

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Addressing Key Scaling Challenges in Fixed-Frequency Transmon Architectures

Permalink

<https://escholarship.org/uc/item/6644g65v>

ISBN

9798276025643

Author

Chen, Larry

Publication Date

2025-12-06

Peer reviewed|Thesis/dissertation

Addressing Key Scaling Challenges in Fixed-Frequency Transmon Architectures

by

Larry Chen

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Physics

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Irfan Siddiqi, Chair

Professor K. Birgitta Whaley

Professor Harry Levine

Fall 2025

Addressing Key Scaling Challenges in Fixed-Frequency Transmon Architectures

Copyright 2025

by

Larry Chen

Abstract

Addressing Key Scaling Challenges in Fixed-Frequency Transmon Architectures

by

Larry Chen

Doctor of Philosophy in Physics

University of California, Berkeley

Professor Irfan Siddiqi, Chair

A central discourse in the field of superconducting quantum computing surrounds the question of whether the transmon qubit is the correct unit cell for scaling to fault-tolerant quantum computation. Since the transmon was first introduced, there has been a rapid development of more complex qubit designs that promise better noise protection and higher fidelity control. Despite these innovations, all of the largest superconducting qubit processors that exist today rely on the transmon qubit as the fundamental unit cell, and many use the flux-tunable variant to manage many of the difficult scaling constraints. However, the simplicity of the fixed-frequency, single-junction transmon remains highly appealing for building robust, large-scale systems. In this thesis, we address several key scaling challenges that have historically limited fixed-frequency transmon architectures. First, we present an analysis of the frequency-crowding problem for fixed-frequency architectures. We model the statistical distribution of qubit frequencies resulting from fabrication variance and propose strategies for mitigating spectral collisions. Next, we explore the use of parametrically modulated couplers to engineer fast, high-fidelity entanglement between highly-detuned transmons, which circumvents both frequency crowding and gate-speed limitations imposed by fixed-coupling architectures. Finally, we address the challenge of system instability caused by interactions with two-level system (TLS) defects. Since fixed-frequency transmons cannot be tuned away from these parasitic resonances, we introduce a resource-efficient method to integrate electric-field tuning of TLS defects into a scalable architecture. By addressing these key challenges, this work provides a pathway for scaling fixed-frequency transmons toward fault-tolerant quantum computation.

To my family and friends.

Contents

Contents	ii
List of Figures	iv
List of Tables	vi
1 Introduction	1
1.1 Quantum Information	4
1.2 Qubit Control	7
2 Classical and Quantum Circuits	12
2.1 Electrical Circuits	13
2.2 Lagrangian and Hamiltonian Analysis	14
2.3 Coupling	20
2.4 Voltage and Current Sources	31
2.5 Dissipation	36
2.6 Quantum Circuits	39
2.7 Nonlinear Circuits	43
3 Designer Microwave Circuits	49
3.1 Frequency Domain Analysis	49
3.2 Microwave Networks	52
3.3 The CPW Transmission Line	56
3.4 Coupled Line Models	67
3.5 Modeling CPW Meanders	70
4 Qubit Control	76
4.1 Single Qubit Gates	76
4.2 Two-Qubit Gates	87
4.3 Frequency Crowding	96
4.4 Architectural Tradeoffs	102
5 Parametric Gates	103

5.1	Parametric Modulation	103
5.2	Tunable Coupler Design	105
5.3	Gate Calibration	108
5.4	Benchmarking and Error Budget	110
5.5	Scaling and Outlook	114
6	System Stability	117
6.1	Two-Level System Defects	118
6.2	TLS Mitigation Strategies	119
6.3	Implementation of Local Electric-Field Tuning	121
6.4	TLS Spectroscopy	124
6.5	Electric-Field Tuning of TLS	128
6.6	Stabilizing the TLS Environment	130
6.7	Local Tuning vs Global Tuning	133
6.8	Multi-Qubit Stabilization	133
6.9	Outlook	135
	Bibliography	138
A	Circuits	155
A.1	Resonances	155
A.2	Coupling	156

List of Figures

2.1	LC Oscillator	15
2.2	Floating LC Oscillator	18
2.3	Capacitively coupled LC oscillators	20
2.4	Normal mode frequencies	23
2.5	Coupling geometries	26
2.6	Mutual inductance coupling	29
2.7	Coupling efficiencies for grounded and floating LC oscillators	30
2.8	Voltage and current driving of LC oscillators	32
2.9	Indirect driving of LC oscillators	34
2.10	Driven and damped RLC circuit	37
2.11	Flux quantization with Josephson junctions	44
2.12	Transmon as a nonlinear oscillator	47
3.1	Equivalent circuits for calculating effective resistance	50
3.2	Two-port network	53
3.3	Transmission line model	56
3.4	Transmission line as an impedance transformer	59
3.5	Transmission line resonators	60
3.6	Capacitive loading of a CPW resonator	62
3.7	CPW admittance as a function of coupling location	63
3.8	Purcell loss for lumped and distributed resonators	66
3.9	Coupled transmission line model	67
3.10	Modeling CPW Meanders	70
3.11	Frequency shift due to CPW meander	72
3.12	Experimental validation of meander simulations	74
4.1	Generalized rabi driving of a two-level system	78
4.2	Non-RWA effects on a driven two-level system	83
4.3	Fermionic simulation gates	90
4.4	Frequency optimization for an eight-qubit ring	100
4.5	Scaling frequency optimization to different lattice geometries	101
5.1	Flux-tunable coupling between two transmon qubits	106

5.2	Parametric coupling between two transmons	107
5.3	Calibration of a parametric $\sqrt{i}$ SWAP gate	108
5.4	Benchmarking a $\sqrt{i}$ SWAP gate	112
5.5	Driven decoherence with parametric driving	113
5.6	Cycle benchmarking fidelity over time	114
6.1	Two-level systems in superconducting qubit devices	118
6.2	Circuit model for the TIC-TAQ control line	122
6.3	Control line parameters for TLS tuning	125
6.4	Multi-level corrections to the AC Stark shift	126
6.5	Experimental validation of the multi-level AC Stark shift model	127
6.6	Characterization of a qubit and its TLS environment for a varying voltage bias	129
6.7	Stabilizing qubit performance with TIC-TAQ	131
6.8	Benchmarking TIC-TAQ across a multi-qubit device	134
6.9	Offset-charge dependence of MIST	136

List of Tables

2.1	Coupling efficiencies for Π - and T -circuits	28
4.1	Frequency constraints for a fixed-frequency, fixed-coupling gate architecture	98
5.1	Coherent errors extracted from cross-entropy benchmarking	111
5.2	Simultaneous and isolated two-qubit gate errors	114
5.3	Qubit parameters for the 6-qubit parametric device	115
5.4	Coupler parameters for the 6-qubit parametric device	115

Acknowledgments

This thesis marks the completion of the long road to a PhD, which would not have been possible without the many brilliant and wonderful people that I have had the pleasure of working with. First and foremost I have to thank my advisor, Irfan Siddiqi, for his thoughtful leadership of QNL and his willingness to let us explore the science that interests us. I've learned a great deal from Irfan over the years beyond physics, from cryogenics and plumbing to making slides and choosing the correct font (i.e., never Calibri). I am continually amazed by his ability to bring so many resources and talented people together in the pursuit of incredible science, and it is this group of people that has made my time at Berkeley so great.

Ravi Naik was the very first person I worked with in QNL, the one who proposed the parametric gates project, and the one who taught me much of what I know about circuit QED. Many of the ideas in this thesis came out of our discussions over the years. I'd like to thank Alexis Morvan for his mentorship and guidance, and for teaching me to focus on the right problems and not get lost in the noise. His enthusiasm for new ideas was always motivating and contagious. I am indebted to John Mark Kreikebaum for teaching me everything I know about fabrication, and for spending many long days in the cleanroom with me during the height of the pandemic to get a working parametric device. His methodical approach to problem solving, dedication to note taking, and meticulous work ethic have left a lasting impact, and are characteristics I still try to emulate in my own work. I also thank Will Livingston for teaching me to think outside the box with his clever ideas and unique insights. He taught me many things about qubit measurement, calibration, as well as how to write a group theory based sequencer.

After John Mark left, the responsibilities of managing the QNL cleanroom fell onto me and Christian Jünger. While I would be lying if I said I enjoyed all the cleanroom time, having Christian as my cleanroom buddy made it all the more bearable. He was also a great climbing partner and yoga buddy, helping me stay flexible after spending so many hours sitting tight in the lab. Kan-Heng Lee joined QNL two years later, just after the Great Flood of Plassys, and was immediately handed the difficult task of rebuilding our junction process, installing three new tools, and taking over cleanroom management. His hard work and dedication to improving all of our fabrication processes is commendable. I am immensely grateful to Kan-Heng for keeping me out of the cleanroom, and making all of my qubit devices come to life. None of the work in this thesis would have been possible without him.

Brian Marinelli deserves many thanks. As the other graduate student in my cohort, Brian and I endured almost every challenge in lab together, apart from my fridge-induced concussion. These included all-nighter's soldering bias tees, many late nights closing the fridge, a pandemic, multiple floods, and even a ransomware attack. Although we worked on separate projects, it always felt like we were working together, because Brian would be a sounding board for every one of my ridiculous ideas, and was always thinking of solutions to my experimental problems. I also need to thank Brian for filling the LN2 traps every week and helping me through all my problem sets.

Thank you also to Linus Kim, whose questions about seemingly simple concepts often uncovered my surface-level understanding and forced me to think more deeply about physics. Thank you to Noah Goss, whose optimism and enthusiasm for science was always inspiring. Thank you also for convincing me to go to many conferences together. Thanks to Chuanhong Liu and Ke Wang for many insightful discussions over lunch and for always accommodating my unwillingness to eat at the Berkeley dining halls. I am also grateful to them for keeping me up to date with all the newest Bay Area Chinese restaurants. Thanks to Brad Mitchell, whose selfless dedication to improving lab resources inspired me to do the same. Thanks to Gerwin Koolstra for his help with many of the early measurements with the parametric experiment, and Long Nguyen for teaching me the importance of selling your work. Thank you to everyone else in QNL for the great conversations and many lunch trains.

Thanks to the entire Qubic team (Gang Huang, Yilun Xu, Neelay Fruitwala, Abhi Rajagopala, and Neel Vora), and in particular Neelay, for their dedication to open-source control hardware and for supporting all of my feature requests and gateway builds.

Thank you to all the people in the Physics department who supported our research and kept the lab spaces running: Carlos Bustamante, Anthony Vitan, Espe Munoz-Riddle, and Stephen Pride Raffel. Thanks also to Chris Spitzer for keeping QNL running smoothly.

Thank you to Birgitta Whaley and Harry Levine for taking the time to be on my thesis committee, and Hartmut Haeflner and Joel Moore for serving on my qual committee.

I also want to thank Dave Schuster for taking a chance on me when I was just a first year in undergrad, and introducing me to the wonderful field of circuit QED. Many of the people I have worked with since coming to Berkeley (Ravi, Gerwin, Kan-Heng) have also come through SchusterLab, and their kindness and passion are a reflection of their time spent with Dave. Thank you also to Andrew Oriani, who was my first research mentor and taught me much about electronics, machining, and junctions. His insistence on doing everything the right way continues to influence my approach to research.

Finally, I'd like to thank all of my friends outside of lab who have made life in Berkeley so enjoyable: Christina, Wes, Anlei, Lydia, Kristen, Lizzie, Bel, Akshay, Gloria, and Matthew. Thank you also to my friends outside the Bay for their visits and support: Phillip, Jenny, Kevin, Emily, John, Caitlin, and JZ. Thank you to all the people at ResOak for their support and community. I'm also incredibly grateful for my family and their unconditional love and support through this entire journey. Thank you to my parents, Ling and Xiuxian, for always checking in on me and cheering me on. Thank you to my in-laws, Xianggong and Min, for giving me a second home in the Bay. Thank you to my sister, Morissa, for always reminding me that I'm still the least educated member of our family. I'm glad that I can now say it's her turn to reclaim that title. Most of all, I want to acknowledge my wife, Chelsea Hu, for being my number one supporter, and for always inspiring me to do everything to the best of my ability. Thank you for bringing me to Berkeley.

Chapter 1

Introduction

Information is physical, and the physical world is quantum; therefore, information must be quantum. This key insight, embodied in Feynman’s famous quote [60], catalyzed a decades-long effort to realize a system capable of performing universal quantum computation. At the center of this new computational paradigm is the replacement of the classical bit with its quantum counterpart: the *qubit*. Unlike a register of n -classical bits, describing an entangled system of n -qubits requires storing and manipulating 2^n complex numbers. This illustrates two uniquely quantum phenomena — superposition and entanglement — that give rise to a “quantum parallelism” responsible for the computational power of a quantum computer [46]. The potential of quantum information processing became more apparent with the discovery of several key algorithms that demonstrated a theoretical speedup over their best known classical counterparts. These include Shor’s algorithm [183], which offers an exponential speedup for integer factorization, and Grover’s algorithm [75, 76], which provides a quadratic speedup for searching unstructured databases.

While these developments demonstrate the potential power of quantum computation, the non-local correlations that underpin this power are inherently fragile. Quantum states are susceptible to decoherence, where unwanted interactions with the environment corrupt the stored information stored within these delicate states, introducing computational errors. This problem is compounded by a fundamental engineering dilemma: a qubit must be coupled to the external world for control and measurement, but this same coupling opens a channel for environmental noise to destroy the quantum correlations. Even if we could engineer a perfectly isolated “neutrino qubit”¹, our inability to manipulate it would render the qubit quite useless. Beyond environmental decoherence, the analog nature of a qubit state still makes it susceptible to control errors. Unlike a classical bit, which is digitized so that small fluctuations in a transistor voltage can be ignored, a qubit state is described by a continuous amplitude and phase. As a result, even small control errors can build up over the course of a computation and corrupt the information.

To overcome these challenges, the theory of quantum error correction (QEC) provides

¹This phrase should be attributed to John Martinis.

a framework for realizing fault-tolerant computation from noisy and imperfect qubits. The fundamental idea, first proposed by Peter Shor in 1996 [184], is to redundantly encode a logical qubit within the quantum correlations of a larger number of noisy physical qubits, so that errors on the constituent qubits can be detected and corrected. Shor’s discovery laid the foundation for an active field of research and the development of many QEC codes [5]. Among the most promising QEC schemes is the surface code, which requires only a two-dimensional array of physical qubits with local connectivity [98, 63]. The viability of all such codes is guaranteed by the threshold theorem, which proves that if the error rate of the physical qubits is below a certain critical value (the threshold), increasing the system size yields an exponential suppression of the logical error rate, enabling arbitrarily long and accurate computations [104, 3]. Achieving fault tolerance is therefore fundamentally a challenge of scale². A useful quantum computer will require not only a large number of logical qubits to operate on meaningful problem sizes, but also a large overhead of physical qubits to suppress errors below a functional level. This challenge, coupled with the promise of beyond-classical computation, has prompted billions³ of dollars in investment across both industry and academic research in an effort to realize a fault-tolerant quantum computer.

The pursuit of this goal is being carried out across an array of physical platforms, all of which must satisfy a set of core requirements, known as DiVincenzo’s criteria [52]. These include:

1. A quantum system in which a qubit can be encoded. This can be an electron spin, the polarization of a photon, or two quantum states in an atom.
2. A method to initialize qubits to a known state.
3. Quantum coherence must be maintained over sufficiently long time scales to enable both storage and processing of information.
4. The ability to apply unitary operations to manipulate the quantum state of one or more qubits.
5. A way to measure the state of a qubit.

These criteria can be satisfied on a variety of physical systems and have given rise to many quantum information processing platforms, each with their advantages and drawbacks. In this thesis, we will focus entirely on superconducting circuits, which has emerged as a leading platform for building a scalable quantum computer in the decades since Martinis, Devoret, and Clarke’s first demonstrations of macroscopic quantum behavior in a Josephson tunnel junction [122, 49]. These pioneering experiments paved the way for the first observations of coherent quantum oscillations in a superconducting charge qubit in 1999 [146]. Rapid progress ensued in the following years, when the remaining building blocks for quantum

²Qubit quality is of course important, but we will always need large numbers of high-quality qubits.

³In a panel with John Martinis at AQC 2025, Pedram Roushan recounted a funny story about how the original proposed title of Ref. [63] was “How to make a quantum computer with ten million dollars”.

computation were demonstrated, including longer coherence times in superconducting phase qubits [125], the first two-qubit gates [159, 189], and the circuit quantum electrodynamics (circuit QED) framework for qubit coupling and readout [202]. Despite these successes, the performance of early charge- and phase-qubits was ultimately capped by their extreme sensitivity to both charge and flux noise, which limited coherence times to the sub-microsecond regime. This changed with the introduction of the transmon qubit in 2007 [105] and its subsequent experimental realization [178], which demonstrated the promised exponential charge noise suppression and brought superconducting qubits into the microsecond coherence regime.

This breakthrough was a pivotal moment for the field, making the transmon qubit the dominant platform for decades to come. In the years since, transmon coherence times have improved by another three orders of magnitude [185], crossing into the millisecond regime [19], and systems consisting of ~ 100 transmon qubits have demonstrated the first instances of below-threshold error correction up to a distance-7 surface code [2]. While these successes, coupled with the transmon’s simplicity, make a strong case for the transmon qubit’s privileged position, many still question the practicality of a system that requires hundreds of physical qubits to realize a single logical qubit. This raises an important question: is the transmon qubit the quantum analog to the “transistor”, ready for the scale of fault-tolerant computation, or are we still in the vacuum-tube era, in search of a better fundamental building block?

This question has motivated a significant thrust of research into alternative qubit modalities that offer better noise protection and higher-fidelity control at the cost of system complexity. These include alternative noise-protected superconducting circuit designs [53, 79], as well as a different paradigm of bosonically encoded qubits [195]. The former leverages symmetries in the circuit design to suppress the qubit’s susceptibility to bit-flip noise, phase-flip noise, or both. These include single-mode circuits like the fluxonium [121, 149] or blochonium [161], as well as multi-mode circuits such as the $0-\pi$ [98, 24, 78], $\cos(2\theta)$ [54, 187], and bi-fluxon qubits [14, 91]. In the latter approach, a single logical degree of freedom is encoded redundantly within the large Hilbert space of a linear oscillator to yield a similar type of noise protection. Some examples of such bosonic codes include the GKP qubit [73, 29], cat qubit [136, 155], and binomial codes [133, 86]. The central premise to all of these approaches is that if the hardware unit cell can offer some intrinsic noise protection, we can engineer logical qubits with significantly fewer unit cells. In fact, this was demonstrated recently with a concatenated repetition code consisting of bosonically encoded cat qubits [166].

On the other hand, there is something to be said about the simplicity of the humble transmon, and we continue to see ever larger systems with lower error rates being built. In this thesis, we aim to strengthen the case for the transmon as a scalable unit cell. We will focus in particular on the fixed-frequency variant, by addressing three key scaling challenges that have often led people to favor its flux-tunable counterpart. First, we analyze the problem of frequency crowding inherent to fixed-frequency architectures, by modeling the inherent fabrication variance and proposing optimization strategies to mitigate spectral collisions. Motivated by these studies, we then explore the use of parametrically modulated

couplers to engineer fast two-qubit gates between highly detuned qubits. Finally, we introduce a resource-efficient method to integrate electric-field tuning of TLS defects into a fixed-frequency transmon architecture.

Outline

We will begin in the remainder of this chapter with a brief overview of quantum information theory. This is covered in more detail by many great references. We aim not to reproduce these, but to provide an introduction focused on quantum control from an experimental perspective. In Chapter 2, we introduce the fundamentals of superconducting circuits through a slightly unconventional approach. We introduce many of the techniques seen in the circuit QED literature, and show that many important effects relevant to the design of a superconducting qubit device are fundamentally classical. We then introduce the Josephson junction and discuss the implications of introducing nonlinearity and quantum behavior separately. In Chapter 3, we bring this lumped circuit theory into the practical domain, by guiding the reader through real design problems for the experimental devices used in the subsequent chapters. We will discuss how distributed element models are needed to accurately design superconducting qubit devices, and also propose more scalable approaches to device design with modular simulation.

In Chapter 4, we begin with some extended background on quantum control and single- and two-qubit gates with superconducting qubits. We discuss how the choice of two-qubit gate imposes many system design constraints and is often a dominant influence on the qubit architecture. We then focus on fixed-frequency and fixed-coupling architectures with all-microwave gates, and show how this qubit architecture can lead to untenable frequency allocation constraints as systems scale to larger sizes. In Chapter 5, we show how introducing a tunable coupler can enable fast parametrically-activated two-qubit gates even between far-detuned qubits. We explain the experimental implementation of a high-fidelity $\sqrt{\text{iSWAP}}$ gate using this architecture, and discuss the relevant error mechanisms that could be further improved to yield even lower error rates. Finally, in Chapter 6, we introduce a method for local and in-situ control of TLS defects using integrated electric-field tuning. This allows for control over the noise environment of each individual qubit, enhancing both performance and stability at a system level.

1.1 Quantum Information

A *qubit* is the fundamental unit of quantum information. In its most abstract form, this can be any two-level quantum system, with a Hilbert space that spans the states $|0\rangle$, $|1\rangle$. The Hamiltonian for such a system is typically written as

$$H_q = -\frac{\hbar\omega_q}{2} \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} = -\frac{\hbar\omega_q}{2} \sigma_z \quad (1.1)$$

where $\sigma_z = |0\rangle\langle 0| - |1\rangle\langle 1|$ and

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix} \quad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad (1.2)$$

are the eigenstates of the Hamiltonian⁴. Unlike a classical bit, the quantum bit can be in a *superposition* of the basis states. This can be expressed as an arbitrary quantum state

$$|\psi\rangle = \cos\theta |0\rangle + e^{i\phi} \sin\theta |1\rangle \quad (1.3)$$

where the state is both normalized ($\langle\psi|\psi\rangle = \cos^2\theta + \sin^2\theta = 1$) and contains only a relative phase ϕ between the two basis states (global phases are not experimentally measurable). Here, we notice that the set of unique states can be spanned by $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi]$ so that each unique state maps onto a point on the surface of a 3-dimensional unit sphere where $|0\rangle$ and $|1\rangle$ map onto the “north” and “south” poles respectively. This is typically referred to as the Bloch sphere.

Since the Hamiltonian is both time-independent and diagonal, the time evolution for the system is easily computed as

$$U(t) = e^{-iH_q t/\hbar} = e^{i\frac{1}{2}\omega_q t} |0\rangle\langle 0| + e^{-i\frac{1}{2}\omega_q t} |1\rangle\langle 1| \quad (1.4)$$

where each state gains a time-dependent phase that depends on its energy⁵ (or frequency). In particular, the relative phase between the two states will evolve at a frequency ω_q corresponding to their energy difference. Although the Hamiltonian is time-independent here, the dynamics are not time-independent. This makes it inconvenient to use these basis states as-is for encoding quantum information, since the logical state becomes time-dependent, rather than stationary, even when no additional operations are applied.

Unitary Transformations

To address this, we introduce the notion of a “rotating frame” or, more generally, a unitary transformation, which can simplify the system dynamics. In contrast, the original Hamiltonian, H is typically referred to as being in the “laboratory” frame. Given the lab frame Hamiltonian, H , the Schrodinger equation describes the resulting time evolution:

$$i\hbar \frac{\partial}{\partial t} |\psi(t)\rangle = H(t) |\psi(t)\rangle \quad (1.5)$$

subject to some initial condition $|\psi(t_0)\rangle = |\psi_0\rangle$. We now consider how we would move into a frame defined by a unitary U . Suppose instead that we replace $|\psi\rangle \rightarrow |\psi'\rangle = U |\psi\rangle$, we want

⁴We add a minus sign to the Hamiltonian so that the $|1\rangle$ state is the higher-energy state.

⁵With superconducting qubits, frequencies are the experimental quantity that are measured directly, so we will typically choose units such that $\hbar = 1$ and drop the explicit $\hbar$.

to find the corresponding hamiltonian H' that satisfies the Schrodinger equation. Making the substitutions, we find

$$i\hbar \frac{\partial}{\partial t} |\psi'\rangle = H' |\psi'\rangle \quad (1.6a)$$

$$i\hbar \frac{\partial}{\partial t} (U |\psi\rangle) = H' U |\psi\rangle \quad (1.6b)$$

$$i\hbar \dot{U} |\psi\rangle + i\hbar U \frac{\partial}{\partial t} |\psi\rangle = H' U |\psi\rangle \quad (1.6c)$$

where we use the product rule in the second line. We multiply both sides by $U^\dagger$ and rearrange, which gives

$$i\hbar \frac{\partial}{\partial t} |\psi\rangle = \left(U^\dagger H' U - i\hbar U^\dagger \dot{U} \right) |\psi\rangle \quad (1.7)$$

Finally, we see that we will recover the original equations of motion if

$$H = U^\dagger H' U - i\hbar U^\dagger \dot{U} \Rightarrow H' = U H U^\dagger + i\hbar \dot{U} U^\dagger \quad (1.8)$$

What this tells us is that even for the same Hamiltonian, the effective dynamics of a system can look very different depending on the states that we choose as our basis states.

The Qubit Frame

We now apply this to our qubit Hamiltonian and make a unitary transformation $U_{\text{rot}} = e^{iH_q t/\hbar}$ to move into the “rotating” frame. Since $[H_q, U] = 0$, we see that this transforms the Hamiltonian as

$$H'_q = U_{\text{rot}} H_q U_{\text{rot}}^\dagger + i\hbar \dot{U}_{\text{rot}} U_{\text{rot}}^\dagger = H_q U_{\text{rot}} U_{\text{rot}}^\dagger - H_q U_{\text{rot}} U_{\text{rot}}^\dagger = \mathbf{0} \quad (1.9)$$

and the basis states are modified as $|0'\rangle = U_{\text{rot}} |0\rangle$ and $|1'\rangle = U_{\text{rot}} |1\rangle$. In this rotating frame, the dynamics are governed by the propagator

$$U'(t) = e^{-iH'_q t/\hbar} = \mathbf{I} \quad (1.10)$$

where $\mathbf{I} = \sum_n |n\rangle\langle n|$ is the identity operator and the new basis states no longer have any time-dependence.

We refer to this unitary transformation as moving into the rotating frame because in the Bloch sphere picture, the time-evolution (in the lab frame) of any state not at the poles looks like rotation around the z-axis at a frequency ω_q . Thus, in order for the time-evolution to appear stationary, the dynamics must be observed from the perspective of a co-rotating observer. Experimentally, this means that the frequency of the qubit must be known to a high-degree of precision, since the logical rotating frame is actually defined by a external frequency reference and not the qubit itself. We refer to this technique as *frame tracking*.

1.2 Qubit Control

We now discuss the coherent control of a superconducting qubit. In order to manipulate quantum information, we must be able to take a qubit from an initial state $|\psi_i\rangle$ to an arbitrary final state $|\psi_f\rangle$. For a single qubit, this amounts to implementing rotations on the Bloch sphere.

Single Qubit Rotations

To implement a rotation around the X-axis, we consider adding a drive Hamiltonian of the form

$$H_d = \Omega(t) \cos(\omega_d t + \phi) \sigma_x \quad (1.11)$$

where $\Omega(t)$ describes a time-dependent drive-envelope, ω_d is a driving frequency, and ϕ is a drive phase. For simplicity, we assume that only the drive amplitude can be time-dependent and keep the drive phase ϕ time-independent. This is equivalently written as

$$H_d = \frac{\Omega(t)}{2} [e^{i(\omega_d t + \phi)} + e^{-i(\omega_d t + \phi)}] (\sigma_+ + \sigma_-) = \frac{1}{2} [\tilde{\Omega}(t)e^{i\omega_d t} + \tilde{\Omega}^*(t)e^{-i\omega_d t}] (\sigma_+ + \sigma_-) \quad (1.12)$$

where we have folded the phase ϕ into the now complex amplitude $\tilde{\Omega}(t) = \Omega(t)e^{i\phi}$. We now move into the rotating frame that defines the logical qubit, defined by $U = \exp[-\frac{1}{2}i\omega_q t \sigma_z]$ as above. We can easily calculate the following the following terms:

$$\dot{U}U^\dagger = -\frac{1}{2}i\omega_q \sigma_z e^{-\frac{1}{2}i\omega_q t \sigma_z} e^{\frac{1}{2}i\omega_q t \sigma_z} = -\frac{1}{2}i\omega_q \sigma_z \quad (1.13a)$$

$$U\sigma_z U^\dagger = \left(e^{-\frac{1}{2}i\omega_q t} |0\rangle\langle 0| + e^{\frac{1}{2}i\omega_q t} |1\rangle\langle 1| \right) \sigma_z \left(e^{\frac{1}{2}i\omega_q t} |0\rangle\langle 0| + e^{-\frac{1}{2}i\omega_q t} |1\rangle\langle 1| \right) = \sigma_z \quad (1.13b)$$

$$U\sigma_+ U^\dagger = \left(e^{-\frac{1}{2}i\omega_q t} |0\rangle\langle 0| + e^{\frac{1}{2}i\omega_q t} |1\rangle\langle 1| \right) \sigma_+ \left(e^{\frac{1}{2}i\omega_q t} |0\rangle\langle 0| + e^{-\frac{1}{2}i\omega_q t} |1\rangle\langle 1| \right) = \sigma_+ e^{i\omega_q t} \quad (1.13c)$$

$$U\sigma_- U^\dagger = \left(e^{-\frac{1}{2}i\omega_q t} |0\rangle\langle 0| + e^{\frac{1}{2}i\omega_q t} |1\rangle\langle 1| \right) \sigma_- \left(e^{\frac{1}{2}i\omega_q t} |0\rangle\langle 0| + e^{-\frac{1}{2}i\omega_q t} |1\rangle\langle 1| \right) = \sigma_- e^{-i\omega_q t} \quad (1.13d)$$

so the transformed Hamiltonian is

$$H' = U(H_q + H_d)U^\dagger + iUU^\dagger = UH_q U^\dagger + iUU^\dagger + UH_d U^\dagger \quad (1.14a)$$

$$= \frac{1}{2} [\tilde{\Omega}(t)e^{i\omega_d t} + \tilde{\Omega}^*(t)e^{-i\omega_d t}] (\sigma_+ e^{i\omega_q t} + \sigma_- e^{-i\omega_q t}) \quad (1.14b)$$

$$= \frac{1}{2} [\tilde{\Omega}(t)e^{i\Sigma t} + \tilde{\Omega}^*(t)e^{i\Delta t}] \sigma_+ + \frac{1}{2} [\tilde{\Omega}(t)e^{-i\Delta t} + \tilde{\Omega}^*(t)e^{-i\Sigma t}] \sigma_- \quad (1.14c)$$

where $\Delta = \omega_q - \omega_d$ and $\Sigma = \omega_q + \omega_d$. If the drive is on resonance ($\omega_d = \omega_q$), we see that $\Delta = 0$ and $\Sigma = 2\omega_q = 2\omega_d$, so that two of the four terms have a much slower time dependence than the rest. In this limit, we can make the well-known *rotating-wave approximation* (RWA), where we drop the fast terms, to simplify the Hamiltonian to

$$H' \approx \frac{1}{2} [\tilde{\Omega}(t)^* \sigma_+ + \tilde{\Omega}(t) \sigma_-] = \frac{1}{2} [\tilde{\Omega}(t) |0\rangle\langle 1| + \tilde{\Omega}^*(t) |1\rangle\langle 0|] \quad (1.15)$$

so that the only time dependence that remains is in the drive amplitude. Since we assumed earlier that the complex phase has no time dependence, the Hamiltonian will commute with itself at all times ($[H'(t), H'(t')] = 0$) and the propagator is given by

$$U_d(t) = \exp \left[-i \int_0^t H'(t') dt' \right] = \exp \left[-\frac{1}{2} i \theta(t) (\cos(\phi) \sigma_x - \sin(\phi) \sigma_y) \right] \quad (1.16a)$$

$$= \begin{pmatrix} \cos \frac{\theta(t)}{2} & -i e^{i\phi} \sin \frac{\theta(t)}{2} \\ -i e^{-i\phi} \sin \frac{\theta(t)}{2} & \cos \frac{\theta(t)}{2} \end{pmatrix} \quad (1.16b)$$

$$= R_{xy}(\theta, \phi) \quad (1.16c)$$

where $\theta(t) = \int_0^t \Omega(t') dt'$ is the total area under the drive envelope $\Omega(t)$. We recognize this as a rotation matrix on the Bloch sphere, with a rotation angle determined by $\theta(t)$ and axis defined by the angle ϕ in the XY-plane. For $\phi = 0$ this becomes a rotation about the X-axis and for $\phi = \frac{\pi}{2}$ this becomes a rotation about the Y-axis.

What about Z-rotations? The rotation matrix for a Z-rotation looks like:

$$R_z(\phi) = e^{-i \frac{\phi}{2} \sigma_z} = \begin{pmatrix} e^{-i \frac{\phi}{2}} & 0 \\ 0 & e^{i \frac{\phi}{2}} \end{pmatrix} \quad (1.17)$$

which looks exactly like the time evolution of the qubit in the laboratory frame if we take $\omega_q t \rightarrow -\phi$ (see Eq. 1.4). This suggests that we can implement a Z-rotation by simply detuning the qubit from its frame tracking frequency for a short period of time. Suppose the frequency of the qubit can be tuned such that the instantaneous frequency at a time t is given by $\omega(t) = \omega_q + \Delta(t)$. After moving into the qubit frame, we find

$$H' = U H_q U^\dagger + i U U^\dagger = \frac{\Delta(t)}{2} \sigma_z \quad (1.18)$$

so that the time evolution under this detuning looks like

$$U_d(t) = \exp \left[-i \frac{1}{2} \phi(t) \sigma_z \right] = \begin{pmatrix} e^{-i \frac{\phi(t)}{2}} & 0 \\ 0 & e^{i \frac{\phi(t)}{2}} \end{pmatrix} = R_z(\phi) \quad (1.19)$$

where $\phi(t) = \int_0^t \Delta(t') dt'$ is the integrated detuning.

Virtual-Z Gates

In the previous section we showed that a temporary qubit detuning could be used to realize a Z rotation. But since the detuning is symmetric between the qubit frequency and the frame tracking frequency, this suggests that we could just as well detune the frame tracking frequency to implement the same gate. This is equivalent to adding discrete phase updates to the frame, which has the advantage of taking zero time. Physically, this involves updating

the phases of the XY rotations that encode the qubit frame. To see how this works, we look at how rotations R_z commute across XY rotations R_{xy} . Suppose we have a rotation $R_z(\alpha)$ before and a rotation $R_z(\gamma)$ after a rotation $R_{xy}(\theta, \phi)$ ⁶. The resulting rotation matrix looks like

$$R_z(\gamma)R_{xy}(\theta, \phi)R_z(\alpha) = \begin{pmatrix} e^{-i\frac{\gamma}{2}} & 0 \\ 0 & e^{i\frac{\gamma}{2}} \end{pmatrix} \begin{pmatrix} \cos \frac{\theta}{2} & -ie^{-i\phi} \sin \frac{\theta}{2} \\ -ie^{i\phi} \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{pmatrix} \begin{pmatrix} e^{-i\frac{\alpha}{2}} & 0 \\ 0 & e^{i\frac{\alpha}{2}} \end{pmatrix} \quad (1.20a)$$

$$= \begin{pmatrix} e^{-i\frac{(\alpha+\gamma)}{2}} \cos \frac{\theta}{2} & -ie^{-i(\phi-(\alpha-\gamma)/2)} \sin \frac{\theta}{2} \\ -ie^{i(\phi-(\alpha-\gamma)/2)} \sin \frac{\theta}{2} & e^{i\frac{(\alpha+\gamma)}{2}} \cos \frac{\theta}{2} \end{pmatrix} \quad (1.20b)$$

If we let $\gamma = -\alpha$, we can simplify this to

$$R_z(-\alpha)R_{xy}(\theta, \phi)R_z(\alpha) = \begin{pmatrix} \cos \frac{\theta}{2} & -ie^{-i(\phi-\alpha)} \sin \frac{\theta}{2} \\ -ie^{i(\phi-\alpha)} \sin \frac{\theta}{2} & \cos \frac{\theta}{2} \end{pmatrix} = R_{xy}(\theta, \phi - \alpha) \quad (1.21)$$

and we see that this becomes equivalent to a rotation around a modified angle $\phi' = \phi - \alpha$. This implies that

$$R_{xy}(\theta, \phi)R_z(\alpha) = R_z(\alpha)R_{xy}(\theta, \phi - \alpha) \quad (1.22)$$

so we see that any Z-rotation $R_z(\alpha)$ applied *before* $R_{xy}(\theta, \phi)$ is equivalent to applying $R_z(\alpha)$ *after* $R_{xy}(\theta, \phi - \alpha)$. In this way, we can keep deferring Z gates without changing the total effective unitary, by updating the phase of our rotation axis and pushing all Z rotations to the very end of our circuit. Since the final measurement is always performed in the Z basis, this final deferred Z-rotation before the measurement can simply be dropped with no consequence.

Entanglement

With the ability to do X-, Y-, and Z-rotations, we gain the ability to implement arbitrary rotations of a single-qubit state on the Bloch sphere. However, due to the measurement postulate of quantum mechanics, a single qubit in isolation is no different from a probabilistic classical bit. The true power of quantum information is unlocked when multiple qubits can be brought together and allowed to interact in non-classical ways.

A particularly clean example of the power of entanglement is the idea of *superdense coding*. To explain this we first introduce the two-qubit Bell states

$$|\phi_{\pm}\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle) \quad |\psi_{\pm}\rangle = \frac{1}{\sqrt{2}}(|01\rangle \pm |10\rangle) \quad (1.23)$$

which are maximally entangled. For any of these Bell states, if we measure only a single qubit, we will obtain a 0 or 1 with 50% probability. However given the measurement result of the first qubit, we always know with certainty what the state of the second qubit will be

⁶Note that time goes from right to left with operators.

upon measurement. Now suppose Alice and Bob share an entangled pair of qubits $|\phi_+\rangle$. By performing just a single-qubit operation on her qubit $(I, \sigma_x, \sigma_y, \sigma_z)$, she can transform $|\phi_+\rangle$ into any one of the 4 Bell states (up to a global phase):

$$I |\phi_+\rangle = |\phi_+\rangle \quad (1.24a)$$

$$\sigma_x^A |\phi_+\rangle = |\psi_+\rangle \quad (1.24b)$$

$$\sigma_y^A |\phi_+\rangle = e^{-i\frac{\pi}{2}} |\psi_-\rangle \quad (1.24c)$$

$$\sigma_z^A |\phi_+\rangle = |\phi_-\rangle \quad (1.24d)$$

If she then sends her single qubit to Bob, he now has control over both qubits and can do a Bell basis measurement⁷ to determine which one of the 4 Bell states he possesses. There are 4 possible Bell states, which means that Alice has managed to transmit 2 bits of classical information by transmitting only a single quantum bit! Now one might object by saying that Alice and Bob still had to transmit a single qubit in the beginning to share the entangled resource (or maybe a third party entangled the pair and distributed one qubit to each of them). However, this initial transmission contained no information—the information was encoded into the two qubits only after Alice implements her single-qubit rotation.

Having seen that entanglement is useful, we now consider how to generate it. Building on our single-qubit example, we consider adding a drive Hamiltonian of the form

$$H_d = \Omega(t) \cos(\omega_d t + \phi) \sigma_x^a \sigma_x^b \quad (1.25)$$

$$= \frac{1}{2} \left[\tilde{\Omega}(t) e^{i\omega_d t} + \tilde{\Omega}^*(t) e^{-i\omega_d t} \right] (\sigma_+^a \sigma_+^b + \sigma_+^a \sigma_-^b + \sigma_-^a \sigma_+^b + \sigma_-^a \sigma_-^b) \quad (1.26)$$

where $\Omega(t)$ again describes some time-dependent drive envelope. If we move into the rotating frame of the two qubits, defined by $U = \exp \left[-\frac{i}{2} (\omega_a \sigma_z^a + \omega_b \sigma_z^b) t \right]$, the drive Hamiltonian becomes

$$H'_d = U H_d U^\dagger \quad (1.27a)$$

$$= \frac{1}{2} \left[\tilde{\Omega}(t) e^{i\omega_d t} + \tilde{\Omega}^*(t) e^{-i\omega_d t} \right] \times (\sigma_+^a \sigma_+^b e^{i\Sigma_{ab}t} + \sigma_+^a \sigma_-^b e^{i\Delta_{ab}t} + \text{c.c.}) \quad (1.27b)$$

$$= \frac{1}{2} \tilde{\Omega}(t) (e^{i(\omega_d + \Sigma_{ab})t} + e^{-i(\omega_d - \Sigma_{ab})t}) \sigma_+^a \sigma_+^b \quad (1.27c)$$

$$+ \frac{1}{2} \tilde{\Omega}^*(t) (e^{i(\omega_d + \Delta_{ab})t} + e^{-i(\omega_d - \Delta_{ab})t}) \sigma_+^a \sigma_-^b + \text{c.c.} \quad (1.27d)$$

where we now define $\Delta_{ab} = \omega_a - \omega_b$ and $\Sigma_{ab} = \omega_a + \omega_b$. Here, we see that if the drive frequency is $\omega_d \pm \Delta_{ab}$ or $\omega_d = \pm \Sigma_{ab}$, we obtain a time-independent term in the drive Hamiltonian in the rotating frame, in which case we can make the RWA again and drop all remaining terms. Suppose we choose $\omega_d = \Delta_{ab}$. Then we are left with

$$H'_d \approx \frac{1}{2} \tilde{\Omega}^*(t) \sigma_+^a \sigma_-^b + \frac{1}{2} \tilde{\Omega}(t) \sigma_-^a \sigma_+^b = \frac{1}{2} \left[\tilde{\Omega}(t) |01\rangle\langle 10| + \tilde{\Omega}^*(t) |10\rangle\langle 01| \right] \quad (1.28)$$

⁷To be precise, he can apply a CNOT followed by a Hadamard gate, to map the Bell basis onto the two-qubit Z basis.

which looks a lot like Eq. 1.15, except for the eigenstates involved. Now suppose we start in state $|01\rangle$. For a drive envelope $\tilde{\Omega}(t) = \Omega(t)e^{i\phi_d}$ where $\theta = \int_0^\tau \Omega(t) dt = \frac{\pi}{2}$ with drive phase $\phi_d = \frac{\pi}{2}$, we will end up in a state

$$U_d |10\rangle = e^{-i \int_0^\tau H'_d dt} |10\rangle = \frac{1}{\sqrt{2}} (|01\rangle + |10\rangle) \quad (1.29)$$

which is precisely the maximally entangled Bell state $|\phi_+\rangle$! More generally, we can show that the entire space of two-qubit entangling operations can be written in the form

$$U = \exp \left[-i \frac{\pi}{2} (t_x \sigma_x^a \sigma_x^b + t_y \sigma_y^a \sigma_y^b + t_z \sigma_z^a \sigma_z^b) \right] \quad (1.30)$$

for $t_x, t_y, t_z \in [-1, 1]$, up to single-qubit rotations [42]. This is sometimes called the “canonical” gate or the “KAK” gate, and the coefficients t_x, t_y, t_z are often constrained to a subset

$$\left(0 \leq t_z \leq t_y \leq t_x \leq \frac{1}{2} \right) \cup \left(0 \leq t_z \leq t_y \leq 1 - t_x \leq \frac{1}{2} \right) \quad (1.31)$$

which defines the *Weyl chamber* [215]. Each coordinate (t_x, t_y, t_z) within the Weyl chamber defines a unique two-qubit unitary, that cannot be mapped to a different point in the Weyl chamber with only single-qubit operations.

Chapter 2

Classical and Quantum Circuits

The field of circuit QED is built upon a foundation of nonlinear quantum circuits. The quantum regime is required to enable the manipulation of quantum information in discrete energy levels and maintain coherence between them. And although not strictly required to exhibit quantum effects, it is typically necessary to introduce nonlinear behavior to gain universal control over the discrete energy levels in a quantum circuit¹. As a result, we require a quantum description of these circuits to gain a complete understanding of their behavior

There are many excellent reviews of circuit quantization and quantum-mechanical circuit analysis, and we encourage the interested reader to check out the following references:

- **Introduction to quantum electromagnetic circuits** *Vool and Devoret* [201]: This is a short review of many foundational concepts in quantum circuit theory, consolidated from a set of Les Houches lecture notes.
- **Lecture Notes on Quantum Electrical Circuits** *Ciani, DiVincenzo, and Terhal* [39]: A textbook, complete with exercises, that covers a range of introductory to advanced topics including circuit quantization, black-box quantization, non-reciprocal circuits, and noise.
- **A review of design concerns in superconducting quantum circuits** *Levenson-Falk and Shanto*: A modern review of design concerns for superconducting quantum circuits. The review contains an excellent reference list, making it a good launching point to explore a particular design topic in more depth.

A common trend in much of the circuit QED literature is a heavy emphasis on the machinery required to give a complete quantum-mechanical description of the circuit, which can often obscure the underlying simplicity and classical origin of a given effect. In many situations, a classical and linear description of the circuit is sufficient for understanding the salient properties of a circuit QED system. This is especially true of much of the supporting

¹Since the Josephson junction typically introduces both quantum-ness and nonlinearity, the two are often conflated.

circuitry (readout resonators, filters, control lines, etc.) involved in controlling and measuring a superconducting qubit. From a design perspective, understanding whether a given circuit behavior or effect is quantum, nonlinear, or both, is important, since effects that are both linear and classical can generally be computed or simulated much more efficiently.

In this chapter, we will introduce and review some foundational concepts in circuit QED, but from the perspective of linear and classical circuits. We will then discuss the Josephson junction, which introduces both quantum and nonlinear behavior. This will build the foundation for the subsequent chapter, where we work through several practical examples of superconducting quantum device design.

2.1 Electrical Circuits

Electrical circuits are networks consisting of discrete components, or circuit elements, that describe the collective behavior of electrons in matter and their interaction with electromagnetic fields. These circuit elements can be either lumped or distributed, and describe the relationship between the voltages and currents at each port. While circuit elements are commonly specified in terms of voltage (V) and current (I), they can also be described with respect to their time integrals, flux (Φ) and charge (Q) respectively.

$$\Phi(t) = \int_{-\infty}^t V(t') dt' \quad Q(t) = \int_{-\infty}^t I(t') dt' \quad (2.1)$$

When circuit elements are connected together at their ports to form a network, the voltages and currents are further constrained by Kirchhoff's laws:

$$\sum_v I_v = 0 \quad (2.2)$$

$$\sum_e V_e = 0, \quad (2.3)$$

which state that sum of all currents flowing into a node v is always zero and that the sum of all voltage differences around a closed loop is also zero.

These equations can be seen as a consequence of Maxwell's equations. For Kirchhoff's current law, we start with Ampere's law and integrate this over a surface S enclosing the node.

$$0 = \oint_S \frac{1}{\mu_0} (\nabla \times \mathbf{B}) \cdot \hat{\mathbf{n}} dS = \oint_S \left(\mathbf{J} + \epsilon_0 \frac{\partial \mathbf{E}}{\partial t} \right) \cdot \hat{\mathbf{n}} dS \quad (2.4)$$

Here, the left-hand side is zero from the divergence theorem and the fact that the divergence of a curl is always zero. We then use Gauss's law, which tells us that the integral of electric flux over the closed surface is equal to the total enclosed charge.

$$0 = \oint_S \mathbf{J} \cdot \hat{\mathbf{n}} dS + \frac{\partial}{\partial t} \left[\epsilon_0 \oint_S \mathbf{E} \cdot \hat{\mathbf{n}} dS \right] = \sum_v I_v + \dot{Q} \quad (2.5)$$

If the total charge on each node is conserved (no time-varying voltage), then $\dot{Q} = 0$, and we have recovered Kirchhoff's current law.

Similarly, for Kirchhoff's voltage law, we start with Faraday's law and integrate over the open surface S enclosed by a loop C . The first term can then be expressed as a line integral using Stokes' theorem, and the second term is simply the magnetic flux threaded through the loop.

$$0 = \int_S \left(\nabla \times \mathbf{E} + \frac{\partial \mathbf{B}}{\partial t} \right) \cdot \hat{\mathbf{n}} = \oint_C \mathbf{E} \cdot d\mathbf{l} + \dot{\Phi}_{\text{ext}} = \sum_e V_e \quad (2.6)$$

Again, if the external flux enclosed by the loop is conserved (no time-varying magnetic field), then $\dot{\Phi} = 0$, and we have recovered Kirchhoff's voltage law.

Finally we note that we can also integrate these constraints to express them in terms of charge and flux.

$$\sum_v Q_v = Q_T \quad \sum_e \Phi_e = \Phi_{\text{ext}} \quad (2.7)$$

From the above analysis, we see that Kirchhoff's laws can be interpreted as a discretization of Maxwell's equations. Rather than solving for the electric and magnetic fields over 3D space, the circuit represents a reduced basis over which the voltages/currents or flux/charge can be solved.

Circuit elements relate these quantities to one another, giving us the following relationships for a resistor (R), inductor (L), capacitor (C):

$$V = RI \quad \Phi = LI \quad Q = CV \quad (2.8)$$

that are typically seen in the circuit QED literature². Given a network of such elements, we could analyze the behavior of the circuit directly using these circuit relationships and Kirchhoff's laws as constraints. However, it is often useful to introduce a circuit Lagrangian or Hamiltonian, which we will briefly explain in the following section.

2.2 Lagrangian and Hamiltonian Analysis

Circuit Lagrangians and Hamiltonians are a foundational tool for analyzing circuits, and have been covered extensively in the circuit QED literature. For a comprehensive treatment of the subject, we will refer the reader to Ref. [201] and Ref. [39]. Nonetheless, we will briefly discuss the Lagrangian and Hamiltonian formulations here, using the canonical LC oscillator as an example, to emphasize several points that are often overlooked.

The LC oscillator consists of a single degree of freedom that is typically expressed in terms of the node flux³ ϕ . To construct the Lagrangian, we need to consider the total energy

²The structure of these relationships also implies the existence of another type of component relating the flux to the charge directly. This is known as a memristor [38], which we will not discuss in this thesis.

³We can relate the node fluxes ϕ_i, ϕ_j at either end of a circuit to the branch flux Φ across the circuit as $\Phi = \phi_i - \phi_j$, which chooses a sign convention for the current. The ground node is typically assigned $\phi_g = 0$.




Figure 2.1: **LC Oscillator.** (a) A circuit diagram showing a parallel LC oscillator, consisting of a capacitance C and inductance L in parallel. (b) The circuit variables ϕ and Q can also be expressed as a mode amplitude a . The equations of motion for a trace out a circle in the complex plane at a frequency ω_0 .

stored in the lossless inductor and capacitor, which can be expressed as

$$E_L = \int IV dt = \frac{1}{L} \int \phi \dot{\phi} dt = \frac{\phi^2}{2L} \quad (2.9)$$

$$E_C = \int IV dt = C \int \ddot{\phi} \dot{\phi} dt = \frac{1}{2} C \dot{\phi}^2 \quad (2.10)$$

where we assume that the inductor and capacitor are linear so L and C are constants. Here, we see that E_C looks like a kinetic energy where ϕ is the position, $V = \dot{\phi}$ is the velocity, and the effective mass is C . The potential energy is given by E_L and is a quadratic potential, so we can write the circuit Lagrangian as

$$\mathcal{L} = E_C - E_L = \frac{1}{2} C \dot{\phi}^2 - \frac{1}{2L} \phi^2 \quad (2.11)$$

From here, we determine the equations of motion using the Euler-Lagrange equations and find

$$\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial \dot{\phi}} \right) = C \ddot{\phi} \quad \frac{\partial \mathcal{L}}{\partial \phi} = -\frac{1}{L} \phi \quad \Rightarrow \ddot{\phi} = -\frac{1}{LC} \phi \quad (2.12)$$

which matches the second order differential equation for a harmonic oscillator with a resonant frequency $\omega_0 = \frac{1}{\sqrt{LC}}$. If we recall that $C\ddot{\phi} = C\dot{V} = \dot{Q} = I$ and $-\phi/L = -I$, we see that this equation of motion is simply enforcing Kirchhoff's current law between the two branches.

Once we have the Lagrangian, we can also define the conjugate momentum of ϕ as

$$\frac{\partial \mathcal{L}}{\partial \dot{\phi}} = C \dot{\phi} = CV = Q \quad (2.13)$$

and we can then express the Hamiltonian as

$$H = Q\dot{\phi} - \mathcal{L} = \frac{Q^2}{2C} + \frac{\phi^2}{2L} \quad (2.14)$$

From Hamilton's equations, we now obtain two first order differential equations governing the behavior of the circuit

$$\dot{\phi} = \frac{\partial H}{\partial Q} = \frac{1}{C}Q \quad \dot{Q} = -\frac{\partial H}{\partial \phi} = -\frac{1}{L}\phi \quad (2.15)$$

These equations can be expressed as a matrix equation

$$\frac{d}{dt} \begin{pmatrix} \phi \\ Q \end{pmatrix} = \frac{1}{\sqrt{LC}} \begin{pmatrix} 0 & \sqrt{\frac{L}{C}} \\ -\sqrt{\frac{C}{L}} & 0 \end{pmatrix} \begin{pmatrix} \phi \\ Q \end{pmatrix} = \omega_0 \underbrace{\begin{pmatrix} 0 & Z_r \\ -\frac{1}{Z_r} & 0 \end{pmatrix}}_{\mathbf{M}} \begin{pmatrix} \phi \\ Q \end{pmatrix} \quad (2.16)$$

where we define the impedance of the oscillator as $Z_r = \sqrt{L/C}$. The eigenvalues of $\mathbf{M}$ are $\lambda = \pm i$, so the solutions are

$$\begin{pmatrix} \phi(t) \\ Q(t) \end{pmatrix} = \phi_+ \begin{pmatrix} 1 \\ \frac{i}{Z_r} \end{pmatrix} e^{i\omega_0 t} + \phi_- \begin{pmatrix} 1 \\ -\frac{i}{Z_r} \end{pmatrix} e^{-i\omega_0 t} \quad (2.17)$$

where $\phi_{\pm}$ are constants (in units of flux) that are determined by the initial conditions. The factor of $\pm i$ in the eigenvectors tells us that the oscillations of the charge Q lag behind ϕ by a phase of 90° .

While the coupled first-order differential equations are straightforward to solve, we can also express the Hamiltonian in a more natural way by defining a new variable a , which represents a mode amplitude, in terms of ϕ and Q as

$$a = \frac{1}{\sqrt{2\hbar Z_r}} \phi + i\sqrt{\frac{Z_r}{2\hbar}} Q \quad (2.18)$$

We can easily check here that a has units of square root photon number. We can now express the Hamiltonian in terms of the mode amplitude a , which highlights the fact that there is only a single degree of freedom involved. The Hamiltonian can be written in terms of the mode amplitude and its complex conjugate a^* as

$$H = \hbar\omega_0 a^* a = \hbar\omega_0 |a|^2 = \frac{1}{2} \frac{\omega_0}{Z_r} \phi^2 + \frac{1}{2} \omega_0 Z_r Q^2 = \frac{\phi^2}{2L} + \frac{Q^2}{2C} \quad (2.19)$$

Now that we have the Hamiltonian expressed in terms of a, a^* , can we rewrite Hamilton's equations in terms of a and a^* . To do so, we first note that

$$\phi = \sqrt{\frac{\hbar Z_r}{2}} (a + a^*) \quad Q = i\sqrt{\frac{\hbar}{2Z_r}} (a^* - a) \quad (2.20)$$

so we can express $\dot{a}$ as

$$\dot{a} = \frac{1}{\sqrt{2\hbar Z_r}} \dot{\phi} + i \sqrt{\frac{Z_r}{2\hbar}} \dot{Q} \quad (2.21a)$$

$$= \frac{1}{\sqrt{2\hbar Z_r}} \frac{\partial H}{\partial Q} - i \sqrt{\frac{Z_r}{2\hbar}} \frac{\partial H}{\partial \Phi} \quad (2.21b)$$

$$= -\frac{i}{\hbar} \left[\frac{\partial H}{\partial \Phi} \sqrt{\frac{\hbar Z_r}{2}} + \frac{\partial H}{\partial Q} i \sqrt{\frac{\hbar}{2Z_r}} \right] \quad (2.21c)$$

$$= -\frac{i}{\hbar} \left[\frac{\partial H}{\partial \Phi} \frac{\partial \Phi}{\partial a^*} + \frac{\partial H}{\partial Q} \frac{\partial Q}{\partial a^*} \right] \quad (2.21d)$$

$$= -\frac{i}{\hbar} \frac{\partial H}{\partial a^*} \quad (2.21e)$$

We can do the same for $\dot{a}^*$ to show that

$$\dot{a}^* = \frac{i}{\hbar} \frac{\partial H}{\partial a} \quad (2.22)$$

For the harmonic oscillator hamiltonian above, these equations of motion become

$$\dot{a} = -\frac{i}{\hbar} \frac{\partial H}{\partial a^*} = -i\omega_0 a \quad (2.23)$$

$$\dot{a}^* = \frac{i}{\hbar} \frac{\partial H}{\partial a} = i\omega_0 a^* \quad (2.24)$$

and we see that as promised, we now have two decoupled first order differential equations.

At this point it is worth emphasizing that although we have introduced some factors of $\hbar$, this Hamiltonian is entirely classical – we have not yet promoted Φ and Q to operators, nor have we enforced a non-trivial commutation relation between them. As a result, there is no zero-point energy or vacuum fluctuations in the ground state and both ϕ and Q are simultaneously well-defined, with classical equations of motion that can be derived from Hamilton’s equations. Thus, while the Lagrangian and Hamiltonian methods of analyzing a circuit are most commonly seen in the circuit QED literature, it is equally valid for analyzing classical circuits.

Another point worth mentioning is that while it is common to choose the node flux ϕ as the “position” variable, ϕ is not privileged here and it is equally valid to select Q as the “position” variable in this context. To see that this is true, we first express the inductive and capacitive energies in terms of Q and $\dot{Q}$:

$$E_L = \int IV dt = L \int \dot{Q} \ddot{Q} dt = \frac{1}{2} L \dot{Q}^2 \quad (2.25)$$

$$E_C = \int IV dt = \frac{1}{C} \int \dot{Q} Q dt = \frac{Q^2}{2C} \quad (2.26)$$




Figure 2.2: **Floating LC Oscillator.** (a) The circuit for a floating LC oscillator which consists of two nodes with an inductance L and capacitance C between them. The two nodes are connected to ground via capacitances C_g . (b) A change of basis defined by U allows us to find a basis where the two degrees of freedom are decoupled.

Here we see that the role of E_L and E_C have been swapped and E_L now looks like a kinetic energy where Q is the position, $I = \dot{Q}$ is the velocity, and the effective mass is L . The circuit Lagrangian for the system is equivalently expressed as

$$\mathcal{L} = E_L - E_C = \frac{1}{2}L\dot{Q}^2 - \frac{1}{2C}Q^2 \quad (2.27)$$

and the conjugate momentum to Q is now

$$\frac{\partial \mathcal{L}}{\partial \dot{Q}} = L\dot{Q} = LI = \phi \quad (2.28)$$

It is straightforward to check from here that this Lagrangian yields an equivalent equation of motion for Q , and we recover the same dynamics.

We have shown above that E_L and E_C can equivalently be chosen to represent either the kinetic or potential energies in a Harmonic oscillator. So why is ϕ most commonly chosen to be the “position” variable so that E_L represents the potential energy in the circuit QED literature? The answer lies in the fact that the primary way to introduce a dissipationless nonlinearity is with a Josephson junction that acts as a nonlinear inductance. Since it is more straightforward and common to consider a non-quadratic potential, it is preferential to make ϕ our “position” variable.

Floating LC Oscillator

Now that we have introduced the classical circuit Lagrangian and Hamiltonian, we analyze a more complex single-mode circuit with multiple nodes. In the circuit above, we see that

there are two nodes, with fluxes ϕ_1, ϕ_2 and charges Q_1, Q_2 . The capacitive energy is given by

$$E_C = \frac{1}{2}C_g\dot{\phi}_1^2 + \frac{1}{2}C_g\dot{\phi}_2^2 + \frac{1}{2}C(\dot{\phi}_1 - \dot{\phi}_2)^2 \quad (2.29)$$

while the inductive energy is given by

$$E_L = \frac{1}{2L}(\phi_1 - \phi_2)^2 \quad (2.30)$$

It is often easier to write these energies in matrix form, and we can express the Lagrangian with respect to the capacitance matrix $\mathbf{C}$ and inverse inductance matrix $\mathbf{L}^{-1}$ as

$$\mathcal{L} = \frac{1}{2}\dot{\Phi}^T \mathbf{C} \dot{\Phi} - \frac{1}{2}\Phi \mathbf{L}^{-1} \Phi \quad (2.31)$$

where $\Phi^T = (\phi_1 \ \phi_2)$ and

$$\mathbf{C} = \begin{pmatrix} C_g + C & -C \\ -C & C_g + C \end{pmatrix} \quad \mathbf{L}^{-1} = \begin{pmatrix} \frac{1}{L} & -\frac{1}{L} \\ -\frac{1}{L} & \frac{1}{L} \end{pmatrix} \quad (2.32)$$

If we inspect the circuit, we find that although it consists of two loops, there is only one non-trivial loop that contains both an inductance and a capacitance. As a result, there is only one non-trivial degree of freedom. To make this more clear, we consider changing variables to a symmetric and anti-symmetric combination of the node variables

$$\phi_+ = \phi_1 + \phi_2 \quad \phi_- = \phi_1 - \phi_2 \quad (2.33)$$

and the relationship between the two bases can be written in matrix form

$$\underbrace{\begin{pmatrix} \phi_+ \\ \phi_- \end{pmatrix}}_{\Phi_{\pm}} = \underbrace{\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}}_{\mathbf{U}} \begin{pmatrix} \phi_1 \\ \phi_2 \end{pmatrix} \quad (2.34)$$

Applying this change of variables to the circuit Lagrangian yields


$$\mathcal{L} = \frac{1}{2}\dot{\Phi}_{\pm}^T \tilde{\mathbf{C}} \dot{\Phi}_{\pm} - \frac{1}{2}\Phi_{\pm} \tilde{\mathbf{L}}^{-1} \Phi_{\pm} \quad (2.35)$$

where the capacitance and inductance matrices in this basis are (remembering that $\mathbf{U}^{-1}$ is also symmetric since $\mathbf{U}$ is symmetric):

$$\tilde{\mathbf{C}} = \mathbf{U}^{-1} \mathbf{C} \mathbf{U}^{-1} = \begin{pmatrix} \frac{1}{2}C_g & 0 \\ 0 & C + \frac{1}{2}C_g \end{pmatrix} \quad \tilde{\mathbf{L}}^{-1} = \mathbf{U}^{-1} \mathbf{L}^{-1} \mathbf{U}^{-1} = \begin{pmatrix} 0 & 0 \\ 0 & \frac{1}{L} \end{pmatrix} \quad (2.36)$$

In this new basis, we see that ϕ_+ does not appear in the inductive term of the Lagrangian so the Euler-Lagrange equation implies that

$$\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial \dot{\phi}_+} \right) = \frac{\partial \mathcal{L}}{\partial \phi_+} = 0 \quad (2.37)$$

Figure 2.3: **Capacitively coupled LC oscillators.**

so that $Q_+ = \frac{\partial \mathcal{L}}{\partial \phi_+}$ is a constant. In other words, the total charge Q_+ is conserved and the corresponding flux ϕ_+ is often referred to as a cyclic coordinate or a “free” variable [36], and we can set $Q_+ = 0$ without loss of generality. We also notice that ϕ_+ and ϕ_- are now decoupled, since there are no terms in the Lagrangian that contain both ϕ_+ and ϕ_- ($\tilde{\mathbf{L}}$ and $\tilde{\mathbf{C}}$ are both diagonal), and this allows us to analyze the remaining degree of freedom ϕ_- in isolation. From the Lagrangian, we see that the ϕ_- terms in the Lagrangian are equivalent to the previously analyzed LC oscillator, with just a modified effective capacitance $C_- = C + \frac{1}{2}C_g$. We note that this is the same effective capacitance across the two nodes that we would have obtained by adding the parallel and series capacitors.

This circuit is often referred to as a “floating” LC oscillator in contrast to the previously analyzed “grounded” LC oscillator. While the two are often treated interchangeably in the literature, they are only exactly equivalent in isolation, and there are subtle differences involved when coupling these two circuits to external circuitry. These differences will be discussed in more depth in the subsequent sections.

2.3 Coupling

The circuits analyzed so far have been single-mode circuits with a single degree of freedom. In practice, most circuits that we deal with consist of multiple degrees of freedom, and when their equations of motion are interdependent, these degrees of freedom are said to be “coupled”. One of the simplest such circuits consists of two LC oscillators coupled via a capacitance C_c . To analyze such a circuit, it is often useful to express the relevant equations of motion in matrix form. The (classical) Hamiltonian for this system is given by

$$H = \frac{1}{2} \mathbf{Q}^T \mathbf{C}^{-1} \mathbf{Q} + \frac{1}{2} \mathbf{\Phi}^T \mathbf{L}^{-1} \mathbf{\Phi} \quad (2.38)$$

In the basis defined by

$$\mathbf{\Phi} = \begin{pmatrix} \phi_1 \\ \phi_2 \end{pmatrix} \quad \mathbf{Q} = \begin{pmatrix} Q_1 \\ Q_2 \end{pmatrix} \quad (2.39)$$

where ϕ_i, Q_i are the node fluxes and charges, the capacitance and inductance matrices are given by

$$\mathbf{C} = \begin{pmatrix} C_1 + C_c & -C_c \\ -C_c & C_2 + C_c \end{pmatrix} \quad \mathbf{L} = \begin{pmatrix} L_1 & 0 \\ 0 & L_2 \end{pmatrix} \quad (2.40)$$

From the Hamiltonian, we can determine the equations of motion

$$\dot{\mathbf{Q}} = \mathbf{C}\dot{\mathbf{\Phi}} \quad \mathbf{\Phi} = -\mathbf{L}\dot{\mathbf{Q}} \quad (2.41)$$

where we choose $\mathbf{\Phi}$ to be the “position” variable. We recall again that $V = \dot{\phi}$ and $I = \dot{Q}$.

In this basis, the diagonal terms of the Hamiltonian describe the “self” energy terms, while the off-diagonal terms describe how energy is exchanged or “coupled” between the two degrees of freedom. If we isolate the diagonal terms, we see that the coupling capacitance C_c renormalizes the effective mode capacitance

$$\frac{1}{\tilde{C}_i} = \mathbf{C}_{ii}^{-1} = \frac{C_j + C_c}{C_1 C_2 + C_1 C_c + C_2 C_c} \approx \frac{1}{C_i + C_c} \quad (2.42)$$

where $j = i + 1 \pmod{2}$. Then, we can define an effective renormalized mode frequency and impedance

$$\tilde{\omega}_i = \sqrt{\mathbf{L}_{ii}^{-1} \mathbf{C}_{ii}^{-1}} \quad \tilde{Z}_i = \sqrt{\frac{\mathbf{C}_{ii}^{-1}}{\mathbf{L}_{ii}^{-1}}} \quad (2.43)$$

To better understand the energy-exchange terms, we move into the mode amplitude basis by defining the mode amplitudes

$$a_i = \frac{1}{\sqrt{2\hbar\tilde{Z}_i}}\phi_i + i\sqrt{\frac{\tilde{Z}_i}{2\hbar}}Q_i \quad (2.44)$$

Expressing the Hamiltonian in Eq. 2.38 in terms of a_i and a_i^* , we obtain the expression

$$H/\hbar = \tilde{\omega}_1 a_1^* a_1 + \tilde{\omega}_2 a_2^* a_2 - g_c (a_1^* - a_1)(a_2^* - a_2) \quad (2.45)$$

where the coupling energy g_c is given by

$$g_c = \frac{1}{2} \sqrt{\frac{1}{\tilde{Z}_1 \tilde{Z}_2}} \mathbf{C}_{12}^{-1} = \frac{1}{2} \frac{C_c}{\sqrt{(C_1 + C_c)(C_2 + C_c)}} \sqrt{\tilde{\omega}_1 \tilde{\omega}_2} \quad (2.46)$$

The equations of motion for the system are then given by Eq. 2.21 for each of the a_i , and we now obtain a set of coupled differential equations:

$$\frac{d}{dt} \begin{pmatrix} a_1 \\ a_2 \\ a_1^* \\ a_2^* \end{pmatrix} = -i \underbrace{\begin{pmatrix} \tilde{\omega}_1 & g_c & 0 & -g_c \\ g_c & \tilde{\omega}_2 & -g_c & 0 \\ 0 & g_c & -\tilde{\omega}_1 & -g_c \\ g_c & 0 & -g_c & -\tilde{\omega}_2 \end{pmatrix}}_{\mathbf{M}} \begin{pmatrix} a_1 \\ a_2 \\ a_1^* \\ a_2^* \end{pmatrix} \quad (2.47)$$

We can solve this differential equation by diagonalizing $\mathbf{M}$ which “decouples” the differential equations. We will solve for the specific case where $\tilde{\omega}_1 = \tilde{\omega}_2$, but first we will make a short aside to discuss the important concept of *coupling efficiency*.

Coupling Efficiency

It is often useful to express the coupling as a fraction of the geometric mean of the mode frequencies $g_c = \frac{1}{2}K\sqrt{\tilde{\omega}_1\tilde{\omega}_2}$, in which case we define a dimensionless *coupling efficiency* [148],

$$K = \frac{C_c}{\sqrt{(C_1 + C_c)(C_2 + C_c)}} \quad (2.48)$$

that quantifies how much of the total capacitance in each LC oscillator is shared between the two. The form of the coupling efficiency depends on the specific coupling geometry. When the two modes have the same impedance and frequency so that $L \equiv L_1 = L_2$ and $C \equiv C_1 = C_2$, we can express Eq. 2.42 as

$$\mathbf{C}_{11}^{-1} = \mathbf{C}_{22}^{-1} = \frac{1}{C} \frac{1 + \eta}{1 + 2\eta} \approx \frac{1}{C} \frac{1}{1 + \eta - \eta^2} \quad (2.49)$$

and the renormalized mode frequency and impedance as

$$\tilde{\omega} = \frac{1}{\sqrt{LC}} \sqrt{\frac{1 + \eta}{1 + 2\eta}} = \omega_0 \sqrt{\frac{1 + \eta}{1 + 2\eta}} \quad \tilde{Z} = \sqrt{\frac{L}{C}} \sqrt{\frac{1 + \eta}{1 + 2\eta}} = Z_r \sqrt{\frac{1 + \eta}{1 + 2\eta}} \quad (2.50)$$

Here, we define $\eta = \frac{C_c}{C}$. In this case, the coupling efficiency becomes

$$K = \frac{C_c}{C + C_c} = \frac{\eta}{1 + \eta} \approx \eta - \eta^2 + \eta^3 \quad (2.51)$$

We see that the coupling efficiency asymptotes to 1 as $C_c \rightarrow \infty$, which points to the fact that the shared energy between the two modes cannot exceed the self energy. In the limit of small η , we can approximate $K \approx \eta$, which is equivalent to saying that the coupling capacitance does not renormalize the total mode capacitance.




Figure 2.4: **Normal mode frequencies.** (a) The normalized mode splitting as a function of coupling efficiency for the full equations of motion, and the equations of motion under RWA. The relative error between the RWA mode splitting and the full mode splitting is shown in green. (b) The mode frequencies $\omega_{\pm}$ for two values of g_c , when $\tilde{\omega}_2$ is swept across $\tilde{\omega}_1$. This leads to an avoided crossing where the minimum gap scales with the coupling strength g_c . $\tilde{\omega}_1$ is fixed at 5 GHz.

Normal Modes

We now return to the coupled differential equations in Eq. 2.47, for the specific case where $\tilde{\omega} \equiv \tilde{\omega}_1 = \tilde{\omega}_2$. In this case, $g_c = \frac{1}{2}K\tilde{\omega}$. The eigenvalues for the matrix $\mathbf{M}$ in Eq. 2.47 are $\pm\omega_{\pm}$ and $\pm\omega_{\mp}$ where

$$\omega_{\pm} = \sqrt{\tilde{\omega}^2 \pm 2g_c\tilde{\omega}} = \tilde{\omega}\sqrt{1 \pm \frac{2g_c}{\tilde{\omega}}} = \tilde{\omega}\sqrt{1 \pm K} \quad (2.52)$$

This leads to a mode splitting given by

$$\omega_{+} - \omega_{-} = \tilde{\omega} \left[\sqrt{1+K} - \sqrt{1-K} \right] \quad (2.53a)$$

$$\approx \tilde{\omega} \left[\left(1 + \frac{1}{2}K - \frac{1}{8}K^2 + \frac{1}{16}K^3 \right) - \left(1 - \frac{1}{2}K - \frac{1}{8}K^2 - \frac{1}{16}K^3 \right) \right] \quad (2.53b)$$

$$= K\tilde{\omega} + \frac{1}{8}K^3\tilde{\omega} \quad (2.53c)$$

$$= 2g_c \left(1 + \frac{1}{8}K^2 \right) \quad (2.53d)$$

and we see that the frequency difference between the normal modes is approximately $2g_c$ to second order in K . Here, we see that while it is commonly stated that the mode splitting is

$2g_c$, this is in fact only an approximation. This approximation is equivalent to making the rotating wave approximation (RWA), where the terms proportional to $a_1 a_2$ and $a_1^* a_2^*$ in the coupling Hamiltonian are dropped. Then, the Hamiltonian in Eq. 2.45 becomes

$$H/\hbar = \tilde{\omega}_1 a_1^* a_1 + \tilde{\omega}_2 a_2^* a_2 + g_c(a_1 a_2^* + a_1^* a_2) \quad (2.54)$$

which amounts to dropping the terms on along the anti-diagonal of $\mathbf{M}$ [174]. In this case, $\mathbf{M}$ becomes

$$\mathbf{M} = \begin{pmatrix} \tilde{\omega}_1 & g_c & 0 & 0 \\ g_c & \tilde{\omega}_2 & 0 & 0 \\ 0 & 0 & -\tilde{\omega}_1 & -g_c \\ 0 & 0 & -g_c & -\tilde{\omega}_2 \end{pmatrix} \quad (2.55)$$

and the eigenvalues of $\mathbf{M}$ now have a simple form for all $\tilde{\omega}_1, \tilde{\omega}_2$. They are given by

$$\pm \omega_{\pm}^{(\text{RWA})} = \pm \left[\frac{1}{2} \Sigma \pm \sqrt{g_c^2 + \left(\frac{\Delta}{2} \right)^2} \right] \quad (2.56)$$

where $\Sigma = \tilde{\omega}_1 + \tilde{\omega}_2$ and $\Delta = \tilde{\omega}_1 - \tilde{\omega}_2$ are the sum and difference frequencies. Now, the mode splitting is given by

$$\omega_+^{(\text{RWA})} - \omega_-^{(\text{RWA})} = 2 \sqrt{g_c^2 + \left(\frac{\Delta}{2} \right)^2} \quad (2.57)$$

so that the mode splitting is *exactly* $2g_c$ when $\Delta = 0$. In Fig. 2.4, we plot the mode splitting as a function of K with for both the exact solution and the RWA solution, as well as the relative error between the two solutions. We also show the mode frequencies ω_+ and ω_- as we sweep $\tilde{\omega}_2$, while keeping $\tilde{\omega}_1$ fixed. This illustrates the *avoided crossing* that occurs between the eigenmodes when there exists a finite coupling. Note that the avoided crossing is purely a classical effect arising from the mode hybridization.

Alternatively, we could have solved for the system in the Φ basis by solving the second order differential equation that arises from Eq. 2.41:

$$\ddot{\Phi} = -(\mathbf{LC})^{-1} \Phi \quad (2.58)$$

In the case where $L \equiv L_1 = L_2$ and $C \equiv C_1 = C_2$, diagonalizing $(\mathbf{LC})^{-1}$ also results in the same eigenbasis defined by symmetric and antisymmetric linear combinations of Φ_1, Φ_2 . The eigenvalues are

$$\omega_+ = \frac{1}{\sqrt{LC}} \quad \omega_- = \frac{1}{\sqrt{L(C + 2C_c)}} \quad (2.59)$$

Here we see that ω_+ is constant as a function of C_c , while in Eq. 2.52, ω_+ also has a dependence on K . This is because as C_c increases, the renormalized frequency $\tilde{\omega}$ also decreases at a rate equal to half the mode splitting. In terms of the renormalized frequencies, we find

$$\omega_+ = \tilde{\omega} \sqrt{\frac{1+2\eta}{1+\eta}} = \tilde{\omega} \sqrt{1+K} \quad \omega_- = \tilde{\omega} \sqrt{\frac{1}{1+\eta}} = \tilde{\omega} \sqrt{1-K} \quad (2.60)$$

where we recall that $K = \frac{\eta}{1+\eta}$, and we see that these solutions are equivalent.

Coupling Geometries

As mentioned in the previous section, the coupling efficiency K depends on the precise coupling geometry and whether the coupling is mediated by an inductance or capacitance. In this section, we will discuss several different coupling geometries, and show how the coupling efficiency changes.

In contrast to the circuit shown in Fig. 2.3, where the capacitors form a Π -circuit, we can also consider an alternative two mode-circuit where the capacitors form a T -circuit (see Fig. 2.5(b)). In this case, there is an additional node in the middle, to which we assign a node flux ϕ_c . The capacitance and inverse inductance matrix for this circuit can be written down as

$$\mathbf{C} = \begin{pmatrix} C_1 & -C_1 & 0 \\ -C_1 & C_1 + C_c + C_2 & -C_2 \\ 0 & -C_2 & C_2 \end{pmatrix} \quad \mathbf{L}^{-1} = \begin{pmatrix} \frac{1}{L_1} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \frac{1}{L_2} \end{pmatrix} \quad (2.61)$$

The potential energy term in the Lagrangian is $\frac{1}{2}\mathbf{\Phi}^T\mathbf{L}^{-1}\mathbf{\Phi}$, so we see that once again we have free variable ϕ_c that does not appear in the Lagrangian. To eliminate the free variable, we apply the constraint from the Euler-Lagrange equations, which tells us that

$$\frac{\partial \mathcal{L}}{\partial \dot{\phi}_c} = (C_1 + C_c + C_2)\dot{\phi}_c - C_1\dot{\phi}_1 - C_2\dot{\phi}_2 = Q_c \quad (2.62)$$

where Q_c is the conserved conjugate charge to ϕ_c . From here, we can express $\dot{\phi}_c$ in terms of the remaining node variables to eliminate it from the Lagrangian:

$$\dot{\phi}_c = \frac{1}{C_\Sigma} (C_1\dot{\phi}_1 + C_2\dot{\phi}_2 + Q_c) \quad (2.63)$$

For brevity, we will just show the resulting Lagrangian⁴, which looks like

$$\mathcal{L} = \frac{1}{2}\dot{\tilde{\mathbf{\Phi}}}^T \tilde{\mathbf{C}} \dot{\tilde{\mathbf{\Phi}}} + \frac{Q_c^2}{2C_\Sigma} - \frac{1}{2}\tilde{\mathbf{\Phi}}^T \tilde{\mathbf{L}}^{-1} \tilde{\mathbf{\Phi}} \quad (2.64)$$


with the new capacitance and inverse inductance matrices

$$\tilde{\mathbf{C}} = \frac{1}{C_\Sigma} \begin{pmatrix} C_1C_c + C_1C_2 & -C_1C_2 \\ -C_1C_2 & C_2C_c + C_1C_2 \end{pmatrix} \quad \tilde{\mathbf{L}}^{-1} = \begin{pmatrix} \frac{1}{L_1} & 0 \\ 0 & \frac{1}{L_2} \end{pmatrix} \quad (2.65)$$

where the center node flux ϕ_c has been eliminated ($\tilde{\mathbf{\Phi}}^T = (\phi_1 \ \phi_2)$). If we compare this to our previous result for the Π -circuit, we see that the Lagrangian is equivalent to a Π -circuit with the modified capacitances

$$C'_1 = \frac{C_1C_c}{C_\Sigma} \quad C'_2 = \frac{C_2C_c}{C_\Sigma} \quad C'_c = \frac{C_1C_2}{C_\Sigma} \quad (2.66)$$

⁴See Appendix A.2 for the full derivation.

Figure 2.5: **Coupling geometries.**

and the astute reader may notice that we have just derived the $Y - \Delta$ transform for a capacitive network. From here, we can use our previous results to write the renormalized mode frequencies and impedances as

$$\tilde{\omega}_i = \sqrt{\frac{C_\Sigma}{L_i C_i C_c}} \quad \tilde{Z}_i = \sqrt{\frac{L_i C_\Sigma}{C_i C_c}} \quad (2.67)$$

and the coupling efficiency is

$$K_{c,T} = \frac{C_1 C_2}{\sqrt{(C_1 C_c + C_1 C_2)(C_2 C_c + C_1 C_2)}} = \sqrt{\frac{C_1 C_2}{(C_1 + C_c)(C_2 + C_c)}} \quad (2.68)$$

If we compare the two coupling configurations, we find that the role of C_c is flipped. In the Π -circuit, the coupling capacitance serves as a link between the two resonant modes, so a larger capacitance (lower impedance) leads to stronger coupling. On the other hand, in the T -circuit, the coupling capacitance acts as a voltage divider to ground, so a larger capacitance (lower impedance) becomes more like a short to ground, leading to a weaker coupling.

In comparison to capacitive coupling, the inductive coupling case is analyzed less frequently in the literature. Furthermore, there are some more subtleties involved in implementing an inductive coupling in practice, so we will work through several examples here. We start by considering what looks like a direct analogy to the capacitive coupling case,

where we have an inductor L_c between the two nodes of ϕ_1, ϕ_2 , as shown in Fig. 2.5(c). Intuitively, we expect the two modes to have stronger coupling when L_c is small, since the impedance of the connection decreases with smaller L_c .

In this case, the capacitance and inductance matrices are given by

$$\mathbf{C} = \begin{pmatrix} C_1 & 0 \\ 0 & C_2 \end{pmatrix} \quad \mathbf{L}^{-1} = \begin{pmatrix} \frac{1}{L_1} + \frac{1}{L_c} & -\frac{1}{L_c} \\ -\frac{1}{L_c} & \frac{1}{L_2} + \frac{1}{L_c} \end{pmatrix} \quad (2.69)$$

From here it is relatively straightforward to read off the renormalized frequencies and impedances

$$\tilde{\omega}_i = \frac{1}{\sqrt{\tilde{L}_i C_i}} = \sqrt{\frac{L_i + L_c}{L_i L_c C_i}} \quad \tilde{Z}_i = \sqrt{\frac{\tilde{L}_i}{C_i}} = \sqrt{\frac{L_i L_c}{(L_i + L_c) C_i}} \quad (2.70)$$

Since the coupling term is now inductive, we can express it as

$$g_\ell = -\frac{1}{2} \frac{1}{L_c} \sqrt{\tilde{Z}_1 \tilde{Z}_2} = -\frac{1}{2} \frac{1}{L_c} \sqrt{\tilde{L}_1 \tilde{L}_2} \sqrt{\tilde{\omega}_1 \tilde{\omega}_2} = -\frac{1}{2} K_{\ell, \Pi} \sqrt{\tilde{\omega}_1 \tilde{\omega}_2} \quad (2.71)$$

where the inductive coupling efficiency is now

$$K_{\ell, \Pi} = \frac{1}{L_c} \sqrt{\tilde{L}_1 \tilde{L}_2} = \sqrt{\frac{L_1 L_2}{(L_1 + L_c)(L_2 + L_c)}} \quad (2.72)$$

We make two important observations from the above equations. First, the sign of the coupling g_ℓ is now negative, in contrast to the capacitive coupling g_c . Second, the expression for $K_{\ell, \Pi}$ looks strikingly similar to the expression for $K_{c, T}$, but with the role of the capacitances and inductances flipped. This highlights the duality between inductive and capacitive coupling. From here, we would expect the inductive T -coupling to match the form of the capacitive Π -coupling, which we will proceed to show.

As in the capacitive circuit, we now have an extra node with flux ϕ_c . The capacitance and inductance matrices are

$$\mathbf{C} = \begin{pmatrix} C_1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & C_2 \end{pmatrix} \quad \mathbf{L}^{-1} = \begin{pmatrix} L_1^{-1} & -L_1^{-1} & 0 \\ -L_1^{-1} & L_1^{-1} + L_c^{-1} + L_2^{-1} & -L_2^{-1} \\ 0 & -L_2^{-1} & L_2^{-1} \end{pmatrix} \quad (2.73)$$

and we see that this time it is the kinetic energy E_C that does not depend on $\dot{\phi}_c$, so we refer to ϕ_c as a “frozen” variable. In this case, the Euler-Lagrange equations enforce the condition

$$\frac{\partial \mathcal{L}}{\partial \phi_c} = \frac{d}{dt} \left[\frac{\partial \mathcal{L}}{\partial \dot{\phi}_c} \right] = 0 \quad (2.74)$$

so we can write ϕ_c in terms of the other variables as

$$\phi_c = \frac{L_2 L_c}{L_\Sigma} \phi_1 + \frac{L_1 L_c}{L_\Sigma} \phi_2 \quad (2.75)$$

	Π -coupled	T -coupled
Capacitive	$K_{c,\Pi} = \frac{C_c}{\sqrt{(C_1+C_c)(C_2+C_c)}}$	$K_{c,T} = \sqrt{\frac{C_1 C_2}{(C_1+C_c)(C_2+C_c)}}$
Inductive	$K_{\ell,\Pi} = \sqrt{\frac{L_1 L_2}{(L_1+L_c)(L_2+L_c)}}$	$K_{\ell,T} = \frac{L_c}{\sqrt{(L_1+L_c)(L_2+L_c)}}$

 Table 2.1: **Coupling efficiencies for Π - and T -circuits.**

where $L_\Sigma = L_1 + L_2 + L_c$. We again make the relevant substitutions, from which we can derive the effective capacitance and inverse inductance matrix

$$\tilde{\mathbf{C}} = \begin{pmatrix} C_1 & 0 \\ 0 & C_2 \end{pmatrix} \quad \tilde{\mathbf{L}}^{-1} = \frac{1}{L_1 L_2 + L_1 L_c + L_2 L_c} \begin{pmatrix} L_2 + L_c & -L_c \\ -L_c & L_1 + L_c \end{pmatrix} \quad (2.76)$$

Here, we are working again in the reduced basis defined by $\Phi^T = (\phi_1 \ \phi_2)$. Then, we can check that the inductance matrix is

$$\mathbf{L} = \begin{pmatrix} L_1 + L_c & L_c \\ L_c & L_2 + L_c \end{pmatrix} \quad (2.77)$$

which, as expected, is the inductive analog of the capacitance matrix for the T-circuit. One difference, however, is that the off-diagonal term again has the opposite sign compared to its capacitive counterpart. We can now read off the renormalized mode capacitances and inductances from Eq. 2.76 to get the coupling rate

$$g_\ell = -\frac{1}{2} \frac{L_c}{L_1 L_2 + L_1 L_c + L_2 L_c} \sqrt{\tilde{L}_1 \tilde{L}_2} \sqrt{\tilde{\omega}_1 \tilde{\omega}_2} = -\frac{1}{2} \frac{L_c}{\sqrt{(L_1 + L_c)(L_2 + L_c)}} \sqrt{\tilde{\omega}_1 \tilde{\omega}_2} \quad (2.78)$$

Finally, we note that we can again compare the inductance matrices for the inductive Π - and T -circuits to derive the equivalent Y - Δ transform for a set of 3 inductors.

We summarize the coupling efficiencies in Table 2.1 for the four coupling geometries shown in Fig. 2.5.

Mutual Inductance

The circuits analyzed thus far have involved direct inductive connections between two nodes; however, circuits can also be coupled via a mutual inductance. The two port inductance matrix for a mutual inductance is defined by $\dot{\Phi} = \mathbf{L}\dot{\mathbf{I}}$, where

$$\mathbf{L} = \begin{pmatrix} L_1 & M \\ M & L_2 \end{pmatrix} \quad (2.79)$$




Figure 2.6: **Mutual inductance coupling.** A mutual inductance between two inductors (a) can equivalently be represented as an inductive T -coupler (b).

Here we see that the diagonal terms L_1, L_2 define the voltage induced by a current at their corresponding nodes, while the off-diagonal terms M define the voltage induced at the opposite node due to a change in the current at the corresponding node. Unlike the previous cases involving a direct inductance, M can be either positive or negative depending on the direction of the induced current in the second loop relative to the current in the first loop, but the magnitude of M is constrained by $0 \leq |M| \leq 1$.

The inverse inductance matrix is now given by

$$\mathbf{L}^{-1} = \frac{1}{L_1 L_2 - M^2} \begin{pmatrix} L_2 & -M \\ -M & L_1 \end{pmatrix} \quad (2.80)$$

from which we can read off the renormalized mode inductance as

$$\frac{1}{\tilde{L}_i} = \mathbf{L}_{ii}^{-1} = \frac{L_j}{L_1 L_2 - M^2} \quad (2.81)$$

This again leads to a renormalized mode frequency and impedance as given by Eq. 2.43, and we can then write the coupling energy as

$$g_M = -\frac{1}{2} \frac{M}{\sqrt{L_1 L_2}} \sqrt{\tilde{\omega}_1 \tilde{\omega}_2} = -\frac{1}{2} K \sqrt{\tilde{\omega}_1 \tilde{\omega}_2} \quad (2.82)$$

where $K = \frac{M}{\sqrt{L_1 L_2}}$ is the same as the commonly defined “coupling coefficient” for a transformer. At this point, we may also recognize that the inductance matrix for the mutual inductance coupling looks exactly like the inductance matrix for the T -circuit, where we take $L_c \rightarrow M$ and $L_i \rightarrow L_i - M$. This yields another common circuit identity for replacing a mutual inductance with a directly shared inductance.

Floating vs. Grounded Oscillators

Earlier, we introduced the floating LC oscillator, and showed how it is equivalent to the grounded LC oscillator with a renormalized capacitance $C \rightarrow C + \frac{1}{2}C_g$ in isolation. However,




Figure 2.7: **Coupling efficiencies for grounded and floating LC oscillators.** (a) Circuit diagram for a floating LC oscillator coupled to a grounded LC oscillator. (b) Circuit diagram for two capacitively coupled floating LC oscillators. (c) Coupling efficiencies for the different oscillator configurations: grounded-floating (G-F), floating-floating (F-F) and grounded-grounded (G-G). The lower plot shows K_{gf} as a function of the ground capacitance ratio $1 - C_g/(2C)$. The dotted gray line shows the coupling efficiency K_{ff} as a comparison.

when we couple a floating LC oscillator to another circuit element, the difference in circuit structure can have a meaningful effect on the coupling. To illustrate this, we can compare the effective capacitive coupling efficiencies between a grounded and floating LC oscillator K_{gf} , and two floating LC oscillators K_{ff} . The full derivation of K_{gf} and K_{ff} can be found in Appendix A.2, but we will quote the results here. For a grounded and floating LC oscillator, we have

$$K_{gf} = \frac{\gamma C_-}{(C_a + \frac{1}{2}C_g + \frac{1}{2}\gamma C_+)(C_b + 2\gamma C_+)} \quad (2.83)$$

where we define $\gamma = \frac{C_g}{C_g + C_+}$ and $C_{\pm} = (C_1 \pm C_2)/2$. On the other hand, for two floating oscillators, we have

$$K_{ff} = \frac{1}{2} \frac{C_{\Delta} - \frac{C_{\alpha}^2 - C_{\beta}^2}{C_{\Sigma} + \frac{1}{2}C_g}}{\sqrt{\left(C'_a + \frac{1}{2}C_{\Sigma} - \frac{1}{2}\frac{(C_{\alpha} + C_{\beta})^2}{C_{\Sigma} + \frac{1}{2}C_g}\right)\left(C'_b + \frac{1}{2}C_{\Sigma} - \frac{1}{2}\frac{(C_{\alpha} - C_{\beta})^2}{C_{\Sigma} + \frac{1}{2}C_g}\right)}} \quad (2.84)$$

where we define $C_{\Sigma} = (C_{11} + C_{22} + C_{12} + C_{21})/2$, $C_{\Delta} = (C_{11} + C_{22} - C_{12} - C_{21})/2$, $C_{\alpha} = (C_{11} - C_{22})/2$, and $C_{\beta} = (C_{12} - C_{21})/2$. The renormalized self-capacitances are $C'_a = C_a + \frac{1}{2}C_g$

and $C'_b = C_b + \frac{1}{2}C_g$.

To understand these expressions better, we compare them to the coupling efficiency we derived earlier for the case of two grounded LC oscillators in a Π -coupled geometry (see Eq. 2.48)

$$K_{gg} = \frac{C_c}{\sqrt{(C_a + C_c)(C_b + C_c)}} \quad (2.85)$$

Here, we notice that the K_{ff} shares a very similar form to K_{gg} , up to some additional constants, while K_{gf} depends strongly on the ratio $\gamma \propto C_g$. This is because the coupling between the differential mode of the floating oscillator and the common mode of the grounded oscillator is mediated via the ground capacitance C_g .

In Fig. 2.7(c), we plot the coupling efficiencies for these three configurations as a function of the coupling capacitance C_c . Here, we set $C_1 = C_c$ and $C_2 = 0$ for the grounded-floating circuit, and $C_{11} = C_{22} = C_c$ and $C_{12} = C_{21} = 0$ for the floating-floating circuit. For the floating oscillators, we keep the total mode capacitance $C \equiv C_i + \frac{1}{2}C_g$ constant, but we can vary the ratio between the ground capacitance C_g and differential capacitance C_i . We find that for small $\eta = C_g/C$, the grounded-floating and floating-floating circuits yield similar coupling efficiencies that approach $K_{gg}/2$. However for large η , K_{gf} saturates, while K_{ff} approaches K_{gg} . In the lower plot, we also show K_{gf} as a function of $C_i/C = 1 - \frac{C_g}{2C}$. Here, we see that the coupling efficiency K_{gf} depends strongly on the relative contribution of the ground capacitance to the total mode capacitance. This shows that we cannot always assume that the grounded and floating oscillators are effectively interchangeable, and that we must be careful when analyzing how the two configurations couple to external circuitry.

2.4 Voltage and Current Sources

So far we have analyzed the behavior of circuits in isolation, but in order to manipulate these circuits, we must connect them to voltage or current sources. In this section, we will discuss how to analyze such circuits.

We first consider driving a series LC oscillator with an ideal voltage source, as shown in Fig. 2.8(a). The Lagrangian for the system is

$$L = \frac{1}{2}C(\dot{\phi} - \dot{\phi}_d)^2 - \frac{1}{2L}\phi^2 \quad (2.86)$$

where $\dot{\phi}_d = V_d(t)$ is constrained by the voltage source. The conjugate charge to ϕ is $Q = \frac{\partial L}{\partial \dot{\phi}} = C(\dot{\phi} - \dot{\phi}_d)$, so the Hamiltonian is

$$H = Q\dot{\phi} - L = \frac{1}{2C}Q^2 + \frac{1}{2L}\phi^2 + QV_d \quad (2.87)$$

and the resulting equations of motion are

$$\dot{\phi} = \frac{\partial H}{\partial Q} = \frac{1}{C}Q + V_d \quad \dot{Q} = -\frac{\partial H}{\partial \phi} = -\frac{1}{L}\phi \quad (2.88)$$




Figure 2.8: **Voltage and current driving of LC oscillators** (a) A voltage source in parallel with a series LC circuit. (b) The response amplitude of ϕ and Q as a function of the drive frequency ω_d . On resonance ($\omega_d = \omega_r$), both ϕ and Q diverge. At DC, the charge becomes an offset charge $Q = CV_0$, where V_0 is the voltage amplitude. (c) A current source driving a parallel LC circuit. (d) The response amplitude of ϕ and Q . Here, we see that the behavior of ϕ and Q are flipped compared to the voltage driving case in (b). The DC response of the flux becomes $\phi = LI_0$, where I_0 is the current amplitude.

From here, we can obtain two second order differential equations for ϕ and Q :

$$\ddot{\phi} = -\omega_r^2 \phi + \dot{V}_d \quad \ddot{Q} = -\omega_r^2 Q + \frac{V_d}{L} \quad (2.89)$$

where $\omega_r = \frac{1}{\sqrt{LC}}$ is the resonance frequency. The second equation can be interpreted as Kirchhoff's voltage law if we recognize that $V_\ell = L\ddot{Q} = L\dot{I}$ is the voltage across the inductor, while $V_c = Q/C$ is the voltage across the capacitor, so the equations of motion become $V_\ell + V_c = V_d$. To obtain the time-dependent behavior of $\phi(t)$ and $Q(t)$, we recognize that Eq. 2.89 can be mapped to the equations of motion for a driven harmonic oscillator. For $V_d = V_0 \cos(\omega t)$,

the solutions⁵ are:

$$\phi(t) = -V_0 \frac{\omega}{\omega_r^2 - \omega^2} \sin \omega t \quad (2.90)$$

$$Q(t) = \frac{V_0}{L} \frac{1}{\omega_r^2 - \omega^2} \cos \omega t = \frac{CV_0}{1 - \frac{\omega^2}{\omega_r^2}} \cos \omega t \quad (2.91)$$

In Fig. 2.8(b), we plot the amplitude of ϕ and Q as a function of the drive frequency ω . We see that the amplitude is suppressed away from ω_r and diverges as $\omega \rightarrow \omega_r$. This is the resonance condition. For a DC voltage ($\omega \rightarrow 0$), the equations of motion are exactly the same as the equations for a parallel LC oscillator, except that the charge $Q(t)$ gains a constant offset $Q_0 = CV_0$.

Next we look at the corresponding case where a parallel LC oscillator is driven with an ideal current source, shown in Fig. 2.8(c). To account for the current branch, we must include the potential energy of the branch that includes the current source. This is given by

$$U = \int_{-\infty}^t dt' I(t') V(t') = I(t) \Phi - \int_{-\infty}^t dt' \dot{I}(t') \Phi(t') \quad (2.92)$$

so we see that we must add a term $I_d \Phi$ to the potential energy to represent the potential energy supplied by the current source. Here Φ is the branch flux across the current source, which for the circuit shown in Fig. 2.8(c) is equal to the node flux ϕ . The full Lagrangian is given by

$$\mathcal{L} = \frac{1}{2} C \dot{\phi}^2 - \frac{1}{2L} \phi^2 - I_d \phi \quad (2.93)$$

and the conjugate charge to ϕ is simply $Q = C\dot{\phi}$. The Hamiltonian is

$$H = Q\dot{\phi} - \mathcal{L} = \frac{1}{2C} Q^2 + \frac{1}{2L} \phi^2 + I_d \phi \quad (2.94)$$

and Hamilton's equations give:

$$\dot{\phi} = \frac{\partial H}{\partial Q} = \frac{1}{C} Q \quad \dot{Q} = -\frac{\partial H}{\partial \phi} = -\frac{1}{L} \phi - I_d \quad (2.95)$$

From here we can obtain two second order differential equations for ϕ and Q :

$$\ddot{\phi} = -\omega_r^2 \phi - \frac{1}{C} I_d \quad \ddot{Q} = -\omega_r^2 Q - \dot{I}_d \quad (2.96)$$

This time we multiply the first equation by C , and we identify the terms as $I_c = C\ddot{\phi} = C\dot{V} = \dot{Q}_c$, the current flowing through the capacitor, and $I_\ell = \frac{1}{L}\phi$, the current flowing through the inductor. Thus we can rewrite Eq. 2.96 as $I_c + I_\ell + I_d = 0$, which we recognize as Kirchhoff's

⁵See Appendix A.1 for a derivation of the classical driven and damped harmonic oscillator.




Figure 2.9: **Indirect driving of LC oscillators.** (a) A capacitor C_c is placed in between the voltage source and the LC oscillator, which acts as a voltage divider. (b) A mutual inductance M couples a current source to the LC oscillator. Here, the inductive circuit acts as an effective current divider.

current law. Again, we invoke our result from Appendix A.1 to obtain the inhomogenous solutions for a current drive $I_d = I_0 \cos(\omega t)$:

$$\phi(t) = -\frac{I_0}{C} \frac{1}{\omega_r^2 - \omega^2} \cos \omega t = \frac{LI_0}{1 - \frac{\omega^2}{\omega_r^2}} \cos \omega t \quad (2.97)$$

$$Q(t) = I_0 \frac{\omega}{\omega_r^2 - \omega^2} \sin \omega t \quad (2.98)$$

As expected, the behavior of $\phi(t)$ and $Q(t)$ look similar to the voltage driving case with their behavior flipped; however, we point out that both $\phi(t)$ and $Q(t)$ are $\pi/2$ out of phase for a current drive compared to a voltage drive. While this is simply a negligible reference phase when only a single type of driving is present, this should not be overlooked when considering a system with both voltage and current driving.

While the direct voltage and current driving cases are instructive, this typically results in a coupling to the environment that is too strong for a realistic circuit QED system. In practice, we will usually couple an external voltage (current) drive with a small capacitance (mutual inductance), as shown in Fig. 2.9. In the voltage driving case, we have the following Lagrangian

$$\mathcal{L} = \frac{1}{2} C \dot{\phi}^2 + \frac{1}{2} C_c (\dot{\phi} - \dot{\phi}_d)^2 - \frac{1}{2L} \phi^2 \quad (2.99)$$

where $V_d = \dot{\phi}_d$ is the driving voltage due to the source. Here, the conjugate charge to ϕ is $Q = (C + C_c)\dot{\phi} - C_c V_d$, so the Hamiltonian is

$$H = Q\dot{\phi} - \mathcal{L} = \frac{1}{2(C + C_c)} (Q + C_c V_d)^2 + \frac{1}{2L} \phi^2 - \frac{1}{2} C_c V_d^2 \quad (2.100)$$

so we see that $Q_d = C_c V_d$ acts as an offset charge. From the Hamiltonian, we derive the new equations of motion

$$\dot{\phi} = \frac{\partial H}{\partial Q} = \frac{1}{C + C_c}(Q + C_c V_d) \quad \dot{Q} = -\frac{\partial H}{\partial \phi} = -\frac{1}{L}\phi \quad (2.101)$$

and we see that these are equivalent to the equations of motion in the direct driving case, except with $V_d \rightarrow \frac{C_c}{C+C_c} V_d$. The capacitors act as a voltage divider and renormalize the effective drive seen by the oscillator.

We get a similar effect with a current drive coupled via a small mutual inductance. Here the inductance matrix is given by

$$\mathbf{L} = \begin{pmatrix} L_d & M \\ M & L_1 \end{pmatrix} \Rightarrow \mathbf{L}^{-1} = \frac{1}{L_d L_1 - M^2} \begin{pmatrix} L_1 & -M \\ -M & L_d \end{pmatrix} \quad (2.102)$$

so the Lagrangian is

$$\mathcal{L} = \frac{1}{2} C_1 \dot{\phi}_1^2 - \frac{1}{2} \frac{L_d}{|\mathbf{L}|} \phi_1^2 + \frac{M}{|\mathbf{L}|} \phi_1 \phi_d - \frac{1}{2} \frac{L_1}{|\mathbf{L}|} \phi_d^2 + I_d \phi_d \quad (2.103)$$

We note that the sign of I_d is flipped from the previous example because the convention for the inductance matrix $\mathbf{L}$ is to have both currents pointing into the transformer section. We can use the Euler-Lagrange equations to eliminate ϕ_d since the absence of $\dot{\phi}_d$ in the Lagrangian implies that

$$\frac{\partial \mathcal{L}}{\partial \phi_d} = \frac{1}{|\mathbf{L}|} (M \phi_1 - L_1 \phi_d) + I_d = 0 \quad (2.104)$$

We can also obtain the same constraint $L_1 \phi_d = M \phi_1 + |\mathbf{L}| I_d$ by taking the inductance matrix $\mathbf{L}$, which relates the currents and node fluxes:

$$\begin{aligned} \phi_1 &= L_1 I_1 + M I_d \\ \phi_d &= M I_1 + L_d I_d \end{aligned} \quad (2.105)$$

and eliminating I_1 to solve for ϕ_d in terms of the remaining variables. We then eliminate ϕ_d from the Lagrangian to get

$$\mathcal{L} = \frac{1}{2} C_1 \dot{\phi}_1^2 - \frac{1}{2 L_1} \phi_1^2 + \frac{M}{L_1} \phi_1 I_d + \frac{1}{2} \frac{|\mathbf{L}|}{L_1} I_d^2 \quad (2.106)$$

which looks like Eq. 2.99 up to a constant, where we take $I_d \rightarrow \frac{M}{L_1} I_d$. Thus we see that the mutual inductance acts as a current divider in conjunction with the self inductance L_1 that suppresses the current driving.

Before we move on, we make one final point about an inductively driven circuit. The non-zero mutual inductance M implies that the circuit in question will also have some back-action on the driving circuit. In most cases, however, we can assume that this back-action

is negligible, so we work in the limit where $M = L_d \rightarrow 0$, but $\Phi_{\text{ext}} = MI_d$ is non-zero such that it still has an impact on the circuit. Then, Eq. 2.105 reduces to

$$\begin{aligned}\phi_1 &= L_1 I_1 + MI_d \rightarrow L_1 I_1 + \Phi_{\text{ext}} \\ \phi_d &= MI_1 + L_d I_d \rightarrow \Phi_{\text{ext}}\end{aligned}\tag{2.107}$$

and we are left with a single constraint $\phi_1 - \Phi_{\text{ext}} = LI_1$. Then, the potential energy for this circuit branch is given by

$$E_L = \int_{-\infty}^t dt' I(t') V(t') = \frac{1}{L_1} \int_{-\infty}^t (\phi_1 - \Phi_{\text{ext}}) \dot{\phi}_1 = \frac{1}{2L_1} (\phi_1 - \Phi_{\text{ext}})^2\tag{2.108}$$

Thus, we see that we can equivalently model the effect of such a drive on the circuit in this limit by taking $\phi_1 \rightarrow \phi_1 - \Phi_{\text{ext}}$ [39].

2.5 Dissipation

The Lagrangian and Hamiltonian formalism deals with conservative forces and is generally unable to handle dissipation in a straightforward manner. For classical circuits, this is generally fine as we can derive the equations of motion directly by enforcing Kirchhoff's voltage and current laws.

As an example, we can consider the case of a driven parallel LCR circuit, where we add a resistor to the driven parallel LC oscillator discussed previously. Kirchhoff's current law states that the sum of the currents through the three branches must equal the total current introduced by the current source. This yields the equation:

$$I_d = I_c + I_\ell + I_r = C\dot{V}_c + \frac{1}{L}\Phi_l + \frac{V_r}{R}\tag{2.109}$$

while Kirchhoff's voltage law tells us that $\dot{\phi} = V \equiv V_\ell = V_c = V_r$. As a result, we can write the Eq. 2.109 as

$$\ddot{\phi} + \frac{1}{RC}\dot{\phi} + \frac{1}{LC}\phi = \frac{I_d}{C}\tag{2.110}$$

This is exactly the equations of motion for a driven damped harmonic oscillator, where the “position” variable is now the node flux ϕ , the driving “force” is the current $I_d(t)$, and the effective “mass” is the capacitance C ⁶. We then identify the resonance frequency as $\omega_r = \frac{1}{\sqrt{LC}}$ and the decay rate $\gamma = \kappa_p = \frac{1}{RC}$. The transient solutions of the system are given by $\phi(t) = \phi e^{st}$, where

$$s = -\frac{\kappa_p}{2} \pm \sqrt{\left(\frac{\kappa_p}{2}\right)^2 - \omega_0^2} = -\frac{\kappa_p}{2} \pm i\omega_0 \sqrt{1 - (2Q)^{-2}}.\tag{2.111}$$

This leads to an oscillatory component at a frequency $\Omega_0 = \omega_0 \sqrt{1 - (2Q)^{-2}}$ and a decay at a rate $\kappa_p/2$ when the system is underdamped ($Q > \frac{1}{2}$).

⁶See Appendix A.1 for a derivation of the classical driven and damped harmonic oscillator.




Figure 2.10: **Driven and damped RLC circuit.** (a) A current source driving a parallel RLC circuit. (b) The frequency response functions of ϕ and Q for various values of $Q = \omega_r/\kappa_p$. The dissipative term shifts the resonance frequency and broadens the peak.

Here, we have introduced the quality factor, defined as $Q \equiv \omega_0/\kappa_p$. Writing this in terms of the circuit parameters, we find that the quality factor can also be expressed as a ratio of the resistance to the characteristic impedance of the resonant circuit

$$Q \equiv \frac{\omega_0}{\kappa_p} = \frac{RC}{\sqrt{LC}} = \frac{R}{Z_r}. \quad (2.112)$$

It is worth noting that since $Q \gg 1$ in order to observe quantum mechanical effects, the shift in the resonance frequency due to the damping term is for the most part negligible (typically we have $Q > 10^3$).

The driven solutions of the system are now $\phi(t) = \text{Re} [\phi(\omega)e^{i\omega t}]$ where

$$\phi(\omega) = \frac{I_d}{C} \frac{1}{(\omega_r^2 - \omega^2) + i\gamma\omega} \quad (2.113)$$

which is now a Lorentzian with a FWHM approximately $2\kappa_p$. The damping rate $\kappa_p = (RC)^{-1}$ has the effect of broadening the response function of the LC oscillator, which we see in Fig. 2.10(b).

We can similarly consider a series RLC circuit driven by a voltage source. In this case the Kirchhoff's voltage law yields

$$V_d = V_c + V_\ell + V_r = \frac{Q_c}{C} + L\dot{I}_r + I_r R \quad (2.114)$$

while Kirchhoff's current law enforces the condition $\dot{Q} = I \equiv I_\ell = I_c = I_r$. This yields a similar equation of motion

$$\ddot{Q} + \frac{R}{L}\dot{Q} + \frac{1}{LC}Q = \frac{V_d}{L} \quad (2.115)$$

Here, the “position” variable is now the charge Q , the driving “force” is now the voltage $V(t)$, and the effective “mass” is now the inductance L . While the intrinsic resonance frequency is still $\omega_0 = \frac{1}{\sqrt{LC}}$, the decay rate is now $\gamma = \kappa_s = \frac{R}{L}$, and the quality factor is

$$Q = \frac{\omega_0}{\kappa_s} = \frac{L}{R\sqrt{LC}} = \frac{Z_r}{R} \quad (2.116)$$

While these equations of motion are valid in the classical limit, we cannot easily generalize this to a quantum circuit. An alternative method of introducing dissipation into the Hamiltonian formalism is known as the Caldeira-Leggett model [27]. We refer the reader to Ref. [201] or Ref. [40]⁷ for a proper derivation, but we will roughly explain the central ideas. The model involves replacing the resistor (or any dissipative real-valued impedance/admittance) with an infinite array of LC oscillators each at a different frequency $\omega_m = 1/\sqrt{L_m C_m}$. The admittance $y_m = \sqrt{C_m/L_m}$ of the oscillator is matched to the real component of the admittance $\text{Re}[Y(\omega_m)]$, which characterizes the effective dissipation of the element⁸. We can then write down a Hamiltonian for the system of oscillators coupled to the circuit variable ϕ :

$$H_Y = \sum_m \left[\frac{1}{2C_m} Q_m^2 + \frac{1}{2L_m} (\phi_m - \phi)^2 \right] \quad (2.117)$$

This Hamiltonian H_Y can be added to the system Hamiltonian, and subsequently quantized. In the limit of an infinite bath, information in the circuit of interest can be irreversibly lost into bath, leading to an effective dissipation.

The Caldeira-Leggett model also leads to an important result known as the fluctuation-dissipation theorem. The idea here is that if the bath is thermalized to a finite temperature T , we can replace each of the terms in the Hamiltonian above with the average energy $\frac{1}{2}k_B T$, according to the equipartition theorem. This allows us to obtain the correlation function of the current fluctuations in each oscillator and sum these contributions over the entire bath of oscillators to derive their spectral density [201]:

$$S_{II}(\omega) = 2k_B T \text{Re}[Y(\omega)] \quad (2.118)$$

The spectral density of voltage fluctuations due to a dissipative element is similarly given by:

$$S_{VV}(\omega) = 2k_B T \text{Re}[Z(\omega)] \quad (2.119)$$

These fluctuations are known as Johnson-Nyquist noise [153]. The fluctuation-dissipation theorem is important in that it explains the equivalence between noise and dissipation, both of which give rise to irreversible dynamics. There is also a quantum analog of the fluctuation-dissipation theorem that exhibits some notable differences from its classical counterpart, and we recommend the reader consult Ref. [40] for a complete review of the subject.

⁷The classical version of Caldeira-Leggett is discussed in Appendix C.

⁸For a resistor, $Y(\omega) = R^{-1}$ is uniform over all frequencies, however, this model is capable of handling an arbitrary dissipative element.

2.6 Quantum Circuits

So far we have discussed electrical circuits entirely in the classical limit, although much of the circuit analysis we have worked through is equally applicable to quantum circuits. So what makes an electrical circuit quantum? The quantum-to-classical transition is largely characterized by dissipation, which leads to irreversible dynamics. Closed quantum systems, on the other hand, are described by unitary time evolution which is by definition reversible. The reality, however, is that no experimental quantum system that we can interact with is truly closed. The circuit typically interacts with an environment that is uncontrolled and consists of a large number of degrees of freedom, leading to noise or equivalently dissipation. This ultimately leads to a loss of quantum coherence [219], which suppresses the quantum computational power that we would like to leverage.

Practically, for quantum effects to be relevant, we require our circuits to have low dissipation so that the quality factor Q of the circuit modes interest are large ($Q \gg 1$). We also require $k_B T \ll \hbar \omega$, where ω is the frequency of the circuit modes, since thermal noise can also lead to excess decoherence.

Canonical Quantization

In the quantum limit of low dissipation, the superconducting charge and phase can be promoted to operators $\hat{Q}$, $\hat{\phi}$ with the commutator

$$[\hat{\phi}, \hat{Q}] = i\hbar \quad (2.120)$$

Then, the mode amplitudes a and a^* for the quantum harmonic oscillator become

$$a \rightarrow \hat{a} = \frac{1}{\sqrt{2\hbar Z_r}} \hat{\phi} + i\sqrt{\frac{Z_r}{2\hbar}} \hat{Q} \quad (2.121)$$

$$a^* \rightarrow \hat{a}^\dagger = \frac{1}{\sqrt{2\hbar Z_r}} \hat{\phi} - i\sqrt{\frac{Z_r}{2\hbar}} \hat{Q} \quad (2.122)$$

with the commutation relation

$$[a, a^\dagger] = \frac{1}{2\hbar Z_r} [\hat{\phi}, \hat{\phi}] - \frac{i}{2\hbar} [\hat{\phi}, \hat{Q}] + \frac{i}{2\hbar} [\hat{Q}, \hat{\phi}] + \frac{Z_r}{2\hbar} [\hat{Q}, \hat{Q}] = \frac{1}{i\hbar} [\hat{\phi}, \hat{Q}] = 1 \quad (2.123)$$

Then, the Hamiltonian in Eq. 2.19 becomes

$$\hat{H} = \hbar\omega_0 \left(a^\dagger a + \frac{1}{2} \right) \quad (2.124)$$

and now includes the zero-point energy $\hbar\omega_0/2$. Here we find the first two consequences of quantum mechanics: in the quantization of the discrete energy levels of the quantum LC oscillator and the vacuum fluctuations that stem from the non-commuting observables $\hat{\phi}$ and $\hat{Q}$.

Superconductivity

What is perhaps most remarkable about circuit QED is that quantum coherence can exist macroscopically within the reduced degrees of freedom in a superconducting circuit [123]. According to Ginzburg–Landau theory, the Cooper pairs in a superconductor can be described by a single wavefunction

$$\Psi(\mathbf{r}, t) = \sqrt{n_s(\mathbf{r}, t)} e^{i\theta(\mathbf{r}, t)}, \quad (2.125)$$

where $n_s(\mathbf{r}, t)$ is the Cooper pair density at a given point, and $\theta(\mathbf{r}, t)$ is the phase of the wavefunction. Here, the wavefunction is normalized so that $\int d^3\mathbf{r} |\Psi(\mathbf{r}, t)|^2 = \int d^3\mathbf{r} n_s(\mathbf{r}, t) = N_s$ is the total number of Cooper pairs. This wavefunction should obey the Schrodinger equation for charged particles:

$$i\hbar \frac{\partial \Psi(\mathbf{r}, t)}{\partial t} = \frac{1}{2m_s} [-i\hbar \nabla - q_s \mathbf{A}(\mathbf{r}, t)]^2 \Psi(\mathbf{r}, t) + q_s V(\mathbf{r}, t) \Psi(\mathbf{r}, t) \quad (2.126)$$

where, $V(\mathbf{r}, t)$ and $\mathbf{A}(\mathbf{r}, t)$ are the scalar and vector potentials, and m_s and $q_s = -2e$ are the mass and charge of the Cooper pairs. From this Schrodinger equation, we can derive the probability current for the wave function, which becomes a physical current of Cooper pairs when multiplied by their charge $q_s = -2e$.

$$\mathbf{J}_s = q_s \frac{i\hbar}{2m_s} (\Psi^* \nabla \Psi - \Psi \nabla \Psi^*) - \frac{q_s^2}{2m_s} \Psi \Psi^* \mathbf{A} \quad (2.127a)$$

$$= q_s n_s(\mathbf{r}, t) \frac{1}{m_s} [\hbar \nabla \theta(\mathbf{r}, t) - q_s \mathbf{A}(\mathbf{r}, t)] \quad (2.127b)$$

$$= \frac{q_s^2 n_s(\mathbf{r}, t)}{m_s} \frac{\hbar}{q_s} \left[\nabla \theta(\mathbf{r}, t) + \frac{q_s}{\hbar} \mathbf{A}(\mathbf{r}, t) \right] \quad (2.127c)$$

$$= -\frac{1}{\Lambda} \frac{\Phi_0}{2\pi} \gamma \quad (2.127d)$$

where we introduce the superconducting flux quantum $\Phi_0 = 2e/h$ and the London coefficient $\Lambda = \frac{m_s}{n_s q_s^2}$, and the *gauge-invariant phase gradient* γ . To understand why γ is gauge-invariant, we recall from classical E&M that Maxwell's equations are invariant under a gauge transformation

$$\mathbf{A} \rightarrow \mathbf{A} + \nabla \chi \equiv \mathbf{A}' \quad V \rightarrow V - \frac{\partial \chi}{\partial t} \equiv V' \quad (2.128)$$

Now suppose we write the current $\mathbf{J}_s$ in terms of the transformed vector potential $\mathbf{A}'$:

$$\mathbf{J}_s = q_s n_s \frac{1}{m} [\hbar \nabla \theta' - q_s \mathbf{A} - q_s \nabla \chi] \quad (2.129)$$

However, since currents are physically observable quantities, this must also be gauge invariant, so we must have $\hbar \nabla \theta' - q_s \nabla \chi = \hbar \nabla \theta$ and θ must transform as

$$\theta' = \theta + \frac{q_s}{\hbar} \chi = \theta - \frac{2\pi}{\Phi_0} \chi \Rightarrow \Psi'(\mathbf{r}, t) = \Psi(\mathbf{r}, t) e^{-i \frac{2\pi}{\Phi_0} \chi} \quad (2.130)$$

which leads to an unmeasurable phase shift on the superconducting wavefunction.

Finally, we note that for the wave function $\Psi(\mathbf{r}, t) = \sqrt{n_s}e^{i\theta}$, the Schrodinger equation (Eq. 2.126) yields the relationship

$$-\hbar \frac{\partial \theta(\mathbf{r}, t)}{\partial t} + \frac{1}{2n_s(\mathbf{r}, t)} \frac{\partial n_s(\mathbf{r}, t)}{\partial t} = \frac{1}{2n_s} \Lambda \mathbf{J}_s^2 + q_s V(\mathbf{r}, t) \quad (2.131)$$

Generally, we can assume that the charge density is constant and the same within all the superconductors comprising a circuit, so the second term on the LHS becomes zero. Then, given a general two-terminal superconducting circuit element, we obtain the expression

$$\frac{\partial \varphi}{\partial t} = \frac{\partial \theta_1}{\partial t} - \frac{\partial \theta_2}{\partial t} = \frac{2e}{\hbar} (V_1 - V_2) = \frac{2\pi}{\Phi_0} \Delta V \quad (2.132)$$

As a result, we see that the superconducting wavefunction phase φ is proportional to the branch flux Φ across the circuit element, and that the circuit variables $\hat{Q}$ and $\hat{\Phi}$ are now encoded in the amplitude and phase of the macroscopic wavefunction of the superconducting wavefunction⁹.

Flux Quantization

An important consequence of superconductivity is that the magnetic flux threaded through a superconducting loop must be quantized in units of the superconducting flux quantum $\Phi_0 = \frac{h}{2e}$. To see why, we consider the supercurrent $\mathbf{J}_s$ and integrate it around a closed loop C . We find

$$\begin{aligned} \Lambda \oint_C \mathbf{J}_s \cdot d\mathbf{l} &= \frac{\hbar}{q_s} \oint_C d\mathbf{l} \cdot (\nabla \theta) - \oint_C \mathbf{A} \cdot d\mathbf{l} \\ &= \frac{\hbar}{q_s} \oint_C d\mathbf{l} \cdot (\nabla \theta) - \int_S \mathbf{B} \cdot d\mathbf{S} \end{aligned} \quad (2.133)$$

We note that the first term on the right hand side can be evaluated for a path starting at $\mathbf{x}_1$ and ending at $\mathbf{x}_2$ as

$$\int_{\mathbf{x}_1}^{\mathbf{x}_2} (\nabla \theta) \cdot d\mathbf{l} = \theta(\mathbf{x}_2, t) - \theta(\mathbf{x}_1, t) \quad (2.134)$$

Since the phase $\theta(\mathbf{x}, t)$ must be singled-valued modulo 2π , we conclude that around a closed contour C , we must have

$$\lim_{\mathbf{x}_1 \rightarrow \mathbf{x}_2} [\theta(\mathbf{x}_2, t) - \theta(\mathbf{x}_1, t)] = 2\pi n \quad (2.135)$$

⁹One should be careful to identify whether the phase operator $\hat{\varphi}$ is periodic or extended, which depends on the context. We refer the reader to Ref. [48] for a discussion on the compactness of $\hat{\varphi}$.

Since the current inside a superconductor is zero, the left hand side must also vanish, and we are left with

$$\Phi = \int_S \mathbf{B} \cdot d\mathbf{S} = 2\pi n \frac{\hbar}{q_s} = n\Phi_0 \quad (2.136)$$

For discrete superconducting circuits, flux quantization manifests as an additional constraint on the sum of the branch fluxes making up a loop

$$\sum_{(i,j) \in E} \Phi_{ij} = \sum_{(i,j) \in E} (\phi_i - \phi_j) = \Phi_{\text{ext}} \quad (2.137)$$

where E is the set of edges (i, j) making up the circuit loop. We then ask what implications this has on the quantization on a linear circuit. Suppose we have a circuit consisting of two inductors L_1 and L_2 that form a loop. The total current across the two inductors will be given by

$$I = I_1 + I_2 = \frac{1}{L_1} \Phi_1 + \frac{1}{L_2} \Phi_2 \quad (2.138)$$

where Φ_1 and Φ_2 are the branch fluxes across the two inductors. The flux quantization condition is given by $\Phi_1 - \Phi_2 = \Phi_{\text{ext}}^{10}$, and we can choose to define the node flux ϕ as

$$\Phi_1 = \phi \quad \Phi_2 = \phi - \Phi_{\text{ext}} \quad (2.139)$$

so the total potential energy is given by

$$E_L = \int I \dot{\phi} dt = \frac{1}{2L_1} \phi^2 + \frac{1}{2L_2} (\phi - \Phi_{\text{ext}})^2 \quad (2.140a)$$

$$= \frac{1}{2} \left(\frac{1}{L_1} + \frac{1}{L_2} \right) \left[\phi - \frac{L_1}{L_1 + L_2} \Phi_{\text{ext}} \right]^2 + \frac{1}{2} \frac{1}{L_1 + L_2} \Phi_{\text{ext}}^2 \quad (2.140b)$$

where in the second line we complete the square. The first term is the just the potential energy for an inductor with inductance L , where $L^{-1} = L_1^{-1} + L_2^{-1}$ and a translated node flux $\phi' = \phi - \frac{L_1}{L_1 + L_2} \Phi_{\text{ext}}$, while the second term is a constant offset energy. Since the node flux does not have periodic boundary conditions here, this results in just a gauge transformation that does not change the system dynamics. Note that we could have also assigned the external flux such that

$$\Phi_1 = \phi + \Phi_{\text{ext}} \quad \Phi_2 = \phi \quad (2.141)$$

in which case the potential energy is given by

$$E_L = \frac{1}{2} \left(\frac{1}{L_1} + \frac{1}{L_2} \right) \left[\phi + \frac{L_2}{L_1 + L_2} \Phi_{\text{ext}} \right]^2 + \frac{1}{2} \frac{1}{L_1 + L_2} \Phi_{\text{ext}}^2 \quad (2.142)$$

The two are equivalent up to a definition of ϕ' , and we note that we can translate between the two Hamiltonians via the unitary transformation [212]:

$$U = e^{-\frac{i}{\hbar} Q \Phi_{\text{ext}}} \quad (2.143)$$

¹⁰We choose to define both branch fluxes Φ_1 and Φ_2 in the same direction, so the sign of Φ_2 in the flux-quantization condition is flipped.

2.7 Nonlinear Circuits

Although we can describe a linear circuit quantum-mechanically, it is quite difficult to leverage the quantum properties of circuit to do very much without a source of nonlinearity. The quantum LC oscillator without any nonlinearity will be in the correspondence limit, and any attempt to drive the system on resonance will result in a coherent state since the energy levels are evenly spaced. Flux quantization in the presence of external flux leads only to a gauge transformation that cannot be measured. This is not to say that we cannot observe quantum-mechanical effects, as the vacuum fluctuations and flux quantization can still be observed. However, in order to construct a device capable of processing quantum information, we need universal control over some set of basis states encoded in the circuit, which requires the circuit to be nonlinear. Since dissipation will suppress quantum coherence, we also need the source of nonlinearity to be dissipationless.

Josephson Junctions

The Josephson junction provides exactly the dissipationless nonlinearity that we need, and is consequently the foundation upon which modern superconducting qubit devices are built. Josephson junctions are formed when a potential barrier is created between two superconductors so that quantum tunneling of the supercurrent can occur. While there are many ways to create such a barrier, we will focus on superconductor-insulator-superconductor (SIS) junctions, since these are most commonly seen in superconducting qubit devices. The Josephson junction is characterized by the two Josephson relations:

$$I(t) = I_c \sin(\varphi(t)) \quad (2.144)$$

$$\dot{\varphi}(t) = \frac{2e}{\hbar} V(t) = \frac{2\pi}{\Phi_0} V(t) \quad (2.145)$$

where φ is the superconducting phase difference across the junction, and $\Phi = \frac{2\pi}{\Phi_0} \varphi$ is the circuit branch flux across the junction. The critical current I_c is a property of the junction itself, and typically depends exponentially on the junction barrier thickness and linearly in the junction area. These equations lead to many surprising consequences, the first being that the junction can sustain a non-zero current even when $V = 0$ across the junction, which is known as the *DC Josephson effect*. On the other hand, for a non-zero constant voltage V , the phase will grow linearly in time and we obtain an oscillating current with frequency $2\pi V/\Phi_0$. This is known as the *AC Josephson effect*.

As a circuit element, the Josephson junction behaves like a nonlinear inductor. In particular, we can see that

$$\dot{I}(t) = \dot{\varphi} I_c \cos(\varphi(t)) = \frac{2\pi I_c}{\Phi_0} V(t) \cos(\varphi(t)) \quad (2.146)$$

$$V(t) = \frac{1}{L_J \cos(\varphi(t))} \dot{I}(t) \quad (2.147)$$




Figure 2.11: **Flux quantization with Josephson junctions.** (a) A circuit diagram showing two Josephson junctions enclosing a superconducting loop. A magnetic flux Φ_{ext} threads the loop. (b) The normalized $I_{c, \text{eff}}(\Phi_{\text{ext}})$ for different values of the SQUID asymmetry parameter d . As d increases and the SQUID becomes more asymmetric, the effective range over which the total critical current can be tuned is reduced.

where we define $L_J = \frac{\Phi_0}{2\pi I_c}$ as the Josephson inductance. Surprisingly, this inductance can even be negative. We can also compute the stored energy within the junction as

$$\int_{-\infty}^t I(t')V(t') dt' = \frac{\Phi_0 I_c}{2\pi} \int_{-\infty}^t \dot{\varphi} \sin \varphi dt' = -E_J \cos \varphi \quad (2.148)$$

where we define the Josephson energy E_J as $E_J = \frac{\Phi_0 I_c}{2\pi}$. If we express E_J in terms of L_J and Taylor expand the cosine term, we find that the total energy can be written as

$$-\frac{1}{L_J} \left(\frac{\Phi_0}{2\pi} \right)^2 \cos(\varphi) \approx -\frac{1}{L_J} \left(\frac{\Phi_0}{2\pi} \right)^2 \left[1 - \frac{1}{2}\varphi^2 + \frac{1}{24}\varphi^4 + \dots \right] \quad (2.149)$$

As a result, we see that to second order in φ , and ignoring the constant factor, the energy stored in the junction looks like the inductive energy $\frac{1}{2L_J}\Phi^2$.

Flux Quantization with Josephson Junctions

Now that we have introduced the Josephson junction, we return to the problem of flux quantization. We suppose now that we have two junctions that form a superconducting loop, a structure known as the DC SQUID¹¹. The net current flowing through the SQUID is given by the sum of the currents through each junction:

$$I_s = I_{c1} \sin \varphi_1 + I_{c2} \sin \varphi_2, \quad (2.150)$$

¹¹SQUID stands for superconducting quantum interference device.

where φ_1 and φ_2 are the phase differences across the SQUID. We then redefine the phases in terms of the sum and difference phases

$$\delta = \varphi_1 - \varphi_2 \quad \varphi = \frac{1}{2}(\varphi_1 + \varphi_2) \quad (2.151)$$

and we can express φ_1, φ_2 in terms of δ, φ as

$$\varphi_1 = \varphi + \frac{1}{2}\delta \quad \varphi_2 = \varphi - \frac{1}{2}\delta \quad (2.152)$$

Using the sum and difference trigonometric identities, we can write

$$\sin \varphi_1 = \sin \left(\varphi + \frac{\delta}{2} \right) = \sin \varphi \cos \frac{\delta}{2} + \cos \varphi \sin \frac{\delta}{2} \quad (2.153a)$$

$$\sin \varphi_2 = \sin \left(\varphi - \frac{\delta}{2} \right) = \sin \varphi \cos \frac{\delta}{2} - \cos \varphi \sin \frac{\delta}{2} \quad (2.153b)$$

At this point, we also enforce the flux quantization condition

$$\delta = \varphi_1 - \varphi_2 = \frac{2\pi\Phi_{\text{ext}}}{\Phi_0} \quad (2.154)$$

where Φ is the total flux threaded through the loop. Then, we can write the total current as

$$I_s = (I_{c1} + I_{c2}) \cos \frac{\pi\Phi_{\text{ext}}}{\Phi_0} \sin \varphi + (I_{c1} - I_{c2}) \sin \frac{\pi\Phi_{\text{ext}}}{\Phi_0} \cos \varphi \quad (2.155a)$$

$$= I_\Sigma \left[\cos \frac{\pi\Phi_{\text{ext}}}{\Phi_0} \sin \varphi + d \sin \frac{\pi\Phi_{\text{ext}}}{\Phi_0} \cos \varphi \right] \quad (2.155b)$$

where we define

$$I_\Sigma = I_{c1} + I_{c2} \quad d = \frac{I_{c1} - I_{c2}}{I_{c1} + I_{c2}} \quad (2.156)$$

We then define a new angle φ_0 such that

$$\sin \varphi_0 = \frac{d \sin \frac{\pi\Phi_{\text{ext}}}{\Phi_0}}{\sqrt{\cos^2 \frac{\pi\Phi_{\text{ext}}}{\Phi_0} + d^2 \sin^2 \frac{\pi\Phi_{\text{ext}}}{\Phi_0}}} \quad (2.157a)$$

$$\cos \varphi_0 = \frac{\cos \frac{\pi\Phi_{\text{ext}}}{\Phi_0}}{\sqrt{\cos^2 \frac{\pi\Phi_{\text{ext}}}{\Phi_0} + d^2 \sin^2 \frac{\pi\Phi_{\text{ext}}}{\Phi_0}}} \quad (2.157b)$$

$$\Rightarrow \tan \varphi_0 = d \tan \frac{\pi\Phi_{\text{ext}}}{\Phi_0} \quad (2.157c)$$

Substituting this into the equation for I_s , we conclude that

$$I_s = I_\Sigma \sqrt{\cos^2 \frac{\pi \Phi_{\text{ext}}}{\Phi_0} + d^2 \sin^2 \frac{\pi \Phi_{\text{ext}}}{\Phi_0}} \sin(\varphi + \varphi_0) \quad (2.158)$$

which leads to an effective critical current

$$I_{c,\text{eff}}(\Phi_{\text{ext}}) = I_\Sigma \sqrt{\cos^2 \frac{\pi \Phi_{\text{ext}}}{\Phi_0} + d^2 \sin^2 \frac{\pi \Phi_{\text{ext}}}{\Phi_0}} \quad (2.159)$$

In the simple case where $I_{c1} = I_{c2} = I_c$, we see that $d = 0$ and this reduces to

$$I_{c,\text{eff}}(\Phi_{\text{ext}}) = 2I_c \left| \cos \frac{\pi \Phi_{\text{ext}}}{\Phi_0} \right| \quad (2.160)$$

From the above analysis, we see that the nonlinearity of the current phase relationship leads to non-trivial behavior, unlike the case with two linear inductors. The SQUID now acts like a Josephson junction, but with a flux-tunable critical current. This allows us to use the SQUID as tunable nonlinear inductor to modify the properties of a circuit in-situ with an applied magnetic flux. Finally, we note that while this analysis is valid in the DC limit, one must be careful when the magnetic flux $\Phi_{\text{ext}}(t)$ has a time dependence. We refer the reader to Ref. [212] and Ref. [168] for a proper treatment of this case.

Transmon as a Nonlinear Oscillator

As our final circuit, we consider what happens when we replace the linear inductor in the quantum LC oscillator with a nonlinear Josephson junction [105]. The Lagrangian for the system is given by

$$\mathcal{L} = \frac{1}{2} C \dot{\Phi}^2 + E_J \cos \left(\frac{2\pi \Phi}{\Phi_0} \right) \quad (2.161)$$

which leads to a classical equation of motion for the system

$$C \ddot{\Phi} + I_c \sin \left(\frac{2\pi \Phi}{\Phi_0} \right) = 0 \quad (2.162)$$

This looks like the equations of motion for a pendulum in a gravitational field, where we let $2\pi\Phi/\Phi_0$ be the angle of the pendulum and we let

$$\frac{2\pi I_c}{\Phi_0} \frac{1}{C} = \frac{1}{\sqrt{L_J C}} \quad (2.163)$$

be the effective inverse length g/l .

The quantum Hamiltonian for this system is given by

$$\hat{H} = \frac{1}{2C} \hat{Q}^2 - E_J \cos \left(\frac{2\pi \hat{\Phi}}{\Phi_0} \right) = 4E_c \hat{n} - E_J \cos \hat{\varphi} \quad (2.164)$$




Figure 2.12: **Transmon as a nonlinear oscillator.** (a) The circuit for a transmon, which consists of a capacitance shunting a Josephson junction. (b) The nonlinear potential due to the Josephson junction can be mapped to the gravitational potential for a rigid pendulum. For small angles the pendulum is approximately a harmonic oscillator, but nonlinear effects appear at larger angles. (c) The anharmonic potential and energy levels for a transmon (solid) with E_C and E_J , plotted against the harmonic potential and its energy levels (dashed) for an oscillator with $L = \frac{1}{E_J} \left(\frac{\Phi_0}{2\pi}\right)^2$ and $C = \frac{e^2}{2E_C}$. The energy levels of the harmonic oscillator are evenly spaced, while the spacing between the energy levels in the transmon decrease for higher energies. The ω_{01} transition of the transmon is also reduced by approximately E_C .

where we have replaced the charge $\hat{Q}$ and flux $\hat{\Phi}$ with normalized operators $\hat{n} = \hat{Q}/(2e)$ and $\hat{\varphi} = 2\pi\hat{\Phi}/\Phi_0$. We also define the charging energy E_C as $E_C = \frac{e^2}{2C}$. The behavior of this system depends heavily on the ratio of the charging energy E_C to the Josephson energy E_J . In the low E_J/E_C limit, the system is typically referred to as a Cooper pair box and acts as a charge qubit, while in the extremely high E_J/E_C limit ($E_J/E_C \sim 10^4$) the system acts as a phase qubit. For the last two decades, however, this system has primarily been operated in the *transmon* regime, where $E_J/E_C \sim 50$ – 150 [105]. In this limit, the charge dispersion of the eigenstates is exponentially suppressed while still maintaining sufficient intrinsic nonlinearity in the system.

While the energy levels of the Hamiltonian in Eq. 2.164 can be solved exactly in terms of Mathieu functions [47], the weak nonlinearity of the transmon allows us to perturbatively solve for the energy levels and eigenstates of the system by expanding the cosine potential. Then, the frequency and anharmonicity of the Transmon can be expressed as [51]:

$$\omega_{01} = \sqrt{8E_J E_C} - E_C \left[1 + \frac{1}{4}\eta + \frac{21}{128}\eta^2 + \mathcal{O}(\eta^3) \right] \quad (2.165)$$

$$\alpha = \omega_{12} - \omega_{01} = -E_C \left[1 + \frac{9}{16}\eta + \frac{81}{128}\eta^2 + \mathcal{O}(\eta^3) \right] \quad (2.166)$$

Here, we include the terms up to second order in the parameter $\eta = \frac{2E_C}{E_J}$, since the more

commonly seen zero-th order expressions are only accurate to a few percent in the Transmon regime.

The simplicity of the transmon, its sufficient nonlinearity, and its relative insensitivity to both charge and flux noise, all combine to make it an ideal candidate for implementing a superconducting qubit. Furthermore, the weak nonlinearity of the transmon also makes it straightforward to design the supporting circuitry for qubit control and readout. In most transmon-based systems, we can analyze the circuits in the harmonic oscillator basis, treating the nonlinearity of the Josephson junctions perturbatively. This makes the circuit analysis discussed in this chapter readily applicable, and greatly simplifies the design process for such systems. In the following chapter, we will discuss some additional techniques for designing practical superconducting qubit systems in the microwave regime.

Chapter 3

Designer Microwave Circuits

In the previous chapter, we introduced the basic principles of circuit analysis using lumped-element models. However, at microwave frequencies, lumped circuits may not always be good approximations to physical structures. In this chapter we will introduce better models for common superconducting qubit structures with distributed-elements, and show how these models can lead to more accurate predictions of device parameters.

3.1 Frequency Domain Analysis

Thus far we have primarily been analyzing circuits in the time-domain; however with microwave circuits it is often useful to analyze circuits in the frequency domain instead. In general, the linearity of the circuit elements discussed so far allows these circuits to be treated as *linear time-invariant* (LTI) networks. For a more in-depth treatment on such systems, we refer the reader to Chapter 5 of Ref. [39], but in the following sections we will briefly introduce the important concepts. For the typical two-terminal elements that we use, we have the following current voltage relationships:

$$V(t) = RI(t) \quad V(t) = L\dot{I}(t) \quad V(t) = \frac{1}{C} \int_{-\infty}^t I(t) dt \quad (3.1)$$


and we can perform a Fourier transform to obtain

$$V(\omega) = \frac{R}{\sqrt{2\pi}} \int_{-\infty}^{\infty} I(t)e^{-i\omega t} dt = RI(\omega) \quad (3.2a)$$

$$V(\omega) = \frac{L}{\sqrt{2\pi}} \int_{-\infty}^{\infty} \dot{I}(t)e^{-i\omega t} dt = -\frac{i\omega L}{\sqrt{2\pi}} \int_{-\infty}^{\infty} I(t)e^{-i\omega t} dt = i\omega LI(\omega) \quad (3.2b)$$

$$V(\omega) = \frac{1}{C} \int_{-\infty}^{\infty} Q(t)e^{-i\omega t} dt = -\frac{1}{i\omega C} \int_{-\infty}^{\infty} I(t)e^{-i\omega t} dt = \frac{1}{i\omega C} I(\omega) \quad (3.2c)$$

Here, we perform an integration by parts for the inductor and capacitor. It is also common to analyze these circuits in the Laplace domain where we take $i\omega \rightarrow s$.

Figure 3.1: **Equivalent circuits for calculating effective resistance.**

As an example of how the impedances can be useful, we go back to the example of the RLC oscillator. We can express the admittance of the parallel circuit in terms of $s = i\omega$ as

$$Y(s) = Z(s)^{-1} = sC + \frac{1}{sL} + \frac{1}{R} = C \left(\frac{s^2 + \frac{1}{RC}s + \frac{1}{LC}}{s} \right) = C \left(\frac{s^2 + \kappa_p s + \omega_0^2}{s} \right) \quad (3.3)$$

Here, we see that the complex zeros of the admittance (poles of the impedance) are simply the transient solutions to the differential equation given in Eq. 2.110, allowing us to easily identify the resonance frequency and decay rate.

Similarly, we can consider a series RLC circuit. In this case, the impedance of the series circuit is

$$Z(s) = \frac{1}{sC} + sL + R = \frac{L}{s} \left(s^2 + \frac{R}{L}s + \frac{1}{LC} \right) = L \left(\frac{s^2 + \kappa_s s + \omega_0^2}{s} \right) \quad (3.4)$$

and we see that the complex zeros of the impedance (poles of the admittance) yield the resonance frequency and decay rate of the circuit.

Equivalent Circuits

This type of impedance analysis is also useful for simplifying more complex circuits. In general, if we know the impedance response of a system, we can construct a simpler equivalent circuit that has the same impedance, and thus the same behavior. A typical use case involves calculating decay rates or loaded quality factors of a resonant mode that is coupled indirectly to an external port. If we coupled the port directly to the a parallel LC circuit, the total parallel resistance of the circuit would be $R_{\text{env}} \parallel R_{\text{int}} \approx R_{\text{env}}$ for a typical high-Q circuit coupled to a $50 \, \Omega$ port, leading to strong damping. Thus, we will typically couple the resonator circuit to the transmission line via a small capacitance or mutual inductance.

In the case of a capacitance, we can write the total admittance of the coupling capacitance C_κ and environmental impedance Z_0 at the resonator frequency ω_r as

$$Y_{\text{env}}(\omega_r) = \left(\frac{1}{i\omega_r C_\kappa} + Z_0 \right)^{-1} = \frac{\omega_r^2 C_\kappa^2 Z_0}{1 + \omega_r^2 C_\kappa^2 Z_0^2} + i\omega_r \frac{C_\kappa}{1 + \omega_r^2 C_\kappa^2 Z_0^2}. \quad (3.5)$$

This leads to an additional loading capacitance

$$C^* = \frac{C_\kappa}{1 + \left(\frac{Z_0}{Z_r}\right)^2 \left(\frac{C_\kappa}{C}\right)^2} \approx C_\kappa \left[1 - \left(\frac{Z_0}{Z_r}\right)^2 \left(\frac{C_\kappa}{C}\right)^2\right] \quad (3.6)$$

and an external resistance

$$R_{\text{ext}} = Z_0 \left[1 + \left(\frac{Z_r}{Z_0}\right)^2 \left(\frac{C}{C_\kappa}\right)^2\right] \quad (3.7)$$

The external decay rate due to the R_{ext} is given by

$$\kappa_{\text{ext}} = \frac{1}{R_{\text{ext}}(C + C^*)} \quad (3.8a)$$

$$= \frac{1}{Z_0 C} \left(\frac{Z_0}{Z_r}\right)^2 \left(\frac{C_\kappa}{C}\right)^2 \left[1 + \frac{C_\kappa}{C} + \left(\frac{Z_0}{Z_r}\right)^2 \left(\frac{C_\kappa}{C}\right)^2\right]^{-1} \quad (3.8b)$$

$$= \frac{Z_0 \omega_r^2 C_\kappa^2}{C} \left[1 + \frac{C_\kappa}{C} + \left(\frac{Z_0}{Z_r}\right)^2 \left(\frac{C_\kappa}{C}\right)^2\right]^{-1} \quad (3.8c)$$

$$\approx \frac{Z_0 \omega_r^2 C_\kappa^2}{C} \left[1 - \frac{C_\kappa}{C} + \left(\frac{C_\kappa}{C}\right)^2 - \left(\frac{Z_0}{Z_r}\right)^2 \left(\frac{C_\kappa}{C}\right)^2\right] \quad (3.8d)$$

We can consider a similar case of a resonator coupled via a mutual inductance M to the resonator. Here, the total admittance of the resistive and inductive elements is given by

$$Y_{\text{env}}(\omega_r) = \frac{Z_0 M^2}{Z_0^2 L^2 + \omega_r^2 (LL_0 - M^2)^2} + \frac{1}{i\omega_r} \frac{Z_0^2 L + \omega_r^2 L_0 (LL_0 - M^2)}{Z_0^2 L^2 + \omega_r^2 (LL_0 - M^2)^2} \quad (3.9)$$

Again, we see can extract an effective resistance and inductance from the expression for the admittance. We can write the external resistance as

$$R_{\text{ext}} = Z_0 \frac{L^2}{M^2} \left[1 + \left(\frac{Z_r}{Z_0}\right)^2 \left(\frac{L_0}{L} - \frac{M^2}{L^2}\right)^2\right] \quad (3.10)$$

and the modified inductance as

$$L' = L \left[\frac{Z_0^2 + Z_r^2 \left(\frac{L_0}{L} - \frac{M^2}{L^2}\right)^2}{Z_0^2 + Z_r^2 \frac{L_0}{L} \left(\frac{L_0}{L} - \frac{M^2}{L^2}\right)} \right] \quad (3.11)$$

We can then express the external decay rate due to the R_{ext} as

$$\kappa_{\text{ext}} = \frac{1}{R_{\text{ext}} C} = \frac{1}{Z_0} \frac{\omega_r^2 M^2}{L} \left[1 + \left(\frac{Z_r}{Z_0}\right)^2 \left(\frac{L_0}{L} - \frac{M^2}{L^2}\right)^2\right]^{-1} \quad (3.12)$$

3.2 Microwave Networks

For more complex circuits, it is often useful to break up the circuit into a set of smaller circuits that can be connected together at ports to form a larger network. In this case, each subsystem can be described completely by the voltage/current relationships at each port of the network, making it possible to construct equivalent circuits that exhibit the same behavior for simplified modeling. The simplest such network is a 1-port network, which in our context typically consists of some two-terminal element. This element is described by an impedance $Z(\omega)$ or admittance $Y(\omega) = Z^{-1}(\omega)$, which relates the voltage $V(\omega)$ across the two terminals to the current $I(\omega)$ through the element, via Ohm's law:

$$V(\omega) = Z(\omega)I(\omega) \quad I(\omega) = Y(\omega)V(\omega) \quad (3.13)$$

Here we are explicit about the frequency dependence of the relevant quantities, to remind ourselves that we are analyzing the network in the frequency domain, but we will keep the frequency dependence implicit for the remainder of the chapter.

For an N -port network, we now need N^2 parameters to fully specify the current and voltage relationship between each port, the network is now described by a $N \times N$ impedance matrix $\mathbf{Z}(\omega)$ or admittance matrix $\mathbf{Y}(\omega) = \mathbf{Z}^{-1}(\omega)$, where the inverse is the full matrix inverse. The basis for these matrices, is usually taken to be the same as the ports, such that

$$\mathbf{V} = \begin{pmatrix} V_1 \\ \vdots \\ V_N \end{pmatrix} \quad \mathbf{I} = \begin{pmatrix} I_1 \\ \vdots \\ I_N \end{pmatrix} \quad (3.14)$$

where V_i is the voltage across the i th port and I_i is the current going *into*¹ the i th port. Then we have a system of N linear equations given by $\mathbf{V} = \mathbf{Z}\mathbf{I}$ or $\mathbf{I} = \mathbf{Y}\mathbf{V}$ defining the behavior of the network. Each element Z_{ij} of the impedance matrix is defined as

$$Z_{ij} = \left. \frac{V_i}{I_j} \right|_{I_k=0 \text{ for } k \neq j} \quad (3.15)$$

which is the impedance between port i and port j when all other ports are an open circuit. Similarly, we can define Y_{ij} as

$$Y_{ij} = \left. \frac{I_i}{V_j} \right|_{V_k=0 \text{ for } k \neq j} \quad (3.16)$$

which is equivalent to the admittance between port i and port j when all other ports are short-circuited.

The utility of the network analysis formalism is that given a set of several independent networks, we can connect them at their ports and easily compute the network parameters for

¹This determines the sign convention, which is usually taken to be +1 for a current flowing into the port.




Figure 3.2: **Two-port network.** The voltage at each port is defined with respect to the two terminals at each port, and the current is defined as going into the port.

the combined system by enforcing the relevant boundary conditions. For example, suppose we have a two-port network with impedance matrix

$$\mathbf{Z} = \begin{bmatrix} Z_{11} & Z_{12} \\ Z_{21} & Z_{22} \end{bmatrix} \quad (3.17)$$

and a one-port network with impedance Z_L . We can connect the one-port network, or load impedance, to port 2 to create an effective one-port network looking into port 1. To determine the resulting port impedance Z'_{11} , we must enforce the additional constraints $V_2 = V_L$ and $I_L = -I_2$, which implies that $V_2 = -Z_L I_2^2$. With this additional constraint, we can write I_2 as a function of I_1

$$V_2 = Z_{21}I_1 + Z_{22}I_2 \quad (3.18)$$

$$-Z_L I_2 = Z_{21}I_1 + Z_{22}I_2 \quad (3.19)$$

$$I_2 = -\frac{Z_{21}}{Z_L + Z_{22}} I_1 \quad (3.20)$$

and then eliminate I_2 , so that the voltage-current relation for port 1 becomes

$$V_1 = Z_{11}I_1 + Z_{12}I_2 = Z_{11}I_1 - \frac{Z_{12}Z_{21}}{Z_L + Z_{22}} I_1 = Z'_{11}I_1 \quad (3.21)$$

More generally, we can reduce an N -port network to an $(N - 1)$ -port network by connecting a load impedance Z_L across port i . In this case, the impedance matrix $\mathbf{Z}$ defines a system of N linear equations

$$V_j = \sum_k Z_{jk} I_k \quad (3.22)$$

where Z_{jk} is the entry in the j -th row and k -th column of the impedance matrix. We can eliminate V_i and I_i from these equations, as in the two-port example. We add the additional

²The minus sign here accounts for the different directions of the currents, which are both defined as going into their respective ports.

constraint

$$V_i = -Z_L I_i = \sum_k Z_{ik} I_k \quad (3.23)$$

where we can pull I_i out of the sum so that it is expressed as

$$I_i = -\frac{1}{Z_L + Z_{ii}} \sum_{k \neq i} Z_{ik} I_k \quad (3.24)$$

Now that we have an expression for I_i , we can substitute this into each equation for V_j so that

$$V_j = \sum_{k \neq i} Z_{jk} I_k + Z_{ji} I_i \quad (3.25)$$

$$= \sum_{k \neq i} Z_{jk} I_k - \frac{Z_{ji}}{Z_L + Z_{ii}} \sum_{k \neq i} Z_{ik} I_k \quad (3.26)$$

$$= \sum_{k \neq i} \left(Z_{jk} - \frac{Z_{ji} Z_{ik}}{Z_L + Z_{ii}} \right) I_k \quad (3.27)$$

and we have an expression for the voltage-current relationship between the remaining ports. We can apply a similar analysis when connecting two multi-port networks together at one or more of their ports. The general prescription is that each connection adds two additional constraint equations that can be used to eliminate the voltage/current variables at those ports. In the case of cascaded two-port networks, there is a particular simple way to derive the resulting network parameters using the so-called ABCD matrices [164], which can also be generalized to cascaded 2N-port networks. However, analyzing the resulting parameters for arbitrary connections between multi-port networks is most easily done with computer-aided RF design software such as `scikit-rf` [8] or `JosephsonCircuits.jl` [162].

This type of microwave network analysis can be a valuable tool for simplifying the design and simulation of superconducting quantum circuits in a modular way. Smaller subsystems can be simulated in isolation with more expensive FEM methods and the resulting network parameters can then be connected together to yield the behavior of the full system. As long as all interactions between the subsystems occur at the well-defined port boundaries, the behavior of the overall system will be well-described using such a method. Second, many of the distributed microwave structures present on superconducting qubit devices do not have exact lumped circuit representations, and require a full impedance matrix to properly describe their behavior over a wide range of frequencies.

However, circuit quantization is most easily done with lumped circuits. In order to get around this and account for the nonlinearity of the Josephson junction, we typically turn to *black-box quantization* (BBQ) [151] or the related *energy-participation ratio* (EPR) [135] methods. Here, the idea is to treat the linear portion of the circuit as a N -port network where each port is connected to one of the N Josephson junctions in the system. The behavior of the network can then be obtained from full-wave FEM simulation or circuit analysis, and

the resulting linear network can be synthesized into a equivalent lumped circuit model that is amenable to quantization.

This network synthesis is an important but often misunderstood portion of these procedures, and different approaches have been taken to implement this step over the years. A modified version of Foster's synthesis [62, 16], valid for small dissipation, was initially used in Ref. [151] to synthesize an equivalent circuit for subsequent quantization. Here, each pole of an arbitrary impedance function $Z(s)$ is approximated as a parallel RLC circuit with resonance frequency $\omega_p \approx (L_p C_p)^{-\frac{1}{2}}$ and decay rate $\kappa_p = (R_p C_p)^{-1}$, so that we have

$$Z(s) \approx Z_F(s) = \sum_{p=1}^M Y_p(s)^{-1} \quad Y_p(s) = \frac{1}{R_p} + sC_p + \frac{1}{sL_p} \quad (3.28)$$

For s near a pole $s_n = -\frac{\kappa_n}{2} \pm i\omega_n$, the term $Y_n(s)^{-1}$ dominates and we can approximate $Z_F(s) \approx Y_n(s)^{-1}$ or equivalently $Y_F(s) \approx Y_n(s)$. Then, we see that

$$Y'_F(\omega_n) \approx Y'_n(\omega_n) = iC_n - \frac{1}{i\omega_n^2 L_n} = i \left[C_n + \frac{1}{\omega_n^2 L_n} \right] \approx i \left[C_n + \frac{L_n C_n}{L_n} \right] = 2iC_n \quad (3.29)$$

and also that

$$\text{Re}[Y_F(\omega_n)] \approx \text{Re}[Y_n(\omega_n)] = \frac{1}{R_n} \quad (3.30)$$

This implies that we can determine the effective resistance and capacitance of the equivalent RLC circuit corresponding to a given pole of the impedance function as

$$R_n \approx \frac{1}{\text{Re}[Y(\omega_n)]} \quad C_n \approx \frac{1}{2} \text{Im}[Y'_F(\omega_n)] \quad (3.31)$$

which gives us a very straightforward method to determine the effective lumped circuit parameters corresponding to the resonance directly from the admittance function $Y(\omega)$. We now make several comments about this method.

First, it is important for $Q = R_n/Z_n$ to be large since ω_n will differ from $(L_n C_n)^{-\frac{1}{2}}$ at order $1/Q^{23}$. Second, it is important to ensure that the resulting equivalent circuit yields an impedance $Z_F(s)$ that is a good approximation of the true impedance $Z(s)$, at least near the poles. Otherwise, the extracted parameters will not be accurate. In these cases, we can resort to other network synthesis methods that have been developed, such as Brune's synthesis [25], which was used by Ref. [188] to implement black-box quantization. Finally, we stress that these circuit synthesis methods do not require one to obtain the original impedance function via expensive full-wave simulations, and are equally applicable to impedances derived via circuit analysis⁴. This common misconception is likely due to its use in many early papers to

³See Eq. 2.111. This is generally fine, since even for $\kappa/2\pi = 50$ MHz on a resonator with frequency $\omega/2\pi = 5$ GHz yields an error of just 62 kHz.

⁴For example, this is used in the superconducting circuit analysis package **QuCAT** [70].




Figure 3.3: **Transmission line model.** (a) Lumped-element approximation for a 1D transmission line. The full transmission line model is derived by taking the continuum limit $\Delta x \rightarrow 0$ for a fixed total length $N\Delta x = d$. (b) Given the parameters Z_0 and v (or equivalently ℓ and c), the behavior of a transmission line segment of length d can be modeled as a two-port network. (c) A 2D cross-section of a coplanar-waveguide (CPW) transmission line, parameterized by a width w and a gap g . The dark gray denotes the substrate, and the dark blue denotes the superconducting metal. The transmission line parameters ℓ and c can be simulated from the 2D structure using tools like Ansys Q2D.

derive equivalent circuit models and perform black-box quantization of 3D cavity structures [151, 188, 135] that do require such simulations.

Once we have an equivalent circuit model representing a best fit (or exact synthesis) to the original impedance response of the network, we can proceed to quantize the circuit with any Josephson junctions included. We will not explain the process in detail here, but we encourage the reader to consult Chapter 5 of Ref. [39] for a review of the subject. Instead, we will spend the remainder of this chapter explaining how we can more accurately and quickly obtain the impedance response of a complex distributed element circuit for subsequent quantization.

3.3 The CPW Transmission Line

The coplanar waveguide (CPW) transmission line is ubiquitous in superconducting qubit devices—used to deliver signals across a chip, to construct filters, and to form resonators for both readout and storage of quantum information. In this section, we describe how to model such structures in a realistic system that may include connections to other circuit elements.

We consider a 1D transmission line, which can be modeled as an infinite chain of series inductors $L = \ell\Delta x$ and parallel capacitances $C = c\Delta x$ with inductance per unit length ℓ and capacitance per unit length c . We can write the circuit Lagrangian for this system as

$$\mathcal{L} = \sum_i \frac{1}{2} C \dot{\Phi}^2 - \frac{1}{2L} (\Phi_{i+1} - \Phi_i)^2 = \sum_i \Delta x \left[\frac{1}{2} c \dot{\Phi}^2 - \frac{1}{2\ell} \left(\frac{\Phi_{i+1} - \Phi_i}{\Delta x} \right)^2 \right] \quad (3.32)$$

If we take the continuum limit ($\Delta x \rightarrow 0$), we obtain the continuous version of the Lagrangian,

$$\mathcal{L} = \int dx \left[\frac{c}{2} (\partial_t \Phi)^2 - \frac{1}{2\ell} (\partial_x \Phi)^2 \right]. \quad (3.33)$$

Using the Euler-Lagrange equations to solve for the equations of motions:

$$\sum_{\mu=x,t} \partial_\mu \left[\frac{\partial \mathcal{L}}{\partial (\partial_\mu \Phi)} \right] = \frac{\partial \mathcal{L}}{\partial \Phi} \quad (3.34)$$

we then obtain a 1-D wave equation

$$c \partial_t^2 \Phi - \frac{1}{\ell} \partial_x^2 \Phi = \partial_t^2 \Phi - \frac{1}{\ell c} \partial_x^2 \Phi = 0. \quad (3.35)$$

The solutions to this equation are of the form

$$\Phi(x, t) = \Phi_\pm e^{i(\pm kx - \omega t)} \quad (3.36)$$

where $+$ indicates a rightward traveling wave and $-$ indicates a leftward traveling wave. Plugging this into the wave equation, we find that

$$-\omega^2 + \frac{1}{\ell c} k^2 = 0 \rightarrow \omega = \pm k\nu \quad (3.37)$$

where $\nu = \frac{1}{\sqrt{\ell c}}$ is the phase velocity. The conjugate momentum to the flux variable $\Phi(x, t)$ can also be derived from the Lagrangian as

$$Q(x, t) = \frac{\partial \mathcal{L}}{\partial (\partial_t \Phi)} = \int c \partial_t \Phi dx = \int cV(x, t) dx = \int q(x, t) dx \quad (3.38)$$

where we define the charge density as

$$q(x, t) = c \partial_t \Phi(x, t) = -i\omega c \Phi_\pm e^{i(\pm kx - \omega t)} \quad (3.39)$$

Then, we can integrate this over the length of the transmission line to find the total charge as a function of position:

$$Q(x, t) = \int q(x', t) dx' = \mp \frac{\omega c}{k} \phi_\pm e^{i(\pm kx - \omega t)} = Q_\pm e^{i(\pm kx - \omega t)} \quad (3.40)$$

The voltage and current along the transmission line can also be derived from the flux and charge by simply taking a time derivative.

The impedance of wave is defined by the ratio of the voltage to the current (or equivalently flux to charge) such that

$$Z_0 = |\phi_{\pm}/Q_{\pm}| = \frac{k}{\omega c} = \frac{1}{\nu c} = \sqrt{\frac{\ell}{c}}, \quad (3.41)$$

and we see that this give us the standard definition of the characteristic impedance Z_0 for the transmission line.

Transmission lines are an incredibly useful tool for designing microwave resonators and filters, because they are quasi-1D structures. We can describe the behavior of an arbitrary length transmission line very well with only two parameters, Z_0 and ν (or equivalently ℓ and c), that can be simulated very accurately using 2D simulators like Ansys Q2D. This is not only faster than simulating a transmission with a full-wave FEM simulator, but also more accurate.

Transmission Line as a Two-Port Network

Once the transmission line parameters Z_0 and ν are known, we can model a transmission line section of length d as a two-port network with an impedance matrix $\mathbf{Z}(d)$. Since the transmission line is both symmetric and reciprocal, we must have $Z_{11} = Z_{22}$ and $Z_{12} = Z_{21}$, so there are two parameters to compute. To do so, we consider the voltage and current profiles along the transmission line, which are given by

$$V(x) = V_+ e^{ikx} + V_- e^{-ikx} \quad (3.42)$$

$$I(x) = \frac{V_+}{Z_0} e^{ikx} - \frac{V_-}{Z_0} e^{-ikx} \quad (3.43)$$

where port 1 is at $x = -d$ and port 2 is at $x = 0$. By definition, $Z_{11} = V(-d)/I(-d)$ when port 2 is an open circuit, so we enforce the boundary condition $I(0) = 0$, and we find

$$I(0) = \frac{1}{Z_0} [V_+ - V_-] = 0 \quad \Rightarrow \quad V_+ = V_- \quad (3.44)$$

and the reflection coefficient is $\Gamma = V_-/V_+ = 1$. The impedance at port 1 is given by $Z_{11}(d) = V(-d)/I(-d)$, so we have

$$Z_{11}(d) = \frac{V(-d)}{I(-d)} = Z_0 \frac{e^{ikd} + e^{-ikd}}{e^{ikd} - e^{-ikd}} = Z_0 \frac{2 \cos(kd)}{2i \sin(kd)} = -iZ_0 \cot(kd) \quad (3.45)$$

Similarly, we consider the voltage-current relation for Z_{21} . Here the boundary conditions is still $I(0) = 0$, and we find

$$Z_{21}(d) = \frac{V(0)}{I(-d)} = Z_0 \frac{1 + 1}{e^{ikd} - e^{-ikd}} = \frac{Z_0}{i \sin(kd)} = -iZ_0 \csc(kd) \quad (3.46)$$




Figure 3.4: **Transmission line as an impedance transformer.** A transmission line section transforms a (possibly frequency dependent) load impedance Z_L to a load impedance $Z'_L(\omega)$ given by Eq. 3.49. For $\omega_0 = \frac{\pi\nu}{2d}$ ($d = \lambda/4$), the transmission line section acts as an impedance inverter, taking $Z_L = \infty \rightarrow Z'_L(\omega_0) = 0$ and vice versa.

and we can write the full impedance matrix as

$$\mathbf{Z}(\theta) = -iZ_0 \begin{pmatrix} \cot(\theta) & \csc(\theta) \\ \csc(\theta) & \cot(\theta) \end{pmatrix} \quad (3.47)$$

where we define $\theta \equiv kd$. The admittance matrix can be found by inverting this matrix, and is given by

$$\mathbf{Y}(\theta) = \mathbf{Z}^{-1}(\theta) = -iY_0 \begin{pmatrix} \cot \theta & -\csc \theta \\ -\csc \theta & \cot \theta \end{pmatrix} \quad (3.48)$$

where $Y_0 = Z_0^{-1}$ is the characteristic admittance.

Transmission Line as an Impedance Transformer

Using our previous result from Eq. 3.21, we can now analyze the effective input impedance of a transmission line terminated by an arbitrary load impedance Z_L . This is given by

$$Z'_L(\theta) = Z_{11} - \frac{Z_{12}Z_{21}}{Z_L + Z_{22}} = \frac{Z_{11}Z_L + Z_{11}Z_{22} - Z_{12}Z_{21}}{Z_L + Z_{22}} \quad (3.49a)$$

$$= \frac{-iZ_0Z_L \cot \theta - Z_0^2 \cot^2 \theta + Z_0^2 \csc^2 \theta}{Z_L - iZ_0 \cot \theta} \quad (3.49b)$$

$$= Z_0 \frac{Z_0 - iZ_L \cot \theta}{Z_L - iZ_0 \cot \theta} \quad (3.49c)$$

$$= Z_0 \frac{Z_L + iZ_0 \tan \theta}{Z_0 + iZ_L \tan \theta} \quad (3.49d)$$

It is instructive to look at the two cases where $d = \lambda/4 \Rightarrow \theta = \pi/2$ and $d = \lambda/2 \Rightarrow \theta = \pi$. In the first case, we have

$$Z'_L(\pi/2) = \frac{Z_0^2}{Z_L} \quad (3.50)$$




Figure 3.5: **Transmission line resonators.** (a,c) A circuit diagram depicting a $\lambda/2$ (a) and $\lambda/4$ (c) resonator. The voltage and current profiles for the first 4 modes of the resonator are shown above and below the transmission line. The open boundary condition ($Z_L = \infty$) at $x = d$ enforces $I(d) = 0$, while the shorted boundary condition ($Z_L = 0$) at $x = d$ enforces $V(d) = 0$ (b,d) The imaginary component of the admittance $\text{Im}[Y(\omega)]$ as seen through the port at $x = 0$. The zeros correspond to the resonance frequencies. The admittance function can be fit to a lossy Foster circuit to extract equivalent lumped circuit parameters for each mode.

so the $\lambda/4$ section transforms a short ($Z_L = 0$) to an open circuit ($Z_L = \infty$) and vice versa at a frequency $\omega = \frac{\pi\nu}{2d}$. This is why a quarter-wave transmission line section is often referred to as a quarter-wave impedance transformer. On the other hand, we find that $Z'_L(\pi) \rightarrow Z_L$, so the load impedance is unchanged at $\omega = \frac{\pi\nu}{\lambda}$.

Boundary Conditions and Transmission Line Resonators

When both ends of a transmission line section are terminated with a load impedance, the transmission line section can no longer support a continuous spectrum of modes. Instead, the spectrum is modified such that there exists only a countable number of discrete resonances. We start by considering a transmission line segment of length $\theta = kd$, terminated in an open

circuit on both ends. Looking in from one port, the impedance and admittance are given by

$$Z_{\text{in}}(\theta) = -iZ_0 \cot \theta \quad Y_{\text{in}}(\theta) = iY_0 \tan \theta \quad (3.51)$$

The imaginary zeros of the admittance occur at $\theta_n = k_n d = n\pi$ for $n \in \mathbb{Z}$ and we can fit these zeros to the lossy Foster circuit described in the previous section. This yields the following lumped circuit parameters for the equivalent LC circuit [179]:

$$\omega_n = \nu k_n = \frac{n\pi\nu}{d} \quad (3.52a)$$

$$C_n = \frac{1}{2} \text{Im}[Y'(\omega_n)] = \frac{d}{2\nu Z_0} \sec^2 \left(\frac{\omega_n d}{\nu} \right) = \frac{1}{2} \frac{d}{\nu Z_0} = \frac{1}{2} cd \quad (3.52b)$$

$$L_n = \frac{1}{\omega_n^2 C_n} = \frac{2Z_0}{k_n^2 \nu d} = \frac{2Z_0 d}{\nu n^2 \pi^2} = \frac{2\ell d}{n^2 \pi^2} \quad (3.52c)$$

$$Z_{r,n} = \sqrt{\frac{L_n}{C_n}} = \sqrt{\frac{4Z_0^2}{k_n^2 d^2}} = \frac{2Z_0}{k_n d} = \frac{2Z_0}{n\pi} \quad (3.52d)$$

Importantly, we see that since the equivalent capacitance is constant for each resonance, the impedance of each resonance scales inversely with the mode number. Finally, we note that because the fundamental resonance ($n = 1$) occurs for a length $d = \frac{\pi}{k_1} = \frac{\lambda}{2}$, this is typically referred to as a $\lambda/2$ or a half-wave resonator.

We can also consider a transmission line segment that is terminated by a short circuit on one end. From Eq. 3.49, we find that for $Z_L = 0$, the input impedance for the transmission line section is:

$$Z_{\text{in}}(\theta) = iZ_0 \tan \theta \quad Y_{\text{in}}(\theta) = -iY_0 \cot \theta \quad (3.53)$$

and the imaginary zeros of the admittance now occur at $\theta_n = k_n d = (n + \frac{1}{2})\pi$. Since the fundamental resonance ($n = 0$) now occurs for a length $d = \frac{\pi}{2k_0} = \frac{\lambda}{4}$, this is typically referred to as a $\lambda/4$ or a quarter-wave resonator. For this resonance, we can derive the following equivalent circuit parameters:

$$\omega_n = \nu k_n = \frac{(2n + 1)\pi\nu}{2d} \quad (3.54a)$$

$$C_n = \frac{1}{2} \text{Im}[Y'(\omega_n)] = \frac{d}{2\nu Z_0} \csc^2 \left(\frac{\omega_n d}{\nu} \right) = \frac{1}{2} \frac{d}{\nu Z_0} = \frac{1}{2} cd \quad (3.54b)$$

$$L_n = \frac{2Z_0}{k_n^2 \nu d} = \frac{8Z_0 d}{\nu(2n + 1)^2 \pi^2} = \frac{8\ell d}{(2n + 1)^2 \pi^2} \quad (3.54c)$$

$$Z_{r,n} = \frac{2Z_0}{k_n d} = \frac{4Z_0}{(2n + 1)\pi} \quad (3.54d)$$

Having derived the equivalent circuit models for the modes in both a $\lambda/4$ and a $\lambda/2$ transmission line resonator, we now make several comments. A common mistake when constructing an equivalent lumped circuit for a transmission line resonator is to assume that the




Figure 3.6: **Capacitive loading of a CPW resonator.** (a) A $\lambda/4$ CPW resonator, with a capacitance C_L , loading the open end of the CPW. (b) The mode frequencies, denoted by the white squares, are shifted down by the capacitance. Here we show the mode frequencies for $d_L = C_L/c = 0.1$. In the limit $C_L \rightarrow \infty$, the resonator becomes a $\lambda/2$ resonator with shorted boundary conditions. (c) The voltage profile of the fundamental mode as a function of the capacitance $C_L = cd_L$. The capacitance adds an effective length to the CPW, denoted by the dotted lines.

impedance of the resonance is the same as the characteristic impedance of the transmission line. However, we have shown that this is in fact incorrect. Rather, the impedance of the resonance is scaled by $2/\theta_n$ and scales inversely with the mode number. Moreover, we find that for the same frequency ω , the fundamental mode of the quarter-wave resonator has an impedance that is twice that of the effective impedance of the half-wave resonator. Since the capacitive coupling between two resonators scales with the impedance of the modes, this implies that the coupling between two quarter-wave resonators will be twice as high for the same coupling capacitance C_c . Finally, we note that a $\lambda/2$ resonator can also be constructed by shorting both ends of the transmission line segment.

While the transmission line resonances discussed above have simple solutions, they are often not good approximations for realistic circuit structures. In practice, most transmission line resonators are loaded by additional lumped capacitances at the points where they are coupled to the rest of the circuit. To illustrate the loading effect of this capacitance, we consider adding a capacitor C_L in parallel to the open end of a $\lambda/4$ resonator. The total input admittance for the loaded transmission line is now

$$Y_{\text{in}}(\theta) = i\omega C_L - iY_0 \cot \theta = iY_0 \theta \frac{d_L}{d} - iY_0 \cot \theta \quad (3.55)$$

where we define $d_L = C_L/c$ to be the effective length of transmission line needed to yield




Figure 3.7: **CPW admittance as a function of coupling location.** (a) The admittance response of a CPW resonator changes depending on the coupling location along the CPW length. (b) The admittance for two different values of x . We see that when the admittance is viewed from the midpoint of the half-wave $x = 0.5$ (green), there are no zeros in the admittance $Y(\omega)$ at $n\pi$ for odd n , because these modes have zero voltage at this location. Similarly, for $x = 0.25$ (blue), the mode at $\theta = 2\pi$ is decoupled from the port. The slope of the admittance at each zero also changes depending on the coupling point x .

the same capacitance. The mode frequencies are given by $Y_{\text{in}}(\theta_n) = 0$, which is now a transcendental equation, and can no longer be solved analytically. However, as we can see from Fig. 3.6, the solutions for each mode are now shifted down in frequency. This is expected since capacitances add in parallel, and the effective lumped capacitance of each mode is now

$$C_n = \frac{1}{2} \text{Im}[Y'_{\text{in}}(\omega_n)] = \frac{1}{2}(C_L + cd) \quad (3.56)$$

It is worth noting that this type of an analysis is not restricted to just capacitive loading. The impedance $Y_L = i\omega C_L$ could be the admittance of any one-port network that is added in parallel with the $\lambda/4$ CPW resonator.

As another illustrative example, we now examine another consequence of the transmission line. Because the voltage and current profiles of a CPW resonator have a spatially varying envelope, both the effective coupling strength and the loaded quality factor now depend on both the precise location and form of the coupling (capacitive or inductive). Because the capacitive energy is proportional to CV_1V_2 , the coupling will be scaled with the magnitude of the voltage profile at a given location, and will become zero at a voltage node. Similarly, the inductive energy is proportional to LI_1I_2 , so the coupling will scale with the magnitude of the current profile. For example, we can consider coupling to a $\lambda/2$ resonator with open boundary conditions, but exactly at the midpoint. Here, we can model the system as two transmission line sections in parallel, each with length $d/2$, so the total admittance looking

in through the coupling point is

$$Y_{\text{in}}(\theta) = iY_0 \tan \frac{\theta}{2} + iY_0 \tan \frac{\theta}{2} = 2iY_0 \tan \frac{\theta}{2} \quad (3.57)$$

and we see that the zeros now occur at $\theta_n = 2\pi n$ instead of πn . The fundamental mode of the CPW resonator has not changed and is still at $\omega_n = \frac{\pi\nu}{d}$; however, the odd- n modes are all uncoupled from the perspective of a port at the midpoint! More generally, we could consider coupling to a port at a distance xd from one end ($0 \leq x \leq 1$). Then, the total admittance from the perspective of the coupling point is

$$Y_{\text{in}}(\theta) = iY_0 \tan(x\theta) + iY_0 \tan((1-x)\theta) = \frac{iY_0 \sin \theta}{\cos(x\theta) \cos((1-x)\theta)} \quad (3.58)$$

Then we see that the zeros of Y_{in} remain at $\theta_n = n\pi$, except when $\cos(x\theta_n) = 0$ or $\cos((1-x)\theta_n) = 0$, in which case the mode is completely decoupled. The effective capacitance of the n -th mode is given by

$$C_n(x) = \frac{1}{2} \text{Im}[Y'_{\text{in}}(\omega_n)] = \frac{1}{2} cd \left[\frac{x}{\cos^2(x\theta_n)} + \frac{1-x}{\cos^2((1-x)\theta_n)} \right] = \frac{1}{2} \frac{cd}{\cos^2(x\theta_n)} \quad (3.59)$$

where in the last equality we use the fact that $\theta_n = n\pi \Rightarrow \cos((1-x)\theta_n) = \pm \cos(x\theta_n)$. This tells us that the effective capacitive coupling to the mode as a function of the coupling point scales as

$$g(x) \propto \frac{1}{\sqrt{C_n(x)}} \propto |\cos(xn\pi)| \quad (3.60)$$

which is simply the shape of the voltage profile for the n -th mode.

Here, we stress that the effective capacitance and inductance of the equivalent circuit models in Eq. 3.52a and Eq. 3.54a are only well-defined up to a specified port location. The values of these lumped equivalent models are selected to yield an impedance or admittance response that is a best-fit model to the true impedance or admittance response. Since changing the location of the port changes the admittance response, this also changes the equivalent lumped circuit as viewed from the port. As a result, we must be careful when replacing distributed elements with lumped equivalent models, as doing so improperly can drastically change the derived circuit or Hamiltonian parameters for a given design.

Analyzing Purcell Loss

Now that we understand how to analyze and fit transmission line models to effective lumped models, we will work through another example that illustrates how interference effects in a transmission line can be leveraged to our benefit. In most superconducting qubit systems, qubits are dispersively coupled to a resonator in order to measure a qubit state. These read-out resonators are typically strongly coupled to the environment to enable fast measurement; however, this can cause excess qubit relaxation mediated by the resonator's coupling to the

environment. This *Purcell effect* [165], as it is typically referred to, is really a classical effect and can be modeled using the network analysis methods we have discussed so far [85].

In the simplest model, we have a qubit capacitively coupled to a lumped parallel LC resonator. The resonator decay rate is mediated by a capacitive coupling to a $Z = 50 \, \Omega$ semi-infinite transmission line, which we replace with a $R = 50 \, \Omega$ resistance. In this case, the effective decay rate of the qubit has been shown to be [211, 111]:

$$\Gamma_P = \kappa \left(\frac{g^2}{\omega_q - \omega_r} \right)^2 \left(\frac{\omega_q}{\omega_r} \right)^3 \left(\frac{2\omega_q}{\omega_q + \omega_r} \right) \quad (3.61)$$

However, it is also straightforward to numerically compute the decay rate $\Gamma_P = (R_q C_q)^{-1}$ by fitting the pole of the impedance $Z(\omega)$ from the perspective of the qubit node to its Foster equivalent circuit. This is given by

$$Y(\omega) = i\omega C_q + \frac{1}{i\omega L_q} + \left[\frac{1}{i\omega C_c} + Z_r(\omega) \right]^{-1} \quad (3.62)$$

and enables a simple comparison of different readout resonator circuitry by simply changing the resonator impedance response $Z_r(\omega)$ ⁵. To illustrate this, we compare the normalized Purcell loss $\Gamma_p/\kappa_r = 1/(\kappa_r T_1)$ for both the simple lumped element resonator as well as a $\lambda/2$ CPW resonator. For the two systems, we have

$$Z_{r,\text{lumped}}(\omega) = \frac{1}{i\omega C_r + \frac{1}{i\omega L_r} + Y_\kappa(\omega)} \quad (3.63)$$

$$Z_{r,\text{cpw}}(\omega) = Z_0 \frac{Z_\kappa(\omega) + iZ_0 \tan\left(\frac{\omega d}{\nu}\right)}{Z_0 + iZ_\kappa(\omega) \tan\left(\frac{\omega d}{\nu}\right)} \quad (3.64)$$

$$Z_\kappa(\omega) = \frac{1}{Y_\kappa(\omega)} = \frac{1}{i\omega C_\kappa} + Z_{\text{env}} \quad (3.65)$$

where we use the expression from Eq. 3.49 to derive the impedance of the loaded CPW transmission line.

In Fig. 3.8, we show Γ_p/κ_p as a function of the qubit frequency. We see that the decay rate of the qubit is much higher for the CPW resonator at frequencies above the fundamental mode, where the qubit can couple to higher modes of the CPW resonator. This was first studied in Ref. [85]. However, we recall from the last section that the admittance response of the CPW resonator can vary significantly depending on where the coupling ports are. In Fig. 3.8, we show an alternative coupling geometry, where the resonator couples to the $50 \, \Omega$ environment at a point d_s away from the open boundary. From the perspective of the qubit port, we now have a open stub of length d_s in parallel with the environmental impedance that

⁵In practice we also drop the $\frac{1}{i\omega L_q}$ term since it does not contribute the $\text{Re}[Y(\omega)]$. In this case, the total effective capacitance of the qubit mode is modified to $C'_q = \text{Im}[Y'(\omega)]$, since the inductive term no longer contributes.




Figure 3.8: **Purcell loss for lumped and distributed resonators.** We simulated two circuits representing a $\omega_r/(2\pi) \approx 6$ GHz resonator capacitively coupled to a qubit port via a coupling capacitance $C_c = 10$ fF, which results in a coupling $g/2\pi \approx 96$ MHz at $\omega_q = \omega_r$. The qubit capacitance is $C_q = 100$ fF, and the lumped resonator parameters were chosen to approximate the fundamental mode of the CPW resonator. The resonator is coupled to a 50Ω environment via a capacitance $C_k = 12.5$ fF, which results in an effective decay rate $\kappa_r/(2\pi) \approx 2$ MHz. The simulated circuits are shown for the (a) lumped resonator and (b) CPW resonator. (c) The Purcell enhancement/suppression factor Γ_p/κ_p is plotted for both the lumped and CPW resonators.

is introducing a resistive loss. Since a $\lambda/4$ open stub looks like a short circuit at $\omega_s = \frac{\pi\nu}{2d_s}$, this effectively hides the environment from the rest of the circuit and acts as a notch filter near ω_s . This circuit was first introduced in Ref. [191], and is referred to as an *intrinsic* Purcell filter, since the filter relies on interference effects within the resonator itself and does not introduce any additional structures. Fig. 3.8 also shows the resulting Purcell suppression factor for the intrinsic Purcell filter compared to both the lumped resonator and the $\lambda/2$ CPW resonator that is coupled to the environment at its voltage antinode. We see that this substantially improves the Purcell loss at frequencies near the resonator frequency, at the cost of worse Purcell suppression at lower frequencies⁶. This type of $\lambda/4$ stub filter can

⁶The other drawback with this design is that it requires a much larger C_k to achieve the same resonator decay rate κ_r , since the coupling capacitor is now closer to the voltage node of the $\lambda/2$ resonator.




Figure 3.9: **Coupled transmission line model.** (a) Lumped-element approximation for two coupled 1D transmission lines. The full model is derived by taking the continuum limit $\Delta x \rightarrow 0$ for a fixed total length $N\Delta x = d$. (b) The coupled lines are most easily described by the impedance $Z_{0,i}$ and phase velocity ν_i in the eigenbasis of the lines. Given the parameters $Z_{0,o}$, ν_o and $Z_{0,e}$, ν_e , the behavior of a transmission line segment of length d can be modeled as a two-port network. These parameters can be derived from the bare circuit parameters c , c_c , ℓ , m . (c) A 2D cross-section of two CPW transmission lines in close proximity. The individual lines are separated by a center-to-center distance s . In general, the two lines do not need to have the same CPW width w and gap g . The transmission line parameters can again be simulated from the 2D structure using tools like Ansys Q2D.

also be engineered into the resonator feedline to add a second notch, and this strategy was employed in the “Everglades” design that was used for the experiments in Chapter 6.

3.4 Coupled Line Models

While distributed transmission lines structures can be coupled to other circuits via lumped capacitances and inductances, it is often more natural to use a distributed form of coupling where two transmission lines are brought into close proximity, leading to a mutual capacitance inductance per unit length that couples the two lines. For an exhaustive treatment on coupled-line analysis, see Ref. [160]; however, we will walk through a derivation for a simple case where we have only two coupled transmission lines.

We model the coupled lines in the same as we do for the single transmission line in Sec. 3.3, but with an additional mutual inductance $M = m\Delta x$ cross capacitance $C_c = c_c\Delta x$. The

Lagrangian in the continuum limit can be written as

$$\mathcal{L} = \int dx \left[\frac{1}{2} \dot{\Phi}^T \mathbf{C} \dot{\Phi} - \frac{1}{2} \partial_x \Phi^T \mathbf{L}^{-1} \partial_x \Phi \right], \quad (3.66)$$

where the node fluxes, capacitance matrix, and inductance matrix are given by

$$\Phi = \begin{pmatrix} \phi_1 \\ \phi_2 \end{pmatrix} \quad \mathbf{C} = \begin{pmatrix} c_1 + c_c & -c_c \\ -c_c & c_2 + c_c \end{pmatrix} \quad \mathbf{L} = \begin{pmatrix} L_1 & M \\ M & L_2 \end{pmatrix}$$

We once again derive the equations of motion from the Euler-Lagrange equation, which now yields a pair of coupled differential equations,

$$\mathbf{C} \partial_t \Phi - \mathbf{L}^{-1} \partial_x \Phi = \partial_t \Phi - (\mathbf{LC})^{-1} \partial_x \Phi = 0 \quad (3.67)$$

We then look at solutions of the form

$$\Phi_{\pm} = \begin{pmatrix} \phi_{1,\pm} \\ \phi_{2,\pm} \end{pmatrix} e^{i(\pm kx - \omega t)}, \quad (3.68)$$

which results in the eigenvalue equation

$$[-\omega^2 \mathbf{I} + k^2 (\mathbf{LC})^{-1}] \begin{pmatrix} \phi_{1,\pm} \\ \phi_{2,\pm} \end{pmatrix} = [(\mathbf{LC})^{-1} - \nu^2 \mathbf{I}] \begin{pmatrix} \phi_{1,\pm} \\ \phi_{2,\pm} \end{pmatrix} = 0 \quad (3.69)$$

We can interpret the eigenvectors as the ratio of the flux amplitudes on each transmission line for the two normal modes, and the eigenvalues ν as the corresponding phase velocities for the normal modes. In the case of identical transmission lines ($\ell = \ell_1 = \ell_2$ and $c = c_1 = c_2$), the matrix $(\mathbf{LC})^{-1}$ reduces to

$$(\mathbf{LC})^{-1} = \frac{1}{(\ell^2 - m^2)(c^2 + 2cc_c)} \begin{pmatrix} \ell(c + c_c) - mc_c & \ell c_c - m(c_c + c) \\ \ell c_c - m(c_c + c) & \ell(c + c_c) - mc_c \end{pmatrix} \quad (3.70)$$

From the symmetry of the matrix $(\mathbf{LC}^{-1})$, we see that there should be an even (symmetric) or odd (anti-symmetric) eigenmode, with eigenvalues

$$\nu_e = \sqrt{\frac{(\ell - m)(c + 2c_c)}{(\ell + m)(\ell - m)(c + 2c_c)c}} = \frac{1}{\sqrt{(\ell + m)c}} \quad (3.71)$$

$$\nu_o = \sqrt{\frac{(\ell + m)c}{(\ell + m)(\ell - m)(c + 2c_c)c}} = \frac{1}{\sqrt{(\ell - m)(c + 2c_c)}} \quad (3.72)$$

If we take the limit where $m, c_c \rightarrow 0$, we see that $\nu_e = \nu_o = \frac{1}{\sqrt{\ell c}}$, and we just have two identical and uncoupled transmission lines.

Now we look at the impedances for the two modes. To do this, we first write out the two charge fields as

$$\mathbf{Q} = \nabla_{\partial_t \Phi} \mathcal{L} = \int dx \mathbf{C} \partial_t \Phi = \mp \frac{\omega}{k} \mathbf{C} \Phi \quad (3.73)$$

and similarly to the single transmission line case, we find

$$\mathbf{Z} = \mp \nu^{-1} \mathbf{C}^{-1} \quad (3.74)$$

For the case of two symmetrically coupled lines, we can write the inverse capacitance matrix $\mathbf{C}^{-1}$ in the eigenbasis

$$\tilde{\mathbf{C}}^{-1} = \mathbf{u}^T \mathbf{C}^{-1} \mathbf{u} \quad (3.75)$$

$$= \frac{1}{2c(c+2c_c)} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} c+c_c & c_c \\ c_c & c+c_c \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \quad (3.76)$$

$$= \begin{pmatrix} \frac{1}{c} & 0 \\ 0 & \frac{1}{c+2c_c} \end{pmatrix} \quad (3.77)$$

and we see that

$$\tilde{\mathbf{Z}} = \begin{pmatrix} \sqrt{\frac{\ell+m}{c}} & 0 \\ 0 & \sqrt{\frac{\ell-m}{c+2c_c}} \end{pmatrix} = \begin{pmatrix} Z_{0,e} & 0 \\ 0 & Z_{0,o} \end{pmatrix} \quad (3.78)$$

Coupled Lines as a 4-Port Network

Given a segment consisting of two coupled lines, we can treat the ends of each individual transmission line as a port and the coupled-line segment forms a 4-port network (see Fig. 3.9(b)). In the even/odd mode basis, the two “lines” are decoupled, so we can take our result from the single-transmission line and apply it to each mode. Then, the impedance matrix should be block diagonal

$$\tilde{\mathbf{Z}} = \begin{pmatrix} \tilde{\mathbf{Z}}_e & \mathbf{0} \\ \mathbf{0} & \tilde{\mathbf{Z}}_o \end{pmatrix} \quad (3.79)$$

where the even and odd mode impedance matrices are

$$\tilde{\mathbf{Z}}_i = \begin{pmatrix} -iZ_{0,i} \cot(\theta_i) & -iZ_{0,i} \csc(\theta_i) \\ -iZ_{0,i} \csc(\theta_i) & -iZ_{0,i} \cot(\theta_i) \end{pmatrix} \quad (3.80)$$

for $i = e, o$, where $\theta_i = k_i d$ is the electrical length for each mode (these are generally not the same). Here, the rows/columns correspond to $e_{\text{in}}, e_{\text{out}}, o_{\text{in}}, o_{\text{out}}$, so we can perform a basis transformation to express the impedance matrix in the physical port basis shown in Fig. 3.9(b) as $\mathbf{Z} = \mathbf{u} \tilde{\mathbf{Z}} \mathbf{u}^T$, where

$$\mathbf{u} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & -1 & 0 \\ 0 & 1 & 0 & -1 \end{pmatrix} \quad (3.81)$$




Figure 3.10: **Modeling CPW Meanders.** (a) A meandered CPW transmission line. The center-to-center spacing is denoted by s . The tight spacing relative to the CPW dimensions leads to cross-coupling between different sections of the line and modifies the impedance response compared to an ideal transmission line. (b) We model the meandered transmission line by connecting adjacent ports of a set of N coupled transmission lines.

and we obtain the following terms for Z_{ij} :

$$Z_{11} = Z_{22} = Z_{33} = Z_{44} = -\frac{1}{2}i(Z_{0,e} \cot(\theta_e) + Z_{0,o} \cot(\theta_o)) \quad (3.82)$$

$$Z_{12} = Z_{21} = Z_{34} = Z_{43} = -\frac{1}{2}i(Z_{0,e} \csc(\theta_e) + Z_{0,o} \csc(\theta_o)) \quad (3.83)$$

$$Z_{13} = Z_{31} = Z_{24} = Z_{42} = -\frac{1}{2}i(Z_{0,e} \cot(\theta_e) - Z_{0,o} \cot(\theta_o)) \quad (3.84)$$

$$Z_{14} = Z_{41} = Z_{23} = Z_{32} = -\frac{1}{2}i(Z_{0,e} \csc(\theta_e) - Z_{0,o} \csc(\theta_o)) \quad (3.85)$$

3.5 Modeling CPW Meanders

Another particularly useful application of coupled-line theory is for understanding the effect of tightly meandered segments on the frequency of a CPW resonator. As shown in Fig. 3.10, long CPW segments are typically meandered tightly to reduce the overall circuit footprint in a device. However, when there are no airbridges and vias between the neighboring transmission line segments, this leads to cross-coupling that changes the impedance response of the transmission line. To first order, we consider the simplest case where we have 2 coupled lines, and ports 2 and 4 are directly connected. Making this connection enforces the two

constraints $V_2 = V_4$ and $I_2 = -I_4$. Then, the voltage current relations can be written as

$$V_1 = Z_{11}I_1 + (Z_{12} - Z_{14})I_2 + Z_{13}I_3 \quad (3.86a)$$

$$V_3 = Z_{31}I_1 + (Z_{32} - Z_{34})I_2 + Z_{33}I_3 \quad (3.86b)$$

$$V_2 = Z_{21}I_1 + (Z_{22} - Z_{24})I_2 + Z_{23}I_3 \quad (3.86c)$$

$$= Z_{41}I_1 + (Z_{42} - Z_{44})I_2 + Z_{43}I_3 = V_4 \quad (3.86d)$$

From here, we can write I_2 in terms of I_1 and I_3 as

$$I_2 = \frac{(Z_{31} - Z_{21})I_1 + (Z_{43} - Z_{23})I_3}{Z_{22} + Z_{44} - Z_{24} - Z_{42}} \quad (3.87)$$

Now all that remains is to substitute this expression for I_2 to obtain the new two port impedances Z'_{ij} for $i, j \in \{1, 3\}$. Since the network is completely symmetric and reciprocal, we must have $Z'_{11} = Z'_{33}$ and $Z'_{13} = Z'_{31}$. We then find

$$Z'_{11} = Z'_{33} = Z_{11} + \frac{(Z_{12} - Z_{14})(Z_{41} - Z_{21})}{Z_{22} + Z_{44} - Z_{24} - Z_{42}} \quad (3.88a)$$

$$Z'_{13} = Z'_{31} = Z_{13} + \frac{(Z_{12} - Z_{14})(Z_{43} - Z_{23})}{Z_{22} + Z_{44} - Z_{24} - Z_{42}} \quad (3.88b)$$

and we compute each of the terms below:

$$Z_{22} + Z_{44} - Z_{24} - Z_{42} = 2(Z_{22} - Z_{24}) = -2iZ_{0,o} \cot \theta_o \quad (3.89a)$$

$$Z_{12} - Z_{14} = -iZ_{0,o} \csc \theta_o \quad (3.89b)$$

$$Z_{41} - Z_{21} = -(Z_{12} - Z_{14}) = iZ_{0,o} \csc \theta_o \quad (3.89c)$$

$$Z_{43} - Z_{23} = -iZ_{0,o} \csc \theta_o \quad (3.89d)$$

Substituting this into the Eqs. 3.88 above, we find:

$$Z'_{11} = Z'_{33} = Z_{11} + \frac{1}{2}i \frac{Z_{0,o} \csc^2 \theta_o}{\cot \theta_o} \quad (3.90a)$$

$$= -\frac{1}{2}i \left(Z_{0,e} \cot \theta_e + Z_{0,o} \cot \theta_o - Z_{0,o} \frac{\csc^2 \theta_o}{\cot \theta_o} \right) \quad (3.90b)$$

$$= -\frac{1}{2}i (Z_{0,e} \cot \theta_e - Z_{0,o} \tan \theta_o) \quad (3.90c)$$

$$Z'_{13} = Z'_{31} = Z_{13} - \frac{1}{2}i \frac{Z_{0,o} \csc^2 \theta_o}{\cot \theta_o} \quad (3.91a)$$

$$= -\frac{1}{2}i \left(Z_{0,e} \cot \theta_e - Z_{0,o} \cot \theta_o + Z_{0,o} \frac{\csc^2 \theta_o}{\cot \theta_o} \right) \quad (3.91b)$$

$$= -\frac{1}{2}i (Z_{0,e} \cot \theta_e + Z_{0,o} \tan \theta_o) \quad (3.91c)$$




Figure 3.11: **Frequency shift due to CPW meander.** (a) Normalized resonance frequencies for the first 4 modes of a short-circuit $\lambda/4$ (dotted lines) and an open-circuit $\lambda/2$ (solid lines) CPW resonator as a function of the CPW meander spacing s . We normalize the spacing by the total CPW width $w + 2g$, where w is the trace width and g is the gap width. We plot the results for a varying number N of meander segments. (b) Normalized resonance frequencies as a function of the number of meander segments N . We see that the resonance frequencies corresponding to $\theta_m = m\pi/2$ do not shift when m is a multiple of N .

To check these equations, we can examine the limit in which the two lines are completely decoupled ($m = c_c = 0$), in which case $Z_{0,e} = Z_{0,o} = Z_0$ and $\theta_e = \theta_o = \theta$. Then, the expressions above simplify to:

$$Z'_{11} = Z'_{33} = -\frac{1}{2}iZ_0(\cot\theta - \tan\theta) = -iZ_0 \cot(2\theta) \quad (3.92)$$

$$Z'_{13} = Z'_{31} = -\frac{1}{2}iZ_0(\cot\theta + \tan\theta) = -iZ_0 \csc(2\theta) \quad (3.93)$$

which matches our impedance matrix for a standard transmission line with electrical length $2\theta = 2kd$.

In practice, we often meander a CPW section with many turns, in which case it is simpler to analyze the circuit numerically. To do so, we simulate a 2D cross section of N CPW traces

using Ansys Q2D, allowing us to compute the $N \times N$ inductance and capacitance matrices for a given CPW width w , CPW gap g , and center-to-center separation s . For this study, we swept s between 40 μm and 150 μm for the fixed CPW dimensions $w = 17.1 \mu\text{m}$ and $g = 10 \mu\text{m}$, which yields a characteristic impedance of approximately 50 Ω . We then diagonalize the system and construct the 2N port network corresponding to length d section of N coupled lines, similar to the $N = 2$ case from the previous section⁷. Finally, we perform a port reduction to reduce the 2N-port network to a 2-port network by enforcing the boundary conditions $V_i = V_{i+2}$ and $I_i = -I_{i+2}$, as shown in Fig. 3.10. In Fig. 3.11, we show the impedance response for the resulting meandered transmission line terminated by either a short or open circuit on one end. We then sweep the CPW spacing s for a varying number of turns, and we see that the frequency of the poles will shift as the spacing s decreases. Notably, we find that resonances corresponding to $\theta_m = m\pi/2$, where m is a multiple of N , do not shift at all, and that the sign of the frequency shift alternates each time we cross a multiple of N .

Another subtlety involved with modeling the meandering CPW is that the adjacent CPW traces are connected with a bend of finite length (Fig. 3.12(a)). Within the bend, the spacing between the two lines is no longer constant and the lines are no longer parallel, making it difficult to exactly model the behavior of this section without falling back to a full 3D simulation. To address this, we introduce a bend ratio r_{bend} that divides the length of the bend ($l_{\text{bend}} = \pi s$) between the coupled line segments and an uncoupled line that connects the adjacent ports. Specifically, we increase the length of the coupled line segments by $r_{\text{bend}}\pi s/2$ at each connection point, and connect adjacent ports via an uncoupled transmission line section with length $(1 - r_{\text{bend}})\pi s$, as shown in Fig. 3.12(b). For $r_{\text{bend}} = 1$, the adjacent ports are directly connected. Fig. 3.12(c) shows the frequency shift for a $d = 5250 \mu\text{m}$ $\lambda/4$ resonator with $N = 10$ for both $r_{\text{bend}} = 0$ and $r_{\text{bend}} = 1$ as function of the separation s . Here we see that for a typical resonator frequency $\omega_r \approx 6 \text{ GHz}$, the frequency can shift by over 20% depending on the meander spacing s , making it critical that we properly model this effect.

To validate our analysis, we designed a device with 8 $\lambda/4$ CPW resonators, all with a total length of exactly $d = 5250 \mu\text{m}$, shown in Fig. 3.12(d). Rather than adjusting the length of the resonator, we evenly spaced the resonator frequencies by adjusting the meander spacing s , aiming for a target spacing of approximately 100 MHz. Measurements of the fabricated device (Fig. 3.12(e)) demonstrated remarkably accurate resonator frequencies, with a mean resonator detuning of $89 \pm 3 \text{ MHz}$. In order to determine the value of r_{bend} that most accurately reflects the true behavior of the system, we fit the measured resonator frequencies f_{meas} to a model where we both vary r_{bend} and include a correction factor α to the simulated frequencies f_{sim} . We minimize the quantity

$$\sum_{i=1}^8 \left| f_{\text{meas}}^{(i)} - (1 + \alpha) f_{\text{sim}}^{(i)}(r_{\text{bend}}) \right|^2 \quad (3.94)$$

⁷See Chapter 7 of Ref. [160] for more details on the exact procedure.




Figure 3.12: **Experimental validation of meander simulations.** (a) A meandered section of a CPW transmission line. The dotted line separates the bend section, with a total length πs . (b) A circuit model of the bend. The horizontal transmission line sections represent the coupled sections, while the vertical section represents an uncoupled section. We divide the bend length πs between the coupled and uncoupled sections specified by a ratio r_{bend} between 0 and 1. (c) The simulated frequency of the fundamental mode of a $\lambda/4$ resonator with a total length $d = 5250 \mu\text{m}$, for $r_{\text{bend}} = 0$ and $r_{\text{bend}} = 1$. The dashed gray line represents the ideal mode frequency $\omega/2\pi = \nu/(4d)$. We determine $nu = 0.404c$, where c is the speed of light in vacuum for a CPW with $w = 17.1 \mu\text{m}$ and $g = 10 \mu\text{m}$ on a silicon substrate using Ansys Q2D. (d) A rendering of the resonator design, which consists of 8 $\lambda/4$ resonators all with $d = 5250 \mu\text{m}$. The resonator frequencies were targeted by adjusting the meander spacing using the simulated curve for $r_{\text{bend}} = 1$ from (c). (e) The phase response of the microwave reflectometry measurement of the experimental device. (f) A CDF of the frequency deviations, specified in the sum Eq. 3.94. The deviations are minimized for $r_{\text{bend}} = 0.3$ and $\alpha = -0.0072$.

across the 8 resonators, and find a best fit value of $\alpha = -0.0072$ and $r_{\text{bend}} = 0.3$. We believe that the correction is primarily due to a combination of the kinetic inductance of the niobium thin film, as well as the silicon overetch of approximately 120 nm, both of which were not accounted for when constructing the network model used to simulate the resonator frequencies. Subsequent 2D cross-sectional simulations of a trenched CPW showed that the phase velocity increases by $\sim 1\%$, which would result in a corresponding increase in the frequency.

The remarkable uniformity of the resonator detunings validates this method of modeling tightly meandered CPW traces, enabling drastic speedups in the simulation of more complex microwave structures. With only 2d cross-sectional simulations (which can in principle be run once and saved to a database) coupled with microwave network analysis, we were able to simulate resonator frequencies to a high degree of accuracy in only a fraction of the time compared to a full 3D-simulation of the 8 resonator system. The remaining sources of potential deviation (kinetic inductance, silicon overetch etc.), are also typically not accounted for when performing 3D FEM simulations, but are straightforward to model with the network analysis method.

This study highlights just one example of how circuit models can be refined to more accurately reflect the true device structure, while still retaining computational efficiency. We expect this type of model refinement and modular simulation to become increasingly important as superconducting qubit devices scale with respect to both size and complexity [134]. On the flip side, we can also design individual unit cells with better shielding and isolation—leveraging fabrication processes like through-silicon vias [82] and superconducting bump bonding [171, 64]—to ensure that electromagnetic interactions between different unit cells occur only at well-defined ports that can be accurately modeled with microwave network analysis. Together, these parallel efforts in model refinement and modular design can make microwave network analysis a viable alternative to full-wave 3D FEM simulation, enabling faster iteration cycles and more scalable design workflows.

Chapter 4

Qubit Control

In Chapter 1, we introduced the basic concepts of qubit control in an abstract setting. In this chapter, we discuss the practical implementation of both single- and two-qubit gates, aided by a better understanding of quantum circuits. This will provide the theoretical background for the work presented in Chapter 5 on parametric gates. While engineering high-fidelity quantum gates is full of complex subtleties, what we find is that much of this behavior can be understood by extrapolating the behavior of a driven two-level system. Resonant coupling leads to population swaps, and off-resonant coupling leads to frequency shifts. We then discuss the architectural implications of different two-qubit gate schemes and analyze some of the drawbacks of all-microwave gates with both fixed-frequency transmons and static coupling.

4.1 Single Qubit Gates

Charge and Flux Drives

To drive a quantum circuit with Hamiltonian H_0 , we can either charge couple ($\hat{Q}$) or flux couple ($\hat{\Phi}$) to the system, in which case the driven hamiltonian looks like

$$H = H_s + H_d = H_s - I(t)\Phi_{\text{zpf}}(a + a^\dagger) + iV(t)Q_{\text{zpf}}(a^\dagger - a) \quad (4.1)$$

where $I(t)$ and $V(t)$ are classical drive currents and voltages that couple via the flux or charge degrees of freedom. Here we note that $I(t)$ and $V(t)$ are defined at the circuit node and are typically rescaled by a coupling quality factor $\sqrt{Q_c}$ with respect to the drive current or voltage at the control line. If we define

$$\varepsilon_\ell(t) = \Phi_{\text{zpf}}I(t) \quad \varepsilon_c(t) = Q_{\text{zpf}}V(t) \quad (4.2)$$

as the effective inductive and capacitive drives with units of energy, we see that the total effective drive is

$$H_d = -[\varepsilon_\ell(t) + i\varepsilon_c(t)]a - [\varepsilon_\ell(t) - i\varepsilon_c(t)]a^\dagger = -\varepsilon(t)a - \varepsilon^*(t)a^\dagger \quad (4.3)$$

where $\varepsilon(t) = \varepsilon_\ell(t) + i\varepsilon_c(t)$ is the combined drive amplitude. Here we see that current and voltage driving act out of phase with respect to each other, which reflects their status as canonically conjugate variables. In most cases, we drive the system with a periodic drive, so we consider $\varepsilon_\ell(t)$ and $\varepsilon_c(t)$ of the form

$$\varepsilon_\ell(t) = \Omega_\ell(t) \cos(\omega_d t + \phi_\ell) = \frac{1}{2} \left[\tilde{\Omega}_\ell(t) e^{i\omega_d t} + \tilde{\Omega}_\ell^*(t) e^{-i\omega_d t} \right] \quad (4.4)$$

$$\varepsilon_c(t) = \Omega_c(t) \cos(\omega_d t + \phi_c) = \frac{1}{2} \left[\tilde{\Omega}_c(t) e^{i\omega_d t} + \tilde{\Omega}_c^*(t) e^{-i\omega_d t} \right] \quad (4.5)$$

where $\Omega_\ell(t)$ (Ω_c) represents a real-valued envelope amplitude for the flux (charge) drive, ω_d is the drive frequency, and ϕ_ℓ (ϕ_c) is the phase of the flux (charge) drive. In the second equality we fold the phase of the drive into a complex amplitude $\tilde{\Omega}_{\ell/c}(t) = \Omega(t) e^{i\phi_{\ell/c}}$.

At this point we can apply a unitary transformation $U_s = \exp[iH_s t]$ to move into the rotating frame of the system, and the interaction Hamiltonian becomes:

$$H' = U_s H U_s^\dagger + iU_s U_s^\dagger = H_s + U_s H_d U_s^\dagger - H_s U_s U_s^\dagger = U_s H_d U_s^\dagger \quad (4.6)$$

At this point, we must pick a basis to evaluate the Hamiltonian in. Since we are primarily concerned with the dynamics between the different eigenstates of the system hamiltonian H_0 , it makes the most sense to work in the basis where $H_s = \sum_n \omega_n |n\rangle\langle n|$ and $U_s = \sum_n e^{i\omega_n t} |n\rangle\langle n|$ are both diagonal. Then we have

$$U_s a U_s^\dagger = \sum_{mn} a_{nm} e^{-i\omega_{nm} t} |n\rangle\langle m| \quad (4.7)$$

where $a_{nm} = \langle n_s | a | m_s \rangle$ are the matrix elements a in the system eigenbasis and $\omega_{nm} = \omega_m - \omega_n$ is the transition frequency between $|m_s\rangle$ and $|n_s\rangle$. This results in the full drive Hamiltonian:

$$H'_d = U_s H_d U_s^\dagger \quad (4.8a)$$

$$= -\frac{1}{2} \left[\tilde{\Omega}_\ell(t) + i\tilde{\Omega}_c(t) \right] \sum_{mn} a_{nm} e^{-i\Delta_{nm} t} |n_s\rangle\langle m_s| \quad (4.8b)$$

$$- \frac{1}{2} \left[\tilde{\Omega}_\ell^*(t) + i\tilde{\Omega}_c^*(t) \right] \sum_{mn} a_{nm} e^{-i\Sigma_{nm} t} |n_s\rangle\langle m_s| \quad (4.8c)$$

$$+ \text{c.c.} \quad (4.8d)$$

where $\Delta_{nm} = \omega_{nm} - \omega_d$ and $\Sigma_{nm} = \omega_{nm} + \omega_d$ are the sum and difference frequencies between the drive and each transition.

While this is a rather complex Hamiltonian, it can generally be simplified in practice with proper system engineering. For the transmon and other weakly anharmonic systems, the eigenstates are approximately the Fock states $|n\rangle$, so $a_{nm} = \sqrt{n} \delta_{n+1,m}$ ¹, and is fairly sparse.

¹Here, δ_{nm} is the Kronecker delta function




Figure 4.1: **Generalized Rabi driving of a two-level system** (a) The Bloch sphere trajectories of a driven two-level system starting in $|0\rangle$, for several different detunings Δ . We see that as the detuning increases, the axis of rotation tilts closer toward the poles. (b) Level diagram of a two-level system showing the detuned drives. (c) A 2-D plot showing the population transfer between the $|0\rangle$ and $|1\rangle$ states as a function of both drive detuning and time (both are normalized to the drive rate Ω_Δ). This results in the well-known “chevron” pattern. The oscillation rate increases, and the contrast decreases, as we go to higher detunings. (d) The expectation values along σ_x , σ_y , and σ_z for different drive detunings.

For highly anharmonic systems such as a fluxonium qubit, this is no longer true and the drive operator a can contain additional non-zero matrix elements in the system eigenbasis. However, the symmetries in the Hamiltonian will typically suppress many of the possible matrix elements. Furthermore, we can typically truncate the drive Hamiltonian and isolate a particular transition $|i\rangle \rightarrow |j\rangle$, provided that the other allowed transition frequencies involving $|i\rangle$ and $|j\rangle$ are sufficiently detuned from ω_{ij} . This justifies the importance of studying gates using the two-level approximation, since the dominant effects of flux or charge driving can be captured with a truncated Hamiltonian involving only two levels, with additional corrections added back in perturbatively. In the following sections, we will also study these corrections.

Two Level Approximation

We start by reviewing the two-level approximation, and we can choose $m = 1$, $n = 0$ without loss of generality. We drop all other terms by arguing either that the matrix element is sufficiently small, the term does not involve either $|0_s\rangle$ or $|1_s\rangle$, or the transition frequency is sufficiently detuned from ω_{01} . If we let $a_{01} = a_{10}^* = g$ be the scaling factor due to the matrix

element ($g = 1$ for the lowest transition in the fock basis), the resulting drive Hamiltonian in the system frame looks like:

$$H'_d = -\frac{g}{2} \left[\tilde{\Omega}_\Delta(t) e^{-i\Delta t} + \tilde{\Omega}_\Sigma(t) e^{-i\Sigma t} \right] |0_s\rangle\langle 1_s| + \text{c.c.} \quad (4.9)$$

where $\tilde{\Omega}_\Delta(t) = \tilde{\Omega}_\ell(t) + i\tilde{\Omega}_c(t)$ is the co-rotating (or beamsplitter) drive amplitude and $\tilde{\Omega}_\Sigma(t) = \tilde{\Omega}_\ell^*(t) + i\tilde{\Omega}_c^*(t)$ is the counter-rotating (or squeezing) drive amplitude. We can write $\tilde{\Omega}_\Delta = \Omega_\Delta e^{i\phi_\Delta}$ and $\tilde{\Omega}_\Sigma = \Omega_\Sigma e^{i\phi_\Sigma}$, where

$$\Omega_\Delta = \sqrt{\Omega_\ell^2 + \Omega_c^2 + 2\Omega_c\Omega_\ell \sin(\phi_\ell - \phi_c)} \quad \tan \phi_\Delta = \frac{\Omega_\ell \sin \phi_\ell + \Omega_c \cos \phi_c}{\Omega_\ell \cos \phi_\ell - \Omega_c \sin \phi_c} \quad (4.10)$$

$$\Omega_\Sigma = \sqrt{\Omega_\ell^2 + \Omega_c^2 - 2\Omega_c\Omega_\ell \sin(\phi_\ell - \phi_c)} \quad \tan \phi_\Sigma = \frac{\Omega_c \cos \phi_c - \Omega_\ell \sin \phi_\ell}{\Omega_c \sin \phi_c + \Omega_\ell \cos \phi_\ell} \quad (4.11)$$

and we see that the co-rotating and counter-rotating terms may have different amplitudes and phases. For now, we will drop the counter-rotating terms and consider a resonant drive with constant amplitude. The drive Hamiltonian simplifies to

$$H'_d = -\frac{g}{2} \left[\tilde{\Omega}_\Delta |0_s\rangle\langle 1_s| + \tilde{\Omega}_\Delta^* |1_s\rangle\langle 0_s| \right] \quad (4.12)$$

and the resulting unitary evolution becomes

$$U'_d(t) = \exp[-iH'_d t] = \begin{pmatrix} \cos\left(\frac{g}{2}\Omega_\Delta t\right) & -ie^{i\phi_\Delta} \sin\left(\frac{g}{2}\Omega_\Delta t\right) \\ -ie^{-i\phi_\Delta} \sin\left(\frac{g}{2}\Omega_\Delta t\right) & \cos\left(\frac{g}{2}\Omega_\Delta t\right) \end{pmatrix} \quad (4.13)$$

Given an initial state $|0_s\rangle$, the population P_0 in the state will be

$$P_0 = |\langle 0_s | U'_d(t) | 0_s \rangle|^2 = \cos^2\left(\frac{g}{2}\Omega_\Delta t\right) = \frac{1}{2} + \frac{1}{2} \cos(g\Omega_\Delta t) \quad (4.14)$$

and we see that this leads to oscillations at a rate $g\Omega_\Delta$ between the two states.

Detuned Driving: AC Stark Effect

Now, we relax the first assumption that the drive is resonant, and consider what happens when the detuning $\Delta_{01} = \omega_{01} - \omega_d$ is non-zero. Then the drive Hamiltonian becomes

$$H'_d = -\frac{g}{2} \left[\tilde{\Omega}_\Delta e^{-i\Delta t} |0_s\rangle\langle 1_s| + \tilde{\Omega}_\Delta^* e^{i\Delta t} |1_s\rangle\langle 0_s| \right] \quad (4.15)$$

where we continue to ignore the counter-rotating terms. To solve for the dynamics exactly, we can remove the time-dependence by transforming into the frame of the drive, keeping

in mind that we are now leaving the system frame. We apply the unitary transformation $U_{\text{drive}} = \exp [i\frac{1}{2}\Delta t (|0_s\rangle\langle 0_s| - |1_s\rangle\langle 1_s|)]$ to obtain a new effective Hamiltonian H_d'' :

$$H_d'' = -\frac{g}{2} \left[\tilde{\Omega}_\Delta |0_s\rangle\langle 1_s| + \tilde{\Omega}_\Delta^* |1_s\rangle\langle 0_s| \right] - \frac{\Delta}{2} [|0_s\rangle\langle 0_s| - |1_s\rangle\langle 1_s|] \quad (4.16a)$$

$$= -\frac{1}{2} \begin{pmatrix} \Delta & g\Omega_\Delta e^{i\phi_\Delta} \\ g\Omega_\Delta e^{-i\phi_\Delta} & -\Delta \end{pmatrix} \quad (4.16b)$$

$$= -\frac{1}{2} g\Omega_\Delta (\cos \phi_\Delta \sigma_x - \sin \phi_\Delta \sigma_y) - \frac{1}{2} \Delta \sigma_z \quad (4.16c)$$

$$= -\frac{\Lambda}{2} \hat{n} \cdot \vec{\sigma} \quad (4.16d)$$

Here, we see that the new Hamiltonian H_d'' looks like a rotation operator along an axis $\hat{n} = (g\Omega_\Delta \cos \phi_\Delta, -g\Omega_\Delta \sin \phi_\Delta, \Delta) / \Lambda$, where $\Lambda = \sqrt{g^2\Omega_\Delta^2 + \Delta^2}$. The rotation axis is now tilted such that its angle β from the Z-axis is

$$\tan \beta = \frac{g\Omega_\Delta}{\Delta} \quad (4.17)$$

and the precession frequency is modified as Λ . In the limit $\Delta \rightarrow 0$, $\beta \rightarrow \pi/2$, $\Lambda \rightarrow g\Omega_\Delta$ and we recover the resonant case where the drive leads to a rotation around an axis in the XY-plane. In the other far-detuned limit ($\Delta \gg g\Omega_\Delta$), $\beta \rightarrow 0$ and we see that the drive now leads to a Z-rotation. The time evolution under this Hamiltonian is given by

$$U_d''(t) = \exp [-iH_d''t] \quad (4.18a)$$

$$= \cos \left(\frac{\Lambda t}{2} \right) \mathbf{I} - \frac{i}{\Lambda} \sin \left(\frac{\Lambda t}{2} \right) \begin{pmatrix} \Delta & g\Omega_\Delta e^{i\phi_\Delta} \\ g\Omega_\Delta e^{-i\phi_\Delta} & -\Delta \end{pmatrix} \quad (4.18b)$$

However, we recall that this propagator is acting on the modified states $U_{\text{drive}} |\psi\rangle$, so the time evolution operator acting on the original basis states is given by $U_d'(t) = U_{\text{drive}}^\dagger U_d''(t) U_{\text{drive}}$ and we have

$$U_d'(t) = \begin{pmatrix} \cos \left(\frac{\Lambda t}{2} \right) - i\frac{\Delta}{\Lambda} \sin \left(\frac{\Lambda t}{2} \right) & -i\frac{g\Omega_\Delta}{\Lambda} \sin \left(\frac{\Lambda t}{2} \right) e^{i(\phi_\Delta - \Delta t)} \\ -i\frac{g\Omega_\Delta}{\Lambda} \sin \left(\frac{\Lambda t}{2} \right) e^{-i(\phi_\Delta - \Delta t)} & \cos \left(\frac{\Lambda t}{2} \right) + i\frac{\Delta}{\Lambda} \sin \left(\frac{\Lambda t}{2} \right) \end{pmatrix} \quad (4.19a)$$

Following the previous section, we can once again look at the population P_0 as a function of time given an initial state $|0_s\rangle$:

$$P_0 = |\langle 0_s | U_d'(t) | 0_s \rangle|^2 = \cos^2 \left(\frac{\Lambda t}{2} \right) + \frac{\Delta^2}{\Lambda^2} \sin^2 \left(\frac{\Lambda t}{2} \right) \quad (4.20a)$$

$$= \frac{1}{2} (1 + \cos(\Lambda t)) + \frac{\Delta^2}{2\Lambda^2} (1 - \cos(\Lambda t)) \quad (4.20b)$$

$$= \frac{1}{2} \left(1 + \frac{\Delta^2}{\Lambda^2} \right) + \frac{1}{2} \left(1 - \frac{\Delta^2}{\Lambda^2} \right) \cos(\Lambda t) \quad (4.20c)$$

$$= 1 + \left(\frac{g\Omega_\Delta}{\Lambda} \right)^2 \left[\frac{\cos(\Lambda t) - 1}{2} \right] \quad (4.20d)$$

From the above equation, we see that the off-resonant driving leads to faster oscillations in P_0 but with a reduced amplitude that scales as $\Omega_\Delta^2/\Lambda^2$. In the limit where $g\Omega_\Delta \ll \Delta$, these oscillations become negligible and we can write the propagator $U'_d(t)$ as

$$U'_d(t) \approx \begin{pmatrix} e^{-i\frac{1}{2}(\Lambda-\Delta)t} & 0 \\ 0 & e^{i\frac{1}{2}(\Lambda-\Delta)t} \end{pmatrix} \quad (4.21)$$

where we approximate $\Delta/\Lambda \approx 1$ and $g\Omega_\Delta/\Lambda \approx 0$. This now looks like a Z-rotation with an angle $\theta(t) = (\Lambda - \Delta)t$, which is equivalently interpreted as a drive-induced detuning

$$\Lambda - \Delta = \sqrt{g^2\Omega_\Delta^2 + \Delta^2} - \Delta \approx \frac{g^2\Omega_\Delta^2}{2\Delta} \quad (4.22)$$

Since $\Delta = \omega_{01} - \omega_d$, we see that this leads to a positive frequency shift when $\omega_d < \omega_{01}$ and a negative frequency shift when $\omega_d > \omega_{01}$. This frequency shift is known as the AC stark shift.

The above derivation relied on the fact that $g\Omega_\Delta \ll \Delta$. In the following, we will show that the AC stark shift is still present in the strong-driving limit, depending on the choice of basis. To do so, we return to the drive Hamiltonian H_d'' in the frame of the drive (Eq. 4.16). We then diagonalize this Hamiltonian to find the new eigenenergies, Λ , and eigenstates under the driven Hamiltonian:

$$\Lambda_\pm = \pm \frac{1}{2} \sqrt{g^2\Omega_\Delta^2 + \Delta^2} = \pm \Lambda \quad (4.23)$$

$$|+\rangle_\Lambda = \begin{pmatrix} -\sin \frac{\beta}{2} \\ e^{-i\phi_\Delta} \cos \frac{\beta}{2} \end{pmatrix} \quad |-\rangle_\Lambda = \begin{pmatrix} \cos \frac{\beta}{2} \\ e^{-i\phi_\Delta} \sin \frac{\beta}{2} \end{pmatrix} \quad (4.24)$$

where $\cot \beta = \frac{\Delta}{g\Omega_\Delta}$ defines the mixing angle β . We choose to define the angle with respect to $\cot \beta^2$ so that we can restrict β to the domain $[0, \pi]$. Here, we see that for $\Delta > 0$, $\beta < \pi/2$ and $|+\rangle_\Lambda$ will have more overlap with $|0_s\rangle$. For $\Delta < 0$ this is flipped. Thus the effective detuning from the drive in this eigenbasis is

$$\tilde{\omega}_{01} - \omega_d = \begin{cases} \omega_+ - \omega_- = \Lambda & (\Delta > 0) \\ \omega_- - \omega_+ = -\Lambda & (\Delta < 0) \end{cases} \quad (4.25)$$

We then find the effective frequency shift relative to the original transition frequency as

$$\tilde{\omega}_{01} - \omega_d = \text{sgn}(\Delta)\Lambda \quad (4.26a)$$

$$\tilde{\omega}_{01} - \omega_{01} = \text{sgn}(\Delta)\Lambda - \omega_{01} + \omega_d \quad (4.26b)$$

$$\tilde{\omega}_{01} - \omega_{01} = \text{sgn}(\Delta)\Lambda - \Delta \quad (4.26c)$$

Here, we once again see that the transition frequency $\tilde{\omega}_{01}$ is shifted down for $\Delta < 0 \Rightarrow \omega_{01} < \omega_d$, and it is shifted up for $\Delta > 0 \Rightarrow \omega_{01} > \omega_d$, which agrees with our previous derivation of the AC stark shift in the far-detuned limit.

²Note that β ends up being the same as the previous section modulo π .

Here, we have shown how the same system, an off-resonantly driven two-level system, can exhibit two types of behavior. In the small detuning limit, which one is relevant depends on the the eigenbasis. If our computational states are the dressed eigenstates $|\pm\rangle_\Lambda$, we will observe an effective frequency shift. If our computational states are the original system eigenstates $|0_s\rangle, |1_s\rangle$, we will observe small off-resonant oscillations. However, we can map between these sets of eigenstates if the drive envelope $\Omega_\Delta(t)$ is varied slowly. According to the adiabatic theorem, the system will follow the instantaneous eigenstates up to a phase provided that

$$\langle -_\Lambda | \dot{+}_\Lambda \rangle \ll |\omega_+ - \omega_-| \quad (4.27)$$

Here, we observe that $|\dot{+}_\Lambda\rangle = -\frac{1}{2}\dot{\beta}|-_\Lambda\rangle$, so the adiabatic condition becomes

$$\left| -\frac{1}{2}\dot{\beta} \right| = \left| -\frac{1}{2} \frac{d}{dt} \left[\cot^{-1} \left(\frac{\Delta}{g\Omega_\Delta} \right) \right] \right| \quad (4.28a)$$

$$= \left| -\frac{1}{2} \left(-\frac{g^2\Omega_\Delta^2}{g^2\Omega_\Delta^2 + \Delta^2} \right) \left(\frac{g\Omega_\Delta\dot{\Delta} - g\dot{\Omega}_\Delta\Delta}{g^2\Omega_\Delta^2} \right) \right| \quad (4.28b)$$

$$= \frac{1}{2} \frac{|g\Omega_\Delta\dot{\Delta} - g\dot{\Omega}_\Delta\Delta|}{\Lambda^2} \ll \Lambda \quad (4.28c)$$

For a fixed detuning Δ , this means that we cannot ramp Ω_Δ too fast relative to the detuning if we want to adiabatically realize a stronger AC stark shift.

Beyond RWA

Now that we understand the dynamics well under the RWA, we consider the effect of the counter-rotating terms. For a resonant drive, the Hamiltonian in the system frame looks like

$$H'_d = -\frac{g}{2} \left[\tilde{\Omega}_\Delta(t) |0_s\rangle\langle 1_s| + \text{c.c.} \right] - \frac{g}{2} \left[\tilde{\Omega}_\Sigma(t) e^{-2i\omega_d t} + \text{c.c.} \right] = H_\Delta + H_\Sigma(t) \quad (4.29)$$

where we treat the time dependent terms $H_\Sigma(t)$ as a perturbation on top of the time-independent RWA hamiltonian H_Δ . When the drive strength Ω_Σ is no longer small compared to Σ , this leads to a secondary oscillation at a frequency $\Sigma = 2\omega_d$ with amplitude Ω_Σ/Σ . Fig. 4.2 shows the simulated trajectories of a driven two-level system for different values of Ω_Σ/Σ . We see that as Ω_Σ approaches Σ , the amplitude of the secondary oscillations increases and can cause a meaningful deviation from the desired trajectory of the system. The Fourier transforms of the expectation value $\langle \sigma_z \rangle$, shown in Fig. 4.2(c) and Fig. 4.2(f), reveal additional peaks at $\Sigma \pm \Omega_\Sigma$ which validates the frequency of the secondary oscillation. Finally, we note that in the adiabatic limit, we can also treat the counter-rotating drive as an off-resonant drive that leads to an effective frequency shift $\delta\omega \approx \frac{\Omega_\Sigma^2}{2\Sigma}$. This frequency shift is known as the Bloch-Siegert shift [20].




Figure 4.2: **Non-RWA effects on a driven two-level system.** The counter-rotating terms in the driven Rabi model affect the system dynamics when Ω_Σ is no longer small compared to Σ . (a) The Bloch sphere trajectory of a driven two-level system starting in the $|0\rangle$ state when $\Sigma = 10\Omega_\Sigma$. We see an additional oscillation on top of the trajectory that occurs at a frequency Σ . (b) The projected σ_x , σ_y , and σ_z expectation values of the trajectory. We see beating in the dominant Rabi oscillation dynamics. (c) The FFT of the σ_z trajectory shows peaks at Ω_Δ , Σ , and $\Sigma \pm \Omega_\Sigma$. This confirms the frequency of the secondary oscillation. (d)-(f) The same as (a)-(c), but for $\Sigma = 50\Omega_\Sigma$.

Drive Polarization

One important point that we have overlooked so far is that the effective drive amplitudes of the co-rotating and counter-rotating terms depend on the relative amplitudes and phases of the inductive and capacitive drive strengths. From Eq. 4.10, we see that when the drive strengths are the same, $\Omega_c = \Omega_\ell = \Omega$, and $\phi_\ell - \phi_c = \pi/2$

$$\Omega_\Delta = \Omega \sqrt{2 + 2 \sin\left(\frac{\pi}{2}\right)} = 2\Omega \quad \Omega_\Sigma = \Omega \sqrt{2 - 2 \sin\left(\frac{\pi}{2}\right)} = 0 \quad (4.30)$$

and the counter-rotating terms completely vanish. Here, the flux-driving field leads the charge-driving field by a phase $\pi/2$, which is reminiscent of the phase relationship between the magnetic and electric fields in left-hand circularly polarized light³. Under these conditions,

³Similarly, the co-rotating terms cancel if $\phi_\ell - \phi_c = -\pi/2$, corresponding to right-hand circularly polarized light.

RWA is no longer an approximation, but becomes exact. This can be useful in cases where the drive strengths are no longer small compared to the drive frequency ($\Omega \sim \omega_d$). This type of circularly polarized driving was first demonstrated in NV centers [118], and more recently in superconducting circuits with a fluxonium qubit [173]. It was also discussed in Ref. [174] under the name *balanced driving*.

While circularly polarized driving can be a clever method to eliminate error contributions from the counter-rotating terms, it comes with the hardware overhead of an additional control line⁴. Moreover, it is not strictly necessary to fully remove the counter-rotating terms at all times. What matters is that the final integrated unitary is correct — the exact trajectory of the qubit state is typically irrelevant, since dynamical phases from the trajectory can be easily corrected. This enables for alternative solutions such as commensurate pulses, where the pulse start and end times are set to integer multiples of $(2\omega_{01})^{-1}$ [173], or pulse-shaping techniques such as DRAG, which can suppress such non-adiabatic drive terms [141].

Multi-level Systems

While we try our best to engineer non-linear circuits with pairs of states that can be well approximated by two-level systems, it is important to account for errors in the target unitary due to these neighboring and off-resonant levels. We will address most of these potential errors in the context of the transmon, which typically has a weak nonlinearity $|\alpha| \sim 0.05\omega_{01}$, but the analysis applies to arbitrary systems. In analyzing these systems, we will find that many of the corrections can be understood to lowest order as resulting from off-resonant driving of another two-level system. For simplicity, we will also assume that we are capacitively driving so that $\Omega_\ell = 0$ and $\Omega_\Delta = i\Omega_c = \Omega$ and $\Omega_\Sigma = -i\Omega_c = -\Omega$.

We begin by looking at the presence of the $|2\rangle$ state in a transmon, where the anharmonicity is negative ($\alpha < 0$) and defined as $\alpha = \omega_{12} - \omega_{01}$. For the transmon, we have non-zero $a_{nm} = \sqrt{n}\delta_{n+1,m}$, so when we drive the system at a frequency ω_{01} , the $|1\rangle$ - $|2\rangle$ transition will experience an off-resonant drive with detuning α . From our previous analysis, we know that off-resonant driving in the diabatic limit will lead to unintended population transfer at the level of $P_{\text{leakage}} \propto \Omega^2/\Lambda^2 \sim \Omega^2/\alpha^2$. For typical transmon qubits, where $\Omega/(2\pi) \lesssim 10$ MHz and $\alpha/(2\pi) \sim -250$ MHz, this amounts to a leakage error on the order of 0.1%, which might seem small until we remember that we are trying to achieve gate errors below 10^{-4} . While this is a coherent process, any population in the $|2\rangle$ state will gain a phase at a rate α relative to the $|1\rangle$ state, which can prevent its return to the $|1\rangle$ state and cause a buildup of population in the leakage state. This can be seen in experiments where a variable time delay between consecutive gates is swept, and coherent constructive or destructive interference appears at time delays that multiples of $1/\alpha$ ⁵ [120].

⁴It is possible, in principle, to implement circular driving with a single control line and a quarter-wave delay, as noted in Ref. [174], but this imposes a narrow and fixed band of operation that can be difficult to target for a fixed-frequency qubit.

⁵The Ramsey error filter (REF) discussed in Ref. [120] becomes a convenient way to amplify leakage for calibration.

The off-resonant driving of a neighboring state will also cause an AC stark shift that changes the frequency of the target transition, leading to rotation errors [129]. Again, we can approximate the Stark shift using our two-level system derivations above and we find that the Stark shift will be approximately $\tilde{\Delta}_{01} = -\tilde{\Delta}_{12}/2 \approx -\Omega^2/\alpha^6$. For the typical transmon parameters quoted above, this represents a frequency shift of $\tilde{\Delta}_{01} \sim +400$ kHz, leading to a tilt in the rotation axis of $\sim 3^\circ$ and a corresponding phase error.

As mentioned in the previous section, whether the off-resonant driving of the neighboring transition leads to phase errors or leakage errors depends on whether the pulse shape $\Omega(t)$ allows for the adiabatic tracking of the driven eigenstates. From Eq. 4.28, we know that this non-adiabatic correction is of order $\dot{\Omega}\Delta/\Lambda^2 \approx \dot{\Omega}/\Delta$. To overcome this error, a wide variety of pulse-shaping techniques have been developed over the years under a framework known as *shortcuts to adiabaticity* (STA), which all endeavor to minimize the non-adiabatic terms in the time evolution with clever pulse shaping. The first of these techniques and most commonly used of these techniques is DRAG [142, 66], in which an out of phase component proportional to $i\dot{\Omega}(t)/\Delta$ is added to the control pulse to cancel the non-adiabatic terms.

Once the non-adiabatic terms have been canceled to the desired order, the remaining effect of the neighboring levels are phase errors and detuning errors from the stark shift or a DRAG correction. One way to compensate for such errors is to add an equal detuning to the drive frequency, so that the effective rotations are no longer tilted [35]. The downside of this approach, however, is that the optimal pulse frequency now depends on the amplitude $|\Omega|$ of the drive. This makes it more difficult to decouple the pulse parameters for sequential calibration. A better approach is to leverage Z gates before and after a rotation to compensate for the axis tilt, which are conveniently zero-cost with virtual-Z gates [129]. We will analyze the conditions under which this is possible below.

If we look at the diagonal elements of the unitary in Eq. 4.19, we see that these can be expressed as

$$\cos\left(\frac{\Lambda t}{2}\right) \pm i\frac{\Delta}{\Lambda}\sin\left(\frac{\Lambda t}{2}\right) = \cos\left(\frac{\theta}{2}\right)e^{\pm i\alpha} \quad (4.31)$$

where the angles θ and α are defined by

$$\cos\frac{\theta}{2} \equiv \sqrt{\cos^2\left(\frac{\Lambda t}{2}\right) + \frac{\Delta^2}{\Lambda^2}\sin^2\left(\frac{\Lambda t}{2}\right)} \quad (4.32)$$

$$\tan\alpha \equiv \frac{\Delta}{\Lambda}\tan\left(\frac{\Lambda t}{2}\right) = \cos\beta\tan\left(\frac{\Lambda t}{2}\right) \quad (4.33)$$

Using some basic trigonometric identities, we can show also that

$$\sin\frac{\theta}{2} = \sqrt{1 - \cos^2\frac{\theta}{2}} = \frac{g\Omega_{\Delta}}{\Lambda}\sin\left(\frac{\Lambda t}{2}\right) = \sin\beta\sin\left(\frac{\Lambda t}{2}\right) \quad (4.34)$$

⁶Since it is the $|1\rangle$ - $|2\rangle$ transition that is Stark-shifted, only half of this shift will contribute to the $|0\rangle$ - $|1\rangle$ shift, and the sign will be flipped.

allowing us to rewrite the unitary in Eq. 4.19 as

$$U'_d(t) = \begin{pmatrix} e^{-i\alpha} \cos \frac{\theta}{2} & -ie^{i(\phi_\Delta - \Delta t)} \sin \frac{\theta}{2} \\ -ie^{-i(\phi_\Delta - \Delta t)} \sin \frac{\theta}{2} & e^{i\alpha} \cos \frac{\theta}{2} \end{pmatrix} = R_z(\alpha) R_{xy}(\theta, \phi_\Delta - \Delta t) R_z(\alpha) \quad (4.35)$$

The result is that detuned driving is equivalent to a rotation of θ about an axis in the XY-plane sandwiched by two Z-rotations. One caveat, however, is that the range of accessible rotation angles θ is reduced such that $|\theta| \leq 2\beta$. For a resonant drive, $\beta = \frac{\pi}{2}$, so there is no such restriction. The geometric picture for this limitation is that the tilted rotation must cross through the plane defined by $z = \cos \theta$ in order for us to reach the target state with only Z-rotations. However, as long as this condition is met, we can always correct for the effect of the detuning by applying two Z-rotations $R_z(-\alpha)$ before and after the detuned drive and adjusting the phase ϕ_Δ to compensate for the detuning:

$$R_{xy}(\theta, \phi'_\Delta) = R_z(-\alpha) U'_d(t) R_z(-\alpha) \quad (4.36)$$

While this may seem like good news, it comes at the cost of a longer gate times for a fixed drive strength Ω . For a given rotation angle θ , we can implement the rotation resonantly $\Delta = 0$ in a time $\tau_r = \theta/\Omega$. Then, the time it would take to implement the same rotation with a detuned drive (τ_d) is given by

$$\tau_d = \frac{2}{\Lambda} \arcsin \left[\sin \frac{\theta \Lambda}{2 \Omega} \right] \quad (4.37)$$

While it is not straightforward to simplify this any further, we can prove a bound on τ_d for the range $0 < \theta \leq \pi$. In this range, we know that $\sin(\theta/2)$ is monotonically increasing. We also now that $\Lambda/\Omega \geq 1$, where the bound is saturated for zero detuning. Since the derivative of $\arcsin(x)$ is strictly greater than 1, this implies that

$$\arcsin \left[\sin \frac{\theta \Lambda}{2 \Omega} \right] \geq \frac{\theta \Lambda}{2 \Omega} = \frac{\tau_r \Lambda}{2} \quad (4.38)$$

and we conclude that $\tau_d \geq \tau_r$. This bound is clearly saturated for $\Delta = 0$, where $\Lambda = \Omega$ and the above equation becomes an equality. But as the detuning increases, τ_d must also increase relative to τ_r . To see how much this increases, consider a common situation where we want to do a $\pi/2$ rotation. On resonance, this will take a time $\tau_r = \pi/(2\Omega)$. The largest detuning that still allows us to implement such a rotation occurs for $\Delta = \Omega$, which implies that $\Lambda = \sqrt{2}\Omega$. Then, we find $\tau_d = \sqrt{2}\pi/(2\Omega) = \sqrt{2}\tau_r$, which is approximately 41% longer than the resonant case. The above analysis reinforces the notion that resonant interactions are always faster for a fixed interaction strength. However, in most cases, the effective detuning caused by neighboring transitions is sufficiently small so that the increase in gate time is negligible. In these instances, the simplicity of correcting phase errors with virtual-Z gates significantly outweighs the minor increase in gate time.

4.2 Two-Qubit Gates

Capacitive and Inductive Coupling

In order to engineer entanglement between two qubit modes, these modes must be coupled so that they can share energy. In analogy to the microwave flux and charge drives that allow us to drive single-qubit gates, this coupling can be either inductive or capacitive in nature. Here, one can think of the quantum field of one qubit mode as driving the other mode, with the main difference being that this also leads to backaction on the first qubit.

Given two system Hamiltonians H_1 , H_2 , a general form of the coupling Hamiltonian via a capacitance or mutual inductance is of the form

$$H = H_a + H_b - g_\ell(a + a^\dagger)(b + b^\dagger) - g_c(a - a^\dagger)(b - b^\dagger) \quad (4.39)$$

where g_ℓ and g_c are the inductive and capacitive coupling strengths derived in Section 2.3. Again, we can group into energy conserving terms $ab^\dagger$, $a^\dagger b$ and two mode squeezing terms ab and $a^\dagger b^\dagger$, so that the coupling Hamiltonian is

$$H_c = -g_\Delta(ab^\dagger + a^\dagger b) - g_\Sigma(ab + a^\dagger b^\dagger) \quad (4.40)$$

where $g_\Delta = g_\ell - g_c$ is the beam splitter coupling strength and $g_\Sigma = g_\ell + g_c$ is the squeezing coupling strength. Like in the driving case, this can lead to an addition or cancellation of the two effective coupling strengths depending on the signs and strengths of the physical coupling strengths. As a result, “balanced” coupling can be achieved when $g_c = -g_\ell$ [174].

Again, we can move into our effective computation frame by making a unitary transformation that cancels the bare system dynamics. We let $H_s = H_a + H_b$ and we work in full system eigenbasis where both H_s and $U_s = e^{iH_s t}$ are diagonal. If H_a and H_b are fully decoupled such that $[H_a, H_b] = 0$, we can write $U_s = U_a U_b = U_b U_a$, where $U_j = e^{iH_j t}$, and the interaction Hamiltonian will contain terms that look like

$$U_s(a \otimes b)U_s^\dagger = (U_a a U_a^\dagger) \otimes (U_b b U_b^\dagger) \quad (4.41a)$$

$$= \sum_{n_a, n_b, m_a, m_b} a_{n_a m_a} b_{n_b m_b} e^{-i(\omega_{n_a m_a} + \omega_{n_b m_b})} |n_a n_b\rangle \langle m_a m_b| \quad (4.41b)$$

$$U_s(a \otimes b^\dagger)U_s^\dagger = (U_a a U_a^\dagger) \otimes (U_b b^\dagger U_b^\dagger) \quad (4.41c)$$

$$= \sum_{n_a, n_b, m_a, m_b} a_{n_a m_a} b_{m_b n_b}^* e^{-i(\omega_{n_a m_a} - \omega_{n_b m_b})} |n_a n_b\rangle \langle m_a m_b| \quad (4.41d)$$

Here, the matrix elements $a_{n_a m_a} = \langle n_a | a | m_a \rangle$ and $b_{n_b m_b} = \langle n_b | b | m_b \rangle$ are defined in their respective eigenbases, which may or may not be the Fock basis. This results in an effective coupling Hamiltonian that looks reminiscent of the drive hamiltonian we saw in the single

qubit case:

$$H'_c = U_s H_c U_s^\dagger \quad (4.42a)$$

$$= -g_\Delta \sum_{n_a, n_b, m_a, m_b} a_{n_a m_a} b_{m_b n_b}^* e^{-i(\omega_{n_a m_a} - \omega_{n_b m_b})t} |n_a n_b\rangle \langle m_a m_b| \quad (4.42b)$$

$$-g_\Sigma \sum_{n_a, n_b, m_a, m_b} a_{n_a m_a} b_{n_b m_b} e^{-i(\omega_{n_a m_a} + \omega_{n_b m_b})t} |n_a n_b\rangle \langle m_a m_b| + \text{c.c.} \quad (4.42c)$$

where we have sum and difference frequencies between energy levels that are connected via a non-zero matrix element.

Resonant Coupling

For a transmon qubit, the system eigenstates are approximately the fock states, allowing us to simplify the interaction Hamiltonian to

$$H'_c = - \begin{pmatrix} 0 & 0 & 0 & g_\Sigma e^{-i\Sigma t} & 0 & 0 & \dots \\ 0 & 0 & g_\Delta e^{-i\Delta t} & 0 & 0 & 0 & \dots \\ 0 & g_\Delta e^{i\Delta t} & 0 & 0 & 0 & 0 & \dots \\ g_\Sigma e^{i\Sigma t} & 0 & 0 & 0 & \sqrt{2}g_\Delta e^{i(\Delta - \alpha_b)t} & \sqrt{2}g_\Delta e^{-i(\Delta + \alpha_a)t} & \dots \\ 0 & 0 & 0 & \sqrt{2}g_\Delta e^{-i(\Delta - \alpha_b)t} & 0 & 0 & \dots \\ 0 & 0 & 0 & \sqrt{2}g_\Delta e^{i(\Delta + \alpha_a)t} & 0 & 0 & \dots \\ \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots \end{pmatrix} \quad (4.43)$$

where $\Delta = \omega_{01}^{(a)} - \omega_{01}^{(b)}$ and $\Sigma = \omega_{01}^{(a)} + \omega_{01}^{(b)}$ are the difference and sum frequencies between the two transmons, and α_i are their anharmonicities. Here, we notice that due to the nonlinearity of the transmon, the matrix elements connecting each pair of levels evolve at a different frequency, which allows us to select for a transition. If the qubits are tunable, we can engineer time dependent detuning $\Delta(t)$ that allows us to selectively turn on an interaction. Suppose we bring the two transmons into resonance such that $\omega_a = \omega_b \Rightarrow \Delta = 0$. Then, the matrix elements $\langle 01 | H'_c | 10 \rangle$ and $\langle 10 | H'_c | 01 \rangle$ will lose their time dependence and be the only “slow” terms in the Hamiltonian. The resulting Hamiltonian will then be

$$H'_c \approx g_\Delta [|01\rangle \langle 10| + |10\rangle \langle 01|] \quad (4.44)$$

which looks exactly like the Hamiltonian for a driven single qubit (see Eq. 4.12). The resulting unitary will also look similar, and if we embed this back into the full two qubit

Hilbert space, we obtain

$$U'_d(t) = \exp[-iH'_c t] = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(g_\Delta t) & -i \sin(g_\Delta t) & 0 \\ 0 & -i \sin(g_\Delta t) & \cos(g_\Delta t) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (4.45)$$

This again looks like a rotation, but with a rate $2g_\Delta$ and about a fixed rotation, since there is no phase associated with the coupling. However, since it spans a non-local subspace of the two-qubit Hilbert space, it leads to swapping of population between the two qubits, and can be used to generate entanglement. In particular, we obtain the unitary for an iSWAP ($\sqrt{\text{iSWAP}}$) if we set $g_\Delta \tau = \pi/2$ ($g_\Delta t = \pi/4$).

Due to the frequency selectivity of these transitions, we could also choose to do a rotation in the $|11\rangle$ - $|02\rangle$ (set $\omega_a = \omega_b + \alpha_b$) or $|11\rangle$ - $|20\rangle$ (set $\omega_a + \alpha_a = \omega_b$) subspaces. Here, the resulting rotation leads to an excursion outside the computational subspace, but if a full 2π rotation is performed, both the $|11\rangle$ and $|02\rangle$ ($|20\rangle$) states gain a phase $\cos(\pi) = -1$, and the resulting unitary in the computational subspace looks like a CZ gate. This geometric phase arises from the total solid angle area enclosed under the closed trajectory⁷. On the other hand, suppose we do a full rotation such that $g_\Delta \tau = \pi$. Then we find that although $U'_d(t)$ leads to a phase $e^{i\pi} = -1$ on both $|01\rangle$ and $|10\rangle$, this is equivalent to a $R_z(\pi)$ rotation on each individual qubit and no entanglement is generated⁸.

Given that we can generate 3 different classes of maximally entangling two qubit gates with these resonant subspace rotations, a natural question to ask is which one is better. From the perspective of gate speed, it would seem that the $\sqrt{\text{iSWAP}}$ is the clear winner, since it the smallest rotation, it can be implemented theoretically in the shortest time $g_\Delta \tau = \pi/4$. For a fixed coupling rate, the iSWAP gate takes twice as long as the $\sqrt{\text{iSWAP}}$ and the CZ takes $\sqrt{2} \times$ longer than the iSWAP⁹. On the other hand, the iSWAP-like gates generally suffer from more control error, due to off-resonant interactions in the two-photon subspace. Just as with single qubit gates, this off-resonant coupling can lead to either population transfer to leakage states or a frequency shift in the relevant transition, depending on whether it is adiabatic. In the two photon subspace, the Hamiltonian looks like

$$H_{2p} = \begin{pmatrix} 0 & \sqrt{2}g_\Delta & \sqrt{2}g_\Delta \\ \sqrt{2}g_\Delta & -\alpha_b + \Delta & 0 \\ \sqrt{2}g_\Delta & 0 & -\alpha_a - \Delta \end{pmatrix} \quad (4.46)$$

⁷The fact that a 2π rotation leads to a phase -1 is related to the double-covering of $\text{SU}(2)$ over $\text{SO}(3)$. It takes a 4π rotation to obtain a true identity operation.

⁸It is often said that the $|11\rangle$ - $|20\rangle$ rotation can generate entanglement because it involves an intermediate state that is outside the computational subspace. However, this cannot be correct since the $|11\rangle$ - $|20\rangle$ is still entangling for a system of two qutrits, despite the trajectory being contained within the two qutrit Hilbert space.

⁹We gain a factor of $\sqrt{2}$ since the coupling rate in the subspace is $\sqrt{2}g_\Delta$, but we need to do a full 2π rotation instead of a π rotation for the iSWAP.




Figure 4.3: **Fermionic simulation gates.** The phase space of the fSim or fermionic simulation gates, which are parameterized by a swap angle θ and a conditional phase ϕ . Many different canonical two-qubit gates lie within this space of gates, including the iSWAP, $\sqrt{\text{iSWAP}}$, CZ/CNOT gate and the Berkeley gate [216].

where we remove the time dependence in the Hamiltonian by making a unitary transformation $U(t) = e^{i(\alpha_b - \Delta)t} |02\rangle\langle 02| + e^{i(\alpha_a + \Delta)t} |20\rangle\langle 20|$. In general, this is difficult to analyze analytically, but when the two qubits are on resonance ($\Delta = 0$) and their anharmonicities are equal $\alpha \equiv \alpha_a = \alpha_b$, we can block diagonalize the Hamiltonian by moving into the basis defined by $|11\rangle$, $|B\rangle = (|02\rangle + |20\rangle)/\sqrt{2}$ and $|D\rangle = (|02\rangle - |20\rangle)/\sqrt{2}$. In this basis, H_{2p} is transformed to

$$H_{2p} = \begin{pmatrix} 0 & 2g_\Delta & 0 \\ 2g_\Delta & -\alpha & 0 \\ 0 & 0 & -\alpha \end{pmatrix} \quad (4.47)$$

and we see that the state $|D\rangle$ decouples¹⁰ completely from the computational state $|11\rangle$. The remainder of the system looks like the off-resonantly driven system, so this will lead to off-resonant population swapping between $|11\rangle$ and $|B\rangle$ at rate Λ with a maximum leakage population $P_{\text{leakage}} = (4g_\Delta/\Lambda)^2$, where $\Lambda = \sqrt{16g_\Delta^2 + \alpha^2}$ according to Eq. 4.20. If the gate is non-adiabatic and the coupling g_Δ is also tunable¹¹, this leakage population can be minimized by carefully matching g_Δ to α such that $\Lambda = 4ng$ for $n \in \mathbb{Z}$ [12]. Otherwise, $\Delta(t)$ must be ramped slowly enough and g_Δ must be small enough such that the eigenstates in the two-photon manifold follow an adiabatic trajectory where there is no population transfer and the states only gain a phase.

Even if leakage is suppressed for the iSWAP-like gates, the remaining phase accumulation due to the off-resonant coupling can still be problematic. Unlike with single-qubit gates,

¹⁰The states $|B\rangle$ and $|D\rangle$ are often referred to as “bright” and “dark” states, since the dark state $|D\rangle$ decouples from the rest of the system.

¹¹Even if g_Δ is not tunable in the traditional sense, it typically scales linearly with frequency ($g \propto \omega$). Then, the “fixed” value of g can be fine-tuned by selecting the interaction frequency $\omega \equiv \omega_a = \omega_b$.

the frequency shift on the $|11\rangle$ state leads to an entangling phase, which cannot be easily corrected with virtual phase updates. This leads a modified unitary

$$U'_d = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(\theta) & -i \sin(\theta) & 0 \\ 0 & -i \sin(\theta) & \cos(\theta) & 0 \\ 0 & 0 & 0 & e^{i\phi} \end{pmatrix} = \text{fSim}(\theta, \phi) \quad (4.48)$$

where ϕ is the extra entangling phase, and this unitary is often referred to as the fermionic simulation gate [99, 65]. Generally speaking, an extra conditional phase on an iSWAP gate will reduce the entangling power of the gate. In fact, if $\phi = \pi$, the unitary becomes locally equivalent to a SWAP gate, which is fully classical. Small values of ϕ , however, can generally be tolerated by adapting compilation algorithms to account for the “non-standard” gate in many NISQ experiments. In fact, Google Quantum AI’s 2019 supremacy experiment used a “Sycamore gate” ($\text{fSim}(\theta = \pi/2, \phi = \pi/6)$) as their primary entangling gate for the random circuit sampling [4]. In other cases, when a true iSWAP is needed, the g_Δ can be lowered at the cost of a longer gate time, since the phase $\phi \propto g_\Delta^2/\alpha$. Then, any residual conditional phase can be “balanced” with single qubit rotations to redistribute the phase with virtual Z gates to the $|01\rangle$ and $|10\rangle$ states, which lowers the effective coherent error on the gate [56, 132].

The CZ gate by construction does not suffer from these conditional phase errors, and as a result, it is often preferred to the iSWAP-like gates. This is especially true in the context of QEC circuits, which up until recently were thought to be efficiently implemented only with CNOT/CZ-like gates [126]. However, the diabatic CZ gate naturally suffers from more leakage, since the $|20\rangle$ (or $|02\rangle$) state is explicitly fully populated during the gate. One solution to this leakage problem is to engineer an adiabatic trajectory such that the $|20\rangle$ ($|02\rangle$) state is never explicitly populated.

Off-Resonant Coupling

While the resonant coupling discussed so far is a simple and straightforward way to implement fast two qubit interactions, we can still obtain an entangling in the detuned limit, just as with off-resonant driving. This was alluded to when discussing conditional phase errors in the context of the iSWAP-like gates, and it is often called a ZZ -interaction, since $\sigma_z \otimes \sigma_z$ is the generator of a conditional phase rotation. We can alternatively think of the conditional phase or ZZ -interaction as state dependent frequency shift, since the state of one qubit will add a $\pm\sigma_z$ term to the Hamiltonian, which is proportional to the qubits frequency, and we define the ZZ rate ζ as

$$\zeta = (\omega_{11} - \omega_{10}) - (\omega_{01} - \omega_{00}) \quad (4.49a)$$

$$= (\omega_{11} - \omega_{01}) - (\omega_{10} - \omega_{00}) \quad (4.49b)$$

$$= \omega_{00} + \omega_{11} - \omega_{01} - \omega_{10} \quad (4.49c)$$

which is the frequency difference between the target qubit when it's spectator qubit is in the $|1\rangle$ state versus the $|0\rangle$ state. It is worth noting that the definition is symmetric with respect to which qubit is designated the spectator and which is the target. This is problematic for single qubit gates, since a resonant drive on one qubit that was calibrated when the neighboring qubits were in the ground state will now be off-resonant if the neighboring qubits are in the excited state. As a result, we typically want to keep $\zeta \lesssim \Gamma_2$, since the dephasing time Γ_2 sets a limit on the frequency precision of the qubit.

We now calculate the ZZ rate ζ as a function of the coupling strengths g_Δ, g_Σ . While this is often calculated in the literature using a Schrieffer-Wolff transformation, we believe that this effect is more clearly understood through the lens of perturbation theory. Suppose we have two transmons with fixed fundamental frequencies ω_a, ω_b such that their detuning $\Delta = \omega_a - \omega_b$ is also fixed. If we also have a non-zero coupling term, like in Eq. 4.40, that is always present, the computational states become the eigenstates of the joint Hamiltonian, with shifted eigenfrequencies. We can calculate the perturbed eigenfrequencies with perturbation theory, treating H_c as a perturbation on top of H_a and H_b . Since $\langle nm | H_c | nm \rangle = 0$ for all states $|nm\rangle$ (H_c is zero along the diagonal), the first order frequency correction vanishes and we must go to second order in perturbation theory. If we only consider the lowest two states of each transmon, the second order corrections to the energies are

$$\omega_{00}^{(2)} = -\frac{g_\Sigma^2}{\Sigma} \quad \omega_{01}^{(2)} = \omega_b - \frac{g_\Delta^2}{\Delta} \quad \omega_{10}^{(2)} = \omega_a + \frac{g_\Delta^2}{\Delta} \quad \omega_{11}^{(2)} = \Sigma + \frac{g_\Sigma^2}{\Sigma} \quad (4.50)$$

so the total ZZ rate ζ is given by

$$\zeta = \Sigma + \frac{g_\Sigma^2}{\Sigma} - \frac{g_\Sigma^2}{\Sigma} - \omega_b + \frac{g_\Delta^2}{\Delta} - \omega_a - \frac{g_\Delta^2}{\Delta} = 0 \quad (4.51)$$

Surprisingly, although both qubit frequencies are shifted by the perturbation, there is no conditional frequency shift if both systems are true two-level systems.

If we take into account the higher levels of the system, we find the additional corrections

$$\omega_{01}^{(2)} = \omega_b - \frac{g_\Delta^2}{\Delta} - \frac{2g_\Sigma^2}{\Sigma + \alpha_b} \quad \omega_{00}^{(2)} = -\frac{g_\Sigma^2}{\Delta} \quad (4.52a)$$

$$\omega_{10}^{(2)} = \omega_a + \frac{g_\Delta^2}{\Delta} - \frac{2g_\Sigma^2}{\Sigma + \alpha_a} \quad \omega_{11}^{(2)} = \Sigma + \frac{g_\Sigma^2}{\Sigma} - \frac{2g_\Delta^2}{\alpha_a + \Delta} - \frac{2g_\Delta^2}{\alpha_b - \Delta} \quad (4.52b)$$

and the ZZ rate is now

$$\zeta = -2g_\Delta^2 \left[\frac{1}{\alpha_a + \Delta} + \frac{1}{\alpha_b - \Delta} \right] + 2g_\Sigma^2 \left[\frac{1}{\Sigma + \alpha_a} + \frac{1}{\Sigma + \alpha_b} \right] \quad (4.53)$$

The second term is typically dropped¹², since $\Sigma = \omega_a + \omega_b$ is usually 10–100× larger than $\Delta = \omega_a - \omega_b$, making it small with respect to the first term. It is instructive now to examine

¹²Although some recent works, such as Ref. [113], argue for retaining this term.

the various limits of the first term in Eq. 4.53. As α_a, α_b become very large compared to Δ and g_Δ the systems involved will closer approximate at true two-level system and we see that $\zeta \rightarrow 0$ in agreement with our previous analysis. This tells us that the ZZ interaction is completely a higher-level effect. In the opposite limit, where $\alpha_a, \alpha_b \rightarrow 0$ the level shifts cancel out as well, so the ZZ -interaction is also a nonlinear effect. For practical systems that are only weakly anharmonic, the simplest way to minimize the ZZ -rate is to reduce the coupling rate g_Δ or increase the detuning Δ , both of which suppress the off-resonant interaction. However, if neither of these is possible due to other constraints, we can also employ more exotic strategies to suppress the interaction. One such method involves coupling a transmon to a capacitively shunted flux qubit, which has a frequency tunable anharmonicity that can be made positive in contrast to the negative anharmonicity of the transmon [109, 218]. If the anharmonicities α_a, α_b are equal magnitude but opposite in sign ($\alpha_a = -\alpha_b$), we see from Eq. 4.53 that $\zeta = 0$.

Finally, we note that even when both systems are transmons such that $\alpha_a, \alpha_b < 0$, all hope is not lost. Under these circumstances, we can still cancel ζ when g_Δ/Δ is not small if the two qubits are in the “straddling regime” ($|\Delta| < |\alpha|$), and are coupled via an intermediate higher-frequency mode. This was demonstrated in Refs. [145, 193] with frequency tunable qubits/couplers, and later in Ref. [92] with a fully fixed-frequency system. We will not show the full calculation here, since it requires going to 4th order in perturbation theory, but the key insight is that when the two transmons are in the straddling regime, the sign of their contribution to ζ is positive. In other words, the $|200\rangle$ and $|002\rangle$ states now push the $|101\rangle$ states up in frequency (the middle index now refers to the coupler state). Since the coupler frequency is above both qubits, the $|110\rangle$ and $|011\rangle$ states will push the $|101\rangle$ state down in frequency, and these contributions can be tuned to cancel the total ZZ rate ζ , even when g_Δ/Δ is not small. This example helps to illustrate the utility of viewing the ZZ -interaction through the lens of energy level repulsion.

Finally, we point out that although the ZZ -interaction is often viewed as a parasitic and unwanted interaction in the context of many gate architectures, we can also leverage the ζ to implement a CZ gate between far-detuned qubits if we have control over the coupling g_Δ . This involves simply increasing the coupling $g_\Delta(t)$ adiabatically to increase $\zeta(t)$ for a period of time [190]. The acquired conditional phase is then given by $\phi_{11} = \int_0^T \zeta(t) dt$. The advantage of doing an adiabatic CZ gate like this is that the larger detuning make it easier to maintain an adiabatic trajectory. However, this comes at a cost of slower interaction rate.

Microwave-Activated Coupling

Although flux-tunable qubit frequencies and coupling rates enable strong interaction rates, this typically adds both more control complexity and additional decoherence channels. As a result, there was a strong focus through most of the 2010’s to develop and optimize two-qubit gates based on microwave driving in a fixed coupling architecture. Because single qubit gates already require a microwave drive on each qubit, this adds no extra overhead. The most common of these gate schemes is known as the cross-resonance gate, which involves driving

a “control” qubit at the frequency a dispersively coupled “target” qubit [167, 37]. We will not derive the effective unitary here, but it can be shown that this leads to an effective Hamiltonian in the interaction frame

$$H_{\text{CR}} = \tilde{\varepsilon}_0(t) |\overline{01}\rangle\langle\overline{00}| + \tilde{\varepsilon}_1(t) |\overline{11}\rangle\langle\overline{10}| + \text{c.c.} \quad (4.54)$$

where $\tilde{\varepsilon}_0(t)$ ($\tilde{\varepsilon}_1(t)$) is the effective drive strength on the target qubit when the control qubit is in the $|0\rangle$ ($|1\rangle$) state [198]. Here, the overline on the states reminds us that the computational states are “dressed”, or hybridized, by a static coupling strength g_Δ , as discussed in the previous section. This leads to both a state-dependent Rabi rate, $\mu(t) = (\tilde{\varepsilon}_0(t) - \tilde{\varepsilon}_1(t))/2$, in addition to a state-independent Rabi rate, $\nu(t) = (\tilde{\varepsilon}_0(t) + \tilde{\varepsilon}_1(t))/2$, on the target qubit. As a result, Eq. 4.54 is often expressed as

$$H_{\text{CR}} = \mu(t)\sigma_z^{(c)} \otimes \sigma_x^{(t)} + \nu(t)I^{(c)} \otimes \sigma_x^{(t)} \quad (4.55)$$

The effective drive strengths, including the effect of the higher levels of the transmon, are given by

$$\tilde{\varepsilon}_0(t) = -\frac{g_\Delta}{\Delta}\varepsilon(t) \quad \tilde{\varepsilon}_1(t) = -\frac{g}{\Delta}\frac{\Delta + \alpha_c}{\Delta - \alpha_c}\varepsilon(t) \quad (4.56)$$

$$\Rightarrow \mu(t) = \frac{\tilde{\varepsilon}_0(t) - \tilde{\varepsilon}_1(t)}{2} = \frac{2g_\Delta\alpha_c}{\Delta(\alpha_c - \Delta)}\varepsilon(t) \quad (4.57)$$

where $\varepsilon(t)$ is the bare drive strength on the control qubit, g_Δ is the coupling strength between the qubits, and $\Delta = \omega_c - \omega_t$ is the frequency difference between the control and target qubits [198]. If this drive is turned on for a duration τ_g such that $\int_0^{\tau_g} \mu(t) dt = \frac{\pi}{2}$, we can realize an effective unitary that is locally equivalent¹³ to

$$U_{\text{CR}} = e^{-i\frac{1}{2}\sigma_z^{(c)} \otimes \sigma_x^{(t)} \int_0^{\tau_g} \mu(t) dt} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i & 0 & 0 \\ -i & 1 & 0 & 0 \\ 0 & 0 & 1 & i \\ 0 & 0 & i & 1 \end{pmatrix} \quad (4.58)$$

and this can be made to look like a CNOT gate by adding an additional $R_z(-\pi/2)$ rotation on the control qubit and $R_x(-\pi/2)$ rotation on the target qubit:

$$(R_z(-\pi/2) \otimes R_x(-\pi/2)) U_{\text{CR}} = \begin{pmatrix} 1 & i & 0 & 0 \\ i & 1 & 0 & 0 \\ 0 & 0 & 1 & i \\ 0 & 0 & i & 1 \end{pmatrix} \begin{pmatrix} 1 & -i & 0 & 0 \\ -i & 1 & 0 & 0 \\ 0 & 0 & -i & 1 \\ 0 & 0 & 1 & -i \end{pmatrix} \quad (4.59a)$$

$$= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} = \text{CNOT} \quad (4.59b)$$

¹³We drop the $\exp[-i\sigma_z^{(c)} \otimes \sigma_x^{(t)} \int_0^{\tau_g} \mu(t) dt]$ term, since this commutes with entangling ZX term and can easily be canceled with a single qubit drive on the target qubit.

While this generates a convenient unitary, the CR gate has many drawbacks, the biggest of which is a tension between minimizing idle ZZ errors and maximizing the gate speed. The effective CR interaction is maximized when the two transmons are in the straddling regime (Δ between $\alpha/2$ and α), and reaches a maximum rate¹⁴ $\mu \lesssim 1.5g_\Delta$ in this regime [198]. For typical parameters of $\alpha/(2\pi) \approx 300$ MHz, $g_\Delta/(2\pi) \approx 4$ MHz, and $\Delta/(2\pi) \approx 200$ MHz, this leads to a ZZ -rate $\zeta \approx 400$ kHz and a theoretical minimum CNOT gate time of $\tau_g \approx 80$ ns [181]. Since the coupling is fixed, this ζ is always present and will cause not only single-qubit gate errors, but also errors in the two-qubit gate, since the ZX -interaction engineered by the cross-resonance effect does not commute with the ZZ -interaction. Moreover, even for the theoretically maximum gate rate, these parameters achieve an effective on/off ratio of $\mu/\zeta \approx 10$, and practical gate speeds are usually further limited by other issues such as leakage or crosstalk. More recently, this on/off ratio has been improved by leveraging the multipath coupler discussed in the previous section, and on/off ratios of $\mu/\zeta \approx 130$ have been achieved [92]. However, this still falls short of the on/off ratios achievable with flux-tunable couplers.

These challenges led to the development of a related off-resonant version of the cross-resonance gate, in which both transmons are driven at a detuned frequency to create a conditional Stark shift [138, 206]. Since this interaction commutes with the static ZZ interaction, this can be used to either enhance or suppress the static ZZ rate ζ_s . The total driven ZZ rate, ζ_d , as a function of the drive strength is approximately

$$\zeta_d = \zeta_s + \frac{2g_\Delta\alpha_a\alpha_b\Omega_a\Omega_b\cos(\phi_a - \phi_b)}{\Delta_{a,d}\Delta_{b,d}(\Delta_{a,d} + \alpha_a)(\Delta_{b,d} + \alpha_b)} \quad (4.60)$$

to first order in g_Δ [206]. Here, $\Delta_{i,j} = \omega_i - \omega_j$, where ω_a, ω_b are the frequencies of the two qubits and ω_d is the frequency of the drives on both qubits. We see also that the sign of the interaction depends on the relative phase $\phi_a - \phi_b$ between the drives on each qubit. By turning on this interaction for a time $\zeta_d\tau_g = \pi$, a CZ gate can be implemented up to local phases. While the driven ZZ rate, ζ_d , is comparable to the effective $\sigma_z \otimes \sigma_x$ rate μ for the same coupling strength, it is possible to achieve much larger on-off ratios with the conditional Stark shift, since the static ZZ rate ζ_s can be actively nulled during single-qubit operations. This also allows larger static coupling rates to be used, which can further increase the entanglement rate. While these properties make the conditional Stark shift gate favorable compared to the cross-resonance interaction, it is still necessary for the two transmons to be in the straddling regime to achieve a reasonable ZZ rate. Unfortunately, this restricts qubit-qubit detunings to $|\Delta| \lesssim 300$ MHz, which imposes challenging frequency targeting requirements for fixed-frequency transmons [139].

¹⁴Eq. 4.56 is valid only for small $\varepsilon(t)$.

4.3 Frequency Crowding

When choosing between the various two-qubit gate schemes discussed so far, it is important to not only consider the achievable gate fidelities in isolation, but also how well the architecture scales with many simultaneously executed gates. Crosstalk and imperfect signal isolation can lead to increased gate errors under parallel operation even in current state-of-the-art systems. While it is possible to mitigate classical crosstalk with active cancellation protocols [138], it is typically still necessary to ensure that targeted transitions in nearby qubits are spectrally detuned for sufficient addressability of individual transitions. Furthermore, smaller detunings can cause elevated ZZ -interactions between coupled qubits. With flux-tunable qubits, qubit frequencies can be dynamically tuned over a large range ($\mathcal{O}(1)$ GHz) and optimized in-situ to mitigate these error sources and ensure addressability [102, 101]. Fixed-frequency architectures, on the other hand, are not afforded this luxury, and even have stricter frequency targeting requirements for optimal gates. This, coupled with large, fabrication-limited, uncertainties in manufactured qubit frequencies, make it extremely challenging to determine practical target frequencies for fixed-frequency devices at scale. In this section, we analyze how these fabrication limits in frequency targeting affect the scaling properties of fixed-frequency, all-microwave, gate architectures. The work in the subsequent sections was first reported in Ref. [139], and was performed jointly with Alexis Morvan and Jeff Larson.

Critical Current Dispersion in Josephson Junctions

The frequency and anharmonicity of a transmon qubit are determined by the Josephson energy E_J (or equivalently the critical current $I_c = E_J/(2\pi\Phi_0)$), of the junction and the total capacitive charging energy E_C across the junction. The total capacitance for most transmon qubits is dominated by capacitor pads at the scale of tens to hundreds of microns, and can typically reproducibly defined with submicron precision. The junction critical current, however, is defined by the area of a small ($\sim 100 \times 100$ nm) Al/AlOx junction and even state-of-the-art fabrication processes struggle to achieve a relative precision below 2% at wafer-scale [107, 156, 199, 140]. This makes the relative spread in junction critical currents $\sigma_{I_c}/\langle I_c \rangle$ the dominant contribution to the frequency dispersion. In particular, we have $\sigma_f/\langle f \rangle \approx \frac{1}{2}\sigma_{I_c}/\langle I_c \rangle$, since $f_{01} \propto \sqrt{E_J}$. For a typical transmon with $f_{01} \sim 5$ GHz, this represents a best case spread of $\sigma_f \sim 50$ MHz, which may only be 4–6 $\times$ smaller than the transmon anharmonicity.

The junction critical current can be estimated from room-temperature probing of the junction normal state resistance R_n , which is related to I_c via the Ambegoakar-Baratoff relation [7]:

$$I_c = \frac{\pi\Delta(T)}{2eR_n} \quad (4.61)$$

where $\Delta(T)$ is the BCS superconducting gap¹⁵. As a result, devices can be pre-screened prior to device packaging to determine the best candidates. Recently, several groups have also demonstrated post-fabrication trimming protocols, based on either local heating via lasers [84, 95, 213] or voltage annealing [158, 204], to adjust junction critical currents at room temperature. These techniques can further improve the frequency targeting precision down to 15–20 MHz ($\sim 15\%$), which becomes primarily limited by the correlation between R_n and empirically measured frequencies at cryogenic temperatures [84, 204]. While this represents a significant improvement, scaling to devices with more than a few hundred qubits still poses a challenge due to the low expected yield of frequency-optimal devices.

Frequency Constraints

Frequencies collisions occur when an unwanted degeneracy leads to a degradation in control fidelity, typically due to unwanted couplings or crosstalk within the device. For example, when two nearby¹⁶ qubits have a similar $|0\rangle\text{--}|1\rangle$ transition frequencies, driving one qubit will induce a frequency shift on the neighboring qubit that elevates the gate error. To avoid such a situation, we can establish a lower bound, δ , on the minimum detuning between f_{01} for neighboring qubits. The bound δ can be estimated by simulating the relevant gate dynamics with an measured or estimated crosstalk coefficient, and enforcing a minimum tolerable gate error. These frequency requirements lead to linear constraints of the form

$$|f_d - f_t| \geq \delta, \quad (4.62)$$

where f_d is the frequency of a drive frequency, and f_t is the target frequency to avoid. Since the connectivity or lattice structure of the qubit device will determine the number of constraints, we formalize this by defining a directed graph G , where each vertex $v \in V$ represents a qubit and each edge $(v_i, v_j) \in \vec{E}$ represents a possible entangling interaction between two qubits. The orientation of the edge is necessary to account for interactions such as the CR gate that are not symmetric with respect to the two qubits. Table 4.3 lists the set of constraints that we consider in this thesis, which we explain below.

Addressability

The first category of constraints that we consider involve single-qubit drives, and are thus widely applicable to almost all superconducting qubit architectures. To avoid unwanted Stark shifts on qubit j when applying a pulse at the $|0\rangle\text{--}|1\rangle$ frequency of qubit i , we must ensure that f_i is sufficiently detuned from both f_j and $f_j + \alpha_j$. These constraints are referred to as type *A* in Table 4.3.

¹⁵Well below T_c we can approximate $\Delta(T) \approx \Delta(0)$. We typically use a value of $\Delta(0) \approx 170 \mu\text{eV}$ [108].

¹⁶Crosstalk, unfortunately may not always be local. See Ref. [2] for an example of device-wide crosstalk.

		Definition	Participants	Bounds
Addressability	A1	$ f_i - f_j \geq \delta_{A1}$	$(i, j) \in E$	17 MHz
	A2	$ f_i - f_j - \alpha_j \geq \delta_{A2}$	$(i, j) \in E$	30 MHz
Entanglement	C1	$f_i + \alpha_i \leq f_d \leq f_i$	$(i, j) \in \vec{E}$	—
	E1	$ f_d - f_i \geq \delta_{E1}$	$(i, j) \in \vec{E}$	17 MHz
	E2	$ f_d - f_i - \alpha_i \geq \delta_{E2}$	$(i, j) \in \vec{E}$	30 MHz
	D1	$ f_d - f_i - \alpha_i/2 \geq \delta_{D1}$	$(i, j) \in \vec{E}$	2 MHz
Spectators	S1	$ f_d - f_k \geq \delta_{S1}$	$(i, j, k) \in N$	17 MHz
	S2	$ f_d - f_k - \alpha_k \geq \delta_{S2}$	$(i, j, k) \in N$	25 MHz
	T1	$ f_d + f_k - 2f_i - \alpha_i \geq \delta_{T1}$	$(i, j, k) \in N$	17 MHz

Table 4.1: **Frequency constraints for a fixed-frequency, fixed-coupling gate architecture.** We group these constraints into three categories (described in the main text). The index refers to the participants and is defined by the system graph. $\vec{E}$ represents the set of oriented graph edges, and E represents the corresponding unoriented edges, meaning that each edge is taken into account twice in each direction. The spectator constraints have different participants for the CR and CZ cases. In the CR case, the pulse is applied only to the control qubit, whereas in the CZ case, the pulse is applied to both entangled qubits. The variable f_d represents the frequency of the applied microwave drive. For the CZ case, the drive frequency is no longer fixed. The bounds are taken from Ref. [84].

Entangling Gates

The constraints for entangling gates depends on the choice of two-qubit gate. For this study, we chose to focus on the cross-resonance (CR) gate and the conditional Stark shift (CZ) gate, since these had demonstrated the highest gate fidelities at the time for a fixed-frequency and fixed-coupling architecture. However, this analysis can easily be extended for other gate schemes. Both CR and CZ gates perform optimally when the qubits involved are in the straddling regime ($|\Delta_{ij}| < |\alpha_i|, |\alpha_j|$), which enforces a constraint on the frequency difference of each edge (C1). In addition, the drive frequency for the two qubit gate must avoid both the $|0\rangle\text{--}|1\rangle$ (E1) and $|1\rangle\text{--}|2\rangle$ (E2) transitions of both qubits involved. In the case of the CR gate, the drive frequency will be at the frequency of the target qubit, so this leads to a redundant constraint with the single-qubit constraints; however, the drive frequency for the CZ gate can be freely varied. In addition, since the two-qubit drive strengths are typically larger than for single-qubit drives, we also include the two-photon transition at $f_{01} + \alpha/2$ as another frequency to avoid (D1). Finally, we must also consider the effect of a drive on neighboring qubits not involved in the entangling operation, which we refer to as *spectators* of the entangling gate [192]. These constraints are listed as S1, S2 (single-photon spectator collisions) and T1 (two-photon spectator collision) in Table 4.3.

Optimization Procedure

The goal of our optimization procedure is to find a set of frequencies for a given connectivity graph. In particular, we use the set of constraints shown in Table 4.3. Since the set of parameters in this case is disjoint, we use mixed-integer programming with the Gurobi solver [77] and the Python package PYOMO [80].

Initially, we found that a naive objective function yielded solutions that are on the edges of the collision-free regions, meaning that small perturbations in frequencies would result in a considerable number of collisions. As a result, this naive optimization yields a design that is not robust against the frequency dispersion introduced by the fabrication process. To circumvent this, we developed a three-step approach that serves as a proxy for the yield. In particular, we seek to maximize the distance between the target frequency and the collision regions, since this reduces the likelihood of collisions.

To move the solution frequencies away from the boundaries of the collision regions, we introduce a new variable for each set of constraints that corresponds to the distance of the left-hand side of the constraint definitions in Table 4.3. We then add this to the threshold δ for each constraint, and introduce an additional constraint that forces these distances to be equal for each constraint type and edge. This finds the largest hypersphere of radius R that can fit inside the zero-collision region. In the second step, we relax the requirement that the threshold of each type (rows in Table 4.3) be equal, but we still require that they be equal for each edge. We also enforce a constraint specifying that all of these distances are larger than the radius R calculated in the previous step. This amounts to allowing for a hyperellipse with radius $R_i \leq R$ per constraint type. Finally, in the last step, we allow each edge to differ. In addition, for the CZ gate architecture, we add an additional variable to each edge that corresponds to the now continuous drive frequency, and optimize this for each edge. See Ref. [139] for more details on the optimization procedure.

Given a solution set, we evaluate it by estimating the expected yield as a function of the frequency dispersion. We model the yield of zero-collision devices with Monte Carlo simulations, following the methods in Ref. [84]. We assume that the fabricated junction frequencies follow a normal distribution with a mean around the target frequency and a variable standard deviation σ_f . Then, for each lattice geometry and target frequency map, we sample 1000 frequency configurations to determine the estimated yield as a function of σ_f . The yield is computed by counting the number of sample configurations that have zero-collisions based on the specified constraints.

Small-Scale Systems

To build some intuition about the results of frequency optimization, we first applied the procedure to an eight-transmon ring [21]. This linear topology is the minimal connectivity possible for a quantum processor and sets the upper limit on the fabrication yield that one can expect. In Fig. 4.4, we present the frequency pattern obtained with optimization described in the preceding section, along with the allowed regions for the frequencies for each transmon. We also present the yield as a function of the frequency dispersion, with state-of-the-art fabrication spreads with and without post-processing denoted to give a sense of achievable yields [107, 84].

For the qubit CR-based architecture, we find that even a linear topology with state-of-the-art dispersions of 50 MHz leads to a low yield of $\sim 10\%$. This yield can be increased substantially with post-processing enabled dispersion of ~ 14 MHz, which highlights the need to integrate such




Figure 4.4: **Frequency optimization for an eight-qubit ring.** (a) The optimized frequencies for an eight-qubit ring with CR constraints. The collision regions for the $|0\rangle \rightarrow |1\rangle$ transition are marked by the shaded gray regions. The solid black line denotes the $|0\rangle \rightarrow |1\rangle$ transition, the dashed black line the $|1\rangle \rightarrow |2\rangle$ transition, and the dotted black line the two-photon transition $|0\rangle \rightarrow |2\rangle$. The colored distributions indicate normal distributions centered around the target frequency for the state-of-the-art frequency dispersions in Refs. [107] (dashed) and [84] (solid). (b) The yield of zero-collision devices for the two different gate architectures as a function of the average frequency dispersion. (c) The frequency of each collision type — single qubit (1Q), entangling (Ent.), and spectator (Spec.) — for the two different gate architectures given optimized qubit frequencies.

techniques to minimize the frequency dispersion. In Fig. 4.4c, we see the distribution of errors among the different types of collisions. The majority of the collisions come from the entanglement constraints associated with the CR gate — these significantly reduce the overall yield.

This motivated us to also consider the yield for a CZ-based architecture, which has the added flexibility of a variable drive frequency. As a result, the yield is higher and we see a larger plateau region where the yield is approximately unity. To calculate the yield for this specific architecture, we sampled random frequencies around the target layout frequency, and for each random sample, we used the optimizer to find a set of optimal driving frequencies for the entangling gate. For this architecture, dispersions of 50 MHz are already sufficient to reach a high yield of $\sim 90\%$, and an almost perfect yield can be attained with additional post-processing. Compared with the CR-based architecture, the collision rates decrease for both single- and two-qubit gates due to the relaxation of the drive frequency constraints for entangling operations.

Scaling to Lattices

As the number of transmons on a processor increases, so does the number of collisions. For large systems, individual optimization of all the frequencies across an entire chip becomes intractable. Instead, we must optimize a smaller unit cell of frequencies, which can then be tiled to generate a larger lattice. To construct such a solution, we start by defining the unit cell. To ensure that tiling is possible, we optimize the frequency layout with periodic boundary conditions, so that the solution will also be valid for the tiled construction. Since calculating the yield is quick for the CR case,




Figure 4.5: **Scaling frequency optimization to different lattice geometries.** Common lattice geometries for superconducting qubit systems are shown in (a)-(d). Each circle represents a qubit, and the links between them represent a possible entangling interaction between the two qubits. To estimate the yield up to large system sizes, we simulate each geometry with periodic boundary conditions. (a) We simulate the linear chain with a unit cell of 8 sites. (b) A square lattice with 16 sites. (c) A hexagonal lattice with 20 sites. (d) A heavy-hexagonal lattice with 15 sites. (e) We show the minimum frequency dispersion necessary to achieve a 10% yield for each lattice type and gate architecture. The dashed line indicates the best frequency spread achieved with post-fabrication trimming [84].

we can directly compute the zero-collision yield for these large lattices. However, for the CZ case, the additional optimization on each instance of the sampled frequencies make it computationally intractable to directly calculate the yield. Instead, we use a scaling argument to estimate a yield for the total system up to boundary effects. In particular, we compute

$$y = y_m^{N/n_m} = y_0^N \quad (4.63)$$

where y is the yield of the larger system, y_m is the yield of the unit cell subject to periodic boundary conditions, N is the total number of sites in the system, and n_m is the number of sites in the unit cell. Since the boundary effects will only reduce the number of potential collisions, this provides a lower bound on the yield, even for small system sizes. This method allows us to compare the scalability of different lattice geometries for both the CR and CZ gate architectures. Fig. 4.5 shows the different lattices with varying connectivity: a square lattice for the surface code, a heavy hexagonal lattice used by IBM, and a hexagonal lattice for comparison with its heavy counterpart. We show the frequency dispersion required to achieve a 10% yield for a 1000-qubit lattice for each geometry. Unsurprisingly, the more densely connected lattice has a lower yield. We note that, given the thresholds listed in Table 4.3 for a CR-based qubit architecture, none of these lattices can be fabricated with a 10% yield given even the best-demonstrated frequency dispersions with post-fabrication trimming. This shows that a collision-free 1000-qubit device is likely unachievable for the CR-based architecture without significant improvements in fabrication precision or a relaxation of the constraints. In contrast, the conditional Stark shift-based qubit architecture is able to scale to 1000 qubits for two of the four lattices with an already attainable frequency dispersion.

While we did not analyze other gate architectures beyond the CR- and CZ-based architectures, the methods described here can be readily applied to other architectures and qubit modalities.

However, we expect yields to be similarly low for the other fixed-frequency and fixed-coupling gate schemes, since they share the most restrictive requirement that qubit-qubit detunings fall within the straddling regime. This points to a need for alternative gate schemes for fixed-frequency qubits. Finally, we note that this yield problem can also be addressed by moving from monolithic chips to a more modular architecture where smaller chips with fewer qubits can be individually selected and connected to form a larger processor. This has become an increasing area of focus in recent years [72, 61, 152], and will likely be necessary for realizing a large-scale quantum processor.

4.4 Architectural Tradeoffs

What we have seen so far is that although single qubit gates are almost universally implemented with microwave driving, there is a large spectrum of possible two-qubit gate implementations with different architectural requirements. As a result, the desired two-qubit gate scheme typically has the strongest influence on the architecture of a large multi-qubit device. In this section, we will briefly discuss some of the tradeoffs involved in selecting a particular two-qubit gate scheme.

From the perspective of hardware requirements, we can divide the space of two qubit gates along two axis based on whether or not they require flux-tunable qubits and flux-tunable coupling. Generally speaking, we can engineer faster gates with more tunability. Since this allows us to selectively reduce the qubit-qubit detuning and increase the coupling rate only during the gate, we are less constrained by a need to mitigate parasitic ZZ -coupling. However, this comes at a cost of extra hardware overhead to support stable flux biases and additional dephasing and decoherence channels. Moreover, the baseband flux pulses used to implement these gates suffer from substantial pulse distortion and require complex pulse pre-compensation filters to suppress non-markovian errors [13, 169]. In contrast, the microwave-activated coupling schemes require no flux-tunability of either frequency or qubit-qubit coupling rates, generally leading to lower dephasing and decay rates. The tradeoff, however, is lower on/off ratios for the entangling interaction rate, which limits the maximum speed of a two-qubit gate.

In the following chapter, we describe an experimental implementation of an alternative gate architecture that aims to address all of these concerns. By adding a limited amount of flux-tunability only in the coupler, we can not only activate strong two-qubit interactions between far-detuned transmons, but also suppress stray entanglement while idling or during single-qubit gates. Restricting the tunability to the coupler also helps to minimize excess decoherence from the additional control channel, and activating interactions with parametric modulation avoids baseband pulse distortion. While these advantages come at the cost of lower entanglement rates compared to resonant coupling, we will show that this can be a worthwhile compromise for a more scalable architecture.

Chapter 5

Parametric Gates

In the previous chapter, we discussed various gate architectures and their advantages and drawbacks. In this chapter, we will advocate for an alternative gate architecture that involves parametric modulation of a tunable coupler to engineer two-qubit gates between fixed-frequency qubits. This approach to two-qubit gates resolves the most stringent frequency constraints of a static coupling architecture, by enabling fast-entanglement generating between far-detuned transmons.

5.1 Parametric Modulation

Parametric gates are somewhat of a hybrid between the flux-tunable gates and microwave-activated gates discussed in the previous chapter. Like the flux-tunable gates, these involve parametrized terms that already exist in the system Hamiltonian, typically frequencies (ω) or coupling rates (g), rather than adding additional terms with an external classical drive. However, they work by modulating the system parameters at microwave frequencies, rather than slowly tuning them at near DC.

Frequency Modulation

Given a flux-tunable qubit, we can ask what happens if we modulate the qubit frequency $\omega_q(t)$ at a drive frequency ω_d so that

$$\omega_q(t) = \omega_0 + \varepsilon_d \sin(\omega_d t + \phi_d) \quad (5.1)$$

where ω_0 is the average frequency and ε_d and ϕ_d are the drive amplitude and phase. The resulting time dependent Hamiltonian for the qubit becomes:

$$H(t) = \frac{\omega_q(t)}{2} \sigma_z = \underbrace{\frac{\omega_0}{2} \sigma_z}_{H_0} + \underbrace{\frac{\varepsilon_d}{2} \cos(\omega_d t + \phi_d) \sigma_z}_{H_d} \quad (5.2)$$

We can then express the time-dependent propagator as

$$\begin{aligned} U(t) &= e^{-i \int_0^t H(t') dt'} = e^{-i H_0 t} e^{i \frac{\varepsilon_d}{2\omega_d} \sin(\omega_d t + \phi_d) \sigma_z} \\ &= e^{-i \frac{1}{2} \omega_0 \sigma_z} \sum_{n=-\infty}^{\infty} J_n \left(\frac{\varepsilon_d}{\omega_d} \right) e^{i \frac{1}{2} n \omega_d t \sigma_z} \end{aligned} \quad (5.3)$$

where we have used the Jacobi-Anger identity to rewrite the propagator as an infinite sum, $J_n(z)$ is the n -th Bessel function of the first kind [119]. Here, we find that the qubit is effectively “split” across an infinite but discrete spectrum of sidebands spaced ω_d apart, where the weight in each sideband is given by $J_n \left(\frac{\varepsilon_d}{\omega_d} \right)$.

Coupling Modulation

Now suppose we could directly modulate the coupling between two qubits so that

$$g_{\Delta}(t) = g_{\Delta,0} + A_{\Delta} \cos(\omega_d t + \phi_d) \quad (5.4)$$

$$g_{\Sigma}(t) = g_{\Sigma,0} + A_{\Sigma} \cos(\omega_d t + \phi_d) \quad (5.5)$$

where we assign an independent mean coupling rate and modulation amplitude, but keep the same frequency and phase, for the two coupling rates. This is because a modulation of the underlying coupling mechanism typically will affect both terms. Then, the effective Hamiltonian for the system becomes

$$H = -\frac{1}{2} \omega_a \sigma_z^{(a)} - \frac{1}{2} \omega_b \sigma_z^{(b)} - g_{\Delta}(t) \left[\sigma_+^{(a)} \sigma_-^{(b)} + \sigma_-^{(a)} \sigma_+^{(b)} \right] - g_{\Sigma}(t) \left[\sigma_+^{(a)} \sigma_+^{(b)} + \sigma_-^{(a)} \sigma_-^{(b)} \right] \quad (5.6)$$

If we make the standard unitary transformation to obtain the effective Hamiltonian in the system frame, we obtain

$$H' = -g_{\Delta}(t) \left[\sigma_+^{(a)} \sigma_-^{(b)} e^{i\Delta t} + \sigma_-^{(a)} \sigma_+^{(b)} e^{-i\Delta t} \right] - g_{\Sigma}(t) \left[\sigma_+^{(a)} \sigma_+^{(b)} e^{i\Sigma t} + \sigma_-^{(a)} \sigma_-^{(b)} e^{-i\Sigma t} \right] \quad (5.7)$$

where $\Delta = \omega_a - \omega_b$ and $\Sigma = \omega_a + \omega_b$, and we notice that this looks exactly like Eq. 4.43 except for the time dependent coupling rates. Here, we see that if the time-dependence of the coupling term matches one of the frequencies $\pm\Delta$ or $\pm\Sigma$, we obtain a resonant coupling term even if Δ (Σ) is non-zero. In particular, we see that if $\omega_d = \Delta$, the Hamiltonian becomes

$$\begin{aligned} H' &= -\frac{1}{2} A_{\Delta} e^{-i\phi_d} \sigma_+^{(a)} \sigma_-^{(b)} - \frac{1}{2} A_{\Delta} e^{i\phi_d} \sigma_-^{(a)} \sigma_+^{(b)} e^{2i\Delta t} - g_{\Delta,0} \sigma_+^{(a)} \sigma_-^{(b)} e^{i\Delta t} \\ &\quad - \frac{1}{2} A_{\Sigma} e^{-i\phi_d} \sigma_+^{(a)} \sigma_+^{(b)} e^{i(\Sigma-\Delta)t} - \frac{1}{2} A_{\Sigma} e^{i\phi_d} \sigma_+^{(a)} \sigma_+^{(b)} e^{i(\Sigma+\Delta)t} - g_{\Sigma,0} \sigma_+^{(a)} \sigma_+^{(b)} e^{i\Sigma t} + \text{c.c.} \end{aligned} \quad (5.8)$$

where we see that only the $\sigma_+^{(a)} \sigma_-^{(b)}$ and $\sigma_-^{(a)} \sigma_+^{(b)}$ terms are time independent. Assuming $A_{\Delta}, g_{0,\Delta} \ll \Delta$ and $A_{\Sigma}, g_{0,\Sigma} \ll \Sigma$, we can drop the remaining oscillating terms to obtain an effective time-independent Hamiltonian

$$H' \approx -\frac{1}{2} A_{\Delta} \left[e^{-i\phi_d} |10\rangle\langle 01| + e^{i\phi_d} |01\rangle\langle 10| \right] \quad (5.9)$$

that looks like the resonant coupling Hamiltonian in Eq. 4.44, but with a phase ϕ_d that is determined by the drive. As a result, we find that not only do we get a rotation in the $|01\rangle$ – $|10\rangle$ subspace, but also a phase that allows us to control the axis of rotation. This is a crucial insight — we gain control over the phase ϕ_d of the coupling term because the interaction is occurring in the rotating frame, in contrast to the lab frame interaction that occurs when we bring two qubits directly into resonance. This additional control knob allows us to natively integrate virtual- Z rotations within entangling operations [1] and break time-reversal symmetry to engineer synthetic magnetic fields [172, 170] and non-reciprocal coupling [112].

The interaction Hamiltonian in Eq. 5.9 gives rise to a family of iSWAP-like gates. However, we could also choose to let $\omega_d = \Sigma$, in which case we obtain the two-photon coupling terms:

$$H' \approx -\frac{1}{2}A_\Sigma \left[e^{-i\phi_d} |11\rangle\langle 00| + e^{i\phi_d} |00\rangle\langle 11| \right] \quad (5.10)$$

In the case of two coupled transmons, we could also set the modulation frequency to $\Delta + \alpha_a$ or $\Delta - \alpha_b$ to engineer interactions within the two-photon subspace, and more generally, we can bring any two levels into effective resonance as long as there is a non-zero matrix element connecting them. This highlights the relative versatility of parametric coupling, compared to resonant coupling in the laboratory frame, which is unable to generate a swapping interaction that is not excitation preserving.

Finally, we note that when modulating the coupling, the modulation frequency ω_d does not enter directly into the effective coupling strength $\frac{1}{2}A_d$. In contrast, the effective coupling is suppressed by $J_n\left(\frac{\varepsilon_d}{\omega_d}\right)$ in the case of direct frequency modulation of the qubit. This allows us to maintain large detunings between qubits while still allowing for high interaction rates, which significantly relaxes constraints on frequency targeting and helps to suppress contextual errors. However, for a tunable coupler with the finite coupling range $(g_{\min}, g_{\max})$, parametric coupling will still yield interaction rates that are approximately $4\times$ smaller than direct resonance. In particular, we lose a factor of two because the peak-to-peak coupling amplitude $g_{\max} - g_{\min} = 2A_d$ corresponds to a modulation amplitude $A_d = (g_{\max} - g_{\min})/2$, and another factor of two from dropping the counter-rotating terms which do not participate meaningfully in the interaction. However, despite this reduction in gate speed, parametric coupling modulation retains other advantages over resonant coupling that make it a worthwhile pursuit. While this approach to two-qubit gates has been demonstrated before [130, 68, 90], most prior implementations of parametric gates have been slow, with gate times > 100 ns, or use a coupler design that does not easily scale to a 2D lattice of qubits. In the following sections, we will discuss the design and implementation of a multi-transmon device that addresses both of these shortcomings to demonstrate high-fidelity parametric entangling gates.

5.2 Tunable Coupler Design

The device we designed for this implementation consists of 6 fixed-frequency transmons in a line, with flux-tunable couplers between each pair of adjacent transmons. Here the qubits are floating, while the couplers are grounded SQUIDS (Fig. 5.1(a)). While the coupler design is very similar to tunable couplers used for resonant coupling [148, 210], parametric driving imposes more stringent constraints on the coupler parameters. In particular, we want the effective coupling $g(t)$, between




Figure 5.1: **Flux-tunable coupling between two transmon qubits.** (a) The flux-tunable coupling circuit used in this experiment. The coupler circuit was first proposed in [148], and consists of a grounded SQUID that can be used to change the effective impedance between the two qubits. The additional direct capacitance between the two qubits helps cancel the residual coupling between the qubits when the SQUID is near the $\Phi_{\text{ext}} = 0$ state. (b) Qubit-qubit static coupling between a representative pair of qubits on the device. The coupling g_{Δ} is inferred from measurements of the static ZZ -rate ζ between the two qubits.

the two transmons to be effectively a classical parameter. This means that the intermediate coupler mode must remain in its ground state at all times, and adiabatically follow the instantaneous ground state despite the fast modulation of the coupler frequency. This requires a large detuning $\Delta_i = \omega_i - \omega_c$ between each qubit and the coupler, so our goal was to minimize the parasitic self-capacitance of the coupler mode.

Fig. 5.1(b) shows the measured static coupling rate g_{Δ} for a representative pair, showing that the tunable coupler design is capable of reaching a minimum coupling $g_{\text{min}}/2\pi = 2$ MHz and a maximum coupling $g_{\text{max}}/2\pi = 100$ MHz. Since the qubits are detuned by a frequency $\Delta = 511$ MHz, we cannot measure the static coupling directly. Instead, we measure the ZZ rate ζ using the JAZZ (joint amplification of ZZ) protocol [69, 194], which uses a conditional Ramsey interferometry experiment to extract the ZZ rate while simultaneously canceling the local frequency drifts. We then compute g_{Δ} from Eq. 4.53 where we drop the g_{Σ} contribution.

Finally, we demonstrate the parametric interaction in the $|01\rangle$ – $|10\rangle$ subspace of the two transmons, by modulating the coupler at a frequency near $(\omega_0 - \omega_1)/2$ to activate a swapping interaction. The resulting chevron pattern is shown in Fig. 5.2(a)–(b). Here, we modulate at $\Delta/2$ rather than Δ , because $g_{\Delta}(\Phi_{\text{ext}})$ is an even function about $\Phi_{\text{ext}} = 0$. As a result, when we modulate $\Phi_{\text{ext}}(t)$ symmetrically about the zero flux point at a frequency ω_d , we introduce a frequency doubling such that $g_{\Delta}(\Phi_{\text{ext}}(t))$ is at a frequency $2\omega_d$ ¹. Alternatively, we could choose to modulate the coupler near the detuning frequency, by modulating $\Phi_{\text{ext}}(t)$ only on the positive or negative side. However, this would result in a DC offset in the voltage/current that increases the overall pulse bandwidth and leads to more pulse distortion.

¹This is equivalent to the trigonometric identity $\cos^2 \theta = \frac{1 + \cos 2\theta}{2}$




Figure 5.2: **Parametric coupling between two transmons.** (a),(b) A parametrically-activated interaction between $|10\rangle$ and $|01\rangle$ within the single-photon subspace for a representative pair of qubits on the device (Q3-Q4). The population in each state exhibits the same dynamics as the driven two-level system in Fig. 4.1. (c) The drive frequency $\omega_d = (\omega_{01} - \omega_{10})/2$ that results in maximum population contrast as a function of the drive amplitude. The frequency of the transition shifts down with larger amplitude, due to the DC flux offset, which leads to a lower average frequency of the coupler mode under modulation. This shifts the $|0\rangle$ - $|1\rangle$ transition of the higher frequency qubit more than the lower frequency qubit and reduces their relative detuning. (d) The parametric coupling rate $\tilde{g}$ as a function of the drive amplitude. This is approximately $4\times$ lower than the static coupling g , due to both the counter-rotating terms and the maximum modulation amplitude.




Figure 5.3: **Calibration of a parametric $\sqrt{i}$ SWAP gate.** (a) The pulse parametrization used to implement the gate. We use a flattop pulse envelope with cosine ramps. We fix the total gate time τ_g and the ramp time τ_r . The remaining parameters — amplitude A , frequency ω_d , and phases β , ζ_a , ζ_b — are calibrated to engineer the proper unitary. Pulse sequences for each calibration step are shown in (b)-(f). The parameters at the center denote the primary sweep parameters.

5.3 Gate Calibration

Having demonstrated the parametric interaction, we now want to calibrate a specific entangling gate instance. With only low frequency drives (< 1 GHz), our choices for a maximally entangling two-qubit gate include the CZ, iSWAP, and $\sqrt{i}$ SWAP gates. In order to implement the bSWAP-like family of gates, we would need a drive at the sum or average frequency² (Σ or $\Sigma/2$). For the following demonstration, we chose to implement the $\sqrt{i}$ SWAP gate. This has the advantage of being the fastest of the three, and it was also demonstrated recently that the $\sqrt{i}$ SWAP is more efficient than both the iSWAP and CZ gates for compiling arbitrary Haar-random unitaries [87]. In addition, the control we gain over the rotation phase in the single-photon subspace makes the parametric $\sqrt{i}$ SWAP fully compatible with virtual-Z gates, which eliminates one of the typical drawbacks of using this gate. Finally, although the $\sqrt{i}$ SWAP gate is not a Clifford gate and is thus not well suited for most stabilizer-based QEC circuits, it is straightforward to implement an iSWAP gate with just two repeated applications of the $\sqrt{i}$ SWAP.

Because the parametric $\sqrt{i}$ SWAP gate is essentially just a 90° rotation in the $|01\rangle$ – $|10\rangle$ subspace, the calibration protocol for the gate is effectively identical to the procedure for tuning up a single-qubit X90 gate. The pulse parametrization for the $\sqrt{i}$ SWAP gate is shown in Fig. 5.3(a). We choose a cosine ramp with a flat top to maximize the total interaction area, while still ensuring sufficient adiabaticity to minimize leakage to the $|20\rangle$ and $|02\rangle$ states. We chose the total gate time of $\tau_g = 30$ ns by first finding the largest parametric coupling strength, $\tilde{g}_{\max}$ that could be achieved across all pairs without inducing spurious transitions. We then choose $\tau_g \gtrsim 4/\tilde{g}_{\max}$ to account for the lower coupling during the ramp up and ramp down. Keeping the gate time τ_g fixed

²Flux drives at these frequencies are possible, however our filtering scheme included a 1 GHz low-pass filter.

two-qubit gate instances across all pairs simplifies the compilation for multi-qubit circuits with parallel two-qubit gates.

Next, we must calibrate the amplitude A and frequency ω_d for the gate. In the case of single-qubit gates, the transition frequency is only weakly dependent on the drive amplitude, so we choose to correct for detunings with just Z gates before and after the gate. However, as we saw in the previous section, the $|01\rangle$ – $|10\rangle$ transition frequency has a much stronger dependence on the drive amplitude. As a result, we must perform a joint optimization of the two parameters. To obtain an initial starting point, we prepare the system in the $|10\rangle$ state followed by two applications of the coupler pulse. We then perform a 2D sweep of the amplitude and frequency of the coupler pulse to find the set of parameters (A, ω_d) that maximizes the population transfer to $|01\rangle$. Once we have an initial set of parameters, we proceed to an error amplification step using floquet calibration.

Floquet Calibration

To amplify coherent errors and obtain a more accurate calibration, we use a variant of the floquet calibration in Ref. [147]. There are two insights that make this method effective. The first is that errors can be amplified with repeated applications of the gate. Second, oscillation frequencies are more robust to decoherence and we can accurately determine these frequencies with a sensitivity limited only by the total circuit depth.

To calibrate the amplitude, we prep the qubit in the $|01\rangle$ state and measure the resulting population in $|10\rangle$ after m applications of the coupler pulse. This results in a periodic signal as a function of m , whose frequency can be extracted easily with a FFT or by fitting to a sinusoidal function. Because this signal frequency depends linearly on the pulse amplitude (up to aliasing), we can easily determine the optimal pulse amplitude corresponding to a period $m = 4$ with an amplitude sweep consisting of only a few points. The linear dependence also enables more accurate extrapolation, so the calibration is relatively insensitive to the sweep range.

The frequency calibration is similar, and is based on a similar calibration protocol used in Ref. [35] to calibrate a single-qubit pulse detuning. Here, we repeat a circuit unit that consists of two back to back applications of the coupler pulse, with a 90° phase offset³ to echo the trajectory back to the original state. When there is no detuning, the population in the prep state $|01\rangle$ should remain constant as a function of m . However, any detuning between the pulse frequency and the transition frequency will lead to an oscillation in the $|01\rangle$ population, whose frequency increases linearly with the detuning. Once again, we can fit this signal frequency and fit to an absolute value function to determine the target transition frequency.

Finally, we iterate between the amplitude and frequency calibration protocols to account for their co-dependence. A typical calibration run requires only 2-3 iterations to converge provided that the initial starting parameters are within 10% of the correct values. Once these values have converged, we move on to calibrating the three phases β_0 , ζ_a , and ζ_b .

³The phase offset on the AWG pulse is also doubled by modulating symmetrically about $\Phi_{\text{ext}} = 0$, so a 90° phase offset is equivalent to reversing the rotation direction.

Two-Qubit Frame Tracking

We first make a quick digression to explain frame tracking with parametric modulation [1]. The ideal unitary resulting from the interaction Hamiltonian in Eq. 5.10 is given by

$$R_{XY}(\theta, \beta) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \frac{\theta}{2} & ie^{i\beta} \sin \frac{\theta}{2} & 0 \\ 0 & ie^{-i\beta} \sin \frac{\theta}{2} & \cos \frac{\theta}{2} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (5.11)$$

where we have replaced the phase ϕ_d in Eq. 5.10 with β , and the rotation angle is given by $\theta = A_{\Delta}\tau_g$. Given two qubits, we can show that individual Z-rotations on the qubits act on the two qubit Hilbert space like

$$R_z(\gamma_a) \otimes R_z(\gamma_b) = \begin{pmatrix} e^{-i(\gamma_a+\gamma_b)/2} & 0 & 0 & 0 \\ 0 & e^{-i(\gamma_a-\gamma_b)/2} & 0 & 0 \\ 0 & 0 & e^{i(\gamma_a-\gamma_b)/2} & 0 \\ 0 & 0 & 0 & e^{i(\gamma_a+\gamma_b)/2} \end{pmatrix} \quad (5.12)$$

and we see that this looks like a Z-rotation of $\gamma_1 - \gamma_2$ ($\gamma_1 + \gamma_2$) in the $|01\rangle$ - $|10\rangle$ ($|00\rangle$ - $|11\rangle$) subspace. Then, we can show that

$$R_{XY}(\theta, \beta) R_z(\gamma_a) \otimes R_z(\gamma_b) = R_z(\gamma_a) \otimes R_z(\gamma_b) R_{XY}(\theta, \beta + \gamma_a - \gamma_b) \quad (5.13)$$

and we see that we can again commute the Z gates across the coupler pulse by simply updating the drive phase. In addition, this tells us that the coupler phase β is not independent, but must be referenced to the initial phases ϕ_a, ϕ_b that define the reference frames for the two qubits involved.

Phase Calibration

The phases ζ_a and ζ_b are included primarily to compensate for the effective frequency shift of the two qubits during the coupler modulation. We can measure the resulting phase shift with a Ramsey interferometry-like experiment, by preparing the target qubit in a superposition $|X\rangle = R_x(\pi/2)|0\rangle = \frac{1}{\sqrt{2}}(|0\rangle + i|1\rangle)$, applying the coupler pulse, and measuring in the X-basis. We prepare $|X0\rangle$ to measure ζ_a and $|0X\rangle$ to measure ζ_b . By sweeping the relevant Z-correction, we can fit the resulting signal to a sinusoidal function that allows us to determine the correct phase.

Finally, we can measure β_0 , which is the relative phase difference between the reference frames for the two qubits. We do this by preparing the qubit in the $|0X\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |01\rangle)$ state followed by two applications of the coupler pulse to transfer any population in $|01\rangle$ fully to $|10\rangle$. The resulting state will be $\frac{1}{\sqrt{2}}(|00\rangle + e^{i\beta}|10\rangle)$, allowing us to measure phase β by measuring the first qubit in the X-basis.

5.4 Benchmarking and Error Budget

After calibrating the gate, we use a suite of benchmarking techniques to characterize the gate error. We will not go into the theory behind these techniques here, but the interested reader can

Parameter	θ	ϕ	β
Target	45°	0°	0°
Measured	46.07°	-0.83°	0.75°
Error	1.07°	-0.83°	0.75°

Table 5.1: **Coherent errors extracted from cross-entropy benchmarking.**

consult Ref. [81] for a comprehensive review. To characterize the overall gate error, we use both cycle benchmarking (CB) [57] and cross-entropy benchmarking (XEB) [22, 4, 12], which involve interleaving the gate of interest with cycles of random single qubit unitaries. In CB, these random unitaries are restricted to the Pauli group, while for XEB, the unitaries are sampled from a Haar random distribution. In either case, the random single-qubit unitaries effectively average the gate errors into a single-depolarizing channel, and result in an exponentially decaying signal that can be measured to extract an error-per-cycle. This is commonly referred to as a “dressed” error, that includes errors from the twirling gates. We can then compare this error-per cycle to a “reference” experiment where we replace the two-qubit gate of interest with the identity gate to extract the error of the twirling gates. Then, the interleaved gate error can be estimated as [81, 12]:

$$\epsilon_g = 1 - \frac{1 - \epsilon_d}{1 - \epsilon_r} \quad (5.14)$$

Fig. 5.4 shows the benchmarking results for a representative pair on the sample. From fitting the CB decay curves, we find a dressed Pauli error $\epsilon_d = 7.3(1) \times 10^{-3}$ and a reference Pauli error $\epsilon_r = 3.2(1) \times 10^{-3}$. This results in a Pauli error for the $\sqrt{\text{iSWAP}}$ gate of $\epsilon_g = 4.1(1) \times 10^{-3}$. With XEB, we find a total Pauli error of $\epsilon_{\text{XEB}} = 7.5(3) \times 10^{-3}$, which agrees with our result from CB. XEB also allows us to estimate the purity decay error [12], and we find an error-per-cycle of $\epsilon_{p,d} = 6.4(3) \times 10^{-3}$ for the dressed cycle and a reference purity decay error of $\epsilon_{p,r} = 2.6(1) \times 10^{-3}$ for the single qubit twirling gates. This results in an estimated purity error of $\epsilon_{p,g} = 3.7(3) \times 10^{-3}$ for the two-qubit gate alone.

The remaining error should be accounted for by coherent control errors, which we can also verify with our XEB data. Because the calculating the cross-entropy fidelity of the sampled bitstrings involves simulating the target unitary, we can parametrize our target unitary and optimize for the gate parameters that yield the highest XEB fidelity. Following Ref. [12], we use a Nelder-Mead optimizer and parametrize the rotation angles θ , β , and ϕ , where ϕ is the CPhase angle. The fit parameters are shown in Table 5.1, which yields an optimized XEB error of $\epsilon_{\text{XEB,opt}} = 6.7(7) \times 10^{-3}$. We can calculate the resulting coherent contribution to the Pauli error by computing the Pauli Transfer Matrix (PTM) representation of the fitted unitary to the target unitary:

$$\epsilon_{\text{coh}} = 1 - \frac{1}{d^2} \text{Tr} [\Lambda_G \Lambda_U^{-1}] \quad (5.15)$$

where Λ_G is the PTM for the fitted unitary [81]. Using the above equation, we find an error of $\epsilon_{\text{coh}} = 0.26 \times 10^{-3}$, which agrees within error bars to the purity contribution and the total gate error.




Figure 5.4: **Benchmarking a $\sqrt{i}$ SWAP gate.** (a) Cycle benchmarking results for the reference and dressed (interleaved $\sqrt{i}$ SWAP) cycles. The extracted Pauli error for the two qubit gate is estimated to be $\epsilon_g = 4.1(1) \times 10^{-3}$. (b) Cross-entropy benchmarking results using the $\sqrt{i}$ SWAP gate. The optimized fidelity was obtained by parameterizing the target unitary and choosing the gate parameters θ , β , and ϕ that yielded the highest XEB fidelity. (c) The total leakage population as a function of cycle depth during an XEB experiment. We find a total leakage error of $\epsilon_{l,g} = 1.2(3) \times 10^{-3}$. (d) The error contributions from decoherence, leakage, and unitary control error.

Finally, since leakage and decoherence both contribute to a loss of purity, we attempt to separate the two by employing a leakage benchmarking technique where we monitor the $|2\rangle$ state population in both qubits during a randomized benchmarking (RB) or XEB circuit [35, 209, 12]. The resulting decay curves can be fit to a rate equation [137]:

$$p_2 = C(1 - \Gamma)^m + p_\infty \quad (5.16)$$

where $\Gamma = \gamma_\uparrow + \gamma_\downarrow$ is the total leakage rate ($\gamma_\uparrow$ is the rate of leakage into $|2\rangle$ and $\gamma_\downarrow$ is the rate of *seepage* back into the computational subspace), and $p_\infty = \gamma_\uparrow/\Gamma$ is the steady-state $|2\rangle$ -state population. We find a total leakage error of $\epsilon_{l,d} = 2.9(4) \times 10^{-3}$ and a combined reference leakage error of $\epsilon_{l,r} = 1.6(3) \times 10^{-3}$. This results in an estimated 2Q leakage contribution of $\epsilon_{l,g} = 1.2(3) \times 10^{-3}$.

The remaining purity error should be accounted for by decoherence. We can estimate the




Figure 5.5: **Driven decoherence with parametric driving.** (a) The effective T_1 decay of the photon in the single-photon subspace as a function of parametric drive amplitude. The dashed horizontal lines indicate the idle T_1 times for Q3 (blue) and Q4 (yellow). We see a lower driven T_1 time at low amplitudes that trends down sharply as we approach the $\Phi_{\text{ext}} = \Phi_0/2$ point, which we hypothesize is due to the stronger participation of the qubit mode in the coupler SQUID. The smaller fluctuations in T_1 are likely due to changes in the TLS distribution at the different amplitude-dependent average frequencies. (b) A T_1 trace in the single-photon subspace at $A_d \approx 0.56$ (AWG units), corresponding to the amplitude used for the $\sqrt{\text{iSWAP}}$ gate. We fit the total decay of the $P_{|01\rangle} + P_{|10\rangle}$ signal to obtain an exponential decay in the single-photon subspace of $T_1 = 22.7 \mu\text{s}$.

expected decoherence error from the T_1 and $T_{2,\text{echo}}$ times for each qubit as

$$\epsilon_{\text{decoh}} = 1 - \frac{1}{d^2} \prod_{q=a,b} \left[2 + 2\sqrt{1 - \gamma_1^{(q)}} \sqrt{1 - \gamma_2^{(q)}} - \gamma_1^{(q)} \right] \quad (5.17)$$

where d is again the Hilbert space dimension, $\gamma_1 = 1 - e^{-\Gamma_1 \tau_g}$ and $\gamma_2 = 1 - e^{-2\Gamma_2 \tau_g}$ are the estimated relaxation and dephasing contributions for a gate time τ_g , and $\Gamma_1 = 1/T_1$ and $\Gamma_2 = 1/T_{2,\text{echo}} - 1/2T_1$ are the relaxation and pure dephasing rates respectively [96, 44]. Using our estimated idle coherence times from Table 5.3, we estimate a decoherence contribution of only $\epsilon_{\text{decoh}} = 1.3 \times 10^{-3}$, which does not seem to account for the remaining purity decay error. Instead, we must use the relevant coherence times under the action of the parametric drive. Fig. 5.5 shows an effective T_1 measurement where we measure the decay of the total population in the single-photon subspace under a parametric drive with the same amplitude as the $\sqrt{\text{iSWAP}}$, where we find an effective driven $T_1 = 22.7 \mu\text{s}$. In general, we find that the T_1 trends down as we increase the drive amplitude, likely to do greater hybridization between the qubit and bare coupler modes. We expect the coupler mode to have a lower T_1 , since the coupler SQUIDs have junctions that are $\sim 25\times$ larger than the qubit junctions, leading to a greater density of two-level systems (TLS). We can do a similar measurement for the driven T_2 , the primary difference being that we apply a 90° phase shift on the drive halfway through the idle time to reverse the trajectory and return

Pair	Q2–Q3	Q4–Q5
Isolated	$12.1(4) \times 10^{-3}$	$12.5(4) \times 10^{-3}$
Simultaneous	$15.3(4) \times 10^{-3}$	$12.3(4) \times 10^{-3}$

Table 5.2: **Simultaneous and isolated two-qubit gate errors.** Pauli errors for both isolated and simultaneous benchmarking of a $\sqrt{\text{iSWAP}}$ on qubit pairs Q2–Q3 and Q4–Q5.




Figure 5.6: **Cycle benchmarking fidelity over time.** We show the Pauli errors for $\sqrt{\text{iSWAP}}$ gate, measured via CB over a total span of ~ 38 hours. The gate was calibrated only once at the beginning of this time period. We see no significant degradation in gate fidelity, which suggests that the calibration parameters are stable. Small fluctuations in fidelity are likely due to transient TLS effects.

the photon to the initial starting state [130]. With this measurement, we find a driven T_2 time of $T_2 = 17.2 \mu\text{s}$, leading to an effective decoherence error $\epsilon_{\text{decoh}} = 2.4 \times 10^{-3}$ that agrees more closely with the expected value.

5.5 Scaling and Outlook

Having demonstrated a high-fidelity $\sqrt{\text{iSWAP}}$ gate between two detuned fixed-frequency qubits and a thorough understanding of the primary error mechanisms, we now turn to some scaling considerations. One such concern that was discussed earlier in Chapter 4 is crosstalk between the microwave-flux drives. In Table 5.2, we show the isolated and simultaneous CB decay errors for the $\sqrt{\text{iSWAP}}$ gate performed on Q2–Q3 and Q4–Q5. Interestingly, we find that the change in gate error is relatively small, especially compared to similar experiments with single qubit gates. This indicates that the microwave flux-crosstalk is much smaller than the microwave charge-crosstalk, likely due to the much smaller magnetic dipole moment of the coupler SQUID compared to the electric dipole moment of the much larger qubit pads.

We also monitor the two-qubit error rate over an extended period of approximately 38 hours to understand the calibration stability of the parametric gate. In Fig. 5.6, we show the dressed,

Qubit	$\omega_{01}/2\pi$ (GHz)	$\alpha/2\pi$ (GHz)	T_1 (μ s)	$T_{2,r}$ (μ s)	$T_{2,e}$ (μ s)
Q0	4.859	-0.186	37.2	9.9	24.3
Q1	5.370	-0.176	32.0	37.9	40.6
Q2	5.845	-0.175	32.1	6.7	31.0
Q3	5.396	-0.176	29.1	31.8	33.6
Q4	5.033	-0.176	31.2	39.8	42.6
Q5	5.635	-0.181	36.3	25.6	46.3

Table 5.3: **Qubit parameters for the 6-qubit parametric device.** The anharmonicities for Q0 and Q5 are lower due to extra ground capacitance on the boundary qubits. Coherence times were averaged over a period of 1 hour.

Pair	$\Delta/2\pi$ (MHz)	$\zeta/2\pi$ (kHz)	ϵ_g ($\times 10^{-3}$)
Q0–Q1	510	18	3.2
Q1–Q2	524	16	12.7
Q2–Q3	551	16	7.1
Q3–Q4	363	52	4.2
Q4–Q5	602	31	8.1

Table 5.4: **Coupler parameters for the 6-qubit parametric device.**

reference, and entangling for the $\sqrt{i}$ SWAP gate over this time period. Here, we see that the overall trend is flat, with fluctuations that are most likely due to coherence fluctuations from the spectral diffusion of TLS defects [100]. While this measurement demonstrates the stability of the system calibration for a single pair of qubits, the stability of the overall system across 6 qubits was less impressive. At any given moment in time, we typically saw a strongly coupled TLS affecting the relaxation rate of least one of the 6 qubits, making it difficult to simultaneously calibrate the entire system to a high fidelity. This directly motivated our subsequent studies, in Chapter 6, on mitigating TLS-induced dropouts to improve multi-qubit system stability.

Beyond overall system stability, we also see room for continued improvement on the two qubit gates themselves. From the error analysis, we know that the dominant error mechanism for the gate stems from a loss of purity. Leakage is responsible for approximately one-third of this error, and should be well suppressed by simple counterdiabatic pulse-shaping techniques like DRAG [196]. The remaining gate error is incoherent error, much of which is due to driven decoherence. We believe this is predominantly limited by the large $\sim 1 \mu\text{m} \times 1 \mu\text{m}$ SQUID junctions, which are likely to host more defects in both the oxide barrier [124] as well as the surfaces of the shadow evaporated aluminum [67, 116]. While the size disparity between the coupler junctions and the qubit junctions make it difficult to fabricate them in a single step, preliminary studies show that we can successfully realize a $3\times$ reduction in the coupler junction area fabricating them in a separate step with a $3\times$ increase in the critical current density ($\sim 2 \mu\text{A}/\mu\text{m}^2$). Furthermore, the coherence times observed

in this device are approximately an order of magnitude below the current state-of-the-art. As a result, we expect continued improvements in fabrication to bring these two-qubit gate errors into the $\mathcal{O}(10^{-4})$ regime.

Ultimately, we have demonstrated that parametrically modulated couplers are capable of implementing fast two-qubit gates between highly-detuned fixed-frequency transmons, and we do so using a design that can easily scale to a large, multi-qubit system. Compared to all-microwave architectures with static coupling, parametric gates enable significantly higher on-off ratios and allow neighboring qubits to be well-outside the straddling regime. They are also simpler to calibrate than baseband flux pulses, require no complex pulse pre-compensation filters [169], and allow for simple two-qubit frame tracking protocols. Finally, parametric coupling modulation removes the need for DC flux lines on each qubit, and new proposals for gradiometric flux-locking may also remove the need for local flux on the couplers [15], which opens the door to multiplexed all-microwave control in the future. Taken together, these features establish parametric coupling as a compelling gate architecture for building a scalable quantum processor.

Chapter 6

System Stability

State-of-the-art superconducting qubit systems currently contain up to a few hundred qubits, and although these system sizes are not yet large enough to reach 10^{-6} logical error rates, the calibration requirements for operating a system at this scale already poses a daunting challenge. In current transmon-based architectures, each qubit may have over 100 independent parameters that must be calibrated to ensure high-fidelity operation [4], and although many of these calibrations can be parallelized, it may take up to 30 minutes to fully calibrate a device at the scale required to implement a logical qubit [101]. Furthermore, complex experiments and calculations may take hundreds of hours of wall clock time to execute [106, 71], despite the fast clock speeds of superconducting qubits. Thus, it is critical to not only optimize the performance of individual qubit operations, but to also maintain such performance over a sufficiently long duration.

The two primary sources of uncontrolled system instability in superconducting qubit systems are quasiparticle bursts [200, 208, 128] and fluctuating two-level system (TLS) defects [100], and a significant amount of effort has been devoted to addressing both of these problems [83, 89, 30, 102]. Quasiparticle bursts can cause simultaneous decay events across a large number of qubits, breaking a fundamental assumption of quantum error correction that errors are uncorrelated. However, passive strategies such as gap-engineering have been successful at mitigating the primary detrimental effects of such quasiparticle bursts [127]. Resonant TLS defects similarly enhance T_1 decay on individual qubits, and although localized, these elevated decay rates can persist for minutes to days. These temporarily impaired qubits are an important class of qubit “dropouts”, which decrease the effective error-correcting code distance, despite improved protocols for handling unusable qubits or couplers [45].

In this chapter, we introduce in-situ electric-field tuning of TLS defects as an effective mitigation strategy for TLS-induced dropouts in superconducting qubit devices. We will start by giving a brief overview of TLS defects and their effect on superconducting qubit devices¹, followed by a survey of existing mitigation strategies for coping with TLS defects. We will then make a case for local-electric field tuning as the preferred method for addressing this problem. We show how clever design can enable the resource efficient integration of TLS control into the same control line used for single qubit gates. Using this control architecture, we demonstrate dynamic manipulation of the each qubit’s local TLS spectrum to engineer robustness against excess decoherence and reduce

¹We highly recommend Ref. [143], which provides a comprehensive review of the subject.




Figure 6.1: **Two-level systems in superconducting qubit devices.** (a) Illustrations of some proposed physical mechanisms for TLS in superconducting qubit devices, including (left to right): adsorbed molecules on the surface [110], defects in the silicon substrate [214], and defects in amorphous oxides [114, 9]. (b) The double well potential for a TLS according to the Standard Tunneling Model [157, 163, 143]. The potential is characterized by an asymmetry energy ε , and a tunneling energy Δ_0 that depends on the barrier height. (c) The dependence of the TLS frequency on the applied electric field, $\vec{E}$, and the TLS electric dipole moment, $\vec{p}$ is shown. The minimum frequency is set by the tunneling energy Δ_0 .

fluctuations in both qubit lifetimes and single qubit gate errors. Much of the work discussed in this chapter was first reported in Ref. [33] and was performed jointly with Kan-Heng Lee.

6.1 Two-Level System Defects

Two-level systems were first introduced as a way to explain the low-temperature behavior of amorphous glassy materials [157, 163]. Under the proposed *standard tunneling model* (STM), individual defects are modeled as an effective particle in a double-well potential that can occupy one of two low-lying energy states. The Hamiltonian for such a system is

$$H_{\text{TLS}} = \frac{1}{2} \begin{pmatrix} \varepsilon & \Delta_0 \\ \Delta_0 & -\varepsilon \end{pmatrix} = \frac{1}{2} (\varepsilon \sigma_z + \Delta_0 \sigma_x) \quad (6.1)$$

where ε characterizes the well asymmetry and $\Delta_0 \propto e^{-\lambda}$ is the tunneling rate between the two wells, which depends exponentially on a physical tunneling parameter. At sufficiently low temperatures

($k_B T \ll \Delta_0$), thermal activation is exponentially suppressed and quantum tunneling dominates in the behavior of this system. Given an ensemble of such defects, we expect ε and λ to be independent and uniformly distributed. This leads to a constant density of states with respect to the TLS energy and explains several key properties of amorphous materials at low-temperature.

The TLS potential is not static, but can be modified by environmental coupling to both a strain- and electric-field via the well asymmetry ε [143]:

$$\varepsilon = 2\gamma \cdot \mathbf{S} + 2\mathbf{p} \cdot \mathbf{E} + \varepsilon_0. \quad (6.2)$$

In the context of superconducting qubits, this electric dipole coupling mediates interactions with a qubit or resonator's electric field, while the phonon coupling describes the interaction between the TLS and a surrounding phonon bath, which is their primary source of dissipation. TLS's may also interact with one another, leading to the phenomenon of spectral diffusion [100]. The prevailing model for this process involves a dispersive coupling between the high-frequency TLS's with $\hbar\omega \gg k_B T$, and a bath of low-frequency, thermally-activated fluctuators ($\hbar\omega \ll k_B T$) [58, 144]. This leads to both a diffusive random walk as well as discrete telegraphic jumps in the TLS frequency over time. In either case, the unstable time dynamics of the TLS bath make the stable operation of a superconducting quantum processor a challenging task.

6.2 TLS Mitigation Strategies

Given how problematic TLS defects are for superconducting qubit systems, a variety of strategies have emerged in the effort to mitigate their detrimental effects. The first and arguably most difficult approach involves materials engineering to develop higher-quality fabrication processes. The amorphous oxides at the surfaces and interfaces of the superconducting electrodes host a large number of TLS defects that exchange energy with the qubit via an electric dipole-dipole coupling. Since this is the dominant source of dielectric loss for superconducting qubits [124], it is of the utmost importance to minimize unwanted oxides as much as possible at every step of the fabrication process. At protected interfaces between the substrate and superconducting metal, oxides can be minimized with proper substrate cleaning [6]. However, oxides on the exposed surfaces are harder to mitigate, since any exposure to ambient oxygen will cause immediate oxidation of the superconducting metal². This has motivated efforts to passivate or cap surfaces with lower-loss materials [9, 32] as well as hermetically seal device packaging [131]. Recent experiments have also pointed to defects in the bulk silicon substrate as a possible source of TLS defects [214], which suggests that higher purity silicon wafers may also be a necessary. However, fabrication improvements primarily aim to reduce the *average* defect density, and dropouts induced by strongly-coupled defects continue to affect superconducting qubits despite these improvements.

For a given fabrication stack and resulting TLS density, the next line of defense is to minimize the coupling strength between the TLS defects and the qubit with proper design. This strategy aims to reduce the electric field participation at material surfaces, a goal typically achieved by designing larger circuit components and minimizing their area-to-perimeter ratios [207, 203, 50]. Fabrication techniques such as etching deeper trenches into CPW gaps have also been shown to reduce the

²There have been efforts to use less reactive superconductors such as TiN or NbN, but these alternative materials have not yielded meaningfully better results.

effective TLS density [207, 28]. Finally, it is worth mentioning recent works that attempt to improve the coherence times of the TLS defects themselves, by fabricating qubit or resonator structures on a phononic bandgap structure [34, 154]. However, this increases the non-markovianity of the TLS bath, and can complicate noise mitigation techniques that typically rely on short correlation times of the environment.

The strategies discussed so far are primarily concerned with reducing *average* defect density and the likelihood of encountering strongly-coupled defects. But despite our best efforts, TLS defects will continue to be an unavoidable presence in the foreseeable future, so we must consider *in-situ* control strategies to actively suppress the impact of excess TLS noise during qubit operation. This is typically done in most state-of-the-art processors by carefully orchestrating the operating frequencies of individual qubits to avoid relaxation hot-spots caused by strongly-coupled TLS defects [102, 101, 217]. However, qubit operating frequencies must account for a variety of frequency-dependent error sources beyond the TLS spectrum. Crosstalk and stray coupling between neighboring qubits require their operating and idle frequencies to be sufficiently detuned. Two qubit gates often require the qubits involved to interact at the same frequency. Readout errors are often high-sensitive to the qubit-resonator detuning. All of these factors combine to make orchestrating qubit frequency trajectories a highly-constrained and non-local optimization problem. Furthermore, updates to the qubit frequency due to drifting TLS lead to additional calibration or modeling overhead to minimize coherent control errors, since qubit control parameters are also sensitive to the operating frequency [169, 88].

The success of qubit frequency orchestration begs the question of whether it is possible to instead tune the TLS frequency. Thankfully, the coupling of the TLS to an external strain- or electric-field, as seen in Eq. 6.2, provides a mechanism to do so. In fact, experiments dating back to 2012 have demonstrated such tuning of TLS defects, first with piezo-electric actuators to apply mechanical strain to the substrate [74, 117] and later with DC electrodes to apply a global electric field [115]. These experiments showed that the frequencies of individual TLS defects could be precisely controlled, and this technique was later used to optimize qubit T_1 times [114]. The biggest advantage with TLS frequency tuning compared to qubit frequency tuning is that optimal TLS frequencies are far less constrained from the perspective of device performance. The TLS environment for each qubit is physically localized to a small region where the electric field of the qubit mode is sufficiently strong to mediate coupling, so every qubit will be sensitive to a different set of defects. Furthermore, the precise frequencies of the strongly coupled defects are no longer important — we only need to detune them sufficiently³ from the relevant qubit transition frequencies. As a result, we expect TLS frequency tuning to be a simpler and more scalable approach to addressing the TLS problem.

Having recognized the potential benefits of TLS frequency tuning over qubit frequency orchestration, we now turn to its practical implementation in a multi-qubit device. First, there is the choice of either electric-field tuning or strain-tuning to as the control parameter for the TLS frequency. There is an argument to be made for strain tuning, as several experiments have shown that TLS residing in the junction tunnel barrier do not respond to a DC electric field [115, 18, 17]. While there are a number of TLS defects reported in these experiments that do not respond to an

³Sufficiently, here, means that $\Delta = \omega_q - \omega_{\text{TLS}} \gg g$, where g is the coupling strength between the qubit and the TLS. The effective interaction scales as $\frac{g^2}{\Delta}$.

applied electric field⁴, it is worth pointing out that the reported devices intentionally contain large μm^2 -scale parasitic junctions that are typically shorted out in real processors. In practice, most qubit tunnel junctions are $\sim 300\times$ smaller. This, coupled with estimates of the TLS density in the AlOx tunnel junction to be $\sim 0.3\text{-}2$ defects/ $\mu\text{m}^2/\text{GHz}$ [124, 11, 205], suggest that the number of TLS's that will be shielded from an electric field is relatively small for most current devices. Furthermore, recent work has suggested that a large number of the problematic defects are actually present at the surfaces of the shadow evaporated leads near the junction,

Electric field tuning is also much simpler to integrate *locally* into a superconducting quantum processor, which is an important requirement for scalability. The fact that each qubit couples to a separate and independent bath of TLS defects makes it unlikely that a single degree of freedom will be sufficient to maintain an optimal TLS spectrum for multiple qubits at once. For a large-scale processor with hundreds or thousands of qubits, it is almost certainly impossible. This rules out the large piezoelectric actuators used in prior experiments to strain the entire chip, and would likely require unconventional fab processes to integrate on-chip MEMs actuators for applying local strain. In contrast, a local electric-field bias can be implemented for each qubit in essentially the same way as the existing control lines used for single-qubit control. In the following section, we describe in detail the design of a multiplexed control line for simultaneous qubit and TLS control, which is the primary contribution of this chapter.

6.3 Implementation of Local Electric-Field Tuning

One of our main goals when designing an architecture with integrated local TLS control was to do so with no additional wiring overhead, unlike prior proposals for local electric TLS tuning [114]. Superconducting qubit systems already have a scaling problem when it comes to wire count, and introducing an extra wire for each qubit does not help the cause. The natural step then, is to think about whether we can use the existing qubit control lines, which typically deliver microwave signals to the qubit, to also apply a DC electric field bias to the qubit's TLS bath. There are two potential challenges to this approach. First, since a qubit control line necessarily couples an AC voltage across the junction, a DC voltage applied to the same line will induce a charge offset across the junction. However, this charge offset does not affect the qubit operation, because the transmon is charge-insensitive⁵. The second challenge is that qubit control lines are typically engineered to be weakly coupled to the qubit to limit radiative decay through the port, which makes it difficult to apply a sufficiently large electric field near the junction where the majority of TLS defects are present. To overcome this problem, we leverage symmetries in the circuit structure to decouple the electric field that tunes the TLS defects from the electric field that drives qubit transitions.

As mentioned in Section 2.2, the transmon can be constructed as either a *floating* or *grounded* oscillator. While these circuits are effectively equivalent in isolation, some important differences are highlighted when we couple them to external circuitry. For a floating transmon, the qubit is encoded in a differential mode, so it is sensitive only to a differential voltage bias and decoupled

⁴Below the critical current I_c , cooper pairs will tunnel across the junction to compensate for an applied voltage bias.

⁵This is not quite true, as we will see later. It is typically only the lowest transitions that are charge insensitive.




Figure 6.2: **Circuit model for the TIC-TAQ control line.** (a) An effective circuit diagram for the qubit/control-line system. The DC voltage bias for TLS control is combined with the RF signal for qubit XY control with an off-chip bias tee. (b) A circuit model for the reduced qubit-control line system showing the $Y-\Delta$ transform. When the capacitances to the environmental impedance are balanced ($\delta = 0$), there is no radiative loss through the control line.

from a common-mode bias. Meanwhile, the TLS defects have much dipole moments that are many orders of magnitude smaller and are scattered around the surfaces of the qubit electrodes. As a result, they will be sensitive to electric fields induced by both a common-mode and differential voltage bias. This allows us to engineer a very strong common-mode coupling the TLS defects on the qubit pads, without inducing excess radiative decay on the qubit, by designing a nearly-symmetric control line pad with similar capacitances to both qubit pads. Since this allows for the *targeted in-situ control of TLS and qubits*, we call this architecture TIC-TAQ for short.

Fig. 6.2(a) shows a lumped element circuit model of a floating qubit coupled to the control line (a semi-infinite CPW transmission line). The control line is modeled as an effective impedance to ground,

$$Z_{\text{env}} = \frac{1}{Z_0^{-1} + i\omega C_{g,c}} \quad (6.3)$$

where $Z_0 = 50 \, \Omega$ is the characteristic impedance of the CPW transmission line and $C_{g,c}$ is the extra capacitance to ground due to the coupling pad at the open end of the control line. We can make a $Y-\Delta$ transform with the control-line circuitry (teal) to yield an equivalent circuit model, shown in

Fig. 6.2(b), where

$$Z_{a,c} = \frac{1 + i\omega Z_{\text{env}}(C_{1,c} + C_{2,c})}{i\omega C_{2,c}} \quad (6.4a)$$

$$Z_{b,c} = \frac{1 + i\omega Z_{\text{env}}(C_{1,c} + C_{2,c})}{i\omega C_{1,c}} \quad (6.4b)$$

$$Z'_{\text{env}} = \frac{1 + i\omega Z_{\text{env}}(C_{1,c} + C_{2,c})}{-\omega^2 Z_{\text{env}} C_{1,c} C_{2,c}}. \quad (6.4c)$$

Now we can use the standard rules for simplifying parallel and series impedances to obtain the admittance $Y_{\text{eff}}(\omega)$ shunting the qubit capacitor and junction. This adds an effective capacitance so that the total qubit capacitance C_{Σ} becomes

$$C_{\Sigma} = C_q + \text{Im} [Y'_{\text{eff}}(\omega_q)] \quad (6.5)$$

as well as an effective resistance $R = 1/\text{Re}[Y_{\text{eff}}(\omega_q)]$, where ω_q is the qubit frequency. The analytic expression for C_{Σ} is given by

$$C_{\Sigma} = C_q + \frac{1}{2} \left(\frac{C_{1,c}C_{2,c}}{C_c} + \frac{C_{g,1}C_{g,2}}{C_g} \right) \left[\frac{1 + \beta \left(\alpha + \frac{1}{2}(1 + \alpha^2) \frac{1}{Z_0\omega} \frac{(C_{1,c} + C_{g,1})(C_{2,c} + C_{g,2})}{C_{1,c}C_{2,c}C_g + C_{g,1}C_{g,2}C_c} \right)}{1 + \beta^2} \right] \quad (6.6)$$

where $C_g = (C_{g,1} + C_{g,2})/2$ and $C_c = (C_{1,c} + C_{2,c})/2$ are the average ground and coupling capacitances, and we define the following dimensionless factors:

$$\alpha = C_{g,c}Z_0\omega \quad (6.7)$$

$$\beta = \alpha + \frac{1}{2}(1 + \alpha^2) \frac{1}{Z_0\omega} \frac{C_c + C_g}{C_c C_g} \quad (6.8)$$

The effective resistance is given by

$$R_{\text{eff}} = \frac{1}{\text{Re}[Y_{\text{eff}}(\omega)]} = Z_0 \left[\frac{1 + \beta^2}{\delta^2(1 + \alpha^2)} \right] \quad (6.9)$$

where we define the asymmetry factor δ as

$$\delta = \frac{C_{1,c}C_{g,2} - C_{2,c}C_{g,1}}{(C_{1,c} + C_{2,c})(C_{g,1} + C_{g,2})} \quad (6.10)$$

Finally, this leads to a qubit relaxation rate

$$\Gamma = \frac{1}{R_{\text{eff}}C_{\Sigma}} = \frac{\delta^2(1 + \alpha^2)}{Z_0C_{\Sigma}(1 + \beta^2)} \quad (6.11)$$

To simplify these expressions, we first note that when the asymmetry factor δ is small (i.e. $C_{g,1} \approx C_{g,2}$ and $C_{c,1} \approx C_{c,2}$), the scaling factor in Eq. 6.6 can be approximated as

$$\alpha + \frac{1}{2}(1 + \alpha^2) \frac{1}{Z_0\omega} \frac{(C_{1,c} + C_{g,1})(C_{2,c} + C_{g,2})}{C_{1,c}C_{2,c}C_g + C_{g,1}C_{g,2}C_c} \approx \alpha + \frac{1}{2}(1 + \alpha^2) \frac{1}{Z_0\omega} \frac{C_c + C_g}{C_c C_g} = \beta \quad (6.12)$$

and C_Σ can be simplified to

$$C_\Sigma \approx C_q + \frac{1}{2} \left(\frac{C_{1,c}C_{2,c}}{C_c} + \frac{C_{g,1}C_{g,2}}{C_g} \right), \quad (6.13)$$

which is the expression we would have found for C_Σ if the control line pad was just an isolated node ($Z_{\text{env}} \rightarrow \infty$). Next, we can rewrite α as $C_{g,c}Z_0\omega_q = \frac{C_{g,c}Z_0}{C_\Sigma Z_q}$. In the typical transmon regime, $Z_q \sim 10^3 \Omega$, and $C_{g,c}/C_\Sigma \sim 15\%$, so $\alpha < 1\%$ is a negligible contribution. Taking $\alpha \approx 0$, we can simplify β to

$$\beta \approx \frac{1}{Z_0\omega} \left(\frac{1}{C_c} + \frac{1}{C_g} \right) \quad (6.14)$$

and if $C_c \ll C_g$, we can further simplify this to $\beta \approx (Z_0\omega C_c)^{-1} = \frac{C_\Sigma Z_q}{C_c Z_0} \gg 1$. Finally, putting all this together, we find that

$$\Gamma \approx \frac{Z_0\omega_q^2(\delta C_c)^2}{C_\Sigma} \quad (6.15)$$

which looks like the expression we derived previously for a grounded qubit⁶ with a coupling capacitance $C_c \rightarrow \delta C_c$ that is rescaled by the asymmetry factor.

This expression confirms our intuition about the nearly-symmetric control line. In most cases, $C_{g,1} \approx C_{g,2}$, so $\delta C_c \approx \frac{1}{2}(C_{1,c} - C_{2,c})$, and the perfectly symmetric case $C_{c,1} = C_{c,2} \Rightarrow \delta = 0$ leads to no radiative decay, but also no coupling. More generally, we can view the circuit as a capacitive voltage divider that sets the voltages V_1 and V_2 on the qubit pads, relative to an applied voltage on the control line. Thus, even when $C_{g,1} \neq C_{g,2}$, we can compensate by adjusting $C_{1,c}$ and $C_{2,c}$ to symmetrize the circuit and make $V_1 = V_2$. In practice, this gives us enough degrees of freedom to make $C_{1,c}, C_{2,c}$ arbitrarily large while independently controlling $\delta \ll 1$ to reach the desired level of radiative decay and qubit-control line coupling.

Fig. 6.3(a) shows the simulated coupling quality factor for the control line design we use. At a pad distance of 15 μm , we are able to achieve coupling capacitances $C_c \sim 0.5$ fF, which is more than an order of magnitude larger than the typical coupling capacitance of ~ 25 aF for a grounded transmon, while maintaining a radiation-limited T_1 time of well above 1 ms [10]. We also simulate the electric field distribution on the surfaces of the qubit using Maxwell 3D, which shows that a 1V voltage bias applied to the control line can induce electric fields $|\mathbf{E}| > 130$ V/m even at the far edge of the qubit pad. This corresponds to a tuning sensitivity, $\gamma_E \sim 4$ MHz/V for a typical dipole moment $p_{\text{TLS}} \sim 3$ debye [176, 23].

6.4 TLS Spectroscopy

Before we continue, we will take a slight detour to explain in more detail how we characterize the TLS spectral landscape for each qubit. In the early days, TLS defects were identified primarily through avoided crossings in the qubit's frequency spectrum, which occur due to hybridization between the qubit and a TLS. However, identifying TLS through resonant enhancements of the qubit relaxation rate offers a more sensitive probe for observing TLS defects with weaker coupling

⁶See Eq. 3.8




Figure 6.3: **Control line parameters for TLS tuning.** (a) Simulated coupling quality factors for the control line design as a function of the distance between the control line pad and the qubit pads, and the pad asymmetry. The zero-coupling point (infinite coupling Q) occurs at a non-zero asymmetry due to the capacitive loading from the readout resonator pad (RR). This breaks the symmetry between the two effective ground capacitances. (b) A rendering of a single qubit on the experimental device showing the qubit capacitor pads (QB), the readout resonator coupling pad (RR), and the TIC-TAQ control line (CL). The simulated electric field distribution on the surface of the device is shown for a 1 V DC bias applied at the control line. The inset shows the dependence of the TLS frequency on the applied electric field, $\vec{E}$, and the TLS electric dipole moment, $\vec{p}$.

to the qubit [11]. This method, often referred to as *swap spectroscopy* or simply *TLS spectroscopy*, is essentially a qubit T_1 measurement performed over a range of frequencies. The qubit is prepared in the excited state, and the $|1\rangle$ -state population is measured as a function of the delay time, during which the qubit frequency is swept. It relies on the fact that when a qubit is resonant with the TLS, any excitations will swap into the TLS where they are quickly dissipated into a lossy phonon bath. As a result, the qubit relaxation rate Γ_1 gains a frequency difference [11]:

$$\Gamma_1 = \frac{2g^2\Gamma}{\Gamma^2 + \Delta^2} + \Gamma_{1,q} \quad (6.16)$$

where $\Gamma_{1,q}$ is the intrinsic qubit relaxation rate, $\Gamma = \Gamma_{2,\text{TLS}} + \Gamma_{2,q}$ is the sum of the qubit and TLS decoherence rates, g is the qubit-TLS coupling strength, and $\Delta = \omega_q - \omega_{\text{TLS}}$ is the qubit-TLS detuning. We see that this creates a Lorentzian peak in the frequency-dependent qubit relaxation rate $\Gamma_1(\omega)$, whose width scales with $\Gamma \approx \Gamma_{2,\text{TLS}}$ and height scales with g^2 . Thus, given $\Gamma_1(\omega)$, we can fit each peak to extract both the coupling rate and T_2 for each TLS.

If we only need to identify the TLS frequency, we can simplify the procedure by measuring the excited state population P_1 after only a single fixed-time delay, while sweeping the qubit frequency [31]. This speeds up the measurement by a factor n , where n is the number of time we would have used to measure the qubit's energy relaxation rate Γ_1 , at the cost of reducing the bandwidth of TLS coupling strengths that can be detected. Here, the time delay, τ , between the qubit excitation pulse and the measurement must be chosen carefully, since the resulting spectrum will be most sensitive to TLS defects with $2g^2/\Gamma \sim \tau^{-1}$ [93]. We must also be careful that any SPAM errors are




Figure 6.4: **Multi-level corrections to the AC Stark shift.** (a) A level diagram for the lowest three-levels of a transmon qubit. A drive on the transmon will always couple to every transition in the system, leading to corrections from the two-level system model for the AC Stark shift. (b) A comparison of the numerical (solid) and analytical (solid) calculations for the multi-level AC Stark shift as a function of drive frequency, on both the $|0\rangle-|1\rangle$ and $|1\rangle-|2\rangle$ transitions of the transmon. The two methods agree up to the two-photon corrections, that are ignored by the analytical model.

sufficiently frequency-independent, since we are now using the excited state population as a proxy for the qubit decay rate.

Multi-level Corrections to the AC Stark Shift

During the delay time, we must tune the qubit frequency in order to sample the spectral landscape, and there are generally two ways to do this. If the qubit is flux-tunable, like the split-junction transmon or fluxonium, this is straightforward, but one of the primary goals of TLS tuning is to alleviate the need for DC flux control of the qubit frequency. Instead, we rely on the AC stark effect to gain spectral control over a single-junction or “fixed-frequency” transmon [31]. To do so, we apply an off-resonant microwave drive to the qubit to adiabatically map the qubit onto a floquet qubit with a shifted frequency [150].

However, since the transmon is a multi-level system, the formulas we derived in Sec. 4.1 are only valid in the weak driving limit. To accurately model the frequency shift over a sufficiently wide bandwidth, we must now take into account the additional contributions from the higher-levels [177]. While the exact frequency shift in this case is typically numerically calculated using Floquet analysis [150], we can still construct an approximate analytical formula for each transition by summing the contributions from the adjacent two-level subspaces.

We recall from Sec. 4.1 that the Stark shift due to an off-resonant tone for each two-level subspace in isolation is given by

$$\delta_{n,n-1}^{(0)} = \tilde{\omega}_n - \omega_n = \text{sgn}(\Delta_n)\Lambda_n - \Delta_n \quad (6.17)$$

where $\Lambda_n = \sqrt{n\Omega^2 + \Delta_n^2}$ and the superscript reminds us that this is the lowest order approximation. Here, ω_n is the frequency of the $|n-1\rangle-|n\rangle$ transition, and $\Delta_n = \omega_n - \omega_d$ is the detuning between




Figure 6.5: **Experimental validation of the multi-level AC Stark shift model.** (a) The experimental pulse sequence used to measure the effective qubit frequency under an off-resonant drive. The Stark shift drive is sandwiched between a Ramsey interferometry measurement. (b)-(c) Experimental measurements of the AC Stark shift as a function of the drive amplitude at frequencies near the $|0\rangle$ - $|1\rangle$ (b) and $|1\rangle$ - $|2\rangle$ (c) transitions. The experimental data points are plotted as circles, and overlaid on top of both the analytical and numerical predictions.

the transition frequency and the drive frequency ω_d . Then, we can make the assumption that the off-resonant drive acts on each level independently, so we divide the Stark shift on each transition evenly between the two eigenstates involved. However, since the $|n\rangle$ state is involved in both the $|n-1\rangle$ - $|n\rangle$ and $|n\rangle$ - $|n+1\rangle$ transitions, we must add this contribution to the Stark shift for both transitions. Thus, the corrected Stark shift is given by

$$\delta_n \approx \delta_{n,n-1}^{(0)} - \frac{1}{2}\delta_{n-1,n-2}^{(0)} - \frac{1}{2}\delta_{n+1,n}^{(0)} \quad (6.18)$$

where the minus sign on the contributions from adjacent transitions indicates that they shift the levels in the opposite direction. Fig. 6.4(b) compares the numerically calculated Stark shift with the analytical formula in Eq. 6.18 for the $|0\rangle$ - $|1\rangle$ transition of a transmon as a function of the drive frequency. We see that the two are in good agreement except near $\omega_{02}/2$ and $\omega_{13}/2$, where multi-photon effects come into play. We also validate this experimentally in Fig. 6.5, where we show the measured Stark shift on both the $|0\rangle$ - $|1\rangle$ and $|1\rangle$ - $|2\rangle$ transitions as a function of the drive strength, plotted against both the numerical and analytical curves. While we see that this analytical method gives good agreement up to much stronger drive strengths, we can also recover the simpler expression given in Ref. [31] by taking the limit of small Ω for the lowest energy level of the transmon:

$$\delta_1 \approx \delta_{01}^{(0)} - \frac{1}{2}\delta_{12}^{(0)} \approx \frac{1}{2} \left(\frac{\Omega^2}{\Delta_1} - \frac{\Omega^2}{\Delta_2} \right) = \frac{1}{2} \frac{\alpha \Omega^2}{\Delta_1(\Delta_1 + \alpha)} \quad (6.19)$$

where $\alpha = \omega_{01} - \omega_{12}$ is the transmon anharmonicity.

Although this is a useful method to frequency tune a single-junction transmon, an often cited drawback is that the AC Stark shift only offers a practical tuning range of ~ 100 MHz, compared to the > 1 GHz tuning range with a flux-tunable transmon. For strong Stark drives, it also becomes harder to maintain adiabaticity when mapping between the bare qubit and Floquet qubit frames, and drive-induced heating can become an issue [150]. First, the accessible tuning range is more than sufficient for the purpose of mitigating dropouts, since we are primarily trying to characterize the TLS spectrum near the qubit transition. Second, we can use the higher levels of the transmon to probe lower and lower frequency ranges, since the transmon anharmonicity is negative⁷. Thus, we see that we can avoid many of the downsides of flux-tuning at the cost of some added complexity to the TLS spectroscopy protocol.

6.5 Electric-Field Tuning of TLS

We now discuss our experimental demonstrations of local electric-field tuning with the TIC-TAQ architecture. The experimental device, shown in Fig. 6.6(a), consists of 16 uncoupled single-junction transmons. This allows us to demonstrate the relevant physics of TLS tuning without the additional complexity of operating a fully-coupled device. The qubits and resonators are separated into two nominally identical groups (A and B), with each group sharing a multiplexed readout bus. The readout circuitry and Purcell filtering were described earlier in Sec. 3.3. Each qubit also contains an individual control line, as described in Sec. 6.3. The TLS control voltages were generated with a Yokogawa GS200 voltage source⁸, and since the control line ends in an open circuit, we were able to apply DC voltages as high as ± 30 V to the control line without any heating at the sample. However, in order to avoid heating elsewhere in the fridge, we bypass the cryo-attenuators by adding a Mini-Circuits ZX85-12G bias tee to each line at the MXC plate. This allows us to also filter the DC and AC lines separately, in addition to maintaining a sufficiently large load impedance for the DC voltage source so that there is no current draw.

We can now characterize the TLS spectrum for each qubit as a function of the applied voltage bias. Fig. 6.6(b) shows the transition frequencies of TLS defects coupled to an example transmon over a 30 V bias range, where we observe a range of tuning sensitivities γ_E between 9 MHz/V and 330 MHz/V. Here we measure P_1 at a fixed time delay of 20 μ s. The horizontal dashed line denotes the fundamental frequency, f_{01} , of the transmon qubit with no Stark shift, and the corresponding line cut at constant frequency is shown in Fig. 6.6(c). Here, we see dips wherever the voltage bias brings a TLS into resonance with the qubit. The vertical dashed lines mark two example voltage biases: $V_1 = -0.15$ V (yellow triangle) and $V_2 = 3.252$ V (dark blue triangle). The corresponding line cuts in Fig. 6.6(d) illustrate the drastic changes in the qubit's dissipative environment under the two bias conditions. We also notice an abrupt change in the TLS spectrum at $V_b = 9$ V, which we hypothesize is due to a charge scrambling event that occurred during the total acquisition time of ~ 11 hours [197].

⁷When using the higher levels of the transition, we must remember that the energy relaxation rate scales with n . However, the qubit-TLS coupling rate also scales $\sqrt{n}$, so we recover the same sensitivity.

⁸It is not clear that such a low noise voltage source is strictly necessary for the TLS control, although it certainly does not hurt.




Figure 6.6: **Characterization of a qubit and its TLS environment for a varying voltage bias.** (a) A schematic of the 16-qubit “Everglades” device used in the experiment. (b) T_1 Stark shift spectroscopy as a function of applied voltage. The qubit $|1\rangle$ -state population P_1 is measured at a fixed time of 20 μs . The horizontal dashed line denotes the fundamental frequency, f_{01} , of the transmon qubit with no Stark shift. The vertical dashed lines mark two example voltage biases: $V_1 = -0.15$ V (yellow triangle) and $V_2 = 3.525$ V (dark blue triangle). Each acquisition of P_1 as a function of qubit frequency takes approximately 1 minute for a total acquisition time of approximately 11 hours. At $V_b = 9$ V, we see an abrupt change in the TLS spectrum, which may be due to a charge scrambling event [197]. (c) Linecut of P_1 as a function of the bias voltage. At V_1 (yellow), we find a dip in P_1 corresponding to a TLS that is resonant with the qubit. In contrast, there is no resonant TLS at V_2 (dark blue). (d) Linecuts of P_1 as a function of frequency at V_1 (yellow) and V_2 (dark blue). (e) The qubit frequency, f_{01} , measured via ramsey interferometry as a function of the bias voltage. (f) Single qubit Clifford randomized benchmarking (RB) as a function of the bias voltage. (g) Qubit T_1 , $T_{2,\text{echo}}$, and RB error rates (ϵ_P) measured at $V_1 = -0.15$ V (yellow) and $V_2 = 3.525$ V (dark blue).

Next, we confirm that the target qubit is not directly affected by a DC voltage bias on its control line. As mentioned earlier, the voltage bias will also induce an offset charge, n_g , across the junction. However, we expect the exponentially suppressed charge dispersion of the transmon to ensure that the qubit frequency remains insensitive to the voltage bias [105]. To verify this, we interleave a Ramsey interferometry experiment to measure the qubit frequency at each voltage bias. In Fig. 6.6(e), we see that the qubit frequency remains largely unchanged at different voltage biases, apart from avoided level crossings caused by resonantly coupled TLS defects. These observed frequency shifts due to TLS represent an additional error mechanism for qubits in the presence, since any gates performed under these conditions will be off-resonant and pick up additional phase errors. However, this also provides an alternative method for extracting the qubit-TLS coupling rates from different defects. Finally, we also interleave single qubit Clifford randomized benchmarking (RB) [103, 81] measurements to verify that the qubit control is not adversely impacted by the applied DC voltage. In Fig. 6.6(f), we see that the RB error spikes at voltage biases where a TLS is resonant with the qubit, but otherwise does not show any coherent voltage dependence. Rather, the applied voltage bias primarily impacts the qubit by coupling to and modifying its dissipative environment.

This independent control of the qubit's noise environment with respect to its intrinsic properties allows for the unconstrained optimization of important performance metrics, including the qubit T_1 time, $T_{2,e}$ time, and single qubit gate fidelities. In Fig. 6.6(g), we compare these metrics at the two example voltage biases. At $V_1 = -0.15$ V, where the qubit is approximately resonant with a TLS, we see a 60% (35%) reduction in the qubit's T_1 ($T_{2,e}$) time compared to $V_2 = 3.525$ V, where there is no resonant TLS. We also find a 35-fold increase in the single qubit RB error when the TLS is on-resonance, from both the reduced coherence times and coherent errors caused by the TLS-induced frequency shift. This highlights not only the mechanism by which TLS-induced errors can cause dropout qubit, but also demonstrates how we can mitigate such errors with precise control over the TLS defects.

6.6 Stabilizing the TLS Environment

Next, we use TIC-TAQ to address the instability of the TLS background over time, which is the primary driver of performance instabilities for superconducting qubit systems [100, 26, 197]. As mentioned earlier, the TLS landscape for given qubit can change drastically on a timescale of several hours, due to either spectral diffusion [100] or charge scrambling events [197]. These changes can bring a strongly coupled TLS into resonance with a qubit for a prolonged period of time, which renders the qubit effectively inoperable and leads to a temporary dropout. To mitigate such errors, we can periodically monitor the TLS environment for each qubit and apply a reactive voltage bias in response to changes in the spectrum. This was first demonstrated in Ref. [114], where Lisenfeld *et. al.* optimize a global electric field bias to tune the dominant strongly-coupled defects away from the qubit frequency. We perform a similar two-step optimization procedure which we explain below.

The general procedure for determining the optimal voltage biases, V_i^{opt} , is outlined in Fig. 6.7(a). For each time step, t_i , we first perform a coarse voltage scan (-10 V to 10 V) where we do an TLS spectroscopy measurement to characterize the TLS landscape over a small frequency window ($f_{01} \pm 5$ MHz). From these measurements, we obtain a $|1\rangle$ -state population $P_{1,i}(f, V)$ after




Figure 6.7: **Stabilizing qubit performance with TIC-TAQ.** (a) A diagram illustrating the optimization procedure applied on qubit B1. (b) The optimal voltages V_i^{opt} (blue) as determined by the optimization algorithm as a function of time over a period of ~ 38 hours (top panel). The static baseline voltage, $V_b = 0$, (orange) is shown as a reference. The RB error evaluated at V_i^{opt} (blue) and $V_b = 0$ (orange) are shown in the bottom panel. (c) Cumulative distributions of 500 repeated T_1 measurements taken over a 22 hour time period. The gray traces show the distributions measured at 51 static voltages, while the dark blue trace shows the distribution measured at the optimal voltage V_i^{opt} for each iteration. The shaded gray region spans the minimum and maximum mean T_1 time for all static voltage configurations, and the vertical blue line denotes the mean T_1 time for the dynamic voltage configuration. The cumulative distributions of the (d) means (μ) and (e) standard deviations (σ) for the static voltage biases are also plotted. The vertical line denotes the mean (standard deviation) for the dynamic voltage configuration.

a fixed time delay, which we use as a proxy for the frequency- and voltage-dependent T_1 . We then average the traces over the sampled frequencies to compute an averaged signal $\bar{P}_{1,i}(V)$, which effectively smooths the signal over voltage⁹. Next, we apply a weighted average over the $\bar{P}_{1,i}(V)$ from prior iterations to compute the loss function as

$$\mathcal{P}_{1,i}(V) = \frac{1}{N} \sum_{j=0}^i W(t_i, t_j) \bar{P}_{1,j}(V), \quad (6.20)$$

where the normalization factor N is given by $N = \sum_{j=0}^i W(t_i, t_j)$. This reduces statistical noise and accounts for historical context to bias the loss function toward more stable voltage regions. We then compute V_i^* , for each iteration as

$$V_i^* = \operatorname{argmax}_V \mathcal{P}_{1,i}(V) \quad (6.21)$$

from the coarse-grained loss function to determine an initial estimate for the optimal voltage. Finally, we refine the initial estimate V_i^* by measuring the qubit T_1 time in a small window around V_i^* and selecting the voltage that yields the highest T_1 as V_i^{opt} .

In Fig. 6.7(b), we show the optimal voltages V_i^{opt} determined via this optimization procedure as a function of time over a period of ~ 38 hours. For our choice of hyperparameters, the process of determining V_i^{opt} takes approximately 200 seconds per iteration. Between these iterations, we evaluate the single qubit gate errors at both V_i^{opt} and a baseline voltage, $V_b = 0$ V. Over the full measurement period, we measure a mean gate error of 1.8×10^{-3} for the baseline configuration, which is $2.6\times$ larger than the mean gate error of 6.8×10^{-4} that we find for the optimized voltages. Moreover, at $V_b = 0$, we see many time periods with elevated single qubit gate errors, which we attribute to the presence of a TLS that intermittently moves into resonance with the qubit frequency. In contrast, the single qubit gate errors are relatively stable for the optimized voltages.

To better quantify these dropout events, we define an outlier threshold, $\varepsilon_{P,\text{outlier}}$, based on the distribution of gate errors for the optimized configuration,

$$\varepsilon_{P,\text{outlier}} = Q_3 + 1.5 \times (Q_3 - Q_1), \quad (6.22)$$

where Q_1 (Q_3) is the first (third) quartile and $Q_3 - Q_1$ is the interquartile range (IQR). For the dataset shown in Fig. 6.7(c), we calculate an outlier threshold of $\varepsilon_{P,\text{outlier}} = 1.26 \times 10^{-3}$. Using this threshold, we detect an outlier 36% of the time for the baseline configuration, compared to only 1.3% of the time for the optimized configuration.

While the improvement with dynamic optimization is strikingly clear, one could argue that the performance improvement here was just due to the particular choice of the baseline voltage. To preempt this argument, we perform a separate experiment where the qubit T_1 time is measured at *many* static voltage biases (51 points between -10 V and 10 V) over a period of ~ 22 hours. This allows us to simulate the effect of many independent cooldowns, which have been shown to change the precise frequencies of TLS defects, but not their overall spectral densities [180, 41]. In Fig. 6.7(c), we show the resulting T_1 distributions for each of the static bias configurations in grey.

⁹Since the dependence of TLS frequency on voltage is locally linear, measuring $P_1(f_q + \Delta f, V)$ is equivalent to measuring $P_1(f_q, V + \Delta V)$, $f_q + \Delta$, where $\gamma_E = \Delta f / \Delta V$ is the tuning sensitivity of the TLS in question. This allows us to sample intermediate voltages more quickly.

The corresponding T_1 distribution for the optimized voltage biases is shown in dark blue. Here, we find that the optimized voltage biases resulted in a mean T_1 of 52.5 μs , which represents a 14% (38%) improvement over the best (worst) mean T_1 across all static biases for the same time period. Moreover, the standard deviation $\sigma_{\text{dynamic}} = 8.4 \mu\text{s}$ for the optimized voltages is smaller than most of the standard deviations σ_{static} for the static biases ($\sigma_{\text{static}} \in [7.7, 14.4] \mu\text{s}$). This highlights the improvement, in both the absolute coherence times and also their relative stability, enabled by dynamic TLS mitigation. However, this is to be expected, since a qubit is almost guaranteed to suffer from a resonant TLS defect given sufficient time.

6.7 Local Tuning vs Global Tuning

While we have already discussed the importance of *local* TLS control, we now present some data to validate our reasoning. To do this, we perform an experiment similar to the previous section, where we compare the T_1 distributions using both global and local TLS control on a set of 12 qubits. To simulate a single global control knob, we fix the voltage bias for every qubit to the same value and scan over many different voltage values. For the local configuration, we select the optimal voltage bias for each qubit, using the same optimization procedure as before. We continuously monitor each qubit for ~ 26 minutes, and select the V_i^{opt} from the final iteration to benchmark the local TLS control.

Fig. 6.8(b) shows the distributions of the average T_1 times (over ~ 105 minutes) for 12 qubits. We compare the distributions for 51 global bias values between -10 V and 10 V (grey) to the distribution of average T_1 times for the locally optimized voltage biases (dark blue). The local configuration had a mean T_1 of 44.3 μs across the 12 qubits, compared to a best (worst) mean T_1 of 41.6 μs (34.1 μs) for the global configurations, representing a 6.5 – 30% improvement in mean T_1 compared to all global configurations.

However, while we find that the average T_1 across the qubits is always better with local control, this does not always strictly improve the T_1 for each individual qubit. This can be explained in part by TLS drift post-optimization, but we found that it was mainly the result of some DC crosstalk between a subset of the qubit control lines. Crosstalk can be detrimental because it breaks the assumption that each control line independently controls only a single qubit’s TLS environment, and this is not accounted for in our parallelized optimization. We confirm this by measuring a DC crosstalk matrix, where we use sweep the voltage bias on one control line at a time while performing TLS spectroscopy on every qubit. Surprisingly, of the 132 potential qubit-control pairs that we measured, only 7 pairs exhibited detectable levels of crosstalk. Closer inspection reveals that this primarily occurs when a control line passes directly by a neighboring qubit. Since the control line is shielded above by airbridge crossovers, this suggests that the coupling is mediated primarily through the silicon substrate, and could be suppressed easily with the integration of superconducting through-silicon vias in future devices [82].

6.8 Multi-Qubit Stabilization

As a final benchmark of local and dynamic TLS control, we repeat the experiment discussed in Section 6.6 simultaneously on 6 qubits (selected to minimize crosstalk) for a total duration of ~ 40




Figure 6.8: **Benchmarking TIC-TAQ across a multi-qubit device.** (a) A schematic illustrating the locally optimized (left) and global (right) voltage bias configurations. The 4 hatched squares represent qubits that were not measured. (b) A comparison of the cumulative distributions of T_1 times across 12 qubits with local optimization (dark blue) compared to global optimization (gray). The yellow trace highlights the global configuration where $V_b = 0$ for all qubits. The dotted lines denote the means of the corresponding distributions, and the grey shaded region spans the minimum and maximum mean T_1 times for all global configurations. The inset shows the distribution of mean T_1 times across the chip, with the mean for the $V_b = 0\text{ V}$ (yellow) and locally optimized (dark blue) configurations marked by vertical lines. (c), (d) Repeated measurements of single qubit gate error (T_1) on qubit A5, measured at $V_b = 0\text{ V}$ (orange) and V_i^{opt} (blue). The black dashed line in (c) denotes the outlier threshold for this qubit, defined in Sec. 6.6. (e) The number of outlier qubits across the 6 qubit subset over 40 hours. (f), (g) The distributions of single qubit gate errors (T_1 times) for each qubit. The circles denote outliers in the distribution and the horizontal lines denote the mean RB error (T_1 times) across all 6 qubits.

hours. Between each iteration, we measure the qubit T_1 and perform isolated single qubit RB on each qubit at both the periodically optimized voltages $V_{i,q}^{\text{opt}}$ (where q represents the qubit index) and $V_b = 0$ V. Fig. 6.8(c)((d)) shows the time-series of RB errors (T_1 times) for a representative qubit (QA5). For the static configuration, $V_b = 0$ V (orange), we find a sudden drop in the qubit T_1 and a corresponding spike in the single qubit gate error beginning at approximately $t = 21$ hours. In contrast, the qubit performance remains stable and unaffected by changes in the TLS spectrum when the voltage bias is periodically optimized (blue).

For each qubit, we again calculate an outlier threshold¹⁰ to screen for TLS-induced outlier events. Fig. 6.8(e) shows the number of outlier qubits over time for both voltage configurations. Over a period of ~ 40 hours, we detect one or more outlier qubits only 8.2% of the time with periodic optimization, compared to 67% of the time at a static voltage. Moreover, we detect no iterations with two or more simultaneous outliers with voltage optimization, compared to 18% of the time at $V_b = 0$. Finally, we compare the distributions of both average gate error (Fig. 6.8(f)) and average T_1 (Fig. 6.8(g)) for each of the 6 qubits over the entire experiment. Here, we see a 36% improvement in the average gate error with local and dynamic optimization ($\bar{\epsilon}_P = 7.6 \times 10^{-4}$) compared to the static voltage $V_b = 0$ ($\bar{\epsilon}_P = 11.9 \times 10^{-4}$). Similarly, we see a 17% improvement in the average T_1 time with periodic local optimization ($\bar{T}_1 = 43.3 \mu\text{s}$) compared to the static baseline voltage ($\bar{T}_1 = 36.9 \mu\text{s}$).

These results demonstrate the advantage of using TIC-TAQ to locally optimize the noise environment across a multi-qubit device and confirm that the spectral diffusion of TLS defects is a dominant contributor to fluctuations in qubit performance. Notably, there was no recalibration of single qubit gate parameters over the entire 40 hour time period, highlighting the qubit's insensitivity to the TLS control parameter.

6.9 Outlook

In this chapter, we have introduced a scalable quantum control architecture, TIC-TAQ, to address the threat of TLS defects in a multi-qubit device. Our method enables local and independent control over each qubit's dissipative environment with no additional on-chip overhead, making it simple to integrate into current superconducting quantum processor designs. Leveraging in-situ spectral tuning of TLSs, we demonstrate improvements to qubit coherence times and single qubit gate performance that remain robust to changes in the TLS landscape. We also demonstrate the scalability of our method, showing that we can simultaneously stabilize single-qubit performance across 6 qubits over a 40 hour time period, with an overhead that does not scale with device size. This work represents a substantial step towards resolving spatial and temporal performance instabilities in superconducting quantum processors, and also opens the door to several interesting avenues of research.

While our approach to performance stabilization relied on periodic optimization, several recent works have also demonstrated error stabilization by slowly modulating a TLS control field such that the qubit sees a different instantaneous noise environment on a *shot-by-shot* basis over an extended measurement [97, 43]. This has the effect of averaging different TLS configurations and the resulting noise, trading some excess relaxation for a more stable noise environment. We believe that this

¹⁰See Eq. 6.22




Figure 6.9: **Offset-charge dependence of MIST.** (a) An FFT of Ramsey interferometry measurements, on the $|1\rangle$ – $|2\rangle$ transition of a transmon qubit, as a function of the applied bias voltage. The two sinusoidal signals correspond to the even and odd charge parity states. We fit the periodic signal to extract a gate charge period of 13.9 mV. (b) Offset charge dependence of measurement-induced state transitions. We prepare the qubit in the $|0\rangle$ or $|1\rangle$ state, and sweep the power of a 1 μ s readout tone to bring the readout resonator to a steady state photon number $\bar{n}$. The measured probability of the qubit remaining in the same state is plotted as a function of both $\bar{n}$ and the offset charge.

can be a promising approach to complement the periodic optimization strategy discussed above — characterization of the spectral landscape can help identify regions free from strongly-coupled defects, and local averaging can be leveraged to stabilize the weakly coupled background. Another approach is to simply keep up with the changing TLS background via real time characterization and calibration. Notably, a recent experiment demonstrated the real-time calibration of a quantum processor by training a reinforcement learning agent to directly leverage error detection events as a learning signal for updating control parameters in-situ [186]. It is easy to imagine integrating the TLS bias voltage as another control parameter in such a calibration framework, that would enable parameter updates without interrupting a long-running algorithm.

The TIC-TAQ architecture also enables control of the transmon’s offset charge, which has interesting implications for measurement-induced state transitions or transmon ionization, which have recently regained attention as a key limitation for high-fidelity dispersive readout [175, 182, 94, 55]. Although the transmon is typically described as offset-charge insensitive, this is only true for the lowest energy states. For a typical transmon with $E_J/E_C \sim 80 - 100$, we find MHz levels of charge dispersion by the time we reach the $|5\rangle$ – $|6\rangle$ transition. While this may not seem like a problem at first glance, it can lead to offset-charge dependent readout errors, because these unwanted state transitions are primarily mediated by the higher levels of the transmon [59]. In Fig. 6.9(a), we show a measurement of the offset-charge dispersion of the $|1\rangle - |2\rangle$ state as a function of the applied voltage bias, where we see a gate charge period ($2e$) of 13.9 mV. We find the same periodic dependence in Fig. 6.9(b), where we measure the state-transition probability as a function of mean photon number starting in both the $|0\rangle$ and $|1\rangle$ state. Notably, we see a small offset-charge period relative to the range of typical bias voltages we apply for TLS optimization, making it possible to independently optimize the TLS noise environment and stabilize the offset

charge to a value that allows for higher power readout.

A final direction of interest involves leveraging TLS control to better understand and characterize fabrication processes, in hopes of accelerating efforts to further reduce the overall defect density. Spectral diffusion of TLS defects make individual coherence measurements—at a single qubit frequency and single point in time—incredibly poor predictors of the long-term average coherence times of a superconducting qubit [31]. However, by measuring coherence times over a range of applied voltages, we can very quickly sample the average TLS background [43]. We expect this to accelerate research on materials and fabrication processes by enabling faster screening and process characterization of both test devices and full-scale processors. TLS control also allows us to determine the distribution of TLS dipole moments and tunneling energies Δ_0 in a given material, which can yield a more holistic picture of the dielectric loss in a particular material platform [115]. This can be further extended by placing multiple control electrodes around the qubit, allowing one to triangulate the precise location, and potentially the full vector dipole moment, of individual defects. This technique has already yielded evidence that the defect density is predominantly concentrated at the surfaces of the shadow-evaporated aluminum leads [116], making the case for lift-off free junction fabrication even stronger. We also envision future work combining local TLS control with fluxonium qubits to enable the characterization of TLS defects over a much wider bandwidth, from MHz to GHz frequencies.

These future directions highlight the dual capacity of local TLS control: it is at once a calibration tool for actively stabilizing the performance of a quantum processor, but also a diagnostic probe for better understanding the sources of decoherence. These two functions can work in tandem to continue advancing the platform of superconducting qubits.

Bibliography

- [1] Deanna M. Abrams et al. “Implementation of XY entangling gates with a single calibrated pulse”. In: *Nature Electronics* 3.12 (Dec. 2020), pp. 744–750. DOI: 10.1038/s41928-020-00498-1.
- [2] Rajeev Acharya et al. “Quantum error correction below the surface code threshold”. In: *Nature* 638.8052 (Feb. 2025), pp. 920–926. DOI: 10.1038/s41586-024-08449-y.
- [3] Dorit Aharonov and Michael Ben-Or. *Fault-Tolerant Quantum Computation With Constant Error Rate*. 1999.
- [4] Google Quantum AI. “Quantum supremacy using a programmable superconducting processor”. In: *Nature* 574.7779 (Oct. 2019), pp. 505–510. DOI: 10.1038/s41586-019-1666-5.
- [5] Victor V. Albert and Philippe Faist, eds. *The Error Correction Zoo*. 2025.
- [6] M. Virginia P. Altoé et al. “Localization and Mitigation of Loss in Niobium Superconducting Circuits”. In: *PRX Quantum* 3 (2 Apr. 2022), p. 020312. DOI: 10.1103/PRXQuantum.3.020312.
- [7] Vinay Ambegaokar and Alexis Baratoff. “Tunneling Between Superconductors”. In: *Phys. Rev. Lett.* 10 (11 June 1963), pp. 486–489. DOI: 10.1103/PhysRevLett.10.486.
- [8] Alexander Arsenovic et al. “scikit-rf: An Open Source Python Package for Microwave Network Creation, Analysis, and Calibration [Speaker’s Corner]”. In: *IEEE Microwave Magazine* 23.1 (2022), pp. 98–105. DOI: 10.1109/MMM.2021.3117139.
- [9] Mustafa Bal et al. “Systematic improvements in transmon qubit coherence enabled by niobium surface encapsulation”. In: *npj Quantum Information* 10.1 (Apr. 2024), p. 43. DOI: 10.1038/s41534-024-00840-x.
- [10] Joseph C. Bardin et al. “Quantum Computing: An Introduction for Microwave Engineers”. In: *IEEE Microwave Magazine* 21.8 (2020), pp. 24–44. DOI: 10.1109/MMM.2020.2993475.
- [11] R. Barends et al. “Coherent Josephson Qubit Suitable for Scalable Quantum Integrated Circuits”. In: *Phys. Rev. Lett.* 111 (8 Aug. 2013), p. 080502. DOI: 10.1103/PhysRevLett.111.080502.

- [12] R. Barends et al. “Diabatic Gates for Frequency-Tunable Superconducting Qubits”. In: *Phys. Rev. Lett.* 123 (21 Nov. 2019), p. 210501. DOI: 10.1103/PhysRevLett.123.210501.
- [13] R. Barends et al. “Superconducting quantum circuits at the surface code threshold for fault tolerance”. In: *Nature* 508.7497 (Apr. 2014), pp. 500–503. DOI: 10.1038/nature13171.
- [14] M. T. Bell et al. “Spectroscopic Evidence of the Aharonov-Casher Effect in a Cooper Pair Box”. In: *Phys. Rev. Lett.* 116 (10 Mar. 2016), p. 107002. DOI: 10.1103/PhysRevLett.116.107002.
- [15] Denis Bénâtre et al. “Simultaneous sweet-spot locking of gradiometric fluxonium qubits”. In: *Phys. Rev. Appl.* 24 (5 Nov. 2025), p. 054031. DOI: 10.1103/gjks-ctvm.
- [16] E. R. Beringer et al. *Principles of Microwave Circuits*. Ed. by C. G. Montgomery, R. H. Dicke, and E. M. Purcell. Vol. 8. MIT Radiation Laboratory Series. Reprint editions also exist, such as the 1964 Dover edition (ISBN 9780486614038) and the 1987 IET edition (ISBN 9780863411007). New York, NY: McGraw-Hill Book Company, Inc., 1948.
- [17] Alexander Bilmes et al. “Quantum sensors for microscopic tunneling systems”. In: *npj Quantum Information* 7.1 (Feb. 2021), p. 27. DOI: 10.1038/s41534-020-00359-x.
- [18] Alexander Bilmes et al. “Resolving the positions of defects in superconducting quantum bits”. In: *Scientific Reports* 10.1 (Feb. 2020), p. 3090. DOI: 10.1038/s41598-020-59749-y.
- [19] Matthew P. Bland et al. “Millisecond lifetimes and coherence times in 2D transmon qubits”. In: *Nature* 647.8089 (Nov. 2025), pp. 343–348. DOI: 10.1038/s41586-025-09687-4.
- [20] F. Bloch and A. Siegert. “Magnetic Resonance for Nonrotating Fields”. In: *Phys. Rev.* 57 (6 Mar. 1940), pp. 522–527. DOI: 10.1103/PhysRev.57.522.
- [21] M. S. Blok et al. “Quantum Information Scrambling on a Superconducting Qutrit Processor”. In: *Phys. Rev. X* 11 (2 Apr. 2021), p. 021010. DOI: 10.1103/PhysRevX.11.021010.
- [22] Sergio Boixo et al. “Characterizing quantum supremacy in near-term devices”. In: *Nature Physics* 14.6 (June 2018), pp. 595–600. DOI: 10.1038/s41567-018-0124-x.
- [23] Jan David Brehm et al. “Transmission-line resonators for the study of individual two-level tunneling systems”. In: *Applied Physics Letters* 111.11 (Sept. 2017), p. 112601. DOI: 10.1063/1.5001920.
- [24] Peter Brooks, Alexei Kitaev, and John Preskill. “Protected gates for superconducting qubits”. In: *Phys. Rev. A* 87 (5 May 2013), p. 052306. DOI: 10.1103/PhysRevA.87.052306.

- [25] Otto Brune. “Synthesis of a finite two-terminal network whose driving-point impedance is a prescribed function of frequency”. PhD thesis. Massachusetts Institute of Technology, 1931.
- [26] Jonathan J. Burnett et al. “Decoherence benchmarking of superconducting qubits”. In: *npj Quantum Information* 5.1 (June 2019), p. 54. DOI: 10.1038/s41534-019-0168-5.
- [27] A.O Caldeira and A.J Leggett. “Quantum tunnelling in a dissipative system”. In: *Annals of Physics* 149.2 (1983), pp. 374–456. DOI: [https://doi.org/10.1016/0003-4916\(83\)90202-6](https://doi.org/10.1016/0003-4916(83)90202-6).
- [28] G. Calusine et al. “Analysis and mitigation of interface losses in trenched superconducting coplanar waveguide resonators”. In: *Applied Physics Letters* 112.6 (Feb. 2018), p. 062601. DOI: 10.1063/1.5006888.
- [29] P. Campagne-Ibarcq et al. “Quantum error correction of a qubit encoded in grid states of an oscillator”. In: *Nature* 584.7821 (Aug. 2020), pp. 368–372. DOI: 10.1038/s41586-020-2603-3.
- [30] L. Cardani et al. “Reducing the impact of radioactivity on quantum circuits in a deep-underground facility”. In: *Nature Communications* 12.1 (May 2021), p. 2733. DOI: 10.1038/s41467-021-23032-z.
- [31] M. Carroll et al. “Dynamics of superconducting qubit relaxation times”. In: *npj Quantum Information* 8.1 (Nov. 2022), p. 132. DOI: 10.1038/s41534-022-00643-y.
- [32] Ray D. Chang et al. “Eliminating Surface Oxides of Superconducting Circuits with Noble Metal Encapsulation”. In: *Phys. Rev. Lett.* 134 (9 Mar. 2025), p. 097001. DOI: 10.1103/PhysRevLett.134.097001.
- [33] Larry Chen et al. *Scalable and Site-Specific Frequency Tuning of Two-Level System Defects in Superconducting Qubit Arrays*. 2025. arXiv: 2503.04702 [quant-ph].
- [34] Mo Chen et al. “Phonon engineering of atomic-scale defects in superconducting quantum circuits”. In: *Science Advances* 10.37 (2024), eado6240. DOI: 10.1126/sciadv.ad06240.
- [35] Zijun Chen et al. “Measuring and Suppressing Quantum State Leakage in a Superconducting Qubit”. In: *Phys. Rev. Lett.* 116 (2 Jan. 2016), p. 020501. DOI: 10.1103/PhysRevLett.116.020501.
- [36] Sai Pavan Chitta et al. “Computer-aided quantization and numerical analysis of superconducting circuits”. In: *New Journal of Physics* 24.10 (Nov. 2022), p. 103020. DOI: 10.1088/1367-2630/ac94f2.
- [37] Jerry M. Chow et al. “Simple All-Microwave Entangling Gate for Fixed-Frequency Superconducting Qubits”. In: *Phys. Rev. Lett.* 107 (8 Aug. 2011), p. 080502. DOI: 10.1103/PhysRevLett.107.080502.

- [38] L. Chua. “Memristor-The missing circuit element”. In: *IEEE Transactions on Circuit Theory* 18.5 (1971), pp. 507–519. DOI: 10.1109/TCT.1971.1083337.
- [39] Alessandro Ciani, David P. DiVincenzo, and Barbara M. Terhal. *Lecture Notes on Quantum Electrical Circuits*. TU Delft OPEN Books, July 2025. DOI: 10.59490/tb.85.
- [40] A. A. Clerk et al. “Introduction to quantum noise, measurement, and amplification”. In: *Rev. Mod. Phys.* 82 (2 Apr. 2010), pp. 1155–1208. DOI: 10.1103/RevModPhys.82.1155.
- [41] Dante Colao Zanuz et al. “Mitigating losses of superconducting qubits strongly coupled to defect modes”. In: *Phys. Rev. Appl.* 23 (4 Apr. 2025), p. 044054. DOI: 10.1103/PhysRevApplied.23.044054.
- [42] Gavin E. Crooks. *Gates, States, and Circuits: Notes on the circuit model of quantum computation*. 2019.
- [43] Andrew Dane et al. *Performance Stabilization of High-Coherence Superconducting Qubits*. 2025. arXiv: 2503.12514 [quant-ph].
- [44] Hillary Dawkins, Joel Wallman, and Joseph Emerson. “Combining T_1 and T_2 estimation with randomized benchmarking and bounding the diamond distance”. In: *Phys. Rev. A* 102 (2 Aug. 2020), p. 022220. DOI: 10.1103/PhysRevA.102.022220.
- [45] Dripto M. Debroy et al. *LUCI in the Surface Code with Dropouts*. 2024. arXiv: 2410.14891 [quant-ph].
- [46] David Deutsch. “Quantum theory, the Church–Turing principle and the universal quantum computer”. In: *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences* 400.1818 (July 1985), pp. 97–117. DOI: 10.1098/rspa.1985.0070.
- [47] M. H. Devoret, A. Wallraff, and J. M. Martinis. *Superconducting Qubits: A Short Review*. 2004. arXiv: cond-mat/0411174 [cond-mat.mes-hall].
- [48] Michel H. Devoret. “Does Brian Josephson’s Gauge-Invariant Phase Difference Live on a Line or a Circle?” In: *Journal of Superconductivity and Novel Magnetism* 34.6 (June 2021), pp. 1633–1642. DOI: 10.1007/s10948-020-05784-9.
- [49] Michel H. Devoret, John M. Martinis, and John Clarke. “Measurements of Macroscopic Quantum Tunneling out of the Zero-Voltage State of a Current-Biased Josephson Junction”. In: *Phys. Rev. Lett.* 55 (18 Oct. 1985), pp. 1908–1911. DOI: 10.1103/PhysRevLett.55.1908.
- [50] Oliver Dial et al. “Bulk and surface loss in superconducting transmon qubits”. In: *Superconductor Science and Technology* 29.4 (Mar. 2016), p. 044001. DOI: 10.1088/0953-2048/29/4/044001.

- [51] Nicolas Didier et al. “Analytical modeling of parametrically modulated transmon qubits”. In: *Phys. Rev. A* 97 (2 Feb. 2018), p. 022330. DOI: [10.1103/PhysRevA.97.022330](https://doi.org/10.1103/PhysRevA.97.022330).
- [52] David P. DiVincenzo. “The Physical Implementation of Quantum Computation”. In: *Fortschritte der Physik* 48.9-11 (2000), pp. 771–783. DOI: [https://doi.org/10.1002/1521-3978\(200009\)48:9/11<771::AID-PROP771>3.0.CO;2-E](https://doi.org/10.1002/1521-3978(200009)48:9/11<771::AID-PROP771>3.0.CO;2-E).
- [53] B Douçot and L B Ioffe. “Physical implementation of protected qubits”. In: *Reports on Progress in Physics* 75.7 (June 2012), p. 072001. DOI: [10.1088/0034-4885/75/7/072001](https://doi.org/10.1088/0034-4885/75/7/072001).
- [54] Benoit Douçot and Julien Vidal. “Pairing of Cooper Pairs in a Fully Frustrated Josephson-Junction Chain”. In: *Phys. Rev. Lett.* 88 (22 May 2002), p. 227005. DOI: [10.1103/PhysRevLett.88.227005](https://doi.org/10.1103/PhysRevLett.88.227005).
- [55] Marie Frédérique Dumas et al. “Measurement-Induced Transmon Ionization”. In: *Phys. Rev. X* 14 (4 Oct. 2024), p. 041023. DOI: [10.1103/PhysRevX.14.041023](https://doi.org/10.1103/PhysRevX.14.041023).
- [56] Alec Eickbusch et al. “Demonstration of dynamic surface codes”. In: *Nature Physics* (Oct. 2025). DOI: [10.1038/s41567-025-03070-w](https://doi.org/10.1038/s41567-025-03070-w).
- [57] Alexander Erhard et al. “Characterizing large-scale quantum computers via cycle benchmarking”. In: *Nature Communications* 10.1 (Nov. 2019), p. 5347. DOI: [10.1038/s41467-019-13068-7](https://doi.org/10.1038/s41467-019-13068-7).
- [58] Lara Faoro and Lev B. Ioffe. “Internal Loss of Superconducting Resonators Induced by Interacting Two-Level Systems”. In: *Phys. Rev. Lett.* 109 (15 Oct. 2012), p. 157005. DOI: [10.1103/PhysRevLett.109.157005](https://doi.org/10.1103/PhysRevLett.109.157005).
- [59] Mathieu Féchant et al. “Offset Charge Dependence of Measurement-Induced Transitions in Transmons”. In: *Phys. Rev. Lett.* 135 (18 Oct. 2025), p. 180603. DOI: [10.1103/PhysRevLett.135.180603](https://doi.org/10.1103/PhysRevLett.135.180603).
- [60] Richard P. Feynman. “Simulating physics with computers”. In: *International Journal of Theoretical Physics* 21.6 (June 1982), pp. 467–488. DOI: [10.1007/BF02650179](https://doi.org/10.1007/BF02650179).
- [61] Mark Field et al. “Modular superconducting-qubit architecture with a multichip tunable coupler”. In: *Phys. Rev. Appl.* 21 (5 May 2024), p. 054063. DOI: [10.1103/PhysRevApplied.21.054063](https://doi.org/10.1103/PhysRevApplied.21.054063).
- [62] Ronald M. Foster. “A Reactance Theorem”. In: *Bell System Technical Journal* 3.2 (1924), pp. 259–267. DOI: <https://doi.org/10.1002/j.1538-7305.1924.tb01358.x>.
- [63] Austin G. Fowler et al. “Surface codes: Towards practical large-scale quantum computation”. In: *Phys. Rev. A* 86 (3 Sept. 2012), p. 032324. DOI: [10.1103/PhysRevA.86.032324](https://doi.org/10.1103/PhysRevA.86.032324).
- [64] B Foxen et al. “Qubit compatible superconducting interconnects”. In: *Quantum Science and Technology* 3.1 (Nov. 2017), p. 014005. DOI: [10.1088/2058-9565/aa94fc](https://doi.org/10.1088/2058-9565/aa94fc).

- [65] B. Foxen et al. “Demonstrating a Continuous Set of Two-Qubit Gates for Near-Term Quantum Algorithms”. In: *Phys. Rev. Lett.* 125 (12 Sept. 2020), p. 120504. DOI: 10.1103/PhysRevLett.125.120504.
- [66] J. M. Gambetta et al. “Analytic control methods for high-fidelity unitary operations in a weakly nonlinear oscillator”. In: *Phys. Rev. A* 83 (1 Jan. 2011), p. 012308. DOI: 10.1103/PhysRevA.83.012308.
- [67] Suhas Ganjam et al. “Surpassing millisecond coherence in on chip superconducting quantum memories by optimizing materials and circuit design”. In: *Nature Communications* 15.1 (May 2024), p. 3687. DOI: 10.1038/s41467-024-47857-6.
- [68] M. Ganzhorn et al. “Benchmarking the noise sensitivity of different parametric two-qubit gates in a single superconducting quantum computing platform”. In: *Phys. Rev. Res.* 2 (3 Sept. 2020), p. 033447. DOI: 10.1103/PhysRevResearch.2.033447.
- [69] J.R. Garbow, D.P. Weitekamp, and A. Pines. “Bilinear rotation decoupling of homonuclear scalar interactions”. In: *Chemical Physics Letters* 93.5 (1982), pp. 504–509. DOI: [https://doi.org/10.1016/0009-2614\(82\)83229-6](https://doi.org/10.1016/0009-2614(82)83229-6).
- [70] Mario F Gely and Gary A Steele. “QuCAT: quantum circuit analyzer tool in Python”. In: *New Journal of Physics* 22.1 (Jan. 2020), p. 013025. DOI: 10.1088/1367-2630/ab60f6.
- [71] Craig Gidney. *How to factor 2048 bit RSA integers with less than a million noisy qubits*. 2025. arXiv: 2505.15917 [quant-ph].
- [72] Alysso Gold et al. “Entanglement across separate silicon dies in a modular superconducting qubit device”. In: *npj Quantum Information* 7.1 (Sept. 2021), p. 142. DOI: 10.1038/s41534-021-00484-1.
- [73] Daniel Gottesman, Alexei Kitaev, and John Preskill. “Encoding a qubit in an oscillator”. In: *Phys. Rev. A* 64 (1 June 2001), p. 012310. DOI: 10.1103/PhysRevA.64.012310.
- [74] Grigorij J. Grabovskij et al. “Strain Tuning of Individual Atomic Tunneling Systems Detected by a Superconducting Qubit”. In: *Science* 338.6104 (2012), pp. 232–234. DOI: 10.1126/science.1226487.
- [75] Lov K. Grover. “A fast quantum mechanical algorithm for database search”. In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. New York, NY, USA: Association for Computing Machinery, 1996, pp. 212–219. DOI: 10.1145/237814.237866.
- [76] Lov K. Grover. *A fast quantum mechanical algorithm for database search*. 1996. arXiv: quant-ph/9605043 [quant-ph].
- [77] Gurobi Optimization, LLC. *Gurobi Optimizer Reference Manual*. 2021.

- [78] András Gyenis et al. “Experimental Realization of a Protected Superconducting Circuit Derived from the $0-\pi$ Qubit”. In: *PRX Quantum* 2 (1 Mar. 2021), p. 010339. DOI: 10.1103/PRXQuantum.2.010339.
- [79] András Gyenis et al. “Moving beyond the Transmon: Noise-Protected Superconducting Quantum Circuits”. In: *PRX Quantum* 2 (3 Sept. 2021), p. 030101. DOI: 10.1103/PRXQuantum.2.030101.
- [80] William E Hart, Jean-Paul Watson, and David L Woodruff. “Pyomo: modeling and solving mathematical programs in Python”. In: *Math. Program. Comput.* 3.3 (2011), pp. 219–260.
- [81] Akel Hashim et al. “Practical Introduction to Benchmarking and Characterization of Quantum Computers”. In: *PRX Quantum* 6 (3 Aug. 2025), p. 030202. DOI: 10.1103/PRXQuantum.6.030202.
- [82] T. M. Hazard et al. “Characterization of superconducting through-silicon vias as capacitive elements in quantum circuits”. In: *Applied Physics Letters* 123.15 (Oct. 2023), p. 154004. DOI: 10.1063/5.0170055.
- [83] Fabio Henriques et al. “Phonon traps reduce the quasiparticle density in superconducting circuits”. In: *Applied Physics Letters* 115.21 (Nov. 2019), p. 212601. DOI: 10.1063/1.5124967.
- [84] Jared B. Hertzberg et al. “Laser-annealing Josephson junctions for yielding scaled-up superconducting quantum processors”. In: *npj Quantum Information* 7.1 (Aug. 2021), p. 129. DOI: 10.1038/s41534-021-00464-5.
- [85] A. A. Houck et al. “Controlling the Spontaneous Emission of a Superconducting Transmon Qubit”. In: *Phys. Rev. Lett.* 101 (8 Aug. 2008), p. 080502. DOI: 10.1103/PhysRevLett.101.080502.
- [86] L. Hu et al. “Quantum error correction and universal gate set operation on a binomial bosonic logical qubit”. In: *Nature Physics* 15.5 (May 2019), pp. 503–508. DOI: 10.1038/s41567-018-0414-3.
- [87] Cupjin Huang et al. “Quantum Instruction Set Design for Performance”. In: *Phys. Rev. Lett.* 130 (7 Feb. 2023), p. 070601. DOI: 10.1103/PhysRevLett.130.070601.
- [88] Eric Hyypä et al. “Reducing Leakage of Single-Qubit Gates for Superconducting Quantum Processors Using Analytical Control Pulse Envelopes”. In: *PRX Quantum* 5 (3 Sept. 2024), p. 030353. DOI: 10.1103/PRXQuantum.5.030353.
- [89] V. Iaia et al. “Phonon downconversion to suppress correlated errors in superconducting qubits”. In: *Nature Communications* 13.1 (Oct. 2022), p. 6425. DOI: 10.1038/s41467-022-33997-0.
- [90] X. Y. Jin et al. *Fast, tunable, high fidelity cZ-gates between superconducting qubits with parametric microwave control of ZZ-coupling*. 2024. arXiv: 2305.02907 [quant-ph].

- [91] Konstantin Kalashnikov et al. “Bifluxon: Fluxon-Parity-Protected Superconducting Qubit”. In: *PRX Quantum* 1 (1 Sept. 2020), p. 010307. DOI: 10.1103/PRXQuantum.1.010307.
- [92] A. Kandala et al. “Demonstration of a High-Fidelity cnot Gate for Fixed-Frequency Transmons with Engineered ZZ Suppression”. In: *Phys. Rev. Lett.* 127 (13 Sept. 2021), p. 130501. DOI: 10.1103/PhysRevLett.127.130501.
- [93] J. Kelly et al. “Optimal Quantum Control Using Randomized Benchmarking”. In: *Phys. Rev. Lett.* 112 (24 June 2014), p. 240504. DOI: 10.1103/PhysRevLett.112.240504.
- [94] Mostafa Khezri et al. “Measurement-induced state transitions in a superconducting qubit: Within the rotating-wave approximation”. In: *Phys. Rev. Appl.* 20 (5 Nov. 2023), p. 054008. DOI: 10.1103/PhysRevApplied.20.054008.
- [95] Hyunseong Kim et al. “Effects of laser-annealing on fixed-frequency superconducting qubits”. In: *Applied Physics Letters* 121.14 (Oct. 2022), p. 142601. DOI: 10.1063/5.0102092.
- [96] Yosep Kim et al. “High-fidelity three-qubit iToffoli gate for fixed-frequency superconducting qubits”. In: *Nature Physics* 18.7 (July 2022), pp. 783–788. DOI: 10.1038/s41567-022-01590-3.
- [97] Youngseok Kim et al. “Error mitigation with stabilized noise in superconducting quantum processors”. In: *Nature Communications* 16.1 (Sept. 2025), p. 8439. DOI: 10.1038/s41467-025-62820-9.
- [98] A.Yu. Kitaev. “Fault-tolerant quantum computation by anyons”. In: *Annals of Physics* 303.1 (2003), pp. 2–30. DOI: [https://doi.org/10.1016/S0003-4916\(02\)00018-0](https://doi.org/10.1016/S0003-4916(02)00018-0).
- [99] Ian D. Kivlichan et al. “Quantum Simulation of Electronic Structure with Linear Depth and Connectivity”. In: *Phys. Rev. Lett.* 120 (11 Mar. 2018), p. 110501. DOI: 10.1103/PhysRevLett.120.110501.
- [100] P. V. Klimov et al. “Fluctuations of Energy-Relaxation Times in Superconducting Qubits”. In: *Phys. Rev. Lett.* 121 (9 Aug. 2018), p. 090502. DOI: 10.1103/PhysRevLett.121.090502.
- [101] Paul V. Klimov et al. “Optimizing quantum gates towards the scale of logical qubits”. In: *Nature Communications* 15.1 (Mar. 2024), p. 2442. DOI: 10.1038/s41467-024-46623-y.
- [102] Paul V. Klimov et al. *The Snake Optimizer for Learning Quantum Processor Control Parameters*. 2020. arXiv: 2006.04594 [quant-ph].
- [103] E. Knill et al. “Randomized benchmarking of quantum gates”. In: *Phys. Rev. A* 77 (1 Jan. 2008), p. 012307. DOI: 10.1103/PhysRevA.77.012307.

- [104] Emanuel Knill, Raymond Laflamme, and Wojciech H. Zurek. “Resilient Quantum Computation”. In: *Science* 279.5349 (1998), pp. 342–345. DOI: 10.1126/science.279.5349.342.
- [105] Jens Koch et al. “Charge-insensitive qubit design derived from the Cooper pair box”. In: *Phys. Rev. A* 76 (4 Oct. 2007), p. 042319. DOI: 10.1103/PhysRevA.76.042319.
- [106] Jin Ming Koh et al. “Measurement-induced entanglement phase transition on a superconducting quantum processor with mid-circuit readout”. In: *Nature Physics* 19.9 (Sept. 2023), pp. 1314–1319. DOI: 10.1038/s41567-023-02076-6.
- [107] J M Kreikebaum et al. “Improving wafer-scale Josephson junction resistance variation in superconducting quantum coherent circuits”. In: *Superconductor Science and Technology* 33.6 (Apr. 2020), 06LT02. DOI: 10.1088/1361-6668/ab8617.
- [108] John Mark Kreikebaum. “Superconducting Qubit Enabled Single Microwave Photon Detection”. PhD thesis. University of California, Berkeley, 2020.
- [109] Jaseung Ku et al. “Suppression of Unwanted ZZ Interactions in a Hybrid Two-Qubit System”. In: *Phys. Rev. Lett.* 125 (20 Nov. 2020), p. 200504. DOI: 10.1103/PhysRevLett.125.200504.
- [110] P. Kumar et al. “Origin and Reduction of $1/f$ Magnetic Flux Noise in Superconducting Devices”. In: *Phys. Rev. Appl.* 6 (4 Oct. 2016), p. 041001. DOI: 10.1103/PhysRevApplied.6.041001.
- [111] Pavel D. Kurilovich et al. *High-frequency readout free from transmon multi-excitation resonances*. 2025. arXiv: 2501.09161 [quant-ph].
- [112] Randy Kwende, Theodore White, and Ofer Naaman. “Josephson parametric circulator with same-frequency signal ports, 200 MHz bandwidth, and high dynamic range”. In: *Applied Physics Letters* 122.22 (May 2023), p. 224001. DOI: 10.1063/5.0150427.
- [113] Eli M Levenson-Falk and Sadman Ahmed Shanto. “A review of design concerns in superconducting quantum circuits”. In: *Materials for Quantum Technology* 5.2 (June 2025), p. 022003. DOI: 10.1088/2633-4356/ade10d.
- [114] Jürgen Lisenfeld, Alexander Bilmes, and Alexey V. Ustinov. “Enhancing the coherence of superconducting quantum bits with electric fields”. In: *npj Quantum Information* 9.1 (Jan. 2023), p. 8. DOI: 10.1038/s41534-023-00678-9.
- [115] Jürgen Lisenfeld et al. “Electric field spectroscopy of material defects in transmon qubits”. In: *npj Quantum Information* 5.1 (Nov. 2019), p. 105. DOI: 10.1038/s41534-019-0224-1.
- [116] Jürgen Lisenfeld et al. *Mapping the positions of Two-Level-Systems on the surface of a superconducting transmon qubit*. 2025. arXiv: 2511.05365 [quant-ph].
- [117] Jürgen Lisenfeld et al. “Observation of directly interacting coherent two-level systems in an amorphous material”. In: *Nature Communications* 6.1 (Feb. 2015), p. 6182. DOI: 10.1038/ncomms7182.

- [118] P. London et al. “Strong driving of a single spin using arbitrarily polarized fields”. In: *Phys. Rev. A* 90 (1 July 2014), p. 012302. DOI: 10.1103/PhysRevA.90.012302.
- [119] Yao Lu. “Parametric control of flux-tunable superconducting circuits”. PhD thesis. University of Chicago, 2019. DOI: 10.6082/uchicago.1973.
- [120] Erik Lucero et al. “High-Fidelity Gates in a Single Josephson Qubit”. In: *Phys. Rev. Lett.* 100 (24 June 2008), p. 247001. DOI: 10.1103/PhysRevLett.100.247001.
- [121] Vladimir E. Manucharyan et al. “Fluxonium: Single Cooper-Pair Circuit Free of Charge Offsets”. In: *Science* 326.5949 (2009), pp. 113–116. DOI: 10.1126/science.1175552.
- [122] John M. Martinis, Michel H. Devoret, and John Clarke. “Energy-Level Quantization in the Zero-Voltage State of a Current-Biased Josephson Junction”. In: *Phys. Rev. Lett.* 55 (15 Oct. 1985), pp. 1543–1546. DOI: 10.1103/PhysRevLett.55.1543.
- [123] John M. Martinis, Michel H. Devoret, and John Clarke. “Experimental tests for the quantum behavior of a macroscopic degree of freedom: The phase difference across a Josephson junction”. In: *Phys. Rev. B* 35 (10 Apr. 1987), pp. 4682–4698. DOI: 10.1103/PhysRevB.35.4682.
- [124] John M. Martinis et al. “Decoherence in Josephson Qubits from Dielectric Loss”. In: *Phys. Rev. Lett.* 95 (21 Nov. 2005), p. 210503. DOI: 10.1103/PhysRevLett.95.210503.
- [125] John M. Martinis et al. “Rabi Oscillations in a Large Josephson-Junction Qubit”. In: *Phys. Rev. Lett.* 89 (11 Aug. 2002), p. 117901. DOI: 10.1103/PhysRevLett.89.117901.
- [126] Matt McEwen, Dave Bacon, and Craig Gidney. “Relaxing Hardware Requirements for Surface Code Circuits using Time-dynamics”. In: *Quantum* 7 (Nov. 2023), p. 1172. DOI: 10.22331/q-2023-11-07-1172.
- [127] Matt McEwen et al. “Resisting High-Energy Impact Events through Gap Engineering in Superconducting Qubit Arrays”. In: *Phys. Rev. Lett.* 133 (24 Dec. 2024), p. 240601. DOI: 10.1103/PhysRevLett.133.240601.
- [128] Matt McEwen et al. “Resolving catastrophic error bursts from cosmic rays in large arrays of superconducting qubits”. In: *Nature Physics* 18.1 (Jan. 2022), pp. 107–111. DOI: 10.1038/s41567-021-01432-8.
- [129] David C. McKay et al. “Efficient Z gates for quantum computing”. In: *Phys. Rev. A* 96 (2 Aug. 2017), p. 022330. DOI: 10.1103/PhysRevA.96.022330.
- [130] David C. McKay et al. “Universal Gate for Fixed-Frequency Qubits via a Tunable Bus”. In: *Phys. Rev. Appl.* 6 (6 Dec. 2016), p. 064007. DOI: 10.1103/PhysRevApplied.6.064007.

- [131] M. Mergenthaler et al. “Ultrahigh vacuum packaging and surface cleaning for quantum devices”. In: *Review of Scientific Instruments* 92.2 (Feb. 2021), p. 025121. DOI: 10.1063/5.0034574.
- [132] Xiao Mi et al. “Information scrambling in quantum circuits”. In: *Science* 374.6574 (2021), pp. 1479–1483. DOI: 10.1126/science.abg5029.
- [133] Marios H. Michael et al. “New Class of Quantum Error-Correcting Codes for a Bosonic Mode”. In: *Phys. Rev. X* 6 (3 July 2016), p. 031006. DOI: 10.1103/PhysRevX.6.031006.
- [134] Zlatko K. Mineev et al. *Circuit quantum electrodynamics (cQED) with modular quasi-lumped models*. 2021. arXiv: 2103.10344 [quant-ph].
- [135] Zlatko K. Mineev et al. “Energy-participation quantization of Josephson circuits”. In: *npj Quantum Information* 7.1 (Aug. 2021), p. 131. DOI: 10.1038/s41534-021-00461-8.
- [136] Mazyar Mirrahimi et al. “Dynamically protected cat-qubits: a new paradigm for universal quantum computation”. In: *New Journal of Physics* 16.4 (Apr. 2014), p. 045014. DOI: 10.1088/1367-2630/16/4/045014.
- [137] Bradley K. Mitchell. “Investigating Microwave-Activated Entangling Gates on Superconducting Quantum Processors”. PhD thesis. University of California, Berkeley, 2022.
- [138] Bradley K. Mitchell et al. “Hardware-Efficient Microwave-Activated Tunable Coupling between Superconducting Qubits”. In: *Phys. Rev. Lett.* 127 (20 Nov. 2021), p. 200502. DOI: 10.1103/PhysRevLett.127.200502.
- [139] Alexis Morvan et al. “Optimizing frequency allocation for fixed-frequency superconducting quantum processors”. In: *Phys. Rev. Res.* 4 (2 Apr. 2022), p. 023079. DOI: 10.1103/PhysRevResearch.4.023079.
- [140] Daria A. Moskaleva et al. “Wafer-scale uniformity improvement of Dolan-bridge Josephson junction by shadow evaporation bias correction”. In: *Scientific Reports* 15.1 (Aug. 2025), p. 28097. DOI: 10.1038/s41598-025-08787-5.
- [141] F. Motzoi and F. K. Wilhelm. “Improving frequency selection of driven pulses using derivative-based transition suppression”. In: *Phys. Rev. A* 88 (6 Dec. 2013), p. 062318. DOI: 10.1103/PhysRevA.88.062318.
- [142] F. Motzoi et al. “Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits”. In: *Phys. Rev. Lett.* 103 (11 Sept. 2009), p. 110501. DOI: 10.1103/PhysRevLett.103.110501.
- [143] Clemens Müller, Jared H Cole, and Jürgen Lisenfeld. “Towards understanding two-level-systems in amorphous solids: insights from quantum circuits”. In: *Reports on Progress in Physics* 82.12 (Oct. 2019), p. 124501. DOI: 10.1088/1361-6633/ab3a7e.

- [144] Clemens Müller et al. “Interacting two-level defects as sources of fluctuating high-frequency noise in superconducting circuits”. In: *Phys. Rev. B* 92 (3 July 2015), p. 035442. DOI: 10.1103/PhysRevB.92.035442.
- [145] Pranav Mundada et al. “Suppression of Qubit Crosstalk in a Tunable Coupling Superconducting Circuit”. In: *Phys. Rev. Appl.* 12 (5 Nov. 2019), p. 054023. DOI: 10.1103/PhysRevApplied.12.054023.
- [146] Y. Nakamura, Yu. A. Pashkin, and J. S. Tsai. “Coherent control of macroscopic quantum states in a single-Cooper-pair box”. In: *Nature* 398.6730 (Apr. 1999), pp. 786–788. DOI: 10.1038/19718.
- [147] C. Neill et al. “Accurately computing the electronic properties of a quantum ring”. In: *Nature* 594.7864 (June 2021), pp. 508–512. DOI: 10.1038/s41586-021-03576-2.
- [148] Charles Neill. “A path towards quantum supremacy with superconducting qubits”. PhD thesis. University of California, Santa Barbara, 2017.
- [149] Long B. Nguyen et al. “High-Coherence Fluxonium Qubit”. In: *Phys. Rev. X* 9 (4 Nov. 2019), p. 041041. DOI: 10.1103/PhysRevX.9.041041.
- [150] Long B. Nguyen et al. “Programmable Heisenberg interactions between Floquet qubits”. In: *Nature Physics* 20.2 (Feb. 2024), pp. 240–246. DOI: 10.1038/s41567-023-02326-7.
- [151] Simon E. Nigg et al. “Black-Box Superconducting Circuit Quantization”. In: *Phys. Rev. Lett.* 108 (24 June 2012), p. 240502. DOI: 10.1103/PhysRevLett.108.240502.
- [152] Jingjing Niu et al. “Low-loss interconnects for modular superconducting quantum processors”. In: *Nature Electronics* 6.3 (Mar. 2023), pp. 235–241. DOI: 10.1038/s41928-023-00925-z.
- [153] H. Nyquist. “Thermal Agitation of Electric Charge in Conductors”. In: *Phys. Rev.* 32 (1 July 1928), pp. 110–113. DOI: 10.1103/PhysRev.32.110.
- [154] Mutasem Odeh et al. “Non-Markovian dynamics of a superconducting qubit in a phononic bandgap”. In: *Nature Physics* 21.3 (Mar. 2025), pp. 406–411. DOI: 10.1038/s41567-024-02740-5.
- [155] Nissim Ofek et al. “Extending the lifetime of a quantum bit with error correction in superconducting circuits”. In: *Nature* 536.7617 (Aug. 2016), pp. 441–445. DOI: 10.1038/nature18949.
- [156] Amr Osman et al. “Mitigation of frequency collisions in superconducting quantum processors”. In: *Phys. Rev. Res.* 5 (4 Oct. 2023), p. 043001. DOI: 10.1103/PhysRevResearch.5.043001.
- [157] B. I. Halperin P. w. Anderson and c. M. Varma. “Anomalous low-temperature thermal properties of glasses and spin glasses”. In: *The Philosophical Magazine: A Journal of Theoretical Experimental and Applied Physics* 25.1 (1972), pp. 1–9. DOI: 10.1080/14786437208229210.

- [158] David P. Pappas et al. “Alternating-bias assisted annealing of amorphous oxide tunnel junctions”. In: *Communications Materials* 5.1 (Aug. 2024), p. 150. DOI: 10.1038/s43246-024-00596-z.
- [159] Yu. A. Pashkin et al. “Quantum oscillations in two coupled charge qubits”. In: *Nature* 421.6925 (Feb. 2003), pp. 823–826. DOI: 10.1038/nature01365.
- [160] Clayton R. Paul. *Analysis of Multiconductor Transmission Lines*. Second. Wiley-IEEE Press, 2007.
- [161] Ivan V. Pechenezhskiy et al. “The superconducting quasicharge qubit”. In: *Nature* 585.7825 (Sept. 2020), pp. 368–371. DOI: 10.1038/s41586-020-2687-9.
- [162] Kaidong Peng et al. “X-parameter based design and simulation of Josephson traveling-wave parametric amplifiers for quantum computing applications”. In: *2022 IEEE International Conference on Quantum Computing and Engineering (QCE)*. 2022, pp. 331–340. DOI: 10.1109/QCE53715.2022.00054.
- [163] W A Phillips. “Two-level states in glasses”. In: *Reports on Progress in Physics* 50.12 (Dec. 1987), p. 1657. DOI: 10.1088/0034-4885/50/12/003.
- [164] David M. Pozar. *Microwave engineering*. Fourth. Wiley, 2012.
- [165] E. M. Purcell. “Spontaneous Emission Probabilities at Radio Frequencies”. In: *Proceedings of the American Physical Society*. Vol. 69. American Physical Society, June 1946, p. 681. DOI: 10.1103/PhysRev.69.674.2.
- [166] Harald Putterman et al. “Hardware-efficient quantum error correction via concatenated bosonic qubits”. In: *Nature* 638.8052 (Feb. 2025), pp. 927–934. DOI: 10.1038/s41586-025-08642-7.
- [167] Chad Rigetti and Michel Devoret. “Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies”. In: *Phys. Rev. B* 81 (13 Apr. 2010), p. 134507. DOI: 10.1103/PhysRevB.81.134507.
- [168] R.-P. Riwar and D. P. DiVincenzo. “Circuit quantization with time-dependent magnetic fields for realistic geometries”. In: *npj Quantum Information* 8.1 (Mar. 2022), p. 36. DOI: 10.1038/s41534-022-00539-x.
- [169] M. A. Rol et al. “Time-domain characterization and correction of on-chip distortion of control pulses in a quantum processor”. In: *Applied Physics Letters* 116.5 (Feb. 2020), p. 054001. DOI: 10.1063/1.5133894.
- [170] Ilan T. Rosen et al. “A synthetic magnetic vector potential in a 2D superconducting qubit array”. In: *Nature Physics* 20.12 (Dec. 2024), pp. 1881–1887. DOI: 10.1038/s41567-024-02661-3.
- [171] D. Rosenberg et al. “3D integrated superconducting qubits”. In: *npj Quantum Information* 3.1 (Oct. 2017), p. 42. DOI: 10.1038/s41534-017-0044-0.

- [172] P. Roushan et al. “Chiral ground-state currents of interacting photons in a synthetic magnetic field”. In: *Nature Physics* 13.2 (Feb. 2017), pp. 146–151. DOI: 10.1038/nphys3930.
- [173] David A. Rower et al. “Suppressing Counter-Rotating Errors for Fast Single-Qubit Gates with Fluxonium”. In: *PRX Quantum* 5 (4 Dec. 2024), p. 040342. DOI: 10.1103/PRXQuantum.5.040342.
- [174] Daniel Sank et al. “Balanced coupling in electromagnetic circuits”. In: *Phys. Rev. Appl.* 23 (2 Feb. 2025), p. 024012. DOI: 10.1103/PhysRevApplied.23.024012.
- [175] Daniel Sank et al. “Measurement-Induced State Transitions in a Superconducting Qubit: Beyond the Rotating Wave Approximation”. In: *Phys. Rev. Lett.* 117 (19 Nov. 2016), p. 190503. DOI: 10.1103/PhysRevLett.117.190503.
- [176] B. Sarabi et al. “Projected Dipole Moments of Individual Two-Level Defects Extracted Using Circuit Quantum Electrodynamics”. In: *Phys. Rev. Lett.* 116 (16 Apr. 2016), p. 167002. DOI: 10.1103/PhysRevLett.116.167002.
- [177] Andre Schneider et al. “Local sensing with the multilevel ac Stark effect”. In: *Physical Review A* 97.6 (June 2018). DOI: 10.1103/physreva.97.062334.
- [178] J. A. Schreier et al. “Suppressing charge noise decoherence in superconducting charge qubits”. In: *Phys. Rev. B* 77 (18 May 2008), p. 180502. DOI: 10.1103/PhysRevB.77.180502.
- [179] David Schuster. “Circuit Quantum Electrodynamics”. PhD thesis. Yale University, 2007.
- [180] Yoni Shalibo et al. “Lifetime and Coherence of Two-Level Defects in a Josephson Junction”. In: *Phys. Rev. Lett.* 105 (17 Oct. 2010), p. 177001. DOI: 10.1103/PhysRevLett.105.177001.
- [181] Sarah Sheldon et al. “Procedure for systematically tuning up cross-talk in the cross-resonance gate”. In: *Phys. Rev. A* 93 (6 June 2016), p. 060302. DOI: 10.1103/PhysRevA.93.060302.
- [182] Ross Shillito et al. “Dynamics of Transmon Ionization”. In: *Phys. Rev. Appl.* 18 (3 Sept. 2022), p. 034031. DOI: 10.1103/PhysRevApplied.18.034031.
- [183] P.W. Shor. “Algorithms for quantum computation: discrete logarithms and factoring”. In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. 1994, pp. 124–134. DOI: 10.1109/SFCS.1994.365700.
- [184] P.W. Shor. “Fault-tolerant quantum computation”. In: *Proceedings of 37th Conference on Foundations of Computer Science*. 1996, pp. 56–65. DOI: 10.1109/SFCS.1996.548464.
- [185] Irfan Siddiqi. “Engineering high-coherence superconducting qubits”. In: *Nature Reviews Materials* 6.10 (Oct. 2021), pp. 875–891. DOI: 10.1038/s41578-021-00370-4.

- [186] Volodymyr Sivak et al. *Reinforcement Learning Control of Quantum Error Correction*. 2025. arXiv: 2511.08493 [quant-ph].
- [187] W. C. Smith et al. “Superconducting circuit protected by two-Cooper-pair tunneling”. In: *npj Quantum Information* 6.1 (Jan. 2020), p. 8. DOI: 10.1038/s41534-019-0231-2.
- [188] Firat Solgun, David W. Abraham, and David P. DiVincenzo. “Blackbox quantization of superconducting circuits using exact impedance synthesis”. In: *Phys. Rev. B* 90 (13 Oct. 2014), p. 134504. DOI: 10.1103/PhysRevB.90.134504.
- [189] Matthias Steffen et al. “Measurement of the Entanglement of Two Superconducting Qubits via State Tomography”. In: *Science* 313.5792 (2006), pp. 1423–1425. DOI: 10.1126/science.1130886.
- [190] J. Stehlik et al. “Tunable Coupling Architecture for Fixed-Frequency Transmon Superconducting Qubits”. In: *Phys. Rev. Lett.* 127 (8 Aug. 2021), p. 080505. DOI: 10.1103/PhysRevLett.127.080505.
- [191] Y. Sunada et al. “Fast Readout and Reset of a Superconducting Qubit Coupled to a Resonator with an Intrinsic Purcell Filter”. In: *Phys. Rev. Appl.* 17 (4 Apr. 2022), p. 044016. DOI: 10.1103/PhysRevApplied.17.044016.
- [192] Neereja Sundaresan et al. “Reducing Unitary and Spectator Errors in Cross Resonance with Optimized Rotary Echoes”. In: *PRX Quantum* 1 (2 Dec. 2020), p. 020318. DOI: 10.1103/PRXQuantum.1.020318.
- [193] Youngkyu Sung et al. “Realization of High-Fidelity CZ and ZZ-Free iSWAP Gates with a Tunable Coupler”. In: *Phys. Rev. X* 11 (2 June 2021), p. 021058. DOI: 10.1103/PhysRevX.11.021058.
- [194] Maika Takita et al. “Experimental Demonstration of Fault-Tolerant State Preparation with Superconducting Qubits”. In: *Phys. Rev. Lett.* 119 (18 Oct. 2017), p. 180501. DOI: 10.1103/PhysRevLett.119.180501.
- [195] B M Terhal, J Conrad, and C Vuillot. “Towards scalable bosonic quantum error correction”. In: *Quantum Science and Technology* 5.4 (July 2020), p. 043001. DOI: 10.1088/2058-9565/ab98a5.
- [196] L. S. Theis et al. “Counteracting systems of diabaticities using DRAG controls: The status after 10 years(a)”. In: *Europhysics Letters* 123.6 (Oct. 2018), p. 60001. DOI: 10.1209/0295-5075/123/60001.
- [197] Ted Thorbeck et al. “Two-Level-System Dynamics in a Superconducting Qubit Due to Background Ionizing Radiation”. In: *PRX Quantum* 4 (2 June 2023), p. 020356. DOI: 10.1103/PRXQuantum.4.020356.
- [198] Vinay Tripathi, Mostafa Khezri, and Alexander N. Korotkov. “Operation and intrinsic error budget of a two-qubit cross-resonance gate”. In: *Phys. Rev. A* 100 (1 July 2019), p. 012301. DOI: 10.1103/PhysRevA.100.012301.

- [199] J. Van Damme et al. “Advanced CMOS manufacturing of superconducting qubits on 300 mm wafers”. In: *Nature* 634.8032 (Oct. 2024), pp. 74–79. DOI: 10.1038/s41586-024-07941-9.
- [200] Antti P. Vepsäläinen et al. “Impact of ionizing radiation on superconducting qubit coherence”. In: *Nature* 584.7822 (Aug. 2020), pp. 551–556. DOI: 10.1038/s41586-020-2619-8.
- [201] Uri Vool and Michel Devoret. “Introduction to quantum electromagnetic circuits”. In: *International Journal of Circuit Theory and Applications* 45.7 (2017), pp. 897–934. DOI: <https://doi.org/10.1002/cta.2359>.
- [202] A. Wallraff et al. “Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics”. In: *Nature* 431.7005 (Sept. 2004), pp. 162–167. DOI: 10.1038/nature02851.
- [203] C. Wang et al. “Surface participation and dielectric loss in superconducting qubits”. In: *Applied Physics Letters* 107.16 (Oct. 2015), p. 162601. DOI: 10.1063/1.4934486.
- [204] Xiqiao Wang et al. *Precision frequency tuning of tunable transmon qubits using alternating-bias assisted annealing*. 2024. arXiv: 2407.06425 [quant-ph].
- [205] S. Weeden et al. *Statistics of Strongly Coupled Defects in Superconducting Qubits*. 2025. arXiv: 2506.00193 [quant-ph].
- [206] K. X. Wei et al. “Hamiltonian Engineering with Multicolor Drives for Fast Entangling Gates and Quantum Crosstalk Cancellation”. In: *Phys. Rev. Lett.* 129 (6 Aug. 2022), p. 060501. DOI: 10.1103/PhysRevLett.129.060501.
- [207] J. Wenner et al. “Surface loss simulations of superconducting coplanar waveguide resonators”. In: *Applied Physics Letters* 99.11 (Sept. 2011), p. 113513. DOI: 10.1063/1.3637047.
- [208] C. D. Wilen et al. “Correlated charge noise and relaxation errors in superconducting qubits”. In: *Nature* 594.7863 (June 2021), pp. 369–373. DOI: 10.1038/s41586-021-03557-5.
- [209] Christopher J. Wood and Jay M. Gambetta. “Quantification and characterization of leakage errors”. In: *Phys. Rev. A* 97 (3 Mar. 2018), p. 032306. DOI: 10.1103/PhysRevA.97.032306.
- [210] Fei Yan et al. “Tunable Coupling Scheme for Implementing High-Fidelity Two-Qubit Gates”. In: *Phys. Rev. Appl.* 10 (5 Nov. 2018), p. 054062. DOI: 10.1103/PhysRevApplied.10.054062.
- [211] Alec Yen et al. “Interferometric Purcell suppression of spontaneous emission in a superconducting qubit”. In: *Phys. Rev. Appl.* 23 (2 Feb. 2025), p. 024068. DOI: 10.1103/PhysRevApplied.23.024068.

- [212] Xinyuan You, J. A. Sauls, and Jens Koch. “Circuit quantization in the presence of time-dependent external flux”. In: *Phys. Rev. B* 99 (17 May 2019), p. 174512. DOI: 10.1103/PhysRevB.99.174512.
- [213] Eric J. Zhang et al. “High-performance superconducting quantum processors via laser annealing of transmon qubits”. In: *Science Advances* 8.19 (2022), eabi6690. DOI: 10.1126/sciadv.abi6690.
- [214] Zi-Huai Zhang et al. “Acceptor-Induced Bulk Dielectric Loss in Superconducting Circuits on Silicon”. In: *Phys. Rev. X* 14 (4 Oct. 2024), p. 041022. DOI: 10.1103/PhysRevX.14.041022.
- [215] Jun Zhang et al. “Geometric theory of nonlocal two-qubit operations”. In: *Phys. Rev. A* 67 (4 Apr. 2003), p. 042313. DOI: 10.1103/PhysRevA.67.042313.
- [216] Jun Zhang et al. “Minimum Construction of Two-Qubit Quantum Operations”. In: *Phys. Rev. Lett.* 93 (2 July 2004), p. 020502. DOI: 10.1103/PhysRevLett.93.020502.
- [217] Peng Zhao et al. “Combating fluctuations in relaxation times of fixed-frequency transmon qubits with microwave-dressed states”. In: *Phys. Rev. A* 105 (6 June 2022), p. 062605. DOI: 10.1103/PhysRevA.105.062605.
- [218] Peng Zhao et al. “High-Contrast ZZ Interaction Using Superconducting Qubits with Opposite-Sign Anharmonicity”. In: *Phys. Rev. Lett.* 125 (20 Nov. 2020), p. 200503. DOI: 10.1103/PhysRevLett.125.200503.
- [219] Wojciech H. Zurek. “Decoherence and the transition from quantum to classical”. In: *Physics Today* 44 (1991), pp. 36–44.

Appendix A

Circuits

A.1 Resonances

Damped and driven harmonic oscillators are found in many physical systems beyond electrical circuits. To build some physical intuition for such a system, we can consider a 1-dimensional driven damped harmonic oscillator with mass m , intrinsic resonance frequency $\omega_0 = \sqrt{\frac{k}{m}}$, and damping coefficient γ . The differential equation governing the motion of this system is given by

$$\ddot{x} + \gamma\dot{x} + \omega_0^2 x = \frac{F(t)}{m} \quad (\text{A.1})$$

For $F(t) = 0$, we can guess a solution of the form Ae^{st} , so we have

$$Ae^{st} (s^2 + \gamma s + \omega_0^2) = 0. \quad (\text{A.2})$$

Solving for s , we find

$$s = -\frac{\gamma}{2} \pm \sqrt{\left(\frac{\gamma}{2}\right)^2 - \omega_0^2} = -\frac{\gamma}{2} \pm i\omega_0 \sqrt{1 - \left(\frac{\gamma}{2\omega_0}\right)^2} \quad (\text{A.3})$$

In the limit of small damping ($\gamma < 2\omega_0$), $x(t)$ is of the form

$$x(t) = (x_+ e^{i\Omega_0 t} + x_- e^{-i\Omega_0 t}) e^{-\frac{\gamma}{2}t} \quad (\text{A.4})$$

where the new resonance frequency is

$$\Omega_0 = \omega_0 \sqrt{1 - \frac{\gamma^2}{4\omega_0^2}} \approx \omega_0 \left[1 + \frac{1}{8} \frac{\gamma^2}{\omega_0^2} \right]. \quad (\text{A.5})$$

Here we see that the homogeneous solution leads to transient behavior, which will tend to $x(t) \rightarrow 0$ in the limit $t \gg \gamma^{-1}$.

In the case of a periodic driving force, we define $\text{Re}[z(t)] = x(t)$, and we can write the equation of motion as

$$\ddot{z} + \gamma\dot{z} + \omega_0^2 z = \frac{F_0}{m} e^{i\omega t} \quad (\text{A.6})$$

and guess a solution of the form $z(t) = Ae^{i\omega t}$. Plugging in this guess, we find

$$A(-\omega^2 + i\gamma\omega + \omega_0^2)e^{i\omega t} = \frac{F_0}{m}e^{i\omega t} \quad (\text{A.7})$$

from which we see that A is entirely fixed. We can solve for the amplitude and phase as

$$A = \frac{F_0}{m} \frac{1}{(\omega_0^2 - \omega^2) + i\gamma\omega} = x_d e^{-i\phi} \quad (\text{A.8})$$

$$x_d = \frac{F_0}{m} \frac{1}{\sqrt{(\omega_0^2 - \omega^2)^2 + \gamma^2\omega^2}} \quad (\text{A.9})$$

$$\phi = \arctan\left(\frac{\gamma\omega}{\omega_0^2 - \omega^2}\right)$$

What is worth noting is that the maximum amplitude response here occurs at a resonant driving frequency

$$\omega_d = \sqrt{\omega_0^2 - \frac{\gamma^2}{2}} = \omega_0 \sqrt{1 - \frac{\gamma^2}{2\omega_0^2}} \quad (\text{A.10})$$

A.2 Coupling

Capacitive T-Coupling

The Lagrangian for the system is

$$L = \frac{1}{2}\dot{\Phi}^T \mathbf{C} \dot{\Phi} - \frac{1}{2}\Phi^T \mathbf{L}^{-1} \Phi \quad (\text{A.11})$$

where the capacitance and inverse inductance matrices are

$$\mathbf{C} = \begin{pmatrix} C_1 & -C_1 & 0 \\ -C_1 & C_1 + C_c + C_2 & -C_2 \\ 0 & -C_2 & C_2 \end{pmatrix} \quad \mathbf{L}^{-1} = \begin{pmatrix} \frac{1}{L_1} & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \frac{1}{L_2} \end{pmatrix} \quad (\text{A.12})$$

From here, it is clear that ϕ_c does not appear in the Lagrangian so it is a free-variable that must have a corresponding conjugate charge Q_c that is conserved. The Euler-Lagrange equations tell us that

$$\frac{\partial L}{\partial \dot{\phi}_c} = (C_1 + C_c + C_2)\dot{\phi}_c - C_1\dot{\phi}_1 - C_2\dot{\phi}_2 = Q_c \quad (\text{A.13})$$

so, we can express $\dot{\phi}_c$ in terms of the remaining node variables to eliminate it from the Lagrangian:

$$\dot{\phi}_c = \frac{1}{C_\Sigma} (C_1\dot{\phi}_1 + C_2\dot{\phi}_2 + Q_c) \quad (\text{A.14})$$

where $C_\Sigma = C_1 + C_2 + C_c$ is the sum of all capacitances at ϕ_c . We can multiply out the kinetic energy term of the Lagrangian to get

$$E_C = \frac{1}{2}\dot{\Phi}^T \mathbf{C} \dot{\Phi} = \frac{1}{2}C_\Sigma\dot{\phi}_c^2 + \frac{1}{2}C_1\dot{\phi}_1^2 + \frac{1}{2}C_2\dot{\phi}_2^2 - C_1\dot{\phi}_1\dot{\phi}_c - C_2\dot{\phi}_2\dot{\phi}_c \quad (\text{A.15})$$

Making the relevant substitutions, we find

$$\frac{1}{2}C_\Sigma\dot{\phi}_c^2 = \frac{1}{C_\Sigma} \left[\frac{1}{2}C_1^2\dot{\phi}_1^2 + C_1C_2\dot{\phi}_1\dot{\phi}_2 + \frac{1}{2}C_2^2\dot{\phi}_2^2 + \frac{1}{2}Q_c^2 + C_1\dot{\phi}_1Q_c + C_2\dot{\phi}_2Q_c \right] \quad (\text{A.16})$$

$$C_1\dot{\phi}_1\dot{\phi}_c = \frac{1}{C_\Sigma} \left[C_1^2\dot{\phi}_1 + C_1C_2\dot{\phi}_1\dot{\phi}_2 + C_1\dot{\phi}_1Q_c \right] \quad (\text{A.17})$$

$$C_2\dot{\phi}_2\dot{\phi}_c = \frac{1}{C_\Sigma} \left[C_1C_2\dot{\phi}_1\dot{\phi}_2 + C_2^2\dot{\phi}_2 + C_2\dot{\phi}_2Q_c \right] \quad (\text{A.18})$$

so that

$$E_C = \frac{1}{2}C_1\dot{\phi}^2 + \frac{1}{2}C_2\dot{\phi}_2^2 + \frac{Q_c^2}{2C_\Sigma} - \frac{1}{C_\Sigma} \left[\frac{1}{2}C_1^2\dot{\phi}_1^2 + C_1C_2\dot{\phi}_1\dot{\phi}_2 + \frac{1}{2}C_2^2\dot{\phi}_2^2 \right] \quad (\text{A.19})$$

We then group the corresponding terms and simplify the coefficients

$$\begin{aligned} C_1 - \frac{C_1^2}{C_\Sigma} &= \frac{C_1C_2 + C_1C_c}{C_\Sigma} \\ C_2 - \frac{C_2^2}{C_\Sigma} &= \frac{C_1C_2 + C_2C_c}{C_\Sigma} \end{aligned} \quad (\text{A.20})$$

to get a final kinetic energy term

$$E_C = \frac{1}{2} \frac{C_1C_c + C_1C_2}{C_\Sigma} \dot{\phi}_1^2 + \frac{1}{2} \frac{C_2C_c + C_1C_2}{C_\Sigma} \dot{\phi}_2^2 - \frac{C_1C_2}{C_\Sigma} \dot{\phi}_1\dot{\phi}_2 + \frac{Q_c^2}{2C_\Sigma} \quad (\text{A.21})$$

This corresponds to a capacitance matrix

$$\mathbf{C} = \frac{1}{C_\Sigma} \begin{pmatrix} C_1C_c + C_1C_2 & -C_1C_2 \\ -C_1C_2 & C_2C_c + C_1C_2 \end{pmatrix} \quad (\text{A.22})$$

where the kinetic energy term is now $E_C = \frac{1}{2}\dot{\Phi}^T \mathbf{C} \dot{\Phi} + \frac{Q_c^2}{2C_\Sigma}$.

Inductive T-Coupling

The Lagrangian for the system is

$$L = \frac{1}{2}\dot{\Phi}^T \mathbf{C} \dot{\Phi} - \frac{1}{2}\dot{\Phi}^T \mathbf{L}^{-1} \dot{\Phi} \quad (\text{A.23})$$

where the capacitance and inverse inductance matrices are

$$\mathbf{C} = \begin{pmatrix} C_1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & C_2 \end{pmatrix} \quad \mathbf{L}^{-1} = \begin{pmatrix} L_1^{-1} & -L_1^{-1} & 0 \\ -L_1^{-1} & L_1^{-1} + L_c^{-1} + L_2^{-1} & -L_2^{-1} \\ 0 & -L_2^{-1} & L_2^{-1} \end{pmatrix} \quad (\text{A.24})$$

Here we see that $\dot{\phi}_c$ does not participate in the Lagrangian, so ϕ_c is a frozen variable that cannot change in time. The Euler-Lagrange equations tell us that

$$\frac{\partial L}{\partial \phi_c} = \frac{1}{L_\Sigma}\phi_c - \frac{1}{L_1}\phi_1 - \frac{1}{L_2}\phi_2 = 0 \quad (\text{A.25})$$

so we conclude that ϕ_c is fixed in terms of the other two variables

$$\phi_c = \frac{L_\Sigma}{L_1}\phi_1 + \frac{L_\Sigma}{L_2}\phi_2 \quad (\text{A.26})$$

where L_Σ is now the reciprocal sum defined as $L_\Sigma^{-1} = L_1^{-1} + L_2^{-1} + L_c^{-1}$. The potential energy term of the Lagrangian can be multiplied out as

$$E_L = \frac{1}{2}\mathbf{\Phi}^T \mathbf{L}^{-1} \mathbf{\Phi} = \frac{\phi_1^2}{2L_1} + \frac{\phi_2^2}{2L_2} + \frac{\phi_c^2}{2L_\Sigma} - \frac{1}{L_1}\phi_1\phi_c - \frac{1}{L_2}\phi_2\phi_c \quad (\text{A.27})$$

Next, we group terms involving ϕ_c and make the relevant substitutions

$$\frac{\phi_c^2}{2L_\Sigma} = \frac{L_\Sigma}{2} \left[\frac{1}{L_1^2}\phi_1^2 + \frac{2}{L_1L_2}\phi_1\phi_2 + \frac{1}{L_2^2}\phi_2^2 \right] \quad (\text{A.28})$$

$$\frac{1}{L_1}\phi_1\phi_c = L_\Sigma \left[\frac{1}{L_1^2}\phi_1^2 + \frac{1}{L_1L_2}\phi_1\phi_2 \right] \quad (\text{A.29})$$

$$\frac{1}{L_2}\phi_2\phi_c = L_\Sigma \left[\frac{1}{L_1L_2}\phi_1\phi_2 + \frac{1}{L_2^2}\phi_2^2 \right] \quad (\text{A.30})$$

and we find that

$$\frac{\phi_c^2}{2L_\Sigma} = \frac{1}{2} \left[\frac{1}{L_1}\phi_1\phi_c + \frac{1}{L_2}\phi_2\phi_c \right] = \frac{L_\Sigma}{2} \left[\frac{1}{L_1^2}\phi_1^2 + \frac{2}{L_1L_2}\phi_1\phi_2 + \frac{1}{L_2^2}\phi_2^2 \right] \quad (\text{A.31})$$

We then substitute these terms into Eq. A.27, which results in the potential energy

$$E_L = \frac{1}{2} \frac{1}{L_1}\phi_1^2 + \frac{1}{2} \frac{1}{L_2}\phi_2^2 - \frac{1}{2} \frac{L_\Sigma}{L_1^2}\phi_1^2 - \frac{1}{2} \frac{L_\Sigma}{L_2^2}\phi_2^2 - \frac{L_\Sigma}{L_1L_2}\phi_1\phi_2 \quad (\text{A.32})$$

Finally, we group the corresponding terms and simplify the coefficients to find

$$\begin{aligned} \frac{1}{L_1} - \frac{L_\Sigma}{L_1^2} &= \frac{1}{L_1} \left[1 - \frac{L_\Sigma}{L_1} \right] \\ &= \frac{1}{L_1} \left[\frac{L_1L_2 + L_1L_c + L_2L_c - L_2L_c}{L_1L_2 + L_1L_c + L_2L_c} \right] \\ &= \frac{L_2 + L_c}{L_1L_2 + L_1L_c + L_2L_c} \end{aligned} \quad (\text{A.33})$$

and similarly

$$\frac{1}{L_2} - \frac{L_\Sigma}{L_2^2} = \frac{L_1 + L_c}{L_1L_2 + L_1L_c + L_2L_c} \quad (\text{A.34})$$

The resulting potential energy is

$$E_L = \frac{1}{2} \frac{1}{L_1L_2 + L_1L_c + L_2L_c} \left[(L_2 + L_c)\phi_1^2 - 2L_c\phi_1\phi_2 + (L_1 + L_c)\phi_2^2 \right] \quad (\text{A.35})$$

which corresponds to an inverse inductance matrix

$$\mathbf{L}^{-1} = \frac{1}{L_1L_2 + L_1L_c + L_2L_c} \begin{pmatrix} L_2 + L_c & -L_c \\ -L_c & L_1 + L_c \end{pmatrix} \quad (\text{A.36})$$

From here it is relatively straightforward to check that

$$\mathbf{L} = \begin{pmatrix} L_1 + L_c & L_c \\ L_c & L_2 + L_c \end{pmatrix} \quad (\text{A.37})$$

and we notice that this looks very similar to the capacitance matrix $\mathbf{C}$ for the capacitive π -circuit, except the off-diagonal terms have opposite sign.

Coupling Efficiency for a Floating and Grounded LC Oscillator

Here we will derive the coupling efficiency between a floating and grounded LC oscillator. In the circuit (Fig. 2.7(a)), we see that there are three nodes, with fluxes $\phi_{a,1}$, $\phi_{a,2}$, and ϕ_b , which results in the capacitance matrix

$$\mathbf{C} = \begin{pmatrix} C_g + C_a + C_1 & -C_a & -C_1 \\ -C_a & C_g + C_a + C_2 & -C_2 \\ -C_1 & -C_2 & C_b + C_1 + C_2 \end{pmatrix} \quad (\text{A.38})$$

To determine the effective coupling between the two modes, we must make a basis transformation, defined by

$$\begin{pmatrix} \phi_{a,+} \\ \phi_{a,-} \\ \phi_b \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & 1 & 0 \\ 1 & -1 & 0 \\ 0 & 0 & 1 \end{pmatrix}}_{\mathbf{U}} \begin{pmatrix} \phi_{a,1} \\ \phi_{a,2} \\ \phi_b \end{pmatrix} \quad (\text{A.39})$$

so that the effective capacitance matrix becomes $\tilde{\mathbf{C}} = \mathbf{U}^{-1}\mathbf{C}\mathbf{U}^{-1}$. Then, the resulting capacitance matrix becomes:

$$\tilde{\mathbf{C}} = \begin{pmatrix} \frac{1}{2}(C_g + C_+) & \frac{1}{2}C_- & -C_+ \\ \frac{1}{2}C_- & C'_a + \frac{1}{2}C_+ & -C_- \\ -C_+ & -C_- & C_b + 2C_+ \end{pmatrix} \quad (\text{A.40})$$

where we define $C_{\pm} = (C_1 \pm C_2)/2$, and $C'_a = C_a + \frac{1}{2}C_g$. The coupling efficiency is then given by

$$K_{gf} = \frac{\tilde{\mathbf{C}}_{23}^{-1}}{\sqrt{\tilde{\mathbf{C}}_{22}^{-1}\tilde{\mathbf{C}}_{33}^{-1}}} \quad (\text{A.41})$$

The matrix elements can be computed as

$$\det(\tilde{\mathbf{C}})\tilde{\mathbf{C}}_{22}^{-1} = \frac{1}{2}(C_g + C_+)(C_b + 2C_+) - C_+^2 \quad (\text{A.42a})$$

$$= \frac{1}{2}[C_b(C_g + C_+) + 2C_g C_+] \quad (\text{A.42b})$$

$$\det(\tilde{\mathbf{C}})\tilde{\mathbf{C}}_{33}^{-1} = \frac{1}{2}(C_g + C_+) \left(C'_a + \frac{1}{2}C_+ \right) - \frac{1}{4}C_+^2 \quad (\text{A.42c})$$

$$= \frac{1}{2} \left[C'_a(C_g + C_+) + \frac{1}{2}C_g C_+ \right] \quad (\text{A.42d})$$

$$\det(\tilde{\mathbf{C}})\tilde{\mathbf{C}}_{23}^{-1} = -\frac{1}{2}C_+ C_- + \frac{1}{2}(C_g + C_+)C_- = \frac{1}{2}C_g C_- \quad (\text{A.42e})$$

Finally, we can substitute these expressions into Eq. A.41 to obtain

$$K_{gf} = \frac{C_g C_-}{\sqrt{(C_b(C_g + C_+) + 2C_g C_+) (C'_a(C_g + C_+) + \frac{1}{2}C_g C_+)}} \quad (\text{A.43a})$$

$$= \frac{\gamma C_-}{\sqrt{(C'_a + \frac{1}{2}\gamma C_+) (C_b + 2\gamma C_+)}} \quad (\text{A.43b})$$

where we define $\gamma = \frac{C_g}{C_g + C_+}$.

Coupling Efficiency for two Floating LC Oscillators

We will now derive the coupling efficiency between two floating LC oscillators. In the circuit (Fig. 2.7(b)), we see that there are now four nodes, with fluxes $\phi_{a,1}$, $\phi_{a,2}$, $\phi_{b,1}$, and $\phi_{b,2}$ which results in the capacitance matrix:

$$\mathbf{C} = \begin{pmatrix} C_a + C_g + C_{11} + C_{22} & -C_a & -C_{11} & -C_{12} \\ -C_a & C_a + C_g + C_{21} + C_{22} & -C_{21} & -C_{22} \\ -C_{11} & -C_{21} & C_b + C_g + C_{11} + C_{21} & -C_b \\ -C_{12} & -C_{22} & -C_b & C_b + C_g + C_{12} + C_{22} \end{pmatrix} \quad (\text{A.44})$$

We again make a basis transformation to determine the effective coupling between the two modes, which is now given by

$$\begin{pmatrix} \phi_{a,+} \\ \phi_{a,-} \\ \phi_{b,+} \\ \phi_{b,-} \end{pmatrix} = \underbrace{\begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & -1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & -1 \end{pmatrix}}_{\mathbf{U}} \begin{pmatrix} \phi_{a,1} \\ \phi_{a,2} \\ \phi_{b,1} \\ \phi_{b,2} \end{pmatrix} \quad (\text{A.45})$$

so that the effective capacitance matrix becomes $\tilde{\mathbf{C}} = \mathbf{U}^{-1} \mathbf{C} \mathbf{U}^{-1}$. Now, the effective capacitance matrix in the new basis is:

$$\tilde{\mathbf{C}} = \begin{pmatrix} \frac{1}{2}(C_g + C_\Sigma) & \frac{1}{2}(C_\alpha + C_\beta) & -\frac{1}{2}C_\Sigma & -\frac{1}{2}(C_\alpha - C_\beta) \\ \frac{1}{2}(C_\alpha + C_\beta) & C'_a + \frac{1}{2}C_\Sigma & -\frac{1}{2}(C_\alpha + C_\beta) & -\frac{1}{2}C_\Delta \\ -\frac{1}{2}C_\Sigma & -\frac{1}{2}(C_\alpha + C_\beta) & \frac{1}{2}(C_g + C_\Sigma) & \frac{1}{2}(C_\alpha - C_\beta) \\ -\frac{1}{2}(C_\alpha - C_\beta) & -\frac{1}{2}C_\Delta & \frac{1}{2}(C_\alpha - C_\beta) & C'_b + \frac{1}{2}C_\Sigma \end{pmatrix} \quad (\text{A.46})$$

Here we have defined the new variables:

$$C_\Sigma = \frac{1}{2}(C_{11} + C_{22} + C_{12} + C_{21}) \quad (\text{A.47a})$$

$$C_\Delta = \frac{1}{2}(C_{11} + C_{22} - C_{12} - C_{21}) \quad (\text{A.47b})$$

$$C_\alpha = \frac{1}{2}(C_{11} - C_{22}) \quad (\text{A.47c})$$

$$C_\beta = \frac{1}{2}(C_{12} - C_{21}) \quad (\text{A.47d})$$

$$C'_i = C_i + \frac{1}{2}C_g \quad (\text{A.47e})$$

The coupling efficiency is now given by

$$K_{ff} = \frac{\tilde{\mathbf{C}}_{24}^{-1}}{\sqrt{\tilde{\mathbf{C}}_{22}^{-1} \tilde{\mathbf{C}}_{44}^{-1}}} \quad (\text{A.48})$$

and the matrix elements are now given by

$$\det(\tilde{\mathbf{C}}) \tilde{\mathbf{C}}_{22}^{-1} = C_g \left[\frac{1}{2} \left(\frac{1}{2} C_g + C_\Sigma \right) \left(C'_b + \frac{1}{2} C_\Sigma \right) - \frac{1}{4} (C_\alpha - C_\beta)^2 \right] \quad (\text{A.49a})$$

$$\det(\tilde{\mathbf{C}}) \tilde{\mathbf{C}}_{44}^{-1} = C_g \left[\frac{1}{2} \left(\frac{1}{2} C_g + C_\Sigma \right) \left(C'_a + \frac{1}{2} C_\Sigma \right) - \frac{1}{4} (C_\alpha + C_\beta)^2 \right] \quad (\text{A.49b})$$

$$\det(\tilde{\mathbf{C}}) \tilde{\mathbf{C}}_{24}^{-1} = \frac{1}{4} C_g \left[C_\Delta \left(\frac{1}{2} C_g + C_\Sigma \right) - (C_\alpha^2 - C_\beta^2) \right] \quad (\text{A.49c})$$

We then substitute these expressions into Eq. A.48 to obtain the coupling efficiency

$$K_{ff} = \frac{1}{2} \frac{C_\Delta - \frac{C_\alpha^2 - C_\beta^2}{C_\Sigma + \frac{1}{2} C_g}}{\sqrt{\left(C'_a + \frac{1}{2} C_\Sigma - \frac{1}{2} \frac{(C_\alpha + C_\beta)^2}{C_\Sigma + \frac{1}{2} C_g} \right) \left(C'_b + \frac{1}{2} C_\Sigma - \frac{1}{2} \frac{(C_\alpha - C_\beta)^2}{C_\Sigma + \frac{1}{2} C_g} \right)}} \quad (\text{A.50})$$