern is meant to be detected by the algorithm. (b) Swap spectroscopy experiment with octave sampling. Starting from a bin spanning the full measurement range, at each subsequent octave the bin width is halved and the bin length is doubled. The color of each bin represents the average value of the measured P_e over $n_s = 5$ samples. The red boxes indicate the resonances reported by the analysis explained in Sec. 5.4.2.

bins in Fig. 5.3 (a)]. On the other hand, a more weakly coupled resonance mode appears later in time, with a narrower frequency width and a longer period. In this case, the bins are longer and narrower [see the top rows of bins in Fig. 5.3 (a)]. In order to cover the entire measurement space, all bins must be adjacent (without overlapping). This condition constrains the ratio of number of bins in consecutive rows to be an integer. We choose this integer to be two because it is the only ratio for which a bin containing the first oscillation minimum [at $t = 1/4g$] does not contain any other later minima. For example, a factor of three would contain both the first and second [at $t = 3/4g$] oscillation minima, as can be deduced from Fig. 5.3 (a). We refer to this method as *octave* sampling because consecutive bin rows are suited to detect resonances with a coupling strength ratio of two.

In order to make this bin division systematic, we introduce the concept of a coupling octave with coupling strength g_m , where m is the octave number ranging from 0 to

m_f [256]. The final octave number m_f is determined by the desired frequency or time resolution, as discussed below. For each octave, the full frequency spectrum to be analyzed ranges between a minimum and maximum frequency f_{\min} and f_{\max} . This range is divided into 2^m bins of equal size, with frequency width $\Delta f = 2g_m$ and time length $\Delta t = 1/4g_m$. The location in time of the bins' lower edge is $t = 1/4g_m$. One such bin, with $g_m = g$, is highlighted in red in Fig. 5.3 (a). Note that the highlighted bin is not centered on the oscillation minimum. This is because a low- P_e measurement in that area corresponds to a *range* of possible coupling strengths, namely, those for which $g_m/2 \leq g \leq g_m$. The resonance mode plotted in Fig. 5.3 (a) is at the upper end of this range, and, hence, at the lower edge of the bin.

The execution of the algorithm is determined by the total bandwidth $B = f_{\max} - f_{\min}$, which is the frequency width of the single bin spanning the whole spectrum at the zeroth octave. This bandwidth corresponds to a coupling octave $g_0 = B/2$ and, therefore, to a time length $\Delta t = 1/4g_0$. For the next octave, we divide the width of the bins by two such that the subsequent row has twice as many bins as in the previous step. The length of the bins in time is correspondingly doubled. An example of this division is shown by the orange grid in Fig. 5.3 (a).

Thus, if we are given as inputs f_{\min} , f_{\max} , and m_f , the execution of the algorithm, starting from the zeroth octave $m = 0$, goes as such:

1. Divide the frequency range in 2^m bins, each with width $2g_m = B/2^m$.
2. Each bin spans the time values $1/4g_m \leq t \leq 1/2g_m$.
3. Take n_s swap spectroscopy samples within each of the 2^m bins, sampling uniformly at random in frequency and inverse time. That is, for bin $k = 1, \dots, 2^m$, draw f_p and t as

$$\begin{aligned} f_p &\sim \mathcal{U}(f_{\min} + 2(k-1)g_m, f_{\min} + 2kg_m), \\ t &\sim \mathcal{U}(2g_m, 4g_m)^{-1}, \end{aligned}$$

where the notation $X \sim \mathcal{U}(a, b)$ signifies that X is drawn randomly from a continuous uniform distribution \mathcal{U} between a and b .

4. Increment m and start over for the next octave.

The total number of bins to be measured, N_{bins} , depends both on the size of the bandwidth B and the final octave number m_f . To set m_f , we can choose either a maximum time $t_{m_f} = 1/2g_{m_f}$ or a final frequency resolution $\Delta f_{m_f} = 2g_{m_f}$, or, and perhaps most useful, a minimum coupling strength $g_{\min} = g_{m_f}/2$. Any of these quantities determine the number of bins for the final octave through the octave coupling g_{m_f} , and must be picked according to the goal of the experimenter. Then, following from the equation for

the bin width above, we find $m_f = \lceil \log_2(B/\Delta f_{m_f}) \rceil = \lceil \log_2(B/2g_{m_f}) \rceil$ (the result of the logarithm is rounded up). Accordingly, N_{bins} can be calculated by summing the number of bins per octave,

$$\begin{aligned} N_{\text{bins}} &= \sum_{m=0}^{m_f} 2^m \\ &= 2^{m_f+1} - 1 \\ &= \frac{2B}{\Delta f_{m_f}} - 1. \end{aligned} \tag{5.3}$$

The total number of points is thus $N_{\text{oct}} = n_s N_{\text{bins}}$.

In comparison, with the same frequency resolution, traditional swap spectroscopy divides the frequency axis in $B/\Delta f_{m_f}$ points, and the time axis in $t_{m_f}/\Delta t_{\text{min}}$ points, where Δt_{min} is the time resolution². While octave sampling reaches a time resolution of $1/B$, it would be unfair to the traditional method to use that number directly. Instead, we assume that $1/\Delta t_{\text{min}}$ is on the order of hundreds of megahertz, allowing the detection of strong couplings such as those to other qubits or resonators. The total number of points is then

$$\begin{aligned} N_{\text{trad}} &= \frac{t_{m_f}}{\Delta t_{\text{min}}} \frac{B}{\Delta f_{m_f}} \\ &= \frac{1}{\Delta t_{\text{min}}} \frac{B}{\Delta f_{m_f}^2}. \end{aligned} \tag{5.4}$$

The number of points scales as $\mathcal{O}(1/\Delta f_{m_f}^2)$ for the traditional method, whereas it scales as $\mathcal{O}(1/\Delta f_{m_f})$ for octave sampling: a *square-root* improvement.

As a last remark, we note that while the octave sampling strategy was designed according to the dynamics of *coherent* resonances, it is also well suited to handle *incoherent* resonances, such as the one at 4.35 GHz in Fig. 5.2 (a). Although an incoherent resonance does not show clear oscillations, it still increases the qubit relaxation rate. This relaxation appears as a low-excitation region in the spectrum, where frequency width and time position obey similar scaling rules as explained above. Such regions of low excitation can be detected just as well as oscillation minima.

5.4.2 Experimental Results

The result of an experimental implementation of the octave sampling algorithm is shown in Fig. 5.3 (b), which is the efficient version of Fig. 5.2 (a). Each bin is colored

²Although the concept of an octave number m does not apply to the traditional method, we use the value of Δf_{m_f} and t_{m_f} for comparison

according to the average excitation probability \bar{P}_e measured over $n_s = 5$ samples. We are able to discern the same features as in Fig. 5.2 (a), i.e., distinct low excitation regions, while acquiring much less data.

The frequency ranges from $f_{\min} = 4.146$ GHz to $f_{\max} = 5.170$ GHz such that $B = 1024$ MHz. We choose the final octave number to be $m_f = 8$, allowing us to detect TLSs with $g \geq 1$ MHz. This corresponds to $g_{m_f} = 2$ MHz and a minimum bin width of $\Delta f_{m_f} = 4$ MHz. Using Eq. (5.3), $N_{\text{bins}} = 511$ and $N_{\text{oct}} = 2555$.

Referring to Eq. (5.4) and given the resolution to detect oscillations up to $1/\Delta t_{\min} = 200$ MHz (corresponding to $g = 100$ MHz) in a $t_{m_f} = 250$ ns time interval, traditional swap spectroscopy requires $N_{\text{trad}} = 25600$ ³.

For the parameters used in this experiment, octave sampling requires 1 order of magnitude fewer points than traditional swap spectroscopy. It is worth noting that $N_{\text{oct}} \neq \sqrt{N_{\text{trad}}}$ due to the prefactors in Eqs. (5.3) and (5.4). Since octave sampling provides a square-root *scaling* improvement, the reduction in the number of measurements grows for experiments with higher resolution. For example, if we increase m_f to 9, N_{oct} roughly doubles to 5115, whereas N_{trad} is quadrupled to 102400.

Given the octave sampling results, we intend to determine if there are one or more resonance modes interacting with the qubit. If there are resonances, we also want an estimate of their coupling parameters f_{RM} and g . If there are no resonance modes at all, the qubit does not undergo any swap, and we should always measure it to be in $|e\rangle$ with $\bar{P}_e = 1$. Hence, a measurement of $\bar{P}_e < 1$ indicates energy loss due to a resonance mode interacting with the qubit.

In practice, however, other spurious experimental effects can lower the measured \bar{P}_e below the theoretical value of one, even in the absence of a resonance mode. Those include, for instance, the “bare” energy relaxation rate of the qubit, state preparation, measurement visibility, bin averaging, and statistical fluctuations. We therefore require an analysis method that will reliably detect and extract resonances from the octave data, while minimizing false positives.

The method used to analyze the octave data primarily relies on a peak-finding function meant to detect low-excitation bins in the spectrum. To avoid duplicate detections of the same resonance mode, we apply a procedure to combine peaks corresponding to the same resonance found in different octaves. We configure the sensitivity of the analysis by setting the minimum *prominence* value used for peak finding. Setting the prominence to a lower value (more sensitive) will detect more peaks, potentially leading to detection of fainter resonances. However, a low value could also generate false positives if the data is noisy. More details on this analysis method can be found in Appendix C.2.

³A time resolution $\Delta t_{\min} = 2.5$ ns corresponds to the 400 MHz sampling rate used for sampling a signal up to 200 MHz

Table 5.1: Resonance modes detected after analyzing the octave sampling data; the corresponding bins are boxed in red in Fig. 5.3 (b). The parameters listed result from a prominence value of 0.39, except for RM4. For RM4, a lower value of 0.09 (more sensitive) is necessary. For values above 0.72, no resonance is detected. The prominence threshold indicates the maximum tested value for which the resonance is detected; at that value, the reported parameters are slightly different.

Parameter	RM1	RM2	RM3	RM4
Octave number m	7	8	3	8
Bin center frequency (GHz)	4.806	4.832	5.106	4.364
Bin coupling range (MHz)	[2, 4]	[1, 2]	[32, 64]	[1, 2]
Prominence threshold	0.39	0.50	0.72	0.09

The result of the octave analysis consists of the bin location—frequency and octave number—for each detected resonance. Since the octave number corresponds to a coupling strength range, the task of detecting resonances and finding their approximate coupling parameters is achieved.

We analyze the octave sampling data shown in Fig. 5.3 (b) with different prominence values to provide an understanding of the sensitivity required to detect resonances. In that experiment, a qubit interacts with three known modes: two are synthesized with a microwave source and one is the on-chip readout resonator. The analysis detects those three modes at a prominence value of 0.39. When the prominence is decreased to 0.09, an additional resonance is detected at 4.35 GHz. Even at that sensitivity, no false positives are reported. The coupling parameters resulting from the analysis are reported in Table 5.1. The bins corresponding to those detections are boxed in red in Fig. 5.3 (b).

We purposely choose the two synthesized modes, RM1 and RM2, to be close in frequency to illustrate an important feature of our method: two distinct resonances are detected separately only if their frequency spacing is sufficiently large. In particular, the frequency separation must be at least twice as large as the largest of the two coupling strengths. This ensures that the oscillation minima are separated by one bin width. If that condition is not met, the two resonances are located either in the same bin or in neighboring bins, resulting in the detection of a single peak. Here, RM1 and RM2 are separated by 20 MHz and the coupling strength of RM1 is at most 4 MHz. This means that there is at least one bin separating the two modes, allowing them to be detected independently.

The resonator mode, RM3, is detected at a low octave number. This is because it is characterized by a large coupling strength to the qubit (see Table 3.1 for the

coupling capacitance). The last detected mode, RM4, is very faint. It corresponds to an incoherent weakly coupled TLS, as we also see in the traditional swap spectroscopy data of Fig. 5.2 (a). This mode can only be detected by setting a sufficiently low prominence value. It is likely that RM4 could be detected with a less sensitive analysis if an additional octave was sampled.

For the three known modes, we coarsely estimate the frequency and coupling strength with a minimal amount of data. Obtaining more precise and accurate results necessitates the online estimation algorithm, to be explained in the next section.

5.5 Online Bayesian Learning Algorithm

The offline octave sampling algorithm is data efficient and can be performed from a state of zero knowledge of the qubit's spectrum. However, it does not provide a very accurate estimate of the coupling parameters of a resonance mode. To improve accuracy, we can use the coarse estimate given by the offline method to execute an *online Bayesian learning algorithm* and refine the parameters in a very short time. This process relies on measuring a few dozen points of the qubit oscillation in the swap spectrum, using again the pulse sequence of Fig. 5.1. Note that for the online algorithm to work, the qubit must undergo an oscillation. Therefore, this method cannot be used to estimate the coupling parameters of an incoherent resonance mode. If it is unknown whether a mode is coherent or not, a traditional swap spectroscopy experiment has to be run first.

Given an initial probability distribution over the coupling parameters with a resonance mode, the online algorithm successively selects measurement settings to increase knowledge. After the result of a measurement is recorded, the distribution is updated according to Bayes' theorem and a new measurement setting is produced. This procedure is repeated iteratively until the distribution converges as desired. The source code developed for this work can be found online [263].

5.5.1 Theoretical Method

The online estimation algorithm is the experimental implementation of the theory proposed in Ref. [257]. It employs a particle filter method to efficiently represent the *prior* and *posterior* distributions and compute Bayes' theorem at each iteration.

A particle distribution is a discretized representation of a probability distribution. The denser the distribution in a particular region of the parameter space, the higher the probability of those parameters. Here, each particle represents a two-tuple of the coupling parameters (f_{RM}, g) of a resonance mode. At the beginning of an iteration,




Figure 5.4: Illustration of a simulated iteration of the online particle filter algorithm with 40000 particles. (a) Given a prior distribution, we heuristically generate measurement settings f_p and t meant to increase information. (b) Following the measurement, the likelihood \mathcal{L} of the result is computed for each particle. Note that the values shown on the scale are normalized. (c) We apply Bayes' theorem to determine the posterior distribution. This task is achieved by resampling the particles according to their likelihoods. The distribution is split in two "clouds." After resampling, the posterior distribution can be used as the next iteration's prior.

we compute the means, $\mu(f_{\text{RM}}) = \langle f_{\text{RM}} \rangle$ and $\mu(g) = \langle g \rangle$, and standard deviations $\sigma(f_{\text{RM}}) = \sqrt{\langle f_{\text{RM}}^2 \rangle - \langle f_{\text{RM}} \rangle^2}$ and $\sigma(g) = \sqrt{\langle g^2 \rangle - \langle g \rangle^2}$ of the prior distribution.

The next step is to perform a single measurement to determine the excited population P_e at a particular probe frequency f_p and time t . These *measurement settings* are heuristically selected to increase information gain [257]. In practice, t should scale inversely with $\sigma(g)$, while f_p should be within a factor of $\mu(g)$ on either side of $\mu(f_{\text{RM}})$.

We choose the following measurement settings:

$$f_p = \begin{cases} \mu(f_{\text{RM}}) + r_1 \mu(g) & \text{for } M \leq M_0, \\ \mu(f_{\text{RM}}) + c r_1 \sigma(f_{\text{RM}}) & \text{for } M > M_0, \end{cases} \quad (5.5)$$

$$t = \begin{cases} r_2 \tanh \left[\frac{a}{\sigma(g) t_{\text{max}}} \right] t_{\text{max}} & \text{for } M \leq M_0, \\ \frac{1+r_2}{2} \tanh \left[\frac{a}{\sigma(g) t_{\text{max}}} \right] t_{\text{max}} & \text{for } M > M_0, \end{cases} \quad (5.6)$$

where r_1 is picked from $\mathcal{U}(-1/2, 1/2)$, r_2 is picked from $\mathcal{U}(0, 1)$, $a = \pi/2$, $c = 5$, M is the iteration number, and we set $M_0 = 25$ (see below). These parameters are

empirical constants determined in Ref. [257], although they are slightly adjusted for this experiment in order to have a larger distribution for f_p and t when $M > M_0$. Note that, unlike the method proposed in Ref. [257], we choose to limit t to a maximum value well under T_1 . This is done to prevent measurements from occurring after the oscillation has substantially decayed. For this purpose, we use the hyperbolic tangent function as it has a linear behavior for small arguments, such that $\tanh(a/\sigma(g)t_{\max})t_{\max} \simeq a/\sigma(g)$ when $\sigma(g)$ is large.

After M_0 iterations, we modify the heuristic to accelerate convergence. Initially, we choose probe frequencies coarsely according to the value of g . Then, as our knowledge improves, $\sigma(f_{\text{RM}})$ decreases and can be used to select frequencies in a narrower range around $\mu(f_{\text{RM}})$. The factor c is used to avoid choosing measurement frequencies *too* narrowly. The time t is always weighted by $\sim 1/\sigma(g)$, but we bias the selection to larger values after M_0 iterations.

The last step in the iteration is to apply Bayes' theorem to update our knowledge of the coupling parameters. We want to obtain the posterior distribution based on the measurement result $P_e(f_p, t)$. This is achieved in two sub-steps: (1) We compute the likelihood of obtaining the measurement value given each particle's (f_{RM}, g) parameters. (2) We resample the distribution according to these likelihoods.

We compute the likelihood from the measurement result $P_e = n_e/n$, which is the proportion of n_e excited state outcomes for n individual measurement shots. Since the theoretical fraction we expect to measure is $\tilde{P}_e(f_p, t, f_{\text{RM}}, g)$ [given by Eq. (5.2) in the decoherence-free case⁴], we know that the result is a binomial random variable $n_e \sim \mathcal{B}(n, \tilde{P}_e)$. Accordingly, the likelihood of obtaining a particular measurement result given the measurement settings (f_p, t) and a particle with parameters (f_{RM}, g) is the probability mass function

$$\mathcal{L}(n_e | f_p, t, f_{\text{RM}}, g) = \binom{n}{n_e} \tilde{P}_e^{n_e} (1 - \tilde{P}_e)^{(n-n_e)}, \quad (5.7)$$

where $\binom{n}{n_e}$ is the binomial coefficient.

In effect, we are computing the probability that the measurement result corresponds to a resonance mode with coupling parameters (f_{RM}, g) . The next step is to resample the distribution to keep only those parameters that are most probable. Although this task can be achieved in a variety of ways, the general idea is to pick particles from the prior at random, weighted by the likelihood. To avoid duplicate particles in the posterior distribution, we add normally distributed random noise proportional to the

⁴In the experiment, we use a model for \tilde{P}_e that accounts for relaxation and measurement visibility. For relaxation, we use Eq. (17) in Ref. [257]. To account for measurement visibility, we clamp the theoretical probability between 0.05 and 0.95.

covariance of the prior. The procedure chosen for this experiment is described in Appendix C.4.

The iteration process is visualized in Fig. 5.4, allowing us to understand more intuitively how the particle filter technique works. If the measurement is useful, i.e., the resulting likelihood favors a subset of the prior, the posterior distribution is shrunk or filtered, improving knowledge of the parameters. Otherwise, if the likelihood does not discriminate the particles, the distribution is not modified significantly. After resampling, the next iteration starts with the last iteration’s posterior as prior.

The task of the online Bayesian learning algorithm is simpler than octave sampling because we already have a coarse estimate of the interaction parameters. The particle filter can therefore “fit” to the most likely parameters given the measurements. At the end of the final iteration, the parameters are given by the mean of the particle distribution. If the algorithm converges, the final particle “cloud” is small, resulting in an estimate that is accurate: both true and precise. If the algorithm does not converge, meaning that the final particle cloud is not tightly concentrated in a single region, it might be necessary to run the experiment again. Regenerating the initial particle distribution via the octave sampling method could also improve estimation performance.

Ideally, the uncertainty on the estimated parameters is given by the standard deviation of the final particle distribution. In practice, however, directly taking the standard deviation is generally not valid in an experimental context. This is because the algorithm does not take into account potential errors on the value of the measurement settings or any model inaccuracies. To obtain an uncertainty on the parameters, we instead re-run the full inference steps of the algorithm multiple times on the already-gathered data, starting from the same initial particle distribution. If the model and measurement settings were accurate, the results of re-running the inference would be identical. However, since that is not the case, the estimates obtained by re-running the inference vary. The error we report is then the standard deviation of the different estimates. Since this procedure does not require acquiring new data it can be performed offline, after the experiment.

We note that one possible cause of failure is the overestimation of g by an integer multiple. In that case, crests in the oscillation for the different frequencies partially overlap. In order to prevent such failures, the experiment can be run multiple times.

5.5.2 Experimental Results

We run the online Bayesian learning algorithm on three distributions generated from the octave data, one for each detected resonance mode. The generation of those distributions is discussed in Appendix C.3. We do not execute the algorithm on RM4 since the resonance is incoherent and does not undergo the oscillations necessary to estimate the coupling parameters with this method.

Table 5.2: Estimated parameters for the three resonance modes detected after running the online Bayesian learning algorithm. RM1 and RM2 are synthesized modes and RM3 corresponds to the qubit’s measurement resonator.

Parameter	RM1	RM2	RM3
f_{RM} (GHz)	4.8091(3)	4.8297(2)	5.033(12)
g (MHz)	2.78(3)	1.62(8)	37.9(3)

For each mode, we perform 35 iterations of the algorithm, at which point the distribution has converged. The runtime of the algorithm for a single resonance mode is approximately 23 s. Afterwards, we re-run the full inference 200 times—enough for the resulting statistics to stop changing—and report the final means and standard deviations in Table 5.2. As expected, the parameters of the synthesized modes and of the measurement resonator are correctly identified. Note that the errors shown do not include systematic errors caused by an inaccurate flux-amplitude-to-frequency calibration (see Appendix C.1).

Here, the sources are set at 4.810 and 4.8305 GHz. These values are higher than the ones found in Table 5.2 by ~ 1 MHz, likely due to a systematic calibration error. In principle, it is possible to relate the power emitted by a synthesizer to the “coupling strength” of the mode. However, the attenuation and reflection of the signal between the source and qubit make it impossible to accurately find such a relation. Instead, we fit the swap spectroscopy measurement of Fig. 5.2 (b) with a non-linear least squares regression, giving $g = 2.852(1)$ and $1.472(1)$ MHz for RM1 and RM2, approximately two standard deviations away from the results in Table 5.2. The anticrossing frequency and the coupling strength of the resonator are estimated to be 5.032 GHz and 37.4 MHz by a full Hamiltonian fit (see Section 3.4 and Table 3.1).

To test the performance of the estimation algorithm, we run it 1000 times on RM2 with slightly different initial distributions. The mean of each of the 1000 initial particle distributions is distributed uniformly at random within a 10 and 1.5 MHz range for f_{RM} and g , respectively. Each individual particle distribution is uniform, with a width of 15 MHz in f_{RM} and 2.5 MHz in g .

We plot the convergence of the parameters in Fig. 5.5. As shown by the histograms, more than 99% of the runs converge successfully to properly estimate the frequency and coupling strength, with just a few failures. The average of the parameters after the 35th iteration is $f_{\text{RM}} = 4.8301(4)$ GHz and $g = 1.45(9)$ MHz.

The estimated true parameters of the resonance, which we compute by fitting the dataset combining all 35000 measurements, are $f_{\text{RM}} = 4.83008$ GHz and $g = 1.445$ MHz. These values are shown with a red line on the histograms. We compare the




Figure 5.5: Performance of the online estimation algorithm over 1000 runs with RM2. (a-b) Histograms over the means of the initial and final particle distributions. (c-d) Mean of the posterior particle distribution computed after each iteration’s measurement. As more iterations are made, the particle distribution converges toward the true value of the parameters. The 0th iteration corresponds to the initial distribution. The true values of the parameters are identified by a red line in the histograms. The total runtime of the experiment is 6.4 h. Each individual run executes 35 iterations and, thus, 35 measurements, taking ≈ 23 s. Most of the time (60%) is spent acquiring data. In fact, each measurement comprises 786 shots at a repetition rate of 2 kHz. The leftover time is attributable to data transfer and processing. The total computation time for the estimation algorithm is ≈ 1 s.

results to a fit instead of the synthesized frequency to circumvent potential calibration inaccuracies. It is worth noting that the power for this experiment is slightly different from that of the experiment summarized in Table 5.2.

Several experimental errors could cause the variation of a parameter in the likelihood model of Eq. (5.7). While the binomial likelihood accounts for variance in the qubit measurement, it assumes that all parameters are constant. Consequently, any parameter fluctuation or drift causes a discrepancy between the inference model and the physical system. For example, the qubit probe frequency f_p and energy relaxation

time T_1 can be modified by strongly coupled TLSs [254, 261, 260, 262]; f_p can additionally be perturbed by a nonconstant flux pulse amplitude (see Appendix C.1) or flux noise in the SQUID loop generated by, e.g., voltage sources or on-chip magnetic defects. Depending on the nature of the resonance mode, f_{RM} and possibly g may vary due to, e.g., changing temperature or strong coupling to TLSs. If the mode is another qubit, its frequency may be affected in the same manner as f_p . Finally, the measurement visibility itself could vary due to possible instrumentation issues.

The errors introduced above can be classified according to their timescale relative to the length of the online estimation experiment. A parameter that changes during the execution of the algorithm would lead to inconsistencies with subsequent measurements. This effect is expected to be taken into account by the variance of the repeated inferences, thereby increasing the calculated standard error. For example, if f_{RM} changes during the experiment, repeating the inference multiple times would likely lead to a bimodal distribution, i.e., a spread in the estimated value of f_{RM} .

If a parameter changes over a longer timescale, in our case $\gtrsim 30$ s, measurements taken during an experiment remain consistent and the standard error does not increase. If such an error is suspected, the calibration should be repeated at regular intervals, as mentioned in Sec. 5.2.

5.6 Discussion

Both algorithms presented above depend on a few parameters that are crucial to their function. For the octave sampling algorithm, the choice of the frequency range to be measured is naturally determined by the properties of the device: superconducting qubits have a limited frequency range within which they operate optimally. For the device in this work, the upper end of the measurement range f_{max} simply corresponds to the maximum attainable frequency. The lower limit f_{min} is chosen to be as low as desired, keeping in mind that TLS far below the operating frequency range of the qubit are not a cause for concern. In addition, given that the purpose of detecting interactions is to then *select* optimal operating frequencies, it might be sufficient to set f_{min} to the lowest frequency where high-fidelity control and readout is achievable. Since we use a resonator for readout, the farther away the qubit is in frequency, the lower the fidelity of the measurement. Other constraints, e.g., pulse control bandwidth, might dictate even tighter limits.

A second important parameter for the octave sampling method is n_s . In principle, a single high-quality (many shots) measurement of P_e at the center of the bin should be sufficient. This would be analogous to traditional swap spectroscopy. However, because of the efficiency of the octave method, we can afford to take a few more measurements per bin. This is what we have chosen to do by randomly distributing $n_s = 5$ measure-

ments per bin. This redundancy increases the detection sensitivity and protects against possible statistical fluctuations in the measurement.

One more parameter worth discussing is the number of octaves to be measured m_f , which corresponds to t_{m_f} , or equivalently, Δf_{m_f} , as explained in Sec. 5.4.1. Generally, this parameter should be determined by the requirements of the experiment for which the calibration is made. If a long gate sequence is needed, e.g., for randomized benchmarking, detecting weakly coupled resonance modes is important. This would not necessarily be the case for shorter experiments, such as process tomography. A total time $t_{m_f} \sim t_{\text{exp}}$, where t_{exp} is the length of the experimental gate sequence, is therefore generally a good choice.

For the particle filter algorithm, the choice of a , c , and M_0 is discussed in Ref. [257]. Other parameters of interest include t_{max} and the number of particles to be used. The time t_{max} is used in Eq. (5.6) to restrict the maximum measurement time. This is necessary because the qubit eventually decays to the ground state. To obtain reliable results, t_{max} should be set well below T_1 . Note that another way to limit the maximum measurement time would be to replace $\tanh(a/\sigma(g)t_{\text{max}})t_{\text{max}}$ with $a/\sigma(g)$ in Eq. (5.6) (as in the original proposal) and simply stop the algorithm once a sufficiently small $\sigma(g)$ is reached.

The number of particles to be used is constrained mainly by the performance of the computer running the resampling procedure and, potentially, numerical accuracy issues [264]. As a rule of thumb, at least 10000 particles should be used; in this work, we have used 40000.

In Sec. 5.2, we explain that our method can be used not only for a simple qubit swap spectroscopy experiment, but also to look for resonances with the $|e\rangle \leftrightarrow |f\rangle$ transition with a double-excitation protocol. In fact, the algorithms discussed in this work are very general and apply even to systems that do not involve a resonance mode. We can use the online and offline algorithms to efficiently detect the location and estimate the parameters of any qubit dynamics akin to a chevron pattern. This is the case, for example, with a whole class of parametric two-qubit gates, where instead of varying f_p we vary the frequency of a flux drive applied to the SQUID of a qubit or tunable coupler [265, 266, 267].

Finally, we briefly discuss the problem of *choosing* qubit operating frequencies. Once the calibration showcased in this work is accomplished and all resonant couplings are identified, the next step is to use this information to optimize the performance of a quantum computer. This process depends on the quantum computing architecture. For an array of directly-coupled superconducting qubits, we want to avoid crosstalk between neighboring qubits and minimize interactions with TLSs. We therefore need to choose the idle frequencies of all qubits at the same time, taking into account both wanted and unwanted couplings. Additional concerns apply for choosing the operating frequencies of two-qubit gates: we must consider the frequency *path* that the qubits

will take during the gate. For example, it is undesirable for a qubit to cross through a resonance with a TLS. If such crossings are unavoidable, e.g., if the device is afflicted by many TLSs, knowing the coupling strength of each TLS helps select the optimal set of qubit interaction frequencies.

While this work does not explain the process needed to perform this optimization (see, e.g., Ref. [268] for an example), we emphasize that the runtime improvement of the offline and online algorithms when compared to traditional swap spectroscopy presents several advantages: First, the calibration may be run more often. Second, the calibration is affordable enough to be run on a larger spectrum, giving the frequency optimization process more information to work with.

5.7 Conclusions

In conclusion, we explain two methods for the Hamiltonian parameter estimation of resonant couplings in the context of tunable superconducting qubits. Both algorithms work well and are able to successfully identify and accurately estimate the parameters of various resonance modes. The octave sampling technique can be run without prior knowledge about potential resonances in the environment of the qubit and allows efficient detection of coupled modes within a chosen parameter range. The online Bayesian algorithm can be performed or omitted depending on whether a more accurate estimate of the coupling parameters is desired. Using these algorithms reduces the number of measurements needed from $\mathcal{O}(1/\Delta f_{m_f}^2)$ to $\mathcal{O}(1/\Delta f_{m_f})$. This translates to a reduction in runtime by one order of magnitude in typical conditions.

We experimentally demonstrate both techniques on a superconducting Xmon transmon qubit and evaluate their performance. We are able to detect the resonance with the qubit's measurement resonator, as well as with synthesized resonance modes and a naturally occurring weakly coupled TLS. We determine that the methods are efficient, reliable, and readily automated. We expect this type of calibration to be critical to the operation of large-scale quantum computers, superconducting or otherwise. Future work includes integrating the information we acquire by our methods into a comprehensive optimization process for selecting the operating frequency of each qubit in a quantum computer and implementing the calibration of a two-qubit gate with the online Bayesian algorithm.

Time Fluctuations

Amorphous dielectric materials have been known to host two-level systems (TLSs) for more than four decades. Recent developments on superconducting resonators and qubits enable detailed studies on the physics of TLSs. In particular, measuring the loss of a device over long time periods (a few days) allows us to investigate stochastic fluctuations due to the interaction between TLSs. In this chapter¹, we measure the energy relaxation time of a frequency-tunable planar superconducting qubit over time and frequency. The experiments show a variety of stochastic patterns that we are able to explain by means of extensive simulations. The model used in our simulations assumes a qubit interacting with high-frequency TLSs, which, in turn, interact with thermally activated low-frequency TLSs. Our simulations match the experiments and suggest the density of low-frequency TLSs is about three orders of magnitude larger than that of high-frequency ones.

6.1 Introduction

Superconducting devices operated in the quantum regime [99] are ideal tools to study the properties of amorphous dielectric materials [252]. These materials are known to be characterized by defects that can be modeled as two-level systems (TLSs) [269]. TLSs can interact with superconducting resonators or qubits, resulting in dissipation channels that are particularly prominent in planar devices. Such devices are fabricated by depositing superconducting films made from metals, e.g., aluminum (Al) or niobium, on silicon (Si) or sapphire substrates. A few examples of planar devices can be found in our works of Refs. [114] and [249], where we have investigated coplanar waveguide (CPW) resonators [234] as well as Xmon transmon qubits [92].

¹This chapter was largely adapted from Ref. [262]. The list of author contributions can be found in the Statement of Contributions within the front matter of this thesis. © 2021 American Physical Society

A large body of work on CPW resonators and qubits has shown that TLSs are likely hosted in native oxide layers [270, 112, 271, 272, 105, 273, 274, 171, 253, 275] at the substrate-metal (SM), substrate-air (SA), or metal-air (MA) interfaces [276, 228, 114, 277]. TLSs originate within these layers because naturally occurring oxides deviate from crystalline order. This deviation may result in trapped charges, dangling bonds, tunneling atoms, or collective motion of molecules.

It is convenient to distinguish between two categories of TLSs based on their energy E and the device operating temperature T . When $E > k_B T$, the corresponding TLSs reside in the quantum ground state; these TLSs are hereafter referred to as quantum-TLSs (Q-TLSs). When $E < k_B T$, the TLSs are thermally activated and are referred to as thermal-TLSs (T-TLSs). Typically, superconducting resonators are characterized by a resonance frequency f_r and qubits by a transition frequency f_q , with $f_r \sim f_q \sim 5$ GHz, and are operated at $T \sim 50$ mK. Hence, the energy threshold between Q- and T-TLSs is $E/h \sim 1$ GHz.

Superconducting quantum devices interact (semi-)resonantly with Q-TLSs [278], affecting the internal quality factor of resonators, Q_i , or the energy relaxation time of qubits, T_1 . Several authors have hypothesized that Q-TLSs additionally interact with T-TLSs [279, 280, 281], leading to experimentally observed stochastic fluctuations in Q_i and f_r [282, 279, 171, 253] as well as T_1 and f_q [283, 252]. The model proposed by these authors depart from the TLS standard tunneling model (STM), where TLS interactions are neglected [269]. The interacting model is sometimes called the *generalized tunneling model* (GTM).

It has recently been shown that planar fixed-frequency transmon qubits exhibit random fluctuations in both T_1 and f_q over very long time periods [261, 260, 284]. Frequency-tunable transmon qubits, as the Xmon, show TLS-induced fluctuations predominantly in T_1 [254]. TLS-induced f_q fluctuations are present but are overshadowed by additional noise processes such as flux noise². These findings serve as the main motivation for the experiments and simulations presented in this chapter.

In this chapter, we present the experimental measurement of spectrotemporal charts for an Xmon transmon qubit as well as the results of detailed simulations corresponding to these experiments. In the spectrotemporal charts, T_1 is measured and simulated for time periods up to 48 h and for f_q ranges up to 300 MHz. Our main objective is to validate the Q-TLS-T-TLS interaction hypothesis in the GTM by comparing experiments and simulations. In our simulations, a qubit interacts with an ensemble of Q-TLSs, the frequencies of which undergo stochastic fluctuations due to the interaction with T-TLSs. For every Q-TLS we consider a set of interacting T-TLSs, where the dynamics of each T-TLS state are governed by a random telegraph signal (RTS). The

²The flux noise experienced by tunable transmon qubits is caused by the increased flux sensitivity due to the SQUID and the flux line coupled to it.

Q-TLS frequency fluctuation process, which is broadly referred to as *spectral diffusion*, is responsible for the random fluctuations in T_1 .

The comparison between experiments and simulations reveals that the Q-TLS–T-TLS interaction likely exists, as proposed in the GTM. In particular, our simulations reproduce well the spectral-diffusion patterns presented in the experiments. Our model suggests that the density of T-TLSs is significantly higher than that of Q-TLSs. We find a T-TLS density of approximately $6 \times 10^5 \text{ GHz}^{-1} \mu\text{m}^{-3}$, which is about three orders of magnitude larger than the Q-TLS density.

Finally, we show that certain statistical analyses, such as the Allan deviation, are not able to capture the fluctuation characteristics of a given time series (e.g., the number of T-TLSs contributing to the stochastic process). Instead, a direct analysis of the time series provides a more accurate description of the stochastic processes due to TLSs.

The chapter is organized as follows. In Sec. 6.2, we review the theory necessary to describe the stochastic fluctuations of T_1 . In Sec. 6.3, we explain the methods required to perform experiments and simulations. In Sec. 6.4, we present our main results. In Sec. 6.5, we provide an in-depth discussion on some of our main results. Finally, in Sec. 6.6, we summarize our findings and suggest a roadmap for future work.

6.2 Theory

In this section, we introduce physical models of TLSs (Sec. 6.2.1); we then describe the qubit–Q-TLS and Q-TLS–T-TLS interaction (Subsecs. 6.2.2 and 6.2.3); finally, we amalgamate the previous concepts in order to explain qubit stochastic fluctuations (Sec. 6.2.4).

6.2.1 Physical Models of TLSs

The STM is a phenomenological model describing defects in amorphous dielectric materials. The defects are commonly assumed to be quantum-mechanical double-well potentials, or TLSs, with energy barrier V . In the STM, the TLS tunneling energy Δ_0 is calculated by means of the WKB approximation,

$$\Delta_0 \simeq h\Omega_0 \exp\left(-\frac{d}{\hbar}\sqrt{2mV}\right). \quad (6.1)$$

In this equation, Ω_0 is the attempt frequency (assumed to be the same for both wells), d is the spatial distance between the two wells, and m is the mass of the physical entity associated with the TLS (e.g., a molecular mass) [285].

The unperturbed Hamiltonian of a TLS reads $\hat{H}_{\text{TLS}} = (\Delta\hat{\sigma}_z + \Delta_0\hat{\sigma}_x)/2$, where Δ is the asymmetry energy between the two wells of the TLS; $\hat{\sigma}_z$ and $\hat{\sigma}_x$ are the usual

Pauli matrices in the so-called diabatic (“left” and “right”) basis. By diagonalizing this Hamiltonian we obtain $\hat{H}_{\text{TLS}} = E \hat{\sigma}_z/2$, where

$$E = \sqrt{\Delta^2 + \Delta_0^2} \quad (6.2)$$

is the TLS energy and $\hat{\sigma}_z = [\hat{\sigma}_z \cos(\theta) + \hat{\sigma}_x \sin(\theta)]/2$ is the Pauli matrix in the energy basis; $\theta = \arctan(\Delta_0/\Delta)$ is the rotation angle used to perform the diagonalization.

One of the hypothesis in the STM is that Δ and Δ_0 are uncorrelated quantities with joint probability density

$$f_{\Delta, \Delta_0} = \begin{cases} \frac{D}{\Delta_0}, & \text{for } \Delta \geq 0 \text{ and } \Delta_0 \geq \mathcal{E}_{\min}; \\ 0, & \text{otherwise.} \end{cases} \quad (6.3)$$

In this equation, D is the TLS density in units of inverse energy and volume and \mathcal{E}_{\min} is the minimum tunneling energy. A further hypothesis is that interactions between TLSs are very weak and, thus, *negligible*.

The hypotheses behind the STM prevent this model from explaining a variety of features observed in devices affected by TLS defects. Among other phenomena, the STM cannot explain the temperature dependence of the frequency noise of superconducting resonators [279] as well as the strong temperature dependence of the relaxation rate of Q-TLSs measured with qubits [286]. Most importantly, the STM cannot explain the spectral diffusion dynamics observed both in the work of Ref. [254] and in our experiments.

In order to resolve these shortcomings, it is necessary to extend the STM to the GTM by making the following modifications:

1. Interactions between TLSs are not neglected.
2. The joint probability density is assumed to be nonuniform with respect to Δ ,

$$f_{\Delta, \Delta_0} = \begin{cases} \frac{1 + \mu}{\Delta_0} \left(\frac{\Delta}{\mathcal{E}_{\max}} \right)^\mu, & \text{for } 0 \leq \Delta \leq \mathcal{E}_{\max} \\ & \text{and } \mathcal{E}_{\min} \leq \Delta_0 \leq \mathcal{E}_{\max}; \\ 0, & \text{otherwise.} \end{cases} \quad (6.4)$$

In this equation, $\mu < 1$ is a small positive parameter and \mathcal{E}_{\max} is a maximum energy cutoff dictated by the energy scales of the system under consideration (see Sec. 6.3.2).

The interaction energy between any pairs of TLSs is assumed to be a function of their spatial separation r ,

$$U(r) = \frac{U_0}{r^3}, \quad (6.5)$$

where U_0 is a material-dependent parameter associated with electric or elastic interactions. It is worth noting that interactions can occur between pairs of Q-TLSs or T-TLSs as well as between a T-TLS and a Q-TLS.

In the study of superconducting planar qubits, both f_q and T_1 are affected by the interactions hypothesized in the GTM. These type of qubits interact semi-resonantly with an ensemble of Q-TLSs, where each Q-TLS can strongly interact with one or more T-TLSs. Such interactions lead to stochastic fluctuations in T_1 and f_q .

6.2.2 Qubit-Q-TLS Interaction

The interaction between a qubit and a single Q-TLS leads to perturbations in T_1 and f_q . These perturbations depend on the coupling strength between the qubit and Q-TLS, g , and on the difference between the Q-TLS transition frequency $f_{\text{Q-TLS}}$ and f_q , $\Delta f = f_q - f_{\text{Q-TLS}}$. In this work, we consider only T_1 fluctuations because, for a tunable qubit, f_q fluctuations are dominated by other noise processes such as flux noise.

In the rotating frame of the qubit and after a rotating wave approximation, the Hamiltonian of the qubit coupled to the Q-TLS reads

$$\hat{H}_{\text{q,Q-TLS}} = h\Delta f \hat{\sigma}_q^+ \hat{\sigma}_q^- + hg \left(\hat{\sigma}_q^+ \otimes \hat{\sigma}_{\text{Q-TLS}}^- + \text{H.c.} \right), \quad (6.6)$$

where $\hat{\sigma}_q^\mp$ and $\hat{\sigma}_{\text{Q-TLS}}^\mp$ are the qubit and Q-TLS lowering and raising operators in the energy basis and H.c. is the Hermitian conjugate of the first term in parentheses. The coupling strength g is due to the electric dipole moment \vec{p} of the Q-TLS and the electric field \vec{E}_q of the qubit³, $hg = \vec{p} \cdot \vec{E}_q$.

The master equation in Lindblad form of a qubit-Q-TLS system reads

$$\frac{d\hat{\rho}}{dt} = -\frac{i}{\hbar} [\hat{H}_{\text{q,Q-TLS}}, \hat{\rho}] + \sum_j \left(\hat{L}_j \hat{\rho} \hat{L}_j^\dagger - \frac{1}{2} \{ \hat{L}_j^\dagger \hat{L}_j, \hat{\rho} \} \right), \quad (6.7)$$

where $\hat{\rho}(t)$ is the density matrix, $\hat{H}_{\text{q,Q-TLS}}$ is given by Eq. (6.6), $j \in \{\text{q, Q-TLS}\}$, and \hat{L}_j and \hat{L}_j^\dagger are Lindblad operators.

³The electric field \vec{E}_q is the field associated with the qubit capacitor, which is described in Sec. 3.1 and Appendix D.3.

To account for the energy relaxation rates of the qubit and Q-TLS, we introduce the following Lindblad operators

$$\hat{L}_q = \sqrt{\tilde{\Gamma}_1^q} \hat{\sigma}_q^- \quad (6.8)$$

$$\hat{L}_{\text{Q-TLS}} = \sqrt{\Gamma_1^{\text{Q-TLS}}} \hat{\sigma}_{\text{Q-TLS}}^- \quad (6.9)$$

where $\Gamma_1^{\text{Q-TLS}}$ is the energy relaxation rate of the Q-TLS due to phononic interactions with the environment, $\tilde{\Gamma}_1^q$ is the bare energy relaxation rate of the qubit⁴.

The quantity $\hat{L}_j \hat{\rho} \hat{L}_j^\dagger = 0$ at all times because there is at most one excitation in a qubit-Q-TLS coupled system. With this assumption, and by defining the effective non-Hermitian Hamiltonian [287]

$$\hat{H}_{\text{eff}} = \hat{H}_{\text{q,Q-TLS}} - \frac{i}{2} \left(\tilde{\Gamma}_1^q \hat{\sigma}_q^+ \hat{\sigma}_q^- + \Gamma_1^{\text{Q-TLS}} \hat{\sigma}_{\text{Q-TLS}}^+ \hat{\sigma}_{\text{Q-TLS}}^- \right), \quad (6.10)$$

the Lindbladian of Eq. (6.7) can be written as a simple Schrödinger equation with a “decaying wave function” $|\Psi(t)\rangle = \alpha(t)|e\rangle + \beta(t)|1\rangle$, where $\alpha(t)$ and $\beta(t)$ are the time-dependent complex amplitudes associated with the excited state $|e\rangle$ of the qubit and $|1\rangle$ of the Q-TLS.

The exact result of the Schrödinger equation for $\alpha(t)$ given that $\alpha(t=0) = 1$ and $\beta(t=0) = 0$ is

$$\alpha(t) = \frac{1}{2\Lambda} \left[a \exp\left(-\frac{\Lambda}{4}t\right) + b \exp\left(\frac{\Lambda}{4}t\right) \right] \times \exp\left(-\frac{\tilde{\Gamma}_1^q + \Gamma_1^{\text{Q-TLS}}}{4}t\right), \quad (6.11)$$

where Λ is given by Eq. (6.14), $a = \Lambda - (\Gamma_1^{\text{Q-TLS}} - \tilde{\Gamma}_1^q) + 4\pi i \Delta f$, and $b = \Lambda + (\Gamma_1^{\text{Q-TLS}} - \tilde{\Gamma}_1^q) - 4\pi i \Delta f$.

Since we are calculating a decay, we are only interested in the envelope of $\alpha(t)$, $\tilde{\alpha}(t)$. We thus set $\text{Im}[\Lambda] = 0$ in the two exponential terms of Eq. (6.11) and calculate the envelope probability $\tilde{P}_e(t) = |\tilde{\alpha}(t)|^2$ for the qubit to be in $|e\rangle$,

$$\begin{aligned} \tilde{P}_e(t) &= \left| \frac{a}{2\Lambda} \right|^2 \exp\left[-\frac{\Gamma_1^{\text{Q-TLS}} + \tilde{\Gamma}_1^q + \text{Re}[\Lambda]}{2}t\right] + \left| \frac{b}{2\Lambda} \right|^2 \exp\left[-\frac{\Gamma_1^{\text{Q-TLS}} + \tilde{\Gamma}_1^q - \text{Re}[\Lambda]}{2}t\right] \\ &\quad + \frac{ab^* + a^*b}{|2\Lambda|^2} \exp\left[-\frac{\Gamma_1^{\text{Q-TLS}} + \tilde{\Gamma}_1^q}{2}t\right]. \end{aligned} \quad (6.12)$$

⁴This is the rate caused by all dissipation sources other than TLSs.

When $\Gamma_1^{\text{Q-TLS}} > \tilde{\Gamma}_1^{\text{q}}$, which is the regime of interest in our experiments, the term proportional to $|b|^2$ in Eq. (6.12) is dominant. Therefore, in order to find an approximate expression for the Q-TLS contribution only, we subtract the qubit contribution $\tilde{\Gamma}_1^{\text{q}}$ from the rate in the exponential proportional to $|b|^2$. As a result, the contribution to the energy relaxation rate of the qubit due to the Q-TLS can be approximated by

$$\Gamma_1^{\text{q,Q-TLS}} = \frac{\Gamma_1^{\text{Q-TLS}} - \tilde{\Gamma}_1^{\text{q}} - \text{Re}[\Lambda]}{2}, \quad (6.13)$$

where

$$\Lambda = \sqrt{\left(\tilde{\Gamma}_1^{\text{q}} + 2i(2\pi\Delta f) - \Gamma_1^{\text{Q-TLS}}\right)^2 - 16(2\pi g)^2}, \quad (6.14)$$

with $i^2 = -1$.

In presence of amorphous dielectric materials, the qubit is coupled to an ensemble of Q-TLSs. In this case, Eq. (6.13) represents the individual contribution to the energy relaxation rate of the qubit due to the k -th Q-TLS, $\Gamma_1^{\text{q,Q-TLS}} \rightarrow \Gamma_1^{\text{q},k}$; each Q-TLS is now characterized by its own coupling strength g_k , frequency f_k , and energy relaxation rate Γ_1^k . The effective qubit relaxation rate is therefore given by

$$\Gamma_1^{\text{q}} = \frac{1}{T_1} = \tilde{\Gamma}_1^{\text{q}} + \sum_k \Gamma_1^{\text{q},k}. \quad (6.15)$$

6.2.3 Q-TLS-T-TLS Interaction

We intend to calculate the frequency shift experienced by a Q-TLS due to the interaction with a T-TLS. We assume that the unperturbed energy and eigenstates are $E = E_{\text{T-TLS}}$ and $\{|-\rangle, |+\rangle\}$ for the T-TLS and $E = E_{\text{Q-TLS}} \gg E_{\text{T-TLS}}$ and $\{|0\rangle, |1\rangle\}$ for the Q-TLS. These two TLSs form a quantum-mechanical system with Hamiltonian given by Eq. (11) in the work of Ref. [280]. Assuming the interaction energy U between the T-TLS and Q-TLS is given by Eq. (6.5), the four eigenenergies of the system are

$$E_0^{\mp} = -\frac{E_{\text{Q-TLS}}}{2} \mp \sqrt{\left(\frac{E_{\text{T-TLS}}}{2}\right)^2 + U\Delta + U^2} \quad (6.16a)$$

and

$$E_1^{\mp} = +\frac{E_{\text{Q-TLS}}}{2} \mp \sqrt{\left(\frac{E_{\text{T-TLS}}}{2}\right)^2 - U\Delta + U^2}, \quad (6.16b)$$

where Δ is the asymmetry energy of the T-TLS.

The frequency shift δf^\mp of the Q-TLS due to the interaction with the T-TLS reads

$$h \delta f^\mp = E_1^\mp - E_0^\mp - E_{\text{Q-TLS}}, \quad (6.17)$$

which is negative when the T-TLS is in $|-\rangle$ and positive otherwise.

A T-TLS is thermally activated because of the condition $E_{\text{T-TLS}} < k_B T$ and, thus, switches state in time. This causes the sign of δf^\mp to change, affecting the time evolution of the frequency of the Q-TLS coupled to it.

6.2.4 Qubit Stochastic Fluctuations

We assume that the state of a T-TLS over time is modeled by an RTS with switching rate

$$\gamma = \gamma_0 \exp\left(-\frac{V}{k_B T}\right), \quad (6.18)$$

where γ_0 is a heuristic proportionality constant and V is implicitly given by Eq. (6.1).

A Q-TLS is generally coupled to several T-TLSs, where the ℓ -th T-TLS is characterized by a certain value of γ_ℓ and δf_ℓ^\mp . Given the state ($|\mp\rangle$) of each T-TLS at a time t , we can approximate the effective frequency shift of the Q-TLS by summing the individual values of $\delta f_\ell^\mp(t)$. Since the T-TLS state is modeled by an RTS, the effective shift varies with t leading to a time series

$$f_{\text{Q-TLS}}(t) = \frac{E_{\text{Q-TLS}}}{h} + \sum_{\ell} \delta f_\ell^\mp(t). \quad (6.19)$$

For the k -th Q-TLS, f_k fluctuates in time according to Eq. (6.19). As a consequence, Γ_1^q fluctuates because of its dependence on $\Gamma_1^{q,k}$, which, in turn, depends on f_k through Eq. (6.14). The stochastic fluctuations of $\Gamma_1^q = 1/T_1$ are the main subject of this chapter.

6.3 Methods

In this section, we describe the methods used to perform the experiments on T_1 fluctuations (Sec. 6.3.1) and the corresponding simulations (Sec. 6.3.2).

6.3.1 Experiments

In this work, we use an Xmon transmon qubit to probe TLS defects. The main goal of our experiments is to characterize fluctuations in T_1 over long time periods and for different values of f_q . We measure T_1 by means of a standard energy relaxation experiment, a “ T_1 experiment.” Details on the qubit and setup are given in Chapter 3.

Table 6.1: Experimental parameters for the three datasets introduced in Sec. 6.4. Number of frequency points, N_f . Qubit frequency, f_q . Repetition period, Δt . Observation time, t_{obs} .

Dataset	N_f (-)	f_q range (GHz)	Δt (s)	t_{obs} (h)
1	16	[4.369, 4.669]	640	42.5
2	31	[4.500, 4.560]	1000	47.2
3	31	[4.500, 4.530]	1000	48.1

In a T_1 experiment, we prepare the qubit in the excited state $|e\rangle$ by means of a π pulse. We then measure the average population of $|e\rangle$, P_e , for many values of a delay time spaced logarithmically between 1 ns and 200 μs . We read out the qubit state over 655 high-power single-shot measurements (see Section 3.3.3) for each delay time to find P_e with a visibility $\gtrsim 90\%$. Due to the various relaxation channels affecting the qubit, including TLS interactions, P_e decays exponentially in time. We obtain T_1 by fitting the exponential decay and acquire between 36 and 38 points for each T_1 experiment.

We measure T_1 for different values of f_q by setting a quasi-static flux bias ϕ_Z^{qs} applied to the qubit. The correspondence between ϕ_Z^{qs} and f_q is obtained from a qubit parameter calibration. Depending on the experiment, we set f_q over different bandwidths varying between 30 and 300 MHz. We select N_f linearly spaced values of f_q for each T_1 experiment. The T_1 measurements are repeated continuously at a repetition period Δt over an observation time t_{obs} , leading to matrices of data points as detailed in Appendix D.1. These matrices constitute the spectrotemporal charts of T_1 presented in Sec. 6.4. The experimental parameters for the three datasets shown in this work are reported in Table 6.1.

6.3.2 Simulations

The procedure to simulate the effect of TLSs on the stochastic fluctuations in T_1 is composed of three main steps: (1) Generate an ensemble of Q-TLSs interacting with the qubit. (2) Generate several T-TLSs interacting with each Q-TLS. (3) Generate a time series for each T-TLS and propagate the effect of the T-TLSs' switching state to each Q-TLS, and, finally, to the qubit.

Before detailing each step of the procedure, it is worth introducing a few general assumptions:

- We consider that all TLSs are distributed uniformly in the oxide layers at the SA and MA interfaces of the qubit device. The thickness of these layers is assumed to be $t_{\text{ox}} = 3$ nm for both interfaces, a reasonable amount for a device exposed to ambient air pressure [276, 288, 289]. Note that device fabrication affects the oxide layers. For example, without thorough substrate cleaning, it would be expected that there would be an oxide layer at the SM interface⁵.
- All the TLS parameters used in this procedure are assumed to be fixed for the entire duration of each simulation (see Appendix D.2); for each T-TLS, for example, γ is constant in time.
- We assume that all T-TLSs belong to a single species (see Sec. 6.5.1).
- We set $\tilde{\Gamma}_1^{\text{q}} = 1/27$ MHz, which is the value estimated according to the highest observed T_1 values for our device.
- For all distributions used in this work, we determine the probability density function (PDF) by normalizing a given distribution [e.g., that represented by Eq. (6.4)] over the chosen boundary values; we also find the cumulative density function (CDF). In order to pick a random value from a distribution, we generate a random quartile value between 0 and 1. We then calculate the random value corresponding to the generated quartile either by inverting the CDF or via root finding.

For step (1), we follow a similar procedure as in the work of Ref. [92]. Each Q-TLS is characterized by a 3-tuple of fundamental parameters, $(f_{\text{Q-TLS}}, g, \Gamma_1^{\text{Q-TLS}})$. We pick $f_{\text{Q-TLS}}$ uniformly at random from a frequency range relevant to our experiments. Since $f_{\text{q}} \sim 4.5$ GHz, we generate Q-TLSs with $f_{\text{Q-TLS}} \in [4, 5]$ GHz.

In order to generate g , we need a numerical value for both the effective electric dipole moment \tilde{p} ⁶ and $\|\vec{E}_{\text{q}}\|$ at the position of the Q-TLS.

We pick \tilde{p} from a known probability density that has been experimentally measured, e.g., in the work of Ref. [270],

$$f_{\tilde{p}} = \begin{cases} \frac{1}{\tilde{p}} \sqrt{1 - \left(\frac{\tilde{p}}{\tilde{p}_{\text{max}}}\right)^2}, & \text{for } \tilde{p}_{\text{min}} \leq \tilde{p} \leq \tilde{p}_{\text{max}}; \\ 0, & \text{otherwise.} \end{cases} \quad (6.20)$$

⁵The substrate in this work was cleaned before deposition of the aluminum layer; the fabrication recipe is shown in Appendix A.1.

⁶The angle η between \vec{p} and \vec{E}_{q} is integrated in the distribution for \tilde{p} , i.e., $\tilde{p} = \|\vec{p}\| \cos \eta$ [270].

In this equation, we set the minimum and maximum value of \tilde{p} to be $\tilde{p}_{\min} = 0.1$ debye and $\tilde{p}_{\max} = 6$ debye; we choose \tilde{p}_{\max} as in Ref. [92] and \tilde{p}_{\min} assuming that any smaller dipole moment is negligible.

The position of a Q-TLS can be randomly picked at any point within the qubit oxide layers. We may then determine \vec{E}_q at each of these points by means of a conformal mapping technique. This technique allows us to transform the electric field of the qubit capacitor, \vec{E}_q , into the known field of a parallel-plate capacitor. Details on this procedure are given in Appendix D.3.

Finally, we assume that $\Gamma_1^{\text{Q-TLS}} \propto \Delta_0^2$ [269], where the tunneling energy of the Q-TLS, Δ_0 , is picked from an inverse probability distribution. We choose the bounds such that the resulting decay rates range between 1 and 100 MHz, with most rates at the low end of this range.

In order to complete step (1), we need to know the total number of Q-TLSs, $N_{\text{Q-TLS}}$, and their associated 3-tuple parameters. The Q-TLSs are hosted within an interaction region with volume determined by the length of the two CPW segments forming the qubit Al island and the same cross-sectional area used to pick \vec{E}_q (see Appendix D.3), $V_{\text{int}} = 96 \mu\text{m} \times 3 \text{nm} \times 376 \mu\text{m} \times 2$. Given a Q-TLS bandwidth $B_{\text{Q-TLS}} = 1 \text{GHz}$, assuming a Q-TLS density $D = 200 \text{GHz}^{-1} \mu\text{m}^{-3}$ (see Sec. 6.5.2), and disregarding all Q-TLSs with $g < 70 \text{kHz}$, we obtain $N_{\text{Q-TLS}} \sim 570$.

In step (2), each T-TLS is characterized by a 2-tuple of fundamental parameters, $(\delta f^{\mp}, \gamma)$. We generate δf^{\mp} from Eq. (6.17), where Δ and Δ_0 are picked from the GTM distribution of Eq. (6.4). We assume $\mathcal{E}_{\min} = 125 \text{MHz}$, $\mathcal{E}_{\max} = 1 \text{GHz}$, and $\mu = 0.3$ [280]. The interaction energy $U(r)$ is calculated from Eq. (6.5), where $U_0 = k_B \times 10 \text{K nm}^3$ and r is the Q-TLS-T-TLS distance; this distance must be picked at random. Given a cylindrical region with radius r and height t_{ox} centered on the Q-TLS and a uniform T-TLS density, the CDF for the number of T-TLSs is proportional to r^2 . As a consequence, the PDF is linear in r , $f_r \propto r$. We pick r from f_r assuming $r_{\min} = 15 \text{nm}$ and $r_{\max} = 60 \text{nm}$ as bounds (see Sec. 6.5.2 for a discussion on r_{\max}).

We then generate γ from Eq. (6.18). In addition to the parameters used to generate δf^{\mp} , we need $T = 60 \text{mK}$, $\gamma_0 \approx 0.4 \text{Hz}$, $\Omega_0 = 1 \text{GHz}$, $m = 16 \text{u}$, and $d = 2 \text{\AA}$ (see Sec. 6.5.1 for a discussion on the physical meaning of these parameters). Note that the effective qubit temperature $T = 60 \text{mK}$ corresponds to a qubit ground state population of 2.7 %, which is approximately the value observed in our experiments.

Similarly to step (1), in order to complete step (2) we need to select the number of T-TLSs interacting with each Q-TLS, $N_{\text{T-TLS}}$. We generate a set of $N_{\text{T-TLS}} = 10$ T-TLSs, ensuring that each of them additionally fulfills the condition $E_0^+ - E_0^- = \sqrt{E_{\text{T-TLS}}^2 + 4U(\Delta + U)} < E_{\max} = k_B T/2$. We choose half of the thermal energy as our activation threshold, although similar values would work as well.

In step (3), we generate the simulated spectrotemporal charts for Γ_1^q (and, thus, T_1).

Stochastic fluctuations are due to a T-TLS switching state randomly between the left and right well. We simulate these fluctuations as an RTS with a single γ for both the left and right well, i.e., assuming a symmetric noise process. For an RTS, the probability of spending a time t in a certain state is given by the PDF $f_t = \gamma \exp(-\gamma t)$. Starting from a random state, we produce a list of times spent in each T-TLS state until reaching t_{obs} . In order to generate a time series for the T-TLS state, we sample the time list at Δt intervals. The values of both Δt and t_{obs} used in the simulations are the same as for the experiments and are reported in Table 6.1.

The T-TLS state corresponds to a particular δf^\mp . Therefore, as explained in Sec. 6.2.4, the time series $f_{\text{Q-TLS}}(t)$ for each Q-TLS can be calculated by means of Eq. (6.19). Finally, we evaluate Eq. (6.15) for all values of interest of f_{q} ; in order to match the spectrotemporal charts measured in the experiment, we choose f_{q} for the ranges and N_f values reported in Table 6.1.

The simulations are performed using the Julia Programming Language [290]. The computer code `QubitFluctuations.jl` can be obtained from a GitLab repository [291].

6.4 Results

The main results of this work are presented in Fig. 6.1, which shows the experimental and simulated spectrotemporal charts of T_1 . Details on the experiments and simulations are described in Subsecs. 6.3.1 and 6.3.2, respectively, with parameters reported in Table 6.1. Each realization of a simulation is random due to the very nature of the method (because, e.g., $f_{\text{Q-TLS}}$ is distributed uniformly). We thus choose to display simulated spectrotemporal charts that resemble the experiments.

A visual inspection of the T_1 stochastic fluctuations in Fig. 6.1 reveals three distinct spectral-diffusion patterns:

1. Band-limited diffusive.
2. Fast narrowband telegraphic.
3. Slow wideband telegraphic.

Generally, it is also possible to observe combinations of such patterns.

The three patterns can be qualitatively explained by performing ad hoc simulations using a similar method as in Sec. 6.3.2. However, instead of randomly generating the 3- and 2-tuple of steps (1) and (2), we set these tuples by hand. We simulate the effect of several T-TLSs on one Q-TLS, considering three T-TLS sets with different ranges of δf^\mp and γ . For clarity, we choose three Q-TLSs with distinct values of $f_{\text{Q-TLS}}$ — Q-TLS 1, 2, and 3 — one for each set of T-TLSs.

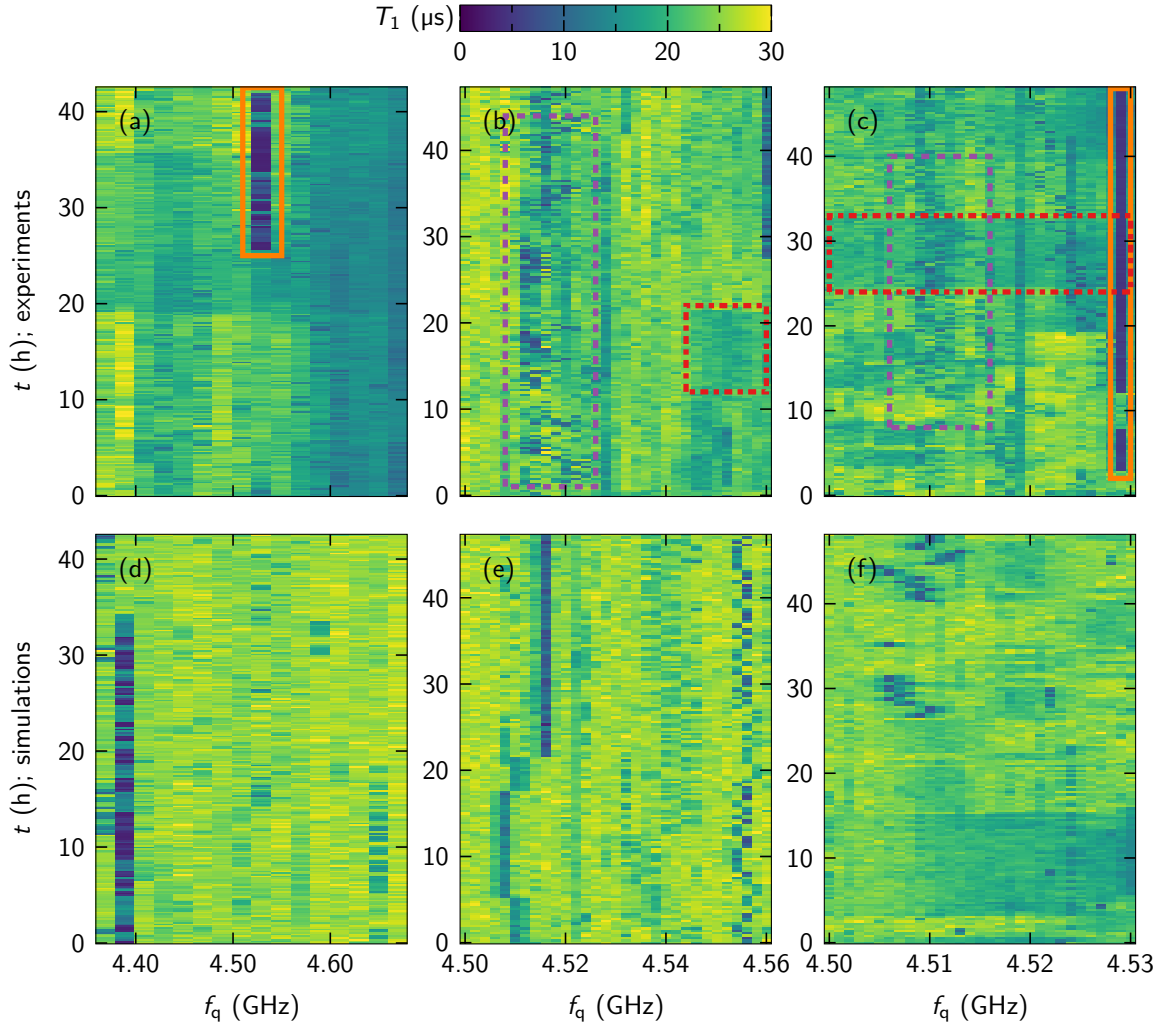


Figure 6.1: Experimental [(a), (b), and (c); datasets 1, 2 and 3, respectively, in Table 6.1] and simulated [(d), (e), and (f)] spectrotemporal charts of T_1 vs. f_q and t , where the panels in each column display an experiment and the corresponding simulation. Spectral-diffusion patterns in the experiments are highlighted with boxes. Band-limited diffusive: dashed purple boxes. Fast narrowband telegraphic: solid orange boxes. Slow wideband telegraphic: dash-dotted red boxes. In the simulations, we add a background time series of Gaussian white noise with a standard deviation of 2 kHz, which is comparable to the fitting error of our T_1 experiments.




Figure 6.2: Three spectral-diffusion patterns. (a) Q-TLS frequency f_{Q-TLS} vs. t for Q-TLS 1 (left purple line), 2 (middle orange line), and 3 (right red line). (b) and (c) Simulated spectrotemporal charts of T_1 vs. f_q and t for $g = 50$ and 100 kHz, respectively. The color map for T_1 is the same as in Fig. 6.1.

In broad strokes, the band-limited diffusive process is reproduced by simulating the effect of many (~ 10) T-TLSs on Q-TLS 1; we select T-TLSs with low values of $1/\gamma$ (ranging between tens of minutes and hours) and small values of δf^\mp (< 1 MHz). The fast narrowband telegraphic process, instead, is generated by considering a few ($\lesssim 3$) T-TLSs acting on Q-TLS 2; in this case, we select high values of $1/\gamma$ (on the order of hours) as well as small values of δf^\mp (< 1 MHz). Similarly to the case of the fast narrowband process, the slow wideband telegraphic process is created assuming also a few ($\lesssim 3$) T-TLSs, this time coupled to Q-TLS 3; in this instance, however, we select very high values of $1/\gamma$ (on the order of days) and large values of δf^\mp (≤ 20 MHz).

Figure 6.2 illustrates the results of the simulation of the three patterns. Figure 6.2 (a) exemplifies the effect of the three different sets of T-TLSs on Q-TLS 1, 2, and 3. Figures 6.2 (b) and (c) demonstrate the impact of each Q-TLS on the spectrotemporal chart of T_1 for a small (a) and large (b) value of g . The T-TLS and Q-TLS parameters used in the simulations are reported in Table 6.2.

Q-TLS 1 is affected by many T-TLSs that switch continuously within the observation time. The T-TLSs act additively on the Q-TLS, resulting in a diffusive shift of f_{Q-TLS} [see Eq. (6.19)]. Different from Brownian diffusion, the shift in f_{Q-TLS} does not exceed the sum of the individual frequency shifts induced by each T-TLS at any observation time. The diffusive process is thus characterized by a limited frequency bandwidth, as shown in Fig. 6.2 (a). The spectrotemporal chart of T_1 displays a similar behavior; T_1 fluctuates

in time over a finite frequency range, exhibiting moderate and strong variations in Figs. 6.2 (b) and (c), respectively.

Q-TLS 2, which is affected by a few T-TLSs, switches mainly between two values of $f_{\text{Q-TLS}}$ (low and high); for both states, much smaller fluctuations at higher switching rates are noticeable. The telegraphic nature of this process affects dramatically the spectrotemporal chart of T_1 when $f_{\text{Q-TLS}} \simeq f_q$. This is the case in the example of Fig. 6.2 (a) when Q-TLS 2 dwells in the low frequency position. In this state, T_1 becomes largely reduced compared to when the Q-TLS resides in the high frequency position, as displayed in Figs. 6.2 (b) and (c). The low value of $\Gamma_1^{\text{Q-TLS}}$ leads to a narrowband process, with more pronounced T_1 variations in Fig. 6.2 (b) compared to Fig. 6.2 (c).

It is worth noting that, in our example, the high frequency position lies between two values of f_q [vertical solid light-gray lines in Fig. 6.2 (a)] but is too far from either of them to significantly impact T_1 . This effect shows that the frequency resolution of our experiments [i.e., the x -axis “pixeling” in Figs. 6.2 (b) and (c)] affects the spectrotemporal chart of T_1 .

Q-TLS 3 behaves similarly to Q-TLS 2, although one of the T-TLSs has a significantly larger value of δf^\mp . Due to low values of γ , Q-TLS 3 undergoes telegraphic frequency shifts only a couple of times during observation. The high value of $\Gamma_1^{\text{Q-TLS}}$ strongly damps the effect on T_1 , resulting in a wideband process. In fact, the effect is barely visible in Fig. 6.2 (b), even when the Q-TLS is almost on resonance with the qubit. In presence of a strong coupling, however, the impact on the spectrotemporal chart of T_1 is clearly identifiable; as shown in Fig. 6.2 (c), the effect extends over a large frequency range.

6.5 Discussion

In this section, we discuss the physical characteristics of a T-TLS (Sec. 6.5.1); we then discuss the density of TLSs (Sec. 6.5.2); finally, we provide insight on the interpretation of the Allan deviation and power spectral density (Sec. 6.5.3).

6.5.1 Physical Characteristics of a T-TLS

The two quantities required to represent T-TLSs in the simulations shown in Fig. 6.1 are δf^\mp and γ of Eqs. (6.17) and (6.18), respectively. The former is determined only by parameters chosen according to the GTM. The latter requires the knowledge of additional physical characteristics of T-TLSs: m and d , as well as Ω_0 ; explicitly,

$$\gamma = \gamma_0 \exp \left[- \left(\frac{\hbar}{d} \right)^2 \frac{1}{2m} \left(\ln \frac{\Omega_0}{\Delta_0} \right)^2 / (k_B T) \right]. \quad (6.21)$$

Table 6.2: T-TLS and Q-TLS parameters used in the simulations of Fig. 6.2.

Q-TLS	$f_{\text{Q-TLS}}$ (GHz)	$\Gamma_1^{\text{Q-TLS}}$ (MHz)	γ (Hz)	δf^{\mp} (MHz)
1	4.510	10	2×10^{-5}	0.9
			5×10^{-5}	0.7
			8×10^{-5}	0.7
			1×10^{-4}	0.6
			2×10^{-4}	0.6
			3×10^{-4}	0.5
			4×10^{-4}	0.3
2	4.531	5	1×10^{-3}	0.1
			3×10^{-5}	0.8
			8×10^{-5}	0.2
3	4.570	90	2×10^{-4}	0.1
			6×10^{-6}	20
			8×10^{-6}	3

The T-TLS mass m must be between that of a very light particle such as an electron and that of heavier elements such as atoms and molecules. That is, it can vary over several orders of magnitude. The interwell distance d should be on the order of angstroms. Electrons and atoms cannot get displaced by more than the interatomic bond length. In the case of molecules, the commonly accepted fluctuation model involves the collective motion of atoms, where each individual atom also cannot move more than the interatomic bond length [285, 252].

In our simulations, we assume a single species of T-TLSs. In order to obtain simulated spectrotemporal charts that resemble the experimental ones, the product d^2m in Eq. (6.21) must lie within one order of magnitude of $1 \times 10^{-45} \text{ m}^2 \text{ kg}$. Considering that d is confined within a few angstroms, the value of m cannot be chosen arbitrarily. If there was clear evidence of multiple T-TLS species characterized by different ranges of γ , they could be modeled assuming different values of m and d . For example, lighter particles would have higher values of γ . Those different species of T-TLS may also be characterized by different parameters determining the various distributions described in Section 6.2. Nonetheless, we determined heuristically that the simulated distributions of γ and δf^{\mp} is sensitive to the parameter choice and that the parameters presented in this work could not be changed significantly without sacrificing closeness to the experimental data.

We assume that TLSs, and thus T-TLSs, are hosted in *oxide* layers at the SM, SA, or MA interfaces (see Sec. 6.1). The oxide layers are composed of molecules with an oxygen (O) atom bound to a pair of neighboring atoms. A T-TLS can be modeled as an O atom with mass $m = 16$ u tunneling between two wells (i.e., states) at a distance d from each other. It is reasonable to assume that d is comparable to the bond length between the O atom and a neighboring atom [269]. In many applications, using Si or sapphire substrates and Al as a metal results in amorphous Si or Al oxide interfacial layers. The bond length between the O and Si or Al atoms is on the order of 2 \AA [292, 293]; this is why in our simulations we choose $d = 2 \text{ \AA}$.

Equation (6.1) is valid only when $V \geq 0$. Accordingly, it must be that $\Omega_0 \geq \Delta_0$ for all values of Δ_0 picked from the GTM distribution. On the one hand, choosing a value $\Omega_0 \sim \Delta_0$ leads to $V \sim 0$, which would correspond to a single- rather than a double-well potential. On the other hand, we cannot choose Ω_0 to be arbitrarily large due to its relationship to γ in Eq. (6.21). In fact, there is a small range of values of Ω_0 that results in a distribution of γ similar to that empirically inferred from the spectrotemporal charts of Fig. 6.1. We choose $\Omega_0 = 1$ GHz to match the experimental range $\gamma \in [10^{-6}, 10^{-2}]$ Hz (i.e., a period from days to minutes) as closely as possible. In this case, we obtain T-TLSs with $V \gtrsim 1.8$ GHz.

6.5.2 Density of TLSs

The TLS density D is estimated by counting the number N of TLSs within a certain interaction region with volume V_{int} and bandwidth B , $D = N/(V_{\text{int}} B)$.

In the case of Q-TLSs, their number $N_{\text{Q-TLS}}$ can be readily obtained by counting the interactions between a qubit and a Q-TLS in spectroscopy experiments [92, 249, 275, 294]. For qubits where Q-TLSs are hosted in a volume of native oxide, the estimated density is $D_{\text{Q-TLS}} \sim 100 \text{ GHz}^{-1} \mu\text{m}^{-3}$. In order to reproduce well our experimental spectrotemporal charts, in the simulations we choose $D_{\text{Q-TLS}} = 200 \text{ GHz}^{-1} \mu\text{m}^{-3}$.

Spectroscopic methods cannot be used to count the number of T-TLSs because, at such low frequencies, the qubit is in an incoherent thermal state. The experimental spectrotemporal charts reveal that Q-TLSs are generally affected by multiple sources of telegraphic noise, as clearly shown by the band-limited diffusive pattern in Fig. 6.1. This observation makes it possible to infer the number of T-TLSs coupled to each Q-TLS, $N_{\text{T-TLS}}$; in the simulations, we choose $N_{\text{T-TLS}} = 10$. These T-TLSs are assumed to be contained inside an interaction region with volume V_{int} centered on their host Q-TLS. It is worth pointing out that our choice of $N_{\text{T-TLS}} = 10$ can still result in both the fast narrowband and slow wideband telegraphic patterns in Fig. 6.1; this is because δf^{\mp} and γ are distributed over a large parameter range possibly leading to a single predominant T-TLS.

The experiment of Fig. 6.1 (c) allows us to resolve T-TLSs with interaction strengths $U(r) \geq 1$ MHz. According to Eq. (6.5), this condition corresponds to a maximum interaction distance $r_{\max} = 60$ nm. Notably, this condition is similar to that hypothesized in the work of Ref. [280]. As explained in Sec. 6.3.2, the T-TLS interaction region is a cylinder with radius r_{\max} and a height of t_{ox} ; the volume associated with this region is $V_{\text{int}} \approx 3.4 \times 10^{-5} \mu\text{m}^3$ ⁷.

Given $B = (E_{\max} - E_{\min})/h = 500$ MHz, we finally get $D_{\text{T-TLS}} \approx 6 \times 10^5 \text{ GHz}^{-1} \mu\text{m}^{-3}$. This value is much larger than $D_{\text{Q-TLS}}$, suggesting that D varies significantly in frequency and is higher at lower frequencies. This finding is in contrast with the typical assumption made by the STM practitioners that TLSs are uniformly distributed in frequency. It is worth noting that a result similar to ours has been recently reported in the work of Ref. [254], although our value for $D_{\text{T-TLS}}$ is even larger than in that work.

6.5.3 On the Interpretation of the Allan Deviation and Power Spectral Density

Time series experiments similar to those reported here are frequently studied by means of statistical analyses such as the Allan deviation (AD) or the power spectral density (PSD), or both. For example, this approach has been pursued in the work of Refs. [261, 260]. It is tempting to ascribe simple models to these statistical estimators

⁷For the Q-TLS density used in our simulations, $D_{\text{Q-TLS}} = 200 \text{ GHz}^{-1} \mu\text{m}^{-3}$, we can find a Q-TLS area density $\sigma_{\text{Q-TLS}} = D_{\text{Q-TLS}} \times 1 \text{ GHz} \times 3 \text{ nm} = 0.6 \mu\text{m}^{-2}$. The average area per Q-TLS is therefore $1/\sigma_{\text{Q-TLS}}$. Assuming each Q-TLS is contained within a square, the radius of the circle inscribed in each square is $r_{\text{Q-TLS}} = \sqrt{1/\sigma_{\text{Q-TLS}}}/2 \approx 600$ nm. Since $r_{\max} \ll r_{\text{Q-TLS}}$, the T-TLS interaction regions do not overlap on average and, thus, we are not double counting T-TLSs.

Table 6.3: Time-series simulation parameters used in Fig. 6.3. The simulations are performed as described in Sec. 6.3.2; however, instead of randomly picking all relevant parameters, we manually specify them. Note that $\gamma = 1/(2\tau_0)$.

$M_{\text{T-TLS}}$	$f_{\text{Q-TLS}}$ (GHz)	g (MHz)	$\Gamma_1^{\text{Q-TLS}}$ (MHz)	γ (μHz)	δf^{\mp} (MHz)
1	4.5011	0.04	15	100	0.6
	4.5011	0.02	10	75	0.8
4	4.5015	0.02	10	70	0.6
	4.4989	0.02	10	140	0.8
	4.4986	0.02	10	75	0.4

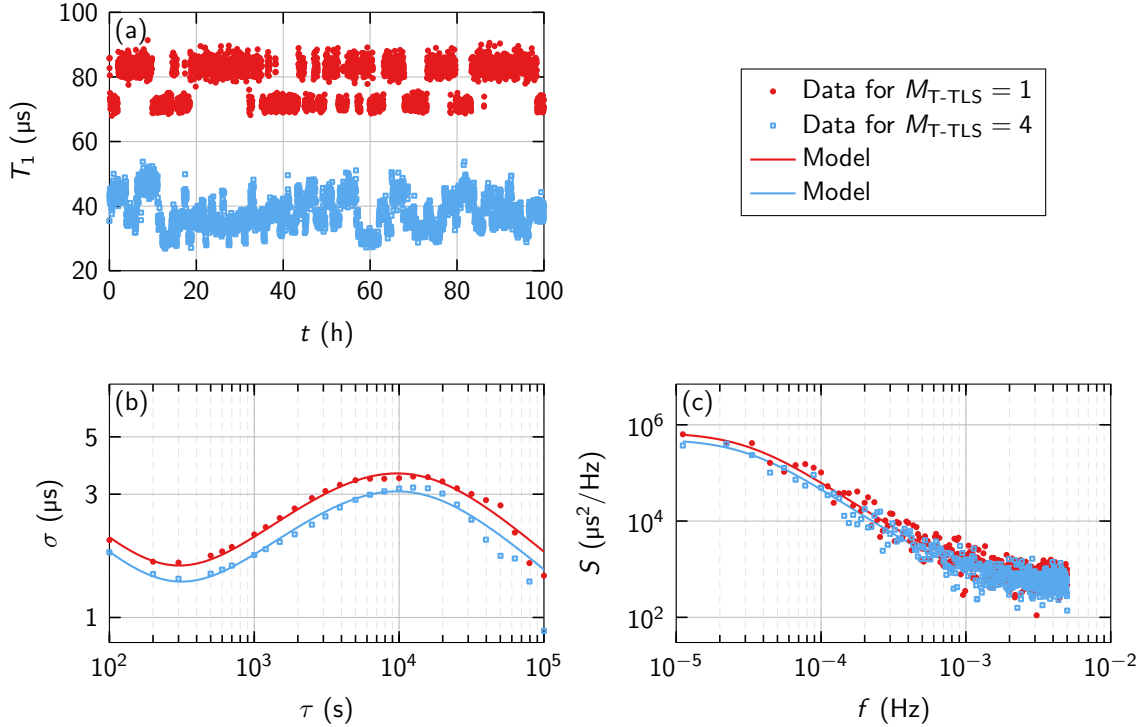


Figure 6.3: Comparison between the statistical analyses of two simulated times series. (a) Simulated time series of T_1 vs. t . The series for $M_{\text{T-TLS}} = 1$ is vertically offset by $40 \mu\text{s}$ for clarity. (b) Estimated overlapping AD σ vs. τ and associated fitting curves from Eq. (6.22). We find $A_0 = 5.84(5)$ and $4.98(6) \mu\text{s}$, $h_0 = 749(111)$ and $583(119) \mu\text{s}^2 \text{Hz}^{-1}$, and $1/\tau_0 = 195(7)$ and $195(10) \mu\text{Hz}$ for $M_{\text{T-TLS}} = 1$ and 4 , respectively. Note that we are fitting σ^2 with the Levenberg–Marquardt algorithm, but plotting σ . The overlapping AD is computed at logarithmically-spaced points. (c) Estimated PSD S vs. f . We use the fitting parameters from (b) to overlay the model of Eq. (6.23) to the data. The PSD is estimated using the Welch’s method with 25 h overlapping segments (rectangular window). The value of τ_0 fitted for $M_{\text{T-TLS}} = 1$ matches (within the confidence interval) that chosen in the simulations and reported in Table 6.3; the fitted τ_0 for $M_{\text{T-TLS}} = 4$, instead, does not match any of the values in Table 6.3.

in order to extract T-TLS parameters such as their switching rate γ and number $M_{\text{T-TLS}}$; in this case, $M_{\text{T-TLS}}$ is the total number of T-TLSs affecting the qubit by interacting with a single or multiple Q-TLSs. It is common, however, to encounter scenarios where these models are misleading.

Figure 6.3 presents two distinct scenarios that illustrate this issue. The time series

in Fig. 6.3 (a) are obtained by simulating one scenario with $M_{\text{T-TLS}} = 1$ and another with $M_{\text{T-TLS}} = 4$. The simulations parameters are reported in Table 6.3. As expected, there is a stark visual difference between the two time series: In the first scenario, it is possible to clearly identify one RTS; this is impossible in the second scenario. However, this difference is not reflected in either the overlapping AD or PSD. In both simulated scenarios, we observe a pronounced peak in the overlapping AD and a lobe in the PSD. These features are indicative of Lorentzian noise. However, they appear to be practically the same for the two scenarios. In fact, it is possible to fit the overlapping AD or PSD using a simple model based on a single source of Lorentzian noise, along with white noise. The model reads

$$\sigma^2 = \frac{h_0}{2\tau} + \left(\frac{A_0\tau_0}{\tau}\right)^2 \left(4e^{-\tau/\tau_0} - e^{-2\tau/\tau_0} - 3 + \frac{2\tau}{\tau_0}\right) \quad (6.22)$$

for the AD and

$$S = h_0 + \frac{4A_0^2\tau_0}{1 + (2\pi f\tau_0)^2} \quad (6.23)$$

for the PSD, where τ and f are the analysis interval and frequency, h_0 and A_0 are the white and Lorentzian noise amplitudes, and τ_0 is the Lorentzian characteristic time [295].

Although the two simulated time series are associated with entirely different scenarios, the simple models of Eqs. (6.22) and (6.23) fit accurately both the overlapping AD and PSD for very similar values of τ_0 ; we obtain $1/\tau_0 = 195(7)$ μHz when $M_{\text{T-TLS}} = 1$ and $1/\tau_0 = 195(10)$ μHz when $M_{\text{T-TLS}} = 4$. This conclusion can be qualitatively understood by noticing that multiple physical sources of Lorentzian noise combine to form a single wideband peak in the overlapping AD (or lobe in the PSD). As a consequence, this feature can be mistakenly fitted with a model comprising a single Lorentzian term. For this reason, we elect *not to analyze* our experimental results by ascribing simple models to the AD (or PSD).

6.6 Conclusions

We study the physics of TLSs by means of a frequency-tunable planar superconducting qubit. We show that simulations based on the TLS interacting model (or GTM) can explain the spectrotemporal charts of T_1 observed in the experiments over long time periods. We find that the density of T-TLSs is much larger than that of Q-TLSs, meaning TLSs are nonuniformly distributed over large frequency bandwidths. Our finding corroborates the results reported in the work of Ref. [254].

Our experiments demonstrate that the additional dimension provided by frequency tunability makes tunable qubits a better probe to study spectral diffusion compared

to fixed-frequency devices. Hence, we suggest that future work on TLS stochastic fluctuations should explore even wider frequency bandwidths. A large bandwidth would increase the chances to encounter a scenario where a pair of Q-TLSs interacts with a single T-TLS, resulting in a synchronous fluctuation of the two Q-TLSs. Such an experiment would conclusively prove the validity of the TLS–TLS interaction hypothesis in the GTM.

It is well-known that external strain or electric fields applied to a qubit chip modify the Q-TLSs' characteristic energies, Δ or Δ_0 , or both [275]. Therefore, we suggest to apply external fields while exploring long time qubit fluctuations. Such an experiment may make it possible to indirectly observe a similar change in the characteristic energies of the T-TLSs. In fact, both Δ and Δ_0 contribute to changes in δf^\mp , whereas γ is affected only by Δ_0 . In principle, this procedure would allow us to perform an indirect spectroscopic study of T-TLSs as a function of external fields.

It is also worth noting that recent advances on the coupling of superconducting devices to bulk acoustic waves [296] may pave the way to the acoustic characterization of TLS-induced qubit loss and fluctuations.

Lastly, we expect that performing experiments at different operating temperatures would provide one more knob to modify the frequency bandwidth of thermally activated TLSs. This approach would allow us to characterize the TLS density for different frequency ranges. In particular, it would be expected that the amount of thermally activated would decrease with decreasing temperature, thus leading to smaller fluctuations in T_1 .

Conclusion

It is hopefully now clear that the advances presented in this thesis were and still are part of three major challenges currently affecting circuit QED devices, those being control scalability, calibration, and decoherence. We also want to remind that, though the mainstream media largely focuses on the quantum computing aspect of cQED, the work of this thesis is applicable to many other purposes, including fundamental physics. Indeed, cQED is an extremely powerful platform for light-matter research, which is relevant to a multitude of physical phenomena, such as photosynthesis [297].

Chapter 4 addresses the scalability challenge inherent to the planar layout of superconducting circuits. Above a certain number of qubits, control lines cannot be routed on-chip from the edges. An additional dimension is required. The quantum socket is a true 3D wiring solution, as opposed to a “stacked chips” methods where the routing is done on a different substrate located above or below the main circuit substrate, but still has the wiring inputs coming from the edge [225, 298, 299]. The latest work by groups with large number of qubits still makes use of this non-scalable wiring method [96, 97], it is therefore clear that a solution is still needed, and it will likely be based on a miniaturized version of the quantum socket. To our knowledge, the only other true 3D wiring method was proposed by the Leek group, in Oxford [300].

Chapter 5 proposes two efficient methods to calibrate resonant couplings between various systems in circuit QED. These methods become essential as the size of circuits grows because the number of couplings between qubits, resonators, and often coherent TLSs grows with it. The need for such methods is evident when reading about the sophisticated calibrations of large-scale experiments, e.g., see Ref. [268] and the supplementary material of Refs. [96, 97].

Finally, Chapter 6 explores what is probably the most urgent problem with superconducting architectures: decoherence caused by two-level systems [301]. We provide strong evidence that interactions between low and high-energy TLS are the root cause

of the large coherence time fluctuations observed in qubits. Solving this problem is a hard challenge, because the true origin of those defects is still uncertain. In all likelihood, better fabrication methods will be needed to eliminate the various physical systems that may be responsible for TLS in amorphous dielectrics. Some success has already been achieved, as in Ref. [171], which cleaned electric spins on the surface of a sample, leading to a reduction in noise.

Bibliography

- [1] A. Kenny, *Ancient Philosophy, A New History of Western Philosophy*, Vol. 1 (Oxford University Press, Oxford, 2004).
- [2] E. M. Purcell, Spontaneous Emission Probabilities at Radio Frequencies (B10), in *Proceedings of the American Physical Society*, Vol. 69 (American Physical Society, 1946) p. 681.
- [3] K. H. Drexhage, Interaction of Light with Monomolecular Dye Layers, in *Progress in Optics*, Vol. 12, edited by E. Wolf (Elsevier, 1974) pp. 163–232.
- [4] M. Brune, F. Schmidt-Kaler, A. Maali, J. Dreyer, E. Hagley, J. M. Raimond, and S. Haroche, Quantum Rabi Oscillation: A Direct Test of Field Quantization in a Cavity, *Physical Review Letters* **76**, 1800–1803 (1996).
- [5] M. Brune, E. Hagley, J. Dreyer, X. Maître, A. Maali, C. Wunderlich, J. M. Raimond, and S. Haroche, Observing the Progressive Decoherence of the “Meter” in a Quantum Measurement, *Physical Review Letters* **77**, 4887–4890 (1996).
- [6] I. Georgescu, Nobel Prize 2012: Haroche & Wineland, *Nature Physics* **8**, 777–777 (2012).
- [7] D. van Delft and P. Kes, The discovery of superconductivity, *Physics Today* **63**, 38–43 (2010).
- [8] S. Reif-Acherman, Heike Kamerlingh Onnes and the Nobel Prize in Physics for 1913: The Highest Honor for the Lowest Temperatures, *Physics in Perspective* **15**, 415–450 (2013).
- [9] W. Meissner and R. Ochsenfeld, Ein neuer Effekt bei Eintritt der Supraleitfähigkeit, *Naturwissenschaften* **21**, 787–788 (1933).

BIBLIOGRAPHY

- [10] J. Bardeen, L. N. Cooper, and J. R. Schrieffer, Microscopic Theory of Superconductivity, [Physical Review](#) **106**, 162–164 (1957).
- [11] J. Bardeen, L. N. Cooper, and J. R. Schrieffer, Theory of Superconductivity, [Physical Review](#) **108**, 1175–1204 (1957).
- [12] L. N. Cooper, Bound Electron Pairs in a Degenerate Fermi Gas, [Physical Review](#) **104**, 1189–1190 (1956).
- [13] L. Hoddeson, John Bardeen: The only person to win two Nobel Prizes in physics, [Physics Education](#) **46**, 661–668 (2011).
- [14] B. Josephson, Possible new effects in superconductive tunnelling, [Physics Letters](#) **1**, 251–253 (1962).
- [15] B. D. Josephson, The discovery of tunnelling supercurrents, [Reviews of Modern Physics](#) **46**, 251–254 (1974).
- [16] P. W. Anderson and J. M. Rowell, Probable Observation of the Josephson Superconducting Tunneling Effect, [Physical Review Letters](#) **10**, 230–232 (1963).
- [17] J. M. Rowell, Magnetic Field Dependence of the Josephson Tunnel Current, [Physical Review Letters](#) **11**, 200–202 (1963).
- [18] S. Shapiro, Josephson Currents in Superconducting Tunneling: The Effect of Microwaves and Other Observations, [Physical Review Letters](#) **11**, 80–82 (1963).
- [19] H. M. Konczykowski, F. Holtzberg, and P. Lejay, Physics Nobel prize 1973, [Physics Bulletin](#) **24**, 713–713 (1973).
- [20] R. C. Jaklevic, J. Lambe, A. H. Silver, and J. E. Mercereau, Quantum Interference Effects in Josephson Tunneling, [Physical Review Letters](#) **12**, 159–160 (1964).
- [21] J. Clarke and A. I. Braginski, eds., [Fundamentals and Technology of SQUIDs and SQUID Systems](#), The SQUID Handbook, Vol. 1 (Wiley, Weinheim, Germany, 2004).
- [22] M. Tinkham, [Introduction to Superconductivity](#), 2nd ed. (Dover Publications, Mineola, New York, 2004).
- [23] R. Voss and R. Webb, Macroscopic Quantum Tunneling in 1-Mm Nb Josephson Junctions, [Physical Review Letters](#) **47**, 265–268 (1981).
- [24] J. M. Martinis, M. H. Devoret, and J. Clarke, Energy-Level Quantization in the Zero-Voltage State of a Current-Biased Josephson Junction, [Physical Review Letters](#) **55**, 1543–1546 (1985).

BIBLIOGRAPHY

- [25] D. P. DiVincenzo, Quantum Computation, *Science* **270**, 255–261 (1995).
- [26] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th ed. (Cambridge University Press, 2010).
- [27] D. P. DiVincenzo, The Physical Implementation of Quantum Computation, *Fortschritte der Physik* **48**, 771–783 (2000).
- [28] N. D. Mermin, *Quantum Computer Science: An Introduction* (Cambridge University Press, Cambridge, 2007).
- [29] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O’Brien, Quantum computers, *Nature* **464**, 45–53 (2010).
- [30] C. A. Pérez-Delgado and P. Kok, Quantum computers: Definition and implementations, *Physical Review A* **83**, 012303 (2011).
- [31] A. Montanaro, Quantum algorithms: An overview, *npj Quantum Information* **2**, 15023 (2016).
- [32] P. Benioff, The computer as a physical system: A microscopic quantum mechanical Hamiltonian model of computers as represented by Turing machines, *Journal of Statistical Physics* **22**, 563–591 (1980).
- [33] P. Benioff, Quantum mechanical hamiltonian models of turing machines, *Journal of Statistical Physics* **29**, 515–546 (1982).
- [34] A. Peres, Reversible logic and quantum computers, *Physical Review A* **32**, 3266–3276 (1985).
- [35] R. P. Feynman, Simulating physics with computers, *International Journal of Theoretical Physics* **21**, 467–488 (1982).
- [36] D. Deutsch, Quantum theory, the Church–Turing principle and the universal quantum computer, *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* **400**, 97–117 (1985).
- [37] R. P. Feynman, Quantum Mechanical Computers, *Optics News* **11**, 11–20 (1985).
- [38] D. Deutsch, Quantum Computational Networks, *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* **425**, 73–90 (1989).
- [39] A. Chi-Chih Yao, Quantum circuit complexity, in *Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science* (IEEE, Palo Alto, CA, USA, 1993) pp. 352–361.

BIBLIOGRAPHY

- [40] D. Deutsch and R. Jozsa, Rapid Solution of Problems by Quantum Computation, [Proceedings: Mathematical and Physical Sciences](#) **439**, 553–558 (1992).
- [41] P. Shor, Algorithms for quantum computation: Discrete logarithms and factoring, in [Proceedings 35th Annual Symposium on Foundations of Computer Science](#) (IEEE, Santa Fe, NM, USA, 1994) pp. 124–134.
- [42] P. W. Shor, Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer, [SIAM Journal on Computing](#) **26**, 1484–1509 (1997).
- [43] R. L. Rivest, A. Shamir, and L. Adleman, A method for obtaining digital signatures and public-key cryptosystems, [Communications of the ACM](#) **21**, 120–126 (1978).
- [44] W. Diffie and M. Hellman, New directions in cryptography, [IEEE Transactions on Information Theory](#) **22**, 644–654 (1976).
- [45] D. Coppersmith, [An Approximate Fourier Transform Useful in Quantum Factoring](#), Tech. Rep. RC 19642 (IBM Research Division, New York, NY, USA, 1994) [arXiv:quant-ph/0201067](#) .
- [46] L. K. Grover, A fast quantum mechanical algorithm for database search, in [Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing](#) (ACM Press, Philadelphia, Pennsylvania, United States, 1996) pp. 212–219.
- [47] L. K. Grover, Quantum Mechanics Helps in Searching for a Needle in a Haystack, [Physical Review Letters](#) **79**, 325–328 (1997).
- [48] A. Ambainis, Quantum search algorithms, [ACM SIGACT News](#) **35**, 22–35 (2004).
- [49] J. Watrous, [The Theory of Quantum Information](#), 1st ed. (Cambridge University Press, Cambridge, 2018).
- [50] P. Kaye, R. Laflamme, and M. Mosca, [An Introduction to Quantum Computing](#), 1st ed. (Oxford University Press, Oxford, 2006).
- [51] W. K. Wootters and W. H. Zurek, A single quantum cannot be cloned, [Nature](#) **299**, 802–803 (1982).
- [52] P. W. Shor, Scheme for reducing decoherence in quantum computer memory, [Physical Review A](#) **52**, R2493–R2496 (1995).
- [53] A. Steane, Multiple-Particle Interference and Quantum Error Correction, [Proceedings: Mathematical, Physical and Engineering Sciences](#) **452**, 2551–2577 (1996).

BIBLIOGRAPHY

- [54] R. Laflamme, C. Miquel, J. P. Paz, and W. H. Zurek, Perfect Quantum Error Correcting Code, [Physical Review Letters](#) **77**, 198–201 (1996).
- [55] E. Knill and R. Laflamme, Theory of quantum error-correcting codes, [Physical Review A](#) **55**, 900–911 (1997).
- [56] A. Y. Kitaev, Fault-tolerant quantum computation by anyons, [Annals of Physics](#) **303**, 2–30 (2003).
- [57] D. Gottesman, An introduction to quantum error correction and fault-tolerant quantum computation, in [Quantum Information Science and Its Contributions to Mathematics](#) (American Mathematical Society, 2010) pp. 13–58, [arXiv:0904.2557](#) .
- [58] R. Raussendorf and J. Harrington, Fault-Tolerant Quantum Computation with High Threshold in Two Dimensions, [Physical Review Letters](#) **98**, 190504 (2007).
- [59] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, Surface codes: Towards practical large-scale quantum computation, [Physical Review A](#) **86**, 032324 (2012).
- [60] C. Horsman, A. G. Fowler, S. Devitt, and R. V. Meter, Surface code quantum computing by lattice surgery, [New Journal of Physics](#) **14**, 123011 (2012).
- [61] K. Kechedzhi, V. Smelyanskiy, J. R. McClean, V. S. Denchev, M. Mohseni, S. Isakov, S. Boixo, B. Altshuler, and H. Neven, Efficient Population Transfer via Non-Ergodic Extended States in Quantum Spin Glass, in [Leibniz International Proceedings in Informatics \(LIPIcs\)](#), Vol. 111, edited by S. Jeffery (Schloss Dagstuhl, Dagstuhl, Germany, 2018) pp. 9:1–9:16, [arXiv:1807.04792](#) .
- [62] R. D. Somma, S. Boixo, H. Barnum, and E. Knill, Quantum Simulations of Classical Annealing Processes, [Physical Review Letters](#) **101**, 130504 (2008).
- [63] E. Farhi and H. Neven, Classification with Quantum Neural Networks on Near Term Processors, [arXiv:1802.06002](#) (2018).
- [64] J. R. McClean, S. Boixo, V. N. Smelyanskiy, R. Babbush, and H. Neven, Barren plateaus in quantum neural network training landscapes, [Nature Communications](#) **9**, 4812 (2018).
- [65] I. Cong, S. Choi, and M. D. Lukin, Quantum convolutional neural networks, [Nature Physics](#) **15**, 1273–1278 (2019).

BIBLIOGRAPHY

- [66] S. Bravyi, D. Gosset, and R. König, Quantum advantage with shallow circuits, *Science* **362**, 308–311 (2018).
- [67] A. Aspuru-Guzik, Simulated Quantum Computation of Molecular Energies, *Science* **309**, 1704–1707 (2005).
- [68] A. Peruzzo, J. McClean, P. Shadbolt, M.-H. Yung, X.-Q. Zhou, P. J. Love, A. Aspuru-Guzik, and J. L. O’Brien, A variational eigenvalue solver on a photonic quantum processor, *Nature Communications* **5**, 4213 (2014).
- [69] C. Hempel, C. Maier, J. Romero, J. McClean, T. Monz, H. Shen, P. Jurcevic, B. P. Lanyon, P. Love, R. Babbush, A. Aspuru-Guzik, R. Blatt, and C. F. Roos, Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator, *Physical Review X* **8**, 031022 (2018).
- [70] P. Kok, W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn, Linear optical quantum computing with photonic qubits, *Reviews of Modern Physics* **79**, 135–174 (2007).
- [71] J. L. O’Brien, A. Furusawa, and J. Vučković, Photonic quantum technologies, *Nature Photonics* **3**, 687–695 (2009).
- [72] C. Monroe and J. Kim, Scaling the Ion Trap Quantum Processor, *Science* **339**, 1164–1169 (2013).
- [73] D. G. Cory, Spin Based Test-beds for Quantum Information Processing, in *Frontiers in Optics* (OSA, Rochester, New York, 2006) p. LTu13.
- [74] N. C. Jones, R. Van Meter, A. G. Fowler, P. L. McMahon, J. Kim, T. D. Ladd, and Y. Yamamoto, Layered Architecture for Quantum Computing, *Physical Review X* **2**, 031007 (2012).
- [75] R. Hanson, L. P. Kouwenhoven, J. R. Petta, S. Tarucha, and L. M. K. Vandersypen, Spins in few-electron quantum dots, *Reviews of Modern Physics* **79**, 1217–1265 (2007).
- [76] B. M. Maune, M. G. Borselli, B. Huang, T. D. Ladd, P. W. Deelman, K. S. Holabird, A. A. Kiselev, I. Alvarado-Rodriguez, R. S. Ross, A. E. Schmitz, M. Sokolich, C. A. Watson, M. F. Gyure, and A. T. Hunter, Coherent singlet-triplet oscillations in a silicon-based double quantum dot, *Nature* **481**, 344–347 (2012).
- [77] F. A. Zwanenburg, A. S. Dzurak, A. Morello, M. Y. Simmons, L. C. L. Hollenberg, G. Klimeck, S. Rogge, S. N. Coppersmith, and M. A. Eriksson, Silicon quantum electronics, *Reviews of Modern Physics* **85**, 961–1019 (2013).

BIBLIOGRAPHY

- [78] A. P. Higginbotham, F. Kuemmeth, M. P. Hanson, A. C. Gossard, and C. M. Marcus, Coherent Operations and Screening in Multielectron Spin Qubits, [Physical Review Letters](#) **112**, 026801 (2014).
- [79] J. O’Gorman, N. H. Nickerson, P. Ross, J. J. Morton, and S. C. Benjamin, A silicon-based surface code quantum computer, [npj Quantum Information](#) **2**, 15019 (2016).
- [80] Y. Nakamura, C. Chen, and J. Tsai, Spectroscopy of Energy-Level Splitting between Two Macroscopic Quantum States of Charge Coherently Superposed by Josephson Coupling, [Physical Review Letters](#) **79**, 2328–2331 (1997).
- [81] V. Bouchiat, D. Vion, P. Joyez, D. Esteve, and M. H. Devoret, Quantum Coherence with a Single Cooper Pair, [Physica Scripta](#) **T76**, 165 (1998).
- [82] Y. Nakamura, Y. Pashkin, and J. Tsai, Coherent control of macroscopic quantum states in a single-Cooper-pair box, [Nature](#) **398**, 786–788 (1999).
- [83] D. Vion, A. Aassime, A. Cottet, P. Joyez, H. Pothier, C. Urbina, D. Esteve, and M. H. Devoret, Manipulating the Quantum State of an Electrical Circuit, [Science](#) **296**, 886–889 (2002).
- [84] G. Ithier, E. Collin, P. Joyez, P. J. Meeson, D. Vion, D. Esteve, F. Chiarello, A. Shnirman, Y. Makhlin, J. Schrieffer, and G. Schön, Decoherence in a superconducting quantum bit circuit, [Physical Review B](#) **72**, 134519 (2005).
- [85] J. E. Mooij, T. P. Orlando, L. Levitov, L. Tian, C. H. van der Wal, and S. Lloyd, Josephson Persistent-Current Qubit, [Science](#) **285**, 1036–1039 (1999).
- [86] T. P. Orlando, J. E. Mooij, L. Tian, C. H. van der Wal, L. S. Levitov, S. Lloyd, and J. J. Mazo, Superconducting persistent-current qubit, [Physical Review B](#) **60**, 15398–15413 (1999).
- [87] J. R. Friedman, V. Patel, W. Chen, S. K. Tolpygo, and J. E. Lukens, Quantum superposition of distinct macroscopic states, [Nature](#) **406**, 43–46 (2000).
- [88] C. H. van der Wal, A. C. J. ter Haar, F. K. Wilhelm, R. N. Schouten, C. J. P. M. Harmans, T. P. Orlando, S. Lloyd, and J. E. Mooij, Quantum Superposition of Macroscopic Persistent-Current States, [Science](#) **290**, 773–777 (2000).
- [89] J. M. Martinis, S. Nam, J. Aumentado, and C. Urbina, Rabi Oscillations in a Large Josephson-Junction Qubit, [Physical Review Letters](#) **89**, 117901 (2002).

BIBLIOGRAPHY

- [90] J. Clarke and F. K. Wilhelm, Superconducting quantum bits, *Nature* **453**, 1031–1042 (2008).
- [91] M. H. Devoret and R. J. Schoelkopf, Superconducting Circuits for Quantum Information: An Outlook, *Science* **339**, 1169–1174 (2013).
- [92] R. Barends, J. S. Kelly, A. Megrant, D. Sank, E. Jeffrey, Y. Chen, Y. Yin, B. Chiaro, J. Mutus, C. Neill, P. O’Malley, P. Roushan, J. Wenner, T. C. White, A. N. Cleland, and J. M. Martinis, Coherent Josephson Qubit Suitable for Scalable Quantum Integrated Circuits, *Physical Review Letters* **111**, 080502 (2013).
- [93] A. Nersisyan, S. Poletto, N. Alidoust, R. Manenti, R. Renzas, C.-V. Bui, K. Vu, T. Whyland, Y. Mohan, E. A. Sete, S. Stanwyck, A. Bestwick, and M. Reagor, Manufacturing low dissipation superconducting quantum processors, in *2019 IEEE International Electron Devices Meeting (IEDM)*, Vol. 94710 (2019) pp. 31.1.1–31.1.4, arXiv:1901.08042 .
- [94] A. P. M. Place, L. V. H. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman, G. Cheng, T. Madhavan, H. K. Babla, X. H. Le, Y. Gang, B. Jäck, A. Gyenis, N. Yao, R. J. Cava, N. P. de Leon, and A. A. Houck, New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds, *Nature Communications* **12**, 1779 (2021).
- [95] C. Macklin, K. O’Brien, D. Hover, M. E. Schwartz, V. Bolkhovskiy, X. Zhang, W. D. Oliver, and I. Siddiqi, A near-quantum-limited Josephson traveling-wave parametric amplifier, *Science* **350**, 307–310 (2015).
- [96] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandão, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. S. Kelly, P. V. Klimov, S. Knysh, A. N. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, Quantum supremacy using a programmable superconducting processor, *Nature* **574**, 505–510 (2019).

BIBLIOGRAPHY

- [97] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong, C. Guo, C. Guo, S. Guo, L. Han, L. Hong, H.-L. Huang, Y.-H. Huo, L. Li, N. Li, S. Li, Y. Li, F. Liang, C. Lin, J. Lin, H. Qian, D. Qiao, H. Rong, H. Su, L. Sun, L. Wang, S. Wang, D. Wu, Y. Xu, K. Yan, W. Yang, Y. Yang, Y. Ye, J. Yin, C. Ying, J. Yu, C. Zha, C. Zhang, H. Zhang, K. Zhang, Y. Zhang, H. Zhao, Y. Zhao, L. Zhou, Q. Zhu, C.-Y. Lu, C.-Z. Peng, X. Zhu, and J.-W. Pan, Strong Quantum Computational Advantage Using a Superconducting Quantum Processor, [Physical Review Letters](#) **127**, 180501 (2021).
- [98] J. Preskill, Quantum Computing in the NISQ era and beyond, [Quantum](#) **2**, 79 (2018).
- [99] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, Superconducting Qubits: Current State of Play, [Annual Review of Condensed Matter Physics](#) **11**, 369–395 (2020).
- [100] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, A quantum engineer’s guide to superconducting qubits, [Applied Physics Reviews](#) **6**, 021318 (2019).
- [101] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, Circuit quantum electrodynamics, [Reviews of Modern Physics](#) **93**, 025005 (2021).
- [102] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation, [Physical Review A](#) **69**, 062320 (2004).
- [103] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics., [Nature](#) **431**, 162–7 (2004).
- [104] J. Koch, T. Yu, J. M. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Charge-insensitive qubit design derived from the Cooper pair box, [Physical Review A](#) **76**, 042319 (2007).
- [105] A. Megrant, C. Neill, R. Barends, B. Chiaro, Y. Chen, L. Feigl, J. S. Kelly, E. Lucero, M. Mariantoni, P. J. J. O’Malley, D. Sank, A. Vainsencher, J. Wenner, T. C. White, Y. Yin, J. Zhao, C. J. Palmstrom, J. M. Martinis, and A. N. Cleland, Planar superconducting resonators with internal quality factors above one million, [Applied Physics Letters](#) **100**, 113510 (2012).

BIBLIOGRAPHY

- [106] R. Barends, J. S. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, Superconducting quantum circuits at the surface code threshold for fault tolerance, [Nature](#) **508**, 500–503 (2014).
- [107] A. D. Córcoles, E. Magesan, S. J. Srinivasan, A. W. Cross, M. Steffen, J. M. Gambetta, and J. M. Chow, Demonstration of a quantum error detection code using a square lattice of four superconducting qubits, [Nature Communications](#) **6**, 6979 (2015).
- [108] J. S. Kelly, *Fault-Tolerant Superconducting Qubits*, [Ph.D. thesis](#), UC Santa Barbara, Santa Barbara, California, USA (2015).
- [109] C. K. Andersen, A. Remm, S. Lazar, S. Krinner, J. Heinsoo, J.-C. Besse, M. Gabureac, A. Wallraff, and C. Eichler, Entanglement stabilization using ancilla-based parity detection and real-time feedback in superconducting circuits, [npj Quantum Information](#) **5**, 69 (2019).
- [110] M. Göppl, A. Fragner, M. Baur, R. Bianchetti, S. Filipp, J. M. Fink, P. J. Leek, G. Puebla, L. Steffen, and A. Wallraff, Coplanar waveguide resonators for circuit quantum electrodynamics, [Journal of Applied Physics](#) **104**, 113904 (2008).
- [111] B. A. Mazin, *Microwave Kinetic Inductance Detectors*, [Ph.D. thesis](#), California Institute of Technology (2004).
- [112] J. Gao, *The Physics of Superconducting Microwave Resonators*, [Ph.D. thesis](#), California Institute of Technology (2008).
- [113] R. Barends, *Photon-Detecting Superconducting Resonators*, [Ph.D. thesis](#), TU Delft, Delft-Leiden, Netherlands (2009).
- [114] C. T. Earnest, J. H. Béjanin, T. G. McConkey, E. A. Peters, A. Korinek, H. Yuan, and M. Mariantoni, Substrate surface engineering for high-quality silicon/aluminum superconducting resonators, [Superconductor Science and Technology](#) **31**, 125013 (2018).
- [115] D. M. Pozar, *Microwave Engineering*, 4th ed. (Wiley, Hoboken, NJ, 2011).
- [116] R. N. Simons, *Coplanar Waveguide Circuits, Components, and Systems* (John Wiley & Sons, Inc., New York, USA, 2001).
- [117] M. H. Devoret, Quantum fluctuations in electrical circuits, in *Les Houches, Session LXIII*, edited by S. Reynaud, E. Giacobino, and J. Zinn-Justin (North Holland/Elsevier, Amsterdam, 1997) 1st ed., pp. 351–386.

BIBLIOGRAPHY

- [118] P. A. M. Dirac, *The Principles of Quantum Mechanics*, 4th ed., The International Series of Monographs on Physics (Clarendon Press, Oxford, 1958).
- [119] J. A. Schreier, A. A. Houck, J. Koch, D. I. Schuster, B. R. Johnson, J. M. Chow, J. M. Gambetta, J. Majer, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Suppressing charge noise decoherence in superconducting charge qubits, [Physical Review B **77**, 180502 \(2008\)](#).
- [120] V. Ambegaokar and A. Baratoff, Tunneling Between Superconductors, [Physical Review Letters **10**, 486–489 \(1963\)](#).
- [121] N. A. Court, A. J. Ferguson, and R. G. Clark, Energy gap measurement of nanostructured aluminium thin films for single Cooper-pair devices, [Superconductor Science and Technology **21**, 015013 \(2007\)](#).
- [122] D. I. Schuster, *Circuit Quantum Electrodynamics*, [Ph.D. thesis](#), Yale (2007).
- [123] J. B. Ketterson, *The Physics of Solids* (Oxford University Press, Oxford, 2016).
- [124] R. H. Koch, D. P. DiVincenzo, and J. Clarke, Model for 1/f Flux Noise in SQUIDs and Qubits, [Physical Review Letters **98**, 267003 \(2007\)](#).
- [125] A. Cottet, *Implementation of a quantum bit in a superconducting qubit*, [Ph.D. thesis](#), Université Pierre et Marie Curie - Paris VI, Paris (2002), [arXiv:1011.1669v3](#).
- [126] F. W. J. Olver, A. B. Olde Daalhuis, D. W. Lozier, B. I. Schneider, R. F. Boisvert, C. W. Clark, B. R. Miller, H. S. C. B. V. Saunders, and M. A. McClain, [NIST Digital Library of Mathematical Functions](#), <https://dlmf.nist.gov/> (2017).
- [127] J. H. Béjanin, [Mathieu.jl](#), <https://github.com/jebej/Mathieu.jl> (2017).
- [128] N. Didier, E. A. Sete, M. P. da Silva, and C. T. Rigetti, Analytical modeling of parametrically modulated transmon qubits, [Physical Review A **97**, 022330 \(2018\)](#).
- [129] L. O. Chua, C. A. Desoer, and E. S. Kuh, *Linear and Nonlinear Circuits* (McGraw-Hill, 1987).
- [130] M. Mariani, F. Deppe, A. Marx, R. Gross, F. K. Wilhelm, and E. Solano, Two-resonator circuit quantum electrodynamics: A superconducting quantum switch, [Physical Review B **78**, 104508 \(2008\)](#).
- [131] U. Vool and M. H. Devoret, Introduction to quantum electromagnetic circuits, [International Journal of Circuit Theory and Applications **45**, 897–934 \(2017\)](#).

BIBLIOGRAPHY

- [132] J. R. Schrieffer and P. A. Wolff, Relation between the Anderson and Kondo Hamiltonians, *Physical Review* **149**, 491–492 (1966).
- [133] S. Bravyi, D. P. DiVincenzo, and D. Loss, Schrieffer–Wolff transformation for quantum many-body systems, *Annals of Physics* **326**, 2793–2826 (2011).
- [134] B. W. Shore and P. L. Knight, The Jaynes-Cummings Model, *Journal of Modern Optics* **40**, 1195–1238 (1993).
- [135] M. Orszag, *Quantum Optics* (Springer, Berlin, Heidelberg, 2008).
- [136] D. A. Steck, *Quantum and Atom Optics* (Self-published, Eugene, Oregon, 2019) available online at <http://steck.us/teaching>.
- [137] A. Blais, J. M. Gambetta, A. Wallraff, D. I. Schuster, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, Quantum-information processing with circuit quantum electrodynamics, *Physical Review A* **75**, 032329 (2007).
- [138] R. Raussendorf and H. J. Briegel, A One-Way Quantum Computer, *Physical Review Letters* **86**, 5188–5191 (2001).
- [139] R. Raussendorf, D. E. Browne, and H. J. Briegel, Measurement-based quantum computation on cluster states, *Physical Review A* **68**, 022312 (2003).
- [140] A. D. King, J. Carrasquilla, I. Ozfidan, J. Raymond, E. Andriyash, A. Berkley, M. Reis, T. M. Lanting, R. Harris, G. Poulin-Lamarre, A. Y. Smirnov, C. Rich, F. Altomare, P. Bunyk, J. Whittaker, L. Swenson, E. Hoskinson, Y. Sato, M. Volkman, E. Ladizinsky, M. Johnson, J. Hilton, and M. H. Amin, Observation of topological phenomena in a programmable lattice of 1,800 qubits, *Nature* [10.1038/s41586-018-0410-x](https://doi.org/10.1038/s41586-018-0410-x) (2018).
- [141] A. D. King, J. Raymond, T. Lanting, S. V. Isakov, M. Mohseni, G. Poulin-Lamarre, S. Ejtemaee, W. Bernoudy, I. Ozfidan, A. Y. Smirnov, M. Reis, F. Altomare, M. Babcock, C. Baron, A. J. Berkley, K. Boothby, P. I. Bunyk, H. Christiani, C. Enderud, B. Evert, R. Harris, E. Hoskinson, S. Huang, K. Jooya, A. Khodabandelou, N. Ladizinsky, R. Li, P. A. Lott, A. J. R. MacDonald, D. Marsden, G. Marsden, T. Medina, R. Molavi, R. Neufeld, M. Norouzpour, T. Oh, I. Pavlov, I. Perminov, T. Prescott, C. Rich, Y. Sato, B. Sheldan, G. Sterling, L. J. Swenson, N. Tsai, M. H. Volkman, J. D. Whittaker, W. Wilkinson, J. Yao, H. Neven, J. P. Hilton, E. Ladizinsky, M. W. Johnson, and M. H. Amin, Scaling advantage over path-integral Monte Carlo in quantum simulation of geometrically frustrated magnets, *Nature Communications* **12**, 1113 (2021).
- [142] A. D. King, C. Nisoli, E. D. Dahl, G. Poulin-Lamarre, and A. Lopez-Bezanilla, Qubit spin ice, *Science* **373**, 576–580 (2021).

BIBLIOGRAPHY

- [143] F. Bloch, Nuclear Induction, [Physical Review](#) **70**, 460–474 (1946).
- [144] F. T. Arecchi, E. Courtens, R. Gilmore, and H. Thomas, Atomic Coherent States in Quantum Optics, [Physical Review A](#) **6**, 2211–2237 (1972).
- [145] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, Efficient Z gates for quantum computing, [Physical Review A](#) **96**, 022330 (2017).
- [146] K. N. Patel, I. L. Markov, and J. P. Hayes, Optimal synthesis of linear reversible circuits, [Quantum Information & Computation](#) **8**, 282–294 (2008).
- [147] J. J. Wallman and J. Emerson, Noise tailoring for scalable quantum computation via randomized compiling, [Physical Review A](#) **94**, 052325 (2016).
- [148] G. Haack, F. Helmer, M. Mariantoni, F. Marquardt, and E. Solano, Resonant quantum gates in circuit quantum electrodynamics, [Physical Review B](#) **82**, 024514 (2010).
- [149] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Demonstration of two-qubit algorithms with a superconducting quantum processor, [Nature](#) **460**, 240–244 (2009).
- [150] T. Yamamoto, M. Neeley, E. Lucero, R. C. Bialczak, J. S. Kelly, M. Lenander, M. Mariantoni, A. D. O’Connell, D. Sank, H. Wang, M. Weides, J. Wenner, Y. Yin, A. N. Cleland, and J. M. Martinis, Quantum process tomography of two-qubit controlled-Z and controlled-NOT gates using superconducting phase qubits, [Physical Review B](#) **82**, 184515 (2010).
- [151] M. Mariantoni, H. Wang, T. Yamamoto, M. Neeley, R. C. Bialczak, Y. Chen, M. Lenander, E. Lucero, A. D. O’Connell, D. Sank, M. Weides, J. Wenner, Y. Yin, J. Zhao, a. N. Korotkov, A. N. Cleland, and J. M. Martinis, Implementing the Quantum von Neumann Architecture with Superconducting Circuits, [Science](#) **334**, 61–65 (2011).
- [152] J. Emerson, R. Alicki, and K. Życzkowski, Scalable noise estimation with random unitary operators, [Journal of Optics B: Quantum and Semiclassical Optics](#) **7**, S347–S352 (2005).
- [153] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, Randomized benchmarking of quantum gates, [Physical Review A](#) **77**, 012307 (2008).

BIBLIOGRAPHY

- [154] C. Dankert, R. Cleve, J. Emerson, and E. Livine, Exact and approximate unitary 2-designs and their application to fidelity estimation, [Physical Review A **80**, 012304 \(2009\)](#).
- [155] J. M. Gambetta, A. D. Córcoles, S. T. Merkel, B. R. Johnson, J. A. Smolin, J. M. Chow, C. A. Ryan, C. T. Rigetti, S. Poletto, T. A. Ohki, M. B. Ketchen, and M. Steffen, Characterization of Addressability by Simultaneous Randomized Benchmarking, [Physical Review Letters **109**, 240504 \(2012\)](#).
- [156] E. Magesan, J. M. Gambetta, and J. Emerson, Characterizing quantum gates via randomized benchmarking, [Physical Review A **85**, 042311 \(2012\)](#).
- [157] E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, M. P. da Silva, G. A. Keefe, M. B. Rothwell, T. A. Ohki, M. B. Ketchen, and M. Steffen, Efficient Measurement of Quantum Gate Error by Interleaved Randomized Benchmarking, [Physical Review Letters **109**, 080505 \(2012\)](#).
- [158] J. J. Wallman, C. Granade, R. Harper, and S. T. Flammia, Estimating the coherence of noise, [New Journal of Physics **17**, 113020 \(2015\)](#).
- [159] A. K. Hashagen, S. T. Flammia, D. Gross, and J. J. Wallman, Real Randomized Benchmarking, [Quantum **2**, 85 \(2018\)](#).
- [160] A. Erhard, J. J. Wallman, L. Postler, M. Meth, R. Stricker, E. A. Martinez, P. Schindler, T. Monz, J. Emerson, and R. Blatt, Characterizing large-scale quantum computers via cycle benchmarking, [Nature Communications **10**, 5347 \(2019\)](#).
- [161] M. D. Bowdrey, D. K. Oi, A. J. Short, K. Banaszek, and J. A. Jones, Fidelity of single qubit maps, [Physics Letters A **294**, 258–260 \(2002\)](#).
- [162] M. A. Nielsen, A simple formula for the average gate fidelity of a quantum dynamical operation, [Physics Letters A **303**, 249–252 \(2002\)](#).
- [163] E. Schrödinger, An Undulatory Theory of the Mechanics of Atoms and Molecules, [Physical Review **28**, 1049–1070 \(1926\)](#).
- [164] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits, [Physical Review Letters **103**, 110501 \(2009\)](#).
- [165] J. M. Chow, L. DiCarlo, J. M. Gambetta, F. Motzoi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, Optimized driving of superconducting artificial atoms for improved single-qubit gates, [Physical Review A **82**, 040305 \(2010\)](#).

BIBLIOGRAPHY

- [166] E. Lucero, J. S. Kelly, R. C. Bialczak, M. Lenander, M. Mariantoni, M. Neeley, A. D. O'Connell, D. Sank, H. Wang, M. Weides, J. Wenner, T. Yamamoto, A. N. Cleland, and J. M. Martinis, Reduced phase error through optimized control of a superconducting qubit, [Physical Review A **82**, 042339 \(2010\)](#).
- [167] J. M. Gambetta, F. Motzoi, S. T. Merkel, and F. K. Wilhelm, Analytic control methods for high-fidelity unitary operations in a weakly nonlinear oscillator, [Physical Review A **83**, 012308 \(2011\)](#).
- [168] I. Pietikäinen, S. Danilin, K. S. Kumar, A. Vepsäläinen, D. S. Golubev, J. Tuorila, and G. S. Paraoanu, Observation of the Bloch-Siegert shift in a driven quantum-to-classical transition, [Physical Review B **96**, 020501 \(2017\)](#).
- [169] O. Astafiev, Y. Pashkin, Y. Nakamura, T. Yamamoto, and J. Tsai, Quantum Noise in the Josephson Charge Qubit, [Physical Review Letters **93**, 267007 \(2004\)](#).
- [170] P. Kumar, S. Sendelbach, M. A. Beck, J. W. Freeland, Z. Wang, H. Wang, C. C. Yu, R. Q. Wu, D. P. Pappas, and R. McDermott, Origin and Reduction of $1/f$ Magnetic Flux Noise in Superconducting Devices, [Physical Review Applied **6**, 041001 \(2016\)](#).
- [171] S. E. de Graaf, L. Faoro, J. Burnett, A. A. Adamyan, A. Y. Tzalenchuk, S. E. Kubatkin, T. Lindström, and A. V. Danilov, Suppression of low-frequency charge noise in superconducting resonators by surface spin desorption, [Nature Communications **9**, 1143 \(2018\)](#).
- [172] T. G. McConkey, J. H. Béjanin, C. T. Earnest, C. R. H. McRae, Z. Pagel, J. R. Rinehart, and M. Mariantoni, Mitigating leakage errors due to cavity modes in a superconducting quantum computer, [Quantum Science and Technology **3**, 034004 \(2018\)](#).
- [173] R. M. Lutchyn, L. I. Glazman, and A. I. Larkin, Kinetics of the superconducting charge qubit in the presence of a quasiparticle, [Physical Review B **74**, 064515 \(2006\)](#).
- [174] J. M. Martinis, M. Ansmann, and J. Aumentado, Energy Decay in Superconducting Josephson-Junction Qubits from Nonequilibrium Quasiparticle Excitations, [Physical Review Letters **103**, 097002 \(2009\)](#).
- [175] K. Serniak, M. Hays, G. de Lange, S. Diamond, S. Shankar, L. D. Burkhardt, L. Frunzio, M. Houzet, and M. H. Devoret, Hot Nonequilibrium Quasiparticles in Transmon Qubits, [Physical Review Letters **121**, 10.1103/PhysRevLett.121.157701 \(2018\)](#).

BIBLIOGRAPHY

- [176] C. D. Wilen, S. Abdullah, N. A. Kurinsky, C. Stanford, L. Cardani, G. D’Imperio, C. Tomei, L. Faoro, L. B. Ioffe, C. H. Liu, A. Opremcak, B. G. Christensen, J. L. DuBois, and R. McDermott, Correlated charge noise and relaxation errors in superconducting qubits, *Nature* **594**, 369–373 (2021).
- [177] A. A. Houck, J. A. Schreier, B. R. Johnson, J. M. Chow, J. Koch, J. M. Gambetta, D. I. Schuster, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Controlling the Spontaneous Emission of a Superconducting Transmon Qubit, *Physical Review Letters* **101**, 080502 (2008).
- [178] E. A. Sete, J. M. Gambetta, and A. N. Korotkov, Purcell effect with microwave drive: Suppression of qubit relaxation rate, *Physical Review B* **89**, 104516 (2014).
- [179] H.-P. Breuer and F. Petruccione, *The Theory of Open Quantum Systems* (Oxford University Press, 2007).
- [180] R. Alicki and K. Lendi, *Quantum Dynamical Semigroups and Applications*, 1st ed., Lecture Notes in Physics (Springer Berlin Heidelberg, Berlin, Heidelberg, 2007).
- [181] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J.-S. Tsai, and W. D. Oliver, Noise spectroscopy through dynamical decoupling with a superconducting flux qubit, *Nature Physics* **7**, 565–570 (2011).
- [182] Ansys, *Q3D Extractor - Simulation & Parasitic Extraction Tool*, <https://www.ansys.com/products/electronics/ansys-q3d-extractor> (2017).
- [183] Fast Field Solvers, *FasterCap*, <https://www.fastfieldsolvers.com/fastercap.htm> (2017).
- [184] Fast Field Solvers, *FastHenry2*, <https://www.fastfieldsolvers.com/fasthenry2.htm> (2017).
- [185] Ansys, *HFSS - 3D High Frequency Simulation Software*, <https://www.ansys.com/products/electronics/ansys-hfss> (2016).
- [186] J. Niemeyer and V. Kose, Observation of large dc supercurrents at nonzero voltages in Josephson tunnel junctions, *Applied Physics Letters* **29**, 380–382 (1976).
- [187] G. J. Dolan, Offset masks for lift-off photoprocessing, *Applied Physics Letters* **31**, 337–339 (1977).
- [188] X. Y. Jin, A. Kamal, A. P. Sears, T. Gudmundsen, D. Hover, J. Miloshi, R. Slattery, F. Yan, J. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver, Thermal and Residual Excited-State Population in a 3D Transmon Qubit, *Physical Review Letters* **114**, 240501 (2015).

BIBLIOGRAPHY

- [189] A. Lupaşcu, S. Saito, T. Picot, P. C. de Groot, C. J. P. M. Harmans, and J. E. Mooij, Quantum non-demolition measurement of a superconducting two-level system, *Nature Physics* **3**, 119–125 (2007).
- [190] J. E. Johnson, C. Macklin, D. H. Slichter, R. Vijay, E. B. Weingarten, J. Clarke, and I. Siddiqi, Heralded State Preparation in a Superconducting Qubit, *Physical Review Letters* **109**, 050506 (2012).
- [191] D. Ristè, J. G. van Leeuwen, H.-S. Ku, K. W. Lehnert, and L. DiCarlo, Initialization by Measurement of a Superconducting Quantum Bit Circuit, *Physical Review Letters* **109**, 050507 (2012).
- [192] J. M. Martinis and K. Osborne, Superconducting Qubits and the Physics of Josephson Junctions, in *Les Houches, Session LXXIX*, Les Houches No. 79, edited by D. Estève, J.-M. Raimond, and J. Dalibard (Elsevier, 2003) 1st ed., p. 34, [arXiv:cond-mat/0402415](https://arxiv.org/abs/cond-mat/0402415).
- [193] M. W. Johnson, P. Bunyk, F. Maibaum, E. Tolkacheva, A. J. Berkley, E. M. Chapple, R. Harris, J. Johansson, T. Lanting, I. Perminov, E. Ladizinsky, T. Oh, and G. Rose, A scalable control system for a superconducting adiabatic quantum optimization processor, *Superconductor Science and Technology* **23**, 065004 (2010).
- [194] R. McDermott, M. G. Vavilov, B. L. T. Plourde, F. K. Wilhelm, P. J. Liebermann, O. A. Mukhanov, and T. A. Ohki, Quantum–classical interface based on single flux quantum digital logic, *Quantum Science and Technology* **3**, 024004 (2018).
- [195] R. Barends, J. Wenner, M. Lenander, Y. Chen, R. C. Bialczak, J. S. Kelly, E. Lucero, P. O’Malley, M. Mariantoni, D. Sank, H. Wang, T. C. White, Y. Yin, J. Zhao, A. N. Cleland, J. M. Martinis, and J. J. A. Baselmans, Minimizing quasiparticle generation from stray infrared light in superconducting quantum circuits, *Applied Physics Letters* **99**, 113507 (2011).
- [196] J. Ekin, *Experimental Techniques for Low-Temperature Measurements*, 1st ed. (Oxford University Press, Oxford, 2006).
- [197] J. Johnson, Thermal Agitation of Electricity in Conductors, *Physical Review* **32**, 97–109 (1928).
- [198] H. Nyquist, Thermal Agitation of Electric Charge in Conductors, *Physical Review* **32**, 110–113 (1928).
- [199] S. J. Orfanidis, *Electromagnetic Waves and Antennas* (ECE Department, Rutgers University, Piscataway, NJ USA, 2016).

BIBLIOGRAPHY

- [200] M. V. Moghaddam, C. W. S. Chang, I. Nsanzineza, A. M. Vadiraj, and C. M. Wilson, Carbon nanotube-based lossy transmission line filter for superconducting qubit measurements, *Applied Physics Letters* **115**, 213504 (2019).
- [201] Laird, Eccosorb CRS, <https://www.laird.com/products/microwave-absorbers/injection-molded-machined-cast-liquids-and-microwave-absorb/eccosorb-crs> (2016).
- [202] D. F. Santavicca and D. E. Prober, Impedance-matched low-pass stripline filters, *Measurement Science and Technology* **19**, 087001 (2008).
- [203] DuPont, Pyralux AP Double-Sided Copper-Clad Laminate, <https://www.dupont.com/products/pyralux-ap.html> (2016).
- [204] D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, J. Majer, B. R. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Resolving photon number states in a superconducting circuit, *Nature* **445**, 515–8 (2007).
- [205] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, Approaching Unit Visibility for Control of a Superconducting Qubit with Dispersive Readout, *Physical Review Letters* **95**, 060501 (2005).
- [206] V. B. Braginsky, Y. I. Vorontsov, and K. S. Thorne, Quantum Nondemolition Measurements, *Science* **209**, 547–557 (1980).
- [207] V. B. Braginsky and F. Khalili, Quantum nondemolition measurements: The route from toys to tools, *Reviews of Modern Physics* **68**, 1–11 (1996).
- [208] E. Jeffrey, D. Sank, J. Mutus, T. C. White, J. S. Kelly, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Megrant, P. J. J. O’Malley, C. Neill, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, Fast Accurate State Measurement with Superconducting Qubits, *Physical Review Letters* **112**, 190504 (2014).
- [209] M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, High-Fidelity Readout in Circuit Quantum Electrodynamics Using the Jaynes-Cummings Nonlinearity, *Physical Review Letters* **105**, 173601 (2010).
- [210] J. M. Fink, L. Steffen, P. Studer, L. S. Bishop, M. Baur, R. Bianchetti, D. Bozyigit, C. Lang, S. Filipp, P. J. Leek, and A. Wallraff, Quantum-To-Classical Transition in Cavity Quantum Electrodynamics, *Physical Review Letters* **105**, 163601 (2010).

BIBLIOGRAPHY

- [211] D. H. Slichter, *Quantum Jumps and Measurement Backaction in a Superconducting Qubit*, [Ph.D. thesis](#), University of California, Berkeley (2011).
- [212] M. Reed, *Entanglement and Quantum Error Correction with Superconducting Qubits*, [Ph.D. thesis](#), Yale (2013), [arXiv:1311.6759](#) .
- [213] J. S. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. G. Fowler, I.-C. Hoi, E. Jeffrey, A. Megrant, J. Mutus, C. Neill, P. J. J. O’Malley, C. M. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Cleland, and J. M. Martinis, Optimal Quantum Control Using Randomized Benchmarking, [Physical Review Letters](#) **112**, 240504 (2014).
- [214] M. A. Rol, C. C. Bultink, T. E. O’Brien, S. R. de Jong, L. S. Theis, X. Fu, F. Luthi, R. F. L. Vermeulen, J. C. de Sterke, A. Bruno, D. Deurloo, R. N. Schouten, F. K. Wilhelm, and L. DiCarlo, Restless Tuneup of High-Fidelity Qubit Gates, [Physical Review Applied](#) **7**, 041001 (2017).
- [215] F. Yoshihara, K. Harrabi, A. O. Niskanen, Y. Nakamura, and J. S. Tsai, Decoherence of Flux Qubits due to $1/f$ Flux Noise, [Physical Review Letters](#) **97**, 167001 (2006).
- [216] E. L. Hahn, Spin Echoes, [Physical Review](#) **80**, 580–594 (1950).
- [217] Ł. Cywiński, R. M. Lutchyn, C. P. Nave, and S. Das Sarma, How to enhance dephasing time in superconducting qubits, [Physical Review B](#) **77**, 174509 (2008).
- [218] R. Koenig and J. A. Smolin, How to efficiently select an arbitrary Clifford group element, [Journal of Mathematical Physics](#) **55**, 122202 (2014).
- [219] G. Feng, J. J. Wallman, B. Buonacorsi, F. H. Cho, D. K. Park, T. Xin, D. Lu, J. Baugh, and R. Laflamme, Estimating the Coherence of Noise in Quantum Control of a Solid-State Qubit, [Physical Review Letters](#) **117**, 260501 (2016).
- [220] J. H. Béjanin, T. G. McConkey, J. R. Rinehart, C. T. Earnest, C. R. H. McRae, D. Shiri, J. D. Bateman, Y. Rohanizadegan, B. Penava, P. Breul, S. Royak, M. Zapatka, A. G. Fowler, and M. Mariantoni, Three-Dimensional Wiring for Extensible Quantum Computing: The Quantum Socket, [Physical Review Applied](#) **6**, 044010 (2016).
- [221] J. M. Martinis, Qubit metrology for building a fault-tolerant quantum computer, [npj Quantum Information](#) **1**, 15005, [npjqi.2015.5](#) (2015).
- [222] T. M. Miller, C. A. Jhabvala, E. Leong, N. P. Costen, E. Sharp, T. Adachi, and D. J. Benford, Enabling large focal plane arrays through mosaic hybridization, in *High*

BIBLIOGRAPHY

- Energy, Optical, and Infrared Detectors for Astronomy V*, Vol. 8453, edited by A. D. Holland and J. W. Beletic (SPIE, Amsterdam, Netherlands, 2012) pp. 721–727.
- [223] D. W. Abraham, J. M. Chow, A. D. C. Gonzalez, G. A. Keefe, M. E. Rothwell, J. R. Rozen, and M. Steffen, Removal of spurious microwave modes via flip-chip crossover, [United-States Patent US20140264287A1](#) (2014).
- [224] T. Brecht, M. Reagor, Y. Chu, W. Pfaff, C. Wang, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf, Demonstration of superconducting micromachined cavities, [Applied Physics Letters](#) **107**, 192603 (2015).
- [225] D. Rosenberg, D.-R. Yost, R. Das, D. Hover, L. Racz, S. Weber, J. Yoder, A. Kerman, and W. Oliver, [3D Integration for Superconducting Qubits](#) (2016).
- [226] P. Roushan, C. Neill, A. Megrant, Y. Chen, R. Babbush, R. Barends, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, A. Fowler, E. Jeffrey, J. Kelly, E. Lucero, J. Mutus, P. J. J. O’Malley, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. White, E. Kapit, H. Neven, and J. Martinis, Chiral ground-state currents of interacting photons in a synthetic magnetic field, [Nature Physics](#) **13**, 146–151 (2017).
- [227] A. Bruno, S. Poletto, N. Haider, and L. DiCarlo, [Extensible circuit QED processor architecture with vertical I/O](#) (2016) Mar., 2016.
- [228] J. M. Gambetta, J. M. Chow, and M. Steffen, Building logical qubits in a superconducting quantum computing system, [npj Quantum Information](#) **3**, 2 (2017).
- [229] T. Brecht, W. Pfaff, C. Wang, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf, Multilayer microwave integrated quantum circuits for scalable quantum computing, [npj Quantum Information](#) **2**, 16002 (2016).
- [230] J. Y. Mutus, T. C. White, R. Barends, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, J. Kelly, A. Megrant, C. Neill, P. J. J. O’Malley, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, K. M. Sundqvist, A. N. Cleland, and J. M. Martinis, Strong environmental coupling in a Josephson parametric amplifier, [Applied Physics Letters](#) **104**, 263513 (2014).
- [231] M. Zapatka and R. Ziser, An introduction to coaxial RF probing solutions for mass-production tests, in [2009 74th ARFTG Microwave Measurement Conference](#) (IEEE, Broomfield, CO, USA, 2009) pp. 1–6.
- [232] R. E. Collin, [Foundations for Microwave Engineering](#) (IEEE, 2001).

BIBLIOGRAPHY

- [233] J. Wenner, M. Neeley, R. C. Bialczak, M. Lenander, E. Lucero, A. D. O'Connell, D. Sank, H. Wang, M. Weides, A. N. Cleland, and J. M. Martinis, Wirebond crosstalk and cavity modes in large chip mounts for superconducting qubits, [Superconductor Science and Technology](#) **24**, 065001 (2011).
- [234] L. Frunzio, A. Wallraff, D. I. Schuster, J. Majer, and R. J. Schoelkopf, Fabrication and characterization of superconducting circuit QED devices for quantum computation, [IEEE Transactions on Applied Superconductivity](#) **15**, 860–863 (2005).
- [235] C. Song, T. W. Heitmann, M. P. DeFeo, K. Yu, R. McDermott, M. Neeley, J. M. Martinis, and B. L. T. Plourde, Microwave response of vortices in superconducting thin films of Re and Al, [Physical Review B](#) **79**, 174512 (2009).
- [236] C. Song, M. P. DeFeo, K. Yu, and B. L. T. Plourde, Reducing microwave loss in superconducting resonators due to trapped vortices, [Applied Physics Letters](#) **95**, 232501 (2009).
- [237] A. D. Córcoles, J. M. Chow, J. M. Gambetta, C. Rigetti, J. R. Rozen, G. A. Keefe, M. Beth Rothwell, M. B. Ketchen, and M. Steffen, Protecting superconducting qubits from radiation, [Applied Physics Letters](#) **99**, 181906 (2011).
- [238] A. G. Fowler and J. M. Martinis, Quantifying the effects of local many-qubit errors and nonlocal two-qubit errors on the surface code, [Physical Review A](#) **89**, 032316 (2014).
- [239] D. Ristè, S. Poletto, M.-Z. Huang, A. Bruno, V. Vesterinen, O.-P. Saira, and L. DiCarlo, Detecting bit-flip errors in a logical qubit using stabilizer measurements, [Nature Communications](#) **6**, 6983 (2015).
- [240] A. Teverovsky, Effect of vacuum on high-temperature degradation of gold/aluminum wire bonds in PEMs, in *2004 IEEE International Reliability Physics Symposium. Proceedings* (IEEE, Phoenix, AZ, USA, 2004) pp. 547–556.
- [241] E. D. Marquardt, J. P. Le, and R. Radebaugh, Cryogenic Material Properties Database, in *Cryocoolers 11*, edited by R. G. Ross (Springer US, Boston, MA, 2002) pp. 681–687.
- [242] E. Bogatin, *Signal Integrity - Simplified*, Prentice Hall Modern Semiconductor Design Series (Prentice Hall, Upper Saddle River, NJ, 2003).
- [243] D. W. Abraham, G. A. Keefe, C. Lavoie, and M. E. Rothwell, Chip mode isolation and cross-talk reduction through buried metal layers and through-vias, [United-States Patent US9397283B2](#) (2014).

BIBLIOGRAPHY

- [244] D. W. Abraham, J. M. Chow, and J. M. Gambetta, Symmetric placement of components on a chip to reduce crosstalk induced by chip modes, [United-States Patent US20140266406A1](#) (2015).
- [245] J. Tournet, D. Gosselink, G.-X. Miao, M. Jaikissoon, D. Langenberg, T. G. McConkey, M. Mariantoni, and Z. R. Wasilewski, Growth and characterization of epitaxial aluminum layers on gallium-arsenide substrates for superconducting quantum bits, [Superconductor Science and Technology](#) **29**, 064004 (2016).
- [246] D. Brock, E. Track, and J. Rowell, Superconductor ICs: The 100-GHz second generation, [IEEE Spectrum](#) **37**, 40–46 (2000).
- [247] O. A. Mukhanov, Energy-Efficient Single Flux Quantum Technology, [IEEE Transactions on Applied Superconductivity](#) **21**, 760–769 (2011).
- [248] D. Gupta, D. E. Kirichenko, V. V. Dotsenko, R. Miller, S. Sarwana, A. Talalaevskii, J. Delmas, R. J. Webber, S. Govorkov, A. F. Kirichenko, I. V. Vernik, and J. Tang, Modular, Multi-Function Digital-RF Receiver Systems, [IEEE Transactions on Applied Superconductivity](#) **21**, 883–890 (2011).
- [249] J. Béjanin, C. Earnest, Y. Sanders, and M. Mariantoni, Resonant Coupling Parameter Estimation with Superconducting Qubits, [PRX Quantum](#) **2**, 040343 (2021).
- [250] A. O. Niskanen, K. Harrabi, F. Yoshihara, Y. Nakamura, S. Lloyd, and J. S. Tsai, Quantum Coherent Tunable Coupling of Superconducting Qubits, [Science](#) **316**, 723–726 (2007).
- [251] M. S. Allman, F. Altomare, J. D. Whittaker, K. Cicak, D. Li, A. Sirois, J. Strong, J. D. Teufel, and R. W. Simmonds, Rf-SQUID-Mediated Coherent Tunable Coupling between a Superconducting Phase Qubit and a Lumped-Element Resonator, [Physical Review Letters](#) **104**, 177004 (2010).
- [252] C. Müller, J. H. Cole, and J. Lisenfeld, Towards understanding two-level-systems in amorphous solids: Insights from quantum circuits, [Reports on Progress in Physics](#) **82**, 124501 (2019).
- [253] M. Moeed, C. T. Earnest, J. Béjanin, A. Sharafeldin, and M. Mariantoni, Improving the Time Stability of Superconducting Planar Resonators, [MRS Advances](#) **4**, 2201–2215 (2019).
- [254] P. V. Klimov, J. S. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Y. Chen, A. Dunsworth, A. G. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, T. Huang, E. Jeffrey, E. Lucero, J. Mutus, O. Naaman, C. Neill,

BIBLIOGRAPHY

- C. M. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, S. Boixo, R. Babbush, V. N. Smelyanskiy, H. Neven, and J. M. Martinis, Fluctuations of Energy-Relaxation Times in Superconducting Qubits, [Physical Review Letters](#) **121**, 90502 (2018).
- [255] M. Mariani, H. Wang, R. C. Bialczak, M. Lenander, E. Lucero, M. Neeley, A. D. O'Connell, D. Sank, M. Weides, J. Wenner, T. Yamamoto, Y. Yin, J. Zhao, J. M. Martinis, and A. N. Cleland, Photon shell game in three-resonator circuit quantum electrodynamics, [Nature Physics](#) **7**, 287–293 (2011).
- [256] Y. R. Sanders, *Characterizing Errors in Quantum Information Processors*, Ph.D. thesis, University of Waterloo (2016).
- [257] M. P. Stenberg, Y. R. Sanders, and F. K. Wilhelm, Efficient Estimation of Resonant Coupling between Quantum Systems, [Physical Review Letters](#) **113**, 210404 (2014).
- [258] C. E. Granade, C. Ferrie, N. Wiebe, and D. G. Cory, Robust online Hamiltonian learning, [New Journal of Physics](#) **14**, 103013 (2012).
- [259] M. Mariani, The Energy of an Arbitrary Electrical Circuit, Classical and Quantum, [arXiv:2007.08519](#) (2020).
- [260] S. Schlör, J. Lisenfeld, C. Müller, A. Bilmes, A. Schneider, D. P. Pappas, A. V. Ustinov, and M. Weides, Correlating Decoherence in Transmon Qubits: Low Frequency Noise by Single Fluctuators, [Physical Review Letters](#) **123**, 190502 (2019).
- [261] J. J. Burnett, A. Bengtsson, M. Scigliuzzo, D. Niepce, M. Kudra, P. Delsing, and J. Bylander, Decoherence benchmarking of superconducting qubits, [npj Quantum Information](#) **5**, 54 (2019).
- [262] J. H. Béjanin, C. T. Earnest, A. S. Sharafeldin, and M. Mariani, Interacting defects generate stochastic fluctuations in superconducting qubits, [Physical Review B](#) **104**, 094106 (2021).
- [263] J. H. Béjanin, [TLSInfer.jl](#), <https://gitlab.com/DQMLab/TLSInfer.jl> (2020).
- [264] L. M. Murray, A. Lee, and P. E. Jacob, Parallel Resampling in the Particle Filter, [Journal of Computational and Graphical Statistics](#) **25**, 789–805 (2016).
- [265] P. Bertet, C. J. P. M. Harmans, and J. E. Mooij, Parametric coupling for superconducting qubits, [Physical Review B](#) **73**, 064512 (2006).

BIBLIOGRAPHY

- [266] D. C. McKay, S. Filipp, A. Mezzacapo, E. Magesan, J. M. Chow, and J. M. Gambetta, Universal Gate for Fixed-Frequency Qubits via a Tunable Bus, [Physical Review Applied](#) **6**, 064007 (2016).
- [267] S. A. Caldwell, N. Didier, C. A. Ryan, E. A. Sete, A. Hudson, P. Karalekas, R. Manenti, M. P. da Silva, R. Sinclair, E. Acala, N. Alidoust, J. Angeles, A. Bestwick, M. Block, B. Bloom, A. Bradley, C. Bui, L. Capelluto, R. Chilcott, J. Cordova, G. Crossman, M. Curtis, S. Deshpande, T. E. Bouayadi, D. Girshovich, S. Hong, K. Kuang, M. Lenihan, T. Manning, A. Marchenkov, J. Marshall, R. Maydra, Y. Mohan, W. O'Brien, C. Osborn, J. Otterbach, A. T. Papageorge, J.-P. Paquette, M. Pelstring, A. M. Polloreno, G. Prawiroatmodjo, V. Rawat, M. Reagor, R. Renzas, N. Rubin, D. Russell, M. Rust, D. Scarabelli, M. Scheer, M. Selvanayagam, R. Smith, A. Staley, M. Suska, N. Tezak, D. C. Thompson, T.-W. To, M. Vahidpour, N. Vodrahalli, T. Whyland, K. Yadav, W. Zeng, and C. Rigetti, Parametrically Activated Entangling Gates Using Transmon Qubits, [Physical Review Applied](#) **10**, 034050 (2018).
- [268] P. V. Klimov, J. S. Kelly, J. M. Martinis, and H. Neven, The Snake Optimizer for Learning Quantum Processor Control Parameters, [arXiv:2006.04594](#) (2020).
- [269] W. Phillips, Two-level states in glasses, [Reports on Progress in Physics](#) **50**, 1657–1708 (1987).
- [270] J. M. Martinis, K. B. Cooper, R. McDermott, M. Steffen, M. Ansmann, K. D. Osborn, K. Cicak, S. Oh, D. P. Pappas, R. W. Simmonds, and C. C. Yu, Decoherence in Josephson Qubits from Dielectric Loss, [Physical Review Letters](#) **95**, 210503 (2005).
- [271] D. S. Wisbey, J. Gao, M. R. Vissers, F. C. S. Da Silva, J. S. Kline, L. Vale, and D. P. Pappas, Effect of metal/substrate interfaces on radio-frequency loss in superconducting coplanar waveguides, [Journal of Applied Physics](#) **108**, 10.1063/1.3499608 (2010).
- [272] 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](#) **109**, 063915 (2011).
- [273] C. J. K. Richardson, N. P. Siwak, J. Hackley, Z. K. Keane, J. E. Robinson, B. Arey, I. Arslan, and B. S. Palmer, Fabrication artifacts and parallel loss channels in metamorphic epitaxial aluminum superconducting resonators, [Superconductor Science and Technology](#) **29**, 064003 (2016).
- [274] A. Dunsworth, A. Megrant, C. M. Quintana, Z. Chen, R. Barends, B. Burkett, B. Foxen, Y. Chen, B. Chiaro, A. G. Fowler, R. Graff, E. Jeffrey, J. S. Kelly, E. Lucero,

BIBLIOGRAPHY

- J. Mutus, M. Neeley, C. Neill, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, and J. M. Martinis, Characterization and reduction of capacitive loss induced by sub-micron Josephson junction fabrication in superconducting qubits, [Applied Physics Letters](#) **111**, 022601 (2017).
- [275] A. Bilmes, S. Volosheniuk, J. D. Brehm, A. V. Ustinov, and J. Lisenfeld, Quantum sensors for microscopic tunneling systems, [npj Quantum Information](#) **7**, 27 (2021).
- [276] J. Wenner, R. Barends, R. C. Bialczak, Y. Chen, J. S. Kelly, E. Lucero, M. Mariantoni, A. Megrant, P. J. J. O'Malley, D. Sank, A. Vainsencher, H. Wang, T. C. White, Y. Yin, J. Zhao, A. N. Cleland, and J. M. Martinis, Surface loss simulations of superconducting coplanar waveguide resonators, [Applied Physics Letters](#) **99**, 2009–2012 (2011).
- [277] W. Woods, G. Calusine, A. Melville, A. Sevi, E. Golden, D. Kim, D. Rosenberg, J. Yoder, and W. Oliver, Determining Interface Dielectric Losses in Superconducting Coplanar-Waveguide Resonators, [Physical Review Applied](#) **12**, 014012 (2019).
- [278] M. Neeley, M. Ansmann, R. C. Bialczak, M. Hofheinz, N. Katz, E. Lucero, A. O'Connell, H. Wang, A. N. Cleland, and J. M. Martinis, Process tomography of quantum memory in a Josephson-phase qubit coupled to a two-level state, [Nature Physics](#) **4**, 523–526 (2008).
- [279] J. Burnett, L. Faoro, I. Wisby, V. L. Gurtovoi, A. V. Chernykh, G. M. Mikhailov, V. A. Tulin, R. Shaikhaidarov, V. Antonov, P. J. Meeson, A. Y. Tzalenchuk, and T. Lindström, Evidence for interacting two-level systems from the $1/f$ noise of a superconducting resonator, [Nature Communications](#) **5**, 4119 (2014).
- [280] L. Faoro and L. B. Ioffe, Interacting tunneling model for two-level systems in amorphous materials and its predictions for their dephasing and noise in superconducting microresonators, [Physical Review B](#) **91**, 014201 (2015).
- [281] C. Müller, J. Lisenfeld, A. Shnirman, and S. Poletto, Interacting two-level defects as sources of fluctuating high-frequency noise in superconducting circuits, [Physical Review B](#) **92**, 035442 (2015).
- [282] C. Neill, A. Megrant, R. Barends, Y. Chen, B. Chiaro, J. S. Kelly, J. Mutus, P. J. J. O'Malley, D. Sank, J. Wenner, T. C. White, Y. Yin, A. N. Cleland, and J. M. Martinis, Fluctuations from edge defects in superconducting resonators, [Applied Physics Letters](#) **103**, 072601 (2013).

BIBLIOGRAPHY

- [283] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, Observation of High Coherence in Josephson Junction Qubits Measured in a Three-Dimensional Circuit QED Architecture, *Physical Review Letters* **107**, 240501 (2011).
- [284] M. Carroll, S. Rosenblatt, P. Jurcevic, I. Lauer, and A. Kandala, Dynamics of superconducting qubit relaxation times, [arXiv:2105.15201](https://arxiv.org/abs/2105.15201) (2021).
- [285] P. Esquinazi, *Tunneling Systems in Amorphous and Crystalline Solids* (Springer, Berlin, Heidelberg, 1998).
- [286] J. Lisenfeld, C. Müller, J. H. Cole, P. Bushev, A. Lukashenko, A. Shnirman, and A. V. Ustinov, Measuring the Temperature Dependence of Individual Two-Level Systems by Direct Coherent Control, *Physical Review Letters* **105**, 230504 (2010).
- [287] P. Meystre and M. Sargent, *Elements of Quantum Optics*, 4th ed. (Springer, Berlin, Heidelberg, 2007).
- [288] A. Kamal, J. L. Yoder, F. Yan, T. J. Gudmundsen, D. Hover, A. P. Sears, P. Welander, T. P. Orlando, S. Gustavsson, and W. D. Oliver, Improved superconducting qubit coherence with high-temperature substrate annealing, [arXiv:1606.09262](https://arxiv.org/abs/1606.09262) (2016).
- [289] C. R. H. McRae, J. H. Béjanin, C. T. Earnest, T. G. McConkey, J. R. Rinehart, C. Deimert, J. P. Thomas, Z. R. Wasilewski, and M. Mariani, Thin film metrology and microwave loss characterization of indium and aluminum/indium superconducting planar resonators, *Journal of Applied Physics* **123**, 205304 (2018).
- [290] J. Bezanson, A. Edelman, S. Karpinski, and V. B. Shah, Julia: A Fresh Approach to Numerical Computing, *SIAM Review* **59**, 65–98 (2017).
- [291] J. H. Béjanin, `QubitFluctuations.jl`, <https://gitlab.com/DQMLab/qubitfluctuations/> (2021).
- [292] A. F. Wells, *Structural Inorganic Chemistry*, 5th ed. (Clarendon Press ; Oxford University Press, Oxford [Oxfordshire] : New York, 1984).
- [293] T. T. Song, M. Yang, J. W. Chai, M. Callsen, J. Zhou, T. Yang, Z. Zhang, J. S. Pan, D. Z. Chi, Y. P. Feng, and S. J. Wang, The stability of aluminium oxide monolayer and its interface with two-dimensional materials, *Scientific Reports* **6**, 29221 (2016).
- [294] H. J. Mamin, E. Huang, S. Carnevale, C. T. Rettner, N. Arellano, M. H. Sherwood, C. Kurter, B. Trimm, M. Sandberg, R. M. Shelby, M. A. Mueed, B. A. Madon, A. Pushp,

BIBLIOGRAPHY

- M. Steffen, and D. Rugar, Merged-Element Transmons: Design and Qubit Performance, *Physical Review Applied* **16**, 024023 (2021).
- [295] W. Riley, *Handbook of Frequency Stability Analysis* (NIST, 2008).
- [296] Y. Chu, P. Kharel, T. Yoon, L. Frunzio, P. T. Rakich, and R. J. Schoelkopf, Creation and control of multi-phonon Fock states in a bulk acoustic-wave resonator, *Nature* **563**, 666–670 (2018).
- [297] A. Potočnik, A. Bargerbos, F. A. Y. N. Schröder, S. A. Khan, M. C. Collodo, S. Gasparinetti, Y. Salathé, C. Creatore, C. Eichler, H. E. Türeci, A. W. Chin, and A. Wallraff, Studying light-harvesting models with superconducting circuits, *Nature Communications* **9**, 904 (2018).
- [298] D. Rosenberg, D. Kim, R. Das, D. Yost, S. Gustavsson, D. Hover, P. Krantz, A. Melville, L. Racz, G. O. Samach, S. J. Weber, F. Yan, J. L. Yoder, A. J. Kerman, and W. D. Oliver, 3D integrated superconducting qubits, *npj Quantum Information* **3**, 42 (2017).
- [299] B. Foxen, J. Y. Mutus, E. Lucero, R. Graff, A. Megrant, Y. Chen, C. Quintana, B. Burkett, J. Kelly, E. Jeffrey, Y. Yang, A. Yu, K. Arya, R. Barends, Z. Chen, B. Chiaro, A. Dunsworth, A. Fowler, C. Gidney, M. Giustina, T. Huang, P. Klimov, M. Neeley, C. Neill, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, and J. M. Martinis, Qubit compatible superconducting interconnects, *Quantum Science and Technology* **3**, 014005 (2017).
- [300] J. Rahamim, T. Behrle, M. J. Peterer, A. Patterson, P. A. Spring, T. Tsunoda, R. Marenti, G. Tancredi, and P. J. Leek, Double-sided coaxial circuit QED with out-of-plane wiring, *Applied Physics Letters* **110**, 222602 (2017).
- [301] I. Siddiqi, Engineering high-coherence superconducting qubits, *Nature Reviews Materials* **6**, 875–891 (2021).
- [302] Diehl Metall, ECOMERICA Alloys, <https://web.archive.org/web/20150730033046/http://www.diehl.com/en/diehl-metall/company/brands/diehl-metall-messing/ecomerica/alloys.html> (2015).
- [303] Otto Fuchs, Copper and Copper Alloys, https://www.otto-fuchs-duelken.de/fileadmin/user_upload/Downloads/OF2285_2014-02_EN.pdf (2015).
- [304] Steel Number, European Steel and Alloy Grades, http://www.steelnumber.com/en/steel_alloy_composition_eu.php?name_id=1310 (2015).

BIBLIOGRAPHY

- [305] R. M. Mueller, C. Buchal, T. Oversluizen, and F. Pobell, Superconducting aluminum heat switch and plated press-contacts for use at ultralow temperatures, *Review of Scientific Instruments* **49**, 515–518 (1978).
- [306] NIST Cryogenic Materials Property Database, *Material Properties: 6061 Aluminum*, http://web.archive.org/web/20180307052018/http://cryogenics.nist.gov/MPropsMAY/6061%20Aluminum/6061_T6Aluminum_rev.htm (2016).
- [307] C. A. Swenson, Recommended Values for the Thermal Expansivity of Silicon from 0 to 1000 K, *Journal of Physical and Chemical Reference Data* **12**, 179–182 (1983).
- [308] C. E. Murray, J. M. Gambetta, D. T. McClure, and M. Steffen, Analytical Determination of Participation in Superconducting Coplanar Architectures, *IEEE Transactions on Microwave Theory and Techniques* **66**, 3724–3733 (2018).

Appendix for Chapter 3

A.1 Qubit Chip Fabrication Recipe

The fabrication recipe for the sample used in Chapters 5 and 6 is described in the steps below. It was developed over the years by C. T. Earnest and C. R. H. McRae. The fabrication process is performed on a 4 in float-zone silicon wafer with $\langle 100 \rangle$ orientation and at least $10 \text{ k}\Omega \text{ cm}$ resistivity. First, the markers, which are used to align the optical layer to the e-beam layer, are etched in the wafer (steps 1-5). Then, the wafer is cleaned and a 100 nm aluminum layer is evaporated (steps 6-8). Optical lithography is used to fabricate the large structures (i.e., everything but the junctions) with a mask covering the whole wafer (steps 9-13). The Josephson junctions are fabricated with electron beam lithography (steps 14-19). Note that the wafer is diced into individual chips right after e-beam patterning, and before development (step 16).

1. Spin photoresist

- Tool: Headway Research spin coater
- Supplies: Shipley S1811 photoresist
- Process: Spin resist with an initial 4 sec spin (low speed) at 500 rpm (100 rpm/s acceleration), then a 60 sec spin (high speed) at 5000 rpm (500 rpm/s acceleration). Bake on hot plate, 90 sec at 120°C .

2. UV exposure of marker mask

- Tool: SUSS MA6 mask aligner
- Process: exposure: 4 sec, alignment gap: $30 \mu\text{m}$, WEC offset: 0, WEC type: contact, exposure type: vacuum, pre vacuum: 5 sec, full: 5 sec, purge vacuum: 20 sec.

3. Develop photoresist

- Tool: UV development wetbench
- Supplies: MF-319 developer, deionized water
- Process: Develop in MF-319 bath for 45 sec, then rinse with DI water 1 min.

4. Dry-etch silicon

- Tool: Oxford Instruments ICP380 DRIE dry silicon etcher
- Process: 15 cycles, for ~ 10 min at 2 $\mu\text{m}/\text{min}$.

5. Strip photoresist

- Tool: Solvent wetbench
- Supplies: Remover PG solvent stripper, IPA
- Process: Prepare two Remover PG baths (heated to 70°C), and one IPA bath. Sonicate for 7 min in each Remover PG bath. IPA bath for 7 min. Rinse with DI water.

6. RCA SC-1 Clean

- Tool: HF wetbench
- Supplies: deionized water, H_2O_2 , NH_4OH
- Process: Add 300 mL DI water, then 60 mL H_2O_2 , then 60 mL NH_4OH to RCA vessel (5:1:1 ratio). Turn on hot plate to 230°C, once vessel has reached 75°C, turn the hot plate down to 160°C and immerse sample for 10 min, monitoring to maintain the temperature at 75°C. Remove vessel from hot plate and give wafer a quick rinse in DI water bath (60 sec) before transferring it to the HF dip (see next step).

7. HF Dip

- Tool: HF wetbench
- Supplies: 10:1 buffered oxide etch (buffered HF), deionized water
- Process: Prepare 3 DI water baths with increasing volumes. Prepare 1% HF solution by adding 50 mL 10:1 buffered HF to 450 mL of DI water . Set timer for 1 min and immerse wafer held by Teflon holder in HF solution. Check wafer to see hydrophilic-phobic visual change. Remove wafer and immerse in first DI bath for 2 min. Repeat for other two DI baths. Dry completely with N_2 gas, then transfer immediately to Plassys holder, and into loadlock. Pump loadlock immediately. Dispose of HF and clean up work area with water.

8. Evaporate Aluminum layer (with 850°C anneal)

- Tool: Plassys MEB 550 SL3-UHV e-beam evaporator
- Process: Ramp heat to 850°C over 30 min and anneal wafer 10 min. Cool down to 25°C overnight. Before evaporation, perform gettering by evaporating Ti on chamber walls. Evaporate 100 nm of Al at 2 nm/s. Pressure after gettering: $\sim 1 \times 10^{-9}$ Torr, during evaporation: $\sim 1 \times 10^{-8}$ Torr.

9. Spin photoresist

- Tool: Headway Research spin coater
- Supplies: Shipley S1811 photoresist
- Process: Spin resist with an initial 4 sec spin (low speed) at 500 rpm (100 rpm/s acceleration), then a 60 sec spin (high speed) at 5000 rpm (500 rpm/s acceleration). Bake on hot plate, 90 sec at 120°C.

10. UV exposure of device mask

- Tool: SUSS MA6 mask aligner
- Process: exposure: 4 sec, alignment gap: 30 μm , WEC offset: 0, WEC type: contact, exposure type: vacuum, pre vacuum: 5 sec, full: 5 sec, purge vacuum: 20 sec.

11. Develop photoresist

- Tool: UV development wetbench
- Supplies: MF-319 developer, deionized water
- Process: Develop in MF-319 bath for 45 sec, then rinse with DI water 1 min. Carefully time development as MF-319 will slightly etch the Al.

12. Dry-etch aluminum

- Tool: Oxford Instruments ICP380 DRIE dry metal etcher
- Process: Prepare DI water bath. Etch Al in ICP mode with added N_2 gas for 60 sec. Take out wafer and immediately immerse wafer in water bath. Rinse thoroughly to remove reactive chlorinated species. Dry with N_2 .

13. Strip photoresist

- Tool: Solvent wetbench
- Supplies: Remover PG solvent stripper, IPA

- Process: Prepare two Remover PG baths (heated to 70°C), and one IPA bath. Sonicate for 7 min in each Remover PG bath. IPA bath for 7 min. Rinse with IPA and dry with N₂ (no water before e-beam resist, which is hydrophobic).
14. Spin e-beam resist (double layer)
- Tool: Headway Research spin coater
 - Supplies: MMA (8.5) EL11 (for the bottom layer, 500 nm), PMMA A4 (for the top layer, 200 nm)
 - Process: First layer, spin MMA, initial 5 sec spin at 500 rpm (100 rpm/sec acceleration), then 50 sec spin at 4000 rpm (1000 rpm/s acceleration), bake 90 sec at 150°C. Second layer, spin PMMA, initial 5 sec spin at 500 rpm (100 rpm/sec acceleration), then 40 sec spin at 4000 rpm (1000 rpm/s acceleration), bake 90 sec at 180°C.
15. E-beam exposure of junctions
- Tool: JEOL JBX-6300FS Electron Beam Lithography System
 - Process: Align to Si markers and pattern junctions.
16. Dice wafer into chips
- Tool: DISCO DAD3240 dicing saw
17. Develop e-beam resist
- Tool: E-beam development wetbench
 - Supplies: MIBK developer, IPA developer
 - Process: 90 sec MIBK/IPA (1:3 ratio), then 15 sec IPA, dry with N₂.
18. Evaporate junctions
- Tool: Plassys MEB 550 SL3-UHV e-beam evaporator
 - Process: Evaporate 50 nm Al at +17°. Oxidize with O₂ at 3 Torr for 35 min. Evaporate 70 nm Al at -17°.
19. Lift-off e-beam resist
- Tool: Solvent wetbench
 - Supplies: Remover PG, IPA
 - Process: Place chip in Remover PG bath at 70°C for 20min with agitation. Peel off aluminum by squirting with pipette on chip. Move to fresh Remover PG bath and sonicate on low power. Rinse with IPA and dry with N₂.

Appendix for Chapter 4

B.1 Wire Compression

In this appendix, we discuss the pressure settings of the three-dimensional wires. In the current implementation of the quantum socket, the pressure exerted by the three-dimensional wires on the chip is controlled by the installation depth of the wire in the lid. This depth depends on the number of rotations used to screw the wire into the M2.5-threaded hole of the lid. Since the wire's tunnel has to be aligned with the corresponding on-chip pad, a discrete number of wire pressure settings is allowed. For the package shown in Fig. 4.1 (b) and Fig. 4.4 (b), the minimum length an unloaded wire has to protrude from the ceiling of the lid's internal cavity to touch the top surface of the chip is $\ell_c = 3.05$ mm (cf. Fig. 4.1 (c)). For a maximum wire stroke $\Delta L = 2.5$ mm, the maximum length an unloaded wire can protrude from the cavity ceiling without breaking when loaded is $\ell_c + \Delta L = 5.55$ mm. The first allowed pressure setting, with wire and pad perfectly aligned, is for $\ell_p = 3.10$ mm. The pitch for an M2.5 screw is 0.45 mm. Hence, five pressure settings are nominally possible, for $\ell_p = 3.10 + 0.45 k$ mm, with $k = 1, 2, \dots, 5$. We found the ideal pressure setting to be for $k = 3$, corresponding to a nominal $\ell_p = 4.45$ mm; the actual average setting for 12 wires was measured to be $\ell_p = 4.48 \text{ mm} \mp 0.28 \text{ mm}$, with standard deviation due to the machining tolerances. For greater depths we experienced occasional wire damage; lesser depths were not investigated. Possible effects on the electrical properties of the three-dimensional wires due to different pressure settings will be studied in a future work.

Table B.1: Chemical composition (weight %) of the two main materials used in the three-dimensional wires. Copper: Cu; tin: Sn; zinc: Zn; lead: Pb; phosphorus: P; aluminum: Al; manganese: Mn.

Material	Cu	Sn	Zn	Fe	Ni	Pb	P	Si	Others
CW724R ^a	73 – 77	0.3	rest	0.3	0.2	≤ 0.09	0.04 – 0.10	2.7 – 3.4	Al = 0.05, Mn = 0.05
CW453K ^b	rest	7.5 – 8.5	≤ 0.2	≤ 0.1	≤ 0.2	≤ 0.02	0.01 – 0.4	-	0.2

^a See [302] and [303].

^b See [304].

B.2 Magnetism

In this appendix, we describe the measurement setup employed to characterize the magnetic properties of the materials used in the quantum socket and present the main measurement results. Additionally, we give an estimate of the strength of the magnetic field caused by one three-dimensional wire inside the microwave package.

The ZGC used in our tests comprises three nested cylinders, each with a lid with a central circular hole; the hole in the outermost lid is extended into a chimney that provides further magnetic shielding. The walls of the ZGC are made of an alloy of Ni and Fe (or mu-metal alloy) with a high relative magnetic permeability μ_r . The alloy used for the chamber is a CO-NETIC® AA alloy and is characterized by a DC relative magnetic permeability at 40 G, $\mu_{DC}^{40} = 80000$, and an AC relative magnetic permeability at 60 Hz and at 40 G, $\mu_{AC}^{40} = 65000$. As a consequence, the nominal magnetic field attenuation lies between 1000 and 1500. The ZGC used in our tests was manufactured by the Magnetic Shield Corporation, model ZG-209.

The flux gate magnetometer used to measure the magnetic field \vec{B} is a three-axis DC milligauss meter from AlphaLab, Inc., model MGM3AXIS. Its sensor is a 38 mm \times 25 mm \times 25 mm parallelepiped at the end of a ~ 1.2 m long cable; the orientation of the sensor is calibrated to within 0.1° and has a resolution of 0.01 mG (i.e., 1 nT) over a range of ∓ 2000 mG (i.e., ∓ 200 μ T).

The actual attenuation of the chamber was tested by measuring the value of the Earth’s magnetic field with and without the chamber in two positions, vertical and horizontal; inside the chamber the measurements were performed a few centimeters from the chamber’s base, approximately on the axis of the inner cylinder. In these and all subsequent tests, the magnetic sensor was kept in the same orientation and position. The results are reported in Table B.2, which shows the type of measurement performed, the magnitude of the measured magnetic field $\|\vec{B}\|$, and the attenuation ratio α . The

maximum measured attenuation was $\alpha \simeq 917$ in the horizontal position.

The ZGC characterization of Table B.2 also serves as a calibration for the measurements on the materials used for the quantum socket. In these measurements, each test sample was positioned approximately 1 cm away from the magnetic sensor. The results, which are reported in Table B.3, were obtained by taking the magnitude of the calibrated field of each sample. The calibrated field itself was calculated by subtracting the background field from the sample field, component by component. Note that the background and sample fields were on the same order of magnitude (between 0.10 mG and 0.80 mG), with background fluctuations on the order of 0.10 mG. Thus, we recorded the maximum value of each x , y , and z component. Considering that the volume of the measured samples is significantly larger than that of the actual quantum socket components, we are confident that the measured magnetic fields of the materials should be small enough not to significantly disturb the operation of superconducting quantum devices. As part of our magnetism tests, we measured a block of approximately 200 g of 5N5 Al in the ZGC; as shown in Table B.3, the magnitude of the magnetic field was found to be within the noise floor of the measurement apparatus ¹.

A simple geometric argument allows us to estimate the actual magnetic field due to one three-dimensional wire, without taking into account effects due to superconductivity (most of the wire is embedded in an Al package, which is superconductive at the qubit operation temperatures). We assume that one wire generates a magnetic field of 0.25 mG (i.e., the maximum field value in Table B.3; this is a large overestimate considering the tested samples had volumes much larger than any component in the wires) and is a magnetic dipole positioned 15 mm away from a qubit. The field generated by the wire at the qubit will then be $B_q \simeq 0.25 r_0^3 / 0.015^3$ mG, where $r_0 \simeq 10$ mm

¹Note that we also performed magnetic tests by exposing all samples to a ultra-high pull neodymium rectangular magnet, with dimensions 25.4 mm \times 25.4 mm \times 9.5 mm and a pull of 10.4 kg. We found magnetic fields with the same order of magnitude as in Table B.3.

Table B.2: ZGC calibration. The margins of error indicated in parentheses were estimated from the fluctuation of the magnetic sensor.

Measurement (-)	$\ \vec{B}\ $ (mG)	α (-)
Vertical position, background field	554(20)	-
Vertical position, with ZGC	0.66(5)	842(34)
Horizontal position, background field	539(20)	-
Horizontal position, with ZGC	0.59(5)	917(44)

is the distance at which the field was measured in the ZGC; thus, $B_q \simeq 0.075$ mG. Assuming an Xmon qubit with a superconducting quantum interference device (SQUID) of dimensions $40 \mu\text{m} \times 10 \mu\text{m}$ (cf. Ref. [92]), the estimated magnetic flux due to the wire threading the SQUID is $\Phi_q \simeq 4 \times 10^{-18}$ Wb. This is approximately three orders of magnitude smaller than a flux quantum $\Phi_0 \simeq 2.07 \times 10^{-15}$ Wb; typical flux values for the Xmon operation are on the order of $0.5 \Phi_0$.

B.3 Thermal Conductance

In this appendix, we describe the method used to estimate the thermal performance of a three-dimensional wire and compare it to that of an Al wire bond. Note that at very low temperature, thermal conductivities can vary by orders of magnitude between two different alloys of the same material. The following estimate can thus only be considered correct to within approximately one order of magnitude. Thermal conductivity is a property intrinsic to a material. To characterize the cooling performance of a three-dimensional wire, we instead use the heat transfer rate (power) per kelvin difference, which depends on the conductivity.

The power transferred across an object with its two extremities at different temperatures depends on the cross-sectional area of the object, its length, and the temperature difference between the extremities. Since the cross-section of a three-dimensional wire is not uniform, we assume the wire is made of two concentric hollow cylinders. The cross-sectional area of the two cylinders is calculated by using dimensions consistent with those of a three-dimensional wire. The inner and outer hollow cylinders are assumed to be made of phosphor bronze and brass alloys, respectively. The thermal

Table B.3: Magnetic field measurements of the materials used for the main components of the quantum socket. The tested samples are significantly larger than any component used in the actual implementation of the three-dimensional wires and microwave package. The margins of error indicated in parentheses were estimated from the fluctuation of the magnetic sensor.

Material	$\ \vec{B}\ $ (mG)
CW724R	0.21(5)
CW453K	0.25(5)
Al 5N5	0.02(5)

Table B.4: Parameters used in the estimate of the heat transfer rate per kelvin difference for a three-dimensional wire and an Al wire bond. In the table are reported: The hollow cylinder inner diameter d_i ; the hollow cylinder outer diameter and wire bond diameter d_o ; the hollow cylinder and wire bond cross-sectional area A ; the thermal conductivity k_t .

	d_i (μm)	d_o (μm)	A (m^2)	k_t ($\text{mW K}^{-1} \text{m}^{-1}$)
Inner conductor (phosphor bronze)	290	380	4.74×10^{-8}	3.7
Outer conductor (brass)	870	1290	7.13×10^{-7}	24.1
Wire bond (Al)	-	50	1.96×10^{-9}	0.01

conductivities of these materials at low temperatures are determined by extrapolating measured data to 25 mK².

The Al wire bonds are assumed to be solid cylinders with diameter 50 μm . In the superconducting state, the thermal conductivity of Al can be estimated by extrapolating literature values [305].

The heat transfer rate per kelvin difference is calculated by multiplying the thermal conductivity k_t with the cross-sectional area A and dividing by the length of the thermal conductor ℓ . The heat transfer rate per kelvin difference of a three-dimensional wire is calculated by summing the heat transfer rate per kelvin difference of the inner conductor to that of the outer conductor and is found to be $\Pi_t \simeq 6 \times 10^{-7} \text{ W K}^{-1}$ at 25 mK. At the same temperature, the heat transfer rate per kelvin of a typical Al wire bond is estimated to be $\Pi_b \simeq 4 \times 10^{-12} \text{ W K}^{-1}$ (cf. Table B.4), much lower than for a single three-dimensional wire. Note that, instead of Al wire bonds, gold wire bonds can be used. These are characterized by a higher thermal conductivity because they remain normal conductive also at very low temperatures. However, Al wire bonds remain the most common choice because they are easier to use.

B.4 Thermo-Mechanical Tests

In this appendix, first we describe the test setup used to find the mean number of cycles before mechanical failure of a three-dimensional wire and show images of the wire after testing; then, we discuss the performance of the springs used in three-dimensional wires at various temperatures.

In order to obtain the mean number of cycles before mechanical failure of a three-dimensional wire we used an automated pneumatic system, which makes it possible to compress the wire a very large number of times. The wire under test was operated at a stroke of 2.0 mm; the test cycle time was 120 strokes per minute; finally, the entire test took place at a temperature of 20 °C. The test was run for approximately 28 h, for a total of 200000 wire compressions. Both the inner and outer conductor of the wire were mechanically functioning properly at the end of the test; mechanical abrasion was visible, even though the overall wire condition was excellent, as shown by the two images in Fig. B.1. From tests on wires with a similar form factor, but made from different materials, we are confident that the number of cycles before DC and microwave electrical failure of our wires will also exceed 200000.

The three types of tested springs are called FE-113 225, FE-112 157, and FE-50 15 and their geometric characteristics are reported in Table B.5. We ran temperature cycle tests by dunking the springs repeatedly in liquid nitrogen and then in liquid helium without any load. At the end of each cycle, we attempted to compress them at room temperature. We found no noticeable changes in mechanical performance after many cooling cycles. Subsequently, the springs were tested mechanically by compressing them while submerged in liquid nitrogen or helium. The setup used for

²Confer http://www.lakeshore.com/Documents/LSTC_appendixI_1.pdf.

Table B.5: Thermo-mechanical tests on hardened BeCu springs. In the table are reported: The outer diameter D of the coil forming the helix structure of the spring; the diameter d of the circular cross-section of the spring (note that the smallest wire diameter is 150 μm); the spring free length L_f , i.e., the spring length at its relaxed position; the number of coils N_c ; the spring force F_c (estimated at all operating temperatures).

Spring type	D (mm)	d (mm)	L_f (mm)	N_c (-)	F_c (N)
FE-113 225	2.30	0.26	11.55	11.25	~ 1.0
FE-112 157	1.30	0.22	18.00	42.00	~ 1.0
FE-50 15	0.60	0.15	31.75	150.00	~ 0.5




Figure B.1: Side and top view images of a three-dimensional wire after 200000 compression cycles.

the compressive loading test of the springs is shown in Movie S4 of the Supplemental Material of Ref. [220], which also shows a properly functioning spring immediately after being cooled in liquid helium. In these tests, we only studied compression forces because in the actual experiments the three-dimensional wires are compressed and not elongated.

The compression force was assessed by means of loading the springs with a mass. The weight of the mass that fully compressed the spring determined the spring compression force F_c . The compression force of each spring is reported in Table B.5. We observed through these tests that the compression force is nearly independent of the spring temperature, increasing only slightly when submerged in liquid helium. Assuming an operating compression $\Delta L = 2.0$ mm, we expect a force between 0.5 N and 2.0 N for the inner conductor and between 2.0 N and 4.0 N for the outer conductor of a three-dimensional wire at a temperature of 10 mK. Note that we chose spring model FE-113 225 for use with the grounding washer.




Figure B.2: Micro images showing three-dimensional wire alignment errors. (a)-(b) Au pads. The pad displayed in (a) is connected to that in (b) by way of a CPW transmission line approximately 11.5 mm long. The die shifted upward between the first (green arrow (1)) and second (green arrow (2)) mating instance, resulting in a lateral misalignment for the bottom pad. The rotational misalignment for the pad in (a) is indicated by a dashed green line. (c) Successful alignment for six Ag pads on the same chip. (d) Peripheral area of an Ag sample (ground plane). The marks are due to contact with the grounding washer.

B.5 Alignment Errors

In this appendix, we provide more details about alignment errors. Figure B.2 shows a set of micro images for Au and Ag samples. The Au pads in panels (a) and (b) were mated two times at room temperature; the three-dimensional wires used to mate these pads featured the smaller tunnel (500 μm width). The pad dimensions were $W_p = 230 \mu\text{m}$ and $T_p = 1000 \mu\text{m}$. Noticeably, for the pad in panel (a) the wire bottom interface

matched the contact pad in both mating instances, even though the matching was affected by a rotational misalignment of approximately 15° with respect to the transmission line longitudinal axis. However, for the pad in panel (b) the lateral misalignment was significant enough that the inner conductor landed on the dielectric gap in the second mating instance.

In our initial design, a perfect match required that the die dimensions should be at most 1 thou smaller than the dimensions of the chip recess, as machined. In the case of the sample holder used to house the Au samples, the chip recess side lengths were 15.028(5) mm, 15.030(5) mm, 15.013(5) mm, and 15.026(5) mm. The Au samples were diced from a Si wafer using a dicing saw from DISCO, model DAD-2H/6, set to obtain a 15 mm \times 15 mm die. Due to the saw inaccuracies, the actual die dimensions were 14.96(1) mm \times 14.96(1) mm, significantly smaller than the chip recess dimensions. This caused the die to shift randomly between different mating instances, causing significant alignment errors.

As described in Sec. 4.3.4, in order to minimize such errors a superior DISCO saw was used, in combination with a DISCO electroformed bond hub diamond blade model ZH05-SD 2000-N1-50-F E; this blade corresponds to a nominal kerf between 35 and 40 μm . Additionally, we used rotational as well as lateral aligning markers; the latter were spaced with increments of 10 μm that allowed us to cut dies with dimensions ranging from 14.97 mm to 15.03 mm, well within the machining tolerances of the sample holder. After machining, the actual inner dimensions of each sample holder were measured by means of a measuring microscope. The wafers were then cut by selecting the lateral dicing markers associated with the die dimensions that fit best the holder being used.

Figure B.2 (c) shows a successful alignment for six Ag pads on the same chip; the chip is mounted in a sample holder with grounding washer. All three main steps for an ideal and repeatable alignment (cf. Sec. 4.3.4) were followed. Figure B.2 (d) shows the distinctive marks left by the grounding washer on an Ag film. The marks are localized towards the edge of the die; the washer covered approximately 500 μm of Ag film. This indicates a good electrical contact at the washer-film interface.

In conclusion, it is worth commenting some of the features in Fig. 4.5 (d) of Sec. 4.3.3. The figure clearly shows dragging of a three-dimensional wire due to cooling contractions. In fact, for the Al chip recess an estimate of the lateral contraction length from room temperature to ~ 4 K can be obtained as

$$\Delta L_{\text{Al}} = \alpha(4)L_{\text{Al}} \simeq (4.15 \times 10^{-3}) (15 \times 10^{-3} \text{ m}) \simeq 62 \mu\text{m} \quad (\text{B.1})$$

where $\alpha(4)$ is the integrated linear thermal expansion coefficient for Al 6061-T6³ at 4 K from Refs. [241] and L_{Al} is the room temperature length of the recess side⁴.

³ISO AlMg1SiCu; UNS A96061.

⁴Note that $\alpha(4)$ can be accurately estimated from the data on Al 6061 at the NIST Cryogenic Materials Property Database [306] (archived version).

Note that the sample holder is actually made from Al alloy 5N5; however, different Al alloys contract by approximately the same quantity. For the Si sample substrate, the lateral contraction length from room temperature to ~ 4 K is approximately given by $\Delta L_{\text{Si}} \approx 3.2 \mu\text{m}$, where the integrated linear thermal expansion coefficient at 4 K was found in Table 2 of Ref. [307]. Below 4 K, the thermal expansion of both materials is negligible for our purposes and, thus, the 4 K estimate can also be considered to be valid at ~ 10 mK.

B.6 Sample Fabrication

In this appendix, we outline the fabrication processes for the samples used to test the quantum socket. A set of samples was made by liftoff of a $\sim 3 \mu\text{m}$ Ag film, which was grown by means of EBPVD (system from Intlvac Canada Inc., model Nanochrome II) on a 3 inch float-zone (FZ) Si (100) wafer of thickness $500 \mu\text{m}$. The superconducting Al on Si samples were made by etching a ~ 120 nm Al film that was deposited by EBPVD on a $500 \mu\text{m}$ FZ Si wafer. The Al on sapphire sample was made by etching a ~ 100 nm Al film that was deposited by UHV EBPVD (system from Plassys-Bestek SAS) on a $500 \mu\text{m}$ *c*-plane single crystal sapphire wafer. Prior to deposition, the wafer was annealed *in vacuo* at approximately 850°C , while being cleaned by way of molecular oxygen; a 1 nm germanium buffer layer was grown at room temperature, before depositing the Al film. Last, two sets of test samples were made by etching Au films of thickness 100 nm and 200 nm, with a 10 nm Ti adhesion underlayer in both sets. The films were grown by EBPVD (Intlvac) on a 3 inch Czochralski (CZ) undoped Si (100) wafer of thickness $500 \mu\text{m}$.

The $3 \mu\text{m}$ Ag samples were required to reduce the series resistance of the CPW transmission lines (cf. Subsecs. 4.4.2, 4.4.3, and 4.4.4). Fabricating such a relatively thick film necessitated a more complex process as compared to that used for the Au and Al samples. The Ag samples were fabricated with a thick resist tone reversal process. The wafer was spun with an AZ P4620 positive tone resist to create a resist thickness of $\sim 14 \mu\text{m}$, then soft baked for 4 min at 110°C . Because the resist layer is so thick, a rehydration step of 30 min was necessary before exposure. Optical exposure was performed for 30 s in a mask aligner from SÜSS MicroTec AG, model MA6, in soft contact with a photomask. After exposure the sample was left resting for at least 3 h so that any nitrogen created by the exposure could dissipate. The tone reversal bake was done for 45 min in an oven set to 90°C , filled with ammonia gas. The sample then underwent a flood exposure for 60 s and was developed in AZ® 400K for 15 min. Subsequently, $3 \mu\text{m}$ of Ag was deposited and liftoff of the resist was performed in acetone for 5 min with ultrasounds.

Appendix for Chapter 5

C.1 Qubit Flux Pulse

Performing swap spectroscopy requires the ability to set the qubit's frequency to a desired value f_p for a particular duration. In our experiment, this is done with a flux pulse applied to the SQUID with an AWG. As shown in Fig. 3.5, the flux pulse reaches the sample after going through multiple stages of filtering, attenuation, and connections. This means that the waveform will be modified compared to what is generated by the AWG.

In addition, while we control the *amplitude* of the pulse, we are ultimately interested in the resulting *frequency* of the qubit. We therefore need a way to convert between the amplitude A of the flux pulse and the qubit probe frequency f_p . This can be done, for example, with pulse spectroscopy, where we send π -pulses to the qubit at different frequencies while it is detuned by a flux pulse of a particular amplitude. The qubit frequency for that amplitude can then be fit. This is repeated for many amplitudes in order to get a map between A and f_p .

The above considerations mean that the measurement settings (f_p, t) that we select may contain multiple kinds of potential errors. This must be taken into account when estimating the coupling parameters with either the offline or online algorithms.

For example, if f_p is higher than the true probe frequency of the qubit due to some systematic error in the amplitude-frequency map, the result f_{RM} reported by the online algorithm is higher than the true value as well. Similarly, an error on the value of t leads to a wrong estimation of g . In practice, this kind of systematic error is not a major problem as long as the error is consistent between experiments. For our application, for example, it does not matter that the estimated frequency of the resonance is not truly 4.8305 GHz. The quantity of interest is the flux amplitude corresponding to the mode.

Another kind of error would occur if the flux pulse amplitude was not constant in time, for example, due to filtering. This would result in a time-dependent probe frequency f_p over the length of the pulse. Assuming that the time-dependence is not so strong that the qubit frequency changes by an amount $\sim g$ in less than a Rabi period, the consequences for the octave algorithm should be negligible. This is because the detection is made via the first minimum of the swap oscillation. The online estimation algorithm would be affected more strongly, since the model assumes that the qubit probe frequency is constant. The error on the parameters estimated by an experiment suffering from this issue, as computed by the repeated inference method discussed in Sec. 5.5.1, would be expected to capture this effect.

To improve accuracy, it is important to properly calibrate the flux pulses prior to the experiment, ensuring that the qubit frequency is stable. This calibration can be done, again, with pulse spectroscopy, by mapping out the qubit frequency over the length of the pulse. Deviations can be subsequently corrected by modifying the shape of the pulse emitted by the AWG.

C.2 Details on Octave Analysis

The task of detecting resonances in data gathered with the octave sampling algorithm (the data plotted in Fig. 5.3) is done in two sub-steps: (1) Find peaks in each octave. (2) Merge the octaves together. The source code implementing this analysis as well as the particle generation procedure can be found online, in the `OctaveAnalysis` module of the `TLSInfer` package [263].

First, the average bin excitation is reversed in order to transform the low-excitation regions into peaks. Then, we find peaks in each octave individually. This is done by looking for points in the data that are *prominent* when compared to neighboring points. Each peak must be taller than the lowest point to the left and right by a certain amount, the *prominence value*. Decreasing this value leads to more peaks being detected and therefore increases the detection sensitivity. The result of the peak finding step for the highest octave of the data is shown in Fig. C.1.

The second step is to check if the peaks detected in different octaves correspond to the same resonance. This is done by comparing the frequency position of the peaks found. If the frequencies are close, we assume that the peaks correspond to the same resonance and the detections are therefore merged. When merging detections from different octaves, the one from the lower octave is preferred. This is because the coupling strength is estimated from the first oscillation minimum, which corresponds to the lowest octave detected.

When the procedure is complete, we are left with a single peak location per resonance, along with the lowest octave where it is found. This gives a coarse estimate of

the frequency and coupling strength of that resonance, as is reported in the main text in Table 5.1.

There are a few possible causes of errors with this procedure. The most likely error occurs when two detections corresponding to a single resonance are reported, or, on the contrary, when two resonances close in frequency are reported as one. The first issue is not very problematic, since further analysis done with the Bayesian online algorithm will likely correctly report that the two detections come from the same resonance. The second is solved by increasing the resolution of the octave data collection.

C.3 Particle Distribution Generation

While a prior particle distribution for the online Bayesian algorithm can be manually specified, it is much more convenient to generate one from available octave data. The first step in the process is to divide the octave data in multiple spectra such that we are left with one resonance mode per spectrum. To split the spectrum, we simply cut the data between each detected resonance (see Appendix C.2). If this cut happens to fall within a bin, we slice the bin into two sub-bins, keeping track of each part's proportion.

Following this, we can restrict ourselves to a frequency-time spectrum containing a single resonance mode. To specify a discrete distribution representing our knowledge of the coupling parameters, we follow a procedure where we pick bins from the spectrum and generate a particle according to the bin parameters. We pick from bins according to a weight, which we choose to be proportional to $1 - \bar{P}_e$, but only for those above a particular threshold. We ignore bins below the threshold. Thus, bins for which the




Figure C.1: Peaks found in the highest octave of Fig. 5.3 (b) with the prominence value set to 0.09. There are many peaks found for RM3. They are merged when the detection analysis proceeds to the lower octaves.

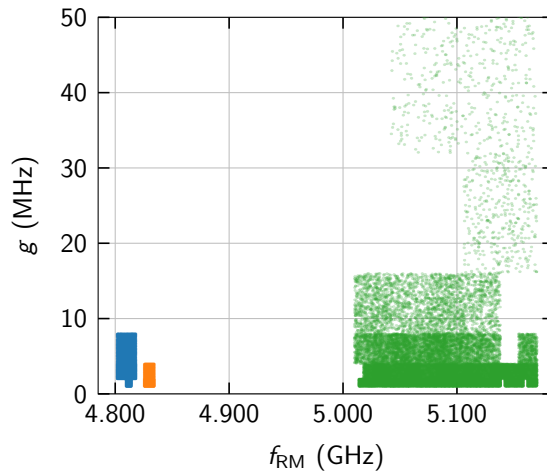


Figure C.2: Initial particle distribution generated from the octave data in Fig. 5.3 (b) after splitting the spectrum in according to the detections. The distribution corresponding to RM1 is plotted in blue, the one for RM2 in orange, and the one for RM3 in green.

average excitation \bar{P}_e is low have a high weight and have a higher chance of being chosen for generating particles.

Each particle represents a two-tuple of frequency–coupling–strength values. To generate a particle, we must choose those two values. The frequency of the generated particle is chosen uniformly at random within the frequency range of the bin drawn. The coupling strength is picked from the uniform distribution $\mathcal{U}(\Delta_m, 2\Delta_m)$, where Δ_m is the bin width. We repeat this procedure (draw a bin, generate a particle) to create as many particles as is desired for the distribution.

Exemplary particle distributions generated for RM1, RM2, and RM3 can be seen in Fig. C.2, with their statistics tabulated in Table C.1. These distributions can be used as the starting point for the online algorithm discussed in Sec. 5.5. In our experiment we use distributions comprising 40000 particles. This number can be adjusted depending on the capabilities of the computer executing the algorithm.

Table C.1: Frequency and coupling parameters statistics of the particle distributions generated for RM1, RM2 and RM3. A plot of those distributions can be seen in Fig. C.2.

Parameter	RM1	RM2	RM3
f_{RM} (GHz)	4.811(4)	4.830(2)	5.088(42)
g (MHz)	3(2)	2(1)	4(7)

C.4 Bayesian Resampling Procedure

The exact resampling procedure is described in detail in Chap. 2 of Ref. [256]. We reproduce a shortened version here for completeness.

As inputs, we require the particle locations, as an array of frequency–coupling–strength tuple, in addition to their likelihoods, as computed in the main text. Note that the likelihoods must be normalized to sum to unity, after which they become *weights*. This normalization ensures that the weights array is a valid discrete probability distribution that can be drawn from.

The algorithm then draws particles from the input distribution according to the weights, and, from the position of the particles drawn, generate new ones by adding “noise.” This prevents having duplicates in the output distribution, even if a particle from the input is drawn multiple times. The resampling algorithm is shown in pseudocode in Algorithm 1. Note that the particle at index i in the $\{\vec{x}_k\}$ array has its corresponding weight at the same index in the $\{w_k\}$ array. In addition, the amount of noise added to the position of the particle drawn is controlled by the spread of the input distribution, quantified by taking the covariance. For more details, see Ref. [256].

Algorithm 1: Particle Resampling

Input: Array of particle positions $\{\vec{x}_k\}$ and weights $\{w_k\}$

Output: New particle positions $\{\vec{y}_k\}$

Function *resample*($\{\vec{x}_k\}, \{w_k\}$)

$a = 0.98$

$\vec{\mu} = \text{mean}(\{\vec{x}_k\})$

$\Sigma = (1 - a^2)\text{cov}(\{\vec{x}_k\})$

for $i \in 1 : n$ **do**

$l = \text{rand}(\text{Discrete}(\{w_k\}))$

$\vec{\mu}_l = a\vec{x}_l + (1 - a)\vec{\mu}$

$\vec{y}_i = \text{rand}(\text{Normal}(\vec{\mu}_l, \Sigma))$

end

return $\{\vec{y}_k\}$

end




Figure D.1: Scatter plot of T_1 vs. f_q and t for dataset 2 [same dataset as in Fig. 6.1 (c)] at the actual measurement time; the color map for T_1 is the same as in Fig. 6.1. Note that the vertical axis is truncated at $t = 5$ h to display the relative measurement times more clearly.

D

Appendix for Chapter 6

D.1 Experimental Details

The spectrotemporal charts displayed in Sec. 6.4 can be interpreted as matrices of T_1 values, with m rows and n columns; m and n represent a time and frequency index, respectively. The $(1, 1)$ entry is the bottom-left element of the matrix, such that time increases from bottom to top. We set f_q from low to high values, completing one row

of each matrix when reaching the highest value of f_q . Subsequent rows are measured restarting always from the lowest value of f_q . Hence, the time $t_{m,n}$ at which each data point (m, n) is taken increases from left to right for the m -th row, starting at $t_{m,1}$ and ending at t_{m,N_f} . The time difference between subsequent rows is a constant value defined as $\Delta t = t_{m+1,1} - t_{m,1}$. Although each measurement in any particular row is taken at a different time, we choose to display the data on a rectangular matrix where each row element is associated with the same time value. As a comparison, Fig. D.1 shows a scatter plot for which each T_1 value is plotted at the actual measurement time. This figure elucidates two limitations of our experiments: (1) The impossibility to measure an entire row at exactly the same time. (2) The fact that $t_{m,N_f} \sim t_{m+1,1}$. It additionally stresses a difference between experiments and simulations, i.e., the fact that in simulations all row elements are calculated at the exact same time.

In order to keep Δt constant we must account for experimental nonidealities. The time required to perform a single T_1 experiment is $t_{\text{exp}} \approx 16$ s and varies slightly between experiments. In addition, latencies in the electronic equipment when setting a new value of f_q result in a short time overhead. To overcome these issues, we measure a test row and record the corresponding measurement time. We then augment this measurement time by a certain buffer time, which we estimate to be sufficiently longer than any possible time variations due to nonidealities. The sum of the measurement time of the test row and the buffer time is Δt . For example, for the dataset shown in Fig. 6.1 (c), the time elapsed to acquire the data of the test row is approximately 992 s. In this case, we choose $\Delta t = 1000$ s. The values of Δt for each dataset shown in Sec. 6.4 are reported in Table 6.1.

D.2 Long-Time Stability

One of the assumptions in Sec. 6.3 is that the TLS parameters do not change in time, i.e., they are considered to be *static*. Thus, the only dynamically varying quantity is the state of a TLS. In order to show that this is a reasonable assumption, in Fig. D.2 we display three experimental time series measured at $f_q = 4.529$ GHz. The first time series corresponds to a column extracted from the spectrotemporal chart of Fig. 6.1 (a); the second series is an additional trace not included in the spectrotemporal charts because too short compared to the other traces; finally, the third series is a column from Fig. 6.1 (c). Each point in the three series is plotted at the actual time at which it is measured relative to the first point of the first series. It is worth noting that the frequency of these time series is not captured in the spectrotemporal chart of Fig. 6.1 (b).

The three time series are measured over the course of approximately three weeks. Despite the large time gap between the first and third series, we observe a similar T_1 -




Figure D.2: Time series of T_1 vs. t showing the relative time between measurements for dataset 1 (dark blue) and 3 (yellow) [same datasets as in Fig. 6.1 (a) and (c)] and an additional dataset (green); all datasets are for $f_q = 4.529$ GHz.

drop pattern: The T_1 times are distributed around two values, 5 and 23 μs . These results indicate a reproducible feature and suggest a static TLS distribution.

It is well known that by cycling the sample temperature, e.g., when warming up and cooling back down a device, results in a strain field that can modify the TLS parameters. However, when operating a sample at a constant temperature and without exceedingly large excitation electric fields (as in the experiments reported in this work), we expect a static TLS distribution.

D.3 Qubit Electric Field

As explained in Section 3.1, the qubit capacitor is a Greek cross formed by two CPW strips of length L . Since $L \gg S + W$, we approximate the qubit capacitor as a CPW segment of infinite length; we additionally assume that the capacitor is made of an infinitesimally thin conducting sheet. When determining \vec{E}_q , we can thus restrict ourselves to points within the CPW vertical cross-section.

We determine \vec{E}_q by means of a conformal mapping technique. A conformal map is a function that locally preserves angles, allowing us to transform the CPW geometry into that of a much simpler infinite parallel-plate capacitor; the map function is given by Eq. (25) in the work of Ref. [308]. We then use this map to transform the electric field of the parallel-plate capacitor into that of the CPW. The electric field is proportional to the qubit electric potential with respect to ground, or zero-point voltage; given the




Figure D.3: Qubit electric field $\|\vec{E}_q\|$ for $\phi_0 = 1$ V vs. width x at one value of the height $z = 1.5$ nm. The origin of the graph is at $x = 0$, corresponding to the middle point of the strip. Due to the symmetry of the CPW segment with respect to its longitudinal axis (i.e., the y -axis; not shown), we display $\|\vec{E}_q(x)\|$ only for half of the CPW segment, for $x \geq 0$. The extent of the conducting sections of the CPW is indicated by the thick blue lines. The dashed black vertical lines are placed at the edge of each conductor; the left line corresponds to the edge of the strip and the right line to the edge of the ground plane.

qubit plasma frequency $f_p = \sqrt{8E_J E_c}/h$, the zero-point voltage reads

$$\phi_0 \simeq \sqrt{\frac{hf_p}{2C_q}} = \frac{e}{C_q} \left(\frac{E_J}{2E_c} \right)^{1/4} \sim 4 \mu\text{V}. \quad (\text{D.1})$$

In order to generate g , we evaluate $\|\vec{E}_q\|$ at randomly picked points (x, z) corresponding to Q-TLS positions. These points are confined within the cross-section region introduced above. The cross-section is centered on the middle point of the strip and has a length of $96 \mu\text{m}$ and a height of 3 nm; the left and right edges of the cross-section extend $12 \mu\text{m}$ into the ground plane and the top edge corresponds to the oxide layer's top edge. Figure D.3 shows $\|\vec{E}_q(x, z)\|$.