

Abstract

Building Blocks for Modular Circuit QED Quantum Computing

Christopher James Axline

2018

The ability to couple isolated quantum systems demonstrates precise and powerful control over natural phenomena. It can be of use in quantum information by connecting systems in a modular network to realize distributed, fault-tolerant quantum computing. This requires highly performing modules and efficient communication channels. We have developed or integrated state-of-the-art hardware in circuit quantum electrodynamics to produce such an elementary module. We explore the dissipation mechanisms that limit their fundamental performance; by reducing overall loss, we can increase the efficiency of a scaled modular network. Further, we demonstrate a method to transmit quantum information from a module, via propagating photons, into another module. Transmitting quantum information requires that locally stored states are converted into propagating fields. Using RF-controlled four-wave mixing, we experimentally demonstrate parametric conversion of quantum states from a superconducting microwave cavity, preserving the quantum information in a propagating form with high fidelity. This process operates identically in either direction, allowing release and capture of the propagating mode by temporal control over the conversion. Such control also permits entanglement, simply by “half-release” of a state, which we confirm by observing non-classical correlations. Thus, state transfer and entanglement are both realized with high fidelity and without compensating by conditioning. This fundamentally deterministic strategy sets new limits on the rate of communication and entanglement generation. By encoding quantum states in multi-photon fields, we show how one can overcome signal transmission loss in such a quantum network, further improving efficiency through the prospect of quantum error correction. These results establish a compelling approach for deterministic, network-based quantum computation, and may serve as the basis for how superconducting quantum circuits can be scaled to even greater numbers.

Building Blocks for Modular Circuit QED Quantum Computing

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by
Christopher James Axline

Dissertation Director: Professor Robert J. Schoelkopf

December, 2018

Copyright © 2019 by Christopher James Axline
All rights reserved.

Contents

Contents	iii
List of Figures	iv
List of Tables	vii
Acknowledgments	ix
1 Introduction	1
1.1 Overview of this thesis	3
2 A Quantum Information Roadmap	6
2.1 Quantum computers	6
2.2 The research roadmap	8
2.3 Encoding and protecting quantum information	11
2.4 Scaling up computing power	17
3 Circuit QED for Modular Quantum Networks	20
3.1 A computation framework inspired by networks	21
3.2 Distributed computing techniques and advantages	23
3.3 Construction of a module	26
3.4 State-of-the-art circuit QED components	28
3.5 Useful cQED operations	42
4 Seamless Hardware for Scalable Integration	51
4.1 Challenges of many-element system coherence	51
4.2 Microwave hygiene	52
4.3 Hardware for coherent 3D integration	56
5 Identifying Dissipation Sources in cQED Modules	87
5.1 Energy participation ratios	89
5.2 Modeling and simulating bulk dielectric participation	91
5.3 Evaluating loss in 3D structures	96
5.4 Loss in planar structures	106
5.5 Varying participation ratios in coaxlines	115
5.6 Measurement variability	133

5.7	Varying participation ratios in transmons	139
5.8	Summary of loss mechanisms and bounds	146
6	Letting Schrödinger’s Cat Out of the Bag	150
6.1	Parametric pumping	151
6.2	Experimental setup	161
6.3	Quantifying conversion performance	168
6.4	Entanglement with flying photons	179
6.5	Conclusion	187
7	Putting Schrödinger’s Cat Back in the Bag: Direct Quantum State Transfer and Entanglement Distribution	189
7.1	Extending conversion to two nodes of a network	190
7.2	Wavepacket shaping	199
7.3	State transfer	203
7.4	Half-release, full-catch, and entanglement	222
7.5	System imperfections	227
7.6	The future of state transfer and stationary–propagating conversion	238
8	Conclusion	240
8.1	Perspectives	240
A	Appendix	243
A.1	Analytical approximation of partially loaded coaxline waveguide	244
A.2	Driven transmon Hamiltonian transformation	245
A.3	Detrimental effects of pumped processes	249
A.4	Input-output theory for multi-system readout	255
	Bibliography	258

List of Figures

2.1	A research roadmap	10
2.2	The Bloch sphere	12
2.3	Two forms of information media	14
3.1	Connected quantum modules	24
3.2	Configurations of scalable devices	27
3.3	The transmon qubit	31
3.4	Dispersive readout trajectories of a transmon	35
3.5	Readout histograms	35
3.6	Basic resonator measurement configurations	36
3.7	Response in different resonator configurations	39
3.8	Resonator power dependence	41
3.9	Feedback cooling process	43
3.10	Number peak resolution of the cavity state	46
3.11	Direct spectroscopy of a storage cavity	47
4.1	Possible hazards in planar architectures	53
4.2	The coaxline architecture	57
4.3	Behavior of a dielectric-loaded waveguide	59
4.4	Coaxline cutoff frequencies	60
4.5	Coaxline as a quasi-stripline geometry	61
4.6	Coaxline electric field profile	62
4.7	Coaxline impedence	63
4.8	Evanescent coupling to 3D enclosures	64
4.9	Effect of a dissipative coupler	65
4.10	Coaxline transmission measurement	66
4.11	Serpentine stripline properties	68
4.12	Serpentine stripline electric field profile	70
4.13	Coaxline–transmon coupling	72
4.14	Coaxline Purcell filter	75
4.15	Purcell filter Q-ratio	76
4.16	Coaxline integrated with coaxial stub cavity	79
4.17	Multiplexed coaxline enclosure	80
4.18	Experimental scheme for testing cross-coupling	82
4.19	A coaxline variation with seams	83

4.20	Seam loss in the seammux	84
4.21	Ideas for complex planar–3D integration	85
5.1	Resonant circuit with many losses	90
5.2	Effective parallel-plate capacitor model for planar structures	95
5.3	Typical 3D cavity geometries	99
5.4	Coaxial stub participation simulations	100
5.5	Bulk metal surface participations	102
5.6	Best-in-class 3D cavity measurements	103
5.7	Experiment testing adverse effects of moisture	104
5.8	Performance of a coaxial stub cavity over repeated cooldowns	107
5.9	Candidate dissipative surface layers	108
5.10	Dividing a 3D simulation into parts	110
5.11	Using 2D electrostatics to scale 3D simulations	112
5.12	Perimeter energy scaling factors in the coaxline	114
5.13	Simulated coaxline surface layer losses	116
5.14	Aggregate coaxline measurements compared to participations	118
5.15	Coaxline performance from combined loss	119
5.16	Variations in coaxline diameter	120
5.17	Variations in coaxline stripline width	120
5.18	Variations in coaxline substrate width	121
5.19	Geometry of multiple-stripline devices	123
5.20	Quality factors from multiple-stripline device measurements	124
5.21	Multiple-stripline performance from combined loss	125
5.22	Varying temperature to extract enclosure kinetic inductance fraction	127
5.23	Measured and simulated values of kinetic inductance fraction	127
5.24	Quality factor histogram, comparing processes	129
5.25	Profile of an etched coaxline	131
5.26	Comparing substrate diced edge quality	133
5.27	Coaxline quality factors over repeated cooldowns	134
5.28	Transmon coherence properties over repeated cooldowns	135
5.29	Coaxline sensitivity to perturbations in placement	136
5.30	Symptoms of mechanical vibration	136
5.31	Modeling mechanical vibration	137
5.32	Quality factor histogram, comparing packages	138
5.33	Reconciling 2D and 3D simulations of transmon leads	140
5.34	Energy participation from junction outward	142
5.35	Transmon lifetime for varied geometry	143
5.36	Comparative scaling of surface losses with geometry	144
5.37	Measured lifetimes of assorted devices, compared to metal–substrate interface participations	148
6.1	Launching a quantum state from a node	155
6.2	Frequency ordering in a conversion experiment	156
6.3	Realization of a basic circuit QED module	159
6.4	Cavity damping analysis	160

6.5	Experimental setup	164
6.6	Calibration of Stark shifts	166
6.7	Calibration of conversion rate	167
6.8	Calibration of detector efficiency	168
6.9	Cavity damping by mode-conversion	170
6.10	Cavity damping for higher Fock states	171
6.11	Measuring the field of converted, propagating states	174
6.12	Traveling multi-photon quantum states	176
6.13	Traveling cat states	177
6.14	Kerr effect during half-release of a cat-state	179
6.15	Generating entanglement between stationary and traveling fields	180
6.16	Entanglement between a stationary and flying photon	181
6.17	Entanglement bound by witness	183
6.18	Conditioned Q-functions after half-release of $ 2\rangle$	185
6.19	Entanglement between a stationary and flying cat state	186
7.1	Connecting two modules for on-demand state transfer	190
7.2	Mode frequencies and alignment	191
7.3	Schematic of the experiment wiring for a full pitch-and-catch experiment	192
7.4	Images of device used in pitch-and-catch experiment	194
7.5	Mode spectroscopy	195
7.6	Transmon measurement independence	197
7.7	Pump leakage between systems	198
7.8	Temporal mode-matching of sender and receiver	205
7.9	Spectroscopy to determine transfer efficiency	206
7.10	Cavity populations after pitch and catch of single and many-photon states	207
7.11	Cavity populations over time	208
7.12	Reconstructing Fock manifold Wigner functions	210
7.13	Bloch sphere representation of Fock manifold transfer	211
7.14	Photon loss process matrix	213
7.15	Reconstructing binomial manifold Wigner functions	215
7.16	Binomial encoding loss trajectories	216
7.17	Comparison of trivial and error-corrected fidelities	217
7.18	Conditioned spectroscopy of transferred binomial codewords	218
7.19	Conditioned tomography of transferred binomial codewords	220
7.20	Modified design to permit error correction	221
7.21	Entangling a photon between remote cavities	223
7.22	Comparison with post-selected efficiency	230
7.23	Transmission inefficiency from sender transmon excitation	233
7.24	Effect of transmon excitation on the wavepacket	234
7.25	Estimation of transmission loss	237
A.1	Transmon heating from applied pumps	254
A.2	Two-cavity readout trajectories	257

List of Tables

5.1	Bulk dielectric loss tangents	92
5.2	Results of water ice measurement	105
5.3	Serpentine resonator simulated surface losses	122
5.4	Results from silicon devices	130
5.5	Results from etched silicon devices	131
5.6	Summary of material bounds	146
6.1	Conversion system Hamiltonian parameters	163
7.1	Measured pitch-and-catch system parameters	193
7.2	Measures of entanglement	226
7.3	Summary of transfer losses	231
A.1	Pump strengths at varied device couplings	255

Acknowledgments

I am grateful to many people without whom this work, and the experience I associate with it, would not have come to exist. They have mentored me in physics, and many other aspects of life, too.

I will begin with my committee, the members of which have been kind enough to guide me throughout these six years, as well as read this document and provide feedback. (I am also grateful to the others who have looked at part or all of this thesis and provided feedback.) Rob Schoelkopf has been a fun and enthusiastic advisor. He is admirable in his commitment to a high quality of research, understanding, and fairness. His eagerness to share knowledge is clearly evident as soon as he steps in the lab; an anecdote, an informational tidbit, or an educational whiff of GE Varnish is certain to result. His position requires an understanding of how individual projects fit in with a grand scheme, and with one other. Rob fills this duty readily, and I am undoubtedly impressed when he extols his vision of the path forward long before I could have put two and two together. It is also thanks to Rob that certain pun-ny names for devices have stuck. Michel Devoret is a model scientist and a deeply caring human being. His passion extends beyond physics, and he will excitedly discuss almost any topic, for hours. He appreciates fundamentally beautiful physics, as well as well-thought-out technical details that bring an experiment to realization. Whenever there is a fork in the road, and I must ask, “Should I brute-force this?” or “Should I take a little more time to come up with an elegant solution?” the response is obvious, if one imagines that Michel is watching. I believe his passion and dedication comes from the heart, and it makes our local and scientific communities better places. Liang Jiang is, seriously, the nicest person. He has a way of making you feel okay about (comparatively) not knowing very much about anything. He’s quick to jump to the core idea in a presentation or colloquium. Working with Liang and his group (notably Mengzhen Zhang and Changling Zou) has been a breeze, given how easy it is to bridge the divide between theoretical and experimental un-

derstanding. I only regret that I did not take enough advantage of the valuable resources they provide. Steve Girvin, Leonid Glazman, and Peter Rakich have provided valuable feedback regarding manuscripts and presentations, and have been a pleasure to work with. Shyam has been the resounding voice of reason, from a calm and assertive “You probably shouldn’t be doing that” when questionable choices are being made, to brilliant physics advice. Hearing Luigi Frunzio’s infectious laughter is a sure-fire way to know you’re on the fourth floor of Becton. I’m grateful for his help in keeping our fabrication facilities running, providing advice for all occasions, and I won’t forget his pizza-as-a-plate or second-best-physicist-in-Italy stories, among others.

I am grateful to many former labmates and members of the physics community for their mentorship, friendship, and all-around good spirit. Matt Reagor, my first mentor in the lab, always took the time to teach me the things I didn’t know, but also gave me plenty of room to tinker. You were my model for a “senior graduate student”, and I am reminded of your positive attitude whenever I think about how to tackle a problem. We have Phil Reinhold and Reinier Heeres to thank for moving us out the Dark Ages of LabView. The software infrastructure and support they provided has been invaluable in enabling nearly every experiment, and most simulations, over the past four years on the floor. And somehow, they had time to do their own physics, too. Evan Zalys-Geller has told us what can be sensibly built, and what, even as physicists, is too ridiculous of an approximation to construct. Chen Wang is a brilliant, humble, versatile, and talented physicist. He always had the answers, and with a smile on his face, when we worked together; I am glad for the opportunity to learn from his way of thinking. I have appreciated Jacob Blumoff’s unique blend of enthusiasm, sass, and wit, as well as his unbeatable catering choices. His shrewd predictions were realized when Lake Jacob, his namesake, rapidly filled with Becton house water. Kevin Chou has been extremely generous in his time, whether it is spent teaching me about quantum measurements, or fixing every chiller in the lab—multiple times. I could usually count on one of these two to be in lab late at night so that they could flip a valve I had forgotten earlier. To my officemates over the years: Eric Jin, Teresa Brecht, Ankit Disa, Kyle Serniak, Katrina Sliwa, Nick Frattini: thank you for your companionship, discussions about physics, and for putting up with all the shenanigans. Thanks go to Teresa and Chan U Lei for taking good care of Lazarus. Thanks to Michael Hatridge and Anirudh Narla for teaching me the ways of a dilution refrigerator, and Nick Frattini and Chris Wang for allowing us to pass that knowledge

on. It's the kind of job that is easy to appreciate when things are going well, but is defined by how well you can recover from the failures... good luck! Nick Dent has been our savior when things turn really bad, really fast. It is a shame that we only spend time with Nick in the midst of a catastrophe; nevertheless, his visits are eagerly anticipated and a guarantee of fun conversation over trips to Blue State and Caseus. Luke Burkhart and Wolfgang Pfaff have been my partners in flying-cat-related crime. Wolfgang taught me how to think about an experiment and formulate a paper. (The secret, apparently, is waking up super early, going to a coffee shop to ponder physics, and then arriving in lab at the same time as everybody else.) He also has a knack for effective management, and I'm glad to have been part of his postdoc experiment. Luke has been, of course, immensely helpful in realizing these experiments. With these two everything becomes smooth sailing, more or less. Having watched Luke quickly take hold of his own projects, one doesn't need a quantum oracle to predict that he will be very successful.

The Yale facilities staff have not only made the research we do possible, but also enjoyable. Vinnie Bernardo of the former Gibbs machine shop is an exemplar of professionalism, and is ever-so-quick to send a message, "Your parts are ready!" just after warning a day or two prior that "this might not be possible with conventional tools." Mike Power has spent countless hours in the clean room with me to test or develop processes, and I am glad for his company and expertise. The Applied Physics and Yale Quantum Institute administrative staff have made life much easier for us all: Devon, Maria, Terri, Nuch, Racquel, and Florian.

Equal thanks must go to the labmates and friends who have kept me sane throughout these years. One arc of my graduate school career began late one night over Thai food, discussing whether a certain piece of software should be named Prism or Sauron. That was the start of a years-long war with Murphy's Law. Or was it quantum mechanics? By suddenly and aggressively monitoring every aspect of our measurement equipment, were we projecting the system into a 'broken' state? Regardless, we were presented with every kind of malfunction and catastrophe imaginable: the biblical plagues of Becton. Fires, floods, and [electrical] famine summoned us to lab, almost predictably between 2–6 AM on Saturday and Sunday mornings. (Rob has called these "growing pains". Since my responsibilities have been passed on, however, all has been quiet and well-behaved. Explain *that*.) As painful and degrading as these experiences sometimes were, they built a sense of camaraderie through a shared sense of responsibility for the lab. Though not necessarily in

this form, I hope that future generations will get to experience a similar feeling.

Beyond physics, my impression of New Haven will form around some fond memories: exploring the frozen wastes of New Haven after a single snowfall of 34" overnight; sledding down the Divinity School hill; flying myself and three other labmates to Baltimore for March Meeting and being able to justify the expense to Yale; trips to Piseco, Hammonasset, Acadia, Montauk, and others; being stuck in airports and airport hotels for many other March Meetings; hunting virtual Pokemon during the sudden revival of a childhood fascination; attending the D.C. March for Science; attending excellent holiday dinner parties with Michael, Cate, and Brady Hatridge; watching friends scatter throughout the world. During graduate school, friends have gotten puppies, married each other, started companies, and begun families. The world has spun quickly in six years.

To Mom and Papa, thank you for your support in these past six years, and the twenty before those. It means more than can be expressed here. Angela, you have made the travails of graduate school all the more worthwhile. I am excited for our adventures and careers to continue together.

It is commonly said the the fourth floor of Becton is a community like few others in academic research. We now extend that community across the street, to the Yale Quantum Institute. I am grateful to have played witness to the increased interaction it has brought: visitors, colloquia, lunches, and coffee breaks; and with them, exposure to more physics from some of the brightest minds. I am constantly reminded of our privilege to do this work, with the resources we have, and to be surrounded by such talent and enthusiasm.

1

Introduction

The notion of human interaction with a quantum mechanical world, much less our control over it, is a bizarre one. The first experiments to unveil the quantum mechanical nature of our world were performed only within the past century. With silicon fabrication techniques, laser optics, superconducting circuits, and Josephson junctions (among other technology developed in the past half-century), we have finally begun to exert control over the quantum world.

Solving the problem of “quantum control” opens up access to new areas of research, but also a host of tangible applications that can be pursued. Recognizing the overlap of fundamental physics with applied research, interdisciplinary programs naturally spawned during this time. The study of “applied physics” formed at the intersection of physics and engineering. In our field of circuit quantum electrodynamics, in particular, concepts from electrical engineering and computer science mix heavily with quantum mechanics.

These disciplines provide the framework in which to use quantum systems for information storage and processing, and ultimately to construct a quantum computer [1]. We choose quantum systems comprising artificial atoms, made from superconducting structures, to store quantum information. To preserve the quantum nature of such a system, our probing must be careful and indirect; strong interaction between a classical measurement apparatus and a quantum system does not easily permit many sequential operations [2].

This kind of interaction with the classical world is critical for learning information about the state of a quantum system. Such control also permits data input (system initialization) and control signals, which

guide the execution of algorithms [3]. An equally important aspect of our treatment of a quantum system, however, is “environmental control”: what happens when these external interactions are turned off.

The experimenter and the “environment” (external parts of the greater system that are off-limits to the experimenter) are both coupled to a quantum system. When the experimenter relinquishes control, environmental interaction persists. Naturally, since the information it learns is inaccessible, we hope that the environment interacts with our quantum system minimally. Successful control is possible only when the the rate of user-driven operation greatly exceeds the rate at which the environment imposes uncertain effects [4, 5].

However, realistic coupling to quantum systems introduces both channels, more or less proportionally. A circuit may be capacitively coupled to a transmission line, which carries control signals, but will also carry noise. If the coupling is designed to more rapidly control the system, the decoherence rate from the noise source may increase proportionally. Further, some sources of decoherence are independent of this coupling, and set hard limits on the timescales within which operations may be performed (before all information is lost). These elements combined require careful consideration of how quantum systems are constructed and interfaced with the classical world.

More powerful machines, quantum or otherwise, require more elements and more operations on and between these elements [6]. For practical devices, interacting with a single quantum system is insufficient. Even if control over a single quantum system can be mastered, interfacing multiple systems will present new challenges. Environmental interaction that may have influenced the one system may now influence multiple components; added control channels will further increase the likelihood of losses. If we cannot control one system at a meaningful rate, relative to environmental loss, then multiple systems will only scale unfavorably in difficulty.

This work will introduce ideas useful for taking control of quantum systems, back from the hands of Nature and into those of the experimenter. One can endow a quantum system with such protection on several scales. First, individual systems must be shielded from environmental noise and loss. Then, we must determine how to couple to these systems without compromising this built-in protection [7]. Invariably, direct coupling between trivial systems will not suffice. Finally, following the path to construct a functional machine [8, 9, 10], these systems must be coupled together. This system-to-system coupling

must also not compromise its performance.

The essence of maintaining control is to use intermediate systems, or “ancillas”, to mediate interactions [11]. In this way, some elements may be strongly protected against sources of decoherence, while ancillas handle fast operations like those exchanging information between multiple systems, or reading out their states. The overall system can thus be equipped with the best features of each component. Now, to reach the protected elements, noise introduced through couplings must first traverse the ancilla.

Our requirements for control over quantum systems have seemed, at first, to present a conundrum — the simultaneous need to have long-lived and quick-response elements. Through careful design of the elements comprising our quantum mechanical system, however, we show in the following work that control can be established to a very high degree. Through design choices, filtering, and natural isolation, we build a quantum system that satisfies these requirements. Further, these concepts are preserved when many elements are coupled. We will show that the control problem *can* be solved.

The divide between classical and quantum may appear large, even unsurmountable. Regardless, we live in a world that is fundamentally quantum mechanical. If we choose to believe the meaningfulness of this non-intuitive, wacky behavior, then Nature is willing to grant us a concession: for a time, however brief, we may wrest control of a part of the quantum world, and use it for our mortal purposes.

1.1 Overview of this thesis

This thesis discusses several aspects of the effort to build a scalable quantum computer using circuit quantum electrodynamics (cQED). They fit into a framework of scalable, interacting systems called “the modular architecture”. Such an architecture has the benefit of using well-isolated modules that can, in principle, be very high-quality. This work first characterizes a new architecture that can be used to construct a single, simple module. Then, it will discuss the principle by which one system can be interfaced with another. It will allude to related, ongoing studies that fill the gaps that remain, ultimately allowing more complex modules, and networks thereof, to be constructed.

Before presenting the results of these recent experiments, I begin Chapter 2 by describing, in greater detail, how these components mesh with our long-term vision for building a quantum computer using

general quantum computing platforms. I introduce the goals and tasks of engineered quantum systems. I discuss the forms in which information can be encoded, including some that protect it against system imperfections. Finally, I extend these concepts to the modular architecture introduced above.

In Chapter 3, I detail the layout of a cQED network based on these concepts. I introduce the physics necessary to implement these ideas in cQED. This includes the type of system elements, control hardware, and operations that must be implemented. I introduce the elements required for a simple module, and separate the elements that are available from those that have yet to be fully developed. We will need elements and techniques for constructing and interfacing the desired modules; I will examine several candidates.

By this point, it should become apparent that while elements of circuit QED have demonstrated many useful purposes, they have not yet been combined into a simple module with the purpose of assembling a modular architecture. I will introduce a *hardware* architecture in Chapter 4 that allows for straightforward integration of a module's necessary elements—a combination of integrable quasi-planar and 3D structures—and perform characterization to show that their behavior is exceptionally robust. This characterization will include the development of a suite of tools that I hope to be considered useful in constructing complex, integrated systems.

Not yet satisfied with the lifetimes possible in such an integrated module, I will investigate the sources of loss that limit them in Chapter 5. I will use some of these new quasi-planar elements, in addition to recently embraced cavity geometries, as testbeds to evaluate limiting dissipation mechanisms. This is accomplished through systematic variation of parameters and many, many, many experiments. To aid with this multitude of sample testing, I will use variants of the designs in Chapter 4 that permit multiplexing. At the end of this chapter, I summarize the limits that these experiments place on the quality of various materials in planar, quasi-planar, and 3D systems we might use.

With an improved understanding of the factors that limit us, we can design an assembled module accordingly. In Chapter 6, I conscript these integrated modules to demonstrate a basic task that will be critical in the modular architecture: quantum communication. In this case, an experiment is performed that converts stationary states (stored within a quantum memory) into propagating ones. This process can be performed very coherently, meaning that quantum information is preserved through this 'change of character'. This satisfies a goal established in our introduction of modular networks: the ability to

communicate between modules.

Such communication is not complete, however, without a module on the other end of such a communication channel ‘picking up the phone’, or catching the flying information smoothly and efficiently. This challenge is addressed in Chapter 7. Here, we perform an experiment that shapes the form of the flying information—in particular, its wavepacket—so as to enable this receipt to occur. We use this process to learn about the form of loss that likely dominates the transfer. Furthermore, we choose a form of information encoding and propose how it can be used to protect against this type of loss. Combined with the efforts in Chapter 5 to reduce intrinsic loss, this error correction protocol will further benefit the robustness of such a modular interface scheme.

I conclude by taking a step back and looking at this work in the context of the larger goal. Building a quantum computer using the modular architecture will require an extensive effort on many fronts. Much recent and ongoing work will be directly compatible with the experiments shown in this thesis. Their combination has the potential to produce exceptionally powerful quantum machines.

2

A Quantum Information Roadmap

In this chapter, I will introduce the concept of quantum information generally, and as it pertains to the long-term aim of this field of research: building a quantum computer. I will draw parallels to classical computing, both in the operation of a computer and in the progression of research. This progression must follow a path that meets certain broad milestones, regardless of the platform used to construct the system. Since this thesis will address several of steps along this path, I introduce them in more detail, although in general, platform-independent terms. Finally, I conduct a brief survey of the methodologies pursued in scaling quantum systems to much larger sizes. No quantum computing platform has yet to master these remaining challenges, and this is where the forefront of research lies.

2.1 Quantum computers

The concept of a quantum computer carries a certain ambiguity, in part imbued by its prevalence as a “buzz word” in popular media. Some even argue that quantum computers have already been built.¹ If this is the case, why do we continue to seek this fabled goal? First, a true quantum computer, which we will attempt to define below, is still perhaps decades off. But even if this were a reasonable characterization, we would still have much to learn. The “quantum machines” that are presently built in research labs around the world, striving towards the greater goal of a quantum computer, produce valuable scientific knowledge and

1. Those who claim to have built a quantum computer are using a different term than that used by the majority of the community. The only kind of “quantum computer” that has been built to date is more appropriately called an “adiabatic quantum annealing machine”, and does not meet the DiVincenzo criteria that are discussed below. Without individual control over quantum elements in the system, this machine instead probes ensemble behavior of a many-element quantum system.

technologies that can be used in many fields. They offer researchers understanding of natural phenomena in the quantum realm, not just a way to build a fancier quantum machine.

So what is a quantum computer, and what must it be able to do? The concept of a quantum computer was proposed and popularized by Deutsch, Manin, Feynman, Shor, Grover, and others, and initially couched as a natural way to simulate natural systems—which, fundamentally, are also governed by the laws of quantum mechanics [12, 13, 14, 15]. Like many massive undertakings in physics, a singular long-term purpose is difficult to state. Much like a classical computer, the quantum computer should be able to solve mathematical problems of broad interest. These may include such things as protein folding, genetic mapping, machine learning, or, naturally, simulating other quantum systems [16]. Of particular interest to security professionals is the prospect of encoding information in an unbreakable manner, by taking advantage of fundamental properties of quantum mechanical systems such as the no-cloning theorem [17]. This falls under the category of “quantum cryptography”, which has somewhat different needs than those of a quantum computer for, for example, simulation. Cryptography necessarily involves the transmission of information, focusing more heavily on principles of quantum *communication* than quantum *processing*. The ensemble of techniques and pursuits of the manipulation of information stored and transmitted to, from, and within quantum systems forms the basis of “quantum information”.

Quantum computers and communication systems must meet certain requirements—the DiVincenzo criteria—in order to fulfill at least some of the set of tasks laid out in Section 2.2.1 [18]. They are guidelines, largely generalisms found to be consistent with all the experimental platforms of the day. They state that a quantum computer must be scalable, contain identifiable quantum bits (information carriers), and that these quantum bits can be controlled to the extent that they can be “initialized” in a known state at some point. They must store this information for long-enough time compared to the operations that will be performed. It must be possible to implement a “universal” set of quantum gates, or operations on and between them, that can be composed to form any possible operation [19]. Finally, it must be possible to measure the resulting state of such quantum bits after the desired operations. These criteria are each satisfied to a certain extent in true quantum computing platforms. The more detailed steps that must be taken to satisfy them will be given in Section 2.2.

The DiVincenzo criteria further lay out two requirements for quantum communication. The com-

munication of quantum information between distant systems requires conversion between “stationary” quantum bits, and “flying” ones that carry the information between systems. Then, we must develop the ability to let the flying quantum bit “fly” without destroy the information it contains. We will present our realization of these criteria in Chapters 6 and 7.

A quantum computer fitting this definition does not yet exist, largely because of the first criterion, “scalability”, in combination with the ability to control and measure. Some groups claim to be nearing the point where quantum simulations of a (relatively small, though inaccessible to classical simulators) number of particles can be simulated [6, 20, 21]. Classical simulators, however, are also catching up, thus obscuring this cross-over threshold sometimes called the point of “quantum supremacy”. To understand these claims, we must delve into the details of how these various tasks will be accomplished: the “roadmap” to constructing a quantum computer.

2.2 The research roadmap

The research effort in quantum computing must follow a similar path to that of classical computing: advancements are cumulative. Development of new classical hardware and technologies continues, as evidenced (on a grand scale) by the paradigm shift to distributed “cloud” computing and storage services. Yet, at the same time, semiconductor foundries continue to improve the process for fabricating the most basic computing elements—transistors—reducing feature sizes from 10 μm to 10 nm between 1971 and 2017, increasing transistor density and computing power, and drawing out the trend characterized by the popular Moore’s Law [22]. Developments occur at all the levels in between, too, improving CPU and integrated circuit architectures, inter-element communication protocols, and the firmware and software that takes control of it all. Connectivity across the globe has become faster and higher-bandwidth, allowing users to take advantage of distributed computing protocols. This growth faces some constraints and challenges, and while it is not likely to become obsolete for many, many generations, there are applications that even scaled-up classical computers could not support. Quantum computers could thus feasibly complement classical ones, by filling these gaps.

Quantum computing also begins with its canonical “transistor”, the qubit. The qubit must be built

as robustly as possible, since it will carry the weight of all that follows. Next, control must be introduced by adding input and output connections. Control hardware allows operations using the qubit, and ultimately using multiple qubits. A measurement probe is also required to observe the result of these manipulations. Quantum “software” will control the qubits and execute the required algorithms. The greater the number of qubits, the more complex algorithms can be performed.

At some point, adding more qubits will produce diminishing returns. This is because the qubits themselves have imperfections, the results of which are amplified by the engineering that facilitates their interactions. The next step must then be to define a “logical qubit”, wherein information can be redundantly encoded, following guidelines set forth by classical information processing. In classical computing, “logical bits” take the form of an array of physical bits, encoding information in 0 or 1. A logical bit could, for example, redundantly encode 0 as 000 and 1 as 111. If a single error occurred on one of the physical bits, for example $000 \rightarrow 010$, the error could be caught by “majority voting”, noting that two of the physical bits remained 0 and the likelihood of a double or triple flip is $3p^2 - 2p^3$ (for the likelihood of a single flip p). For $p < 0.5$, this triple-redundancy produces a favorable encoding. (Typical per-bit ‘soft’ error rates are 100–1000 per billion hours.) Once identified, the error could then be corrected by flipping the offending bit back.

Similarly, with qubits, information can be redundantly encoded in an array thereof. (This progression is shown in Figure 2.1.) With the right choice of logical encoding, this array can form a logical qubit (Section 2.3). Error-correction procedures, such as analogues to the majority voting pertaining to classical bits, can be performed (Section 2.3.3). This can extend the lifetime of information stored in the logical qubit beyond that of any single physical qubit. More complex systems can then be built using logical qubits as the basic building block, sweeping the details of physical qubits beneath a veil of abstraction.

The same kinds of operations that are developed for physical qubits will need to be performed on and between logical qubits. This will extend to larger arrays of logical qubits, enabling more powerful algorithms. Eventually, even logical qubits will suffer from combined imperfections, whether based in the physical imperfections of their constituent elements, or from the operations performed on them. This will demand increasingly complex redundant encoding schemes, essentially wrapping the information in layer upon layer of protection. At some point, a desired algorithm can be performed without an error.




Figure 2.1 | **A research roadmap.** General guidelines that researchers must follow to build a quantum computer capable of meeting the criteria listed in Section 2.1. This diagram includes a suggestion of how a modular architecture, in particular, could be addressed; the logical qubits here are treated as separable modules. (Figure reproduced from [23]; see [Copyright Permissions](#).)

When the chance of a “quantum blue screen of death” (qBSOD) is as infrequent as a BSOD for classical computers, we will be satisfied. Somewhere around this level is called the “fault-tolerant” threshold for quantum computing, a more explicit definition of which will be given in Section 2.3.3.

This thesis tours several landmarks along the route suggested by this research roadmap. I investigate the mechanisms limiting the performance of qubits at the most basic level, discuss methods in which to encode information in a logical qubit that supports error correction, and propose guidelines for the construction of housings and communication pathways in scaled-up systems.

2.2.1 Algorithms

Certain tasks are well-suited to a quantum computer, and some considerably more so than for classical computers. They can be characterized by the relative “speedup” that a quantum computer would provide to them [24]. Further still, some of these tasks convey practical significance on a broader, more worldly scale. Ambitious persons prognosticate that quantum computing will find solutions to world hunger, medicine, climate change, and data security [25, 26].

Beginning at a relatively low level, however, these tasks involve such things as running search algorithms, factoring large prime numbers, computing quantum Fourier transforms, and performing quantum walks [27]. In schemes like quantum key distribution (QKD), some of these concepts of quantum

mechanics have already been applied to cryptography; factoring of small numbers has been demonstrated in relatively simple quantum systems [28, 29]. On a mesoscopic scale, quantum computers are increasingly the platform used in research into quantum machine learning and quantum simulations [30]. Quantum simulations could have broad-reaching effects, such as in quantum chemistry [31]. Using quantum computers to effectively solve for the energy of certain chemical system configurations, one can exactly describe molecules and materials that cannot be simulated classically. At face value, this approach—using quantum systems to simulate other quantum systems—seems like it could be a tenable pursuit.

Control over the energy landscape (the Hamiltonian) of superconducting circuits in cQED and other atomic systems makes them suitable to be conscripted for these purposes. The other important factor to many algorithms, naturally, is raw computing power. The simplest algorithms require a few quantum bits. Current estimates demand something like ~ 100 logical, error-corrected qubits in order to reach a point where classical computers can no longer compete [28, 32]. If one qubit is a challenge to create and to control, then surely many qubits will not be any easier. The challenge of scaling will be discussed later, in Section 2.4. The form and design of these quantum bits will be critical to this pursuit.

2.3 Encoding and protecting quantum information

Choosing a quantum system suitable for information storage and processing is only half of the battle. One must also weigh the options for information encoding within that system.

Information encodings can be characterized in many ways; one is whether they use a few discrete energy levels of a system, or whether they are spread across many in a continuous fashion. Two- or few-level systems are generally easy to control and address, but they are limited in complexity and can be more susceptible to errors than distributed, many-level schemes. Many-level or continuous-variable encodings often take advantage of systems like harmonic oscillators. While these can be physically simple and are often highly coherent, special tools must be developed to interface quantum states with systems that are typically treated classically.




Figure 2.2 | **The Bloch sphere.** A Bloch sphere representation of the quantum state of a two-level system. A pure state lives on the surface of the sphere; we define the eigenstates of σ_z , $|0\rangle$ and $|1\rangle$, to be at the poles.

2.3.1 Encoding in few-level systems

In the same way as bits are defined to be the fundamental information storage mechanism in so-called “classical computers”, qubits can be thought of as the fundamental element of a quantum computer.

Information encoded in the two-level quantum bit, the ‘qubit’, can provide some computational benefit compared to information encoded in a classical bit (when considering certain problems or algorithms). This is because information can be encoded in an arbitrary state that is the superposition between the two basis states, $|0\rangle$ and $|1\rangle$, such as $\alpha|0\rangle + \beta|1\rangle$ (where $|\alpha|^2 + |\beta|^2 = 1$). This encoding is often represented in the Bloch sphere (Figure 2.2), an unmistakable reminder that qubits are distinct from classical bits (which take discrete values confined to one of the poles of the sphere). While qubits are also projected to a particular value when, by measurement, they are interfaced with the classical world, their residence in a quantum world gives them freedom to behave differently at other times.

A qubit’s position on the Bloch sphere, given by the Bloch vector \mathbf{r} , relates to its quantum state ρ as

$$\rho = \frac{1}{2}(\mathbf{I} + \mathbf{r} \cdot \boldsymbol{\sigma}), \quad (2.1)$$

where σ represents the Pauli matrices, $\sigma_x, \sigma_y, \sigma_z$ given by[33]

$$\sigma_x = |1\rangle\langle 0| + |0\rangle\langle 1| \quad (2.2a)$$

$$\sigma_y = i|1\rangle\langle 0| - i|0\rangle\langle 1| \quad (2.2b)$$

$$\sigma_z = |0\rangle\langle 0| - |1\rangle\langle 1|. \quad (2.2c)$$

$$(2.2d)$$

This representation is convenient because it compactly describes any location on the Bloch sphere, and can also be used to represent operations and gates.

For example, the unitary that must be applied to ‘flip’ the qubit from $|0\rangle \leftrightarrow |1\rangle$ is given by

$$\hat{U}_x(\pi) = e^{-i\theta\hat{\sigma}_x/2}|_{\theta=\pi} = e^{-i\hat{\sigma}_x\pi/2}, \quad (2.3)$$

where $\hat{\sigma}_x$ is the operator form of the matrix σ_x . The Hadamard gate, which plays an important role in implementing a controlled-NOT gate, can be expressed as $\hat{H} = (\hat{\sigma}_x + \hat{\sigma}_z)/\sqrt{2}$.

Our choice of qubit for many purposes in this work will be the transmon, further detailed in Section 3.4.2. As an anharmonic oscillator, the lowest two energy levels of the transmon can be used as the ground and excited states of a qubit. Hundreds or thousands of gates can be performed on transmons before their energy is lost. They can be conveniently initialized and measured. In many ways, they satisfy all the necessary criteria for the basic element in a quantum computer.

One qubit cannot build a quantum computer alone. The feasibility of an encoding scheme can only truly be evaluated when considering the storage or processing of more than one bit of quantum information. In most schemes with transmons, this requires coupling more physical devices in order to scale to larger and larger arrays. The rate of information leakage will increase accordingly, mostly under the influence of environmental influences and cross-talk as discussed in Chapter 4. Effective scaling systems with this kind of information encoding will require specially developed error correction schemes (Section 2.3.3) along with very robust engineering and many, many devices.

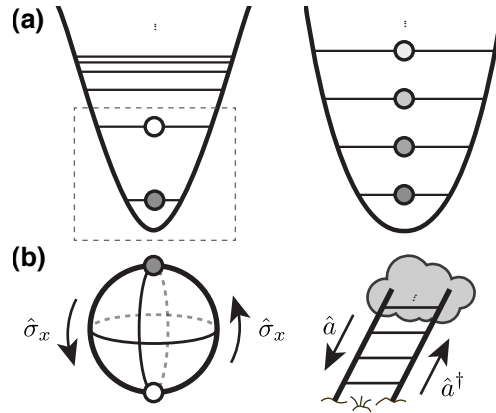


Figure 2.3 | **Two forms of information media.** **a** Energy levels of a two-level system (left, approximated as spanning the lowest two energy levels of an anharmonic oscillator) and a harmonic oscillator (right). **b** A qubit that is a two-level system can be represented within the Bloch sphere; the Hermitian operator $\hat{\sigma}_x$ connects the two states. States of the harmonic oscillator are accessed by applying raising or lowering operators (\hat{a}^\dagger , \hat{a}). Energy eigenstates are those of the operators $\hat{\sigma}_z$ (qubit) or $\hat{a}^\dagger \hat{a}$ (oscillator). In these common representations, operations are not directly comparable, and it is not immediately clear how a qubit can be encoded.

2.3.2 Encoding in many-level systems

Alternative, non-binary classical computation schemes for information encoding have been studied, but never caught on. Ternary bits can encode information in three levels, potentially increasing information density per device. Naturally, operations in ternary computing are more complicated, though in some cases more efficient.

Depending on the quantum system, encoding in higher levels can be more or less difficult. Natural atoms, for example, are often most easily addressed as true qubits or few-level ‘qudits’. Artificial atoms or anharmonic oscillators, however, can be more easily, flexibly addressed in a variety of ways. Storage of information within continuous variable (CV) states of a harmonic oscillator [34] is one popular way of encoding information in cavities. This choice of encoding offers natural, physics-based benefits, but also comes from practical considerations: long-lived, macroscopic harmonic oscillators are often easier to make than many atomic systems. They potentially fill the role of multiple two-level systems in a “hardware-efficient” fashion [35, 36].

Quantum CVs allow versatile and robust encoding of quantum information in high-dimensional

Hilbert spaces. For instance, encoding quantum bits in CV systems can provide the redundancy required to enable quantum error correction (see Section 2.3.3). Non-Gaussian CV states that could be used as quantum-information-processing-enabling resources have been created experimentally in the states of ion motion [37] and atomic spins [38, 39], as well as optical [40, 41] and microwave photons [42, 43, 44, 45]. In particular, microwave cavities in superconducting circuits have recently further enabled the storage [46] and protection [47] of quantum information encoded in non-Gaussian oscillator states.

While many systems are naturally quantum harmonic oscillators, a classically-addressed oscillator does not satisfy the criterion of being a distinct qubit. This is because the spacing between energy levels in a harmonic oscillator is equal; trying to address one energy level will enable excitation movement between them all. To produce distinctly “quantum” states within oscillators, they require a dose of anharmonicity [48]. This is often provided by nonlinear media or elements, which aid in the preparation of quantum states. Without such an element, only classical (coherent) states could be prepared in the cavity.

The chosen form of the encoding is flexible and will depend on the purpose of the device. It can be convenient to encode information in the photon number (Fock) basis of the cavity, in which simple states take the form of, for example, “having one photon” or “having no photons”. More complex encodings, such as in the “even parity” or “odd parity” of a coherent state superposition, can be designated for purposes like error correction. CV encodings are highly compatible with our choice of storage element, the 3D cavity, introduced in Chapter 4.

Result of measurement The way in which information is encoded will determine the kind of measurement that must be performed to access it. Among these and other considerations, the measurement tools available to a particular platform will guide the choice encoding.

Superconducting circuits have access to a number of tools that permit either form of encoding (Section 3.5). The measurement of information stored in qubits can be easily obtained using an appropriately coupled resonator, connected to a standard microwave amplification and measurement chain. Extraction of cavity information typically requires an ancilla qubit and readout resonator; tools like parity measurements and cavity tomography are commonplace [49]. Other platforms may find, depending on the avail-

able tools, that one system or another is best.²

2.3.3 Error correction

Error correction is a concept defined in classical computing and was introduced, in that context, earlier in Section 2.2. Effectively, it requires some redundant encoding of information that, as a whole, will nominally remain impervious to the dominant sources of information loss throughout the computation. Since no system will be perfect, eliminating errors completely is unfeasible. Rather, we strive for a threshold called “fault tolerance,” past which error rates are made tolerable, and the corrective procedure does minimal harm of its own. More explicitly, the fault-tolerant threshold p_{th} depends on a given system and correction procedure. It is said to have been crossed if the combined error rate of many systems can be improved when concatenated.

Error correction, if not autonomous, must begin with detection of the error. Since direct detection can be destructive in nature, auxiliary “ancilla” systems are often used to probe an error syndrome. A choice of encoding depends on the targeted form(s) of dissipation. Typical sources of decoherence are amplitude damping (energy relaxation) and phase damping (dephasing). Since these are orthogonal effects, any qubit error can be broken up into these components. The repetition code, suitable for classical bits, cannot be directly imported for use in qubits because known quantum states cannot be cloned [50]. Using qubits to redundantly encode information, then, such as the arbitrary state $|\psi_0\rangle = \alpha|0\rangle + \beta|1\rangle$, requires a modification to the repetition code.

What we *can* make, following the rules of quantum mechanics, is an entangled state such as $|\psi_3\rangle = \alpha|0\rangle_A|0\rangle_B|0\rangle_C + \beta|1\rangle_A|1\rangle_B|1\rangle_C$. This may seem like a subtle distinction, and appears at first glance as if the state $|\psi_0\rangle$ has, in fact, been cloned. But a cloned state over three physical qubits A , B , and C would appear as

$$(\alpha|0\rangle_A + \beta|1\rangle_A) \otimes (\alpha|0\rangle_B + \beta|1\rangle_B) \otimes (\alpha|0\rangle_C + \beta|1\rangle_C), \quad (2.4)$$

2. For example, hybrid quantum systems often mix mechanical, optical, and microwave resonators, often with purpose of building conversion elements. These systems make it clear, however, that some elements are much easier to interact with than others. If one element is ideally suited for storing a quantum state, but it cannot be directly, easily read out, then storage of information there will be contingent on the existence of a suitable converter.

a fundamentally different state. If one of the physical qubits in $|\psi_3\rangle$ undergoes a bit-flip error, it has the effect $|0\rangle \leftrightarrow |1\rangle$. By measuring each of the resulting qubit states, and assuming no more than one qubit underwent a bit flip, one can identify the offending qubit and back out the original state. The matter is complicated by the existence of many more types of errors, a desire to non-perturbatively measure the system to check for errors, and errors propagated by this detection process itself. Codes of varying complexity exist to address these challenges, like the Shor code and the Steane code, which can correct for both bit-flip and phase-flip errors [5, 51]. Even the simple bit-flip repetition code presented here spotlights the inherent power of entanglement as a resource in quantum circuits.

In cavity CV encodings, the approach to error correction and chosen logical encodings can be quite different. Instead of multiple physical qubits used to increase redundancy, multiple modes of a single physical cavity can be used [34]. The discussion of loss in single qubits made use of simple operators, tied to particular loss syndromes; bosonic modes can have an unlimited set of errors: one-photon loss, two-photon loss, etc. In principle, corrective codes could be constructed to target one or many of the errors thought to be dominant [52]. Seemingly odd constructions may result; for example, bosonic logical encodings can take the form of states not confined to a finite Hilbert space, like superpositions of coherent states in the “cat code” [53]. Some logical codewords are not even strictly orthogonal, such as the cat code in the limit of small photon number. In codes such as the binomial code, encodings can be simply defined in terms of Fock state superpositions; in ones like the GKP code, they are distributed among phononic quadratures [54]. Despite their differences, CV encodings still function to produce a logical qubit, with basis states designated as $|0\rangle_L$ and $|1\rangle_L$.

2.4 Scaling up computing power

To satisfy the voracious appetite of quantum information algorithm designers, we must scale up our systems to control and interact hundreds of logical qubits, if not more. Further, from our discussion of error correction, it is clear that logical qubits must be composed of many physical ones. In this section, we will detail some of the approaches to achieve this using superconducting circuits.

One challenge is that scaling schemes, intended to protect imperfect qubits or other elements, tend to

add more imperfect elements. In order for this approach to succeed, the benefit of having more elements must outweigh the additional overhead, a threshold that is sometimes called the “break-even” point. So far, we have only made a comparison based on “quantum hardware”. But particularly resource-intensive applications of scaling will also levy demands on classical hardware and computational power [6, 55].

Scaling represents a complex form of a simple, canonical problem in the control of quantum systems. Quantum scientists and engineers seek to couple many more elements together, enabling new functionality and to ability to perform more powerful algorithms. Each new coupling introduces more noise to that element—whether via cross-talk, external coupling, or resistive dissipation from more material, more layers, or more complex processing. This often counteracts the intended improvement in function.

Many elements thus conspire to make scaling a difficult matter. But many scaling approaches are also underway to tackle this challenge.

2.4.1 A brief survey of approaches

Quantum systems have been sufficiently developed to the point where tens of qubits, each under individual control, can be interfaced and manipulated. In an effort to scale to even larger numbers, error-correcting codes and architectures have been developed.

The surface code [56, 57], a planar extension of the toric code [58], uses ancilla qubits and a series of stabilizer measurements to form a logical qubit that is protected against certain types of errors. It constitutes another style of correction code that is considered feasible for large-scale scaling approaches. These codes can be used in conjunction with repeated superconducting circuit arrays. Many physical qubits can form a protected plaquette, and thus, one logical qubit. The exact details of many of the steps noted in Section 2.2, such as error correction and logical gates, are still works in progress. Similar schemes have been proposed for arrays with trapped ions [59].

Topologically protected qubits constitute another promising vision for scaling [60, 61, 62]. A topological quantum computer is based on non-abelian anyons that can be “braided” and remain insensitive to perturbations over a certain scale. The challenge for this technology right now comes from low-level tasks, like creating a qubit that can be measured and controlled. Once a qubit is built, then the natural topological protection afforded to it can make the task of scaling much less onerous.

Superconducting circuits, as a platform for scaling, certainly benefit from the experience of the semiconductor industry. Technologies for wafer-scale processing, patterning, and even interconnection can potentially be harnessed for quantum computing. Once qubit and algorithm designs have reached a certain level, a path to scaling circuits can be more readily hewn. Compared to technologies like trapped ions, neutral atoms, and spin qubits, methods to achieve the required higher-level interactions—like assembling and interacting logical qubits—are far more clear. In these platforms, the approach to full-blown scaling remains more vaguely defined, since details are harder to apply without having chosen the technology with which to optimally scale.

2.4.2 A modular scaling architecture

One way to manage the complexity of scaling a quantum computation platform is by employing a modular architecture [8, 10, 63]. This architecture, which seeks to separate computing modules into nodes connected reconfigurably within a network, will be discussed more extensively in Chapter 3. In the context of error correction, and with respect to alternative approaches, however, it is worth briefly highlighting this scheme. Like the error-protected logical qubit of the surface code, a modular node would ideally consist of layers of data protection. Nodes could be identical, and connected in arrays, or could serve unique and specialized purposes.

The modular architecture is effectively a quantum network, but by another name and with a few distinctions in purpose. This means that the communication criteria posed in Section 2.2 will be of greater importance than for more ‘traditional’ schemes. Of the many viable pathways to our ultimate goal—a quantum machine with vast computational power—we will choose to pursue this option in the following chapters. They will introduce the components of a modular quantum network in greater detail, including hardware that will serve as a model for how systems can be scaled, and the techniques that can enable communication between nodes of a modular array.

3

Circuit QED for Modular Quantum

Networks

Circuit quantum electrodynamics (cQED) is a relatively modern field of research stemming from a storied parent, cavity quantum electrodynamics (CQED). It makes use of circuits, clearly, and in our case those circuits are constructed from superconducting materials that reduce dissipation as much as possible. The ‘quantum’ in cQED is achieved by freezing out most uncontrolled degrees of freedom so as to produce a system that can be placed into its lowest-energy, quantum ‘ground state’. The remaining degrees manifest as quantized charge or flux in the circuit.

At its heart, though, cavity and circuit QED both deal with the coupling of confined light, quantized in photons, and its interaction with particles of matter, like atoms. The electrodynamics of cQED are, in many ways, very similar to those of CQED; however, separate technologies are used to interface photons that generally exist on differing energy scales between the two. Another way in which cQED can distinguish itself is in the degree and ease of coupling that can be achieved. Electric and magnetic fields (light) can be made to couple very strongly to circuit elements (matter) [64, 65], and the addition of 3D cavities does not hinder that fact [66, 67]. The goal on the horizon—the quantum computer—requires systems that can be decisively and powerfully controlled. Circuit QED offers facilities that satisfy this requirement.

In this chapter, I will briefly discuss one present vision for how a quantum computer can be constructed. Then, I will detail the tools available to us in circuit QED that can be used to make strides towards that goal.

Superconductivity As mentioned above, circuit QED makes use of superconducting circuits, in which the primary charge carriers are not electrons, but Cooper pairs [68]. (Cooper pairs lend their name to the grandfather of our superstar qubit, the transmon.) The primary superconductor of choice is aluminum, which is relatively inexpensive and easy to work with. Bulk aluminum begins to superconduct below ~ 1.2 K (colder than the temperature of outer space), which leads us to cool our circuits down in a closed-cycle dilution refrigerator for them to become superconducting. Similar devices have been made from alternative superconductors, such as niobium [69], niobium titanium nitride [70], and titanium nitride [71]. Each presents benefits and drawbacks, in terms of dissipation and convenience. Throughout this work, with few exceptions, we will choose to use aluminum in our circuits.

3.1 A computation framework inspired by networks

The modular approach, first introduced briefly in Section 2.4, is based on a decomposable network of computational nodes. It serves as a powerful way to tame complexity while scaling up a quantum system. We use the term ‘network’ in the sense of a ‘local area network’ (LAN), to describe a kind of structured connectivity between computational qubits. Separately, ‘networks’ are used in the context of quantum computation to describe the transmission of quantum information for such purposes as quantum key distribution and quantum cryptography [17, 72].

In this sense, the interconnection of modules that send information back and forth would imply a natural analogue to the Internet. Many of the ideas, and some technologies, are indeed transferable to both kinds of quantum network. The purpose of the Internet is the transfer of information between nodes; often, additional processing occurs at one node or another. For example, say that a client wants to retrieve a piece of information. They submit their request, perhaps a search query or phrase. The request is transmitted across the Internet and reaches a server node. The server processes the request and returns a message.

The path may be more complex, however. The message, addressed to a particular server, will pass through nodes that route the signal. Classical signals will be attenuated over long distances; repeaters will boost their strength. Sometimes, the information will be lost or garbled entirely. Classical error correction

can help to protect against this. Finally, the bandwidth of the transoceanic cables that carry much of the world's high-speed communication is limited. Such communication channels must be shared; information must be encoded densely and cleverly so as to use available resources efficiently.

A quantum Internet or a quantum LAN will face many of these challenges, and others that the Internet does not face at all. But quantum networks, using resources such as entanglement, also have their benefits. If an entangled pair of qubits can be distributed across the network, interesting quantum operations can be performed. Because such entanglement exemplifies a nonlocal correlation, operations on one of the pair of qubits can translate to an operation on the opposite pair (effectively, teleportation) without having to send quantum states along the original path (which might have detrimental effects). This trick can be very powerful, but not so powerful as to violate the universal speed limit—the speed of light. This is because classical communication is still required to interpret the result.

One example of such a transaction is quantum key distribution (QKD). In one such scheme involving entanglement [73], a pair of entangled qubits is generated, split, and transmitted to two parties, Alice and Bob. For each incoming qubit, Alice and Bob measure in a basis from a set of selected bases that differ between the two of them. Once the measurements are complete, the parties can share their choice of bases, and further share the results from the measurements where the bases were different. This group of results detects eavesdropping; the remainder can form the cryptographic secret key. QKD based on similar protocols has been implemented in practice [74, 75].

Entanglement between nonlocal states is also the basis of quantum teleportation [76]. In this application, a pair of entangled qubits (A and B) is again split and transmitted to the two communicating parties, Alice and Bob, establishing a quantum link. Alice wishes to transfer the information contained in a remaining qubit (C). By measuring her system of entangled-and-information qubits (A and C) in the basis of a maximally-entangled state, the Bell basis, she entangles the information qubit with the entangled pair. By sharing the result of this measurement with Bob via a classical channel, Bob will learn the form of the entanglement of the information qubit with his half of the entangled pair (B and C). This amounts to “teleporting” the information contained in C.

Both of these processes rely on a quantum network to distribute quantum information, and a classical network (like the Internet) to support it. For long-distance communication, quantum repeater networks

will be required to support this transmission [77]. Though we have not discussed their exact form, these protocols typically make use of optical networks, in practice, which are less susceptible to loss and better suited for communication over long distances. (Recall that fiber optic networks are the backbone of the telecommunications industry.)

Microwave networks for computation (quantum LANs) will differ from these quantum communication networks (the quantum Internet) in a few ways. Because they are lossier, they will be constrained to a smaller scale than optical networks. Still, computation networks will borrow ideas from quantum repeater schemes, key distribution, and so forth. But in the end, a modular computation network is intended to achieve the computational goals laid out in the DiVincenzo criteria. The essence of such a scheme is the ability to interface quantum states stored and processed in network nodes with propagating states that connect the nodes (Figure 3.1). It should be able to implement the same universal gate set and perform the same algorithms as other proposed architectures. Its very nature is to address the issue of “scalability”; if certain elements exist and criteria are met, then there exist no hard limits to the number of modules that can be integrated. Its correspondence to a long-range quantum communication network is only partly coincidental, and should not be surprising.

In recent years, quantum networks are being recognized as more and more suitable for quantum information processing (QIP) [78, 79]. It has been shown theoretically that there are favorable thresholds for quantum error correction in such modular architectures, even with noisy quantum communication channels [63]. Experiments with multiple platforms are currently underway to realize prototypes of quantum networks that may be used for either purpose [10, 80, 81]. By examining the field of quantum communication, we can learn from existing conversion processes and borrow technologies and algorithms that will be just as important for modular quantum computation networks.

3.2 Distributed computing techniques and advantages

The distribution of modules across a network for computing provides many advantages compared to direct-coupling approaches [8, 10, 63]. In this section, we will explore some of these practical benefits, as well as the algorithms that may be well-adapted to such a scheme.




Figure 3.1 | **Connected quantum modules.** A vision of [a subset of] the “quantum Internet” is supposed, wherein interconnected modules constitute representative cavity QED devices.

3.2.1 Independence of modules

Direct-coupling schemes require that many elements are fabricated or assembled simultaneously. If one element does not have the desired parameters or functionality, then all or a portion of the computational fabric cannot be used. Sometimes, this non-functioning element can be ignored, and the rest of the elements utilized. In other cases, the entire fabric must be discarded and recreated.

Modules, like fuses in a circuit, can be easily replaced if one does not function as desired, without having to destructively access other parts of the network. Further, their independence means that each can be constructed separately from the network and assigned as nodes of the network on an ad-hoc basis. Breaking up the whole into small, testable modules that are connected through well-defined communication channels reduces undesired crosstalk and minimizes the spreading of errors through the system. For example, if a system element has failed catastrophically, it may broadcast harmful noise to its neighbors. In a modular architecture, this module could simply be isolated and ignored by disabling its communication links, and replaced when feasible. In the meantime, information could follow alternate routings—detours—between nodes of the network.

Module independence means that hardware “patches” can be rolled out if necessary, perhaps improving the functionality or even reducing physical footprint (an important, practical requirement for scaling). In this “prototyping” kind of stage of a quantum computer, we work with modularity each and every day.

(Simply put, our refrigerators are not hard-wired to permit exactly one kind of experiment.) We cannot be certain that one type of device will function optimally; with the flexibility that modules provide, we can iteratively and efficiently optimize a modular design without slowing the course of science.

3.2.2 Entanglement distribution

Section 3.1 described the relation of communication networks to the computational network we are trying to build here. Entanglement is a precious resource in any quantum computer, and even more so in a modular, network-based one where the dominant loss will likely reside in the links *between* modules. If entanglement can be generated and distributed among nodes the network, this loss can possibly be circumvented. Entangled states can be more fragile than a single photon; thus, the challenge will arise from how to distribute the entanglement without destroying it.

Two methods are commonly used to generate entanglement between distant elements in a quantum system like ours. The first, a nominally probabilistic method¹, generates two qubits and entangles them using a beamsplitter, and then detects them on one output arm or another [83, 84, 85]. The second, which will be discussed in Chapter 6 and 7, combines generation and transmission into one step [86, 87, 88]. Partial conversion from a nodal stationary state to a flying one can generate entanglement between the two; the latter is “caught” by another node to become stationary once more, thus entangling the two nodes.

Just as in classical computers, certain schemes are well-adapted to parallel processing techniques, or simultaneous use of multiple modules. In one recent example [89], a method for performing gates between distant modules was demonstrated. This achievement satisfies a critical requirement for scalable quantum computation (Section 2.1), the ability to implement a universal gate set, in the context of gates between modules. The first requirement in this scheme is generation of an entangled state between two qubits. In ref. [89], entanglement was generated and managed locally; if robust state transfer were developed between nodes, this procedure could have been demonstrated between two truly remote modules. This comes in the midst of a convergence of contemporary efforts in similar tasks and goals [87, 88, 90], which provide a prime example of the readiness of a modular architecture to accomplish the tasks formerly considered only

1. Probabilistic generation methods can be made to immediately try again until successful entanglement is heralded [82]. While still distinct from some fundamentally deterministic schemes, this kind of heralding can have an indistinguishable end result with high enough fidelity and generation rates.

in directly coupled systems.

3.3 Construction of a module

In order to achieve the alluring goal of modular computation, we must identify the necessary ingredients and figure out how to put them together. At first glance, forming a modular “network” in circuit QED sounds more formidable than in optics, a field that has long since succeeded in connecting remote systems at great distances [91]. Undaunted, I will introduce in this section a primitive module that fits into our concept of modular computing and networking. I will also suggest how this module might become more complex in time. In the following section, the elements that fulfill the construction blueprint for a module will be identified and introduced.

3.3.1 Required elements

Modules that can perform the tasks set out by the DiVincenzo criteria must take one of several particular forms. One way in which physical elements can be scaled and molded into a logical fabric is by directly connecting them (Figure 3.2a). In this approach, physical qubits are directly connected or in close physical proximity to one another (within the extent of a single package). Resonators and qubits generally constitute the elements within a directly coupled device. The lack of isolation between adjacent elements leaves them susceptible to noise and cross-talk. (Some schemes, like those using flux-tunable coupling elements [92], do provide higher amounts of isolation.) Logic constructed in this way does not enjoy the benefits of modularity, including node independence, easy element replacement, and network extensibility. The physical elements that form a logical qubit are treated with equal significance.

A basic modular node, however, begins by considering two very different kinds of elements: those dedicated to data storage, and those tasked with interfacing with, or processing, that data (Figure 3.2b). This distinction is based on the elements’ available coherence times, and the recognition that certain elements perform a specialized function better than others. It manifests in their degree of coupling to the environment, and may be built on radically different hardware. (In our implementation of such a module in Chapter 7, as an example, the disparity in coupling rates between these two elements is roughly three orders of magnitude.) The data storage, or “quantum memory” element, is shown in Figure 3.2b as

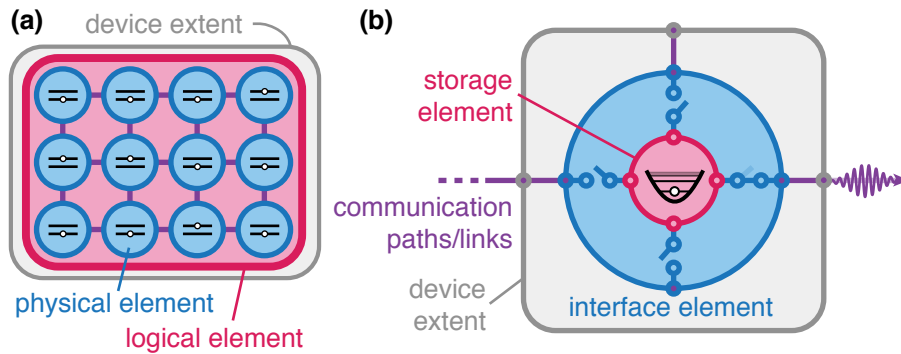


Figure 3.2 | **Configurations of scalable devices.** **a** An alternative means to scaling that does *not* take a modular approach. Logical elements (red) are composed of many identical physical elements, represented here as qubits (blue), connected within a single physical device (gray). **b** A basic, connected quantum module for information storage must contain a few critical elements. The data storage element is shown as a red circle; it may be, for example, a cavity mode. A layer that facilitates interactions with the environment is a concentric blue circle; its ability to manage communication is indicated by switches along the possible pathways. The physical extent (gray box) of a single module is shown. Communication links (purple) connect separate modular nodes within a network.

nested within the communication element. This indicates its increased isolation from the environment (the space beyond the physical device). The communication element facilitates use of the communication link, coupling separate physical devices. To preserve the independence of the modules, it is critical that communication can be enacted on demand. If the link is effectively ‘connected’ at all times, the isolation is spoiled.

While some small amount of loss can be tolerated through error correction, the communication element must generally facilitate coherent transmission of information from within the quantum memory out through the communication path (and the reverse process). It best serves the purpose of a module, in fact, if *all* elements are chosen to be maximally coherent and compatible with error correction schemes, and scalable within a module. An elementary module without redundancy or scaling can demonstrate key aspects of a modular network, but must eventually be made more complex.

3.3.2 Additional layers

Scalability must be addressed on multiple levels. A basic module that is lossless will still suffer from loss in the communication path. A lossy module connected within a network of lossless links will also suffer. Our real-world modules will suffer both afflictions.

Therefore, more developed modules will have many layers in which to redundantly encode information; the “data storage element” of the simple model will take a more complex form. These devices may feature concatenated layers of error correction. (A single layer of error correction of a memory cavity has already been demonstrated [47].) As one looks outward from the lowest level of the hierarchy, the protection of data relative to the communication link grows. At the same time, these outer layers will be more and more directly coupled to the communication element(s).

The complexity of the communication structure may increase as well. Any-to-any connectivity can be made possible within such a network, given enough routing elements, and the communication layer can support this. Control and processing layers, to perform algorithmic tasks not strictly related to error correction or communication, may also appear more numerous. Beginning in Chapter 6, we will construct the most elementary form of a module by following the requirements of this section. We will use the circuit QED components outlined in the following section, and choose them to permit future growth of complexity.

3.4 State-of-the-art circuit QED components

In this section, I will introduce the principal elements that can be used in, and in association with, a circuit QED module. They include 3D cavities, Josephson-junction-based qubits like the transmon, ancillary resonators, and elements that can be used for parametric conversion and amplification at relevant energy scales. In Chapter 4, I will introduce a few new tools to bolster this arsenal.

3.4.1 Quantum memory cavities

An element for quantum information storage must be long-lived, and also easy to incorporate into more complex designs. The important figure of merit is not simply the lifetime, but this time relative to operational (gate and algorithm) times. Furthermore, the energy lifetime of a microwave resonator can be highly power-dependent, probably due to surface and material imperfections (Section 3.4.3). Thus, the “single-photon” decay rate is the most relevant metric, since quantum circuits will generally operate at energy scales from less than one photon up to several photons. Circuit QED boasts exceptional information storage lifetimes, particularly when 3D cavities are used [67].

Since the advent of 3D cQED, several microwave cavity designs have been used extensively. The first of these 3D cavities, the rectangular cavity [66], propelled lifetimes from hundreds of nanoseconds to tens or hundreds of microseconds (for reasons that will be analyzed in Chapters 4 and 5). Even higher- Q cavities followed, with the cylindrical cavity capable of single-photon lifetimes around 10 ms [93]. The cylindrical cavity, however, required significant engineering to be integrated with a transmon qubit. The qubit frequency would have to be chosen to avoid other modes in the crowded cavity spectrum, and to be sufficiently far from these modes to permit wide-bandwidth control pulses. The large cavity volume and central field concentration of the low-loss TE_{011} mode means that coupling a transmon would be highly perturbative (see [67], Section 4.5.4). As a compromise, the coaxial stub cavity then emerged as the state of the art [46]. The stub cavity still offered many-millisecond lifetimes, but also easy and flexible coupling to other elements, including transmon qubits.

Cavities have become so long-lived (at the single-photon powers suitable for quantum information processing) that they have flipped the paradigm for quantum information storage. Before transmons were coupled to 3D cavities, their lifetimes were comparable. The transmon was designated as the qubit, and the cavity was used for readout. With the latest generation of 3D cavities, however, cavity lifetime began to surpass that of transmons. This motivated development of bosonic encodings and measurement techniques to take advantage of cavity modes as memory elements (descriptions of which will follow in this section). The coupling between transmon and cavity modes allows hybridization and inherited loss from one to another (the so-called “Purcell effect”; see Section 4.3.1). Whereas earlier generation cavities typically imposed this loss on transmons, the direction has reversed. The loss from transmons now sometimes limits 3D cavities to which they are strongly coupled. Still, the transmon qubit plays an important role in systems where cavities dominate information lifetimes.

3.4.2 Transmon qubits

As discussed in Section 2.3.2, a system of linearly coupled harmonic oscillators is insufficient to perform complex quantum tasks, such as creating and manipulating interesting, non-Gaussian states. Some source of nonlinearity is required to interface with our quantum memory. An “original” qubit, a two-level quantum system, can introduce this nonlinearity, and allow convenient control and measurement of a cavity

mode. Such a nonlinear medium can also facilitate many-photon conversion processes and execute certain operations more efficiently than with oscillator modes.

A natural choice in circuit QED is a Josephson-junction-based qubit, like the transmon. The transmon can be understood by quantizing a parallel LC circuit, a problem which has been extensively treated [94, 95, 96]. This circuit is extended to include the Josephson junction, which acts so as to add a nonlinear inductance (Figure 3.3a).

The transmon qubit is a type of charge qubit operated in a regime particularly insensitive to charge noise. In this work, our transmons will be of the “fixed-frequency” variety; in other designs, the frequency can be tuned by replacing the junction with two junctions enclosing a flux-capturing loop [97, 98]. The transmon evolved from the Cooper-pair box, in which the number of Cooper pairs crossing the junction is quantized [99], and has a Hamiltonian of the form [100]

$$\hat{H} = 4E_C(\hat{n} - n_g)^2 - E_J \cos(\hat{\varphi}). \quad (3.1)$$

The circuit representation of the transmon qubit in Figure 3.3a relates the circuit elements to the Hamiltonian parameters E_C and E_J . The charging energy $E_C = e^2/(2C_\Sigma)$ is associated with the transfer of an electron between capacitor pads (though, in reality, only Cooper pairs will cross). The junction (inductor) energy is given by $E_J = (\hbar/(2e))^2/L_J$. The equation-of-motion coordinate for the circuit is chosen to be the flux, Φ , which is elevated to a quantum operator by quantizing it such that its commutator with the conjugate variable, capacitor charge Q , is given by $[\hat{\Phi}, \hat{Q}] = -i\hbar$. It is related to the phase across the Josephson junction as $\hat{\varphi} = 2\pi\hat{\Phi}/\Phi_0$. In transmons, the ratio of E_J to E_C is large, on the order of 50–100. This places transmons in a regime where φ does not deviate significantly from zero. In this regime, the offset charge $n_g = -C_{\text{pads}}V/(2e)$ can be removed from the equation (at the expense of a slight shift in the energy spectrum) by applying a unitary gauge transformation [94]. Finally, the flux operator can be expressed in terms of the second-quantization lowering operator \hat{t} as $\hat{\Phi} = \Phi_{\text{ZPF}}(\hat{t} + \hat{t}^\dagger)$, where $\Phi_{\text{ZPF}} = \sqrt{\frac{\hbar Z_c}{2}}$ for the characteristic oscillator impedance Z_c . This means that the Hamiltonian of




Figure 3.3 | **The transmon qubit.** **a** A circuit representation of the transmon, where the boxed cross represents the Josephson junction. In the transmon design, the junction is shunted with a large capacitance and has $E_J/E_C \sim 50$. It is effectively equivalent to a nonlinear inductor (with some small junction capacitance). **b** The energy levels of the transmon are sketched (not to scale). The shapes of each wavefunction's squared-probability amplitude are drawn in red. When used as a qubit, the transition between the transmon's lowest levels, $|g\rangle$ and $|e\rangle$, can be uniquely addressed. **c** Optical micrographs of a fabricated transmon (center), dominated by the capacitor pads. The exact shape of the junction differs (inset) depending on fabrication technique. Scanning electron micrographs are of the bridge-free technique (left) [103] or the Dolan bridge technique (right) [104].

the transmon can now take the form

$$H = \hbar\omega_t \hat{t}^\dagger \hat{t} - E_J \left(\cos(\hat{\varphi}) + \frac{\hat{\varphi}^2}{2} \right), \quad (3.2)$$

where the lower-order terms have been moved out of the cosine to form the harmonic energy term.² One can analytically solve for the eigenstates and eigenenergies using the original Hamiltonian (Figure 3.3b) [101], or expand the cosine to the desired order, truncate the basis dimension, and diagonalize this Hamiltonian to obtain mode frequencies [102]. This latter evaluation is particularly useful when we include terms corresponding with a coupled cavity, for example.

The Hamiltonian of a transmon coupled to a linear resonant mode \hat{r} in the so-called dispersive regime,

² From the transformation $\hat{n} = -i\partial/\partial\hat{\varphi}$, the E_C term in Equation 3.1 and the second-order term in φ when the cosine is Taylor-expanded combine to take the form of an oscillator Hamiltonian with $\hbar\omega = \sqrt{8E_J E_C}$. Offset energy is ignored.

taken to the fourth order in $\hat{\varphi}$, is given as

$$H/\hbar = \omega'_r \hat{r}^\dagger \hat{r} + \omega'_t \hat{t}^\dagger \hat{t} - \chi_{rt} \hat{r}^\dagger \hat{r} \hat{t}^\dagger \hat{t} - \frac{K_r}{2} \hat{r}^\dagger \hat{r} \hat{r}^\dagger \hat{r} - \frac{K_t}{2} \hat{t}^\dagger \hat{t} \hat{t}^\dagger \hat{t}, \quad (3.3)$$

where it is clear which terms couple, and to which order. Here, the frequencies have been “dressed” by the mixing of the coupled eigenstates. By comparing to Equation 3.2, one can relate the self-Kerr terms (K_r , K_t) and cross-Kerr term (χ_{rt}) to the transmon parameters; this process will be further detailed in Chapter 6.

Jaynes-Cummings model approximation An alternate picture through which to understand this begins by approximating the transmon as a two-level system, so that it can be treated as a spin-1/2 particle. Then, the coupling between the transmon and \hat{r} is described by the Jaynes-Cummings Hamiltonian [105] in the rotating wave approximation,

$$H_{\text{JC}}/\hbar = \omega_r \hat{r}^\dagger \hat{r} + \omega_t \frac{\hat{\sigma}_z}{2} + g(\hat{r} \hat{\sigma}_+ + \hat{r}^\dagger \hat{\sigma}_-). \quad (3.4)$$

The operator $\hat{\sigma}_z$ is one of the Pauli operators of a two-level system described in Section 2.3.1, while $\hat{\sigma}_\pm$ are combinations of the orthogonal operators, $\hat{\sigma}_x \pm i\hat{\sigma}_y$. In the dispersive limit, the coupling rate g is small compared to the detuning between modes, and H_{JC} can be further approximated as

$$H/\hbar = \omega_r \hat{r}^\dagger \hat{r} + \omega_t \frac{\hat{\sigma}_z}{2} - \chi_{rt} \hat{\sigma}_z \hat{t}^\dagger \hat{t} \quad (3.5)$$

where $\chi_{rt} = \frac{g^2}{\Delta}$.

Physically, transmons can be simply interfaced with 3D cavities, such as those used as a quantum memory in Section 3.5.5. To do this, a transmon can be inserted into the spatial mode volume of the cavity. For the strongest coupling, the dipole moment of the transmon must be oriented so as to overlap with the field pattern of the desired cavity mode. The shape of the transmon capacitor pads can be designed as an antenna accordingly.

Typical parameters Many parameters can be measured or simulated before the transmon is even integrated into an experiment. Following the Ambegaokar-Baratoff relation [106],

$$E_J = \frac{h\Delta}{8e^2 R_N}, \quad (3.6)$$

one can take a measurement of resistance at room temperature (R_N) to predict the inductance of the junction (L_J , related to E_J above) when the device operates in its superconducting state. The relationship in our thin-film aluminum, where the superconducting gap is $2\Delta \approx 84$ GHz [107], gives approximately $L_J \sim 1.27$ nH per k Ω of measured junction resistance. We typically measure junction resistances corresponding with $1 < L_J < 10$ nH. The charging energy $E_C = e^2/(2C_\Sigma)$, where the total capacitance C_Σ is approximately that of the capacitor pads. This capacitance can be extracted from an electrostatic simulation, and typically produces values $C_{\text{pads}} \sim 100$ fF. From these values, the full energy spectrum can be estimated using simulations and black-box modeling techniques [102].

3.4.3 Operational modes

So far, we have introduced two critical types of elements that can be used to construct a cQED module: the memory cavity, and the transmon qubit. Additional elements are still required to fill supporting roles and, later, to facilitate communication between remote modules. Even more elements will be required for complex, error-correctable modules. But though their purposes may differ, and their number may increase, the ingredients remain the same—resonators and transmons will form the basis of these additional circuit elements.

One of the necessary tasks in a module and in a quantum computer is to be able to perform “readout”, or measure the state of, the transmon, storage cavity, and logical components. The dispersive coupling regime of a transmon–resonator system permits readout of the transmon or cavity states. The details of this measurement process will be described in Section 3.4.3. The process begins with a resonator.

The exact form of the readout resonator changes the physics very little. Since the transmon is a structure that is lithographically patterned on a substrate, it may seem natural to use a planar resonator to couple to it and read it out. Compact, lumped-element resonators [108] as well as distributed, coplanar waveguide

(CPW) resonators [109] are available for this purpose. However, a 3D cavity, like those available to serve as quantum memories, can also be tasked with readout. In each of these cases, the nature of the dispersive coupling (Section 3.4.2) can be used to couple them and read out [64, 66].

Measuring resonators

Transmon readout Direct measurement of a transmon is difficult to perform. For qubits that are part of superconducting circuits, direct coupling to an external element is too weak to be useful, or would introduce error channels. Further, if the transmon state were ‘released’ to become incident on a homodyne or heterodyne detector (as it will be in Chapter 6), the nature of such detectors would impose limitations based on quantum uncertainty [110].

Instead, we circumvent the limitations of direct-measurement techniques and read out using a dispersively coupled resonator as a proxy for the transmon state. The dispersive coupling between a transmon \hat{t} and a readout resonator \hat{a} , $\chi_{at}\hat{a}^\dagger\hat{a}\hat{t}^\dagger\hat{t}$, allows high-fidelity measurement of the state of the transmon qubit. Because the frequency of the output depends on the transmon state, the phase of a transmitted signal differs when the transmon is in its ground ($|g\rangle$) or excited state ($|e\rangle$). Measuring this phase thus corresponds to a measurement of the transmon state. Longer readout pulses and higher powers correspond with stronger, more projective measurements [111]. Depending on the application, various strengths and techniques may be used [96, 112].

Representative trajectories of the demodulated ‘IQ’ signal are shown in Figure 3.4. Each one of these trajectories can be excised and then integrated or averaged to obtain a single measurement shot. By histogramming the resulting IQ-values, we can establish a threshold between the two distributions (Figure 3.5). If this threshold sufficiently distinguishes between the two states, the qubit can be read out on a shot-by-shot basis. Single-shot measurement enables feedback and conditional operations, which can be important for tasks like error correction [47]. Readout of more than one transmon at once is also possible, and will be useful in Chapter 7; an example of this operation is given in Appendix A.4.

Resonator characterization The measurement of a microwave resonator will be useful in other circumstances too, such as to establish its coupling strength and internal quality factor. To fully characterize

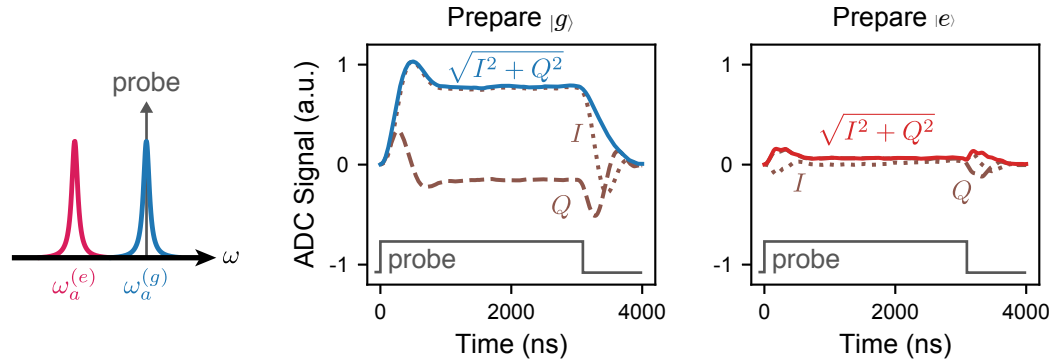


Figure 3.4 | **Dispersive readout trajectories of a transmon.** Frequency arrangement and trajectories of an excited readout resonator that is coupled to a qubit in the dispersive limit. The qubit is prepared in its ground or excited state, and the readout resonator is probed at the frequency corresponding with the qubit in its ground state ($\omega_a^{(g)}$). The insets show the duration of the square probe pulse. The trajectories have been rotated in IQ-space to minimize the Q-quadrature signal contrast.

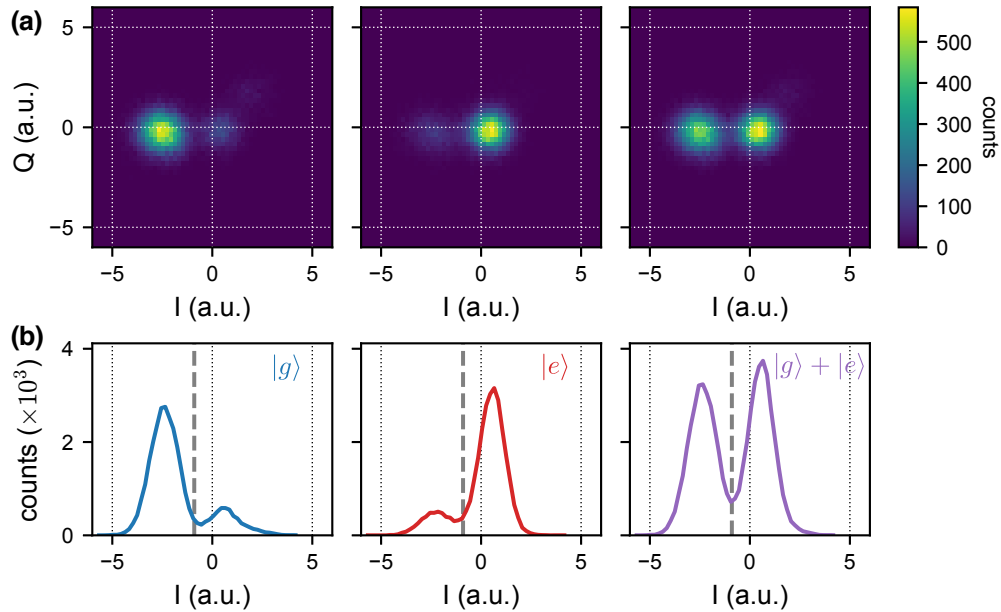


Figure 3.5 | **Readout histograms.** **a** IQ histogram of repeated measurements of a transmon prepared in its ground state, its excited state, or a superposition state. **b** The histograms in **a** averaged along Q to obtain a 1D distribution. This shows that a small population remains in $|g\rangle$ or $|e\rangle$, even when the opposite state is prepared. These distributions permit selection of a threshold (dashed line) that supports state determination on a shot-by-shot basis.




Figure 3.6 | **Basic resonator measurement configurations.** In each panel, the nomenclature associated with the equations in the text is introduced. Direction of the signal is indicated with black arrows. **a** Transmission through a parallel RLC resonator from port 1 to port 2. **b** Reflection uses a circulator to route the reflected signal to the output port. **c** The hanger configuration replaces the circulator in **b** with an unbalanced power divider. Half of the input and reflected power is lost with respect to the desired ports (red arrows).

a resonance, we must measure the complex transmission parameter S_{21} . Fitting this data to a circuit model of the resonance allows extraction of the resonant frequency, “loaded” and “unloaded” quality factors, and other parameters modeled in the circuit.

We typically arrange resonant circuits in one of three ways, designated “transmission”, “reflection”, or “hanger” (Figure 3.6). For the same resonator characteristics, each produces a slightly different response; reflection and hanger offer the additional benefit of being, in some sense, “self-calibrating”. An uncalibrated measurement of these circuits will incorporate the gain and attenuation of all other circuit elements along the signal path; these are difficult to calibrate out in a measurement that passes through a dilution refrigerator and contains active and directional components. The overall “gain factor” in a transmission measurement must be calibrated independently, whereas the gain factor in a reflection or hanger measurement can be treated as a dependent fit parameter.

The response of a resonant circuit in each of these configurations can be easily described by *input-output theory*, among other means [113]. (Classical models can also describe these circuits [114, 115, 116].) Transmission measurements (Figure 3.6a) are common in cavity-qubit experiments because the two ports of the circuit can be asymmetrically coupled, thus directing emission preferentially. Reflection measurement (Figure 3.6b), in combination with a circulator, can be used to appropriately direct the input and output signals, with signal loss dominated by insertion loss of the circulators. However, since this singular

port will be coupled relatively strongly, care must be taken to avoid thermal and dissipative loading of elements that would normally be driven from a weakly-coupled input port. Hanger measurements (Figure 3.6c) are known for their simplicity, wherein an unbalanced tee (25Ω impedance seen from any one port, compared to a matched 50Ω) replaces the circulator of the reflection circuit. In this way, half of the incident energy never makes it to the cavity; similarly, half of the emitted energy travels away from the output port (red signals in Figure 3.6c). However, the hanger configuration requires no special circuit elements and allows for many elements to be multiplexed using the same feedline. It has been a mainstay in the measurement of planar resonators for a long time [116]. The characteristic behavior of each of these circuits is described in Figure 3.7.

Using input–output theory, we derive here the measured signal in the three canonical cases of Figure 3.6. In the case of transmission, there are two external ports, κ_{in} and $\kappa_{\text{out}} \equiv \kappa_c$. We introduce the signal via κ_{in} , which we assume to be weakly coupled compared to the other port. We also model an internal loss port, with decay rate κ_i . The total loss will then be $\kappa_{\text{tot}} \approx \kappa_i + \kappa_c$, and the transmission will be

$$S_{21}^{(\text{tx})}(\omega) = \frac{\sqrt{\kappa_c \kappa_{\text{in}}}}{i(\omega - \omega_0) - \kappa_{\text{tot}}/2} \quad (3.7)$$

where ω_0 is the resonant frequency. The absolute value of the transmission depends on the input coupling. The phase rolls through π radians across the resonance. The insertion loss is set by κ_{in} and κ_c , whereas the total loss determines the bandwidth of the resonance. To extricate the internal loss from the total loss rate, the two coupling rates must be known well. In practical cases, these values within a closed measurement setup are only known within several decibels (dB). In these situations, transmission measurements can be more difficult to accurately calibrate.

Measurement in reflection produces a response

$$S_{21}^{(\text{refl})}(\omega) = \frac{i(\omega - \omega_0) + \kappa_c - \kappa_{\text{tot}}/2}{i(\omega - \omega_0) - \kappa_{\text{tot}}/2}. \quad (3.8)$$

The diameter of the circle when $\text{Im}[S_{21}]$ is plotted against $\text{Re}[S_{21}]$ is $Q/Q_c = \kappa_c/\kappa_{\text{tot}}$. The phase of S_{21} rolls through 2π across the resonance.

By considering the fact that the hanger circuit is identical to the reflection one, when viewed from the

lower port of the tee, we can model the hanger resonator as a combination of half of the unperturbed input signal with the response from the remaining incident half. A hanger resonator's symmetric response then goes as

$$S_{21}^{(\text{hgr})}(\omega) = \frac{1}{2} \left(1 + S_{21}^{(\text{refl})}(\omega) \right). \quad (3.9)$$

The circle diameter represents $Q/(2Q_c)$. The total quality factor can be obtained from the “circle phase”, which parameterizes phase around the S_{21} circle in terms of frequency, and encompasses 2π in all cases. The circle phase obeys a function of the form $\tan^{-1}\left(\frac{2(\omega-\omega_0)}{\kappa_{\text{tot}}}\right)$. From this, and the measurement of the circle diameter (or a multi-dimensional fit to $S_{21}(\omega)$), one can extract both Q_i and Q_c .

In these last two cases, our simplified model excludes any asymmetry from transmission line impedance mismatches, which rotates and scales these circles about $S_{21} = 1$. By considering the complete model, such asymmetry can be accounted for [115]. In practice, particularly in planar resonators, asymmetry can be very significant.

The data and fits shown in Figure 3.7 are for complex data as a function of frequency. The values of S_{21} are plotted in complex space, forming a circle that is characteristic of a microwave resonance. For the most accurate results, it is critical to fit the complex data, rather than the magnitude or phase alone. In the unfortunate case where one intends to extract information about Q_i of resonator, but the resonator is “overcoupled” ($Q_c \ll Q_i$), then fitting the inverse of the complex data, S_{21}^{-1} , which also forms a circle in complex space, can produce better results [109].

The internal (“unloaded”) quality factor represents the lossiness of a resonator, and will be described further in Chapters 4 and 5. It is, however, not the quantity measured at face value; other elements of the resonator measurement circuit “load” the measured resonance bandwidth, permitting only direct measurement of the “total” quality factor,

$$Q_{\text{tot}}^{-1} = Q_c^{-1} + Q_i^{-1} \quad (3.10)$$

where Q_i and Q_c represent the internal (or “unloaded”) and coupling (or “external”) quality factors, respectively. The values of two contributions can be inferred using the circuit models given above. Since the coupling quality factor is a value largely determined by the measurement taker, it behooves one to make this

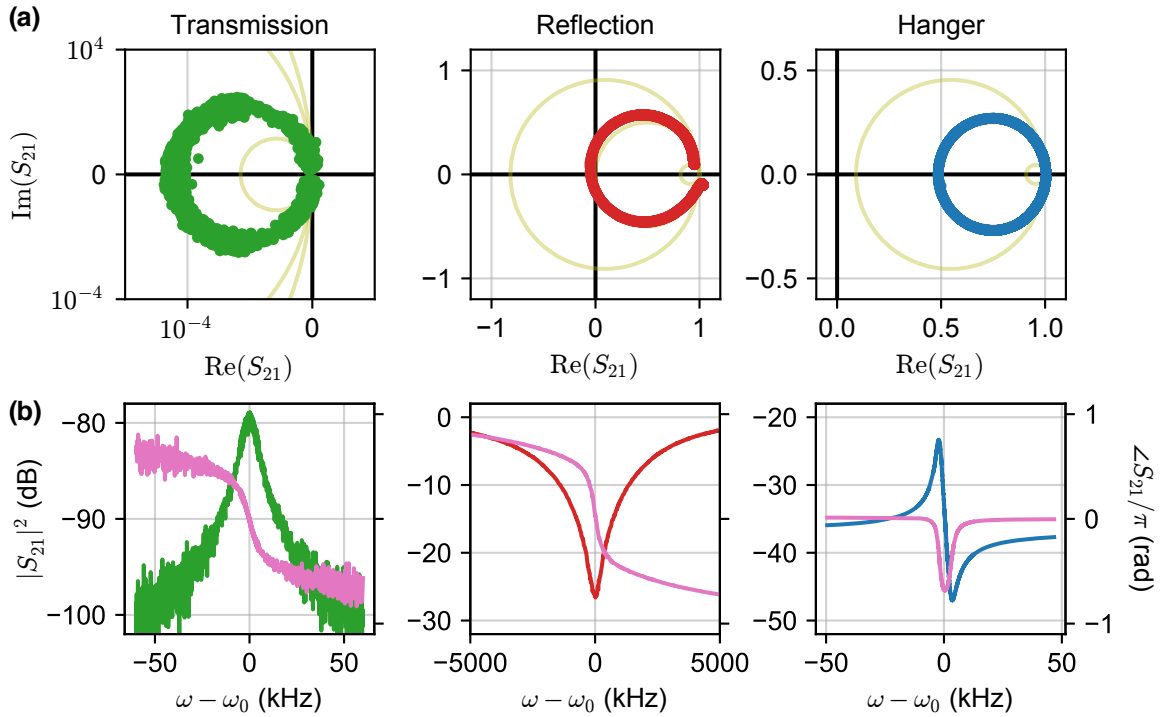


Figure 3.7 | **Response in different resonator configurations.** In each panel, simulated responses (gold) are plotted for $Q_c = 10Q_i$, $Q_c = Q_i$, and $Q_i = 10Q_c$. **a** The complex transmission parameter S_{21} signal from a measurement of a resonant circuit in each configuration, correcting for gain, attenuation, electrical delay, and asymmetry in the cases of reflection and hanger. In transmission, the size of the circle grows as κ_c increases; its actual size depends on κ_{in} . In reflection, an undercoupled resonator's signal will be centered near unity, a critically coupled resonator's signal will pass through the origin, and a strongly overcoupled resonator will intercept $(-1, 0)$. In hanger, the reduction in contrast and mixing with the input signal means that a critically coupled resonance circle will pass through $(0.5, 0)$ and an overcoupled signal will pass through the origin. Both reflection and hanger are subject to asymmetries in impedance as viewed from one direction or the other, which can be included in their models. **b** The raw magnitude (left axis) and phase response (pink, right axis) of the data in **a**.

value relatively large compared to Q_i . In this way, the measured value $Q_{\text{tot}} \approx Q_i$, thus allowing extraction of Q_i with increased accuracy (by reducing the dependence on other parameters). This “undercoupled” regime of $Q_c \gg Q_i$ also benefits a resonator designated for information storage, so that information does not externally decay unnecessarily. Thus, such measurement is proof-of-concept demonstration of a practical regime. However, a Q_c that is too large will reduce measurement signal-to-noise ratio. A healthy undercoupling (for example, $Q_c = 10Q_i$ means the above approximation has $< 10\%$ inaccuracy) can balance trustworthiness with ease of measurement.

Circulating photon number in a hanger configuration is given by

$$\bar{n} = \frac{Q_i^2 Q_c}{\hbar \omega_0^2 (Q_i + Q_c)^2} P_{\text{in}}, \quad (3.11)$$

which allows the experimenter to calibrate the approximate photon number to within a typical factor of 2–5 using the fit resonance data and information about the input lines leading up to the device. Circuit QED typically operates in the regime of ~ 1 photon in resonators. By comparison, the highest- Q 3D cavities are niobium-surface cavities used in particle accelerators with $Q > 10^{10}$; these are operated at very high circulating power [117]. Similar cavity designs operated at or below powers corresponding with single-photon occupancy have remained below $Q = 10^9$.

The presiding theory for this behavior is the presence of the two-level systems (TLSs) at, or near, the surface of the cavity [118, 119]. In this model, some density of polarizable TLSs is present in a location that can be excited by the intracavity electric field. At high powers, an RF drive will immediately “saturate” the TLSs to a mixed state; they rapidly cycle between their two energy levels. At weak powers, the TLSs can decay and decohere, according to their internal dynamics, long before they are likely to be re-excited. When TLSs are unfortunate enough to interact appreciably with our resonant mode of interest (by virtue of a small frequency detuning), this adds dissipation to the system. This power-dependent behavior manifests in the measured Q_i of a resonator as

$$Q_i = \frac{\sqrt{1 + P/P_c}}{p_{\text{TLS}} \cdot \tan \delta_{\text{TLS}} \cdot \tanh\left(\frac{\hbar \omega}{2k_B T}\right)} \quad (3.12)$$

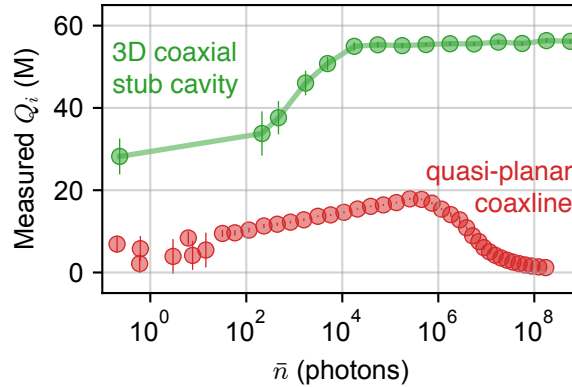


Figure 3.8 | **Resonator power dependence.** Measured resonator quality factor as a function of calculated circulating photon number for two representative resonators. The amount by which Q_i changes from low to high power is well-described by a model of sampling two-level systems in a surface dielectric layer. The lower data corresponds with a quasi-planar design, while the upper is a voluminous 3D cavity. In the quasi-planar design, the drop in Q_i at higher powers is likely due to a breakdown of the superconductivity of the thin film. (Both designs are discussed in Chapter 4.) TLSs generally have a larger fractional effect on planar or quasi-planar structures.

and can be used to determine the loss tangent $\tan \delta_{\text{TLS}}$ and saturation power P_c of the TLS ensemble, where p_{TLS} is the difficult-to-predict participation ratio of the ensemble. Because the participation of *all* surface loss mechanisms tends to scale generally with the field surface-to-volume ratio, this behavior is often more strongly manifested in planar resonators (Figure 3.8).

3.4.4 Parametric conversion elements

Some additional, Josephson-junction-based circuit elements will be helpful to include in our circuit QED toolbox. First, quantum-limited parametric amplifiers can allow amplification of a qubit readout signal while adding just the minimum noise required by fundamental quantum mechanical vacuum fluctuations. These amplifiers, which include the Josephson parametric converter (JPC), are versatile and can be operated to facilitate more than just amplification [120, 121]. They can be envisioned as a critical communication or conversion element that routes or mediates signals between modules, for example.

Other junction-based conversion and coupling elements have also been developed [122]. The Superconducting Nonlinear Asymmetric Inductive eLement (SNAIL) is an intriguing choice for such a converter, in that it can be tuned to reduce conversion imperfections induced by some higher-order nonlinear terms. Hybrid designs of one or many junctions and resonators have been considered. The desired re-

quirements for a modular conversion element will be discussed further in Chapter 6.

Our cQED toolbox This section has introduced a powerful set of tools that will support this thesis, and many more to come. While there is insufficient space to detail the many others of which this toolbox forms of a small subset, it should be noted that Josephson-junction-based elements and harmonic oscillators constitute the basis of most of these circuit constructs. Next, I will discuss the important operations that can be done using these elements—particularly those that aid us in our goal of building a network of computational nodes.

3.5 Useful cQED operations

In describing the operations enabled by circuit QED elements, we also consider the criteria that a module or a scalable quantum computer must satisfy. We will address how each element can be initialized, controlled, and measured. We will also introduce tools that enable more complex operations, such as those required to detect and correct errors, or perform state tomography on the large Hilbert space of cavities. This serves not as a broad overview of all possible operations—for there are many—but rather an overview of tools that will be frequency referenced in the coming chapters.

3.5.1 System initialization

Ideal resonators and cavities made out of aluminum ($T_c \sim 1.2$ K), and cooled to the base temperature of a dilution refrigerator (~ 20 mK), should have negligible populations beyond the zero-excitation ground or zero-photon states. Broken Cooper pairs of the superconductor form electron quasiparticles that obey Fermi–Dirac statistics; at 10 GHz and 20 mK the probability of finding a quasiparticle in thermal equilibrium goes as $(1 + e^{h\nu/(k_B T)})^{-1} \approx e^{-24} \approx 10^{-10}$ [68]. Realistic superconducting circuits, however, are often measured to have anomalously high populations of non-equilibrium quasiparticles [123, 124]. Possible reasons for this population include poor experiment-stage filtering, an uncontrolled magnetic field environment, or weak thermalization of elements such as circulators that comprise part of the measurement chain. Therefore, the initial state of the system is not always guaranteed to be in its collective ground state.

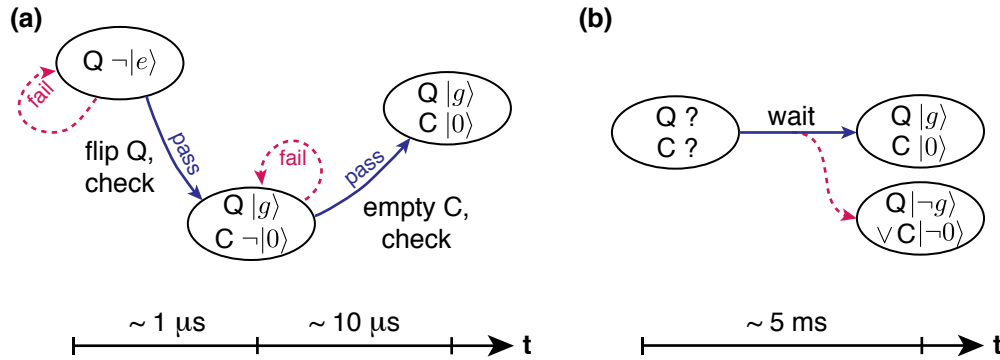


Figure 3.9 | **Feedback cooling process.** **a** A state machine diagram showing the process of feedback cooling for elements in a typical experiment (Q: qubit, C: cavity). More complex state machines can be enacted with multiple qubits and cavities. **b** A process that does not cool using feedback. The relevant timescales are compared. After some fraction of repetitions, the qubit or cavity may still be found in an excited state, since no active control is exercised.

Greater certainty of the initial state can be achieved in two ways. The first, which requires no special hardware except the ability to measure, is by postselection. Postselection enables conditioning of later data based on the state in which the system was measured to have begun. The downside of postselection comes from the resultant measurement and repetition rates—if the system is frequently found not to be initialized, then most of the resulting data is thrown away. In extreme cases, operations that require multiple successive steps to be successful can never be realized through postselection; the probabilities of an occurrence will simply be too low to be useful.

A feedback scheme, on the other hand, uses the result of the initialization measurement to implement a conditional initialization pulse. If the measurement produced a successful result, then the measurement can proceed as normal. If it was unsuccessful, then a “cooling” procedure can be performed to try to remove the excitation from the offending element (Figure 3.9).

In our experiments, we utilize a field-programmable-array-based hardware platform to enable this rapid feedback loop [125]. Whereas typical thermal populations might be as high as $p_{-0} \sim 0.1$, we are able to use quantum non-demolition (QND) measurements [111], cavity-photon-number-discriminating (“selective”) π -pulses to invert qubit populations, and rapid pump-based evacuation sequences on cavities (Chapter 6) to reduce this population to one that is $\lesssim 0.02$.

State preparation

To prepare complex cavity states, we can apply pulses at qubit and cavity frequencies that are designed using a numerically optimized control technique [126]. This state preparation technique can produce non-Gaussian cavity states, like cat (coherent state superposition) states, individual Fock states, or Fock state superpositions, with typical fidelities above 95 %.

This form of state creation only requires the linear control drives, the means of application of which are detailed in the following section, and carefully calibrated Hamiltonian parameters. As detailed in Section 3.4.2, the Hamiltonian of our systems includes nonlinear terms provided by the transmon, such as “Kerr” terms. Such “higher-order” terms can be shown, in combination with linear operations, to allow universal control over a continuous-variable system, in finite time and to arbitrary accuracy [127]. We implement this by modeling the Hamiltonian (by measuring coefficients of arbitrarily high orders of the transmon–cavity system) and determining the drives that are required to produce a desired unitary operation. This system thus contains all the elements to perform universal computation, as long as the proper calibrations have been performed (though for particularly high-order terms, effects like frequency dispersion and mismatched delay lengths will likely dominate). As a result, various efficient ways to implement a universal set of operations have been developed for our cavity–qubit systems [19, 126].

3.5.2 Transmon and cavity control

Transmons and cavities can be driven with microwave pulses, often via weakly-coupled ports. In 3D cQED, these ports can be included easily and non-intrusively (see Section 4.3.1). Parametric processes (which will be introduced in Chapter 6) can use the same lines and ports to drive these modes.

Control over a transmon qubit is achieved with coherent, zero-detuning shaped drives that are calibrated to rotate the state by a defined amount with the qubit’s Bloch sphere. The shape of the pulse determines its spectral content. A pulse with a temporal Gaussian profile, for example, will also deliver spectral content shaped by a Gaussian about the carrier frequency [128]. For the transmon, the anharmonicity (in addition to the pulse power that must be delivered) determines the necessary length of a pulse [129]. A pulse that is too lengthy will be spectrally narrow, but limit the number of performable operations in the

coherence time of the transmon. Sometimes, in spite of this, we must use narrow pulses to select transmon transitions that are dependent on a particular number state of the cavity.

A similar pulse at the cavity frequency has a very different effect, because it is a nearly-harmonic oscillator and one individual transition is not being addressed (at least, the relevant experimental parameters will not allow it). This pulse has the effect of shuffling energy up the photon-number ladder as time goes on, and can only create classical (coherent) states that have a Poisson distribution. Thermal effects, as an example, may perturb this distribution by introducing competing decay rates. In general, however, this process is described by the unitary $\hat{D}(\alpha) = \exp(\alpha\hat{a}^\dagger - \alpha^*\hat{a})$, where α relates to the strength of the drive [94].

3.5.3 Cavity measurement

The readout scheme described in Section 3.4.3 allows us to interrogate the state of the readout resonator in order to learn about the state of the cavity. But how does one learn about the state of a storage cavity—that is, a resonator that is coupled to a transmon but not directly coupled to the readout resonator?

One basis in which the cavity state can be described is the Fock, or photon number, basis. In this description, the cavity state wavefunction is given by

$$|\psi\rangle = \mathcal{N} \sum_n c_n |n\rangle \quad (3.13)$$

where the Fock states $\{|n\rangle\}$ are weighted by amplitudes c_n such that the cavity has the probability $|c_n|^2$ of being found to contain n photons. By knowing all these coefficients, one can reconstruct the full wavefunction of the cavity.

The dispersive coupling between transmon and storage cavity causes a photon-number-dependent frequency shift of the transmon, in the same way that it does between the transmon and the readout resonator [130]. If this coupling is strong enough to discriminate between different frequencies, then the probabilities $|c_n|^2$ can be directly measured via a transmon spectroscopy experiment. We can probe whether the cavity contains n photons by first applying a π -pulse, selective on the transition for the cavity containing n photons, followed by measuring the transmon state. The fidelity of the positive readout result (cavity

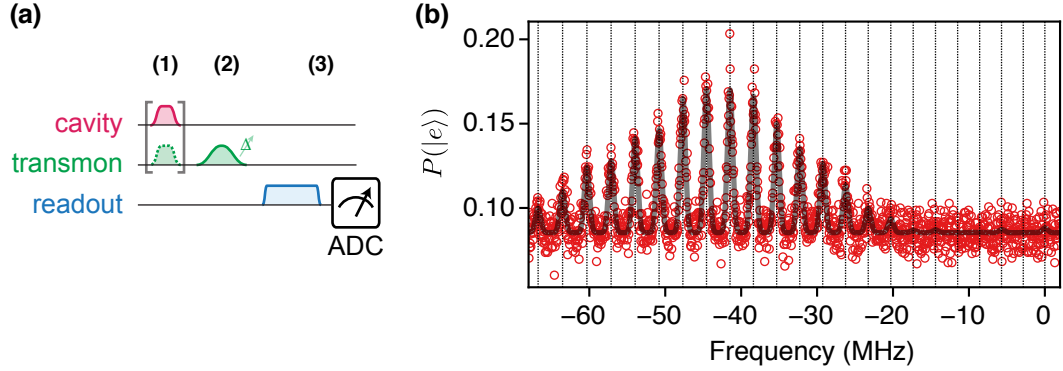


Figure 3.10 | **Number peak resolution of the cavity state.** **a** The pulse sequence of a selective spectroscopy experiment to measure population of a cavity coupled to a transmon. Steps (1), (2), and (3) are state preparation, selective transmon π -pulse, and readout, respectively. **b** Results of spectroscopy produce resolvable number peaks. This information reveals, in part, the cavity state.

containing n photons) is simply the transmon readout fidelity for finding the transmon in its excited state $|e\rangle$. The negative result (cavity *not* containing n photons) is determined by the fidelity for $|g\rangle$ as well as the probability with which the n -photon state has not been mapped successfully onto $|e\rangle$; the combined fidelity for the negative results can thus be limited by qubit decoherence.

The result will be a convolution of Gaussian peaks, according to the dispersive shift, appearing at positions in frequency space parameterized as a function of cavity photon number

$$f_0(n) = n\chi_{at} + n(n-1)\chi'_{at}/2 \quad (3.14)$$

where χ'_{at} is the next-order term coupling the cavity and the transmon in the expanded cosine Hamiltonian.

Figure 3.10 shows the number-state-dependent mapping of a large coherent state $|\alpha = \sqrt{14}\rangle$ onto the qubit state. The Gaussian peaks are distributed according to Equation 3.14, with their amplitudes given by a Poisson distribution, and their widths by the selectiveness (spectral narrowness) of the π -pulse in step (2). We can also directly measure probe the cavity on a single number peak, as in Figure 3.11.

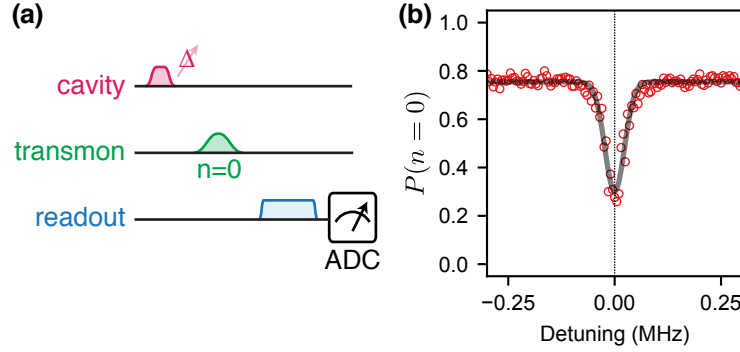


Figure 3.11 | **Direct spectroscopy of a storage cavity.** **a** The pulse sequence of a selective spectroscopy experiment to measure the frequency response of a single cavity number-peak. **b** Results of spectroscopy produces an inverted behavior compared to Figure 3.10 because of the transmon mapping dependent on $n = 0$.

3.5.4 Cavity state tomography

The spectroscopy experiment described above allows access to the photon number probabilities of the cavity state. However, it does not constitute full state tomography, since the phases between the Fock basis states cannot be known.

The same kind of state tomography as is done on a qubit, by measuring the positive operator-valued measure (POVM) associated with the Pauli operators [131], is prohibitively difficult on a cavity due to its massive Hilbert space. However, one can construct more general POVMs to implement measurements that are optimally suited to the detector and the system, including continuous-variable representations [132]. Wigner tomography, a task that is experimentally accessible in our systems, is akin to full state tomography on a cavity. It can be done through a series of displaced parity measurements [105, 133].

This is possible because measuring the Wigner function is as simple as measuring the expectation value of displaced parity and scaling it by a $2/\pi$, [134, 135]

$$\langle \hat{W} \rangle = \left\langle \frac{2}{\pi} \sum_{n=0}^{\infty} (-1)^n \hat{D}(\alpha) |n\rangle \langle n| \hat{D}^\dagger(\alpha) \right\rangle, \quad (3.15)$$

which is a combination of straight-forward operations (displacement and parity measurement) in our system.

Wigner quasiprobability

The Wigner function is one form of a generalized (parameterized by s) quasiprobability function, W_s , in term of measured quadratures x and p :

$$W_s(x, p) = \frac{1}{\pi(s' - s)} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} W_{s'}(x', p') \times \exp\left\{\left[-\frac{(x - x')^2 + (p - p')^2}{s' - s}\right]\right\} dx' dp' \quad (3.16)$$

with $s < s'$. The most common forms of this quasiprobability distribution include the Glauber P-function (W_1), the Wigner function (W_0), and Husimi Q-function (W_{-1}). This phase-space transformation allows us to define more “physical” (lower- s) distributions from ones that are more “abstracted” [136, 137]. However, this self-referencing definition requires a representation of $W_{s'}(x', p')$ from which to begin.

As a quasiprobability distribution that can actually be constructed from measurements, the Wigner function is one such useful representation of non-classical electromagnetic fields. One way it can be expressed is in terms of the Wigner transform of the density matrix $\hat{\rho}$,

$$W(x, p) = W_{-1}(x, p) = \frac{1}{2\pi} \int_{-\infty}^{+\infty} e^{ipu} \langle x - u/2 | \hat{\rho} | x + u/2 \rangle du, \quad (3.17)$$

in terms of conjugate variables x and p . The parity at a displaced location in phase space is effectively sampled. Importantly, this implies that the density matrix can be reconstructed from the Wigner function.

Husimi-Q quasiprobability

The Q-function is obtained by Equation 3.16 with $s' = 0$, $s = -1$, and is equivalent to convolving the Wigner function with a Gaussian as

$$Q(x, p) = W_0(x, p) = \frac{1}{\pi} \int_{-\infty}^{+\infty} \int_{-\infty}^{+\infty} W(x', p') e^{-(x-x')^2 - (p-p')^2} dx' dp'. \quad (3.18)$$

The convolution smooths the Wigner function by adding additional noise and removing all distinct non-classical features, like negativity, and suppressing interference fringes [105]. This will be useful in Chapter 6 when describing states represented by measurements using a phase-preserving amplifier, which adds

noise by sampling both quadratures, yielding the Q-function. (Any additional loss will further obscure the resulting distribution in Equation 3.18 according to Equation 6.27.) Alternative measurement techniques, such as use of a phase-sensitive amplifier, can instead yield the Wigner function [135, 138].

Displaced parity

State tomography begins by preparing or otherwise transferring a state of interest into the chosen cavity. Then, the cavity is displaced by a variable, complex-valued displacement β . At each displacement, the average parity is measured. A best-case measurement then produces the Wigner function (worse measurements add noise, reducing contrast and any negativity). As described above, the cavity density matrix state description can be reconstructed from the Wigner function.

There are several straight-forward measurement methods that can produce the values of displaced parity of a cavity state [138]. In the systems in this work, we will select a somewhat less conventional way—the parity is mapped onto the transmon by applying a set of selective π -pulses simultaneously, on either the even or odd number-peaks, up to $\bar{n} = 25$. Then, taking the difference of the even and odd mapping sequences, we obtain a value proportional to parity (and therefore to the Wigner function), independent of the initial transmon state. This method is equivalent to other representations of the Wigner function [134].

To compensate for finite π -pulse and measurement contrast, we must then normalize the resulting data so that it integrates to unity over all phase space, leading to a physical Wigner function. Displacements are typically performed up to $\beta = 2.5$, a value corresponding with a photon number up to 12.5 photons, and results in Wigner functions that should capture $> 99\%$ of the energy contained in typical relevant cavity states, wherein $\bar{n} \leq 2$.

3.5.5 Mapping and encoding/decoding

The cavity measurements described above are a form of mapping, of the cavity state, onto the transmon state. More complex mappings can be also be performed, like the qcMAP [139]. Such encoding remapping can take advantage of the processing capability of the transmon, for example, while moving that information to the cavity for long-term storage, or to take advantage of error correction protocols [47].

In our system, the same tools that allow us to prepare arbitrary states in a cavity or a qubit [126] also enable straight-forward construction of conditional, arbitrary unitaries. These can enable control over logical states. In addition, we will use such unitaries to, as one example, perform a cavity basis rotation (e.g. $|0\rangle \mapsto |0 + 1\rangle$) and allow straight-forward measurement in other bases.

The tools discussed here constitute only a handful of those available. The convenient control knobs provided by circuit QED, hand in hand with a rich catalog of microwave-frequency technologies, enrich this toolbox. By selecting and assembling these state-of-the-art elements, we intend to produce a highly coherent module. The following chapter will begin this process, introducing an element that is compatible with high- Q 3D cavities, and completing and evaluating an elementary module. Chapters 6 and 7 will then develop and implement a procedure for using this module as a node within a modular quantum network.

4

Seamless Hardware for Scalable Integration

In this chapter, I discuss the fundamental quest to improve to coherence of the elements that make up our quantum systems. Superconducting microwave cavities and transmon qubits are critical members of this contingent. At this time of this writing, their coherence times are long enough to perform many useful tasks and permit research for years to come. Techniques are being developed to address their limited coherence times, by storing information redundantly and performing error-correction. Hand-in-hand with this effort, however, must be a push to improve the basic performance of cavities and qubits. Fundamentally improved coherence times are complementary to these efforts, reducing the number of resources required to achieve the final goal. Even with the best recipe, the resulting cooked meal can only be as good as the ingredients you begin with. Here, we strive to improve both ingredients and recipes.

4.1 Challenges of many-element system coherence

Scaling presents challenges that are not apparent in small systems of one or two qubits. The first, most obvious challenge comes from fabricating the devices required for such a system. Imagine there is some success probability associated with slightly imperfect fabrication of a qubit with the desired parameters, say $p \sim 0.90$. In addition to the likelihood of measuring proper post-fabrication parameters, this may also include survival rate and performance within limits once tested under operating conditions, or repeatability between cycles of an experiment. If only seven qubits are required to be fabricated simultaneously within

these limits, the chance of success is already worse than flipping a single coin.

Perfecting a fabrication process is certainly not out of reach, however, as evidenced by the massively successful semiconductor fabrication industry. We consider this largely an engineering challenge that will be solved in due course, and seek only briefly to identify present limitations of qubit and resonator coherence properties in this chapter to follow. Near-term improvements will facilitate current research. In the mean time, we focus on less obvious challenges faced by many-element systems.

One canonical problem that exists in controlled quantum systems is that more elements demand more couplings between elements and more control signals (external couplings). If not carefully controlled, these inter-element couplings can lead to unwanted behavior. External couplings allow for the introduction of environmental noise. Therefore, it tends to be the case that increased complexity does not directly translate to increased usability (such as the ability for greater computing power or information storage). At some point, determined by engineering prowess (and at some point fundamental limits), a larger, scaled system will perform more poorly than a smaller one. One of a handful of problems endemic to microwave circuits will likely be the cause.

4.2 Microwave hygiene

The problem of unwanted couplings falls into the domain of “microwave hygiene”, the electromagnetic “cleanliness” of the environment in which a system is designed. Since we work with systems that are nominally engineered on every level—from the superconducting box to the Hamiltonian guiding element interactions—full knowledge of the system is paramount. Poor microwave hygiene is characterized by the presence of unwanted, unexpected (spurious) resonant modes, or broadband “leakage” of field from one section of a device to another (Figure 4.1). It has been the focus of classical studies in microwave electronics, and has recently found special importance in the design of experiments with superconducting circuits.

Superconducting circuits, which rely on high-quality material performance, are more susceptible to many of these losses; in normal metal circuits, conductor loss may dominate over these effects, rendering them unnoticeable. These loss mechanisms include unwanted propagation/emission channels, such as slotline or box modes. They may also include physical dissipation, such as PCB circuit board material or




Figure 4.1 | **Possible hazards in planar architectures.** An illustration of the difficulties that can be present in a microwave circuit when trying to maintain a low-loss environment. Here, a lithographically patterned chip is positioned within a circuit board holder, which in turn sits within a metal enclosure (one half shown). In addition to the desired coupling paths (black arrows), unwanted couplings or losses such as cross-talk/parasitic coupling (green dashed lines), slotline/box modes (yellow arrows), material dissipation (red lines), and radiation (pink arrows) may be present.

adhesives that may be present. Circuit properties may be modified by parasitic capacitance to grounded elements or parasitic inductance from the presence of wirebonds.

Some of these effects can be mitigated through careful microwave engineering and design. Box/package modes may be suppressed by “shimming” circuits, or covering them in tightly fitting metal shields. Slotline modes might be suppressed by cross-over elements coupling ground planes within a circuit, such as “air bridge” structures. Circuit elements (like filters) can be introduced to compensate for coupling to predictable modes. Materials can be carefully selected—and even etched and removed—and the system constructed so as to reduce the spatial concentration of energy susceptible to dissipation (Section 5.1). Finally, precise multilayer interconnects can be designed so as to mitigate the unwanted effects of connections like wirebonds.

Still, planar circuits are difficult to fully protect from environmental losses. Many of these “fixes” result in smaller, more isolated circuits and more concentrated electric fields. A more robust response follows the trend toward larger, more dilute 3D resonant modes. Using waveguide structures, a “hygienic” environment can be constructed from the ground up (Section 4.3). Beginning with a blanket suppression of propagating modes and the presence of only superconducting materials and ultra-low-loss dielectrics,

losses can be considered one-by-one as every element is added (Chapter 5). In this way, simple 3D systems allow losses to begin to be reliably modeled.

Despite their relative cleanliness, traditional 3D modes have some fundamental limitations due to their typically seam-based construction, as the following section will show. These design types also face some capability limitations; for example, omitting seams in superconducting cavities makes it difficult to introduce flux necessary for tuning certain elements. The most hygienic and broadly capable designs must ultimately address both the challenges of planar designs and 3D cavities, while incorporating features from both (Section 4.3.1).

4.2.1 The effect of seams in 3D cavities

Seams, for example at the joint between two halves of a cavity shell, are generally a design necessity when traditional machining is used to produce 3D cavities with well-behaved TE or TM resonant modes. Subtractive machining processes, where grooves and holes are milled in metal, are simply incompatible with the requirement of cavities to form a light-tight box for microwave photons. Some processes, such as electron-beam welding, may be used to join the halves of a cavity together. Additive processes like 3D printing might even be able to avoid a seam altogether. However, the effects of these processes on materials used in superconducting quantum information systems can be detrimental and is not fully understood.¹ Therefore, it became important to understand the effect of seams on the performance of superconducting microwave cavity resonators.

The design of these two shells containing cavity halves, and thus the location of the seam, can be engineered so as to reduce (or even eliminate) the presence of current that would cross it. One can even play tricks, such as by using a microwave choke, to create the appearance of a short or open circuit at certain frequencies. However, no machined structure is ideal. Inaccuracy in design and fabrication will lead to an electric field structure of resonant modes that drives current across the seam. “Loading” of a resonant mode can perturb not only its resonant frequency, but also the spatial distribution of its electromagnetic field. In our systems, the tunnels for recessed coupling connectors alone can cause this effect. Often,

1. Both of these processes were investigated as part of this thesis work. The resulting cavities did not present coherence properties exceeding those of traditionally processed materials.

coupling to planar elements such as transmons requires the introduction of a high-permittivity dielectric substrate, such as sapphire, into the cavity. According to simulations, these structures tend to perturb the distribution of the electric current far more than by non-idealities alone. In this way, some current will always find its way across a seam.

So why is current necessarily bad? If two bulk superconductors are connected, what is the mechanism by which dissipation should occur? As with many matters in experimental physics, the practical system does not match the ideal one. Common bulk superconducting materials, such as aluminum or indium, are quick to oxidize under normal atmospheric conditions. This produces a layer of dielectric oxide, which dissipates energy at microwave frequencies, at the surfaces and thus the seam between the two superconducting shells. Even if a mechanical or chemical process is used to remove this oxide, exposure to an oxygen-rich environment will almost immediately cause it to regrow. Avoiding this is difficult without significant investment in new experimental procedures.

Some early designs involving cavities with seams are now thought to have been limited by seam loss. For example, the “vertical transmon” design produced low transmon coherence that was not explained by traditional decoherence channels [35]. Rectangular cavities—even those made from chemically etched high-purity aluminum—have demonstrated lifetimes worse than those consistent with conductor or dielectric loss alone [67, 93].

Recent studies have sought to develop a model based on experimental observations of seam-based performance degradation [140]. This model enables an estimation of the seam loss based on mode and material properties. Finite-element modeling can simulate the current density of the resonant mode. Systematic experiments with controlled seam participation can establish bounds on the magnitude of dissipation (the resistivity) produced by joining different materials of interest. Together, these guidelines can allow cavity designers to predict the effect of current flowing across dissipative seams, influencing design decisions.

There is a possible solution to eliminate lossy joints, however. If seams must be present in a particular designs, the superconducting cavity shells can be plated in (or constructed out of) a malleable superconductor like indium, stripped of oxide within an inert environment, and immediately bonded together. This procedure applied to seamed rectangular cavities was shown to produce significantly higher quality factors [141]. Naturally, this process can be difficult to reverse, thus limiting its practicality. An easier solu-

tion, in some cases, is to pursue a seamless cavity design. We will find these come with additional benefits, such as strong control over the electromagnetic environment to which a circuit is exposed.

4.3 Hardware for coherent 3D integration

4.3.1 A seamless architecture for system integration

One way in which the geometric requirements of 3D cavities can be reconciled with traditional machining techniques is by engineering a “seamless” structure that supports resonant TEM modes. Such a structure seeks to strictly obey the requirements of proper microwave hygiene, avoiding the need for additional structures or filtering [142]. At the same time, it should allow the controlled integration of many elements. Ideally, there should exist a path towards miniaturization and scaling.

Package design

Our response to these requirements is to carefully engineer a 3D waveguide package that we call the “coax-line”². Circuits formed from select, low-loss materials and placed in 3D enclosures can avoid many common forms of loss found in planar cQED systems (discussed further in Section 4.3.1). Our design begins with a seamless circular waveguide that forms the package enclosure and acts as a ground plane (Section 4.3.1). Circuit elements are patterned on a substrate, usually sapphire, to define each mode of the device. The deposited and machined conductors are typically chosen to be aluminum. Where no conductor is present on-chip, the waveguide attenuates signals below its cutoff frequency (typically around 40 GHz). The chip is suspended within the enclosure by far-removed clamps at one or each end, where the fields from critical circuit elements are exponentially attenuated.

The clamps are made from two pieces tightening against the chip from both sides, often with soft indium wire or freshly mixed epoxy adhesive (Stycast 2850FT) in between. Either the broad side or the narrow side (within a slot) can be directly clamped. Once the chip is properly positioned, the clamp is fully tightened and the apparatus mounted to the end of the package. Since these enclosures are prone to long

2. It is a quasi-coaxial transmission stripline with strong similarity to a coaxial transmission line, sometimes also called the coax-line or the co-Axline.




Figure 4.2 | **The coaxline architecture.** **a** A diagram of the coaxline with a section of the bulk enclosure removed for visibility. The stripline (gold; electric field pattern shown) is patterned on a sapphire substrate (turquoise). The substrate is suspended in the enclosure by clamps (shown opened) at either end. The transmission signal S_{21} is introduced via the coupling pin (port 1) and measured through the output pin (port 2). **b** The center conductor of a coaxial transmission line “floating” in vacuum. **c** In the coaxline, the stripline conductor plays the role of the center pin in **b**. (Figure adapted from [7]; see [Copyright Permissions](#).)

chip aspect ratios (with dimensions 0.5 x 1 x 40 mm or more), vibrational coupling is a concern (Section 5.6.3). Therefore, the opposite end of the chip is often clamped by identical means.

In this architecture, complex, many-element designs should be possible. We will first evaluate the performance of this design, however, by considering a simple, illuminating test structure on the chip: a single resonator. We will choose the form of a $\lambda/2$ stripline-like resonator and pattern it on the substrate (Section 4.3.1). The lowest resonance frequency is primarily determined by the length of the conducting strip, but also depends on chip size, chip placement, and enclosure diameter. These dependences are studied in the following sections. Like all cQED elements, this kind of device performs best when the highest-quality superconducting materials are used and the system is cooled into its superconducting state.

The enclosed $\lambda/2$ stripline resonator will act like a section of a transmission line with each end terminated with an open circuit. Without the stripline, the enclosure is simply a dielectric-loaded circular waveguide. The stripline is a TEM mode, and supports propagation with no cutoff frequency. The waveguide alone supports only TE and TM modes, and thus has a finite cutoff frequency. Importantly, the cutoff frequency will determine the level of isolation between operational modes as well as from the environment.

Waveguide enclosure

Looking at a cross section of this system, one encounters a partly dielectric-filled circular tunnel near the ends where no patterned conductor is present. The substrate must be held at some point, so the presence of some dielectric is unavoidable when using conventional machining and assembly techniques. For this design to preserve an electromagnetic environment in which spurious coupling are suppressed, the cutoff frequency of this structure must be sufficiently high.

The cutoff frequency can be determined analytically for some geometries, and by simulation for others. The propagation constant β , which describes the propagation of the electric field according to $E \propto e^{-j\beta z}$, can always be expressed in terms of a cutoff wavenumber k_c as $\beta = \sqrt{k^2 - k_c^2}$. Above the cutoff wavenumber, and therefore above the cutoff frequency, the square root operand is positive and β is real, and thus the wave freely propagates. Below the cutoff frequency, the operand becomes negative, the exponential gains another imaginary factor, and the electric field decays exponentially. This is called the evanescent, rather than propagating, regime of the EM wave behavior. The strength of electric field decay depends on how far beneath the cutoff frequency a particular operating frequency is.

The cutoff frequency of a partially loaded waveguide can be calculated exactly for some simple geometries. We can gain some intuition for the effect of dielectric loading by (1) exact simulations of exemplary structures, and (2) a rough analytic approximation. We find it analytically in the approximate geometry shown in Figure 4.3a, the details of which are solved in Appendix A.1.

This analytical behavior can be compared to that of a numerical finite-element simulation of this approximate geometries, as well as the exact coaxline geometry shown in Figure 4.3b. The simulation measures the transmission per unit length as a function of frequency for different diameters of the enclosure (behavior for the coaxline shown in Figure 4.3c). The frequency above which transmission becomes loss-

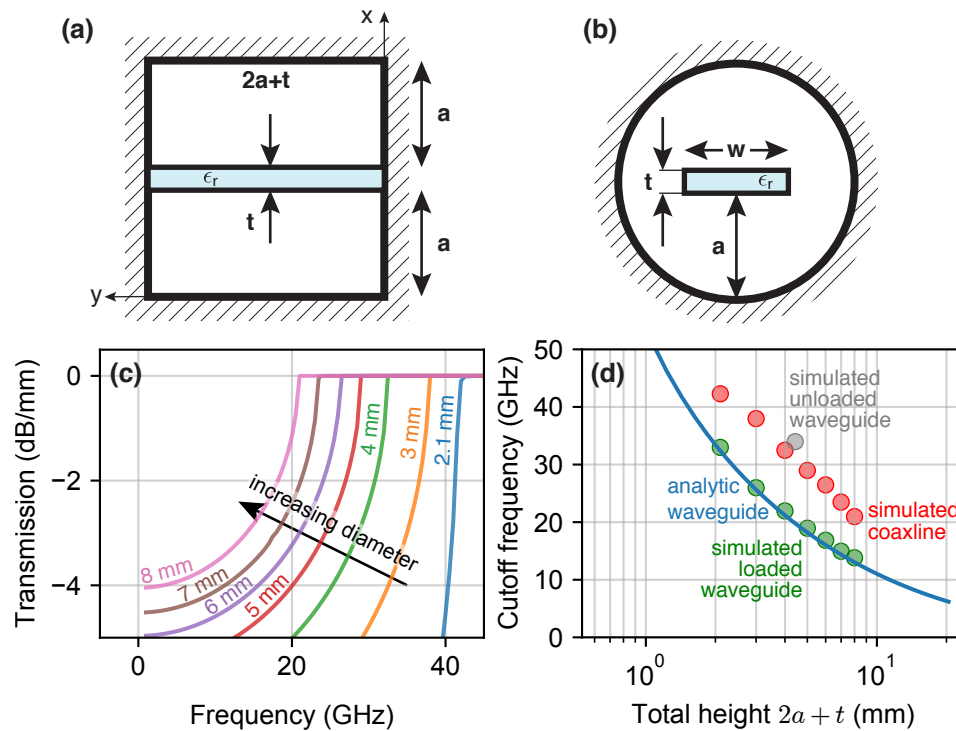


Figure 4.3 | **Behavior of a dielectric-loaded waveguide.** **a** An approximate, analytically solvable geometry representing the loaded coaxline enclosure. A chip of width w , thickness t , and relative dielectric constant ϵ_r is suspended by distance a from the top and bottom walls. The enclosure width is set equal to $2a + t$ so as to not limit the cutoff frequency. **b** A cross-section of the enclosure in Figure 4.2a in a section without the stripline. The enclosure diameter is given by $2a + t$. **c** The behavior of transmission parameter S_{21} per length of dielectric-loaded coaxline enclosure of different diameters. A thickness $t = 0.43$ mm and dielectric constant $\epsilon_r = 10$ are assumed. **d** Cutoff frequencies in different geometries, extracted in part from the data in **c**.




Figure 4.4 | **Coaxline cutoff frequencies.** Simulated coaxline cutoff frequencies in a practical configuration: a 4 mm diameter tunnel, with sapphire (left) or silicon (right) substrates of variable width and thickness. Commonly available thicknesses are simulated (sapphire: 0.43 and 0.53 mm; silicon: 0.275, 0.3, 0.33, 0.5, and 1.0 mm). Apparent discretization in the cutoff frequency is due to simulation resolution.

less is identified as the cutoff frequency. Values of the cutoff frequency for the analytical treatment as well as the two simulated geometries are shown in Figure 4.3d.

As a practical exercise, we also vary the width and thickness of the substrate and calculate the change in cutoff frequency for the coaxline (Figure 4.4). In this case, the cutoff frequency nears 20 GHz (a reasonable bound on the typical operational range) only for relatively thick substrates that nearly fill the tunnel.

Quasi-stripline resonators

The strip of metal patterned on the substrate (Figure 4.5a) is similar in placement and field profile to the inner conductor of a coaxial transmission line (Figure 4.5b), and also that of a traditional stripline (Figure 4.5c). It can be treated as a general transmission line with characteristic impedance Z_0 and the condition of an open circuit at either end, as represented by the circuit diagram in Figure 4.5d. A transmission line frozen in time appears to have the following voltage dependence on distance along it, z :

$$V(z) = V_0^+ \left(e^{-j\beta z} + \Gamma e^{j\beta z} \right). \quad (4.1)$$

where the reflection coefficient $\Gamma = 1$ in our case of an open circuit.

The first resonant mode ($\lambda/2$) that can be supported by such a structure will occur at $\beta l = \pi$, and

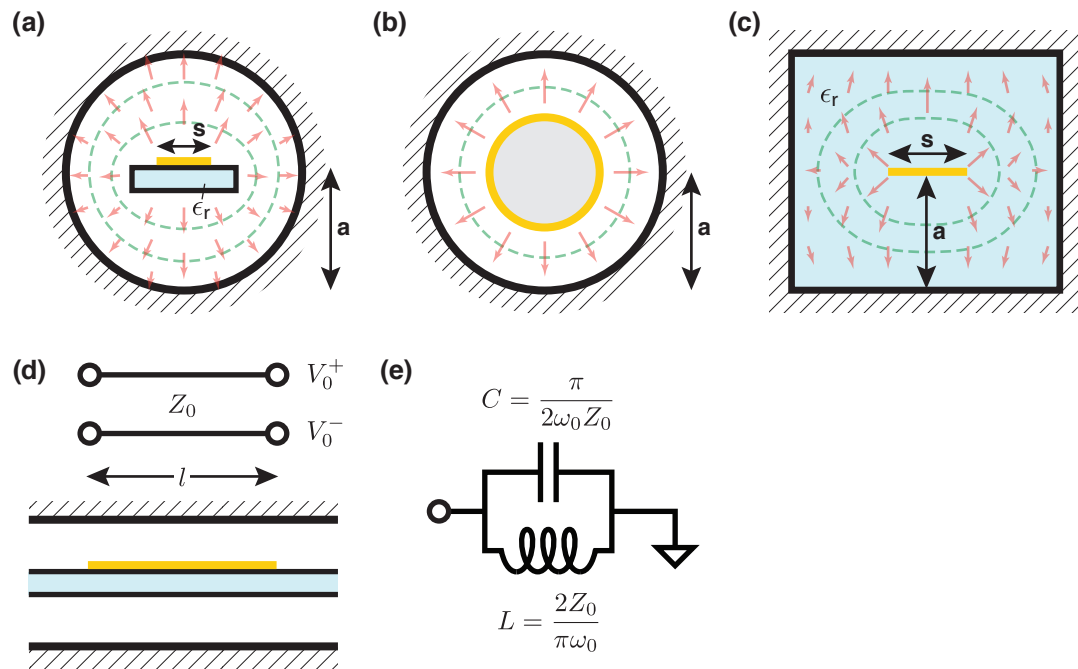


Figure 4.5 | **Coaxline as a quasi-stripline geometry.** **a** A cross-section of the coaxline in a section where the stripline is present. The stripline has width s , and is simulated as an infinitesimally thin sheet. The \vec{E} field vectors (red) and contours of \vec{H} (green, dashed) are also plotted. **b** A coaxial line bears strong resemblance to the coaxline, as does **c** a traditional stripline geometry. **d** The stripline and coaxline can be modeled as a transmission line, with impedance Z_0 and length l , terminated in an open circuit. **e** The equivalent circuit model for the transmission line at its resonant frequency ω_0 .

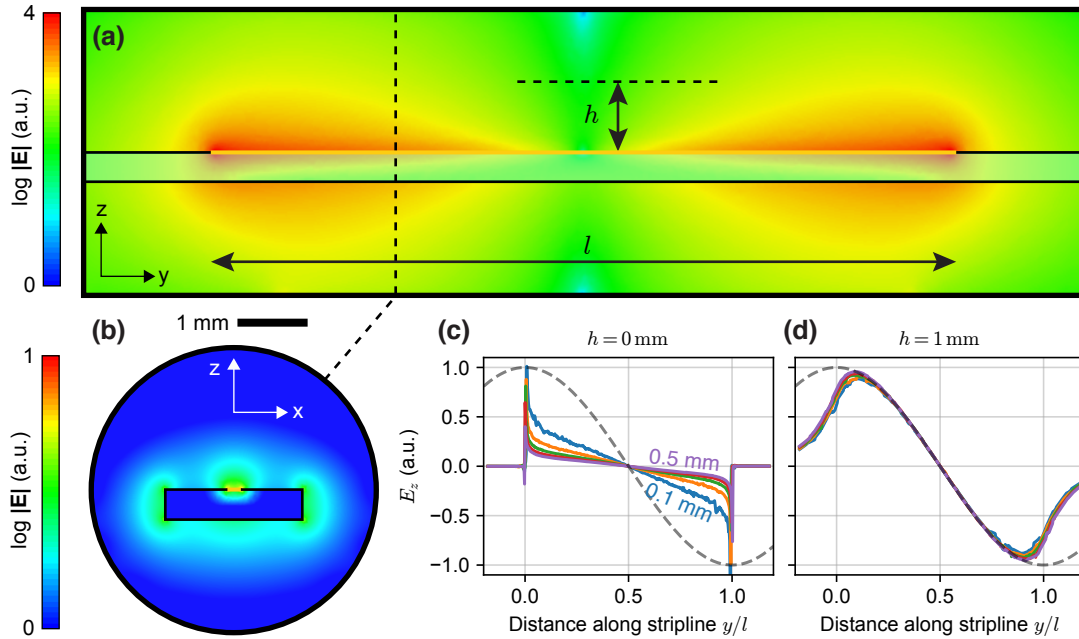


Figure 4.6 | **Coaxline electric field profile.** Simulated electric field magnitude on resonance for **a** a lengthwise cross-section, and **b** a crosswise cross-section at a location $\lambda/8$ from the end of the stripline. **c–d:** Simulated transverse electric field along the length of a coaxline-style 10.7 mm-long stripline in a 4 mm-diameter tunnel, with varied stripline width (0.1–0.5 mm in 0.1 mm steps) on a sapphire chip with 1.95 mm x 0.43 mm cross-section. The electric field is calculated at a height h directly above the center of the stripline **c** or 1.0 mm above the stripline **d**. Electric field measured very close to the stripline is sensitive to width and edge-dependent effects, and does not follow the expected cosine behavior (dashed line). At the position roughly halfway between the substrate and the outer wall, the behavior is width-independent and matches the cosine expectation to a much greater degree.

the voltage profile along the line can thus be expressed (in terms of the amplitude V_0^+)

$$\frac{V(z)}{2V_0^+} = \cos(\pi z/l), \quad (4.2)$$

showing that the voltage, runs over half of a cosine period from end to end. This voltage will be approximately proportional to the magnitude of the radial electric field in our quasi-stripline, quasi-coaxial geometry, which is simulated throughout the system in Figure 4.6a,b. In reality, the exact geometry of the system will perturb this ideal behavior (Figure 4.6c,d).

Characteristic impedance (Z_0) is an important parameter for describing a transmission channel. Here, for one, it will determine how easily the coaxline couples to the external 50Ω environment. The character-

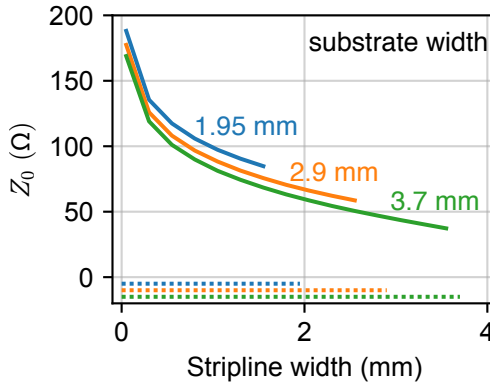


Figure 4.7 | **Coaxline impedance.** Characteristic impedance of a 4 mm diameter coaxline as a function of stripline width s , for three chosen values of substrate width w (their lengths represented by dashed lines).

istic impedance of the coaxline depends strongly on the stripline width s , and somewhat less strongly on the substrate width w (Figure 4.7). Compared to a coaxial line, it is quite difficult to achieve $Z_0 = 50 \Omega$ except at large s and w .

Hygienic coupling

In coupling to this device, we intend to preserve its seamless, hygienic nature. Inspired by the standard, robust coupling method used in 3D cavities, we introduce input and probe output signals using two evanescently-coupled pins within sub-cutoff waveguides that intersect the primary waveguide enclosure (Figure 4.2a). Pins are recessed to an adjustable depth within each coupling enclosure, located above each end of the stripline (Figure 4.8a). These provide variable capacitive coupling C_c to the resonant circuit, “loading” it externally. One or more pins can be arrayed along the extent of the device.

With one pin, reflection or feedline-coupled transmission (“hanger”) measurements are possible using a tee-style circuit (see Figure 4.8b inset). Hanger measurements, where circuit simplicity comes at the cost of abandoning some reflected power, can be transformed into a reflection measurement by replacing the unbalanced tee with a circulator or directional coupler. With two pins or more, transmission measurements can be taken directly between ports (Section 3.4.3). A typical measurement setup often includes two pins, one coupled weakly and one coupled strongly. The asymmetric nature of the coupling helps to naturally direct emission from the system. The weakly coupled port can be used to introduce control




Figure 4.8 | **Evanescent coupling to 3D enclosures.** **a** A diagram showing the coupling mechanism at each port in detail. An SMA flange connector transitions into a cylindrical pin within the coupling waveguide, emulating a $50\ \Omega$ coaxial cable. The pin extends so far into the tunnel as to leave a gap z up to the primary enclosure waveguide. **b** The distance z is varied *ex situ*, thus adjusting C_c , and the coupling quality factor of the measured fundamental resonant mode is extracted by measuring transmission between port 1 and 2. Several diameters of the primary enclosure are tested, all for 12 mm long striplines. Fits (dashed) have only one free parameter; their slopes are fixed to the calculated propagation constant. (Figure adapted from [7]; see [Copyright Permissions](#).)

signals, while the strongly coupled port can be connected to the output amplification-and-measurement chain. This approach yields predictable couplings that can be varied over a wide dynamic range without compromising package integrity.

As a first check, we can demonstrate that external coupling can be made arbitrarily weak. By varying the coupling attenuation distance and measuring Q_c , we see good agreement with the expected exponential scaling over six decades, with no observed upper limit (Figure 4.8b). Due to the natural waveguide decay, this large dynamic range in coupling strength can be achieved simply by modifying the length of a pin. This flexibility is a powerful tool when integration to multiple elements is required: we can achieve very strong coupling ($Q_c \sim 10^3$) to some elements (used for measurement or readout), while at the same time using weak couplings ($Q_c \sim 10^8$) to excite and control long-lived memory elements.

We must also be certain that our coupling pins, which are composed of normal metal for convenience, do not induce a large amount of dissipation. Figure 4.9 simulates this effect for copper coupling pins. The ratio of dissipative coupling strength to external coupling strength, given by $1/(Q_d/Q_{\text{tot}}) = Q_c/(Q_d + Q_c)$, reaches a maximum of 0.3% in the case of strong coupling. In weakly coupled cases, internal dissipation is observed to limit the system. Therefore, coupling strength Q_c can safely be used as a

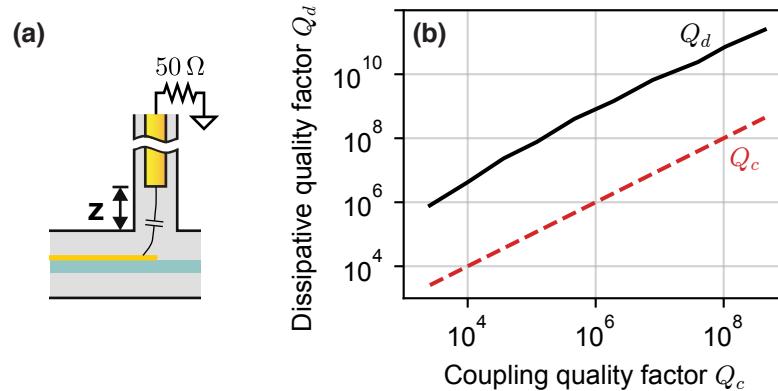


Figure 4.9 | **Effect of a dissipative coupler.** **a** The height of the pin z is adjusted to change its capacitance to the stripline and, indirectly, the coupling quality factor Q_c . Q_c can be measured by simulating the coupler port to introduce the only dissipation in the system, the $50\ \Omega$ characteristic impedance of the transmission line. **b** The quality factor Q_d caused by dissipation in a copper coupling pin is simulated for varied pin insertion. Its effect on the total quality factor is more than two orders of magnitude reduced as compared to the coupling Q_c . (Figure adapted from [7]; see [Copyright Permissions](#).)

proxy for a limiting effect from the coupler apparatus on the total quality factor Q_{tot} .

Just as in the coupling tunnels, we have restricted wave propagation in the primary enclosure to seamless waveguides with cutoff frequencies far greater than operating frequencies. If the system has been properly designed, then there should exist strong isolation between elements that is in agreement with simulated values. If no other modes exist, then, in addition to reliable couplings, the background spectrum should reveal no unexpected features. The transmission parameter S_{21} can be measured over a wide frequency range using a vector network analyzer (VNA), thus capturing both the response of the resonator as well as the background spectrum (Figure 4.10). It reveals a clean resonance and background ripples isolated by at least 60 dB.

Even though the enclosure body is seam-free, we can evaluate whether seams at each end introduce dissipation, spoiling our otherwise hygienic system. We predict their effect using a model of seam loss as a distributed admittance (Section 4.2.1) applied to simulation. The simulation places a conservative bound, $Q_i \geq 10^8$, on typical designs removed at least 7 mm from the waveguide end. These simulations assume a conservative value for seam conductance g_{seam} , and so it is likely that seams have an even smaller effect in reality than predicted. Since typical devices see significantly greater isolation from end seams than the 7 mm in this case, seams seem unlikely to affect performance in devices of this form. Their modeling, at




Figure 4.10 | **Coaxline transmission measurement.** **a** The transmission-type coupling scheme used to measure resonator transmission from port 1 to port 2. **b** Measurements show only the expected resonant modes, here at 7.7 GHz and 15.5 GHz (insets have 300 kHz span). The fundamental mode sees isolation of > 60 dB. The noise floor is due to the frequency dependence of the readout system and amplifier noise. (Figure adapted from [7]; see [Copyright Permissions](#).)

the very least, establishes a known bound.

Performance metrics

By enclosing chip-based circuit elements in such a package, the resulting coaxial-transmission-line-like device can be highly coherent and extendable. Addressing a host of likely losses within the package, including coupling, seams, and materials, we hope to produce chip-based element single photon relaxation rates at the level of the state of the art ($\sim 50 \mu\text{s}$). Good performance will also be important if the package is to be used as a design testbed (Chapter 5).

To verify that these kinds of devices are suited to such complex experiments and further integration, we measure the coherence times of the resonators within. For resonators, this requires achieving a high internal quality factor (Q_i) at sufficiently weak coupling (high coupling quality factor Q_c). This weak coupling serves two purposes: reducing experimental error, since Q_i is extracted from a measurement of Q_{tot} , and demonstrating that these devices can be measured in a practically useful regime, for use as quantum memories (Section 3.4.1). We measure these parameters by cooling coaxline resonators to ~ 20 mK and exciting circulating powers between ~ 1 and $\sim 10^6$ photons. The devices are connected in a feedline-coupled configuration and transmission (S_{21}) is measured using a VNA. Coupling parameters

are extracted from fits to S_{21} (see Section 3.4.3). Measurements are usually performed in an undercoupled configuration ($Q_c \geq Q_i$, with Q_c up to 10^9) so that the total quality factor Q ($1/Q = 1/Q_i + 1/Q_c$) is a direct measure of the internal losses.

We observe Q_i as high as $(8.0 \pm 0.5) \times 10^6$ at single-photon power³. In designs where the lifetime of an undercoupled stripline resonator is measured using a partner transmon (Section 4.3.1), lifetimes up to $T_1 = 250 \mu\text{s}$, or an equivalent $Q_i = 11.2$ million, can be reached. This suggests that quality factors in lithographic devices are not solely dependent on materials or fabrication methods, but are also affected by package contributions. This hygienic package evidently sets new, higher limits on the performance of its contents. Therefore, coaxline devices are well-positioned to act as a testbed for alternative loss mechanisms. Further studies are performed in Chapter 5 to pinpoint the dominant sources of loss. Regardless, the coherence levels achieved here situate the coaxline as a promising element for modular integration. Section 4.3.1 will demonstrate how this integration can be useful.

Space-saving modifications

The quasi-stripline design can be modified to reduce the footprint of a resonator, foremost increasing its prospects for integration or scaling. To do so, the path of the stripline is shortened and bent into a meandering, serpentine path (Figure 4.11a). For reasons that will be examined, simulations show that these modifications can leave impedance and frequency relatively unchanged as the length is shortened (Figure 4.11b). Since this takes the device towards the parameter regime of lumped-element planar resonators [108], energy density increases (Section 5.5). The practical result is that, while space is conserved, the magnitude of dissipation due to dielectric and conductor imperfections (discussed in greater detail in Chapter 5) is increased. Therefore, this approach may only be desirable if the resonator is used in a strongly coupled configuration. Furthermore, the resonance behavior becomes more difficult to analytically predict and intuitively understand as the extremity of meandering increases.

We can understand the effects of this design by examining the nature of the transmission line resonator. The stripline is characterized by a characteristic impedance, which on resonance is given by a capacitance

³. The best reported Q_i at the time of writing for lithographically-defined aluminum-on-sapphire resonators fabricated under similar conditions (e-beam evaporation, no substrate annealing) has been ~ 1 million [109].

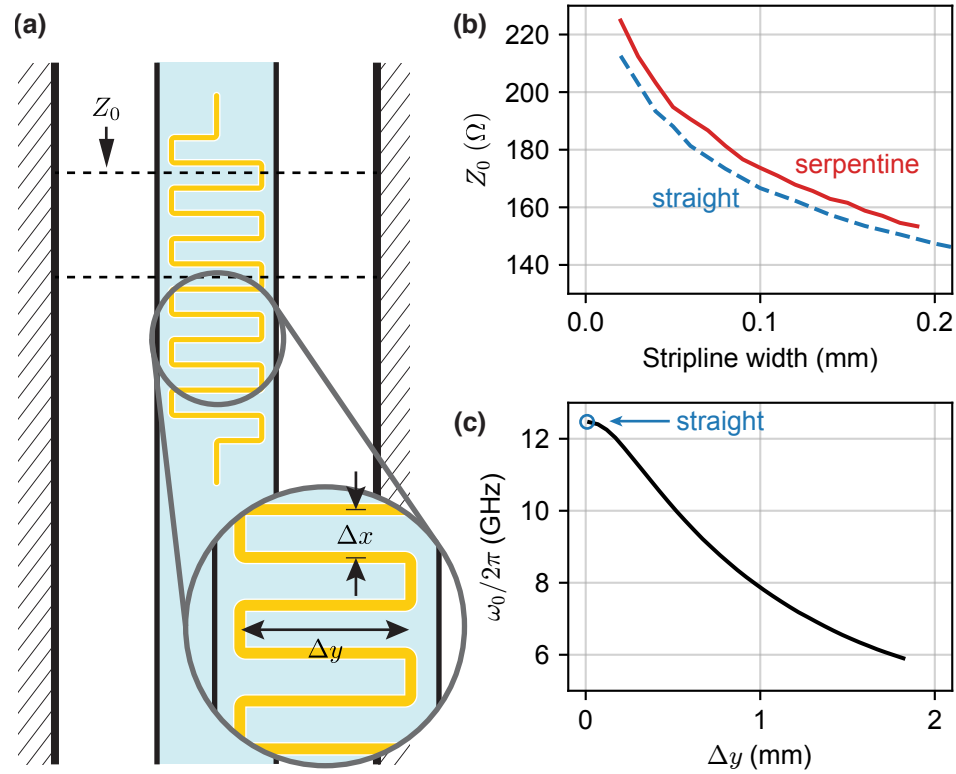


Figure 4.11 | **Serpentine stripline properties.** **a** The structure of a serpentine stripline transmission line resonator. We probe the characteristic impedance Z_0 from a simulated port along a segment of the resonator (dashed lines), simulating a segment of transmission line. **b** Impedance is calculated for a serpentine design of varied conductor width. The distance between strips Δx is kept at a constant 0.24 mm, and the structure rests on a sapphire substrate with 1.95 mm x 0.43 mm cross section in a 4.25 mm diameter enclosure. The serpentine line ($\Delta y = 1.4$ mm) becomes straight as $\Delta y \rightarrow 0$, producing a relative change in impedance $\leq 2\%$. **c** The frequency of a serpentine stripline resonator is calculated for fixed length (6 mm), fixed Δx (0.39 mm), and varied Δy .

and inductance per unit length as $Z_0 = \sqrt{L/C}$. As a straight length of fixed-width transmission line is raveled into a serpentine shape, the inductance per unit length will increase. (Here, it is important to specify that “length” refers to linear distance along the substrate, whereas “arc length” will refer to distance along the winding metal path.) The capacitance per unit length will also increase. The increases both contribute to lowering the resonant frequency ω_0 (Figure 4.11c). Compared to a straight stripline, this will require a shorter linear distance and produce a more spatially confined mode.

Just as in the straight stripline, the conductor width should affect impedance but not resonant frequency. This structure has a strong likeness to a coaxial transmission line, which has $L \propto \ln(b/a)$ and $C \propto 1/\ln(b/a)$ (where a and b are in the inner and outer conductor diameters). Changing the stripline width is roughly equivalent to changing the inner diameter of the coaxial transmission line, a . The dependence on $\ln(b/a)$ cancels in the expression for the resonant frequency, $\omega_0 = 1/\sqrt{LC}$; ω_0 should not depend (to first order) on a . The impedance, however, is expected to vary as $Z_0 \propto \ln^2(b/a)$, and it follows the proper behavior in Figure 4.11.

Coupling sensitivity Besides making resonators more compact, this modification can be used in hybrid form to tailor the sensitivity of resonator coupling relative to an element positioned along the enclosure. Figure 4.12 compares a straight and hybrid straight-serpentine geometry. The hybrid geometry produces a change in the gradient of the electric field midway along the extent of the structure. To satisfy the resonance boundary condition, the field varies more significantly in the serpentine region, while remaining nearly constant over the straight region. In fact, the \vec{E} -field variation occurs over a 35% longer distance compared to the straight stripline, even though the hybrid stripline is shorter overall. If a pin were inserted into the enclosure at a particular position, intended to couple to the stripline with a particular strength, this strength would be more readily achieved using the hybrid geometry (given some inevitable imprecision in aligning the stripline and pin).

These benefits come with a trade-off: the possibility of increased dissipation (Section 5.5). In addition to the capacitance and inductance between the stripline conductor and the ground plane, there also exists some “self-capacitance” (and inductance) that arises from the proximity of charges and currents between distributed segments of the stripline. The ac driving field redistributes charges along the conductor

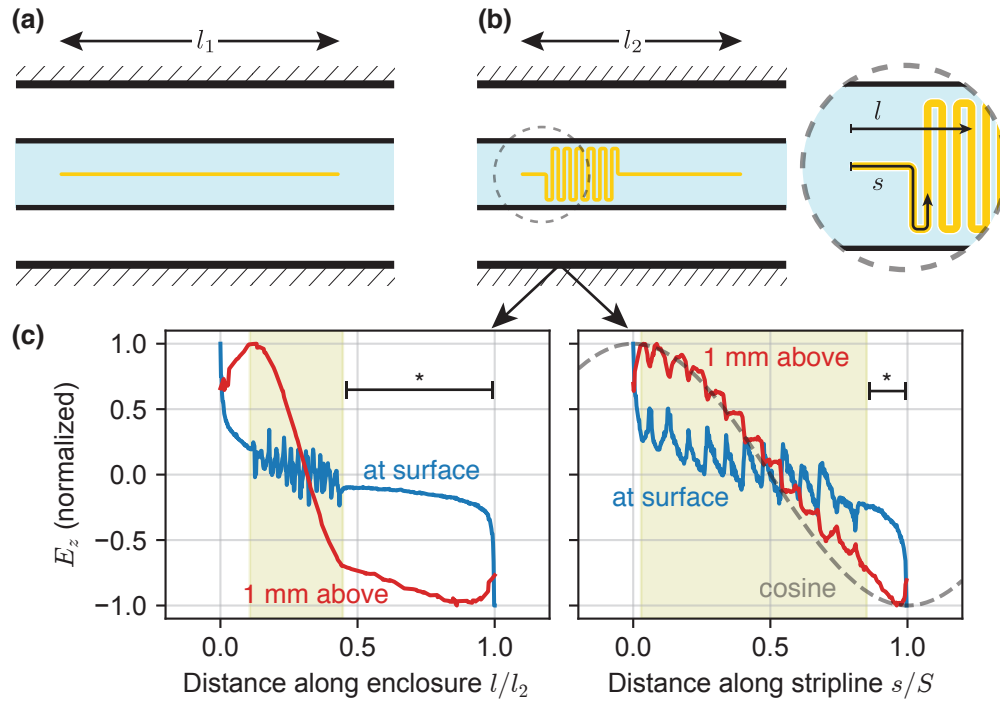


Figure 4.12 | **Serpentine stripline electric field profile.** **a** A straight stripline with length $l_1 = 10.7$ mm and resonant frequency $\omega_0/(2\pi) = 7.8$ GHz in a 4.25 mm diameter enclosure on a 1.95 mm x 0.43 mm sapphire substrate. **b** A hybrid resonator with $l_2 = 6$ mm (4 mm straight and 2 mm serpentine) has the same resonant frequency. Distance can be measured linearly (l) or as arc length along the stripline (s). The stripline has total arc length S from end to end. **c** The radial component of the electric field is simulated immediately above (blue) and 1 mm above (red) the hybrid resonator structure as a function of ‘linear’ length (left) or ‘unraveled’ arc length (right). The serpentine region is shaded. The cosine behavior expected for a straight stripline is plotted for comparison on the right. The large straight region (*) has 3.4 mm length and slowly changing E_z . By comparison, E_z changes by this factor in the straight stripline over just 2.5 mm. The ripples in the right panel reflect the non-uniform distribution of field across the cross-section of a segment of serpentine stripline.

throughout the cycle of a resonance. In a narrow tunnel with a long stripline, charges from one end of the stripline are strongly screened from the other end, so the capacitance to ground will dominate. As the stripline becomes serpentine, however, the relevant distance between unscreened areas of the stripline decreases. This effect can also be thought to contribute to increased participation ratio, as \vec{E} and \vec{H} fields will be more concentrated between areas of the stripline, lingering with greater energy density in the substrate and its surface layers, and less in the vacuum between the stripline and the ground plane. Simulations of the form discussed in Chapter 5 can reveal the likely magnitude of this effect, providing design guidance.

Transmon integration

One benefit of this coaxline design is that resonators, qubits, and even filters can be fabricated on a single chip. Consistent with the waveguide design form, we find that coupling between these elements can be engineered, well-controlled, and adjusted by changes in circuit parameters set with lithographic precision.

In the first step of this integration, we pattern a transmon qubit alongside a resonator (Figure 4.13a) [143]. We characterize the system’s coherence and compare measured parameters to simulation. We control the qubit using a weakly coupled input port—the same port used to drive the resonator. The system is read out via a resonator transmission measurement, using the asymmetric coupling setup described in Section 4.3.1. Transmons patterned and coupled in this fashion exhibit 30–80 μs lifetimes, near to the state-of-the-art values for transmon T_1 ’s. This is equivalent to quality factors as high as 3 million (by $Q = \omega T_1$), not far from those of the resonators measured in Section 4.3.1.

These coaxline resonator-transmon systems are generally well-behaved. Undercoupled resonators were observed to produce equally high Q_i with and without transmons. Important system parameters, such as mode frequencies and qubit anharmonicity, are found to agree well with predictions from finite-element simulations of the design. Limitations thereof tend to arise from inaccuracy in assembly, rather than in fabrication (Section 4.3.1); this limitations can be overcome with greater care. All told, this increase in complexity does not appear to reduce coherence values or parameter control.

In particular, with a lithographic resonator-qubit system, one can test whether on-chip element coupling obeys the same waveguide-attenuated behavior as external coupling. We expect that by varying the distance z between element ends, the chip enclosure should exponentially attenuate electric field $|\vec{E}| \propto$

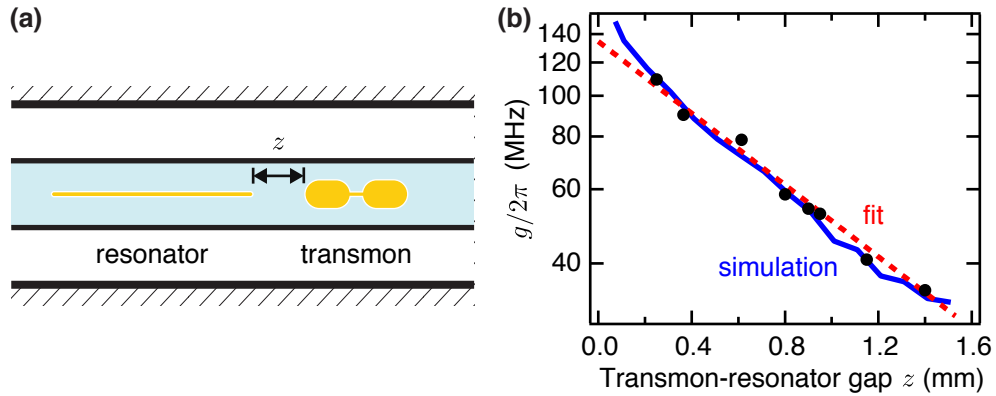


Figure 4.13 | **Coaxline–transmon coupling.** **a** The resonator and transmon are patterned and fabricated on the same substrate. At a fixed frequency detuning, their lithographically determined linear separation z controls the coupling strength. **b** The coupling strength is measured for transmon-resonator systems fabricated with different separations z . The data are fit using an exponential function $Ae^{-\beta z}$ with single free parameter A . The calculated attenuation constant, 8.5 dB/mm, comes from simulation (Figure 4.3c) of a 2.8 mm-diameter waveguide with bare substrate. A full system finite-element simulation using black-box quantization predicts similar scaling. (Figure adapted from [7]; see [Copyright Permissions](#).)

$e^{-\beta z}$ as it does for the pin couplers in Section 4.3.1. The resonator-qubit dispersive shift χ should scale as $\chi \sim |\vec{E}|^2$. Accuracy of coupling parameters like this one is important in nearly every aspect of system control.

To test this scaling, chips with different separations z were prepared. However, since detuning and physical separation both affect coupling, different resonator-qubit detunings Δ between samples makes direct comparison difficult. To relate them consistently to z , we calculate an effective coupling g defined by the relation $\chi = 2g^2/\Delta$, related to the detuned Jaynes-Cummings model [64]. When z is varied experimentally, we find that the measured change in g is consistent with a calculated waveguide attenuation scale length $1/\beta \approx 1.02$ mm, as well as with simulations (Figure 4.13b). This suggests that no unintended coupling is present, and that reasonably small separations between elements can produce a range of qubit-resonator couplings useful for typical cQED applications.

The lifetimes of resonators and transmons in this system can be understood by examining the spatial participation of each mode in dissipative dielectrics and conductors, discussed further in Section 5.1. The large resonator mode volume dilutes lossy material participation—the same effect that increases coherence of 3D cavities relative to traditional planar circuits (Chapter 5). By measuring resonators in waveguide

enclosures with different diameters, we find that higher Q_i generally corresponds with larger diameter (Section 5.4). This scaling behavior is consistent with loss originating from waveguide surface resistance, a waveguide dielectric layer, or on-chip dielectric layers, but does not distinguish between these mechanisms.

Bandpass Purcell filters

The Purcell effect, in the context of circuit or cavity QED, modifies the rate of spontaneous emission from one mode via another. The practical importance of a filter is easiest to understand with an example that demonstrates the Purcell effect in a typical transmon-resonator system. Imagine that two modes are present, a transmon (mode \hat{c}) and a readout resonator (mode \hat{b}). The transmon and readout resonator are detuned by Δ and are coupled with a dispersive coupling strength $\chi_{bc} = 5$ MHz. To efficiently read out information about the mode \hat{c} via mode \hat{b} , an output coupling rate $\kappa_{\text{out}} \sim 5$ MHz from the mode \hat{b} is then desired. Because the two modes are coupled, the Purcell effect will limit the lifetime of \hat{c} from the strong environmental, nominally 50Ω -coupling of \hat{b} . This adds a spontaneous emission channel in parallel (for the single, fundamental mode of \hat{b}) with rate

$$\gamma_1^{\text{Pur}} \approx \frac{\kappa_{\text{out}} \chi_{bc}}{2\Delta}. \quad (4.3)$$

More accurately, the mode \hat{c} with capacitance C is coupled to a generalized admittance $Y(\omega)$ that captures the behavior of all of the modes of \hat{b} to which it couples:

$$\gamma_1^{\text{Pur}} = \frac{\text{Re}[Y_{\text{in}}(\omega)]}{C}. \quad (4.4)$$

Typical frequencies might produce a 1 GHz detuning, thus leading to a limiting $T_1 = 2\pi/\gamma_1^{\text{Pur}} = 80 \mu\text{s}$ from the single mode alone. This is a common transmon lifetime in the absence of a strong Purcell effect; independent losses in parallel would then produce at best a $40 \mu\text{s}$ lifetime. Reducing coupling or increasing detuning is one way to mitigate the Purcell effect, but may have unwanted side effects. One reasonable response, then, is to add a Purcell filter with the goal of maintaining coupling strength and reducing spontaneous emission due to the Purcell effect.

So far, we have established that the enclosure provides good isolation over a wide band of frequencies

beneath its cutoff frequency. Stripline elements interrupt this isolation, allowing resonant modes and a region where the field decay is halted. By strategically patterning particular circuit elements, these concepts can be used to control the spectrum to which elements couple, thus engineering filters. The design of filters can be guided by calculation of the multi-mode Purcell effect using either lumped element circuit models or finite element simulations.

In planar circuits architectures, Purcell filters of various types are frequently used to isolate, for example, the transmon from a strongly coupled readout cavity mode [144, 145, 146]. In a $\lambda/4$ -type filter, for example, a suppressed mode can be positioned a distance $\lambda/4$ (referring to that mode’s wavelength) from a lossy mode, thus creating an impedance transformation that causes that mode to appear significantly isolated. In 3D cavities, waveguide segments can act as high-pass Purcell filters. In the quasi-planar coaxline structure, the geometry restricts the kinds of filters that can be used for this purpose. Difficult access to a “ground plane”, for example, rules out the possibility of $\lambda/4$ filters. A natural choice for a filter design is that of a bandpass filter.

A series bandpass stripline Purcell filter is a natural choice in our geometry. A single-pole bandpass filter possesses the same transfer function and behavior as our $\lambda/2$ stripline resonators. Inserting a second resonator (mode \hat{a}) in series with the other modes acts like a frequency-dependent transmission line, extending the spatial propagation of modes at that frequency, while allowing the waveguide isolation to persist at a relatively high level at other frequencies.

Figure 4.14 shows this circuit arranged in the coaxline architecture, using stripline resonators. Using finite-element simulations, we can adjust coupling strengths by varying the gaps and the filter detuning (via filter length). Changing the gap z_{ab} adjusts the linear coupling between the resonators g_{ab} . Varying the resonator–transmon gap z_{bc} adjusts the dispersive coupling, which has rate χ_{bc} . A combination of both gaps, as well as frequencies of the elements, will set the transmon’s dispersive shift from the filter, χ_{ab} . More filter elements (a repetition of the boxed elements in Figure 4.14a) further isolate the transmon. By changing all of these filter parameters, we can manipulate and probe the protective behavior of the bandpass Purcell filter element.

It is helpful to divide our analysis into two regimes of filter parameters. In one, the filter frequency $f_a = \omega_a/2\pi$ is adjusted such that it hybridizes strongly with the readout mode ($f_a = f_b$). Here, the

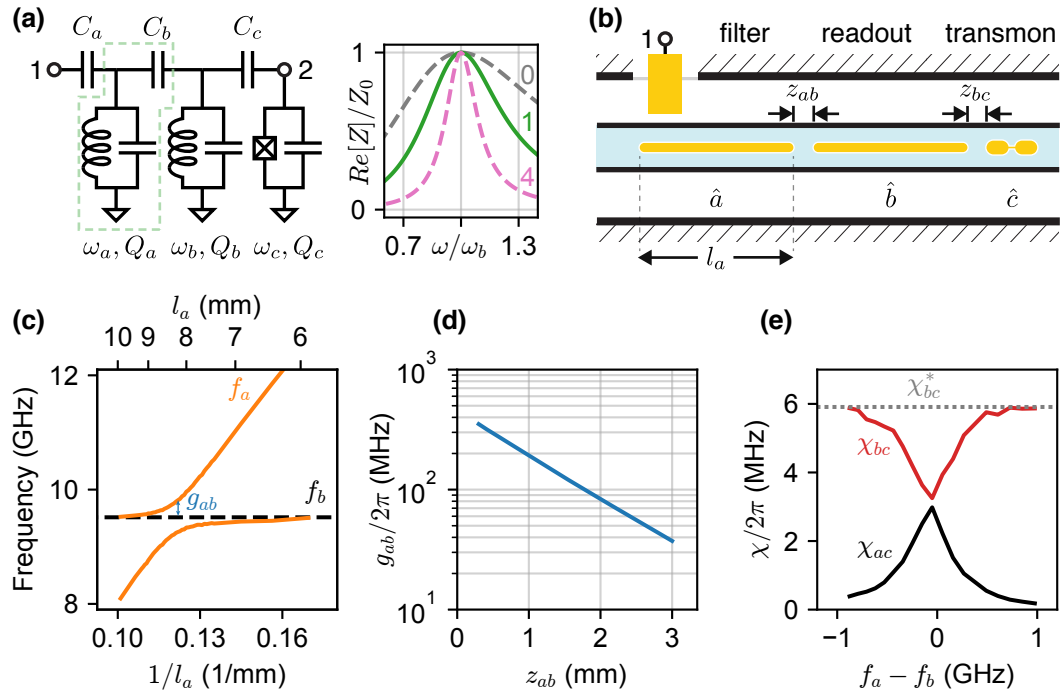


Figure 4.14 | **Coaxline Purcell filter.** **a** Circuit schematic for a transmon–readout system with added Purcell filter (dashed box). Couplings are determined by impedance seen from virtual port 2 to the coupling port 1 in a reflection measurement configuration, which is shown (normalized) for zero, one, and four-pole filters. **b** The circuit is realized with a series bandpass stripline filter in the coaxline architecture. Capacitive couplings C_b (C_c) are set by adjusting the distances z_{ab} (z_{bc}) in lithography. **c** Filter length l_a is swept in simulation; an avoided crossing reveals resonator coupling g_{ab} set by z_{ab} (shown in **d**). **e** For one g_{ab} , dispersive couplings between transmon and filter (χ_{ac}) or readout (χ_{bc}) show strong hybridization near $f_a = f_b$. With filter removed, $\chi = \chi_{bc}^*$.

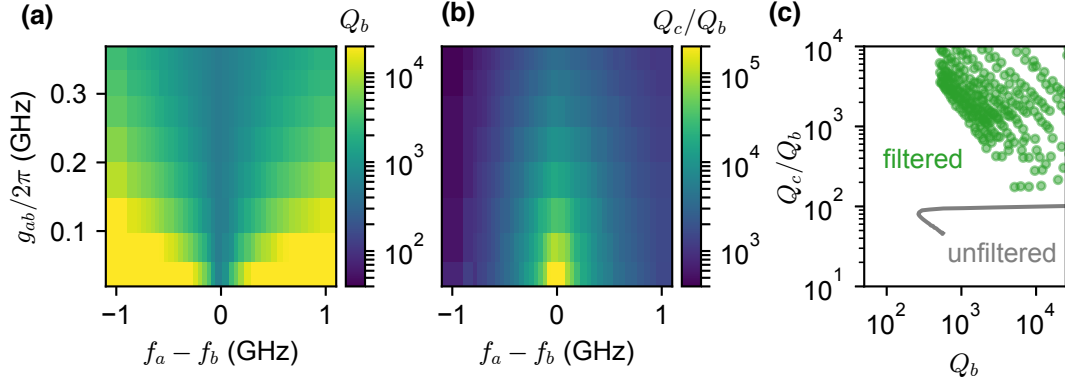


Figure 4.15 | **Purcell filter Q-ratio** (a–b) Filter parameters varied in Figure 4.14 alter readout coupling (Q_b) and transmon–readout Q -ratio (Q_c/Q_b). c Higher Q -ratios can be obtained in the filtered case (at practical values of Q_b).

hybridization extends to produce equal coupling quality factors (Q_a, Q_b) and dispersive coupling rates (χ_{ab}, χ_{bc}) to the transmon. This arises from the geometry: each degenerate mode shows the same \vec{E} -field behavior at the ends of the joint structure, which largely determines these coupling strengths. The assignment of ‘readout’ and ‘filter’ to each mode is blurred; either can be used to read out the transmon.

Importantly, this regime can also produce the lowest Q_b and highest Q -ratio (Figure 4.15). This Q -ratio represents the protection of the transmon with respect to an output coupling strength. Decoupling by increasing C_a by moving structures away from the coupling pin has the effect of increasing all Q values without adjust their ratio. Adjusting the filter parameters, however, can increase Q of the transmon mode while maintaining the strength of the readout coupling. From ref. [147], the expected suppression Q -ratio should be

$$\frac{Q_c}{Q_b} = \frac{1 + [2(\omega_c - \omega_a)/\kappa_a]^2}{1 + [2(\omega_b - \omega_a)/\kappa_a]^2}. \quad (4.5)$$

We find ratios that qualitatively match this expected behavior.

This behavior is most easily understood from the perspective of the transmon, which in this case is well detuned from both modes. At this transmon frequency, the act of adding a filter resonator in-line with readout resonator mostly serves to add many factors of waveguide attenuation length, thus increasing isolation from the coupling port. At the readout frequency, the mode extends with minimal added attenuation across the additional length of the filter, maintaining the output coupling readout of the readout, κ_b .

For clarity in discriminating the modes and reduced sensitivity to parameters, it can also be useful to

operate in a detuned-filter regime. Here, the modes \hat{a} and \hat{b} are detuned slightly from one another to lift their degeneracy and avoid their full hybridization (Figure 4.14c). Conveniently, the notion of a ‘readout’ and ‘filter’ mode is clear. The output coupling rates of the filter and readout will be different, with κ_a generally ~ 10 times faster than κ_b . Similarly, the spatial separation and additional frequency detuning ensures that the dispersive coupling between the transmon and the readout mode will be much greater than for the filter mode (Figure 4.14e). Thus, the system is most efficiently read out at the readout frequency ω_b , as normal.

While the strong hybridization regime shows the highest Q -ratios and the fastest output coupling rates, other factors may cause one to prefer the detuned regime. In the hybrid regime, both modes convey equal information about the quantum state of the transmon. An equilibrium photon number population (\bar{n}) in either of the resonators could thus lead to effects similar to measurement-induced dephasing, which depends on \bar{n} , κ , and χ in a way that may differ between the regimes. In some designs with crowded frequency spectra, considerations of higher modes may further complicate this picture.

In practice, Purcell filters have been implemented in both straight stripline and serpentine stripline geometries. The energy lifetimes of filtered systems appear to be protected from Purcell emission, as predicted by circuit and finite-element simulations. Tunable-frequency transmons, however—which allow circuit response to be probed as a function of frequency—would permit more systematic study of circuit behavior and system loss sources.

Coaxial stub cavity integration

Another powerful tool in the circuit QED architecture has been the 3D coaxial stub cavity. The stub cavity [46] shares features similar to the coaxline, namely, its seamless nature and the TEM resonant mode structure. It is thus naturally integrable; the coaxline and stub cavity enclosures can be formed within the same block of high-purity aluminum. The interface of the two intersecting enclosures will be seamless and straightforward. From a standpoint of microwave hygiene, the only new consideration (besides the added resonant modes) of this integrated system is that the stub cavity enclosure is generally wider than that of the coaxline, and will present a lower cutoff frequency. As a result, the stub cavity enclosure must be terminated (usually with a lid) far away from the resonant structures.

The Hamiltonian of the system consisting of the cavity (\hat{a}) and transmon (\hat{t}) is given by [102]

$$H = \hbar\omega_a\hat{a}^\dagger\hat{a} + \hbar\omega_t\hat{t}^\dagger\hat{t} - E_J \left(\cos(\hat{\varphi}) + \frac{\hat{\varphi}^2}{2} \right). \quad (4.6)$$

The flux across the junction is given by

$$\hat{\varphi} = \sum_{k=a,t} \varphi_k (\hat{k} + \hat{k}^\dagger), \quad (4.7)$$

where φ_k are magnitudes of the normalized zero-point fluctuations of flux across the junction, $\varphi_k \equiv \Phi_k/\varphi_0$, and Φ_k is given by $\sqrt{\hbar Z/2}$, where Z is the characteristic impedance of the \hat{k} in the black-box formalism [94]. Transforming into the rotating frame of all modes (Appendix A.2) and expanding the cosine to $\hat{\varphi}^4$ gives

$$H/\hbar = \sum_{k \neq l} \chi_{kl} \hat{k}^\dagger \hat{k} \hat{l}^\dagger \hat{l} + \sum_k \frac{\chi_{kk}}{2} \hat{k}^{\dagger 2} \hat{k}^2, \quad (4.8)$$

where we have introduced the cross-Kerr coefficient (or dispersive shift) $\chi_{kl} = -E_J \phi_k^2 \phi_l^2$ and the self-Kerr coefficient (or anharmonicity) $\chi_{kk} = -E_J \phi_k^4/2$. Examples of how these parameters can be measured can be found in refs. [35, 44].

Besides those in the Hamiltonian, another important parameter is the cavity's single-photon energy lifetime. Due to its extraordinary lifetime, the coaxial stub cavity typically fulfills a special role in the challenge of scaling: the need for a long-lifetime quantum memory (Section 3.4.1). The resulting basic integrated module combines a stripline resonator, transmon qubit, and this new stub cavity. The capacitor pads of the transmon are positioned to bridge the modes of the two structures (Figure 4.16). This allows the transmon to couple the modes easily, mediating their interaction (Section 3.4.2). The large transmon nonlinearity and linear enclosure attenuation will create much stronger χ coupling between adjacent modes than between next-nearest neighbors.

We characterize parameters of this significantly integrated system, including coupling and coherence values. Both the qubit and the high- Q cavity can perform well, with one system showing qubit $T_1 = 110 \mu\text{s}$, qubit Ramsey decay time $T_2^* = 40 \mu\text{s}$, cavity $T_1 = 2.8 \text{ ms}$, and cavity $T_2^* = 1.5 \text{ ms}$ ⁴. These

⁴ T_1 decay for a cavity is measured by preparing the single-photon Fock $|1\rangle$ state. The cavity T_2 is measured by performing a

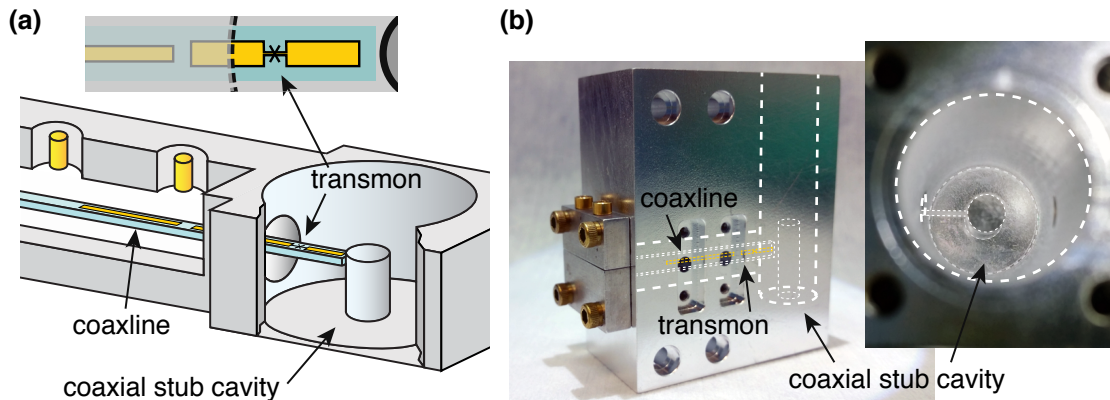


Figure 4.16 | **Coaxline integrated with coaxial stub cavity.** **a** A schematic of a coaxline enclosure terminating into a coaxial stub cavity. Here, a coaxline and transmon are clamped only at one end (not shown) so as to be a suspended bridge between the two structures. **b** Images of the entire device (left) and coaxial stub cavity (right) overlaid with an outline of the features in **a**. (Figure adapted from [7]; see [Copyright Permissions](#).)

qubit lifetimes are among the highest measured generally in 3D cavities, and in this case the coaxial stub resonator T_1 in the presence of the transmon was the same as its equivalent T_1 without (calculated from its internal quality factor). This suggests that no additional sources of dissipation are introduced when these elements are combined into a single, seamless package. Such integrated systems will form the basis of the experiments to follow in this work, particularly those in Chapters 6 and 7.

In the near-term, this kind of coupling allows for significantly more complex many-resonator, many-qubit circuits. When combined with more advanced techniques for fabricating 3D enclosures using lithography and multi-wafer bonding [141, 148], it even offers an attractive route towards long-term scaling. In the near term, however, one way in which scaling of complexity can be immediately beneficial is by increasing experimental throughput by creating multiplexed systems.

4.3.2 Multiplexed measurement with seamless designs

Experiments investigating loss often rely on varying many parameters over a wide range. Therefore, many samples must be tested, and even more to obtain statistically significant results. It is imperative, there-

Ramsey experiment with an effective $\pi/2$ rotation applied using a pulse that maps $|0\rangle \leftrightarrow (|0\rangle + |1\rangle)/\sqrt{2}$ within the one-photon Bloch sphere.




Figure 4.17 | **Multiplexed coaxline enclosure.** **a** A top-down view of multiple chips and enclosures in the same package. A common feedline accesses all of the enclosures. Additional enclosures would extend to the left and right; more enclosures and feedlines in a single package could also be stacked into a 2D array (in/out of the page). **b** A cross-sectional view from the side at the location indicated by the dashed arrow in **a**. **c** An equivalent circuit model of this “hanger” feedline configuration.

fore, that measurement techniques allowing for high sample throughput are used. Multiplexed sample measurement setups involving many samples within a single “package” not only permit high-throughput measurement, but also promote a common, controlled environment. Care must be taken, however, such that the effect of multiplexing does not introduce cross-talk or any means by which the quality of one sample can interfere with another.

Using the seamless coaxline architecture, we designed a system (the “coaxmux”) such that elements designed upon a single chip could be measured in a multiplexed fashion. Each enclosure for a single device appears nearly identical to that in Section 4.3.1, albeit with a change in the coupler design. In this multiplexed case, the seamless enclosures are repeated at regular intervals. A common feedline intersects all of the enclosures (Figure 4.17). This feedline is made from a metal wire (which can be high-purity superconductor), suspended in “vacuum dielectric”, forming what is essentially a $50\ \Omega$ coaxial transmission line. The wire inserts into standard SMA flange couplers at either end. In the standard coaxmux, feedline reflections are generally suppressed beneath 20 dB, forming quite a good transmission line.

Coupling is determined by lithography and assembly. If chip placement were fixed and guaranteed, then the distance from the feedline (the position of which is not adjustable) to each element on every chip would determine its coupling. In reality, the means of assembly allow the axial position of the chip to be adjusted. A desired coupling is typically achieved by careful simulation and accurate assembly; after one measurement, the coupling can be altered by adjusting the axial position of the chip. This can be repeated in an iterative fashion as necessary. Precise assembly will forgo the need for iteration, but possibly reduce the flexibility in case a device needs to be measured with differing coupling strengths.

In order to use the same signal lines and hardware to read out multiple devices, we choose to multiplex them in frequency. We can do this for a transmon-readout system by patterning a wafer of chips with chips that are identical except for stripline readout resonators of slightly different lengths. The lithographic precision with which stripline resonant frequencies can be set means that many devices can be multiplexed in the same device. Since inter-transmon coupling is filtered by the readout modes in individual enclosures, it should be small as long as readout frequencies are sufficiently separated.

In a package with five devices, we established a bound on the inter-transmon coupling by applying a fractional π -pulse to one transmon, and then performing a Ramsey or spectroscopy experiment to determine the resulting frequency shift in a second transmon (Figure 4.18). We performed this experiment on all pairs of transmons within the five transmon system, bounding the coupling strength $g/(2\pi) \leq 1$ kHz. The modes most likely to strongly couple had readout frequencies at 8.75 and 8.78 GHz, with transmon frequencies at 6.52 and 5.81 GHz, respectively. The other three enclosures contained Purcell filters in series (Section 4.3.1), with other frequencies otherwise similar; thus, we expect additional isolation to protect those transmons and for the coupling to be more difficult to measure.

In this multiplexed architecture, assembly is more difficult because—due to space constraints—a common clamp holds all the chips together at each end of the package. This means that some chips may be held more firmly than others. As a result, vibration often plagues the system (Section 5.6.3). This assembly scheme can be altered, and vibration potentially suppressed, at the cost of adding loss: by splitting the cavity into two parts, end clamps can be replaced by internal clamps. Now, unfortunately, this introduces seams back into our seamless design.




Figure 4.18 | **Experimental scheme for testing cross-coupling.** A scheme for measuring the response of transmon B to the excitation of transmon A in a separate enclosure. Transmon A is excited and a Ramsey experiment is performed on transmon B, with a variable inter-pulse spacing τ . The detuning of the resulting signal indicates the frequency shift of the mode and is related to the coupling strength g between them.

The seammux: a compromise

The “seammux” is a two-piece design with a body and a lid, creating a seam that runs along the length of the enclosure (Figure 4.19). The chip is placed on raised pedestals near the end that are integral to the body. Identical pedestals extend from the lid. Soft indium wire of varying diameter is placed between the chip and one or both pedestals to apply variable pressure or accommodate variable chip thicknesses. Like the coaxmux, the pattern of enclosures can be repeated.

The lid can be secured to the body by fasteners along the length of and between enclosures, as well as around the border. This encourages flatness of the joint, and could provide more even pressure from multiple fastening points if compared to the coaxmux. Furthermore, the absence of external clamps means that fewer parts must be relied upon to maintain intended alignment and interface pressure through the course of thermal contraction. For these reasons, possibly, the symptoms of acoustic resonances are *not measurable* in devices in this package.⁵

The vibrational response of chips measured in the coaxline or coaxmux geometries is consistent with a mechanical flexural mode of the substrate (Section 5.6.3). The frequency of this mechanical resonance, typically ~ 10 kHz for our chips with dimension $0.5 \times 2 \times 40$ mm, analogous to the EM resonance of the stripline itself, is largely determined by the length of the chip. If the observed resonance is driven by its proximity to a particular driving frequency inherent to the experimental apparatus, such as the dilution

5. This claim is based on about forty measurements of 0.43 mm-thick, 1.5–1.95 mm-wide sapphire chips.




Figure 4.19 | **A coaxline variation with seams.** **a** The seammux is similar to the coaxmux except that the chip rests on internally constructed pedestals at each end of the enclosure instead of external clamps. **b** A cross-sectional cut at the location of the dashed arrow in **a**. A separate lid (red) forms the other half of the clamp. The seam formed by the lid joining the body is maximally distant from the chip, instead of in the middle. **c** An image of the device without a lid. A single unit is outlined in yellow.

refrigerator's turbomolecular pump, then by adjusting the length of the chip, the response should fall away sharply. When much shorter chips were measured in the coaxline geometry, the vibrational behavior did not significantly change (although the resonance frequency increased, as expected). Therefore, it is reasonable to hypothesize that the resonance is broadly driven, and is a feature of the coupling rather than a particularly unfortunate driving frequency within the apparatus. This is consistent with typical transmon chips in 3D rectangular cavities, which have a markedly similar clamping mechanism compared to the seammux. While these chips are much shorter, acoustic vibration has never been directly attributed to loss in these systems.

Care must be taken to avoid seam loss in the seammux, naturally. The model of seam loss discussed in Section 4.2.1 can be applied to this geometry. Simulations can place bounds on the lifetimes of elements within as a function of distance to the seam (Figure 4.20). Of note, however, is that the longest-lived elements measured in a coaxline-like geometry have come from undercoupled resonators measured in seammux packages. By carefully considering the effect of the seam on a typically seamless package de-

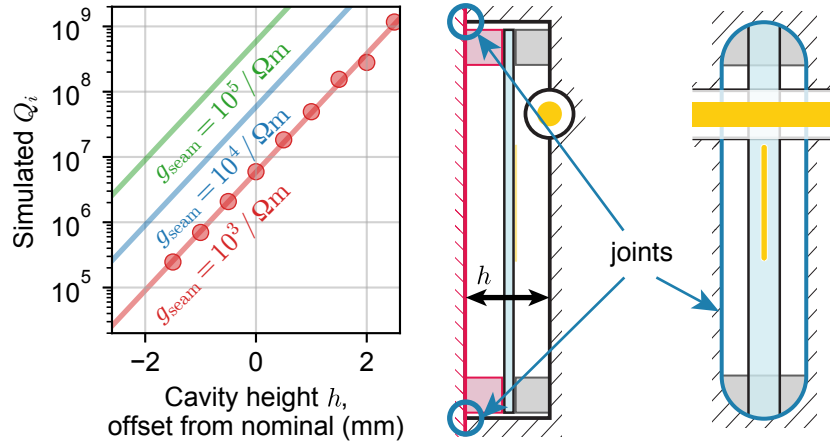


Figure 4.20 | **Seam loss in the seammux.** Simulation of seam current as the position of the seam is varied in simulation. The nominal height is 6.5 mm; the nominal width is 2.9 mm. The simulated Q_i data is shown for a conservative value of seam conductance (per unit length), $g_{\text{seam}} = 10^3 / \Omega\text{m}$. The fit line is shown for this g_{seam} as well as more optimistic values.

sign, the seammux can overcome the complexity of end-clamping and enable highly coherent integrated quasi-planar devices.

4.3.3 Scalable systems in circuit QED

Scalability, as discussed in Chapter 3, refers to the concept of increasing the computational power towards a quantum computer that is capable of performing more complex algorithms. This is accomplished by adding logical (and thereby physical) qubits, as well as the necessary error-correction mechanisms that protect these new elements and the added interactions and couplings they bring. In our hygienic systems, we hope to introduce this protection naturally, to a large extent.

Seamless modules The seamless or near-seamless package designs discussed so far have already demonstrated “clean” integrability, either by multiplexing through use of a common feedline, or combining many elements on a single chip. The feedline-coupled designs, which seek to isolate the separate systems, can also regulate interactions between them if replaced with a bus resonator or bus cavity (Figure 4.21). Single-chip designs also can form the basis of quantum modules, which would require (at the very least) transmon, readout, and memory capabilities. With a particular vision of module in mind, we can combine these tools to enable the necessary interactions.




Figure 4.21 | **Ideas for complex planar-3D integration.** There exist natural extensions of the scalability for systems described in this chapter. In the near term, at least, elements could be interleaved and multiplexed to produce a denser connectivity graph that is useful for many experiments in quantum information. A top view (a) and side view (b) of two multiplexed coaxline resonators that are coupled via a third, perpendicular transmon chip on another layer. This idea could be naturally extended to multiple interconnects, multiple layers, and multiple unit cells.

The relatively large coaxline enclosure size and rudimentary external wiring design sets an ultimate limit on the number of systems that can be interconnected. Particularly when integrated with 3D cavities, so as to enable, for example, a long-lived quantum memory, the size of an individual system will be large. Therefore, these systems are probably suitable for scaling in the near- to mid-term; modifications will probably need to be made to scale to thousands of devices.

Future modifications might not simply increase device density, or reduce element footprints, but also add features that assist in scaling. These capabilities could include the integration of on-chip amplifiers, enabling single-shot multiplexing. Switching, routing, or conversion elements could become critical aspects of devices that would serve as modules within a quantum network (Section 3.3).

The MMIQC A longer-term scaling solution will look similar but will probably be composed of different hardware—hardware that is not only lithographically patterned, but assembled with lithographic precision as well. One such candidate is the multilayer microwave integrated quantum circuit (MMIQC) architecture. The ability to create 3D cavities [140] and transmons [149] with reasonably high coherence has been demonstrated in this architecture. Using fabrication schemes similar to those used to create the 3D MMIQC cavities, trenches could be formed to act like coaxline enclosures. Circuits might be pat-

terned on thin membranes, minimizing dielectric participation and allowing large spatial extents of the mode fields.

A wafer-based solution to scaling like the MMIQC could also address the complexity of interconnection. Coupling pins and tunnels would take an analogous form in the MMIQC, but in a stripline architecture in trench enclosures. With proper seam loss mitigation and ground plane connections between layers, the MMIQC could offer a vastly greater level of complexity and integrability while affording the benefits of lithographic precision and assembly.

This vision for an ultimate scalable architecture serves as valuable motivation for designing, understanding, and making our near-to-mid-term coaxline-style prototype modules robust.

5

Identifying Dissipation Sources in cQED Modules

The previous chapter introduced new, high-performance devices that could be incorporated into a module that satisfies our vision for modular quantum computing. We also emphasized the importance of maintaining coherence, through careful engineering, as coaxline enclosures, striplines, and 3D cavities are combined within this modular architecture. One outstanding challenge in scaling the number of elements (and modules) further, then, will be to minimize the dissipation and thereby extend the coherence times of individual cavities and qubits.

In engineered quantum circuits such as ours, dissipation can take the form of the conversion of microwave electric field phonons into phonons in dielectric materials (heat and vibration), resistive heating in conductive materials, or by electrically or photonically coupling to external paths from which it cannot be recovered. Stray electric coupling can be modeled, simulated, and controlled by microwave engineering and common hygiene practices (many of which were discussed in Section 4.3.1). The presence of dielectric materials in an ac electric field, however, automatically introduces a means for dissipation via phonon excitation, and is difficult to isolate or avoid.

Some amount of dielectric material has so far proven to be necessary in cQED, serving as the support substrate on which circuits are patterned. In the presence of a dielectric substrate, thin surface layers will often also form in which electric field energy can reside and dissipate. Metals, whether deposited into thin layer or in bulk, are also necessary in circuit QED systems. These metals can either host oxides or impurities

that impede conduction.

One or more of these loss mechanisms will limit our systems at any point. That is to say that while all channels of loss contribute to some extent, one or a few will dominate and the others will be inconsequential—at least until those limited sources are removed, and one of the others begins to dominate. Only a system constructed of perfect dielectric (vacuum) and perfect superconducting walls (at 0 K) is fully immune from these losses.¹ Therefore, it is advantageous to design systems such that as much electric energy field energy resides in vacuum as is possible, and that currents are confined to high-quality materials. In order to move forward in the absence of perfect systems, these loss mechanisms must be understood and characterized. The relative spatial energy concentration in different materials and regions, called energy participation ratio, is a useful way to analyze these likely sources of loss.

Two classes of devices are of interest in this study. The first consists of 3D structures and cavities that are largely formed by traditional machining (Section 5.3). The second class consists of planar or quasi-planar devices that are characterized by thin, patterned metal layers on a dielectric substrate that is generally shaped by dicing or cleaving (Section 5.4). These layers are patterned by lithography and may undergo additional processing, such as treatment in solvents or by plasma etching. Both machining and lithography-related processes may alter the intrinsic properties of these materials. Natural exposure to an oxygen-rich environment may also alter interface properties through oxidation. Both of these types of device will be incorporated into a cQED module, and so it is important to understand the loss mechanisms that limit each.

In this chapter, I will introduce the formalism of participation ratios required to treat this problem (Section 5.1). Various forms of potential loss will be identified and studied. Testbed systems are needed to identify the materials and interfaces contributing to the largest dissipation in our cQED systems, as well as

1. A superconductor well below its critical temperature, T_c , is perfect only to a very good approximation, however. The surface resistance, and thus the quality factor of a superconducting cavity resonator, depends on the fraction of quasiparticles in the conductor and their response to oscillating electromagnetic fields. Mattis-Bardeen theory gives the dependence of the surface impedance on frequency and normal-state material quality [67]; it can be expressed in terms of a complex surface impedance $\sigma_1 + i\sigma_2$ that both depend on temperature. The ratio of the components determines the significance of the resistive channel, which is proportional to $((\hbar\omega)^2 / (k_B T)) \ln(4k_B T / (\hbar\omega)) e^{-\Delta / (k_B T)}$, where Δ is the superconducting gap energy in the BCS theory and ω is the frequency of the fields [68, 150]. At an achievable temperature $T \approx 20$ mK, the dimensionless part of this factor evaluates to about 10^{-46} . This should produce cavity quality factors much greater than 10^{10} at microwave frequencies, and yet no such devices have been measured. The temperature-dependence of this behavior can be shown to agree with expectations before saturating at low temperature due to other mechanisms [67, 151].

evaluate them after measures are taken to mitigate the offending losses. The seamless systems discussed in Section 4.3 are well-suited for such studies, which will involve both simulation and measurement of many varied systems within these testbed structures. An understanding of these limitations will be critical to the task of constructing the best possible cQED module.

5.1 Energy participation ratios

Energy participation ratios constitute an intuitive way to characterize and compare the spatial and material-based division of electromagnetic (EM) energy in a system. The participation ratio is indexed by each mode, or configuration of electric and magnetic fields. Single structures, like resonators, for example, will have different arrangements of EM fields at different frequencies.

In a microwave circuit, where EM fields exist in structures that can be more capacitor-like or inductor-like, it can be useful to consider two separate participation quantities. The first, the electric field participation ratio, is defined as $p_k = U_k^{(E)} / U_{\text{tot}}^{(E)}$, the ratio of electric field energy in a material indexed by k to the total electric field energy. The second, the conductor participation ratio, is defined as $\alpha_k = U_k^{(H)} / U_{\text{tot}}^{(H)}$ — a similar comparison, but for magnetic field energy.

It is important to note that on resonance, the electric and magnetic energies are equal: $U_{\text{tot}}^{(E)} = U_{\text{tot}}^{(H)}$. This means that we must be conscientious of sources of dissipation that arise in conductors as well as dielectrics.

5.1.1 Mechanism-dependent quality factor bounds

Among candidate interfaces, bounds can be set based on system coherence over all previous experiments involving this interface. The most stringent, meaningful bounds are set by the experiments in which the interface in question participates strongly, and yet the system demonstrates high coherence. Therefore, a suitable testbed must (1) isolate the desired interface as well as possible, reducing the possibility of another interface participating with equal or greater value, (2) generally retain high coherence, and (3) ideally, allow a parameter to be varied so as to observe a clear dependence related to the interface in question.

The first two requirements become clear from the general expression for the quality factor arising from




Figure 5.1 | **Resonant circuit with many losses.** The total quality factor comes from a combination of the internal quality factor Q_i which encompass “environmental” loss rates Γ_i that lower the measured quality factor. Additionally, the coupling quality factor Q_c also “loads” the total measured Q .

a single dominant loss mechanism, Q_{material} , where the participation is given by p_{material} :

$$Q_{\text{material}} \geq p_{\text{material}} \times Q_{\text{measured}}. \quad (5.1)$$

The third requirement becomes evident through practice, since contributions of many sources obscure the search for a single effect (Section 5.1.2). These individual sources combine in the fashion indicated in Figure 5.1. Loss rates Γ_k are related to quality factors by the definition of the quality factor,

$$Q_{\text{material},k} = \frac{\omega U_{\text{tot}}}{\Gamma_{\text{material},k}} \quad (5.2)$$

which relates the energy stored, U_{tot} to the rate of dissipation. Since loss rates adds, this means that quality factors attributed to N loss mechanisms must be combined in reciprocal and accordingly weighted by their participation ratios, $p_{\text{material},k}$:

$$\frac{1}{Q_i} = \sum_{k=1}^N \frac{p_{\text{material},k}}{Q_{\text{material},k}} = \sum_{k=1}^N \frac{1}{Q_k}. \quad (5.3)$$

In Figure 5.1, the coupling quality factor Q_c is treated separately, leading to a “total” measured quality factor

$$Q_{\text{measured}} = \frac{1}{Q_c^{-1} + \sum_{k=1}^N Q_k^{-1}} = \frac{1}{Q_c^{-1} + Q_i^{-1}}. \quad (5.4)$$

Our candidate testbed, the seamless coaxline architecture, offers high coherence and tunable partici-

pation of modes through geometric modification. It is well suited to evaluate the performance and limitations of planar devices. The seamless coaxial stub cavity also demonstrates high coherence (millisecond lifetimes), and along with other 3D cavities is thus positioned to place bounds on the quality of bulk conductors and their surfaces. Transmons, which localize fields strongly near bulk dielectric, can also reveal insight into the quality of dielectric surfaces. Therefore, tests can be performed in many similar systems differing slightly in parameters affecting the participation ratios of loss mechanisms suspected of limiting these systems. With enough variations, and of course, proper statistical sampling, bounds on many of these suspected sources will arise.

5.1.2 Limits and combined losses

Many of the analyses performed in this chapter will use the assumption of a single loss mechanism to set a limiting bound on the quality of the mechanism in question. But what if N sources contribute, particularly at a similar level? Since we have established that Q s add in reciprocal, so will the participation-weighted Q_{material} s (Equation 5.3). The more sources contributing at equal level, the less stringent our bound will become on any particular source's Q_{material} .

If more than one mechanism is likely present in a loss survey, then the participation of all N should be estimated by simulation for each of the tested variations. Bounds can be evaluated in the normal fashion for each source. However, a multi-parameter fit can also be performed, obtaining best-fit values for each of the parameters $\{Q_{\text{material},1}, \dots, Q_{\text{material},N}\}$. For more than one parameter, it is difficult to intuitively evaluate the goodness of this fit, which falls within an N -dimensional space. The accuracy of this procedure depends on the accuracy of the model used to produce the contributions, potentially leading to large uncertainty. Furthermore, the number of measurements required to produce constraining fits increases exponentially with N . Finally, any variability in sources of competing magnitude that are unaccounted for in the model will obscure all other correlations (Section 5.6).

5.2 Modeling and simulating bulk dielectric participation

We first consider the dissipation induced by macroscopic (or “bulk”) dielectrics that serve as substrates for our circuits. The dissipation of a dielectric material at microwave frequencies is given by its loss tangent,

Table 5.1 | **Bulk dielectric loss tangents.** Best-case loss tangent bounds for dielectrics common in cQED, adapted from ref. [152].

material	$\tan \delta (\times 10^{-6})$	Q_{material}	source
sapphire	0.02–0.3	$(3–50) \times 10^6$	[153, 154]
AlO _x	1500–4000	250–700	[155, 156, 157]
AlN	1100–1800	550–900	[158]
silicon	0.15–200	$5 \times 10^3–7 \times 10^6$	[152, 159]
SiO ₂	320	3000	[158]
SiN _x	100–200	5000–10,000	[158]

$\tan \delta$, which relates the real and imaginary parts of the dielectric permittivity, $\epsilon = \epsilon_1 + i\epsilon_2$:

$$\tan \delta_{\text{material}} = \epsilon_2 / \epsilon_1 = 1 / Q_{\text{material}}. \quad (5.5)$$

The value of $\tan \delta$ is frequency dependent and can even be anisotropic (orientation dependent). Loss tangents for materials common in cQED are presented in Table 5.1. Note that these values are representative of single samples; the loss tangent measured for any one sample may depend on the source material quality, growth type, doping or impurities, and additional processing.

Energy participation fractions can be obtained by simulating the electromagnetic field profile of a resonant structure, for example in a finite element modeling software like Ansys HFSS, COMSOL, or Sonnet. These macroscopic (mm-to-cm-size) dielectric objects can be simulated exactly within the model geometry. An eigenmode-type simulation produces electric and magnetic fields at the resonant frequency, from which the energy density can be evaluated:

$$U_E = \frac{1}{2} \iiint (\mathbf{E} \cdot \mathbf{D}) dV = \frac{1}{2} \iiint \epsilon |\mathbf{E}|^2 dV, \quad (5.6)$$

where ϵ is the electric permittivity of the material. In the case of a resonant cavity partially filled with dielectric (region 1) and vacuum (region 2), the energy participation ratio in the dielectric can be expressed as:

$$p^{\text{dielectric}} = \frac{U_{E,1}}{U_{E,1} + U_{E,2}} = \frac{\epsilon_r \int_1 |\mathbf{E}|^2 dV}{\epsilon_r \int_1 |\mathbf{E}|^2 dV + \int_2 |\mathbf{E}|^2 dV}. \quad (5.7)$$

In practice this quantity easily converges in simulation and is simple to compute. These will be compared to measured values in designs containing dielectrics in Section 5.8. As an approximate point of comparison, though, the coaxline resonators introduced in Section 4.3.1 typically have $p_{\text{sapphire}} \sim 0.4$, whereas transmon qubits participate around $p_{\text{sapphire}} \sim 0.9$. The discrepancy between these structures can also be easily understood by conforming each geometry to that of a simple, analytically addressable toy geometry discussed below.

5.2.1 Model of a parallel-plate capacitor

An idealized structure, a toy parallel-plate capacitor, can be a helpful concept in modeling the electric field participation in surface dielectrics among a variety of geometries. In the case of a rectangular cavity, this model is related in a straightforward way and can be simply understood (see also Section 5.3.3). The TE_{101} mode of a rectangular cavity has nonzero electric field on just two faces, and so the surface fields can be represented by a parallel-plate capacitor (a cross-section of top and bottom plates is shown in Figure 5.2a). A lossy dielectric layer (green) is modeled on top of the surface (black), and has a thickness t (typically assumed to be 3 nm) much less than the cavity height d (usually ~ 7 mm). The distributed energy density along the width of a wall will scale as $\sin^2(\pi x/w)$, but a parallel plate capacitor in electrostatics produces a constant value across the width. In the cavity, the ratio of energy in the dielectric layer to total energy also has a constant distribution, making the parallel-plate model a suitable representation. This surface dielectric participation ratio in this model is then given by

$$p_{\text{dielectric}} = \frac{2U_{E,d}}{U_{E,\text{all}}} = \frac{2(t/\epsilon_d)E_0^2 A/4}{dE_0^2 A/4} = \frac{2t}{\epsilon_d d} \quad (5.8)$$

where $A = w \times L$ is the area of the three-dimensional capacitor plate. Clearly, d serves as the characteristic distance that sets the participation ratio in this model (for the rectangular cavity TE_{101} mode).

This model can be extended to characterize 3D cavities fully filled with a dielectric medium, by changing the filling dielectric in the capacitor model. Planar structures, like transmons or striplines, however, contain conductors that produce fields partly living in regions with varying dielectric permittivity. These can also be modeled as parallel plate capacitors, though partly filled with substrate dielectric. This model

allows the energy to be calculated in all volumes, thereby obtaining an electric energy participation ratio. All structures will have an effective parallel plate distance d' , permitting convenient comparisons.

As an example, the coplanar waveguide (CPW) geometry (Figure 5.2b) features a center conductor and two flanking ground planes. It sits on a dielectric substrate that is assumed to be significantly more extensive than all other dimensions. To model the thin layer of lossy dielectric (ϵ_d), a half-filled parallel plate capacitor represents the division of energy between the two regions. The plates are coated with the lossy layer in question. The capacitance of a CPW (on the boundary of ideal, infinite half-spaces of dielectric and vacuum) is given semi-analytically by

$$C \approx 2(1 + \epsilon_{sub})\epsilon_0 L \frac{K\left(\frac{w}{w+2g}\right)}{K\left(\sqrt{1 - \left(\frac{w}{w+2g}\right)^2}\right)} \quad (5.9)$$

where $K(x)$ is the elliptical integral of the first kind ([160], Equation 2.29). Because the capacitance of the structure in Figure 5.2c is given by

$$C = (1 + \epsilon_{sub})\epsilon_0 L \frac{w'}{2d'} \quad (5.10)$$

the two capacitances can be related to find e.g. $w' \approx 2.44d'$ for the case where $w = g$.

The goal is to formulate a metric d' that is independent of platform or geometry. Thus, an absolute w' must be determined that appropriately represents $p_{\text{dielectric}}$. This is found by noting that in the partially filled capacitor,

$$p_{\text{dielectric}} = \frac{4t}{d'\epsilon_d} \left(\frac{\epsilon_{sub}}{1 + \epsilon_{sub}} \right) \quad (5.11)$$

where t is the dielectric layer thickness. By performing a single 2D electrostatic simulation of the realistic geometry and explicitly calculating $p_{\text{dielectric}}$, the value for w' can be found. We find that interdigitated capacitor (IDC) structures with $w = g$ typically produce $w' \sim g$, whereas coplanar waveguide (CPW) and coplanar capacitor (CPC) structures with $w = g$ typically produce $w' \sim 1.3g$.

In transmons, **E**-field energy is localized between capacitor pads, which have an effective parallel plate distance that is approximately the distance from the closest edges of the two pads. This distance is typically




Figure 5.2 | **Effective parallel-plate capacitor model for planar structures.** **a** A rectangular cavity (cross-sectional dimensions d, w) with a metal–air surface dielectric layers (of thickness $t \ll d$, relative permittivity ϵ_d) can be treated exactly as a parallel plate capacitor to obtain the participation ratio of the dielectric layer. **b** Cross section of a CPW geometry. The conductor has thickness h and a metal–substrate dielectric layer thickness t (not drawn to scale). **c** Structures on the surface of a dielectric substrate (relative permittivity ϵ_{sub}), including the CPW, can be effectively represented as a half-filled parallel plate capacitor. The ratio of w' to d' is set by the geometry (such as g and w) and easily determined analytically, whereas either individual value depends on h and t , and is most easily determined by electrostatic simulation.

much smaller than the distance to the outer cavity or enclosure wall; thus, this capacitance dominates. Energy between the pads would divide equally between upper and lower half-spaces, if not for the dielectric constant of the substrate. (Or, the displacement electric field is continuous across the substrate–vacuum boundary.) The dielectric produces an apparent increase in effective ‘electrical’ volume, so an even distribution of energy in this effective space (where the speed of light remains that in vacuum) will produce a participation ratio that is approximately

$$p_{\text{substrate}} = \frac{U_{(E),\text{dielectric}}}{U_{(E),\text{vacuum}} + U_{(E),\text{dielectric}}} \approx \frac{\epsilon_r}{1 + \epsilon_r} \quad (5.12)$$

which for sapphire is ~ 0.9 . Stripline structures are much more distributed, however; their capacitance is defined between center conductor and the ground plane. Depending on the structure’s aspect ratio, therefore, the energy is no longer equally divided between two electrical half-spaces. Specifically, more energy is expected to exist in vacuum, reducing the bulk substrate participation ratio.

5.3 Evaluating loss in 3D structures

Bare 3D cavities (that is, cavities absent of substrate chips or any other materials except readout couplers) appear to consist of only the bulk material from which they are fabricated. (Planar structures, such as those patterned on substrates, will be treated in Section 5.4). Losses may therefore arise from defects, lattice vacancies, or magnetic impurities within the conductor [161, 162], microscopic layers of oxide that have grown on their surfaces [119, 156, 163, 164], or from gases that condense and “cryo-plate” (adsorb at cold temperatures) on surfaces of the cavity [165, 166].

We have systematically navigated a fraction of the parameter space necessary to identify the source of limiting losses in several 3D geometries. This pursuit was made difficult by the variability between samples that was sometimes present (Section 5.3.6), perhaps due to changing ‘environmental’ conditions. This may be the manifestation of additional, uncontrolled sources of loss, which complicate the interpretation of an incomplete model (Section 5.1.2).

5.3.1 Origin of bulk surface losses

The first step in a study of a loss mechanism is to formulate a model that can be compared to measurable data. In the case of energy dissipation in surface layers of bulk conductors, exact knowledge of the composition and properties of these materials is still unclear. However, a model that assumes certain properties—such as the thickness of a layer, or the permittivity of a material—serves as a basis with which to perform meaningful comparisons. Here, we seek to inform these guesses so that our comparisons using this model will depict the physical relations as accurately as possible.

Two forms of surface loss are typically attributed to bare cavities formed from bulk metal, such as our 3D high-purity aluminum cavities: dielectric oxide, and conductor loss. When these are assumed to be the only contributions, they are characterized by their participation ratios and material quality factors as

$$\frac{1}{Q_{(\text{from surfaces})}} = \frac{p_{\text{diel}}}{Q_{\text{diel}}} + \frac{p_{\text{cond}}}{Q_{\text{cond}}}. \quad (5.13)$$

The nature of these contributions can be understood as follows.

When exposed to air, aluminum grows an oxide. While the exact scale of this surface layer is not known exactly, the growth process is understood to be self-limiting, typically at thickness levels of several nanometers [164, 167]. The chemical composition of this aluminum oxide, Al_2O_3 , is similar to sapphire. This suggests that it is likely to have a comparable value of electrical permittivity, $\epsilon_r \sim 10$. The structure of the crystal is likely much more amorphous, disordered, and inclusive of impurities, however; higher loss tangents are therefore to be expected [168]. Not all surface oxides will grow alike, and may depend highly on the type of substrate. In particular, the nature of surface oxides grown on deposited metals will be discussed further in Section 5.4.1.

Conductor loss is another important form of loss associated with surfaces. Since the current within the superconductor is largely relegated to within a characteristic length scale from the surface, the penetration depth λ —the equivalent of the skin depth in a normal metal—the quality of the surface is more important than the bulk quality, generally. In thick aluminum in the ‘dirty limit’, this depth has been found to be just $\lambda \sim 50$ nm [67]. (The penetration depth λ in this limit will exceed the London length λ_L slightly, by a factor related to the coherence length ξ_0 as $(\xi_0/\lambda_L)^{1/3}$. [68]) Conductor loss can therefore be classified as a surface effect, extant on a scale similar to that of oxide growth. Studies of conductor surface loss for varying surface preparation have shown that the highest-quality microwave cavities are made by using high-purity aluminum and preparing the surface by chemical etching [93].

5.3.2 Simulating surface losses

From our basic understanding of the sources of material dissipation gleaned in Section 5.3.1, we can ascertain reasonable estimates of surfaces losses in thin surfaces simulated with a conventional, finite-element, high-frequency electromagnetic field solver. On the scale of millimeters, a superconductor with nanometer-scale penetration depth can be simulated as a perfect conductor (in which fields will immediately decay). A thin dielectric layer will only minimally perturb the energy distribution. Electric and magnetic field surface energy can be obtained and compared to the energy stored in all system objects to obtain the surface layer participation ratios.

It would be difficult to directly simulate both nanometer-scale surface layers as well as the large, centimeter-scale mode volume of a 3D cavity. At the same time, however, the approximation that resolves this conundrum—

taking the electric field to be constant throughout the layer—should be relatively accurate. Therefore, we proceed to perform simulations in this simple and direct fashion, bypassing the tricks required for more complex features (Section 5.4.1).

5.3.3 Variation of energy participations

Due to the simplicity of 3D cavities, the principal way in which p_{diel} or p_{cond} can be adjusted is by geometric variations. This proves difficult, however, for two reasons. First, while seamed rectangular cavities can be machined with many different dimensions, Section 4.2 explains why they are not likely to set material quality limits compared to other designs. Seam loss can vary from assembly to assembly and run to run, increasing uncertainty. On the other hand, to explore a range of participation ratios in seamless designs that would produce noticeable effects, geometries would have to change significantly. Such drastic changes would task machining capabilities, particularly for designs like the coaxial stub cavity (Section 4.3.1), with aspect ratios that are already difficult to manage. Some geometries have higher-order modes with frequencies within the measurement bandwidth and with a different arrangement of \mathbf{E} and \mathbf{B} that will participate differently. In practice, however, the coupling to these higher modes can be very different than the fundamental mode, making them difficult to measure during the same experimental run.

We can provide analytical expressions for the change in participation ratio for some designs, and perform simulations for others. In rectangular cavities, we expect a very simple relationship for the TE₁₀₁ mode, where the participation ratios

$$p_{\text{diel}} = 2t/(\epsilon_r c) \quad (5.14a)$$

$$p_{\text{cond}} = 2\lambda_L/c \quad (5.14b)$$

depend only on the narrow dimension c (Figure 5.3). In coaxial stub cavities with inner and outer radii a and b , respectively, we expect the participation ratios to be approximately

$$p_{\text{diel}} \approx \frac{t_{\text{diel}}}{\epsilon_r} \left(\frac{a^{-1} + b^{-1}}{\ln(b/a)} \right) \quad (5.15a)$$

$$p_{\text{cond}} \approx \lambda_L \left(\frac{a^{-1} + b^{-1}}{\ln(b/a)} \right) \quad (5.15b)$$




Figure 5.3 | **Typical 3D cavity geometries.** Definitions of the geometry of rectangular cavities (a) and coaxial stub cavities (b) that will be measured in this chapter. c The surface layers of interest for participation ratio calculation in both geometries are a dielectric layer (green) and conduction layer (orange). The black layer contains no volume, merely indicating a boundary.

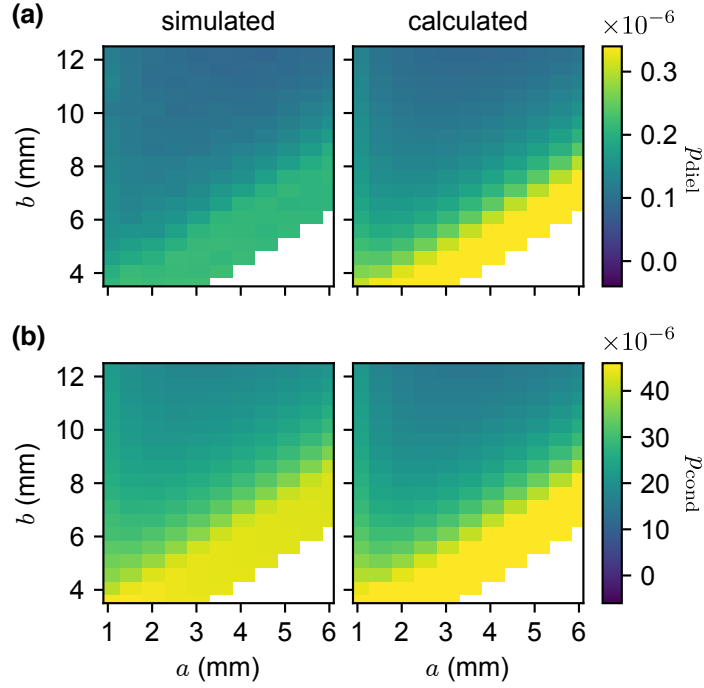


Figure 5.4 | **Coaxial stub participation simulations.** a–b Simulated (left) and analytically calculated (right) values of dielectric participation (a) and conductor participation (b) in a coaxial stub cavity with dimensions defined in Figure 5.3. All dielectric interfaces are taken to have $\epsilon_d = 10$ and $t = 3$ nm.

where the coefficients of this equation can be understood through the discussion in Section 5.3.4. The geometric similarity between the bulk surfaces of coaxline enclosures and the coaxial stub cavity make this a useful first approximation for that geometry.

We compare these analytic expressions to calculated ones for exact coaxial stub cavity geometries in Figure 5.4. They match very closely except where a approaches b , where increased significance of fringing effects is not captured in the analytic expression. Similar simulations can be done for the rectangular cavity, where the dependence only on c holds very strongly. (The curved ends of realistic rectangular cavities, located along at end locations b apart, produce a deviation in p_{cond} and only when $b \sim a$.)

Bare cavities are natural testbeds for measuring dissipation from bare cavity surfaces, of course. But more complicated systems can be used, such as those containing additional dielectrics and resonant structures described in Section 5.4. Since the losses in these structures will be evaluated anyway, they can also be used to obtain comparable bounds on the contribution of bulk surfaces.

This combination of challenges makes it difficult to perform a simple sweep in geometric parameters

that can lead to an identifiable change in loss (Figure 5.4). Furthermore, the inability to change any parameters (except mode number, which even then can be limited by the range of accessible measurement frequencies) *in situ* means that measurement variability will play a role in obscuring already-shallow trends. Nevertheless, we press forward by simulating the participation ratios in common designs, measuring them, and compiling the results to try to place the highest bounds on surface material qualities.

Choosing appropriate external coupling strengths

This chapter will present the internal quality factor Q_i , which will often be shortened to Q , for many resonant devices. Though the total loaded quality factor is more directly measured, the internal Q can be extracted from the loaded quality factor through the fitting procedure discussed in Section 3.4.3. We will strive to measure in an undercoupled regime, where $Q_i \approx Q_{\text{tot}}$. Measuring in a regime that is too over-coupled can be misleading for several reasons. First, it can distort this extracted value, making the internal quality factor appear higher due to noise. Second, a weak coupling is less likely to introduce external radiation or harmful effects that are difficult to model. Third, resonators used as memory cavities will be used in a weakly coupled regime; thus, testing them in this regime is more representative of actual use conditions. Therefore, we try to avoid this regime as much as possible throughout these measurements.

5.3.4 Bounding surface dielectric and conductor losses

Dielectric loss is evaluated by calculating the energy in a thin layer of thickness t_{diel} and permittivity ϵ_r , and comparing it to the total energy. This is almost always used in the regime where the relative dielectric energy is less than 1% of the total energy; therefore, the denominator in the expression for participation ratio omits the energy in this layer. The continuity of the electric displacement field \mathbf{D} across the dielectric boundary mandates that $E_{\perp, \text{diel}} = D_{\perp, \text{vac}}/\epsilon_r = E_{\perp, \text{vac}}/\epsilon_r$, so from Equation 5.6 a calculation of the ratio of electric energies becomes

$$p_{\text{diel}} = \frac{\frac{1}{2}\epsilon_r \iiint_{\text{diel}} |\mathbf{E}_{\text{diel}}|^2 dV}{U_{(E), \text{all}}} \approx \frac{\frac{1}{2}\epsilon_r t_{\text{diel}} \iint_{\text{diel}} (1/\epsilon_r^2) |\mathbf{E}_{\text{vac}}|^2 dA}{U_{(E), \text{vac}}} = \frac{t_{\text{diel}} \iint_{\text{diel}} |\mathbf{E}_{\text{vac}}|^2 dA}{\epsilon_r \iiint_{\text{vac}} |\mathbf{E}_{\text{vac}}|^2 dV} \quad (5.16)$$




Figure 5.5 | **Bulk metal surface participations.** The relationship between dielectric and conductor participation follows the analytic curve (lines) in Section 5.3.3, with deviations from differences between the realistic geometries and the analytical ones. The range of simulated data (points) includes the space of all reasonably producible variations (including those in Figure 5.4). Arrows indicate the direction in $\{p_{\text{diel}}, p_{\text{cond}}\}$ -space towards other devices that have been used to measure participation, such as cylindrical cavities [67, 93] and fully planar resonators [169].

where dA represents a patch of the surface that separates dielectric from vacuum (and the approximation holds for $p_{\text{diel}} \ll 1$).

Conductor loss is evaluated in simulation for bulk cavities in a fashion similar to dielectric losses, since

$$U_B = \frac{1}{2} \iiint_V \mathbf{H} \cdot \mathbf{B} dV = \frac{\mu_0}{2} \iiint_V |\mathbf{H}|^2 dV \quad (5.17)$$

in our case. The quantity of interest is the surface layer magnetic field energy as it compares to the total magnetic field energy. Here, the concept of a ‘layer’ is less concrete. The relevant ‘thickness’ that gives us energy in the proper units is λ , the penetration depth of the superconductor (as discussed in Section 5.3.1). Therefore, the conductor participation ratio becomes

$$\alpha \equiv p_{\text{diel}} = \frac{\frac{1}{2} \iiint_{\text{cond}} |H_{\text{cond}}|^2 dV}{U_{(H),\text{all}}} \approx \frac{\lambda_L \iint_{\text{cond}} |H_{\text{vac}}|^2 dA}{\iiint_{\text{vac}} |H_{\text{vac}}|^2 dV} \quad (5.18)$$

where, like the dielectric, the area integral is performed over the \mathbf{H} -field as simulated at the surface boundary.

The results of simulated aluminum and niobium cavity measured internal quality factors are plotted

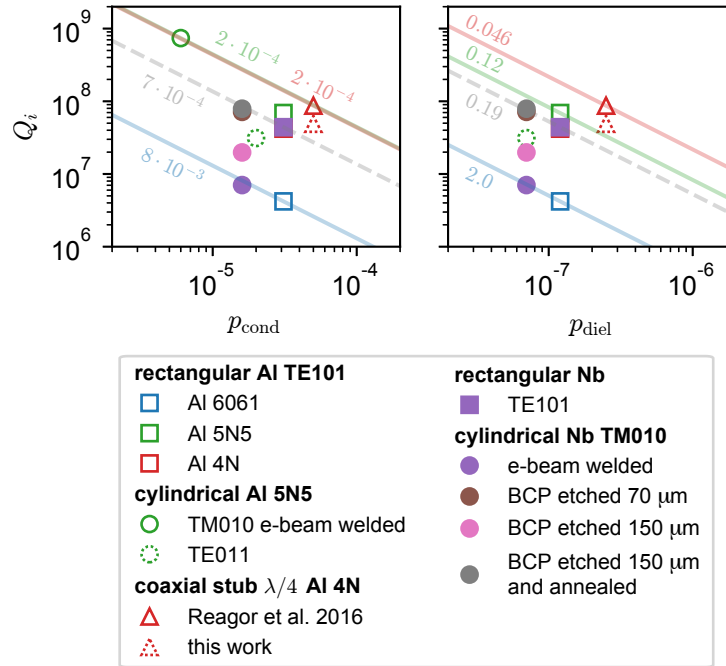


Figure 5.6 | **Best-in-class 3D cavity measurements.** Outer conductor and dielectric for the combination of materials and geometries of bare 3D cavities measured. The lines are bounds for each of the materials, for which the associated loss tangents are marked. The legend lists the material type, quality, cavity geometry, measured mode, and any additional processing.

in Figure 5.6 against surface dielectric and conductor losses. Lines of bounded Q_{material} help to identify which experiments set the highest bounds on the material quality factor for each material. Three types are aluminum are distinguished: the aluminum alloy 6061, aluminum 5N5 (99.995% pure), and aluminum 4N (99.99% pure). In the case of the electron-beam welded niobium cavity, different processing steps that may affect surface quality are also indicated. They consist of combinations of buffered chemical polish (BCP) etches and annealing, which is done by heating the cavity to 800 °C in ultra-high vacuum for two hours. The bounding potential of these measured values will be compared to those in literature and to the quasi-planar resonators and transmons measured later in Section 5.8.

5.3.5 Adsorption losses

Few studies have considered that a surface dissipative layer may not be ‘permanent’. If air or moisture is present in a cavity when it is cooled down from room temperature and ultimately below its superconduct-




Figure 5.7 | **Experiment testing adverse effects of moisture.** **a** A picture of the package and diagram of the experiment performed, where frozen water fills the cavity up to a height indicated by the arrow. **b** The simulated frequencies of the two modes are drawn as solid lines; the two frequencies measured in an experiment without water are plotted as points. The dashed lines are the measured frequencies *with* ice, and both agree well with a filling level of 11.6 mm.

ing transition temperature, that gas and moisture will freeze if it cannot escape. Experiments in which cavities were intentionally filled with water showed that our typical process of assembling cavities does not seal them hermetically. The competing processes of a refrigerator cooldown—simultaneous pumping and freezing—means that a significant quantity of moisture can remain, however.²

Frozen moisture will probably collect as either a bulk volume (thick layer sitting at the bottom of a cavity) or surface dielectric layer (adsorbed evenly on the walls), displacing vacuum and introducing dissipation. We performed an experiment to determine the effect of frozen water on the surfaces of a cavity by attempting to measure its dielectric loss tangent in a seamless coax stub cavity resonator (Figure 5.7a).

Coaxial stub cavities were filled with filtered, deionized water to represent condensed atmospheric moisture. The cavities were “sealed” normally. A turbomolecular pump struggled much more than usual to evacuate the outer vacuum chamber, until the pulse tube cooler was turned on and the dilution refriger-

2. In this experiment, the pressure of the chamber in which the cavity resides drops below 100 nanobar well before its temperature drops below 0 °C. Any liquid water will thus quickly become vapor, skirting the triple point and eventually desublimating into ice around -100 °C if it still remains.

Table 5.2 | **Results of water ice measurement.** Characterization of water ice microwave loss properties by resonator measurements. Mode A suffered from significant vibration, which may have reduced the measured Q_i with respect to Mode B.

mode	f_0 (GHz)	p_{ice}	Q_i	Q_{ice} bound
A	3.025	0.665	2.7×10^6	1.8×10^6
B	8.080	0.455	5.8×10^6	2.6×10^6

ator began to cool. This struggle was indicative of pumping on water at its liquid–vapor phase transition, causing it to slowly evaporate. The refrigerator eventually cooled to reach usual temperatures near 20 mK.

The resonators were measured normally. They were found to have significantly reduced frequencies, consistent with strong dielectric loading in the form of a bulk solid. Assuming water filling the cavity from the bottom up and a dielectric constant $\epsilon_{\text{ice}} = 3.15$ [170], the level of water was estimated based on the frequency (Figure 5.7b) and was found to have decreased from the room temperature quantity. When the refrigerator was warm and opened, approximately this quantity of water was observed to remain in the cavities.

The measurements and bounds in Table 5.2 suggest that frozen water is, in some form, not detrimental to the present level of performance of 3D cavities. We were able to place an upper bound on the loss tangent of this form of ice to be $\tan \delta \leq 3.8 \times 10^{-7}$ at 20 mK and 8 GHz. This is contrary to the reports of the loss tangent of ice at atmospheric pressures, extrapolated to be $\tan \delta = 5 \times 10^{-4}$ at 255 K and 9.4 GHz, at best [170, 171, 172]. It seems likely that water ice cooled to cryogenic temperatures, probably forming the orthorhombic ice phase XI, may have very different electrical properties than normal hexagonal crystalline ice. Realistic adsorbed amounts of water or atmospheric gas, like nitrogen, will be much smaller than those tested here. Thus, the combination of small participation ratio and high material quality factor bound should make the effect of such gases negligible *as a source of bulk dielectric loss* for many orders of magnitude to come. (However, atmospheric moisture has been shown to have an effect on the surfaces and edges of patterned structures [173].) If these results can be confirmed with further studies, frozen liquids or gases could even be incorporated in cavity designs intentionally.

5.3.6 Repeated thermal cycling

Degradation to the quality factor can occur if a sample is subjected to repeated cooling and warming between room temperature and millikelvin temperatures (Figure 5.8). This behavior has previously been observed in some high- Q 3D cavities [151], though not in others [174]. While the cause is not definitively known, and it does not always occur, there exist some theories based on the loss mechanisms discussed previously. Material impurities may not be static over many cycles of thermal contraction, and may migrate or concentrate near the surface where conductor participation is highest. For example, hydrogen diffusion may form hydrides at the surface of the metal or cause surface oxide to rearrange, positioning two-level systems in sensitive locations [175]. Contact with solvents or air containing dirt or dust may form a thin layer of dielectric. The material may also face mechanical damage, altering surface roughness.

Stripping the oxide and/or an outer layer of aluminum by chemical etching may return the surface to close to its original condition. Control over the experimental environment can protect against changes in magnetic field, for example, that might produce variability from thermal cycle to thermal cycle. These remain to be examined further; by identifying and removing sources of variability, we can aid in our identification of the other loss sources discussed throughout this chapter.

5.4 Loss in planar structures

Planar structures typically involve additional processing, different materials, and more complication compared with bare cavities. Therefore, we treat them separately, and primarily to survey the properties of these added materials. The same analyses discussed in Section 5.3 can be applied to the enclosures binding these structures; however, the plethora of possible parallel loss mechanisms will likely inhibit the setting of bounds in this manner.

Similarly to bulk surfaces, loss in planar or quasi-planar structures can come from the dielectric oxides on the surface of metals (Figure 5.9); an extensive review can be found in ref. [176]. Additionally, however, the substrates on which conductors rest and the interfaces between materials contribute, depending on the fabrication procedure. Generally, a single conductor layer (or multiple layers formed *in situ* without atmospheric exposure) is thought to have three possible interfaces. The metal–substrate (MS) interface is




Figure 5.8 | **Performance of a coaxial stub cavity over repeated cooldowns.** **a** Quality factors of the same two coaxial stub resonators, part of the same aluminum piece, measured over successive cooldowns. The state of the cavity prior to each cooldown is listed. “Vented” means that the cavities were fitted with an air-permeable lid; “humidified” indicates that the cavity was not stored in a dry environment in between cooldowns; “ice” means that the cavity was subjected to the conditions described in Section 5.3.5; “baked” indicates that the cavity was heated above 100°C while under vacuum in an attempt to remove moisture. **b–c** The relative frequency shift of each device. Data from cooldown 5 onward is plotted separately, since Cooldown 4.5 induces a large jump in frequency.




Figure 5.9 | **Candidate dissipative surface layers.** Interfaces that may contribute to loss in the coaxline geometry. MS: metal–substrate, MA: metal–air (metal–vacuum), SA: substrate–air (substrate–vacuum). Like bare cavities, the outer surface can also give rise to dissipation. Additionally, the bulk loss tangent of the substrate can participate.

wedged between the metal and its supporting substrate; the metal–air (MA) interface, which will be metal–vacuum during measurement, forms on top of the metal; the substrate–air (SA) interface may form before, during, or after the metal deposition and might be similar to the MS or MA materials.³ The properties of thin-film conductors can also differ from their bulk counterparts, even if they consist of the same material.

In this section, we will investigate the nature of materials-based loss in transmons, stripline resonators, and comparable structures. Quasi-planar structures are evaluated using the packages introduced in Section 4.3.1, whereas transmons are measured in more traditional enclosures (the rectangular geometry of Figure 5.3). I will discuss our model of loss mechanisms, the extent to which approximations can be made in this analysis, and what conclusions about geometry can be deduced.

5.4.1 Accurate planar structure simulations

Simulation of energy participation in electromagnetic resonant structures formed by thin-film conductors, such as stripline resonators, brings special challenges. This is due to a combination of two factors: the wide range of feature sizes in a typical device (from hundreds of nanometers near a Josephson junction, up to cavity sizes of centimeters), and approximations of thin-film metals that prevent simulation convergence. Exact simulation of thin metal or dielectric layers is computationally challenging, and so they are approximated as two-dimensional sheets. This produces diverging electric field at feature edges, for which

3. In some references, “air” is replaced with “vacuum” when labeling these interfaces. We retain the “air” notation since, without exposure to oxygen at some point, these interfaces would probably not exist.

a finite mesh will systematically under-represent the energy fraction stored in interface dielectrics.

We have developed a multi-stage, physically robust approach to address these complications and simulate arbitrary structures. Once appropriate parameter ranges have been determined, such simulations can be applied to a variety of geometries with minimal human supervision. These simulation strategies enable our desired study of designs that vary participation ratios, placing bounds on the properties of surface materials in planar structures. This approach separates and independently simulates nanometer-scale and centimeter-scale features. Structures like the coaxlines of Chapter 4 require one additional electrostatic simulation to accomplish this. When applied to the transmon (Section 5.7.1), however, the junction and shunting capacitors must be treated differently, requiring multiple energy-extraction and scaling steps.

To avoid unphysical divergence, we first divide the lithographically patterned structure and its associated (assumed) MS and MA surfaces into “perimeter regions” and “interior regions” (Figure 5.10) with their boundary fixed at a constant distance (x_0 , typically 1 μm) from the edge.⁴ In a global coarse 3D simulation, electric field in the interior regions does not vary sharply, and therefore converges easily to a trusted value that we may immediately record as $\mathbf{E}_{\text{MA}}(x, y)$ or $\mathbf{E}_{\text{MS}}(x, y)$ at the top and bottom surfaces of the planar structure, respectively. We use these field distributions to calculate the surface participation associated with the interior region of the structure (denoted by the subscript “int”):

$$p_{i,\text{int}} = t \int\int_{\text{int}} \frac{\epsilon}{2} |\mathbf{E}_i(x, y)|^2 dx dy / U_{E,\text{tot}} \quad (5.19)$$

where $i = \text{MS}$ or MA , and $U_{E,\text{tot}}$ is the total electric field energy in the entire space (dominated by energy in the substrate and vacuum). Here we have multiplied the field integral by the assumed thickness of the surface layer, $t = 3 \text{ nm}$, further assuming that the electric field is uniform across that thickness.

Critically, the field given by simulation of a 2D sheet is not exactly \mathbf{E}_i . In the simulation, the field will be physically sampled just above (where $\epsilon = 1$, the relative vacuum permittivity) or just below (where $\epsilon = \epsilon_{\text{sub}}$, the relative substrate permittivity) the sheet. If we denote the sampled field as \mathbf{E}'_i , where $i = \text{MS}$

⁴ The SA surface can be similarly divided by a contour at a constant distance x_0 from the outside of the edge. The treatment of the SA surface is otherwise analogous to that of MS.

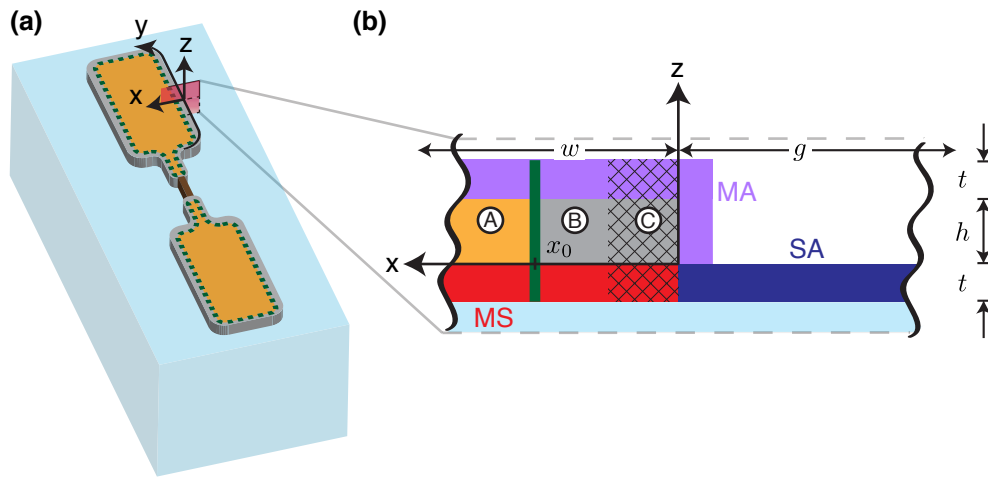


Figure 5.10 | **Dividing a 3D simulation into parts.** **a** The “global” high-frequency 3D simulation includes millimeter-to-centimeter-scale features within the cavity, such as the substrate and lithographic geometry. Here, a transmon stands in for any planar structure; the junction geometry is treated as a lumped element. All conductors are simulated as perfectly conducting sheets, and interface layers are omitted. The inner (orange) and perimeter (gray) regions are separated (dashed green contour). The red window identifies the cross section in **b**. **b** A cross section of edge of the lithographic structure, which is faithfully created within a 2D electrostatic simulation that supplements the global simulation. The three interfaces of interest with thickness t are shown as MS (red), MA (purple), and SA (blue), while the divided regions of the superconductor (thickness h) are shown in orange (Ⓐ, “interior region”) and gray (“perimeter region”) and are separated at position x_0 . The perimeter region is further divided into a cross-hatched region (Ⓒ), that converges only in this 2D simulation, and a region (Ⓑ) that is convergent in both simulations. When simulated, symmetric boundaries are established to represent an interdigitated capacitor (IDC) style device with conductor width w and gap width g ($w, g \ll x_0$). Drawing is not to scale. (Figure adapted from [163]; see [Copyright Permissions](#).)

is below, and $i = \text{MA}$ is above, then the extrapolated field is given by

$$\mathbf{E}_{\text{MS}} = \frac{\epsilon_{\text{sub}}}{\epsilon} \mathbf{E}'_{\text{MS}} \quad (5.20\text{a})$$

$$\mathbf{E}_{\text{MA}} = \frac{1}{\epsilon} \mathbf{E}'_{\text{MA}}. \quad (5.20\text{b})$$

The perimeter regions can be described by a spatial coordinate (x, y, z) as shown in Figure 5.10, where the y -axis winds around the edges of the pads of a transmon, remaining tangent to the perimeter. We further divide the perimeter regions into two halves. Energy in the half adjacent to the edge ($0 < x < x_0/2$) fails to converge, regardless of initial mesh parameters, following the adaptive mesh refinement process. The other half ($x_0/2 < x < x_0$) can be made to converge using a mesh configuration that is computationally accessible. The key concept to this strategy is to employ a constant ratio, or “scaling factor” F_i , to convert the integrated field energy in the convergent half into that of the entire perimeter regions, so that

$$p_{i,\text{per}} = F_i t \int_{x_0/2}^{x_0} dx \oint_y \frac{\epsilon}{2} |\mathbf{E}_i(x, y)|^2 dy / U_{E,\text{tot}} \quad (5.21)$$

The spatial distribution of electric field in the perimeter region (“per”) can be written using separation of variables as $|\mathbf{E}(x, y, z)| = C(y) f(x, z)$ in the limit of $x, z \rightarrow 0$. This is because the electric field near a metal edge should have a local scaling property independent of distant electromagnetic boundary conditions. Here $f(x, z)$ describes the edge scaling that can be applied to any cross section, independent of y . The actual form of $f(x, z)$ depends on material thicknesses and dielectric constants, and is difficult to derive analytically. However, we can compute $f(x, z)$ in a 2D cross-sectional electrostatic simulation of the conducting structure in question, which accounts for the actual thicknesses of each material (Figure 5.10b). The reduced dimensionality allows for accurate computation of the field inside the interface layer with sub-nanometer spatial resolution. We choose boundary conditions representative of the width of the pad (w) and the separation of opposing electrodes g (for a transmon), or the separation between the ground plane (for a stripline). Although such a cross-sectional simulation may not perfectly reflect the boundary condition in 3D space, as we already noted, $f(x, z)$ is independent of the distant boundary conditions as long as $x, z \ll g, w$. As an illustration, $f(x, -t/2)$ is shown in Figure 5.11a for a few very

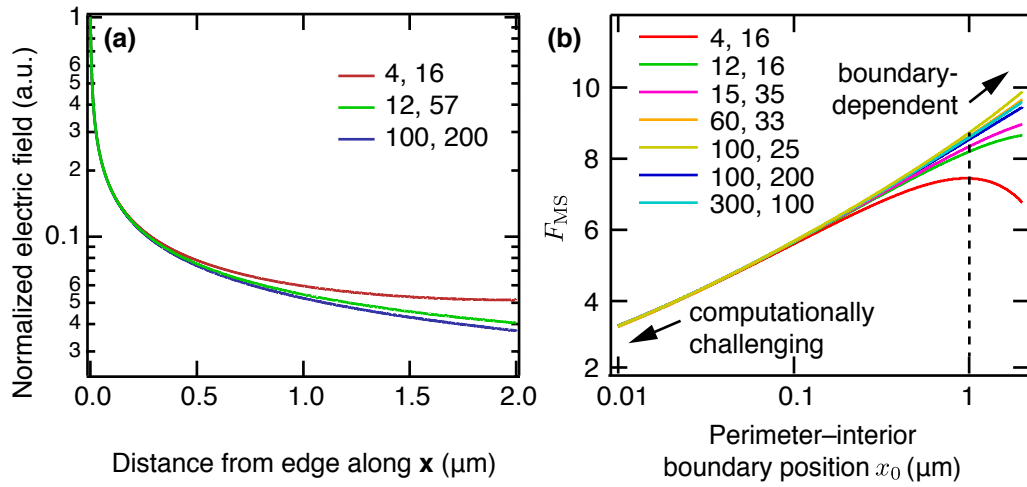


Figure 5.11 | **Using 2D electrostatics to scale 3D simulations.** **a** Normalized distribution of electric field within a cross section of the MS interface near the conductor's edge, $f(x, -t/2)$. It is calculated from the edge inward along a line that bisects the MS interface (Figure 5.10b). Values are plotted for simulations of three sets of boundary conditions, where the first number represents the width of the conductor, w , in micrometers, and the second number is the distance between features, g , in micrometers, as in Figure 5.10b. **b** A selection of MS scaling factors F_{MS} (Equation 5.22a) for varied perimeter region width assignments (x_0), simulated with various boundary conditions and following the labeling convention of **a**. We typically choose $x_0 = 1 \mu\text{m}$, indicated by the dashed black line. This value strikes a balance between boundary condition insensitivity and computational ease (indicated by arrows and the diverging lines). (Figure adapted from [163]; see Copyright Permissions.)

different values of g and w .

For striplines and transmons with feature sizes typically 10–500 μm in their smallest dimension (still ignoring the junction and its leads), and separations on about the same scale, the edge scaling function $f(x, z)$ correctly extrapolates the field within the perimeter region for properly chosen x_0 . From $f(x, z)$ we can calculate the scaling factor F_i based on the ratios of integrated field energy within the cross section:

$$F_{\text{MS}} = \frac{\int_0^{x_0} dx \int_{-t}^0 f^2(x, z) dz}{\int_{x_0/2}^{x_0} dx \int_{-t}^0 f^2(x, z) dz} \quad (5.22a)$$

$$F_{\text{MA}} = \frac{\int_0^{x_0} dx \int_h^{h+t} f^2(x, z) dz + \int_{-t}^0 dx \int_0^{h+t} f^2(x, z) dz}{\int_{x_0/2}^{x_0} dx \int_h^{h+t} f^2(x, z) dz} \quad (5.22b)$$

Scaling factors F_{MS} for various extents of the perimeter region are shown in Figure 5.11b. We limit our method to the regime where $x_0 \ll g, w$, where F_{MS} is insensitive to the values of g and w . In practice, we use $x_0 = 1 \mu\text{m}$ for typical structures (which are at least 10 μm in width and separation). Inserting these simulated scaling factors into Equation (5.21) allows one to arrive at $p_{i,\text{per}}$.

Surface–air scaling factors The scaling factors for the remaining surface, SA, can be determined in almost exactly the same way. The factor F_{SA} references an energy in the SA layer, within $0 \leq x \leq -x_0/2$, to the MS region beneath \textcircled{B} . Much like the interior region of the conductor, where simulations converge easily, energy from the remainder of the SA region can be extracted explicitly. In our simulations, this includes the side and rear faces of the substrate chip.

In coaxlines, it is helpful to have pre-compiled a library of scaling factors for many variations. These simulated factors and their dependences on geometric variation are shown in Figure 5.12. The magnitude of the MA scaling factor depends on whether a dielectric layer covering the sides of the conductor is included in the 2D simulation. Omitting this layer reduces F_{MA} by a factor of ~ 10 . All factors depend somewhat critically on conductor thickness, and in particular the ratio t/h .

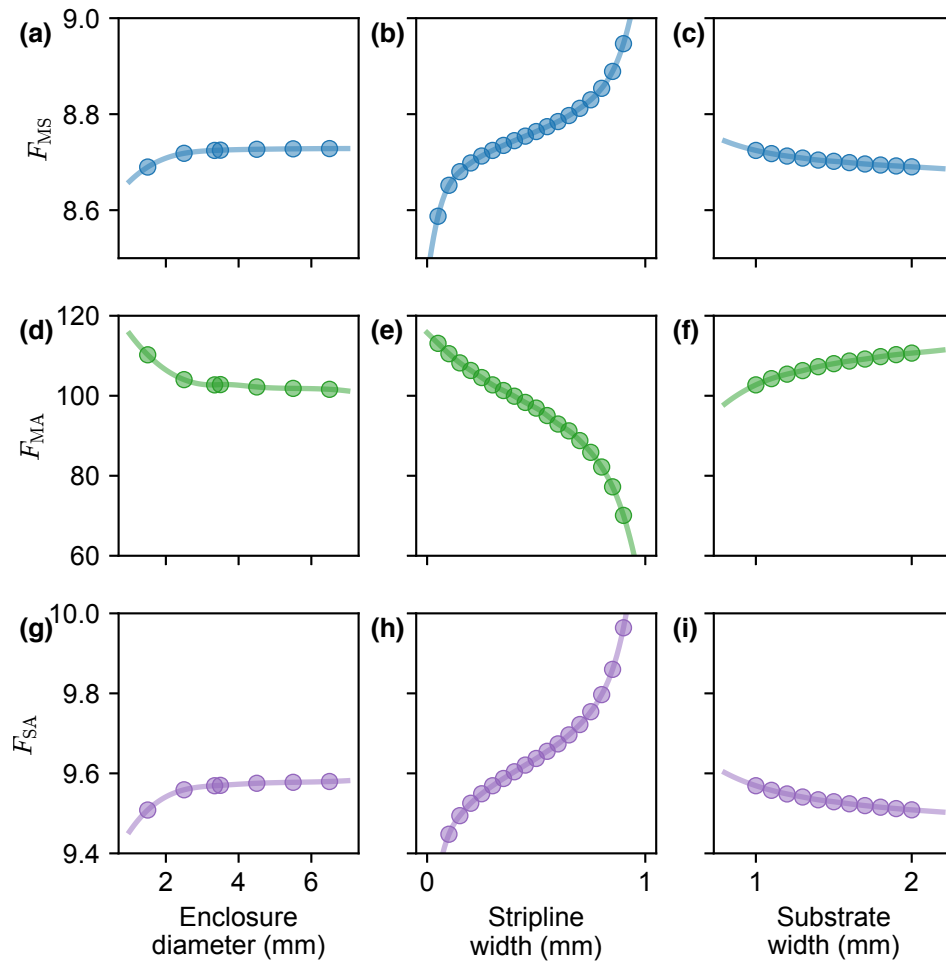


Figure 5.12 | **Perimeter energy scaling factors in the coaxline.** Scaling factors for MS, MA, and SA (rows) as they are calculated to change with selected variations in coaxline geometry (columns). The division between interior and perimeter region is set at $1 \mu\text{m}$. The substrate is sapphire ($\epsilon_{\text{sub}} = 10$) with 0.43 mm thickness. Conductor thickness $h = 80 \text{ nm}$ and dielectrics have $t = 3 \text{ nm}$ and $\epsilon_d = 10$. The lines represent cubic splines that are used to interpolate.

5.5 Varying participation ratios in coaxlines

The effects in Section 5.3, as well as those introduced above for planar devices, can be probed by measuring many striplines fabricated with different parameters. Our coaxline devices, consisting of a single layer of superconductor (aluminum) patterned far from the edges of a diced substrate—typically sapphire—are among the simplest planar devices that can be measured. To simplify matters and compare devices more consistently, a multiplexed design like that in Section 4.3.2 can be used to measure many variations at once.

Whereas participation ratios are difficult to vary over a wide range in bare 3D cavities (Section 5.3.3), planar structures offer more flexibility. Compared to transmons, striplines feature more straightforward fabrication and a simpler loss model. A stripline resonator can be made short, long, straight, serpentine (adjusting all losses), or even thick or thin (adjusting effective penetration depth). By altering the patterned design, a matter easily accomplished through lithography, loss-determining parameters can be varied, as observed in Figure 5.13.

Multiple stripline structures can be patterned on a substrate (Section 5.5.2), creating multiple interacting modes that do not just scale monotonically together, but can vary participations vastly in a diverging fashion. In devices with multiple substrates, the concentration of energy in dielectric or in vacuum depends on the arrangement of these chips. It is thus apparent that many options exist for varying participation ratios; we will survey a number of such variations here.

Some variations have the effect of scaling the expected behavior of multiple mechanisms in more or less the same way. For example, bringing transmon capacitor pads closer increases the participation of metal–substrate, metal–air, and substrate–air interface layers. Increasing the diameter of the coaxline enclosure dilutes nearly all volume participations (but reduces isolation between elements). This behavior should be expected; the same participation ratio model explains the increase in lifetime between coplanar waveguide cavities and 3D cavities. While these multi-mechanism trends are important to identify and implement, it means that geometric variations can be less useful in distinguishing particular sources and setting bounds on their magnitudes.

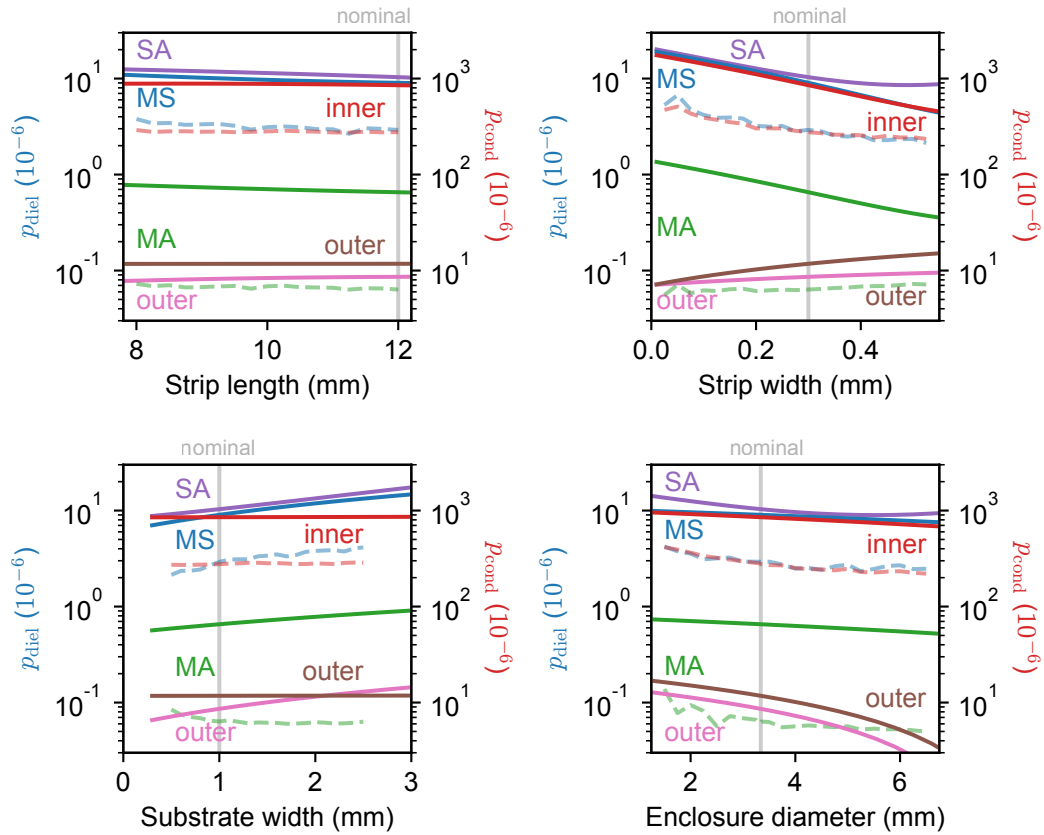


Figure 5.13 | **Simulated coaxline surface layer losses.** Simulated dielectric and conductor participation ratios of a coaxline stripline resonator for (a) variations in length, (b) variations in width, (c) variations in substrate width, and (d) variations of the enclosure diameter. Labels are nearest the axes they represent. Each sweep of an individual variable holds the others fixed at their nominal values (length 12 mm, width 0.3 mm, substrate 1 mm by 0.43 mm, enclosure 3.34 mm). MS (blue) and MA (green) are metal–substrate or metal–air dielectric interface participations, respectively. “Inner” (red) and “outer” (brown) refer to the conductor participation for inner (stripline) or outer (package) conductors, assuming $\lambda_L = (50 \text{ nm}, 250 \text{ nm})$, respectively. “Outer” (pink) refers to the package oxide layer. Solid lines represent data obtained through an exhaustive multi-stage simulation (Section 5.4.1); dashed lines are conventional simulation results. (For the package surfaces, they agree.) All conductors are assumed to be aluminum, with a sapphire substrate. Throughout this chapter, all dielectric interfaces are taken to have $\epsilon_d = 10$ and $t = 3 \text{ nm}$.

5.5.1 Geometry of a single stripline

Any chosen pattern can be inserted into the coaxline enclosure, with different designs determined by lithography (or even adjusted through post-processing). In particular, one can simply vary the substrate width, thickness and positioning, as well as the conductor width, length, and shape. Varying the diameter of the enclosure itself also turns the knob on nearly all participations within the coaxline enclosure (Figure 5.13).

In Figure 5.14, we show the measured internal quality factors of around 60 measurements of coaxlines compared to their simulated values of various layer participations. The bounds that can be set on the loss tangents associated with each of these sources' materials are plotted as lines, and are determined by the highest product of quality factor and participation. Numerous experiments fall beneath, rather than coincident with, these lines. In these cases, a single source of loss does not describe the observed behavior well.

If more than one mechanism contributes at a near-equal level, then the bounds derived from these devices alone will predict a quality factor that is much lower than the one measured. This results in some of the data of Figure 5.15a sitting below the line of equal expected and measured Q . The remaining data, which sits above the line, suffers from loss that is lower than predicted by all bounds combined. Another way to visualize this data dependent on many dimensions is to admit uncertainty, treating the bounds as guesses. The data can then be fit using a many-dimensional model to produce a best-fit loss tangent for each mechanism (Figure 5.15b). This approach is accurate only if the set of loss mechanisms we model is complete.

While the resultant fit does a better job of drawing the line through the data than the bound values do, it cannot account for the fact that nominally identical devices were sometimes measured to have vastly different quality factors. This observation confirms that our loss model is incomplete. Further, the minimal improvement from Figure 5.15a to Figure 5.15b suggests these unknown effects often dominate. (Alternatively, to some extent, material or interface quality may actually differ between samples.) We propose some ideas for what these unaccounted mechanisms may be towards the end of Section 5.5. In the meantime, however, it can be fruitful to try to learn about “best-case” bounding mechanisms by analyzing

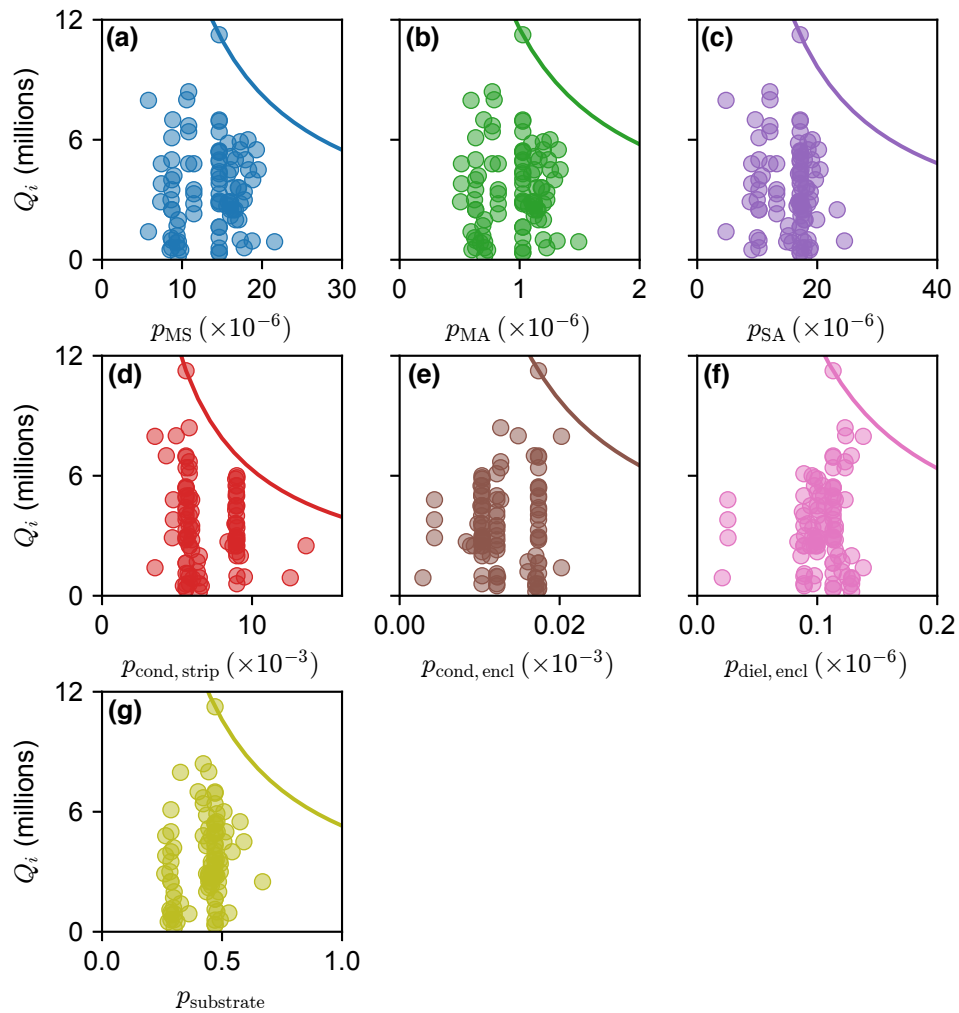


Figure 5.14 | **Aggregate coaxline measurements compared to participations.** Measured internal quality factors are compared with the simulated participation ratios stemming from (a) the metal-substrate layer, (b) the metal-air layer, (c) the substrate-air layer, (d) the stripline conductor quality, (e) the enclosure conductor quality, and (f) a dielectric layer coating the enclosure.




Figure 5.15 | **Coaxline performance from combined loss.** **a** Expected internal quality factors (Q) are determined from the bounds set by the highest- Q and largest-participation samples when compared to each participation type. The bounded loss tangents here are $(6 \times 10^{-3}, 9 \times 10^{-2}, 5 \times 10^{-3}, 2 \times 10^{-5}, 5 \times 10^{-3}, 8 \times 10^{-1}, 2 \times 10^{-7})$ for MS, MA, SA, inner conductor, outer conductor, outer dielectric layers, and bulk dielectric, respectively. **b** Expected Q s are determined based on least-square fitting that minimizes the distance from the dashed line. The loss tangents from this fit are approximately $(0, 0, 5 \times 10^{-7}, 3 \times 10^{-5}, 5 \times 10^{-3}, 0, 0)$.

measurements that appear least affected by variable mechanisms.

Enclosure diameter One variation that can be performed is that of enclosure diameter. To produce controlled conditions, we constructed a single-piece multiplexed enclosure in which identical chips were measured in tunnels of varying diameter. In Figure 5.16, we plot the data of these specially controlled experiments (square datapoints), as well as the compilation of all other data, as a function of enclosure diameter.

The best Q s in these devices are measured for intermediate enclosure size, which supports none of the predicted loss mechanisms, nor any fixed combination of them. From the few experiments in controlled, multiplexed devices, however, the rise in Q seems sharper than that predicted by the inner, substrate-based mechanisms. It is possible that, in these cases, losses from the enclosure walls dominated. This scenario would require that the material quality factors of these layers was lower than the best-known bounds.

Stripline width In one experiment, devices from the same chip and fabrication run were produced to have different widths. Referring to Figure 5.13, this varies the participation of surface layers associated with the substrate and patterned features differently than those of the enclosure. Therefore, it can be useful to distinguish the limiting cause of loss.




Figure 5.16 | **Variations in coaxline diameter.** Measured quality factors versus enclosure diameter. Dotted lines represent the calculated behavior of a loss mechanism associated with a substrate layer or stripline conductor (“inner” effects); dashed lines are for the enclosure losses (“outer” effects). The dotted and dashed black lines are the combination of inner and outer effects, respectively. Square datapoints come from experiments where four identical devices were measured simultaneously in a multiplexed package. The test values of loss tangents used here are $(10^{-2}, 2 \times 10^{-1}, 8 \times 10^{-3}, 10^{-4}, 10^{-2}, 1.3)$ for MS, MA, SA, inner conductor, outer conductor, and outer dielectric layers, respectively.




Figure 5.17 | **Variations in coaxline stripline width.** Measured quality factors versus stripline width. Dotted lines represent the calculated behavior of loss mechanisms associated with a substrate layer or stripline conductor (“inner” effects); dashed lines are for the enclosure losses (“outer” effects). Square datapoints come from an experiment where five devices with varied width were measured simultaneously in a multiplexed package with identical enclosures. The test values of loss tangents used here are $(10^{-2}, 10^{-2}, 7 \times 10^{-3}, 10^{-4}, 10^{-2}, 1.3)$ for MS, MA, SA, inner conductor, outer conductor, and outer dielectric layers, respectively.




Figure 5.18 | **Variations in coaxline substrate width.** Measured quality factors versus stripline width. Dotted lines represent the calculated behavior of loss mechanisms associated with a substrate layer or stripline conductor (“inner” effects); dashed lines are for the enclosure losses (“outer” effects). Square datapoints come from an experiment where four devices with varied substrate width were measured simultaneously in a multiplexed package with identical enclosures. The test values of loss tangents used here are $(10^{-2}, 10^{-2}, 10^{-2}, 10^{-4}, 10^{-2}, 1.3)$ for MS, MA, SA, inner conductor, outer conductor, and outer dielectric layers, respectively.

Figure 5.17 shows the measured quality factors arranged as a function of stripline width. We observe stronger dependence on surface effects at smaller widths, with increasing quality factor as width increases. This is consistent with substrate surface conductor or dielectric loss, and corresponds well with the simulations in Figure 5.13. However, as width increases, the quality factor in this controlled experiment begins to level off. This could indicate the sudden dominance of effects of the wall of the enclosure. The dashed lines in Figure 5.17 roughly follow this behavior. Together, the combination of all the predicted loss mechanisms might explain these observations.

Substrate size Varying substrate size has the effect of changing dielectric participation more than conductor participation. In one experiment, we prepared identical samples on a single chip, but diced the chip to make the substrate wider or narrower. The measured Q s of the resulting experiment are highlighted in Figure 5.18.

This figure shows that the controlled experiment agrees well with the trend of the best values in the combined data from all experiments: Q increases with increased substrate width. The dotted and dashed lines show that, for all else fixed, the electric-field-based losses should display the opposite behavior. Conductor losses, on the other hand, are not predicted to change as a function of substrate width, since no energy in the magnetic field is being repositioned. Therefore, it seems like the best values in each cate-

Table 5.3 | **Serpentine resonator simulated surface losses.** Simulated dielectric and conductor losses for a serpentine resonator of three different pitches, compared to a typical non-serpentine coaxline device (one with length matched to the serpentine device, and one with frequency matched to the serpentine device). The “pitch” here is defined as $\Delta x - w$ in Figure 4.11, where w is the conductor width. Simulations are done in a 4 mm-diameter tunnel with $w = 0.1$ mm and a 1.95 mm by 0.43 mm sapphire substrate.

pitch (mm)	length (mm)	p_{MS} ($\times 10^{-6}$)	p_{MA} ($\times 10^{-6}$)	p_{SA} ($\times 10^{-6}$)	$p_{\text{cond,strip}}$ ($\times 10^{-3}$)	$p_{\text{cond,encl}}$ ($\times 10^{-3}$)	$p_{\text{diel,encl}}$ ($\times 10^{-9}$)	$p_{\text{substrate}}$
-	5.7	30	2.0	23	1.6	7.4	62	0.735
-	12.0	21	1.5	18	1.5	7.4	84	0.517
0.05	5.7	17	2.5	30	1.5	1.6	37	0.557
0.15	5.7	19	2.0	33	1.4	1.5	41	0.513
0.25	5.7	11	2.1	18	0.8	1.6	54	0.462

gory of substrate width are not limited by dielectric losses, but by a combination of conductor losses and another yet-unknown effect.

Serpentine striplines An alternative coaxline structure introduced in Section 4.3.1 is that with serpentine leads. The narrow “pitch” between the meandering segments of conductor causes the resonant mode energy to reside near interfaces and surfaces to a greater extent than the “unraveled” form of the coaxline. Therefore, they should experience higher loss. In practice, serpentine striplines are typically not used as memories, and are instead strongly overcoupled and used as readout resonators. In the interest of evaluating their possible performance, however, the results of a participation ratio simulation are shown in Table 5.3.

5.5.2 Multi-mode stripline patterns

The most extreme variations in participation ratio often come from a drastic design variation, rather than adjustment of a continuous variable. One such variation that can be probed *in situ* is that of multiple modes. End-to-end coupled stripline modes, such as those in Section 4.3.1, are more or less independent and will not produce differing participation ratios. If two striplines are coupled side-by-side, however, two vastly different fundamental modes appear (Figure 5.19). These can be characterized as common or differential modes; the common mode appears most like the mode of one stripline.

From its likeness to a single stripline, the common mode will have moderate bulk dielectric partic-

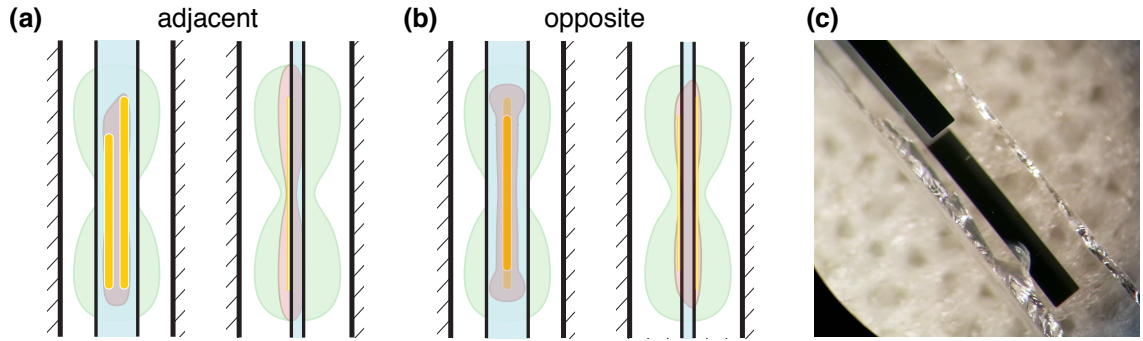


Figure 5.19 | **Geometry of multiple-stripline devices.** Electric field contours are also shown. **a** Adjacent striplines feature a common mode (green), where both resonate in a fashion similar to that of a single stripline. Their differential mode (red) concentrates more field between them and through the substrate. **b** Opposite striplines have a common mode (green), similar to the adjacent striplines, as well as a differential mode (red) that heavily focuses field through the substrate. **c** Optical micrograph of a device with opposite striplines.

ipation ($\sim 25\text{--}40\%$). The proximity of the differential mode, however, produces significantly higher, transmon-like participation in the bulk dielectric ($\sim 90\%$) and surface layers. One can also pattern the two striplines on opposite sides of the chip. This creates two modes, again, but where one participates strongly in the bulk substrate (akin to the transmon). The other acts more like a standard stripline mode on either side.

We calculate the expected participation ratio of such modes of such devices, and measure several. The results of these measurements are shown in Figure 5.20. While this small sample size produces results not as statistically significant as the compilation of measurements above, it does indicate some consistent trends. Further, the range of participation ratios spanned by modes of the same physical device (thus reducing measurement variation) is much larger than for standard coaxline devices, which helps to distinguish trends from other sources.

Looking at the participation ratio trends, devices again seem more consistently limited by MS, MA, and SA interface layers. Comparing the loss tangent bounds that can be set with this experiment to those in the coaxline devices, only the quality of the enclosure dielectric layer is bounded at a higher value. A seven-parameter fit of the form

$$Q_{\text{measured}}^{-1} = \sum_k p_{k,\text{simulated}} \tan \delta_k \quad (5.23)$$

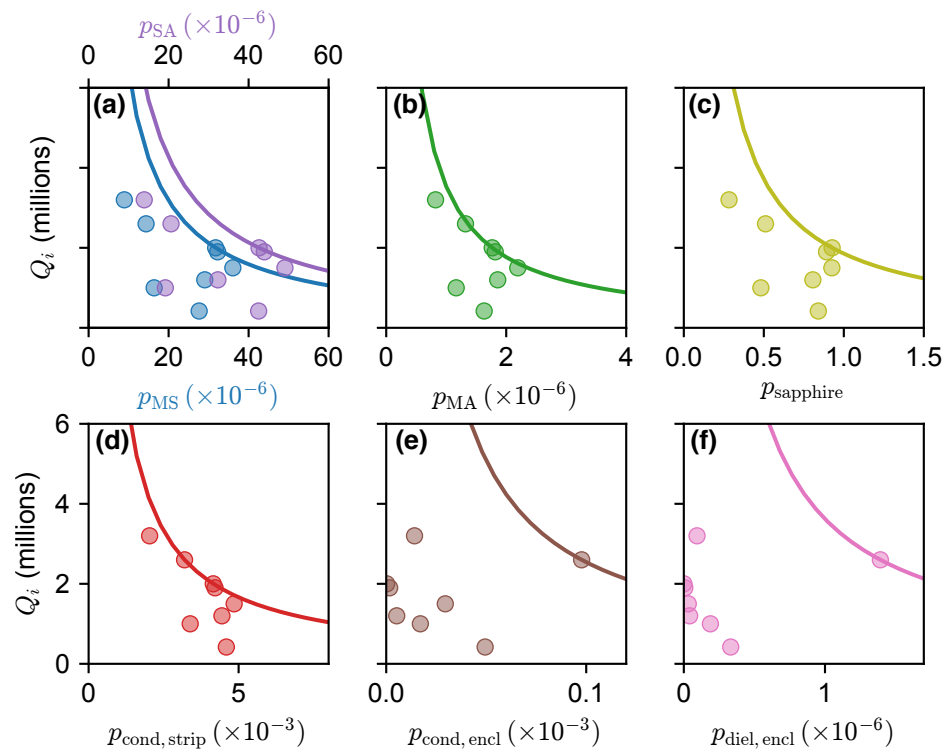


Figure 5.20 | **Quality factors from multiple-stripline device measurements.** Measured internal quality factors of multi-mode samples are compared with the simulated participation ratios stemming from (a) the metal–substrate and substrate–air layers, (b) the metal–air layer, (c) the bulk substrate, (d) the stripline conductor quality, (e) the enclosure conductor quality, and (f) a dielectric layer coating the enclosure.




Figure 5.21 | **Multiple-stripline performance from combined loss.** A multi-source loss model for double-strips. **a** Expected internal quality factors (Q) are determined from the bounds set by the highest- Q and largest-participation samples when compared to each participation type. The bounded loss tangents used to construct these “expected Q s” are $(1.6 \times 10^{-2}, 2.8 \times 10^{-1}, 1.2 \times 10^{-2}, 1.2 \times 10^{-4}, 3.9 \times 10^{-3}, 2.7 \times 10^{-1}, 5.4 \times 10^{-7})$ for MS, MA, SA, inner conductor, outer conductor, outer dielectric, and bulk dielectric, respectively. **b** Expected Q s are determined based on least-square fitting that minimizes the distance from the dashed line. The loss tangents from this fit are $(1 \times 10^{-9}, 3 \times 10^{-9}, 2 \times 10^{-9}, 1.6 \times 10^{-4}, 2 \times 10^{-10}, 0, 0)$ for MS, MA, SA, inner conductor, outer conductor, outer dielectric, and bulk dielectric, respectively.

allowing the loss tangents $\{\tan \delta_k\}$ to vary produces more reasonable agreement than for the coaxlines (Figure 5.21). This fit suggests that dominant inner conductor loss best explains the data. Normalized to the smaller sample size, this kind of experiment is far more effective at reaching new regimes of energy participation ratio and setting material quality bounds.

5.5.3 Direct measurement of kinetic inductance

Using the methods described in refs. [67, 93], we can directly measure the kinetic inductance fraction α by varying the temperature of the sample. The model we use is derived by calculating the surface impedance ($Z_s = R_s + i\delta X_s$) of a superconductor using BCS theory. Varying temperature will cause both the real and imaginary parts of the impedance to change. The change in R_s will change the quality factor of our microwave resonator to change, whereas a change in the differential reactance δX_s will shift the frequency. We can fit the change in frequency according to

$$\frac{f(T) - f(T \approx 0)}{f(T \approx 0)} = \alpha \frac{\delta X_s(T)}{X_s(T \approx 0)} \quad (5.24)$$

where the surface reactance at low temperature is given by $\omega\mu_0\lambda$ and λ is the penetration depth. From this analysis, we can extract α directly from measurements. We can also compare this value to a simulated value to obtain λ , the penetration depth, and learn about the quality of the conductor.

This analysis in coaxlines is complicated somewhat by the fact that the inner and outer conductors may have different kinetic inductances and possibly different material properties. Since the inner and outer conductors of many of the devices surveyed here are both aluminum, they will have similar critical superconducting temperatures. This makes it difficult to disassociate their kinetic inductances without carefully tracking the resonator through both superconducting transitions, or making assumptions about the relative scale of the two effects. One way they can be easily separated, however, is to switch the inner or outer conductors to materials with significantly different critical temperatures.

We performed a measurement using deposited niobium, with a thin-film critical temperature $T_c \sim 9$ K, as the center conductor for a coaxline device in an aluminum enclosure. The resonator was tracked through T_c of the enclosure, and the frequency and quality factor were recorded. In Figure 5.22 the data were fit to the model of BCS complex surface impedance to obtain the kinetic inductance fraction and the critical temperature. Since the critical temperature of our bulk aluminum is well known from prior measurements, we can fix the T_c as necessary to aid the fit. Here, we show data passing through T_c , allowing for extraction of the outer conductor $T_c = 1.168(2)$ K with high certainty. This also returns a reasonably bounded value on the kinetic inductance fraction, $\alpha = 1.81(2) \times 10^{-5}$, of our outer enclosure. (Because T_c of the niobium is so much larger here, we dismiss its effect on the changing frequency of the resonator. We justify this decision by observing the clear and expected behavior of a singular aluminum superconductor, exhibiting negligible deviance up to at least 2 K.)

This measured value for the enclosure α alone can be used to calibrate/verify simulations and extract information about the center conductor α in other measurements (Figure 5.23). We repeat this temperature-dependence measurement to extract α in several other aluminum/aluminum devices. These devices include those with different enclosure sizes, different substrate widths, and even coaxlines patterned with multiple modes. The variation in kinetic inductance is small among them, and does not conclusively make conductor loss culpable.




Figure 5.22 | **Varying temperature to extract enclosure kinetic inductance fraction.** Differential frequency shift measured as a package containing a niobium stripline resonator in an aluminum coaxline enclosure is heated up. From this, we can extract the kinetic inductance fraction and critical temperature of the enclosure package.




Figure 5.23 | **Measured and simulated values of kinetic inductance fraction.** **a** Comparison of the simulated expectation of the kinetic inductance fraction of the stripline to the measured values. Agreement in the slope is good, indicating a reasonable assumption of penetration depth, but the simulation appears to maintain an offset participation that may not be realistic. **b** The measured quality factors as a function of the measured kinetic inductance fractions. The line depicts the value obtained from the simulated bound.

5.5.4 Fabrication methods

The suspicion of substrate-based loss mechanisms prompts the question: can special processing be done to remove these layers entirely? [109] Other works have shown that process changes can be made to adjust the presence of harmful interface layers [164, 173]. To this end, we experimented with different fabrication processes.

In a departure from the norm of a standard “additive” liftoff process, where photoresist or e-beam resist is first patterned on the wafer before metal deposition, we performed “subtractive” lithography on select samples. Here, the aluminum was first deposited on sapphire under conditions that could be more stringently controlled and varied than for the additive process. Then, to remove the aluminum and to form the desired pattern, dry plasma etching or wet chemical etching was performed. In some cases, an oxygen plasma “descum” step was also done.

One can also vary the conditions in which the metal is grown. Depending on growth conditions, one can form epitaxial or amorphous aluminum with small or large grains. Our typical fabrication process makes use of an e-beam deposition tool, which is generally regarded to produce higher-quality aluminum films than by sputtering. We measured a number of devices grown at relatively low temperature (near room temperature) using a molecular beam epitaxy tool. The performance of the specially processed devices is shown in Figure 5.24 are compared to the normally processed ones.

5.5.5 Substrate and conductor materials

In a more drastic departure from normal fabrication processes, we can also consider changing our substrate or conductor materials to find ones of a better quality.

Sapphire is known to be one of the lowest-loss materials available to use as a substrate in our frequency and temperature regime (Table 5.1). In addition, using sapphire can be preferable because it is an excellent conductor of heat at low temperatures [177]. Silicon, however, sports the advantage of greater ease of processing, and the knowledge and backing of a vast semiconductor industry. It is versatile as it can be etched used to form useful microwave structures such as vias and cross-overs. Thus, in recent years, properly treated silicon has become a viable rival substrate to sapphire for superconducting circuits.

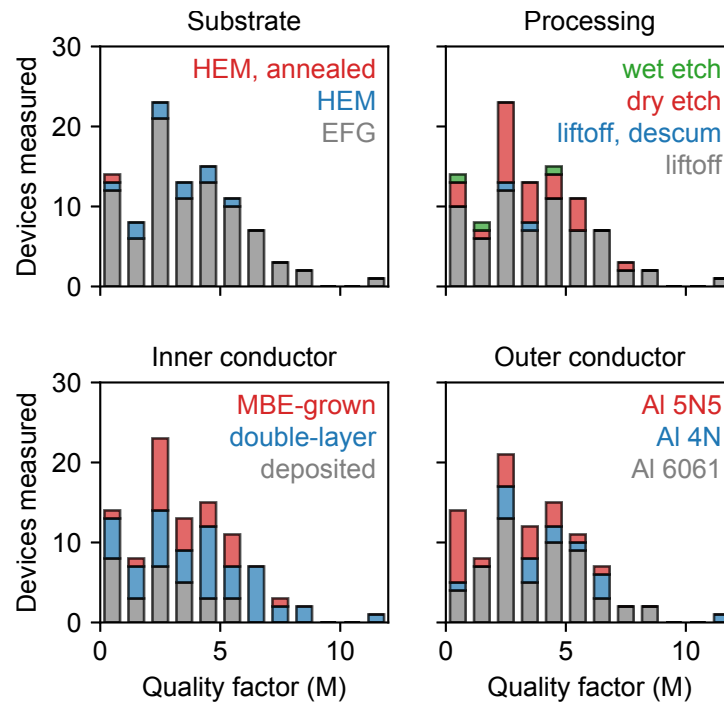


Figure 5.24 | **Quality factor histogram, comparing processes.** Histograms of measured quality factors for stripline resonators in coaxline packages. These measurements correspond with variations in (a) substrate growth type, (b) processing steps, (c) the inner conductor growth method, or (d) the outer conductor material purity. The substrate is edge-defined, film-fed growth (EFG) or heat exchanger method (HEM) sapphire, which in one case has been annealed at $\sim 800^\circ\text{C}$ overnight prior to deposition. The patterning process is done with photoresist and liftoff, followed by an optional oxygen plasma “descum” step (including a dip in hydrofluoric acid). An alternative patterning process involves deposition of the conductor followed by resist patterning and a wet etch (using MF319) or a “dry” plasma etch (using Cl in an Oxford 100 plasma etcher). Double-layer deposition follows the recipe for bridge-free transmon fabrication, whereas MBE growth refers to the low-temperature growth of an amorphous aluminum film on sapphire. 5N5 aluminum is 99.9995% pure, while 4N aluminum is only 99.99% pure; aluminum 6061 is a common alloy.

Table 5.4 | **Results from silicon devices.** Results of the best coaxline devices on a silicon substrate for any given process and material combination. Process 1 corresponds with a standard silicon liftoff fabrication, with an added HF vapor etch between each step [179]. Process 2 is a plasma etch using SF₆, with etch parameters 10 sccm, 13 mTorr, 100 W power and -390 V bias. Process 3 uses a SiO₂ hardmask (later stripped with HF vapor) and a wet etch using NH₄O₄. Process 4 uses an SF₆ plasma etch with 50 sccm flow, 80 W power, and 95 mTorr pressure. Process 5 is a dry etch of 80 nm NbTiN using a mix of SF₆ and O₂ (flow rates 20 sccm and 13.5 sccm, respectively) at 30 mTorr.

Device	Conductor	Process	f_0 (GHz)	$Q_i(\bar{n} \sim 1)$ (M)	Q_c (M)
CT2a	Al	1	7.51	0.738	1.1
CT2b	Al	1	7.39	0.690	19.8
FU14a2	TiN	2	8.392	0.49	0.77
NIST	TiN	3	8.859	0.20	7.0
FV14b1	TiN	4	7.833	0.14	0.6
BX15b	NbTiN	5	8.015	0.64	42

Silicon, however, has generally produced lower lifetimes in microwave resonators and transmons than for sapphire. A study of the lifetime of transmons fabricated on silicon (data reproduced for comparison later in Section 5.8) found that the bulk loss tangent was probably a dominant source in most regimes [97]. This would mark a notable difference from sapphire devices.

In this study, we also fabricated coaxlines with aluminum conductors on a silicon substrate. Silicon devices mimicked the geometry and processing of sapphire ones, but performed more poorly. We also tested different metals, such as niobium, niobium titanium nitride, and titanium nitride. These required different (subtractive) fabrication processes, and could not be compared one-to-one with the sapphire devices. However, they also performed poorly, in general, as shown in Table 5.4—possibly due to the generation and redistribution of lossy materials when etching using the processes that are non-standard for sapphire devices [178].

Since these materials are widely used in other regimes, it should be possible to make them high quality. We failed to develop a successful fabrication recipe except in one case of aluminum-on-silicon devices, discussed in the following section.

Etching for substrate removal Compared to sapphire, silicon can be easily etched. Etching can significantly reduce the influence of loss mechanisms associated with the bulk substrate and metal–substrate layer, while increasing those of the substrate–air and metal–air interfaces. (Changes to geometry alone




Figure 5.25 | **Profile of an etched coaxline** **a** A cross section of an aluminum stripline on a silicon substrate, before and after undergoing the DRIE process. **b** A scanning electron microscope image of the end of the resulting stripline structure.

Table 5.5 | **Results from etched silicon devices.** Results of several aluminum-on-silicon coaxline resonators after undergoing the BOSCH process.

Device	a (μm)	b (mm)	$Q(\bar{n} \sim 1)$ (M)	$Q(\bar{n} \gg 1)$ (M)
control A	0	0	0.58	0.74
A1	20–40	0.110	0.5	0.5
A2	0.7	0.105	5.2	7.3
control B	0	0	0.20	0.30
B1	0.7	0.03	7.2	12
B2	0.7	0.03	5.9	9.6
B3	0.7	0.03	4.2	>6

would scale these participations together, as shown in Figure 5.13.) Following the procedure in ref. [97], we perform a one-step deep reactive ion etching (DRIE) BOSCH process on aluminum striplines on a silicon substrate, producing the profile seen in Figure 5.25a. We then measured the devices in the same fashion as all the other resonators in this chapter; the results are shown in Table 5.5.

Device A1 performed no better than the control device. Though device A1 had a large undercut, microscopy revealed that this undercut was not uniform around the perimeter of the device, and led to significant overhangs and folding in places. In some locations, the etch produced a curved wall as opposed to the intended straight profile. We attribute this behavior to poor substrate cooling due to the long aspect ratio of the chip; temperature variability is known to critically affect the BOSCH process [180].

Device A2 fared better according to electron microscopy (Figure 5.25b). In this device, we measured a single-photon quality factor much larger than for other unetched silicon-based devices in this study. It also showed a strong improvement in quality factor at higher powers, a characteristic of limitation by TLSes

in a surface dielectric layer.

The devices in the B family were etched with a shallower aspect ratio and better temperature control, and also showed good improvement compared to the control device. This drastic jump in performance to the same level as sapphire devices suggests that silicon devices offer a promising way forward if fabrication difficulties are mastered and material properties carefully controlled. It also suggests that the silicon devices in this study were limited by the bulk silicon loss tangent and/or the metal–substrate layer, consistent with the findings in ref. [97].

5.5.6 Substrate processing

The device processing workflow usually ends with dicing or cleaving the patterned substrate to produce a desired shape. The typical method of dicing, which uses a rapidly spinning saw to cut through the material in a method akin to traditional machining, causes rough, chipped edges. This kind of dicing clearly damages the edges, creating chips and fractures. It is unclear exactly how, but reasonable to assume, that stressed and fractured substrate edges could influence the loss mechanisms discussed in this section.

We experimented with other dicing techniques, both involving lasers instead of blades. Our chips, which contain resonators, should not be at as significant of a risk from local heating and other potential ill effects of lasers. Two forms of laser ablation were used: “stealth dicing” (a process that internally modifies the crystal structure and then uses tape expansion to cleave), and a process called “LEAF dicing” (a low-temperature ablation process that sublimates and evaporates material and is rasterized across the cuts before mechanical cleaving) [181]. Figure 5.26 shows microscope images of the edges of chips subjected to different forms of dicing, as well as the result of measured resonators in each case. As with many of the other process alterations discussed in this section, dicing could not be identified as a ‘smoking gun’ in this case. While it is certainly possible that saw-based dicing could, or does, limit us in some instances, it does not seem consistent with being a dominant effect. If that were the case, anyhow, these experiments have shown that alternative dicing techniques are available.




Figure 5.26 | **Comparing substrate diced edge quality.** a–b Optical micrographs of sapphire sample chips following (a) traditional blade dicing and (b) LEAF dicing (which appears qualitatively similar to the stealth-cut chips). c Histogram of device performance, highlighting the devices that have undergone different dicing processes. Both blade-free dicing processes produced quality factors within the standard distribution, thereby not pinpointing blade damage as a certain limiting source of loss.

5.6 Measurement variability

Throughout this chapter, we have lamented the fact that measured quality factors do not seem to be well-described by an all-encompassing model of consistently contributing loss mechanisms. Here, we provide examples of these observations and touch upon a few of the difficult-to-quantify mechanisms that may contribute.

5.6.1 Repeated thermal cycling of transmons and resonators

Much like with 3D cavities, degradation of the quality factor or changes in device parameters can occur in quasi-planar resonators and qubits through repeated cooling and warming [182, 183]. This may arise from changes to material and interface properties, or possibly due to changes in compression of the clamping mechanism, for example via stiffening of compliant indium. Though rare with sapphire, tightly clamped chips can sometimes crack or break entirely.

We performed a series of experiments, repeatedly cooling down the same samples in the same package. They were not cleaned or repackaged between experiments. We measured the resonant frequency and quality factor over time (Figure 5.27). The Q_i of most devices followed a general downward trend, and did not strongly correlate with the conditions in which the device was cooled. The frequencies also decreased




Figure 5.27 | **Coaxline quality factors over repeated cooldowns.** **a** Quality factors of the same five coaxline devices over successive cooldowns measured in a single multiplexed coaxmux package. In cooldown 1, the package is treated normally. In cooldown 2, it is baked (see Figure 5.8 caption). In cooldown 3, it is vented. In cooldown 4, the sample is directly and extensively evacuated. In cooldown 5, the sample is baked and then cooled without breaking vacuum. **b** The relative frequency shift of each device.

as compared to the frequencies of the first cooldown. While variation in frequency could suggest a shift or settling of chips within their mounts and relative to the enclosures, a downward trend may be more suggestive of an added loss, such as growth of dielectric oxide or accumulation of debris, from repeated atmospheric exposure.

We also repeated this repetition experiment for transmons mounted in seammux packages (Figure 5.28). Here, the transmons are cooled down in identical conditions each time. Again, while there is some cooldown-to-cooldown variation, the general trend of all coherence values is downward. This is not necessarily representative of all transmons; some characteristic “gold standard” devices have retained high coherence when tested on a roughly annual basis. Such fluctuation may be more indicative of the measurement apparatus, such as cable and filtering parameters, than device surfaces and interfaces. However, it highlights the variability that makes the evaluation of loss in planar devices challenging.

5.6.2 Assembly imprecision

Since the capacitance of a distributed quasi-planar element depends on its position within the coaxline enclosure, a mispositioned stripline or transmon can produce undesired results. Furthermore, a change in position *in situ*, such as that stimulated by mechanical vibration, can cause a variable frequency and broadened linewidth that may limit its performance (Section 5.6.3). To determine the coaxline’s robust-

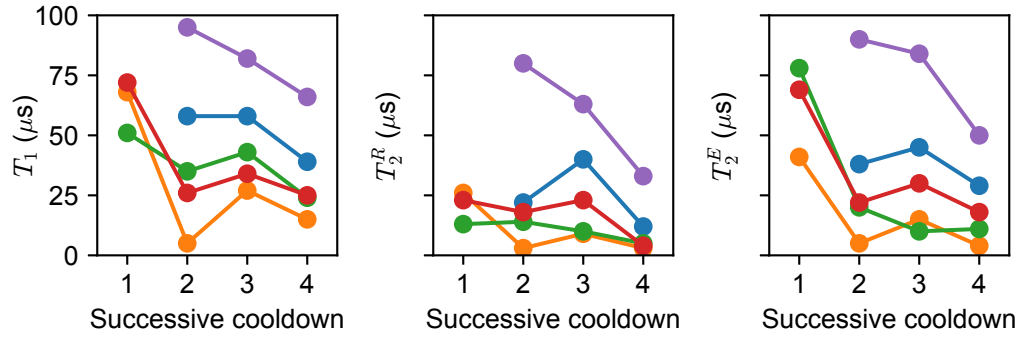


Figure 5.28 | **Transmon coherence properties over repeated cooldowns.** The T_1 (a) and Ramsey and echo T_2 coherence times (b–c) of five transmons measured in a multiplexed seammux package over successive cooldowns.

ness to misalignment during assembly, we simulate the resulting frequency with respect to perturbations in different directions (Figure 5.29). Adjustments to chip placement and orientation will also change the participation ratios of various interfaces. Therefore, unexpected variability in assembly can lead to variability in measured Q s.

5.6.3 Effect of acoustic vibration

Throughout these measurements, kHz-level vibrations were observed in measurement signals, particularly near the microwave resonance frequency of planar devices. This has the effect of reducing measured performance of resonators, setting lower bounds on material qualities than might otherwise be achieved. Since the magnitude of the effect depends on a number of variables that are not tightly controlled, such as the manner in which the chips are held, it also leads to measurement inconsistency. In this subsection, I discuss observations and likely causes and solutions to this undesirable behavior.

The spectrum of the vibration can be determined by taking a Fourier transform of a time-series measurement of the magnitude or phase of resonator transmission (Figure 5.30). For rapid sampling of the resonance peak at a rate much faster than the kHz noise, the movement of the peak as a function of time is resolvable. When the movement is stimulated by an acoustic tone that is largely sinusoidal (peaked at one frequency), sufficiently slow sampling has the effect of producing an seemingly broadened resonance peak.

Frequency-sensitive measurements, such as those performed on a vector network analyzer, are thus

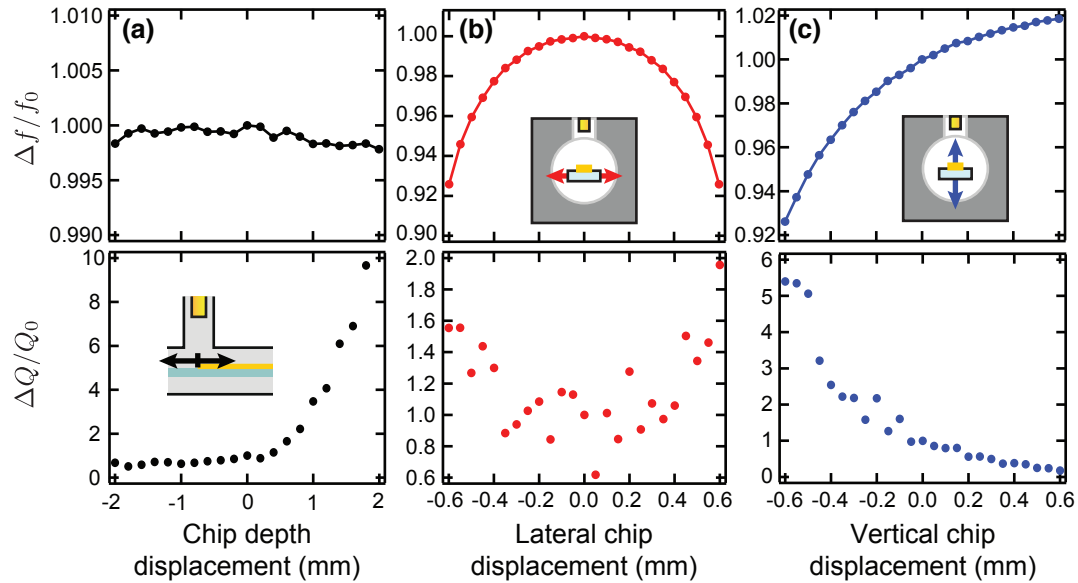


Figure 5.29 | **Coaxline sensitivity to perturbations in placement.** Fractional change in frequency and quality factor, referenced to a nominal frequency f_0 and quality factor Q_0 at zero displacement, for (a) a displacement along the axis of the enclosure, (b) parallel to the substrate face, or (c) perpendicular to the substrate face. (Figure adapted from [7]; see [Copyright Permissions](#).)




Figure 5.30 | **Symptoms of mechanical vibration.** a Power spectral density of $\arg(I + iQ)$ of the demodulated steady-state transmission response for a coaxline device that is clamped at both ends. b The same type of device is measured in a package where it is only clamped at one end. Significant response in magnitude and phase is observed at very low frequencies and near 4 kHz.

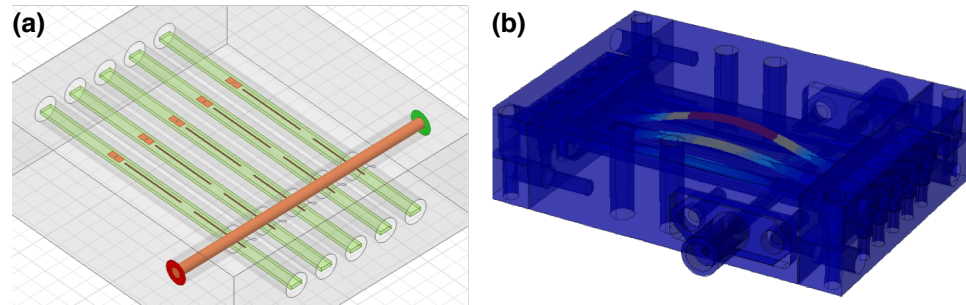


Figure 5.31 | **Modeling mechanical vibration.** **a** A model of a coaxmux device susceptible to vibration, where sapphire chips are colored in green and clamped at both ends. **b** This model is input into a mechanical finite element simulation, which identifies the mechanical resonant modes and their shapes. Here, we see one such mode corresponding with the flexing of one of the sapphire chips (exaggerated). A second chip flexes slightly because it has a nearby mechanical resonant frequency.

sensitive to this effect. A broadened linewidth will translate to reduced quality factor, which is representative of the energy lifetime of the device. Measurements of the energy ring-down can potentially differentiate between energy and phase stability of the resonator. This measurement method, however, can only extract the resonator's loaded quality factor. Further, any power-dependent effects will be folded into the time dynamics and become more difficult to discern, compared to steady-state measurements. Therefore, only undercoupled resonators with particularly power-independent behavior can be used to extract information about the low-power unloaded quality factors of vibrating resonators.

A suite of mechanical simulations and corresponding tests, wherein different device assembly methods were explored, were performed to identify the mechanism through which resonator frequency changes (Figure 5.31). Using the finite element modeler COMSOL, eigenmode simulations of the exact device geometry (using boundary conditions that allowed the ends of the beam-like substrate chip to pivot) reproduced the measured acoustic resonance frequencies very closely. Alternate suspicions, such as vibration of the package as a whole, or resonance of normal or feedline coupler pins, did not produce simulated resonant frequencies anywhere near the observed ones.

Ultimately, these resonators are to be used in conjunction with a transmon. So how will a proximal transmon, or even one patterned alongside on the same chip, react to this kind of noise? Here, the confinement of the transmon field reduces its sensitivity to changing capacitance-to-ground, say, by vibrating on a suspended chip, at least in comparison to distributed lithographic transmission line structures. The




Figure 5.32 | **Quality factor histogram, comparing packages.** A histogram of device performance for stripline resonators in coaxline and coaxmux packages (gray) as well as those in seammux packages (orange).

coupling strength of the transmon to a nearby feature should also vary little, particularly if they are already widely detuned.

Acoustic noise can potentially cause more potent problems, however. Vibration can heat the substrate or the conductor. It could also induce bulk or surface piezoelectricity, which may reduce lifetimes in transmons [184]. Clearly, it is preferable to avoid uncontrolled vibrations, no matter the method through which they inflict loss. This will require improved precision and careful engineering of assembly and clamping techniques.

5.6.4 A seamed, vibration-free alternative enclosure

As we just saw, samples measured in coaxline or coaxmux packages often suffer from significant vibration. Vibration can reduce the signal-to-noise ratio and broaden the resonance curve (depending on measurement bandwidth), and will be discussed further in Section 5.6.3. The seammux package, introduced in Section 4.3.2, does not appear to induce or allow vibrational broadening. We measured a number of devices in this package, the data for which is included among the other plots in this chapter. For clarity, however, we highlight the devices measured in this package in Figure 5.32.

While the best individual Q was, in fact, measured in a seammux package, these packages seem to enable the same general distribution of quality factors as is provided by the coaxline and coaxmux packages. The presence of a seam in the seammux package may introduce variability in measurement that is removed when vibrations are suppressed—an unfortunate trade-off that can be mitigated as shown in Figure 4.20.

It offers a hopeful outlook, however: that by understanding which element is key in suppressing vibration in the seammux, we may inform our construction of improved coaxline and coaxmux packages as well.

5.7 Varying participation ratios in transmons

In addition to a layer of deposited aluminum, transmons consists of a Josephson junction, which could possibly contribute losses through other mechanisms. Therefore, we treat our studies of the transmon separately, including by addressing the Josephson junction and the surrounding features explicitly (described in section 5.4.1). In transmons, the ratio of perimeter to surface area of the shunt capacitor pads is the primary method by which participations are adjusted.

To investigate the relationship between transmon loss mechanisms and their measured performance, we will first calculate the values of MS, MA, and SA interface participation ratios. Then, by performing variations in the geometry that change these simulated values, the lifetime of the transmon will be measured and its relationship to interface energy participation characterized.

5.7.1 Simulating transmon junction leads

The same technique introduced in Section 5.4.1 is applied to the capacitor pad structures of the transmon. The lead and junction area, however, must be treated somewhat differently.

A schematic of the Josephson junction and the leads is shown in Figure 5.33a, where x -axis and y -axis are defined perpendicular and parallel to the leads, respectively. We divide the surfaces associated with the junction leads into three regions based on distance from the junction: the near region ($|y| < 1 \mu\text{m}$), the intermediate region ($1 \mu\text{m} < |y| < 10 \mu\text{m}$), and the far region ($|y| > 10 \mu\text{m}$). The surface participation ratios for these regions are denoted by $p_{i,\text{near}}$, $p_{i,\text{mid}}$ and $p_{i,\text{far}}$ respectively, where $i = \text{MS}, \text{MA}$ or SA . The near region of the leads is not explicitly included in the global simulation. The intermediate and far regions are included in the global simulation, but the surface integration of field energy does not converge due to the influence of edges. A scaling factor solution akin to that in Section 5.4.1 demands $x, z \ll g, w$, but the lead is too narrow and too close to the junction to satisfy this. We use a supplemental “local” 3D simulation of the junction leads as shown in Figure 5.33a, which includes the thicknesses of all materials, to compute the surface participation of all three regions surrounding the leads.

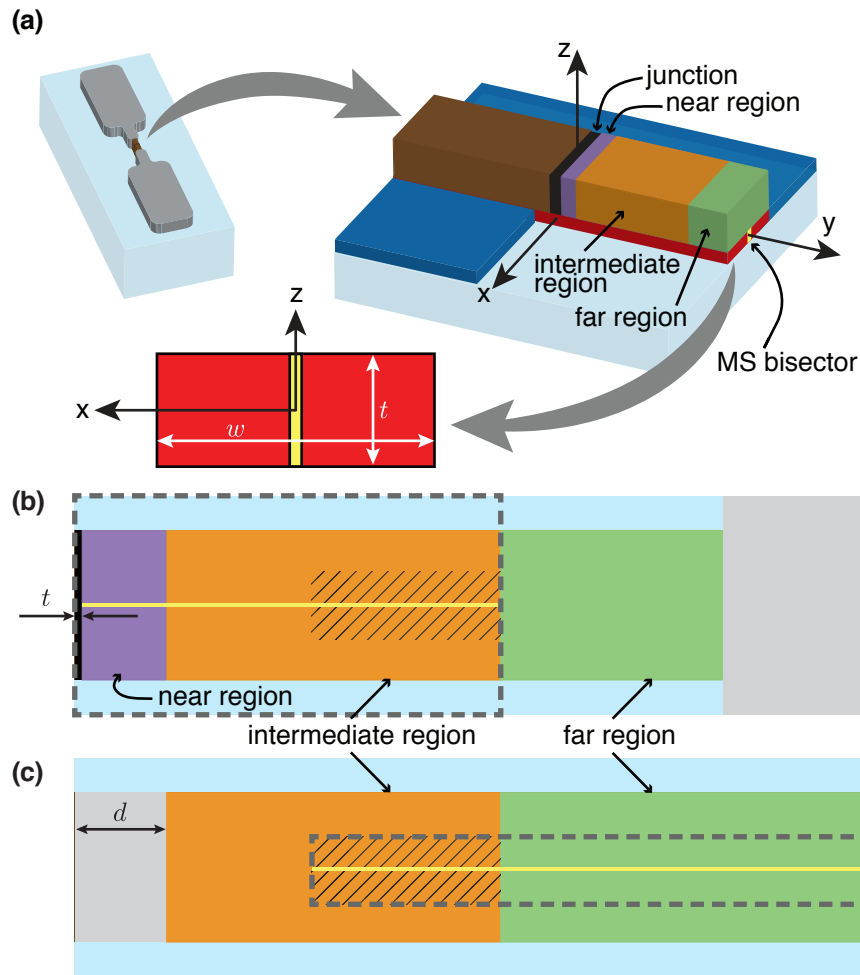


Figure 5.33 | **Reconciling 2D and 3D simulations of transmon leads.** **a** A local, near-junction simulation. Energy is integrated within cross-sections of the MS interface layer (red). Separately, energy measured along the bisector (yellow) is used for “stitching”: comparing local and global simulations in convergent areas to establish a common energy scale. **b** A top-view of the local simulation. The junction (black) is explicitly simulated with thickness t . The “accurate” area here (dashed) includes near and intermediate regions, and some adjacent SA dielectric; the far-region field is too sensitive to boundary conditions. The stitching region (hatching) marks the edge of this area. **c** A top view of the global simulation. The intermediate and far regions are explicit; the near area (extent d) is taken as a lumped element. The accurate portion of the leads (dashed) is near their center; energy near edges or the lumped element diverges. The comparison of stitching extent (hatched) along the bisector (yellow) to that of **b** establishes the constant local–global electric field scaling. Within the stitching region, the cross-sectional energy density is normalized by the bisector energy density to obtain an energy ratio $f(x, z)$ independent of y , useful for finding $p_{i,\text{far}}$ from Equation 5.26. (Figure adapted from [163]; see [Copyright Permissions](#).)

This high-resolution local simulation is performed by applying a static voltage potential between the pair of leads across the junction. The boundary of the local simulation is set sufficiently far (typically 25 μm) to ensure the calculated field distribution $\mathbf{E}_{\text{loc}}(x, y, z)$ in the near and intermediate regions is not affected by the type of boundary condition used. The absolute electric field magnitude in this local simulation is set by the imposed voltage, and must be rescaled by a constant C to be consistent with the field scale of the global simulation from which U_{tot} is obtained.

This constant C can be determined by comparing $\mathbf{E}_{\text{loc}}(x, y, z)$ with the field distribution in the global simulation $\mathbf{E}_{\text{glb}}(x, y, z)$ in a selected overlapping region (“stitching extent”) where both simulations are reliable. In particular, we choose the stitching extent as the center line of the leads in the $5 \mu\text{m} < |y| < 10 \mu\text{m}$ region (Figure 5.33b). Such a choice avoids the numerical imprecision of the global simulation in areas close to the junction or the edges. It also avoids any artificial boundary effects of the local simulation by remaining distant from the boundary. We have confirmed the two simulations show consistent spatial dependence over this stitching extent, $\mathbf{E}_{\text{glb}}(0, y, 0) \propto \mathbf{E}_{\text{loc}}(0, y, 0)$, and the constant C is computed from the ratio of the two.

Surface participation ratios for the near and intermediate regions of the leads can then be immediately calculated by integrating $\mathbf{E}_{\text{loc}}(x, y, z)$ over the volume of interest. For example,

$$p_{i,\text{near}} = C^2 \iiint_{i,\text{near}} \frac{\epsilon}{2} |\mathbf{E}_{\text{loc}}(x, y, z)|^2 dx dy dz / U_{\text{tot}} \quad (5.25)$$

The surface participation ratios from the near and intermediate regions are expected to be independent of the design of the electrodes.

On the other hand, Equation 5.25 does not apply to lead energies in the far region, which is not fully included in the local simulation. To calculate $p_{i,\text{far}}$ we adopt a separation-of-variables approach by noting that $|\mathbf{E}_{\text{glb}}(x, y, z)| = |\mathbf{E}_{\text{glb}}(0, y, 0)|f(x, z)$. Here $f(x, z)$ describes the cross-sectional distribution of electric field in dimensionless units (normalized by the field magnitude at the center line of the lead) (Figure 5.33a). It can be obtained from the local simulation of the junction leads discussed above, which also




Figure 5.34 | **Energy participation from junction outward.** Cumulative energy participation in the MS layer for transmon junctions. While capacitor pad design (selections of extremal geometries shown) alters participation ratio, the near-junction participation is unaffected by distant features. We choose a region $< 1 \mu\text{m}$ (shaded) in which to disregard energy participation. (Figure adapted from [163]; see [Copyright Permissions](#).)

confirms that $f(x, z)$ is independent of y for $y \gg 1 \mu\text{m}$. Therefore,

$$p_{i,\text{far}} = \int_{10 \mu\text{m}}^{y_{\text{far}}} \frac{\epsilon}{2} |\mathbf{E}_{\text{glb}}(0, y, 0)|^2 dy / U_{E,\text{tot}} \iint_i |f(x, z)|^2 dx dz \quad (5.26)$$

where the second integral effectively produces a constant factor that converts the electric field at a single point of the center line into energy per unit length along y . This factor is equal to $7.5 \times 10^{-15} \text{ m}^2$ for the typical lead width of $1 \mu\text{m}$.

Ignoring the junction For typical transmon junction-lead geometries, the calculated energy participation within several μm of the junction (1) is smaller than that across the capacitor pads, usually contributing at most half of the total participation, and (2) varies negligibly across geometric variations of the larger structure. In fact, the largest portion of this contribution comes from the region up to $\sim 10 \text{ nm}$ from the junction. The cumulative MS participation is plotted in Figure 5.34 as a function of distance from the Josephson junction. Because of its geometric invariance, adding this contribution would only reduce the transmon lifetimes of all designs equally.

But what physical argument justifies the dismissal of a possible contribution to loss? One can expect

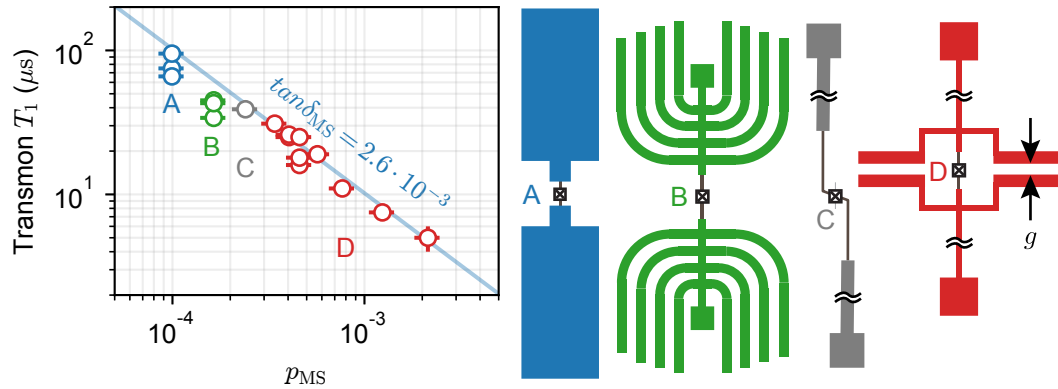


Figure 5.35 | **Transmon lifetime for varied geometry** Measured transmon T_1 (normalized to 6 GHz) for four different transmon geometries. Each is measured in a ‘V2’ rectangular 3D cavity and coupled weakly enough to avoid Purcell limitation. Transmon type A is of the standard 3D variety. Transmon type B is of the ‘spidermon’ type, which maximizes pad perimeter. Transmon type C is a lengthy ‘tiltmon’, with pads distributed across ~ 2 mm and tilted so as to fit on the chip. Transmon type D, the ‘gapmon’, derives most of its capacitance from coplanar-style pads with a varied gap width g . (Figure adapted from [163]; see [Copyright Permissions](#).)

the defects that form two-level systems, thought to dominate dissipation in our oxides, occur with some regular distribution and density. This oxide should be similar in many respects to the oxide that forms the Josephson junction dielectric barrier. Former studies, however, have demonstrated that this junction dielectric (for small junctions, like ours) has a very high quality [124, 185]. These observations help to justify why participation from the hundred-nanometer-scale segment nearest the junction can be ignored at present performance levels.

5.7.2 Transmon performance with geometric variations

To evaluate the loss tangent limitations of transmon interfaces, and compare them with those of quasi-planar coaxline resonators, we probe the energy relaxation times of four geometric designs of transmons, shown in Figure 5.35. The participation ratios in the assumed interfaces (where we also take $\epsilon_d = 10$, $t = 3$ nm for consistency) are simulated according to the above procedure. The measured T_1 values display a strong correlation to the simulated participation ratio of the metal–substrate interface layer, p_{MS} . However, Figure 5.36 shows that other interface and surface layers vary similarly with geometry. However, participation of the bulk substrate dielectric (in this case, sapphire), as well as participation of the package conductor and dielectric layers, does not vary with the observed dependence. In this study [163] and others

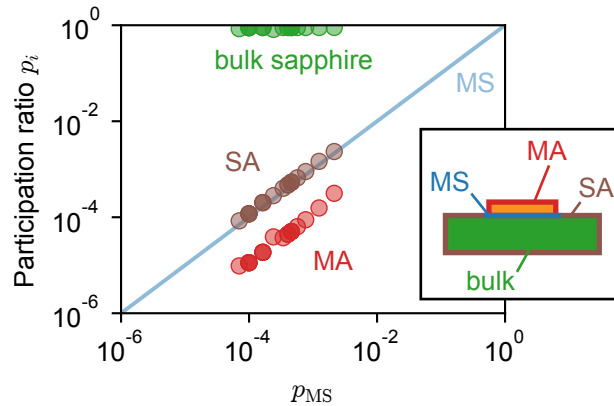


Figure 5.36 | **Comparative scaling of surface losses with geometry** With the variations in transmon geometry from Figure 5.35, metal–air and substrate–air participations scale with metal–substrate participation, while bulk substrate participation does not show this dependence. This suggests that performance limitations in transmons is not limited by bulk, but makes it difficult to immediately determine which interface layer dominates the observed loss. (Figure adapted from [163]; see [Copyright Permissions](#).)

[97, 154, 186, 187, 188, 189], it can be difficult to determine which of the three surfaces are the dominant contributors based on these data alone, because the three loss tangents cannot be extricated from the total measured loss that changes with transmon geometry. However, we can use ancillary observations to help pinpoint the primary culprit.

First, one might assume that the chemistry associated with the aluminum film could lead to similar composition (and thus material quality) of the MA and MS layers. However, the MA layer is naturally inclined to have a participation ratio between ϵ_d and ϵ_d^2 times smaller than that of the MS layer. For the MA to dominate the loss, or at least contribute equally, it would have to have $\tan \delta_{MA} \gtrsim \epsilon_d \tan \delta_{MS}$. Further, one would expect the MA layer to be particularly dependent on the post-fabrication environment. However, repeated measurements of transmons, which would reflect degradation over time, do not vary significantly. More directly, special processing designed to remove the MS interface layer has shown an improvement that was not designed to improve the other layers [109, 164].

Conductor participation, which cannot be accurately simulated using this two-step method (because the 2D simulation of exact edge geometry is an electrostatic one), can still be approximately simulated to obtain its dependence on geometry. From geometric considerations, we expect it to vary similarly to the other interface layers; simulations confirm that it does. The kinetic inductance fraction of the junction (the geometry of which remains fixed), however, dominates significantly over that of the pads (the geom-

etry of which changes) [67, 190, 191]. Further, one study that varied the geometric inductance of planar structures found no difference in performance at the level of $Q = 10^5 - 10^6$ [108]. The magnitudes of change in transmon T_1 can thus not satisfactorily be explained by changing pad kinetic inductance alone.

The SA layer is calculated to have a participation ratio similar to that of the MS layer, and could be thought to contribute equally (Figure 5.36). The same dielectric arguments apply, however: surface quality is much more likely to degrade after fabrication than the MS interface quality. This remains to be tested more rigorously, perhaps by processing designed to target and worsen the material quality of the SA layer.

In fact, producing disparate scaling relationships between the different layers is difficult to do with geometric variations alone. To adequately separate the layers, significant changes in processing or device topology will be required. For example, processes like silicon deep reactive-ion etching (DRIE) can scale the MS and MA participation disproportionately compared to bulk and SA participation (Section 5.5.5). Multi-mode designs can also serve to adjust participation in a non-uniform manner (Section 5.5.2). Finally, special chemistry and processing steps, such as annealing, passivation layers, or changes of material, can be used to directly test the effect on resonator and transmon performance (Section 5.5.4). Future studies like these in controlled environments will be the best way to disentangle these related loss mechanisms.

5.7.3 Quasiparticle sensitivity

Discussion of the transmon's sensitivity to conductor loss has so far been omitted. Much like the surface of 3D cavities and quasi-planar resonators, the transmon structure sees dissipation related to its kinetic inductance fraction. Unlike the other structures, however, the transmon has a Josephson junction. The Josephson junction confers additional sensitivity to the transmon, leading to decoherence or energy relaxation [192]. This added relaxation rate is roughly proportional to the fraction of quasiparticles present (which, in our system, will be all non-equilibrium quasiparticles) [193].

The expected quasiparticle fraction does not depend significantly on geometry (compared to the effect derived from the junction), a fact that helps to justify our focus on dielectric layers in this section. It also supports an interesting assumption: if the non-equilibrium quasiparticle densities are similar in transmons and striplines, then the relatively long lifetimes of transmons should translate to lifetimes in striplines enhanced by a factor of $\sim \alpha_{\text{transmon}}/\alpha_{\text{stripline}}$. The absence of this correspondence in coaxline

Table 5.6 | **Summary of material bounds.** Simulated and measured values for system dissipation sources. Values of Q_{material} are lower bounds and are obtained from single-photon measurements of similar systems. Our set of participations $\{p_k\}$ are obtained from simulation, except for conductor participation (where $p_k = \alpha$, kinetic inductance fraction), which is directly measured (Section 5.5.3). The calculation of our Q_{material} comes largely from the best measured values of Q_i . (Table adapted from [7]; see [Copyright Permissions](#).)

Participation of region	Established Q_{material} bound	Established source	Our Q_{material} bound	Our bounding device
p_{sapphire}	$1-5 \times 10^6$	3D transmons [66, 169], WGMR [153]	$(5.3 \pm 0.5) \times 10^6$	coaxline
p_{silicon}	$0.13-2.7 \times 10^6$	3D transmons [97, 187]	$(0.16 \pm 0.02) \times 10^6$	coaxline
p_{MS}	380	3D transmons [163]	160 ± 20	coaxline
p_{MA}	38	3D transmons [163]	115 ± 25	coaxline
p_{SA}	460	3D transmons [163]	190 ± 40	coaxline
$p_{\text{cond,strip}}$	4,800–110,000	Al WGMR [194, 195]	$63,000 \pm 20,000$	coaxline
$p_{\text{cond,encl}}$	4,400	cylindrical cavity [93], coaxial stub cavity [67]	195 ± 30	coaxline
$p_{\text{diel,encl}}$	22	coaxial stub cavity [67]	360 ± 70	two-strip coaxline

devices supports our conclusions throughout this chapter: more than one source of dissipation is likely responsible for the variable Q in quasi-planar devices housed in 3D enclosures.

5.8 Summary of loss mechanisms and bounds

In this section, I analyze and compare the results from all studies performed and reviewed in this chapter. Discussion of the results is divided according to the categories of loss mechanisms that have been modeled, most of which are summarized in Table 5.6. The experiments producing the highest bound in this study, as well as in the literature, are identified for each mechanism.

Bulk dielectric The transmon mode participates strongly in the bulk dielectric substrate, typically at a level around 0.9 for our substrates, as discussed in Section 5.2.1. The coaxline, with its sapphire participation generally < 0.5 , is still able to set quite a high bound on Q_{sapphire} compared to contemporary work in similar systems, due to its performance more than a factor of two greater than the best transmons.

We investigated the effect of different types of sapphire substrate, and found little correspondence in

general (Figure 5.24). Other studies have drawn more certain conclusions [154]; if bulk substrate quality does play a role here, it is probably obscured by other effects. We suspect that the type of silicon used is of great importance, however. Silicon is more sensitive to doping via handling and processing, and can vary drastically from batch to batch. Other groups have shown significantly higher performance for transmons on silicon [187].

Studies are underway to set even higher bounds on bulk dielectric loss tangents. By inserting a large quantity of raw dielectric into a high- Q cavity using a movable stage, its participation ratio can be adjusted *in situ*. This technique shows the potential for measuring loss tangents of materials like sapphire in the range of $\tan \delta < 10^{-10}$. Such resolution would much more definitively bound loss tangents, perhaps even as a useful tool on a wafer-by-wafer basis.

Deposited conductor quality The quality of deposited aluminum is generally bounded at a higher value than the bulk aluminum from which cavities are machined. This likely indicates high material purity and careful treatment of the growth conditions, generally within an electron beam deposition tool.

The coaxline sets among one of the highest bounds for this conductor quality. Since the kinetic inductance fractions here can be directly extracted from swept temperature measurements, limits set by such means are especially robust.

Bulk conductor quality Coaxline resonators do not “sample” the bulk conductor much more than dielectric-free 3D cavities, but their quality factors are significantly more limited—almost an order of magnitude difference between best-case values. Therefore, we do not expect to set new limits using these devices. Likewise, we did not observe any change in distribution of quality factors based on the purity of the package material (Figure 5.24).

Dielectric interface layers Loss induced by dissipative dielectric interface layers has been established as a dominant source in many of our planar systems, particularly transmons. Useful comparisons between designs and materials by plotting (in Figure 5.37) their energy lifetimes compared to simulated participation of a metal-substrate interface layer. In the limit where this contribution is singular, the bounds set from this analysis become accurate values for the quality of the interface material.

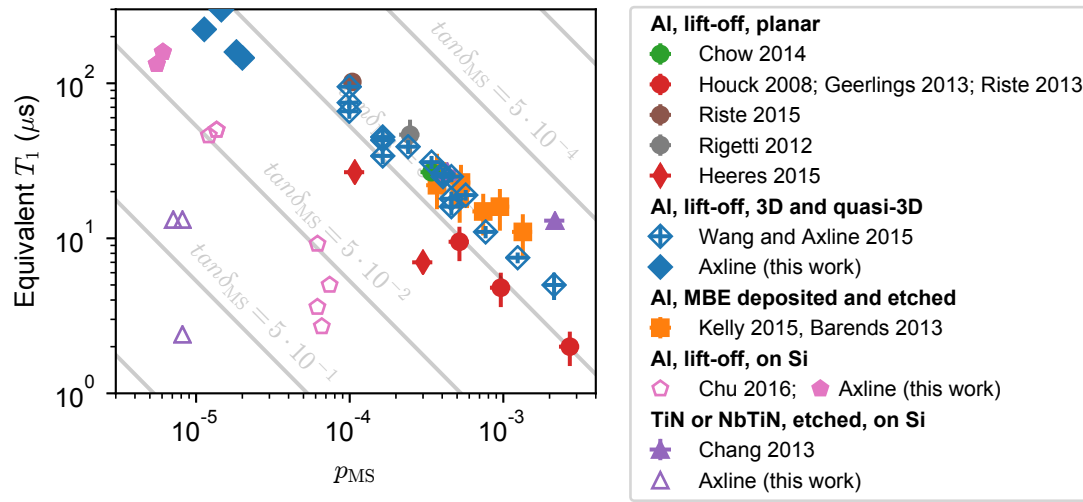


Figure 5.37 | **Measured lifetimes of assorted devices, compared to metal–substrate interface participations.** A comparison of transmon device performance from literature, including those in this work. All of the devices in “this work” are stripline resonators measured in coaxline, coaxmux or seammux packages. (Figure adapted from [163]; see [Copyright Permissions](#).)

In Figure 5.37, we include simulated participations for numerous past and present planar devices, including transmons, CPW, IDC, and CPC structures. Structures in literature are simulated as well as possible given the provided dimensions. We mostly plot “best-in-class” devices from this chapter, which maximize the product of Q and p . We include some devices that deviate from the norm of aluminum deposited on a sapphire substrate. For comparison, then, we provide devices in literature made with similar techniques.

While transmons still produce the most stringent bounds on dielectric interface layers, the coaxline architecture allows for some strategic adjustments to help differentiate the layers. Multiple striplines patterned adjacent to or opposite one another, or etching enabled by silicon, allow for adjustments to participation ratio that are not accessible through continuous changes to geometry. As a result, coaxlines are able to place a slightly higher (but unsurprising⁵) bound on the metal–air interface layer.

5. It is unsurprising because we should expect these materials to be similar, and thus their loss tangent bounds to be related within much less than an order of magnitude.

5.8.1 Consensus

It is difficult to state with certainty that one loss mechanism should be the target of our corrective efforts. Rather, it seems that most of the mechanisms modeled here, as well as some that were not, can play a role in various circumstances.

It is important to be constantly vigilant and aware of the mechanisms that may be present. Even if materials or interfaces are improved, their bounds only recede somewhat into the shadows. Improvements will continue, and these monsters will rear their ugly heads again, in time.

6

Letting Schrödinger’s Cat Out of the Bag

A modular quantum network architecture requires communication between the nodes of the network. Before two-way communication is established, however, we must address the more fundamental task: conversion between stationary information, retained in the storage mode of the node, and a propagating form of that information. Using these locally stored states as resources in an error-protected, network-based quantum information processing architecture hinges on the ability to interface them with traveling signals.

In this chapter, we will develop a parametric process to transmit these local states. To maintain generality, this transfer must be faithful for arbitrary states. This will later permit non-classical, multi-photon encodings—like Schrödinger cat states—to be used for error-correction schemes. Furthermore, temporal control over the process could allow states to be simultaneously launched *and* not-launched—yet another Schrödinger cat-style superposition. Our real-life node will also suffer from losses, imbuing the conversion timescale with particular significance.

Quantum entanglement is a valuable resource, particularly in a modular quantum network. Another goal, beyond complete transmission of a quantum state, will be to distribute entanglement throughout a network. Since this chapter focuses on the conversion between stationary and propagating modes, we will show the natural realization of this idea for our pre-network system: entanglement between stationary and propagating qubits.

In the sections that follow, we will outline the hardware and physics required to permit this conversion. Then, we will characterize the result, using many of the state-of-the-art circuit QED tools introduced in

Chapters 3 and 4. We will show that this conversion process is the first step of a computational, module-based network in circuit QED, an idea that will segue into a demonstration of a completed, elementary network in Chapter 7.

6.1 Parametric pumping

To convert stationary memory states to propagating ones, we propose a parametric, four-wave-mixing-based pumping scheme using the transmon as a conversion element. Understanding the required modes and classical drives will inform our choice of hardware to be included in our basic module.

Several approaches have been used so far to map stationary onto traveling states in superconducting quantum circuits. Tuning the coupling between a superconducting artificial atom and an output mode allows the generation and shaping of single photons [110, 196, 197, 198]. However, the controlled mapping of general multi-photon states between a CV quantum memory and traveling signals has so far remained an outstanding challenge.

Quasi-classical oscillator states have been successfully mapped using dedicated coupling elements such as flux-tunable couplers [199, 200] or parametric converters [201], which require tuning by external magnetic flux. However, to date, we have not yet integrated flux-tuning methods with the kind of high- Q storage cavities we intend to use for synthesizing, storing, and protecting complex quantum states. For the duration of this work, we hold fast to the philosophy endorsed in Section 4.3.1, and avoid coupling elements that may damage the hard-won coherence of elements in our system.

Instead, we aim to couple a storage cavity mode, \hat{a} , and an output mode, \hat{b} , using a nonlinear element that enables photon conversion between them. Since the fixed-frequency transmons in our toolbox are a relatively long-lived and powerful type of qubit, we spurn a flux-tunable design and choose to couple modes \hat{a} and \hat{b} using only the transmon “artificial atom” in the strongly dispersive regime of cavity QED [130]. The single Josephson junction of the transmon provides the required nonlinearity for such conversion, while preserving cavity coherence on the order of milliseconds [46]. These three elements—cavity, transmon, and output mode—thus form the basis for a module capable of this kind of conversion. They will be further detailed in Section 6.2.

6.1.1 Classical driving

We will use classical pumps to invoke the desired parametric interaction. First, let us examine the effect of a single classical drive on the transmon.

We begin with the Hamiltonian of the transmon, given by the quantized LC oscillator including the Josephson junction,

$$\hat{H} = \hbar\omega_t \hat{t}^\dagger \hat{t} - E_J (\cos(\hat{\varphi}_t) - \hat{\phi}_{\text{lin}}) = \hbar\omega_t \hat{t}^\dagger \hat{t} - E_J (\cos(\varphi_t(\hat{t}^\dagger + \hat{t})) - \hat{\phi}_{\text{lin}}) \quad (6.1)$$

where φ_t is the flux through the junction in the transmon mode \hat{t} (Equation 4.7). Since the linear part of the \hat{H} is already captured by $\hbar\omega_t \hat{t}^\dagger \hat{t}$, we remove this portion from the cosine ($\hat{\phi}_{\text{lin}} = 1 - \hat{\varphi}^2/2$). Though it appears convoluted, this separation of terms is particularly useful for characterization in the black-box formalism [102].

Without the influence of dissipation, the Hamiltonian of a driven transmon is

$$\hat{H} = \hbar\omega_t \hat{t}^\dagger \hat{t} - E_J (\cos(\varphi_t(\hat{t}^\dagger + \hat{t})) - \hat{\phi}_{\text{lin}}) + \epsilon(t)(\hat{t}^\dagger + \hat{t}), \quad (6.2)$$

where $\epsilon(t)$ is the time-dependent amplitude of an external drive that is coupled in via flux across the junction. To cast this into a form that describes the drive-based mixing behavior of the transmon, we will transform this expression into a rotating, displaced frame [202]. This transformation explicitly seeks to eliminate the driving term $\epsilon(t)(\hat{t}^\dagger + \hat{t})$.

The transformation of the state $\hat{\rho} \rightarrow \tilde{\rho}$ is achieved by defining a unitary $\hat{U} = \hat{U}_2 \hat{U}_1$ that satisfies $\tilde{\rho} = \hat{U} \hat{\rho} \hat{U}^\dagger$ and acts as $\hat{U} |\psi\rangle \rightarrow |\tilde{\psi}\rangle$. By putting the transformed state through the time-dependent Schrödinger equation, its Hamiltonian \tilde{H} is related to the original Hamiltonian \hat{H} as

$$\tilde{H} = i\hbar \dot{\hat{U}} \hat{U}^\dagger + \hat{U} \hat{H} \hat{U}^\dagger. \quad (6.3)$$

The two unitaries are \hat{U}_1 , a displacement transformation of the form $\hat{D}(\xi(t))$ introduced in Section 3.5.2, and \hat{U}_2 , a rotation transformation, given by $\exp(i\omega_t t \hat{t}^\dagger \hat{t})$. Importantly, we treat the drive ξ classically and in the ‘stiff pump’ regime, such that it is not subject to zero-point fluctuations and can be

taken as a complex scalar variable (rather than an operator). To satisfy the purpose of the transformation, we choose $\xi(t)$ such that the driving term in the resulting transformed Hamiltonian disappears (Appendix A.2).

Applying the displaced frame transformation leaves us with an ordinary differential equation for ξ ,

$$\dot{\xi} = -i\epsilon(t) - i\omega_t\xi, \quad (6.4)$$

which is evidently a driven harmonic oscillator. Since the solution depends on the drive, we can gain some intuition in the simplest case, a constant $\epsilon(t) = \epsilon \exp(-i\omega_d t)$, turned on at $t = 0$. Solving Equation 6.4 gives us

$$\xi(t) = \frac{\epsilon \left(-e^{-i\omega_t t} + e^{-i\omega_d t} \right)}{\Delta}, \quad (6.5)$$

where $\Delta = |\omega_d - \omega_t|$ is the detuning of the drive from resonance. Still as a function of our time-dependent drive $\xi(t)$ from Equation 6.4, our transformed Hamiltonian appears as

$$\hat{H}_1 = \hbar\omega_t \hat{t}^\dagger \hat{t} - E_J \cos \left(\varphi_t (\hat{t}^\dagger + \hat{t} - \xi(t) - \xi^*(t)) \right); \quad (6.6)$$

if desired, the definition of ξ can be reversed so as to make all coefficients positive.

A realistic system will suffer from dissipation, altering these equations of motion. In particular, such drives are typically introduced via some dissipative resonant mode. Therefore, we safely choose to model this dissipation as pure photon loss, the dominant error source in a harmonic oscillator [53]. This behavior will be crucial in accurately modeling any conversion process. One way of including dissipation, in particular the decay rate κ_b of an oscillator mode, is to use the Lindblad equation instead of simply transforming the Hamiltonian as above [203].

The time dynamics in the displaced frame are given by $\dot{\hat{\rho}}_1 = \hat{U}_1 \hat{\rho} \hat{U}_1^\dagger + \hat{U}_1 \hat{\rho} \hat{U}_1^\dagger + \hat{U}_1 \hat{\rho} \hat{U}_1^\dagger$, where the Lindblad equation specifies $\dot{\hat{\rho}} = -i[\hat{H}, \hat{\rho}] + \frac{\kappa_t}{2} D[\hat{t}](\hat{\rho})$. The resulting expression for the drive should not be surprising, as it is similar to Equation 6.4 except for now incorporating dissipation:

$$\dot{\xi} = -i\epsilon(t) - \left(\frac{\kappa_t}{2} + i\omega_t \right) \xi. \quad (6.7)$$

In the case, again, of a constant drive, the solution to Equation 6.7 is

$$\xi(t) = \frac{\epsilon \left(-e^{-i\omega_t t - \kappa_t t/2} + e^{-i\omega t} \right)}{i\kappa_t/2 + \Delta}. \quad (6.8)$$

Finally, we apply the rotating drive transformation, \hat{U}_2 . With respect to our rotated mode, $\tilde{t} = \hat{t}e^{-i\omega_t t}$, the driven–dissipative, rotated, displaced form of the Hamiltonian becomes

$$\tilde{H}_{12} = -E_J \cos \left(\varphi_t(\tilde{t} + \tilde{t}^\dagger - \xi(t) - \xi^*(t)) \right). \quad (6.9)$$

Expanding the cosine within this Hamiltonian produces

$$\begin{aligned} \tilde{H} &= -E_J \cos \left(\varphi_t(\tilde{t}^\dagger + \tilde{t} + \xi + \xi^*) \right) \\ &\approx -E_J \sum_{k,l,m,n} \frac{(-1)^{(k+l+m+n)/2} \varphi_t^{k+l+m+n}}{k!l!m!n!} \xi^k \xi^{*l} \tilde{t}^\dagger{}^m \tilde{t}^n, \end{aligned} \quad (6.10)$$

where the sum is taken over all $(k + l + m + n) \bmod 2 = 0$. The frequency at which any given term in the expansion oscillates is given by $(-k + l)\omega_d + (m - n)\omega_t$. We can make the usual rotating wave approximating by neglecting all terms for which this frequency is non-zero.

6.1.2 Four-wave mixing and memory–output conversion

We will use this formalism of a classical drive on a transmon—namely, Equation 6.10—to select terms of the expansion that produce a desired effect. Because the cosine-coupling enables all four-wave mixing processes that conserve energy, we can create interactions between the strongly detuned resonator modes by applying pump tones. We wish to implement a conversion Hamiltonian of the form

$$H/\hbar = g\hat{a}\hat{b}^\dagger + g^*\hat{a}^\dagger\hat{b}, \quad (6.11)$$

which will convert photons between mode \hat{a} and mode \hat{b} , with rate g . In our experimental realization, these will be the memory and output modes, respectively (Figure 6.1). To ensure that the state swapped from the memory to the output mode does not return to the memory, the output mode is strongly coupled




Figure 6.1 | **Launching a quantum state from a node.** A schematic of two modular network node modes, \hat{a} and \hat{b} , that are interfaced using a nonlinear conversion element \hat{t} with temporally controlled rate $g(t)$. The mode \hat{b} is connected to an external transmission line mode \hat{b}_{out} with rate κ_{out} . (Figure adapted from [86]; see [Copyright Permissions](#).)

to a transmission line. In the parlance of cavity QED, this module can be understood to contain two orthogonal Fabry-Perot cavities that are coupled by an (artificial) atom.

This coupling naturally effects conversion of the state from a “stationary” resonant mode to a propagating wavepacket. In addition to the two modes a and b , two driving modes will be necessary to produce this interaction. The full Hamiltonian, including these two drives, is

$$\begin{aligned}
 H = & \hbar\omega_a \hat{a}^\dagger \hat{a} + \hbar\omega_b \hat{b}^\dagger \hat{b} + \hbar\omega_c \hat{t}^\dagger \hat{t} - E_J \left(\cos(\hat{\varphi}) - \frac{\hat{\varphi}^2}{2} \right) \\
 & + \epsilon_1(t) e^{-i\omega_1 t} (\hat{a}^\dagger + \hat{a}) + \epsilon_2(t) e^{-i\omega_2 t} (\hat{b}^\dagger + \hat{b}),
 \end{aligned} \tag{6.12}$$

where $\epsilon_{1,2}$ are the drive amplitudes. Moving into a displaced, rotating frame as was done above for the single drive, and choosing the convention in which the drives have positive sign, we can rewrite this as

$$\begin{aligned}
 H = & -E_J \cos(\phi_a(\tilde{a}^\dagger + \tilde{a} + \tilde{\xi}_1^*(t) + \tilde{\xi}_1(t)) \\
 & + \phi_b(\tilde{b}^\dagger + \tilde{b} + \tilde{\xi}_2^*(t) + \tilde{\xi}_2(t)) \\
 & + \phi_c(\tilde{t}^\dagger + \tilde{t})).
 \end{aligned} \tag{6.13}$$

The rotating drives $\tilde{\xi}_i$ for modes $i = a, b$ are related to the non-rotating drives (Equation 6.8) and their classical displacement amplitudes according to

$$\tilde{\xi}_i(t) = \xi_i(t) e^{-i\omega_i t} = \frac{\epsilon_i(t) e^{-i\omega_i t}}{i\kappa_i/2 + \Delta_i} \approx \frac{\epsilon_i(t) e^{-i\omega_i t}}{\Delta_i}. \tag{6.14}$$

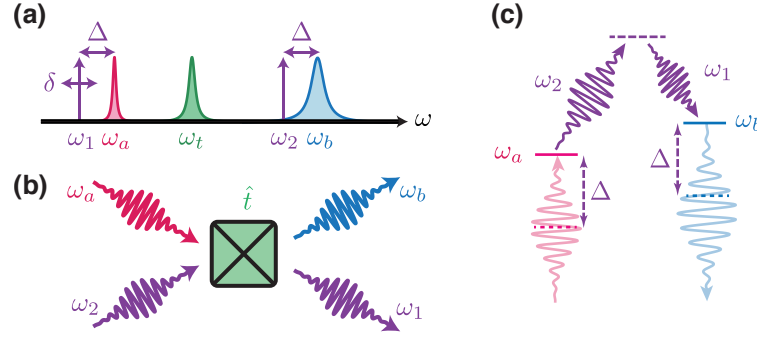


Figure 6.2 | **Frequency ordering in a conversion experiment.** **a** The mode and pump arrangement in frequency space. **b** The nonlinear element “mixes” four waves, preserving the energy conservation relation $|\omega_1 - \omega_2| = |\omega_a - \omega_b|$. **c** This mixing can also be shown in a form akin to a Raman transition, although it does not make use of a real or virtual transition of a higher mode of the system. The detuning Δ indicates the frequency difference $\omega_a - \omega_1 = \omega_b - \omega_2$, as indicated in the text; detunings here are exaggerated for effect and do not represent true frequencies. (Figure adapted from [23]; see [Copyright Permissions](#).)

Here, we have separated the slowly varying envelope of the displacement, $\xi_i(t)$, from its rapidly oscillating component. The detuning Δ_i is given by the detuning of the drive frequency with respect to the closest mode to which the drive is applied, i.e., $\Delta_1 = \omega_1 - \omega_a$ and $\Delta_2 = \omega_2 - \omega_b$. κ_i is the damping rate of the mode to which the pump is applied. A complete representation of the drives would include the fact that the classical drives excite the transmon through many modes simultaneously. However, the relative detuning of a drive from its “host” mode is small compared to the overall detuning between the mode frequencies; therefore, this effect can largely be ignored. The frequency ordering, relative spacing of modes, and how they mix is shown in Figure 6.2.

To obtain a more intuitive form of the conversion Hamiltonian containing parameters that can be directly measured, we apply the rotating wave approximation (RWA) to Equation 6.14 by expanding the cosine up to orders of φ_k^4 and retaining all terms that conserve energy. Besides the terms that are already present in the undriven case, this reveals the additional terms

$$\begin{aligned}
 H_1/\hbar = & |\xi_1|^2 \left(2\chi_{aa}\hat{a}^\dagger\hat{a} + \chi_{ab}\hat{b}^\dagger\hat{b} + \chi_{at}\hat{t}^\dagger\hat{t} \right) \\
 & + |\xi_2|^2 \left(2\chi_{bb}\hat{b}^\dagger\hat{b} + \chi_{ab}\hat{a}^\dagger\hat{a} + \chi_{bt}\hat{t}^\dagger\hat{t} \right) \\
 & + g\hat{a}\hat{b}^\dagger e^{-i\delta t} + g^*\hat{a}^\dagger\hat{b}e^{i\delta t}.
 \end{aligned} \tag{6.15}$$

The first two lines in this expression correspond to Stark shifts that will influence the resonant condition of

the conversion process. However, we can compensate for them by detuning the drives. The last line is the desired conversion Hamiltonian, where $g = \chi_{ab}\xi_1^*\xi_2$. The variable $\delta = \Delta_1 - \Delta_2$ denotes a relatively small detuning from the conversion resonance. This derivation assumes that the pumps are tuned such that the conversion process is close to resonant, and that all other terms are fast-oscillating, vanishing within the RWA.

6.1.3 Stationary-to-propagating conversion and Q-switching

Using input–output theory [113], we approximate the coupling to be constant within the bandwidth of our process and use the Heisenberg equation of motion (EOM) on the field operator for mode \hat{a} , as

$$\partial_t \hat{a} = -i [\hat{a}, H] = -ig\hat{b}, \quad (6.16a)$$

$$\partial_t \hat{b} = -i [\hat{b}, H] - \frac{\kappa_{\text{out}}}{2} \hat{b} = -ig\hat{a} - \frac{\kappa_{\text{out}}}{2} \hat{b}, \quad (6.16b)$$

$$\hat{b}_{\text{out}} = \sqrt{\kappa_{\text{out}}} \hat{b} \quad (6.16c)$$

where κ_{out} represents the coupling of \hat{b} to the transmission line. Since $\hat{b}_{\text{in}} = 0$, the output field \hat{b}_{out} is determined simply by $\hat{b}_{\text{out}} = \sqrt{\kappa_{\text{out}}} \hat{b}$. In solving the EOM, we consider a Hamiltonian that includes only the conversion term in Equation 6.15,

$$H = g(\hat{a}^\dagger \hat{b} e^{-i\delta t} + \hat{a} \hat{b}^\dagger e^{i\delta t}). \quad (6.17)$$

We ignore any internal loss, retaining only coupling to the transmission line, and we disregard input modes $\hat{a}_{\text{in}}, \hat{b}_{\text{in}}$ because they are in the vacuum state during the release. For simplicity, we take g to be real; however, it can be shown that the phase of this coupling term is simply mapped onto the output field. We also assume that the g -moderated coupling rate dominates over the intrinsic cavity loss rate κ_0 , as well as any off-resonant cross-coupling between \hat{a} and \hat{b} .

Our ultimate goal is to establish conversion rate between the storage cavity, \hat{a} , and a propagating wavepacket in \hat{b}_{out} . This should appear like an effectively enhanced decay rate from \hat{a} . Because this has the ability to dynamically “switch” the quality factor (Q) of the cavity from ‘high’ to ‘low’, this process is

sometimes termed a “ Q -switch”. Beyond conversion, it can also be used to deterministically reinitialize a cavity state to vacuum in a much shorter time than the intrinsic Q would typically allow.

To understand the behavior of an effective Q -switch, we can analytically solve the equations of motion in Equation 6.16, assuming some initial condition $\hat{a}(0)$:

$$\hat{a}(t) = \frac{\hat{a}(0)}{\beta} e^{-\frac{\gamma t}{4}} \left(\beta \cosh\left(\frac{t\beta}{4}\right) + \gamma \sinh\left(\frac{t\beta}{4}\right) \right), \quad (6.18a)$$

$$\hat{b}(t) = -i \frac{4g\hat{a}(0)}{\beta} e^{-\frac{t}{4}(\kappa_{\text{out}} - 2i\delta)} \sinh\left(\frac{t\beta}{4}\right), \quad (6.18b)$$

where we have introduced

$$\gamma = \kappa_{\text{out}} + 2i\delta, \quad (6.19a)$$

$$\beta = \sqrt{\gamma^2 - (4g)^2}. \quad (6.19b)$$

We first consider the case $\delta = 0$; i.e., the conversion drives make the process resonant. We are mostly concerned with the case $g \ll \kappa_{\text{out}}$; in this regime the coupling produces exponential damping of the cavity, which can be seen by approximating

$$\hat{a}(t) \approx \hat{a}(0) e^{-2g^2 t / \kappa_{\text{out}}}, \quad (6.20a)$$

$$\hat{b}(t) \approx -i\hat{a}(0) \frac{2g}{\kappa_{\text{out}}} \left(e^{-\kappa t/2} - e^{-\kappa_{\text{out}} t/2} \right), \quad (6.20b)$$

$$\hat{b}_{\text{out}}(t) = \sqrt{\kappa_{\text{out}}} \hat{b}(t) \approx -i \frac{\kappa}{2} \hat{a}(t), \quad (6.20c)$$

where $\kappa = 4g^2/\kappa_{\text{out}}$, a value we can now identify as our effective damping rate. We find that this is equivalent to the Purcell (spontaneous emission) rate for a resonant Jaynes-Cummings interaction enabled by the pumps.

It would be interesting to know by how much this rate changes if, say, experimental uncertainties or higher-order Hamiltonian terms detuned this process from resonance. In this case, we retain the term δ .




Figure 6.3 | **Realization of a basic circuit QED module.** Our architecture consists of a storage cavity mode \hat{a} , an output mode \hat{b} coupled to a transmission channel \hat{b}_{out} with a rate κ_{out} , and a transmon (green) that allows for a tunable coupling $g(t)$ between a and b . The drives ξ_1, ξ_2 can be introduced via the input couplers of their respective modes. This enables the desired conversion process, producing an effective decay rate of \hat{a} via κ_{out} . (Figure adapted from [23]; see [Copyright Permissions](#).)

For $\delta \neq 0$ we can approximate $\beta \approx \gamma - 8g^2/\gamma$. The evolution of \hat{a} is then

$$\hat{a}(t) \approx \hat{a}(0)e^{-\frac{t}{2}\frac{4g^2}{\gamma}} \equiv \hat{a}(0)e^{-\frac{t}{2}\gamma_{\text{conv}}}. \quad (6.21)$$

Since γ_{conv} is complex, we must evaluate its real part to obtain the line width of the conversion resonance.

We have

$$\gamma_{\text{conv}} = \frac{4g^2}{\kappa_{\text{out}} + 2i\delta} = \frac{4g^2(\kappa_{\text{out}} - 2i\delta)}{\kappa_{\text{out}}^2 + (2\delta)^2}, \quad (6.22)$$

the real part of which is a Lorentzian with full-width half-maximum κ_{out} as a function of detuning. Figure 6.4a shows the decay of the cavity versus detuning, and behaves as expected with linewidth κ_{out} .

Using the apparatus shown in Figure 6.3, we can measure $\kappa(\delta)$ by fixing Δ for one of the pumps and sweeping the detuning δ , thus affecting only one pump frequency. (Further experimental details will be given in Section 6.2.) At each point we measure the decay of the storage mode and extract κ via a single exponential fit. The data for two different values of g are shown in Figure 6.4a.

The excellent agreement of the data with a Lorentzian line shape indicates a resonant process. If this process is indeed the predicted conversion, we expect that its damping rate is given by the difference between the maximum of the Lorentzian fit and its offset, $\kappa - \kappa_{\text{ofs}}$. $\kappa_{\text{ofs}} = \kappa_{\text{loss}}$ is the damping due to any other losses, intrinsic or pump-induced. Hence, we expect that the inefficiency of the conversion is given by $\sim \kappa_{\text{loss}}/\kappa$. For $g/2\pi \gtrsim 100$ kHz we find $\kappa_{\text{loss}}/\kappa$ between 0.01 and 0.015.

If the condition $g \ll \kappa_{\text{out}}$ is not fulfilled, the decay of the storage mode is no longer exponential and

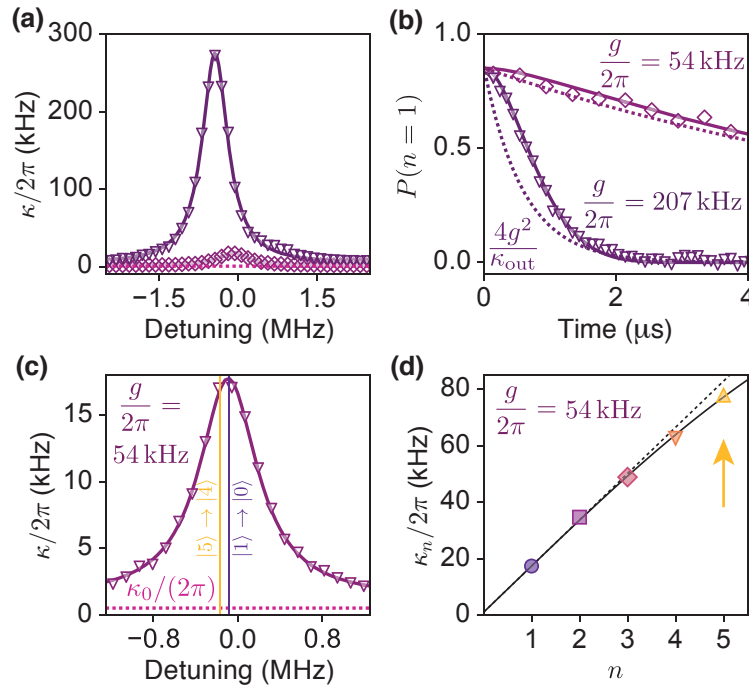


Figure 6.4 | **Cavity damping analysis.** **a** Cavity damping as function of relative pump detuning δ . The two data sets represent $g/2\pi = 54$ and 207 kHz. Each point is obtained by fitting individual decay curves with a single exponential, $\propto \exp(-\kappa t)$. Solid lines are Lorentzian fits. The resonance condition for the conversion process is offset from 0 because the pumps induce a Stark shift on the resonators. Note that for the larger g the decay rate κ is not a very good approximation because the damping is not purely exponential in time. **b** Non-exponential decay for large g , with pump frequencies on resonance with the conversion process. Solid lines: complete model, based on independent calibration of g . Dashed lines: $4g^2/\kappa_{\text{out}}$ approximation. **c** Reduction of κ_n due to the Kerr effect. Data and fit are for $g/2\pi = 54$ kHz. Orange and blue lines mark the resonance condition for one and five photon Fock states, respectively. Even relatively detuned from conversion resonance, the effective decay rate is much larger than the intrinsic loss rate κ_0 . The width of this resonance is κ_{out} , as given in Equation 6.22. **d** Decay of the Fock states $|n\rangle$. Dashed line: extrapolation from κ_1 . Solid line: Correction accounting for Kerr and the Lorentzian profile of the resonance. The deviation between the measured κ_5 (arrow) and $5 \times \kappa_1$ is 6%. (Figure adapted from [23]; see Copyright Permissions.)

can no longer be described accurately by a single decay rate κ . This can be seen in Figure 6.4b, where we show the damping of the storage mode for $g/2\pi = 54$ kHz and 207 kHz. At 207 kHz there is a clear deviation between the exact model and the approximation.

6.1.4 Conversion as a beamsplitter

The release from the storage cavity into the transmission line mode can also be described by an effective beam splitter interaction that does not explicitly take the output mode into account. This interaction is given by the unitary transformation

$$U_{\text{release}} = \exp\left(\frac{\theta}{2}(\hat{a}^\dagger \hat{b}_{\text{out}} - \hat{a} \hat{b}_{\text{out}}^\dagger)\right) \quad (6.23)$$

that acts on the cavity and propagating mode. For $g \ll \kappa_{\text{out}}$, the cavity decays exponentially with a rate $\kappa = 4g^2/\kappa_{\text{out}}$. In that case, the mixing angle θ is given by

$$\theta = 2 \arccos(\exp(-\kappa t/2)). \quad (6.24)$$

Stopping the conversion process at a finite time T , some portion of the original state would remain. For example, a 50:50 beam splitter interaction can then be achieved by a pump duration $T_{1/2} = \ln 2/\kappa$, corresponding with $\theta = \pi/2$. This has the effect of maximally entangling the remainder of the stationary state with propagating one, and will be important in demonstrating entanglement distribution later in this chapter. The next chapter will go far beyond this basic demonstration of temporal control.

6.2 Experimental setup

So far, we have only introduced a schematic of the hardware elements required for conversion in Figure 6.3. Our precise experimental scheme and hardware elements are shown in Figure 6.5a. Our hardware platform consists of the same few state-of-the-art elements introduced in Chapter 4: a 3D memory cavity (\hat{a}), a transmission qubit (\hat{t}), and an output mode (\hat{b}) that doubles as a readout resonator. Our 3D memory mode is an 99.99%-pure aluminum superconducting coaxial stub cavity, of which we use the fundamental $\lambda/4$ mode

(here, $\omega_a/2\pi = 4.1$ GHz), and the output mode is an overcoupled $\lambda/2$ stripline resonator (here, $\omega_b/2\pi = 10.0$ GHz). The stripline and transmon are fabricated on the same chip. The chip is inserted through a waveguide tunnel. Strongly undercoupled input pins (left and middle couplers) allow application of RF control and measurement tones, while signals leave \hat{b} to the transmission line (not shown) through an output coupler pin (rightmost coupler). The output mode is coupled with rate $\kappa_{\text{out}} = 1/240$ ns to a propagating transmission line mode, \hat{b}_{out} , where the emitted signals are amplified and recorded.¹ This configuration uses the hardware and integration schemes discussed in Chapter 4 to produce a long memory lifetime, $\kappa_0 = 1/450$ μ s, while still allowing for fast readout and control of arbitrary quantum states. Parameters of each mode, including frequencies and coherence times, are included in Table 6.1.

The sample is cooled to the base temperature $T_{\text{base}} \approx 15$ mK of a dilution refrigerator. A wiring diagram that shows how signals are introduced to the device is depicted in Figure 6.5b. Each mode of the system is addressed by a separate microwave generator acting as a local oscillator (LO) (cavity and output: *Agilent E8275D*; transmon: *Vaunix LabBrick LMS-103-13*); pulses are generated by IQ modulation (using *Marki Microwave IQ0618LXP* IQ-mixers). Importantly, we generate the pump tones using the same generators used for the (near-)resonant control pulses of the cavity and output mode. Given the conversion Hamiltonian, this guarantees that the signal emitted from the output mode is phase-locked to the output mode LO, which is used for mixing the signal down to low frequencies before digitizing.

Phase-locking is further assured in the following way: each driven mode is accompanied by a phase, which is determined by the quasi-random phase of the generator used to apply each drive. The local oscillator (LO) of the cavity mode \hat{a} is shared between drives and pumps, and the LO of the output mode \hat{b} is shared between drives, pumps, and signal demodulation. We apply the initial state phase ϕ_a , and the phase of the pumps, $\phi_{\xi,1}$ and $\phi_{\xi,2}$, to their respective operators in the Hamiltonian, described in further detail in the following section, as

$$H = g(\xi_1^* e^{i\phi_1} \xi_2 e^{-i\phi_2} \hat{a} e^{-i\phi_a} \hat{b}^\dagger) + \text{h.c.}, \quad (6.25)$$

1. Although the decay lifetime due to internal sources is difficult to directly measure in this overcoupled configuration, the analysis in Chapter 4 suggests the coaxline resonator should have an internal quality factor greater than one million. This gives a decay rate ratio between output and internal decay channels $\kappa_{\text{out}}/\kappa_i \approx 400$, such that internal dissipation contributes $< 1\%$ to overall dissipation in the module.

Hamiltonian term		Value (MHz)
Mode frequency	$\omega_b/2\pi$	9.999×10^3
	$\omega_a/2\pi$	4.073×10^3
	$\omega_t/2\pi$	6.674×10^3
Cross-Kerr	$\chi_{ab}/2\pi$	-0.013 ± 0.001
	$\chi_{at}/2\pi$	-3.825 ± 0.001
	$\chi_{bt}/2\pi$	-1.3 ± 0.1
Self-Kerr	$\chi_{aa}/2\pi$	-0.022 ± 0.002
	$\chi_{bb}/2\pi$	-0.001 ± 0.001
	$\chi_{tt}/2\pi$	-144.0 ± 0.5
Damping term		Value (μs)
Cavity energy decay	$1/\kappa_0$	450 ± 50
Output energy decay	$1/\kappa_{\text{out}}$	0.24 ± 0.02
Transmon qubit relaxation	T_1	50 ± 10
Transmon Ramsey decay	T_{2R}	25 ± 5
Transmon Hahn echo decay	T_{2E}	35 ± 5

Table 6.1 | **Conversion system Hamiltonian parameters.** Measured system parameters. See text for explanation. For the cavity and transmon qubit decay times, the uncertainties given are the typical fluctuations observed in the course of a day.

where $\hat{a}e^{-i\phi_a}$ represents the phase that has been gained during the state creation in \hat{a} . Since $\phi_a = \phi_{\xi,1}$ and $\phi_b = \phi_{\xi,2}$ and the output field is detected and demodulated using the \hat{b} generator, the phase becomes locked, and any phase difference from run to run has no effect.

The output signal is processed by a Josephson parametric converter (JPC) operated in amplification mode [204] (Section 3.5). The amplifier is configured to provide approximately 25 dB of gain with a bandwidth of approximately 15 MHz and a noise visibility ratio around 6 dB. This allows us to detect signals emitted at the frequency of the output mode with a detection efficiency of $\sim 45\%$ (details follow).

As depicted in Figure 6.5, an FPGA controls every aspect of an experimental run, and we take advantage of the tools introduced in Section 3.5. The FPGA performs feedback cooling to initialize the cavity and transmon as well as possible (in this case, we achieve a qubit population *not* in $|g\rangle$ of ~ 0.02 , and a cavity population of $\bar{n} < 0.01$). It decodes and thresholds results of transmon readout, allowing additional branching and postselection if necessary. This experimental setup also permits simple preparation of arbitrary states in the transmon and cavity using numerical optimization techniques (Section 3.5.1).

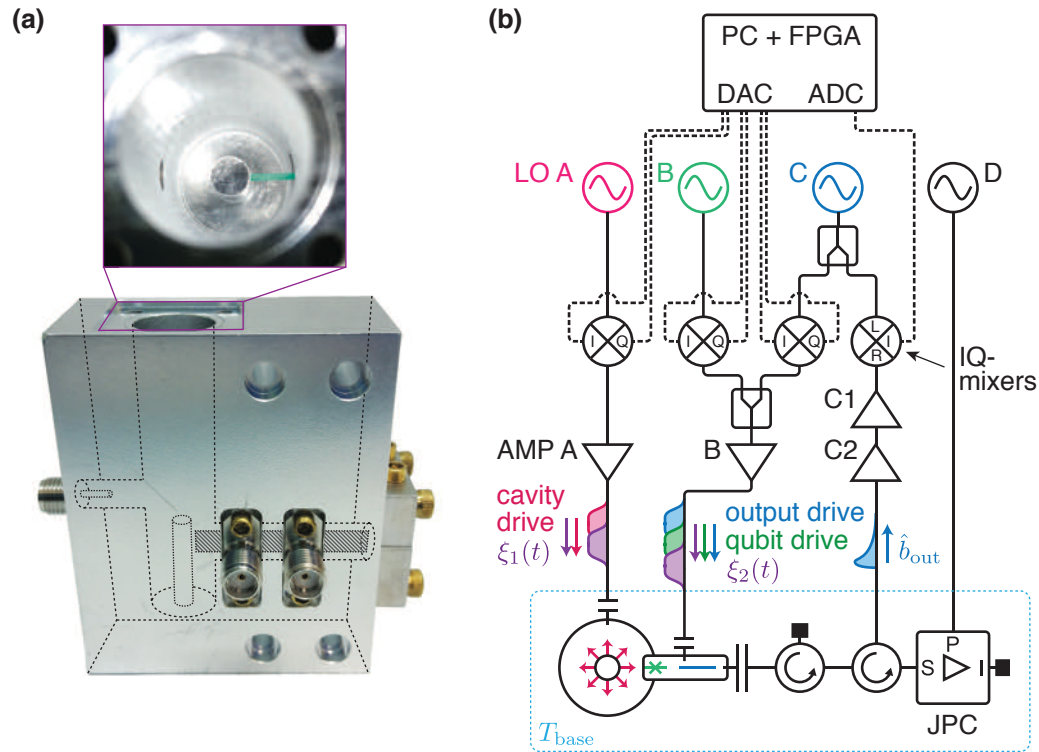


Figure 6.5 | **Experimental setup.** **a** Sample. The locations of the stub cavity as well as the coupler and chip tunnels are indicated by dashed lines; the position of the sapphire chip is indicated by the hatched rectangle. Cavity signals are applied through the SMA connector visible on the left, output mode signals enter and leave through the connectors at the front. The chip is held in place by a clamp on the right side. The inset shows a top-down view into the cavity, with the inserted chip false-colored in green. **b** Wiring diagram (schematic). Each mode has a microwave generator as local oscillator (LO). IQ modulation tones are synthesized by an integrated FPGA system with digital-to-analog converter (DAC) outputs (*Innovative Integration VPXI-ePC*), and mixed with the LO (LOs A, B, and C). Signals are amplified at room temperature (AMP A and B, *MiniCircuits ZVA-183-S+*) and sent into the refrigerator (*Oxford Triton 200*) where the sample is cooled to $T_{\text{base}} \approx 15$ mK. Input signals are applied through weakly coupled pins (depicted as small capacitors). Signals leave the output resonator through a strongly coupled pin (depicted as large capacitor) and is amplified by a JPC that is pumped continuously by a microwave generator (LO D, *Agilent N5183A*). The signal is further amplified at 4 K (AMP C1, *Caltech CIT1-4254-065*) and at room temperature (AMP C2, *Miteq AFS3-00101200-35-ULN*), mixed down with the output LO (LO C, *Marki Microwave IQ0618LXP*), and recorded and demodulated by the FPGA system via analog-to-digital converters (ADC). (Figure adapted from [23]; see [Copyright Permissions](#).)

6.2.1 Calibrations

Before the system can be characterized in earnest, a number of calibrations must be performed. These will determine (1) the strength of our pumps, in photon number, as a function of the DAC amplitude we have control over, (2) the strength and detuning of the conversion process with all pumps applied, and (3) the efficiency of our detector.

Stark shifts

The conversion rate is given by the product of the dispersive shift between storage and output modes and the pump strength, $g = \chi_{ab}\xi_1^*\xi_2$. To calibrate this rate, we need to determine the number of pump photons $|\xi_{1(2)}|^2$ in the storage (output) mode when applying the drives. This can be done by measuring the Stark shift of the transmon mode: when applying the pump tone on the storage (output) resonator, the Stark shift is given by $\delta\omega = \chi_{at(bt)}|\xi_{1(2)}|^2$. Since χ_{at} and χ_{bc} can be determined independently, measuring the Stark shift is a calibration for ξ_1, ξ_2 . For $\Delta/2\pi = -30$ MHz and -40 MHz we show the Stark shift calibration in Figure 6.6. We then use this pump calibration to obtain χ_{ab} . Applying a pump tone on the output mode \hat{b} results in a Stark shift of mode \hat{a} with magnitude $\chi_{ab}|\xi_2|^2$. Since ξ_2 is known, measuring the Stark shift yields χ_{ab} , and the value of g can then be calculated. We find that we can induce appreciable coupling for pump detunings $|\Delta|/2\pi \lesssim 50$ MHz with our experimental setup. As the coupling strength scales as $g \propto 1/\Delta^2$ (cf. (6.15)), larger detunings do not allow us to achieve the desired magnitude of g . To avoid resonant driving of the modes, we generally choose the maximal detuning with which we can achieve a given value of g . However, due to non-idealities in the system, larger detunings and stronger power may sometimes be more favorable (Appendix A.3).

Conversion rate

Though $\chi_{ab}\xi_1^*\xi_2$ predicts our conversion strength on resonance, identifying the exact resonance condition requires a separate experiment. Because the pump-induced Stark shifts are variable, they modify the frequency at which the conversion process is resonant. To probe this, we vary one pump detuning δ (within a range small compared to absolute pump detuning Δ) to identify the center and shape of the conversion process resonance. At each pair of pump amplitudes, we identify the maximum conversion rate $g(\xi_1, \xi_2)$,




Figure 6.6 | **Calibration of Stark shifts.** Pump strength calibration using the transmon's Stark shift. Applying pump tones, controlled by a DAC voltage, on storage (a) and output modes (b) results in a Stark shift of the transmon mode, measured by spectroscopy while applying the pump. Since the Stark shift is proportional to the photon number in the pump (which is proportional to the drive amplitude squared), we obtain a calibration for $\xi_{1,2}$. Solid lines are linear fits. (Figure adapted from [23]; see [Copyright Permissions](#).)

as shown in Figure 6.7. We fit these values to the expected form in Equation 6.22 to ultimately predict a rate $g(\xi_1, \xi_2)$ for any two pump strengths. This constitutes an extension of the procedure shown in Figure 6.4.

Detector efficiency

In order to later calibrate the efficiency of a propagating signal incident on our detector, we must understand how our type of measurement is affected by detector inefficiency. We will use a heterodyne detector to probe the outgoing field \hat{b}_{out} .

Our heterodyne detector intercepts the two quadratures of our signal, I and Q . We use the convention such that these are defined as $I = \frac{1}{2}\langle\hat{a} + \hat{a}^\dagger\rangle$ and $Q = \frac{-i}{2}\langle\hat{a} - \hat{a}^\dagger\rangle$, such that a complex displacement amplitude in this phase space is given by $\alpha = (I + iQ)/2$, and the value of α^2 corresponds with a photon number. Linear loss scales quasiprobability functions in this phase space (like the Wigner function, introduced in Section 3.5), adding additional fluctuations. Linear loss can be modeled as a fictitious beam splitter that operates on our signal mode, \hat{a} , and injects a variable amount of the vacuum state, \hat{a}_v , such that

$$\hat{a}' = \sqrt{\eta}\hat{a} + \sqrt{1-\eta}\hat{a}_v. \quad (6.26)$$

Here, η is the loss coefficient (or the transmissivity of the beam splitter). Quasiprobability distributions

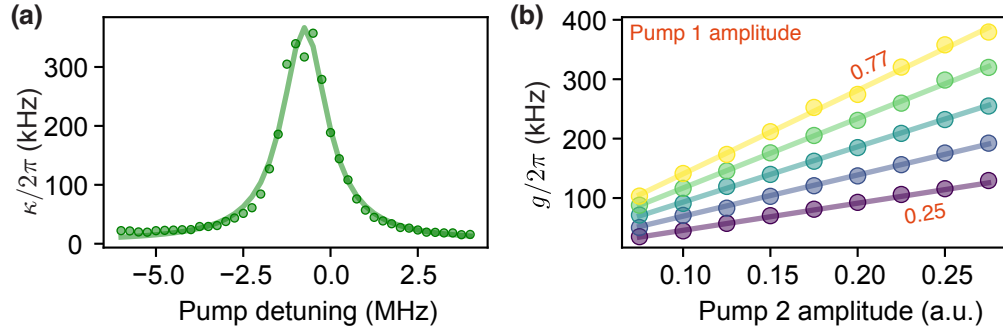


Figure 6.7 | **Calibration of conversion rate.** **a** An example measurement of cavity decay rate versus the frequency of one pump. Each point is from an exponential fit to the decay of the $n = 1$ population of the cavity. Solid line is a fit which gives the maximum decay rate κ . The conversion rate is then $g = \sqrt{\kappa_{\text{out}}\kappa}/2$. **b** Extracted g for various pump amplitudes (in fraction of the maximum amplitude sourced by the DAC). Pump 1 amplitude is varied from 0.25 to 0.77 in linear steps. Lines are a 2D linear fit to the entire data set with one fit parameter.

are transformed like

$$W(I, Q; s) \mapsto \frac{1}{\eta} W\left(\frac{I}{\sqrt{\eta}}, \frac{Q}{\sqrt{\eta}}; \frac{s + \eta - 1}{\eta}\right), \quad (6.27)$$

in terms of the generalized quasiprobability expression given in Equation 3.16. Working out this transformation applied to an incoming Wigner function undergoing such loss, we find that the resulting Wigner function is given by the convolution of the original quasiprobability function $W(I, Q)$ with the Gaussian parameterized as $\exp(-(I^2 + Q^2))/\pi$.

If the whole measurement apparatus has less than unit efficiency, $\eta < 1$, this clearly produces a smoothed quadrature distribution, and thus a smoothed Wigner function [137]. The measurement is the result of loss and added noise from the vacuum fluctuations mandated by quantum mechanics. A portion of this noise (half a quantum) comes from our heterodyne detection scheme and use of a phase-preserving amplifier.

We can determine the amount of noise introduced by our detector by first calibrating the axes of the quadratures it measures, by sampling vacuum (Figure 6.8). Next, we direct an independently-calibrated coherent state, with $|\alpha = 1\rangle$, to intercept the detector. The measured displacement then corresponds with the detector efficiency, η_{det} . The means by which the quadrature measurements are assembled and understood is described further in Section 6.3.3.

Having established a conversion efficiency of close to unity, the details of which will be discussed later,




Figure 6.8 | **Calibration of detector efficiency.** **a** Reconstruction of the Q-function of an incident vacuum state. The width of the expected 2D Gaussian profile of vacuum is used to calibrate both axes. **b** A single-photon coherent state is prepared and released from the output mode, becoming incident on the detector. As a reconstructed Q-function, it displays a sub-unity displacement; the best-case expected distribution would be centered at $I = 1$. The displacement on this scale of the square root of photon number thus corresponds with $\sqrt{\eta_{\text{det}}}$. (Figure adapted from [23]; see [Copyright Permissions](#).)

the position of the coherent state population in the Q-function is a direct measure of our detection efficiency. Assuming no loss in the conversion, we obtain $\eta_{\text{det}} = 0.44 \pm 0.03$, consistent with the expected performance of our parametric amplifier [204].

6.3 Quantifying conversion performance

To enable communication between network nodes with single or multi-photon states, a coherent release must meet several important criteria. First, in order to enable distribution of quantum information with high fidelity, we require a large ‘on–off ratio’. In the ‘off’-state the coherence of the memory must be preserved, while the ‘on’-state should allow fast, on-demand release. Further, successful communication requires faithful state mapping, independent of the number of photons. This means that the conversion Hamiltonian, introduced in Section 6.1, cannot discriminate against photon number (at least, within some operational range of interest). Finally, shaping the wave packet is required to enable capture by a receiving node [205] and to generate entanglement. While explicit, optimized, “impedance-matched” shaping will be instituted in Chapter 7, we only require here that our release process is capable of precise temporal control.

6.3.1 Evaluating evacuation completeness, on–off ratio

We first explore the maximum damping rate we can induce with pump tones of varying strength (Figure 6.9a). We prepare the Fock state $|1\rangle$ in the cavity, and then monitor the cavity population over time while applying the pumps with frequencies ω_1 and ω_2 . The pump frequencies are tuned on resonance with the conversion process of Equation 6.17 using $\omega_a - \omega_1 = \omega_b - \omega_2 = 2\pi \times (30 - 50)$ MHz. Increasing the pump strength allows us to tune the cavity energy decay rate from its intrinsic value of $\kappa_0 = 1/0.45$ ms to $\kappa \approx 1/0.5$ μ s for $g/2\pi = 207$ kHz, the maximum conversion rate achievable with the available pump power (Figure 6.9b). At this point, $g \approx 0.3 \times \kappa_{\text{out}}$, and the decay becomes limited by the bandwidth of the output mode.

This ‘Q-switch’ is very close to an ideal damping of the memory. It cools the cavity close to the vacuum, with a residual population of $\bar{n} \lesssim 0.01$, the noise floor of our measurement. We can therefore use the conversion as a fast reset, which is a useful tool for experiments with long-lived quantum memories [126, 206]. Further, the measured decay of the cavity population is in excellent agreement with predictions based on theory and independent calibrations of the pump strengths. This agreement, together with the absence of any significant heating in the system, suggests a very high conversion efficiency from the storage to the output mode. We estimate that the loss rate into undesired channels, κ_{loss} , is about $0.01 \times \kappa$, corresponding to an expected inefficiency of the conversion of $1 - \eta_{\text{conv}} \approx 0.01$.

In order to preserve the independence and coherence of our module when the conversion process is not active, we require the conversion to demonstrate a large on–off ratio. The fastest achievable damping is given by the bandwidth of the output mode, κ_{out} . The data in Figure 6.9 show that we can achieve a maximum on–off ratio of the decay that exceeds 10^3 .

6.3.2 Evacuating many photons

We also wish to verify that the cavity evacuation is independent of the input state. Using the same experimental pumping scheme as in Figure 6.9, we prepare larger Fock states and monitor the population of the input Fock state, $P(n)$. For state-independent damping of a harmonic oscillator, with only a single-

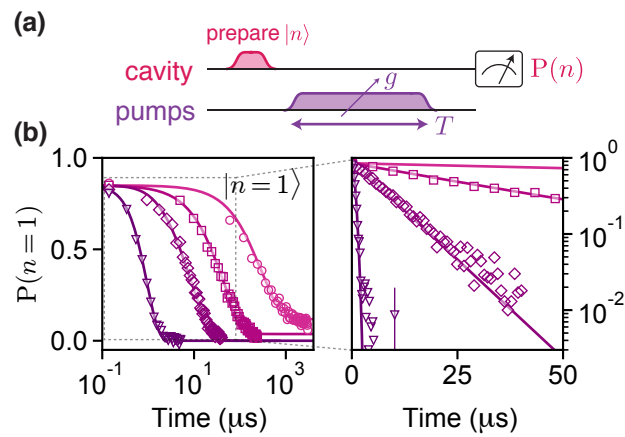


Figure 6.9 | **Cavity damping by mode-conversion.** **a** After preparing the cavity in an n -photon Fock state, we monitor the population as a function of time for different pump strengths. The pump tones have constant amplitude in time, with a smooth ring-up and ring-down. **b** Decay of the single-photon state $|1\rangle$. $g/2\pi = 0$ (circles), 25 kHz (squares), 54 kHz (diamonds), and 207 kHz (triangles). Solid lines: for $g = 0$, exponential fit, yielding the natural decay time; for $g > 0$, theoretical prediction based on independently calibrated pump parameters. For large g the decay is not simply described by a single exponential (Section 6.1.3). The last datapoint for the fastest decay shows the average and standard deviation for the residual cavity population, consistent with the vacuum state ($P(1) = 0.01 \pm 0.01$). For small or vanishing g ($\lesssim 30$ kHz), the equilibrium state of the cavity is a small thermal state. The background $n \neq 1$ population at $T = 0$ comes from sub-unity measurement contrast and initialization. (Figure adapted from [23]; see Copyright Permissions.)

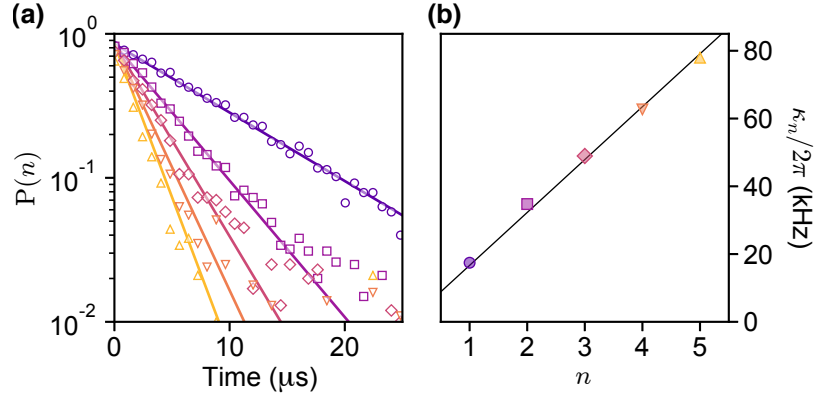


Figure 6.10 | **Cavity damping for higher Fock states.** **a** Decay of number states $|n\rangle$, with n ranging from 1 to 5; $g/2\pi = 54$ kHz. Solid lines are single-exponential fits $P(n) \propto \exp(-\kappa_n t)$. **b** Extracted decay rates κ_n . Solid line is a linear fit to $\kappa_n = \kappa_0 + n\kappa$, where κ_0 is the independently measured loss rate. (Figure adapted from [23]; see [Copyright Permissions](#).)

photon decay operator \hat{a} , we expect the state $|n\rangle$ to decay with a rate

$$\kappa_n = n\kappa. \quad (6.28)$$

This is equivalent to saying that states with more photons, which obey boson statistics, are more likely to suffer from a loss of one of them since there are more to lose. From exponential fits to the decay of $P(n)$ we find very good agreement with this linear behavior (Figure 6.10). For larger n we expect that κ_n will gradually decrease due to the Kerr effect [44]. For $n \leq 5$ we find a deviation of κ_n from $n\kappa_1$ of $\leq 6\%$, and therefore state independence is a good approximation (Section 6.1.3). As described earlier, Kerr and the bandwidth of the process begin to alter this linearity; this can be improved further by reducing the magnitude of the Kerr effect by adjustment of sample parameters.

Further, successful communication requires faithful state mapping, independent of the number of photons. Thus, each photon in the memory should be removed, described by the annihilation operator \hat{a} , while creating an outgoing photon, described by the creation operator \hat{b}^\dagger . We analyze this next.

6.3.3 Faithful stationary–propagating conversion

A crucial requirement for our interface is that non-classical states are preserved faithfully in the conversion process. To determine whether this is the case, we must characterize states before and after conversion. So far, we have only measured the remaining cavity population by performing cavity measurements. Now, we must understand how to analyze the propagating field that is output from the conversion process.

Measuring propagating photons

Reconstructing the state of a cavity via Wigner tomography is relatively straight-forward. Extracting the state of a propagating mode incident on a lossy detector, however, is more complicated. In the following, we will describe the process necessary to obtain the quantum state from this kind of measured signal. Achieving quantitative accuracy will require calibration of the detector efficiency, as in Section 6.2.1.

Because our detector performs a measurement of the form $\langle a + \hat{a} \rangle$, the averaged signal vanishes for most states of interest. To visualize this data, then, we compute a probability distribution in phase space. We can integrate the signal incident on our detector, which provides a time series of I and Q data points for each shot. This contains an “ (I, Q) -bit” of information, essentially, and can be integrated against the expected exponential envelope of cavity decay to obtain it. Repeating this for many shots accumulates an ensemble average measurement of the outgoing state. Assembling a histogram of these shots is equivalent to producing a discretized quasiprobability function, the Husimi Q-function, which is defined for the basis of coherent states as

$$P(z) = \frac{1}{\pi} |\langle \alpha = z | \psi \rangle|^2. \quad (6.29)$$

Then, as discussed in Section 3.5.4, the Q-function can be used to reconstruct the state and the density matrix ρ of the propagating signal. The “fidelity” of the resultant state is given by its overlap on the expected state, ρ_{expect} , as

$$\mathcal{F} = \left(\text{Tr} \sqrt{\sqrt{\rho} \rho_{\text{expect}} \sqrt{\rho}} \right)^2. \quad (6.30)$$

Converting quantum states

We prepare a cavity state and record the field using our heterodyne detection and quantum-limited amplifier (Figure 6.11a). The averaged in-phase signal, $\langle I(t) \rangle$, from releasing a coherent state with average photon number $\bar{n} = 1$ is shown in Figure 6.11b. Because the output mode has a finite bandwidth, we observe an exponential rise of the signal at rate κ_{out} , followed by an exponential decay with the induced decay rate κ . The emitted signal clearly retains coherence with the cavity state, made visible as an oscillation by demodulating with a small detuning from the output frequency. Importantly, the amplitude of the oscillations is consistent with a high conversion efficiency from the cavity to the output mode. By calibrating the signal amplitude in terms of the number of photons emitted by the output resonator, we estimate that the propagating field contains 1 ± 0.15 photons.

This calibration, and its associated uncertainty, are determined through the following process. First, we prepare a reference signal S_{ref} that makes no use of the conversion process. This is done by calibrating the strength of a drive required to release a one-photon coherent state, $|\alpha = 1\rangle$, from mode \hat{b} by measuring the imposed Stark shift (6.2.1). When this state is naturally emitted, the detector registers a steady-state amplitude S_{ref}^0 enclosed in ring-up and ring-down profiles with rate $\kappa_{\text{ref}} = \kappa_{\text{out}}$. A state prepared and released for the purpose of evaluating the conversion efficiency will lack this steady-state component, instead following the shape of the natural exponential emission with effective decay rate κ depending on the strength of the conversion process. The efficiency, given by the energy released compared to that prepared, is proportional to the ratio of integrated detector signals; that is,

$$\eta = \frac{\langle b_{\text{out}}^{\text{pitch}} \rangle}{\langle b_{\text{out}}^{\text{ref}} \rangle} = \frac{\int S^{\text{pitch}}(t) dt}{\int S^{\text{ref}}(t) dt} \quad (6.31)$$

since $\langle b_{\text{out}} \rangle \propto \int S(t) dt$. Specifically, the input–output relation and the quantum Langevin equation dictation that

$$b_{\text{out}} = b(0)\sqrt{\kappa}e^{-\kappa t/2} \quad (6.32)$$

for an emitted wavepacket in the rotating frame of the resonator. Thus, our total expression for efficiency

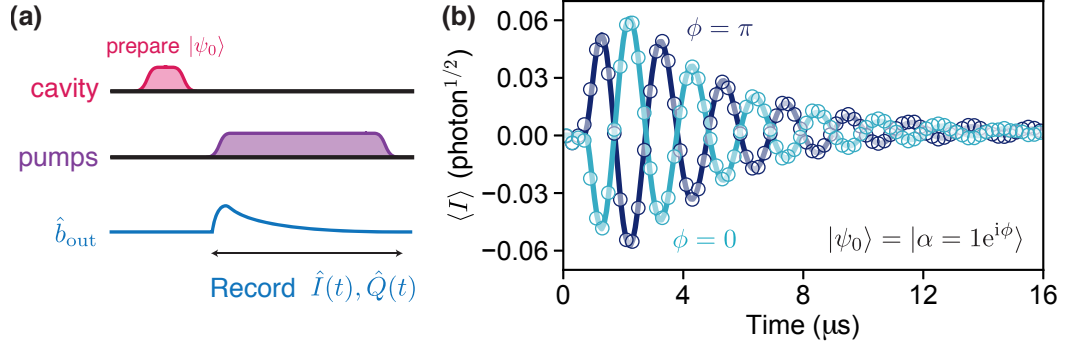


Figure 6.11 | **Measuring the field of converted, propagating states.** **a** After preparing a cavity state, we monitor the output field in heterodyne detection while applying the pump tones. **b** Averaged in-phase signal for two coherent states with $\bar{n} = 1$ and opposite phases. $g/2\pi = 125$ kHz for this data. Solid lines: fit to a sum of two exponentials with sinusoidal oscillation, $\langle I(t) \rangle \propto (\exp(-\kappa t) - \exp(-\kappa_{\text{out}} t)) \times \cos(2\pi f t + \phi_{\pm})$. f is the finite difference between the signal and demodulation frequencies, and ϕ_{\pm} are the phases of the resulting signals, corresponding to prepared states $|\pm\alpha\rangle$; here, $\phi_- = \pi - \phi_+$. (Figure adapted from [23]; see [Copyright Permissions](#).)

becomes

$$\eta = \frac{\langle b_{\text{out}}^{\text{pitch}}(0) \rangle S_0^{\text{pitch}}}{\langle b_{\text{out}}^{\text{ref}}(0) \rangle S_0^{\text{ref}}} \sqrt{\frac{\kappa_{\text{out}}}{\kappa}}, \quad (6.33)$$

where S_0 is the initial amplitude of the decay curve measured at the detector. For our reference pulse, we use a coherent state $|\alpha = 1\rangle$. For our calibration pulse (the one that is converted), we prepare a Fock state superposition, $|\psi\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$. Our ratio of expectation values is thus $\frac{\langle b_{\text{out}}^{\text{pitch}}(0) \rangle}{\langle b_{\text{out}}^{\text{ref}}(0) \rangle} = 1/2$. We measure and fit the decay of the calibration signal, $b_{\text{out}}^{\text{pitch}}$, allowing us to evaluate this expression. We do so for many conversion strengths, finding $\eta > 0.9$. Uncertainty arises from many sources, beginning with our measurement of the Hamiltonian parameter χ_{bt} , which affects the error bar on S_0^{ref} . Uncertainty associated with the time at which the ring-down begins in S_0^{pitch} affects its uncertainty, since the signal is oscillating. Additional uncertainty is associated with κ . Within these uncertainties, and from our former measurement of detector efficiency, this is in agreement with our expectation of a small inefficiency in the conversion.

By sending a classical state like $|\alpha = 1\rangle$ that we have the ability to directly measure, we can extract and retain the envelope of the decay. This envelope is then used to integrate the field from each shot and extract and histogram shot-by-shot (I, Q) values for any arbitrary state (Section 6.3.3). Although non-classical features are thus blurred by the detector, state-essential signatures are preserved in the raw data;

knowledge of the detection efficiency allows us to quantitatively confirm the faithful release of quantum states. We illustrate this using two classes of non-Gaussian oscillator states.

The first class are Fock state superpositions of the form $(|0\rangle + |n\rangle)/\sqrt{2}$, which display an n -fold symmetry in their quasiprobability distributions. For a set of such states we show the Wigner function of the cavity state, measured directly after preparation (Section 3.5), and the Q-function of the released field (Figure 6.12; from comparison it is clear that the two distributions share the same symmetry. For additional clarity, we integrate the Q-function radially to obtain a probability distribution as a function of angle, $\text{Pr}(\phi)$. In this representation it can be seen that the symmetry is fully preserved; the contrast is as expected, given our detection efficiency.

A second class of states of particular interest for CV quantum information processing are “Schrödinger cat” states of the form $|\mathcal{C}_\alpha^\pm\rangle = \mathcal{N}(|\alpha\rangle \pm |-\alpha\rangle)$, which are eigenstates of photon number parity. We create and release the even (+) and odd (−) parity coherent-state superpositions $|\mathcal{C}_{\sqrt{2}}^\pm\rangle$ with average photon number $|\alpha|^2 = \bar{n} = 2$ (Figure 6.13). Because in heterodyne detection only the Q-function is directly accessible, the characteristic coherence fringes are strongly suppressed in the traveling field data (Section 3.5.4); as a result, the distributions appear fairly similar. However, subtracting the marginals — obtained by integrating the Q-function along one axis — clearly reveals a difference, with a magnitude that is consistent with our detection efficiency and a high degree of state preservation. The heterodyne detection used in this experiment limits the number of photons accessible to about ~ 5 ; from the analysis of Section 6.3.3, we estimate that in this regime the fidelities with the ideal states exceed 90%.

The release of the cavity states shown can enable error-correctable transmission of quantum information. Because we have temporal control over the pump tones, we can shape the wave packet, which enables capture of emitted fields by a receiving module [205]. An inherent challenge for this direct quantum state transmission is that inevitable photon loss in the transmission channel will corrupt the received state. However, by choosing an appropriate encoding, the receiver will be able to detect and correct this error. For example, the states $|2\rangle$ and $(|0\rangle + |4\rangle)/\sqrt{2}$ are codewords of a binomial code [52] that can readily be sent by our system. Single photon loss in the transmission channel will result in a change of parity, which can be detected and corrected by a receiver. These concepts will be detailed further in Chapter 7.

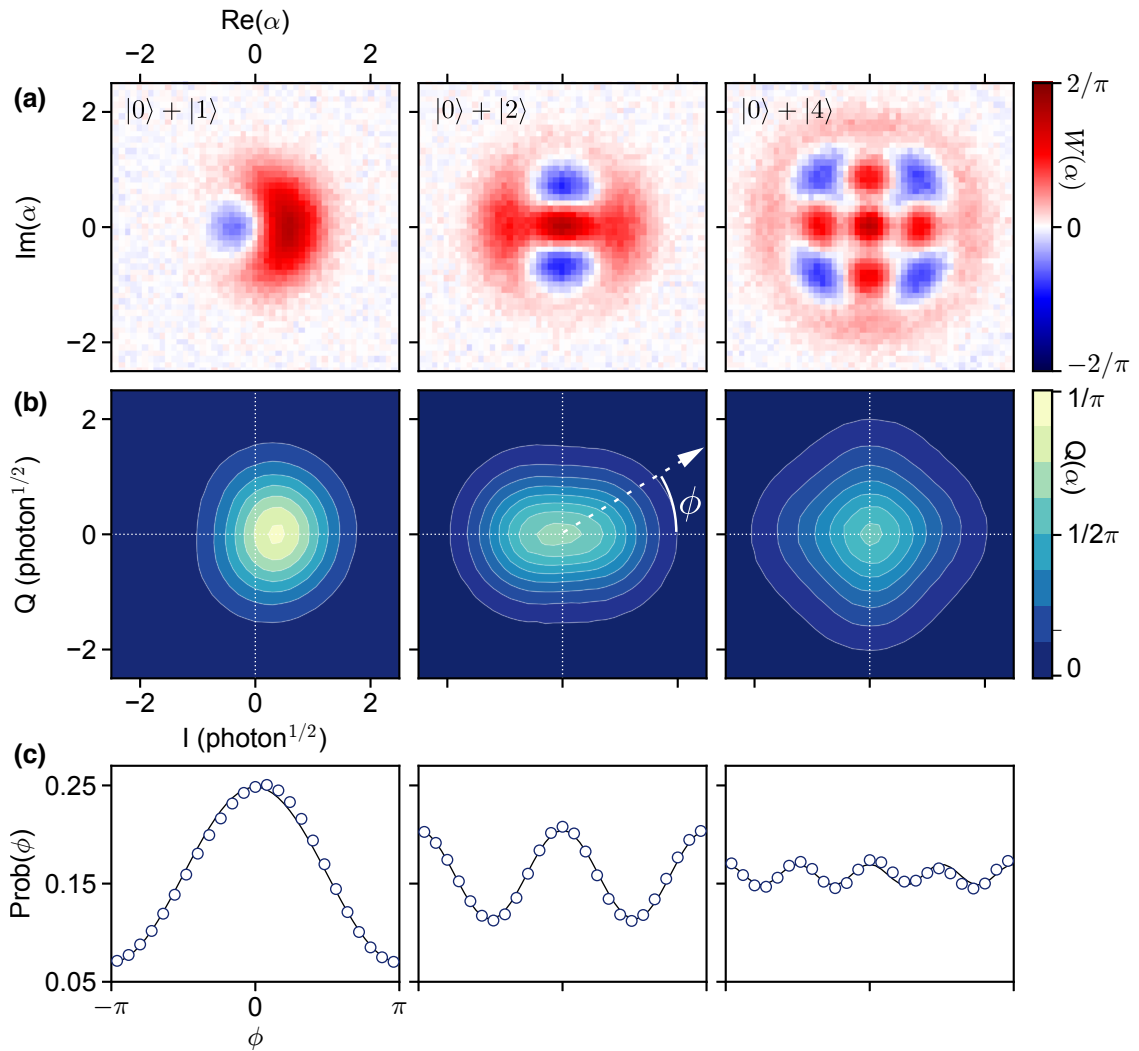


Figure 6.12 | **Traveling multi-photon quantum states.** Fock state superpositions $(|0\rangle + |n\rangle)/\sqrt{2}$, arranged in columns and labeled as insets. **a** The measured Wigner functions of the prepared cavity states (not corrected for imperfect readout). **b** The measured Q-function of the traveling signal, not corrected for detection loss. **c** The radially integrated Q-function, $\text{Pr}(\phi) = \int rQ(r, \phi)dr$, is compared with the expected contrast for the ideal state (solid), taking into account the detection efficiency. The symmetry of $\text{Pr}(\phi)$ let us intuitively compare to the expected state symmetry. All Q-function data have been taken with $g/2\pi = 164$ kHz. Fock state superposition Q-functions were taken with 10^7 samples; Cat state Q-functions with 10^6 samples. (Figure adapted from [23]; see [Copyright Permissions](#).)

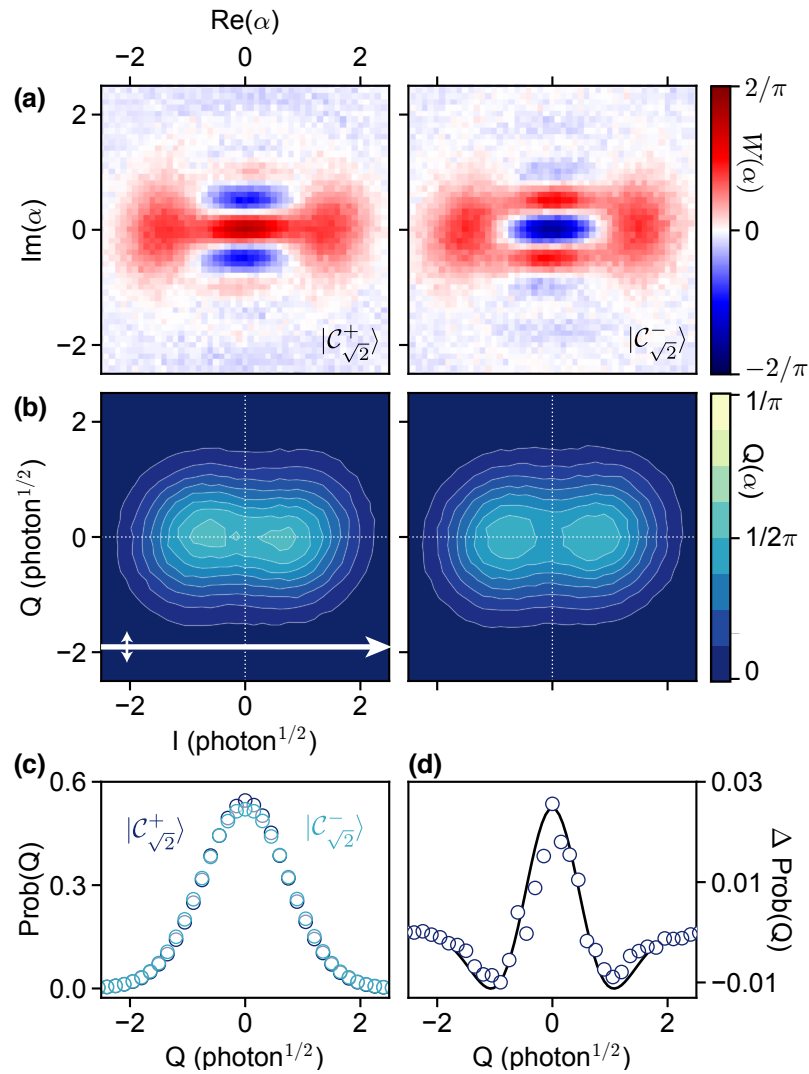


Figure 6.13 | **Traveling cat states.** **a** Wigner and **b** Q-functions of even and odd cat states, $|c_{\sqrt{2}}^{\pm}\rangle$. **c** Marginals $\text{Pr}(Q)$, obtained by integrating over I . The long arrow in **b** depicts an integrated slice; slices are taken for all values of Q . **d** Difference between the marginals (odd subtracted from even). Solid lines: expected signals for the ideal states, taking into account the detection efficiency. (Figure adapted from [23]; see [Copyright Permissions](#).)

Effect of cavity Kerr on multi-photon states

By analyzing how well the process preserves quantum information, we have established that the conversion is faithful within the bounds of our measurement error. Further, multiple photons can be transmitted in a state-independent fashion. This value will be limited by the process bandwidth, arising from variable Stark shifts, as well as the cavity Kerr. It is worth exploring what effect Kerr has, and at what point it will be limiting.

The relatively large self-Kerr of our particular cavity leads to a slight deviation from state-independent decay around $n \geq 5$, as well as a ‘smearing’ of some of the prepared states in Figures 6.12 and 6.13. Here we briefly describe the origin of the term and what parameter alterations could reduce these effects.

The self-Kerr of any mode within the system is given by the fourth-order expansion term of the transmon cosine Hamiltonian (Section 3.4.2) and is independent of the drive. The single-photon Kerr-effect in the cavity [105] is described by the Hamiltonian

$$H_{\text{Kerr}} = \frac{\chi_{aa}}{2} \hat{a}^{\dagger 2} \hat{a}^2. \quad (6.34)$$

Comparison with the cosine expansion yields a relation to the junction energy E_J as $\chi_{kk}/2 = -E_J \phi_k^4/4$.

A state prepared in a cavity with significant self-Kerr will experience dephasing over time. The characteristic state collapse time (at which point the information becomes highly “smeared”) can be expressed as [44]

$$T_{\text{collapse}} = \frac{\pi}{2\sqrt{\bar{n}}\chi_{aa}} \quad (6.35)$$

for a state in e.g. \hat{a} with average photon number \bar{n} . In our system, $\chi_{aa} = 2\pi \times 22$ kHz. With a cavity evacuation rate κ , the product $\kappa T_{\text{collapse}}$ must be much greater than 1 for this Kerr effect to be negligible. Given our experimental parameters, this product can be as low as

$$\kappa T_{\text{collapse}} = (2\pi \times 18.2 \times 10^3 \text{ Hz}) \times (11.4 \times 10^{-6} \text{ s}) = 1.3 \quad (6.36)$$

for states containing $\bar{n} \sim 1$ and for $g = 2\pi \times 54$ kHz. Also, the construction of the Kerr Hamiltonian term dictates that this effect will only be visible when particular multi-photon states are used.




Figure 6.14 | **Kerr effect during half-release of a cat-state.** **a** Cavity state after creating $|\mathcal{C}_{\sqrt{2}}^+\rangle$ followed by $3\ \mu\text{s}$ of delay, for the experiments in Section 6.4. While the state maintains its purity, the Kerr effect leads to a rotation and a ‘smearing’ of the populations. **b** Cavity state after half-release of $|\mathcal{C}_{\sqrt{2}}^+\rangle$. The state now appears fully mixed due to the entanglement with the traveling mode; the additional rotation with respect to the case of **a** comes from the Stark shifts induced by the pumps. Mapping onto the transmon is achieved by displacing one of the ‘blobs’ to the vacuum followed by mapping the vacuum onto the transmon excited state. This is indicated by the circle and arrow; the radius of the circle is $1/2$, which is the standard deviation of a coherent state in the Wigner function. It can readily be seen that the smearing leads to a small imperfection in the mapping because the blob will not be confined to the vacuum after the displacement. (Figure adapted from [23]; see [Copyright Permissions](#).)

In particular, we should see an effect when releasing, for example, $(|0\rangle + |4\rangle)/\sqrt{2}$ or $\mathcal{C}_{\sqrt{2}}^\pm$ (Figure 6.14). While this value of $\kappa T_{\text{collapse}}$ may appear to be prohibitive for the slow release of complex multi-photon cavity states, simple parameter changes (notably in the pump strengths, κ_{out} , and χ_{aa}) can greatly enhance this factor and make the Kerr collapse time irrelevant to future experiments. In fact, the experiment in Chapter 7 will show that such modifications are readily achievable.

6.4 Entanglement with flying photons

We can also show that temporal control over the pumps allows us to generate entanglement between cavity and traveling modes by partial conversion. We use the large on–off ratio over the release process to convert only a part of the energy stored in the cavity (Figure 6.15). This is the analogue of a partially reflective beam splitter, and can thus generate entanglement between the reflected (remaining in the cavity) and transmitted field (in the transmission line). We prepare an input state in the cavity and then release half of its energy while recording the output field. This ‘half-release’ corresponds to a 50:50 beam splitter, for which we expect maximally entangled states (Section 6.1.4). After switching off the conversion process, we immediately perform a single-shot, high-fidelity ($\gtrsim 0.95$) measurement on the cavity. The non-classical




Figure 6.15 | **Generating entanglement between stationary and traveling fields.** **a** We partially release the cavity field by applying the pumps for a reduced amount of time, T . We condition the recorded field on the outcome of a subsequent measurement of the cavity—as having m photons in the rotated basis—to obtain correlations. **b** Partial release is analogous to a field impinging on a partially transmitting mirror, where the transmittance t is set by the pump time T . (Figure adapted from [23]; see [Copyright Permissions](#).)

correlations between recorded field and cavity outcomes measured in different bases are indicative of the generation of entanglement. We demonstrate this using two different state encodings: single photons (Figure 6.16) and cat states (Figure 6.19).

We start in the state $|1\rangle|0\rangle$; the first ket denotes the state of the cavity, and the second that of the traveling mode. Half-conversion using the operation in Equation 6.23 maps this state onto the Bell state

$$\frac{1}{\sqrt{2}}(|1\rangle|0\rangle + |0\rangle|1\rangle) = \frac{1}{\sqrt{2}}(|+\rangle|+\rangle - |-\rangle|-\rangle), \quad (6.37)$$

where we have defined $|\pm\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$. When we measure the cavity in the number basis and find it to be in the state $|0\rangle$ ($|1\rangle$), we expect to find the traveling state in $|1\rangle$ ($|0\rangle$). To perform this measurement, we apply a transmon π -pulse that is selective on either $|0\rangle$ or $|1\rangle$ and then measure in a single shot if the transmon is in the excited state (Section 3.5); due to the asymmetry of the readout errors, we discard the outcomes in which the transmon is found in the ground state ($\sim 50\%$ of the cases). This behavior is confirmed by the near-ideal contrast in the Q-functions from the traveling field, conditioned on cavity outcomes (Figure 6.16a, left column).

To show non-classicality in the correlations, we measure the cavity also in a rotated basis: that spanning the $(|0\rangle \pm |1\rangle)/\sqrt{2}$ or $|+\rangle, |-\rangle$ basis states. We first map $|+\rangle, |-\rangle$ onto $|0\rangle, |1\rangle$ by an optimized control pulse (Section 3.5.5), followed by a measurement in the number basis. Conditioning on the these outcomes, we find that the Q-functions closely resemble those of $(|0\rangle \mp |1\rangle)/\sqrt{2}$, consistent with a high-

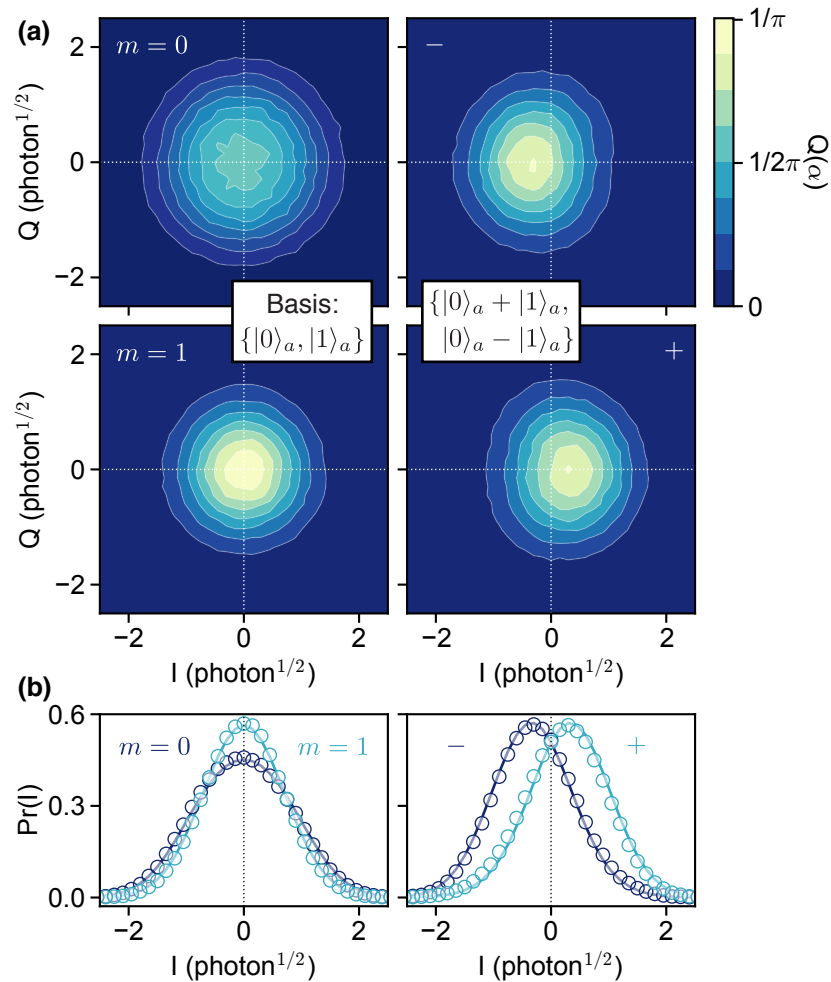


Figure 6.16 | **Entanglement between a stationary and flying photon.** Half-release of $|1\rangle$. **a** Q-functions conditioned on finding the cavity in either $|0\rangle$ or $|1\rangle$ (left), or finding the cavity in either $|0\rangle + |1\rangle$ or $|0\rangle - |1\rangle$ (right). **b** Marginals, obtained by integration along Q . Data have been taken with $g/2\pi = 164$ kHz and 10^6 samples per state and basis. (Figure adapted from [23]; see [Copyright Permissions](#).)

fidelity entangled state. This data allows us to further confirm the presence of entanglement through a simple witness.

Using a simple model of the detector we can place a lower bound on the fidelity of the entangled state by measuring an entanglement witness [207]. An entanglement witness sets a threshold that can distinguish separable states from entangled ones [208, 209]. For a two-qubit Bell state of the form (6.37) a strict lower bound on the entangled state fidelity is given by [210]

$$F \geq \frac{1}{2} (\rho_{22} + \rho_{33} + \tilde{\rho}_{11} + \tilde{\rho}_{44} - \tilde{\rho}_{22} - \tilde{\rho}_{33} - 2\sqrt{\rho_{11}\rho_{44}}), \quad (6.38)$$

where ρ is the density matrix of the joint state in the number basis, and $\tilde{\rho}$ is the density matrix in the $|+\rangle, |-\rangle$ basis. This lower bound can be evaluated with the data shown in Figure 6.16. The diagonal density matrix elements in the above expression can be found from the probabilities with which we find the cavity and propagating modes in a particular state, i.e.,

$$F \geq \frac{1}{2} (P_a(0)P_b(1|0_a) + P_a(1)P_b(0|1_a) - 2\sqrt{P_a(0)P_b(0|0_a)P_a(1)P_b(1|1_a)} \\ + P_a(+)P_b(+|+_a) + P_a(-)P_b(-|-_a) - P_a(-)P_b(+|-_a) - P_a(+)P_b(-|+_a)). \quad (6.39)$$

Here, $P_a(i)$ is the probability to find the cavity in state $|i\rangle$ in a measurement in the $|i\rangle, |\bar{i}\rangle$ basis, and $P_b(j|i_a)$ is the conditional probability to find the propagating state in $|j\rangle$ after having found the cavity in $|i\rangle$.

The probabilities $P_a(\cdot)$ follow directly from the statistics of the cavity measurements, which leaves us with the task to find the conditional probabilities for the traveling state. We compute these by fitting the measured marginals $\Pr(I)$ (Figure 6.17). With known detection efficiency, we can compute the ideal marginals, $D_{0,1}$ in the number basis, and $D_{+,-}$ in the rotated basis. Assuming our detector is linear, we can then fit the conditioned field data in the $|i\rangle, |\bar{i}\rangle$ basis to $\alpha D_i + (1 - \alpha)D_{\bar{i}}$, with $0 \leq \alpha \leq 1$. The conditional probabilities are then determined by the single fit parameter, $P_b(i|i_a) = \alpha$. Assuming again that the conversion is lossless and we can reliably determine the detection efficiency, we find a lower bound for the entangled state fidelity of $F \geq 0.91 \pm 0.02$, uncorrected for any inefficiency in the preparation or measurement of the cavity state. This clearly exceeds the classical bound of 0.5 and confirms that the

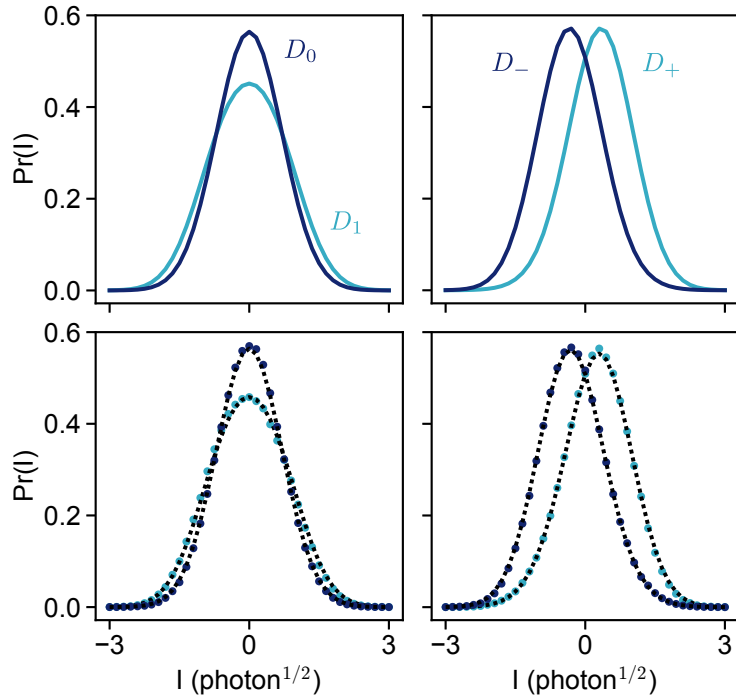


Figure 6.17 | **Entanglement bound by witness.** Estimating the fidelity of the entangled state from half-releasing a single photon. **a** Ideal marginals, assuming only the detection efficiency of 0.40 ± 0.01 we have recorded during the half-release experiment. **b** Marginal data conditioned on the cavity outcomes (same as shown in Figure 6.16b). Dashed black lines are fits to the model described in the text. Data have not been corrected for initialization or the conditioning cavity measurement fidelity. (Figure adapted from [23]; see [Copyright Permissions](#).)

half-release generates quantifiable entanglement.

6.4.1 Entanglement of two photons

While entanglement with single traveling photons can also be observed with two-level systems [196], the conversion method presented here can generate entanglement between non-classical multi-photon states. One such state is the two-photon Fock state, $|2\rangle$.

The entangled states arising from half-releasing Fock states larger than $|1\rangle$ cannot be maximally entangled states in the number basis, but still display non-classical correlations. We illustrate this by half-releasing the Fock state $|2\rangle$, which is mapped onto

$$\frac{1}{\sqrt{2}} |1\rangle |1\rangle + \frac{1}{2} (|2\rangle |0\rangle + |0\rangle |2\rangle) = \frac{1}{\sqrt{2}} |1\rangle |1\rangle + \frac{1}{2} (|+\rangle_2 |+\rangle_2 - |-\rangle_2 |-\rangle_2), \quad (6.40)$$

where now $|\pm\rangle_2 = (|0\rangle \pm |2\rangle)/\sqrt{2}$. To reveal correlations we can thus probe the states $|0\rangle, |1\rangle, |2\rangle$ (by number-selective π -pulses and transmon measurement), and $|+\rangle_2, |1\rangle, |-\rangle_2$ (mapping onto $|0\rangle, |1\rangle, |2\rangle$ by a single combined transmon/cavity pulse, followed by number-selective π -pulses) (Figure 6.18).

6.4.2 Entanglement with flying cats: Schrödinger’s cat-apult

Another interesting state to entangle (particularly for its connection to error correction schemes) is a cat state, or coherent state superposition. While such two-mode entangled cat states have been created previously in locally coupled oscillators [206], and with itinerant optical photons [211], our scheme can realize an interface between stationary and flying cats. We demonstrate this “Schrödinger catapult” by half-releasing the cat state $|\mathcal{C}_{\sqrt{2}}^+\rangle$. To show non-classical correlations we measure the cavity in the coherent state basis, finding it in either $|\pm\alpha\rangle$, or in the parity basis, thus finding it in either the even or odd cat state $|\mathcal{C}_1^\pm\rangle$. The conditioned Q-functions of the flying field are shown in Figure 6.19a. Again, the correlations are consistent with a high-fidelity entangled state. A slight reduction of contrast in the coherent state basis results from state evolution due to the Kerr effect, which reduces the fidelity of the cavity measurement (Section 6.3.3).

Because a coherent state $|\sqrt{2}\alpha\rangle|0\rangle$ is mapped onto $|\alpha\rangle|\alpha\rangle$, we can readily see that the cat state $|\mathcal{C}_{\sqrt{2}}^+\rangle|0\rangle$ results in a two-mode entangled cat state

$$\mathcal{N}_\alpha(|1\rangle_\alpha|1\rangle_\alpha + |-1\rangle_\alpha|-1\rangle_\alpha) = \mathcal{N}_c(|\mathcal{C}_1^+\rangle|\mathcal{C}_1^+\rangle + |\mathcal{C}_1^-\rangle|\mathcal{C}_1^-\rangle). \quad (6.41)$$

Here, $\mathcal{N}_{\alpha,c}$ are normalization factors, and we use the notation $|\beta\rangle_\alpha$ to denote a coherent state with complex amplitude β . We expect correlations in the coherent state phase and photon number parity bases, because the cat states $|\mathcal{C}_1^\pm\rangle$ are eigenstates of even (+) and odd (−) parity, respectively. To measure in the coherent state basis, we displace the cavity state by $\alpha = 1$, followed by a transmon π -pulse that is selective on the vacuum; $|-1\rangle_\alpha$ is thus mapped onto the the transmon excited state, while $|1\rangle_\alpha$ is mapped onto the ground state. Subsequent transmon measurement thus gives the result in the $|\alpha\rangle, |\bar{\alpha}\rangle$ basis. To measure parity we apply a series of π -pulses that are selective on the even photon numbers on the transmon; this maps even parity onto the excited state, and odd parity onto the ground state.

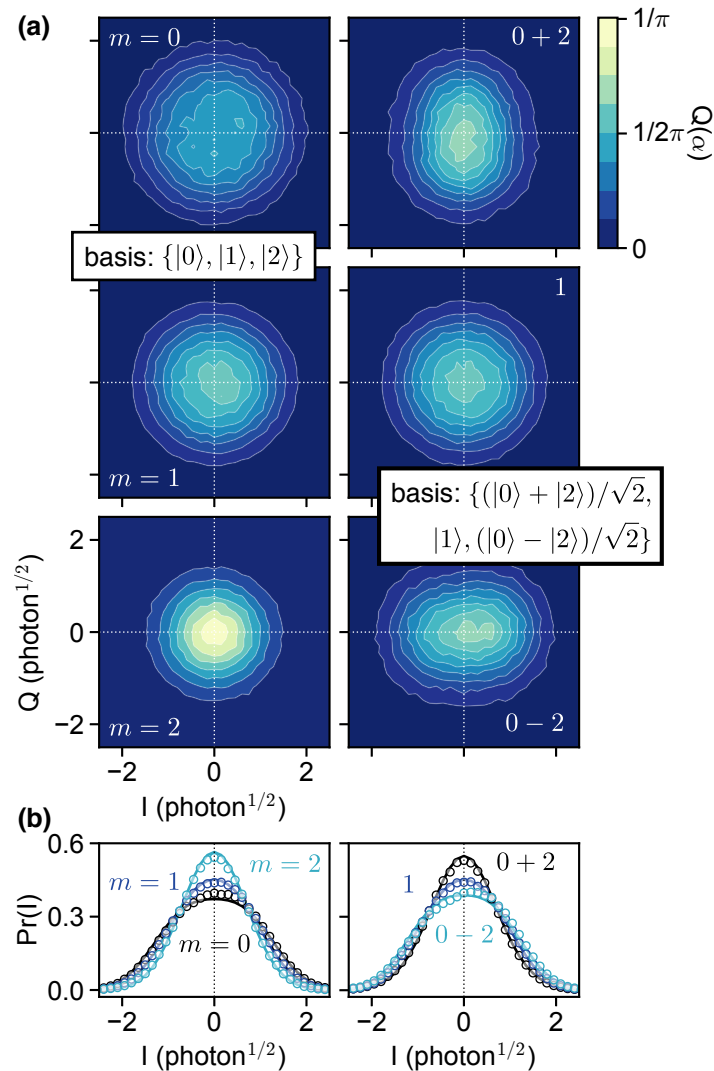


Figure 6.18 | **Conditioned Q-functions after half-release of $|2\rangle$.** **a** Left column: Q-functions conditioned on finding in the cavity in either $|0\rangle$, $|1\rangle$, or $|2\rangle$; for these case we expect the field to be in either $|2\rangle$, $|1\rangle$, or $|0\rangle$, respectively. Right column: conditioned on finding the cavity in $|+\rangle_2$, $|1\rangle$, or $|-\rangle_a$. **b** Marginal probability in each basis. Solid lines: ideal case, with only detection efficiency for the traveling field taken into account. (Figure adapted from [23]; see [Copyright Permissions](#).)

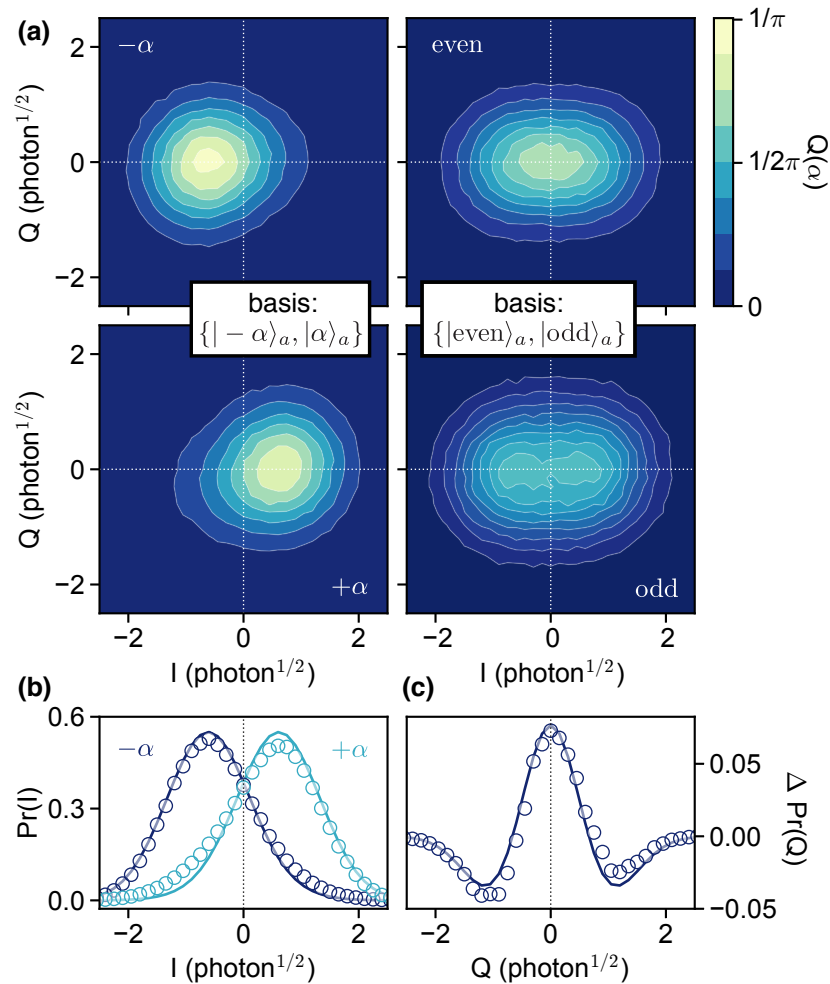


Figure 6.19 | **Entanglement between a stationary and flying cat state.** Half-release of $|\mathcal{C}_{\sqrt{2}}^+\rangle$. **a** Left column: Q-functions conditioned on finding the cavity in either $|\alpha\rangle$ or $|\alpha\rangle$. Right column: Q-functions conditioned on finding either an even or odd number of photons in the cavity. **b** Marginals, obtained by integration along I . **c** For the parity-conditioned data, we show the difference in the marginals. Solid lines: ideal case (perfect entangled state and perfect cavity measurement), taking into account only the detection efficiency in the Q-functions. Data have been taken with $g/2\pi = 164$ kHz and 10^6 samples per state and basis. (Figure adapted from [23]; see [Copyright Permissions](#).)

We note that the states $|\alpha = \pm 1\rangle$ have a non-zero overlap ($\sim 4\%$). This means that the coherent state basis is not perfectly orthogonal for cat states like $|\mathcal{C}_1^\pm\rangle$ with small photon number. This overlap, along with the Kerr effect acting during the $3\ \mu\text{s}$ it takes to create the cat state and map it onto the transmon for measurement (Section 6.3.3), explains why the conditioned Q-functions are slightly asymmetric after measuring the cavity in the $|\alpha\rangle, |\bar{\alpha}\rangle$ basis. It also prevents us from using the simple method we have employed to bound the entangled state fidelity in the single-Fock-state case.

This entanglement between stationary and traveling cats can enable error-correctable distribution of entanglement. Capture of the wave packet emitted by half-release enables the creation of remote entanglement between stationary parties. For the cat states used in this work, any photon loss in the transmission channel will corrupt the state because it results in change of parity. However, photon loss becomes detectable when we half-release a cat state of the form $|\alpha\rangle + |i\alpha\rangle + |-\alpha\rangle + |-i\alpha\rangle$ [212]. Such states are eigenstates of ‘superparity’ with modulo 4 photons, and even/odd (modulo 2) parity measurements can be used to detect and correct single-photon loss [47]. Thus, measuring and comparing the parity between the remote parties will allow for detection and correction of single-photon loss in the transmission line during remote entanglement generation. This proposed scheme will be analyzed further in the next chapter.

6.5 Conclusion

We have shown the deterministic, coherent release of quantum states from a microwave cavity memory. This release is enabled by parametric up-conversion utilizing the non-linearity of our single Josephson junction transmon. This conversion scheme fulfills the requirements we set out for an interface between stationary and traveling oscillator states in a microwave quantum network—we can dynamically control the conversion rate, releasing cavity states almost 1,000 times faster than the intrinsic lifetime; the interface maintains a large on–off ratio; it is parametrically and temporally controllable. Further, and of great importance to error correction, this conversion rate is state-independent for states containing up to a few photons (extendable to up to tens of photons by simple hardware adjustments). The release process is equivalent to a beam splitter interaction, and cavity states are mapped faithfully onto traveling states. This interaction can be controlled precisely and rapidly, enabling the generation of entanglement between cav-

ity and traveling modes.

Our interface can serve as the backbone in a microwave quantum network in which quantum information is stored in cavities. Since the conversion process is controllable in amplitude and phase, we will next show that it allows quantum state transfer and entanglement between remote cavities.

7

Putting Schrödinger’s Cat Back in the Bag: Direct Quantum State Transfer and Entanglement Distribution

In 2004, a panel of quantum computing experts met to discuss progress in the field, comparing milestones reached using various platforms. On considering the two DiVincenzo criteria for off-chip quantum communication, they rated superconducting quantum computation as showing “no viable approach” [213].

Much has changed and advanced in just over a decade. Though it remains true that optical-frequency light suffers lower attenuation than that at microwave frequencies, recent experiments [86, 87, 214] definitively show that superconducting circuits can satisfy the DiVincenzo communication criteria—and well enough so as to be of practical use, as we will show in this chapter.

This process has been given many colloquial names, among them being ‘send and receive’, ‘pitch and catch’, or ‘release and capture’. They all share the same goal: to directly and deterministically transfer arbitrary states sent between distant modules. This “distance” will be short by standards of light. The modules will be connected with less than a meter of cable, and sit even closer to one another in physical space. We will show, however, that it is highly unlikely that direct interactions produce the results we observe, and that we do realize the protocol we set out to.




Figure 7.1 | **Connecting two modules for on-demand state transfer.** Identically-constructed circuit quantum electrodynamics (QED) *sender* and *receiver* modules are connected with transmission line and a circulator, allowing straight-forward measurement of both systems. Each module contains a high- Q memory cavity (\hat{a}), transmon qubit (\hat{t}), and communication mode (\hat{b}) with strong output coupling κ_{out} . Parametric drives (purple waveforms, ξ) enable conversion between memory and communication modes. Below each module is a schematic showing the parametric conversion process, enabled by the transmon nonlinearity, which runs left-to-right in the sender and right-to-left in the receiver. (Figure adapted from [86]; see [Copyright Permissions](#).)

7.1 Extending conversion to two nodes of a network

In this chapter, we set out to fully implement the deterministic state transfer protocol by Cirac *et al.* [205], employing the same hardware in Chapter 6 to serve as remote quantum memory endpoints in a simple network (Figure 7.1). The work in ref. [205] is a prescription for what is possible: under the right circumstances, complete and deterministic state transfer can be achieved. The “circumstances” include the conversion process detailed in Chapter 6. They also include the right system, or hardware platform, that must be highly coherent. Luckily, we have demonstrated this as well.

Each module can be understood to contain two orthogonal cavity modes (memory and communication) that are coupled by an artificial atom (Figure 6.1). The communication modes—implemented as on-chip stripline resonators, as before—are strongly coupled to either end of a transmission line. Realizing on-demand state transfer requires tunable conversion between memory and communication modes within each module, such that (i) the sender emits the state contained in the memory into the transmission




Figure 7.2 | **Mode frequencies and alignment.** Frequency ordering of the modes ($\omega_{a,b,t}$) and pumps ($\omega_{1,2}$) in each module. Pumps are detuned Δ from each mode and an additional relative detuning δ between them to account for Stark shifts. (Figure adapted from [86]; see [Copyright Permissions](#).)

line as a wavepacket with a specified temporal profile, and (ii) the receiver absorbs this wavepacket (Section 7.2). The layout of mode frequencies here (Figure 7.2) is similar to that in Figure 6.2. In this case, notably, the communication frequencies must be matched. The details of this additional layer of complexity are described in Section 7.1.2.

7.1.1 Experimental details

The experimental setup (Figure 7.3) is very similar to that in Chapter 6, with a few notable exceptions. The first is that instead of a direct connection to the amplification chain and measurement/digitization hardware, the output signal from each node must pass through a circulator at least once. This circulator connects the two nodes, ensuring directionality, but also permitting a hardware-efficient means to read out each module without adding further modes.

Parameters of each mode within the two devices, including frequencies and coherence times, are included in Table 7.1. Both samples are within a magnetic shield, while all circulators are outside of the shield. Images of the parts of the device are shown in Figure 7.4.

7.1.2 Mechanical tuning

One stark difference in the composition of measurement hardware between this experiment and that of the last chapter is determined by the requirement that the output signal (\hat{b}_{out}) is frequency-matched to the

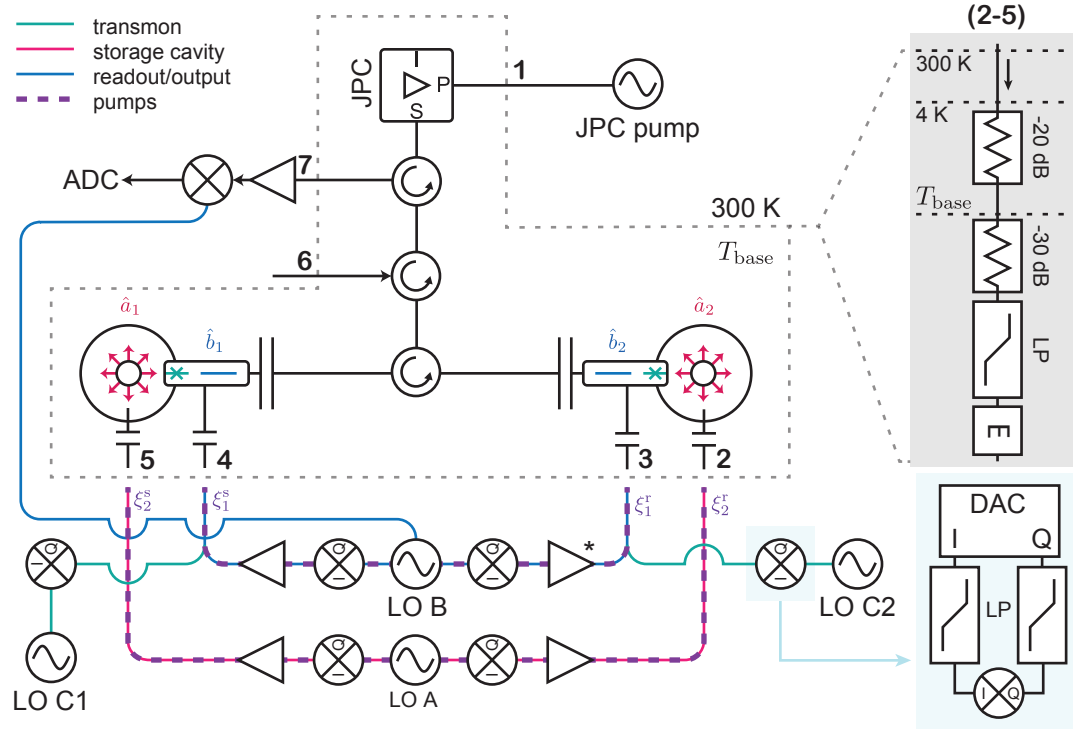


Figure 7.3 | **Schematic of the experiment wiring for a full pitch-and-catch experiment.** Microwave signals are IQ-modulated, amplified, and transmitted to the cold stage of the dilution refrigerator (*Oxford Instruments Triton 200*) where the sample is cooled to $T_{\text{base}} \approx 20$ mK. IQ modulation is performed with an integrated FPGA system with digital-to-analog converter (DAC) outputs (*Innovative Integration VPXI-ePC*), and mixed with the LO (*Marki Microwave IQ0307LXP and IQ0618LXP* IQ-mixers). Inset shows filtering on low-frequency IQ channels. The generators acting as local oscillators (LOs) are *Agilent N5183A* (LOs A and B) and *Vaunix LabBrick LMS-103-13* (LOs C1 and C2). Tones with similar frequencies share IQ channels. Cavity and readout input signals are amplified at room temperature (*MiniCircuits ZVA-183-S+*). Input amplifiers are followed by switches (*Hittite HMC-C019*) except for (*) in the diagram. Input signals are introduced via weakly, capacitively-coupled pins. Signals leave the output resonators through strongly coupled pins (large capacitors) and are amplified by a Josephson parametric converter (JPC) that is pumped by a microwave generator (“JPC pump”, *Agilent N5183A*). The signal is further amplified at 4 K (*Low Noise Factory LNF-LNC7_10A*) and at room temperature (*Miteq AMF-5F-04001200-15-10P*), mixed down with the output LO (*Marki Microwave IQ0618LXP*), and recorded and demodulated by the FPGA system with an analog-to-digital converter (ADC). Lines entering or leaving the refrigerator carry the following signals: (1) the JPC pump tone, (2–5) sample input tones, (6) a diagnostic tone with which to probe samples in reflection and tune the JPC, and (7) the measurement output line. Inset shows attenuation and filtering on lines 2–5, including a low-pass filter below 12 GHz (LP) and an Eccosorb filter (E). Line 6 is similar but with an additional -20 dB of attenuation at the 20 mK stage, while line 1 omits the low-pass filter. (Figure adapted from [86]; see [Copyright Permissions](#).)

Hamiltonian parameter		Sender value (MHz)	Receiver value (MHz)
Frequency	$\omega_a/2\pi$	4219.3	4269.6
	$\omega_b/2\pi$	10031.5	10031.5*
	$\omega_t/2\pi$	6156.1	6417.6
Cross-Kerr	$\chi_{ab}/2\pi$	-16×10^{-3}	-12×10^{-3}
	$\chi_{at}/2\pi$	-2.86	-2.29
	$\chi_{bt}/2\pi$	-2.4	-2.18
Self-Kerr	$\chi_{aa}/2\pi$	-8×10^{-3}	-5×10^{-3}
	$\chi_{bb}/2\pi$	-8×10^{-3}	-6×10^{-3}
	$\chi_{tt}/2\pi$	-183.43	-196.17
Damping parameter		Sender value (μs)	Receiver value (μs)
Cavity single-photon energy decay time	T_1^a	460 ± 10	770 ± 10
Cavity Ramsey decay time	T_{2R}^a	102 ± 3	130 ± 4
Output mode energy decay time	T_1^b	0.14 ± 0.01	0.11 ± 0.01
Transmon relaxation time	T_1^t	26 ± 3	27 ± 3
Transmon Ramsey decay time	T_{2R}^t	12 ± 2	12 ± 2
Transmon Hahn echo decay time	T_{2E}^t	15 ± 2	15 ± 2
Steady-state excitation		Sender value	Receiver value
Transmon	$1 - P(g)$	0.195	0.209
Cavity	\bar{n}	0.166	0.172

Table 7.1 | **Measured pitch-and-catch system parameters.** Uncertainties of measured Hamiltonian parameters are $< 0.1\%$ except when indicated by fewer significant digits. For the cavity and transmon decay times, the uncertainties given are the typical fluctuations observed over the course of one day. *This is the frequency set by tuning and is used during all phases of the experiment.




Figure 7.4 | **Images of device used in pitch-and-catch experiment.** **Left:** the two systems are packaged and assembled in close proximity. The circulator and lines connecting the two systems are outside of the field of the image. **Upper right:** inside the package, dashed lines depict the cavity and tunnel housing the communication mode and transmon. **Lower right:** Close-up view of the chip, which is held by a clamp that is affixed to the package. (Figure adapted from [86]; see [Copyright Permissions](#).)

receiving cavity. To facilitate this, a tuning mechanism is implemented to match the output (or ‘communication’) modes. A low-temperature mechanical micropositioning stage (*Attocube ANPz101-A4*) is used to introduce a small superconducting pin through a hole in the side of the device. The pin perturbs the field of the output mode, lowering that mode’s frequency at greater pin insertion. The pin can be inserted over a range of 5 mm and the mode frequency can be tuned by ~ 300 MHz.

We measure the devices individually and outfit the tuning mechanism on the system with the higher (unperturbed) output mode frequency (in order to guarantee that the output modes can be tuned into resonance). Once the resonance condition is met, we disconnect the positioning controller to maintain frequency stability. We expect that in future versions of modules, fabrication and mounting errors can be minimized such that *in situ* tuning will not be necessary.

Meanwhile, the cavity and transmon modes of the sender system are far-detuned (relative to their linewidths) from the respective modes in the receiver system, since there is no special requirement that relates them. The combined, measured frequency spectra are shown in Figure 7.5.

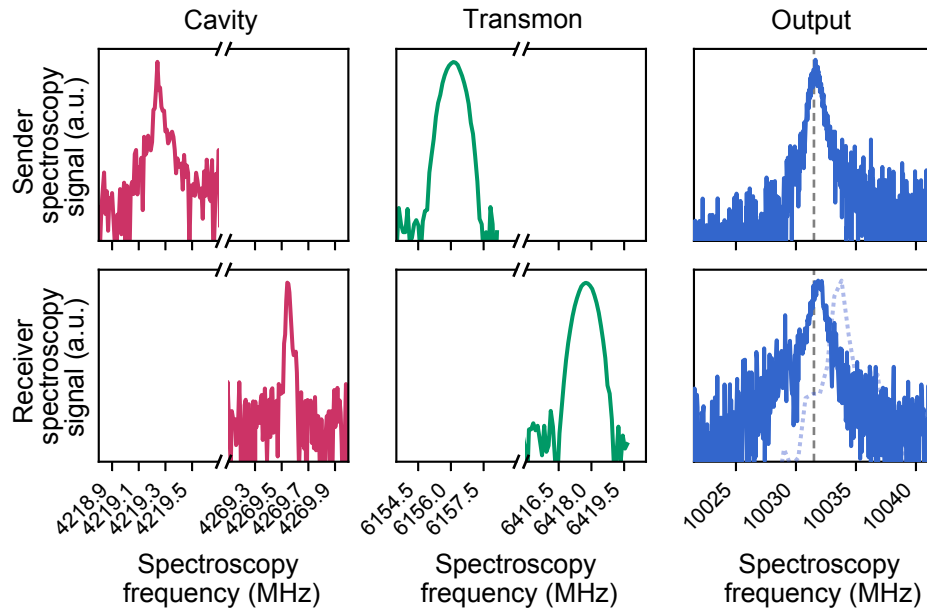


Figure 7.5 | **Mode spectroscopy.** Cavity spectroscopy is performed by applying a displacement at variable modulation frequency and measuring whether the transmon is in the cavity-number-state-dependent state corresponding with photon number peak $n = 0$. Qubit spectroscopy is performed by applying a selective (spectrally narrow) π -pulse at variable frequency near the transmon mode, and reading out. Output mode spectroscopy for System A is performed using a vector network analyzer (VNA) to measure the transmission parameter (S_{21}) between the readout/transmon input port and the measurement line. The dotted lines indicate data taken while the output frequency has been purposefully detuned. Output mode spectroscopy for System B is done by stepping the spectroscopy frequency and reading out normally. The vertical dashed line indicates the on-resonance value. Secondary peaks in transmon and readout spectroscopy arise from finite thermal population in the cavity and transmon, respectively. All results are expressed in normalized logarithmic units.

7.1.3 Cavity state tomography

Wigner tomography on a single cavity will be done through a series of displaced parity measurements [105, 133], and follows Section 3.5.4. As a reminder, in this particular system, the parity is mapped onto the transmon by applying a set of selective π -pulses simultaneously, on either the even or odd number-peaks, up to $\bar{n} = 25$. Taking the difference of the even and odd mapping sequences, we obtain a value proportional to parity (and therefore to the Wigner function), independent of the initial transmon state. To compensate for finite π -pulse and measurement contrast, we then normalize the resulting data so that it integrates to unity over all phase space, leading to a physical Wigner function. Displacements are typically performed up to $\beta = 2.5$, resulting in Wigner functions that should capture $> 99\%$ of the energy contained in the mode for our states with $\bar{n} \leq 2$. We have determined a maximum error associated with this normalization and reconstruction process to be $\sim 6\%$.

7.1.4 Two-system calibrations

Before proceeding with the experiment, we must again calibrate and characterize our conversion process — this time, at both ends.

We verify the independence of the transmon state assignment by performing a simultaneous Rabi experiment: we apply a variable rotation on both transmons, followed by joint measurement of the whole system. The results (Figure 7.6) show that the assigned state of one transmon depends only very weakly on the rotation applied to the other, and that there is minimal crosstalk between transmon control pulses.

To verify that the pump tones applied to one system do not produce unwanted effects on the opposite system, we measure leakage by applying one pump at a constant amplitude and measuring each transmon’s Stark shift (Figure 7.7). We find that measurable leakage occurs only in the direction permitted by the circulator (from sender to receiver), and only between communication modes (since their frequencies are very nearly matched). Based on these measurements, we estimate that this leakage contributes to “miscalibration” error on each system at the level of $\lesssim 1\%$.

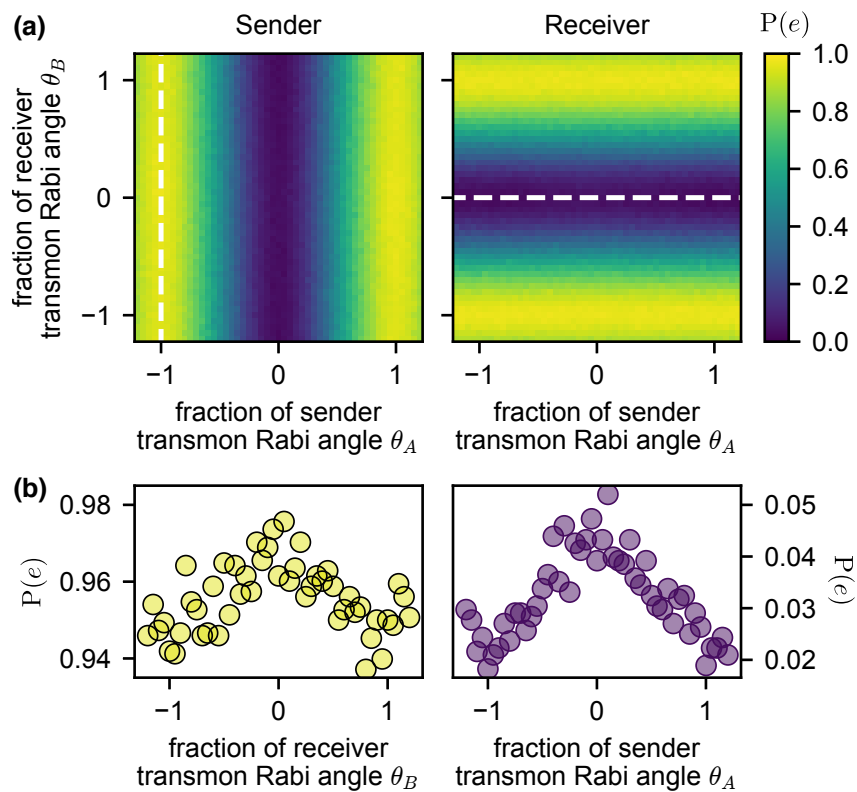


Figure 7.6 | **Transmon measurement independence.** Rabi pulses are applied with varying angles (relative to a unit π -pulse) on the sender and the receiver transmons simultaneously. **a** Both transmon states are measured as a function of both angles. **b** Line cuts (corresponding to dashed lines). Unwanted correlations are $\leq 2.5\%$. (Figure adapted from [86]; see [Copyright Permissions](#).)




Figure 7.7 | **Pump leakage between systems.** We measure the Stark shift (expressed in number of circulating pump photons) on both transmons while applying a single pump to the sender mode (upper row) or the receiver mode (lower row) near either the output/communication mode (left column) or the cavity mode (right column). Solid lines are linear fits to the data, weighted by the statistical error of each point (bars not visible are smaller than marker size). Error bars represent the standard deviation of the fit to the spectroscopy peak at each point. We calculate error from experimental sources to exceed 0.01 photons, dominating the statistical fit error in most cases. (Figure adapted from [86]; see [Copyright Permissions](#).)

7.2 Wavepacket shaping

Using an unshaped pulse as we did in Chapter 6 is, unfortunately, not an option by which to capture all of the released energy. The natural cavity decay profile, a falling exponential in time, is not “temporally impedance-matched” to an unshaped conversion process at the receiver. One can make sense of this by “recording”, as a thought experiment, the outgoing field as a time series $E(-t) \propto e^{-t}$ for $0 < t < t_0$. If allowed to run to infinite time, this process will release all of the cavity energy perfectly. By symmetry, running this process in reverse—directing a field with a rising exponential profile $E(t) \propto e^{t-t_0}$ into an empty cavity—will cause this state to be perfectly absorbed. Unfortunately, we have no way to easily implement this time-reversal on the propagating, falling-exponential wavepacket. With constant coupling at the receiver, attempting to capture a falling exponential will succeed in capturing only a fraction of the energy.

One way to solve this using our control abilities, particularly if identical systems are present at both ends, is to implement a shaped conversion process $g(t) = g(-t)$ at either end. This will mandate a time-symmetric propagating wavepacket form. However, as the following sections will show, the time-symmetry of this wavepacket shape is not strictly necessary; optimizations can be performed to transfer a fraction of the energy that quickly approaches unity.

7.2.1 Calculating waveforms

In the rotating frame of the pumps, the Hamiltonian enabling conversion between memory and output modes a and b is

$$H_{\text{conv}}(t) = i(g(t)\hat{a}\hat{b}^\dagger - g^*(t)\hat{a}^\dagger\hat{b}) \quad (7.1)$$

with

$$g(t) = g(\xi_1(t), \xi_2). \quad (7.2)$$

The correspondence producing the conversion rate $g(t)$ as a function of the two applied pump amplitudes ξ_1 and ξ_2 is calibrated experimentally (Section 6.2.1). The phase convention of the conversion Hamiltonian is chosen to make later computation easier. In the lowest-order approximation, $g_j \propto \xi_1^j \xi_2^j$ for System

j.

Because the conversion always depends on both pumps, only one of the two pumps needs to vary in time to produce any particular $g(t)$. For experimental convenience, we designate ξ_1 to vary in time. The pump ξ_2 is held constant in time, with a smooth (~ 200 ns) ring-up and ring-down profile.

One significant effect of the application of pumps are ac-Stark shifts that shift the frequency of the modes. The Stark shifts during the transfer process are a function of both pump amplitudes:

$$H_{\text{Stark}}(t) = \delta_a(t)\hat{a}^\dagger\hat{a} + \delta_b(t)\hat{b}^\dagger\hat{b} \quad (7.3)$$

Each Stark shift depends on each pump amplitude:

$$\delta_a(t) = \delta_a(\xi_1(t), \xi_2) \quad \delta_b(t) = \delta_b(\xi_1(t), \xi_2) \quad (7.4)$$

This shift is calibrated independently. We find very good agreement with the expected dependence, which is linear in the sum of the pump powers:

$$\begin{aligned} \delta_a(t) &= 2\chi_{aa}|\xi_1(t)|^2 + \chi_{ab}|\xi_2|^2 \\ \delta_b(t) &= 2\chi_{bb}|\xi_2(t)|^2 + \chi_{ab}|\xi_1|^2 \end{aligned} \quad (7.5)$$

Having established the dependence of both the conversion rate $g(t)$ and the Stark shifts δ_a, δ_b , we can write down the equations of motion for the modes \hat{a} and \hat{b} of the sender:

$$\dot{a}(t) = -g(\xi_1(t), \xi_2)b(t) - i\delta_a(\xi_1(t), \xi_2)a(t) \quad (7.6a)$$

$$\dot{b}(t) = g^*(\xi_1(t), \xi_2)a(t) - i\delta_b(\xi_1(t), \xi_2)b(t) - \frac{\kappa_{\text{out}}}{2}b(t) \quad (7.6b)$$

$$b(t) = b_{\text{out}}(t)/\kappa_{\text{out}} \quad (7.6c)$$

Because these equations are linear, we can consider the evolution of the expectation values of the field operators \hat{a} and \hat{b} , so we have dropped the operator notation. This will also allow us to solve the problem classically, which is computationally simpler compared to a full quantum simulation.

Equation 7.6c is the input–output relation, taking as an assumption that there is no incoming field. In this case, for a chosen $b_{\text{out}}(t)$ and ξ_2 , the only undetermined quantities are $a(t)$ and $\xi_1(t)$. In what follows we suppress the time dependence and the explicit dependence on the static ξ_2 for simplicity.

The goal is to eliminate a , leaving an equation for g in terms of b that can be solved numerically. We first note that Equation 7.6b can be written

$$g^*(\xi_1)a = \dot{b} + i\delta_b(\xi_1)b + \frac{\kappa_{\text{out}}}{2}b \quad (7.7)$$

with derivative

$$\dot{g}^*(\xi_1)a + g^*(\xi_1)\dot{a} = \ddot{b} + i\dot{\delta}_b(\xi_1)b + i\delta_b(\xi_1)\dot{b} + \frac{\kappa_{\text{out}}}{2}\dot{b} \quad (7.8)$$

We can multiply Equation 7.6a by $g^*(\xi_1)$ and substitute it into Equation 7.8 to write

$$g^*(\xi_1)a - |g(\xi_1)|^2b - i\delta_a(\xi_1)g^*(\xi_1)a = \ddot{b} + i\dot{\delta}_b(\xi_1)b + i\delta_b(\xi_1)\dot{b} + \frac{\kappa_{\text{out}}}{2}\dot{b} \quad (7.9)$$

Finally, multiplying Equation 7.9 by $g^*(\xi_1)$ and substituting in Equation 7.7 yields

$$\begin{aligned} [g^*(\xi_1) - i\delta_a(\xi_1)g^*(\xi_1)] \left[\dot{b} + i\delta_b(\xi_1)b + \frac{\kappa_{\text{out}}}{2}b \right] - g^*(\xi_1)|g(\xi_1)|^2b \\ = g^*(\xi_1)\ddot{b} + i\dot{\delta}_b(\xi_1)g^*(\xi_1)b + i\delta_b(\xi_1)g^*(\xi_1)\dot{b} + \frac{\kappa_{\text{out}}}{2}g^*(\xi_1)\dot{b} \end{aligned} \quad (7.10)$$

Equation 7.10 is solved numerically to give the correct $\xi_1(t)$ for a given $b_{\text{out}}(t)$. The initial condition $\xi_1(t=0)$ comes from Equation 7.7. This approach has a few important features.

The first is that it inherently accounts for the Stark shifts, in two ways. First, $\xi_1(t)$ will have a phase that varies in time. This dynamic frequency control ensures that $b_{\text{out}}(t)$ can have a fixed frequency, even when the mode a does not. Secondly, the amplitude of $\xi_1(t)$ will change in time in a way that accounts for the frequency shift of the output mode b : the amplitude will increase to compensate for the fact that the conversion process is effectively off-resonant.

By scaling the output field $b_{\text{out}}(t)$ to specify the amount of energy contained therein, we calculate different pump waveforms for full and partial release via the same procedure. While the equations of motion

are linear in a and b , Equation 7.10 is clearly nonlinear in ξ_1 . This is why the pulses for full and partial releases are *not* simply scaled versions of one another, even though the released wavepackets $b_{\text{out}}(t)$ are. For this reason, the capture pulse will remain unchanged and independent of the release pulse.

Wavepacket capture

The calculation of the pump waveform required to capture the propagating wavepacket is very similar to the above. The equations of motion for the receiver are

$$\dot{a}(t) = -g(\xi_1(t), \xi_2)b(t) - i\delta_a(\xi_1(t), \xi_2)a(t) \quad (7.11a)$$

$$\dot{b}(t) = g^*(\xi_1(t), \xi_2)a(t) - i\delta_b(\xi_1(t), \xi_2)b(t) - \frac{\kappa_{\text{out}}^r}{2}b(t) + \sqrt{\kappa_{\text{out}}^r}b_{\text{in}}^r(t) \quad (7.11b)$$

$$b(t) = b_{\text{out}}^r(t)/\kappa_{\text{out}}^r + b_{\text{in}}^r(t)/\kappa_{\text{out}}^r \quad (7.11c)$$

which is identical to Equation 7.6, with the difference that there now exists an input field term $b_{\text{in}}^r(t)$. For clarity we now restore the superscripts s and r for sender and receiver, respectively. To calculate the capture waveform, we specify that this input field has the shape of the released wavepacket: $b_{\text{in}}^r(t) = b_{\text{out}}^s(t)$, and that the field reflected off the receiver is zero: $b_{\text{out}}^r(t) = 0$, which corresponds to perfect absorption. Taken together, these constraints imply

$$\dot{a}(t) = -g(\xi_1(t), \xi_2)b(t) - i\delta_a(\xi_1(t), \xi_2)a(t) \quad (7.12a)$$

$$\dot{b}(t) = g^*(\xi_1(t), \xi_2)a(t) - i\delta_b(\xi_1(t), \xi_2)b(t) + \frac{\kappa_{\text{out}}^r}{2}b(t) \quad (7.12b)$$

$$b(t) = b_{\text{out}}^s(t)/\kappa_{\text{out}}^r \quad (7.12c)$$

which looks just like Equation 7.6, but with the sign of κ_{out} changed. The procedure for obtaining $\xi_1(t)$ is the same as for the sender. The only major difference is that this equation is solved in reverse, with the final condition specifying the occupation of a at the end of the protocol. This corresponds to the fraction of incoming energy that is absorbed, $\eta_{\text{trunc}}^{(r)}$ (to be further detailed in Section 7.5.2). Increasing this fraction corresponds to increasing the pump strength beyond what is achievable in our system.

Importantly, the capture waveform is the same for both full and partial release; this is due to the linear-

ity of the equations of motion (Equation 7.11) on $b_{\text{in}}^r(t)$; in other words, the capture is state-independent. Therefore, the capture waveform depends only on the shape of the incoming wavepacket, not its amplitude.

Choice of wavepacket shape

The envelope of the wavepacket $|b_{\text{out}}(t)|$ is an arbitrary choice, up to constraints on the bandwidth of the conversion process. For experimental convenience, we choose $|b_{\text{out}}(t)| \propto 1 - \cos^2\left(\frac{\pi t}{T}\right)$, where $T = 6 \mu\text{s}$ is the total transfer time. We find empirically that this smooth shape reduces the maximum pump amplitudes required for a given transfer time as compared to other shapes tested. The frequency of the wavepacket is also free to be varied. We choose a frequency $\sim 1 \text{ MHz}$ below the static frequencies of the communication modes, to account for Stark shifts (which are always negative) while the pumps are applied.

7.2.2 Parameter measurement and refinement

For the computation of the correct waveforms we need to supply empirical values of $g(t)$ and the Stark shifts, both as a function of the pump strengths and frequencies that are applied. The value of g is estimated from the rate with which photons leave the storage cavity when pumps are applied. The Stark shifts can be measured directly from spectroscopy while applying the pumps, as in Figure 7.7. We generally follow the protocols introduced in Chapter 6 to perform these calibrations.

7.3 State transfer

Following the derivation of Section 7.2, we compute the shape of the pumps used in this process so as to best match the temporal profile of the traveling wavepacket. System parameters enable the effective coupling strengths between memories and the transmission line, $\kappa^{s,r}(t)/2\pi$, to be tuned dynamically up to 400 kHz—much larger than the intrinsic single photon decay rates of the memories, $\kappa_0^{s,r}/2\pi < 0.4 \text{ kHz}$.

Following the original proposal [205], we insert a circulator into the transmission channel, which enforces the directionality of emission from the sender. The circulator also directs signals reflected off the receiver into an output port, which allows readout of both systems using a single parametric amplifier and

heterodyne detection chain. Omission of the circulator would present different physics; depending on the length of line and resultant mode spacing, the propagating wavepacket could be envisioned as being swapped through an intermediate real or virtual resonant mode at either end [215]. While the memory resonance frequencies need not match, efficient transfer requires that the communication modes be close to resonant compared to their bandwidths ($\kappa_{\text{out}}^{\text{s,f}}/2\pi \sim 1$ MHz). The tuning mechanism from Section 7.1.2 compensates for a small offset in resonance frequency between the communication modes.

7.3.1 Quantifying energy transfer

We begin by characterizing the process by which photons in the sender are emitted, transferred, and absorbed into the receiver memory. First, we quantify the efficiency of absorption alone by preparing a small coherent state in the sender memory, and then executing the protocol under one of two conditions (Figure 7.8a). In one case, we omit the capture pulses and monitor reflection from the receiver. Here, the emitted wavepacket is fully reflected and recorded by our heterodyne detector (Figure 7.8b,c). In contrast, if we apply the complete set of pulses, this reflection is strongly suppressed. By measuring the relative photon flux at the detector, we determine that the receiver absorbs $(93 \pm 1)\%$ of the energy contained in the incident wavepacket (details in Section 7.5).

Cavity spectroscopy

Before or after the transfer process, the relative population of each cavity number peak can be obtained directly from the transmon spectrum, which depends on the photon number in the cavity (Section 3.5). The spectrum is fit to a series of Gaussians (with total area normalized to one), which directly yields the relative photon number occupations when the transmon is in its ground state. The spectroscopy data from a typical experiment are shown in Figure 7.9.

Energy transfer efficiency

To measure the overall transfer efficiency, we prepare few-photon states and apply both release and capture pulses. We measure cavity populations before and after the transfer using the photon number-dependent transmon spectroscopy technique just described, which directly provides the relative populations of the

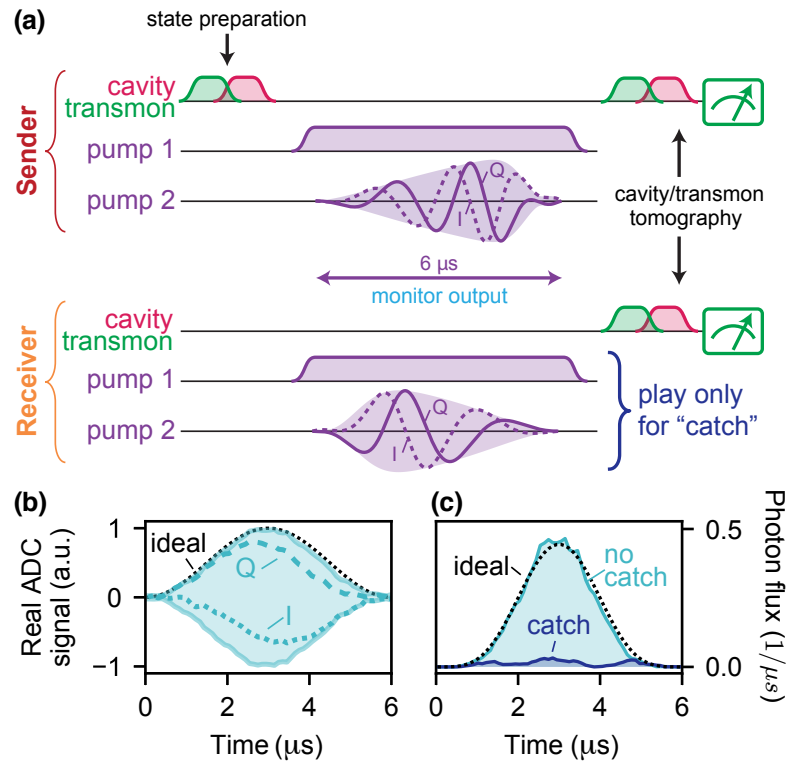


Figure 7.8 | **Temporal mode-matching of sender and receiver.** **a** The transfer protocol consists of state preparation, pumping, and measurement. The sender and receiver are pumped simultaneously because the propagating wavepacket’s time-of-flight (< 3 ns) is negligible compared to its width ($6 \mu\text{s}$). The in-phase (I) and quadrature (Q) components as well as absolute value (shaded) of the pump modulation tones are shown. **b** Measured field emission components (I, Q) and envelope (shaded) upon release of an $\alpha = 1$ coherent state from the sender. **c** Photon flux with receiver pumps omitted (“no catch”) or applied (“catch”), demonstrating absorption of the wavepacket. Detector photon flux calibration is inferred by measuring cavity population after release of a one-photon state. The “ideal” lines mark the shape and amplitude expected for a complete, uncaught release. (Figure adapted from [86]; see [Copyright Permissions](#).)

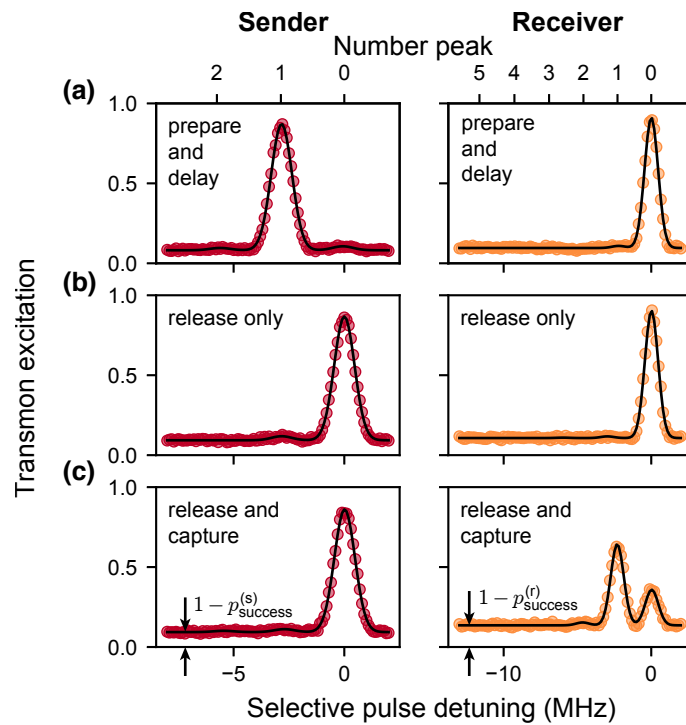


Figure 7.9 | **Spectroscopy to determine transfer efficiency.** We execute the transfer protocol after preparation of a Fock state with $|1\rangle$. The transmon spectrum yields the population of cavity number states. Measurement follows different combinations of pumps on the two systems (**a**: the state is prepared in the sender and $6 \mu\text{s}$ elapse before measurement; **b**: pulses are played on the sender but not the receiver; **c**: the standard transfer protocol is run). Additionally, these data illustrate the rise in background excitation level associated with pumping either system, related to the success probabilities $p_{\text{success}}^{(s,r)}$ (Appendix A.3). (Figure adapted from [86]; see [Copyright Permissions](#).)

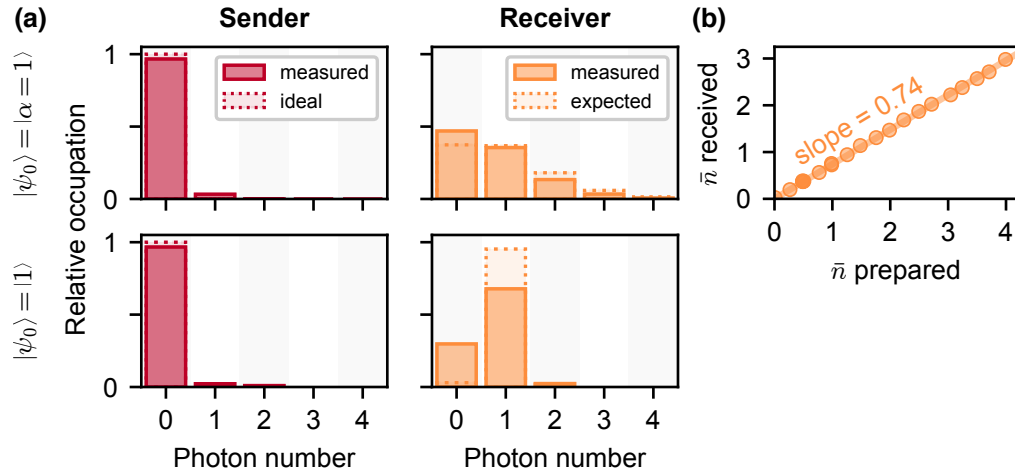


Figure 7.10 | **Cavity populations after pitch and catch of single and many-photon states.** **a** Photon number state occupations of the sender (left) and receiver (right) memory cavities after transfer of a $\bar{n} = 1$ coherent state (upper) or $n = 1$ Fock state (lower). “Expected” occupations in the receiver are those measured in the sender after substituting the transfer process with a delay of the same length. **b** States with $\bar{n} > 1$ are transferred with efficiency near those of single-photon states. The line corresponds with the efficiency $\eta = 0.74$ measured for single-photon states. (Figure adapted from [86]; see [Copyright Permissions](#).)

cavity number states [130]. We define the transfer efficiency η as the average photon number received divided by the average photon number prepared in the sender. Figure 7.10a presents the populations of both memories after transferring states with mean photon number $\bar{n} = 1$, which are obtained from the raw data in Section 7.3.1. From this, we calculate an efficiency $\eta = 0.74 \pm 0.03$.

Our experimental scheme is designed to be independent of input state. We can verify this by preparing a selection of Fock states ($|n\rangle$), Fock state superpositions ($(|n_1\rangle + |n_2\rangle)/\sqrt{2}$), and coherent states ($|\alpha = n\rangle$), with photon number up to $\bar{n} = 4$. By transferring these states and performing the same analysis of cavity populations, we are assured that this single-photon efficiency is, in fact, a photon-number-independent efficiency (Figure 7.10b). (There is no explicit limitation to states with $\bar{n} = 4$; rather, it becomes inconvenient to prepare them using our optimal control techniques, in the case of Fock states, or it becomes difficult to perform spectroscopy on large coherent states because of the large spread in frequency.)

While this transfer efficiency is high, understanding the origin of process imperfections is critical to select optimal error correction protocols and to correct imperfections in future experiments. We can identify




Figure 7.11 | **Cavity populations over time.** We prepare and transfer the Fock state with $n = 1$. Populations $P(n)$ for the lowest two cavity number states ($n = 0, 1$) are measured in both cavities as a function of time. The measured shape of the trajectory in the receiver reflects the inefficiency of the transfer. (Figure adapted from [86]; see [Copyright Permissions](#).)

several factors that contribute significantly to transfer inefficiency: undesired transmon excitation, imperfectly shaped pump pulses, and loss in the transmission path. Each will be addressed in greater detail in Section 7.5.

7.3.2 Time-dependent populations

We can verify that we understand the dynamics of the system by also measuring the cavity population as a function of time.

Through the numerical calculation of pulse shapes ξ_1 and ξ_2 on the sender and receiver (Section 7.2), we know the ideal intracavity field expectation values \hat{a} , \hat{b} in each system as a function of time. These are directly related to the average photon number we would expect to measure at that time, $\bar{n} = a^2$.

We measure the photon number present in the sender and the receiver by truncating the protocol prematurely: the non-dynamic pulses, $\xi_2^{(s)}$ and $\xi_2^{(r)}$, are turned off over 200 ns at a swept measurement time T . We then measure the population $a^{(s)}(T)$ and compare with the calculated value. Similarly, for the receiver population, the full complement of pulses is played on the sender system. At the desired measurement time T , the normally-constant pulse $\xi_2^{(r)}$ is set to ring down, thereby obtaining $a^{(r)}(T)$. The results from these measurements are in good agreement with the expected values (Figure 7.11).

7.3.3 Quantum communication: coherent state transfer

The demonstrated transfer efficiency should allow for quantum communication between the sender and the receiver memories. This means, more precisely, that should be able to prove that we have established a channel capable of transmitting quantum information, not simply a partial reconstruction via classical means. Measurements of cavity population so far, which amount to measurements of energy, are not yet sufficient to show that our prepared states have survived the transmission intact and with their quantum properties retained.

We confirm this explicitly by transferring an overcomplete set of qubit states in the manifold spanned by the Fock states $|0\rangle$ and $|1\rangle$, and by then performing Wigner tomography on the receiver. Maximum-likelihood estimation (MLE) reconstruction is performed on measured Wigner tomography data to return the density matrix found to have most probably produced the data. We use a numerical, iterative, convex optimization algorithm. The algorithm is supplied with normalized (physical) Wigner tomography data, and the optimal reconstructed density matrix is constrained to remain physical throughout the minimization procedure.¹ In this way, we are able to reconstruct the density matrix from each prepared and measured Wigner function (Figure 7.12).

Another useful representation of the reconstructed density matrices, consist with the transfer of a qubit of information, is as vectors on the Bloch sphere (Figure 7.13). In the ideal, lossless case, the cardinal points of the manifold should have Bloch vectors corresponding with $| -Z \rangle$, $| +Z \rangle$, $| +X \rangle$, $| -X \rangle$, $| +Y \rangle$, and $| -Y \rangle$. In the case of our lossy channel, they will fall short; we institute a model to explain this.

Our model assumes a loss channel Ψ that suffers from pure photon loss only, and is characterized by the single-photon efficiency $\eta = 0.74$. In the experiment, this loss potentially occurs throughout the transfer process, distributed spatially and temporally. We can generate density matrices expected in the presence of Ψ in the following way. First, we apply a beamsplitter operation to the product of two states, our ideal state and an ancillary vacuum state, at a variable beamsplitter incidence angle θ . This angle is

1. Through systematic simulation as well as comparison with other reconstruction methods, we find that this reconstruction technique misrepresents the collected data $< 1\%$.

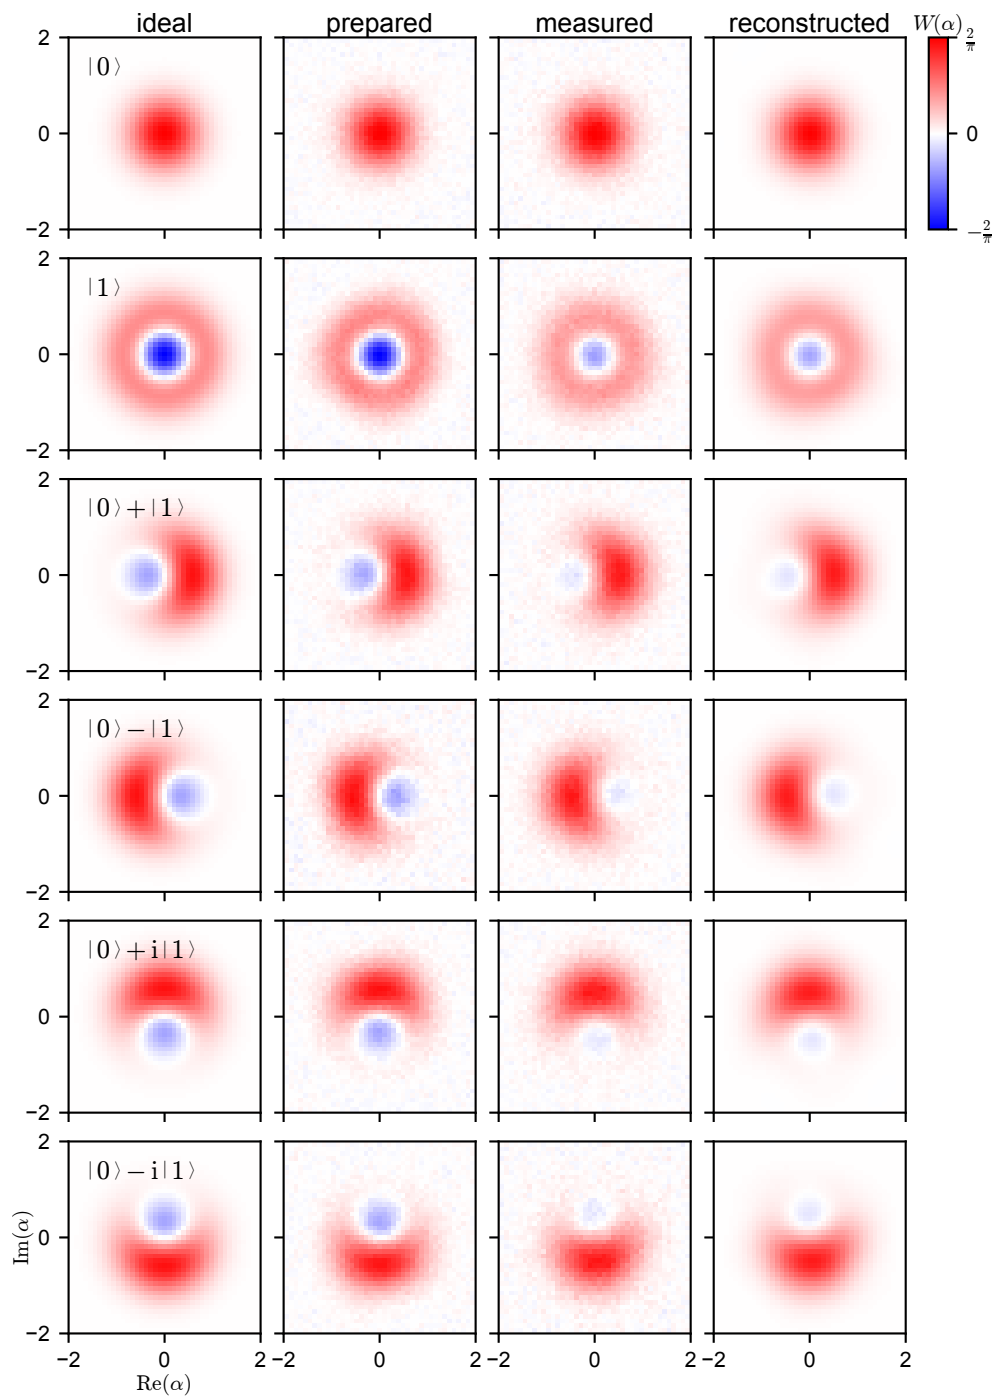


Figure 7.12 | **Reconstructing Fock manifold Wigner functions.** Ideal and measured Wigner functions of the six cardinal states encoded in the single-photon manifold $\{|0\rangle, |1\rangle\}$ as prepared in the sender and received and reconstructed following the transfer protocol. (Figure adapted from [86]; see [Copyright Permissions](#).)




Figure 7.13 | **Bloch sphere representation of Fock manifold transfer.** The prepared (red) and measured (orange) points correspond with six cardinal states in the single-photon encoding, and are connected to form an octahedron representative of the manifold. The green lines extending from each prepared-state point towards $|0\rangle$ are loss trajectories for varying transfer efficiency from $\eta = 0$ to $\eta = 1$. A three-dimensional perspective (shaded yellow) as well as third-angle projections (centered on the XZ plane) are all representations of the same data. (Figure adapted from [86]; see [Copyright Permissions](#).)

related to our measured inefficiency, $p_{\text{loss}} = 1 - \eta$, as

$$\theta = 2 \arccos \sqrt{1 - p_{\text{loss}}}. \quad (7.13)$$

The beamsplitter partly entangles our Hilbert space of interest with an environmental degree of freedom. By tracing over the ancillary system, we then obtain a lossy mixed state that should represent the effect of photon loss. In Figure 7.13, the path predicted by this model for each Bloch vector is plotted in green. In the Fock encoding, photon loss has the effect of reducing all cardinal states to $|0\rangle$ monotonically.

Comparing each received state to the corresponding ideal prepared state, we obtain a fidelity related to the overlap of the ideal and received-and-reconstructed density matrices. From these, we determine an average fidelity $\mathcal{F}_{\text{avg}} = 0.87 \pm 0.04$. In general, the arrows in the 3D representation of the Bloch sphere in Figure 7.13 indicate that this process is consistent with the measured fidelity. (We should expect $\mathcal{F}_{\text{avg}} = 0.91 \pm 0.03$ from pure loss alone.) Some of the measurement entires comprising these data, however, lack any useful information, due to a stochastic failure of the detector (the receiver transmon). This sub-unity “success probability” lowers our “fully deterministic” fidelity to a level $\mathcal{F}_{\text{avg,d}} \geq p_s \times \mathcal{F}_{\text{avg}} = 0.76 \pm 0.04$. The origin of this behavior will be discussed extensively in Section 7.5.1. Both the conditioned and deterministic fidelities significantly exceed the classical bound of $\frac{2}{3}$, the maximum attainable fidelity with which one can reconstruct an unknown qubit state using only classical communication [216]. In the remaining measurements, conditioned data (discounting this effect) will be shown, while for metrics, both values will be given.

More quantitatively, we find that the measured transfer has a process fidelity of 0.95 to this single-source model, bounding the errors not described by photon loss at the 5% level. We can test the accuracy of our photon loss model by reconstructing a process matrix. The process matrix χ defines the mapping from the ideal state density matrix ρ_i to the final, received density matrix ρ_f . In this definition, ρ_i is a two-dimensional logical qubit state in the manifold $\{|0\rangle_L, |1\rangle_L\}$, while ρ_f is in the high-dimensional physical Hilbert space of the cavity: $\{|0\rangle, |1\rangle, \dots, |d-1\rangle\}$. For our purposes, we truncate this space at dimension $d = 5$; for the Fock encoding, no significant population is found outside this space upon reconstruction.

The process matrix is then defined as the map from 2×2 logical density matrices to $d \times d$ phys-

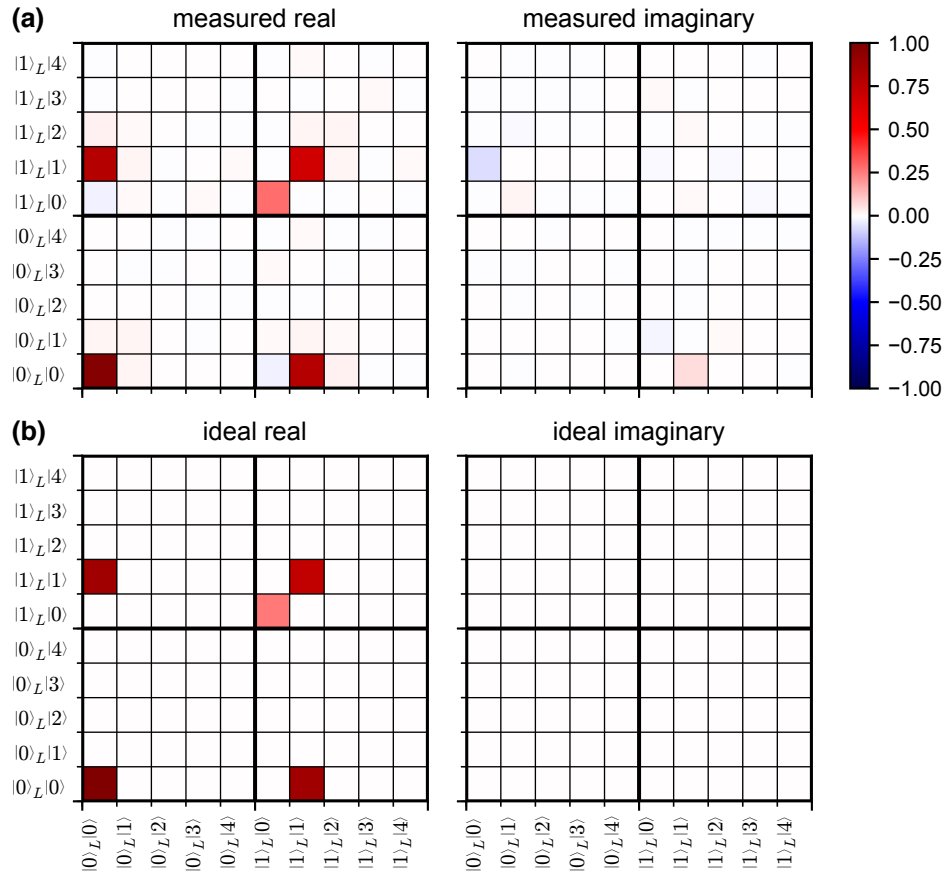


Figure 7.14 | **Photon loss process matrix.** **Top:** Reconstructed process matrix for state transfer in the Fock encoding. The subscript L denotes the logical qubit space, while the kets with no subscript reside in the physical space of the cavity. **Bottom:** Ideal process matrix assuming transmission loss of the measured $\eta = 0.74$. (Figure adapted from [86]; see [Copyright Permissions](#).)

ical density matrices. The reconstructed process matrix χ_m is visualized in Figure 7.14 and compared to the ideal process χ_i assuming only the measured loss $\eta = 0.74$. The process fidelity $\mathcal{F}_{\text{process}} = \frac{1}{4} \text{Tr}(\sqrt{\sqrt{\chi_i} \chi_m \sqrt{\chi_i}})^2 = 0.95$, implying very close agreement with our model of photon loss alone. We believe the source of this 0.05 infidelity also produces the ~ 0.05 average state infidelity (relative to what is expected from the measured inefficiency; as stated previously, $\mathcal{F}_{\text{avg}} = 0.87 \pm 0.04$, while we expect 0.91.)

7.3.4 Multi-photon state transfer

Because it appears that our infidelity of our state transfer protocol is dominated by errors of a single type—photon loss—the scheme can be improved by selecting an appropriate error-correcting code. The use of

cavity memories grants access to tools already developed for correcting loss in stationary states, such as redundantly encoding a qubit within a larger Hilbert space and using photon number parity as an error syndrome [47].

We choose a simple error-correctable code with minimal overhead, the $L = 1$ binomial encoding, which has logical basis states $\{|0\rangle_L = |2\rangle, |1\rangle_L = (|0\rangle + |4\rangle)/\sqrt{2}\}$ [52]. Our transfer scheme is number-state independent, and so with no other modifications we prepare and transmit the cardinal states of this encoding, again measuring the received state with Wigner tomography (Figure 7.15). However, in this case where multi-photon states are used, the results are skewed by a higher-order nonlinearity pertaining only to these states, the Kerr effect.

Kerr effect on binomial code states

The Kerr effect, first observed in the conversion experiment of Chapter 6, was introduced in detail in Section 6.3.3. It only manifests in the Wigner function for states that are not radially symmetric (and contain more than one photon number). Because the Kerr evolution is deterministic², it does not in and of itself lead to a loss of information, so long as its value is known. In order to produce a fidelity representative of the effect of the transfer protocol on the prepared states, and not merely the effect of Kerr at one end or the other, we compare the reconstructed density matrix of the prepared state with a reconstructed density matrix of the received state to which an in-software correction has been applied. The Wigner functions of the resulting ‘de-Kerred’ state are shown in the rightmost column of Figure 7.15.

This operation is simply the Kerr unitary associated with H_{Kerr} in Equation 6.34, with an effective Kerr χ_{aa}^{eff} , equal to approximately the average of χ_{aa} in each system. In order to determine the nominal value of χ_{aa}^{eff} to use for the software correction unitary, we perform a global minimization of the average infidelity over all six states prepared in the manifold with respect to a single Kerr evolution applied to each state. We find $\chi_{aa}^{\text{eff}} = 2\pi \times 10.8$ kHz for our transfer time of 6 μs .

2. Since Kerr evolution does not commute with photon loss, the uncertainty in exactly when the loss occurred does lead to a small unrecoverable loss of information, which manifests as dephasing of the cavity state. Because the Kerr is fairly small and the transfer time quite short, this produces a small effect on the measured average fidelity, around 1–2%.

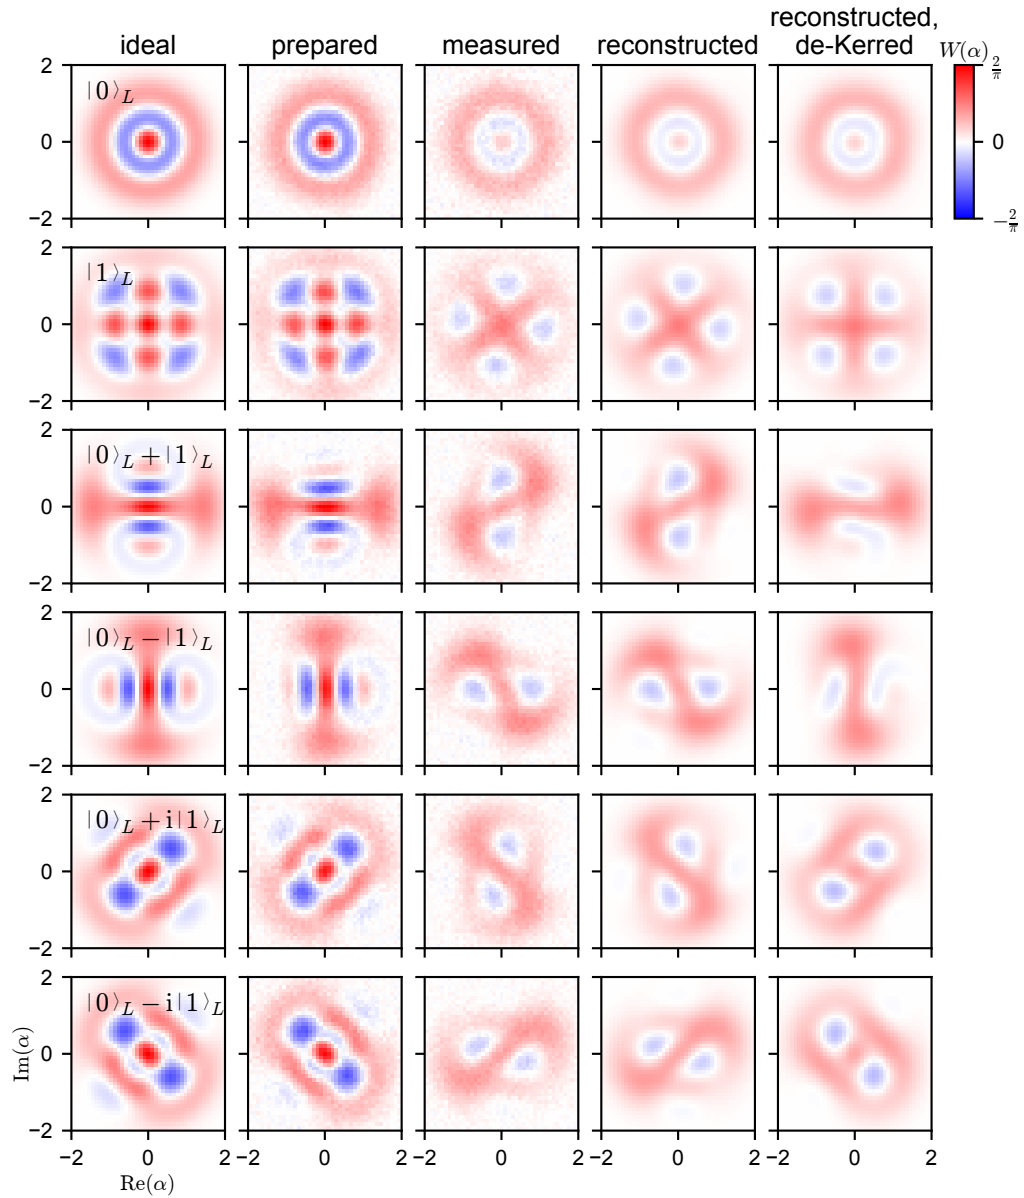


Figure 7.15 | **Reconstructing binomial manifold Wigner functions.** Ideal and measured Wigner functions of the six cardinal states encoded in the binomial code manifold $\{|0\rangle_L, |1\rangle_L\}$ as prepared in the sender and received and reconstructed following the transfer protocol. Kerr evolution during the transfer results in a distortion in the received state Wigner functions, an effect that is accounted for by “de-Kerring” the reconstructed received state. (Figure adapted from [86]; see [Copyright Permissions](#).)




Figure 7.16 | **Binomial encoding loss trajectories.** The prepared (red) and measured (orange) points correspond with six cardinal states in the binomial encoding, and are connected to form an octahedron representative of the manifold. The green lines extending from each prepared-state point towards $(0, 0, 0.5)$ are loss trajectories for varying transfer efficiency from $\eta = 0$ to $\eta = 1$. In the three-dimensional perspective (shaded yellow), green dotted arrows indicate the expected shrinkage given the measured efficiency. Third-angle projections (centered on the XZ plane) reflect the same data. (Figure adapted from [86]; see [Copyright Permissions](#).)

Binomial states within the Bloch sphere

We can also depict the prepared and de-Kerred data on the Bloch sphere (Figure 7.16). This time, the Bloch sphere represents the logical qubit, rather than the two-level system of a single photon. In the binomial encoding, photon loss reduces cardinal points towards the center of the Bloch sphere uniformly, because single-photon errors cause a portion of the state to leave the even-parity codespace. For larger η , the cardinal points begin to turn towards the point $(0, 0, 0.5)$ on the Bloch sphere, since the $|1\rangle_L = (|0\rangle + |4\rangle) / \sqrt{2}$ state overlaps with the vacuum state and $|0\rangle_L = |2\rangle$ does not. Note that our received states appear in the uniform-shrinkage regime, which corresponds to loss small enough that single photon loss errors dominate over higher-order effects.

The states of the binomial code have larger average photon number ($\bar{n} = 2$) relative to the single-photon encoding ($\bar{n} = 0.5$), representing an additional “overhead”. From this increased sensitivity to




Figure 7.17 | **Comparison of trivial and error-corrected fidelities.** Projected performance of the two demonstrated encodings as a function of transfer efficiency. Calculated average infidelity (lines) of Fock and binomial encodings, and measured values (circles) at the present inefficiency (arrow). The performance of an ideal parity detection-and-correction scheme is simulated for the binomial encoding. The resulting infidelities are shown for perfectly prepared states subjected to varying transfer inefficiency (line), and for the measured received states (diamond). Points lie above their respective lines due to small additional infidelity not described by photon loss. The shaded region indicates the regime beyond the break-even point, displaying the growing improvement between the encodings as efficiency improves. Elimination of either transmission loss or release and capture inefficiency would push the corrected infidelity further into this regime ($\eta \sim 0.86$, dotted line). (Figure adapted from [86]; see [Copyright Permissions](#).)

photon loss we predict a mean fidelity of 0.60 and measure $\mathcal{F}_{\text{avg}} = 0.54 \pm 0.04$ in this manifold ($\mathcal{F}_{\text{avg,d}} \geq 0.47 \pm 0.04$) relative to the ideal states. Though its mean fidelity is lower than that of the single-photon encoding, the binomial encoding permits the use of parity as an error syndrome measurement.

7.3.5 Using binomial states to protect against photon loss

This feature can enable detection and deterministic correction of single-photon loss errors. Above some transfer efficiency threshold, an error-corrected qubit would be transmitted with higher average fidelity than a qubit encoded in the single-photon manifold. Our efficiency ($\eta = 0.74$) exceeds that of this “break-even” threshold ($\eta \sim 0.67$), defined as the crossing of simulated mean fidelities in each case (Figure 7.17).

The green line in Figure 7.17 is calculated by applying our photon-loss-only model to the six cardinal states of the binomial encoding. At each efficiency value, the fidelity of the resulting (partially mixed) density matrices are compared to those of each original ideal state, and the average infidelity is reported. The dashed green line is then calculated from these density matrices by applying the error-correction procedure described in Section 7.3.5.

Beyond break-even, error correction can overcome the overhead associated with the binomial state




Figure 7.18 | **Conditioned spectroscopy of transferred binomial codewords.** Spectroscopy is performed following three conditioning measurements on even parity and transmon state.

encoding. Error correction is possible using high-fidelity quantum non-demolition (QND) parity measurements [47], which could be effectively realized by adding a dedicated readout channel to each module. Following this measurement, fast feedback could be used to apply a conditional correction pulse. Along with modest improvements to the release and capture efficiencies, error correction should place the transfer firmly within this advantageous regime. Optimization over a wider set of bosonic codes could produce even larger improvement. These error correction concepts can also be extended to improve entanglement fidelity without sacrificing the determinism of the protocol.

Testing and improving error correction

We can attempt to perform the first step of error correction — error detection — using the system in this experiment. We perform a parity measurement, preceded and succeeded by additional conditioning measurements that verify the transmon remains in the ground state, since our parity measurement critically relies on this being the case. We perform spectroscopy on the received cavity following transmission of each of the binomial codewords; data that has been post-selected on these conditions is shown in Figure 7.18 and is compared to unconditioned data.

The conditioned data shows lower background, indicative of the exclusion of background transmon excitations. In both cases, it also shows a significant reduction in the height of the odd-number peaks, like $|1\rangle$ and $|3\rangle$. In the case of the $|2\rangle$ state, the height at $n = 2$ actually increases; this suggests that fidelity of

the state may have improved.

To be certain of the effect on state coherence, we perform the same test but now measure the conditioned and unconditioned Wigner functions of the two states. Figure 7.19a shows the Wigner functions of non-conditioned data from an experiment where the conditioning measurement *has not been performed*. The absence of three additional measurements appears to preserve the shape of the Wigner functions. In Figure 7.19b, conditioned data is compared to what is expected by simulating the effect of post-selection. Consistent with the spectroscopy results in Figure 7.18, the $|0\rangle_L$ state improves its contrast; the $|1\rangle_L$ state becomes shrunken and smeared. Neither agree fully with the predicted conditioning result, indicating that our measurements are less than perfect. Notably, the radial symmetry of the $|0\rangle_L$ state means that it is more impervious to the cavity dephasing that the conditioning measurements may induce.

We conclude that error correction of the transfer cannot be performed effectively in the current experimental sample. Because a single transmon simultaneously provides the functions of state preparation and cavity tomography as well as supplying the nonlinearity needed for the conversion process, it is not possible to optimize the hardware parameters for all of these essential features. Most notably, a large conversion rate between each memory and output mode requires a large cross-Kerr χ_{ab} .

However, this means that resonantly driving the output mode into a coherent state (as is done during transmon readout) dephases the cavity due to the dispersive shift combined with uncertainty in the occupation of the output mode [217]. Therefore, parity measurement without significant dephasing of the cavity is impossible in the current configuration. As indicated in Figure 7.20, the addition of a dedicated, separate chip with ancilla transmon and readout mode, also with small cross-Kerr between readout and cavity, would allow for high-fidelity QND parity measurement [47], while maintaining the ability to rapidly convert between memory and propagating modes. This modular distribution of functionality may also allow modifications to the conversion side of the device to mitigate some of the non-idealities discussed above; for example, some nonlinear element other than a transmon can be used to facilitate conversion [122, 201].

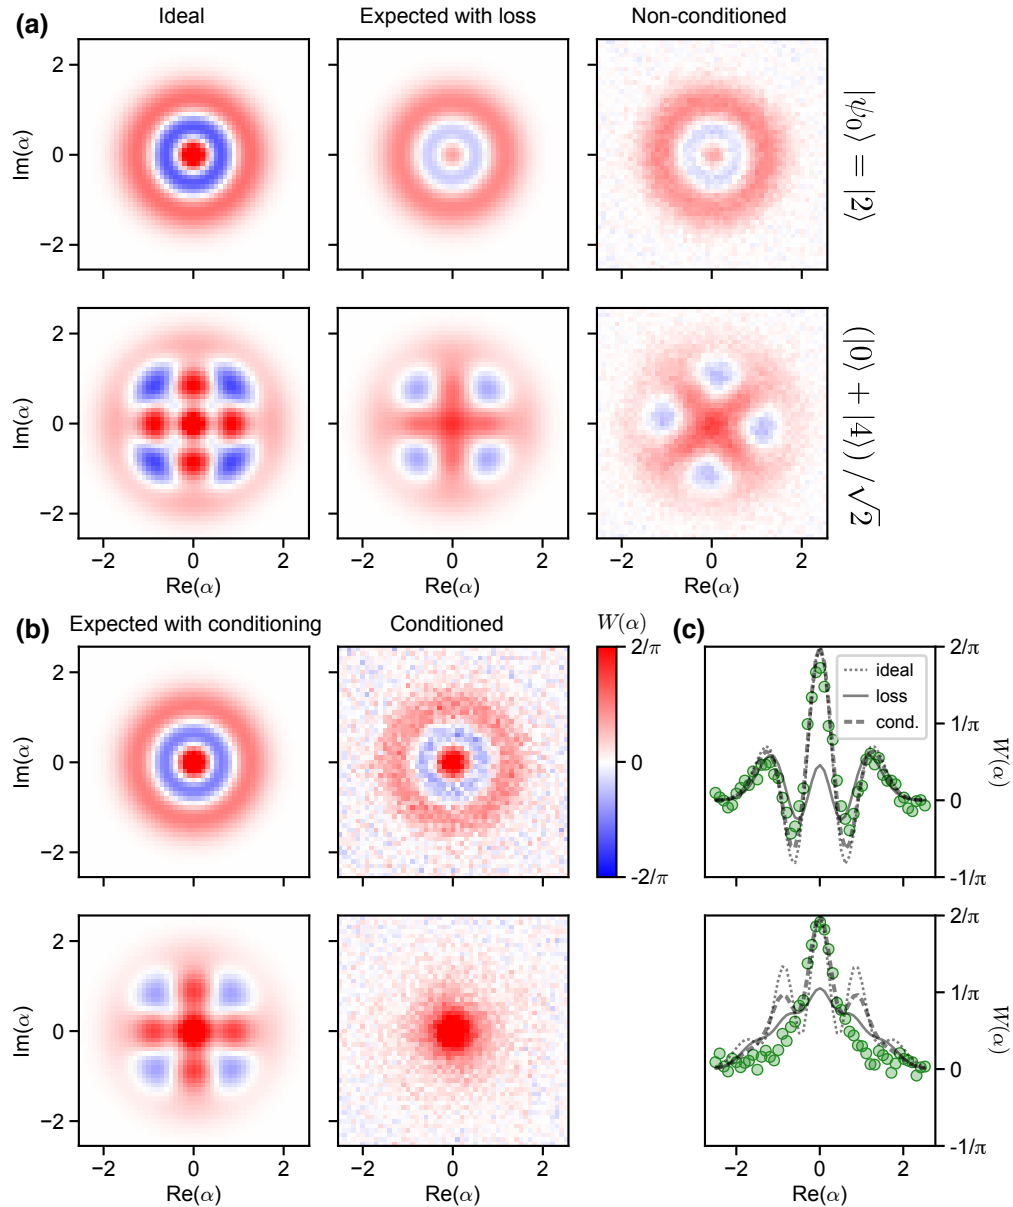


Figure 7.19 | **Conditioned tomography of transferred binomial codewords.** **a** Ideal, simulated lossy, and non-conditioned Wigner functions of the two binomial codewords. **b** The simulated results of lossy and parity-measurement-conditioned data, compared to the actual conditioned data. **c** Line cuts of ideal, predicted lossy, and predicted lossy-and-conditioned Wigner functions compared to data. The data generally falls between the lossy and conditioned simulations.




Figure 7.20 | **Modified design to permit error correction.** In this modified design of the experimental apparatus, an independent transmon and readout chip performs the tasks of state preparation, cavity tomography, and parity measurements. The communication chip features a dedicated conversion element and communication mode. Parameters can be designed to suit each set of tasks appropriately.

Implementation of binomial code error correction

To implement this error-correction procedure, the density matrix for a given state is split into two components: the even and odd photon number parity subspaces. A different correction unitary is applied to each subspace. To correct for photon loss, the odd-parity subspace is corrected via the operation:

$$|1\rangle \rightarrow |2\rangle, \quad |3\rangle \rightarrow \frac{1}{\sqrt{2}} (|0\rangle + |4\rangle) \quad (7.14)$$

The correction in the even-parity subspace takes the form of a rotation that adjusts the relative weights of the $|0\rangle$ and $|4\rangle$ states:

$$\begin{aligned} |2\rangle &\rightarrow |2\rangle, \\ |0\rangle &\rightarrow \cos \theta_c |0\rangle + \sin \theta_c |4\rangle, \\ |4\rangle &\rightarrow -\sin \theta_c |0\rangle + \cos \theta_c |4\rangle \end{aligned} \quad (7.15)$$

This is to account for the fact that the no-parity-jump event alters the relative probability amplitude of the state $(|0\rangle + |4\rangle) / \sqrt{2}$ (essentially, Bayesian interference [218]). The two components of the density matrix are then recombined with their respective probabilities unchanged, resulting in a higher-purity corrected density matrix. The optimum rotation angle θ_c depends on the probability of loss, and thus on

inefficiency. For each value of the transfer inefficiency, we calculate the angle of rotation that minimizes the mean infidelity over the manifold. (Ignoring this optimization is simulated to produce a 1% reduction in mean fidelity at $\eta = 0.85$, a 3% reduction at $\eta = 0.74$, and a 9% reduction at $\eta = 0.5$.) This mean infidelity is that which is reported in Figure 7.17.

This procedure is identical to the experimental implementation that will be needed to perform this correction in hardware: measure parity; apply one of two correction unitaries, conditioned on the result of the measurement; and perform unconditional tomography. This process would follow [47] closely, only with the substitution of the appropriate unitaries specific to this correction scheme. Other photon-loss-robust encodings may be more optimal under certain conditions, including at different values of efficiency, and those encodings and corrective unitaries could be substituted in a similar manner [54].

Code optimization

There exist a number of reasonable photon-loss-robust encodings, besides the $L = 1$ binomial code, at our present value of inefficiency [54]. Depending on several system parameters, higher-level bosonic codes might improve correction performance. As an example, the cat code performance can exceed that of the $L = 1$ binomial code at particular values of efficiency. Since the continuous variable nature of the code means that there will be an optimal photon number (α^2), one must compare using the photon number that will maximize the error-corrected mean fidelity. The effect of photon-number-dependent cavity self-Kerr will also be a factor to consider. These encoding optimizations will become important when physical optimizations have been exhausted.

7.4 Half-release, full-catch, and entanglement

Our experimental scheme readily enables us to generate on-demand remote entanglement by applying a pump sequence on the sender that releases half of its stored energy (Figure 7.21a). If the initial state is a single photon, this results in entanglement between the memory and the emitted radiation (Section 6.4).

The interaction between the cavity and the transmission line can be described by an effective beam splitter interaction (Equation 7.16), where the conversion strength and time determine the coupling θ

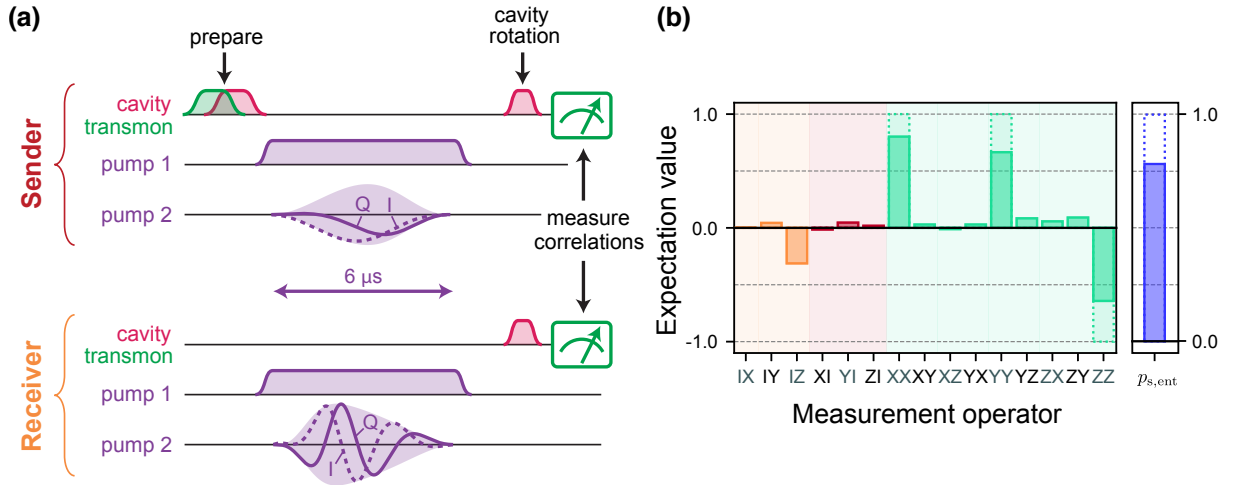


Figure 7.21 | **Entangling a photon between remote cavities.** **a** Single-photon entanglement is generated with a half-release of $|1\rangle$, followed by unchanged capture pulses (shown in the same style as Figure 7.8a). The absolute-valued envelope (shaded) is shown along with in-phase (dashed) and quadrature (solid) components. The capture pulse is state-independent and is thus unchanged. Cavity photon number correlations are measured following rotations into the appropriate measurement basis. **b** Reconstructed expectation values (shaded) of two-qubit Pauli operators compared to those of the maximally-entangled Bell state $|01\rangle + |10\rangle$ (dashed). This measurement, like the others in this chapter, is effectively conditioned by renormalizing data where either transmon did not remain in the ground state (Section 7.5.1). Scaling these Pauli bars by the value of $p_{s,\text{ent}}$ approximates unconditioned data; details are given in the following two sections. (Figure adapted from [86]; see [Copyright Permissions](#).)

between the modes (Section 6.1.4),

$$H_{\text{bs}} = \frac{-i\theta}{2} (\hat{a}_1^\dagger \hat{a}_2 - \hat{a}_1 \hat{a}_2^\dagger) \quad (7.16)$$

In particular, releasing half the energy stored in the cavity initially corresponds to a ‘50:50 beam splitter’ with $\theta = \pi/2$. As specified in Section 7.2, the pulse required to realize this can be calculated precisely by specifying that one-half of the prepared state remains. This leads to a release pulse with a different shape and a lower amplitude. The capture pulse remains the same, because the temporal shape of the traveling wavepacket is the same.

Absorption of the wavepacket by the receiver ideally results in the Bell state $(|10\rangle + |01\rangle) / \sqrt{2}$ shared between the memories. Immediately after switching off the pumps, we apply a mapping pulse that rotates the state of one or both cavities into a particular basis. Within that basis, we then measure the likelihood

of finding the cavities in a particular joint photon-number state by applying selective- π pulses on each transmon and reading out.

In particular, for each module we perform number-selective measurements on $n = 0$ and $n = 1$, as well as $n = -1$ (a background measurement that evokes no cavity response). These $3 \times 3 = 9$ probabilities are combined to produce four values for each rotation: the probability of finding joint photon numbers 00, 01, 10, or 11 (the diagonal elements of the density matrix in this joint basis). We choose three basis rotations for each cavity: $\{I, Y(\frac{\pi}{2}), X(\frac{\pi}{2})\}$, leading to measurements in the $\{z, x, y\}$ bases, respectively. This produces correlation probabilities in $3 \times 3 = 9$ joint bases, enough to reconstruct the full joint state. The data for each rotation are then supplied to a maximum likelihood estimation (MLE) reconstruction program to produce the density matrix describing the joint system with the largest likelihood. From this reconstructed density matrix, we then assemble the Pauli bars in Figure 7.21b, which reveal non-classical correlations between sender and receiver.

Success rate conditioning

The measurement of joint cavity operators is affected by the success probability of each transmon in a way similar to that of spectroscopy and Wigner tomography measurements (Section 7.5.1). Our density matrix and Pauli bars are thus naturally conditioned on the case where both transmons remain in the ground state following the half-transfer and before the measurement is performed. The unconditioned values (and resulting entanglement metrics) can be determined by accounting for the probability with which each transmon becomes excited or the system leaves the observed Hilbert space. While they cannot be directly measured, they can be approximated by scaling the data in Figure 7.21b by the value of $p_{s,\text{ent}}$. (As $p_{s,\text{ent}} \rightarrow 1$, the approximate unconditioned data would begin to equal the conditioned data.)

As opposed to other measurements throughout this work, the excitation of *both* systems factors in. To estimate the unconditioned entanglement measurement operator expectation values, we add an extra dimension to the reconstructed density matrix representing the measured joint cavity state. This dimension represents a space outside of the observable space. It is populated according to the product of the two transmon excitation probabilities. The density matrix is then renormalized and truncated to exclude the ‘unobservable’ dimension, producing a density matrix reasonably estimating what we believe an uncondi-

tioned measurement would produce.

Entanglement fidelity

The resulting entanglement success probability $p_{s,\text{ent}} = 0.78 \pm 0.04$ is lower than for the state transfer for the same reason—that both system losses are sampled. The fidelity of the joint state to the ideal Bell state is $\mathcal{F}_{\text{Bell}} = 0.77 \pm 0.02$ (deterministic: $\mathcal{F}_{\text{Bell,d}} \geq p_{s,\text{ent}} \times \mathcal{F}_{\text{Bell}} = 0.61 \pm 0.02$), confirming the successful generation of entanglement.

One important metric is the entanglement generation rate R , the rate of experimental repetition. In the conditioned case, this omits additional time required to avoid failed runs (which could in principle be heralded against), and is inclusive of all experimental runs in the fully deterministic case. We are able to achieve a net entanglement generation rate of $(140 \mu\text{s})^{-1}$ (for fidelity 0.77; equivalently $(110 \mu\text{s})^{-1}$ for fidelity 0.61, fully deterministic), limited by the average time it takes to reset the system ($\sim 100 \mu\text{s}$). This rate exceeds the single-photon loss in either memory ($< (450 \mu\text{s})^{-1}$), a strict requirement for scaling up the network size. We note that memory lifetimes about five times higher have been measured with identical hardware [126]; we should not yet be limited by this particular rate.

More entanglement metrics

Using the reconstructed density matrix, we are able to calculate a number of additional metrics to describe the entangled state. Values “unconditioned” and “conditioned” on success probability are shown in Table 7.2. Some of these are useful for comparison with other experiments or when considering, for example, choices among purification schemes. Entanglement fidelity, concurrence, and purity have been calculated according to standard definitions [219]. The generation rate of maximally entangled qubit pairs, the ebit rate, represents the rate at which an optimal purification scheme [220] could generate a maximally entangled pair from many pairs identical to ours, and can be calculated using the logarithmic negativity [221], $E_{\mathcal{N}}$:

$$R_e = p_{s,\text{ent}} \cdot R \cdot E_{\mathcal{N}}. \quad (7.17)$$

Success probability $p_{s,\text{ent}}$ enters in both the logarithmic negativity (where lower $E_{\mathcal{N}}$ demands more pairs in the unconditioned case) and the generation rate (where overall generation of higher-fidelity entanglement

Value		Conditioned (measured)	Unconditioned (estimated)
Success probability	$p_{s,\text{ent}}$	0.78 ± 0.04	1
Fidelity	$\mathcal{F}_{\text{Bell}}$	0.77 ± 0.02	0.61 ± 0.04
Concurrence	\mathcal{C}	0.66 ± 0.03	0.51 ± 0.04
Purity	γ	0.66 ± 0.03	0.40 ± 0.04
Generation rate (μs^{-1})	R	$1/(140 \pm 10)$	$1/(110 \pm 10)$
Logarithmic negativity	$E_{\mathcal{N}}$	0.66 ± 0.03	0.30 ± 0.04
ebit rate (kebit/s)	R_e	4.7 ± 0.5	2.7 ± 0.4

Table 7.2 | **Measures of entanglement.** All comparison values are with respect the maximally-entangled Bell state, $(|10\rangle + |01\rangle)/\sqrt{2}$. “Unconditioned” values are those metrics taken for the estimated unconditioned density matrix reconstructed according to the process in Section 7.4.

is slower).

7.4.1 Error correction of entangled states

Error correction of entangled states will be slightly different from the correction of states fully transferred in the binomial code, as an example. When given the option of using many-photon states to generate entanglement, we are not confined to using one scheme in particular. As a convenient and intuitive example, let us examine the process of detecting and correcting photon losses during distribution of a four-legged cat state.

Entanglement distribution in this case would begin with the state

$$|\psi_0\rangle = \mathcal{N} (|\alpha\rangle + |-\alpha\rangle + |i\alpha\rangle + |-i\alpha\rangle)_A |0\rangle_B, \quad (7.18)$$

a four-legged Schrödinger cat state with information encoded in the even parity subspace, and, furthermore, in a subspace that is modulo-4 parity.

Applying the conversion Hamiltonian results in entanglement between states in the sender (A) and

receiver (B) that is of the form

$$\begin{aligned} &|0 \bmod 4\rangle_A |0 \bmod 4\rangle_B + |2 \bmod 4\rangle_A |2 \bmod 4\rangle_B \\ &+ |1 \bmod 4\rangle_A |3 \bmod 4\rangle_B + |3 \bmod 4\rangle_A |1 \bmod 4\rangle_B \end{aligned} \quad (7.19a)$$

$$\begin{aligned} &|0 \bmod 4\rangle_A |3 \bmod 4\rangle_B + |2 \bmod 4\rangle_A |1 \bmod 4\rangle_B \\ &+ |1 \bmod 4\rangle_A |2 \bmod 4\rangle_B + |3 \bmod 4\rangle_A |0 \bmod 4\rangle_B \end{aligned} \quad (7.19b)$$

in the case of a zero or one photon loss, respectively. Note that, for each state, the *total* photon number between A and B adds up to a multiple of 4. Further, in the lossless case, measurement of parity (to clarify, standard 2-parity) on A or B will produce a result $(\mathcal{P}(A), \mathcal{P}(B)) = (0, 0)$ or $(1, 1)$.

After a photon loss, the states in B are all lowered by one. If we measured parity on either side it will now produce $(0, 1)$ or $(1, 0)$ — a mismatch. Depending on having measured the former or the latter, respectively, the projected state remains entangled in either of

$$|0 \bmod 4\rangle_A |3 \bmod 4\rangle_B + |2 \bmod 4\rangle_A |1 \bmod 4\rangle_B, \text{ or} \quad (7.20a)$$

$$|1 \bmod 4\rangle_A |2 \bmod 4\rangle_B + |3 \bmod 4\rangle_A |0 \bmod 4\rangle_B. \quad (7.20b)$$

As in ref. [47], it can be sufficient to track which of these two states has been established, or one can apply a conditional corrective unitary to transfer the state into a more desirable form.

7.5 System imperfections

While our transfer efficiency in this experiment ($\eta = 0.74$) is high, understanding the origin of process imperfections is critical to select optimal error correction protocols and to correct imperfections in future experiments. We can identify several factors that contribute significantly to transfer inefficiency: undesired transmon excitation, imperfectly shaped pump pulses, and loss in the transmission path.

7.5.1 Heralding and determinism

For ideal operation of our protocol, the transmons would remain in their ground states during the transfer. However, we observe non-negligible stochastic excitation during the transfer process due to thermalization and pump-induced transitions to higher levels [222, 223].

Unwanted transmon excitation has two important consequences. For one, an excitation leads to a shift of the resonator frequencies due to their dispersive couplings to the transmon. This abruptly changes the transfer frequency-matching conditions, manifesting as off-resonant emission by the sender, or imperfect wavepacket absorption by the receiver. We estimate these effects to lead to an inefficiency of about 2% for emission, and 6% for absorption. This effect is thus likely the dominant cause of the measured absorption inefficiency.

Secondly, transmon excitation precludes effective measurement of the cavity state. In this case, cavity measurement indiscriminately returns ‘yes’ to a query of any photon number. Excitations thus have the effect of reducing average measurement contrast. By normalizing our measurement data to correct for this, cavity tomography is implicitly conditioned on the transmon having remained in its ground state. It is therefore useful to view the transmon excitation probability as a “failure probability” of the protocol, i.e., we make the conservative assumption that each excitation masks an unsuccessful transfer. The efficiency η quoted throughout this chapter is then, if not stated otherwise, conditioned on the receiver transmon remaining in the ground state, with success probability $p_s = 0.87 \pm 0.03$. The conditioned value can be interpreted as the efficiency that would be measured (i) with a perfectly cold transmon or (ii) by heralding on a transmon measurement after the protocol. The “deterministic efficiency” given the transmon temperature observed here is estimated by the product of the conditioned efficiency and the success probability, $\eta_d \geq p_s \times \eta = 0.87 \times 0.74 = 0.64 \pm 0.03$. This deterministic value represents a lower bound on the quantity; since these failure events are assumed to be maximally destructive, this is the worst-case scenario.

Another way to think about transmon excitations “reducing measurement contrast” is that data that is effectively ignored by normalization conveys no information. We confirm this by performing a spectroscopy experiment, measuring the transmon state after the transfer process. We post-select on the measurement results in each system when the transmon is found in the ground state. Within fit and mea-

surement error, we find identical photon number content in the unselected (including all data) and post-selected cases (Figure 7.22). In other words, the average efficiency η does not change whether a measurement to detect transmon state is performed, and the discarded measurements contain no relevant information. Therefore, a measurement of the transmon ground state can be thought of as heralding on a successful state transfer (a transfer not precluded by an endpoint failure, yet still subject to miscalibrations and intermediate photon loss).

These deterministic fidelity quantities, however, are only lower bounds. In our calculation of the efficiency, we assume that no energy was captured in the event that the transmon is found excited. In reality, some energy was likely absorbed in the time before the transmon became excited. For each fidelity bound, we take the worst-case assumption: that the entire system has left the codespace entirely in an unrecoverable way, and these events contribute zero fidelity to the average. In fact, if the failure completely destroys the information, but leaves the system in a state which can be detected and reset (for example, the transmon in $|e\rangle$), then these events can result in an average fidelity 0.5 by resetting the system to an arbitrary state within the codespace. The same is true for the entanglement (but with fidelity 0.25, which is the best that can be prepared with local resources). Because some of the failure events are in principle recoverable in this fashion, we therefore state our worst-case deterministic fidelities with the caveat that they are lower bounds.

7.5.2 Release and capture efficiency

Reductions to transfer inefficiency come from photon loss in the transmission path, which we estimate adds 15%, as well as imperfect pulse shapes affecting state release and capture, each with an effect around 2%. We note that the bulk of the described imperfections are not fundamental; in particular, improvements to the transmon equilibrium temperature and thermalization rate as well as parameter engineering to avoid pump-induced higher order transitions (Appendix A.3) can substantially reduce the inefficiencies resulting from transmon excitation.

The efficiency of the transfer process can be obtained directly by measurements before and after the transfer protocol. Its origin is investigated here, by breaking down the “total process efficiency” η into segments of the process, such that future incarnations of the experiment may be more improved. A summary

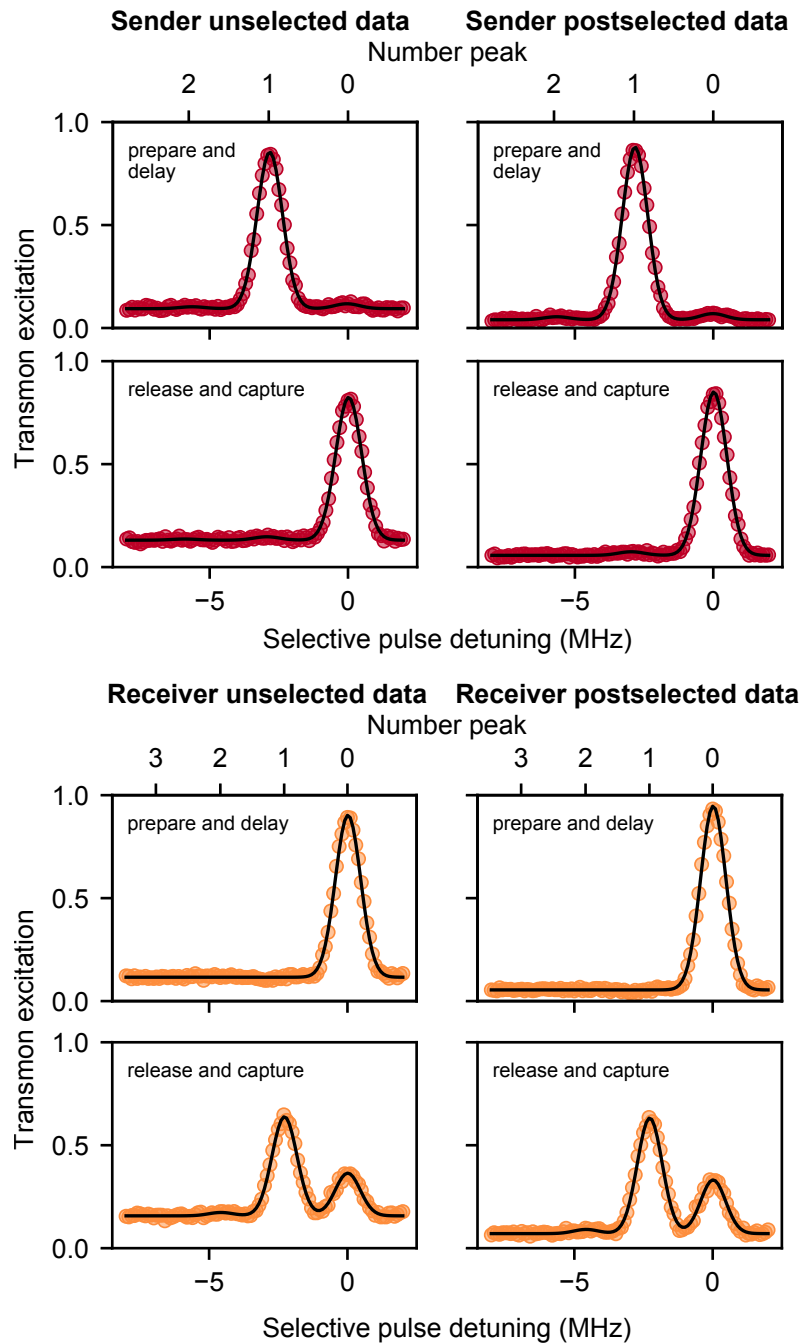


Figure 7.22 | **Comparison with post-selected efficiency.** Transmon spectroscopy is performed with an added measurement of transmon state following the transfer process. All of the data (left column) or only data post-selected on the conditioning measurement (right column) are shown after preparation (top row) or after release and capture (bottom row). Black lines indicate fits from which relative photon number populations are obtained. While background offsets are clearly lower in the post-selected data, the extracted ratio between prepared and received populations change by $< 1\%$ (within uncertainty). (Figure adapted from [86]; see [Copyright Permissions](#).)

Category		Value	Uncertainty	Notes
<i>Measured values</i>				
Unreleased energy fraction	$n_{\text{remain}}^{(s)}$	0.033	0.005	
Uncaught energy fraction	$n_{\text{reflect}}^{(r)}$	0.068	0.005	
Transmission loss	$\eta_{\text{tx,meas}}$	0.80	0.15	
Total efficiency	η	0.74	0.03	
<i>Modeled values</i>				
Truncation of sender pulse	$\eta_{\text{trunc}}^{(s)}$	0.99	-	specified in pulse creation
Sender transmon excitation	$\eta_{\text{excite}}^{(s)}$	0.98	0.01	
Calibration of sender pulse	$\eta_{\text{miscal}}^{(s)}$	0.98	0.005	
Release efficiency	η_{release}	0.95	0.01	
Transmission line loss	η_{tx}	0.85	0.04	estimation from “missing energy”
Truncation of receiver pulse	$\eta_{\text{trunc}}^{(r)}$	0.95	-	specified in pulse creation
Truncation of receiver pulse, revised	$\eta_{\text{trunc,corr}}^{(r)}$	0.99	0.005	
Receiver transmon excitation	$\eta_{\text{excite}}^{(r)}$	0.94	0.02	
Calibration of receiver pulse	$\eta_{\text{miscal}}^{(r)}$	≥ 0.99		
Capture efficiency	η_{capture}	0.92	0.02	

Table 7.3 | **Summary of transfer losses.** A break-down of loss mechanisms thought to contribute to the transfer efficiency. For some categories the individual constituent efficiencies are measurable, while for others only the total loss can be measured, and the components must be deduced. The model is designed such that the product of the three category subtotals, $\eta_{\text{release}} \times \eta_{\text{tx}} \times \eta_{\text{capture}}$, equals the measured total efficiency η .

of the different sources of inefficiency can be found in Table 7.3.

Truncation

Amplitude and time constraints of the transfer process drives are imposed by the physical implementation of the experimental hardware. With finite time and amplitude constraints on the conversion rate $g(t)$, releasing and capturing the entirety of the energy in the sender is not possible. To limit the applied pump power (and thus the amount of heating induced to the system), we implemented the full transfer process specifying waveforms that are expected to release the fraction $\eta_{\text{trunc}}^{(s)} = 0.99$ of the sender state, and capture $\eta_{\text{trunc}}^{(r)} = 0.95$ of the energy in the receiver.

The release pulse calculated by the method described has several constraints imposed on it by hardware

and by the protocol itself. The pulse amplitude and its derivative must reach zero at the endpoints, and the maximum amplitude must remain below some threshold. These impositions produce a remaining energy in the sending cavity, $a_s^2(t)$, that decays asymptotically in time, towards zero. We choose to truncate this decay such that 1% of the energy will remain in the sending cavity ($\eta_{\text{trunc}}^{(s)} = 0.99$). If more pump power were applied, or more time were taken, a larger fraction of the prepared state's energy would be released. However, these adjustments would also affect the prevalence of other errors (such as transmon excitation or readout infidelity). Similarly, the receiving pulse is constrained in time and amplitude such that $\eta_{\text{trunc}}^{(r)} = 0.95$ (a fact visible in the predicted populations over time in Figure 7.11).

Sender transmon excitation

During the transfer, we observe the sender transmon excitation from 3.1% to 9.6%, a process from which we approximate a worst-case transfer inefficiency of $\sim 6\%$. This accounts for varying failure-probability-per-photon, since both transmon excitation and photon population change in the cavity over time as it empties. We perform a control experiment, populating the sender transmon with known probability $P(e)$ into a superposition state between $|g\rangle$ and $|e\rangle$ before executing the transfer, which allows us to refine our approximation. A fully-excited sender transmon still transfers somewhat successfully (Figure 7.23), introducing an inefficiency of only $\sim 39\%$. The failure rate varies linearly with transmon polarization—from which we calculate that a *total efficiency* of only ~ 0.46 could be reached if the sender transmon began in $|e\rangle$. The efficiency particular to the sender transmon excitation *alone*, excluding all other sources of loss, is thus $\eta_{\text{fully-excited}}^{(s)} = 0.61$. We can use this value to correct the worst-case assumption of release efficiency above,

$$\eta_{\text{excite}}^{(s)} = 1 - (1 - \eta_{\text{excite,worst}}^{(s)}) \times (1 - \eta_{\text{fully-excited}}^{(s)}) = 0.976. \quad (7.21)$$

Our understanding of the effect of transmon excitation on shape/frequency mismatch is improved by monitoring the emitted wavepacket as a function of excitation (Figure 7.24). If the transmon in the sender leaves its ground state during the protocol, this will change the resonance condition of the transfer. Depending on the parameters of each system, this will produce different effects. Near resonance, a waveform will be produced with a shape and/or frequency that are unmatched with respect to the receiving pulse. Far from resonance, the prepared state will be converted into a propagating one at a much slower rate than




Figure 7.23 | **Transmission inefficiency from sender transmon excitation.** We measure both cavities as a function of sender transmon excitation probability $P(e)$. **Left:** The measured remaining photon number in the sender cavity is inherently conditioned on the sender transmon being in its ground state. The energy remaining appears therefore unaffected. **Right:** The received number of photons decreases with sender transmon excitation, reflecting that the state has not been released with proper conditions for capture. The line is a linear fit to the data, from which the total efficiency for a completely excited receiver transmon is extracted (cross). Removing the excitation-independent portion then produces the excitation-dependent inefficiency $\approx 39\%$. (Figure adapted from [86]; see [Copyright Permissions](#).)

expected. These inefficiencies will manifest in the sender either as unreleased energy or unmatched (“uncatchable”) propagating energy. The particular manifestation of an excitation-induced error is dependent on system parameters. In our case, the values are χ_{at} and χ_{bt} in each system are similar, leading to Stark shifts that enable a conversion resonance at a different absolute detuning from the communication mode. With alternate system parameters, however, the failure may have resulted in a failure to release the energy at all.

Receiver transmon excitation

From the data in Figure 7.24 we infer that full excitation of the receiver transmon leads to reflection of $\sim 69\%$ of the incident energy. Combining this with the excitation of the transmon—increasing from 4.5% to 13.5% during the transfer—yields a worst-case efficiency associated with transmon excitation of the receiver, $\eta_{\text{excite,worst}}^{(r)} \approx 0.912$. We infer the value of $\eta_{\text{fully-excited}}^{(r)} \approx 0.3$ from the data in Figure 7.24, where we find that the integrated reflected energy in this case was 69% of the maximum. Parameters between the sender and receiver are very similar, so we expect this value to be near that of the sender. Combining these two values as for the sender, we obtain

$$\eta_{\text{excite}}^{(r)} = 1 - (1 - \eta_{\text{excite,worst}}^{(r)}) \times (1 - \eta_{\text{fully-excited}}^{(r)}) = 0.938. \quad (7.22)$$

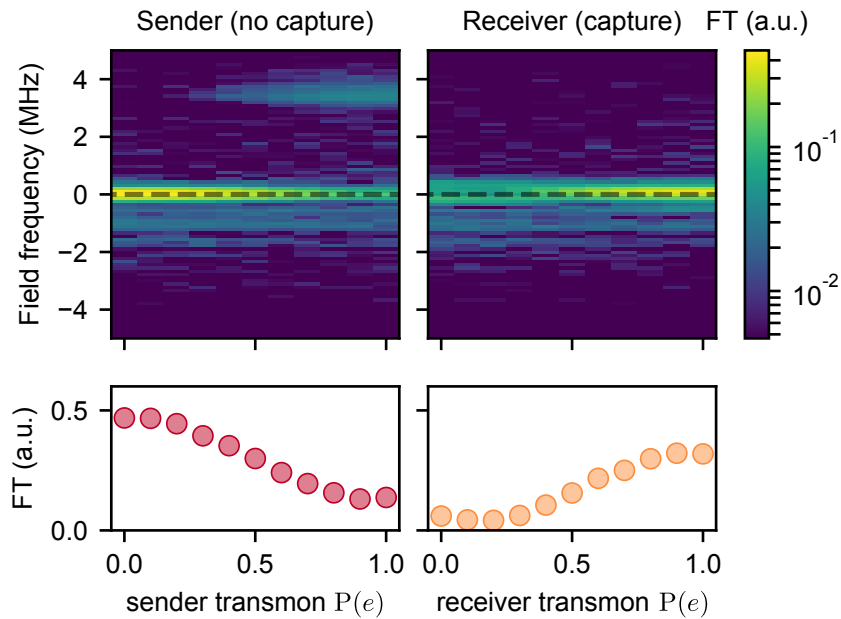


Figure 7.24 | **Effect of transmon excitation on the wavepacket.** To gain insight on how the propagating wavepacket is affected by transmon excitation, we measure the emitted field when the sender transmon is prepared in $|g\rangle$ or $|e\rangle$ and the cavity state $|1\rangle$ is subsequently released. The magnitude of the Fourier transform (FT) is shown. **Left:** The capture pulse is omitted to monitor the released wavepacket. A cut taken at the nominal propagating frequency (dashed line) indicates a decrease in emitted field at the transmission frequency. At large excitations, field appears to be emitted around 3.5 MHz above the nominal frequency. **Right:** The capture pulses are played, and the reflected field is observed. At the nominal propagating frequency, a line cut (dashed) shows that reflected field intensity increases with increased receiver excitation probability (lower). The reflected energy received at maximal receiver excitation is 69% of that received from a normal release with no capture attempt. (Figure adapted from [86]; see [Copyright Permissions](#).)

Taking into account pulse truncation (Section 7.5.2) gives an expected capture efficiency of 0.89. This is in reasonable agreement with the observed reflection of $n_{\text{reflect}}^{(r)} = 6.8\% \pm 0.5\%$ (obtained by integrating the photon flux of the incident wavepacket shown in Figure 7.8c). The remaining mismatch could be explained by a better value of truncation efficiency than expected due to experimental fine-tuning, $\eta_{\text{trunc,corr}}^{(r)} = 0.994$; this quantity cannot be directly measured in our experiment.

This comparison of two energies does not depend on sender or transmission loss values, since signals with or without the capture pulses are affected by those losses equally. Therefore we expect this value to agree closely with our model, which predicts

$$\eta_{\text{trunc}}^{(r)} \times \eta_{\text{excite}}^{(r)} = 0.89,$$

leaving a reflected energy fraction of 0.11.

Miscalibrations

Any unidentified source of inefficiency could be due to imperfectly calibrated pulses or uncertainty in the system parameters. Effective pulses require that accurate system parameters are supplied during the generation process. These parameters come from many calibration measurements, and the error from these calibration values can propagate into errors in the amplitude and frequency of a pulse. Assuming that missing factors in the capture or release error budgets arise from miscalibration error, we estimate an upper bound.

We can estimate the miscalibration efficiency associated with the sender or the receiver based on the measured fraction of energy remaining in the sender, $n_{\text{remain}}^{(s)}$, or reflected from the receiver, $n_{\text{reflect}}^{(r)}$, as well as the other efficiencies. In the sender, we calculate the efficiency due to miscalibration as

$$\eta_{\text{miscal}}^{(s)} = \left(1 - n_{\text{remain}}^{(s)}\right) / \eta_{\text{trunc}}^{(s)} = 0.98$$

since we observe a remaining population fraction $n_{\text{remain}}^{(s)} = 0.032 \pm 0.005$. Note that $\eta_{\text{excite}}^{(s)}$ does not appear as a contribution to the remaining population, since $n_{\text{remain}}^{(s)}$ reflects a conditioned measurement.

For the receiver, miscalibrations seem to yield a slightly better capture efficiency than expected (see

above); this leaves no gap in the error budget for miscalibration. We conclude that $\eta_{\text{miscal}}^{(r)}$ is immeasurably close to unity. Given the uncertainty in the reflection measurement, we place a bound, $\eta_{\text{miscal}}^{(r)} \geq 0.99$. Other possible sources of uncertainty include higher-order nonlinearities and frequency-dependent dissipation, which we have not considered in our model. (For example, parametric drives can excite the transmon beyond just $|e\rangle$, for example, and the transmon T_1 process can cause relaxation when steady-state excitation has been assumed.)

We can define the total release inefficiency from all components in our model to be

$$\eta_{\text{release}} = \eta_{\text{miscal}}^{(s)} \times \eta_{\text{trunc}}^{(s)} \times \eta_{\text{excite}}^{(s)} = 0.945 \quad (7.23)$$

and similarly for the receiver as

$$\eta_{\text{capture}} = \eta_{\text{miscal}}^{(r)} \times \eta_{\text{trunc,corr}}^{(r)} \times \eta_{\text{excite}}^{(r)} = 0.924. \quad (7.24)$$

7.5.3 Transmission loss

The energy dissipated in the transmission line is difficult to measure directly in our experiment. We can, however, assume that any ‘missing’ energy comes from loss; this gives a value for the transmission line efficiency of $\eta_{\text{tx}} = 0.85 \pm 0.04$. We perform the following control experiment to corroborate this value: we apply a constant tone detuned from the output resonator of the sender, and measure the Stark shift of both sending and receiving communication modes ($\hat{b}^{s,r}$) at variable amplitudes of the tone (Figure 7.25). From these Stark shifts and knowledge of the readout-transmon cross-Kerr χ_{bt} and output rate κ_{out} , one can calculate the loss between systems in a way that is independent of the sending and receiving portions of the transfer protocol.

One expects to receive

$$\bar{n}_r = \bar{n}_s \kappa_s \kappa_r \eta_{\text{tx,meas}} / ((\kappa_r/2)^2 + \delta_r^2) \quad (7.25)$$

for a transmission loss $\eta_{\text{tx,meas}}$ between the two communication resonators, and a detuning δ_r of mode \hat{b}^r relative to the drive. We perform these measurements at several detunings and extract a value of $\eta_{\text{tx,meas}} = 0.80 \pm 0.15$. This agrees well with the ‘missing energy’ estimation of η_{tx} above.




Figure 7.25 | **Estimation of transmission loss.** We prepare a steady-state population of photons in the sender and measure the resulting population in the receiver by means of the Stark shift. This measurement allows us to estimate the fraction of photons transmitted from sender to receiver. The line is a linear fit. (Figure adapted from [86]; see [Copyright Permissions](#).)

Reduction of physical loss

The physical losses (such as coaxial cable attenuation, circulator insertion loss, and connector reflections) that we associate with “transmission loss” may be mitigated in part by obtaining higher-quality circulators (probably in more narrow frequency bands), or by omitting the circulator completely. Omitting the circulator would require alternate means of measuring the system or characterizing state transfer. It would also prompt us to consider the resonant modes of the terminated transmission line, which the circulator currently allows us to ignore. Connectors and the presence of reflections could be carefully evaluated to determine whether they contribute dominantly, for example by adding cold switch arrays and reference lines. Such analysis could also be applied *in situ* by identifying optimal frequencies at which to perform the transfer.

7.5.4 Expanding the model beyond photon loss

While our model of photon loss agrees significantly with the measured process (Section 7.3.3), other error mechanisms may explain the remaining uncharacterized infidelity. We have investigated two such mechanisms: decoherence and thermalization.

For simplicity, we chose to apply our models for decoherence and thermalization to the single-photon encoding. To test decoherence, we construct a model for a completely positive and trace-preserving (CPTP) map that includes the action of both photon loss and decoherence. In this representation of a quantum

channel, we express decoherence using the Pauli operator σ_z and photon loss using $\sigma_+ = \sigma_x + i\sigma_y$. Each loss term is weighted by a variable probability of occurrence, and the weighted model is numerically compared to the measured data. We minimize the average infidelity between the measured states and states resulting from this model, over the probability space of these two mechanisms. This routine returns a minimum corresponding with a small amount of decoherence that changes the average fidelity (and the fidelity of any single state) by $< 1\%$. Therefore, this effect alone does not account for the $\sim 5\%$ unexplained by the non-unity process fidelity.

The potential effect of thermalization or “photon gain” (equilibrating, with some probability, to a thermal state at an unspecified point during the transfer process) is evaluated using a similar comparison. Much like the model of photon loss imposes a beamsplitter on the transmitted state with vacuum at the opposite input port, the thermalization process is modeled using a thermal state as the joining state. Including photon gain in the loss model would likely require a reduction in the assumed efficiency due to photon loss, since the same measured states must be ultimately matched. The model we construct allows for a uniform photon number of the thermal bath, \bar{n} , as well as some weight with which the process acts. Minimizing under this model to best match the measured states, we find an improvement in mean fidelity $< 1\%$. Again, the effect of photon gain alone cannot account for the infidelity of the transfer to a pure photon loss process.

7.6 The future of state transfer and stationary–propagating conversion

In stark contrast to contemporary state transfer implementations, the scheme presented in this chapter supports many-dimensional Hilbert spaces, thus providing a clear route towards error-correctable distribution of quantum information and entanglement. The on-demand generation of arbitrary, traveling multi-photon quantum states further provides exciting new opportunities for hybrid quantum systems. For instance, the efficient capture of traveling microwave fields by mechanical oscillators has been demonstrated experimentally [224]. An extension of this system could be used to transmit and receive highly non-classical mechanical states. Mechanical systems can act as transducers with radically different degrees of freedom, such as light in the optical domain [225, 226]. The combination of our system with such a

transducer would thus enable the distribution of exotic continuous variable quantum states in heterogeneous networks. The compatibility of multi-photon states with powerful bosonic error correction codes offers to protect the transmission and distribution of such states.

Using this device and conversion scheme, we have realized a high-fidelity, deterministic quantum state transfer protocol between remote microwave cavity memories using tools available in superconducting cavity QED. Importantly, our implementation is capable of transferring both single- and multi-photon quantum states. This allows use of a multi-photon qubit encoding that, while initially producing larger overhead and reducing mean fidelity, extends the techniques of quantum error correction to enable the mitigation of photon loss in transmission. This achievement taps into the body of work already developed for correcting errors in stationary memories [47, 53, 96, 227] to address the challenge of scalable quantum communication. The demonstration of remote entanglement generation at a rate exceeding the memory loss rates satisfies an essential requirement for scalable quantum communication and distributed computation [228]. Entanglement is a critical resource in quantum networks, and its rapid and on-demand generation will enable high-level operations between remote modules such as non-local gates [229] and entanglement distillation [230, 231]. Our experimental results thus demonstrate precise and tunable coupling of modular superconducting circuit systems that can be used to develop large-scale fault-tolerant quantum computing and communication networks.

8

Conclusion

In the chapters of this thesis, we have discussed and addressed limitations to scaling, both at the foundational level, with devices and materials, and at the much higher level, where we employed these devices to accomplish a critical task in scaling quantum information devices. Just a decade ago, state transfer with such high fidelity would probably not have been foreseen using superconducting circuits. Extensive control over couplings in circuit QED, however, gives us access to versatile parametric operations. We have been lucky to harness powerful tools that form the foundation on which this work is built. As hardware has become more coherent and reliable, algorithms have been developed to make complex operations straight-forward. If these algorithms demonstrated in this work are still relevant to devices a decade from now, they will likely have been enshrouded in many more layers of complexity.

8.1 Perspectives

8.1.1 Scaling

So what comes next? We have demonstrated communication between two modules; that primitive network must be expanded. A major requirement for error correction in this channel was satisfied by transmitting multi-photon states, but error detection or correction was not explicitly performed. As described in Chapter 7, relatively simple modifications to the hardware and conversion element should make this possible with high fidelity. These changes should improve uncorrected fidelity, too, allowing error correction to push fidelities even higher. Understanding the source of pump-induced heating, which may also be

linked to the conversion element, could increase our success probability the remaining distance to unity. These improvements would boost state transfer fidelity and entanglement generation rate; if the memory modes also became longer-lived, the scaling of the comparative times would enable many more operations in the lifetime of the system.

Now that proof-of-concept transfer between cQED modules has been established, the modules must be outfitted with all the ‘bells and whistles’ required to scale the network. We must incorporate the facilities for stationary error correction, which will probably share those required to perform correction of the captured state. Coupling memory modes within a module would satisfy the need for local logical operations. Using a scheme such as that referenced in Section 3.2.2 would allow remote gates to be performed. Distributed entanglement will also need to be error-corrected or purified, mandating the need for additional elements in each module [229, 230].

In the meantime, the basic hardware modules, and the transmon qubits that reside in them, should continue to improve. As coherence times increase, more limiting loss mechanisms will be identified. Whereas at shorter coherence times single sources might have dominated, it seems that multiple mechanisms may now play near-equal roles. This added challenge will require doing multiple things correctly, carefully considering materials, assembly, processing, and more. Before that, though, these remaining limitations will need to be identified. The process outlined in Chapter 5 will hopefully serve as a starting point from which to conduct these tests.

One major benefit of a modular architecture is that these challenges can continue to be addressed in parallel, even as modules are connected together into increasingly complex networks. Identical modules or classes of modules will all benefit. Improved modules could be substituted into existing networks whenever appropriate.

8.1.2 Integrating and adapting

As promising technologies like the multilayer microwave integrated quantum circuit (MMIQC) improve, we will likely see modules decrease in size and increase in robustness and reproducibility (as fabrication techniques play a greater role) [148]. While the form of modules will change, they will manifest the same basic elements and techniques required by a network for quantum computation: storage elements, con-

version elements, and elements for control and measurement.

But if bespoke 3D cavities are ever abandoned, it will not be for some time yet. Before that happens, new technologies that change our paradigm of measurement and control may already have been adapted. To reduce the number of RF drive and readout lines, input and outputs may be multiplexed using planar or 3D structures as in Chapter 4. If flux could be safely introduced without spoiling coherence times, then wide-bandwidth amplifiers could be incorporated on-chip [204, 232]. These adaptations demonstrate just a few possible ways to take advantage of the scalable nature of the cQED elements at our disposal.

The immense control that is possible over cQED systems, combined with the high quality of elements available, makes using these tools a quantum engineer's dream. It should be especially revealing that the experiments in this work made use of relatively simple elements in our arsenal. Far more complex experiments are easy to envision. Many subsequent and supporting steps are already being studied, and will form the basis of future theses. Circuit QED can be a flexible field, and it can and will adapt and apply this knowledge as necessary. The end goal is a fascinating one; but we can not underestimate the physics to be learned, and the technologies and applications developed, along the way.

A

Appendix

A.1 Analytical approximation of partially loaded coaxline waveguide

The lowest supported mode in the analytical geometry of Figure 4.3a, used to approximate the partially loaded waveguide of our coaxline design, will be the TE₁₀ mode. This mode has $E_Z = 0$, and thus we express the wave equations h_z , where $H_Z(x, y, z) = h_z(x, y, z)e^{-j\beta z}$, as follows:

$$\left(\frac{\partial^2}{\partial x^2} + k_L^2\right)h_z = 0, \quad 0 \leq x \leq a \quad (\text{A.1a})$$

$$\left(\frac{\partial^2}{\partial x^2} + k_D^2\right)h_z = 0, \quad a \leq x \leq a + t \quad (\text{A.1b})$$

$$\left(\frac{\partial^2}{\partial x^2} + k_R^2\right)h_z = 0, \quad a + t \leq x \leq 2a + t. \quad (\text{A.1c})$$

The propagation constant β must be phase-matched in each region, thus requiring

$$\beta = \sqrt{\epsilon_r k^2 - k_D^2} = \sqrt{k^2 - k_L^2} = \sqrt{k^2 - k_R^2}. \quad (\text{A.2})$$

Equations A.1 have solutions of the form $A \cos(k_0 x - x_0) + B \sin(k_0 x - x_0)$ to which we apply the boundary conditions $e_y(x = 0) = 0$ and $e_y(x = 2a + t) = 0$ (for the y-component electric field $E_y = e_y e^{-j\beta z}$, where $e_y = \frac{j\omega\mu}{k_i^2} \frac{\partial h_z}{\partial x}$) to obtain the following form for h_z :

$$h_z = A \cos(k_L x), \quad 0 \leq x \leq a \quad (\text{A.3a})$$

$$h_z = C \cos(k_D(x - a)) + D \sin(k_D(x - a)), \quad a \leq x \leq a + t \quad (\text{A.3b})$$

$$h_z = E \cos(k_R(2a + t - x)), \quad a + t \leq x \leq 2a + t. \quad (\text{A.3c})$$

Applying the condition of continuity of tangential fields, we match the two equations valid at each of the boundaries for both h_z and e_y . This means setting Equation A.1a equal to Equation A.1b, the same with Equations A.1b and A.1c, as well as both of the related expressions for e_y . This gives four expressions for the four variables (A, C, D, E) in Equation A.3, and by noting that $k_L = k_R$ the resulting relation

appears:

$$\tan k_{Ra} = -\frac{(k_R/k_D) \sin k_D t + \tan k_{Ra} \cos k_D t}{\cos k_D t - (k_D/k_R) \tan k_{Ra} \sin k_D t}, \quad (\text{A.4})$$

a transcendental equation that must be solved numerically. We express k_D and k_R in terms of β and the chosen wavenumber k . The cutoff for this mode occurs where $\beta = 0$, and so we can numerically solve for the wavenumber k and resulting cutoff frequency, $2\pi f_c = ck$.

A.2 Driven transmon Hamiltonian transformation

Section 6.1.1 makes use of the Hamiltonian of a driven transmon qubit, transformed in a displaced, rotating frame. Here we work out the details of this transformation that were omitted in that discussion.

Our displacement unitary $\hat{U} = \hat{U}_1 = \hat{D}(\xi(t))$ will transform our Hamiltonian

$$\hat{H} = \hbar\omega_t \hat{t}^\dagger \hat{t} - E_J \cos(\varphi_t(\hat{t}^\dagger + \hat{t})) + E_J \left(1 - \frac{\phi_t^2}{2}(\hat{t}^\dagger + \hat{t})^2\right) + \epsilon(t)(\hat{t}^\dagger + \hat{t}), \quad (\text{A.5})$$

as

$$\tilde{\hat{H}} = i\hbar\dot{\hat{U}}\hat{U}^\dagger + \hat{U}\hat{H}\hat{U}^\dagger. \quad (\text{A.6})$$

We treat each term in \hat{H} separately, beginning with $\hbar\omega_t \hat{t}^\dagger \hat{t}$. Here we use the Hadamard lemma, which states

$$e^{\hat{x}} \hat{y} e^{-\hat{x}} = \hat{y} + [\hat{x}, \hat{y}] + \frac{1}{2!} [\hat{x}, [\hat{x}, \hat{y}]] + \dots \quad (\text{A.7})$$

In our case, $\hat{x} = \xi \hat{t}^\dagger - \xi^* \hat{t}$. We will establish the values of several useful commutators that will make the calculation easier:

$$[\hat{t}, \hat{t}^\dagger] = 1 \quad (\text{A.8})$$

$$[\hat{x}, \hat{t}^\dagger \hat{t}] = -\xi \hat{t}^\dagger - \xi^* \hat{t} \quad (\text{A.9})$$

$$[\hat{x}, \hat{t}^\dagger] = -\xi^* \quad (\text{A.10})$$

$$[\hat{x}, \hat{t}] = -\xi \quad (\text{A.11})$$

The second term of the transformation in Equation A.6 applied to this Hamiltonian term then gives

$$\hbar\omega_t(\hat{t}^\dagger - \xi^*)(\hat{t} - \xi). \quad (\text{A.12})$$

The Hadamard lemma can similarly be applied to the third and fourth terms of Equation A.5. The third term gives

$$\hat{U}_1 E_J \left(1 - \frac{\phi_t^2}{2} (\hat{t}^\dagger + \hat{t})^2 \right) \hat{U}_1^\dagger = E_J - \frac{E_J}{2} \phi_t^2 \left(1 + ((\xi + \xi^*) - (\hat{t}^\dagger + \hat{t}))^2 \right). \quad (\text{A.13})$$

The fourth term gives

$$\hat{U}_1 \epsilon(t) (\hat{t}^\dagger + \hat{t}) \hat{U}_1^\dagger = \epsilon(t) (\hat{t}^\dagger + \hat{t} - \xi^\dagger - \xi). \quad (\text{A.14})$$

The second term (the cosine term) can be evaluated by multiple applications of the Baker–Campbell–Hausdorff (BCH) formula, which states that

$$e^{\hat{x}+\hat{y}} = e^{\hat{x}} e^{\hat{y}} e^{-[\hat{x},\hat{y}]/2} \quad (\text{A.15})$$

in particular cases like ours, where the next commutator $[\hat{x}, [\hat{x}, \hat{y}]]$ evaluates to zero. We will be interested in particular in the three-operator case,

$$e^{\hat{x}} e^{\hat{y}} e^{\hat{z}} \quad (\text{A.16})$$

which simplifies to

$$e^{[\hat{x},\hat{y}]/2} e^{[\hat{x},\hat{z}]/2} e^{[\hat{y},\hat{z}]/2} e^{\hat{x}+\hat{y}+\hat{z}} \quad (\text{A.17})$$

upon multiple applications of BCH. We are interested in the special case where $\hat{x} = -\hat{z}$, simplifying further to

$$e^{[\hat{x},\hat{y}]} e^{\hat{x}+\hat{y}+\hat{z}}. \quad (\text{A.18})$$

In the case of the cosine term, this gives

$$\hat{U}_1(-E_J \cos(\varphi_t(\hat{t}^\dagger + \hat{t})))\hat{U}_1^\dagger = \frac{-E_J}{2}\hat{U}_1 \left(e^{i\phi_t(\hat{t}^\dagger + \hat{t})} + e^{-i\phi_t(\hat{t}^\dagger + \hat{t})} \right) \hat{U}_1^\dagger \quad (\text{A.19})$$

$$= \frac{-E_J}{2} \left(e^{i\phi_t(\hat{t}^\dagger + \hat{t})} e^{-i\phi_t(\xi^* + \xi)} + \text{h.c.} \right) \quad (\text{A.20})$$

$$= \frac{-E_J}{2} \left(e^{i\phi_t(\hat{t}^\dagger + \hat{t} - \xi^* - \xi)} + e^{-i\phi_t(\hat{t}^\dagger + \hat{t} - \xi^* - \xi)} \right) \quad (\text{A.21})$$

$$= \frac{-E_J}{2} 2 \cos \phi_t(\hat{t}^\dagger + \hat{t} - \xi^* - \xi). \quad (\text{A.22})$$

With all four terms of the Hamiltonian processed, we evaluate the first term in Equation A.6:

$$i\hbar\dot{\hat{U}}_1\hat{U}_1^\dagger = i\hbar\left(\frac{d}{dt}e^{\hat{x}}\right)e^{-\hat{x}} = i\hbar\dot{\hat{x}}e^{\hat{x}}e^{-\hat{x}}. \quad (\text{A.23})$$

We again use Baker–Campbell–Hausdorff to show that $e^{\hat{x}}e^{-\hat{x}} = 1$. We also calculate

$$\dot{\hat{x}} = \dot{\xi}\hat{t}^\dagger - \dot{\xi}^*\hat{t} - \frac{1}{2}\frac{d}{dt}|\xi|^2. \quad (\text{A.24})$$

These five terms thus complete the transformation of Equation A.5 according to Equation A.6. We wish to express the drive $\xi(t)$ in a form that leaves only the cosine term. To do this, we will select (among those in \hat{H}_1) only the terms that operate between neighboring transmon levels, neglecting energy offsets and higher-order terms. This reduces \hat{H}_1 to

$$\hat{H}_1 = i\hbar(\dot{\xi}\hat{t}^\dagger - \dot{\xi}^*\hat{t}) + \hbar\omega_t(-\hat{t}\xi^* - \hat{t}^\dagger\xi) \quad (\text{A.25})$$

$$+ \epsilon(t)(\hat{t}^\dagger + \hat{t}) - E_J \cos\left(\phi_t(\hat{t}^\dagger + \hat{t} - \xi^* - \xi)\right) \quad (\text{A.26})$$

and then mandates that

$$i\hbar(\dot{\xi}\hat{t}^\dagger - \dot{\xi}^*\hat{t}) + \hbar\omega_t(-\hat{t}\xi^* - \hat{t}^\dagger\xi) = -\epsilon(t)(\hat{t}^\dagger + \hat{t}). \quad (\text{A.27})$$

By selecting the coefficients associated with \hat{t}^\dagger or \hat{t} , we arrive at the equation of motion in Equation 6.4.

Our displaced Hamiltonian has become

$$\hat{H}_1 = \hbar\omega_t \hat{t}^\dagger \hat{t} - E_J \cos\left(\varphi_t(\hat{t}^\dagger + \hat{t} - \xi(t) - \xi^*(t))\right), \quad (\text{A.28})$$

to which we wish to apply the rotating transformation $\hat{U}_2 = \exp(i\omega_t t \hat{t}^\dagger \hat{t})$. The first term is left untouched (which can be confirmed using the Hadamard lemma). The cosine term is not explicitly time-dependent, and thus we have is the very definition of the *interaction picture*. This means that, equivalently, we can transform the operators according to

$$\tilde{t}(t) = \hat{U}_2 \hat{t} \hat{U}_2^\dagger \quad (\text{A.29})$$

which in our case can be solved in a brute-force way by applying the Hadamard lemma (there are surely cleverer ways):

$$e^{i\omega_t t \hat{t}^\dagger \hat{t}} \hat{t} e^{-i\omega_t t \hat{t}^\dagger \hat{t}} = \hat{t} + [i\omega_t t \hat{t}^\dagger \hat{t}, \hat{t}] + \frac{1}{2} [i\omega_t t \hat{t}^\dagger \hat{t}, [i\omega_t t \hat{t}^\dagger \hat{t}, \hat{t}]] + \dots \quad (\text{A.30})$$

$$= \hat{t} + (i\omega_t t) \hat{t} + \frac{1}{2!} (i\omega_t t)^2 \hat{t} + \dots \quad (\text{A.31})$$

$$= \hat{t} e^{i\omega_t t} = \tilde{t}(t) \quad (\text{A.32})$$

since we can recognize this pattern as an exponential series. The same analysis for the raising operator gives $\tilde{t}^\dagger = \hat{t}^\dagger e^{-i\omega_t t}$. Substituting these into the cosine term in the Hamiltonian gives

$$- E_J \cos\left(\varphi_t(\tilde{t}^\dagger + \tilde{t} - \xi(t) - \xi^*(t))\right) \quad (\text{A.33})$$

whereas ξ , not being an operator, is unaffected.

The first (time-derivative) term in the Hamiltonian transformation becomes

$$i\hbar \frac{d}{dt} (i\omega_t \hat{t}^\dagger \hat{t}) = -\hbar\omega_t \hat{t}^\dagger \hat{t}. \quad (\text{A.34})$$

This conveniently cancels with the untouched first term of \hat{H}_1 , as a rotational transformation is intended

to do, thus resulting in only the cosine Hamiltonian term, given in Equation 6.9.

A.3 Detrimental effects of pumped processes

Pump-induced transmon heating or break-down (often measured as nonzero ‘background’ signal) has been observed at high pump photon numbers [86]. Depending on detuning from the resonator modes, and nearness to other less obvious frequency collisions, either heating of the qubit or readout bright-stating may occur with some probability. It is not yet clear which parameter of merit serves as a threshold for this obtrusive behavior, though we of course wish to maximize the pumped conversion rate g_{bs} with respect to this effect.

In this section, we will generally conform to the model given by refs. [233, 234], wherein heating from a driven Jaynes-Cummings system in the dispersive regime can provoke qubit excitation. This model may not be accurate in the strongly pumped regime of our work, but serves as a basis for our understanding until more advanced models are developed. In fact, recent work may be closer to describing the limitations we may face [235, 236]. We will present references to analyses of this phenomenon in related systems, and describe how we may try to reconfigure our experiments to avoid its consequences.

Bishop et al. [222] address the bright-stating behavior and introduce a critical photon number threshold, $N_{\text{crit}} = \Delta^2/4g_{\text{JC}}^2$, above which the perturbative expansion (in photon number) of the driven Jaynes-Cummings Hamiltonian fails to converge. (Here, g_{JC} is the coupling rate described in the undriven form of the Jaynes-Cummings Hamiltonian, as given in Equation 3.4.) This corresponds with a large intracavity intensity and bifurcation behavior and is where the dispersive approximation breaks down. (See also ref. [65].)

Mavrogordatos et al. [237] discuss this somewhat more quantitatively and show experimental agreement throughout the intermediate-power regime. The dependence of N_{crit} remains the same. The discussion of Sank et al. [223] speaks to the origin of these transitions as non-photon-number-conserving transitions that are typically excluded by the rotating wave approximation. This sets a framework by which such effects can be simulated (although it proves to be very computationally challenging given the many degrees of freedom in the system).

Another related explanation for this effect is that the transmon wavefunction is no longer confined to a single well of the cosine potential, and begins to tunnel between them. Simply, the device may no longer act like a transmon. In this case, a more linear conversion element that could handle more photons might provide an appealing solution.

In the following, we consider the beamsplitter coupling strength g_{bs} as defined in Equation 7.16 more explicitly and look at its scaling behavior with quantities (namely Δ and g_{JC}) that affect N_{crit} . Though N_{crit} may not be the figure of merit most relevant to our choice of conversion element here (the transmon), since it is derived from the Jaynes-Cummings Hamiltonian, we can use it as a guide to evaluate whether the expected scalings match observed behaviors. For example, we might wish to maximize g_{bs} and N_{crit} within the constraints of a chosen sample design.

A.3.1 Critical photon number scaling

The critical photon number can be expressed in terms of $\Delta = |\omega_x - \omega_t|$, the detuning of the mode in question from the transmon, and the dispersive coupling of mode x . In the strong dispersive limit, $g_{\text{JC}} = \sqrt{\chi_{\text{xt}}\Delta}$ and so (ignoring factors of two)

$$N_{\text{crit}} \propto \Delta/\chi_{\text{xt}} \quad (\text{A.35})$$

for a drive that is resonant with the cavity mode. In this limit, it is advantageous for modes to be more widely separated from the qubit. Other factors like the desire for large conversion strength will drive the desire to maintain χ .

A.3.2 Conversion strength

The beamsplitter Hamiltonian in Equation 6.13 can be generalized, in the presence of two pumps 1 and 2, and nearest modes x and y , respectively, to take the form

$$g_{\text{bs}} = E_J \phi_a \phi_b \phi_x \phi_y \xi_x^{(1)} \xi_y^{(2)*} \quad (\text{A.36})$$

which can be expressed in terms of transmon anharmonicity α and dispersive shifts χ_{it} (for mode i) as

$$g_{bs} = \frac{1}{2\alpha} \sqrt{\chi_{at}\chi_{bt}\chi_{xt}\chi_{yt}} \xi_x^{(1)} \xi_y^{(2)*}. \quad (\text{A.37})$$

Harmonic transmon

By making the transmon more harmonic, α can be decreased while having no effect on N_{crit} . Alternatively, we can hold g_{bs} fixed by reducing both α and each χ . By reducing χ , we increase N_{crit} . Reducing anharmonicity below about 60 or 70 MHz may be problematic due to the junction capacitance, although moving away from the transmon design entirely (and towards an amplifier-style, low-anharmonicity, junction-based device) may bypass this constraint.

Reordering modes

By moving the qubit to have the lowest frequency among the storage, qubit and readout modes, we could double the detuning to the readout mode and increase N_{crit} , potentially driving it more strongly. The trade-off in this case arises via reduction of Purcell filtering provided to the storage cavity. Careful engineering could balance the need for filtering.

Dump modes

By pumping on alternative modes x and y very far detuned from the junction, so-called photon “dump modes”, N_{crit} increases with Δ . We wish to at least maintain our conversion strength g_{bs} , and therefore maintain χ_{xt} and χ_{yt} near the level previously held by χ_{at} and χ_{bt} . (There is no reason χ_{at} and χ_{bt} should have to change.) For a fixed $\chi = g_{\text{JC}}^2/\Delta$, the coupling strength g_{JC} must be increased with a square-root dependence to hold χ fixed. Both terms scale with the same power in the expression for N_{crit} , such that N_{crit} will still increase even for χ held fixed.

Another advantage of using far-detuned modes to load photons is that they can be made to couple very strongly to the pump without Purcell-limiting high- Q elements. The Purcell limit,

$$\kappa_{\text{lim}} = \frac{g_{\text{JC}}^2}{\Delta^2} \kappa_x = \frac{\kappa_x}{N_{\text{crit}}} \quad (\text{A.38})$$

actually scales with the inverse relation as to the critical photon number.

So how will the mean photon number in mode x look for a CW drive labeled (1)? It will scale as

$$\bar{n}_x^{(1)} = \frac{\epsilon_x^2}{\kappa_x^2 + \Delta_x^2} \quad (\text{A.39})$$

where $\Delta_x = |\omega_1 - \omega_x|$. A dump mode could be almost-resonantly driven. For more photons to enter the mode given the same drive strength, κ would have to remain quite small—on the order of Δ , in fact. This seems non-intuitive. In this resonant limit, however, we have an understanding for \bar{n} should scale based on total and coupling Q and incident power, particularly

$$\bar{n}_x^{(\omega_1=\omega_x)} = 4P_{\text{in}}Q_{\text{tot}}^2/Q_c\hbar\omega_x^2. \quad (\text{A.40})$$

Measurement-induced dephasing

Measurement-induced dephasing [217, 238] describes the process of the environment learning information about the qubit state via quantum fluctuations of excitations in the readout (or some other) resonator. The more populated the resonator is, and the more distinguishable (number-split) those photons are, produces greater dephasing. In the case of dump modes, we are concerned with population in the dump modes or the readout resonator dephasing the qubits. Dephasing rate due to a drive near mode x with decay rate κ scales as

$$\Gamma_m = \frac{\bar{n}_x\kappa_x\chi_{xt}}{\kappa_x^2/4 + \chi_{xt}^2 + \Delta_r^2} \quad (\text{A.41})$$

where Δ_r is the detuning between the drive and mode x . The dump modes would see a change in regime such that $\Delta_r \approx \chi_{xt} \ll \kappa_x$, such that dephasing rate simplifies to

$$\Gamma_m \approx 4\bar{n}_x\chi_{xt}/\kappa_x. \quad (\text{A.42})$$

For fixed χ_{xt} , \bar{n} must scale with (or slower than) κ to remain safe. Any reduction in χ is beneficial. If we raise \bar{n} up as N_{crit} increases, then dephasing scales as

$$\Gamma_m \propto \Delta_{xt}/\kappa_x, \quad (\text{A.43})$$

a small number since Δ will increase by a small factor and κ likely by many orders of magnitude. Note that κ_x is the dominant rate throughout because of the scale ordering established above.

The readout resonator is unlikely to be affected by measurement dephasing for such a large-detuned pump because factors of Δ^2 appear in both the denominator of Γ_m as well as in the denominator of the expression for \bar{n}_r . Both χ_{bt} and κ_b are small compared to Δ_b , which will dominate. In fact, the relative rates scale as

$$\Gamma_m^{(b)}/\Gamma_m^{(x)} \propto \frac{\kappa_b^3 \kappa_x \chi_{bt}}{\Delta_b^4 \chi_{xt}} \quad (\text{A.44})$$

which is about 10^{-7} for tentative dump mode parameters. As expected, the large relative detuning dominates, even in the case of extremely small χ_{xq} , and is aided by the fact that the gain in pump photons compares prior qubit input coupling ($Q_{\text{in}} \sim 10^6$) with the new coupling Q_x , whereas the readout dephasing was always limited by the already-low value of $Q_{\text{out}} \sim 10^3$.

Estimates

In the devices of Chapters 6 (A) and 7 (B), a typical photon number product of 60 and above should be achievable. Device B has mode frequencies at 4.3 GHz, 5.5–6.1 GHz, and 7.5 GHz. This seems mostly limited, and should scale by, N_{crit} . A dump mode approach should increase the detuning from (at most) 2 GHz to 4 GHz (placing mode x at 9.5 GHz). To maintain χ , we must increase g_{JC} by a factor of $\sqrt{2}$, but N_{crit} still increases by a factor of 1.4. In this case, g_{bs} does not change, save for the 1.4 times photon number increase that should be possible.

If we are not so lucky, however, and can't maintain χ , let's see what happens. If χ is ten times smaller to the x and y modes, then g_{bs} is ten times smaller (total). This implies that g_{JC}^2 has been reduced by $2^2/10$. The overall increase in N_{crit} is larger than for fixed χ , by ten times, in fact. The only gain will be in the increased mean mode photon number than can be applied with the drives, which would actually be quite




Figure A.1 | **Transmon heating from applied pumps.** Different devices are measured, labeled A–F, for which the parameters are shown in Table A.1. The measurement of transmon excited state probability $P(e)$ after the application of both pumps on \hat{a} (storage) and \hat{b} (readout), with variable amplitude and for $10 \mu\text{s}$.

substantial.

So can we compensate for a reduction in $\sqrt{\chi_{xt}}$ with an increase in $\sqrt{\frac{c_x^2}{\kappa_x^2 + \Delta_x^2}}$? Unfortunately, no. While it is true that N_{crit} will increase proportionally to χ_{xt}^{-1} , the rate g is proportional to $\sqrt{\chi_{xt}\chi_{yt}}$ as well as $\xi_1\xi_2 \sim \sqrt{N_{\text{crit},x}}\sqrt{N_{\text{crit},y}}$. Therefore, the χ does not seem to matter in the case where one increases the drive strength proportionally with the increase in N_{crit} . A small χ can still be beneficial to tune N_{crit} to within a regime where driving is practical at normal excitation powers, however, and also to control Purcell and other linearity-sensitive effects.

A.3.3 Experimental results

We wanted to probe the nature of this pump-induced break-down, since its origin is still opaque. We measured devices with many different sets of parameters, including one experiment with “dump modes”. Though the behavior did not change overwhelmingly from one device to another, there were some differences; none matched definitively with any of the proposed behavior. The results are summarized in Figure A.1 and Table A.1.

Device	$\frac{\chi_{\text{at}}}{2\pi}$ (MHz)	$\frac{\chi_{\text{bt}}}{2\pi}$ (MHz)	$\frac{\kappa_{\text{out}}}{2\pi}$ (MHz)	Best $\bar{n}_a \bar{n}_b$	$\frac{g}{2\pi}$ (kHz)	$\frac{\kappa_{\text{eff}}}{2\pi}$ (kHz)
A	8.3	0.9	1.9	500	460	450
B	2.8	0.5	1.1	1500	150	80
C	2.3	2.5	0.7	290	270	420
D	2.5	1.9	1.0	150	160	100
E	1.2	0.6	3.0	3500	101	14
F	2.8	2.0	2.5	540	340	180

Table A.1 | **Pump strengths at varied device couplings.** Separate devices are cooled down, with the frequencies of transmon and readout elements, and the chip position relative to the storage mode, varied so as to alter χ_{at} and χ_{bt} . The labeling of these devices follows that of Figure A.1. For each of these devices, different pump detunings were used to enact the conversion Hamiltonian of Chapters 6–7. The detuning that and pump settings that produced the best product of \bar{n} , while keeping transmon excitation probably below 20%, were identified. From this, the maximum conversion rate g and the effective decay rate κ_{eff} of \hat{a} is calculated at this detuning.

A.4 Input-output theory for multi-system readout

The device in Chapter 7 can be read out in several ways. The naive way is sequential — the sender is measured, and then the receiver is measured. Each measurement result is thresholded to provide a single bit of information: the state in $|g\rangle$ or $|e\rangle$. A faster and more clever way, however, is to use two excitation pulses, applied to the readout resonator of each device, with variable amplitude and phase so as to discriminate between four states: $\{|gg\rangle, |ge\rangle, |eg\rangle, |ee\rangle\}$.

Here, we outline the input–output theory analysis that can be used to simulate the expected trajectories (and resulting “blobs”, or histograms of IQ values) in IQ space. Contrary to Equation 3.7, where steady-state transmission is analyzed, we will want to perform this analysis in the time domain. We will do so by integrating the quantum Langevin equation and applying the input–output relation.

A transmission cavity with two ports (neglecting internal loss, over which all coupling rates will dominate) will have two incoming and two outgoing signals: $a_{\text{in}}^{(\text{pulse})}$, $a_{\text{in}}^{(\text{sig})}$ and $a_{\text{out}}^{(\text{pulse})}$, $a_{\text{out}}^{(\text{sig})}$, respectively. This

is representative of our ‘sender’ system in Chapter 7. The intracavity field will obey

$$a(t) = e^{-(\kappa_{\text{tot}}/2+i\omega)t} \left(- \int_0^t ds e^{(\kappa_{\text{tot}}/2+i\omega)s} (\sqrt{\kappa_c} a_{\text{in}}^{(\text{pulse})} + \sqrt{\kappa} a_{\text{in}}^{(\text{sig})}) \right) \quad (\text{A.45})$$

$$= - \int_{-\infty}^t ds e^{(\kappa_{\text{tot}}/2+i\omega)(s-t)} (\sqrt{\kappa_c} a_{\text{in}}^{(\text{pulse})}(s) + \sqrt{\kappa} a_{\text{in}}^{(\text{sig})}(s)) \quad (\text{A.46})$$

where $\kappa_{\text{tot}} = \kappa_c + \kappa$, the sum of the coupling-in-port and signal-port decay rates.

From the input–output relation $a_{\text{out}}^{(\text{sig})} - a_{\text{in}}^{(\text{sig})} = \sqrt{\kappa} a(t)$,

$$a_{\text{out}}^{(\text{sig})} = a_{\text{in}}^{(\text{sig})} - \sqrt{\kappa} \int_{-\infty}^t ds e^{(\kappa_{\text{tot}}/2+i\omega)(s-t)} (\sqrt{\kappa_c} a_{\text{in}}^{(\text{pulse})}(s) + \sqrt{\kappa} a_{\text{in}}^{(\text{sig})}(s)). \quad (\text{A.47})$$

For the sender cavity, $a_{\text{in}}^{(\text{sig})} = 0$. If the receiver cavity is driven, then it will also obey this equation (with different parameters, possibly), but $a_{\text{in}}^{(\text{sig},\text{receiver})}$ will be the output of the first cavity, $a_{\text{out}}^{(\text{sig},\text{sender})}$.

The effect of a transmon in either cavity can be taken so as to change the cavity frequency ω . The trajectory of the readout resonator can be simulated under the set of four conditions possible in a two-cavity configuration: $|gg\rangle$, $|ge\rangle$, $|eg\rangle$, $|ee\rangle$, where the first and second letter refer the state of the sender and receiver transmon, respectively. By numerically integrating and cascading the output of the first cavity onto the input of the second in all four cases, we can produce these trajectories (Figure A.2). The relative drive amplitudes, phases, and absolute detuning can all be adjusted. Such simulations can be a useful tool in optimizing the separation of the response of readout resonators, particularly in the case where we intend to readout multiple systems simultaneously.

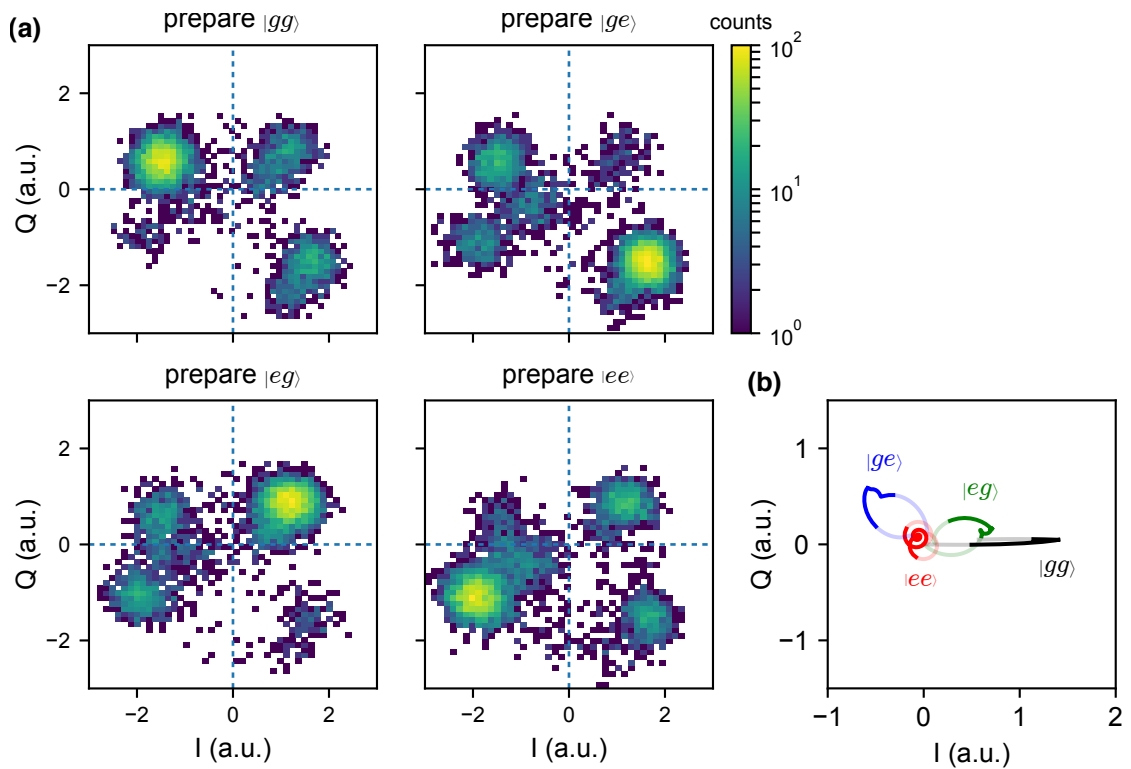


Figure A.2 | **Two-cavity readout trajectories.** **a** The sender–receiver system in Chapter 7 is prepared with all four combinations of each system’s transmon in $|g\rangle$ or $|e\rangle$. IQ-histograms are taken for many repeated measurements, for each preparation. **b** Simulation of trajectories in time of two driven readout resonators in a cascaded system. The four lines represent the same prepared states as in **a**, but with somewhat different drive and system parameters. The lightly shaded regions of each trajectory are the ring-up and ring-down of each resonator; this captures 30% of the total time.

Bibliography

- [1] S. Haroche and J.-M. Raimond, *Physics Today*, 51 (1996).
- [2] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, *Rev. Mod. Phys.* **82**, 1155 (2010), [arXiv:0810.4729](#) .
- [3] P. W. Shor, (2000), [arXiv:0005003 \[quant-ph\]](#) .
- [4] D. P. DiVincenzo and P. W. Shor, *Physical Review Letters* **77**, 3260 (1996).
- [5] A. M. Steane, *Physical Review A* **68**, 42322 (2003).
- [6] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, R. Barends, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, A. Fowler, B. Foxen, R. Graff, E. Jeffrey, J. Kelly, E. Lucero, A. Megrant, J. Mutus, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, H. Neven, and J. M. Martinis, (2017), [arXiv:1709.06678](#) .
- [7] C. Axline, M. Reagor, R. Heeres, P. Reinhold, C. Wang, K. Shain, W. Pfaff, Y. Chu, L. Frunzio, and R. J. Schoelkopf, *Applied Physics Letters* **109**, 042601 (2016).
- [8] L. Jiang, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, *Physical Review A* **76**, 062323 (2007).
- [9] N. H. Nickerson, J. F. Fitzsimons, and S. C. Benjamin, *Physical Review X* **4**, 041041 (2014).
- [10] C. Monroe, R. Raussendorf, A. Ruthven, K. R. Brown, P. Maunz, L. M. Duan, and J. Kim, *Physical Review A* **89**, 1 (2014), [arXiv:1208.0391](#) .
- [11] J. Z. Blumoff, *Multiqubit experiments in 3D circuit quantum electrodynamics*, Ph.D. thesis, Yale University (2017).
- [12] R. P. Feynman, *International Journal of Theoretical Physics* **21**, 467 (1982).
- [13] D. Deutsch, *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences* **400**, 97 (1985).
- [14] P. Shor, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (IEEE Comput. Soc. Press, 1994) pp. 124–134.
- [15] L. K. Grover, *Physical review letters* **79**, 4709 (1997).
- [16] I. Kassal, J. D. Whitfield, A. Perdomo-Ortiz, M.-H. Yung, and A. Aspuru-Guzik, *Annual Review of Physical Chemistry* **62**, 185 (2011).

- [17] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden, *Reviews of Modern Physics* **74**, 145 (2002).
- [18] D. P. DiVincenzo, *Fortschritte der Physik* **48**, 771 (2000).
- [19] R. W. Heeres, B. Vlastakis, E. Holland, S. Krastanov, V. V. Albert, L. Frunzio, L. Jiang, and R. J. Schoelkopf, *Physical Review Letters* **115**, 137002 (2015).
- [20] Jeremy Hsu, “CES 2018: Intel’s 49-Qubit Chip Shoots for Quantum Supremacy,” (2018).
- [21] M. J. Bremner, A. Montanaro, and D. J. Shepherd, *Quantum* **1**, 8 (2016), arXiv:1610.01808 .
- [22] Rachel Courtland, “Intel Now Packs 100 Million Transistors in Each Square Millimeter,” (2017).
- [23] W. Pfaff, C. J. Axline, L. D. Burkhardt, U. Vool, P. Reinhold, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, *Nature Physics* **13**, 882 (2017).
- [24] T. F. Rønnow, Z. Wang, J. Job, S. Boixo, S. V. Isakov, D. Wecker, J. M. Martinis, D. A. Lidar, and M. Troyer, *Science (New York, N.Y.)* **345**, 420 (2014).
- [25] Barry Briggs, “Things to Come: Could the cloud and quantum computing help feed the world?” (2016).
- [26] Amelia Heathman, “Microsoft quantum computer: the things quantum computing could solve,” (2017).
- [27] S. Jordan, “Quantum Algorithm Zoo,” (2018).
- [28] B. P. Lanyon, J. D. Whitfield, G. G. Gillett, M. E. Goggin, M. P. Almeida, I. Kassal, J. D. Biamonte, M. Mohseni, B. J. Powell, M. Barbieri, A. Aspuru-Guzik, and A. G. White, *Nature Chemistry* **2**, 106 (2010).
- [29] D. Lu, B. Xu, N. Xu, Z. Li, H. Chen, X. Peng, R. Xu, and J. Du, *Physical Chemistry Chemical Physics* **14**, 9411 (2012).
- [30] A. Perdomo-Ortiz, N. Dickson, M. Drew-Brook, G. Rose, and A. Aspuru-Guzik, *Scientific Reports* **2**, 571 (2012).
- [31] I. Georgescu, S. Ashhab, and F. Nori, *Reviews of Modern Physics* **86**, 153 (2014).
- [32] D. S. Abrams and S. Lloyd, *Physical Review Letters* **83**, 5162 (1999).
- [33] N. D. Mermin, *Quantum Computer Science An Introduction* (Cambridge University Press, 2007).
- [34] D. Gottesman, A. Kitaev, and J. Preskill, *Phys. Rev. A* **64**, 12310 (2001), arXiv:0510107 [quant-ph] .
- [35] B. Vlastakis, G. Kirchmair, Z. Leghtas, S. E. Nigg, L. Frunzio, S. M. Girvin, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, *Science* **342**, 607 (2013).
- [36] Z. Leghtas, G. Kirchmair, B. Vlastakis, R. J. Schoelkopf, M. H. Devoret, and M. Mirrahimi, *Physical Review Letters* **111**, 120501 (2013).

- [37] C. Monroe, D. M. Meekhof, B. E. King, and D. J. Wineland, *Science* **272**, 1131 (1996).
- [38] A. Signoles, A. Facon, D. Grosso, I. Dotsenko, S. Haroche, J.-M. Raimond, M. Brune, and S. Gleyzes, *Nat. Phys.* **10**, 715 (2014).
- [39] R. McConnell, H. Zhang, J. Hu, S. Ćuk, and V. Vuletić, *Nature* **519**, 439 (2015), [arXiv:1508.03056](#).
- [40] A. Ourjoumtsev, R. Tualle-Brouri, J. Laurat, and P. Grangier, *Science* **312**, 83 (2006).
- [41] J. S. Neergaard-Nielsen, B. M. Nielsen, C. Hettich, K. Mølmer, and E. S. Polzik, *Phys. Rev. Lett.* **97**, 83604 (2006), [arXiv:0602198 \[quant-ph\]](#).
- [42] S. Deléglise, I. Dotsenko, C. Sayrin, J. Bernu, M. Brune, J.-M. Raimond, and S. Haroche, *Nature* **455**, 510 (2008).
- [43] M. Hofheinz, H. Wang, M. Ansmann, R. C. Bialczak, E. Lucero, M. Neeley, A. D. O’Connell, D. Sank, J. Wenner, J. M. Martinis, and A. N. Cleland, *Nature* **459**, 546 (2009).
- [44] G. Kirchmair, B. Vlastakis, Z. Leghtas, S. E. Nigg, H. Paik, E. Ginossar, M. Mirrahimi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, *Nature* **495**, 205 (2013).
- [45] L. Bretheau, P. Campagne-Ibarcq, E. Flurin, F. Mallet, and B. Huard, *Science* **348**, 776 (2015).
- [46] M. Reagor, W. Pfaff, C. Axline, R. W. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, K. Chou, M. J. Hatridge, L. Frunzio, M. H. Devoret, L. Jiang, and R. J. Schoelkopf, *Physical Review B* **94**, 014506 (2016).
- [47] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, *Nature* **536**, 441 (2016).
- [48] S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, *Physica Scripta* **T137**, 014012 (2009).
- [49] J. Blumoff, K. Chou, C. Shen, M. Reagor, C. Axline, R. Brierley, M. Silveri, C. Wang, B. Vlastakis, S. Nigg, L. Frunzio, M. Devoret, L. Jiang, S. Girvin, and R. Schoelkopf, *Physical Review X* **6**, 031041 (2016).
- [50] W. K. Wootters and W. H. Zurek, *Nature* **299**, 802 (1982).
- [51] P. W. Shor, *Phys. Rev. A* **52**, R2493 (1995).
- [52] M. H. Michael, M. Silveri, R. Brierley, V. V. Albert, J. Salmilehto, L. Jiang, and S. Girvin, *Physical Review X* **6**, 031006 (2016).
- [53] M. Mirrahimi, Z. Leghtas, V. V. Albert, S. Touzard, R. J. Schoelkopf, L. Jiang, and M. H. Devoret, *New Journal of Physics* **16**, 045014 (2014).
- [54] V. V. Albert, K. Noh, K. Duivenvoorden, R. T. Brierley, P. Reinhold, C. Vuillot, L. Li, C. Shen, S. M. Girvin, B. M. Terhal, and L. Jiang, (2017), [arXiv:1708.05010](#).
- [55] A. Fowler, in *Bulletin of the American Physical Society* (American Physical Society, 2018).

- [56] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, *Journal of Mathematical Physics* **43**, 4452 (2002).
- [57] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, *Physical Review A* **86**, 032324 (2012).
- [58] A. Kitaev, *Annals of Physics* **303**, 2 (2003).
- [59] B. Lekitsch, S. Weidt, A. G. Fowler, K. Mølmer, S. J. Devitt, C. Wunderlich, and W. K. Hensinger, *Science Advances* **3**, e1601540 (2017).
- [60] R. Walter Ogburn and J. Preskill, in *Quantum Computing and Quantum Communications*, edited by C. P. Williams (Springer Berlin Heidelberg, Berlin, Heidelberg, 1999) pp. 341–356.
- [61] A. Y. Kitaev, *Physics-Uspekhi* **44**, 131 (2001).
- [62] H. Zhang, C.-X. Liu, S. Gazibegovic, D. Xu, J. A. Logan, G. Wang, N. van Loo, J. D. S. Bommer, M. W. A. de Moor, D. Car, R. L. M. Op het Veld, P. J. van Veldhoven, S. Koelling, M. A. Verheijen, M. Pendharkar, D. J. Pennachio, B. Shojaei, J. S. Lee, C. J. Palmstrøm, E. P. A. M. Bakkers, S. D. Sarma, and L. P. Kouwenhoven, *Nature* (2018), 10.1038/nature26142.
- [63] N. H. Nickerson, Y. Li, and S. C. Benjamin, *Nature Communications* **4**, 1756 (2013).
- [64] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, *Nature* **431**, 162 (2004).
- [65] D. I. Schuster, *Circuit Quantum Electrodynamics*, Ph.D. thesis, Yale University (2007).
- [66] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, *Physical Review Letters* **107**, 240501 (2011).
- [67] M. Reagor, *Superconducting Cavities for Circuit Quantum Electrodynamics*, Ph.D. thesis, Yale University (2015).
- [68] M. Tinkham, *Introduction to Superconductivity: Second Edition*, 2nd ed. (Dover Publications, 2004).
- [69] A. A. Houck, J. Koch, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Quantum Inf Process* **8**, 105 (2009).
- [70] Y. Nakamura, H. Terai, K. Inomata, T. Yamamoto, W. Qiu, and Z. Wang, *Applied Physics Letters* **99**, 212502 (2011).
- [71] J. B. Chang, M. R. Vissers, A. D. Córcoles, M. Sandberg, J. Gao, D. W. Abraham, J. M. Chow, J. M. Gambetta, M. B. Rothwell, G. A. Keefe, M. Steffen, D. P. Pappas, M. Beth Rothwell, G. A. Keefe, M. Steffen, and D. P. Pappas, *Applied Physics Letters* **103**, 12602 (2013).
- [72] E. Diamanti, H.-K. Lo, B. Qi, and Z. Yuan, *npj Quantum Information* **2**, 16025 (2016).
- [73] A. K. Ekert, *Physical Review Letters* **67**, 661 (1991).

- [74] M. Peev, C. Pacher, R. Alléaume, C. Barreiro, J. Bouda, W. Boxleitner, T. Debuisschert, E. Diamanti, M. Dianati, J. F. Dynes, S. Fasel, S. Fossier, M. Fürst, J.-D. Gautier, O. Gay, N. Gisin, P. Grangier, A. Happe, Y. Hasani, M. Hentschel, H. Hübel, G. Humer, T. Länger, M. Legré, R. Lieger, J. Lodewyck, T. Lorünser, N. Lütkenhaus, A. Marhold, T. Matyus, O. Maurhart, L. Monat, S. Nauerth, J.-B. Page, A. Poppe, E. Querasser, G. Ribordy, S. Robyr, L. Salvail, A. W. Sharpe, A. J. Shields, D. Stucki, M. Suda, C. Tamas, T. Themel, R. T. Thew, Y. Thoma, A. Treiber, P. Trinkler, R. Tualle-Brouri, F. Vannel, N. Walenta, H. Weier, H. Weinfurter, I. Wimberger, Z. L. Yuan, H. Zbinden, and A. Zeilinger, *New Journal of Physics* **11**, 075001 (2009).
- [75] R. Ursin, F. Tiefenbacher, T. Schmitt-Manderbach, H. Weier, T. Scheidl, M. Lindenthal, B. Blauensteiner, T. Jennewein, J. Perdigues, P. Trojek, B. Ömer, M. Fürst, M. Meyenburg, J. Rarity, Z. Sodnik, C. Barbieri, H. Weinfurter, and A. Zeilinger, *Nature Physics* **3**, 481 (2007).
- [76] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2010).
- [77] L. Jiang, J. M. Taylor, K. Nemoto, W. J. Munro, R. Van Meter, and M. D. Lukin, *Physical Review A* **79**, 032325 (2009).
- [78] H. J. Kimble, *Nature* **453**, 1023 (2008), [arXiv:0806.4195](https://arxiv.org/abs/0806.4195).
- [79] W. Dür, R. Lamprecht, and S. Heusler, *European Journal of Physics* **38**, 043001 (2017).
- [80] S. Ritter, C. Nölleke, C. Hahn, A. Reiserer, A. Neuzner, M. Uphoff, M. Mücke, E. Figueroa, J. Bochmann, and G. Rempe, *Nature* **484**, 195 (2012), [arXiv:1202.5955](https://arxiv.org/abs/1202.5955).
- [81] H. Bernien, B. Hensen, W. Pfaff, G. Koolstra, M. S. Blok, L. Robledo, T. H. Taminiou, M. Markham, D. J. Twitchen, L. Childress, and R. Hanson, *Nature* **497**, 86 (2013).
- [82] P. C. Humphreys, N. Kalb, J. P. J. Morits, R. N. Schouten, R. F. L. Vermeulen, D. J. Twitchen, M. Markham, and R. Hanson, (2017), [arXiv:1712.07567](https://arxiv.org/abs/1712.07567).
- [83] N. Roch, M. E. Schwartz, F. Motzoi, C. Macklin, R. Vijay, A. W. Eddins, A. N. Korotkov, K. B. Whaley, M. Sarovar, and I. Siddiqi, *Physical Review Letters* **112**, 170501 (2014).
- [84] A. Narla, S. Shankar, M. Hatridge, Z. Leghtas, K. Sliwa, E. Zalys-Geller, S. Mundhada, W. Pfaff, L. Frunzio, R. Schoelkopf, and M. Devoret, *Physical Review X* **6**, 031036 (2016).
- [85] C. Dickel, J. J. Wesdorp, N. K. Langford, S. Peiter, R. Sagastizabal, A. Bruno, B. Criger, F. Motzoi, and L. DiCarlo, *Physical Review B* **97**, 064508 (2018).
- [86] C. J. Axline, L. D. Burkhardt, W. Pfaff, M. Zhang, K. Chou, P. Campagne-Ibarcq, P. Reinhold, L. Frunzio, S. M. Girvin, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, *Nature Physics* (2018), [10.1038/s41567-018-0115-y](https://arxiv.org/abs/10.1038/s41567-018-0115-y), [arXiv:1712.05832](https://arxiv.org/abs/1712.05832).
- [87] P. Campagne-Ibarcq, E. Zalys-Geller, A. Narla, S. Shankar, P. Reinhold, L. D. Burkhardt, C. J. Axline, W. Pfaff, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret, (2017), [arXiv:1712.05854](https://arxiv.org/abs/1712.05854).
- [88] P. Kurpiers, P. Magnard, T. Walter, B. Royer, M. Pechal, J. Heinsoo, Y. Salathé, A. Akin, S. Storz, J.-C. Besse, S. Gasparinetti, A. Blais, and A. Wallraff, (2017), [arXiv:1712.08593](https://arxiv.org/abs/1712.08593).

- [89] K. S. Chou, J. Z. Blumoff, C. S. Wang, P. C. Reinhold, C. J. Axline, Y. Y. Gao, L. Frunzio, M. H. Devoret, L. Jiang, and R. J. Schoelkopf, (2018), [arXiv:1801.05283](#) .
- [90] B. Vogell, B. Vermersch, T. E. Northup, B. P. Lanyon, and C. A. Muschik, [Quantum Science and Technology](#) **2**, 045003 (2017).
- [91] T. Wilk, S. C. Webster, A. Kuhn, and G. Rempe, [Science](#) **317**, 488 (2007).
- [92] J. Wenner, Y. Yin, Y. Chen, R. Barends, B. Chiaro, E. Jeffrey, J. Kelly, A. Megrant, J. Mutus, C. Neill, P. O'Malley, P. Roushan, D. Sank, A. Vainsencher, T. White, A. N. Korotkov, A. Cleland, and J. M. Martinis, [Physical Review Letters](#) **112**, 210501 (2014).
- [93] M. Reagor, H. Paik, G. Catelani, L. Sun, C. Axline, E. Holland, I. M. Pop, N. A. Masluk, T. Brecht, L. Frunzio, M. H. Devoret, L. Glazman, and R. J. Schoelkopf, [Applied Physics Letters](#) **102**, 192604 (2013).
- [94] S. M. Girvin, in *Quantum Machines: Measurement Control of Engineered Quantum Systems*, edited by M. H. Devoret, B. Huard, R. J. Schoelkopf, and L. F. Cugliandolo (Oxford University Press, 2011).
- [95] K. S. Chou, *Teleported operations between logical qubits in circuit quantum electrodynamics*, Ph.D. thesis, Yale University (2018).
- [96] M. Reed, *Entanglement and Quantum Error Correction with Superconducting Qubits*, Ph.d. thesis, Yale University (2013).
- [97] Y. Chu, C. Axline, C. Wang, T. Brecht, Y. Y. Gao, L. Frunzio, and R. J. Schoelkopf, [Applied Physics Letters](#) **109**, 112601 (2016).
- [98] M. Hutchings, J. Hertzberg, Y. Liu, N. Bronn, G. Keefe, M. Brink, J. M. Chow, and B. Plourde, [Physical Review Applied](#) **8**, 044003 (2017).
- [99] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai, [Nature](#) **398**, 786 (1999).
- [100] V. Bouchiat, D. Vion, P. Joyez, D. Esteve, and M. H. Devoret, [Physica Scripta](#) **T76**, 165 (1998).
- [101] L. Bishop, *Circuit Quantum Electrodynamics*, Ph.D. thesis, Yale University (2010).
- [102] S. E. Nigg, H. Paik, B. Vlastakis, G. Kirchmair, S. Shankar, L. Frunzio, M. H. Devoret, R. J. Schoelkopf, and S. M. Girvin, [Physical Review Letters](#) **108**, 240502 (2012).
- [103] F. Lecocq, I. M. Pop, Z. Peng, I. Matei, T. Crozes, T. Fournier, Cécile Naud, W. Guichard, and O. Buisson, [Nanotechnology](#) **22**, 315302 (2011).
- [104] G. J. Dolan, [Applied Physics Letters](#) **31**, 337 (1977).
- [105] S. Haroche and J.-M. M. Raimond, *Exploring the Quantum* (Oxford University Press, Oxford, 2006).
- [106] V. Ambegaokar and A. Baratoff, [Physical Review Letters](#) **10**, 486 (1963).
- [107] D. H. Douglass and R. Meservey, [Physical Review](#) **135**, A19 (1964).

- [108] K. Geerlings, S. Shankar, E. Edwards, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret, [Applied Physics Letters](#) **100**, 192601 (2012).
- [109] A. Megrant, C. Neill, R. Barends, B. Chiaro, Y. Chen, L. Feigl, J. Kelly, E. Lucero, M. Mariantoni, P. J. J. O'Malley, D. Sank, A. Vainsencher, J. Wenner, T. C. White, Y. Yin, J. Zhao, C. J. Palmstrøm, J. M. Martinis, and A. N. Cleland, [Applied Physics Letters](#) **100**, 113510 (2012).
- [110] A. A. Houck, D. I. Schuster, J. M. Gambetta, J. A. Schreier, B. R. Johnson, J. M. Chow, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, [Nature](#) **449**, 328 (2007).
- [111] M. Hatridge, S. Shankar, M. Mirrahimi, F. Schackert, K. Geerlings, T. Brecht, K. M. Sliwa, B. Abdo, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, [Science](#) **339**, 178 (2013).
- [112] M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, [Physical Review Letters](#) **105**, 173601 (2010).
- [113] D. A. Steck, *Quantum and Atom Optics* (available online at <http://steck.us/teaching>; revision 0.10.2, 16 October 2015).
- [114] D. M. Pozar, *Microwave engineering* (Wiley, 2012).
- [115] M. S. Khalil, M. J. A. Stoutimore, F. C. Wellstood, and K. D. Osborn, [Journal of Applied Physics](#) **111**, 54510 (2012).
- [116] J. Gao, *The Physics of Superconducting Microwave Resonators*, Ph.d. thesis, California Institute of Technology (2008).
- [117] H. Padamsee, [Supercond. Sci. Technol.](#) **14** (2001).
- [118] D. P. Pappas, M. R. Vissers, D. S. Wisbey, J. S. Kline, and J. Gao, [IEEE Transactions on Applied Superconductivity](#) **21**, 871 (2011).
- [119] J. Gao, M. Daal, A. Vayonakis, S. Kumar, J. Zmuidzinas, B. Sadoulet, B. A. Mazin, P. K. Day, and H. G. Leduc, [Applied Physics Letters](#) **92**, 152505 (2008).
- [120] A. Kamal, J. Clarke, and M. H. Devoret, [Nature Physics](#) **7**, 311 (2011).
- [121] B. J. Chapman, E. I. Rosenthal, J. Kerckhoff, B. A. Moores, L. R. Vale, J. Mates, G. C. Hilton, K. Lalumière, A. Blais, and K. Lehnert, [Physical Review X](#) **7**, 041043 (2017).
- [122] N. E. Frattini, U. Vool, S. Shankar, A. Narla, K. M. Sliwa, and M. H. Devoret, [Applied Physics Letters](#) **110**, 222603 (2017).
- [123] D. Ristè, C. C. Bultink, M. J. Tiggelman, R. N. Schouten, K. W. Lehnert, and L. DiCarlo, [Nature Communications](#) **4**, 1913 (2013).
- [124] I. M. Pop, K. Geerlings, G. Catelani, R. J. Schoelkopf, L. I. Glazman, and M. H. Devoret, [Nature](#) **508**, 369 (2014).
- [125] Y. Liu, S. Shankar, N. Ofek, M. Hatridge, A. Narla, K. Sliwa, L. Frunzio, R. Schoelkopf, and M. Devoret, [Physical Review X](#) **6**, 011022 (2016).

- [126] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, *Nature Communications* **8**, 94 (2017), [arXiv:1608.02430](#) .
- [127] S. Lloyd and S. L. Braunstein, *Physical Review Letters* **82**, 1784 (1999).
- [128] J. M. Chow, L. DiCarlo, J. M. Gambetta, F. Motzoi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, *Physical Review A* **82**, 040305 (2010).
- [129] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Physical Review A* **76**, 042319 (2007).
- [130] D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, B. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Nature* **445**, 515 (2006), [arXiv:0608693 \[cond-mat\]](#) .
- [131] J. Řeháček, B.-G. Englert, and D. Kaszlikowski, *Physical Review A* **70**, 052321 (2004).
- [132] C. Shen, K. Noh, V. V. Albert, S. Krastanov, M. H. Devoret, R. J. Schoelkopf, S. M. Girvin, and L. Jiang, *Physical Review B* **95**, 134501 (2017).
- [133] L. Sun, A. Petrenko, Z. Leghtas, B. Vlastakis, G. Kirchmair, K. M. Sliwa, A. Narla, M. Hatridge, S. Shankar, J. Blumoff, L. Frunzio, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, *Nature* **511**, 444 (2014).
- [134] A. Royer, *Physical Review A* **15**, 449 (1977).
- [135] K. Banaszek, C. Radzewicz, K. Wódkiewicz, and J. S. Krasinski, *Physical Review A* **60**, 674 (1999).
- [136] U. Leonhardt and H. Paul, *Progress in quantum electronics* **19**, 89 (1995).
- [137] U. Leonhardt, *Measuring the quantum state of light* (Cambridge University Press, Cambridge, 1997).
- [138] P. Bertet, A. Auffèves, P. Maioli, S. Osnaghi, T. Meunier, M. Brune, J. M. Raimond, and S. Haroche, *Phys. Rev. Lett.* **89**, 200402 (2002).
- [139] Z. Leghtas, G. Kirchmair, B. Vlastakis, M. H. Devoret, R. J. Schoelkopf, and M. Mirrahimi, *Physical Review A* **87**, 042315 (2013).
- [140] T. Brecht, M. Reagor, Y. Chu, W. Pfaff, C. Wang, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf, *Applied Physics Letters* **107**, 192603 (2015).
- [141] T. Brecht, *Micromachined Quantum Circuits*, Ph.D. thesis, Yale University (2017).
- [142] M. Sandberg, M. R. Vissers, T. A. Ohki, J. Gao, J. Aumentado, M. Weides, and D. P. Pappas, *Applied Physics Letters* **102**, 72601 (2013).
- [143] M. Devoret, S. Girvin, and R. Schoelkopf, *Annalen der Physik* **16**, 767 (2007).
- [144] M. Göppl, A. Fragner, M. Baur, R. Bianchetti, S. Filipp, J. M. Fink, P. J. Leek, G. Puebla, L. Steffen, and A. Wallraff, *Journal of Applied Physics* **104**, 113904 (2008).

- [145] M. D. Reed, B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf, *Applied Physics Letters* **96**, 203110 (2010).
- [146] J. M. Gambetta, A. A. Houck, and A. Blais, *Physical Review Letters* **106**, 030502 (2011).
- [147] E. A. Sete, J. M. Martinis, and A. N. Korotkov, *Physical Review A* **92**, 012325 (2015).
- [148] T. Brecht, W. Pfaff, C. Wang, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf, [arXiv:1509.01127](https://arxiv.org/abs/1509.01127) [cond-mat, physics:quant-ph] **2**, 16002 (2015), [arXiv:1509.01127](https://arxiv.org/abs/1509.01127).
- [149] T. Brecht, Y. Chu, C. Axline, W. Pfaff, J. Blumoff, K. Chou, L. Krayzman, L. Frunzio, and R. Schoelkopf, *Physical Review Applied* **7**, 044018 (2017).
- [150] J. P. Turneaure, J. Halbritter, and H. A. Schwettman, *Journal of Superconductivity* **4**, 341 (1991).
- [151] J. P. Turneaure and I. Weissman, *Journal of Applied Physics* **39**, 4417 (1968).
- [152] J. M. Martinis and A. Megrant, [arXiv:1410.5793](https://arxiv.org/abs/1410.5793) [cond-mat, physics:quant-ph] (2014).
- [153] D. L. Creedon, Y. Reshitnyk, W. Farr, J. M. Martinis, T. L. Duty, and M. E. Tobar, *Applied Physics Letters* **98**, 222903 (2011).
- [154] O. Dial, D. T. McClure, S. Poletto, G. A. Keefe, M. B. Rothwell, J. M. Gambetta, D. W. Abraham, J. M. Chow, and M. Steffen, *Superconductor Science and Technology* **29**, 044001 (2016), [arXiv:1509.03859](https://arxiv.org/abs/1509.03859).
- [155] C. Deng, M. Otto, and A. Lupascu, *Applied Physics Letters* **104**, 043506 (2014).
- [156] 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, *Physical Review Letters* **95**, 210503 (2005).
- [157] L. Grünhaupt, U. von Lüpke, D. Gusenkova, S. T. Skacel, N. Maleeva, S. Schlör, A. Bilmes, H. Rotzinger, A. V. Ustinov, M. Weides, and I. M. Pop, *Applied Physics Letters* **111**, 072601 (2017).
- [158] A. D. O'Connell, M. Ansmann, R. C. Bialczak, M. Hofheinz, N. Katz, E. Lucero, C. McKenney, M. Neeley, H. Wang, E. M. Weig, A. N. Cleland, and J. M. Martinis, *Applied Physics Letters* **92**, 112903 (2008).
- [159] J. Krupka, J. Breeze, A. Centeno, N. Alford, T. Claussen, and L. Jensen, *IEEE Transactions on Microwave Theory and Techniques* **54**, 3995 (2006).
- [160] R. Simons, *Coplanar waveguide circuits, components, and systems* (2001).
- [161] W. Hartwig, *Proceedings of the IEEE* **61**, 58 (1973).
- [162] N. Samkharadze, A. Bruno, P. Scarlino, G. Zheng, D. P. DiVincenzo, L. DiCarlo, and L. M. K. Vandersypen, *Physical Review Applied* **5**, 044004 (2016).
- [163] C. Wang, C. Axline, Y. Y. Gao, T. Brecht, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf, *Applied Physics Letters* **107**, 162601 (2015).

- [164] 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, (2016), [arXiv:1606.09262](#) .
- [165] Q. Shu, K. Gendreau, W. Hartung, J. Kirchgessner, D. Moffat, R. Noer, H. Padamsee, D. Rubin, and J. Sears, [IEEE Transactions on Magnetics](#) **25**, 1868 (1989).
- [166] R. Geng and H. Padamsee, in *Proceedings of the 1999 Particle Accelerator Conference (Cat. No.99CH36366)*, Vol. 2 (IEEE, 1999) pp. 983–985.
- [167] J. Gao, J. Zmuidzinas, B. A. Mazin, H. G. LeDuc, and P. K. Day, [Applied Physics Letters](#) **90**, 102507 (2007).
- [168] W. D. Oliver and P. B. Welander, [MRS Bulletin](#) **38**, 816 (2013).
- [169] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Córcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, [Physical Review B](#) **86**, 100506 (2012).
- [170] W. A. Cumming, [Journal of Applied Physics](#) **23**, 768 (1952).
- [171] M. R. Vant, R. B. Gray, R. O. Ramseier, and V. Makios, [Journal of Applied Physics](#) **45**, 4712 (1974).
- [172] S. G. Warren, [Applied Optics](#) **23**, 1206 (1984).
- [173] C. J. K. Richardson, N. P. Siwak, J. Hackley, Z. K. Keane, J. E. Robinson, B. Arey, I. Arslan, and B. S. Palmer, [Superconductor Science and Technology](#) **29**, 064003 (2016).
- [174] K. Shepard, C. Scheibelhut, P. Markovich, R. Benaroya, and L. Bollinger, [IEEE Transactions on Magnetics](#) **15**, 666 (1979).
- [175] J.-M. Vogt, O. Kugeler, and J. Knobloch, [Physical Review Special Topics - Accelerators and Beams](#) **16**, 102002 (2013).
- [176] C. Müller, J. H. Cole, and J. Lisenfeld, (2017).
- [177] F. Pobell, *Matter and methods at low temperatures* (Springer, 2007).
- [178] M. Sandberg, M. R. Vissers, J. S. Kline, M. Weides, J. Gao, D. S. Wisbey, and D. P. Pappas, [Applied Physics Letters](#) **100**, 262605 (2012).
- [179] A. J. Keller, P. B. Dieterle, M. Fang, B. Berger, J. M. Fink, and O. Painter, [Applied Physics Letters](#) **111**, 042603 (2017).
- [180] R. Chutani, M. Hasegawa, V. Maurice, N. Passilly, and C. Gorecki, [Sensors and Actuators A: Physical](#) **208**, 66 (2014).
- [181] DISCO Corporation, “[Laser Dicing Solutions | DISCO Laser Features | Methods](#),” .
- [182] L. Frunzio, A. Wallraff, D. Schuster, J. Majer, and R. Schoelkopf, [IEEE Transactions on Applied Superconductivity](#) **15**, 860 (2005).

- [183] R. Simmonds, K. Lang, D. Hite, S. Nam, D. Pappas, and J. Martinis, [Physical Review Letters](#) **93**, 077003 (2004).
- [184] Y. Chu, P. Kharel, W. H. Renninger, L. D. Burkhardt, L. Frunzio, P. T. Rakich, and R. J. Schoelkopf, [Science \(New York, N.Y.\)](#) **358**, 199 (2017).
- [185] Z. Kim, B. Suri, V. Zaretsky, S. Novikov, K. D. Osborn, A. Mizel, F. C. Wellstood, and B. S. Palmer, [Physical Review Letters](#) **106**, 120501 (2011).
- [186] J. M. Gambetta, C. E. Murray, Y.-K.-K. Fung, D. T. McClure, O. Dial, W. Shanks, J. W. Sleight, and M. Steffen, [IEEE Transactions on Applied Superconductivity](#) **27**, 1 (2017).
- [187] A. Dunsworth, A. Megrant, C. Quintana, Z. Chen, R. Barends, B. Burkett, B. Foxen, Y. Chen, B. Chiaro, A. Fowler, R. Graff, E. Jeffrey, J. Kelly, E. Lucero, J. Y. Mutus, M. Neeley, C. Neill, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, and J. M. Martinis, [Applied Physics Letters](#) **111**, 022601 (2017).
- [188] D. Zoepfl, P. R. Muppalla, C. M. F. Schneider, S. Kasemann, S. Partel, and G. Kirchmair, [AIP Advances](#) **7**, 085118 (2017).
- [189] G. Calusine, A. Melville, W. Woods, R. Das, C. Stull, V. Bolkhovskiy, D. Braje, D. Hover, D. K. Kim, X. Miloshi, D. Rosenberg, A. Sevi, J. L. Yoder, E. Dauler, and W. D. Oliver, [Applied Physics Letters](#) **112**, 062601 (2018).
- [190] G. Catelani, S. E. Nigg, S. M. Girvin, R. J. Schoelkopf, and L. I. Glazman, [Physical Review B](#) **86**, 184514 (2012).
- [191] M. Lenander, H. Wang, R. C. Bialczak, E. Lucero, M. Mariani, M. Neeley, A. D. O'Connell, D. Sank, M. Weides, J. Wenner, T. Yamamoto, Y. Yin, J. Zhao, A. N. Cleland, and J. M. Martinis, [Physical Review B](#) **84**, 024501 (2011).
- [192] C. Wang, Y. Y. Gao, I. M. Pop, U. Vool, C. Axline, T. Brecht, R. W. Heeres, L. Frunzio, M. H. Devoret, G. Catelani, L. I. Glazman, and R. J. Schoelkopf, [Nature Communications](#) **5**, 5836 (2014).
- [193] G. Catelani, R. J. Schoelkopf, M. H. Devoret, and L. I. Glazman, [Physical Review B](#) **84**, 064517 (2011).
- [194] Z. K. Mineev, I. M. Pop, and M. H. Devoret, [Applied Physics Letters](#) **103**, 142604 (2013).
- [195] Z. Mineev, K. Serniak, I. Pop, Z. Leghtas, K. Sliwa, M. Hatridge, L. Frunzio, R. Schoelkopf, and M. Devoret, [Physical Review Applied](#) **5**, 044021 (2016).
- [196] C. Eichler, C. Lang, J. M. Fink, J. Govenius, S. Filipp, and A. Wallraff, [Physical Review Letters](#) **109**, 240501 (2012).
- [197] S. J. Srinivasan, N. M. Sundaresan, D. Sadri, Y. Liu, J. M. Gambetta, T. Yu, S. M. Girvin, and A. A. Houck, [Physical Review A](#) **89**, 33857 (2014).
- [198] M. Pechal, L. Huthmacher, C. Eichler, S. Zeytinoglu, A. Abdumalikov, S. Berger, A. Wallraff, and S. Filipp, [Physical Review X](#) **4**, 041010 (2014).

- [199] Y. Yin, Y. Chen, D. Sank, P. J. J. O'Malley, T. C. White, R. Barends, J. Kelly, E. Lucero, M. Mariantoni, A. Megrant, C. Neill, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, *Physical Review Letters* **110**, 107001 (2013).
- [200] M. Pierre, I.-M. Svensson, S. Raman Sathyamoorthy, G. Johansson, and P. Delsing, *Appl. Phys. Lett.* **104**, 232604 (2014), arXiv:1406.2005.
- [201] E. Flurin, N. Roch, J. D. Pillet, F. Mallet, and B. Huard, *Physical Review Letters* **114**, 090503 (2015), arXiv:1401.5622.
- [202] Z. Leghtas, S. Touzard, I. M. Pop, A. Kou, B. Vlastakis, A. Petrenko, K. M. Sliwa, A. Narla, S. Shankar, M. J. Hatridge, M. Reagor, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret, *Science* **347**, 853 (2015).
- [203] M. Genkin and E. Lindroth, *Journal of Physics A: Mathematical and Theoretical* **41**, 425303 (2008).
- [204] N. Bergeal, F. Schackert, M. Metcalfe, R. Vijay, V. E. Manucharyan, L. Frunzio, D. E. Prober, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret, *Nature* **465**, 64 (2010).
- [205] J. I. Cirac, P. Zoller, H. J. Kimble, and H. Mabuchi, *Physical Review Letters* **78**, 3221 (1997), arXiv:9611017 [quant-ph].
- [206] Z. Wang and A. H. Safavi-Naeini, *Nature Communications* **8**, 15886 (2017), arXiv:1608.05946.
- [207] C. A. Sackett, D. Kielpinski, B. E. King, C. Langer, V. Meyer, C. J. Myatt, M. Rowe, Q. A. Turchette, W. M. Itano, D. J. Wineland, and C. Monroe, *Nature* **404**, 256 (2000).
- [208] M. Horodecki, P. Horodecki, and R. Horodecki, *Physics Letters A* **223**, 1 (1996).
- [209] B. M. Terhal, *Theoretical Computer Science* **287**, 313 (2002).
- [210] W. Pfaff, *Quantum Measurement and Entanglement of Spin Quantum Bits in Diamond*, Phd thesis, TU Delft (2013).
- [211] A. Ourjoumtsev, F. Ferreyrol, R. Tualle-Brouiri, and P. Grangier, *Nat. Phys.* **5**, 189 (2009).
- [212] A. Roy, A. D. Stone, and L. Jiang, *Physical Review A* **94**, 032333 (2016).
- [213] R. Hughes, G. Doolen, D. D. Awschalom, C. Caves, M. Chapman, R. Clark, D. Cory, D. DiVincenzo, A. Ekert, P. C. Hammel, P. Kwait, S. Lloyd, G. Milburn, T. Orlando, D. Steel, U. Vazirani, K. B. Whaley, and D. Wineland, *Quantum Computation Report of the Quantum Information Science and Technology Experts Panel*, Tech. Rep. (Los Alamos National Laboratory, 2004).
- [214] P. Kurpiers, T. Walter, P. Magnard, Y. Salathe, and A. Wallraff, *EPJ Quantum Technology* **4**, 8 (2017).
- [215] N. Leung, Y. Lu, S. Chakram, R. K. Naik, N. Earnest, R. Ma, K. Jacobs, A. N. Cleland, and D. I. Schuster, (2018).
- [216] S. Massar and S. Popescu, *Physical Review Letters* **74**, 1259 (1995).

- [217] J. Gambetta, A. Blais, D. I. Schuster, A. Wallraff, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Physical Review A* **74**, 42318 (2006).
- [218] C. M. Caves, C. A. Fuchs, and R. Schack, (2001), [10.1103/PhysRevA.65.022305](https://arxiv.org/abs/10.1103/PhysRevA.65.022305), [arXiv:0106133](https://arxiv.org/abs/0106133) [quant-ph] .
- [219] M. B. Plenio and S. Virmani, (2006), [arXiv:0504163](https://arxiv.org/abs/0504163) [quant-ph] .
- [220] R. Horodecki, P. Horodecki, M. Horodecki, and K. Horodecki, *Reviews of Modern Physics* **81**, 865 (2009).
- [221] M. B. Plenio, *Physical Review Letters* **95**, 090503 (2005).
- [222] L. S. Bishop, E. Ginossar, and S. M. Girvin, *Physical Review Letters* **105**, 100505 (2010).
- [223] D. Sank, Z. Chen, M. Khezri, J. Kelly, R. Barends, B. Campbell, Y. Chen, B. Chiaro, A. Dunsworth, A. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Mutus, M. Neeley, C. Neill, P. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, T. White, J. Wenner, A. N. Korotkov, and J. M. Martinis, *Physical Review Letters* **117**, 190503 (2016).
- [224] R. W. Andrews, A. P. Reed, K. Cicak, J. D. Teufel, and K. W. Lehnert, *Nat. Commun.* **6**, 10021 (2015).
- [225] J. Bochmann, A. Vainsencher, D. D. Awschalom, and A. N. Cleland, *Nat. Phys.* **9**, 712 (2013).
- [226] R. W. Andrews, R. W. Peterson, T. P. Purdy, K. Cicak, R. W. Simmonds, C. A. Regal, and K. W. Lehnert, *Nat. Phys.* **10**, 321 (2014), [arXiv:1310.5276](https://arxiv.org/abs/1310.5276) .
- [227] J. Chiaverini, D. Leibfried, T. Schaetz, M. D. Barrett, R. B. Blakestad, J. Britton, W. M. Itano, J. D. Jost, E. Knill, C. Langer, R. Ozeri, and D. J. Wineland, *Nature* **432**, 602 (2004).
- [228] D. Hucul, I. V. Inlek, G. Vittorini, C. Crocker, S. Debnath, S. M. Clark, and C. Monroe, *Nature Physics* **11**, 37 (2015), [arXiv:1403.3696](https://arxiv.org/abs/1403.3696) .
- [229] D. Gottesman and I. L. Chuang, *Nature* **402**, 390 (1999).
- [230] C. H. Bennett, G. Brassard, S. Popescu, B. Schumacher, J. A. Smolin, and W. K. Wootters, *Physical Review Letters* **76**, 722 (1996).
- [231] D. Deutsch, A. Ekert, R. Jozsa, C. Macchiavello, S. Popescu, and A. Sanpera, *Physical Review Letters* **77**, 2818 (1996).
- [232] C. Macklin, K. O'Brien, D. Hover, M. E. Schwartz, V. Bolkhovskiy, X. Zhang, W. D. Oliver, and I. Siddiqi, *Science* **350**, 307 (2015).
- [233] M. Boissonneault, J. M. Gambetta, and A. Blais, *Physical Review A* **79**, 013819 (2009).
- [234] M. Boissonneault, A. C. Doherty, F. R. Ong, P. Bertet, D. Vion, D. Esteve, and A. Blais, *Physical Review A* **85**, 022305 (2012).
- [235] L. Verney, R. Lescanne, M. H. Devoret, Z. Leghtas, and M. Mirrahimi, (2018), [arXiv:1805.07542](https://arxiv.org/abs/1805.07542) .

-
- [236] R. Lescanne, L. Verney, Q. Ficheux, M. H. Devoret, B. Huard, M. Mirrahimi, and Z. Leghtas, (2018), [arXiv:1805.05198](#).
- [237] T. K. Mavrogordatos, G. Tancredi, M. Elliott, M. Peterer, A. Patterson, J. Rahamim, P. Leek, E. Ginossar, and M. Szymańska, *Physical Review Letters* **118**, 040402 (2017).
- [238] D. I. Schuster, A. Wallraff, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. M. Girvin, and R. J. Schoelkopf, *Physical Review Letters* **94**, 123602 (2005).

Copyright Permissions

- Figures 5.10, 5.11, 5.33, 5.34, 5.35, 5.36, and 5.37 are adapted and reproduced here in accordance with the Author Copyright Agreement. The original publication is Copyright 2015 by the American Physical Society.
- Figures 4.2, 4.8, 4.9, 4.10, 4.13, 4.16, 5.29, and Table 5.6 are adapted and reproduced here in accordance with the Author Copyright Agreement. The original publication is Copyright 2016 by the American Physical Society.
- Other figures throughout are adapted from published versions as stated in their respective captions; the author retains non-exclusive rights to reproduce these works.