Abstract

Microwave Beamsplitters for Oscillator-Based Quantum Information Processing

Stijn J. de Graaf

2025

A single high-Q harmonic oscillator with a fixed 'dispersive' coupling to an ancillary qubit provides a remarkably hardware-efficient platform for a wide range of quantum technologies, capable of acting as a dark matter detector, a simulator of quantum chemistry or a quantum memory with a lifetime longer than its underlying components. The strength of this platform lies in the linearity and favorable decoherence rates of the high-Q oscillator mode. The question then arises: how can we scale this oscillator-based platform to practically useful sizes without compromising a) the speed of operations or b) the properties that make oscillators an attractive platform in the first place? The addition of a tunable oscillator-oscillator coupling, equivalent to an optical beamsplitter, has extended the power of this platform to enable multi-mode entanglement, a key element for quantum computation, but until now, implementations have been limited to low interaction strengths and introduced unwanted oscillator nonlinearity. Inspired by advances in parametric amplification, we show how a three-wave mixing element solves this challenge by acting as a switch, with beamsplitter interaction strengths exceeding those of the dispersive coupling when turned on, and the ability to fully decouple the modes when turned off. We then demonstrate how this regime unlocks a powerful new toolbox of high-fidelity multimode operations which are the analogs of established single-mode control techniques. In particular, we show how these techniques can be leveraged to perform a mid-circuit erasure check, the vital building block for a newly-proposed quantum computer made out of superconducting cavity dual-rail qubits.

Microwave Beamsplitters for Oscillator-Based Quantum Information Processing

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

by
Stijn J. de Graaf

Dissertation Director: Robert J. Schoelkopf

May 2025

Contents

List of Figures

List of Tables

Nomenclature

Abbreviations

ATS	Asymmetrically Threaded SQUID
cSWAP	controlled SWAP
eSWAP	exponential SWAP
GKP	Gottesman-Kitaev-Preskill
JP	Joint Parity
JPNS	Joint-Photon-Number Selective
MCED	Mid-circuit erasure detection
MIST	Measurement Induced State Transitions
qRAM	quantum Random Access Memory
SNAIL	Superconducting Nonlinear Asymmetric Inductive eLement
SNAP	Selective Number Arbitrary Phase
SQUID	Superconducting QUantum Interference Device

Erasure qubits

$p_{\mathsf{leakage}}^{(\mathsf{induced})}$	Leakage probability per MCED above and beyond underlying cavity decoherence
$p_{\mathsf{Pauli}}^{(\mathsf{induced})}$	Pauli error probability per MCED above and beyond underlying cavity decoherence
$p_{\mathsf{leakage}}^{(\mathsf{intrinsic})}$	Leakage probability per MCED due to underlying cavity decoherence
$p_{\mathsf{Pauli}}^{(\mathsf{intrinsic})}$	Pauli error probability per MCED due to underlying cavity decoherence
$p_{\mathsf{erasure}}^{(\mathsf{MCED})}$	Erasure probability per MCED
$p_{\mathsf{leakage}}^{(\mathsf{MCED})}$	Leakage probability per MCED
$p_{\mathsf{Pauli}}^{(\mathsf{MCED})}$	Pauli error probability per MCED

$p_{\mathsf{Pauli}}^{(\mathsf{transmon})}$	Pauli error probability per MCED due to transmon errors
p_{erasure}	Erasure probability per error correction cycle
p_{FN}	False negative probability per MCED
p_{FP}	False positive probability per MCED
p_{leakage}	Leakage probability per error correction cycle
p_{miss}	Missed leakage error probability per MCED
p_{Pauli}	Pauli error probability per error correction cycle
R_{e}	Erasure fraction, $p_{\mathsf{erasure}}/(p_{\mathsf{erasure}} + p_{\mathsf{Pauli}})$

Acknowledgments

Reaching the end of this PhD comes as a bittersweet moment. The pride I feel in completing this work owes everything to the incredible people here at Yale, with whom every conversation taught me something new. At the same time, my sadness in moving on from this place is down to the immense warmth and friendliness of those same people. There are many, both here at Yale and beyond, whom I wish to acknowledge for their important roles in making this moment possible.

First and foremost, I owe a massive thank you to my advisor, ROBERT SCHOELKOPF, a scientist whose relentless appetite for physics is an inspiration. Rob has managed to set the highest standards while remaining incredibly kind-spirited, a trait that is reflected in the culture of his lab. He has always emphasized the importance of communicating the significance of ones work, and his lessons in this regard will stand me in good stead wherever my path leads me. I am also incredibly grateful to my other commitee members, MICHEL DEVORET and STEVEN GIRVIN, whose friendliness and passion for learning shines through in everything they do and whose office doors are always open to curious students. My research has been greatly influenced by the work of YVONNE GAO, both before my arrival at Yale and now also with her own lab at NUS. I am therefore incredibly thankful for her offer to serve as the external reader for this thesis. I also have to thank SHRUTI PURI for her support, for her advice on everything erasure-qubits and for fostering such a collaborative environment between the YQI and Becton.

When I first started in the lab, I had the immense fortune of finding myself under the wing of BEN CHAPMAN. Working together throughout my first four years at Yale, he taught me everything there was to know about setting up and running a microwave experiment. More importantly, his boundless energy and positivity, flying in the face of many grad school stereotypes, made it a joy to come to work every day. Since his departure, I have been lucky to continue that experience working closely with SOPHIA XUE, whose excitement to understand things deeply and keen eye for detail elevates any project she is involved in. I am very excited to see where her experiments take her next. More recently, I have also had the great pleasure of working on a next generation of experiments with JENNY GIAMPALMO and PATRICK WINKEL, as well as with TAKAHIRO TSUNODA and BILLY KALFUS, which has made this last period at Yale very exciting.

Many other members of Qulab, RSL and the YQI have had an important role in this journey, in particular: CHRIS WANG, who guided me through my first SI project and taught me how to operate a dilution fridge, JAMES TEOH, whose theory work on dual-rail qubits gave me an amazing list of experiments to try out, ANIKET MAITI, JOHN GARMON and YAO LU, whose parallel work on beamsplitters taught me a huge amount, ALEX READ, JACOB CURTIS, ALEC EICKBUSCH, ANDY DING, AKSHAY KOOTTANDAVIDA, SANDRO MIANO, KATIE CHANG, ROMAN BASKOV, XU XIAO, LELE XU and YAXING ZHANG, whom I had the pleasure of collaborating with, as well as my other classmates in Becton, VIDUL JOSHI, WEI DAI, and NAIJUN JIN, who have remained good friends throughout. My deskmates, ZHIXIN WANG and SUMERU HAZRA, were generous both in sharing ideas they thought might help my experiments and for being a great sounding board for my own ideas. I owe particular thanks to NICK FRATTINI who was incredibly willing to answer questions about his invention, the SNAIL, which played such an important role in this thesis.

The thesis itself would probably not have been written on time were it not for the support of LUIGI FRUNZIO, who provided feedback on every page. He has been a guardian angel throughout my time at Yale, helping not just with the scientific journey but the life journey too. Others who have saved my life several times include NUCH GRAVES, THERESA EVANGELISTE, ALEX BOZZI, MARIA RAO, FLORIAN CARLE, HEATHER ONG, LIYU QIN, JESSICA TEFFT, GISELLE MAILLET, RACQUEL MILLER, and AMY BADNER.

My journey to this point did not begin when I arrived at Yale. I especially want to thank several amazing Mathematics teachers throughout my school years, in particular JOHN BARLOW, JON STONE and KARTHIK TADINADA, for making me fall in love with solving difficult problems, something from which I can draw a direct line to where I am today.

I also need to thank many friends, whom I can't all name here, for helping me enjoy life outside of the lab here in New Haven. In particular, TOM RIMMER, for exploring the US with me in my first year, RAHUL DEV and CHARLIE COLTHORPE, for coming to visit me all the way from London, PATRICK HULIN and MAGGIE PAVLOVICH, for being the world's best hosts, TALIA BOYLAN, with whom I rode out the depths of the pandemic, as well as YANNI DAHMANI, JOE LEVIN, ZEYU KUANG, DOMINIC TREVISAN, BÜLENT KANMAZ, ANDREAS CARLSSON, LINDA ZUCKERMAN, NICK MILLER, and the Yale Cycling team.

Finally and most importantly, I need to thank those closest to me. My partner YASMIN, undoubtedly the greatest addition to my life since arriving in America, my brothers FRISO and DOUWE, who never ever fail to cheer me up, and my parents PETER AND ELLEN, without whom I would never have reached the start-line, let alone the finish, of this PhD and who have always been just a phone call away, no matter the timezone.

Chapter 1

Introduction

The power of quantum computers to tackle certain challenging problems more efficiently than their classical counterparts relies on their use of *superposition* and *entanglement*, properties that are not accessible to a classical computer. A neat way to demonstrate the effect of the first of these, superposition, is by using a pair of beamsplitters, in a version of Young's famous double-slit experiment (Young, 1804). This setup, known as a Mach-Zehnder interferometer (Mach, 1892; Zehnder, 1891), is shown in Fig. 1.1(a). By shining a laser onto one face of a 50/50 beamsplitter, the beam is divided into two, with half of the light continuing onwards and the other half being reflected. We can then use mirrors to recombine this light in a second 50/50 beamsplitter, such that it emerges out of its two different faces, which we monitor. This setup can reveal two seemingly contradictory observations:

1. By altering the length of one of the two paths connecting the beamsplitters, the intensity of outgoing light displays constructive and destructive interference, with the intensity of light emerging from each face oscillating as a function of the relative path difference (Fig. 1.1(b)). This interference observation is natural if one treats light as a wave.

2. By decreasing the intensity of the light source, the intensity of the emerging light can be made to decrease until eventually light is detected in discrete bursts (Fig. 1.1(c)). This discrete-burst observation is natural if one treats light as consisting of particles, or

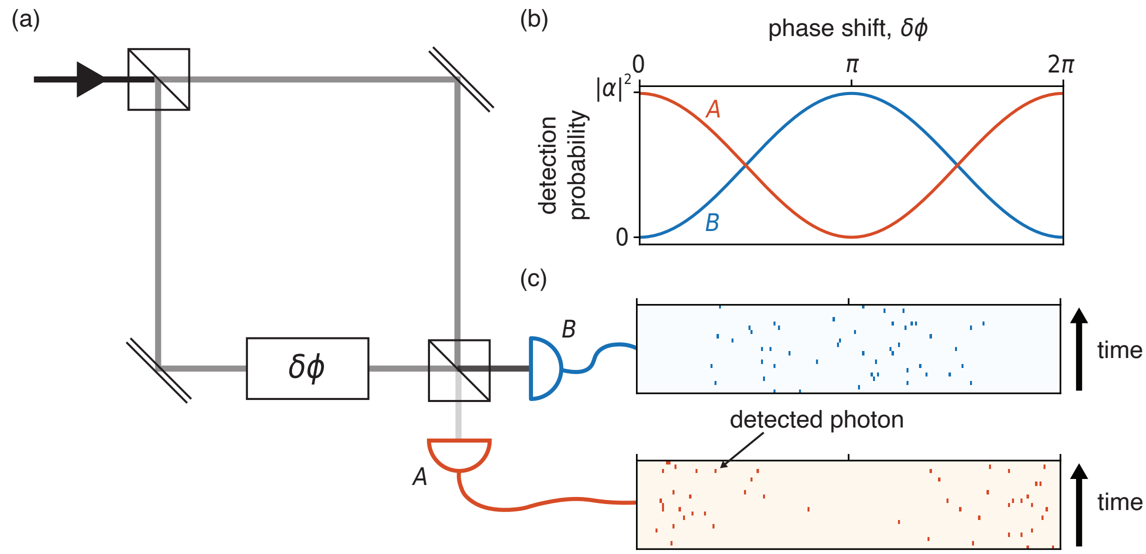

Figure 1.1: **Mach-Zehnder interferometer.** (a) A schematic of the interferometer, in which a coherent beam of light is split by a 50/50 beamsplitter before being recombined in another beamsplitter, whose outputs are monitored with photon detectors A and B. One path of the interferometer contains an optical phase shifter which can add an additional phase $\delta\phi$ relative to the other path. (b) The averaged probability of measuring a photon at detector A or B as a function of the phase shift, $\delta\phi$. For concreteness, $|\alpha|^2 \ll 1$ describes the average photon number (per spatiotemporal mode) in the in-going beam. (c) An example (simulated) detector trace across the full range of $\delta\phi$. Even as the light arrives in individual bursts, the long-time averaged statistics will show constructive and destructive interference.

photons, of light.

Surprisingly, both of these effects can be observed in the same experiment (although not in a single shot). While the individual photons may appear one-by-one, over time their averaged statistics will build up to show the interference patterns in Fig. 1.1(b).

The resolution to this particle-wave problem is that *both* can be viewed as valid complementary descriptions: we can treat the light as composed of individual particles, but their probability amplitudes of traveling along each path of the interferometer can be described by a *superposition*, i.e. a linear combination of both paths, possessing a relative *phase* just as a wave would. In this picture, interference effects in the measured destinations of individual photons become possible. While the description of light as a wave possessing a phase had long been discussed (Shapiro, 1989), what underpinned the first quantum revolution of the early 20$^{\text{th}}$

century was the understanding that matter (at sufficiently small action scales) also obeyed the same wave-particle duality (De Broglie, 1923). The same interference effects first seen in light have subsequently been observed in systems of increasingly large mass, with electrons (Davisson and Germer, 1927), atoms (Estermann and Stern, 1930), molecules (Arndt *et al.*, 1999), and now too, in electromagnetic circuits (Bertet *et al.*, 2001).

Whereas the first quantum revolution provided a unified description of these wave-like and particle-like properties, the second quantum revolution (Dowling and Milburn, 2003) of the late 20$^{\text{th}}$ and early 21$^{\text{st}}$ centuries has established that with sufficient *control* over these superpositions, we may be able to perform useful tasks. Of particular interest is the prospect of performing quantum computation (Nielsen and Chuang, 2010), where information is stored in the superposition states of a large quantum system (typically imagined as an array of quantum bits) and manipulated such that interference effects lead to a final measurement outcome that encodes the result of the computation. Shor's 1994 algorithm for efficient prime factorization (Shor, 1994) remains the foremost example (among several prominent ones (Grover, 1996; Harrow *et al.*, 2009)) of a quantum computer providing a *theoretical* advantage in an area where even the world's most powerful classical computers are inadequate. However, in order to perform useful computation, control over quantum superpositions is not by itself sufficient – we specifically require the ability to generate *entanglement* (Josza and Linden, 2003), a property of multipartite quantum systems where the state of their constituent parts cannot be described independently of one another.

While the example of the Mach-Zehnder interferometer showed the ability to control the superposition state of outgoing photons by carefully controlling the phase shift $\delta\phi$, these linear optical components alone (i.e. beamsplitters, phase shifters and mirrors) cannot generate entanglement (Kok *et al.*, 2007). In order to do so, we require some *non*-linear component (Lloyd, 1992)[1]. A modified version of our Mach-Zehnder apparatus (Fig. 1.2) however, forms the basis of one of the earliest proposals for constructing a 'simple' quantum computer (Chuang and Yamamoto, 1995). In this proposal, information is encoded in which of two paths a photon

[1]This non-linear component can also come in the form of single-photon sources and detectors, such as in the KLM protocol (Knill *et al.*, 2001).

takes, which is hence named a 'dual-rail' qubit (quantum bit). Combining the aforementioned beamsplitters with a 'Kerr' nonlinear medium (Boyd, 2008; Kerr, 1875), which takes the place of the phase shifter, changing the phase of light in one rail conditioned on the presence of a photon in another, allows adjacent qubits to interact with one another (Milburn, 1989), which in turn allows for the generation of quantum entanglement. Adding also the ability to generate and detect single photons provides the control necessary for a 'universal' quantum computer, capable of performing any quantum computation. A handful of optical elements therefore provides a blueprint for enabling quantum advantage over classical methods.

The catch though, with this and with all other proposals, is that while we are manipulating the quantum system, so too is its environment. Qubits, unlike their classical counterparts, are inherently analog systems and are therefore subject to noise which can degrade the fidelity of a computation and whose effect accumulates over time (Nielsen and Chuang, 2010). Quantum computers are therefore always in a race against the clock – operations must be completed faster than noise can accumulate, and the speed limit for these operations is set by the strength of the nonlinear interaction (i.e. how fast the photon-number-dependent phase accumulates). Unfortunately for this particular proposal, a source of strong, low-loss, optical nonlinearity has remained elusive (O'Brien, 2007).

As established earlier however, the quantum mechanical principles from which quantum technologies derive their power do not just apply in the optical domain. A leading approach is to instead build quantum computers out of superconducting microwave circuits, where a strong dissipationless source of Kerr nonlinearity *does* exist, in the form of the Josephson junction (Josephson, 1962). The superconducting analog of the nonlinear optical processor is based on bosonic superconducting qubits (Joshi *et al.*, 2021) where, rather than storing information in propagating optical photons (at ~ 200 THz), we store it in standing-wave microwave photons in a resonator (at ~ 5 GHz). These oscillators are by themselves purely linear, but coupling them to Josephson junction-based circuits (typically a 'transmon' circuit; Koch *et al.*, 2007) provides them with the requisite nonlinearity to perform universal control of their state (Heeres *et al.*, 2015; Krastanov *et al.*, 2015). Long-lived superconducting resonators coupled to nonlinear cir-

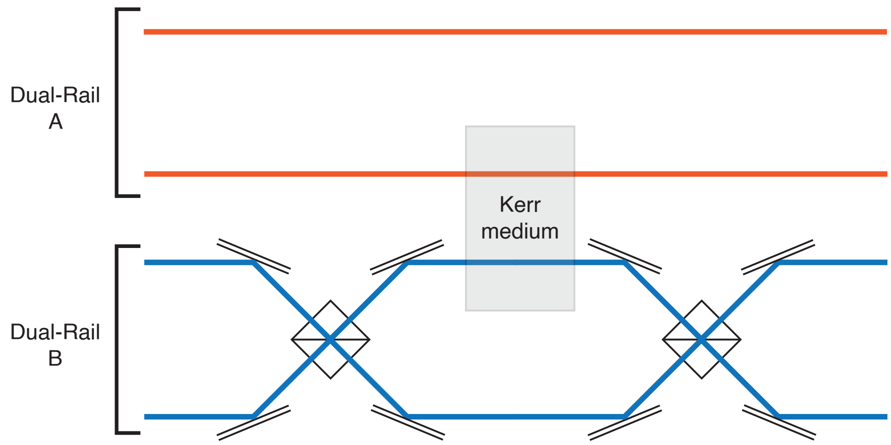

Figure 1.2: **Dual-rail CNOT gate based on optical cSWAP.** Two dual-rail qubits, each consisting of a single photon whose path along one of two fibers encodes 1 bit of information (Chuang and Yamamoto, 1995). The replacement of the phase shifter in the Mach-Zehnder interferometer with a Kerr medium (Kerr, 1875), which shifts the phase conditioned on the presence of a photon in the adjacent rail, forms a cSWAP (or Fredkin) gate for optical states, exchanging the states in the blue (B) rails, conditioned on the presence of a photon in the lower orange (A) rail (Milburn, 1989). When applied to dual-rail qubits, this serves as a CNOT gate, flipping the state of qubit B, conditioned on the state of qubit A.

cuits (where the nonlinear interaction strength greatly exceeds the dissipation rate) have been used to generate non-classical states of light including Fock (Hofheinz *et al.*, 2008), Schrödinger-cat (Kirchmair *et al.*, 2013), cubic-phase (Eriksson *et al.*, 2024) and Gottesman-Kitaev-Preskill (GKP) (Campagne-Ibarcq *et al.*, 2020) states, which in turn have enabled searches for dark matter (Backes *et al.*, 2021; Brubaker *et al.*, 2017; Dixit *et al.*, 2021), and demonstrations of 'beyond-breakeven' quantum error correction (Brock *et al.*, 2024; Ni *et al.*, 2023; Ofek *et al.*, 2016; Sivak *et al.*, 2023).

These proof-of-principle experiments demonstrate the power of bosonic superconducting circuits to be competitive in the race against the clock. However, the power of quantum control scales with the size of system (related to the number of different states we can access) (Cross *et al.*, 2019). In order to extend these achievements to practically useful system sizes, we need a means of controlling multiple resonators, while also isolating them from each other when desired. In other words, we need a switch. A powerful way to do so comes from looking to our original Mach-Zehnder interferometer for inspiration. The work of Zakka-Bajjani *et al.* (2011) showed that the microwave analog of a beamsplitter, which in this case swaps photons between a pair of resonator modes, can be generated by taking a similar Kerr-nonlinear transmon circuit[2], capacitively coupling it to both resonators, and applying strong microwave drives to it. This same idea was subsequently extended to beamsplitter interactions between standing-wave and propagating microwave photons (Flurin *et al.*, 2015; Pfaff *et al.*, 2017), between photons in long- and short-lived resonators (Sirois *et al.*, 2015), and crucially, between photons in pairs of long-lived resonators (Gao *et al.*, 2018).

Access to beamsplitters in the microwave domain allows us to imagine a multi-purpose system comprising a network of microwave resonators connected by tunable beamsplitters, with the required nonlinearity provided by Kerr nonlinear circuits coupled to individual resonators (Teoh *et al.*, 2023). Such a network provides a framework for a wide class of powerful operations based on their optical analogs. By themselves, beamsplitters provide a highly-efficient way of routing quantum information around a network to emulate circuit topologies more complex than the underlying physical substrate. Combining the beamsplitter with a Kerr nonlinearity coupled to a single mode unlocks more powerful operations, such as the controlled-SWAP, a key enabling operation for quantum random access memory (qRAM) (Hann, 2021; Weiss *et al.*, 2024) and exponential-SWAP (Gao *et al.*, 2019), a gate that entangles qubits in two resonators independently of the choice of encoding (Lau and Plenio, 2016). As we shall see later in this thesis, this network also provides all the necessary operations (single- and two-qubit gates, as well as measurements) for a quantum computer consisting of *microwave* dual-rail qubits (Teoh *et al.*, 2023), a promising platform for hardware-efficient error-corrected quantum computing.

In choosing a quantum computing platform, a quantum engineer must pick their poison. In the case of the nonlinear optics approach, the challenge is presented by developing a strong Kerr nonlinearity. In bosonic systems, it is rather the beamsplitter that presents the greater

[2]For completeness, in this particular experiment, the specific circuit was a superconducting quantum interference device, or SQUID, where the microwave drives modulated the magnetic flux through a central loop in the SQUID. In the Gao *et al.* (2018) experiment implementing a beamsplitter between long-lived resonators, the coupling element was a standard transmon circuit where the microwave drives modulated the charge in the circuit. Both ideas will be discussed in more detail in Chapter 3.

technical challenge. As mentioned before, operating a quantum computer is a race against the clock, and these beamsplitters based on Kerr-nonlinear transmon couplers are relatively slow, operating an order of magnitude slower than single-resonator operations. Further increasing the strength of the microwave drive to increase the speed does so at the cost of increased noise, thereby negating any benefit. As such, while the addition of a microwave beamsplitter enables a powerful set of operations, its relatively slow speed limits their practical utility.

This thesis sets out to resolve this issue, replacing the existing Kerr-nonlinear coupling circuit used to generate the beamsplitter interaction with a Kerr-free circuit, inspired by the success of a similar approach in improving the dynamic range of parametric amplifiers (Frattini, 2021). This could provide bosonic networks with not just a switch, but a **'good'** switch – one with an order-of-magnitude higher beamsplitter rate and the ability to completely isolate nodes from each other when desired. This has the power to unlock a new regime in which the beamsplitter interaction strength matches that of the nonlinear 'dispersive' interaction - a regime in which we can devise new high-fidelity multi-mode schemes, paving the way to large-scale quantum technologies based on many bosonic modes.

1.1 Outline of this thesis

The overarching theme of this thesis is the interplay between two elements of oscillator control: the dispersive shift (mediated by a Kerr nonlinearity) and the tunable beamsplitter. In particular, it concerns how we can engineer a beamsplitter interaction whose strength is comparable to that of the dispersive interaction and, since this interaction strength sets the timescale for other operations, the new capabilities that are enabled in this regime. Chapter 2 begins by introducing these two main characters, their physical origins, and the existing toolbox of operations they enable. The single-oscillator dispersive operations demonstrated in that chapter will be especially relevant for later chapters where I, one-by-one, generate the two-oscillator analogs of these operations via the simultaneous application of a strong beamsplitter interaction. While many of the lessons contained in this thesis apply to any system composed of long-lived harmonic oscillators and nonlinear elements, in that chapter I will also introduce the particular experimental platform used, namely 3D superconducting cavities and transmon qubits.

In Chapter 3, I will hone in on the underlying factors that have limited the speed and noise performance of tunable microwave beamsplitters based on Kerr-nonlinear couplers, and propose a solution based on a Kerr-free nonlinear element called the SNAIL (superconducting nonlinear asymmetric inductive element) (Frattini *et al.*, 2017). The SNAIL has its origins as a means of increasing the dynamic range of parametric amplifiers (Frattini *et al.*, 2018; Sivak *et al.*, 2019, 2020), and so I discuss the similarities and differences between this problem and the problem of improving beamsplitters between high-Q oscillators. This also provides an opportunity to compare the merits of this solution to other more recent approaches using, for example, magnetic-flux-pumped superconducting circuits (Lu *et al.*, 2023).

In Chapter 4, I will put these ideas into practice, showing how a system of two microwave resonators coupled by a SNAIL can enable beamsplitter rates in excess of typical dispersive shifts, without introducing additional noise or any always-on nonlinear coupling between the resonators. I describe the considerations when engineering a device to achieve this, paying particular attention to the delivery of magnetic flux, the delivery of a strong charge drive, and the circuit parameters of the SNAIL. I will also lay out the techniques used to characterize the properties of the SNAIL element itself, as well as the performance of the tunable beamsplitter.

The remaining chapters describe the power that can be derived from combining dispersive and beamsplitter interactions, both in theory and experiment. In Chapter 5, I will consider operations where the beamsplitter and dispersive terms are used alternately, including a new technique for performing a fast controlled-SWAP, whereby the state of a nonlinear ancilla controls whether or not two resonator states are exchanged, and an error-detected measurement of the joint-photon-number parity in two oscillators. These results provide experimental evidence for the utility of the operator Bloch sphere model (Tsunoda *et al.*, 2023) in constructing operations where the beamsplitter and dispersive interactions strengths are on par with each other. This model, based on the Schwinger angular momentum description of two coupled modes (Schwinger, 1952), extends ideas used in the context quantum optical beamsplitters (Campos *et al.*, 1989) to the regime of superconducting circuits.

In Chapter 6 I will then introduce a complementary framework for understanding operations where the beamsplitter and dispersive terms are applied simultaneously, which we call the joint-photon number-splitting regime. We will see how the spectrum of a nonlinear ancilla coupled to an oscillator can be modified in the presence of a strong beamsplitter such that it depends only on the *combined* photon number in the two coupled oscillators. By mapping the physics of these coupled oscillators onto that of a single spin, I obtain an intuitive model for constructing multi-cavity operations that generalizes to large photon numbers in the oscillators.

Chapter 7 ties all of these previous pieces together by applying this joint-photon number-splitting regime to demonstrate an essential operation for a dual-rail qubit consisting of 3D superconducting cavities (Teoh *et al.*, 2023), namely a mid-circuit erasure check – a way of catching events where a photon is lost from either one of the rails without disturbing the quantum state. Crucially, we will do so with a circuit layout that extends readily to a large surface code (Kitaev, 2003) of dual-rail qubits.

Finally, in Chapter 8 I will discuss next steps, focusing in particular on the questions that remain for scaling up dual-rail qubits, possible approaches for further improving beamsplitter performance and the potential scope of capabilities enabled by the beamsplitter-coupled microwave resonators.

Acknowledgments

Chapters 3 - 5, proposing and implementing the SNAIL beamsplitter and demonstrating the cSWAP primarily relate to Chapman *et al.* (2023), on which Benjamin Chapman and I were co-first authors. While Ben was exclusively responsible for fabrication, and I led the design and simulation of the experimental package and chips, the collection and analysis of the data was performed jointly, with notable assistance from Sophia Xue.

Chapters 6 - 7, describing the joint-photon number-splitting regime and using it to demonstrate a dual-rail mid-circuit erasure check primarily relate to de Graaf *et al.* (2025), on which I was co-first author with Sophia Xue. I was responsible for the theoretical treatment of the

joint-photon number-splitting regime and the analysis of different pulse shapes, while the collection and analysis of data verifying the predictions and demonstrating the erasure check was performed jointly with Sophia.

Chapter 2

Elements of high-Q oscillator control

In the introduction, I described the possibility of manipulating quantum information in large networks of long-lived linear oscillators, whether as a computer operating on encoded qubits (Joshi *et al.*, 2021), as a binary tree enabling quantum random access memory (qRAM) (Hann, 2021; Naik *et al.*, 2017; Weiss *et al.*, 2024), or as an analog simulator of quantum chemistry (Hu *et al.*, 2018; Katz and Monroe, 2023; Owens *et al.*, 2018; Wang *et al.*, 2020). One can construct a network capable of these applications using just two types of Hamiltonian interaction: a tunable beamsplitter interaction between linear oscillators, and a static dispersive interaction that couples linear oscillators to nonlinear ancillary oscillators. Fig. 2.1 shows an example of an oscillator network, as well as a minimal unit-cell in which a central linear oscillator participates in both Hamiltonian interactions. The Hamiltonian of this unit cell, consisting of two linear oscillators (Alice and Bob, described by lowering operators $\hat{a}$ and $\hat{b}$) and a nonlinear oscillator (truncated to its lowest two energy levels, $|g\rangle$ and $|e\rangle$), can ideally be written as:

$$\hat{\mathcal{H}}_{\chi\mathsf{bs}} = \hat{\mathcal{H}}_{\mathsf{bs}} + \hat{\mathcal{H}}_{\mathsf{disp}}, \tag{2.1}$$

where

$$\frac{\hat{\mathcal{H}}_{\text{bs}}}{\hbar} = \frac{g_{\text{bs}}(t)}{2}\hat{a}^\dagger\hat{b} + \frac{g^*_{\text{bs}}(t)}{2}\hat{b}^\dagger\hat{a}, \tag{2.2}$$

$$\frac{\hat{\mathcal{H}}_{\text{disp}}}{\hbar} = \chi\hat{b}^\dagger\hat{b}\,|e\rangle\langle e|\,. \tag{2.3}$$

This unit cell provides a platform for probing the competition between these two Hamiltonian terms, particularly in the regime where the magnitude of the beamsplitter interaction ($g_{\text{bs}}(t)$) becomes comparable to that of the dispersive interaction (χ), and will form the basis of the experiments in this thesis.

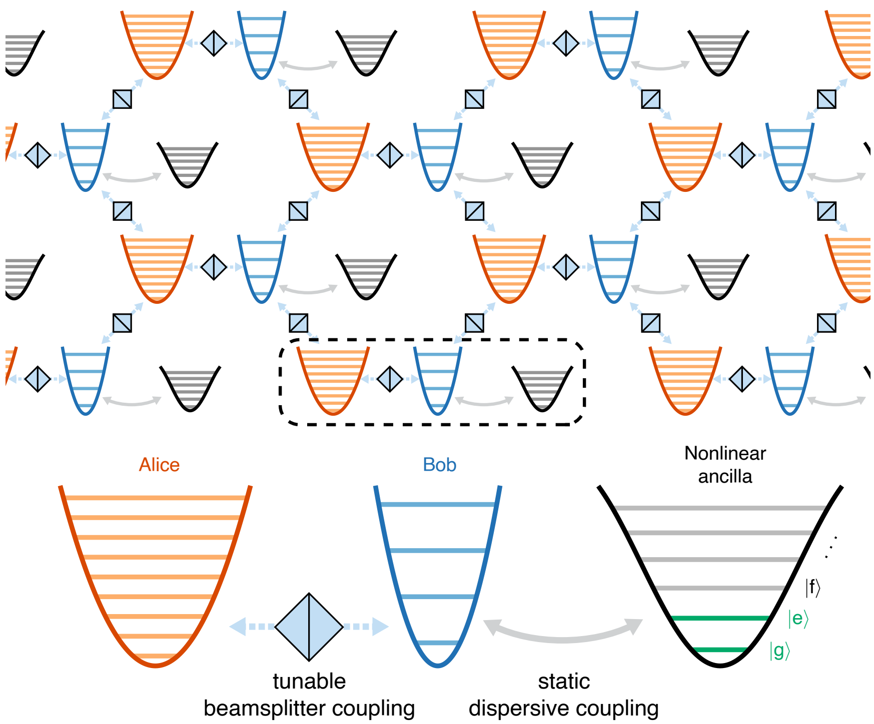

Figure 2.1: **Building blocks of a quantum information processor based on a bosonic network.** (a) Section of a possible 2D tiling of linear oscillators (orange and blue) linked by tunable beamsplitter interactions, with every other oscillator coupled to a nonlinear oscillator (black) via a dispersive interaction. (b) The minimal unit cell of this tiling, consisting of just two linear oscillators (Alice and Bob) and a nonlinear oscillator, whose lowest two energy levels (green) form an ancillary qubit.

The goal of this chapter is to review the capabilities afforded by this unit cell, in the existing regime where $g_{\mathsf{bs}} \ll \chi$. In particular, I shall describe:

1. key properties of the linear and nonlinear superconducting oscillators,

2. how this informs their respective roles, and

3. the toolbox of operations enabled by the combined Hamiltonian, $\hat{\mathcal{H}}$.

Not only will this provide context for what is and isn't possible in the $g_{\mathsf{bs}} \ll \chi$ regime, but the description of established single-oscillator control techniques will lay the groundwork for developing analogous multi-oscillator techniques when $g_{\mathsf{bs}} \gtrsim \chi$.

2.1 A hybrid discrete-continuous variable system

The network we have been describing is an example of a *hybrid* oscillator-qubit system (Liu *et al.*, 2024) consisting both of elements described by continuous-variable (in the case of linear oscillators) and by discrete-variable (in the case of ancillary qubits) quantum degrees of freedom. Here I will very briefly review some key properties of the former, allowing us to more formally motivate the requirements for the beamsplitter and dispersive couplings.

In the context of superconducting circuits, a linear oscillator may be represented as an LC circuit, whose Hamiltonian,

$$\hat{\mathcal{H}}_{\mathsf{osc}} = \frac{\hat{\Phi}^2}{2L} + \frac{\hat{Q}^2}{2C} = \hbar\omega_a \left(\hat{a}^\dagger \hat{a} + \frac{1}{2}\right) \tag{2.4}$$

can be expressed in terms of two **continuous-valued** conjugate variables: the branch flux across the inductor, $\hat{\Phi} = \Phi_{\mathsf{ZPF}} \left(\hat{a} + \hat{a}^\dagger\right)$, and the charge on the capacitor, $\hat{Q} = -iQ_{\mathsf{ZPF}} \left(\hat{a} - \hat{a}^\dagger\right)$, where $\Phi_{\mathsf{ZPF}} = (\hbar^2 L/4C)^{1/4}$ and $Q_{\mathsf{ZPF}} = (\hbar^2 C/4L)^{1/4}$ represent the zero-point fluctuations of each variable when the oscillator is in its ground state (Vool and Devoret, 2017), and $\omega_a = 1/\sqrt{LC}$ is the resonance frequency of the oscillator. These conjugate variables are non-commuting, satisfying $\left[\hat{\Phi}, \hat{Q}\right] = 2i\Phi_{\mathsf{ZPF}} Q_{\mathsf{ZPF}} = i\hbar$. Alternatively, we can normalize these quantities by their zero-point fluctuations ($\times\sqrt{2}$) to obtain the dimensionless flux and charge 'quadratures':

$$\hat{\phi} \equiv \frac{\hat{\Phi}}{\sqrt{2}\Phi_{\text{ZPF}}} = \frac{\hat{a}+\hat{a}^\dagger}{\sqrt{2}}, \tag{2.5}$$

$$\hat{q} \equiv \frac{\hat{Q}}{\sqrt{2}Q_{\text{ZPF}}} = -i\frac{\hat{a}-\hat{a}^\dagger}{\sqrt{2}}, \tag{2.6}$$

which instead satisfy $\left[\hat{\phi}, \hat{q}\right] = i$, where $\hat{\phi}$ and $\hat{q}$ can be interpreted as the real and imaginary components of the $\hat{a}$ operator ($\times\sqrt{2}$)[1] (Leonhardt, 1997).

The state of this oscillator is described by its density matrix, $\hat{\rho}$. While there are many bases in which we could represent $\hat{\rho}$, an especially useful way of describing it is via the Wigner quasiprobability distribution (or 'Wigner function') (Wigner, 1932), a real-valued function in 2-dimensional phase space, spanned by the real continuous variables ϕ and q:

$$W(\phi, q) = \frac{1}{2\pi}\int_{-\infty}^{\infty} \left\langle \phi - \frac{\phi'}{2} \middle| \hat{\rho} \middle| \phi + \frac{\phi'}{2} \right\rangle e^{iq\phi'} d\phi'. \tag{2.7}$$

This function is analogous to a probability distribution in classical phase space, with the marginals providing the probability distributions along ϕ,

$$\int_{-\infty}^{\infty} W(\phi, q) dq = |\psi(\phi)|^2. \tag{2.8}$$

and along q,

$$\int_{-\infty}^{\infty} W(\phi, q) d\phi = |\psi(q)|^2. \tag{2.9}$$

Following from the roles of $\hat{\phi}$ and $\hat{q}$ as real and imaginary components of $\hat{a}$, we will often see the Wigner function written as a function of a single complex quantity $\alpha \equiv (\phi + iq)/\sqrt{2}$ (Cahill and Glauber, 1969). The Wigner function will be used throughout this thesis as a way of visualizing the evolution of oscillator states, with several features that can be used to gain intuition (see

[1]Why the factors of $\sqrt{2}$? This is a convention that allows us to write the commutation relation in its current form. One is free to remove or change the factor of $\sqrt{2}$ but note that then $[\hat{\phi}, \hat{q}] \neq i$.

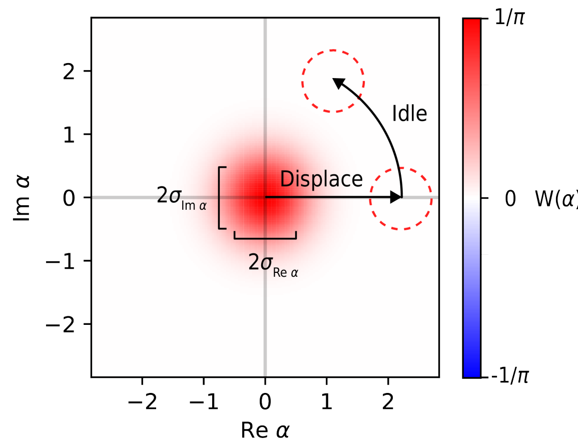

Figure 2.2: **Example Wigner function of a vacuum state.** The Wigner function of an oscillator vacuum state appears as a symmetric 2-dimensional Gaussian centered at the origin, with standard deviations $\sigma_{\text{Re }\alpha} = \sigma_{\text{Im }\alpha} = 1/2$. Displacements (effected by driving the oscillator's charge degree of freedom) translate the Wigner function, while phase space rotations (effected by idling for some duration) rotate it about the origin of phase space.

the excellent book by Leonhardt (1997) for a complete list).

As an example of this intuition, a coherent state $|\alpha_0\rangle$ appears as a symmetric 2-dimensional Gaussian distribution centered at $\alpha = \alpha_0$,

$$W_{|\alpha_0\rangle}(\alpha) = \frac{1}{\pi}\exp\left(-2|\alpha - \alpha_0|^2\right) \tag{2.10}$$

with equal standard deviations along the Re(α) and Im(α) axes, $\sigma_{\text{Re}(\alpha)} = \frac{\phi_{\text{ZPF}}}{\sqrt{2}} = \frac{1}{2}$ and $\sigma_{\text{Im}(\alpha)} = \frac{q_{\text{ZPF}}}{\sqrt{2}} = \frac{1}{2}$ set by the Heisenberg uncertainty relation (Heisenberg, 1927). Applying a charge drive to the LC oscillator ($\hat{\mathcal{H}}_{\text{drive}} \propto \hat{q}$) effects a displacement unitary $D(\delta\alpha)\,|\alpha_0\rangle = |\alpha_0 + \delta\alpha\rangle$, which can simply be viewed as a translation of the Wigner function by $\delta\alpha$. Similarly, idling ($\hat{\mathcal{H}} = \hbar\omega_a \hat{a}^\dagger\hat{a}$) for some duration t enacts the unitary $e^{-i\omega_a t\hat{a}^\dagger\hat{a}}\,|\alpha_0\rangle = \left|\alpha_0 e^{-i\omega_a t}\right\rangle$, corresponding to a 'phase space rotation' of the entire Wigner function about the origin, by an angle $\omega_a t$.

Both displacements and phase space rotations belong to the class of **Gaussian operations**, so-called because they map Gaussian states, whose Wigner function is a multivariate Gaussian distribution[2], to other Gaussian states (Weedbrook *et al.*, 2012). These states, which include

[2]The Wigner quasiprobability distribution can be extended to two (or more) oscillators, where we refer to it as the 'joint Wigner distribution' of the pair (or set) of modes, as we will use later in the thesis.

coherent states ($\sigma_{\text{Re}\,\alpha} = \sigma_{\text{Im}\,\alpha} = 1/2$), thermal states ($\sigma_{\text{Re}\,\alpha} = \sigma_{\text{Im}\,\alpha} > 1/2$) and squeezed states ($\sigma_{\text{Re}\,\alpha} < \sigma_{\text{Im}\,\alpha}$, or for any other pair of perpendicular axes in phase space), can be completely described with a small number of coordinates. For example, a single-mode Gaussian state can be completely specified by its center location, the axis along which the Gaussian is squeezed (if any) and the standard deviations parallel and perpendicular to this axis. The evolution under a Gaussian operation can therefore be efficiently classically calculated using linear algebra – all one needs to do is keep track of these few coordinates.

Gaussian operations are generated by Hamiltonians containing terms with at most two raising and lowering operators, with the full set consisting of displacements ($\hat{a} + \hat{a}^\dagger$), phase space rotations ($\hat{a}^\dagger \hat{a}$), squeezing ($(\hat{a})^2 + (\hat{a}^\dagger)^2$), beamsplitters ($\hat{a}^\dagger \hat{b} + \hat{b}^\dagger \hat{a}$), and two-mode-squeezing ($\hat{a}\hat{b} + \hat{a}^\dagger \hat{b}^\dagger$). To perform tasks that are not simulable classically, we therefore cannot rely on linear continuous-variable systems alone. In fact, it turns out that adding just a single non-Gaussian operation acting on a single mode is sufficient to achieve universal multi-mode control (Lloyd and Braunstein, 1999), which can be obtained by engineering a nonlinear Hamiltonian term containing at least 3 raising or lowering operators. In superconducting circuits, this is most typically introduced via a 4$^{\text{th}}$-order, or Kerr, nonlinearity in the form of a transmon circuit (Koch *et al.*, 2007; Schreier *et al.*, 2008).

2.1.1 Transmons as a source of Kerr nonlinearity

Replacing the linear inductor in the LC oscillator with a nonlinear inductor in the form of a Josephson junction modifies the circuit Hamiltonian to the following,

$$\hat{\mathcal{H}} = \frac{\left(\hat{Q} - Q_{\text{ofs}}\right)^2}{2C} - \frac{(\Phi_0/2\pi)^2}{L_{\text{J}}} \cos\left(2\pi \frac{\hat{\Phi}}{\Phi_0}\right), \tag{2.11}$$

where L_{J} is the Josephson inductance of the junction and $\Phi_0 \equiv \frac{h}{2e}$ is the magnetic flux quantum. This replacement introduces an ‘island’ in the circuit with no completely superconducting path to ground (see Fig. 2.3(b)), and the energy of charges on this island can be affected by changes in the electric field of its environment. This is accounted for by introducing an offset charge

Q_{ofs}. In making the switch from a linear to a nonlinear inductor, the flux variable is now continuous but periodic, whereas the charge variable is discrete and can take on only integer values. Therefore, the continuous-variables Wigner function picture cannot capture the physics of this circuit in its full generality. To reflect the underlying physics of the Josephson junction, the Hamiltonian of the circuit is more typically written in terms of a discrete number of Cooper pairs on the junction 'island', $\hat{N} = \frac{\hat{Q}}{2e}$, and the continuous superconducting phase drop across the junction, $\hat{\varphi} = 2\pi\frac{\hat{\Phi}}{\Phi_0}$, as:

$$\hat{\mathcal{H}} = 4E_{\mathrm{C}}\left(\hat{N} - N_{\mathsf{ofs}}\right)^2 - E_{\mathrm{J}}\cos\hat{\varphi}, \tag{2.12}$$

where the prefactors are the charging energy, $E_{\mathrm{C}} \equiv \frac{e^2}{2C}$, and the Josephson energy, $E_{\mathrm{J}} \equiv (\Phi_0/2\pi)^2/L_{\mathrm{J}}$.

What defines this circuit as a *transmon* (as opposed to a generic Cooper pair box circuit (Bouchiat *et al.*, 1998)) is the large ratio $E_{\mathrm{J}}/E_{\mathrm{C}} \gg 1$, achieved by using a large shunting capacitance C. As predicted by Koch *et al.* (2007) and shown by Schreier *et al.* (2008), the fluctuations in the transmon eigenfrequencies as a function of the environment-controlled N_{ofs} (i.e., charge noise fluctuations) are exponentially suppressed as we increase this ratio, allowing us to neglect the role of N_{ofs} in the Hamiltonian. In this regime, the zero-point fluctuations of the phase coordinate are small relative to the periodicity of the Josephson potential, $\varphi_{\mathsf{ZPF}} = \left(\frac{2E_{\mathrm{C}}}{E_{\mathrm{J}}}\right)^{\frac{1}{4}} \ll 2\pi$, allowing us to perform a Taylor series expansion about the minimum

$$\hat{U}(\hat{\varphi}) = -E_{\mathrm{J}}\cos\hat{\varphi} = E_{\mathrm{J}}\frac{\hat{\varphi}^2}{2!} - E_{\mathrm{J}}\frac{\hat{\varphi}^4}{4!} + \dots \tag{2.13}$$

In this limit, we recover the harmonic oscillator description, allowing us to once more write the Hamiltonian in terms of two continuous degrees of freedom, $\hat{N} = -iN_{\mathsf{ZPF}}(\hat{a} - \hat{a}^\dagger)$ and

$\hat{\varphi} = \varphi_{\mathsf{ZPF}}(\hat{a} + \hat{a}^\dagger)$, or equivalently, in terms of raising and lowering operators:

$$\hat{\mathcal{H}} = 4E_{\mathrm{C}}\hat{N}^2 + \frac{E_{\mathrm{J}}}{2}\hat{\varphi}^2 - \frac{E_{\mathrm{J}}}{24}\hat{\varphi}^4 + \dots \tag{2.14}$$

$$= \hbar\omega_a \hat{a}^\dagger \hat{a} + \hbar g_4 \left(\hat{a} + \hat{a}^\dagger\right)^4 + \dots \tag{2.15}$$

where $\hbar\omega_a = \sqrt{8E_{\mathrm{J}}E_{\mathrm{C}}}$ and the 4$^{\text{th}}$ order nonlinearity $\hbar g_4 = -E_{\mathrm{J}}\varphi_{\mathsf{ZPF}}^4/4! = -E_{\mathrm{C}}/12$.

We can simplify this nonlinear Hamiltonian by performing a rotating-frame transformation, $\hat{\mathcal{H}} \to \hat{U}\hat{\mathcal{H}}\hat{U}^\dagger + i\hbar\dot{\hat{U}}\hat{U}^\dagger$, where $\hat{U} = e^{-i\omega_a t \hat{a}^\dagger \hat{a}}$ (see Appendix A of Reinhold (2019)) to move into the 'interaction picture':

$$\frac{\hat{\mathcal{H}}}{\hbar} = g_4 \left(\hat{a}e^{i\omega_a t} + \hat{a}e^{-i\omega_a t}\right)^4 + \dots. \tag{2.16}$$

From this viewpoint, we can see that terms in this expansion that do not conserve energy (i.e. which do not preserve the excitation number in the oscillator) will have an exponential prefactor rotating at a rate ω_a or greater. Since the dynamics of interest will occur on timescales much longer than $2\pi/\omega_a$, the effect of these non-energy-conserving terms will approximately average to zero. This 'rotating-wave approximation' (RWA) (Walls and Milburn, 2008) allows us to rewrite the Hamiltonian as that of a Kerr nonlinear oscillator

$$\frac{\hat{\mathcal{H}}}{\hbar} \approx \frac{K_a}{2}\hat{a}^\dagger\hat{a}^\dagger\hat{a}\hat{a}, \tag{2.17}$$

where the Kerr nonlinearity (or 'anharmonicity'), $K_a = 12g_4 = -E_{\mathrm{C}}$, can be interpreted (up to corrections due to higher orders in the Taylor expansion of the potential) as the difference between adjacent energy level transitions, e.g. $((E_2 - E_1) - (E_1 - E_0))/\hbar$. As evidence that this Kerr nonlinearity permits non-Gaussian operations, we can see that lifting the frequency degeneracy of these transitions allows a microwave drive with a frequency bandwidth $\sigma_f \ll |K_a|$ to selectively address the $|0\rangle \to |1\rangle$ transition. Rather than displacing the oscillator state from vacuum to produce a coherent state, as we saw for a purely linear oscillator, we can now prepare the non-Gaussian Fock $|1\rangle$ state.

Truncating the Hilbert space of the nonlinear oscillator to these lowest two energy levels (now referred to as $|g\rangle$ and $|e\rangle$, as shown in Fig. 2.1b) reduces this continuous variable system to a **discrete-variable** system, with a density matrix $\hat{\rho}$ that can be represented as a 2×2 complex-valued matrix. This truncation provides a natural way of encoding a single qubit of information, with 'logical' 0 and 1 represented by $|0\rangle_L \equiv |g\rangle$ and $|1\rangle_L \equiv |e\rangle$, respectively. Depending on the circumstance, both the oscillator-like and the qubit-like representation of nonlinear Josephson-junction-based circuits will be used throughout the thesis. To date, the most popular approach for constructing a superconducting quantum computer exclusively encodes information in these transmon qubits (see Acharya et al. (2024) and Kim *et al.* (2023) for some recent, high-profile examples). However, this is not the only way of using the Kerr nonlinearity to process quantum information.

2.1.2 Encoding information in 3D cavities

An alternative approach is to store the information in linear oscillators and only use the nonlinear degrees of freedom to manipulate this information. This approach adds complexity to the processor design but storing information in a linear resonator provides two significant advantages that can be leveraged:

1. they experience low, highly-biased intrinsic noise, and
2. they provide access to many quantum states in a single physical device.

The first advantage relates to the challenge of engineering a nonlinear resonator relative to a linear one, with more constraints imposed on its fabrication. Engineering a sufficiently strong nonlinearity requires a large current density (and therefore energy density) near the Josephson junction, making the nonlinear mode more prone to energy relaxation through loss channels in this region (Wang *et al.*, 2015). Furthermore, depending on the choice of materials used for the junction (e.g. Al/AlOx/Al, as is typically used), one may be restricted to using particular substrate materials or to using particular chemical processing techniques that do not damage the junction. While there has been considerable progress in recent years to reduce loss in nonlinear superconducting circuits while working within these constraints (Ganjam *et al.*, 2024; Place

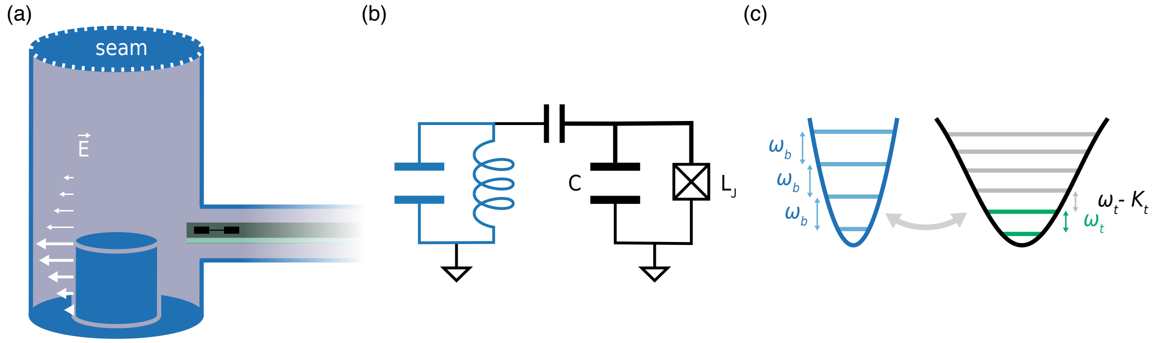

Figure 2.3: **Three representations of a 3D cavity resonator coupled to a transmon qubit.** (a) Illustration of physical implementation of a 3D cavity (blue) consisting of a superconducting coaxial stub (Reagor *et al.*, 2016), whose fundamental mode has a wavelength $4\times$ the length of the stub. The electric field distribution of this mode (shown in white) decays exponentially with distance above the post so that the energy participation in the lossy seam is minimal. The interior of the cavity (light blue) is vacuum. Into this cavity, one inserts a substrate (teal) supporting a 3D transmon (Paik *et al.*, 2011) (black), consisting of a Josephson junction (not shown) shunted by a superconducting capacitor. (b) Electrical circuit schematic with the fundamental cavity mode represented by a harmonic LC oscillator, capacitively coupled to a transmon, where the inductor is replaced by a Josephson junction (with Josephson inductance L_J). The thicker black line indicates the superconducting 'island' of the transmon formed by the introduction of the junction. (c) Quantum optics picture, showing the potentials of the linear (blue) and nonlinear (black) modes and the energy eigenvalues of the lowest eigenstates (horizontal lines). In the cavity, these are (approximately) equally spaced, whereas in the transmon, each transition has a different energy difference. The lowest two levels of the nonlinear mode form a qubit.

et al., 2021), the design of linear resonators provides much more flexibility.

A design for a linear superconducting oscillator that successfully exploits this flexibility is the 3D stub cavity (Reagor *et al.*, 2016), consisting of a coaxial section shorted at one end and made entirely of a superconductor such as high-purity aluminum, forming a distributed $\lambda/4$ microwave resonator (Pozar, 2012) (Fig. 2.3(a)). This design minimizes the mode participation (defined as the fraction of the mode's stored energy) in lossier regions or degrees of freedom, including in the oxide layer on the surface of the superconductor. It does so by conversely maximizing the mode participation in the vacuum, which (at least to the level currently observable in experiment) is completely dissipationless. Crucially, since the electric field decays exponentially away from the top of the stub, the mode participation in the 'seam' between the body and lid of the cavity can be made negligible, suppressing current flowing across this lossy interface. The suppression of these different loss channels permits single photon lifetimes $T_1 > 1$ ms (Reagor *et al.*, 2016; Rosenblum *et al.*, 2018), a few times higher than for transmon qubits (Place *et al.*, 2021). Recent work has extended this even further to $T_1 > 10$ ms by making use of aggressive chemical processing techniques (e.g. hot hydrofluoric acid etching to remove lossy dielectric oxide layers) that are incompatible with transmon fabrication (Oriani *et al.*, 2024), and as high as $T_1 > 25$ ms by modifying the 3D cavity design to even further maximize mode participation in the vacuum (Milul *et al.*, 2023). In fact, even when part of the resonator is hosted on a dielectric substrate, as opposed to purely constructed out of bulk superconductors and vacuum, the greater design flexibility allows for $T_1 > 1$ ms (Ganjam *et al.*, 2024). Importantly, all of these quoted values are for cavities integrated with a transmon qubit to allow for non-Gaussian control – when the ancillary transmon is removed however, cavity lifetimes can be in excess of 2 s (Romanenko *et al.*, 2020).

While this demonstrates the possibility of lower energy relaxation rates in linear superconducting resonators, the more powerful feature is that their intrinsic dephasing rates are yet much higher. Dephasing noise results from fluctuations in the mode frequency. For a transmon, we have already discussed how charge noise can be suppressed, but critical current noise (i.e. fluctuations in E_J) (Van Harlingen *et al.*, 2004) and nonlinear coupling to two-level systems (TLSs – presumed to originate from spin defects near the junction) (Gao *et al.*, 2008; de Leon *et al.*, 2021; Lisenfeld *et al.*, 2016; Martinis *et al.*, 2005; Müller *et al.*, 2019) are both microscopic factors that can modulate the qubit frequency and degrade the pure dephasing time, T_ϕ, of the resonator. (The pure dephasing time can be related to the measured coherence time T_2 of a resonator, via $\frac{1}{T_\phi} = \frac{1}{T_2} - \frac{1}{2T_1}$). However, for distributed cavity resonators, the frequency is primarily set by the macroscopic dimensions of the device, such as the length of the stub in the case of the 3D coaxial stub cavity. Since the relative fluctuations in these macroscopic dimensions are negligible, so too is the mode dephasing. 3D linear resonators routinely demonstrate $T_2 \approx 2T_1$, indicating a strong intrinsic noise bias $T_\phi \gg T_1$ (Rosenblum *et al.*, 2018). The word *intrinsic* should be emphasized, since once the linear oscillator is coupled to a nonlinear element, we introduce an extrinsic source of dephasing which is typically dominant. Nonetheless, as we shall see, one can design nonlinear control sequences to protect against these types of errors.

The second useful feature of linear oscillators is that the user does not need to limit their attention to the lowest two energy levels but has access to the larger Hilbert space. This can be leveraged to either encode more than one bit of information in a single resonator or, more commonly, encode a single bit of information with more redundancy in order to implement error correction at the individual physical qubit level (Girvin, 2023). Qubit encodings that make use of the larger Hilbert space of a linear mode are known as 'bosonic codes' (Cai *et al.*, 2021). Control schemes for bosonic codes typically make use of the fact that oscillator is linear, with equally-spaced energy levels. Therefore, it is often vital (especially so as the photon number in the encoded states grows larger) that the oscillator remain as linear as possible to avoid distorting the encoded logical states, with a resulting drop in fidelity (Kirchmair *et al.*, 2013; Vlastakis *et al.*, 2015).

Bosonic codes: some case studies

Many different bosonic encodings have been proposed (Gottesman *et al.*, 2001; Michael *et al.*, 2016; Mirrahimi *et al.*, 2014; Puri *et al.*, 2017; Teoh *et al.*, 2023), most of which exploit one or both of a) the low nonlinearity of the oscillator modes and b) their strong bias towards relaxation errors over dephasing errors. One encoding that makes use of **both of these properties** is the smallest member of the family of binomial codes (Michael *et al.*, 2016), often referred to as the 'kitten code.' Here, information is encoded in superpositions of different photon number ('Fock') states $|n\rangle$ of an oscillator, with the two logical codewords defined as:

$$|0\rangle_L \equiv \frac{|0\rangle + |4\rangle}{\sqrt{2}}, \qquad\qquad |1\rangle_L \equiv |2\rangle\,, \tag{2.18}$$

both of which contain an even number of photons. All superpositions of these states therefore also contain an even number of photons. On timescales short compared to the lifetime of oscillator, $t \ll T_1$, energy relaxation can be modeled as single-photon loss, described by a Kraus 'jump operator' $\mathcal{E}_{\text{loss}} \approx \sqrt{\kappa t}\ \hat{a}$, where $\kappa = 1/T_1$. Applying this jump operator to our logical states (and normalizing them) maps them onto error states possessing an odd number of photons:

$$|0\rangle_E \equiv \frac{\hat{a}\,|0\rangle_L}{\sqrt{{}_L\langle 0|\,\hat{a}^\dagger\hat{a}\,|0\rangle_L}} = |3\rangle\,, \qquad |1\rangle_E \equiv \frac{\hat{a}\,|1\rangle_L}{\sqrt{{}_L\langle 1|\,\hat{a}^\dagger\hat{a}\,|1\rangle_L}} = |1\rangle\,. \tag{2.19}$$

As a result, a measurement of the photon number parity (n mod 2) can distinguish between logical states and error states so that, in the event of an 'odd' measurement outcome, the application of an appropriate recovery operation (a unitary that maps $|3\rangle \to (|0\rangle + |4\rangle)/\sqrt{2}$ and $|1\rangle \to |2\rangle$) restores the initial quantum state. Crucially though, a measurement that only probes the photon number parity cannot distinguish between different logical states or between different error states. The normalization factors in Eq. 2.19 tell us about the relative sensitivities of each state to loss errors. The fact that ${}_L\langle 0|\,\hat{a}^\dagger\hat{a}\,|0\rangle_L = {}_L\langle 1|\,\hat{a}^\dagger\hat{a}\,|1\rangle_L$ tells us that errors are equally likely from each logical state so that the act of detection itself does not inform us about the initial state. This is crucial to avoid the measurement distorting the encoded logical information.

For concreteness, the Knill-Laflamme conditions (Knill *et al.*, 2000) succinctly provide necessary and sufficient criteria to determine whether a set of errors $\{\mathcal{E}_l\}$ (where $\mathcal{E}_0 = \mathbb{I}$ to account for 'no-jump' events) is *correctable*, requiring

$${}_L\langle i|\,\mathcal{E}_l^\dagger \mathcal{E}_{l'}\,|j\rangle_L = \alpha_{ll'}\delta_{ij}, \tag{2.20}$$

where $\alpha_{ll'}$ is a Hermitian matrix. While I will not go into detail deriving these conditions here, the original reference does so excellently. One can verify that the error set $\{\mathbb{I}, \mathcal{E}_{\text{loss}}\}$ satisfies these conditions for the kitten code.

Cavity dephasing errors on the other hand, can be modeled at very short times t (assuming that the noise is Markovian) by the jump operator $\mathcal{E}_{\text{dephasing}} \approx \sqrt{\gamma_\phi t}\ \hat{a}^\dagger\hat{a}$, where $\gamma_\phi = 1/T_\phi$. These errors map the logical states to

$$|0\rangle_E^{(\text{dephasing})} \equiv \frac{\hat{a}^\dagger\hat{a}\,|0\rangle_L}{\sqrt{{}_L\langle 0|\,\hat{a}^\dagger\hat{a}\hat{a}^\dagger\hat{a}\,|0\rangle_L}} = |4\rangle\,, \quad |1\rangle_E^{(\text{dephasing})} \equiv \frac{\hat{a}^\dagger\hat{a}\,|1\rangle_L}{\sqrt{{}_L\langle 1|\,\hat{a}^\dagger\hat{a}\hat{a}^\dagger\hat{a}\,|1\rangle_L}} = |2\rangle\,. \tag{2.21}$$

Unfortunately, unlike for single photon loss, there does not exist a measurement that can both detect whether a dephasing error has occurred without also being able to distinguish between the different error states. Even the error set excluding single photon loss, $\{\mathbb{I}, \mathcal{E}_{\text{dephasing}}\}$, does not satisfy Eq. 2.20. One way to see this is that ${}_L\langle 0|\, \hat{a}^\dagger \hat{a} \hat{a}^\dagger \hat{a}\, |0\rangle_L \neq {}_L\langle 1|\, \hat{a}^\dagger \hat{a} \hat{a}^\dagger \hat{a}\, |1\rangle_L$, with dephasing errors more likely to occur when in $|0\rangle_L$ than in $|1\rangle_L$. As such, dephasing errors are uncorrectable for this code. The intrinsic protection that 3D cavities offer against dephasing is therefore critical to preserving the stored quantum information. The kitten code makes use of **both** the noise bias of 3D cavities (to ensure a single dominant error type) and their linearity (to provide access to energy levels higher than $|1\rangle$, without distortion). Other bosonic codes that fall in the same category include the Gottesman-Kitaev-Preskill (GKP) code, listed in Table 2.1.

Codes that **only require the larger Hilbert space** include larger binomial codes, such as

$$|0\rangle_L \equiv \frac{|0\rangle + \sqrt{3}\,|6\rangle}{2}, \qquad |1\rangle_L \equiv \frac{\sqrt{3}\,|3\rangle + |9\rangle}{2}, \tag{2.22}$$

as well as the dissipatively-stabilized four-component cat code (Mirrahimi *et al.*, 2014), which can actively correct against both photon loss and dephasing errors. This highlights the benefit of being able to access a larger Hilbert space to redundantly encode information, however the trade-off is that this also places more stringent requirements on the linearity of the oscillator mode. Codes that **only require the beneficial noise structure** include the dual-rail encoding (Chuang and Yamamoto, 1995; Teoh *et al.*, 2023) mentioned in Chapter 1. This encoding only uses the two lowest-lying energy levels in each of two oscillators to define a single logical qubit:

$$|0\rangle_L \equiv |0\rangle_A\,|1\rangle_B \qquad |1\rangle_L \equiv |1\rangle_A\,|0\rangle_B \tag{2.23}$$

Here, the action of single photon loss in either cavity takes both states to a common vacuum state:

$$\hat{b}\,|0\rangle_L = \hat{a}\,|1\rangle_L = |0\rangle_A\,|0\rangle_B \tag{2.24}$$

In this case, measurements of the combined photon number in both oscillators are what distin-

Encoding	**Low nonlinearity required?**	**Intrinsic $T_\phi \gg T_1$ noise bias required?**
GKP	Yes	Yes
4-component cat	Yes	Yes
Binomial (Kitten)	Yes	Yes
Binomial (Larger)	Yes	No
Dissipative cat	Yes	No
Dual-rail	No	Yes
Kerr cat	No*	No

Table 2.1: **Requirements for bosonic encodings.** The properties leveraged in order to obtain an advantage over discrete variable (DV) systems for GKP (Gottesman *et al.*, 2001), Binomial (Michael *et al.*, 2016), Cavity Dual-Rail (Teoh *et al.*, 2023), Kerr Cat (Puri *et al.*, 2017), (2- and 4-Component) Dissipative Cat (Mirrahimi *et al.*, 2014) and 4-Component (non-Stabilized) Cat qubits (Ofek *et al.*, 2016). It should be noted that the degree to which the oscillator is required to be linear varies substantially between encodings. Whereas unwanted Kerr nonlinearity is particularly costly for GKP codes, dissipative cat qubits can achieve substantial error suppression even with moderate Kerr (Putterman et al., 2024).

guishes between logical and error states, without distinguishing logical states from one another. Unlike in the kitten code however, the two error states are not distinguishable and so in the event of an error, the initial quantum state is not recoverable. Nonetheless, the ability to *detect* (if not correct) errors in an efficient way, without perturbing the quantum state, can still be of great value. When the qubit itself forms part of a larger, redundantly-encoded memory, this error-detection information can be used to much more efficiently pinpoint where errors have occurred in the larger system. Such a qubit is known as an erasure qubit (Grassl *et al.*, 1997; Kubica *et al.*, 2023) and the measurement that enables this, the mid-circuit erasure check, will be the focus of Chapter 7.

Once more, however, cavity dephasing errors are more problematic, leading to undetectable errors within the logical subspace (e.g. mapping $(|0\rangle_L + |1\rangle_L)/\sqrt{2} \to (|0\rangle_L - |1\rangle_L)/\sqrt{2}$). The dual-rail encoding thus requires the cavities to possess single photon loss as the dominant error channel but places no requirements on the linearity of the oscillator. This is exemplified by the fact that one can also encode dual-rail qubits in a pair of strongly coupled transmons (Kubica *et al.*, 2023; Levine et al., 2024), which also exhibits noise that is biased against dephasing.

Bosonic qubits encoded in 3D cavities thus leverage these two key properties (linearity and noise bias) to either detect or correct errors, such that encoded information is less prone to noise than the underlying hardware. In Table 2.1 we categorize several popular bosonic codes by the requirements they make of the linearity and intrinsic noise bias of the oscillator[3]. So far we have only considered oscillators in isolation but now we must turn our attention to the interactions required to couple these oscillators to other modes and to manipulate this stored information. In designing a *general-purpose* bosonic processor, it is vital that these interactions are engineered in a way that preserves both the linearity and noise bias that give bosonic codes their value.

2.2 Dispersive control

The workhorse of single-oscillator control in circuit-QED is the 'dispersive' interaction between the linear oscillator and a nonlinear ancillary transmon, engineered by *weakly* coupling the two modes via a mutual capacitance $\tilde{C}$. To derive the form of this interaction, we can use the description of a transmon as a weakly nonlinear oscillator, starting with the Hamiltonian of two completely linear modes, before later introducing the Kerr nonlinearity as a perturbation. In this framework, the capacitive coupling between the 'bare' oscillator ($\hat{B}$) and transmon ($\hat{T}$) modes generates the following interaction:

$$\hat{\mathcal{H}}_{\text{lin}} = \hbar\omega_B \hat{B}^\dagger \hat{B} + \hbar\omega_T \hat{T}^\dagger \hat{T} + \hat{\mathcal{H}}_{\text{coupling}}, \tag{2.25}$$

[3]The potential exception to the statement that bosonic codes rely either on high linearity or high intrinsic noise bias is the Kerr-cat qubit (Puri *et al.*, 2017) which, while using many levels of an oscillator to encode its information, does not place stringent requirements on the underlying hardware to possess an error bias (since it is designed to protect against dephasing errors) or a high degree of linearity (relying on a Kerr nonlinearity to generate the noise protection). Nonetheless, the size of the Kerr nonlinearity relative to the drive strength does eventually limit the number of photons in the cat state, and therefore the noise suppression that can be achieved.

where

$$\hat{\mathcal{H}}_{\text{coupling}} = \frac{\hat{Q}^{(B)}\hat{Q}^{(T)}}{2\tilde{C}}, \tag{2.26}$$

$$= -\frac{Q^{(B)}_{\text{ZPF}}Q^{(T)}_{\text{ZPF}}}{2\tilde{C}}\left(\hat{B}^\dagger - \hat{B}\right)\left(\hat{T}^\dagger - \hat{T}\right), \tag{2.27}$$

$$\equiv -\hbar g\left(\hat{B}^\dagger - \hat{B}\right)\left(\hat{T}^\dagger - \hat{T}\right). \tag{2.28}$$

In the regime that the two mode frequencies are close together, $|\Delta| = |\omega_B - \omega_T| \ll |\omega_B + \omega_T|$, we may apply the rotating-wave approximation (RWA) and neglect terms in Eq. 2.28 that do not preserve photon number. The resulting Hamiltonian, consisting of a bilinear coupling between the two 'bare' modes, may be expressed in the following matrix form

$$\frac{\hat{\mathcal{H}}_{\text{lin}}}{\hbar} \approx \begin{pmatrix} \hat{B}^\dagger & \hat{T}^\dagger \end{pmatrix} \begin{pmatrix} \omega_B & g \\ g & \omega_T \end{pmatrix} \begin{pmatrix} \hat{B} \\ \hat{T} \end{pmatrix}. \tag{2.29}$$

Diagonalizing this matrix provides us with the 'dressed' modes (corresponding to the classical normal modes of the system), allowing the Hamiltonian to be written in the form

$$\frac{\hat{\mathcal{H}}_{\text{lin}}}{\hbar} \approx \omega_b \hat{b}^\dagger \hat{b} + \omega_t \hat{t}^\dagger \hat{t}. \tag{2.30}$$

In the regime that the coupling is weak ($g \ll \Delta$), the two dressed modes largely preserve the character of the bare modes:

$$\omega_b \approx \omega_B + \frac{g^2}{\Delta} \qquad \hat{b} \approx \hat{B} + \frac{g}{\Delta}\hat{T} \tag{2.31}$$

$$\omega_t \approx \omega_T - \frac{g^2}{\Delta} \qquad \hat{t} \approx \hat{T} - \frac{g}{\Delta}\hat{B}. \tag{2.32}$$

We can similarly invert these expressions to obtain the bare modes in terms of the new dressed modes:

$$\hat{B} \approx \hat{b} - \frac{g}{\Delta}\hat{t} \tag{2.33}$$

$$\hat{T} \approx \hat{t} + \frac{g}{\Delta}\hat{b}. \tag{2.34}$$

We will often refer to the ratio (g/Δ) as the participation p of a coupled oscillator mode in a bare nonlinear mode.

To the Hamiltonian describing the two coupled linear modes, we can add in the term from Eq. 2.17 that captures the intrinsic Kerr nonlinearity of the (bare) transmon mode, $\hat{\mathcal{H}}_{\mathsf{Kerr}} = \frac{\hbar K_T}{2}\hat{T}^\dagger\hat{T}^\dagger\hat{T}\hat{T}$. Provided that $|\Delta| \gg |K_T|$, we can treat $\hat{\mathcal{H}}_{\mathsf{Kerr}}$ as a perturbation to the Hamiltonian. Therefore we can write $\hat{\mathcal{H}}_{\mathsf{Kerr}}$ in terms of the dressed modes and keep only the terms within the RWA to obtain (in a frame rotating at ω_b and at ω_t),

$$\frac{\hat{H}_{\mathsf{Kerr}}}{\hbar} \approx \frac{K_t}{2}\hat{t}^\dagger\hat{t}^\dagger\hat{t}\hat{t} + \underbrace{\chi\hat{b}^\dagger\hat{b}\hat{t}^\dagger\hat{t}}_{\hat{\mathcal{H}}_{\mathsf{disp}/\hbar}} + \frac{K_b}{2}\hat{b}^\dagger\hat{b}^\dagger\hat{b}\hat{b}. \tag{2.35}$$

The first and last terms represent the Kerr nonlinearity of the dressed modes, with the inherited nonlinearity in the oscillator-like mode given by $K_b \approx \left(\frac{g}{\Delta}\right)^4 K_T$. The second term is the cross-Kerr or dispersive interaction that we have been seeking, with $\chi \approx 2\left(\frac{g}{\Delta}\right)^2 K_T$, and can be interpreted as a shift of the ancilla (angular) frequency by χ per photon in the oscillator. If we truncate the transmon Hilbert space to its lowest $|g\rangle$ and $|e\rangle$ energy levels, we recover the form of the dispersive Hamiltonian from Eq. 2.3, $\hat{\mathcal{H}}_{\mathsf{disp}} = \hbar\chi\hat{b}^\dagger\hat{b}\,|e\rangle\langle e|$.

When performing microwave spectroscopy of the ancilla qubit to determine its dressed frequency, the linewidth of the observed resonance is approximately set by $\max(1/T_2^{\mathsf{ancilla}}, 1/T_{\mathrm{p}})$, where T_{p} is the pulse duration of the spectroscopy tone and T_2^{ancilla} is the coherence time of the ancilla qubit. As a result, in the ‘strong’ dispersive coupling limit, where $|\chi| \gg 1/T_2^{\mathsf{ancilla}}$, a spectroscopy pulse $T_{\mathrm{p}} > 1/|\chi|$ allows us to resolve peaks associated with different photon number states in the oscillator. This photon number split spectrum provides the basis for photon-number-dependent gates and measurements of the oscillator mode with pulses on the ancilla.

Combining the dispersive coupling Hamiltonian $\mathcal{H}_{\text{disp}}$ with the ability to readout the state of the ancilla (Mallet *et al.*, 2009) provides us with a toolbox of operations for manipulating and measuring oscillator states (Reinhold, 2019; Vlastakis, 2015). In particular, the photon-number-selective measurement (Ch. 2.2.1), the photon-number-parity measurement (Ch. 2.2.1) and the SNAP gate (Ch. 2.2.2) form a set of photon-number-selective single-cavity operations (Fig. 2.4), whose two-cavity analogs will be developed in Chapters 5 and 6 of this thesis.

2.2.1 Quantum non-demolition oscillator measurements

Photon-number-selective measurements

As illustrated in Fig. 2.4(c) and (f), driving the ancilla at a frequency $\omega_{\text{p}} = \omega_t + N_{\text{meas}}\chi$, with a pulse duration $T_{\text{p}} \gg 1/|\chi|$, excites the ancilla if and only if the photon number in the oscillator $N = N_{\text{meas}}$ (Johnson *et al.*, 2010). This photon-number-selective drive is described by the Hamiltonian

$$\frac{\hat{\mathcal{H}}_{N_{\text{meas}}}}{\hbar} = \left(\frac{\epsilon e^{-i\phi}}{2}\left|g\right\rangle\left\langle e\right| + \frac{\epsilon e^{i\phi}}{2}\left|e\right\rangle\left\langle g\right|\right) \otimes \left|N_{\text{meas}}\right\rangle\left\langle N_{\text{meas}}\right|, \tag{2.36}$$

where ϵ and ϕ are its amplitude and phase, so that after time t, we enact a photon-number-selective ancilla qubit rotation, $R_\phi(\epsilon t)$, by angle ϵt about an axis $(\cos\phi, \sin\phi, 0)$ on the Bloch sphere (Nielsen and Chuang, 2010). When $\epsilon t = \pi$ and (say) $\phi = 0$, this enacts the unitary:

$$\hat{U}_{N_{\text{meas}}} = \left(\hat{\mathbb{I}} - \left|N_{\text{meas}}\right\rangle\left\langle N_{\text{meas}}\right|\right) \otimes \hat{\mathbb{I}} + \tag{2.37}$$

$$\left|N_{\text{meas}}\right\rangle\left\langle N_{\text{meas}}\right| \otimes \hat{X}_\pi. \tag{2.38}$$

where $\hat{X}_\pi = \left|g\right\rangle\left\langle e\right| + \left|e\right\rangle\left\langle g\right|$, and which, when the ancilla is initialized in $\left|g\right\rangle$, maps the 'is $N = N_{\text{meas}}$?' information onto the ancilla state. Subsequently performing a projective measurement of the ancilla allows us to probe the observable $\hat{O} = \left|N_{\text{meas}}\right\rangle\left\langle N_{\text{meas}}\right|$ of the oscillator. Importantly, since $\hat{O}$ commutes with both the system Hamiltonian ($\hat{\mathcal{H}}_{\text{disp}}$) and the measurement

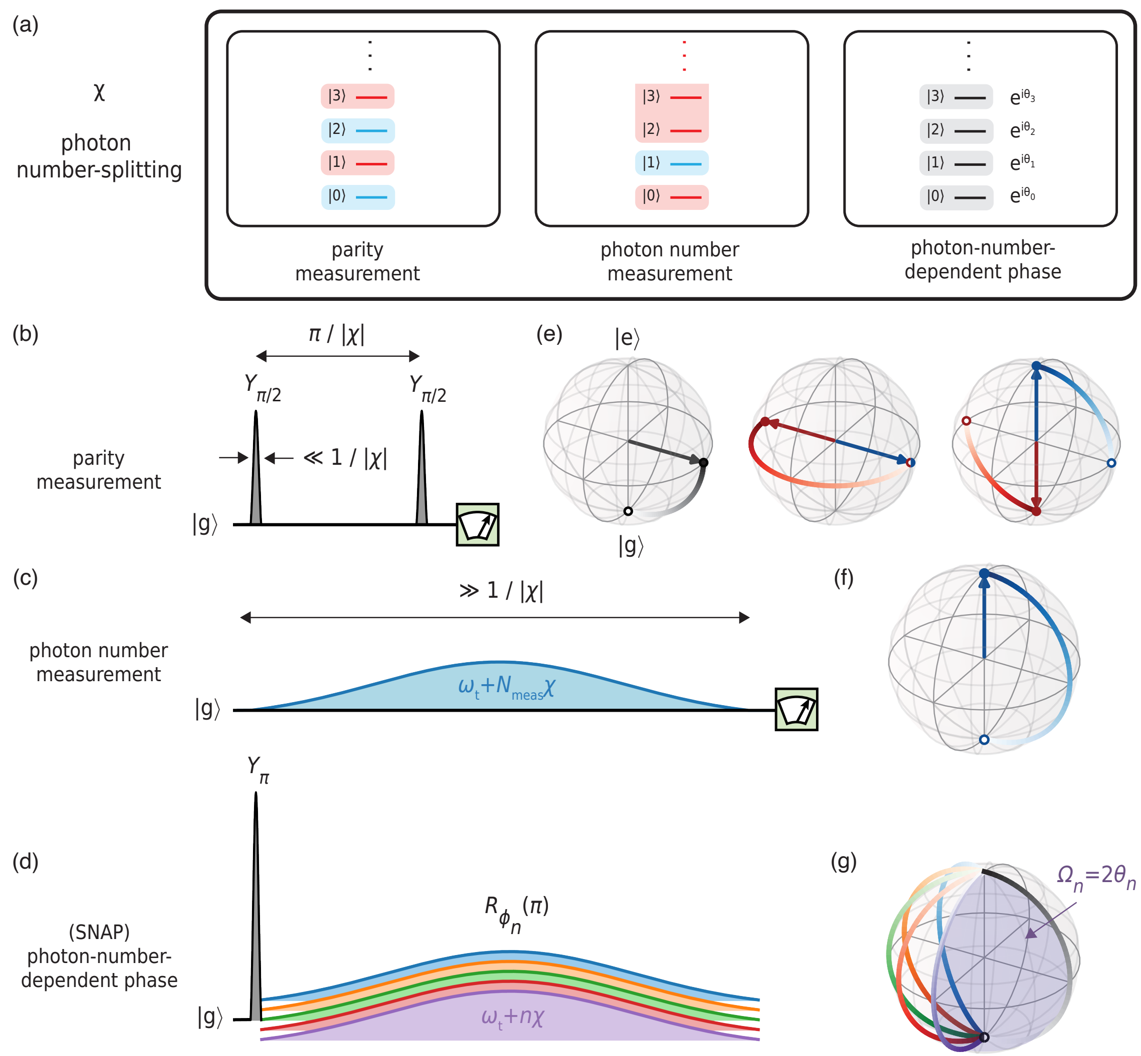

Figure 2.4: **A toolbox of dispersive single-mode control operations.** (a) The toolbox contains a joint-parity measurement, which excites an ancilla qubit only if the photon number is even, a photon-number-selective measurement, which excites the ancilla only if the photon number matches a specific N_{meas}, and a photon-number-selective phase (or 'SNAP'), which applies an arbitrary phase conditioned on the oscillator photon number. (b-d) The pulse sequences for each operation. The short pulses in black possess a frequency bandwidth much wider than $|\chi|$ and so are unconditional on the photon number, whereas the long pulses in color possess a narrow frequency bandwidth much less than $|\chi|$ and so perform photon-number-selective operations. In the case of SNAP, a comb of photon-number-selective pulses is applied. In principle, a comb of pulses could be used to measure a subset of photon numbers too. The relative vertical amplitude of the pulses is not to scale. (e-f) The qubit state trajectories for each operation. In the parity measurement, the qubit state is dependent on the oscillator photon number parity (blue for even; red for odd). In the SNAP gate, the qubit trajectory encloses an arbitrary solid angle conditioned on the oscillator photon number.

Hamiltonian ($\hat{\mathcal{H}}_{N_{\text{meas}}}$), repeated measurements of the same observable will yield the same eigenvalue. This quantum non-demolition (QND) property (Braginsky and Khalili, 1996), which ensures that the measurement outcome accurately describes the post-measurement state, is important whenever we wish to perform feedback conditioned on a measurement, such as in quantum error correction.

Parity measurements

A single measurement of the two-level ancilla can only return 1 bit of information, but by modifying the pulse sequence we can alter the 1-bit 'yes/no' question queried of the oscillator. For example, a frequency comb consisting of components at different values of $N_i\chi$ can be used to ask if the photon number N belongs to the set $\{N_i\}$ (Teoh, 2023). Another key example of this idea is the photon number parity measurement, which asks whether N is even or odd (Sun *et al.*, 2014).

In the time domain, this sequence (Fig. 2.4(b)) consists of two very short $\epsilon t = \pi/2$ pulses on the ancilla separated by a delay of $T_{\text{wait}} = \pi/|\chi|$. These short pulses, for which $1/T_{\text{p}} \gg |\chi|$, are too broad in frequency to distinguish the oscillator photon number and so enact an *unconditional* $\pi/2$ rotation of the ancilla state, independent of N. The qubit state trajectories for this sequence are shown in Fig. 2.4(e). The first $\pi/2$ pulse puts the ancilla in the state $(|g\rangle + |e\rangle)/\sqrt{2}$, the unitary evolution under $\hat{\mathcal{H}}_{\text{disp}}$ imparts a rotation $e^{i\hat{n}\pi|e\rangle\langle e|}$ such that the transmon lies in $(|g\rangle + |e\rangle)/\sqrt{2}$ for even and $(|g\rangle - |e\rangle)/\sqrt{2}$ for odd photon numbers, and the final $\pi/2$ pulse maps the ancilla state onto either $|g\rangle$ or $|e\rangle$ so that a projective measurement of the ancilla state determines the photon number parity, $\hat{P} = (-1)^{\hat{N}}$. Since the parity observable also commutes with both $\hat{\mathcal{H}}_{\text{disp}}$ and the measurement Hamiltonian ($\hat{\mathcal{H}}_{\text{meas}} \propto |g\rangle\langle e| + |e\rangle\langle g|$), its measurement is also QND.

The time duration of the parity sequence ($\pi/|\chi|$) represents the fastest that one can distinguish between adjacent photon numbers, but shorter sequences can be used to measure states spaced further apart. For example, by halving T_{wait}, we can instead probe the 'four-parity' (i.e.

N mod 4), $\hat{P}_4 = (-1)^{\hat{N}/2}$ (Curtis *et al.*, 2021) [4]. These generalized photon number parity measurements $\hat{P}_n$ have particular importance as they serve as error syndrome measurements for a broad class of bosonic codes known as rotation-symmetric codes (Grimsmo *et al.*, 2020), of which cat codes (Cochrane *et al.*, 1999; Mirrahimi *et al.*, 2014; Puri *et al.*, 2017) and binomial codes (Michael *et al.*, 2016) are prominent examples.

2.2.2 SNAP gates

The photon number splitting regime also provides a powerful way to perform gates on the oscillator state. The selective number arbitrary phase (SNAP) gate (Heeres *et al.*, 2015; Krastanov *et al.*, 2015) relies on the ability to perform photon-number-selective ancilla pulses to produce a different unitary evolution of the ancilla associated with each photon number state in the oscillator, $\hat{U} = \sum_n |n\rangle \langle n| \otimes \hat{U}_n$. At the end of the sequence, we require the oscillator and ancilla states to be disentangled so that noise on the ancilla cannot subsequently propagate to the stored logical information. Therefore all of these state-dependent trajectories should return the ancilla to $|g\rangle$, implementing the identity operator, up to some phase factor $e^{i\theta_n}$. Crucially, θ_n has a path-dependent component (the geometric, or Berry phase (Berry, 1984)) that depends on the solid angle enclosed by the path of the ancilla and so by engineering the photon-number-dependent trajectories on the ancilla Bloch sphere, we can impart an arbitrary phase to each photon-number manifold of the oscillator

$$\hat{U}_{\mathsf{SNAP}} = \sum_n |n\rangle \langle n| \, e^{i\theta_n}. \tag{2.39}$$

A convenient choice of pulse sequence (Fig. 2.4(d)) consists of an unconditional π-pulse followed by photon-number selective π-pulses with a different drive phase, ϕ_n for each frequency component. The 'slice' formed by these trajectories (see purple trajectory in Fig. 2.4(g)) encloses a photon-number-dependent solid angle, Ω_n, that can be related to the acquired phase via $\theta_n = \Omega_n/2$.

[4] Note that the state must already have definite, known 2-parity in order for this measurement to work. A single measurement of the ancilla qubit can still only extract 1 bit of information.

SNAP gates provide a parameterizable non-Gaussian operation which, when complemented with Gaussian oscillator displacements, enable universal control of an oscillator state (Krastanov *et al.*, 2015). Sequences that alternate between these two operations can approximate any unitary to arbitrary precision, and have been used to generate Fock states (Heeres *et al.*, 2015), as well as perform recovery operations for 4-component cat qubits (Ofek *et al.*, 2016).

2.2.3 Ancilla fault-tolerance

As established earlier, one of the two key motivations for processing information in linear oscillators is the favorable noise that they experience, with longer relaxation times and a large noise bias $T_\phi \gg T_1$. While the ancillary degree of freedom provides the necessary control, it is vital that it does so without compromising either noise property. In particular, it is important that errors on the ancilla do not propagate to the oscillator state. The ability to tolerate n of these errors is known as n^{th}-order 'hardware' fault-tolerance. In advance of considering fault-tolerant multi-mode schemes later in this thesis, it is necessary to briefly review the main single-mode techniques here.

Parity measurement

Both relaxation and dephasing errors on the ancilla are problematic. As an example of the former, we can consider the effect of an ancilla relaxation error during the parity measurement. Since the jump operator that describes it, $\mathcal{E}^{\text{ancilla}}_{\text{loss}} = \sqrt{\kappa^{\text{ancilla}} t}\ \ |g\rangle\langle e|$ does not commute with $\hat{\mathcal{H}}_{\text{disp}}$, $\left[|g\rangle\langle e|\,, \hat{b}^\dagger\hat{b}\,|e\rangle\langle e|\right] \neq 0$, the effect of an error on the final state depends on the (unknown) time t_{err} at which the error occurred. The parity map sequence then proceeds as follows:

$$|g\rangle\,|\psi\rangle \xrightarrow{\hat{Y}_{\pi/2}} \left[\frac{|g\rangle+|e\rangle}{\sqrt{2}}\right]|\psi\rangle \tag{2.40}$$

$$\xrightarrow{\text{evolve for } t_{\text{err}}} \frac{|g\rangle}{\sqrt{2}}\,|\psi\rangle + \frac{|e\rangle}{\sqrt{2}}e^{i\chi t_{\text{err}}\hat{a}^\dagger\hat{a}}\,|\psi\rangle \tag{2.41}$$

$$\xrightarrow{|g\rangle\langle e|} |g\rangle\, e^{i\chi t_{\text{err}}\hat{a}^\dagger\hat{a}}\,|\psi\rangle \tag{2.42}$$

$$\xrightarrow{\text{evolve for } \frac{\pi}{|\chi|}-t_{\text{err}}} |g\rangle\, e^{i\chi t_{\text{err}}\hat{a}^\dagger\hat{a}}\,|\psi\rangle \tag{2.43}$$

$$\xrightarrow{\hat{Y}_{\pi/2}} \left[\frac{|g\rangle+|e\rangle}{\sqrt{2}}\right] e^{i\chi t_{\text{err}}\hat{a}^\dagger\hat{a}}\,|\psi\rangle \tag{2.44}$$

This highlights two problems. Firstly, there is a 50% chance of an incorrect measurement outcome as the ancilla ends in an equal superposition independently of the initial oscillator state, and secondly, the oscillator state acquires an unknown rotation in phase space, corresponding to a dephasing error.

The impact of a single ancilla relaxation error can be mitigated by using a ‘three-level’ ancilla, using the $|g\rangle$ and $|f\rangle$ levels of the transmon as the ancilla qubit (Ma *et al.*, 2020; Rosenblum *et al.*, 2018). In this case, ancilla relaxation of the form $|e\rangle\langle f|$ leaves the ancilla in $|e\rangle$ where it remains unaffected by the final $\hat{Y}_{\pi/2}$ pulse. The $|f\rangle$ level has its own dispersive coupling to the oscillator $\chi_f \hat{b}^\dagger\hat{b}\,|f\rangle\langle f|$, where in general $\chi_f \neq \chi_e$. In this case, the final ancilla-oscillator state prior to measurement is $|e\rangle \otimes e^{i(\chi_f t_{\text{err}} + \chi_e(\pi/|\chi| - t_{\text{err}}))\hat{b}^\dagger\hat{b}}\,|\psi\rangle$. The first problem has therefore been resolved - the ancilla state will deterministically be found in $|e\rangle$, which acts as a flag state indicating that an error has occurred. This permits error detection.

However, the second issue still persists, with the oscillator acquiring an unknown phase space rotation - an error that cannot be corrected. This can be resolved by matching the dispersive shifts $\chi_e = \chi_f = \chi$ (so-called χ-matching), so that the jump operator *does* commute with the system Hamiltonian: $[\mathcal{E}, \hat{\mathcal{H}}] = \left[|e\rangle\langle f|, \hat{a}^\dagger\hat{a}\,(|e\rangle\langle e| + |f\rangle\langle f|)\right] = 0$. This ‘error-transparency’ condition (Kapit, 2018) ensures that the oscillator deterministically acquires a phase $e^{i\pi\hat{a}^\dagger\hat{a}}$, regardless of t_{err}. On detecting an error, this phase can be corrected by performing a phase rotation (or more straightforwardly by an update in the phase of the oscillator drive ‘in software’), before resetting the ancilla and repeating the measurement. Since this scheme can correct for a single ancilla relaxation error, it ensures 1^{st}-order fault tolerance.

Implementing this fault-tolerance comes at the expense of some extra technical complexity. Firstly, since the $|g\rangle \rightarrow |f\rangle$ transition is 'forbidden' via a single-photon drive on a transmon ($\langle g|\,\hat{Q}\,|f\rangle \approx 0$), either a two-photon Raman transition (Kumar *et al.*, 2016) or two successive $|g\rangle \rightarrow |e\rangle$ and $|e\rangle \rightarrow |f\rangle$ pulses are required. Secondly, the transmon readout must be able to separately distinguish the $|g\rangle$, $|e\rangle$ and $|f\rangle$ states, leading to a reduction in the signal-to-noise ratio (SNR) given the same readout hardware. Finally, χ-matching requires the application of a parametric drive on the transmon ancilla (Rosenblum *et al.*, 2018). As we shall shortly see, applying parametric drives to a transmon without inducing errors or unwanted Hamiltonian terms is a notable challenge.

Protecting against ancilla dephasing requires significantly less complexity. The jump operators describing dephasing, proportional to $|g\rangle\langle g|$, $|e\rangle\langle e|$ and $|f\rangle\langle f|$, all commute with H_{disp}, whether or not a three-level ancilla or χ-matching is employed. This ensures that the parity operation is error-transparent to ancilla dephasing, so that an error at any time can be treated the same as an error happening at the very end of the delay period. An ancilla phase-flip at this point in the sequence results in an incorrect measurement outcome but has no immediate impact on the oscillator state. Nonetheless, unflagged incorrect measurement outcomes can have knock-on impacts when the results are used to decide which subsequent operations to perform on the oscillator. To correct this to 1^{st}-order, we can repeat the measurement 3 times and majority vote, since the measurement is QND. If we are only interested in detecting the error, measuring twice and flagging events where the two results disagree is sufficient.

SNAP

The error transparency condition, while a sufficient condition for ancilla fault tolerance, is not a necessary one, as highlighted by the fault-tolerant implementation of SNAP (Reinhold *et al.*, 2020). Just as for the parity measurement, the SNAP Hamiltonian is error transparent to ancilla relaxation from $|f\rangle$ when using a three-level ancilla. Ancilla dephasing errors, on the other hand, do not commute with the always-on microwave drives, $[|f\rangle\langle f|, |g\rangle\langle f| + |f\rangle\langle g|] \neq 0$.

However, while this clearly violates error-transparency, this condition is in fact sufficient but not necessary. More generally, the scheme must (and does) satisfy path-independence (Ma *et al.*, 2020). This states that depending on the final state of the ancilla (which we can measure), the unitary enacted should always be the same, regardless of whether (and when) any errors occurred. A key factor that enables this is that the amplitude of the microwave comb addressing each photon number peak is equal, with the Hamiltonian in each photon-number subspace differing only by the phase of the drive. This ensures that for any photon number in the oscillator, the ancilla has the same probability of being found in $|f\rangle$, and so dephasing errors are not able to distinguish between different oscillator states. The constraint that the amplitude of the drive be equal on each photon number peak highlights why the photon-number-selective measurement (e.g. is $N = 0$?), with a drive applied to just a single photon-number-subspace cannot be made 1$^{\text{st}}$-order fault-tolerant to ancilla dephasing errors.

2.2.4 What sets the size of χ?

Besides stipulating that the coupling must be dispersive ($g \ll \Delta$), we have not until now discussed what the value of χ should be - an important quantity for defining the transition between the $g_{\text{bs}} \ll |\chi|$ and $g_{\text{bs}} \gtrsim |\chi|$ regimes. The optimal choice will depend on the competition between increased speed and increased degradation of the linearity or noise bias.

A stronger coupling χ allows for separation between the number-split peaks using a shorter ancilla pulse which, all else being equal, should result in higher-fidelity operations. However, stronger coupling also compromises the linearity of the oscillator mode. As seen in Eq. 2.35 the dispersive Hamiltonian is accompanied by an **inherited Kerr nonlinearity** in the oscillator $K_b \approx (g/\Delta)^4 K_T \approx \chi^2/(4K_T)$. Similarly, stronger coupling also results in **inherited relaxation** in the oscillator mode. In the dressed basis (see Eq. 2.34), the jump operator for relaxation on the bare transmon mode is transformed as

$$\begin{aligned}\sqrt{\kappa_T t}\ \hat{T} &\approx \sqrt{\kappa_T t}\ \left(\hat{t} + \frac{g}{\Delta}\hat{b}\right) \\ &= \sqrt{\kappa_T t}\ \hat{t} + \sqrt{\left(\frac{g}{\Delta}\right)^2 \kappa_T t}\ \hat{b},\end{aligned} \tag{2.45}$$

with an inherited relaxation rate $1/T_1^{\text{inherited}} = \left(\frac{g}{\Delta}\right)^2 \times 1/T_1^{\text{ancilla}} = \frac{\chi}{2K_T} \times 1/T_1^{\text{ancilla}}$ that scales linearly with the magnitude of χ (given a fixed transmon nonlinearity K_T).

The degree to which transmon dephasing leads to **inherited dephasing** in the oscillator depends on the spectrum of the noise, making a general rule less easy to come by. In the case that the dephasing noise on the transmon is Markovian, with a white spectrum, $S[\omega] =$ constant, we can use the same technique to find the dressed jump operator for ancilla dephasing:

$$\begin{aligned}\sqrt{\gamma_\phi^T t}\ \hat{T}^\dagger\hat{T} &\approx \sqrt{\gamma_\phi^T t}\ \left(\hat{t}^\dagger + \frac{g}{\Delta}\hat{b}^\dagger\right)\left(\hat{t} + \frac{g}{\Delta}\hat{b}\right) \\ &= \sqrt{\gamma_\phi^T t}\ \hat{t}^\dagger\hat{t} + \sqrt{\left(\frac{g}{\Delta}\right)^2 \gamma_\phi^T t}\ \left(\hat{t}^\dagger\hat{b} + \hat{b}^\dagger\hat{t}\right) + \sqrt{\left(\frac{g}{\Delta}\right)^4 \gamma_\phi^T t}\ \hat{b}^\dagger\hat{b},\end{aligned} \tag{2.46}$$

with the final term yielding an inherited dephasing rate $1/T_\phi^{\text{inherited}} \approx \left(\frac{g}{\Delta}\right)^4 \times 1/T_\phi^{\text{ancilla}} = \left(\frac{\chi}{2K_T}\right)^2 \times 1/T_\phi^{\text{ancilla}}$ that scales quadratically with χ. Interestingly, the central terms also lead to additional inherited decay, since photons from the long-lived oscillator mode are exchanged into the lossier transmon mode. Following Milul *et al.* (2023), this modifies the expression for inherited decay to $1/T_1^{\text{inherited}} = \left(\frac{g}{\Delta}\right)^2 \times 1/T_2^{\text{ancilla}} = \frac{\chi}{2K_T} \times 1/T_2^{\text{ancilla}}$.

Typically a more important source of transmon-induced dephasing is **photon shot noise** due to thermal fluctuations in the transmon population. This imparts a dephasing rate of

$$\gamma_\phi = \frac{\bar{n}\kappa_t\chi^2}{\kappa_t^2 + \chi^2}, \tag{2.47}$$

where $\bar{n}$ is the average thermal population in the transmon (Gambetta *et al.*, 2006). In the strong dispersive regime ($|\chi| \gg \kappa_t$), this reduces to $\gamma_\phi = \bar{n}\kappa_t$ (in other words, $1/T_\phi^{\text{inherited}} = \bar{n} \times 1/T_1^{\text{ancilla}}$) . Unlike the dephasing inherited directly from the transmon mode, this sort of dephasing can be mitigated either passively (by ensuring that the transmon is well-thermalized in the dilution refrigerator) or actively (by continuously monitoring the ancilla state to ensure it remains in $|g\rangle$) (Goldblatt *et al.*, 2024).

Given that the tolerable degree of nonlinearity and decoherence is application-specific, so too is the value of χ. For example, in GKP experiments (Eickbusch *et al.*, 2022), the large photon numbers in the codewords impose a tight constraint on the Kerr nonlinearity $K_b \lesssim 2\pi \times 1$ Hz. Likewise, in experiments designed to detect an intrinsic single-photon lifetime in the cavity that is m times larger than the transmon T_2^{ancilla}, one requires $(g/\Delta)^2 \ll 1/m$. For example, in a recent experiment by Milul *et al.* (2023), a cavity lifetime $T_1 = 25.6$ ms was measured using an ancilla with $T_2^{\text{ancilla}} = 80$ μs (i.e. $m = 320$), by ensuring that $(g/\Delta)^2 = 1.6 \times 10^{-4} \approx 1/6000$ leading to a dispersive shift of only $|\chi|/2\pi = 43$ kHz.

On the other hand, for bosonic encodings with modest photon number (e.g., binomial or dual-rail), a dispersive shift $\chi/2\pi \approx -1$ MHz is typical. With usual parameters, $T_1^{\text{ancilla}} = T_{2E}^{\text{ancilla}} \approx 100$ μs, $\bar{n} \approx 10^{-3}$ and $K_t/2\pi \approx 200$ MHz, this ensures that $1/T_1^{\text{inherited}} \approx 40$ ms, $1/T_\phi^{\text{inherited}} \approx 100$ ms and $K_b/2\pi \approx 1.25$ kHz. While the tools developed in this thesis are not specific to any particular encoding, we will use this latter $\chi/2\pi = -1$ MHz value as our benchmark.

2.3 Parametric beamsplitter control

The dispersive coupling provides us with non-Gaussian single-mode control. To extend this to non-Gaussian control over a network of many oscillators, we can supplement the dispersive coupling with a Gaussian multi-mode interaction – the time-dependent equivalent of the photonic beamsplitter, $\hat{\mathcal{H}}_{\text{bs}}$, described by Eq. 2.2.

This engineered beamsplitter interaction is an example of a ‘parametric’ process which is generated by coupling both linear oscillators to a nonlinear coupler mode and applying off-resonant microwave drives whose amplitude and phase determine $g_{\text{bs}}(t)$. The defining ‘parametric’ feature is that the coupler is not excited out of its ground state (Boyd, 2008), with the Hamiltonian containing no terms explicitly describing the coupler degrees of freedom. This feature is extremely attractive as there is no opportunity for 'jump errors' due to decoherence in the coupler to propagate to the pristine oscillator modes, obviating the need for fault-tolerant control schemes[5].

Prior implementations of this interaction, for which a detailed derivation will be provided in Chapter 3, used a transmon as a Kerr nonlinear coupler. Just as in the field of nonlinear optics, the application of two drives, whose frequencies obey $|\omega_1 - \omega_2| = |\omega_a - \omega_b|$ generates the desired beamsplitter Hamiltonian via a four-wave mixing interaction (Boyd, 2008). As we shall see, the use of a Kerr nonlinear coupler is by no means the only (or optimal) way of achieving this interaction, with beamsplitter amplitudes until now limited to $g_{\text{bs}} \lesssim 100$ kHz $\ll \chi$ (Gao *et al.*, 2018).

The time-dependent beamsplitter drive in Eq. 2.2 may be parameterized as $g_{\text{bs}}(t) = g_{\text{bs}} e^{i(\Delta t + \varphi)}$, with g_{bs} and φ as a real-valued amplitude and phase, and Δ as a frequency detuning from the resonance condition (e.g. in the case of four-wave mixing, $\Delta = (\omega_2 - \omega_1) - (\omega_b - \omega_a)$, assuming $\omega_2 > \omega_1$ and $\omega_b > \omega_a$). Applying a rotating-frame transformation to the beamsplitter Hamiltonian (Eq. 2.2) then allows us to interpret the drive detuning as a frequency shift of one of the oscillators[6],

$$\frac{\hat{\mathcal{H}}_{\text{bs}}}{\hbar} = \frac{g_{\text{bs}}}{2}\left(e^{i\varphi}\hat{a}^\dagger\hat{b} + e^{-i\varphi}\hat{b}^\dagger\hat{a}\right) - \Delta\hat{b}^\dagger\hat{b}. \tag{2.48}$$

Applying a resonant ($\Delta = 0$) drive allows us to emulate the optical beamsplitter from Chapter 1. For a single-photon dual-rail input state $|0, 1\rangle$ and a drive phase $\varphi = 0$, this enacts

[5]One caveat to this is when the coupler mode has a non-negligible thermal occupation, $\bar{n}$. While the ground state cannot suffer from relaxation ($\hat{a}$) or dephasing ($\hat{a}^\dagger\hat{a}$), it can suffer from heating ($\hat{a}^\dagger$), at the (small) rate $\bar{n}/T_1^{(\text{ancilla})}$.

[6]Note that the factor of $1/2$ in the definition of the beamsplitter Hamiltonian matches the convention in Tsunoda *et al.* (2023) but differs from that in Gao *et al.* (2018), Chapman *et al.* (2023) and Lu *et al.* (2023). This choice will greatly simplify the notation when combining the beamsplitter and dispersive Hamiltonians in Chapter 5 and places the beamsplitter Hamiltonian on equal footing with the transmon Rabi drive in Eq. 2.36.

the unitary evolution:

$$\hat{U}_{\mathsf{bs}}(t)\,|0,1\rangle = \mathsf{exp}\left(\frac{-ig_{\mathsf{bs}}t}{2}\left(\hat{a}^\dagger\hat{b}+\hat{a}\hat{b}^\dagger\right)\right)|0,1\rangle \tag{2.49}$$

$$= \cos\left(\frac{g_{\mathsf{bs}}t}{2}\right)|0,1\rangle + i\sin\left(\frac{g_{\mathsf{bs}}t}{2}\right)|1,0\rangle\,. \tag{2.50}$$

After time $t_{\mathsf{bs}} = \frac{\pi}{2g_{\mathsf{bs}}}$, this enacts the equivalent of the 50/50 beamsplitter, yielding the state $(|0,1\rangle + i\,|1,0\rangle)/\sqrt{2}$ and after $t_{\mathsf{SWAP}} = 2t_{\mathsf{bs}} = \frac{\pi}{g_{\mathsf{bs}}}$, we enact a SWAP operation, exchanging the two oscillator states, $|0,1\rangle \rightarrow i\,|1,0\rangle$.

Modulating the phase φ allows us to control the phase of the superposition generated, allowing us to reach any state on the dual-rail qubit. Furthermore, while seemingly trivial, having native access to a SWAP operation allows us to efficiently reorganize the topology of a circuit (Kivlichan *et al.*, 2018), without requiring three successive CNOT gates (as is the case for qubits).

2.4 Combined beamsplitter and dispersive operations

Having discussed the capabilities afforded by the dispersive and beamsplitter interactions by themselves, we can now turn to the combined Hamiltonian of Eq. 2.1, shown in Fig. 2.1, and described by the Hamiltonian

$$\frac{\hat{\mathcal{H}}_{\chi\mathsf{bs}}}{\hbar} = \frac{g_{\mathsf{bs}}}{2}\left(e^{i\varphi}\hat{a}^\dagger\hat{b} + e^{-i\varphi}\hat{a}\hat{b}^\dagger\right) - \Delta'\hat{b}^\dagger\hat{b} \tag{2.51}$$

where $\Delta' = \Delta - \chi\,|e\rangle\langle e|$ is a qubit-state-dependent detuning from the beamsplitter resonance.

2.4.1 Controlled SWAP

A straightforward and practically useful combined dispersive-beamsplitter operation is the controlled SWAP, or Fredkin gate (Fredkin and Toffoli, 1982; Milburn, 1989), which exchanges arbitrary quantum states in two oscillator modes, conditioned on the state of a control qubit:

$$\text{c}\hat{\text{SWAP}} \equiv |g\rangle\langle g| \otimes \hat{\mathbb{I}} + |e\rangle\langle e| \otimes \hat{\text{SWAP}}. \tag{2.52}$$

As discussed in the introduction, this operation provides a building block for conditionally routing quantum states through a network to access quantum or classical information stored in quantum Random Access Memory (Hann *et al.*, 2021; Weiss *et al.*, 2024).

In fact, when considering the effect of the dispersively-coupled qubit, the SWAP operation described in Ch. 2.3, with $\Delta = 0$, is *already* a controlled-SWAP, as was demonstrated by Gao *et al.* (2019). Exciting the qubit to $|e\rangle$ shifts Bob's oscillator frequency, and in turn the resonance frequency for the beamsplitter drive, by χ. In the $g_{\text{bs}} \ll \chi$ regime, the beamsplitter drive is therefore far off-resonant and has (almost) no impact on the oscillator states. While this sequence performs a SWAP conditioned on the transmon $|g\rangle$ state, it can be made conditional on the $|e\rangle$ state instead by choosing a beamsplitter drive detuning $\Delta = \chi$.

The fidelity under this approach is severely limited, however. For an arbitrary control state, the shorter-lived transmon must remain in a superposition for the duration of the slow beamsplitter pulse. We can estimate the infidelity due to transmon errors alone as

$$(1 - \mathcal{F})_{\text{ancilla}} \approx \frac{\pi}{g_{\text{bs}}} \times \frac{1}{T_2^{\text{ancilla}}} \gg \frac{\pi}{\chi T_2^{\text{ancilla}}} \approx 1\%. \tag{2.53}$$

An alternative approach that yields higher fidelity (under certain conditions) is to use the microwave analog of the photonic controlled-SWAP (1.2), which is shown in Fig. 2.5(b). Here, microwave-activated 50/50 beamsplitter gates take the place of physical beamsplitters and the cross-Kerr nonlinearity between the transmon and one of the oscillators enacts a controlled-phase shift. A key difference in the microwave case, however, is that the cross-Kerr nonlinearity is always on by default. An important implication of this, as we have already seen, is that the 50/50 beamsplitter is conditional on the qubit state. As a result, this state must be known before the start of every beamsplitter pulse. Provided this, we can perform transmon pulses to keep it in $|g\rangle$ during the beamsplitter pulse – in effect turning off the Kerr nonlinearity for its duration.

2.4.2 SWAP-test

An important operation in which the qubit does start in a known state is the SWAP-test, which measures the overlap between quantum states in two modes, a key primitive for proposals in quantum fingerprinting (Buhrman *et al.*, 2001) and state purification (Barenco *et al.*, 1997; Childs *et al.*, 2024). The pulse sequence (shown in Fig.2.5(c)) consists of a cSWAP sandwiched between two qubit rotations, followed by a qubit readout. To see how the SWAP-test works, we can consider its action on a generic pure state $|\Psi\rangle = |\psi\rangle\,|\phi\rangle$ up until the final measurement:

$$\begin{aligned}
\hat{Y}_{-\frac{\pi}{2}}\ \text{cSWAP}\ \hat{Y}_{\frac{\pi}{2}}\,|g\rangle\,|\psi\rangle\,|\phi\rangle &= \hat{Y}_{-\frac{\pi}{2}}\ \text{cSWAP}\ \left(\frac{|g\rangle + |e\rangle}{\sqrt{2}}\right)|\psi\rangle\,|\phi\rangle \\
&= \hat{X}_{-\frac{\pi}{2}}\ \frac{|g\rangle\,|\psi\rangle\,|\phi\rangle + |e\rangle\,|\phi\rangle\,|\psi\rangle}{\sqrt{2}} \\
&= |g\rangle\,\frac{|\psi\rangle\,|\phi\rangle + |\phi\rangle\,|\psi\rangle}{2} + |e\rangle\,\frac{|\psi\rangle\,|\phi\rangle - |\phi\rangle\,|\psi\rangle}{2} \\
&\equiv \frac{|g\rangle\,|\Psi\rangle_{\text{sym}} + |e\rangle\,|\Psi\rangle_{\text{anti-sym}}}{\sqrt{2}}.
\end{aligned} \tag{2.54}$$

The probability of measuring the ancilla in $|g\rangle$ is given by

$$\begin{aligned}
P_g &= \frac{{}_{\text{sym}}\langle\Psi|\Psi\rangle_{\text{sym}}}{2} \\
&= \frac{1}{2} + \frac{|\langle\psi|\phi\rangle|^2}{2},
\end{aligned} \tag{2.55}$$

and therefore provides a measure of the state overlap. If the two states are identical, then $P_g = 1$. Furthermore, the state after measuring $|g\rangle$ is projected into a symmetric superposition of the two initial oscillator states. This provides a means of state purification, producing a single less-noisy state from two noisier copies, since the desired (noise-free) component of the two states lives in $|\Psi\rangle_{\text{sym}}$, whereas noise (if uncorrelated between the two oscillators) is evenly divided between $|\Psi\rangle_{\text{sym}}$ and $|\Psi\rangle_{\text{anti-sym}}$ (Barenco *et al.*, 1997; Childs *et al.*, 2024; O'Brien et al., 2023).

Since the ancilla starts in a known state, $|g\rangle$, the theoretical gate sequence can be decomposed into the form shown in Fig. 2.5(c), removing the need for unconditional beamsplitters (Gao *et al.*, 2019). The final active reset of the ancilla and 50/50 beamsplitter (inside the dotted box) are only required if one wants to use the state after measurement.

2.4.3 Exponential SWAP

The same cSWAP construction forms an integral part of an entangling gate between bosonic modes, known as the exponential SWAP, or eSWAP (Gao *et al.*, 2019):

$$\text{eSWAP}\,(\theta) = \exp\left(i\theta \text{SWAP}\right) = \cos\left(\theta\right)\hat{\mathbb{I}} + i\sin\left(\theta\right)\text{SWAP} \tag{2.56}$$

This operation is particularly versatile as it entangles two qubits, each encoded in a individual bosonic mode, *independently* of the choice of encoding. Furthermore, the code-agnostic nature of both the eSWAP and the SWAP-test has been leveraged in a proposal for universal quantum computation with arbitrary multi-mode bosonic encodings (Lau and Plenio, 2016).

This sequence makes use of an important tool called the 'exponentiation gadget' (Nielsen and Chuang, 2010) that can be used to generate exponentials of involutory operators (i.e. those that square to the identity, $\hat{A}^2 = \hat{\mathbb{I}}$) from their ancilla-controlled versions. This gadget is shown in Fig. 2.5(a), into which the photonics-inspired cSWAP construction can be inserted (see Fig. 2.5(d)). Similarly to the case of SWAP-test, the ancilla state is (barring any errors) known throughout the sequence, which can therefore be reorganized to generate a practically-realizable series of gates.

The controlled SWAP on an arbitrary state was subject to both relatively slow operating speeds, set by g_{bs}, and relatively high decoherence rates, set by the transmon decoherence, therefore yielding a high infidelity (see Eq. 2.53). However, for both the SWAP-test and the eSWAP, the infidelity is now only set by the sum of the infidelities of an individual SWAP (which is slow but only subject to relatively low cavity decoherence rates) and a single-mode dispersive operation (which is faster but sensitive to transmon decoherence). Nonetheless, the limitation

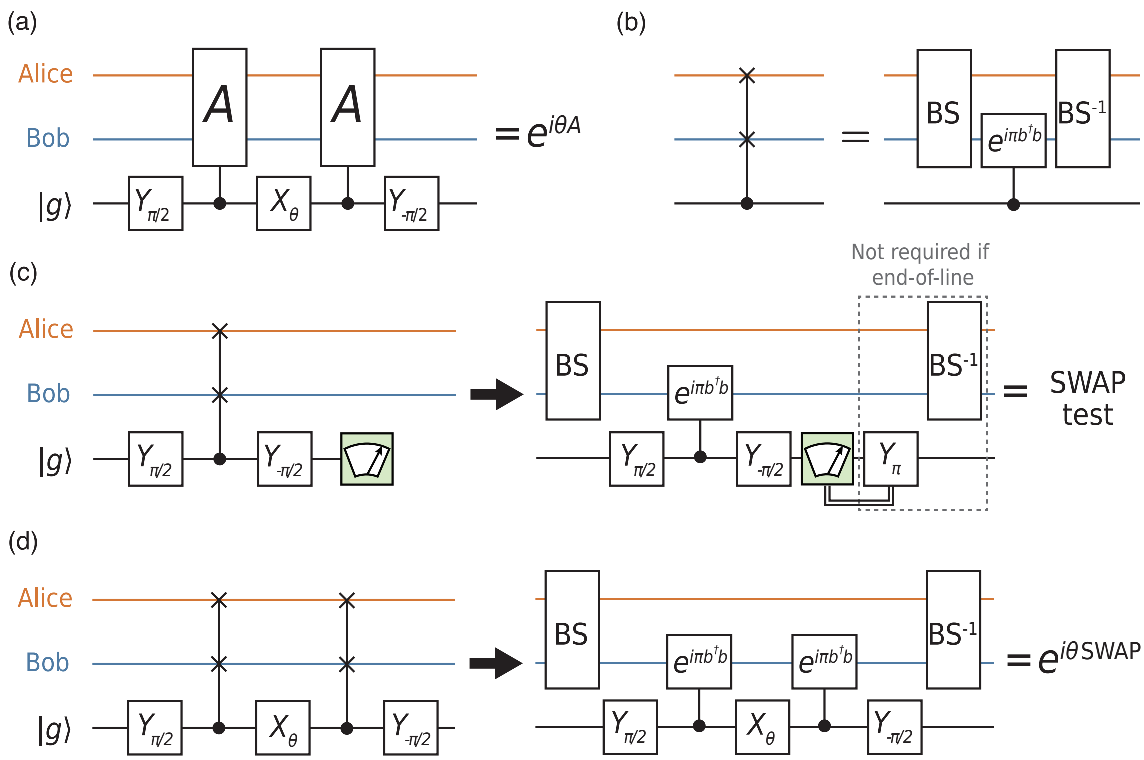

Figure 2.5: **Combined beamsplitter and dispersive operations.** (a) 'Exponentiation gadget' for generating exponentials of involutory operators from their ancilla-controlled counterparts. (b) Photonics-inspired construction of controlled-SWAP gate. (c) Basic SWAP-test circuit on the left, and broken down into its constituent parts on the right, using the construction from (b). Active feedback to reset the ancilla to $|g\rangle$ prior to performing the final 50/50 beamsplitter is only required if one wants to preserve the post-SWAP-test state. (d) Construction of eSWAP, with reordering of constituent gates to minimize errors and ensure ancilla is in $|g\rangle$ during the 50/50 beamsplitter operations.

on the beamsplitter speed sets a harsh speed limit on multi-cavity operations ($T_{\text{operation}} \gg 1\ \mu$s). **The hopeful goal is that we can achieve the best of both worlds**: with access to a fast beamsplitter that can achieve $g_{\text{bs}} > |\chi|$, as well as fault-tolerant schemes to protect against ancilla decoherence, we can achieve fast, high-fidelity, multi-cavity, non-Gaussian operations. I will start addressing the first of these points in the next chapter.

Chapter 3

Designing a three-wave mixing beamsplitter for linear oscillators

We have established that

1. combining beamsplitters and dispersive interactions allows us to perform multi-mode non-Gaussian control of bosonic modes, and

2. a general purpose bosonic processor requires operations that preserve the pristine unitary dynamics of the oscillators (introducing no unwanted nonlinear terms) while also avoiding any extra non-unitary decoherence.

Previous beamsplitter implementations using Kerr-based transmon couplers have been limited to $g_{\mathrm{bs}} \ll |\chi| \approx 2\pi \times 1$ MHz before the added noise becomes dominant, a regime in which the beamsplitter operations set a bound on the multimode gate fidelity. Finding a coupling element that allows for $g_{\mathrm{bs}} \gtrsim |\chi|$ will therefore unlock high-fidelity multimode non-Gaussian operations.

In this chapter, I explain what underlies the limitations of the transmon approach and show theoretically how another charge-driven dipole element, the SNAIL (Frattini *et al.*, 2017) (initially used for quantum limited amplifiers (Frattini *et al.*, 2018; Sivak *et al.*, 2019), but recently also to mediate interactions between transmons (Zhou *et al.*, 2023)) can overcome this. I will then provide some guidelines for how to specifically optimize a SNAIL for high-fidelity operations on high-Q linear oscillators.

3.1 Deriving a beamsplitter from a charge-driven dipole

Before we can understand the pros and cons of different couplers, we must first understand how the beamsplitter Hamiltonian emerges from a generic charge-driven dipole element (describing both the transmon and its replacement, the SNAIL). In particular, we will consider the case where the coupler potential $\hat{U}(\varphi)$ can be written as a function of a single phase variable φ and where the shunting capacitance is sufficiently large that the zero-point phase fluctuations in this variable are small, $\varphi_{\mathsf{ZPF}} \ll 1$. This allows us to Taylor expand $\hat{U}(\varphi)$ around its minimum φ_{min} and write the coupler Hamiltonian as that of a weakly anharmonic oscillator (with corresponding annihilation operator $\hat{C}$) centered at this point, to which we can add the two linear modes ($\hat{A}$ and $\hat{B}$) to which it will couple:

$$\frac{\hat{\mathcal{H}}_{\mathsf{bare}}}{\hbar} = \underbrace{\omega_A \hat{A}^\dagger \hat{A} + \omega_B \hat{B}^\dagger \hat{B} + \omega_C \hat{C}^\dagger \hat{C}}_{\hat{\mathcal{H}}_0/\hbar} + \underbrace{\sum_{n=3}^{\infty} g_n \left(\hat{C} + \hat{C}^\dagger\right)^n}_{\hat{\mathcal{H}}_{\mathsf{nl}}/\hbar}, \tag{3.1}$$

where ω_A, ω_B and ω_C are the bare frequencies of the three modes.

When we then introduce a capacitive coupling between the coupler and each of the two linear oscillators, which ideally have no direct mutual coupling between them, it yields a Hamiltonian term

$$\frac{\hat{\mathcal{H}}_{\mathsf{coupling}}}{\hbar} = -g_a \left(\hat{A} - \hat{A}^\dagger\right)\left(\hat{C} - \hat{C}^\dagger\right) - g_b \left(\hat{B} - \hat{B}^\dagger\right)\left(\hat{C} - \hat{C}^\dagger\right). \tag{3.2}$$

When $|\Delta_a| \equiv |\omega_A - \omega_C| \ll |\omega_A + \omega_C|$ and $|\Delta_b| \equiv |\omega_B - \omega_C| \ll |\omega_B + \omega_C|$, we may keep only the slowly-rotating terms in the RWA:

$$\frac{\hat{\mathcal{H}}^{(\mathsf{RWA})}_{\mathsf{coupling}}}{\hbar} = g_a \left(\hat{A}^\dagger \hat{C} + \hat{C}^\dagger \hat{A}\right) + g_b \left(\hat{B}^\dagger \hat{C} + \hat{C}^\dagger \hat{B}\right). \tag{3.3}$$

Just as when we were deriving the dispersive interaction in Chapter 2, we can then suggestively write the linear part of the Hamiltonian as a matrix,

$$\frac{\hat{\mathcal{H}}_{\mathsf{lin}}}{\hbar} = \frac{\hat{\mathcal{H}}_0}{\hbar} + \frac{\hat{\mathcal{H}}^{(\mathsf{RWA})}_{\mathsf{coupling}}}{\hbar} = \begin{pmatrix} \hat{A}^\dagger & \hat{B}^\dagger & \hat{C}^\dagger \end{pmatrix} \begin{pmatrix} \omega_A & 0 & g_a \\ 0 & \omega_B & g_b \\ g_a & g_b & \omega_C \end{pmatrix} \begin{pmatrix} \hat{A} \\ \hat{B} \\ \hat{C} \end{pmatrix} \tag{3.4}$$

which can be diagonalized to give us the dressed mode operators ($\hat{a}$, $\hat{b}$ and $\hat{c}$) in terms of their bare counterparts, as well as their eigenfrequencies (ω_a, ω_b and ω_c). In this situation, we typically wish to remain in the dispersive limit ($g_a \ll |\Delta_a|$, $g_b \ll |\Delta_b|$) so that the cavities do not inherit too much of the nonlinearity or noise of the Josephson-junction-based coupling element, in which case

$$\hat{A} \approx \hat{a} - p_a \hat{c}, \tag{3.5}$$

$$\hat{B} \approx \hat{b} - p_b \hat{c}, \tag{3.6}$$

$$\hat{C} \approx \hat{c} + p_a \hat{a} + p_b \hat{b}, \tag{3.7}$$

and

$$\omega_a \approx \omega_A + p_a^2 \Delta_a, \tag{3.8}$$

$$\omega_b \approx \omega_B + p_b^2 \Delta_b, \tag{3.9}$$

$$\omega_c \approx \omega_C - p_a^2 \Delta_a - p_b^2 \Delta_b, \tag{3.10}$$

where $p_a = g_a/\Delta_a$ and $p_b = g_b/\Delta_b$.

We then also add 1 or more drives to the coupler mode, each with an amplitude ϵ_i and a frequency ω_i. Provided that the pumps are 'stiff' such that the interaction does not alter ϵ_i, and that $|\omega_i - \omega_c| \ll |\omega_i + \omega_c|$, they can be expressed as classical drives of the form

$$\frac{\hat{\mathcal{H}}_{\mathsf{drive}}}{\hbar} = \sum_i \epsilon_i \left(\hat{c} e^{i\omega_i t} + \hat{c}^\dagger e^{-i\omega_i t} \right). \tag{3.11}$$

Following the prescription in Reinhold (2019), we can write out the full Hamiltonian as

$$\frac{\hat{\mathcal{H}}}{\hbar} = \frac{\hat{\mathcal{H}}_{\mathsf{lin}}}{\hbar} + \frac{\hat{\mathcal{H}}_{\mathsf{drive}}}{\hbar} + \frac{\hat{\mathcal{H}}_{\mathsf{nl}}}{\hbar} \tag{3.12}$$

$$= \omega_a \hat{a}^\dagger \hat{a} + \omega_b \hat{b}^\dagger \hat{b} + \omega_c \hat{c}^\dagger \hat{c} + \sum_i \epsilon_i \left(\hat{c} e^{i\omega_i t} + \hat{c}^\dagger e^{-i\omega_i t} \right) + \hat{\mathcal{H}}_{\mathsf{nl}} \left(\hat{a}, \hat{a}^\dagger, \hat{b}, \hat{b}^\dagger, \hat{c}, \hat{c}^\dagger \right) / \hbar, \tag{3.13}$$

and perform a series of unitary transformations, $\hat{\mathcal{H}} \to \hat{U}\hat{\mathcal{H}}\hat{U}^\dagger + i\hbar\dot{\hat{U}}\hat{U}^\dagger$ to remove everything but the nonlinear coupler Hamiltonian. The first transformation, $\hat{U} = e^{-i(\omega_a \hat{a}^\dagger \hat{a} + \omega_b \hat{b}^\dagger \hat{b} + \omega_c \hat{c}^\dagger \hat{c})t}$, moves us into a frame co-rotating at the three mode frequencies:

$$\frac{\hat{\mathcal{H}}}{\hbar} = \sum_i \epsilon_i \left(\hat{c} e^{i\Delta_i t} + \hat{c}^\dagger e^{-i\Delta_i t} \right) + \hat{\mathcal{H}}_{\mathsf{nl}} \left(e^{i\omega_a t}\hat{a}, e^{-i\omega_a t}\hat{a}^\dagger, e^{i\omega_b t}\hat{b}, e^{-i\omega_b t}\hat{b}^\dagger, e^{i\omega_c t}\hat{c}, e^{-i\omega_c t}\hat{c}^\dagger \right) / \hbar, \tag{3.14}$$

where $\Delta_i \equiv \omega_i - \omega_c$ is the detuning of each drive from the coupler resonance frequency. Subsequently, we can move into a 'displaced' frame via a series of unitaries

$$\hat{U}_i = \exp\left(\xi_i e^{-i\Delta_i t} \hat{c}^\dagger - \xi_i^* e^{i\Delta_i t} \hat{c} \right), \tag{3.15}$$

with $\xi_i \equiv \frac{\epsilon_i}{\Delta_i}$ to remove the drive terms:

$$\hat{\mathcal{H}} = \hat{\mathcal{H}}_{\mathsf{nl}} \left(e^{-i\omega_a t}\hat{a}, e^{i\omega_a t}\hat{a}^\dagger, e^{-i\omega_b t}\hat{b}, e^{i\omega_b t}\hat{b}^\dagger, e^{-i\omega_c t}\hat{c} + \sum_i e^{-i\omega_i t}\xi_i, e^{i\omega_c t}\hat{c}^\dagger + \sum_i e^{i\omega_i t}\xi_i^* \right). \tag{3.16}$$

The overall effect of coupling this nonlinear element to two linear oscillators (Eq. 3.2) and introducing multiple charge drives (Eq. 3.11) is therefore to leave us with the nonlinear Taylor-expanded coupler Hamiltonian $\hat{\mathcal{H}}_{\mathsf{nl}}$ in Eq. 3.1, in terms of a transformed coupler mode operator:

$$\hat{C} \to \hat{C}' \equiv e^{-i\omega_c t}\hat{c} + p_a e^{-i\omega_a t}\hat{a} + p_b e^{-i\omega_b t}\hat{b} + \sum_i e^{-i\omega_i t}\xi_i. \tag{3.17}$$

The key to obtaining a beamsplitter (or other parametric terms) is to look for terms whose phase prefactors are slowly-rotating and so cannot be removed via the RWA. When expanding the generic nonlinear Hamiltonian in Eq. 3.1, we obtain all products of three or more of the mode operators and drive amplitudes ξ_i in $\hat{C}'$. Provided that the coupler possesses an n^{th}-order nonlinearity, g_n, we can obtain a beamsplitter term proportional to $\hat{a}^\dagger\hat{b}$ by applying $n-2$ microwave drives that satisfy the frequency-matching condition

$$\sum_i s_i\omega_i = |\omega_b - \omega_a|, \tag{3.18}$$

where the sign $s_i \in \{+1, -1\}$.

With a transmon coupler, for which we can approximate the nonlinear Hamiltonian as that of a Kerr nonlinear oscillator (see Eq. 2.15),

$$\frac{\hat{\mathcal{H}}^{(4)}}{\hbar} = g_4\left(\hat{C}' + \hat{C}'^\dagger\right)^4, \tag{3.19}$$

applying 2 microwave drives such that $|\omega_2 - \omega_1| = |\omega_b - \omega_a|$ generates a Hamiltonian term $24\hbar g_4 p_a p_b \xi_1 \xi_2^* \hat{a}^\dagger\hat{b}$, where the prefactor of $24 = 4!$ counts the multiplicity of this term in the expansion. By comparing this term to beamsplitter term in Eq. 2.2, we can obtain the drive-dependent beamsplitter rate,

$$g_{\mathsf{bs}}^{(4)}(t) = 48 g_4 p_a p_b |\xi_1(t)||\xi_2(t)| e^{i(\phi_1(t)-\phi_2(t))}, \tag{3.20}$$

where $\xi_i(t) = |\xi_i(t)|e^{i\phi_i(t)}$. If one of the drives is kept fixed, the amplitude $|\xi_i|$ and phase ϕ_i of the other drive controls the amplitude and phase of g_{bs}. This is the approach taken in Gao *et al.* (2018).

Alternatively, we can choose the drive frequencies such that $|\omega_1 + \omega_2| = |\omega_b - \omega_a|$ to obtain the same $g_{\mathsf{bs}}^{(4)}(t) = 48 g_4 p_a p_b \xi_1(t)\xi_2(t)$, where the phase of $g_{\mathsf{bs}}^{(4)}$ now depends on the sum, rather than the difference, of the drive phases. There is no requirement that the two drives be different in frequency and so a (single) monochromatic drive at $\omega_1 = \omega_2 = |\omega_b - \omega_a|/2$ will also suffice. Both of these are examples of 'four-wave-mixing' processes, analogous to the four-wave-mixing nonlinear processes that are present in an optical Kerr medium with a $\chi^{(3)}$ nonlinearity (Boyd, 2008).

3.2 Limitations of a transmon-based coupler

The challenge with the four-wave mixing approach is that many other *non*-beamsplitter terms also emerge from the coupler Hamiltonian

$$\frac{\hat{\mathcal{H}}^{(4)}}{\hbar} = \frac{g_{\text{bs}}\hat{a}^\dagger\hat{b} + g_{\text{bs}}^*\hat{a}\hat{b}^\dagger}{2} + \boxed{\frac{\hat{\mathcal{H}}^{(4)}_{\text{Stark}}}{\hbar}} + \boxed{\frac{\hat{\mathcal{H}}^{(4)}_{\text{Kerr}}}{\hbar}} + \boxed{\frac{\hat{\mathcal{H}}^{(4)}_{\text{disp}}}{\hbar}} + \boxed{\ldots}, \tag{3.21}$$

where the blocks in red group together similar problematic terms, which will be described one-by-one in the following sections. $\hat{\mathcal{H}}^{(4)}_{\text{Stark}}$ contains drive-induced frequency shifts of the modes, $\hat{\mathcal{H}}^{(4)}_{\text{Kerr}}$ contains induced self-Kerr and cross-Kerr nonlinear terms in the previously-linear oscillators, $\hat{\mathcal{H}}^{(4)}_{\text{disp}}$ contains coupler-state-dependent frequency shifts of the linear oscillators, and the remaining non-beamsplitter terms are indicated by the ellipsis. While some of the problematic terms affect operations coupling qubits (see the discussion in Zhou (2024)), several are unique to, or more problematic for, multiphoton bosonic states (Zhang *et al.*, 2019).

Unwanted multiphoton transitions

Whereas most terms emerging from the expansion of Eq. 3.19 will not be resonant (i.e. slowly-rotating) for a given choice of mode and drive frequencies, an unwanted resonant term can lead to unitary and non-unitary dynamics, both of which could be harmful. As an example of the former, a resonant $\xi_1\xi_2^*(\hat{a}^\dagger)^2$ term (when $\omega_1 + \omega_2 = 2\omega_a$) will lead to squeezing of the $\hat{a}$ mode, distorting the stored state. As an example of the latter, a resonant $\xi_1\hat{a}(\hat{l}^\dagger)^2$ term (when $\omega_1 = 2\omega_l - \omega_a$) will exchange one photon in the long-lived Alice mode for two photons in a lossy mode, $\hat{l}$ (e.g., a low-Q mode in the drive line.) These photons are then rapidly dissipated, so this interaction can be viewed as an enhancement of the Alice loss rate.

It is therefore vital to keep track of and avoid all unwanted transitions. However, there are a number of factors that complicate this for a transmon. When using four-wave mixing, the expansion of Eq. 3.19 yields all permutations of four elements picked from the two microwave drives (ξ_1 and ξ_2) and the three system modes ($\hat{a}$, $\hat{b}$ and $\hat{c}$), as well as all of their complex conjugates. This results in a large number of transitions to keep track of, with $(2 \times 5)^4 = 10^4$ possible Hamiltonian terms. When the weaker terms at higher order in the Taylor expansion of the transmon potential are included, this number increases significantly further. On top of this, as described in Xiao *et al.* (2023), there are also 'cascaded processes' where the output photon(s) from one mixing process act as the input for a secondary mixing process. There are therefore a forest of frequencies to avoid.

Stark shifts

More importantly, this 'forest' of frequencies is also moving. Regardless of the choice of the system frequencies, certain terms are always resonant – those in which operators always appear alongside their conjugate. In particular, this includes a Stark shift, or drive-power-dependent frequency shift, of all the mode frequencies:

$$\frac{\hat{\mathcal{H}}^{(4)}_{\text{Stark}}}{\hbar} = \Delta^{(a,4)}_{\text{Stark}} \hat{a}^\dagger \hat{a} + \Delta^{(b,4)}_{\text{Stark}} \hat{b}^\dagger \hat{b} + \Delta^{(c,4)}_{\text{Stark}} \hat{c}^\dagger \hat{c} \tag{3.22}$$

$$\approx 24 g_4 \left(|\xi_1|^2 + |\xi_2|^2\right) \left(p_a^2 \hat{a}^\dagger \hat{a} + p_b^2 \hat{b}^\dagger \hat{b} + \hat{c}^\dagger \hat{c}\right) \tag{3.23}$$

While the shifts of the cavity frequencies, suppressed by a factor p_i^2, are relatively small, the shift of the coupler frequency can be significant. As the beamsplitter drive amplitude is ramped up, this presents a chance for the coupler frequency to become resonant with an unwanted transition. We can relate the magnitude of the transmon Stark shift (Eq. 3.23) to the beamsplitter rate (Eq. 3.20):

$$\frac{\Delta^{(c,4)}_{\text{Stark}}}{\left|g^{(4)}_{\text{bs}}\right|} = \frac{|\xi_1|^2 + |\xi_2|^2}{2|\xi_1||\xi_2|} \frac{1}{p_a p_b} \geq \frac{1}{p_a p_b} \gtrsim 400. \tag{3.24}$$

To achieve $g_{\text{bs}} \gtrsim |\chi| \approx 2\pi \times 1$ MHz, there is therefore a large region of frequency-space (> 400 MHz) that must be clear of frequency collisions. In the experiments of Gao *et al.* (2018), where the drive and mode frequencies were closely spaced, this combination of Stark shift and an unwanted $\xi_1 \hat{a} (\hat{c}^\dagger)^2$ resonance limited $g_{\text{bs}} < 2\pi \times 80$ kHz before drive-photon assisted absorption of cavity photons degraded the single-photon lifetime to limit the operation fidelity.

Inherited Kerr terms

The Stark shift is not the only always-resonant term that causes problems though! The four-wave mixing interaction also gives rise to static Kerr terms in the linear oscillator modes[1]:

$$\frac{\hat{\mathcal{H}}^{(4)}_{\mathsf{Kerr}}}{\hbar} = \frac{K_a^{(4)}}{2}\hat{a}^\dagger\hat{a}^\dagger\hat{a}\hat{a} + \frac{K_b^{(4)}}{2}\hat{b}^\dagger\hat{b}^\dagger\hat{b}\hat{b} + \chi_{ab}^{(4)}\hat{a}^\dagger\hat{a}\hat{b}^\dagger\hat{b} \tag{3.25}$$

$$\approx \left(\frac{12g_4p_a^4}{2}\right)\hat{a}^\dagger\hat{a}^\dagger\hat{a}\hat{a} + \left(\frac{12g_4p_b^4}{2}\right)\hat{b}^\dagger\hat{b}^\dagger\hat{b}\hat{b} + \left(24g_4p_a^2p_b^2\right)\hat{a}^\dagger\hat{a}\hat{b}^\dagger\hat{b} \tag{3.26}$$

The first two 'self'-Kerr terms, $K_a^{(4)}$ and $K_b^{(4)}$, we already saw in Chapter 2 as a byproduct of the dispersive interaction. The coupler transmon therefore provides an extra contribution to the nonlinearity of the linear modes on top of what a coupled ancilla transmon would provide. The cross-Kerr term is especially problematic though, and can be viewed as an always-on ZZ coupling between the modes, applying a phase update to one of the modes dependent on the photon number in the other, even when idling. Residual ZZ couplings between conventional qubits have proven to be an important contribution to two-qubit error rates, necessitating the use of tunable couplers to mediate operations between transmon qubits (Stehlik *et al.*, 2021; Yan *et al.*, 2018) and between fluxonium qubits (Ding *et al.*, 2023). Moving to encoded multiphoton states in oscillators, this problem becomes more acute, with the strength of the entangling coupling scaling with the product of the photon number in each mode.

Oscillator-coupler dispersive shift

Finally, we also acquire a conventional dispersive coupling term between the coupler mode and each of the linear modes

$$\frac{\hat{\mathcal{H}}^{(4)}_{\mathsf{disp}}}{\hbar} = \chi_{ac}^{(4)}\hat{a}^\dagger\hat{a}\hat{c}^\dagger\hat{c} + \chi_{bc}^{(4)}\hat{b}^\dagger\hat{b}\hat{c}^\dagger\hat{c} \tag{3.27}$$

$$\approx \left(24g_4p_a^2\right)\hat{a}^\dagger\hat{a}\hat{c}^\dagger\hat{c} + \left(24g_4p_b^2\right)\hat{b}^\dagger\hat{b}\hat{c}^\dagger\hat{c}. \tag{3.28}$$

[1] The coupler itself will also retain its Kerr nonlinearity, $K_c \approx 12g_4$, with its resonance frequency shifted more and more when excited to states $|e\rangle$ and higher. Since we do not explicitly intend to access these coupler states, this term is not *per se* as problematic as others, although these shifted frequencies do add to the 'forest' of multiphoton resonance frequencies we wish to avoid.

When idling, any heating of the coupler mode due to its finite temperature can therefore propagate to *correlated* dephasing of both oscillator modes. Furthermore, when the coupler is driven, these terms will shift the beamsplitter resonance condition (Eq. 3.18) by an amount $\chi_{\text{DR}} = \chi_{ac} - \chi_{bc}$ (labelled this way because it represents the frequency shift of a dual rail (DR) qubit encoded in the two oscillators). Unless $\chi_{\text{DR}} \ll g_{\text{bs}}$, this shift will ensure the beamsplitter is off-resonance and photons stop exchanging, leading to a SWAP-like error in the bosonic modes. Importantly, even at zero temperature, relaxation and dephasing of the coupler mode can lead to hopping between states when driven, an effect known as 'quantum heating' (Zhang *et al.*, 2019).

Crucially, for a Kerr coupler, both the term we want (the beamsplitter) and the terms we do not want depend on the same parameter, g_4. It is therefore not possible for us to suppress the problematic terms by reducing the g_4 nonlinearity without simultaneously suppressing g_{bs}.

3.3 Three-wave mixing in the abstract

The challenges presented are not unique to beamsplitters between bosonic modes. The problem of multiphoton resonances in strongly driven Kerr oscillators, for example, has been well-studied in the context of both fast transmon readout (Cohen *et al.*, 2023; Khezri *et al.*, 2023; Sank *et al.*, 2016; Shillito *et al.*, 2022; Xiao *et al.*, 2023) and Kerr cat qubits (Frattini *et al.*, 2024; Venkatraman *et al.*, 2022), informing the optimal choice of mode frequencies.

Three-wave mixing in amplifiers

Similarly, producing a strong parametric interaction without spurious Kerr terms is central to developing quantum-limited parametric amplifiers with large dynamic range, where one wants to generate a strong squeezing interaction $\propto (\hat{c}^2 + (\hat{c}^\dagger)^2)$ for as large an input field amplitude as possible. A charge-driven single-junction or a DC SQUID can act as a Josephson parametric amplifier (JPA), generating the desired squeezing via a four-wave-mixing interaction ($\omega_1 + \omega_2 \approx 2\omega_c$) (Yurke *et al.*, 1989), but also generates a Stark shift that limits the maximum input amplitude that can be used before the mode is shifted off resonance (Kochetov and Fedorov, 2015).

An alternative, more successful approach, is to modulate the external flux through the loop of the (DC SQUID-based) JPA with a monochromatic drive to achieve the same squeezing interaction via a three-wave mixing interaction ($\omega_p \approx 2\omega_c$) (Yamamoto *et al.*, 2008). Similarly, the Josephson parametric converter (JPC)(Bergeal *et al.*, 2010a,b) provides a quadrupolar multi-mode circuit mediating a *charge-driven* three-wave mixing interaction. Nonetheless, while these circuits use a different degree of nonlinearity to generate the gain, the 4$^{\text{th}}$-order nonlinearity still remains, with the Stark shift limiting the achievable dynamic range (Liu *et al.*, 2017).

One solution leveraged by both the SNAIL parametric amplifier (SPA; Frattini *et al.* (2018)) and the Josephson array mode parametric amplifier (JAMPA; Sivak *et al.* (2020)), and which will inspire the solution to our beamsplitter problem, is to use a charge-driven element[2] described by a single phase degree of freedom φ, that enables three-wave-mixing interactions via a non-zero g_3 while also suppressing $g_{n>4}$. In this case, to first order in perturbation theory, there will be *no* always-on RWA terms (since this requires an even-order nonlinearity) and so no Stark shift. In practice, when considering the g_3 term up to second order in perturbation theory, it generates both an anharmonicity (frequency shift per coupler excitation)

$$K_c^{(3)} \approx 12\left(g_4 - 5\frac{g_3^2}{\omega_c}\right), \tag{3.29}$$

and a Stark shift (frequency shift per normalized drive power) [3]

$$\Delta_{\text{Stark}}^{(c,3)} \approx 24|\xi|^2\left(g_4 - \frac{9g_3^2}{2\omega_c}\right). \tag{3.30}$$

Nonetheless, with sufficient control over g_3 and g_4, we can get one or the other of these terms to completely cancel, allowing us to decouple the term we want from those we do not.

[2]Other solutions to the beamsplitter problem based on flux-driven elements have also recently been developed (Lu *et al.*, 2023; Maiti *et al.*, 2024). These flux-driven approaches will be discussed and compared to our charge-driven solution in Ch. 3.6.

[3]The reader should beware that the form of the drive in Eq. 3.11 is not valid for an SPA since the drive is very off-resonant. Correcting for this will change the expression for ξ.

Three-wave mixing beamsplitter

The same idea may be applied to the beamsplitter, where a single drive at $\omega_p = |\omega_b - \omega_a|$ generates a beamsplitter, with a rate approximated at low drive amplitudes as

$$g^{(3)}_{\mathsf{bs}}(t) \approx 12 g_3 p_a p_b \xi(t). \tag{3.31}$$

With only a single drive and a lower-order three-wave mixing expansion $\hat{\mathcal{H}}^{(3)}/\hbar = g_3(\hat{C}' + \hat{C}'^{\dagger})^3$, this yields 'only' $(2\times 4)^3 = 216$ possible terms, much fewer than with 4-wave mixing.

Besides just α_c and $\Delta^{(c)}_{\mathsf{Stark}}$ of the coupler mode (whose expressions are still given by Eq. 3.29 and Eq. 3.30), we are now also able to suppress the 4$^{\text{th}}$-order terms in the *multimode* Hamiltonian by setting $g_4 \approx 0$. However, it is again important to include the second-order-perturbation corrections due to the required g_3 nonlinearity. We find an oscillator self-Kerr:

$$K^{(3)}_j \approx p_j^4 \left(12 g_4 - 18 g_3^2 \left(\frac{2\omega_c}{4\omega_j^2 - \omega_c^2} + \frac{4}{\omega_c}\right)\right) + K^{(\mathsf{t})}_j, \tag{3.32}$$

where $K^{(\mathsf{t})}_j$ is the fixed additional contribution to mode j $(= a$ or $b)$ from a coupled transmon ancilla (as discussed in Ch. 2), an always-on ZZ-like cross-Kerr:

$$\chi^{(3)}_{ab} \approx p_a^2 p_b^2 \left(24 g_4 + 36 \frac{g_3^2}{\tilde{\omega}}\right) \tag{3.33}$$

where

$$\tilde{\omega} \equiv \left(\frac{1}{\omega_a - \omega_b - \omega_c} + \frac{1}{-\omega_a + \omega_b - \omega_c} + \frac{1}{\omega_a + \omega_b - \omega_c} + \frac{1}{-\omega_a - \omega_b - \omega_c}\right)^{-1}, \tag{3.34}$$

and dispersive shifts between the coupler and the individual oscillators

$$\chi^{(3)}_{jc} = 24 p_j^2 \left(g_4 + \frac{6 g_3^2 \omega_c}{\omega_j^2 - 4\omega_c^2}\right). \tag{3.35}$$

These corrections indicate that setting $g_3 \neq 0$ and $g_4 = 0$ is not exactly the condition required to remove all unwanted terms. We have however broken the linear relationship between the desired and undesired terms, and so by tuning g_3 and g_4 in tandem we should be able to achieve large g_{bs} while suppressing particular spurious terms. We should note though that the g_3 enters with a different prefactor in each term. As we shall see, tuning g_3 and g_4 will therefore not precisely null all terms *simultaneously* for a generic choice of mode frequencies, although it does present the opportunity to switch between different operating points that optimize for different properties.

3.4 SNAIL as a three-wave-mixer

The SNAIL element that permits the required 3rd order nonlinear potential (shown in Fig. 3.1(b)) consists of a single small junction with Josephson energy αE_{J} in parallel with $M > 1$ larger shunting junctions (E_J), enclosing a loop through which an external magnetic field Φ_{ext} may be applied. In order to optimize this circuit for use as a coupler, we must know how the different circuit parameters affect the different orders of the Taylor expansion in Eq. 3.1. We therefore briefly follow Frattini (2021) to relate these parameters.

In the usual regime where the junction capacitances are small and the resulting self-resonance frequencies very high, we can express the SNAIL potential in terms of a single degree of freedom, the phase drop $\hat{\varphi}_s$ across the entire element, which is divided across each junction equally:

$$\hat{U}_{\mathsf{SNAIL}}\left(\hat{\varphi}_s\right) = -\alpha E_{\mathrm{J}} \cos \hat{\varphi}_s - M E_{\mathrm{J}} \cos \frac{\varphi_{\mathsf{ext}} - \hat{\varphi}_s}{M} = E_{\mathrm{J}} \sum_{n=2} \frac{c_n}{n!} \left(\hat{\varphi}_s - \varphi_{\mathsf{min}}\right)^n , \tag{3.36}$$

where $\varphi_{\mathsf{ext}} = 2\pi \times \frac{\Phi_{\mathsf{ext}}}{\Phi_0}$. In the final equality, we expand the potential about its minimum φ_{min}, which satisfies the transcendental equation

$$\alpha \sin \varphi_{\mathsf{min}} + \sin \frac{\varphi_{\mathsf{min}} - \varphi_{\mathsf{ext}}}{M} = 0 \tag{3.37}$$

Generally φ_{min} must be found numerically but $\varphi_{\mathsf{min}} = 0$ at $\varphi_{\mathsf{ext}} = 2n\pi$ with $n \in \mathbb{Z}$ and $\varphi_{\mathsf{min}} = \pi$

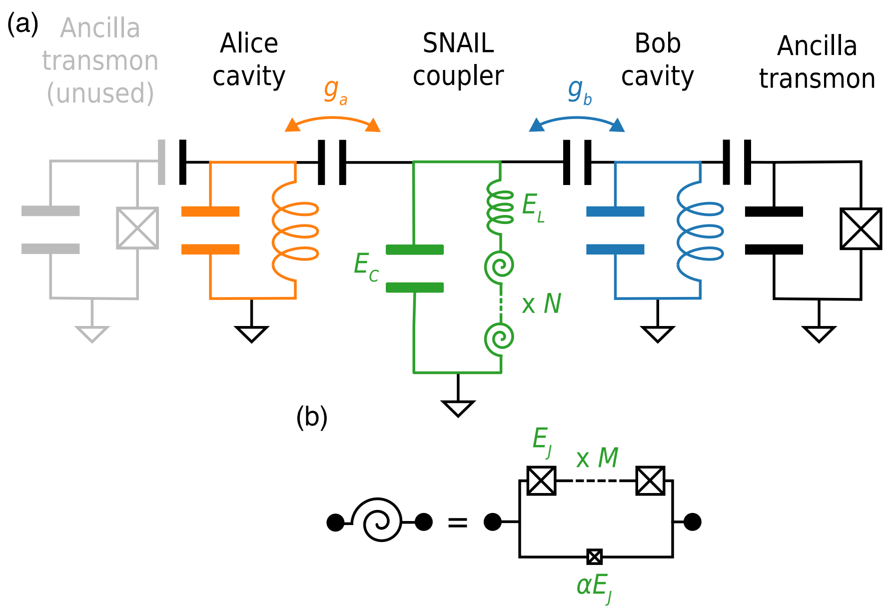

Figure 3.1: **SNAIL coupler circuit** (a) Circuit element representation of experimental setup containing a SNAIL coupler (green) capacitively coupled to two superconducting cavities, Alice (orange) and Bob (blue), with coupling strengths g_a and g_b, respectively. Each cavity is capacitively coupled to a transmon ancilla, although Bob's will be predominantly used. The coupler circuit contains a capacitor with charging energy E_{C}, in parallel with a linear inductance (energy E_L) and N arrayed SNAIL elements. In our experiment, $N = 1$. (b) SNAIL element, consisting of M large junctions (Josephson energy E_{J}) in parallel with 1 small junction (αE_J). In our experiment, $M = 3$. The loop formed by these junctions is threaded by a magnetic flux Φ_{ext}.

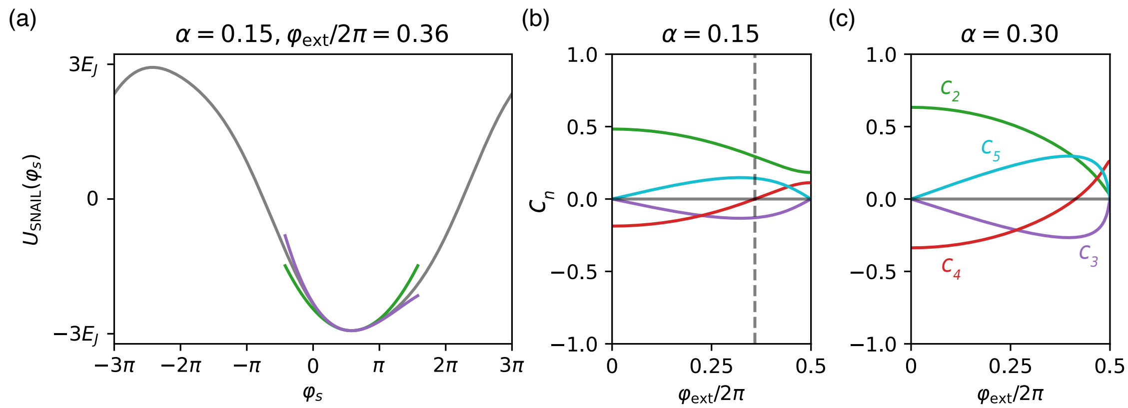

Figure 3.2: **Nonlinear coefficients of an individual SNAIL.** (a) Example potential $U_{\text{SNAIL}}(\varphi_s)$ for a single SNAIL with $M = 3$ shunting junctions and a junction ratio $\alpha = 0.15$, as a function of the superconducting phase across the whole element φ_s (grey). The external flux φ_{ext} is chosen so $c_4 = 0$. Taylor expansion around φ_{min} is overlaid up to c_2 (green) and up to c_3 (purple). (b) Nonlinear coefficients c_n for $\alpha = 0.15$. (c) Nonlinear coefficients c_n for $\alpha = 0.30 \approx 1/M$, showing large values and sharper variations close to $\varphi_{\text{ext}} = \pi$.

at $\varphi_{\text{ext}} = (2n+1)\pi$. The Taylor coefficients $c_n = \frac{1}{E_J}\frac{d^nU}{d\hat{\varphi}_s^n}\big|_{\varphi_s=\varphi_{\text{min}}}$ can then be found:

$$c_2 = +\alpha\cos\varphi_{\text{min}} + \frac{1}{M}\cos\frac{\varphi_{\text{min}} - \varphi_{\text{ext}}}{M} \qquad c_3 = -\alpha\sin\varphi_{\text{min}} - \frac{1}{M^2}\sin\frac{\varphi_{\text{min}} - \varphi_{\text{ext}}}{M} \tag{3.38}$$

$$c_4 = -\alpha\cos\varphi_{\text{min}} - \frac{1}{M^3}\cos\frac{\varphi_{\text{min}} - \varphi_{\text{ext}}}{M} \qquad c_5 = +\alpha\sin\varphi_{\text{min}} + \frac{1}{M^4}\sin\frac{\varphi_{\text{min}} - \varphi_{\text{ext}}}{M} \tag{3.39}$$

These coefficients are shown in Fig. 3.2(b-c) for $M = 3$, with $\alpha = 0.15$ and $\alpha = 0.30$. These highlight a few general properties:

- Odd c_n are antisymmetric about $\varphi_{\text{ext}} = m\pi$ and vanish at these points,
- Even c_n are symmetric about $\varphi_{\text{ext}} = m\pi$,
- For $\alpha > 1/M^{n-1}$, even c_n pass through 0 for some φ_{ext}.
- For fixed α, odd c_n have the same shape, multiplied by a scale factor that alternates sign.

Generally speaking, we can consider an array of N of these SNAIL elements in series, along with a series linear (or 'geometric') inductance L (with associated energy $E_L = \frac{\varphi_0^2}{2L}$), and a large shunting capacitance C (with associated energy $E_C = \frac{e^2}{2C}$). A key insight of Frattini *et al.*

(2018) was that introducing a linear inductance, such that the phase drop across the entire coupler φ is divided between the linear and Josephson inductances, does not simply apply a scale factor to the c_n but changes their flux dependence. Perhaps surprisingly, the minimum of the potential with respect to φ remains φ_{min}. The Taylor coefficients $\tilde{c}_n$ of the expansion around this point,

$$\hat{U}(\hat{\varphi}) = E_{\mathrm{J}} \sum_{n=2} \frac{\tilde{c}_n}{n!} \left(\hat{\varphi} - \varphi_{\text{min}}\right)^n \tag{3.40}$$

can be related to the Taylor coefficients for a single SNAIL element c_n via

$$\tilde{c}_2 = \frac{p}{N} c_2 \qquad \tilde{c}_3 = \frac{p^3}{N^2} c_3 \tag{3.41}$$

$$\tilde{c}_4 = \frac{p^4}{N^3} \left(c_4 - \frac{3c_3^2}{c_2} (1-p) \right) \qquad \tilde{c}_5 = \frac{p^5}{N^4} \left(c_5 - \frac{10 c_4 c_3}{c_2} (1-p) + \frac{15 c_3^3}{c_2^2} (1-p)^2 \right), \tag{3.42}$$

where $p = NL_s/(L + NL_s)$ is the nonlinear inductive participation, or the fraction of the inductance coming from the junctions, $NL_s = NL_J/c_2$. We can relate the $\tilde{c}_n$ to the nonlinear coefficients in Eq. 3.1 via

$$\hbar g_n = \frac{E_J \tilde{c}_n \left(\varphi_{\text{ZPF}}\right)^n}{n!}, \tag{3.43}$$

where

$$\varphi_{\text{ZPF}} = \left(\frac{2E_{\mathrm{C}}}{\tilde{c}_2 E_{\mathrm{J}}} \right)^{\frac{1}{4}}, \tag{3.44}$$

and to the 'bare' coupler frequency in Eq. 3.1 via

$$\omega_C = \sqrt{8 E_{\mathrm{C}} E_{\mathrm{J}} \tilde{c}_2}. \tag{3.45}$$

Importantly, the measured frequency (excluding any dressing with other modes) will also be Lamb shifted by the g_3 and g_4 terms in the Hamiltonian,

$$\omega_C^{(\text{meas})} \approx \omega_C + \underbrace{12 g_4 - \frac{30 g_3^2}{\omega_C}}_{\omega_C^{(\text{Lamb})}} \tag{3.46}$$

The capacitively shunted SNAIL (from now on referred to just as SNAIL) therefore provides a flux-tunable circuit with operating points where g_3 (controlling g_{bs}) is large but where g_4 (contributing to unwanted terms) can at the same time be made positive, negative or zero.

What might limit g_{bs}?

Knowing how the nonlinear parameters relate to the physical parameters of the SNAIL, we can ask, provided that we do not run into any multiphoton transitions, what could limit the g_{bs} we can achieve, or lead to a nonlinear relationship between drive amplitude and the desired interaction?

One set of unavoidable terms constitute higher-order corrections to the beamsplitter rate

$$\frac{\hat{\mathcal{H}}_{\text{corr}}}{\hbar} = \sum_{m=2} g_{\text{bs}}^{(2m+1)} \hat{a}^\dagger \hat{b} + \left(g_{\text{bs}}^{(2m+1)}\right)^* \hat{a}\hat{b}^\dagger \tag{3.47}$$

where

$$g_{\text{bs}}^{(2m+1)} = \frac{2(2m+1)!}{m!(m-1)!} g_{2m+1} p_a p_b |\xi|^2 \xi. \tag{3.48}$$

In particular, the lowest order correction $g_{\text{bs}}^{(5)} = 120 g_5 p_a p_b |\xi|^2 \xi$ contributes either negatively or positively (depending on the sign of g_5 relative to g_3[4]) with an amplitude that scales as the drive amplitude cubed. The point at which the linear and cubic contributions become equal can be thought of as a 3rd-order intercept point (IP_3), which occurs at

$$|\xi|_{\text{IP}_3} = \sqrt{\frac{|g_3|}{10|g_5|}} = \sqrt{\frac{2|\tilde{c}_3|}{|\tilde{c}_5|\varphi_{\text{ZPF}}^2}} \approx \sqrt{\frac{2|c_3|N^2}{p^2|c_5|\varphi_{\text{ZPF}}^2}} \approx \frac{\sqrt{2}N}{p\varphi_{\text{ZPF}}} \tag{3.49}$$

where in the second-to-last equality, we have assumed that $p \approx 1$. This provides a crude approximation for the critical drive amplitude $|\xi|_{\text{crit}}$, similar to that provided in Frattini (2021). If we substitute this value into Eq. 3.31, we can very roughly approximate the maximum possible scale of g_{bs} as

$$g_{\text{bs}}^{\text{IP}_3} \equiv 12 g_3 p_a p_b |\xi|_{\text{IP}_3} \approx \sqrt{2}\omega_c p \left|\frac{c_3}{c_2}\right| p_a p_b, \tag{3.50}$$

a value that is independent of N, and where $|c_3/c_2| = \mathcal{O}(1)$. Note that this expression is likely

[4] When $p = 1$, g_3 and g_5 always have opposite sign, but when $p \neq 1$ this need not be the case.

an *overestimate* of the maximum g_{bs} since it extrapolates the linear contribution out to the point ($|\xi|_{\mathsf{IP}_3}$) where it has been completely overcome by the cubic contribution.

3.5 Optimizing a SNAIL circuit for high-Q bosonic modes

In the context of parametric amplification, Frattini (2021) provides clear guidelines for optimizing the choice of the different circuit parameters to maximize the strength of the parametric interaction, among them:

- Minimize linear inductance L to keep p as close to 1 as possible,
- Array N SNAILs to dilute the nonlinearity and commensurately pump harder,
- Impedance-match the pump port for optimal pump delivery at the pump frequency while protecting against leakage at the signal frequency.

While these guidelines remain mostly valid for beamsplitting, there are some new considerations when the coupled modes must remain highly coherent.

Optimizing N: Arraying

Besides using three- rather than four-wave mixing, a key technique for engineering high dynamic-range standing-wave amplifiers is to array the nonlinear elements, whether they are SQUIDs (Eichler and Wallraff, 2014; Planat *et al.*, 2019; Winkel *et al.*, 2020) or SNAILs (Frattini *et al.*, 2018; Sivak *et al.*, 2020). From the perspective of achieving a linear relationship between drive amplitude and the desired Hamiltonian interaction, our ideal coupler is as pure as possible a g_3, with successively higher g_n much smaller than those before it. Looking at Eq. 3.43, and noting that (away from special flux points) $|c_n| = \mathcal{O}(1)\ \forall\ n$

$$\frac{g_{n+1}}{g_n} \propto \frac{\varphi_{\mathsf{ZPF}}}{N}, \tag{3.51}$$

provided that $p \approx 1$.

One way of suppressing the higher-order nonlinearities is to increase $E_{\mathrm{J}}/E_{\mathrm{C}}$ while keeping their product fixed so that φ_{ZPF} reduces and ω stays the same. Physically, this corresponds to increasing the shunt capacitance and decreasing the inductance of the junctions, and is a good strategy provided that we can keep $p \approx 1$. Alternatively, we can go from a single SNAIL to an array of N SNAILs while simultaneously changing the junction sizes such that $E_{\mathrm{J}} \to NE_{\mathrm{J}}$ in order to keep both ω_C and φ_{ZPF} fixed. This approach (again, provided that $p \approx 1$) reduces unwanted terms proportional to g_4 by a factor of N^2 (and those proportional to g_5 by N^3) while only reducing g_3 by a factor of N. In order to achieve the same beamsplitter rate however, one must simultaneously increase the pump amplitude by a factor of N (or the power by a factor of N^2).

Unlike in an amplifier however, the presence of high-Q oscillator modes coupled to the SNAIL makes it more challenging to drive the coupler more strongly without either a) Purcell-limiting the oscillator modes (Purcell, 1946; Reed *et al.*, 2010), or b) delivering too high an active heat load to the mixing chamber of the dilution refrigerator. As we shall see in Chapter 4, engineering a good on-chip pump port filter is therefore key to achieving a high-fidelity beamsplitter. Nonetheless, increasing the pump power presents the separate issue that while the SNAIL g_n are reduced, the g_4 of the ancilla transmons remains high. Higher pump powers therefore make crosstalk to these modes more likely.

In initial experimental attempts, a device with an $N = 3$ array was used to suppress the g_4 nonlinearity by a factor of 9. However, due to a weak spurious coupling between the pump port and the ancilla transmon, the required 9-fold increase in pump power caused significant frequency shifts of the ancilla mode. Switching to an unarrayed device with $N = 1$ mitigated this issue. In conclusion, while there is a theoretical benefit to arraying, one needs to take special care of microwave hygiene to be able to access this benefit.

Optimizing the number of SNAIL shunt junctions, M

As in the SPA, the choice of M is primarily guided by fabrication considerations. The ratio $|c_3/c_5| = M^2/(M^2 + 1)$, which we would ideally like to be as large as possible, depends incredibly weakly on M, with only a 20% difference between $M = 2$ and the $M \to \infty$ limit. Since $\alpha < 1/M$, for larger M, the ratio between the two junction inductances needs to be made increasingly small. Since both the small and large junctions are fabricated in the same run, they will be made with the same critical current density. As a result, increasingly small α requires increasingly extreme ratios in the overlap areas of the small and large junctions, making fabrication challenging below $\alpha \lesssim 0.05$. Meanwhile, angle evaporation puts a constraint on the symmetry of the device, requiring an odd M. One possible reason to avoid large M is that each extra junction in the array introduces a new high-frequency mode in the spectrum to which we might lose energy. We therefore choose the smallest odd number of junctions greater than 1 [5], $M = 3$.

Optimizing α

Another important parameter to choose is the ratio of junction sizes, α. This value must remain below $1/M$ for the coupler potential to contain a single minimum, and avoid the hysteretic behavior of the C-shunt flux qubit (Yan *et al.*, 2016). As $\alpha \to 1/M$, the SNAIL behaves similarly to a symmetric SQUID, with a steeper frequency response with respect to flux, as well as stronger nonlinearities (see Fig. 3.2).

In the context of qubit readout, amplifiers need to provide low-Kerr and large gain at a very specific frequency, the readout resonator frequency, which may be difficult to predict precisely in advance. This encourages the use of a relatively small $\alpha \approx 1/M^2$ to smooth out the variation in SNAIL parameters with frequency. Couplers for beamsplitting however, do not need to park at a very specific frequency to produce the required interaction. This provides more flexibility in the choice of α.

Provided that we are at the flux point where $g_4 \approx 0$, one way of increasing the value of g_{bs} for the same pump amplitude $|\xi|$ and oscillator participations p_a and p_b is to use a larger α to increase g_3. One important drawback is that the corresponding steeper frequency response $|d\omega_C/d\varphi_{\text{ext}}|$ makes the coupler mode more sensitive to flux noise, which will in turn be inherited

[5] $M = 1$ corresponds to a DC SQUID which does not possess a 3rd-order nonlinearity.

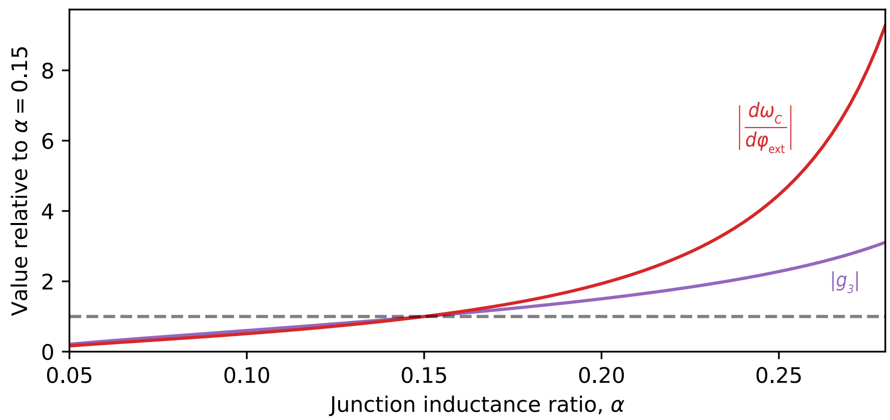

Figure 3.3: **Dependence of SNAIL properties on α.** Simulated variation of SNAIL frequency sensitivity to external magnetic flux $|d\omega_C/d\varphi_{\text{ext}}|$ and third order nonlinearity $|g_3|$, as a function of α at the Kerr-free point $K_c^{(3)} = 0$. Circuit parameters $N = 1$, $M = 3$, $E_{\mathrm{C}}/h = 0.177$ MHz and $E_L/h = 64$ GHz are chosen to match device in Chapter 4, with exception of E_{J}, which is varied to ensure $\omega_C = 2\pi \times 4.5$ GHz at the Kerr-free point. Both values are plotted relative to their value at $\alpha = 0.15$ and highlight a worsening trade-off between nonlinearity and flux noise as $\alpha \to 1/M$.

by the cavity modes and start to compromise their intrinsic noise bias.

To see how these two effects trade off against each other, we can compare $|g_3|$ and $|d\omega_C/d\varphi_{\text{ext}}|$ at the Kerr-free point ($K_c^{(3)} = 0$), as a function of α for the circuit parameters used in the device described in Chapter 4 ($E_C/h = 0.177$ MHz, $E_L/h = 64$ GHz, $N = 1$ and $M = 3$) . The value of E_{J} is adjusted to ensure that $\omega_C/2\pi = 4.5$ GHz for each choice of α. This choice also ensures φ_{ZPF} is fixed for all α. The values in Fig. 3.3 show a roughly $1:1$ trade-off until $\alpha = 0.15$, after which the flux noise increases substantially more rapidly than $|g_3|$.

Based on this, the optimal choice of α will depend on the duty cycle. If the duty cycle is low, and we care a lot about maximizing the cavity T_ϕ while idling then minimizing the sensitivity to flux noise encourages a low $\alpha \approx 0.05$. If instead, we care more about the fidelity while the beamsplitter is activated, then the roughly even tradeoff between $|g_3| \propto 1/t_{\text{bs}}$ and $|d\omega_C/d\varphi_{\text{ext}}| \propto 1/T_\phi$ up to $\alpha = 0.15$ should result in similar infidelities due to flux noise dephasing during the beamsplitter in this range. Choosing $\alpha = 0.15$, as we do, ensures the fastest possible g_{bs} in this range.

Optimizing p

As in amplifiers, maximizing p so that the majority of the inductance is Josephson inductance is ideal. One design challenge (as we shall see in Chapter 4) is that we would like a strong capacitive coupling between the coupler and each of the cavity modes but the cavities themselves are physically separated in the package. One way of coupling to two physically separated modes is to use long leads between the SNAIL's junctions and its capacitor pads. However, these can introduce a large amount of geometric inductance $L_{\mathsf{geom}} \sim \mu_0 l$, where l is the length of the leads. It is therefore paramount to keep the two cavities located as close together as possible. A double-post cavity (Gertler *et al.*, 2023; Koottandavida *et al.*, 2024), where the fields of the two dressed modes overlap is one way of satisfying this. However, for the same reason, coupling an ancilla transmon to these modes without inducing a significant χ_{ab} is difficult.

Optimizing p_a and p_b

Increasing the participation of the oscillator modes in the coupler allows for higher beamsplitter rates for the same coupler nonlinearity and drive amplitude. However, if the participation is too great, it can also limit the coherence of the oscillator modes. When discussing the optimal value of dispersive shift χ for an ancilla transmon (Ch. 2.2.4), we found that the linear oscillator mode inherits decay from the coupled mode as $\kappa_i^{\mathsf{inherited}} \approx p_i^2 \kappa_c$, where $i = a$ or b. We also found that for the case of white noise, the inherited dephasing was proportional to p_i^4. However, since the coupler is a flux-tunable device, its dephasing noise is likely dominated by $1/f$ flux noise, for which $\gamma_\phi^c \propto |d\omega_c/d\varphi_{\mathsf{ext}}|$ (Krantz *et al.*, 2019). We can relate the dephasing experienced by the coupler to the inherited dephasing of the oscillators by comparing the relative sensitivities of the mode frequencies to fluctuations in the external flux,

$$\gamma_\phi^{i,\ \mathsf{inherited}} = \left|\frac{d\omega_i}{d\omega_c}\right| \gamma_\phi^c \approx p_i^2 \gamma_\phi^c, \tag{3.52}$$

where Eq. 3.8 and Eq. 3.9 were used to relate the frequency fluctuations. In order to preserve the low-dephasing noise bias of the oscillator modes, we require

$$\gamma_\phi^{i,\text{ inherited}} \ll \kappa_i^{\text{intrinsic}} + \kappa_i^{\text{inherited}} \tag{3.53}$$

$$\rightarrow \gamma_\phi^{i,\text{ inherited}} \ll \kappa_i^{\text{intrinsic}} \tag{3.54}$$

$$\rightarrow p_i^2 \ll \frac{\kappa_i^{\text{intrinsic}}}{\gamma_\phi^c}, \tag{3.55}$$

where in going from Eq. 3.53 to Eq. 3.54, we have used $T_{1,c} \gtrsim T_{\phi,c}$. For typical values of $T_{\phi,c} \approx 20$ μs and $T_{1,i} \approx 1$ ms, this encourages $p_i^2 \ll 0.02$ and so we opt for $p_i^2 \approx 0.0025$ (corresponding to $\frac{g_i}{\Delta_i} \approx \frac{1}{20}$).

Optimizing frequency stack

In a three-wave mixing beamsplitter, the choice of oscillator frequencies immediately determines the required pump frequency, $\omega_p = |\omega_b - \omega_a|$ (unlike in four-wave mixing). One benefit of choosing a small ω_p is that multiphoton processes involving the lossy modes above the waveguide cutoff frequencies ($\omega_{\text{wg}}/2\pi \approx 20$ GHz) of the tunnels in the 3D package require many pump photons. These processes rely on a high-order nonlinearity, whose strength is suppressed.

On the other hand, a challenge of a very small ω_p ($\ll 3$ GHz) is that the drive is strongly detuned from the SNAIL mode resonance ω_c, whose frequency (as well as those of the cavities) should be kept sufficiently high to ensure that its residual Bose-Einstein thermal population $p_{\text{th}} \lesssim 1\%$. While we might expect the mode temperatures to be thermalized to the mixing chamber stage at 10 mK, in practice, state-of-the-art 3D cavity mode temperatures are around 35 mK (Chou *et al.*, 2024; Milul *et al.*, 2023), requiring mode frequencies that satisfy $\omega/2\pi \gtrsim 3$ GHz. The SNAIL mode provides a resonant embedding structure around the driven Josephson nonlinearity so that driving near its resonant frequency allows for a large normalized drive amplitude $|\xi|$ for the same applied field $|\epsilon|$. This can be seen from the approximate expression for the normalized drive amplitude, $\xi \approx \epsilon/(\omega_p - \omega_c)$. A strongly-detuned drive therefore makes the job of filtering the drive more challenging, although recent work driving a SNAIL through a low-pass filter has shown this is indeed possible (Zhou *et al.*, 2023).

In order to satisfy these requirements, we choose cavity frequencies $\omega_a/2\pi \approx 3$ GHz and $\omega_b/2\pi \approx 7$ GHz, to give a drive frequency $\omega_p/2\pi \approx 4$ GHz, near which we can place the SNAIL frequency. This ensures that at least 5 drive photons are required to drive a multiphoton process involving the waveguide modes. Both of the ancilla frequencies $\omega_{\text{anc}}/2\pi \approx 5.4$ GHz are then chosen to avoid intermodulation products of the oscillator frequencies.

3.6 Comparison to flux-pumped approaches

Before we conclude this chapter, it is important to note that a flux-pumped approach to generating a clean beamsplitter interaction, inspired by the flux-pumped JPA, *is* possible.

While we heard earlier that the JPA suffers from 4$^{\text{th}}$-order nonlinear terms, many can be removed by enforcing a symmetry in the drive field such that the SQUID is perfectly differentially driven (i.e. $\theta_1 = -\theta_2$, where the θ_i refer to the phase across each junction in the SQUID due to the drive) (Lu *et al.*, 2023). When two such purely-differential drives ($\phi_1(t)$ and $\phi_2(t)$) are applied, only terms proportional to $(\phi_1(t) + \phi_2(t))^n \left(p_a(\hat{a} + \hat{a}^\dagger) + p_b(\hat{b} + \hat{b}^\dagger) + \hat{c} + \hat{c}^\dagger\right)^m$, where n and m are both even, are permitted, suppressing up to half of the possible multiphoton resonances compared to a charge-driven transmon. Nonetheless, the permitted terms still include a Stark shift, as well as inherited oscillator self- and cross-Kerr interactions. Despite this, Lu *et al.* (2023) showed that with fewer multiphoton resonances, one can tolerate a significant Stark shift while achieving $g_{\text{bs}}/2\pi \approx 4.32$ MHz with very high-fidelity operations ($\mathcal{F} = 99.96\%$) on single-photon states[6]. However, the inherited nonlinearities of the cavity modes makes this approach suitable only for codes with low photon numbers, such as the dual-rail encoding.

Similarly, bisecting the SQUID loop with a linear inductance and threading one half of the loop with a DC external flux of $\varphi_{\text{ext}} = 0$ and the other with $\varphi_{\text{ext}} = \pi$ produces the asymmetrically threaded SQUID (ATS). When driven with a differential AC flux $\phi_1(t)$, this permits $\phi_1^n(t)\left(p_a(\hat{a} + \hat{a}^\dagger) + \hat{c} + \hat{c}^\dagger\right)^m$, where n and m are both *odd* (Lescanne *et al.*, 2020). While not

[6]The maximum beamsplitter rate achieved in this experiment was as high as $g_{\text{bs}}/2\pi \approx 11$ MHz, although the operating fidelity decreased beyond 4.32 MHz.

useful for beamsplitters, this has been extensively used in the context of dissipative cat qubits to generate a clean two-photon dissipation term $\hat{a}^2\hat{c}^\dagger$, which transfers photons two-at-a-time from the pristine $\hat{a}$ mode into the lossy $\hat{c}$ coupler mode, without introducing static Kerr nonlinearities (provided the two SQUID junctions have the same E_J). It has since been proposed to adapt this circuit to operate with a symmetrically-threaded flux (such that the DC flux through both halves of the loop is $\varphi_{\text{ext}} = \pi$), in order to generate $\phi_1^n(t)\left(p_a(\hat{a}+\hat{a}^\dagger)+p_b(\hat{b}+\hat{b}^\dagger)+\hat{c}+\hat{c}^\dagger\right)^m$, where n is odd and m is even (Maiti *et al.*, 2024). This symmetrically-threaded SQUID (STS), or linear inductive coupler (LINC), should permit a three-wave-mixing beamsplitter without Stark shifts or other Kerr terms to leading order. Nonetheless, these flux-pumped approaches require the delivery of highly-differential AC flux into a 3D superconducting package, where charge drives can be more straightforwardly engineered.

In conclusion, while charge-driven SNAILs do not present the only way of generating a clean beamsplitter interaction, they can do so while suppressing Stark shifts and Kerr terms without the need for highly-controlled AC flux delivery.

Chapter 4

Demonstrating a high on-off ratio microwave beamsplitter

The promise of a SNAIL-based coupler can be summarized as offering a high-fidelity, high on-off ratio beamsplitter, satisfying

$$\chi_{\mathsf{ab}} \ll \tau_{\mathsf{bs}}^{-1} \ll g_{\mathsf{bs}}, \tag{4.1}$$

where τ_{bs} captures the decoherence timescale of the oscillators when the beamsplitter is applied. The equivalent representative timescale for operations is the time to perform a 50/50 beamsplitter, $t_{\mathsf{bs}} = \pi/2g_{\mathsf{bs}}$. The inequality above points to two figures of merit for beamsplitter performance: the number of operations that can be performed in a coherence time, $\tau_{\mathsf{bs}}/t_{\mathsf{bs}}$, and the on-off ratio, $g_{\mathsf{bs}}/\chi_{ab}$, which compares the beamsplitter rate to the strongest always-on interaction between the oscillators. Having looked at the theory, it is now our job to make it work in practice.

We will begin (in Ch. 4.1) by discussing the experimental hardware required to make this possible, with a focus on the two main technical challenges: delivering DC flux and a strong charge-drive to the SNAIL coupler. We will then discuss how we characterize the properties of the static system Hamiltonian (Ch. 4.2), before turning our attention to characterizing the beamsplitter operation itself (Ch. 4.3).

4.1 Engineering a high-Q package for a charge-driven SNAIL

The superconducting package designed and fabricated for the experiments in this thesis (shown in Fig. 4.1) integrates design ideas from various applications in cQED. The package, which implements the circuit schematic in Fig. 3.1, is a modification of the 'Y-mon' coupler (Gao, 2018) used to mediate beamsplitters with a transmon coupler. It is constructed from a block of 99.999% purity aluminum, out of which are machined two 3D $\lambda/4$ stub cavities (see Ch. 2.1.2), whose fundamental modes will play the roles of 'Alice' and 'Bob' . To provide dispersive control of each of these cavity modes, we machine an elliptical tunnel into the side of each cavity, in which we suspend a 'standard' sapphire cavity control chip (Axline, 2018; Reinhold, 2019), and for beamsplitter control, another elliptical tunnel[1] intersecting both cavities houses a chip supporting the coupling element. In the original 'Y-mon' this (transmon) coupling element and its capacitor pads form a 'Y' shape, hence its name.

Switching from a transmon to a SNAIL coupler introduces two engineering questions:

- How can we deliver a DC flux bias φ_{ext} in a superconducting enclosure?
- How can we deliver a strong off-resonant charge-drive ξ?

Furthermore, can these questions be answered without compromising the pristine cavity modes? To do so we integrate two new structures: a DC flux transformer (Mundhada, 2019), and a buffer mode, respectively. While adding these structures, we also remove the readout resonator previously used to probe the coupler state, since the SNAIL's vanishing dispersive shift near the Kerr-free point disallows standard dispersive readout. However as we shall see in Ch. 4.2 this does not prevent us from characterizing the coupler properties.

Before discussing these new structures however, I will briefly summarize the elements that mediate the dispersive control, highlighting areas that might differ from usual implementations.

[1] An ellipse whose major axis lies along with the width of the chip allows for a smaller cross-sectional tunnel area for a given chip width. This results in a smaller 'perforation' of the cavity walls and less distortion of the cavity fields. This is especially true for the wide central tunnel housing the coupler chip. It should be noted that an elliptical tunnel is, however, more challenging to machine that a circular one.

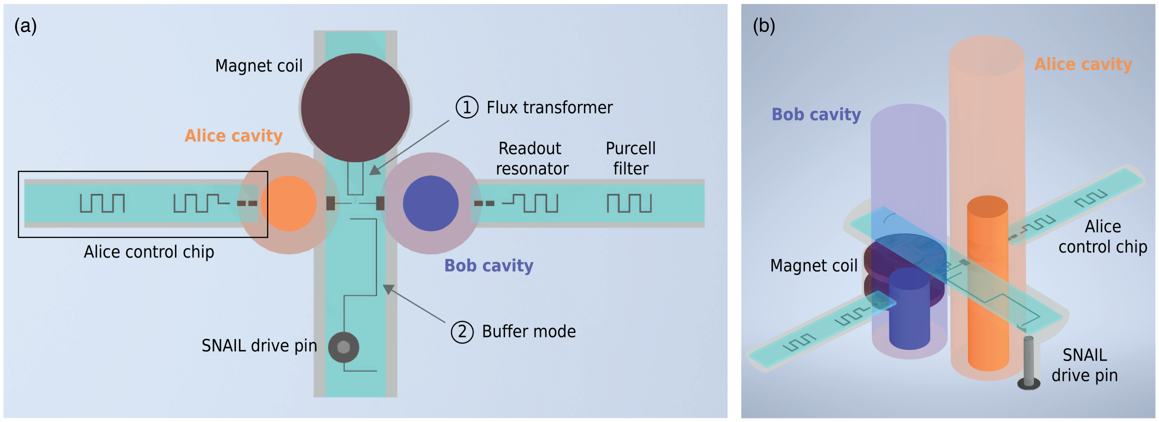

Figure 4.1: **Schematic of experimental package with outer walls removed.** (a) Top-down and (b) isometric view of the experimental package design. The stubs of Alice and Bob post cavities are shown in solid orange and blue, respectively, surrounded by vacuum (grey). Inserted into each cavity is a sapphire dispersive control chip (teal) onto which Al transmon, readout resonator and Purcell filter are patterned. The central chip hosts the SNAIL at its center, a flux transformer and magnet coil above, and a buffer mode and drive pin below. The coupling pins on the readout resonators and Purcell filters are omitted for clarity. In the isometric view, the vacuum inside the cavities is shaded orange or blue to match the stub. The capacitors pads of the SNAIL and of the ancilla transmon are inserted further into the cavity on Alice's side than on Bob's, since for a lower cavity frequency, a larger capacitance is required to obtain the same linear coupling strength, g.

4.1.1 Dispersive control hardware

Each dispersive control chip supports, from nearest the cavity to furthest: a 3D Al/AlOx/Al transmon (Paik *et al.*, 2011), a meandered Al stripline readout resonator, and a similarly meandered Al stripline Purcell filter. The Purcell filter sits beneath a terminated $50\ \Omega$ coaxial transmission line, whose outer conductor is formed by the walls of a cylindrical tunnel machined into the aluminum package and whose inner conductor is a non-magnetic cylindrical coupling pin. The length of this pin can be varied to control the loss rate κ_{Purcell} to the (overcoupled) transmission line, down which the readout signals are sent and received.

The Purcell filter (Reed *et al.* (2010); see Axline (2018) for a detailed description of this implementation) has a frequency ω_{Purcell} close to that of the readout resonator, $\omega_{r,a}$ (or $\omega_{r,b}$), but far from those of the ancilla transmon, $\omega_{t,a}$ ($\omega_{t,b}$), or the cavity, ω_a (ω_b). As a result, it acts as a single-pole filter (Pozar, 2012), passing readout signals at $\omega_{r,a}$ ($\omega_{r,b}$)), while suppressing loss from the high-Q modes down the transmission lines due to the inverse Purcell effect (Houck *et al.*, 2008). The inclusion of such a Purcell filter allows us, in principle, to achieve $\chi_{tr} \approx \kappa_r \approx 2\pi \times 1$ MHz, where χ_{tr} is the dispersive shift between the ancilla transmon and its readout resonator, and κ_r is the loss rate of the readout resonator, without limiting the lifetimes of the high-Q modes. Achieving $\chi_{tr} \approx \kappa_r$ with both values as large as possible maximizes the rate at which the transmon state can be read out (Gambetta *et al.*, 2006), an important parameter to optimize for single- and multi-mode measurements, as in Ch. 7.

Separate transmission lines are used to provide the Rabi drive on the transmon and the displacement drive on the cavity, and are designed to be substantially undercoupled ($Q_{\text{int}} \ll Q_{\text{coupling}}$) so as not to significantly suppress the lifetimes of these modes below their intrinsic values. Meanwhile, the cavity coupling pin is inserted into the side of the post cavity near the top of the stub, where the electric field of the oscillator mode is strongest. The transmon coupling pin is located above the readout resonator, at an electric field antinode of the transmon mode. This allows us to drive the transmon while keeping the pin as far as possible from the cavity mode to avoid limiting its relaxation time.

HFSS finite-element simulations, in concert with pyEPR (Minev *et al.*, 2021), are used to validate these linear (the loss rates of the modes and their frequencies) and nonlinear (the dispersive shifts) circuit parameters, respectively. The dispersive shifts χ_{at} and χ_{bt} between the transmon and the cavity are designed to be $2\pi \times 1$ MHz, following the discussion in Ch. 2.2.4. Given a fixed transmon Kerr nonlinearity $K_t/2\pi \approx 180$ MHz (independently optimized to suppress charge noise while allowing fast gates) and the frequency stack chosen in Ch. 3.5, we optimize these dispersive shifts by varying the capacitive coupling between the two modes, thereby varying their linear coupling g_{bt} (or g_{at}). While this is straightforward for Bob at $\omega_b/2\pi \approx 7$ GHz, achieving a sufficiently large g_{at} to the Alice mode at $\omega_a \approx 3$ GHz is made challenging by its low frequency. For two oscillators with frequencies ω_a and ω_t and characteristic impedances Z_a and Z_t, weakly coupled via a capacitance C_{coupling}, the linear coupling can be approximated as

$$g_{at} \approx \frac{\sqrt{Z_a Z_t}}{2} \omega_a \omega_t C_{\text{coupling}}. \tag{4.2}$$

Therefore, assuming fixed impedances, achieving $g_{at} = g_{bt}$ requires a C_{coupling} which is $\omega_b/\omega_a \approx 7/3$ times larger for Alice than for Bob. To do so, we insert the Alice control chip much further into the cavity, increasing the overlap between the capacitor pad of the transmon mode and the electric field of the cavity mode. One implication of this is that there is a direct dispersive coupling between Alice's transmon and the SNAIL, which is significantly larger than the coupling between Bob's transmon and the SNAIL. This will play a role in the characterization of the SNAIL frequency in Ch. 4.2.1.

The control chips themselves are held on only one end (furthest from the cavity) by a Beryllium-Copper (BeCu) leaf spring which is compressed between two halves of an aluminum clamp. Whereas the Al clamp is stiff and contracts as the temperature is decreased, the BeCu clip is flexible at room temperature while maintaining a high mechanical strength at 20mK. This prevents the chip from cracking due to the thermal-cycle-induced stresses.

4.1.2 Flux delivery

In planar superconducting circuits, such as those containing flux-tunable transmons, DC flux is typically applied via a coplanar waveguide (CPW) loop, whose ends are connected to a DC current source, and which is inductively coupled on-chip to the loop of the transmon or other circuit (DiCarlo *et al.*, 2009). A similar approach in 3D cavities, where the ground plane is not located on-chip but rather is formed by the outer walls of the enclosure, is challenging as this requires making a reliable galvanic connection between the replaceable sapphire chip and the outer walls.

In situations where the modes are not very high-Q and only a global field is required, this can be remedied by moving the loop off chip. This can be done by wrapping a current-carrying wire around a spool attached to the outside of the package, above the location of the loop (Grimm *et al.*, 2020). This requires the package (or at least the part near the spool) to be made of copper (or another normal metal) since a superconducting package will expel magnetic fields. However, the use of normal metal incurs conductor losses which limit the lifetimes of the modes.

Applying a DC flux to a circuit inside a superconducting package without a galvanic connection to ground, and without limiting the coherences of the modes is therefore desirable. The solution used here is a DC flux transformer, proposed in Mundhada (2019), and reminiscent of earlier work by Zimmerman (1971). It consists of a closed on-chip superconducting loop which passes near the SNAIL at one end and at the other passes underneath an aperture which connects to a coil wound around a copper spool. The current in this coil produces a DC magnetic flux through the transformer loop, inducing a supercurrent, which in turn induces a DC magnetic flux through the inductively coupled SNAIL loop. This circuit acts like a regular AC transformer, but by using a superconductor, can do so at DC. The design still requires puncturing the superconducting enclosure but can do so while keeping the aperture (containing the lossy magnet coil and its bobbin) far from the SNAIL and cavity modes. The high-Q modes also maintain a low participation in the normal metal spool around which the wire is wound, limiting the contribution from conductor loss.

Guided by the recommendations in Mundhada (2019), we can optimize the various design parameters. A primary consideration is to ensure sufficient flux through the SNAIL loop $\Phi_{\mathsf{SNAIL}} > \Phi_0$ for a coil current I_{coil} that is sufficiently low to avoid heating the mixing chamber stage of the dilution refrigerator. With care taken to ensure good thermalization of the joint between the superconducting wire used below 4K and the normal wire used above 4K, one can achieve $I_{\mathsf{coil}} > 200$ mA without any change in fridge temperature (see Appendix A of Chapman *et al.* (2023) for details). Maximizing the delivered flux can then be broken down into two components: maximizing the amount of flux threaded through the transformer loop for a given coil current and, given a certain flux in the transformer loop, maximizing the amount of flux threaded through the SNAIL loop.

The first component can be increased by maximizing the area of the loop beneath the coil and the number of turns of the wire around the coil. We wind ~ 500 turns, although a higher number could be used provided a sufficiently large spool (and dedication of time). The second component is set by the relative ratio of the transformer-SNAIL mutual inductance M_{ts} and the

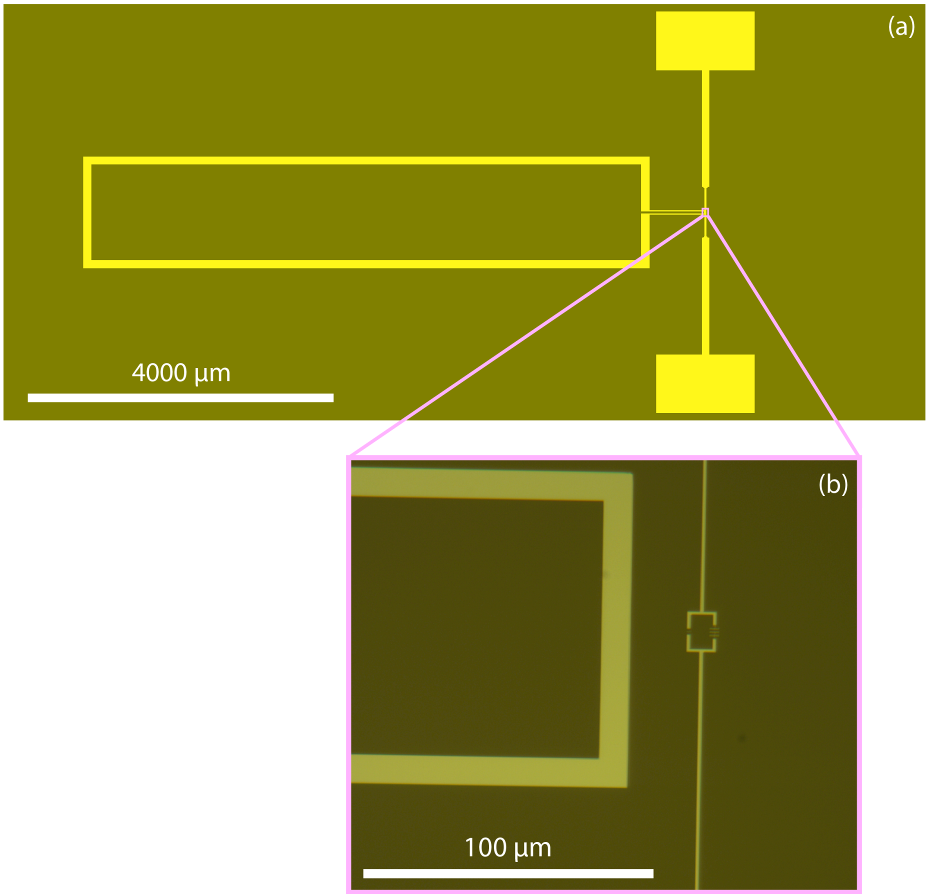

Figure 4.2: **SNAIL with flux transformer loop.** (a) *Schematic* showing the SNAIL and its flux transformer, with the yellow showing areas with aluminum and the olive indicating the sapphire substrate on which it sits. (b) *Optical image* of a real device with the same layout as the one used in Chapman et al. (2023), fabricated and imaged by Benjamin Chapman, in the vicinity of the SNAIL junctions. The structures on the right are the SNAIL, with its single small junction and three large shunting junctions. Thin leads extend vertically before widening out and eventually connecting to large capacitor pads. The structure on the left is the flux transformer, with a trace width of $10\ \mu$m in the narrow portion near the SNAIL and $100\ \mu$m for the rest of the loop.

transformer self-inductance L_t, since

$$\Phi_{\text{SNAIL}} = M_{ts} I_{\text{transformer}} = M_{ts} \frac{\Phi_{\text{transformer}}}{L_t} \rightarrow \frac{\Phi_{\text{SNAIL}}}{\Phi_{\text{transformer}}} = \frac{M_{ts}}{L_t}. \tag{4.3}$$

The ideal strategy from this perspective is to have a wide trace width far from the SNAIL loop (in order to suppress the overall loop inductance) and a narrow trace width near the SNAIL loop (in order to boost the mutual inductance). The lower limit on the SNAIL-end trace width is set by requiring that the critical current through it is not exceeded, prompting a choice of $10\ \mu$m. Meanwhile, the coil-end trace width (chosen to be $100\ \mu$m) should not be so large that the capacitive coupling between the arms of the transformer reduces the frequency of resonant modes hosted by the transformer structure into the range of the high-Q cavity modes, the highest-frequency of which is at 7 GHz in our system. We can further increase M_{ts} by reducing the distance between the transformer loop and the SNAIL loop. Balancing this effect with the proximity effect of a nearby transformer loop interfering with the fabrication of the SNAIL junctions, we opt for a gap of $20\ \mu$m. L_t can also be further reduced by minimizing the overall length of the loop, however this is a weak effect.

We can simulate $I_{\text{coil}}/\Phi_{\text{SNAIL}}$ using a magnetostatics solver (e.g. ANSYS Maxwell), breaking the problem up into its two components. Calculating $\Phi_{\text{transformer}}/I_{\text{coil}}$ can be done by treating the magnet coil windings as a superconducting annulus containing a current equal to the current per wire times the number of wire turns and then simulating the net magnetic flux through the transformer loop. Calculating $\Phi_{\text{SNAIL}}/\Phi_{\text{transformer}}$ (equivalently M_{ts}/L_t) can instead be done by fixing a current in the transformer loop and measuring the ratio of the magnetic flux through the SNAIL loop to the flux through the transformer loop. For the nominal design, with 500 wire turns on the coil, we find $\Phi_{\text{transformer}}/I_{\text{coil}} = 139\ \Phi_0/\text{mA}$ and $\Phi_{\text{SNAIL}}/\Phi_{\text{transformer}} = 3.3 \times 10^{-5}$, giving us $I_{\text{coil}}/\Phi_{\text{SNAIL}} = 218\ \text{mA}/\Phi_0$. This allows us to reach half-flux well within our allowable current range. However, simulations also show that $\Phi_{\text{transformer}}/I_{\text{coil}}$ is (in particular) sensitive to changes in the chip position relative to the bottom of the magnet coil. For this particular device, the measured $I_{\text{coil}}/\Phi_{\text{SNAIL}} = 39\text{mA}/\Phi_0$ is explainable by a shift in the chip position by 1 mm.

Another important concern when it comes to the length of the flux transformer is ensuring the common and differential electromagnetic modes of the transformer remain at a higher frequency than high-Q modes in the system, encouraging a short transformer length. This prevents these modes from hybridizing with the transformer (and inheriting its conductor loss), and helps minimize their participation in multiphoton processes. This should be done while

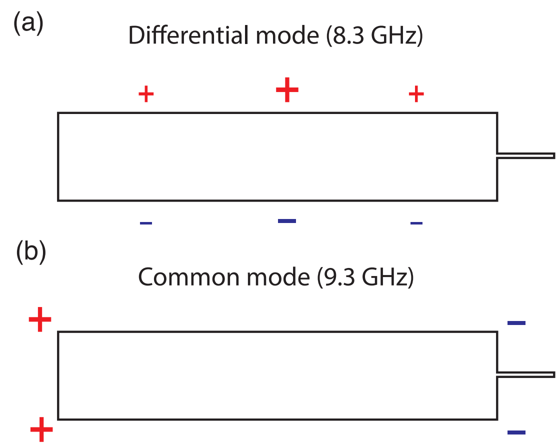

Figure 4.3: **Electric field distributions of resonant modes in Flux transformer.** Schematic showing out-of-plane electric field magnitude of the lowest-frequency resonant modes of the flux transformer, labeled 'differential' and 'common', respectively.

keeping the mode participations in the spool low. In practice, a lower limit on the length is set by the smallest distance the magnet coil can be placed from the SNAIL, which sits directly between the cavities, without intersecting the cavities. In our system, this is ~ 8 mm, which gives simulated resonant mode frequencies of 8.3 GHz (differential) and 9.3 GHz (common), with their electric field distributions sketched in Fig. 4.3. The differential mode couples more strongly to the SNAIL mode owing to the orientation of its electric field relative to the SNAIL dipole.

A final important consideration is that of light-tightness. The entry of infrared photons into the superconducting package contributes to both the generation of non-equilibrium quasiparticles and their photon-assisted tunneling across Josephson junctions (Diamond *et al.*, 2022), leading to increased energy relaxation in the transmon and SNAIL modes (Catelani *et al.*, 2011), and so must be mitigated. Allowing the superconducting leads of the magnetic coil to leave the package can introduce a small opening through which infrared radiation can enter. We manage this by applying a layer of double-sided copper tape underneath the copper spool, between it and the aperture into the coupler tunnel, which blocks light but admits magnetic fields. It is beneficial to have the barrier be as thin as possible, to keep the coil as close to the chip as possible without the barrier introducing conductor losses.

The final design of the flux transformer next to the SNAIL coupler (both a zoomed-out schematic, as well as a zoomed-in optical image) is shown in Fig. 4.2. This setup produces a single flux quantum, Φ_0, through the SNAIL loop for ~ 40 mA of current, comfortably within the limit for refrigerator heating, as verified by the periodicity of the mode frequencies in Ch. 4.2. Importantly, the calibration between applied current and φ_{ext} remains stable over cooldowns lasting several months, even when regularly adjusting the flux bias point (although we have been cautious to ramp the bias current at a low rate of $10^{-}5$ A/s to avoid trapping flux.) In these experiments, a very small bias current $I_0 < 1$ mA (the exact value varied per cooldown) was required to thread exactly no flux through the SNAIL loop, indicating the presence of a small residual magnetic field in the superconducting package when cooling down through T_c.

4.1.3 Pump delivery

To achieve as large as possible g_{bs}, it is necessary to obtain $|\xi|$ approaching $|\xi|_{\text{IP3}}$ (Eq. 3.49), but one wants to do so without Purcell limiting the high-Q modes (including the SNAIL and transmon modes) due to their capacitive coupling to the input drive pin. Reducing this capacitive coupling by retracting the coupling pin also reduces its coupling to the SNAIL mode at ω_p, thereby requiring more pump power to achieve the same $|\xi|$. A lower bound on this coupling strength is therefore set by the point at which the required pump power dissipates enough heat in the microwave attenuators at the base of the refrigerator to meaningfully heat the package modes.

Balancing the desire to couple strongly at one frequency, while reducing the coupling at the frequencies of one or many high-Q modes is also key to qubit readout, motivating the introduction of the Purcell filter discussed in Ch. 4.1.1. Inspired by this, we introduce a 'buffer mode' consisting of a $\lambda/2$ stripline resonator (seen in Fig. 4.1) with a resonance frequency near the desired drive frequency for beamsplitting, $\omega_{\text{buffer}} \approx \omega_b - \omega_a \approx 4$ GHz, located between the input port and the SNAIL. This structure acts as a bandpass filter at $\omega_b - \omega_a$, allowing the beamsplitter drive through while suppressing relaxation at the mode frequencies.

The design of the filter is meandered, with a hooked end so that the electric field excitation of the buffer mode aligns with the dipole moment of the SNAIL mode[2]. The meander is located further from Alice's cavity than Bob's, as the smaller frequency detuning between the buffer mode frequency and Alice's resonant frequency makes it more sensitive to Purcell decay. The SNAIL drive pin is not located at an antinode of the buffer mode (where it would achieve maximum coupling to the SNAIL) but nearer a node of the cavity fields, to minimize the coupling to these modes.

With these considerations accounted for, an HFSS eigenmode simulation of the system with the input port as the only source of dissipation predicts that the Purcell limit on the cavity lifetimes is in excess of 5.3 ms and 2.3 ms for Alice and Bob, respectively, above the expected intrinsic lifetimes. This validates that for this choice of coupling pin length, the beamsplitter drive port should not Purcell limit the cavity modes.

To predict what applied drive amplitude at room temperature is required to achieve a certain value of the Hamiltonian parameter $|\xi|$, we can perform an HFSS driven modal simulation and extract the impedance matrix $\mathbf{Z}$ of the system as a two-port network, with the port 1 being the input port and port 2 being defined across the SNAIL junctions. The entries of $\mathbf{Z}$ can then be related to the $|\xi|$ via:

$$|\xi| = \left|\frac{Z_{21}}{Z_{11} + R_S}\right| \frac{\sqrt{PR_S}\Phi_c^{\mathsf{ZPF}}}{\hbar\omega_p L_s(\omega_p - \omega_c)}, \tag{4.4}$$

where R_S is the source impedance of the generator, P is the RMS power delivered to the input port, L_s is the SNAIL inductance defined in the previous chapter, and Φ_c^{ZPF} is the zero point phase fluctuations of the coupler mode. This last parameter is not provided to us directly by HFSS, but can be obtained with further processing, for example with pyEPR (Minev *et al.*, 2021). A derivation of this formula is given in Appendix B.

Using a value $P = 100$ pW at the input port of the package (a value within the limits imposed by the ~ 25dB attenuation at base - see Fig. 4.5), we use Eq. 4.4 to predict $|\xi|$ as a function of ω_p when the SNAIL flux is at $\varphi_{\mathsf{ext}}/2\pi = 0.35$ (see Fig. 4.4). The feature at ~ 4

[2] The SNAIL itself is aligned perpendicular to the length of the chip to align with the cavity mode electric field. By aligning the SNAIL dipole with a slight angle (as in Grimm *et al.* (2020)) the buffer mode could be designed without a meander, but at the cost of requiring larger SNAIL pads or further insertion into the cavities to achieve the same g_a and g_b.

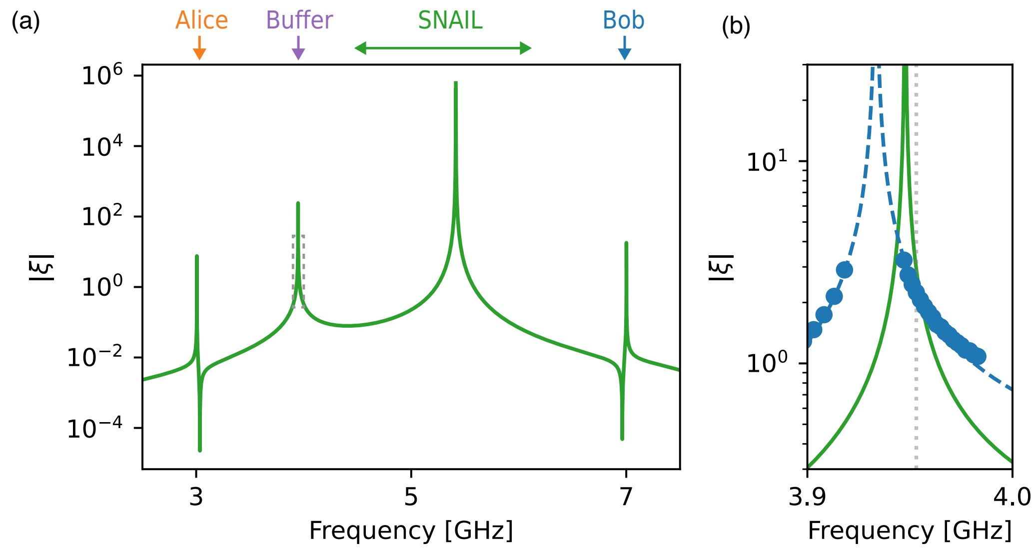

Figure 4.4: **Predicted and measured drive amplitude $|\xi|$.** (a) Predicted $|\xi|$ from Eq. 4.4 (green) as a function of pump frquency, with resonances corresponding to system modes labeled. The SNAIL frequency tunes from 4.4 GHz to 6 GHz with external flux. The dashed grey box around the beamsplitter drive frequency is shown in (b). Blue dots show measured $|\xi|$ values extracted from Stark shift and estimation of cold attenuation in fridge. Data lines up with predicted values when shifted in 14.3 MHz in frequency and by a factor of 2.84 in power (blue dashed line). Vertical dotted gray line shows resonant beamsplitter frequency $(\omega_b^{(0)} - \omega_a^{(0)})/2\pi$. Figure modified with permission from Chapman *et al.* (2023).

GHz is the buffer mode resonance, which permits $|\xi| > 1$ over a window of 40 MHz near the desired pump frequency. In Ch. 4.3, we will discuss a method to extract $|\xi|$ experimentally from the Stark shift of the coupler frequency. Using this technique, we can obtain measured values (blue dots) which, up to a frequency shift of 14.3 MHz and a factor of 2.84 in power, match the predicted ones. These shifts are likely to due to our uncertainty in the attenuation between our room-temperature setup and the package when cold, and the precise location of the coupler chip (relative to the pin / cavities) after mounting.

The curve also highlights another advantage over relying on driving close to the SNAIL resonance frequency, which is that the frequency of the buffer mode resonance is independent of flux, allowing us to drive the SNAIL over the entire range of external flux values.

4.1.4 Wiring and mounting

The wiring diagram for the system is shown in Fig. 4.5. The Alice and Bob output lines are equipped with quantum-limited amplifiers – a JAMPA (Sivak *et al.*, 2020) and SPA (Frattini *et al.*, 2018), respectively – enabling single-shot readout of the transmon state. Importantly, the local oscillator (LO) for the beamsplitter drive is synthesized by mixing together the LO for the Alice and Bob cavity drives using an image-reject mixer. Furthermore, the frequency at which the beamsplitter drive is single-sideband-modulated is set equal to the difference between those for Alice and Bob, $\omega_{\mathsf{SSB}}^{(\mathsf{bs})} = \omega_{\mathsf{SSB}}^{(\mathsf{b})} - \omega_{\mathsf{SSB}}^{(\mathsf{a})}$. These conditions ensure that states originating in one cavity acquire the correct phase when swapped into the other, even if the phase of one cavity's microwave generator drifts relative to the other. In order to nonetheless minimize this drift, we source the LOs from two channels on the same SignalCore SC5510A generator. Following the mixer, the beamsplitter line is heavily bandpass filtered around $\omega_b - \omega_a$ to remove intermodulation products which, at either the frequency of a system mode or an unwanted mixing process in the coupler, could degrade the beamsplitter performance.

The package is mounted inside a light-tight Cryoperm magnetic shield at the mixing chamber stage of a Bluefors XLD400sl dilution refrigerator, with the DC and microwave lines entering the shield via light-tight feedthroughs. The package itself is machined with large $\#10$ thru holes, allowing us to strongly clamp it to a gold-plated copper bracket, to maximize the thermal conductance at this interface and ensure the package is well-thermalized to the mixing chamber of the fridge (Richardson and Smith, 1988). Furthermore, to ensure that the input microwave lines do not contribute to the temperatures of the system modes, the attenuators and infrared filters at the mixing chamber stage (see Fig. 4.5) are clamped to purpose-built copper brackets (inspired by those in Krinner *et al.* (2019)). This keeps these dissipative elements as close as possible to the mixing-chamber temperature to minimize the Johnson-Nyquist noise they inject into the lines. It also helps (particularly for the strong parametric drive line) to transfer heat away from these elements when strong drives are applied. Finally, the package is wrapped in a layer each of Eccosorb HR and Eccosorb LS-30 foam, which absorbs stray infrared photons in the magnetic shield that may generate quasiparticles in the superconducting devices.

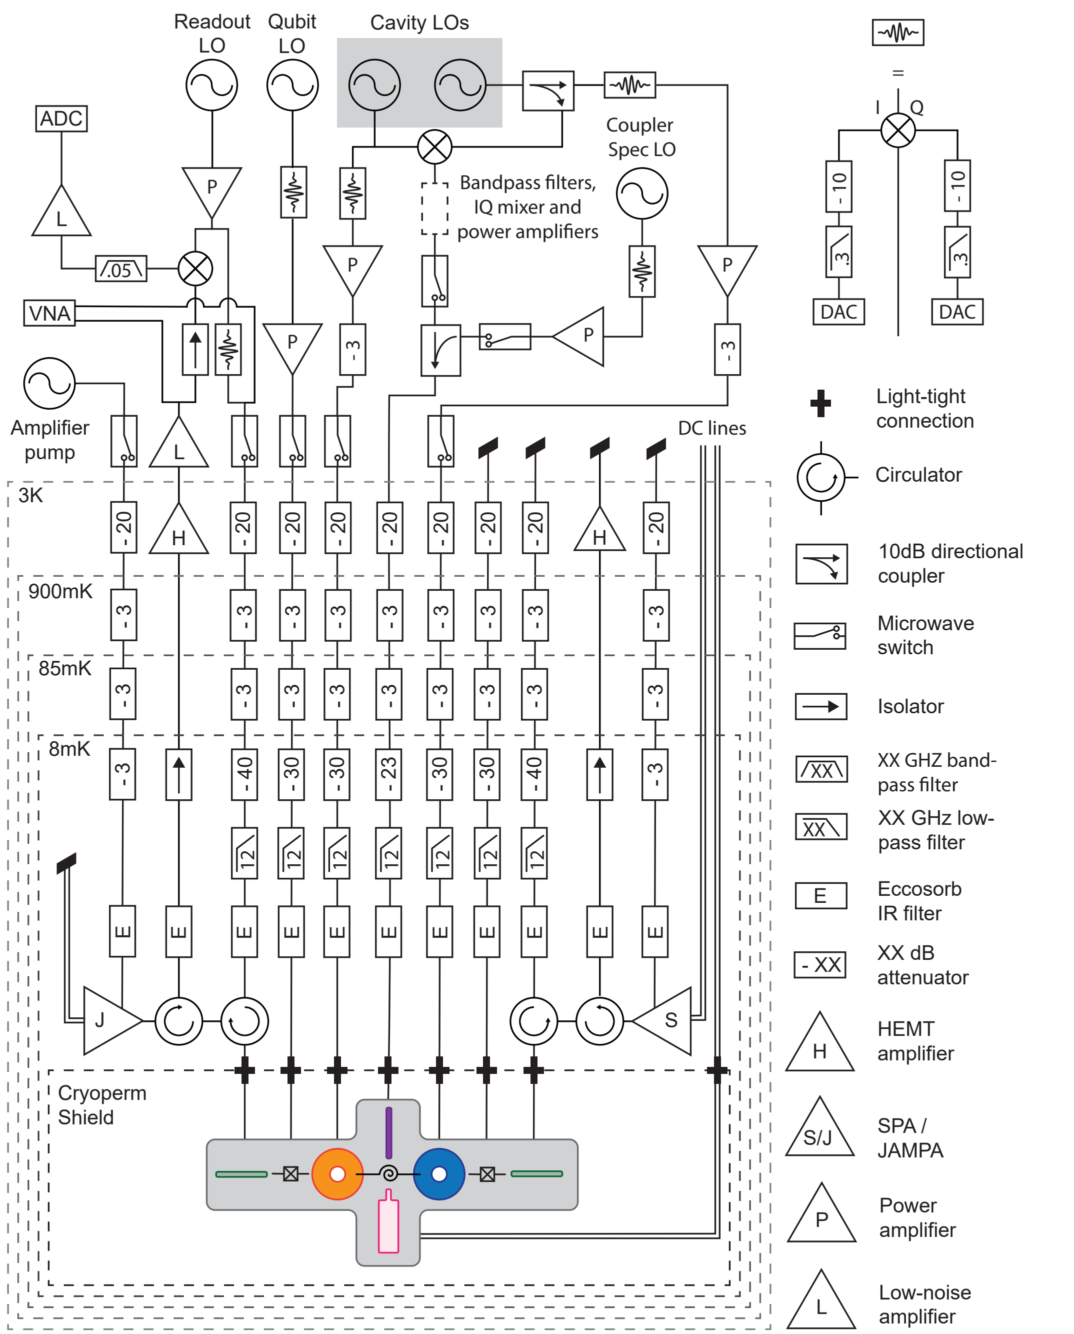

Figure 4.5: **Wiring diagram.** Experimental package is shown in gray, containing Alice (orange) and Bob (blue) cavities, transmon ancilla qubits, readout resonators (green), as well as a SNAIL coupler, a buffer mode (purple) and a flux transformer (pink). Purcell filters are omitted from the diagram for clarity. Figure reproduced with permission from Chapman *et al.* (2023).

4.1.5 System parameters

Integrating a flux transformer and buffer mode has allowed us to achieve stable DC flux bias with reasonable currents and a strong drive without heating the experimental setup. However, before looking at the SNAIL properties or confirming that we can generate the beamsplitter interaction, it is necessary to briefly discuss the coherence and thermal populations, obtained using standard methods (Geerlings *et al.*, 2013; Vlastakis, 2015). These values are shown in Table 4.1.

The oscillators demonstrate the necessary large error bias $T_\phi \gg T_1$, despite the introduction of the control hardware. However, while Alice's relaxation time is reasonable, Bob's is lower than typical for high-purity Al cavities. The bare lifetimes of the cavities, measured to be 2.3 ms and 460 μs, respectively, before the chips were inserted, do not explain the discrepancy. An obvious question to ask is whether the new features are responsible. However, HFSS simulations including Purcell loss from *all* the system ports, as well as conductor loss due to participation in the magnet coil, suggest a limit on T_1 no lower than 2.66 ms and 810 μs. Furthermore, as we shall see in Ch. 4.2.3, the SNAIL mode itself, which is significantly more strongly coupled to both the buffer mode and the flux transformer has coherences in line with typical tunable transmons (Acharya et al., 2024).

One clue to the lower-than-expected coherence is that large variations in cavity T_1 are observed when seemingly-unrelated changes are made to the experimental package from one cooldown to the next, such as changes to readout coupling pin lengths or the re-insertion of a chip. For example a modest increase in the length of Bob's readout pin length to obtain a larger κ_r and so a faster readout duration coincided with a decrease in Alice's T_1 to 347 ± 2 μs. Another cooldown in which all control chips were inserted saw Bob's cavity T_1 at 265 ± 27 μs. The difficulty of predicting this behavior points to unexpected package modes as a possible culprit.

While understanding the cause of this discrepancy is an important matter for future work, we will see that it does not prevent us from improving on previous transmon-based beamsplitter implementations by over an order of magnitude.

		Alice	Bob
Transmon $e \to g$ relaxation time	T_{1q}	$127.2 \pm 1.9\ \mu s$	$57.1 \pm 0.6\ \mu s$
Transmon T_2	T_{2q}	$114.4 \pm 2.9\ \mu s$	$56.8 \pm 1.5\ \mu s$
Transmon dephasing time	$T_{\phi q}$	$208 \pm 10\ \mu s$	$113 \pm 6\ \mu s$
Oscillator relaxation time	T_{1c}	$482 \pm 16\ \mu s$	$91 \pm 4\ \mu s$
Oscillator dephasing time	$T_{\phi c}$	$2010 \pm 220\ \mu s$	$840 \pm 200\ \mu s$
Transmon thermal $\|e\rangle$ population	P_e	0.70 ± 0.14 %	1.02 ± 0.20 %
Cavity thermal occupation	n_{thermal}	0.96 ± 0.19 %	0.11 ± 0.02 %

Table 4.1: **Measured system coherences and thermal populations.** Values as in Chapman *et al.* (2023). This table is reproduced in Appendix A, along with the system parameters for all other experiments in this thesis.

4.2 Characterizing the static Hamiltonian

Before turning on the beamsplitter, the dispersive cavity control and the flux tunability of the SNAIL allows us to probe the Hamiltonian's idling properties. In Chapter 3, we made predictions based on the SNAIL's nonlinear parameters g_n. Here, we provide a way to extract the g_n from just the (linear) frequency of the SNAIL mode, and compare the resulting estimates of the nonlinear properties to experiment.

4.2.1 Extracting g_n from the SNAIL frequency

Despite the lack of readout resonator coupled to the SNAIL, we are able to probe its frequency using the ancilla transmons. In principle, we can use a three-tone-spectroscopy sequence (shown in Fig. 4.6(b)) consisting of three consecutive microwave pulses: one on the SNAIL, one long (i.e. narrow-bandwidth) displacement pulse on Bob's cavity at ω_b, and finally a long π-pulse on Bob's ancilla at ω_{bt}, followed by a readout of the ancilla state. If the first spectroscopy tone is resonant with the SNAIL, it will be excited to $|e\rangle$. This shifts Bob's cavity frequency by χ_{bc} so that the second tone is off-resonant, leaving the cavity in $|0\rangle$. The final pulse on the ancilla is therefore resonant, exciting it to $|e\rangle$, before its state is then read out. Following the same procedure when the initial spectroscopy tone is off-resonant from the SNAIL frequency, instead leaves the ancilla in $|g\rangle$. Sweeping the frequency of the spectroscopy tone and monitoring the ancilla state therefore allows us to determine the SNAIL frequency.

This procedure suffers from a 'dead-spot' when $|\chi_{bc}| < 1/T_{2,c} \approx 2\pi \times 50$ kHz. However, for typical parameters, this excludes a small fraction of external flux values. Furthermore, if Alice's cavity is also equipped with an ancilla, it may also be used and since, generally, χ_{ac} and χ_{bc} vanish at different points, it will be typically be possible to cover the full range using both ancillas.

In our experiment, the small residual *direct* dispersive coupling $\chi_{t_a c}$ between the SNAIL and Alice's transmon (see Fig. 4.6(a)) is in fact sufficient to perform the two-tone spectroscopy sequence shown in Fig. 4.6(c), removing the need for a cavity pulse. Since $\chi_{t_a c}$ is predominantly derived from the nonlinearity of the ancilla transmon, there is no 'dead-spot' as we sweep φ_{ext}. While practically convenient in this case, one generally wants to avoid a large spurious coupling as it may compromise the photon-number selective cavity readout, mistaking a thermal photon in the SNAIL for one in the cavity.

Using the two-tone spectroscopy method, we can measure the SNAIL frequency from $\varphi_{\text{ext}} = 0$ to $\varphi_{\text{ext}} = \pi$, as shown in Fig. 4.6(d). The expression for $\omega_c(\varphi_{\text{ext}})$ in Eq. 3.46 can be used to fit this curve and extract the SNAIL circuit parameters. However, one constraint must be applied to find a unique solution. Given that N and M are known, the coupler frequency (barring the small Lamb shift just for the sake of this explanation) can be written as $\hbar\omega_c(\varphi_{\text{ext}}) = \sqrt{8E_{\text{J}}E_{\text{C}}\tilde{c}_2(p, \alpha; \varphi_{\text{ext}})}$, where p and α determine the shape of the curve as φ_{ext} is varied, and E_{J} and E_{C} are scale factors. Changing $E_{\text{J}} \to \beta E_{\text{J}}$, $E_L \to \beta E_L$, $E_{\text{C}} \to E_{\text{C}}/\beta$ for some β therefore gives exactly the same frequency dependence with flux. The necessary constraint is provided by using the measured room-temperature resistance R_N^T of the SNAIL array to fix its low-temperature kinetic inductance. This can be done using a combination of the Ambegaokar-Baratoff relation (Ambegaokar and Baratoff, 1963) between the low-temperature kinetic inductance, L_{K}, and the low-temperature normal state resistance, R_N^0:

$$L_{\text{K}} = \frac{\hbar}{\pi\Delta} R_N^0, \tag{4.5}$$

where Δ is the superconducting gap of thin-film aluminum, and the relation between the room-temperature and low-temperature resistance from Gloos *et al.* (2000). This gives a combined relation for our devices given by

$$L_{\mathsf{K}} = N \times \frac{L_{\mathsf{J}}}{\alpha + 1/M} \approx 1.27 \times 10^{-12} R_N^T. \tag{4.6}$$

Given $N = 1$, $M = 3$ and $R_N^0 = 2.98$ kΩ, the fit to Eq. 3.46 (shown as a solid line in Fig. 4.6(d)) yields $E_\mathrm{J}/h = 90.0 \pm 0.3$ GHz, $E_L/h = 64 \pm 2$ GHz, $E_\mathrm{C}/h = 177 \pm 2$ MHz, and $\alpha = 0.147 \pm 0.001$. From these circuit parameters, we may extract the predicted nonlinear Hamiltonian parameters, g_n, from Eq. 3.43, as shown in Fig. 4.6(e).

4.2.2 Measuring the coupler Kerr nonlinearity

We may test the validity of the fit by comparing the predicted expression for the Kerr nonlinearity of the coupler, K_c, from Eq. 3.29 to experiment. Two methods can be used to do so.

Once the coupler frequency ω_c and pulse amplitude required to perform a $|g\rangle \to |e\rangle$ π-pulse are known, we can modify the 2-tone (or 3-tone) spectroscopy sequence. We do so by inserting a SNAIL $|g\rangle \to |e\rangle$ π-pulse before and after the spectroscopy pulse. If the spectroscopy tone is resonant with the $|e\rangle \to |f\rangle$ transition frequency $\omega_c^{|e\rangle\to|f\rangle}$, the SNAIL is left in the $|f\rangle$ state, which we can map onto the transmon state. Meanwhile, if the spectroscopy tone is off-resonant, the second π-pulse returns the SNAIL to $|g\rangle$. We can then find the nonlinearity via $K_c = \omega_c^{|e\rangle\to|f\rangle} - \omega_c$. One consideration with this method is that the duration of the SNAIL π-pulses must remain much smaller than K_c in order to perform a π-pulse and not a displacement.

In the close vicinity of the Kerr-free point, another, more technically-straightforward approach is to modify the 2-tone (or 3-tone) spectroscopy sequence so that both the frequency and amplitude of the pulse are varied. When K_c is small, we can see not only the oscillations associated with the single-photon $|g\rangle \to |e\rangle$ transition, but also those associated with the two-photon $|g\rangle \to |f\rangle$ transition at $\omega_c^{|g\rangle\to|f\rangle}/2$ and more generally, the n-photon $|0\rangle \to |n\rangle$ transitions). For small n, the spacing between these transitions corresponds to $K/2$, as is nicely shown for an arrayed SNAIL (different to the one used in the following experiments) for which

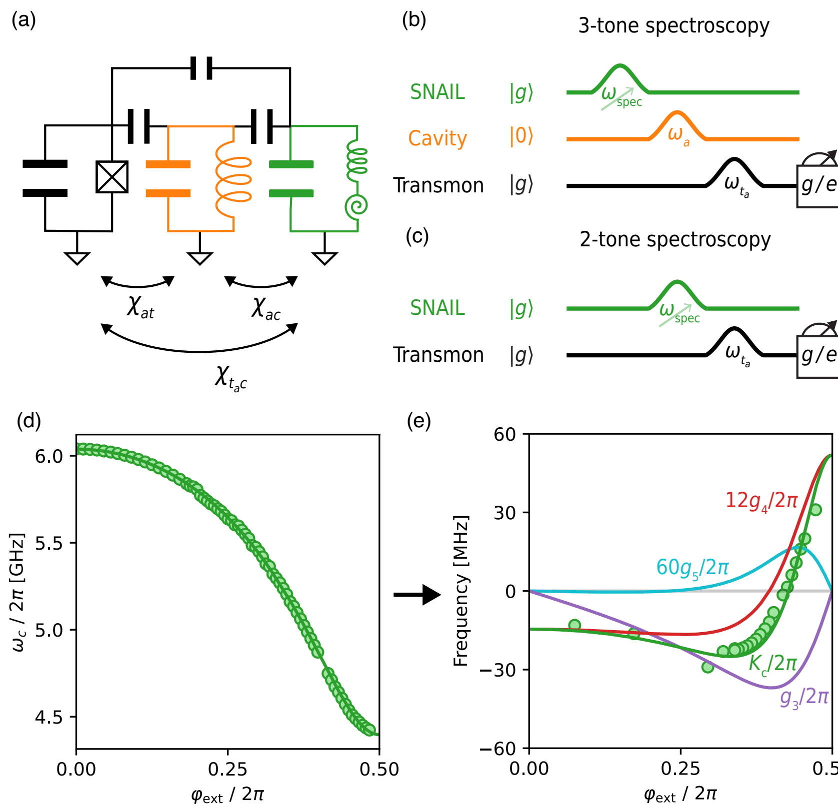

Figure 4.6: **Extracting g_n from SNAIL spectroscopy.** (a) Circuit diagram of an ancilla transmon (black), cavity (orange) and SNAIL (green), showing their mutual capacitances, including a weak direct coupling between the SNAIL and ancilla. These result in dispersive shifts χ between each pair of modes. (b) 3-tone spectroscopy sequence using frequency-selective microwave pulses on the SNAIL, cavity and ancilla, in turn, before an ancilla readout. (c) Alternative 2-tone spectroscopy sequence, which leverages the direct SNAIL-ancilla coupling $\chi_{t_a c}$, using frequency-selective pulses on the SNAIL and ancilla only. (d) Measured SNAIL frequency as a function of external flux φ_{ext} using 2-tone method (circles) and fit to Eq. 3.46 (solid line). The small region of missing data is where the SNAIL is resonant with Alice's ancilla and the two modes hybridize. (e) Extracted nonlinear parameters g_n from fit, as well as estimated Kerr nonlinearity (solid lines). Measured Kerr nonlinearity shown in circles. Subfigures (d) and (e) modified with permission from Chapman *et al.* (2023).

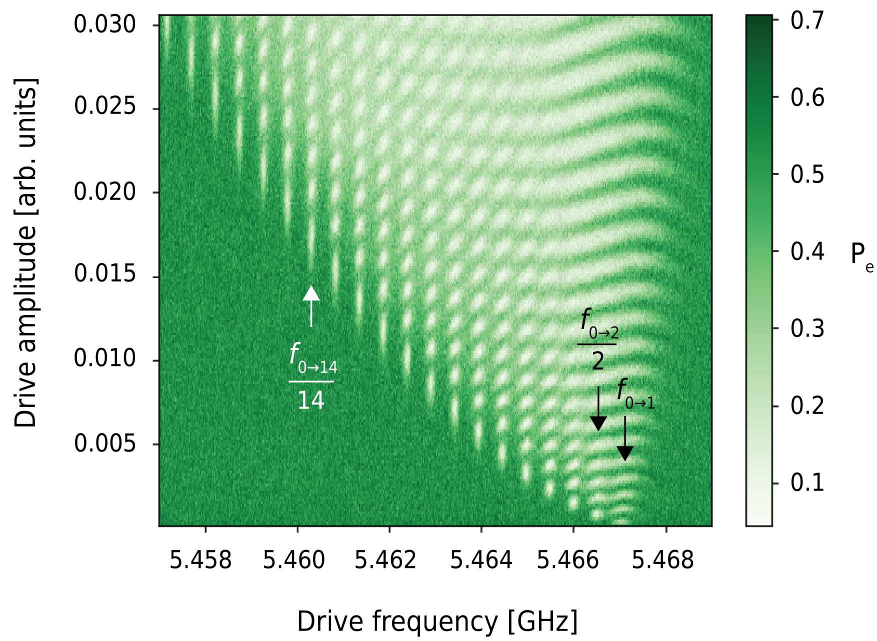

Figure 4.7: **Multiphoton transitions in low-Kerr SNAIL.** Spectroscopy of an arrayed SNAIL (different to the one used in the following experiments) while sweeping the drive amplitude. Besides the $|g\rangle \rightarrow |e\rangle$ ($f_{0\rightarrow 1}$) transition, one can see n-photon $|0\rangle \rightarrow |n\rangle$ transitions at $f_{0\rightarrow n}/n$, each separated by $K_c/2$, which for this device is $K_C/2\pi = -1.0$ MHz. The data shown were taken together with Benjamin Chapman.

many of these transitions can be observed at once in Fig. 4.7. This approach has previously been used to characterize the Kerr nonlinearity of cavity modes in the regime where the Kerr nonlinearity is too strong to use the method in Ch. 4.2.5 (Vlastakis, 2015).

The results are shown in Fig. 4.6(e), showing reasonable agreement with the estimated value, lending credibility to both the perturbative expression in Eq. 3.29 and the extracted g_n. Based on the extracted values, one can see a clear deviation from $12g_4$, the perturbative expression when considering only the 4$^{\text{th}}$-order nonlinearity.

4.2.3 Measuring SNAIL coherences

Armed with both a means of exciting and measuring the SNAIL, we can apply standard transmon coherence measurements to the coupler mode. One caveat is that as the anharmonicity approaches zero, the duration of the excitation pulse must increase to ensure that the SNAIL stays within its $\{|g\rangle, |e\rangle\}$ manifold. This allows us to extract the T_1, T_ϕ (Ramsey) and $T_{\phi E}$ (echo), shown in Fig. 4.8(a).

The measured coherences are state-of-the-art for SNAILs and are on par with typical coherences for tunable transmons in the field (Acharya et al., 2024), with a T_1 that is relatively steady across flux between 50 and 100 μs. The T_2 is around 3 μs, due to $1/f$ flux noise in the SNAIL loop. The dephasing rate from flux noise is expected to scale linearly with the gradient of ω_c with respect to external flux (shown in Fig. 4.8(b)) and indeed it appears to follow this trend, with an increase in T_2 at the half-flux point ($\varphi_{\text{ext}} = \pi$). Furthermore, the T_{2E} is substantially higher, as one expects for $1/f$ noise.

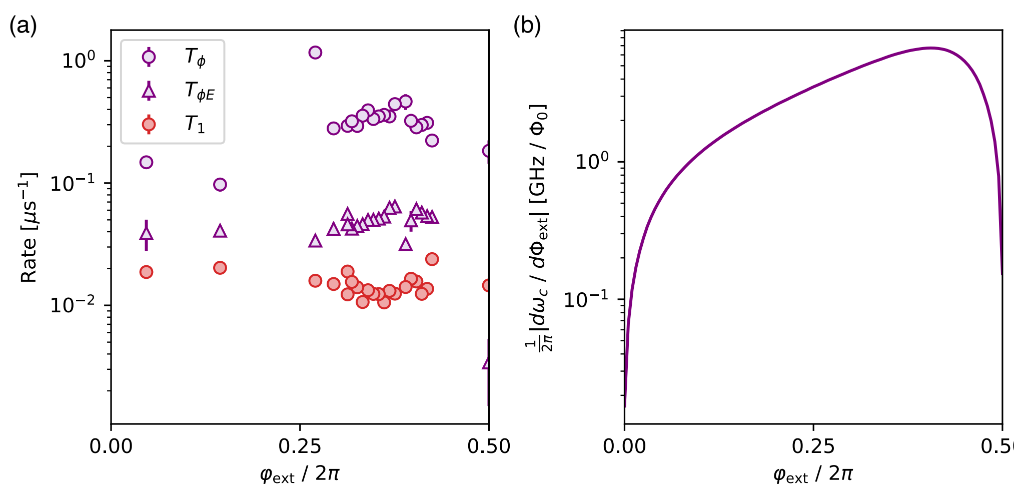

Figure 4.8: **SNAIL coupler coherences.** (a) SNAIL coupler T_1, T_ϕ and $T_{\phi E}$ times, showing little variation with φ_{ext} of T_1 and an increase in T_ϕ and $T_{\phi E}$ at the half-flux ($\varphi_{\text{ext}} = \pi$) sweet-spot. (b) Gradient of SNAIL frequency ω_c with external magnetic flux Φ_{ext}. One expects SNAIL dephasing due to flux noise to be proportional to this value. Figure modified with permission from Chapman *et al.* (2023).

4.2.4 Extracting the linear couplings g_a and g_b

So far we have just characterized the properties of the coupler itself but we would also like to characterize the self-Kerr and cross-Kerr terms, in particular χ_{ab}, which determines our 'off'-performance. These properties depend on the participations p_a and p_b, or equivalently given that the detunings $\Delta_a(\varphi_{\text{ext}})$ and $\Delta_b(\varphi_{\text{ext}})$ are known, the linear couplings g_a and g_b.

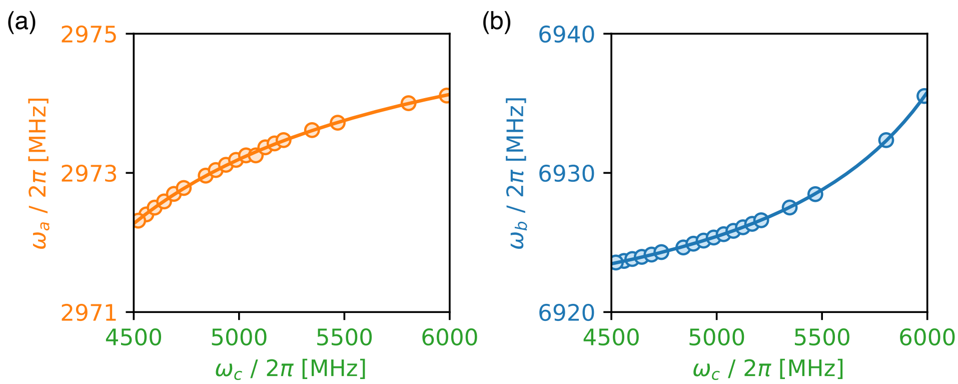

Figure 4.9: **Fitting linear coupling strengths.** Parametric plot of Alice (a) and Bob (b) cavity frequency as a function of the measured coupler frequency. Lines show fits to Eqs. 4.7 and 4.8 from which we extract linear coupling strengths, $g_a/2\pi = 75.6 \pm 0.2$ MHz and $g_b/2\pi = 134.9 \pm 0.1$ MHz, and bare mode frequencies, $\omega_A/2\pi = 2976.018(16)$ MHz and $\omega_B/2\pi = 6915.945(17)$ MHz. Figure modified with permission from Chapman *et al.* (2023).

These can be extracted by looking at the frequency shift of the oscillators as φ_{ext} is varied. We can approximate the inherited frequency shifts (assuming a fixed g_a and g_b, as well as the rotating-wave approximation) using Eqs. 3.8-3.10. Since the frequency shift due to the dressing is much smaller than the mode detuning, we can approximate the measured oscillator frequencies as a parametric function of the measured couple frequency as:

$$\omega_a(\omega_c) \approx \omega_A + \frac{g_a^2}{\omega_A - \omega_c}, \tag{4.7}$$

$$\omega_b(\omega_c) \approx \omega_B + \frac{g_b^2}{\omega_B - \omega_c}. \tag{4.8}$$

These expressions can be used to fit the measured frequencies while varying φ_{ext} in Fig. 4.9 to extract the fixed parameters – the bare mode frequencies, $\omega_A/2\pi = 2976.018(16)$ MHz and $\omega_B/2\pi = 6915.945(17)$ MHz, and the linear couplings, $g_a/2\pi = 75.6 \pm 0.2$ MHz and $g_b/2\pi = 134.9 \pm 0.1$ MHz. We find $g_a < g_b$ despite the SNAIL chip being further inserted into the Alice cavity, due to its lower frequency, highlighting the challenge of coupling strongly to a low frequency mode (Eq. 4.2).

4.2.5 Characterizing the self- and cross-Kerrs

The inherited self- and cross-Kerr terms in the oscillators tend to be weak and so to obtain a more sensitive measurement, we can populate the cavities with a large, calibrated number of photons to amplify their effect.

The self-Kerr of the oscillators (shown in Fig. 4.10(a)) are extracted using an interferometric method (Chou, 2018), applying a pair of detuned displacements (with detuning δ and amplitude α_{d}) separated by a variable delay t and monitoring the final vacuum state population $P_0(t)$. During the delay, the displaced state will rotate in phase space at a frequency that depends weakly on the mean photon number $|\alpha_d|^2$, acquiring a total phase space rotation $\phi_K = (\delta + K/2|\alpha_d|^2)t$. When $\phi_K = (2n+1)\pi$ where n is an integer, the second displacement (which has the same phase as the first) returns the center of the coherent state to the origin (vacuum state). These oscillations can be fit to

$$P_0(t) = A_0 \exp\left[-2|\alpha_d|^2 \left(1 + \cos\phi_K\right)\right] + C \tag{4.9}$$

where A_0 and C capture imperfections in the cavity population measurement, to extract K. This expression neglects the distortion caused by different photon number components rotating at different rates, however provided that the relative rotation accumulated across the spread of photon numbers is small ($|\alpha_{\mathsf{d}}|Kt \ll 1$), this effect can be ignored.

The cavity-coupler dispersive shifts χ_{ac} and χ_{bc} (shown in Fig. 4.10(b)) are instead obtained via a spectroscopic method – prepending the 2-tone SNAIL spectroscopy sequence with a calibrated displacement α_{d} on the oscillator. Dividing the shift in the fitted SNAIL frequency from spectroscopy by the oscillator photon number $|\alpha_{\mathsf{d}}|^2$ gives the desired dispersive shift.

The cross-Kerr χ_{ab} sets the on-off ratio and so we would like to measure it to as low as value as possible. Both spectroscopic and interferometric methods can be used, both of which involve populating one cavity with an increasingly large, calibrated number of photons $|\alpha_d|^2$ and measuring the frequency shift of the other. The interferometric method (Fig. 4.11(a)) determines this frequency using the same Ramsey sequence used to obtain the self-Kerr, but with a fixed

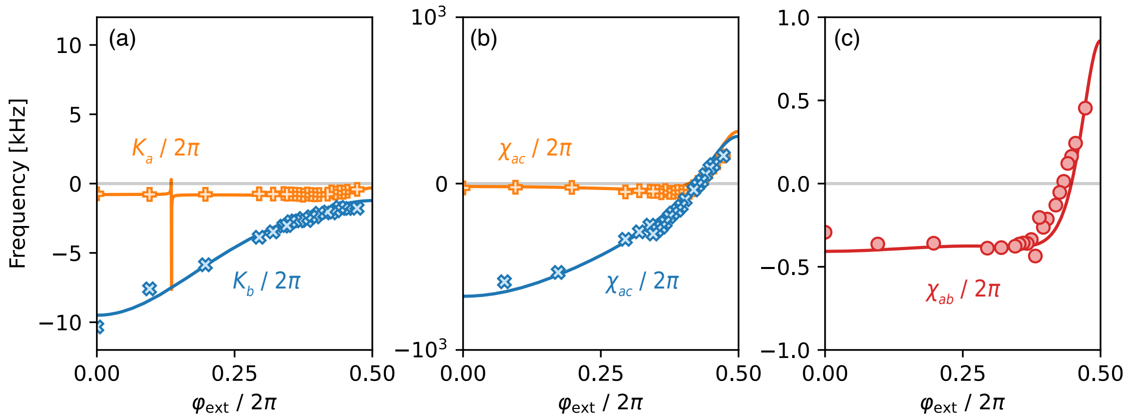

Figure 4.10: **Extracted nonlinear multi-mode Hamiltonian terms.** (a) Measured oscillator self-Kerr in Alice (orange) and Bob (blue) as external flux through the SNAIL loop is varied. Solid lines show prediction from Eq. 3.32. (b) Measured coupler-oscillator dispersive shift on Alice (orange) and Bob (blue). Solid lines show prediction from Eq. 3.35. (c) Measured cavity-cavity cross-Kerr, and prediction from Eq. 3.33 (solid line). Figure modified with permission from Chapman *et al.* (2023).

unit displacement, while the spectroscopic method (Fig. 4.11(b)) uses cavity spectroscopy. The interferometric approach can also be fit to Eq. 4.9 but with ϕ_K replaced by $\phi_{\chi_{ab}} = (\delta + \chi_{ab}|\alpha_d|^2)t$. The results of the spectroscopic and interferometric methods, measuring with each cavity, are shown to have good agreement with each other (Fig. 4.11(e)). The model in Eq. 4.9 ignores the effect of cavity decay, however the fact that the measurements taken with Alice and with Bob agree, despite the differences in their photon loss rates, suggests that this is an acceptable assumption to make. As the examples in Fig. 4.11(c-d) show, there exists a flux point where χ_{ab} is measured to vanish (up to the ~ 10 Hz resolution of our measurement). This confirms a key advantage of a SNAIL tunable coupler – that the ZZ interaction between the two oscillators can be made to turned off, even when the stored photon numbers are very large.

For all of these quantities, the perturbative approximations using the g_n from the fitted SNAIL spectrum provide a good prediction, capturing, for example, the fact that in this device, the inherited negative self-Kerr from the transmon ancillas overcomes the positive contribution from the SNAIL coupler at $\varphi_{\text{ext}} = \pi$, such that neither K_a or K_b pass through zero (Fig. 4.10(a)).

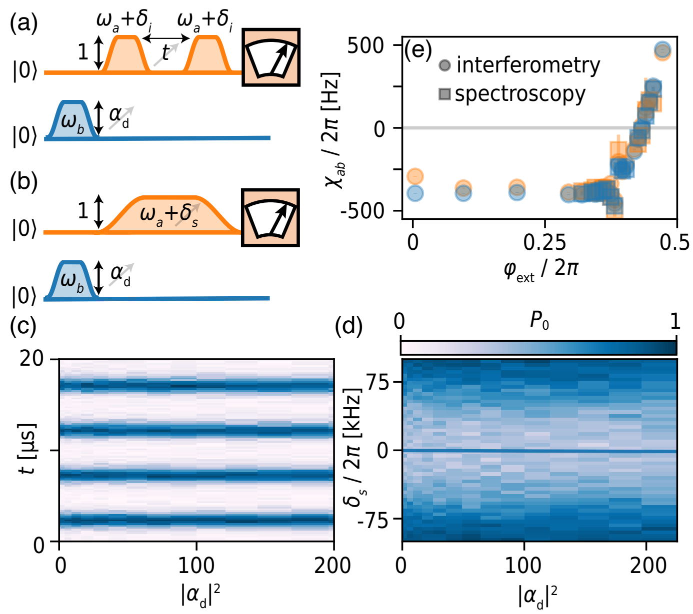

Figure 4.11: **Two methods for extracting χ_{ab}.** (a-b) Methods for determining χ_{ab} by populating one cavity with $|\alpha_{\mathrm{d}}|^2$ photons and measuring the frequency of the other. The frequency measurement can be performed with an interferometric Ramsey-like method (a), or by performing spectroscopy (b). The results of these measurements at $\chi_{ab} \approx 0$ for the two methods are shown in in (c) and (d). (e) The four different meaurement results using the interferometric (circles) and spectroscopic (squares) methods, while measuring Alice (orange) and Bob (blue) mode frequency. These methods agree well with one another, and capture a resonance near $\varphi_{\mathrm{ext}}/2\pi \approx 0.4$. Figure reproduced with permission from Chapman *et al.* (2023).

It should be emphasized that by slightly reducing the anharmonicity of the transmon ancilla, or its coupling to the cavity, both quantities could be made to vanish. We can also see the limitations of the perturbative approach by noting the resonance in χ_{ab} at $\varphi_{\mathrm{ext}}/2\pi \approx 0.40$. Capturing these beyond-perturbative effects is the subject of future work.

4.3 Characterizing the beamsplitter performance

4.3.1 Characterizing beamsplitter speed

Having established that we are able to turn off the static interaction between the cavities, we must now verify that the SNAIL coupler mediates a fast beamsplitter between them when driven. The performance in the single-photon manifold of the two cavities can be characterized by initializing a single photon in Alice, $|\psi\rangle_{AB} = |1\rangle_A |0\rangle_B$, applying a fixed-amplitude microwave drive to the SNAIL whose duration t and frequency ω is varied, and then simultaneously measuring whether each cavity contains 0 photons. Provided there is no heating, we can infer from a negative measurement outcome that $n = 1$. At flux points where $g_3 \neq 0$, plotting the probability of finding a photon in Bob, $P_b = P_{01}$ reveals a chevron pattern centered on $\omega_{\Delta=0} = \left(\omega_b^0 - \omega_a^0\right) + \left(\Delta_{\text{Stark}}^{(b)} - \Delta_{\text{Stark}}^{(a)}\right)$, where ω_a^0 and ω_b^0 denote the oscillator frequencies when the beamsplitter drive is off:

$$P_{01/10} = A\left[\cos^2\left(\frac{\Omega t}{2} + \phi\right) + \left(\frac{\omega - \omega_{\Delta=0}}{\Omega}\right)^2 \sin^2\left(\frac{\Omega t}{2} + \phi\right)\right] + c, \tag{4.10}$$

where the detuned oscillation rate is given by

$$\Omega = \sqrt{g_{\text{bs}}^2 + \left(\omega - \omega_{\Delta=0}\right)^2}, \tag{4.11}$$

ϕ represents the phase offset acquired while the beamsplitter drive is turned up to its maximum amplitude, and A and c capture imperfections in the photon-number selective measurement. Fig. 4.12(a) shows an example of such data overlaid with its fit to Eq. 4.10, from which we can extract both g_{bs} and $\omega_{\Delta=0}$[3].

To compare the resulting g_{bs} to the expression in Eq. 3.31, we must calibrate the normalized drive amplitude per unit applied power, $|\xi|/\sqrt{P_{\text{applied}}}$, which we obtain by measuring the Stark shift of the coupler frequency $\Delta_{\text{Stark}}^{(c)}$. At low drive powers, $\Delta_{\text{Stark}}^{(c)}$ scales linearly with P_{applied}, and so we can fit to find their ratio (Fig. 4.12(c) shows this calibration for $\varphi_{\text{ext}}/2\pi = 0.35$). The expressions for the Stark shift and the coupler's Kerr nonlinearity are closely related (see Eqs. 3.29 and 3.30), and so we can approximate $|\xi|$ by

$$\frac{|\xi|}{\sqrt{P_{\text{applied}}}} \approx \sqrt{\frac{\Delta_{\text{Stark}}^{(c)}/P_{\text{applied}}}{2K_c}}. \tag{4.12}$$

[3] While the difference in the oscillator Starks shifts is small, and can be made to vanish at certain values of φ_{ext}, it is nonetheless important to track if one wants to perform high-fidelity beamsplitter operations.

While this approximation is imperfect very close to the Kerr-free point, the deviation is less than 3% for $\varphi_{\text{ext}}/2\pi < 0.36$.

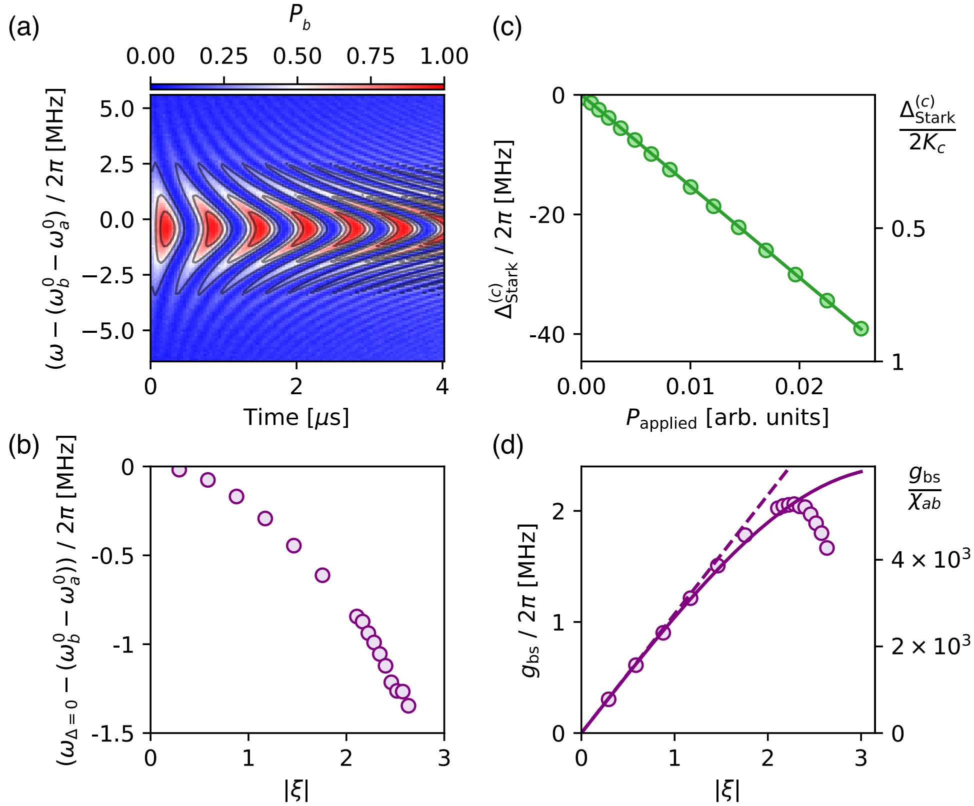

Figure 4.12: **Extracting g_{bs} at short times.** (a) Probability of measuring a photon in Bob, p_b, after initializing in $|\psi\rangle_{AB} = |0\rangle_A |1\rangle_B$ and applying a beamsplitter drive for a variable time and frequency ω. Overlaid contours show a fit to Eq. 4.10 to capture g_{bs} and $\omega_{\Delta=0}$. (b) Fitted chevron center frequency (relative to resonance condition at $|\xi| \to 0$) as $|\xi|$ is increased. (c) Calibration of normalized amplitude ξ as a function of applied power P_{applied} by measuring coupler Stark shift, Δ_{Stark}. Right axis shows Stark shift in dimensionless units. (d) Fitted beamsplitter rate g_{bs} versus normalized drive amplitude $|\xi|$. Right axis shows the on-off ratio, comparing this to the always-on χ_{ab} interaction between the oscillators. Dashed line shows lowest-order linear RWA approximation to g_{bs} from Eq. 3.31, while solid line includes higher-order RWA corrections due to g_5 and higher (Eq. 3.48). Subfigure (d) reproduced with permission from Chapman *et al.* (2023).

Fig. 4.12(d) shows the extracted g_{bs} as a function of $|\xi|$ at $\varphi_{\text{ext}}/2\pi = 0.35$, showing a linear increase at low $|\xi|$, in accordance with Eq. 3.31, and a deviation at higher amplitudes which will be discussed in more detail in Ch. 4.3.4. The maximum $g_{\text{bs}} > 2\pi \times 2$ MHz $\gg |\chi_{\text{ab}}|$, even away from the χ_{ab}-free point, strongly satisfying the requirement for a high-on-off-ratio coupler. This value is also well in excess of the value achievable with a transmon coupler, corresponding to a time per 50/50 beamsplitter operation $t_{\text{bs}} < 125$ ns. Most importantly, this puts us in a regime where $g_{\text{bs}} > |\chi|$, unlocking new multimode non-Gaussian operations, as we will explore in the remaining chapters of this thesis. However, in order for this increased speed to be valuable, we must also confirm that the coherence of the oscillators is not significantly degraded by the beamsplitter drive.

4.3.2 Characterizing driven decoherence

A qualitative inspection of the population oscillations at short times (see Fig. 4.13) indicates that some imperfections are present. As the normalized drive amplitude is increased, other features besides the expected 'chevrons' appear, especially once g_{bs} starts to decrease with $|\xi|$. This includes a high-population streak at a drive frequency of approximately $\omega = \omega_b^0 - \omega_a^0 + 2\pi \times 1$ MHz that becomes noticeable at increasingly short drive times, and a low-population region at drive frequencies below resonance that becomes problematic over an increasingly wide range of frequencies. (We should note that effects that populate or de-populate the measurement transmon or readout resonator may also cause distortions in the observed data.)

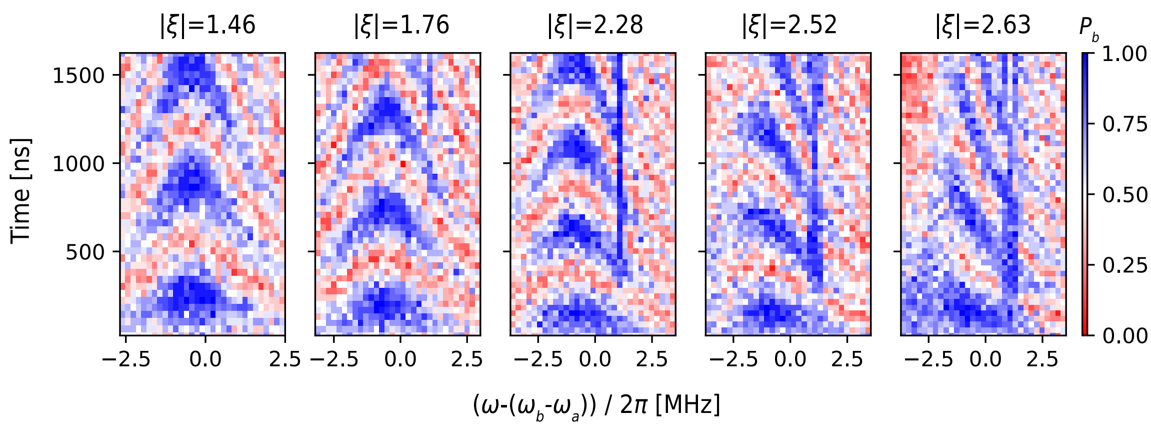

Figure 4.13: **Beamsplitter 'chevrons' at large drive amplitudes.** A selection of individual datasets used to extract g_{bs} in Fig. 4.12(d), showing oscillations of the population in the Bob oscillator when driving the coupler at a frequency ω for a variable duration, having initializing the oscillators in $|1\rangle_A |0\rangle_B$. The middle dataset at $|\xi| = 2.28$ is the one for which g_{bs} is maximized.

To more quantitatively characterize decoherence induced by the beamsplitter, we can look at resonant ($\omega_p = \omega_{\Delta=0}$) oscillations in the single-photon manifold over much longer timescales, where the expected behavior is

$$P_{01/10} = \frac{1}{2} e^{-t/\tau} \left(1 \pm e^{-t/\tau_\phi} \cos\left(g_{\text{bs}} t + \phi\right)\right) + \mathcal{O}\left(\frac{\kappa_a - \kappa_b}{g_{\text{bs}}}\right), \tag{4.13}$$

where τ captures the driven average relaxation rate of the oscillators, and τ_ϕ captures the driven dephasing rate. From these values, we can obtain an overall timescale for errors on a single-photon state of $\tau_{\text{bs}} = 1/\tau + 1/2\tau_\phi$ (Lu *et al.*, 2023).

To reduce the aquisition time required to capture both the fine individual oscillations and the timescale for decoherence, we acquire the data in 41 short windows, each containing 21 points in a duration $t_{\text{window}} \approx 6t_{\text{bs}}$, spread over $3\tau_{\text{bs}}$. A snapshot of these data are shown for $\varphi_{\text{ext}}/2\pi = 0.32$ and $|\xi| = 1.83$ in Fig. 4.14(a), showcasing that the oscillations retain coherence even after $1000t_{\text{bs}}$. Zooming out to the full dataset in Fig. 4.14(b), we cannot make out individual oscillations on the timescale of decoherence, hence the desire to take data in short windows. Within each of these windows, $t_{\text{window}} \ll \tau_{\text{bs}}$, we can capture the mean and amplitude of the oscillations, which are plotted in Fig. 4.14(c, d), respectively. From the exponential fits to the mean, which decays on a timescale τ, and the amplitude, which decays on a timescale $(1/\tau + 1/\tau_\phi)^{-1}$, we can extract both τ and τ_ϕ, across a range of φ_{ext}. Since this can be performed on both the P_{10} and P_{01} data separately, their agreement to within error bar provides a sanity check[4].

The extracted relaxation rate (Fig. 4.14(e)) decreases very little, barely deviating from the unpumped average oscillator T_1 (given by the solid horizontal line), even as $g_{\text{bs}}/2\pi$ approaches 2 MHz. The driven dephasing times, while visibly deviating from the unpumped value, remain relatively constant as g_{bs} is increased, and much longer than the relaxation time, $\tau_\phi \gg \tau$, with single-photon loss still the dominant error.

[4]Since the SPAM errors may differ on each oscillator, the P_{01} and P_{10} data are normalized according to the measured SPAM error. The full description of this method, devised by Jacob Curtis, may be found in App. G of Chapman *et al.* (2023)

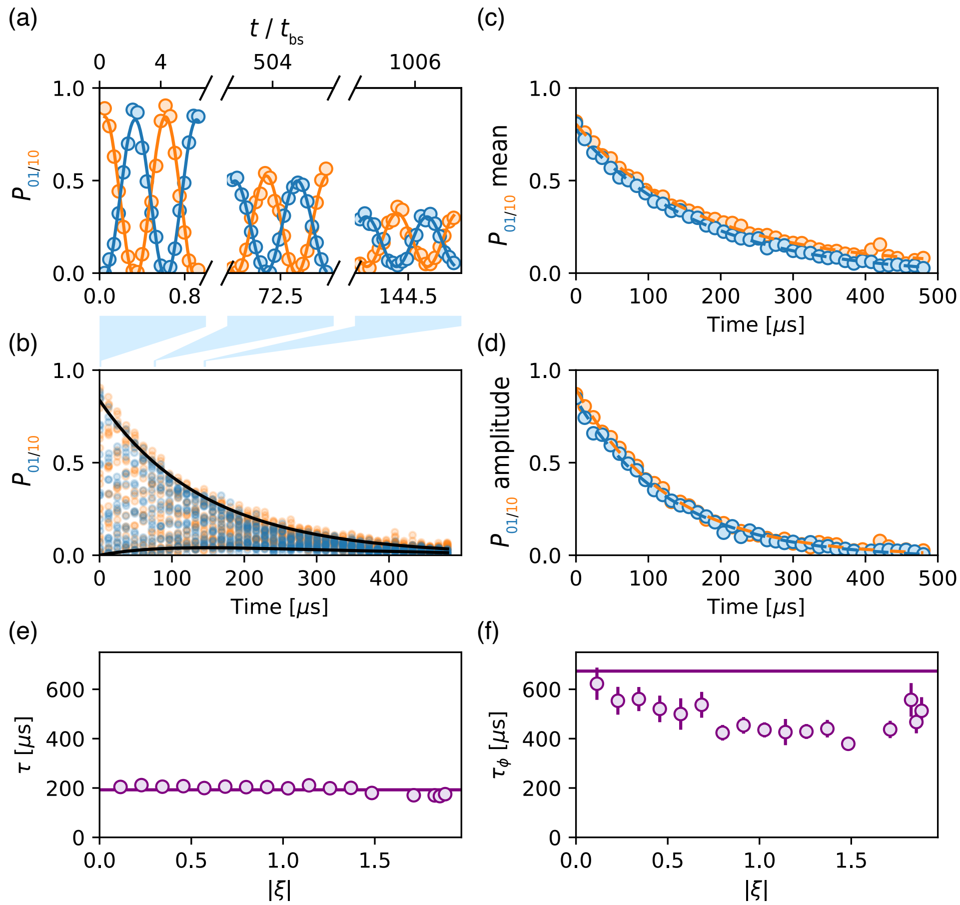

Figure 4.14: **Extracting τ_{bs} at long times.** (a-b) Probability of measuring the state $|0\rangle_A |1\rangle_B$ (blue) or $|1\rangle_A |0\rangle_B$ (orange) after initializing in$|0\rangle_A |1\rangle_B$ and applying a resonant beamsplitter for time t. Solid blue and orange lines show fits to Eq. while black lines shows the fitted envelope of the oscillations. Zoomed-in oscillations in (a) highlight that coherence is preserved, even after $t = 1000t_{\text{bs}}$, where t_{bs} is the timescale for a 50:50 beamsplitter. Data is acquired in separated windows of finely-sampled points. (c-d) The mean (c) and amplitude (d) of fitted oscillations within each window, extracted from the P_{01} (blue) and P_{10} (orange) data. Dashed lines show exponential fits. (e-f) The extracted driven relaxation timescale τ (e) and driven dephasing timescale τ_ϕ (f) as a function of normalized drive amplitude, at $\varphi_{\text{ext}}/2\pi = 0.324$. Figures modified with permission from Chapman *et al.* (2023).

4.3.3 Figures of merit

These measurements provide the two figures of merit, $\tau_{\text{bs}}/t_{\text{bs}}$ and g_{bs}/χ_{ab}, which are plotted across a range of external fluxes φ_{ext} in Fig. 4.15. The coherence of the beamsplitter oscillations remains high across a wide range of φ_{ext}, even when $g_4 \neq 0$. This highlights the benefit of a

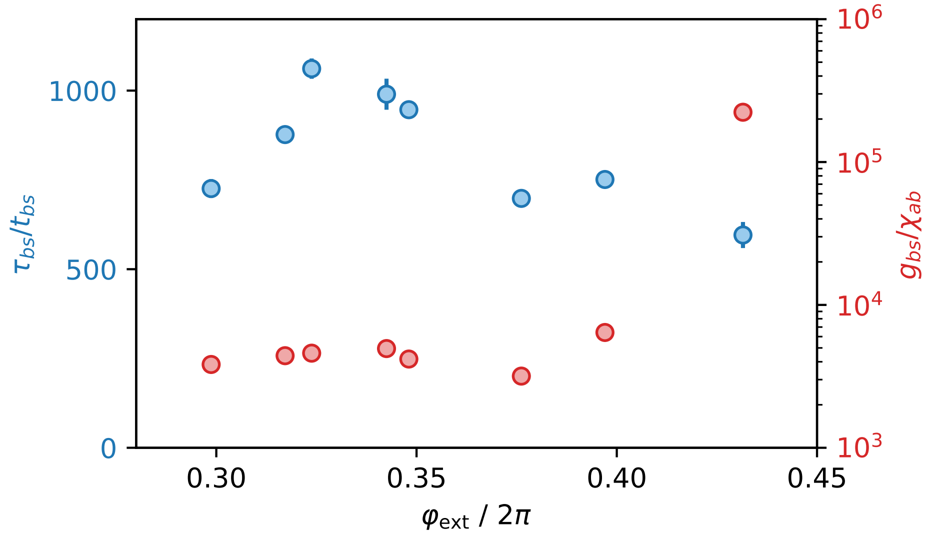

Figure 4.15: **Beamsplitter figures of merit.** Left axis: driven coherence time τ_{bs} compared to the time for one 50/50 beamsplitter, maximized over drive amplitudes $|\xi|$ across a range of external fluxes φ_{ext}. Right axis: on-off ratio $g_{\mathrm{bs}}/\chi_{ab}$ across a range of φ_{ext}, at the value of $|\xi|$ that maximizes $\tau_{\mathrm{bs}}/t_{\mathrm{bs}}$. Figure modified with permission from Chapman *et al.* (2023).

three-wave-mixing approach where the Hamiltonian parameter governing the desired interaction (g_3) can be made large relative to the one governing the undesired interactions (g_4). Meanwhile, the on-off ratio exceeds 10^3 across the range, and can be measured in excess of 10^5, limited by the sensitivity of the measurement. The plot indicates that different operating points may be desirable for different applications. For applications that are very sensitive to Kerr terms, such as for GKP qubits (Gottesman *et al.*, 2001), the χ_{ab}-free point may be optimal. Meanwhile, for applications that are only minimally sensitive to this, such as in a dual-rail qubit (Teoh *et al.*, 2023) where χ_{ab} is only (weakly) felt during two-qubit gates, the point that optimizes beamsplitter coherence is likely better.

4.3.4 Interpreting the turnaround in g_{bs}

We saw that in Fig. 4.14, coherence times did not continue to decrease, even as $g_{\mathrm{bs}}/2\pi$ approached 2 MHz. This indicates that the ultimate limitation on $\tau_{\mathrm{bs}}/t_{\mathrm{bs}}$ may actually set by the turnaround in g_{bs}. In Ch. 3.4, we discussed the fact that higher-order nonlinearities, in particular g_5, can contribute negatively to g_{bs} at large $|\xi|$. However, the predicted g_{bs} including these corrections[5] from Eq. 3.47, plotted as the solid line in Fig. 4.12(d) does not capture the sharpness of the turnaround, predicting a slightly higher maximum g_{bs}. (Using the extracted values for g_3 and g_5, the 'IP3' point from Eq. 3.49, at which the g_5 contribution to g_{bs} becomes equal to the g_3 contribution, is expected to come at $|\xi|_{\text{IP3}} = 3.2$ at this flux point.)

The likely culprit in this case is, as in the case of the transmon coupler, a 4$^{\text{th}}$-order multiphoton process. As $|\xi|$ is increased, the coupler frequency Stark shifts downwards until $2\omega_p = \omega_a + (\omega_c^0 + \Delta_{\text{Stark}}^{(c)})$, leading to a $\xi^2\hat{a}^\dagger\hat{c}^\dagger$ process. Given that $\omega_p = \omega_b - \omega_a$, there is simultaneously also a 5-wave-mixing process, $\xi^3\hat{b}^\dagger\hat{c}^\dagger$. Based on our predictions for g_n and our expression for $\Delta_{\text{Stark}}^{(c)}$, we can predict that these processes become resonant at $|\xi| = 2.4$, around the value at which the turnaround is seen. Understanding exactly how these (and other resonances) affect the beamsplitter interaction is the subject of ongoing work (Baskov, 2024).

The likely impact of this transition on the beamsplitter performance points to more careful consideration of the frequency stack. An effort to account for all possible 4-wave-mixing processes is done in Zhou (2024), but a systematic approach along the lines of Xiao *et al.* (2023) may more easily allow us to capture higher-order effects. Nonetheless, even with these possible multiphoton transitions, $|\xi|$ does approach close to $|\xi|_{\text{IP3}}$, with limited room to keep driving harder. Looking at Eq. 3.50, we can see that increasing g_{bs} beyond this point (for a fixed coupler frequency ω_c) would require increases in p, p_a, p_b or $|c_3/c_2|$ (perhaps by increasing the junction asymmetry α closer to $1/M$). While the first could be safely achieved by reducing the geometric inductance of the coupler leads (in our device, $p \approx 0.75$ at the Kerr-free point), the other changes would require greater care to avoid the cavities inheriting decoherence from the coupler.

[5] Higher-order nonlnearities $g_{n>5}$ are numerically extracted from NINA (Nonlinear Inductive Network Analyzer) tool (Miano *et al.*, 2023), rather than from analytical expressions.

Chapter 5

Ancilla dependent (and independent) parametric control

Now that we have access to a strong beamsplitter interaction between high-Q oscillators which preserves the coherence and linearity of these modes, it is time to revisit combined operations making use of both the beamsplitter and dispersive Hamiltonians. In the previous chapter, we engineered $g_{\text{bs}} \gg 1/\tau_{\text{bs}}$, enabling high-fidelity multi-mode Gaussian operations. In the process, another important threshold was reached, namely $g_{\text{bs}} \approx \chi \gg 1/T_2^{\text{ancilla}}$ - a regime in which high-fidelity multi-mode *non-Gaussian* operations become possible. Previous combined beamsplitter-dispersive operations (see Sec. 2.4) required that the nonlinear ancilla remained in $|g\rangle$ when the beamsplitter drive was applied, to minimize the impact of transmon errors during the slow beamsplitter interactions. Now that parametric control can be as fast as dispersive control, we are no longer subject to this constraint, and can expand the toolbox to include high-fidelity subroutines where the ancilla may be in any state during the beamsplitter operation.

The main goal of this chapter is to demonstrate a framework for designing and calibrating combined operations consisting of alternating applications of ancilla and beamsplitter drives, somewhat like the combination of SNAP and displacement pulses in single-oscillator control, in the $g_{\text{bs}} \approx \chi$ regime. In doing so, I will demonstrate a key example where the higher-fidelity beamsplitter offers a substantial increase in performance - an improved cSWAP, a key element

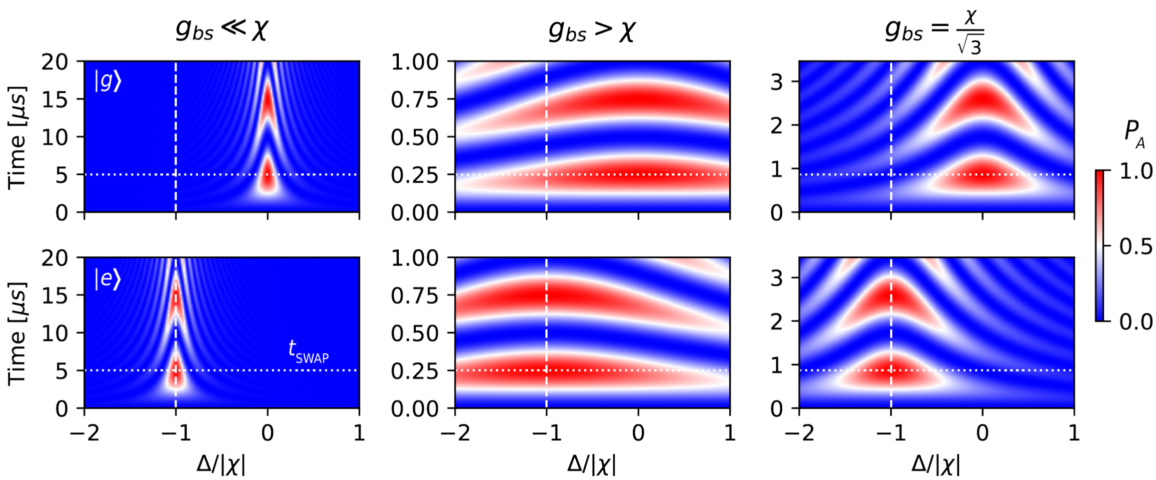

Figure 5.1: **Oscillator SWAPs in different g_{bs}/χ regimes.** Simulated probability P_A of finding a photon in Alice's oscillator after initializing a photon in Bob's oscillator and applying a beamsplitter drive at a detuning Δ. In the top (bottom) row, the ancilla is initialized in $|g\rangle$ ($|e\rangle$). From left to right, the columns show $g_{\text{bs}}/|\chi| = 0.1$, 2, and $1/\sqrt{3}$. The horizontal dotted lines indicate the time to do 1 SWAP, and the vertical lines indicate the detuning used for a cSWAP gate ($\Delta = \chi$, assuming a negative value for χ).

for qRAM. This demonstration (Ch. 5.3) also provides a way to evaluate the performance of the SNAIL-enabled beamsplitter beyond single-photon states, where protecting against Kerr nonlinear starts to be important. I will then demonstrate how this same framework unlocks new capabilities including an *unconditional* SWAP (Ch. 5.4) and a measurement of the joint-photon number parity of two modes (Ch. 5.5).

5.1 Revisiting cSWAP

While the $g_{\text{bs}} \approx |\chi|$ regime offers higher performance, it also requires increased care when designing pulse sequences. To see this, we can compare the method for cSWAP introduced in Sec. 2.4, in both the $g_{\text{bs}} \ll \chi$ and $g_{\text{bs}} \approx \chi$ regimes. As a reminder, the scheme works by applying a beamsplitter with a detuning $\Delta = \chi$, so that the oscillators undergo resonant SWAPs when the ancilla is in $|e\rangle$. Fig. 5.1 shows the time dependence of oscillator photon populations when they are initialized in $|0, 1\rangle$. In the left-most panels of Fig. 5.1, when $g_{\text{bs}} \ll \chi$ and the transmon is in $|g\rangle$, the applied beamsplitter drive is far-detuned, enacting the identity regardless of the exact choice of g_{bs} (up to a coherent error $\mathcal{O}((g_{\text{bs}}/\chi)^2)$). However in the $g_{\text{bs}} \approx \chi$ regime (i.e. the right-most panel of Fig. 5.1), the beamsplitter drive is no longer selective on the transmon state and excitations are exchanged between the two oscillators even with the ancilla in $|g\rangle$. In this regime, it is therefore important to separately consider the evolution of the oscillator states for the $|g\rangle$, $|e\rangle$ (and $|f\rangle$!) states.

In the particular example of cSWAP, we can actually find a solution by looking at these single-photon oscillations. By driving at $\Delta = \chi$ for a duration $t_{\text{SWAP}} = \pi/g_{\text{bs}}$, we perform a SWAP if the ancilla is in $|e\rangle$. At the same time, if the ancilla is in $|g\rangle$, the oscillators undergo off-resonant SWAPs with a smaller contrast, at a rate of $\sqrt{g_{\text{bs}}^2 + \chi^2}$ (see Eq. 4.11.) This can be seen from the fainter off-central white bands in the chevrons in Fig. 5.1, where the white indicates that photons have been partially exchanged and the deep blue indicates that photons have returned to the oscillator in which they were initialized. The trick is to choose g_{bs} such that at $\Delta = \chi$ and $t = \pi/g_{\text{bs}}$, with the ancilla in $|g\rangle$, the oscillators undergo exactly two off-resonant SWAPs. This ensures that states are fully swapped when in $|e\rangle$ (deep red in Fig. 5.1) and return where they began when in $|g\rangle$ (up to a phase space rotation; the deep blue in Fig. 5.1.) The condition that satisfies this can be found by setting

$$\frac{\pi}{g_{\text{bs}}} = \frac{2\pi}{\sqrt{g_{\text{bs}}^2 + \chi^2}} \tag{5.1}$$

$$\rightarrow g_{\text{bs}}^{\text{cSWAP}} = \frac{|\chi|}{\sqrt{3}} \approx 0.58\ |\chi| \tag{5.2}$$

This condition is shown in the rightmost column of Fig. 5.1, where a cSWAP is enacted at a time $t_{\text{cSWAP}} = \pi/g_{\text{bs}} = \sqrt{3}\pi/|\chi|$.

While it is possible to guess at a solution from the single-photon oscillations in this case, constructing gates on general multi-oscillator multi-photon states is more challenging than, say, gates on a two-level system where the Bloch sphere provides a useful visual representation of state trajectories. We also lack an intuitive understanding for how phases accumulate during the evolution of a gate. To provide a language for understanding the operations demonstrated in this chapter, I will briefly introduce the notation of the operator Bloch sphere, presented in

Tsunoda *et al.* (2023), on which I was a co-author, which is in turn inspired by ideas from beamsplitters in quantum optics (Campos *et al.*, 1989; Yurke *et al.*, 1986).

5.2 The operator Bloch sphere

The operator Bloch sphere formalism makes use of the Jordan-Schwinger map between angular momentum operators and photon-number-preserving bilinear combinations of creation and annihilation operators of two oscillator modes (Jordan, 1935; Schwinger, 1952):

$$\hat{L}_X \equiv \frac{\hat{a}^\dagger \hat{b} + \hat{a}\hat{b}^\dagger}{2}, \qquad \hat{L}_Y \equiv \frac{\hat{a}^\dagger \hat{b} - \hat{a}\hat{b}^\dagger}{2i}, \qquad \hat{L}_Z \equiv \frac{\hat{a}^\dagger \hat{a} - \hat{b}^\dagger \hat{b}}{2}, \tag{5.3}$$

or more compactly in terms of a two-vector of the mode operators,

$$\hat{L}_i = \begin{pmatrix} \hat{a}^\dagger & \hat{b}^\dagger \end{pmatrix} \frac{\sigma_i}{2} \begin{pmatrix} \hat{a} \\ \hat{b} \end{pmatrix}, \tag{5.4}$$

where $\boldsymbol{\sigma}_i$ are the usual Pauli matrices. We can find an angular momentum quantum number, l, via

$$\hat{L}^2 = \sum_i \hat{L}_i^2 = \frac{\hat{N}}{2}\left(\frac{\hat{N}}{2} + 1\right), \tag{5.5}$$

where the total photon number in both oscillators, $\hat{N} = \hat{a}^\dagger \hat{a} + \hat{b}^\dagger \hat{b}$, is a conserved quantity. Since the eigenvalues of the $\hat{L}^2$ operator are $l(l+1)$, we can see that the angular momentum quantum number is equivalent to half the total photon number, $l = N/2$ [1].

Their insight was that these new operators satisfy the same SU(2) commutation relations as regular angular momentum operators, $[\hat{L}_i, \hat{L}_j] = i\epsilon_{ijk}\hat{L}_k$, allowing insights from one area of quantum mechanics to be ported over to another. Schwinger initially used these operators to more efficiently describe rotations of quantum angular momenta using the better-understood language of quantum harmonic oscillators. We will instead use the map in the reverse direction.

In the beamsplitter Hamiltonian, mode operators only appear as the photon-number-preserving

[1] To maintain consistency with Schwinger's original notation, I will use $\hbar = 1$ throughout.

bilinear combinations in Eq. 5.3:

$$\hat{\mathcal{H}}_{\mathsf{bs}} = \frac{g_{\mathsf{bs}}}{2}\left(e^{i\varphi}\hat{a}^\dagger\hat{b} + e^{-i\varphi}\hat{a}\hat{b}^\dagger\right) - \Delta'\hat{b}^\dagger\hat{b}, \tag{5.6}$$

and so it can be re-expressed as

$$\hat{\mathcal{H}}_{\mathsf{bs}} = -\frac{\hat{N}\Delta'}{2} + \vec{\Omega}\cdot\hat{\vec{L}}, \tag{5.7}$$

where

$$\vec{\Omega} = \begin{pmatrix} g_{\mathsf{bs}}\cos\varphi \\ -g_{\mathsf{bs}}\sin\varphi \\ \Delta' \end{pmatrix}. \tag{5.8}$$

For the case of $l = 1/2$ (equivalently our two oscillators containing $N = 1$ total photon), this Hamiltonian is exactly that of a Rabi driven qubit where the drive detuning is conditional on the state of an ancilla. This gives us hope that we can compactly represent the dispersive-beamsplitter dynamics on a Bloch sphere. However, for states with $N > 1$, the angular momentum state vector is higher dimensional than $d = 2$ and so cannot be represented on a Bloch sphere. Furthermore, for many multiphoton encodings such as for Schrödinger cat states, the total photon number in the two oscillators (and therefore the total angular momentum) is not even a 'good' quantum number. In these cases, what is the object that we *can* represent compactly on a Bloch sphere?

The answer is that rather than considering states in the Schrödinger picture, we should instead consider the evolution of mode operators in the Heisenberg picture. As seen in Eq. 5.4, the Schwinger angular momentum operators can be expressed as 2×2 Pauli matrices in a basis defined by the mode operators. Assuming for now that the g_{bs}, φ and Δ are kept fixed, the time-independent Heisenberg evolution can be expressed as

$$\begin{pmatrix} \hat{a}(t) \\ \hat{b}(t) \end{pmatrix} = e^{i\hat{\mathcal{H}}_{\mathsf{bs}}t}\begin{pmatrix} \hat{a} \\ \hat{b} \end{pmatrix} e^{-i\hat{\mathcal{H}}_{\mathsf{bs}}t} = e^{-i\frac{\Delta'}{2}t}R_{\hat{n}}(\Omega t)\begin{pmatrix} \hat{a} \\ \hat{b} \end{pmatrix}, \tag{5.9}$$

where $R_{\hat{n}}(\theta) = \exp\{-i\theta\hat{n}\cdot\frac{1}{2}\vec{\sigma}\}$ represents a rotation on a Bloch sphere by θ about an axis $\hat{n}$. As in the Rabi case, the angular frequency $\Omega = \left|\vec{\Omega}\right| = \sqrt{g_{\text{bs}}^2 + \Delta'^2}$ and axis of rotation $\hat{n} = \vec{\Omega}/\Omega$ can be extracted from the prefactors in front of the angular momentum operators in the Hamiltonian.

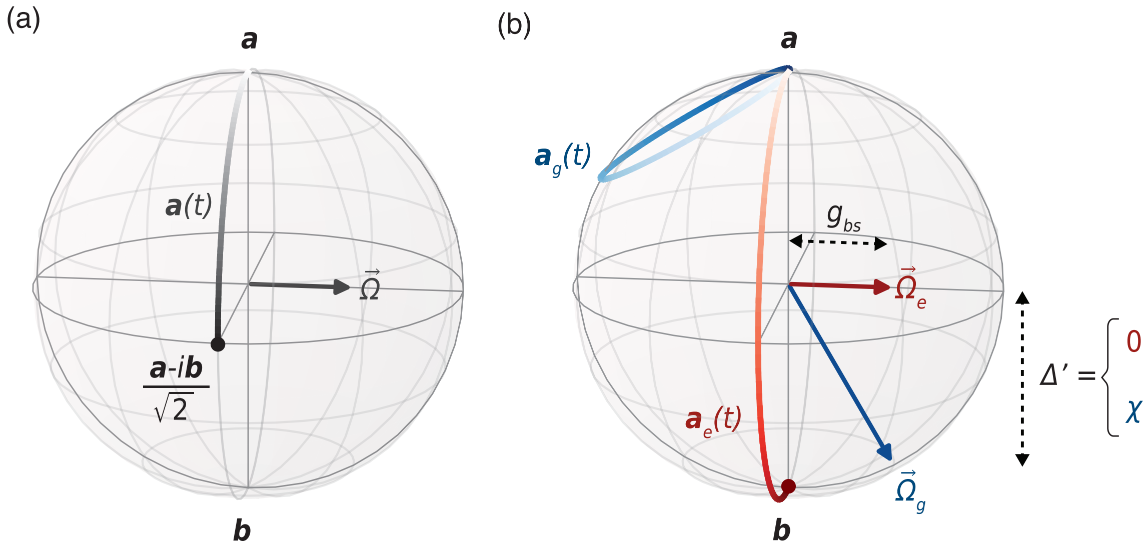

Figure 5.2: **Operator Bloch sphere trajectories.** Trajectories of the mode operator $\hat{a}$ in the Heisenberg picture for (a) 50:50 beamsplitter and (b) controlled-SWAP. The vectors indicate the axis of rotation, set by g_{bs} and Δ' (here we assume $\varphi = 0$). For the controlled-SWAP, the trajectories are conditional on the ancilla state being in $|g\rangle$ (red) or $|e\rangle$ (blue).

This compact expression provides a powerful way to describe the action of the beamsplitter Hamiltonian. If we visualize a Bloch sphere with the Schrödinger (zero time) operators $\hat{a}$ and $\hat{b}$ at the North and South poles respectively, the Heisenberg operators $\hat{a}(t)$ and $\hat{b}(t)$ can be represented as a point on this sphere, whose trajectory is determined by the beamsplitter parameters[2]. For example, any trajectory that takes the operators from one pole to the other is a SWAP, up to a phase space rotation. As another simple example (shown in Fig. 5.2(a)), applying a beamsplitter drive with $\Delta = \varphi = 0$ for a duration $t_{\text{bs}} = \pi/(2g_{\text{bs}})$ enacts a $\pi/2$

[2]While in the calculations above we considered these parameters to be time-independent, they can in principle vary in time.

rotation around the X axis. This 50/50 beamsplitter thus transforms the mode operators as

$$\begin{pmatrix} \hat{a}(t_{\mathsf{bs}}) \\ \hat{b}(t_{\mathsf{bs}}) \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ -i & 1 \end{pmatrix} \begin{pmatrix} \hat{a} \\ \hat{b} \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \hat{a} - i\hat{b} \\ \hat{b} - i\hat{a} \end{pmatrix} \tag{5.10}$$

This language is commonly used in quantum optics when describing the action of beamsplitters and phase shifters (Ou *et al.*, 1987). However, in that case one does not typically consider a continuous and in principle time-dependent beamsplitter evolution, especially occurring simultaneously with a phase shifter.

The effect of the dispersive Hamiltonian enters as an ancilla-state-dependent value of $\Delta' = \Delta - \chi \left|e\right\rangle\left\langle e\right|$, where Δ is the drive detuning when the transmon is in $\left|g\right\rangle$ and is fixed by the experimental parameters. The combined system Hamiltonian is block-diagonal in the ancilla subspace:

$$\begin{aligned} \hat{\mathcal{H}}_{\chi\mathsf{bs}} &= \hat{\mathcal{H}}_g \left|g\right\rangle\left\langle g\right| + \hat{\mathcal{H}}_e \left|e\right\rangle\left\langle e\right|, \\ &= \hat{\mathcal{H}}_{\mathsf{bs}}(g_{\mathsf{bs}}, \varphi, \Delta) \left|g\right\rangle\left\langle g\right| + \hat{\mathcal{H}}_{\mathsf{bs}}(g_{\mathsf{bs}}, \varphi, \Delta - \chi) \left|e\right\rangle\left\langle e\right|, \end{aligned} \tag{5.11}$$

allowing us us to visualize two different trajectories on the operator Bloch sphere, one for each ancilla state (or more if we are using the $\left|f\right\rangle$ or higher level of the ancilla).

Returning to the case of cSWAP, the parameter $\vec{\Omega}$ that determines the trajectory for each ancilla state is

$$\vec{\Omega}_g = \chi \begin{pmatrix} \frac{1}{\sqrt{3}} \\ 0 \\ -1 \end{pmatrix}, \qquad \vec{\Omega}_e = \chi \begin{pmatrix} \frac{1}{\sqrt{3}} \\ 0 \\ 0 \end{pmatrix}, \tag{5.12}$$

with these trajectories shown in Fig. 5.2(b). Indeed for this choice of g_{bs} and Δ the operators are exchanged when the transmon is in $\left|e\right\rangle$ and return to their original positions when the transmon is in $\left|g\right\rangle$.

An important distinction between the operator Bloch sphere and the qubit Bloch sphere is the treatment of the global phases, such as the one imparted by the first 'identity' term in the Hamiltonian in Eq. 5.7. Whereas for a qubit, such a term imparts a global phase on the *state* with no physical impact, on the operator Bloch sphere, the phase is applied to the *operators* in the Heisenberg picture $\hat{a}, \hat{b} \to \hat{a}e^{-i\phi/2}, \hat{b}e^{-i\phi/2}$. Returning to the Schrödinger picture, this corresponds to a phase space rotation of both modes, $\exp\Big(i\phi(\hat{a}^\dagger\hat{a} + \hat{b}^\dagger\hat{b})/2\Big)$ which, as we shall see, *is* measurable.

A powerful general rule is that the phase ϕ is equal to the solid angle enclosed by the trajectory on the operator Bloch sphere. For the case of a single-enclosed closed loop on the Bloch sphere with fixed beamsplitter parameters, this phase is equal to

$$\phi = 2\pi\left(1 - \cos\theta\right) = 2\pi\left(1 - \frac{\Delta'}{\Omega}\right), \tag{5.13}$$

where θ is the angle the rotation axis $\hat{n}$ makes with the Z-axis. Applied to the case of cSWAP, we can quickly find ancilla-state dependent phases $\phi_g = 2\pi - \sqrt{3}\pi$ and $\phi_e = 2\pi$. These phases are not something we could have obtained just by looking at the single-photon trajectories in Fig. 5.1, and they are important! While it is straightforward to keep track of global phase space rotations "in software", ancilla-state dependent phases will leave the ancilla and oscillator states entangled. Correcting for these phases requires a gate "in hardware".

5.3 Generating multimode entanglement with a fast cSWAP

This new fast cSWAP can be used to generate a "cat-in-two-boxes state" (Wang *et al.*, 2016) - an entangled Bell state of two single-oscillator cat qubits encoded in the basis $|0\rangle_L \equiv |+\alpha\rangle$, $|1\rangle_L \equiv |-\alpha\rangle$[3]:

$$\left|\Psi^+\right\rangle \equiv \frac{|0\rangle_L |1\rangle_L + |1\rangle_L |0\rangle_L}{\mathcal{N}} \tag{5.14}$$

$$= \frac{|+\alpha\rangle |-\alpha\rangle + |-\alpha\rangle |+\alpha\rangle}{\mathcal{N}}, \tag{5.15}$$

[3] The two 'basis' (coherent) states have finite overlap for small $|\alpha|$ and so this basis is only *quasi*-orthonormal as $|\alpha|$ becomes large.

where $\mathcal{N} \approx \sqrt{2}$. As discussed in Ch. 1, entanglement generation is not possible with beamsplitters alone. However, when combined with dispersive control, we *can* generate multi-mode entanglement, and access to a better beamsplitter allows us to achieve this entanglement with much higher fidelity. Furthermore, for $\alpha = \sqrt{2}$, this state contains on average 4 photons between the two oscillators, providing a way to evaluate the performance of the beamsplitter on multiphoton states. In particular, compared to the single-photon analysis in Ch. 4, this exposes the state to Kerr nonlinear terms in the Hamiltonian and multiphoton transitions involving higher photon numbers.

We generate the Bell state using the circuit shown in Fig. 5.5(a), where the two oscillators are initialized in coherent states with opposite phase and the cSWAP is used to perform a two-mode SWAP-test measurement. As we saw in Ch. 2.4, the SWAP-test circuit can be used measure the similarity of two bosonic states, but also to project onto their symmetric and antisymmetric subspaces:

$$\text{SWA}\hat{\text{P}}\text{-test}\,|g\rangle\,|\alpha\rangle\,|-\alpha\rangle = \frac{|g\rangle}{\sqrt{2}}\frac{|\alpha\rangle\,|-\alpha\rangle + |-\alpha\rangle\,|\alpha\rangle}{\mathcal{N}} + \frac{|e\rangle}{\sqrt{2}}\frac{|\alpha\rangle\,|-\alpha\rangle - |-\alpha\rangle\,|\alpha\rangle}{\mathcal{N}} \tag{5.16}$$

$$= \frac{|g\rangle\,|\Psi^+\rangle + |e\rangle\,|\Psi^-\rangle}{\sqrt{2}} \tag{5.17}$$

By post-selecting on shots where the ancilla is measured in $|g\rangle$, we can obtain the Bell state $|\Psi^+\rangle$.

The experimental hardware contains a transmon coupled to each cavity, giving us a choice about which to use as the control qubit. Typically, the primary consideration would be the expected infidelity due to transmon errors. In the fast g_{bs} regime, the speed of operations is set by χ and in the regime where $T_2^{(\text{cavity})}/2|\alpha|^2 \gg T_2^{(\text{transmon})}$, errors are dominated by transmon errors. In this picture, the transmon to choose is the one that minimizes $1/|\chi|T_2^{\text{transmon}}$, which in our case would be Alice's. However, in this system, with $2|\alpha|^2 = 4$, the cavity relaxation is already comparable to the transmon decoherence, resulting in a larger penalty (in terms of cavity relaxation) for Alice's slower χ. On a more practical level, unlike Bob's transmon, Alice's has a non-negligible coupling to the beamsplitter drive port, owing to its greater insertion into its cavity (see Ch. 4.1.1). This results in a significant Stark shift of Alice's transmon frequency when the beamsplitter drive is applied, complicating the operation of the cSWAP. So, primarily for ease-of-use, Bob's transmon plays the role of the control qubit for the subsequent experiments.

The external flux through the SNAIL, $\varphi_{\text{ext}}/2\pi = 0.32$, is chosen to be the one for which the greatest beamsplitter fidelity was seen in Fig. 4.15. Note that this is not the χ_{ab}-free point, but one of a wide range of operating points where the on-off ratio > 2000. This performance should therefore be representative of how well the cSWAP can work over a range of SNAIL flux points (and therefore frequencies) where χ_{ab} is significantly suppressed relative to a transmon coupler but not necessarily exactly zero.

5.3.1 Calibration of a fast cSWAP

Calibrating the high-fidelity cSWAP starts by tuning up the beamsplitter amplitude g_{bs}, drive frequency ω_d and pulse duration t_{cSWAP}. The calculations that lead to the ideal parameters, $g_{\text{bs}} = |\chi|/\sqrt{3}$, $\omega_{\text{p}} = \omega_b^0 - \omega_a^0 + \chi$ and $t_{\text{cSWAP}} = \sqrt{3}\pi/|\chi|$, ignore three practically complicating factors that must be accounted for:

1. **Cavity Stark shifts.** As we saw in Fig. 4.12, increasing g_{bs} shifts the relative frequency difference between the cavities, $\omega_b - \omega_a$, and so ω_{p} needs to shift to ensure that $\Delta = \chi$, and the beamsplitter is resonant when the control qubit is in $|e\rangle$.

2. **Finite bandwidth.** The implementation of the buffer mode (see Ch 4.1.3) enables a large $g_{\text{bs}}^{\text{max}}$, but also imposes a bandwidth constraint on the beamsplitter pulses. The origin of this constraint is that the buffer mode resonance lies ~ 20 MHz detuned from $\omega_b - \omega_a$, and so if the bandwidth is not much smaller than 20 MHz, the beamsplitter drive populates the buffer mode with a large number of photons. Since the buffer mode is relatively weakly coupled to the beamsplitter drive port, it has a $T_1 \approx 5$ μs, emptying these photons slowly. While it does so, its non-negligible dispersive shift to other system modes results in shot noise induced dephasing. In practice, ramping g_{bs} to its maximum value in $t_{\text{ramp}}^{(\text{bs})} \geq 80$ ns avoids this. While this particular bandwidth-constraint is specific to the system hardware

used here, any pump port filter will impose a finite bandwidth, requiring a ramp up to the maximum $g_{\textsf{bs}}$. During this rise time however, when $g_{\textsf{bs}} < g_{\textsf{bs}}^{\textsf{max}}$, the exchange of cavity states will proceed more slowly and so $g_{\textsf{bs}}^{\textsf{max}}$ and $t_{\textsf{cSWAP}}$ need to be adjusted to account for this.

3. **State-dependent $g_{\textsf{bs}}$.** In our system, the beamsplitter rate $g_{\textsf{bs}}$ depends to a noticeable degree on the state of the control qubit. When Bob's qubit is used as the control, the same room temperature pump amplitude achieves approximately 5% higher $g_{\textsf{bs}}$ when the qubit is in $|e\rangle$ and approximately 10% higher $g_{\textsf{bs}}$ when the qubit is in $|f\rangle$, relative to when the qubit is in $|g\rangle$. The opposite effect is seen when Alice's qubit is used (a decrease in $g_{\textsf{bs}}$ with increasing qubit energy level). This is consistent with a decrease (increase) in the beamsplitter resonance frequency with increasing energy level in Bob's (Alice's) control qubit, bringing it closer (further) from the buffer mode resonance. As can be seen from Fig. 4.4(b), $|\xi|$ changes steeply with frequency at our operating point, with a relative gradient of -5.2% per MHz. This requires an adjustment of parameters to ensure the $|g\rangle$ and $|e\rangle$ curves complete their desired trajectories at the same time.

We can account for these factors using an iterative approach. Given an initial guess of $g_{\textsf{bs}}$, performing the single-photon oscillation experiment from Fig. 4.12 with the control in $|e\rangle$ allows us to extract (from the center frequency of the chevrons) the drive frequency ω_{p} that maximizes the contrast of single-photon oscillations. With this choice of $g_{\textsf{bs}}$ and ω_{p} set, we can measure single-photon oscillations with the control initialized in $|g\rangle$ or in $|e\rangle$ (see Fig. 5.3(a)). These oscillations can be thought of as vertical linecuts of the colorplots in Fig. 5.1. Ideally, in the time that a single SWAP is completed with the control in $|e\rangle$, we should see a full period of oscillation with the control in $|g\rangle$. If we observe over- (or under-) rotation with the control in $|g\rangle$, this indicates that $g_{\textsf{bs}}$ should be increased (decreased). Upon updating $g_{\textsf{bs}}$, the process then repeats (finding a new ω_{p} for this $g_{\textsf{bs}}$) until the desired cSWAP operation is observed, as in Fig. 5.3(b).

With $g_{\textsf{bs}}$ and ω_p fixed, the cSWAP is now calibrated *up to* phase space rotations on each

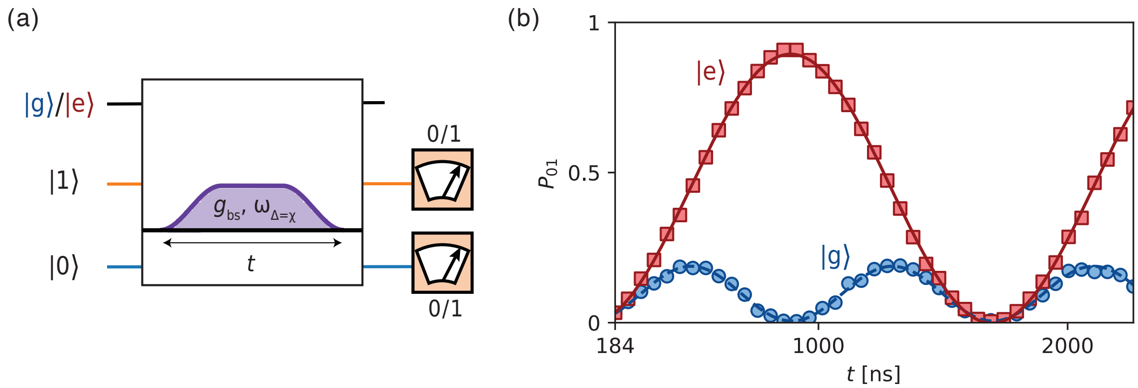

Figure 5.3: **Single-photon oscillations under cSWAP.** (a) Pulse sequence, starting with a single photon in Alice, applying a flat-top beamsplitter pulse with maximum amplitude $g_{\text{bs}} \approx \sqrt{3}/2 \times |\chi|$, detuning $\Delta = |\chi|$ and variable duration (including 92 ns ramps), t, followed by a simultaneous photon number measurement of both oscillators. This sequence can be repeated with the control in $|g\rangle$ or in $|e\rangle$. When g_{bs} and drive frequency ω_p are set correctly, the oscillations appear as in (b). When the control is in $|g\rangle$ the single photon returns to Alice in the same time that it takes for it to SWAP to Bob when the control is in $|e\rangle$. The lack of full contrast in $|e\rangle$ is due to errors in preparing $|1\rangle_A |0\rangle_B$. Figure modified with permission from Chapman *et al.* (2023).

cavity. Besides the phase we expect from the ideal unitary evolution, the experimental setup can introduce some additional phases. The beamsplitter drive can have a phase offset relative to the difference in the phases of the cavity drives (either by deliberating applying a phase offset to the generated microwave signal or because of mismatched electrical lengths), which results in an additional phase space rotation when states swap from one cavity to the other. Furthermore, the cavity frequencies experience a Stark shift when the beamsplitter is applied, by differing amounts on each oscillator.

In general, the unitary (up to this point) is described by four rotation angles,

$$\hat{U} = \left(e^{i\left(\phi_{a,g}\hat{a}^\dagger\hat{a}+\phi_{b,g}\hat{b}^\dagger\hat{b}\right)}|g\rangle\langle g| + e^{i\left(\phi_{a,e}\hat{a}^\dagger\hat{a}+\phi_{b,e}\hat{b}^\dagger\hat{b}\right)}|e\rangle\langle e|\right)\text{cSWAP} \tag{5.18}$$

$$= \left(e^{i\left(\phi_{a,g}\hat{a}^\dagger\hat{a}+\phi_{b,g}\hat{b}^\dagger\hat{b}\right)}e^{i\left((\phi_{a,e}-\phi_{a,g})\hat{a}^\dagger\hat{a}+(\phi_{b,e}-\phi_{b,g})\hat{b}^\dagger\hat{b}\right)|e\rangle\langle e|}\right)\text{cSWAP}, \tag{5.19}$$

where in the final line we have re-written this as a control-state-independent and a control-state-dependent phase on each oscillator.

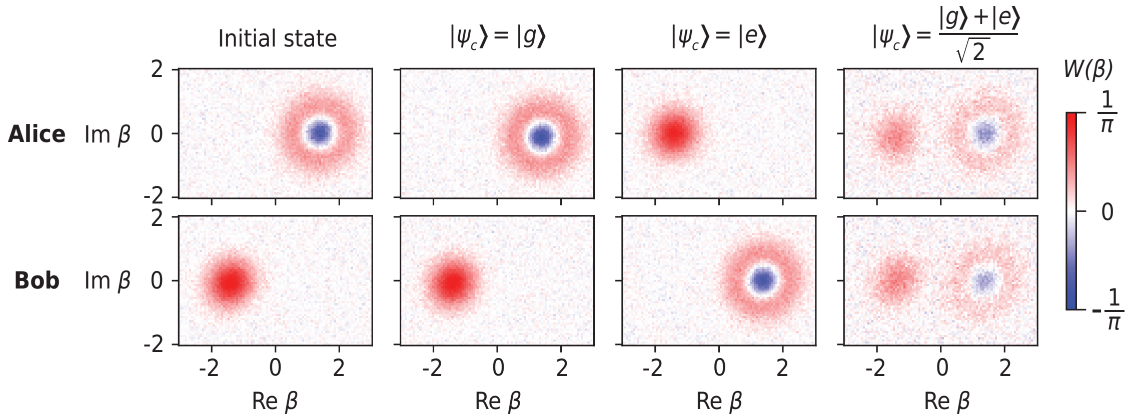

Figure 5.4: **Verifying cSWAP on multiphoton states with Wigner tomography.** Single-oscillator Wigner tomography results after initializing in $D(\sqrt{2})\,|1\rangle \otimes D(-\sqrt{2})\,|0\rangle$ (left column) and after performing a cSWAP with the control qubit in $|g\rangle$, $|e\rangle$ or a superposition $(|g\rangle+|e\rangle)/\sqrt{2}$. The states are neither exchanged or rotated when the control is in $|g\rangle$, and are swapped without rotation with the control in $|e\rangle$. For a superposition control state, the Wigner function shows a 50:50 mixed state of the final states in the central two columns. Figure modified with permission from Chapman *et al.* (2023).

All four phases can be obtained by preparing the oscillator in the state $|\psi_{\text{init}}\rangle = D(\alpha)\,|1\rangle \otimes D(-\alpha)\,|0\rangle$, with the transmon either in $|\psi_c\rangle = |g\rangle$ or in $|\psi_c\rangle = |e\rangle$, performing $\hat{U}$ and then performing Wigner tomography on both cavities. This choice of initial state is both sensitive to phase space rotations and clearly distinguishes between the states initialized in Alice and Bob via the negativity of the Wigner function. The resulting states are given by

$$\hat{U}\left(|\psi_{\text{init}}\rangle \otimes |g\rangle\right) = D(\alpha e^{i\phi_{a,g}})\,|1\rangle \otimes D(-\alpha e^{i\phi_{b,g}})\,|0\rangle \otimes |g\rangle \tag{5.20}$$

$$\hat{U}\left(|\psi_{\text{init}}\rangle \otimes |e\rangle\right) = D(-\alpha e^{i\phi_{a,e}})\,|0\rangle \otimes D(\alpha e^{i\phi_{b,e}})\,|1\rangle \otimes |e\rangle \tag{5.21}$$

From the four tomograms, we can fit to find the center of the displaced state. The angle this makes with the x-axis tells us the value of each phase.

Two of these phases, $\phi_{a,g}$ and $\phi_{b,g}$, (which make up the control-state-independent part of Eq. 5.19) can be corrected “in software” by rotating the axes along which the displacements for the Wigner tomography are performed. The remaining two must be corrected in hardware, which we can do by adding delays t_{pre} and t_{post} before and after the beamsplitter pulse. These delays make use of the always-on dispersive interaction $\hat{\mathcal{H}}_{\mathsf{disp}} = \hbar\chi\hat{b}^\dagger\hat{b}\left|e\right\rangle\left\langle e\right|$ to add phase space rotations conditional on the control being in $|e\rangle$. The pre-delay adds a rotation by an angle χt_{pre} to the state initialized in Bob and ending in Alice after the SWAP. Choosing $t_{\mathsf{pre}} = (\phi_{a,e} - \phi_{a,g})/\chi$ therefore corrects the control-state-dependent phase on Alice. Similarly, choosing a post-delay $t_{\mathsf{post}} = (\phi_{b,e} - \phi_{b,g})/\chi$ fixes the control-state-dependent phase on Bob. The total duration of the cSWAP, including these delays is therefore $1.3\ \mu$s, limited by the value of $|\chi|$.

Fig. 5.4 shows these tomograms after correcting for the phases, with the states exchanged only if the control is in $|e\rangle$ and without any additional phase space rotation. A final check is shown in the right-most column where the control is prepared in a superposition state. These last Wigner functions in each cavity both show mixed states that are a 50/50 mixture of the two initial states. Critically, we see no visible distortion of the displaced states (as we would expect from a large Kerr nonlinearity, for example) and we see that they lie directly opposite each other on the x-axis, indicating that we have successfully removed any control-state-dependent rotations. This verifies the calibrations but does not yet demonstrate entanglement between the two modes.

5.3.2 Probing multi-mode entanglement

Evaluating whether we have truly generated multimode entanglement requires not just local measurements but joint measurements of the two cavities. This can be revealed by measuring the joint Wigner function (Cahill and Glauber, 1969; Wang *et al.*, 2016), which requires a measurement of the joint photon number parity in the two oscillators, or equivalently the product of the local parity in each oscillator[4], $\hat{P}_A \otimes \hat{P}_B \equiv e^{i\pi\hat{a}^\dagger\hat{a}} \otimes e^{i\pi\hat{b}^\dagger\hat{b}}$:

$$W_J(\beta_A, \beta_B) = \langle \hat{D}(\beta_A, \beta_B)\hat{P}_A\hat{P}_B\hat{D}^\dagger(\beta_A, \beta_B)\rangle, \tag{5.22}$$

where β_A and β_B are the quadratures for Alice and Bob, respectively. We can measure the simultaneous local parity by performing a single-shot parity measurement of each oscillator and

[4]This definition, which takes values between -1 and 1, discards the factor of $1/\pi$ used in the single-mode Wigner function (Eq. 2.7).

multiplying the results together shot-by-shot[5] While local measurements collapse the entanglement between the oscillators, this is permissible here since we are not interested in utilizing the post-measurement state.

Since β_A and β_B are both complex, the joint Wigner function is a 4-dimensional object that is hard to visualize. However, evidence of entanglement can be seen in 2-dimensional slices along the $\text{Im}(\beta_A) = \text{Im}(\beta_B) = 0$ (real-real) plane and the $\text{Re}(\beta_A) = \text{Re}(\beta_B) = 0$ (imaginary-imaginary) plane. Interference fringes along the diagonals of the measured slices in Fig. 5.5(b, c) indicate the generation of an entangled non-Gaussian state, and match qualitatively with the ideal plots in Fig. 5.5(f, g).

In principle, joint Wigner tomography allows an experimentalist to reconstruct the full density matrix of two oscillators (Leibfried *et al.*, 1996), for example using a maximum likelihood estimation (MLE) algorithm (Chou *et al.*, 2018; Heeres *et al.*, 2015) or machine learning techniques (Ahmed *et al.*, 2021) to account for SPAM errors. This in turn provides a way to compute the fidelity to an arbitrary two-oscillator state. This approach is extremely costly however, requiring n^4 well-averaged measurements, where n is the number of points along each dimension. For a reasonable resolution ($n \approx 30$ and > 100 shots per measurement), this would take about a week to obtain. This continuous-variable representation of the density matrix is rich with detail but is extremely overcomplete if we are only interested in the evolution within our logical subspace or the logical fidelity. The density matrix in the logical basis of the two encoded qubits only contains $2^4 = 16$ entries, which can be efficiently extracted by sampling the joint Wigner function at 16 specific points, a method known as direct fidelity estimation, proposed in (Flammia and Liu, 2011; da Silva *et al.*, 2011) and demonstrated for cat states in (Vlastakis *et al.*, 2015; Wang *et al.*, 2016). For a single-mode, large-α, encoded two-legged cat, the logical

[5] Transmon relaxation errors can bias the results of the Wigner tomography, leading to an overall offset. To combat this, the results are symmetrized – averaging over all four combinations of mapping even or odd photon number parity in each oscillator to the transmon state $|e\rangle$, following the same approach in single-oscillator tomography (Burkhart *et al.*, 2021).

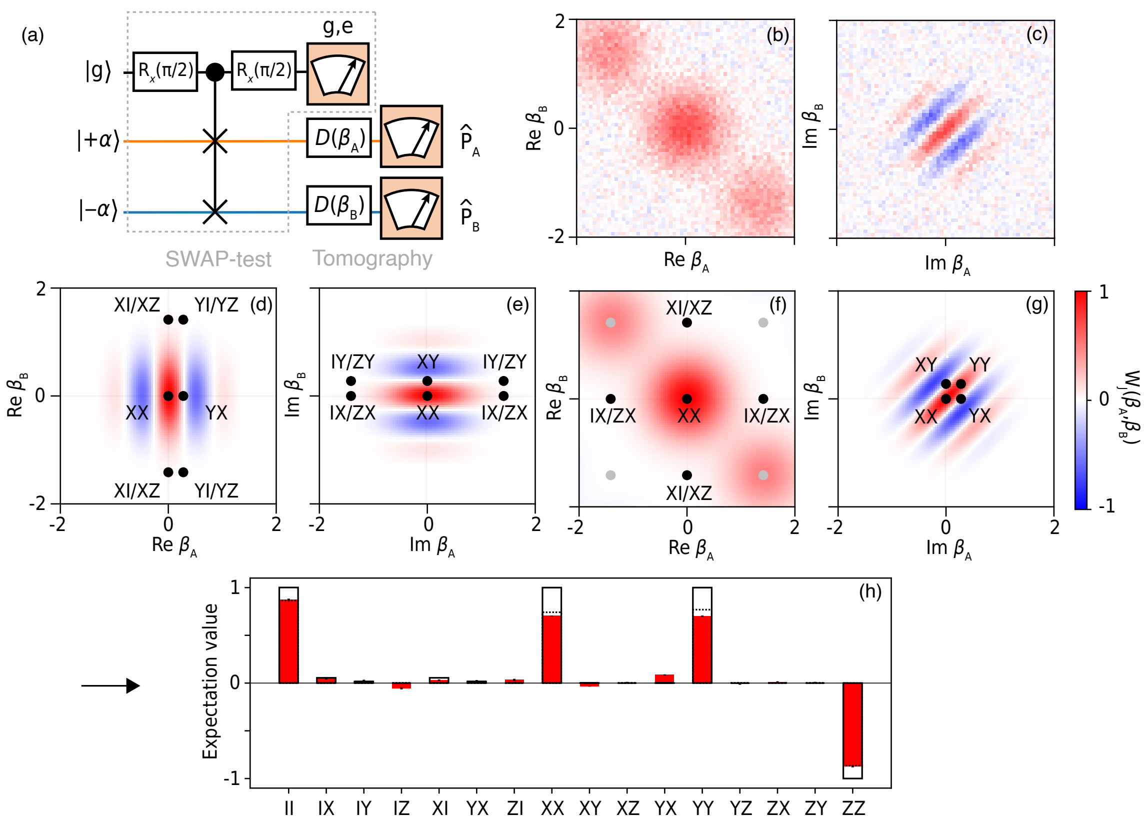

Figure 5.5: **Characterizing fidelity of Bell state.** (a) Pulse sequence used to prepare the Bell state $|\Psi^+\rangle$ (SWAP-test) and perform joint Wigner tomography by simultaneously measuring photon number parity $\hat{P}$ on each cavity. (b-c) Re-Re and Im-Im 2D slices of measured joint Wigner function, displaying interference fringes. (d-g) Joint Wigner function slices for ideal $|\Psi^+\rangle$ state, with dots indicating locations at which to extract two-mode Pauli expectation values in the coherent state basis $\{|+\alpha\rangle, |-\alpha\rangle\}$. Gray dots indicate four different locations that must be measured to extract II and ZZ. (h) Extracted two-mode Pauli expectation values (red), ideal values (solid black) and predicted values from error budget including SPAM errors (dotted black). The YY bar has been boosted by a factor of $e^{\pi^2/16|\alpha|^2} = 1/0.73$ to account for the effects of finite-α when extracting the Pauli values using this method (see Appendix C.) Figure modified with permission from Chapman *et al.* (2023).

expectation values can be extracted from

$$\langle I\rangle \approx W(+\alpha) + W(-\alpha), \qquad \langle X\rangle \approx W(0), \tag{5.23}$$

$$\langle Y\rangle \approx W\left(\frac{j\pi}{8\alpha}\right), \qquad \langle Z\rangle \approx W(+\alpha) - W(-\alpha). \tag{5.24}$$

The two-mode logical expectation operations are then constructed by linear combinations of these points on the two oscillators, so for example

$$\langle IX \rangle \approx W_J\left(+\alpha, 0\right) + W_J\left(-\alpha, 0\right). \tag{5.25}$$

The full set of sampled two-oscillator points are overlaid on the joint Wigner function of a simulated cat-in-two-boxes state in Fig. 5.5(d-g). Appendix C lists the Wigner measurements taken for each two-mode expectation value, along with expressions describing the corrections due to the finite α of the cat state.

Performing these measurements on our prepared state yields the expectation values of all 16 two-mode Pauli operators (see red bars in Fig. 5.5(h)). For a Bell-state, simultaneous measurements of the same Pauli operator on each qubit (i.e. II, XX, YY and ZZ) should give perfectly correlated or anti-correlated results, whereas measuring different Pauli operators on each should yield uncorrelated results. The ideal values for a state with $\alpha = \sqrt{2}$ are shown in solid black rectangles. Note that the YY bar has been re-normalized by a factor of $e^{\pi^2/16|\alpha|^2} = 1/0.73$. This difference is due to the fact that the cat state basis $\{\left|+\alpha\right\rangle, \left|-\alpha\right\rangle\}$ is only completely orthogonal in the limit $\alpha \gg 1$. As a result, for this finite-sized cat, even an ideal decoherence-free state would yield at most $\langle YY \rangle = 0.73$. Meanwhile values obtained from an error budget composed from independently obtained system coherences (see App. N of Chapman *et al.* (2023) for more details) are given by the smaller dashed rectangles. These predictions give a good qualitative agreement with the heights of the measured (red) bars, suggesting that the operation is very close to coherence-limited.

The Bell fidelity can be efficiently extracted from only four of these points, $\mathcal{F} = (\langle II \rangle + \langle XX \rangle + \langle YY \rangle - \langle ZZ \rangle)/4$, from which we obtain a fidelity $\mathcal{F} = (74.1 \pm 0.4)\%$, without normalizing for SPAM errors or finite-α (when applying the finite-α correction factor this fidelity increases to $\mathcal{F}_{\text{corrected}} = (78.6 \pm 0.4)\%$.) This number exceeds the value (50%) that can be obtained from a 'classical' unentangled state $\rho^{(A)} \otimes \rho^{(B)}$, and is close to the error-budget value of $\mathcal{F} = 75.9\% \pm 0.2\%$ – or $\mathcal{F}_{\text{corrected}} = (81.0 \pm 0.3)\%$. However, this also predicts that the

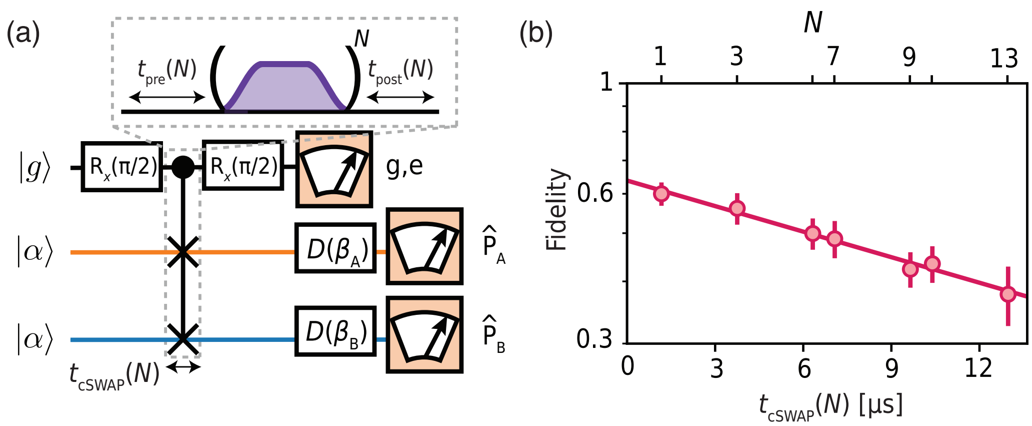

Figure 5.6: **Repeated cSWAP to extract fidelity.** (a) Pulse sequence incorporating odd N number of consecutive cSWAP. A single pre-delay t_{pre} and post-delay t_{post} are include for each N, and so they depend on N. Sampling the Wigner function at different points allows for efficient fidelity extraction. (b) Fidelity for N rounds of cSWAP versus total elapsed time. Solid line shows exponential fit. Figure modified with permission from Chapman *et al.* (2023).

vast majority of the infidelity comes not from errors during the cSWAP itself ($4.2\% \pm 0.1\%$) but from measurement errors during the SWAP-test readout and during the Wigner tomography of each cavity.

To validate the claim that the cSWAP itself is not the limiting factor, with a fidelity much higher than suggested by this single-round Bell fidelity, we can amplify its errors by replacing the single cSWAP with an odd number N of cSWAPs, making use of the fact that $\widehat{\text{cSWAP}}^2 = \hat{\mathbb{I}}$ (see Fig. 5.6). Rather than adding a delay before and after every single beamsplitter pulse, we can simplify the sequence by adding a single delay before and after the train of beamsplitter pulses, which can be calibrated in the same way as for a single cSWAP. One note to make is that since the delay lengths obtained are not constant as a function of N, the total sequence duration is not linear in N (as can be seen from the top x-axis of Fig. 5.6(b)). This figure shows the reduction in the state fidelity with N and from the exponential decay of this curve as a function of sequence length, we can extract a SPAM-corrected fidelity for a single cSWAP of $1 - \mathcal{F}_{\text{cSWAP}} = 4.5\% \pm 0.2\%$, consistent with the error budget value of $4.2\% \pm 0.1\%$, which finds transmon and cavity errors almost equally responsible. This highlights the large increase in cSWAP fidelity that can only be made possible through access to a faster beamsplitter interaction

– with the previous generation of couplers, $t_{\text{cSWAP}} \approx 10\ \mu$s and $T_2^{\text{(transmon)}} \approx 30\ \mu$s, and so even setting aside cavity errors, this would contribute an infidelity per cSWAP on the order of $1 - \mathcal{F}_{\text{cSWAP}} \approx 33\%$.

More importantly, it highlights our ability to execute gates combining beamsplitter and dispersive interactions in the $g_{\text{bs}} \approx |\chi|$ regime where the operator Bloch sphere formalism becomes useful, and to do so very close to the coherence limit. This sets the stage for performing other useful operations which rely on the operator Bloch sphere picture.

5.4 Unconditional SWAP using dynamical decoupling

A high-fidelity *conditional* SWAP provides a non-Gaussian resource that can be used to generate entanglement but it can also be very valuable to have a SWAP (or beamsplitter) that is *unconditional* on the state of the transmon. We already saw how the photonics-inspired construction of the cSWAP in Ch. 2.4 relied on beamsplitters that are implicitly control-state-independent – however in our (superconducting) system, the Kerr nonlinearity is always 'on'. More generally, an unconditional SWAP (uSWAP) allows you to construct sub-circuits in which the control input state may be unknown.

5.4.1 Demonstrating a uSWAP

The uSWAP construction relies on a form of dynamical decoupling, an idea which has been extensively used to both filter out low-frequency noise but also to suppress the effect of crosstalk between adjacent qubits in larger processors when performing simultaneous operations (Tripathi *et al.*, 2022). It is possible to engineer dynamical decoupling in this bosonic system by driving the beamsplitter at a 'symmetric' detuning of $\Delta = \chi/2$ and breaking up the sequence with echo pulses on the ancilla (Tsunoda *et al.*, 2023). For the simplest case of an unconditional SWAP with a single echo, the pulse sequence and the control-state-dependent Bloch sphere trajectories are shown in Fig. 5.7(a, b). Unlike the conditional SWAP presented earlier, the unconditional SWAP can be engineered with any value of g_{bs}. Taking the case of a single echo, provided that

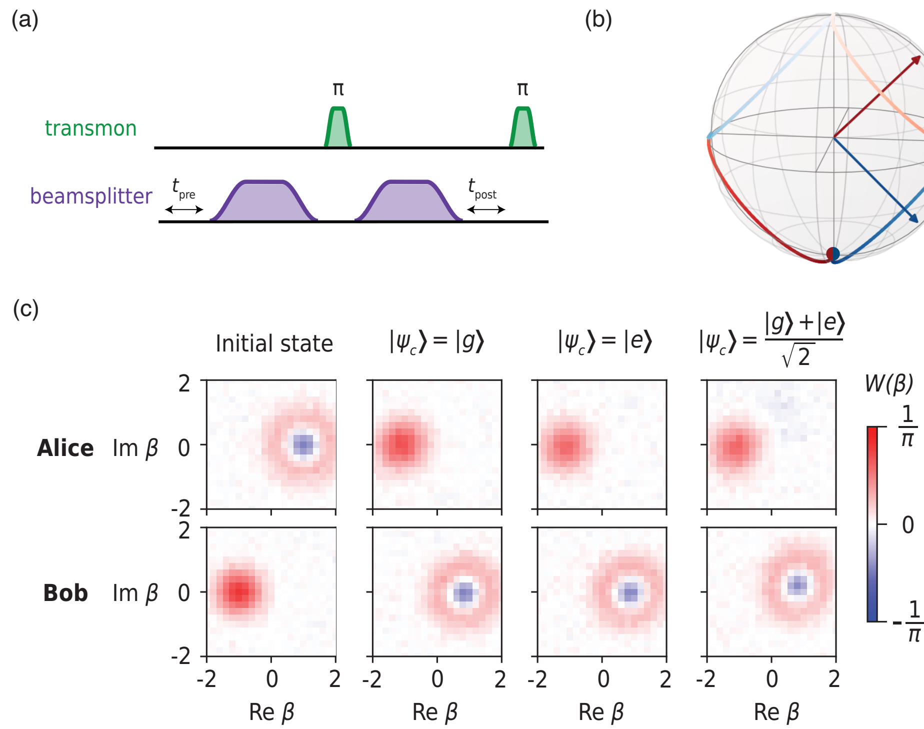

Figure 5.7: **Unconditional SWAP.** (a) Pulse sequence for the transmon and beamsplitter drives (not to scale). The beamsplitter pulse is interrupted halfway by a transmon π-pulse which 'echoes' out the dependence on the transmon state. The final π-pulse, to return the transmon to its initial state, must occur after the final delay to allow phases to be corrected on both oscillators. (b) Control-state-dependent trajectories of the mode operators. If the control is initialized in $|g\rangle$, the mode operator follows the blue curve (about the blue axis, $\hat{n}_g$) until it reaches the equator of the Bloch sphere. At this point, the π-pulse flips the control to $|e\rangle$, and the trajectory then follows the red line (about the red axis, $\hat{n}_e$) until it completes its journey to the South Pole. (c) Local Wigner tomography of both oscillators for different control states $|\psi_c\rangle$ after initializing the oscillators in $D(\sqrt{2})\,|1\rangle \otimes D(-\sqrt{2})\,|0\rangle$. These show the states being exchanged, independent of the control state, even if it starts in a superposition.

$g_{\text{bs}} > |\chi|/4$, the initial trajectory of the mode operators will intersect the equator of the Bloch sphere. Performing an X_π pulse on the qubit at this point exchanges the rotation axes, $\hat{n}_g$ and $\hat{n}_e$, such that the mode operators end up at the opposite pole. For values of $g_{\text{bs}} < |\chi|/4$, we can break the sequence up with more and more echoes to achieve the desired unitary. The unconditional 50/50 beamsplitter (uBS) can be constructed in an analogous way, flipping the qubit state partway through the sequence, so that the trajectories both end on the equator independent of the initial control qubit state.

We can experimentally verify that this sequence dynamically decouples the beamsplitter from the ancilla state for a uSWAP with a single echo and $g_{\mathsf{bs}} = |\chi|/4$. We do so by repeating the single-cavity tomography experiment used to characterize the cSWAP, as shown in Fig 5.7(c). As desired, the cavity states are exchanged independently of the transmon state. In particular, when the transmon is in a superposition state, $|\psi_c\rangle = (|g\rangle + |e\rangle)/\sqrt{2}$, the cavity states are swapped despite taking a superposition of paths on the Bloch sphere to do so. Just as in the case of the cSWAP, delays *still* need to be inserted before and after the pulses. Even though the mode operators reach the same point on the operator Bloch sphere regardless of ancilla state, these two trajectories enclose a different geometric phase and therefore impart some control-state-dependent phase.

A useful advantage of increasing the number of echo pulses is that the two concatenated trajectories lie increasingly close together. This reduces the difference between the control-state-dependent phase rotations and therefore minimizes the length of delays required to disentangle the transmon and cavity states. In the large N limit, the two trajectories are identical, and enclose an identical solid angle.

5.4.2 Re-revisiting cSWAP

The flexibility of the choice of g_{bs} in the unconditional beamsplitter enables a faster alternative to the cSWAP. While the cSWAP demonstrated earlier in this chapter is already an order of magnitude faster than its prior iteration, the value of g_{bs} required is substantially less than the maximum value available to us. With a fast unconditional beamsplitter, it is possible to revisit the photonics-inspired cSWAP approach introduced in Ch. 2.4, consisting of a 50/50 beamsplitter, a parity-map and another 50/50 beamsplitter (see Fig. 5.8(a)). Fig. 5.8(b) shows the duration of this new cSWAP (including delays) as a function of g_{bs}, assuming that the delay times can be made short by introducing sufficient echo pulses during the uBS pulses. What it shows is that for $g_{\mathsf{bs}}/|\chi| > 1.37$, the photonics-inspired approach using uSWAP has a shorter total duration and is therefore likely to be higher-fidelity. In the regime where $g_{\mathsf{bs}} \gg |\chi|$, the duration of the sequence is rate-limited by $|\chi|$ and saturates at $\pi/|\chi|$ (the duration of the phase

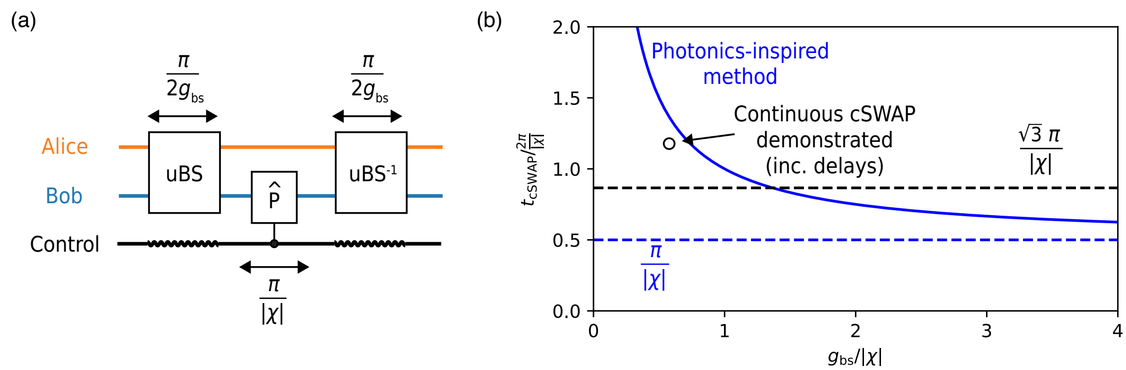

Figure 5.8: **Photonics-inspired cSWAP with using an unselective beamsplitter.** (a) Pulse sequence for the 'photonics-inspired' cSWAP, enabled by an unselective 50:50 beamsplitter (uBS). By applying pulses on the control qubit during the beamsplitter drive (indicated by the oscillating line), its unitary is made independent of the control state. (b) Comparison of the duration of this photonics-inspired cSWAP (with its flexibility in choice of g_{bs}) and the continuous cSWAP demonstrated in this chapter. The horizontal lines indicate the minimum possible duration of each sequence. The circle showing the cSWAP demonstrated in this chapter is above this line due to the addition of pre- and post-delays to correct for control-state-dependent phases. In principle, with sufficient echo pulses during the uBS, the phase correction required for the photonics-inspired approach can be made arbitrarily small. Figure modified with permission from Chapman *et al.* (2023).

shift element).

5.5 Beamsplitter-enabled joint parity measurement

To finish the chapter, I will show a final (but very important) operation that only becomes possible when $g_{\text{bs}} \approx |\chi|$ and which also relies on this operator Bloch sphere picture: a measurement of *joint* photon number parity in two oscillators $\hat{\text{JP}} = (-1)^{\hat{a}^\dagger\hat{a}+\hat{b}^\dagger\hat{b}}$. In the context of joint-Wigner tomography, we already saw that we could measure this operator by performing simultaneous measurements of the single-oscillator parity, $\hat{P}_A = (-1)^{\hat{a}^\dagger\hat{a}}$ and $\hat{P}_B = (-1)^{\hat{b}^\dagger\hat{b}}$, and multiplying the outcomes shot-by-shot. However, by collapsing the individual oscillator wavefunctions, this measurement destroys the entanglement between the two modes. This collapse is permissible in the context of tomography, where the post-measurement state is inconsequential, but not in the context of error correction. As we discovered in Ch. 2.1.2, an example of where a joint-

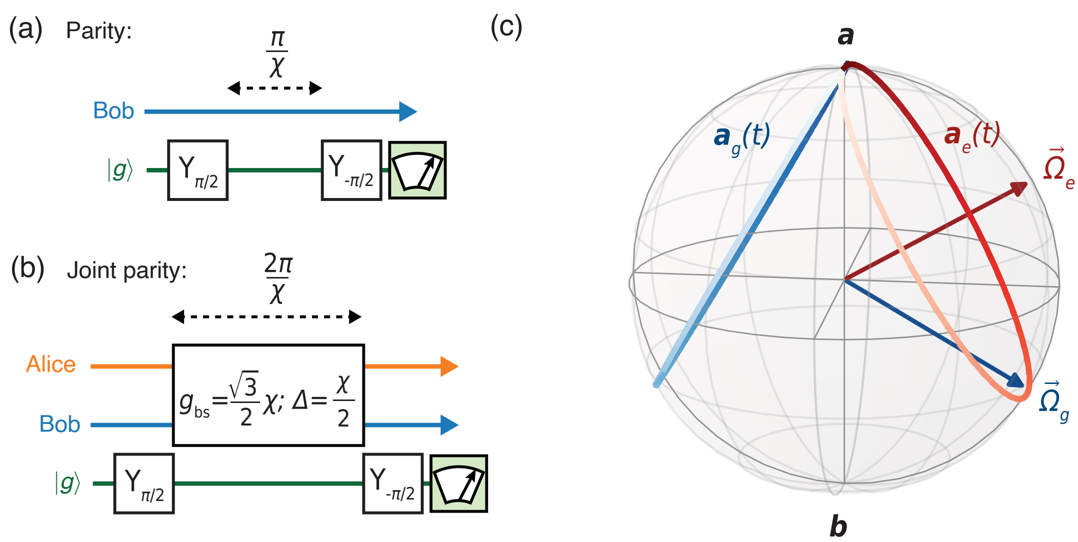

Figure 5.9: **Beamsplitter-enabled joint-parity measurement.** The application of a beamsplitter drive with a particular amplitude and detuning transforms (a) the usual parity measurement sequence into (b) a measurement of the joint photon number parity in two oscillators. (c) The trajectories of the Alice mode operator on the operator Bloch sphere, conditioned on the ancilla state, during the beamsplitter portion of the joint parity measurement sequence.

parity measurement can be extremely useful is in detecting erasure errors in a dual-rail code (something we will explore further in Chapter 7). In this case, it is vital to preserve multi-mode entanglement after the measurement.

One approach to engineering a *mid-circuit* (i.e. one that preserves entanglement) joint parity measurement is to capacitively couple a transmon ancilla to both oscillators (in a 'Y-mon' configuration) and apply a four-wave-mixing parametric drive to the transmon to match the dispersive shift to each mode, $\chi_a = \chi_b$ (Wang *et al.*, 2016). This approach to modulating χ via a microwave drive is the same used to match $\chi_{ge} = \chi_{gf}$ to correct for transmon relaxation errors in Chapter 2.

Instead, the joint parity (JP) measurement can be constructed using an ancilla transmon statically coupled to only *one* of the oscillator modes. By reducing the number of elements statically coupled to another, compared to in a Y-mon approach, one benefits from transmon errors not being able to simultaneously propagate to multiple cavities and from a minimization of crosstalk between these modes. Furthermore, by not having a transmon situated between the two cavities, it frees up space for us to use a dedicated SNAIL coupler to mediate high-fidelity beamsplitter interactions between these modes.

Activating a beamsplitter drive during the wait time of the usual single-oscillator parity measurement scheme transforms it into a joint parity measurement of two oscillators (Fig. 5.9). In a sense, the beamsplitter drive plays the role of the χ-matching drive. This joint parity scheme makes use of the ancilla-state-dependent phase space rotation acquired during a beamsplitter trajectory, which can alternatively be viewed as a joint-photon-number-dependent phase on the ancilla. To find constant beamsplitter parameters that (ideally) enable a joint-parity measurement, we can apply two constraints:

1. Oscillator states should return to where they began. In other words, the mode operators $\hat{a}(t)$ and $\hat{b}(t)$ must complete a closed trajectory on the Bloch sphere, independently of the ancilla state. This can be achieved by fixing $\Omega_g = \Omega_e = \Omega$ and evolving for $T_p = 2\pi/\Omega$. The constraint on Ω requires that $\Delta = \chi/2$.

2. The ancilla $|e\rangle$ state should acquire a phase per total photon, relative to the $|g\rangle$ state, of π. (This can equivalently be thought of as a phase space rotation on the cavity by π radians, conditioned on the ancilla $|e\rangle$ state.) Using Eq. 5.13 for the enclosed solid angle of a loop, we can calculate the phase difference per photon

$$|\phi_e - \phi_g| = \frac{\pi|\chi|}{\Omega} = \pi \tag{5.26}$$

$$\Omega = |\chi| \tag{5.27}$$

$$\rightarrow g_{\mathsf{bs}} = \frac{\sqrt{3}}{2}|\chi| \tag{5.28}$$

The resulting ideal mode operator trajectories can be seen in Fig. 5.9(c).

This sequence is very similar to the cSWAP sequence, and indeed the calibration (shown in Fig. 5.10) can be performed in a similar manner:

1. Given an initial guess for g_{bs} and the beamsplitter pulse frequency ω_{p}, initialize the cavities in $|0\rangle_A |1\rangle_B$ with the transmon in either $|g\rangle$ or $|e\rangle$, play the beamsplitter for a variable

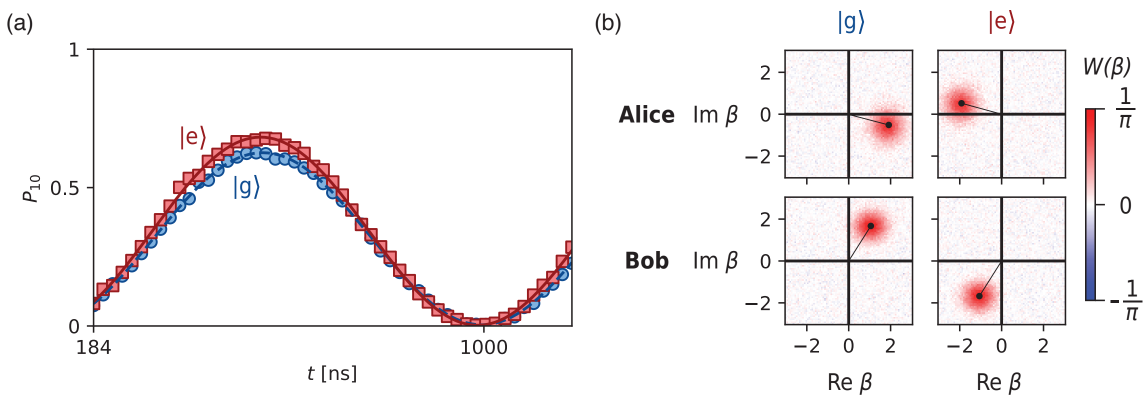

Figure 5.10: **Calibrating a joint-parity map.** (a) Probability of measuring the two-oscillator state $|1\rangle_A |0\rangle_B$ after time t of applying a beamsplitter drive with a detuning $\Delta \approx \chi/2$, having initialized the oscillators in $|0\rangle_A |1\rangle_B$. The traces show the single-photon oscillations conditioned on the transmon state $|g\rangle$ (blue) or $|e\rangle$ (red), which complete a full period at the same time $t \approx 2\pi/|\chi|$. (b) Local Wigner tomography of each oscillator after initializing a coherent state and performing a joint-parity map with the transmon in either $|g\rangle$ (left column) or $|e\rangle$ (right column). Black lines indicate the angle each state makes with the origin. Both Alice and Bob's coherent states are rotated by π conditioned on the control state. In Bob, this is finetuned by adjusting the length of a delay after the beamsplitter pulse. There remains a control-state-independent phase space rotation on each cavity that can be corrected in software.

time and then measure the population of each cavity. Adjust ω_p until the two oscillations proceed at the same rate $\Omega_g = \Omega_e$ and return to their initial state at the same time, T_p (Fig. 5.10(a)).

2. Prepare a coherent state in Alice (the cavity without the control qubit coupled to it), perform the beamsplitter pulse with the control in $|g\rangle$ and $|e\rangle$, and measure the angle the coherent state makes with the origin. Note that this technically does not require full 2D Wigner tomography, just the 1D line of points at a distance $|\alpha|$ from the origin. Ideally, the angle the final coherent state makes when the control is in $|e\rangle$ is different from the angle when the control is in $|g\rangle$ by π. If not, it implies that the geometric phase enclosed by the loop on the Bloch sphere (see Fig. 5.9) is incorrect. The size of this loop can be changed by adjusting g_bs and returning to step 1 (Fig. 5.10(b)).

3. Once the ancilla-state-dependent phase acquired by the coherent state in Alice is π, we can repeat the procedure for Bob. If the phase is not correct, we cannot adjust g_bs, since

this would undo the calibration on Alice's side. However, we can add a delay to apply an $|e\rangle$-state-dependent phase space rotation to Bob – this is not possible for Alice since the state both starts and ends in Alice for both control states. Adding this phase will make the full duration of the sequence equal to $2n\pi/|\chi|$ where $n \in \mathbb{Z}$. The need for this delay originates in the finite bandwidth of the ramp-up. An alternative approach that avoids adding delays is to optimize the beamsplitter pulse shape using OCT, as proposed in Tsunoda *et al.* (2023) (Fig. 5.10(b)).

This validates both of the criteria established when finding the operating parameters.

Finally, we can separately verify that the joint-parity-map, when sandwiched between two transmon $\pi/2$ pulses as in Fig. 5.9(b), enacts a joint-parity measurement on the two-cavity states, $|0,0\rangle$, $|0,1\rangle$, $|1,0\rangle$ and $|1,1\rangle$. We do so while sweeping the phase θ of the second transmon $\pi/2$-pulse in Fig. 5.11(a). The oscillations in the probability of measuring the transmon in $|e\rangle$ with θ (see Fig. 5.11(b)) are out-of-phase for states with opposite joint photon number parity. At $\theta = 0$ and $\theta = \pi$, the odd and even joint parity states are mapped onto one bit of information in the transmon, with the freedom to map either odd or even states to the less error-prone $|g\rangle$ state. This sequence can also act as a form of calibration. The zero phase offset of the $|0,0\rangle$ oscillation confirms that the transmon frequency is well-calibrated, while any left-right movement in the other curves can be used to indicate that the state-dependent phase shift is not yet perfectly calibrated. In our case, the reduction in the contrast can be attributed to state-preparation errors (made worse for higher Fock states), as well as ancilla errors during the readout and the joint-parity map itself.

Ancilla error-detectability

We started the chapter by noting that faster beamsplitters enable $g_{\text{bs}} \gg 1/T_2^{\text{ancilla}}$ and we have since seen that this enables higher-fidelity operations such as cSWAP where ancilla decoherence is a limiting factor. However, as we saw for single cavities, bosonic encodings are often only worthwhile if the ancilla does *not* degrade the oscillator coherence. The tools used to make the single-oscillator parity measurement fault-tolerant can be readily extended to the multi-oscillator

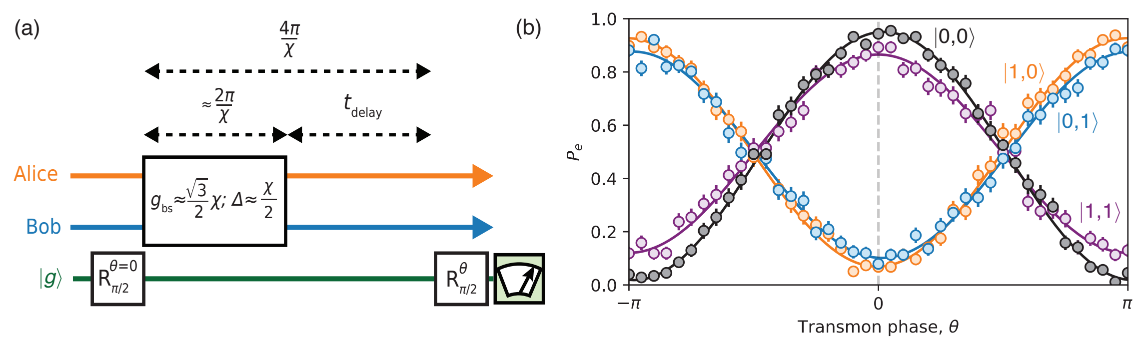

Figure 5.11: **Testing a joint-parity measurement on Fock states.** (a) Pulse sequence where Fock states $|n_A, n_B\rangle$ are initialized in the two oscillators and the joint-parity measurement sequence is performed, including a delay t_{delay} to remove the ancilla-state-dependent phase, with the phase θ of the final $\pi/2$ pulse varied. (b) The probability of measuring the transmon in $|e\rangle$, P_e, oscillates with θ but does so out of phase for states with opposite *joint* photon number. Operating at $\theta = 0$ or $\theta = \pi$ both perform a joint-parity measurement, with either even or odd states mapped to $|e\rangle$.

joint-parity measurement. While full error-correctability is possible, I will briefly describe how to achieve the less-stringent error-*detectability*, which will be more relevant when discussing the erasure check in Chapter 7, as outlined in Tsunoda *et al.* (2023).

First-order protection against ancilla decay is achieved by using a three-level $|g\rangle$-$|f\rangle$ ancilla and replacing the $Y^{ge}_{\pm\pi/2}$ pulses with $Y^{gf}_{\pm\pi/2}$. Ancilla relaxation during the sequence leaves the ancilla in $|e\rangle$ which serves as a flag state and is unaffected by the transmon pulses. Without χ-matching drives on the ancilla (where here χ-matching refers to ensuring that $\chi_{ge} = \chi_{gf}$) ancilla decay will result in a change of the axis about which the mode operators precess, at an unknown time during the sequence. This is an unrecoverable error that can in principle leave the mode operators anywhere on the operator Bloch sphere. However, measuring the ancilla in $|e\rangle$ indicates that that the error has occurred, providing an opportunity to either post-select out the shot, or to reset the oscillators to a known state and treat the event as an *erasure* (as will be discussed in more detail in Ch. 7).

First-order fault-tolerance to ancilla dephasing relies on the adhering to the ‘equal-latitude’ condition which here means that the latitude of the *ancilla* state on its Bloch sphere should be independent of the joint photon number in the oscillators. This condition is already met for the joint-parity measurement as the ancilla stays on the equator until the very end of the measurement. As a result, ancilla dephasing has no impact on the oscillator states. However, it can still result in an incorrect measurement outcome. If this measurement outcome informs subsequent operations on the oscillators, this may still be an issue, but since the measurement is QND, repeating it twice provides more certain information about the true value. Here, ancilla dephasing would result in two measurements that disagree (provided there is no change in the true joint photon number between the two measurements), a result that can be used to post-select data or erase qubits (as for ancilla relaxation).

These same principles have also been used to devise an ancilla-error-detected entangling $ZZ(\theta)$ gate for two dual-rail qubits (consisting of four cavities), constructed by inserting the joint-parity-map into the exponentiation gadget of Ch. 2.4.3 (Tsunoda *et al.*, 2023). By using a three-level ancilla and measuring the ancilla state at the of the sequence, both ancilla dephasing and ancilla relaxation can be detected and converted to erasure errors. This goes to show complete reliance of the cavity dual-rail encoding (Teoh *et al.*, 2023) on access to a fast, high-fidelity beamsplitter, forming the basis for the single-qubit gate, two-qubit gate and mid-circuit erasure check.

Chapter summary

In this chapter, we have seen how the mode operator Bloch sphere language introduced in Tsunoda *et al.* (2023) allows us to construct operations combining the dispersive and beamsplitter interactions in the $g_{\text{bs}} \approx |\chi|$ regime. These have included both higher-fidelity alternatives to previously implemented operations, in the case of cSWAP, as well as newly accessible operations in the case of uSWAP and the mid-circuit joint-parity measurement. While these operations use beamsplitter and ancilla pulses at different times, in the following chapter we will consider the case where they are applied simultaneously, further building out the toolbox of operations enabled in this regime.

Chapter 6

The joint-photon number-splitting regime

The previous chapter showed how the parity measurement for a single oscillator (see Sec. 2.2.1) can be extended to a *joint*-parity measurement for two oscillators by applying a beamsplitter drive during the pulse sequence. This begs the question of whether there is a deeper analogy here – can we generally extend techniques based on dispersive photon number splitting in a single oscillator to two (or more) oscillators, and can we make any statements about the conditions (e.g. on the frequency and amplitude of $g_{\text{bs}}(t)$) under which such an analogy will hold?

In this chapter I will show that indeed there is a deeper connection, revealed by the existence of a joint-photon number splitting regime when g_{bs} becomes comparable in magnitude to χ - a regime in which the spectrum of the nonlinear ancilla depends on the joint-photon number in the coupled oscillators. This regime will provide a complementary way of engineering high-fidelity multi-oscillator control using a combination of static dispersive coupling and tunable beamsplitter coupling (Fig. 6.1(a)). Whereas in the previous chapter, these two Hamiltonian terms were only activated one at a time, here we will consider the simultaneous application of drives on the transmon and the beamsplitter.

In the case of single-cavity control, we saw how there were two ways to obtain universal control of the cavity state: one based on alternating between SNAP drives on the transmon and displacements drives on the cavity, and another based on simultaneous drives on both modes whose forms were obtained via optimal control theory (OCT). The simultaneously-applied control sequences considered in this chapter are therefore similar in vein to this latter class of OCT pulses. However, unlike in the single-cavity case, we retain a notion of interpretability that can be leveraged when designing and troubleshooting pulse sequences. (The application of OCT to the combination of beamsplitter and dispersive interactions, without any attempt at preserving interpretability, may nonetheless be a useful way to obtain new multi-cavity operations.)

6.1 Observing joint-photon number-splitting

We can demonstrate the emergence of the joint-photon number-splitting regime by probing how the usual photon number-split spectrum (Gambetta *et al.*, 2006; Schuster *et al.*, 2007) is modified in the presence of an increasingly strong beamsplitter drive. Fig. 6.1(b) shows the pulse sequence used to obtain this spectrum, initializing the oscillators in states of well-defined total photon number $\hat{N}$, then performing a flat-top beamsplitter pulse simultaneous with a 14.8 μs long chopped-Gaussian transmon pulse (defined in Eq. 6.27), followed by a transmon readout. As in the joint-parity measurement, we initially set $\Delta = \chi/2$, a 'symmetric' choice that ensures that the beamsplitter drive is equally detuned from resonance when the ancilla is in $|g\rangle$ or in $|e\rangle$. Later in the chapter we will investigate what happens for more generic choices of detuning. Since N is a conserved quantity under the beamsplitter Hamiltonian, and the average oscillator lifetime $\bar{T}_1$ greatly exceeds the pulse duration T_{p}, we may separately investigate behavior for different values of N. The resulting spectra for different initial states with $N = 0$, 1 and 2 are shown in Fig. 6.1(c).

These spectra demonstrate the emergence of a joint-photon-number-splitting regime when g_{bs} becomes comparable to $|\chi|$. With the beamsplitter drive turned off ($g_{\mathsf{bs}} = 0$), linecuts show an ancilla spectrum with number-resolved peaks, each shifted by χ per photon in mode $\hat{b}$ but independent of the photon number in mode $\hat{a}$. This is the usual photon number split spectrum. However, at the largest values of g_{bs}[1], where $g_{\mathsf{bs}}/|\chi| = 1.92$, we observe a single dominant

[1]The operating flux point $\varphi_{\mathsf{ext}}/2\pi = 0.334$ was chosen to maximize the range of g_{bs}/χ over which measure-

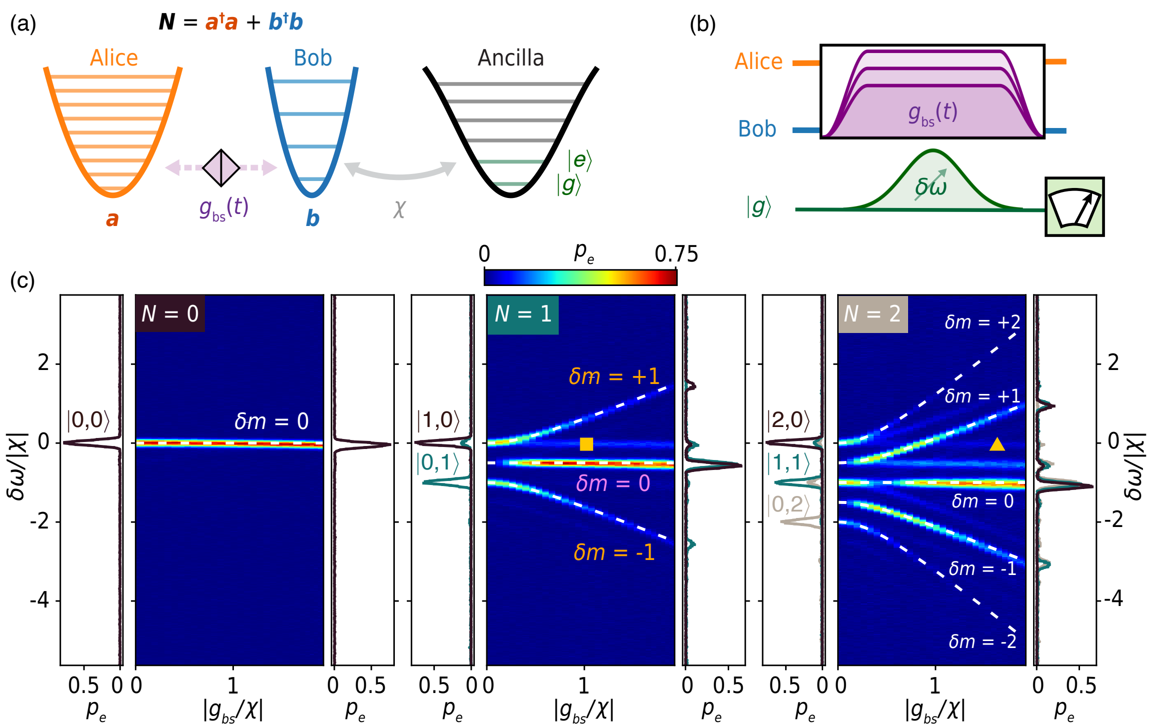

Figure 6.1: **Ancilla spectroscopy in the presence of a beamsplitter drive.** (a) System schematic consisting of two linear oscillators coupled via a tuneable beamsplitter coupling and a nonlinear ancilla coupled to *one* of the linear modes with a static dispersive coupling. The two linear modes contain a total of N photons. (b) Spectroscopy pulse sequence, in which the beamsplitter amplitude g_{bs} is varied. (c) Measured ancilla spectra in the presence of a variable g_{bs} for states with fixed total photon number N. The colorplots show data for initial states $|0,0\rangle$ ($N=0$), $\frac{|0,1\rangle+|1,0\rangle}{\sqrt{2}}$ ($N=1$) and $\frac{|0,2\rangle+\sqrt{2}|1,1\rangle+|2,0\rangle}{2}$ ($N=2$). Predicted transitions (dashed white lines) are labeled by the change in the angular momentum projection quantum number δm (described in Ch. 6.2). Linecuts to the left (right) of each colorplot show spectra for all two-oscillator Fock states in each N-photon manifold at the lowest (highest) value of $|g_{\text{bs}}/\chi|$. Figure

transition for each constant-N subspace, at a transmon frequency detuning of $\delta\omega = N\chi/2$. We can interpret these spectra as the dispersive shift being 'shared' between the two oscillators as photons are swapped rapidly between them. The linecuts at $g_{\text{bs}} \gtrsim |\chi|$ demonstrate the ability to excite the ancilla conditioned on the *joint* oscillator population $\hat{N}$, without learning the *individual* oscillator populations, despite the ancilla only being coupled to one of the modes.

ments could be taken. This required a flux point where the maximum g_{bs} achievable was large, but also where no unwanted transitions were activated as the normalized pump amplitude ξ was increased to reach this value.

The hardware developed in this thesis is key to enabling high-fidelity operations in the joint-photon number-splitting regime. The requirement that $g_{\mathsf{bs}} \approx |\chi|$ must go hand-in-hand with a pulse duration that is both long with respect to $2/|\chi|$ (to spectrally resolve number-split features) and short with respect to T_2^{ancilla} (to ensure high fidelity). Taken together, this requires that

$$g_{\mathsf{bs}} \approx |\chi| \gtrsim \frac{1}{T_{\mathrm{p}}} \gg \frac{1}{T_2^{\mathsf{ancilla}}} \gtrsim 2\pi \times 3 \text{ kHz} \tag{6.1}$$

While we can intuit the behavior at the two extremes of $|g_{\mathsf{bs}}/\chi|$, the transition between these regimes displays a much richer spectrum. Understanding this spectrum better will allow us to answer some practical questions: What is the origin of the 'satellite' peaks either side of the central peaks in the right linecuts? Are these problematic to performing joint operations? Can we engineer joint operations at lower values of g_{bs}? And what are the tradeoffs in doing so? In the following section, I lay out a simple model (similar to the one developed for alternating beamsplitter and ancilla pulses) to help answer these questions.

6.2 Interpreting the spectrum

As in Chapter 5, Schwinger's mapping between bosonic raising and lowering operators and angular momentum operators (Schwinger, 1952) can provide a useful framework for interpreting the observed physics, which scales well even as the combined photon number N becomes large. The combined beamsplitter and dispersive Hamiltonian (see Eq. 5.11) is now supplemented by an ancilla spectroscopy drive:

$$\hat{\mathcal{H}} = \hat{\mathcal{H}}_{\chi\mathsf{bs}} + \hat{\mathcal{H}}_{\mathsf{spec}}, \tag{6.2}$$

where

$$\frac{\hat{\mathcal{H}}_{\mathsf{spec}}}{\hbar} = \frac{\epsilon}{2}\left(e^{i\phi}\left|e\right\rangle\left\langle g\right| + e^{-i\phi}\left|g\right\rangle\left\langle e\right|\right) - \delta\omega\left|e\right\rangle\left\langle e\right|, \tag{6.3}$$

where ϵ, ϕ and $\delta\omega$ are the amplitude, phase and detuning of this drive. In the spectroscopy data shown, the strength of the spectroscopy drive $|\epsilon|$ is kept small relative to min (g_{bs}, χ) (by using a long, frequency-selective pulse) and so $\hat{\mathcal{H}}_{\text{spec}}$ can be treated as a perturbation. This leaves us with an unperturbed Hamiltonian that is the same ancilla-state-dependent Hamiltonian as in Eq. 5.7. We can write this out explicitly, conditioned on each ancilla state, in terms of the Schwinger angular momentum operators:

$$\hat{\mathcal{H}}^{(g)} = \frac{-\hat{N}\Delta}{2} + \Omega_g \left(\hat{n}_g \cdot \vec{\hat{L}}\right) \tag{6.4}$$

$$\hat{\mathcal{H}}^{(e)} = \frac{-\hat{N}\left(\Delta - \chi\right)}{2} + \Omega_e \left(\hat{n}_e \cdot \vec{\hat{L}}\right) \tag{6.5}$$

Previously, we discovered that looking at the evolution of *operators* in the Heisenberg picture was a convenient way of compactly describing the dynamics of bosonic states. However, since we are after the frequencies and matrix elements of transitions seen in our spectra, in this case it is more practical to consider the *eigenstates* in the Schrödinger picture.

6.2.1 Eigenstates and eigenenergies

In the language of angular momenta, the energy eigenstates of the Hamiltonians in Eq. 6.4 and 6.5 are simultaneous eigenstates of two commuting observables: the angular momentum squared, $\hat{L}^2$, with eigenvalues $l(l+1) \equiv \frac{N}{2}\left(\frac{N}{2}+1\right)$, and the angular momentum projection along the $\hat{n}$ axis, $\hat{n} \cdot \vec{\hat{L}}$, with eigenvalues $m = -l, -(l-1), \ldots, +l \equiv -\frac{N}{2}, -\left(\frac{N}{2}-1\right), \ldots, +\frac{N}{2}$. The unit vectors $\hat{n}_g$ and $\hat{n}_e$ (determined by the choice of g_{bs}, φ and Δ) therefore determine the ancilla-state-dependent quantization axes along which the eigenstates are aligned. The eigenenergies of these states (labelled by the total photon number, ancilla state and projection along the quantization axis) can be read off as

$$E^{(g)}_{N,m_g} = -\frac{N\Delta}{2} + m_g\sqrt{g_{\text{bs}}^2 + \Delta^2}, \tag{6.6}$$

$$E^{(e)}_{N,m_e} = -\frac{N\left(\Delta - \chi\right)}{2} + m_e\sqrt{g_{\text{bs}}^2 + \left(\Delta - \chi\right)^2}. \tag{6.7}$$

Since $\hat{N}$ is a conserved quantity, transitions cannot change its value. If we focus our attention on the $N = 1$ subspace and apply a ‘symmetric’ detuning $\Delta = \chi/2$, these expressions simplify to

$$E^{(g)} = -\frac{\chi}{4} \pm \frac{|\chi|}{4}\sqrt{1 + \left(\frac{2g_{\mathsf{bs}}}{\chi}\right)^2} \tag{6.8}$$

$$E^{(e)} = +\frac{\chi}{4} \pm \frac{|\chi|}{4}\sqrt{1 + \left(\frac{2g_{\mathsf{bs}}}{\chi}\right)^2}, \tag{6.9}$$

which are plotted in Fig. 6.2(a). The eigenenergies display a low-field splitting by $\chi/2$ and a high-field splitting by g_{bs}, analogous to a Zeeman effect on a magnetic spin-$1/2$ when a strong transverse magnetic field is applied. Within each constant-N subspace, we can obtain the allowed transition frequencies from the differences between the eigenenergies when the ancilla is in $|g\rangle$ and when then ancilla is in $|e\rangle$:

$$\omega_{N,m_g \to m_e} = \frac{N\chi}{2} + m_e\sqrt{g_{\mathsf{bs}}^2 + (\Delta - \chi)^2} - m_g\sqrt{g_{\mathsf{bs}}^2 + \Delta^2}. \tag{6.10}$$

For a general choice of Δ this leads to $(N+1)^2$ unique transition frequencies in the spectra. However, our intuition had been that a 'symmetric' detuning $\Delta = \chi/2$ would mark a special point and indeed it does. At this choice of the beamsplitter detuning, $\Omega_g = \Omega_e \equiv \Omega$ and the expression for the transition frequencies simplifies considerably:

$$\omega_{N,\delta m} = \frac{N\chi}{2} + \delta m \Omega = \frac{N\chi}{2} + \delta m\sqrt{g_{\mathsf{bs}}^2 + \left(\frac{\chi}{2}\right)^2}, \tag{6.11}$$

where there are now only $2N+1$ transition frequencies, indexed by $\delta m \equiv m_e - m_g = -N, \ldots, N$.

These predicted transition frequencies, plotted as white dashed lines in Fig. 6.1(c), agree very well with our measured spectroscopy data (in the regime that $|\epsilon| \ll \min(g_{\mathsf{bs}}, \chi)$). The smaller peaks either side of the central peak in the joint-photon number-splitting regime are transitions for which $\delta m \neq 0$. We can also verify the expressions for the more numerous non-degenerate transition frequencies when we have a non-symmetric Δ, as shown in Fig. 6.3.

We can try to describe the off-central $\delta m \neq 0$ transitions in the more familiar language of quantum optics, where we can say that the normal modes of the two coupled oscillators are

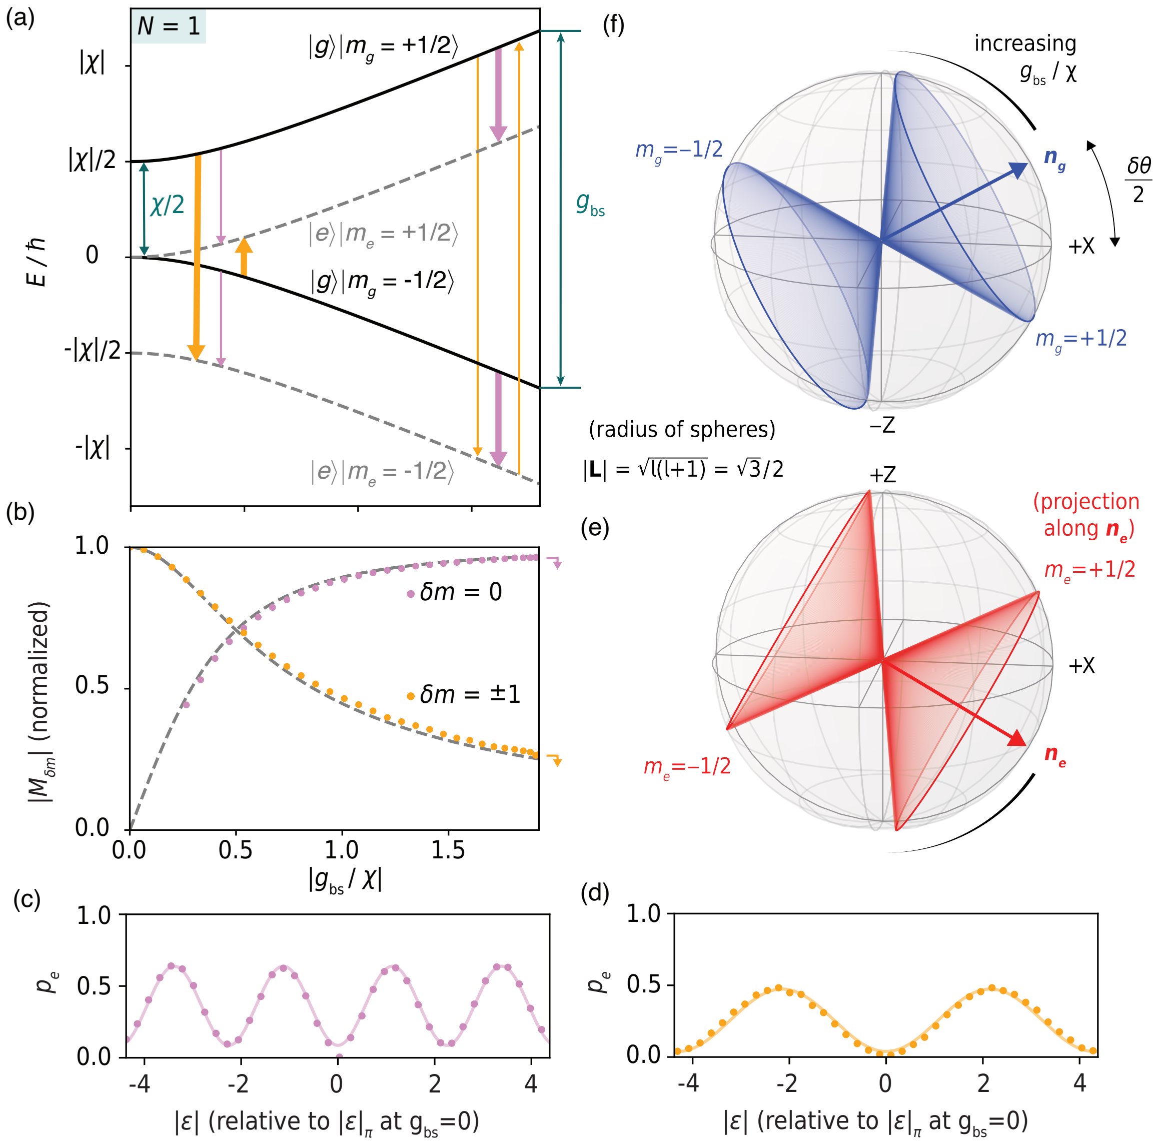

Figure 6.2: **Energy levels and transition matrix elements in $N = 1$ manifold**. (a) Energy level diagram shows two oscillator states (labeled by their eigenvalue m) for each of the two ancilla states, $|g\rangle$ (solid lines) and $|e\rangle$ (dashed lines). At low $|g_{\mathsf{bs}}/\chi|$, the energy levels are split by χ and transitions with $\delta m = \pm 1$ (orange) are strongest. At high fields, the levels are split by g_{bs} and the transitions with $\delta m = 0$ (pink) are strongest. (b) Measured transition matrix elements for transitions with $\delta m = -1$ and $\delta m = 0$, relative to the transition matrix element for $\delta m = -1$ at $g_{\mathsf{bs}} = 0$. Expected value from theory is shown as dashed grey line and are the same for $\delta m = \pm 1$. (c-d) Example power-Rabi oscillations used to extract transition matrix elements. (e-f) Geometric picture of angular momentum description. Each cone represents the state with projection m along the quantization axis $\hat{n}$ but undetermined projection along the other two axes. The sphere has radius $\sqrt{l(l+1)} = \sqrt{3}/2$. As $g_{\mathsf{bs}}/|\chi|$ is increased, $\hat{n}_g$ and $\hat{n}_e$ become more aligned, with smaller $\delta\theta$. Figure modified from de Graaf *et al.* (2025).

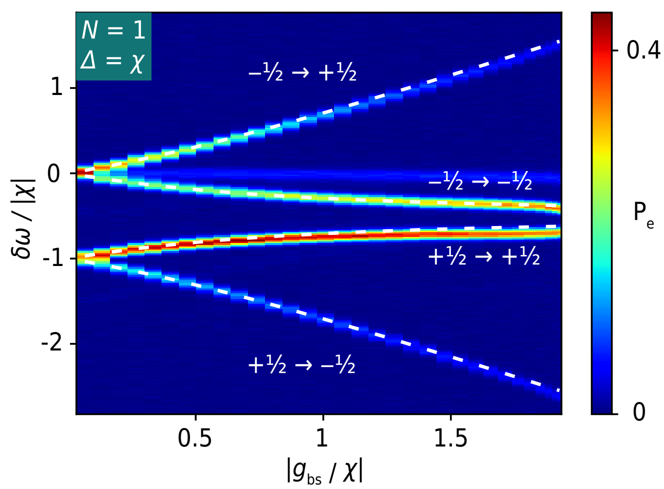

Figure 6.3: **Spectroscopy for non-symmetric detuning** $\Delta \neq \chi/2$. Measured transition frequencies for combined photon number in both cavities $N = 1$ and a non-symmetric beamsplitter detuning $\Delta = \chi$. Method used is the same as for Fig. 6.1. Dashed lines show predicted transition frequencies from Eq. 6.10. Figure modified from de Graaf *et al.* (2025).

different depending on the state of the ancilla. In the large $g_{\mathsf{bs}}/|\chi|$ limit, these normal modes are approximately the symmetric and antisymmetric combinations of the uncoupled cavity modes. However, since the approximately-symmetric mode when the ancilla is in $|g\rangle$ is not completely orthogonal to the approximately antisymmetric mode when the ancilla is in $|e\rangle$ (and vice-versa), exciting the ancilla can also lead to a change in the oscillator state. However, in order to conserve energy when it does so, the excitation drive needs to make up the frequency difference between photons living in the symmetric and antisymmetric modes, with a detuning which increases with increasing coupling strength g_{bs}.

6.2.2 Transition matrix elements

While these expressions accurately capture the frequencies of these transitions, it does not capture their relative prominence. For example, why is it that only the $\delta m = 0$ appear bright in the joint-photon-number-splitting regime? This requires us to know how the matrix elements of the these transitions vary with g_{bs}. This is also crucial to being able to construct high-fidelity operations with simultaneous beamsplitter and ancilla drives.

The transition matrix elements under the action of the spectroscopy Hamiltonian $\hat{\mathcal{H}}_{\mathsf{spec}}$ in

Eq. 6.3 are given by

$$M_{N,m_g\to m_e} = \langle N, m_g; g| \langle g| \left(\frac{\epsilon e^{i\phi}}{2} |g\rangle \langle e| + \frac{\epsilon e^{-i\phi}}{2} |e\rangle \langle g| \right) |e\rangle |N, m_e; e\rangle \tag{6.12}$$

$$= \frac{\epsilon e^{i\phi}}{2} \langle N, m_g; g|N, m_e; e\rangle , \tag{6.13}$$

and are therefore proportional to the state overlap between the initial and final states. The overlap between two angular momentum eigenstates with angular momentum quantum number $l = N/2$, and projections m_g and m_e along two different axes separated by an angle $\delta\theta$, is given by the Wigner (small) d-matrix (Wigner, 1931):

$$\langle N, m_g; g|N, m_e; e\rangle = d^{N/2}_{m_g,m_e}(\delta\theta). \tag{6.14}$$

For this system, the angle between $\hat{n}_g$ and $\hat{n}_e$ is given by

$$\delta\theta = \arctan\left(\frac{\Delta}{g_{\mathsf{bs}}}\right) - \arctan\left(\frac{\Delta - \chi}{g_{\mathsf{bs}}}\right). \tag{6.15}$$

For the case shown in Fig. 6.2(a), where $N = 1$ and $\Delta = \chi/2$, this angle difference is $\delta\theta = 2\arctan(\chi/2g_{\mathsf{bs}})$. This therefore yields a transition matrix element for the degenerate central $\delta m = 0$ transitions of

$$\begin{aligned} |M_{\delta m=0}| &= \frac{\epsilon}{2}\left|d^{\frac{1}{2}}_{\frac{1}{2},\frac{1}{2}}(\delta\theta)\right| = \frac{\epsilon}{2}\left|d^{\frac{1}{2}}_{-\frac{1}{2},-\frac{1}{2}}(\delta\theta)\right| \\ &= \frac{\epsilon}{2}\cos\left(\frac{\delta\theta}{2}\right) \\ &= \frac{\epsilon}{2}\frac{g_{\mathsf{bs}}}{\sqrt{g_{\mathsf{bs}}^2 + \left(\frac{\chi}{2}\right)^2}}, \end{aligned} \tag{6.16}$$

and for the off-central $\delta m = \pm 1$ transitions of

$$
\begin{aligned}
|M_{\delta m=\pm 1}| &= \frac{\epsilon}{2}\left|d^{\frac{1}{2}}_{\frac{1}{2},-\frac{1}{2}}(\delta\theta)\right| = \frac{\epsilon}{2}\left|d^{\frac{1}{2}}_{-\frac{1}{2},\frac{1}{2}}(\delta\theta)\right| \\
&= \frac{\epsilon}{2}\sin\left(\frac{\delta\theta}{2}\right) \\
&= \frac{\epsilon}{2}\frac{\left|\frac{\chi}{2}\right|}{\sqrt{g_{\mathsf{bs}}^2 + \left(\frac{\chi}{2}\right)^2}}.
\end{aligned}
\tag{6.17}
$$

We can compare our predicted expressions to experiment by measuring the ancilla drive amplitude ϵ_π required to perform an ancilla π-pulse, which is inversely proportional to the transition matrix element, for different transitions, over a range of $g_{\mathsf{bs}}/|\chi|$ values. The $\delta m = -1$ transition is chosen over the $\delta m = +1$ transition, since it is further from other transitions, allowing for a cleaner measurement, but should have the same value.

The experimental sequence consists of a power Rabi experiment, in which we initialize the system in $|g\rangle\left(|0,1\rangle + i\,|1,0\rangle\right)/\sqrt{2}$, apply a $T_{\mathsf{p}} = 14.8\ \mu$s Gaussian transmon pulse with variable amplitude ϵ, and measuring the transmon state. Example data is shown in Fig. 6.2(c, d). Fitting these amplitude-oscillations to a sinusoid lets us obtain ϵ_π. To obviate the need for an absolute calibration of the delivered drive amplitude, all measured transition matrix elements are normalized to the value at $g_{\mathsf{bs}} = 0$ for $\delta m = -1$. This measurement requires the transmon drive frequency to be well-calibrated as the matrix element will be underestimated if the drive is off-resonance. However, the measurement is not sensitive to SPAM errors, which only affect the amplitude (not the periodicity) of oscillations.

The extracted transition matrix elements for $N = 1$ are shown in Fig. 6.2(b), agreeing well with the predicted values. The reader will notice that the data for the $\delta m = 0$ transition does not continue all the way to $g_{\mathsf{bs}} = 0$. As the transition matrix element decreases towards zero, an increasingly large ϵ is required to resolve periodicity of the oscillations or, in other words, the product $\epsilon_\pi T_{\mathsf{p}}$ gets larger and larger. The ability to increase T_{p} is limited by decoherence on the ancilla as it approaches $T_2^{(\mathsf{ancilla})}$, and so ϵ must be increased. However, at the point that ϵ approaches χ, the ancilla drive can no longer be treated as a perturbation to the Hamiltonian and the spectrum is modified, with the central transition branching into two separate peaks.

In this example, we can see how at $g_{\mathsf{bs}}/|\chi| = 0$ only the $\delta m = \pm 1$ transitions are allowed and at $g_{\mathsf{bs}} \gg |\chi|$ the opposite is the case, and only the $\delta m = 0$ transitions are allowed. In this latter regime, the strong beamsplitter drive polarizes the direction of the ancilla-state-dependent "angular momentum" ($\hat{n}_g$ and $\hat{n}_e$) along the x-axis, reducing the angle, $\delta\theta$, between the respective quantization axes (see Fig. 6.2(e-f)). This ensures that eigenstates associated with the ancilla in $|g\rangle$ or in $|e\rangle$, with the same value of m (or equivalently the dressed eigenmodes), have ever-increasing overlap. At large g_{bs} we will therefore *only* see the $\delta m = 0$ peaks in the spectrum, which are separated by $\chi/2$ per total photon in the combined oscillators, providing mathematical backing for our earlier intuition that the dispersive shift gets 'shared' between the two oscillators in the presence of a strong beamsplitter drive.

6.3 Considerations for joint-photon number measurements

The joint-photon number-splitting spectrum shows that we can excite the ancilla conditioned on the joint photon number in the two oscillators with an ancilla pulse whose duration $T_{\mathsf{p}} \gtrsim 2\pi/|\chi|$. However, before jumping to a practical realization, there are a few theoretical considerations we have to bear in mind:

- Is the transition rate the same for all states within each photon-number manifold?
- What g_{bs} is required to avoid unwanted transitions?
- Is the two-oscillator state preserved after the measurement?

6.3.1 Transition rates for $N > 1$

We can evaluate the first question by considering the Wigner d-matrix elements for the $\delta m = 0$ (joint-photon-number-split) transitions. The d-matrix is symmetric under negation of both indices,

$$|d^{N/2}_{m_g,m_e}(\theta)| = |d^{N/2}_{-m_g,-m_e}(\theta)|, \tag{6.18}$$

and under exchange of indices

$$|d^{N/2}_{m_g,m_e}(\theta)| = |d^{N/2}_{m_e,m_g}(\theta)|, \tag{6.19}$$

However, there is in general no constraint that matrix elements with the same difference between indices, δm, have the same value. For $\delta m = 0$, the closed form expression is

$$d^{N/2}_{m,m}(\delta\theta) = \cos^{2m}\left(\frac{\delta\theta}{2}\right) P^{(0,2m)}_{N/2-m}\left(\cos\delta\theta\right), \tag{6.20}$$

where $P^{(a,b)}_n(x)$ is the Jacobi polynomial. For $N \leq 3$ these are:

$$\begin{aligned}
&d^0_{0,0}(\delta\theta) = 1, && \\
&d^{\frac{1}{2}}_{\frac{1}{2},\frac{1}{2}}(\delta\theta) = \cos\left(\frac{\delta\theta}{2}\right), && \\
&d^1_{0,0}(\delta\theta) = \cos\left(\delta\theta\right), && d^1_{1,1}(\delta\theta) = \cos^2\left(\frac{\delta\theta}{2}\right), \\
&d^{\frac{3}{2}}_{\frac{1}{2},\frac{1}{2}}(\delta\theta) = \cos\left(\frac{\delta\theta}{2}\right)\frac{3\cos(\delta\theta)-1}{2}, && d^{\frac{3}{2}}_{\frac{3}{2},\frac{3}{2}} = \cos^3\left(\frac{\delta\theta}{2}\right),
\end{aligned} \tag{6.21}$$

which are plotted in Fig. 6.4.

These expressions show that for $N > 1$, the ancilla Rabi rate is different for different oscillator states with the same total photon number. This makes it significantly more challenging to perform joint-photon number selective measurements for $N > 1$ - a simple pulse on the ancilla will not perform a π-pulse for all states in the manifold. One possible solution is to increase g_{bs} further since in the limit that $g_{\mathsf{bs}} \gg |\chi|$, the expressions converge to 1 for all m. However, for the g_{bs} we currently have available (see vertical line in Fig. 6.4), the error due to this difference is still substantial for $N > 1$. A better solution for $N > 1$ may be to eschew interpretability and turn to optimal control pulses. For the remainder of this chapter however, I will consider the cases $N \leq 1$.

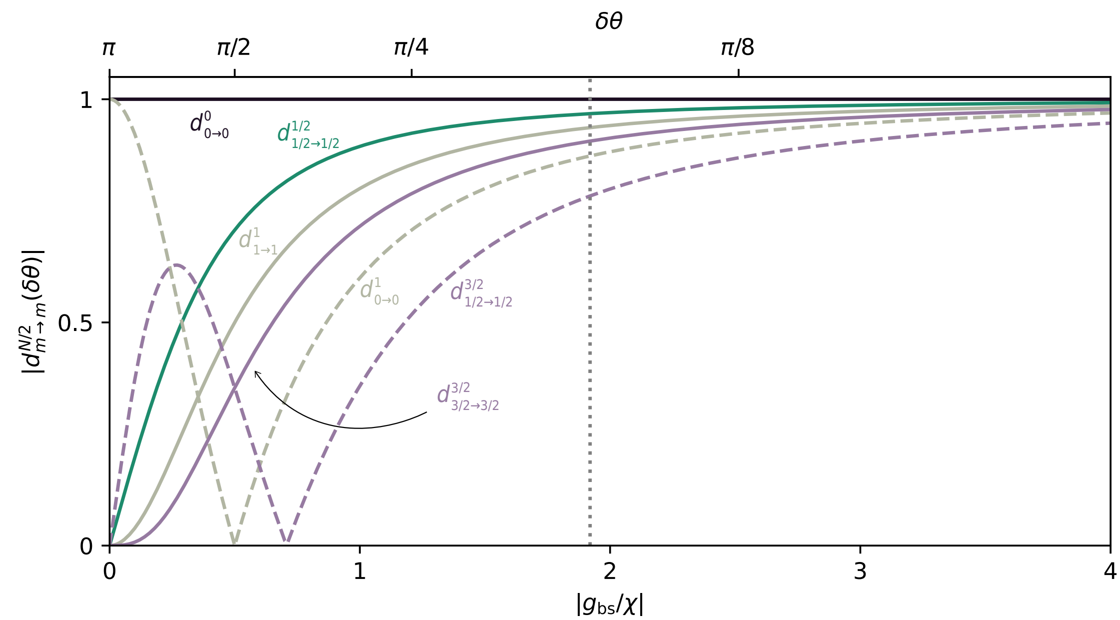

Figure 6.4: **Theoretical matrix elements for (joint-photon-number-split) $\delta m = 0$ transitions**. Wigner d-matrix entries (proportional to transition matrix elements) for transitions with total photon number, $N = 0$ (black), $N = 1$ (green), $N = 2$ (gold) and $N = 3$ (purple). Transitions are shown for which $m_g \to m_e$ is equal to $N/2 \to N/2$ (solid lines) or $N/2-1 \to N/2-1$ (dashed lines), indicating a deviation between the two. Vertical dotted grey line shows the maximum operating available in this device ($|g_{\text{bs}}/\chi| = 1.92$). Note that changing the sign of both m_g and m_e does not change the magnitude of the matrix element.

6.3.2 Ensuring photon number selectivity

Let us consider a pulse sequence where the beamsplitter drive is ramped instantaneously, such that $g_{\text{bs}}/|\chi|$ is constant throughout the pulse. This allows us to use the spectra in Fig. 6.5 as a guide. A long, highly-frequency selective pulse (like the one used to obtain the data) at $\delta\omega = N_{\text{meas}}\chi/2$ will excite the ancilla if the two oscillators have total photon number $N = N_{\text{meas}}$. Depending on the highest total photon number we might expect to see in the two oscillators, N_{max}, we will need to choose g_{bs} so that the $N = N_{\text{meas}}$ transition is sufficiently far detuned from other transitions with $N \leq N_{\text{max}}$. For example, if I want to ask "is $N = 0$?" and my possible input states have $N \leq 2$, then g_{bs} should be chosen to avoid the region where the $(N = 2, \delta m = +1)$ transition intersects the desired $(N = 0, \delta m = 0)$ transition. More precisely,

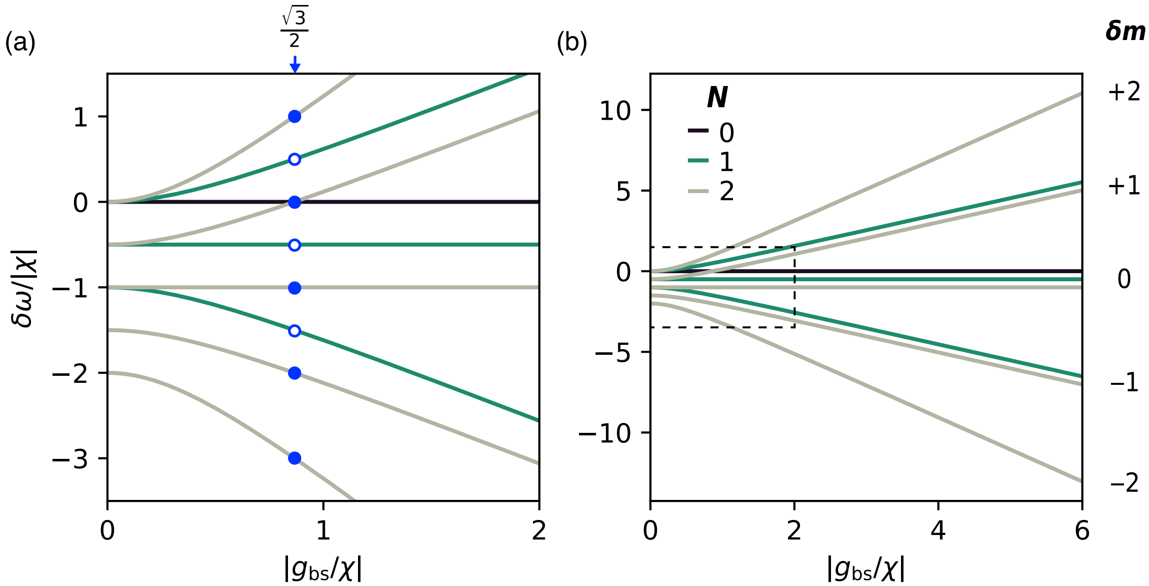

Figure 6.5: **Predicted transition frequencies up to** $N = 2$.(a) Predicted transition frequencies for $N = 0$ (black), $N = 1$ (green) and $N = 2$ (gold). At $|g_{\mathsf{bs}}/\chi| = \sqrt{3}/2$, all even (odd) N transitions lie at even (odd) multiples of $|\chi|/2$, as indicated by the filled (open) blue circles. (b) A wider-scale view of the same plot, showing the grouping of transition frequencies by δm at large $|g_{\mathsf{bs}}/\chi|$.

we would like these unwanted transitions to be at least $|\chi|/2$ away – further away than the nearest expected $\delta m = 0$ transition. For general N, the value of g_{bs} to exceed to avoid any further overlaps with unwanted transitions can be obtained from Eq. 6.11:

$$\omega_{N_{\mathsf{max}},+1} > \omega_{N_{\mathsf{meas}},0} - \frac{\chi}{2} \tag{6.22}$$

$$\frac{N_{\mathsf{max}}\chi}{2} + \sqrt{g_{\mathsf{bs}}^2 + \left(\frac{\chi}{2}\right)^2} > \frac{N_{\mathsf{meas}}\chi}{2} - \frac{\chi}{2} \tag{6.23}$$

$$\rightarrow g_{\mathsf{bs}} > \frac{|\chi|}{2}\sqrt{(N_{\mathsf{max}} - N_{\mathsf{meas}} + 1)^2 - 1} \tag{6.24}$$

6.3.3 Preserving oscillator states

A key feature of a non-destructive measurement is that eigenstates of the measured observable are preserved. For example, when I non-destructively measure "is $N = 0$?", then states with $N = 0$ and states with $N \neq 0$ should be preserved after the measurement. In the context of error-correction, this allows me to check for errors without perturbing the logical information. Therefore, beyond ensuring that the ancilla is excited conditioned on N, I also need to choose an ancilla pulse shape that ensures states return to where they began. This is particularly an issue because the beamsplitter Hamiltonian continuously exchanges photons between the oscillators.

When the ancilla remains in $|g\rangle$ throughout the measurement (i.e., when $N \neq N_{\text{meas}}$ and the ancilla pulse is long, $T_{\text{p}} \gg 2/|\chi|$) the oscillator state evolution can be straightforwardly computed using the operator Bloch sphere picture, undergoing oscillations at a rate $\Omega = \sqrt{g_{\text{bs}}^2 + (\chi/2)^2}$. This means that oscillator states are returned to their initial states when we satisfy

$$\Omega T_{\text{p}} = \sqrt{g_{\text{bs}}^2 + \left(\frac{\chi}{2}\right)^2} T_{\text{p}} = 2 n_{\text{bs}} \pi \text{ for } n_{\text{bs}} \in \mathbb{Z}^+. \tag{6.25}$$

However, for an oscillator state with $N = N_{\text{meas}}$, the ancilla continuously transitions from $|g\rangle$ to $|e\rangle$ during the pulse, and so we cannot compute two individual ancilla-state-dependent trajectories.

One case where the situation simplifies considerably however, is when $N_{\text{meas}} = 0$, since the state $|0,0\rangle$ is invariant under the Hamiltonian. As a means of measuring the joint photon number in two oscillators, the joint-photon-number-splitting approach is therefore best suited to measurements that ask "is $N = 0$?". This is exactly the question we will seek to answer in Chapter 7 when constructing an erasure check for dual-rail qubits. In this case, Eq. 6.25 is valid (in the limit $T_{\text{p}} \gg 2/|\chi|$) and the amplitude of the ancilla π-pulse can be constrained via

$$\int_0^{T_{\text{p}}} \epsilon(t) dt = \pi. \tag{6.26}$$

As such, for a given pulse shape and duration, there is a single ϵ and a discrete range of g_{bs} values that satisfy a joint photon number selective measurement. In practice, the finite duration of the ancilla pulse will slightly modify the values of g_{bs} required, whose experimental calibration will be discussed in detail in Chapter 7.

6.4 State trajectories for different pulse shapes

With the considerations of the previous section in mind, we can take a look at how both the oscillator and ancilla state trajectories evolve for the specific example of a measurement used to distinguish between $N = 0$ and $N = 1$ states, as will be relevant for the erasure check. We will do so for both a longer chopped Gaussian ancilla pulse shape, as well as a shorter square ancilla pulse shape.

6.4.1 Gaussian pulses

For a chopped Gaussian pulse of the form

$$\epsilon(t) = A\left[\exp\left(-\frac{(t - n_{\mathsf{chop}}\sigma)^2}{2\sigma^2}\right) - \exp\left(-\frac{n_{\mathsf{chop}}^2}{2}\right)\right] \text{ for } 0 \le t \le T_{\mathsf{p}}, \tag{6.27}$$

where σ is the RMS width of the pulse and the pulse duration $T_{\mathsf{p}} = 2n_{\mathsf{chop}}\sigma$, we can say as a rough guide that we would like the $N = 1$ transition at $\delta\omega = \chi/2$ to be at least 5 pulse linewidths detuned from the drive on the $N = 0$ transition. Given a time-bandwidth product $\sigma \times \sigma_f = 1/2$, where σ_f is the frequency bandwidth of the pulse, this corresponds to enforcing

$$5\sigma_f \le \frac{|\chi|}{2} \tag{6.28}$$

$$\rightarrow \sigma \ge \frac{5}{|\chi|}, \tag{6.29}$$

or equivalently, for $n_{\mathsf{chop}} = 2$, $T_{\mathsf{p}} \ge \frac{20}{|\chi|} = \frac{10}{\pi} \times \frac{2\pi}{|\chi|}$.

If we saturate this bound, then Eq. 6.25 becomes

$$\left|\frac{g_{\mathsf{bs}}}{\chi}\right| = \sqrt{\left(\frac{n_{\mathsf{bs}}\pi}{10}\right)^2 - \left(\frac{1}{2}\right)^2}. \tag{6.30}$$

If we separately enforce the constraint on the beamsplitter amplitude from Eq. 6.24, which for $N_{\mathsf{max}} = 1$ and $N = 0$ says that $|g_{\mathsf{bs}}/\chi| \ge \sqrt{3}/2$, we require $n_{\mathsf{bs}} \ge 4$. In this particular case, $n_{\mathsf{bs}} = 4$ yields $g_{\mathsf{bs}} \approx 1.15|\chi|$.

The evolution of the ancilla and oscillator states for this choice of Gaussian pulse are shown for Fock states with $N \leq N_{\text{max}} = 1$ in Fig. 6.6. These curves indicate both that the joint photon number information is mapped onto the ancilla state, and that the oscillator states return to their original states (perhaps up to a correctable phase).

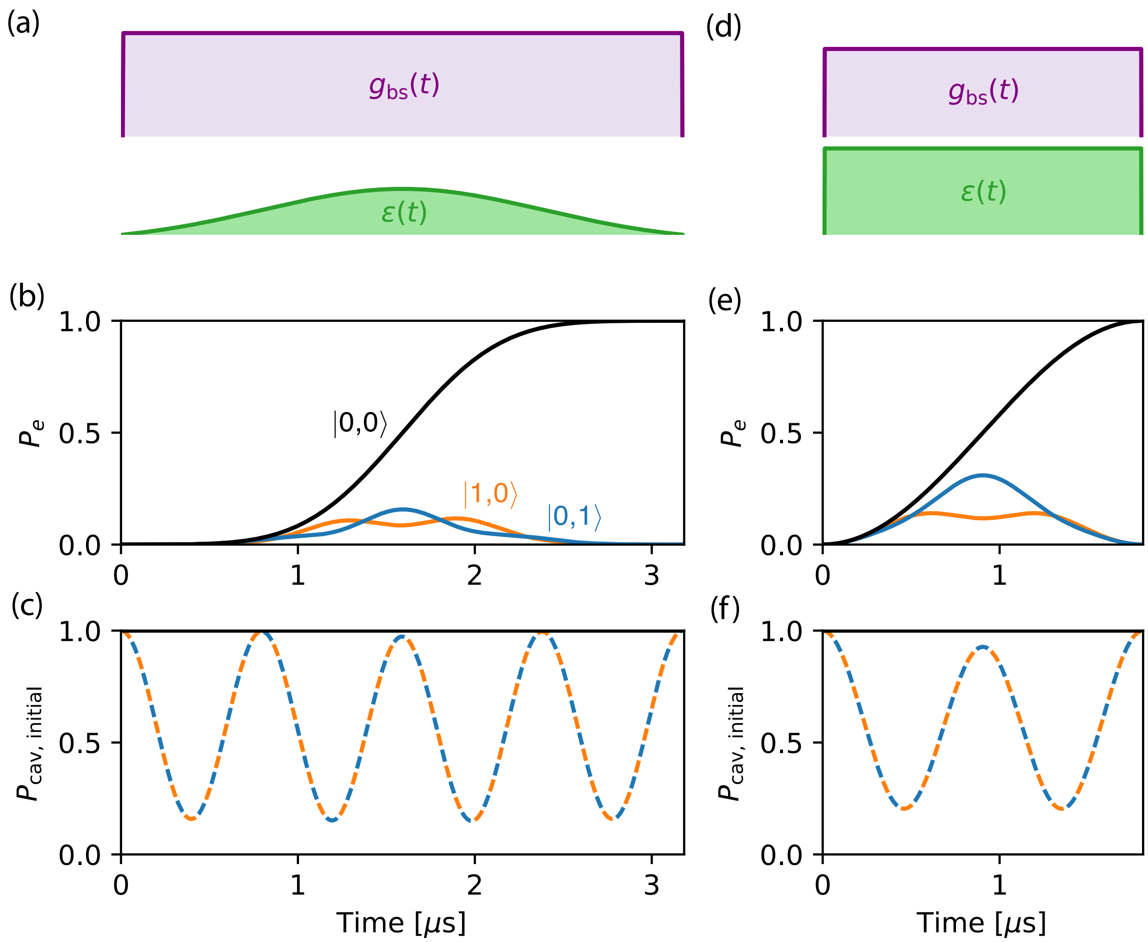

Figure 6.6: **Simulated state trajectories for $N = 0$ selective measurement**. (a) Illustration of beamsplitter and ancilla pulse amplitudes for a chopped Gaussian ancilla pulse shape. (b) Simulated ancilla excited state probability P_e during the pulse sequence for cavity input states $|n_A, n_B\rangle$, showing that the sequence performs a joint photon number measurement. (c) Probability of oscillators being in the same state in which they began during the pulse sequence, showing that they return to their initial states at the end. (d-f) Same plots but for a shorter square ancilla pulse shape, showing greater deviation from the ancilla $|g\rangle$ state for states with $N = 1$.

6.4.2 Square pulses

So far we have considered relatively long chopped Gaussian pulses that have a narrow linewidth around the $N = 0$ transition. However, as we only need to distinguish between $N = 0$ and $N = 1$, a shorter square pulse shape of the form

$$\epsilon(t) = \epsilon \text{ for } 0 \leq t \leq T_p \tag{6.31}$$

is also possible. The constraint on the pulse duration T_p is that in the time the ancilla performs a π-pulse when $N = 0$, it must perform n_ϵ 2π-pulses when $N = 1$, detuned by $\chi/2$,

$$\sqrt{\epsilon^2 \cos^2\left(\frac{\delta\theta}{2}\right) + \left(\frac{\chi}{2}\right)^2} = 2n_\epsilon\epsilon \tag{6.32}$$

$$\sqrt{\left(\frac{\pi}{T_p}\right)^2 \cos^2\left(\frac{\delta\theta}{2}\right) + \left(\frac{\chi}{2}\right)^2} = \frac{2n_\epsilon\pi}{T_p}, \tag{6.33}$$

where

$$\cos^2\left(\frac{\delta\theta}{2}\right) = \frac{g_{\mathsf{bs}}^2}{g_{\mathsf{bs}}^2 + \left(\frac{\chi}{2}\right)^2} \tag{6.34}$$

is the correction to the matrix element from Eq. 6.16, and where in the second line I have used Eq. 6.26. When combining this with Eq. 6.25, we obtain the approximations

$$T_{\mathsf{p}} \approx \frac{2\pi}{|\chi|}\sqrt{4n_\epsilon^2 - 1}\left[1 - \frac{1}{4n_{\mathsf{bs}}^2}\right]^{-\frac{1}{2}} \tag{6.35}$$

$$g_{\mathsf{bs}} \approx |\chi|\sqrt{\frac{n_{\mathsf{bs}}^2 - n_\epsilon^2}{4n_\epsilon^2 - 1}}. \tag{6.36}$$

These expressions are slightly more precise than those in de Graaf *et al.* (2025) since they include the correction to the transition matrix element but they should still not be considered exact since a) the $\delta m \neq 0$ transitions are still ignored and b) Eq. 6.25 approximates that the ancilla remains in $|g\rangle$ when $N = 1$, which is only exactly true for very long pulses. As we will discuss in Chapter 7 when tuning up these pulses experimentally, a Schrödinger equation simulation can be used to quickly finetune g_{bs} and T_{p} near these points.

For the shortest choice of pulse duration (for $n_\epsilon = 1$), the smallest non-zero beamsplitter amplitude occurs for $n_{\mathsf{bs}} = 2$, corresponding to $g_{\mathsf{bs}}/|\chi| \approx 1$, which satisfies the constraint on g_{bs} for avoiding running into unwanted transitions (Eq. 6.24). This gives a pulse duration $T_p \approx 1.8 \times 2\pi/|\chi|$. The ancilla and oscillator state trajectories during this pulse are shown alongside those for the Gaussian pulse in Fig. 6.6. As we can see, this also effectively enacts the desired measurement while returning the oscillator states to where they began. There is therefore flexibility in what pulse shape we use to enact the joint photon number selective measurement, something we will employ experimentally in Chapter 7.

6.5 Reinterpreting Joint Parity

In this chapter and the previous one, we have seen frameworks to understand operations where the beamsplitter and ancilla drives alternate (operator Bloch sphere model) and where these drives occur simultaneously (joint-photon number-splitting spectrum). The two models are not completely disjoint, however, as we can see by reinterpreting the joint parity measurement in this new framework.

At $|g_{\mathsf{bs}}/\chi| = \sqrt{3}/2$, the expression for the transition frequencies (Eq. 6.11) simplifies to

$$\omega_{N,\delta m}^{g_{\mathsf{bs}}/|\chi|=\sqrt{3}/2} = \left(\frac{N}{2} + \delta m\right)\chi. \tag{6.37}$$

For states with N even (odd), the transition frequencies lie at integer (half-integer) multiples of χ. The two ancilla $\pi/2$-pulses, if assumed to be instantaneous, can be expressed in the time domain as

$$f(t) \propto \delta\left(t - \frac{\pi}{|\chi|}\right) + \delta\left(t + \frac{\pi}{|\chi|}\right) \tag{6.38}$$

and in the frequency domain as

$$F(\omega) \propto \cos\left(\frac{\pi\omega}{|\chi|}\right). \tag{6.39}$$

As such, all the transition frequencies for even (odd) total photon number in the two oscillators lie at antinodes (nodes) of the pulse frequency spectrum, as can be seen in Fig. 6.5. This therefore excites the ancilla if and only if N is even, as desired.

While we are able to deduce that this sequence measures joint parity, it is not clear from this framework how the oscillator states are transformed by the measurement, something that is more naturally expressed in the operator Bloch sphere model.

6.6 Joint-photon number-selective control

Just as ancilla π-pulses conditioned on the photon number in a single oscillator enabled SNAP in Ch.2.2.2, joint-photon number-selective pulses also allow us to engineer a 'SNAP-like' entangling gate on two oscillators. By choosing the phases of two back-to-back ancilla π-pulses, so that they enclose a geometric phase on the Bloch sphere of the ancilla, we can choose how much phase to apply to a single joint-photon number manifold. As was discussed in Ch. 6.3, the pulses we can most straightforwardly engineer are those conditioned on $N = 0$ (a manifold containing only a single state, $|0\rangle_A |0\rangle_B$). In this case, the unitary we can apply is

$$\hat{U}(\theta) = e^{i\theta} |0\rangle_A |0\rangle_B \langle 0|_A \langle 0|_B . \tag{6.40}$$

In the case of Fock-encoded qubits in both oscillators, where we encode $|0\rangle_L \equiv |0\rangle$ and $|1\rangle_L \equiv |1\rangle$, this unitary acts as a logical CPHASE(θ) entangling gate on the two encoded qubits, up to single-qubit rotations and a global phase:

$$\text{CPHASE}(\theta) = e^{-i\theta} \left[\left(e^{i\theta |1\rangle_L \langle 1|_L} \right)_A \otimes \left(e^{i\theta |1\rangle_L \langle 1|_L} \right)_B \right] \hat{U}(\theta). \tag{6.41}$$

As an extension, when these two oscillators each play the role of half of a dual-rail qubit (as in Fig. 6.7(a)), where $|0\rangle_L = |0\rangle_A |1\rangle_B$ and $|1\rangle_L = |1\rangle_A |0\rangle_B$, a CPHASE($\theta$) on these two central oscillators also acts as a CPHASE(θ) on the encoded dual-rail qubits.

We can demonstrate the operation on Fock qubits in our two-cavity system, using the pulse sequence shown in Fig. 6.7(b), where the joint-photon-number selective ancilla π-pulses are implemented using two chopped Gaussian pulses with the same frequency detuning $\delta\omega = 0$,

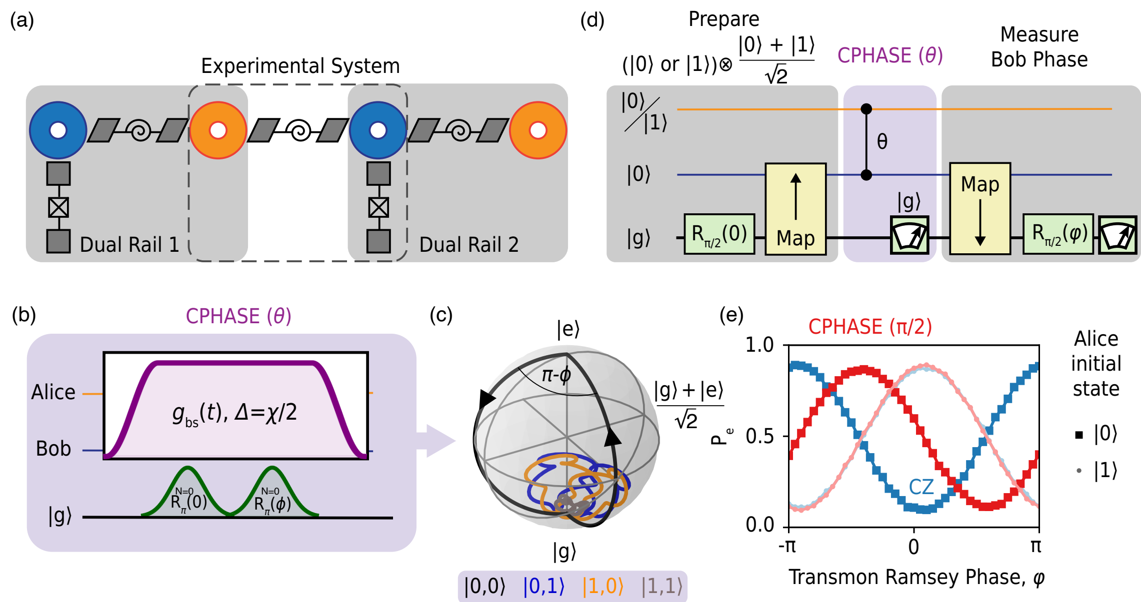

Figure 6.7: **Implementation of a beamsplitter-enabled CPHASE(θ) gate on Fock qubits**. (a) Illustration of two adjacent dual-rail cavity qubits, with beamsplitter couplings enabled by SNAILs. The orange and blue circles represent $\lambda/4$ post cavities viewed from above. The experimental setup used here (within the dashed lines) can be used to operate a gate between these two hypothetical encoded dual-rail qubits. (b) Pulse sequence used for the CPHASE(θ) gate, with the two ancilla pulses applied with a phase differing by ϕ. (c) Simulated ancilla Bloch sphere trajectories during the pulse sequence in (b), conditioned on the photon number in each cavity. The finite duration of the ancilla pulses mean it is very slightly excited out of the ground state for non-$|0,0\rangle$ states. (d) Ramsey sequence used to probe the phase acquired by Bob's oscillator, conditioned on the logical state in Alice. Optimal control theory (OCT) pulses are used to map a superposition state from the ancilla onto Bob's oscillator, and to map the final Bob superposition back onto the ancilla at the end. The initial and final $\pi/2$-pulses differ by a phase φ. (e) Results of Ramsey experiment to probe the phase acquired by Bob's oscillator during the CPHASE(θ) gate with $\theta = \pi/2$ (red) and $\theta = \pi$ (blue, CZ), when Alice oscillator state is initialized in $|0\rangle$ (connected squares) and in $|1\rangle$ (connected dots). Figure reproduced from de Graaf *et al.* (2025).

but with a phase difference ϕ. The beamsplitter amplitude $g_{\text{bs}} \approx 1.6|\chi|$ (shown as a triangle in Fig. 6.1(c)) is chosen to avoid transmon transitions in either the $N = 1$ or $N = 2$ joint-photon-number manifolds. This is required since $|1, 1\rangle$ (with joint photon number $N = 2$) is a valid input state.

The simulated ancilla Bloch sphere trajectories during this pulse sequence, conditioned on the two-cavity state, are shown for the case $\theta = \pi$ (a CZ gate) in Fig. 6.7(c). The two large semicircles of the $|0\rangle_A |0\rangle_B$ trajectory make an angle $\pi - \phi$ at their intersection, imparting a geometric phase on the state. Owing to the finite duration of the ancilla π-pulses, the ancilla is very slightly excited out of its ground state during the sequence for non-$|0\rangle_A |0\rangle_B$ cavity states (although it does return to $|g\rangle$). Unlike in the idealized case of Eq. 6.40, the states $|0\rangle_A |1\rangle_B$, $|1\rangle_A |0\rangle_B$ and $|1\rangle_A |1\rangle_B$ do therefore acquire some geometric phase. However, these phases can be straightforwardly corrected via local rotations on each oscillator to obtain the desired CPHASE(θ) operation.

We can verify the operation of the CPHASE(θ) using the circuit shown in Fig. 6.7(d), where the phase acquired by Bob's oscillator, conditioned on the state in Alice's oscillator, is measured using a Ramsey sequence. First, we prepare either $|0\rangle_L |+\rangle_L$ or $|1\rangle_L |+\rangle_L$ in the oscillators, where $|+\rangle_L = (|0\rangle_L + |1\rangle_L)/\sqrt{2}$. This logical superposition state in Bob is initialized by first preparing a $(|g\rangle + |e\rangle)/\sqrt{2}$ superposition in the ancilla before using an 'encode' OCT pulse which maps the ancilla state onto the oscillator (Heeres *et al.*, 2015). Next, we apply the CPHASE(θ) using the sequence in Fig. 6.7(b), before reading out the ancilla state to verify that it has returned to $|g\rangle$. Post-selecting on measuring $|g\rangle$ catches some of the ancilla errors that occur during or before the CPHASE(θ), and is required to ensure the subsequent measurement is accurate. Finally, we measure the phase of the superposition state in Bob's oscillator by applying a 'decode' OCT pulse, which maps the state back onto the ancilla, before measuring the phase of the ancilla superposition by applying a $\pi/2$-pulse with a variable phase φ and reading out its state.

The results of this experiment (shown in Fig. 6.7(e) for $\theta = \pi/2$ (red) and $\theta = \pi$ (blue)) show our ability to apply a tunable phase on Bob, conditioned on the logical state in Alice. The phase offset of the oscillations in the ancilla population, P_e, encode the final phase of Bob's superposition. When Alice is initialized in $|1\rangle$, the phase oscillations are independent of θ, indicating that the phase on Bob is independent of θ. However, when Alice is initialized in $|0\rangle$, we are able to change the phase on Bob, as evidenced by the leftwards shift of the oscillations. Comparing the two blue (or red) curves shows that a phase of $\theta = \pi$ (or $\pi/2$) is imparted to Bob, conditioned on the state of Alice, hence demonstrating a CZ (or CPHASE($\theta/2$)) gate. The loss in measurement contrast is predominantly set by transmon SPAM errors and cavity photon loss during the OCT 'encode' and 'decode' pulses.

In evaluating the utility of this protocol for error-correction, we are less interested in SPAM errors than the impact of transmon errors during the CPHASE itself. Just as we did for the joint-parity measurement, we can look to the tricks of Ch. 2.2.3 for ways to flag these errors. Replacing the $|g\rangle$-$|e\rangle$ ancilla with a $|g\rangle$-$|f\rangle$ ancilla allows us to catch relaxation errors, as usual. However, key to the fault-tolerant SNAP protocol was the 'equal-latitude' condition – the fact that the ancilla was equally likely to be found in $|e\rangle$ for any oscillator state. This condition is violated for this version of joint-SNAP, as can be seen from the trajectories in Fig. 6.7(c), and so makes the protocol vulnerable to transmon dephasing errors. Performing simultaneous joint-photon-number-selective pulses on multiple number peaks would fix this issue but, as seen in Ch. 6.3, is currently technically challenging.

A similar 'SNAP-like' two-mode gate acting on the $N = 0$ manifold can also be achieved using a 'Y-mon' architecture where a transmon ancilla has a static dispersive coupling to both oscillator modes (Xu *et al.*, 2020). The similarity between these two approaches highlights how a strong tunable beamsplitter can dynamically generate a dispersive coupling between an ancilla and two oscillator modes, with matched χs, despite the ancilla being only statically coupled to one mode.

6.6.1 Completing a two-mode analog of the dispersive toolbox

In Chapter 2, we introduced a toolbox of single-oscillator measurements and gates enabled by photon-number-splitting – the regime in which the dispersive shift greatly exceeds the ancilla decoherence rate, $\chi \gg 1/T_2^{\text{ancilla}}$, and individual oscillator photon number peaks could be resolved in spectroscopy. This (partial) toolbox included the parity measurement, photon number selective measurement, and the SNAP gate. In this chapter, we have shown that there is an analogous joint-photon-number-splitting regime when $g_{\text{bs}} \approx |\chi| \gg 1/T_2^{\text{ancilla}}$ and spectroscopy reveals peaks corresponding to a fixed total photon number in two oscillators. In this regime, the two-oscillator analogs of our dispersive toolbox can be found by taking the regular single-oscillator sequences, doubling their duration and simultaneously applying a beamsplitter drive with a detuning $\Delta = \chi/2$ throughout. This mapping is summarized in Fig. 6.8.

In the next chapter, we will compare how useful the two measurements in this toolbox are at performing a mid-circuit erasure check - a key element of the dual-rail cavity qubit (Teoh *et al.*, 2023).

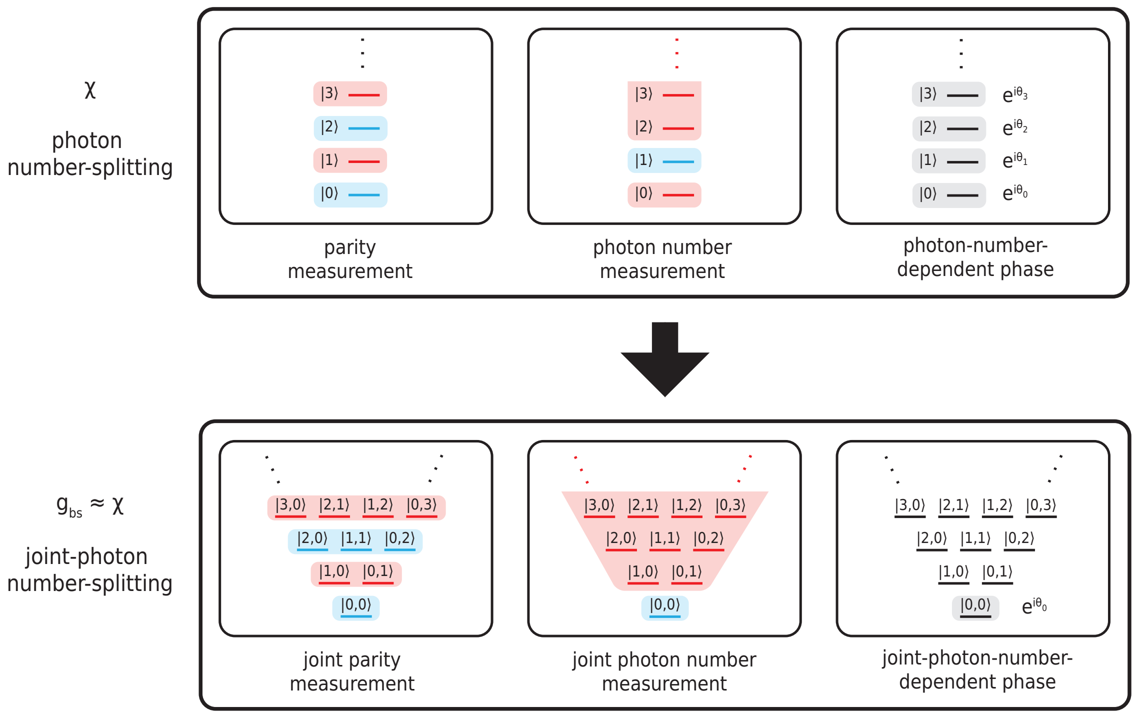

Figure 6.8: **Mapping between single-oscillator and multi-oscillator control techniques**. Top row: partial toolbox of single-oscillator measurements and controls enabled by a dispersively coupled ancilla in the 'photon number-splitting' regime. Bottom row: multi-oscillator analogs of the single-oscillator techniques enabled by applying a strong beamsplitter interaction ($g_{bs} \approx \chi$) during the operation. These operations act on manifolds of states with the same joint photon number N. As we have seen, the multi-oscillator versions of the photon number measurement and SNAP gate are most straightforwardly implemented on the $N = 0$ manifold.

Chapter 7

A hardware-efficient erasure check for dual-rail qubits

The thrust of this thesis has been to provide tools for Gaussian and non-Gaussian operations on pairs of bosonic modes, with a view to creating a network of these modes connected by beamsplitters. In this culminating chapter, I will showcase an application that crucially relies on these elements, namely mid-circuit erasure detection (MCED) – a vital operation for the dual-rail bosonic code.

I will start by providing the salient details regarding the role of the erasure check in the cavity dual-rail encoding (Teoh *et al.*, 2023). Authoritative sources on the dual-rail code can be found in Kubica *et al.* (2023) and the thesis by Teoh (2023), which covers the theory in more substantial detail. The brief summary here will allow us to evaluate the ideal method to use for the erasure check before diving into the experimental implementation.

7.1 The role of an erasure check

The threshold theorem (Knill *et al.*, 1998), which lies at the heart of quantum error correction, states that if physical error rates p_{physical} fall below some threshold $p_{\text{threshold}}$ (determined by the error-correction scheme employed) then one can achieve arbitrarily low *logical* error rates by increasing the number of qubits used to encode the information (quantified by the code distance, d) (Nielsen and Chuang, 2010). This statement can be summarized by the following heuristic:

$$p_{\text{logical}} \propto \left(\frac{p_{\text{physical}}}{p_{\text{threshold}}}\right)^{\frac{d+1}{2}}. \tag{7.1}$$

This expression motivates the search for higher-fidelity operations (reducing p_{physical} substantially below $p_{\text{threshold}}$) and large code distances d (or equivalently the number of physical qubits used to encode a single logical qubit). For circuit-level unbiased noise on qubits encoded in a 2D surface code (Bravyi and Kitaev, 1998; Kitaev, 2003), $p_{\text{threshold}} \approx 1\%$ (Fowler *et al.*, 2012; Wang *et al.*, 2010), depending on the choice of decoder and noise model. This value lies slightly above state-of-the-art error rates for superconducting qubits, with exciting recent work demonstrating that it is possible to suppress p_{logical} with increasing d for a system of 101 physical qubits (Acharya et al., 2024)[1].

While this effort puts the onus on reducing p_{physical}, a parallel effort is focused on *increasing* $p_{\text{threshold}}$. This requires the noise to possess some structure or bias, so that it introduces less entropy to the system per error. Schrödinger cat qubits (Cochrane *et al.*, 1999; Guillaud and Mirrahimi, 2019; Mirrahimi *et al.*, 2014; Puri *et al.*, 2017), with their strong bias towards phase-flips (Z-type errors) over bit-flips (X-type errors), are a classic example of this. Besides having biased noise, it is also necessary to adapt the error correction scheme to make use of this information. Whereas the traditional square 2D surface code devotes equal resources to correcting X- and Z-type errors, modified surface codes (Tuckett *et al.*, 2018, 2020, 2019), such as the XZZX code (Bonilla Ataides *et al.*, 2021; Xu *et al.*, 2022), can be designed to tolerate more of the dominant error type while tolerating fewer of the rarer error type, thereby yielding a higher overall error threshold. In the infinite-noise-bias limit, the XZZX code becomes a classical repetition code, where the 2D code collapses to a 1D line, only correcting for one type of Pauli error at all. This yields a code-capacity threshold of 50% (Girvin, 2023).

[1] These experiments were performed on up to a $d = 7$ system where $7^2 = 49$ 'data' qubits stored the logical information, $7^2 - 1 = 48$ 'measure' qubits detected errors on the data qubits, and an additional 4 'leakage removal' qubits were used to mitigate leakage out of the $|g\rangle - |e\rangle$ manifold.

Erasure qubits (Grassl *et al.*, 1997; Kubica *et al.*, 2023) provide another form of noise structure that can be exploited, as their noise is dominated by detectable leakage errors to known states outside of the logical subspace of the encoded qubit. Detecting these leakage errors and subsequently resetting the qubit to *any* state on the logical Bloch sphere converts these errors to erasure errors. This *generates*, with some probability, a Pauli error (since we have no way of knowing if we reset the qubit to the correct state) but also flags the specific physical qubit on which the erasure qubit occurred and the time-step at which it occurred. The operation which detects leakage errors in a way that does not perturb un-leaked states is known as mid-circuit erasure detection (MCED), or an erasure check.

In the context of the surface code, the most cautious approach is to perform MCEDs after every two-qubit gate[2]. In this case, at the end of each error-correction cycle (consisting of four rounds of two-qubit gates and MCEDs and a single round of stabilizer measurements), the decoder used to determine which correcting operations to perform on the surface code qubits has access to the results of all the usual syndrome measurements plus the time and location of all erasure errors. This extra information makes the job of decoding to find the correct recovery operation much easier. This is reflected in a higher $p_{\text{threshold}}$ (Barrett and Stace, 2010; Delfosse and Zémor, 2020; Kang *et al.*, 2023; Stace *et al.*, 2009), with for example $p_{\text{threshold}} = 4.15\%$ found for circuit-level noise with an erasure fraction, $R_e = p_{\text{erasure}}/(p_{\text{erasure}} + p_{\text{Pauli}}) = 98\%$ (Wu *et al.*, 2022). Systems where the erasures themselves are biased, coming predominantly from one of the qubit states, present the opportunity to reach even higher thresholds (Sahay *et al.*, 2023).

Not only is $p_{\text{threshold}}$ increased, but for a code where erasures are the *only* form of error, and we have perfect detection, the heuristic for the logical error rate from Eq. 7.1 now scales as

$$p_{\text{logical}} \propto \left(\frac{p_{\text{erasure}}}{p_{\text{threshold}}} \right)^d, \tag{7.2}$$

[2]Optimizing the frequency of MCEDs is the subject of a recent work by Gu *et al.* (2024). They find that checking after every two-qubit gate is optimal for large erasure bias, but that with moderate bias and measurement errors, checking after every other gate performs better.

where the exponent is asymptotically twice as large for the same code distance. Erasure qubits therefore promise much faster suppression of logical errors compared to a qubit with unbiased Pauli errors occurring at the same rate, at the cost of adding the hardware for erasure detection. A form of erasure detection that does not introduce extra hardware beyond what is required for existing qubit operations would therefore be hugely valuable.

As discussed in Chapter 2, a dual-rail qubit encoded in cavities (Teoh *et al.*, 2023) is a type of erasure qubit as it transforms cavities' intrinsic noise bias towards relaxation (over dephasing) into a logical noise bias towards leakage errors to the shared $|0, 0\rangle$ state (over Pauli errors within the dual-rail code space). An erasure check, which detects whether the cavities are in $|0, 0\rangle$, is therefore critical to turn a bad, 'leaky' qubit into a good erasure qubit. Furthermore, it must be a *mid-circuit* (as opposed to end-of-line) erasure check so that the logical state of the encoded qubit is preserved in the (hopefully likely) event of no leakage.

7.2 MCED failure mechanisms and their implications

In practice, any MCED will be imperfect and introduce its own errors. These can compromise the advantages of having an erasure qubit in the first place. Different measurement schemes will be subject to different error mechanisms and so knowing how they impact the job of quantum error correction is crucial to evaluating their relative performance.

The four types of error an MCED can suffer are

- **False negatives**: an erasure is not flagged when the qubit is in the leakage space,
- **False positives**: an erasure **is** flagged when the qubit remains in the code space,
- **Leakage**: the qubit suffers a leakage error during the MCED,
- **Pauli errors**: the qubit suffers a Pauli error during the MCED, without flagging an erasure.

These are illustrated in Fig. 7.1.

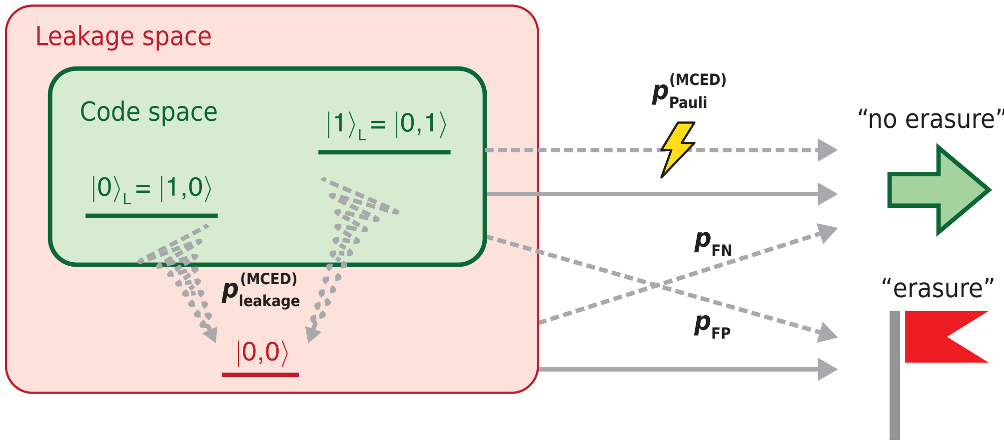

Figure 7.1: **Types of errors during an MCED**. Ideally, the MCED correctly distinguishes between leakage states ("erasure") and code states ("no erasure") without induced any extra errors (solid grey lines). Failure mechanisms (dashed grey lines) include leakage to $|0,0\rangle$ during the measurement (with probability $p^{(\text{MCED})}_{\text{leakage}}$), Pauli errors on code states when no erasure is flagged ($p^{(\text{MCED})}_{\text{Pauli}}$), false negatives that fail to catch leakage states (p_{FN}), and false positives that flag erasures when the dual-rail remained in a code state (p_{FP}).

False negatives

A false negative error allows a leakage state to survive for more than one MCED, so that it is acted on by the subsequent two-qubit gate. In the context of transmon qubits, undetected leakage to states above $|e\rangle$ has proven to be especially problematic as a single error can propagate to multiple correlated errors on adjacent qubits when a two-qubit gate is performed (Varbanov *et al.*, 2020). In this case, the exact leakage state is unknown and so it can be hard to predict the action of a two-qubit gate on a leaked qubit.

In the dual-rail case, since leakage is predominantly to $|0,0\rangle$, the propagation of errors is substantially easier to model. For the $ZZ(\theta)$ gate proposed in Tsunoda *et al.* (2023), leakage on one qubit leads to Pauli errors corresponding to rotations along a single axis (e.g. either I or X errors). Not only does this provide a model for leakage propagation, but also a remarkably simple one. Under the assumption that missed leakage errors do not happen consecutively (i.e. missed leakage errors are caught at the subsequent MCED – a valid assumption when these events are relatively rare), this can be modeled simply as an increase in the Pauli error rate, a corresponding decrease in the erasure fraction R_e and therefore as a decrease in $p_{\text{threshold}}$ (Chang *et al.*, 2024).

One way to understand the relatively low sensitivity towards false negative errors is to recognize that the frequency of missed leakage errors is the product of the false negative probability and the (necessarily small) probability of having a leakage state as input to the MCED in the first place,

$$p_{\text{miss}} = p_{\text{FN}} \times p_{\text{leakage}}. \tag{7.3}$$

Given that missed leakage errors (in this two-qubit gate model) only induce Pauli errors *half* of the time, $p_{\text{miss}}/2$ sets the likelihood of suffering a false-negative-induced Pauli error at each MCED. Provided that this number is less than the existing Pauli error rate in the system, we can say that the false negative rate is not performance-limiting. This target can be expressed as

$$p_{\text{FN}} < 2 \times \frac{p_{\text{Pauli}}}{p_{\text{leakage}}}. \tag{7.4}$$

For state-of-the-art cavities, with a typical physical noise bias $\gamma_\phi/\kappa \approx 10\%$, we expect to see the same logical noise bias, $p_{\text{Pauli}}/p_{\text{leakage}} \approx 10\%$. This results in a target for the false negative rate, $p_{\text{FN}} < 20\%$, that is very lenient.

False positives

The impact of false positives depends on the type of (post-erasure-detection) qubit reset that is available. If the reset is performed via a unitary that maps $|0,0\rangle$ to somewhere on the Bloch sphere, it is impossible to also map every logical state to a point on the Bloch sphere. In the event of a false positive error, the qubit starts in a logical state and the subsequent reset must have some probability of causing a leakage error, which can in principle be engineered to always be to $|0,0\rangle$. In this case, the undetected leakage state can in turn cause Pauli errors in the same way that missed leakages do. However, unlike for missed leakages there is no small pre-factor since $P(\text{no leakage}) \approx 1$. For a unitary reset, we therefore require $p_{\text{FP}} \lesssim p_{\text{Pauli}}$ to retain a Pauli error rate that is set by the intrinsic cavity decoherence – a tight constraint.

Alternatively, if we have access to an unconditional reset that maps all logical states and $|0,0\rangle$ to somewhere on the Bloch sphere, then in the event of a false positive, the erasure conversion is nonetheless performed correctly and it simply adds to the overall erasure rate. Ensuring that this additional contribution does not dominate sets the approximate (and significantly more relaxed) constraint that

$$p_{\text{FP}} < p_{\text{leakage}}. \tag{7.5}$$

An unconditional reset, since it is non-unitary, requires some form of dissipation or measurement. Reset of a dual-rail erasure qubit (unitary or unconditional) has not yet been experimentally demonstrated. For the purposes of this chapter, I will optimize the MCED assuming access to an unconditional reset, but the measurement scheme could be adapted to penalize false positives more severely, as would be necessary with a unitary reset.

Leakage

This category describes all leakage to the ground state during the MCED, including leakage due to the intrinsic relaxation rate of the cavities, and need not imply any enhancement of cavity decoherence during the measurement. Since leakage sets the bulk of the physical errors in Eq. 7.2, they should ideally be kept small relative to the threshold value,

$$p_{\text{leakage}}^{(\text{MCED})} \ll p_{\text{threshold}} \tag{7.6}$$

.

Pauli errors

The rate of Pauli errors occurring during the MCED, $p_{\text{Pauli}}^{(\text{MCED})}$, can be broken down into two sources: underlying cavity errors during the course of the measurement, with probability $p_{\text{Pauli}}^{(\text{intrinsic})}$, and transmon errors which are not flagged as erasures, with probability $p_{\text{Pauli}}^{(\text{transmon})}$. The combined error rate must be kept to a level that preserves the hierarchy of rates ($p_{\text{Pauli}}^{(\text{MCED})} \ll$

$p_{\text{leakage}}^{(\text{MCED})}$) that gives the erasure qubit its benefit. Satisfying

$$p_{\text{Pauli}}^{(\text{transmon})} \lesssim p_{\text{Pauli}}^{(\text{intrinsic})} \tag{7.7}$$

ensures that the erasure fraction R_e is not significantly compromised relative to its idling value, which in turn keeps $p_{\text{threshold}}$ as close as possible to its theoretical maximum value[3].

Synthesis

Ultimately, while I have provided guides for each of these four quantities, they cannot be optimized independently. A holistic optimization should minimize the quantity of interest for error correction,

$$\frac{p_{\text{physical}}}{p_{\text{threshold}}} = \frac{p_{\text{leakage}} + p_{\text{FP}} + p_{\text{Pauli}}}{p_{\text{threshold}}(R_e, p_{\text{FN}})} \approx \frac{p_{\text{leakage}} + p_{\text{FP}}}{p_{\text{threshold}}\,(R_e, p_{\text{FN}})}, \tag{7.8}$$

where in the final approximation we have used the fact that $p_{\text{Pauli}} \ll p_{\text{leakage}}$. A key takeaway is that false negatives are much less harmful than false positives. For example, if we use $p_{\text{leakage}} = 1\%$ and $p_{\text{Pauli}} = 0.1\%$, a false positive rate $p_{\text{FP}} = 1\%$ doubles the fraction above, whereas a false negative rate $p_{\text{FN}} = 1\%$ yields a negligible additional Pauli error probability $\frac{p_{\text{miss}}}{2} = 5 \times 10^{-5}$, with a corresponding negligible impact on the threshold and therefore on the figure-of-merit.

7.3 Choosing a measurement scheme

In our multimode control toolbox (Fig. 6.8) there are two schemes that both enact a mid-circuit measurement that distinguishes the leakage state $|0,0\rangle$ from the code space: the error-detected joint photon number parity measurement (see Ch. 5) and the joint photon number measurement (see Ch. 6) – specifically for joint photon number $N = 0$. With the lessons from the previous section, we can evaluate which approach is better suited to the task.

[3] One caveat is that, per Wu *et al.* (2022) and Chang *et al.* (2024), it appears there is very little increase in either $p_{\text{threshold}}$ or the effective code distance d_{eff} (i.e. the exponent in Eqs. 7.1) beyond $R_e > 0.99$. Therefore, if the intrinsic bias is extremely large, e.g. $R_e \geq 0.999$, it would be fine to tolerate $100 \times p_{\text{leakage}} \gtrsim p_{\text{Pauli}}^{(\text{transmon})} > p_{\text{Pauli}}^{(\text{intrinsic})}$.

7.3.1 Assessing the suitability of a joint-parity measurement

The fully error-detected joint-parity measurement (see Ch. 5.5), using a three-level ancilla and two rounds of measurements, has the advantage that all transmon errors (both dephasing and relaxation) can be detected to first order and distinguished from one another. Thus, ancilla errors can be converted to erasures by resetting the ancilla to $|g\rangle$. This ensures that Pauli errors induced by transmon decoherence and false negative errors are both suppressed to first order, with a residual error rate that scales quadratically with the transmon error rate. This yields increasingly large reductions in these error rates with reductions in transmon decoherence.

This protection comes at the cost of a large number of false positives due to each flagged ancilla error[4]. Not only do all ancilla errors lead to false positives, but the decay and dephasing rates for a $|g\rangle - |f\rangle$ transmon are typically higher than for a $|g\rangle - |e\rangle$ transmon. These two qubits are sensitive to noise at different frequencies (Schoelkopf *et al.*, 2003), but as a rough guide we can assume the same scaling seen in a linear oscillator, where $1/T_1^{f\to e} = 2/T_1^{e\to g}$ and $1/T_\phi^{gf} = 4/T_\phi^{ge}$. While $\chi_{gf} \approx 2\chi_{ge}$, in principle allowing the mapping to proceed twice as fast, this would also require a two-fold increase in g_{bs}, to a level inaccessible in this experiment. For a fairer comparison, I have therefore simulated for the same $\chi_{ge} = \chi_{gf}$, under the assumption that whatever scheme is chosen, one is able to adjust χ to use the same g_{bs} in both cases. If we take these assumptions, along with the transmon decoherence rates in Table A.2, we can perform a Lindblad master equation simulation (using QuTip (Johansson *et al.*, 2013), as we do for all Lindblad master equation simulations in this thesis) of an idealized joint-parity measurement (with infinitely-short ramp times, no delays required to account for ancilla-state-dependent phases and no readout errors). The resulting induced Pauli error rate due to transmon decoherence alone $p_{\mathsf{Pauli}}^{(\mathsf{transmon})} = 0.010\%$ is around an order of magnitude lower than the expected background due to idling errors, $p_{\mathsf{Pauli}}^{(\mathsf{intrinsic})}$. However, the false positive rate, $p_{\mathsf{FP}} = 10.8\%$, is high and would require substantial improvements to transmon coherence to bring much below the level of $p_{\mathsf{threshold}}$.

[4]The term 'false positive' is a bit of a misnomer here since the measurement is able to distinguish between erasures due to ancilla errors and those due to leakage. However, they are still extra erasures not due to leakage errors and so, in the context of the preceding analysis, should be considered part of p_{FP}.

One way to reduce p_{FP} is to compromise on error-detectability of ancilla dephasing errors and perform just a single joint parity measurement. This halves p_{FP} (as well as $p^{(\mathsf{MCED})}_{\mathsf{Pauli}}$ and $p^{(\mathsf{MCED})}_{\mathsf{leakage}}$) and does so at the expense of a higher p_{FN}, to which we have seen we are much less sensitive. Provided that the resulting false negative rate satisfies Eq. 7.4, this is therefore a good trade-off to make. The simulated values for this single-round scheme are shown in the rightmost column of Table 7.1.

We can additionally consider compromising on error-detectability of ancilla relaxation errors and revert to using the $|g\rangle - |e\rangle$ manifold. This more than halves p_{FP} but sacrifices the first-order protection against transmon-induced Pauli errors, with a significant increase in $p^{(\mathsf{transmon})}_{\mathsf{Pauli}}$. However, if this remains less than $p^{(\mathsf{intrinsic})}_{\mathsf{Pauli}}$ (per Eq. 7.7), this could present a tolerable way to reduce the overall erasure rate. We should note that while the simulations employ an infinite-bandwidth instantaneously-ramped beamsplitter pulse, bandwidth constraints (discussed in Ch. 5.3.1) necessitate the use of either OCT pulse shaping of g_{bs} (to ensure the joint parity map completes in $2\pi/|\chi|$ with the correct acquired phases) or delays (which double the duration of the joint-parity map to $4\pi/|\chi|$).

Having considered the different possibilities with a joint parity measurement, can a joint photon number measurement do better?

7.3.2 Evaluating joint photon number selective measurements

What the joint photon number selective (JPNS) measurement can exploit, that none of the joint parity schemes do, is the asymmetry between our sensitivity to false positive and false negative errors. During the joint-parity map, the ancilla state lies on the equator of its Bloch sphere throughout, yielding false negatives and false positives in equal amounts[5]. However, a JPNS measurement only excites the ancilla out of $|g\rangle$ when $N = 0$, an already rare event. Since a transmon in $|g\rangle$ is immune to relaxation and dephasing errors, the result is a bias towards fewer false positives and more false negatives, as desired.

The strength of this bias is set by the shape of the transmon pulse. The longer (and

[5] During the readout portion of the measurement this is not the case, with the ancilla in $|e\rangle$ for leakage states and $|g\rangle$ for non-leaked states. This leads to a bias towards false negatives during the readout portion.

	Joint-photon-number		Joint-parity	
	$\lvert g\rangle - \lvert e\rangle$	$\lvert g\rangle - \lvert f\rangle$	$\lvert g\rangle - \lvert e\rangle$	$\lvert g\rangle - \lvert f\rangle$
p_{FP}	0.61%	2.92%	1.40%	5.43%
$p_{\mathsf{Pauli}}^{(\mathsf{transmon})}$	0.16%	0.10%	0.23%	0.0048%

Table 7.1: **Simulated MCED error rates due to transmon decoherence.** Simulated false positive and induced Pauli error rates for two MCED schemes due to transmon relaxation and dephasing only, using the experimentally obtained system coherences. Both schemes are compared using a $|g\rangle - |e\rangle$ and a $|g\rangle - |f\rangle$ transmon ancilla, where I keep χ fixed in each case. Both schemes use just a **single** measurement round – for two rounds, simply double the numbers above. Finite bandwidth and ancilla heating are not considered for the purposes of this simulation. For the joint-photon-number measurement scheme, a square transmon pulse shape of duration $T_{\mathsf{p}} = 1.7\ \mu$s is used.

smoother) the transmon pulse, the narrower its frequency linewidth relative to $|\chi/2|$, the separation between the $N = 0$ and $N = 1$ peaks in the joint-photon-number splitting regime (see Fig. 6.5). This results in the ancilla state deviating less from $|g\rangle$ during the pulse for a non-leaked input state (with $N = 1$), as we saw in Ch. 6.4. At the same time, for a leakage state input $|0, 0\rangle$, a longer π-pulse provides more opportunity for ancilla errors to induce false negative errors. In order to evaluate the suitability of a joint-photon number measurement, it is therefore worth considering different pulse shapes.

It is not immediately obvious that reducing the deviation from $|g\rangle$ while also extending the duration of the pulse leads to lower false positive rates. To verify this, we can simulate p_{FP} for joint photon number measurements using square transmon pulses of varying duration T_{p}. We measure the probability of finding the ancilla in $|e\rangle$, given that it was initialized in $|\psi\rangle_L \otimes |g\rangle$, averaged over all six cardinal states on the dual-rail Bloch sphere $|\psi\rangle_L$. For each choice of T_{p}, we can measure p_{FP} while sweeping either the transmon dephasing rate $\Gamma_{\phi,t}$ or transmon relaxation rate $\Gamma_{1,t}$ (with no other sources of decoherence). The slopes of these lines (shown in Fig. 7.2(a, c) for pulse duration $T_{\mathsf{p}} = 1.8 \times 2\pi/|\chi|$) show the (linear) susceptibility to each kind of transmon decoherence. Comparing these values across a range of pulse durations confirms that the number of false positive errors due transmon dephasing or relaxation errors decreases with increasing pulse length. The black dashed line in Fig. 7.2(b, d) show a $1/x$ and $1/x^3$ fit (respectively) to the final three points, suggesting approximate relationships (for square-shaped

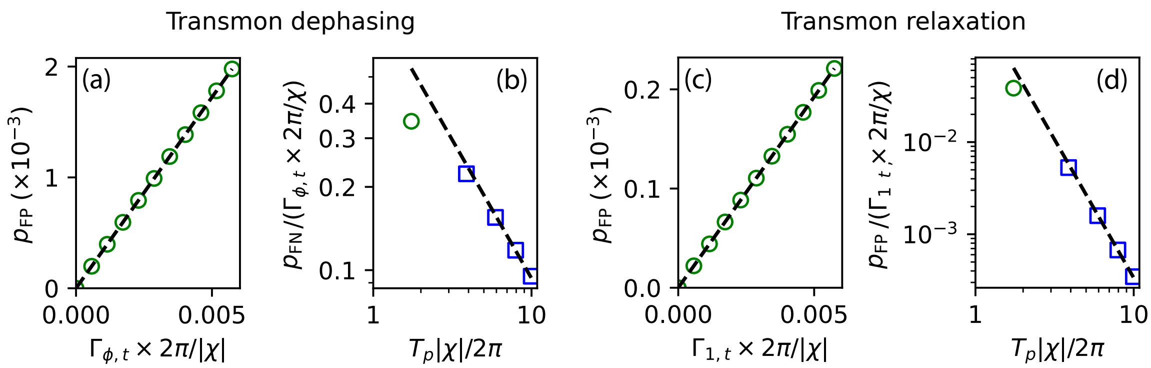

Figure 7.2: **Simulated false positive rate due to ancilla errors during joint photon number measurement**. (a) Simulated false positive rate p_{FP} as a function of transmon dephasing rate $\Gamma_{\phi,\mathsf{t}}$ (normalized to $2\pi/|\chi|$) when only transmon dephasing is included as a collapse operator, for a square transmon pulse of duration $T_{\mathsf{p}} = 1.8\ \mu$s. Dashed line shows a quadratic fit whose slope captures the linear susceptibility to transmon dephasing errors. (b) Fitted susceptibility to transmon dephasing errors, as a function of T_{p}, showing a decrease with longer pulse lengths. Dashed line shows an inverse linear fit to the final three points. (c-d) Same analysis performed for transmon relaxation errors, where the dashed line in (d) is now an inverse cubic fit. Figure reproduced with permission from de Graaf *et al.* (2025).

pulses, at least) of the form

$$p_{\mathsf{FP}}^{\mathsf{dephasing}} \propto \frac{\Gamma_{\phi,t}}{\chi} (T_{\mathsf{p}}\chi)^{-1} \qquad\qquad p_{\mathsf{FP}}^{\mathsf{relaxation}} \propto \frac{\Gamma_{1,t}}{\chi} (T_{\mathsf{p}}\chi)^{-3} . \tag{7.9}$$

Provided that $g_{\mathsf{bs}} \gtrsim |\chi|$, the choice of g_{bs} has a negligible impact on the erasure rate.

An additional vital consideration is how transmon errors propagate to Pauli errors on the dual-rail qubit. This scheme does not use a three-level ancilla, nor is the latitude of the transmon state on the Bloch sphere independent of the dual-rail state (c.f. fault-tolerant SNAP in Ch. 2.2.2), and so it is not first-order fault-tolerant to either transmon relaxation or dephasing – transmon-induced Pauli errors will scale *linearly* with ancilla decoherence rates, e.g. $p_{\mathsf{Pauli}}^{\mathsf{transmon})} = a \times \Gamma_{1,t}$ where a is some prefactor. Nonetheless, the fact that the transmon state remains close to the ground state (for dual-rail logical state inputs) means that this prefactor a can be greatly suppressed. From the same simulations used to probe p_{FP}, we can also extract $p_{\mathsf{Pauli}}^{(\mathsf{transmon})}$, post-selected on the transmon being in $|g\rangle$ (indicating no erasure), as shown in Fig. 7.3. As can be seen in Fig. 7.3(b, e), the susceptibility to transmon relaxation and dephasing rates now

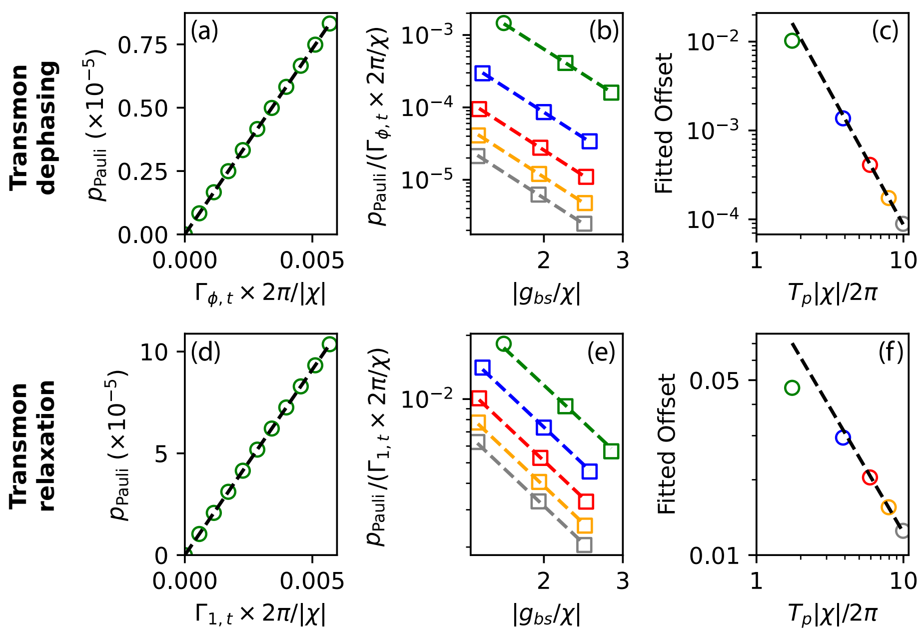

Figure 7.3: **Simulated Pauli error rate due to ancilla errors during joint photon number measurement**. (a) Simulated Pauli error rate p_{Pauli} as a function of transmon dephasing rate $\Gamma_{\phi,\text{t}}$ (normalized to $2\pi/|\chi|$) when only transmon dephasing is included as a collapse operator, for a square transmon pulse of duration $T_{\text{p}} = 1.8$ μs and using a beamsplitter rate $g_{\text{bs}} = 1.6|\chi|$. Dashed line shows a quadratic fit whose slope captures the linear susceptibility to transmon dephasing errors. (b) Fitted susceptibility to transmon dephasing errors for pulses with varying $|g_{\text{bs}}/\chi|$ and varying T_{p} (indicated by the colors), showing a decrease with larger g_{bs} and larger T_{p}. Dashed lines show fits to $A \times (g_{\text{bs}}/\chi)^{-4}$ where the offset A is a free parameter. (c) Fitted offset A for each pulse duration T_{p}, with an inverse cubic fit to the last three points shown in black. (d-f) Same analysis performed for transmon relaxation errors, except the fit in (e) is to $A \times (g_{\text{bs}}/\chi)^{-2}$ and the fit in (f) is to an inverse function. Figure reproduced with permission from de Graaf *et al.* (2025).

reduces both with increasing T_{p} and with increasing g_{bs}. Fits to these susceptibilities suggest approximate relationships (for square-shaped pulses, at least) of the form

$$p_{\text{Pauli}}^{\text{dephasing}} \propto \frac{\Gamma_{\phi,t}}{\chi} \left(\frac{g_{\text{bs}}}{\chi}\right)^{-4} (T_{\text{p}}\chi)^{-3} \propto \frac{(\Gamma_{\phi,t} T_{\text{p}})}{(g_{\text{bs}} T_{\text{p}})^4}, \tag{7.10}$$

$$p_{\text{Pauli}}^{\text{relaxation}} \propto \frac{\Gamma_{1,t}}{\chi} \left(\frac{g_{\text{bs}}}{\chi}\right)^{-2} (T_{\text{p}}\chi)^{-1} \propto \frac{(\Gamma_{1,t} T_{\text{p}})}{(g_{\text{bs}} T_{\text{p}})^2}. \tag{7.11}$$

.

These expressions highlight the benefit of moving to larger g_{bs}, if this can be done without introducing more decoherence. The reason for this is that as $g_{\text{bs}} \times T_{\text{p}}$ (related to the number of beamsplitter oscillations within the pulse) increases, the closer the ancilla trajectories for the input states $|0,1\rangle$ and $|1,0\rangle$ become (see Fig. 6.6). This means that ancilla errors are 'less able' to distinguish between the logical states and so propagate to Pauli errors on the dual-rail qubit.

The expressions also highlight the flexibility that the JPNS measurement has to trade-off our sensitivity to false positive and false negative errors. Switching from a shorter JP scheme to a longer JPNS scheme allows us to reduce p_{FP} at the expense of greater p_{FN}, and increasing T_{p} for the JPNS measurement allows us to further reduce p_{FP}. The ideal point to stop increasing T_{p} is once Eq. 7.5 **and** Eq. 7.7 are satisfied (i.e. $p_{\text{FP}} < p_{\text{leakage}}$ **and** $p_{\text{Pauli}}^{(\text{transmon})} < p_{\text{Pauli}}^{(\text{intrinsic})}$) since beyond this point the intrinsic decoherence of the dual-rail qubit will lead to increasing p_{physical} with increasing T_{p}. Reaching this point (while satisfying Eq. 7.4 for the false negative rate) ensures that the erasure fraction R_e of the MCED is given by its underlying value and not by ancilla relaxation or dephasing. This therefore guides the ideal duration. In the experimental system used here, the difference between the ancilla decoherence rate and the measured oscillator decoherence rates is on the lower end for typical cavity systems. This encourages a shorter square pulse, since the relative penalty for increasing T_{p} is greater. (For systems where this gap is on the higher end, a longer Gaussian pulse might be more appropriate).

The simulated values for a short square pulse ($m = 2$, $n = 1$, in the language introduced in Sec. 6.4) are shown alongside those for a Gaussian pulse in Table 7.1. The $|g\rangle - |e\rangle$ values immediately stand out as attractive, beating the JP scheme in the $|g\rangle - |e\rangle$ manifold and providing p_{FP} **and** $p_{\text{Pauli}}^{(\text{transmon})}$ below their expected background values. Compared to this, the JPNS scheme in the $|g\rangle - |f\rangle$ is unattractive since it accrues substantial extra false positives for little gain in the Pauli error rate.

The conclusion is that **in the context of quantum error correction, a JPNS measurement should be able to outperform a JP measurement as a mid-circuit erasure check**, offering a lower $p_{\text{physical}}/p_{\text{threshold}}$. This is therefore the scheme that will be used in the rest of the chapter. However, there are three main caveats that one should bear in mind:

- **Dual-rail heating.** Unlike the joint-parity measurement, the joint-photon number measurement, as described, is unable to reliably catch heating to the $N = 2$ manifold. While heating is incredibly rare (less than every 100 ms in a typical dual-rail system), we know that long-term undetected leakage can be very harmful. This is also true of other implementations of mid-circuit erasure checks (Koottandavida *et al.*, 2024), which I will discuss later in the chapter. While less straightforward to implement than an "is $N = 0$?" measurement, a simultaneous "is $N = 2$?" measurement could be used to catch heating. One simplifying factor is that we need not preserve oscillator states in the $N = 2$ manifold. Alternatively, we could imagine interspersing JPNS erasure checks with occasional, more-costly JP erasure checks that also serve as leakage reduction units.

- **Dual-rail reset** Performing a two-round MCED (whether JP or JPNS, with a two- or three-level ancilla) ensures that the error syndrome for a leakage state input is distinct from all syndromes for a code state input with a single transmon error. This permits the use of a unitary reset whereby a single photon is loaded into one of the cavities, without having to worry if there is still a photon in the other cavity, in which case the reset excites the dual-rail to the $N = 2$ leakage space. Switching to a lower-error rate single-round MCED requires the use of an unconditional reset which is likely more challenging to implement. If the unconditional reset is especially challenging to implement, pivoting to a two-round MCED may minimize the overall error rate.

- **Non-error-correction contexts.** In the context of short-depth circuits where post-selection rather than error-correction is employed and there is great benefit in having lower Pauli error rates, an ancilla-error-detected JP measurement may be a better approach.

7.4 Tuning-up a mid-circuit erasure check

In Ch. 6.4, we discussed the ideal case of an $N = 0$-selective joint photon number measurement with infinitely fast ramps. Here I will describe the method used to tune-up the pulses in practice, parts of which are reproduced with permission from (de Graaf *et al.*, 2025).

The first practical change needed is the introduction of ramps on both the beamsplitter drive (for the reasons laid out in Ch. 5.3.1) and on the ancilla pulse (in order to prevent coherent errors from driving the $|e\rangle \to |f\rangle$ transition). In both cases, a cosine-ramp is used, taking the form

$$g_{\mathsf{bs}}(t) = \begin{cases} \frac{g_{\mathsf{bs}}}{2}\left(1 - \cos\left(\pi \frac{t}{t^{\mathsf{bs}}_{\mathsf{ramp}}}\right)\right), & 0 < t < t^{\mathsf{bs}}_{\mathsf{ramp}}, \\ g_{\mathsf{bs}}, & t^{\mathsf{bs}}_{\mathsf{ramp}} \leq t \leq T_{\mathsf{p}} - t^{\mathsf{bs}}_{\mathsf{ramp}}, \\ \frac{g_{\mathsf{bs}}}{2}\left(1 + \cos\left(\pi \frac{t - T_{\mathsf{p}} + t^{\mathsf{bs}}_{\mathsf{ramp}}}{t^{\mathsf{bs}}_{\mathsf{ramp}}}\right)\right), & T_{\mathsf{p}} - t^{\mathsf{bs}}_{\mathsf{ramp}} < t < T_{\mathsf{p}}, \end{cases} \tag{7.12}$$

$$\epsilon(t) = \begin{cases} \frac{\epsilon}{2}\left(1 - \cos\left(\pi \frac{t - t^{\mathsf{t}}_0}{t^{\mathsf{t}}_{\mathsf{ramp}}}\right)\right), & t^{\mathsf{t}}_0 < t < t^{\mathsf{t}}_0 + t^{\mathsf{t}}_{\mathsf{ramp}}, \\ \epsilon, & t^{\mathsf{t}}_0 + t^{\mathsf{t}}_{\mathsf{ramp}} \leq t \leq T_{\mathsf{p}} - t^{\mathsf{t}}_0 - t^{\mathsf{t}}_{\mathsf{ramp}}, \\ \frac{\epsilon}{2}\left(1 + \cos\left(\frac{\pi\left(t - T_{\mathsf{p}} + t^{\mathsf{t}}_0 + t^{\mathsf{t}}_{\mathsf{ramp}}\right)}{t^{\mathsf{t}}_{\mathsf{ramp}}}\right)\right), & T_{\mathsf{p}} - t^{\mathsf{t}}_0 - t^{\mathsf{t}}_{\mathsf{ramp}} < t < T_{\mathsf{p}} - t^{\mathsf{t}}_0, \end{cases} \tag{7.13}$$

where $t^{\mathsf{bs}}_{\mathsf{ramp}}$ and $t^{\mathsf{t}}_{\mathsf{ramp}}$ are the duration of the beamsplitter and transmon ramps, respectively, T_{p} is the full duration of the pulse, and $t^{\mathsf{t}}_0 = t^{\mathsf{bs}}_{\mathsf{ramp}}/2 - t^{\mathsf{t}}_{\mathsf{ramp}}/2$ marks the start of the transmon pulse relative to the start of the beamsplitter drive. Setting $t^{\mathsf{t}}_{\mathsf{ramp}} = 24$ ns $\ll 1/|\alpha_t|$ ensures no coherent errors due to the transmon $|e\rangle \to |f\rangle$ transition. An important consideration, discovered from optimizing in simulation for zero false-positive errors and Pauli errors in the absence of decoherence, is that the center (in time) of both the transmon pulse ramps and the beamsplitter pulse ramps should be aligned – ramping up the beamsplitter drive to g_{bs} first and then ramping up the transmon pulse leads unavoidably to false positive errors. This is illustrated in Fig. 7.5(a).

The tune-up procedure consists of the following steps:

1. As in Fig. 4.12, calibrate the following quantities as a function of DAC amplitude:

   - beamsplitter amplitude g_{bs},
   - beamsplitter detuning $\omega_{\Delta=\chi/2}$,
   - ancilla resonance frequency $\omega_{\mathsf{ancilla}}$.

   Both the difference between the cavity frequencies, and the ancilla resonance frequency can be Stark shifted in the presence of the strong beamsplitter drive, and so must be tracked.

2. Identify the maximum beamsplitter amplitude $g_{\text{bs}}^{\text{max}}$ that can be accessed without introducing significantly more errors than when idling. Above certain drive amplitudes, one may see enhanced photon loss, dephasing or heating.

3. Find an initial starting point for the operating parameters: T_{p}, g_{bs}, ϵ, the beamsplitter drive frequency ω, and the detuning of the ancilla drive from its undriven resonance frequency $\delta\omega$. The expressions in Eqs. 6.35-6.36 for the case of infinite pulse bandwidth provide a reasonable starting point for T_{p} and g_{bs}, while Eq. 6.26 is exact for a resonant ancilla pulse.

   A more precise starting point can be obtained by performing a Schrödinger equation simulation which includes the ramp times as well as Stark shifts (i.e., how the ancilla resonance frequency and $\omega_{\Delta=\chi/2}$ change with pump amplitude) and running gradient descent optimization using the infidelity of the state transfers

$$|0,0,g\rangle \to |0,0,e\rangle\,, \tag{7.14}$$

$$|0,1,g\rangle \to |0,1,g\rangle\,, \tag{7.15}$$

$$|1,0,g\rangle \to |1,0,g\rangle \tag{7.16}$$

   as a cost function.

4. Starting with the parameters found above, fine-tune the ancilla drive amplitude ϵ and detuning $\delta\omega$ experimentally by initializing the cavities in $|0,0\rangle$ and performing the erasure check. Choose the parameters that maximize the probability of exciting the transmon, thereby minimizing false negatives. Since the $|0,0\rangle$ state is unaffected by the beamsplitter drive, this calibration depends very weakly on changes in g_{bs} in the following step (only to the extent that the beamsplitter drive Stark shifts ω_{ancilla}.)

5. With T_{p} and ϵ fixed, perform the following experiment while sweeping the values of g_{bs}

and ω:

(a) Initialize the system in $|0, 1, g\rangle$ or $|1, 0, g\rangle$,

(b) Perform N successive erasure checks,

(c) Perform an erasure-detected logical measurement.

From this single experiment, we obtain three key metrics to maximize:

- Probability of passing N checks when initializing in $|0, 1, g\rangle$,
- Probability of passing N checks when initializing in $|1, 0, g\rangle$,
- Probability of returning to $|0, 1, g\rangle$ when initializing in $|0, 1, g\rangle$, (or equivalently the probability of returning to $|1, 0, g\rangle$ when initializing in $|1, 0, g\rangle$.

The first two tell us about p_{FP} while the last metric tells us about $p_{\mathsf{Pauli}}^{(\mathsf{MCED})}$ due to coherent over-/under-rotation errors. With the optimal choice of T_{p}, there should be a point in $(g_{\mathsf{bs}}, \omega_{\mathsf{drive}})$ space that maximizes all three, thereby minimizing both false positives and coherent Pauli errors.

6. If there is no optimal choice for g_{bs} and ω_{drive}, T_{p} must be adjusted. If the operating points that minimized p_{FP} lie at a lower g_{bs} than is required to ensure the dual-rail cavity states return to where they began, the pulse duration T_{p} can be increased (and vice-versa). The previous step can then be repeated.

   We can probe whether the pulse duration is set correctly via a spectroscopy experiment. This involves preparing the cavities in $|0, 0\rangle$, $|0, 1\rangle$ and $|1, 0\rangle$, and performing a single erasure check with a variable detuning $\delta\omega$ on the ancilla pulse. We also post-select on the total photon number in the two cavities remaining the same after the check to remove the effect of photon loss or gain. Fig. 7.4 shows the result for the parameters optimized in this experiment, with the maximum of the $|0, 0\rangle$ peak lining up with the minimum of the $|0, 1\rangle$ and $|1, 0\rangle$ traces. The frequency at which they line up (indicated with the dashed vertical line) is the ancilla frequency detuning used for the check.

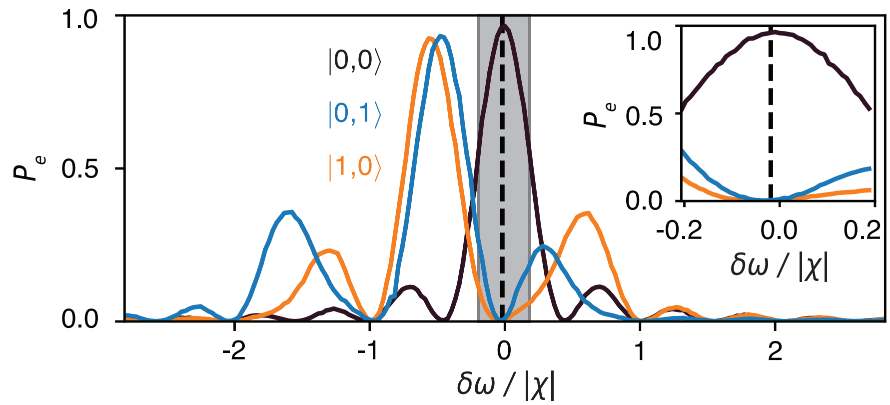

Figure 7.4: **Square Pulse Spectroscopy.** Probability of erasure check reading out $|e\rangle$ as a function of transmon pulse detuning during the check when initializing the dual-rail qubit in $|0,0\rangle$ (black), $|0,1\rangle$ (blue) or $|1,0\rangle$ (orange) state. The results are post-selected on the total photon number in the oscillators remaining the same after the check. Figure reproduced with permission from (de Graaf *et al.*, 2025)

The slight gap between the peaks of the $|0,1\rangle$ and $|1,0\rangle$ curves is expected, even at $\Delta = \chi/2$, and is due to the condition $\epsilon \ll \Omega$ no longer being strictly met. The assumption of treating the ancilla drive as a perturbation that causes transitions between eigenstates with the ancilla in $|g\rangle$ and eigenstates with the ancilla in $|e\rangle$ therefore starts to break down. If the extrema are not aligned at zero detuning, increasing (decreasing) the duration of the square pulse allows us to narrow (widen) the spectrum of the curves such that they align. Doing so will then require re-optimizing g_{bs} and ω_{drive} in this iterative tune-up scheme.

7. Finally, double-check the calibration of ϵ and $\delta\omega$ now that g_{bs} and ω_{drive} have been adjusted.

Besides looking in the frequency domain, we can also measure the ancilla trajectories in the time domain, as we sweep the duration of the applied drive, as shown in Fig. 7.5(b). In agreement with the simulated plots from Fig. 6.6, we see the ancilla state remain relatively close to the ground state during its evolution, when initializing in a dual-rail code state.

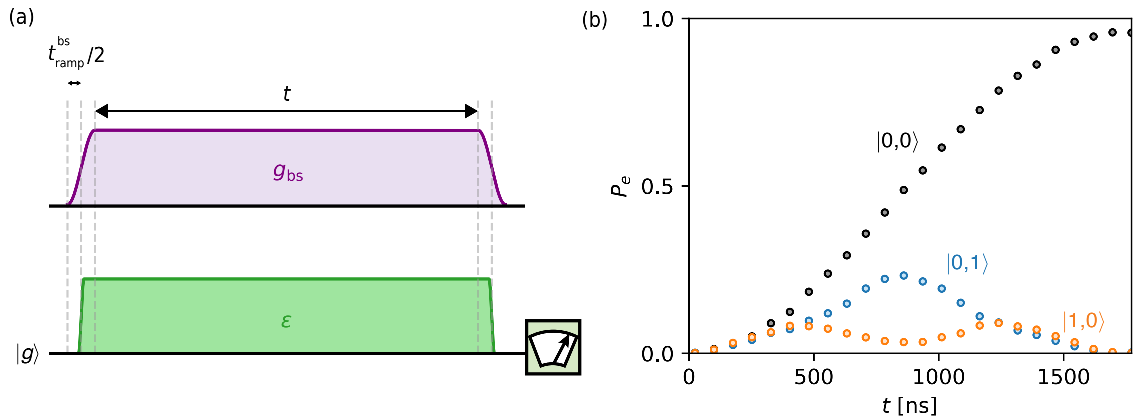

Figure 7.5: **Measured transmon trajectories for different oscillator states.** (a) Pulse sequence used to enact the MCED, consisting of cosine-ramped pulses on both the beamsplitter and transmon, with the centers of the ramps aligned in time. To obtain the trajectories, the flat portion of both pulses are swept together, with t describing the flat duration of the beamsplitter pulse. (b) Measured trajectories showing that the leakage state $|0,0\rangle$ is mapped to the transmon $|e\rangle$ state, indicating an erasure, whereas the dual-rail states $|0,1\rangle$ and $|1,0\rangle$ are mapped to $|g\rangle$ and pass the check. The gap between the $|0,0\rangle$ curve and 1 at the longest time shown (the actual pulse duration) is the false negative rate, p_{FN}. Sub-figure (b) modified with permission from de Graaf *et al.* (2025).

7.5 Evaluating performance

Having argued that the joint photon number measurement can provide the best MCED in our system and calibrated it, we can now evaluate how well it performs with respect to the errors shown in Fig. 7.1.

False negative rate

The false negative rate can be straightforwardly obtained by initializing the cavities in $|0,0\rangle$, performing a single MCED and then performing a final destructive end-of-line measurement of the two-cavity state. Post-selected on finding the cavities in $|0,0\rangle$[6], the probability of measuring the transmon in $|e\rangle$ gives us the false negative rate $p_{\mathsf{FN}} = 3.7 \pm 0.2\%$, well below the $\sim 20\%$ level from Eq. 7.4 at which false negatives start to contribute meaningfully to the Pauli error rate. To be more precise about this statement, simulations performed by Kathleen Chang (see

[6]This allows us to separate out the effect of dual-rail heating which occurs incredibly rarely.

App. F of (de Graaf *et al.*, 2025) for more details), assuming an erasure fraction $R_e = 0.9$ and an MCED after every two-qubit gate, display an error threshold $p_{\text{threshold}} = 3.71 \pm 0.02\%$ with this p_{FN}, as compared to $p_{\text{threshold}} = 3.79 \pm 0.02\%$ with perfect erasure detection. This very small reduction further justifies the decision to opt for a scheme that incurs an increased share of false negative errors.

False positive and induced leakage rates

We can evaluate the types of error affecting states *within* the dual-rail code space by preparing the six cardinal states on the dual-rail Bloch sphere ($|\pm X\rangle, |\pm Y\rangle$ and $|\pm Z\rangle$), performing $4 \leq n \leq 44$ successive MCEDs and then measuring the logical dual-rail operators $\hat{X}_L$, $\hat{Y}_L$ and $\hat{Z}_L$ using photon-number-selective measurements on each cavity (see Fig. 7.6(a)). Importantly, these (destructive) measurements also allow us to distinguish leakage to the $|0,0\rangle$ state. As in the case of the cSWAP (see Ch. 5), repeating the measurement acts as a lever arm to more precisely extract the error rates. The echo Y_L pulse inserted halfway through the sequence (whose role will be described when discussing Pauli errors) does not affect either the rate of leakage or false positive errors.

The overall erasure rate induced per check, $p^{\text{MCED}}_{\text{erasure}} = p^{(\text{MCED})}_{\text{leakage}} + p_{\text{FP}}$, is extracted from the probability of passing n successive MCEDs, unconditioned on the final measured state of the cavities, as shown in Fig. 7.6(b). The exponential decay of the success probability yields $p^{(\text{MCED})}_{\text{erasure}} = 2.92(1)\%$. To separate out $p^{(\text{MCED})}_{\text{leakage}}$, we can separately look at the probability of finding the cavities in $|0,0\rangle$, unconditioned on the results of the MCEDs. The exponential decay of the probability of remaining in the dual-rail code space gives $p^{(\text{MCED})}_{\text{leakage}} = 2.41 \pm 0.02\%$. This value is consistent with the one obtained when replacing each MCED with a delay of the same duration, $p^{(\text{intrinsic})}_{\text{leakage}} = 2.43 \pm 0.02\%$, indicating that the leakage rate is purely due to the intrinsic cavity errors. Meanwhile, we can extract $p_{\text{FP}} = p^{(\text{MCED})}_{\text{erasure}} - p^{(\text{MCED})}_{\text{leakage}} = 0.51 \pm 0.02\%$, slightly below the value simulated for infinite bandwidth pulses in Table 7.1. We have thus satisfied one of our goals, of achieving an erasure rate set predominantly by the underlying cavity relaxation, with far fewer false positive errors than for a joint-parity measurement scheme while also maintaining

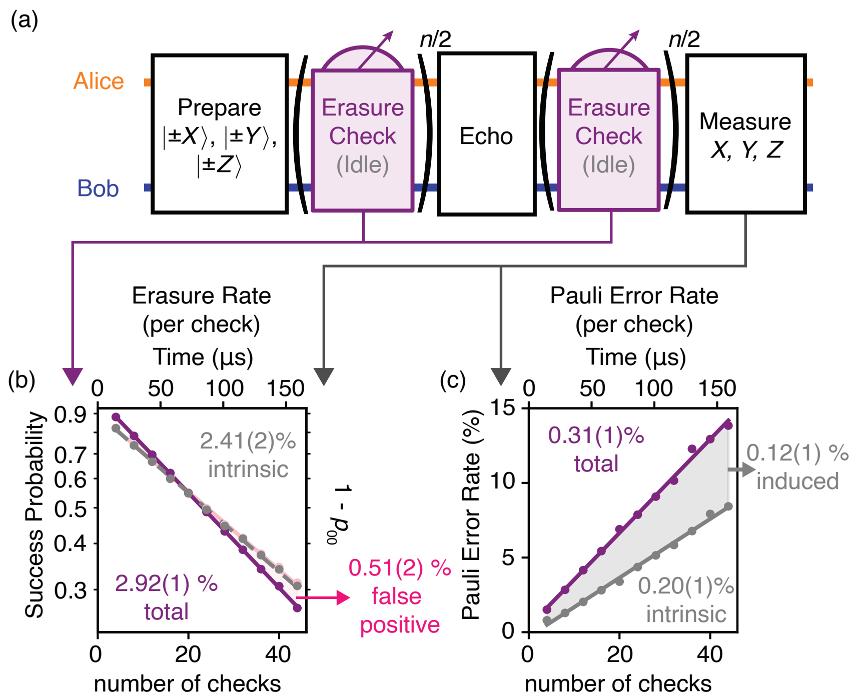

Figure 7.6: **Characterizing errors in dual-rail code space.** (a) Experimental sequence consists of initializing cavities in dual-rail code states, performing n successive MCEDs and measuring $\hat{X}_L$, $\hat{Y}_L$ and $\hat{Z}_L$ on the final state. This final (destructive) measurement can also distinguish leakage to $|0,0\rangle$, as in Chou *et al.* (2024) An echo pulse is inserted halfway to subtract the effect of no-jump backaction arising from post-selecting out leakage in the presence of a large difference in the cavity lifetimes, which leads to a degradation of dual-rail coherence over long timescales (but negligible over the course of a single MCED). (b) From the exponential decay in the probability of passing n successive MCEDs (independent of the final state, and averaged over all input states and measurement axes), we obtain $P_{\text{erasure}}^{(\text{induced})}$ per check (purple). Similarly looking at the probability of measuring $|0,0\rangle$ at the end (independent of the MCED outcomes), we obtain $p_{\text{leakage}}^{(\text{induced})}$ (pink). This is consistent with the same data when idling for the same duration (grey). The difference between $p_{\text{erasure}}^{(\text{induced})}$ and $p_{\text{leakage}}^{(\text{induced})}$ tells us the false positive rate, p_{FP}, indicating that the erasure rate is dominated by intrinsic cavity errors. (c) The decay in the measured state fidelity, post-selected on passing n MCEDs (or after n delays of the same duration) shows us a Pauli error rate that is mostly given by intrinsic cavity errors. Figure reproduced with permission from de Graaf *et al.* (2025).

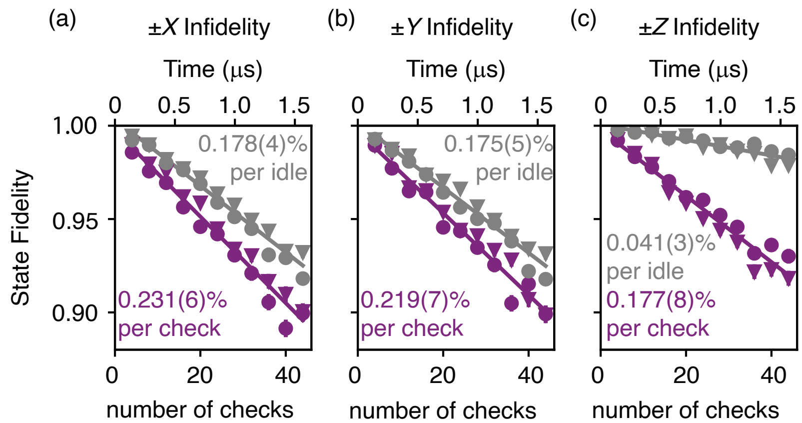

Figure 7.7: **State infidelity after n MCEDs for different initial states.** (a) Measured $\pm\hat{X}_L$ outcomes, post-selected on passing n successive MCEDs, after initializing in $|\pm X\rangle_L$, following sequence in Fig. 7.6(b-c) Same analysis for $|\pm Y\rangle_L$ and $|\pm Z\rangle_L$. Figure reproduced with permission from de Graaf *et al.* (2025).

a false negative rate sufficiently low to have a negligible impact on $p_{\text{threshold}}$.

Intrinsic and induced Pauli error rates

The results of the tomography experiment also allow us to extract the induced Pauli error rate, by looking at the probability of remaining in the initialized state at the end of the sequence, post-selected on passing all of the MCEDs. Looking at the fidelity over many MCEDs allows us to obtain a good estimate of the single-round MCED fidelity, however one effect that can complicate this is no-jump backaction. Post-selecting out leakage events when the loss rate from each cavity is different (as is particularly the case here) deterministically biases the dual-rail state towards one pole of the Bloch sphere. This reduces the fidelity at long times but its effect is quadratically suppressed at the short times relevant for a single MCED (Teoh *et al.*, 2023). In order to remove this effect (for both the idling and MCED data), a single echo pulse is inserted into the QPT sequence, so that cavity photons initialized in each cavity now experience the same time-averaged loss rate over the entire pulse sequence. We should note that while this echo may also remove some of the intrinsic low-frequency dephasing noise on the dual-rail qubit, dephasing due to transmon errors should be unaffected.

These results are shown in Fig. 7.7, alongside the results when idling for the same duration as each MCED, with linear fits showing the reduction in the state fidelity per MCED, for the $|\pm X\rangle_L$, $|\pm Y\rangle_L$ or $|\pm Z\rangle_L$ states[7]. From the state infidelity averaged over the six cardinal states, we can extract an induced Pauli error rate after n MCEDs (or idling for the same duration) via

$$p_{\text{Pauli}}^{\text{(MCED)}} = p_x^{\text{(MCED)}} + p_y^{\text{(MCED)}} + p_z^{\text{(MCED)}} = \frac{3}{2}(1 - \bar{F}), \tag{7.17}$$

which is shown as function of n in Fig. 7.6(c). From the exponential decay, we find $p_{\text{Pauli}}^{\text{(MCED)}} = 0.31 \pm 0.01\%$ per MCED, and $p_{\text{Pauli}}^{\text{(intrinsic)}} = 0.20 \pm 0.01\%$ when idling, and so $p_{\text{Pauli}}^{\text{(induced)}} = 0.12 \pm 0.01\%$ captures the excess of Pauli errors beyond those due to the intrinsic decoherence of the dual-rail qubit. This excess, expected to be due to transmon decoherence, is comparable to the value simulated for infinite bandwidth pulses in Table 7.1 and less than the value due to intrinsic dual-rail errors.

Taken together, these results show the ability of this MCED to effectively protect against transmon decoherence despite not meeting the standard definition of ancilla fault-tolerance. Both $p_{\text{erasure}}^{\text{(MCED)}}$ and $p_{\text{Pauli}}^{\text{(MCED)}}$ are predominantly set by intrinsic cavity decoherence rates and so the erasure fraction R_e is the same as for the idling dual-rail qubit. Combined with a low p_{FN}, this means ancilla errors during the MCED do not compromise $p_{\text{threshold}}$, so that it is still well above the value for a typical (non-erasure) qubit. Nonetheless, $p_{\text{physical}} = 3.23 \pm 0.01\%$ during the MCED itself is too high to currently enable useful quantum error correction since we must also add to this the errors from the two-qubit gate, which are likely higher. A question we would therefore like to answer is how low p_{physical} could be made if the hardware were brought up to the state-of-the-art for coherences.

7.6 How good could it get?

We can predict the performance using Lindblad master equation simulations with transmon coherences $T_1 = T_\phi = 200$ μs and cavity coherences $T_{1,A} = T_{1,B} = 1000$ μs (Ganjam *et al.*,

[7]While the idling dual-rail qubit has a bias against bit-flip errors (since this requires simultaneous relaxation on one cavity and heating on another), applying the beamsplitter drive during the MCED provides significantly less biased errors, although it should be emphasized that this bias is not a useful one for dual-rail qubits.

2024; Oriani *et al.*, 2024; Place *et al.*, 2021). This can be done assuming the same dispersive shift χ and readout duration τ_{readout}, along with pulse parameters (with no ramp times assumed) $g_{\text{bs}}/2\pi = 1.038$ MHz and $T_{\text{p}} = 1.699$ μs, the closest ideal simulated parameters to those used in experiment. Dual-rail dephasing, being harder to simulate without a careful characterization of the noise sources responsible, has not been included.

These simulations predict $p_{\text{leakage}}^{(\text{MCED})} = 0.334\%$ and $p_{\text{FP}} = 0.164\%$, offering a lower p_{physical} substantially below $p_{\text{threshold}}$ while maintaining an erasure rate that is dominated by intrinsic cavity relaxation rates. Meanwhile the simulated $p_{\text{Pauli}}^{(\text{transmon})} = 0.035\%$ is not limiting until the dual-rail dephasing time reaches 10 ms. One factor (besides transmon errors) that starts to become important at this scale is shot-noise dephasing of the dual-rail qubit due to photons in the readout resonator during the MCED, given by the usual Eq. 2.47, where the measured linewidth of the readout resonator $\kappa_r/2\pi = 1.77$ MHz. The dispersive shift between Bob's cavity and the readout resonator, χ_{b,r_b}, can be estimated (assuming no direct capacitive coupling between the two) using

$$\chi_{b,r_b} \approx \frac{\chi_{b,t_b}\chi_b}{\alpha_t} \approx 2\pi \times 2.5 \text{ kHz}. \tag{7.18}$$

Given an average photon number in the readout resonator of $\bar{n} \approx 10$ for a duration of 1 μs, we obtain $p_{\text{Pauli}} = \Gamma_\phi/2 = 0.01\%$ per check. Some form of cloaking (Lledó *et al.*, 2023), to prevent the cavity from interacting with the readout photons, or dynamical decoupling (Ezzell *et al.*, 2023), would therefore be necessary to continue improving beyond this point.

Besides improving the mode coherences, the other big benefit will come from reducing the duration of the MCED. The mapping of the joint photon number onto the transmon state is limited by the magnitude of χ. We saw in Ch. 2.2.4 that increasing $\chi/2\pi > 1$ MHz can lead to unwanted nonlinearity in the cavities, as well as increased dephasing (including thermal shot noise-induced dephasing due to transmon heating). The dual-rail encoding, with on average $1/2$ photon per oscillator is much less sensitive to inherited Kerr nonlinearities than other bosonic encodings, only feeling the effects during two-qubit gates, where two photons may briefly occupy a single oscillator at a time. Provided the transmon can be kept cold (to limit shot noise dephasing), increasing χ to reduce T_{p} may therefore be a good strategy. However, g_{bs} would need to also be increased commensurately. Meanwhile, reducing τ_{readout} requires either increased $\bar{n}$ (requiring care to avoid measurement induced state transitions, or MIST, when doing so (Sank *et al.*, 2016))) or simultaneous increases in the readout linewidth κ_r (requiring improved Purcell filtering to avoid compromising the cavity or transmon lifetimes) and χ_{b,r_b}. Increasing $\bar{n}$ and χ_{b,r_b} both come with increased shot noise induced dephasing of Bob's cavity, placing a bigger onus on the cloaking and dynamical decoupling techniques mentioned earlier.

7.7 Comparison to double-post architecture

The dual-rail architecture considered here, with an ancilla statically coupled to just one of the two oscillators, is not the only possibility. MCEDs have recently been demonstrated using a double-post cavity (Koottandavida *et al.*, 2024), where the two normal modes comprising the dual-rail qubit are common and differential excitations at the two posts. A transmon inserted into the side of this cavity has a static dispersive coupling to both modes, just as in a 'Y-mon' architecture, with a frequency spectrum that distinguishes the leakage state $|0,0\rangle$ from the dual-rail code states. An MCED therefore only requires a frequency-selective transmon pulse, making it both easier to calibrate and (in principle) possible to operate simultaneously with single-qubit gates. However, the integration of a coupler to enable high-fidelity beamsplitter operations between two modes in the same cavity (for single-qubit gates) or two modes in different cavities (for two-qubit gates) has yet to be demonstrated in this architecture. Since the two modes within a dual-rail qubit occupy the same physical space, a coupler bridging between two double-post cavities would both mediate single-qubit and two-qubit gates (see Fig. 7.8). However, this also makes this hardware significantly more prone to crosstalk. As identified in Chapters 3 and 4, avoiding unwanted multiphoton transitions is key to enabling a high-fidelity beamsplitter. A single coupler statically coupled to four cavity modes must therefore take account of multiphoton transitions involving any of their mode frequencies. Furthermore, each of these cavity modes is in turn coupled to three other couplers and a transmon, which can participate in the coupler

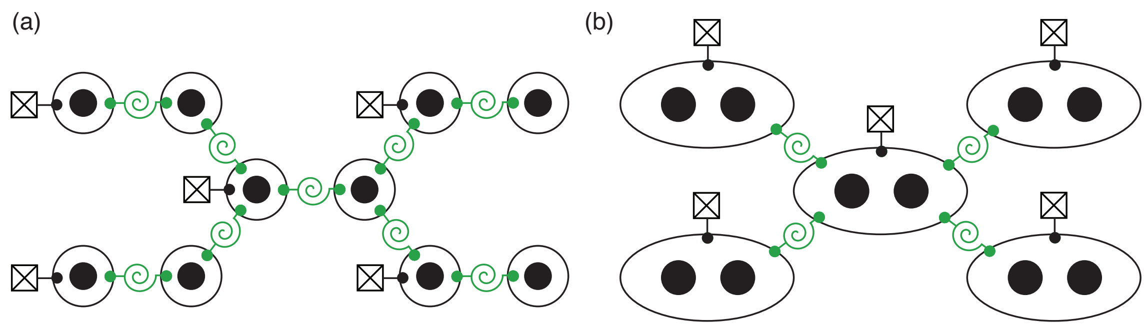

Figure 7.8: **Cavity dual-rail architectures.** (a) Dual-rail layout presented in this work, with concentric circles representing single-mode cavities, coupled to SNAIL couplers for beamsplitters and transmons for nonlinear ancillas. (b) Layout for dual-rail architecture with double-post cavities, where ellipse containing two circles represents a multimode double-post cavity. Double-post architecture allows for fewer couplers and simpler MCED calibration, at the expense of significantly more crosstalk between modes.

Hamiltonian, even if to a weaker degree. The single-post cavity therefore offers better isolation of the different modes but must incorporate more control complexity in the MCED as a result.

Chapter 8

Conclusions and Outlook

8.1 Summary of results

We started this thesis with a vision for extending the successes of quantum operations on individual linear oscillators to a network of linear oscillators via the addition of a 'good' microwave switch. What we have seen is that by adapting the Kerr-free nonlinear SNAIL element devised for parametric amplification to a high-Q environment, it can indeed serve as such a switch. Not only does it provide for the first time a two-mode beamsplitter rate g_{bs} that exceeds the typical rate of the dispersive coupling $|\chi|$ used to control a single mode, it can do so while maintaining high levels of coherence and linearity in the pristine oscillators. This new regime where $g_{\text{bs}} > |\chi|$ has unlocked higher fidelity versions of existing multi-mode schemes, in the case of cSWAP, and new schemes entirely, in the form of non-local joint-parity (JP) and joint-photon-number-selective (JPNS) measurements. In particular, this regime has provided all of the tools necessary for a surface code of cavity dual-rail qubits, a potential platform for hardware-efficient fault-tolerant quantum computation. The JPNS measurement is able to perform a mid-circuit erasure check, the essential ingredient for this new code, detecting photon loss in a pair of oscillators while remaining blind to the photon number in either oscillator individually, only minimally perturbing the encoded logical information. Not only do we achieve error rates approaching the threshold for this code, but remarkably do so with a nonlinear ancilla only statically coupled to one of the two oscillators, showcasing the ability of a high-fidelity beamsplitter to generate non-nearest-neighbor interactions in a minimally-connected network with very low crosstalk.

8.2 Possible next steps

When demonstrating a fast cSWAP on cat states with average total photon number of 4, we saw that even away from the operating point where the cavity-cavity cross-Kerr χ_{ab} completely vanished, the SNAIL induced sufficiently low Kerr nonlinearities that the fidelity was instead limited by the mode coherences. Meanwhile, the mid-circuit erasure detection of a single dual-rail qubit made use of the increased beamsplitter speed g_{bs} but was insensitive to Kerr terms in the Hamiltonian. An exciting next step therefore, would be to showcase the SNAIL beamsplitter on bosonic states with a much larger photon number, where the ability to minimize Kerr nonlinear terms can be leveraged to a greater extent. Beyond-breakeven error correction has been demonstrated in many-photon GKP qubits (Sivak *et al.*, 2023) and qudits (Brock *et al.*, 2024), but performing logical gates would require Gaussian operations, including beamsplitters, with low self-Kerrs, K_a and K_b, and cross-Kerr, χ_{ab} (Noh *et al.*, 2022). In this context it is worth asking what we could do to suppress these terms as far as possible, based on what we have learned here? How could a future SNAIL-oscillator system be modified to allow for as high a photon number as possible while retaining the ability to do fast beamsplitters limited by intrinsic cavity decoherence?

In Ch. 4, we achieved $\chi_{ab} = 0$ but saw that K_a and K_b were prevented from passing through zero by the coupled measurement transmon. In principle, with a more weakly-coupled transmon, K_a and K_b could be nulled by an offsetting contribution from the SNAIL. Indeed, the introduction of echoed conditional displacement (ECD) (Eickbusch *et al.*, 2022) and conditional not displacement (CNOD) (Diringer *et al.*, 2024) techniques allow for fast oscillator control despite a much weaker static transmon coupling. Nonetheless, K_a, K_b and χ_{ab} are unlikely to be completely nulled at the same flux point, which may in turn not be the same flux point that maximizes g_{bs}. Introducing an extra flux loop to the SNAIL to turn it into a gradiometric

SNAIL (Miano *et al.*, 2022) would allow in-situ independent control of g_3 and g_4, which could ease this challenge. A recent alternative to directly coupling a transmon has been to use the SNAIL itself to provide the single-oscillator control, either by using using it to engineer a non-Gaussian trisqueezing drive (Eriksson *et al.*, 2024) or by using it as a Kerr-cat qubit ancilla (Ding *et al.*, 2024). Another alternative would be to use the SNAIL, or another Kerr-free coupler, to mediate a tunable parametric dispersive coupling to the nonlinear transmon ancilla (Maiti *et al.*, 2024). All of these alternative approaches have SNAILs (or similar Kerr-free couplers) as the **only** element statically coupled to the oscillator.

In order to push the limits of how low the Kerr can be suppressed while maintaining a large g_{bs} (even when not operating precisely at the Kerr-free point), it would be incredibly valuable to use a large array of SNAILs, as in an SPA, and apply a much stronger beamsplitter strength $|\xi|$. As discovered here, even with a single-pole on-chip filter, driving the beamsplitter via a capacitively coupled off-chip pin only partially shields the drive from other system modes. The recent development of drive lines that feed directly onto the chip in a 3D package open the door to applying more sophisticated multipole filter design (as in, for example, Putterman *et al.* (2022)) to isolate strongly coupled drive lines from crosstalk and Purcell loss. Besides hardware updates, applying Floquet analysis to further understand the limits on g_{bs} and optimize the frequency stack to avoid multiphoton transitions would greatly help extend the utility of a SNAIL coupler across a wider range of flux operating points.

Another very important consideration when extending to large photon numbers will be understanding what limit the SNAIL coupler places on the dephasing inherited by the cavity modes. A dual-rail qubit provides a useful way of probing this inherited dephasing, since usual qubit noise spectroscopy techniques (Bylander *et al.*, 2011; Yan *et al.*, 2013) can be used while post-selecting out photon loss to $|0,0\rangle$ and, as in the work of Goldblatt *et al.* (2024) with a single-cavity, also post-selecting out transmon errors. One theoretical consideration is to establish how events such as leakage and seepage back into the code space affect the noise spectroscopy dynamics but thankfully with occasional mid-circuit measurements (more frequently than every T_1) it should be possible to make seepage negligible. Noise spectroscopy would also provide a way to assess the viability of an interesting operating point where the SNAIL participation in each dual-rail mode is equal, $p_a = p_b$. At this bias point, both heating of the SNAIL mode and flux noise in the SNAIL loop (to 1$^{\text{st}}$ order) lead to symmetric changes in the oscillator mode frequencies, such that the dual-rail qubit frequency $(\omega_b - \omega_a)$ is unaffected, providing protection against inherited dephasing.

The existence of several interesting operating points for a SNAIL coupler (its $\varphi_{\text{ext}} = n\pi$ flux sweet spots, the dual-rail flux sweet spot, $K_a = 0$, $K_b = 0$, $\Delta_{\text{Stark}} = 0$ and $\chi_{ab} = 0$, maximum g_{bs}, to name a few) highlights the utility of incorporating fast flux control, so that one can optimize the SNAIL for different parts of a quantum circuit – sitting at one operating point when performing beamsplitters and another when idling, for example. There is also the question of how to optimize couplers for different parts of a scaled-up multi-coupler system. The SNAIL offers highly valuable Kerr suppression, but recent SQUID-based couplers have offered the highest beamsplitter fidelities to date, showing 99.92% fidelity at the expense of signficant induced Kerr. In the context of dual-rail qubits, where single qubit gates can be performed without any need to suppress Kerr but where cross-Kerr leads to ZZ errors between adjacent qubits, a heterogeneous architecture may be ideal, with SQUID couplers connecting cavities within a dual-rail and SNAIL couplers connecting cavities in adjacent dual-rail qubits.

These considerations could be applied to a proof-of-principle quantum random access memory (qRAM) (Liu *et al.*, 2023), which has been studied using both Fock-encoded and dual-rail-encoded qubits (Weiss *et al.*, 2024), and is a vital component for quantum algorithms such as Grover's algorithm (Grover, 1996). One key demonstration that is now accessible is to extend the fast transmon-controlled cSWAP shown in this thesis to a cSWAP controlled on the state of an encoded bosonic qubit, which can be achieved using a combination of SWAPs and $ZZ(\theta)$ gates based on the joint-parity map shown in Chapter 5. Combining this with a dual-rail mid-circuit erasure detection (MCED) scheme has been shown to allow for substantially higher query fidelity, where an efficient scheme with a low total erasure rate (as shown in Chapter 7) greatly increases the depth of qRAM circuit that can be developed.

Whereas in the context of qRAM, the MCED is used to post-select out data, a key next step for dual-rail qubits in the context of the surface code is to demonstrate qubit reset conditioned on the MCED outcome, to turn leakage detection into *erasure conversion*. The joint-photon-number-based MCED shown in this thesis demonstrated far fewer erasures than the initially-proposed scheme consisting of two consecutive joint-parity measurements with a $|g\rangle - |f\rangle$ transmon. However unlike this initial scheme, it does not uniquely identify each error syndrome and so requires an unconditional (rather than unitary) system reset. One possible approach within the existing architecture is to do so via a combination of sideband drives to transfer cavity population into the transmons and active reset of these transmons. Another avenue that could leverage the SNAIL's strength as a beamsplitter is to SWAP cavity states into a lossy buffer mode, similarly to how two-photon dissipation in dissipative cat qubits is generated by using a parametric interaction between the storage mode and a lossy buffer (Lescanne *et al.*, 2020; Putterman et al., 2024; Touzard *et al.*, 2018). The ability to incorporate fast flux may then also be useful if there is a SNAIL bias point where it becomes resonant with the lossy mode, thereby hastening the evacuation of photons out of the system (Reed *et al.*, 2010).

The final piece of the dual-rail architecture that would be needed before scaling up is demonstration of the two-qubit gate. The $ZZ(\theta)$ gate of Tsunoda *et al.* (2023), based on the joint-parity measurement demonstrated here, provides a way of doing so while detecting all ancilla errors to first order. However, just as in the case of the measurement, the price paid is a relatively large number of erasures. With improvements in transmon coherences to state-of-the-art levels, the $ZZ(\theta)$ is a viable route, but an interesting question to pursue is whether, inspired by the MCED presented here, there are alternative gate schemes that are less costly. The CPHASE gate discussed in Ch. 6, the most direct analogy to the MCED, indeed reduces the ancilla-induced erasures but is not able to prevent all ancilla dephasing errors from propagating to dual-rail Pauli errors. Another approach that obviates the need to handle ancilla errors at all is to use an induced χ_{ab} within a dual-rail qubit as the requisite nonlinearity for the entangling gate. A heterogeneous system with large χ_{ab} within dual-rails and $\chi_{ab} = 0$ between dual-rails allows for state-dependent phases to accrue when rails are swapped between adjacent dual-rails without ever needing to excite the ancilla – requiring very simple controls while also suppressing idling ZZ errors. Figuring out how to generate a sufficiently *large* χ_{ab} to complete the gate in a short time without inducing any other unwanted nonlinearities that compromise the beamsplitter performance is an interesting problem, but the access we now have to many different kinds of high-fidelity couplers and different coupler operating points makes a much wider range of entangling operations possible.

As this outlook shows, high-speed and high-fidelity beamsplitter operations find use in all areas of bosonic quantum information processing. The ability to engineer them and to combine them with a strong source of nonlinearity in the Josephson junction realizes the requirements of Chuang and Yamamoto (1995) for building a 'simple quantum computer.' What those authors perhaps did not foresee is that it would finally be made possible in a completely different frequency domain (microwave rather than optical) and hardware platform (standing rather than propagating modes). As the field of quantum information processing progresses, we can be hopeful that such cross-pollination of ideas between different sub-fields continues to inspire future research directions.

Bibliography

Acharya et al., (2024), "Quantum error correction below the surface code threshold," arXiv:2408.13687 [quant-ph] .

Ahmed, S., C. Sánchez Muñoz, F. Nori, and A. F. Kockum (2021), Phys. Rev. Res. **3**, 033278.

Ambegaokar, V., and A. Baratoff (1963), Phys. Rev. Lett. **10**, 486.

Arndt, M., O. Nairz, J. Vos-Andreae, C. Keller, G. van der Zouw, and A. Zeilinger (1999), Nature **401** (6754), 680.

Axline, C. J. (2018), *Building Blocks for Modular Circuit QED Quantum Computing*, Ph.D. thesis (Yale University).

Backes, K. M., D. A. Palken, S. A. Kenany, B. M. Brubaker, S. B. Cahn, A. Droster, G. C. Hilton, S. Ghosh, H. Jackson, S. K. Lamoreaux, A. F. Leder, K. W. Lehnert, S. M. Lewis, M. Malnou, R. H. Maruyama, N. M. Rapidis, M. Simanovskaia, S. Singh, D. H. Speller, I. Urdinaran, L. R. Vale, E. C. van Assendelft, K. van Bibber, and H. Wang (2021), Nature **590** (7845), 238.

Barenco, A., A. Berthiaume, D. Deutsch, A. Ekert, R. Jozsa, and C. Macchiavello (1997), SIAM Journal on Computing **26** (5), 1541.

Barrett, S. D., and T. M. Stace (2010), Phys. Rev. Lett. **105**, 200502.

Baskov, R. (2024), In Preparation.

Bergeal, N., F. Schackert, M. Metcalfe, R. Vijay, V. E. Manucharyan, L. Frunzio, D. E. Prober, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret (2010a), Nature **465** (7294), 64.

Bergeal, N., R. Vijay, V. E. Manucharyan, I. Siddiqi, R. J. Schoelkopf, S. M. Girvin, and M. H. Devoret (2010b), Nature Physics **6** (4), 296.

Berry, M. V. (1984), Proceedings of the Royal Society A **392**, 45.

Bertet, P., S. Osnaghi, A. Rauschenbeutel, G. Nogues, A. Auffeves, M. Brune, J. M. Raimond, and S. Haroche (2001), Nature **411** (6834), 166.

Bonilla Ataides, J. P., D. K. Tuckett, S. D. Bartlett, S. T. Flammia, and B. J. Brown (2021), Nature Communications **12** (1), 2172.

Bouchiat, V., D. Vion, P. Joyez, D. Esteve, and M. H. Devoret (1998), Physica Scripta **1998** (T76), 165.

Boyd, R. W. (2008), *Nonlinear Optics, Third Edition*, 3rd ed. (Academic Press, Inc., USA).

Braginsky, V. B., and F. Y. Khalili (1996), Rev. Mod. Phys. **68**, 1.

Bravyi, S. B., and A. Y. Kitaev (1998), "Quantum codes on a lattice with boundary," arXiv:quant-ph/9811052 [quant-ph] .

Brock, B. L., S. Singh, A. Eickbusch, V. V. Sivak, A. Z. Ding, L. Frunzio, S. M. Girvin, and M. H. Devoret (2024), "Quantum error correction of qudits beyond break-even," arXiv:2409.15065 [quant-ph] .

Brubaker, B. M., L. Zhong, Y. V. Gurevich, S. B. Cahn, S. K. Lamoreaux, M. Simanovskaia, J. R. Root, S. M. Lewis, S. Al Kenany, K. M. Backes, I. Urdinaran, N. M. Rapidis, T. M. Shokair, K. A. van Bibber, D. A. Palken, M. Malnou, W. F. Kindel, M. A. Anil, K. W. Lehnert, and G. Carosi (2017), Phys. Rev. Lett. **118**, 061302.

Buhrman, H., R. Cleve, J. Watrous, and R. de Wolf (2001), Physical Review Letters **87** (16), 10.1103/physrevlett.87.167902.

Burkhart, L. D., J. D. Teoh, Y. Zhang, C. J. Axline, L. Frunzio, M. Devoret, L. Jiang, S. Girvin, and R. Schoelkopf (2021), PRX Quantum **2**, 030321.

Bylander, J., S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J.-S. Tsai, and W. D. Oliver (2011), Nature Physics **7** (7), 565.

Cahill, K. E., and R. J. Glauber (1969), Phys. Rev. **177**, 1882.

Cai, W., Y. Ma, W. Wang, C.-L. Zou, and L. Sun (2021), Fundamental Research **1** (1), 50.

Campagne-Ibarcq, P., A. Eickbusch, S. Touzard, E. Zalys-Geller, N. E. Frattini, V. V. Sivak, P. Reinhold, S. Puri, S. Shankar, R. J. Schoelkopf, L. Frunzio, M. Mirrahimi, and M. H. Devoret (2020), Nature **584** (7821), 368.

Campos, R. A., B. E. A. Saleh, and M. C. Teich (1989), Phys. Rev. A **40**, 1371.

Catelani, G., R. J. Schoelkopf, M. H. Devoret, and L. I. Glazman (2011), Phys. Rev. B **84**, 064517.

Chang, K., S. Singh, J. Claes, K. Sahay, J. Teoh, and S. Puri (2024), "Surface code with imperfect erasure checks," arXiv:2408.00842 [quant-ph] .

Chapman, B. J., S. J. de Graaf, S. H. Xue, Y. Zhang, J. Teoh, J. C. Curtis, T. Tsunoda, A. Eickbusch, A. P. Read, A. Koottandavida, S. O. Mundhada, L. Frunzio, M. Devoret, S. Girvin, and R. Schoelkopf (2023), PRX Quantum **4**, 020355.

Childs, A. M., H. Fu, D. Leung, Z. Li, M. Ozols, and V. Vyas (2024), "Streaming quantum state purification," arXiv:2309.16387 [quant-ph] .

Chou, K. (2018), *Teleported operations between logical qubits in circuit quantum electrodynamics*, Ph.D. thesis (Yale University).

Chou, K. S., J. Z. Blumoff, C. S. Wang, P. C. Reinhold, C. J. Axline, Y. Y. Gao, L. Frunzio, M. H. Devoret, L. Jiang, and R. J. Schoelkopf (2018), Nature **561** (7723), 368.

Chou, K. S., T. Shemma, H. McCarrick, T.-C. Chien, J. D. Teoh, P. Winkel, A. Anderson, J. Chen, J. C. Curtis, S. J. de Graaf, J. W. O. Garmon, B. Gudlewski, W. D. Kalfus, T. Keen, N. Khedkar, C. U. Lei, G. Liu, P. Lu, Y. Lu, A. Maiti, L. Mastalli-Kelly, N. Mehta, S. O. Mundhada, A. Narla, T. Noh, T. Tsunoda, S. H. Xue, J. O. Yuan, L. Frunzio, J. Aumentado, S. Puri, S. M. Girvin, S. H. Moseley, and R. J. Schoelkopf (2024), Nature Physics **20** (9), 1454.

Chuang, I. L., and Y. Yamamoto (1995), Phys. Rev. A **52**, 3489.

Cochrane, P. T., G. J. Milburn, and W. J. Munro (1999), Phys. Rev. A **59**, 2631.

Cohen, J., A. Petrescu, R. Shillito, and A. Blais (2023), PRX Quantum **4**, 020312.

Cross, A. W., L. S. Bishop, S. Sheldon, P. D. Nation, and J. M. Gambetta (2019), Physical Review A **100** (3), 10.1103/physreva.100.032328.

Curtis, J. C., C. T. Hann, S. S. Elder, C. S. Wang, L. Frunzio, L. Jiang, and R. J. Schoelkopf (2021), Phys. Rev. A **103**, 023705.

Davisson, C., and L. H. Germer (1927), Nature **119** (2998), 558.

De Broglie, L. (1923), Nature **112** (2815), 540.

Delfosse, N., and G. Zémor (2020), Phys. Rev. Res. **2**, 033042.

Diamond, S., V. Fatemi, M. Hays, H. Nho, P. D. Kurilovich, T. Connolly, V. R. Joshi, K. Serniak, L. Frunzio, L. I. Glazman, and M. H. Devoret (2022), PRX Quantum **3**, 040304.

DiCarlo, L., J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf (2009), Nature **460** (7252), 240.

Ding, A. Z., B. L. Brock, A. Eickbusch, A. Koottandavida, N. E. Frattini, R. G. Cortinas, V. R. Joshi, S. J. de Graaf, B. J. Chapman, S. Ganjam, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret (2024), "Quantum control of an oscillator with a kerr-cat qubit," arXiv:2407.10940 [quant-ph] .

Ding, L., M. Hays, Y. Sung, B. Kannan, J. An, A. Di Paolo, A. H. Karamlou, T. M. Hazard, K. Azar, D. K. Kim, B. M. Niedzielski, A. Melville, M. E. Schwartz, J. L. Yoder, T. P. Orlando, S. Gustavsson, J. A. Grover, K. Serniak, and W. D. Oliver (2023), Phys. Rev. X **13**, 031035.

Diringer, A. A., E. Blumenthal, A. Grinberg, L. Jiang, and S. Hacohen-Gourgy (2024), Phys. Rev. X **14**, 011055.

Dixit, A. V., S. Chakram, K. He, A. Agrawal, R. K. Naik, D. I. Schuster, and A. Chou (2021), Phys. Rev. Lett. **126**, 141302.

Dowling, J. P., and G. Milburn (2003), Phil. Trans. R. Soc. A. **361**, 1655.

Eichler, C., and A. Wallraff (2014), EPJ Quantum Technology **1** (1), 2.

Eickbusch, A., V. Sivak, A. Z. Ding, S. S. Elder, S. R. Jha, J. Venkatraman, B. Royer, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret (2022), Nature Physics **18** (12), 1464.

Eriksson, A. M., T. Sépulcre, M. Kervinen, T. Hillmann, M. Kudra, S. Dupouy, Y. Lu, M. Khanahmadi, J. Yang, C. Castillo-Moreno, P. Delsing, and S. Gasparinetti (2024), Nature Communications **15** (1), 2512.

Estermann, I., and O. Stern (1930), Zeitschrift für Physik **61** (1), 95.

Ezzell, N., B. Pokharel, L. Tewala, G. Quiroz, and D. A. Lidar (2023), Phys. Rev. Appl. **20**, 064027.

Flammia, S. T., and Y.-K. Liu (2011), Phys. Rev. Lett. **106**, 230501.

Flurin, E., N. Roch, J. D. Pillet, F. Mallet, and B. Huard (2015), Phys. Rev. Lett. **114**, 090503.

Fowler, A. G., M. Mariantoni, J. M. Martinis, and A. N. Cleland (2012), Phys. Rev. A **86**, 032324.

Frattini, N. (2021), *Three-wave Mixing in Superconducting Circuits: Stabilizing Cats with SNAILs*, Ph.D. thesis (Yale University).

Frattini, N. E., R. G. Cortiñas, J. Venkatraman, X. Xiao, Q. Su, C. U. Lei, B. J. Chapman, V. R. Joshi, S. M. Girvin, R. J. Schoelkopf, S. Puri, and M. H. Devoret (2024), Phys. Rev. X **14**, 031040.

Frattini, N. E., V. V. Sivak, A. Lingenfelter, S. Shankar, and M. H. Devoret (2018), Phys. Rev. Appl. **10**, 054020.

Frattini, N. E., U. Vool, S. Shankar, A. Narla, K. M. Sliwa, and M. H. Devoret (2017), Applied Physics Letters **110** (22), 222603.

Fredkin, E., and T. Toffoli (1982), International Journal of Theoretical Physics **21** (3), 219.

Gambetta, J., A. Blais, D. I. Schuster, A. Wallraff, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf (2006), Phys. Rev. A **74**, 042318.

Ganjam, S., Y. Wang, Y. Lu, A. Banerjee, C. U. Lei, L. Krayzman, K. Kisslinger, C. Zhou, R. Li, Y. Jia, M. Liu, L. Frunzio, and R. J. Schoelkopf (2024), Nature Communications **15** (1), 3687.

Gao, J., M. Daal, A. Vayonakis, S. Kumar, J. Zmuidzinas, B. Sadoulet, B. A. Mazin, P. K. Day, and H. G. Leduc (2008), Applied Physics Letters **92** (15), 152505.

Gao, Y. (2018), *Multi-Cavity Operations in Circuit Quantum Electrodynamics*, Ph.D. thesis (Yale University).

Gao, Y. Y., B. J. Lester, K. S. Chou, L. Frunzio, M. H. Devoret, L. Jiang, S. Girvin, and R. J. Schoelkopf (2019), Nature **566**, 509.

Gao, Y. Y., B. J. Lester, Y. Zhang, C. Wang, S. Rosenblum, L. Frunzio, L. Jiang, S. M. Girvin, and R. J. Schoelkopf (2018), Phys. Rev. X **8**, 021073.

Geerlings, K., Z. Leghtas, I. M. Pop, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret (2013), Phys. Rev. Lett. **110**, 120501.

Gertler, J. M., S. van Geldern, S. Shirol, L. Jiang, and C. Wang (2023), PRX Quantum **4**, 020319.

Girvin, S. M. (2023), SciPost Phys. Lect. Notes , 70.

Gloos, K., R. S. Poikolainen, and J. P. Pekola (2000), Applied Physics Letters **77** (18), 2915.

Goldblatt, U., N. Kahn, S. Hazanov, O. Milul, B. Guttel, L. M. Joshi, D. Chausovsky, F. Lafont, and S. Rosenblum (2024), "Recovering quantum coherence of a cavity qubit through environment monitoring and active feedback," arXiv:2403.02081 [quant-ph] .

Gottesman, D., A. Kitaev, and J. Preskill (2001), Phys. Rev. A **64**, 012310.

de Graaf, S. J., S. H. Xue, B. J. Chapman, J. D. Teoh, T. Tsunoda, P. Winkel, J. W. O. Garmon, K. M. Chang, L. Frunzio, S. Puri, and R. J. Schoelkopf (2025), npj Quantum Information **11** (1), 1.

Grassl, M., T. Beth, and T. Pellizzari (1997), Phys. Rev. A **56**, 33.

Grimm, A., N. E. Frattini, S. Puri, S. O. Mundhada, S. Touzard, M. Mirrahimi, S. M. Girvin, S. Shankar, and M. H. Devoret (2020), Nature **584** (7820), 205.

Grimsmo, A. L., J. Combes, and B. Q. Baragiola (2020), Phys. Rev. X **10**, 011058.

Grover, L. K. (1996), in *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*, STOC '96 (Association for Computing Machinery, New York, NY, USA) p. 212–219.

Gu, S., Y. Vaknin, A. Retzker, and A. Kubica (2024), "Optimizing quantum error correction protocols with erasure qubits," arXiv:2408.00829 [quant-ph] .

Guillaud, J., and M. Mirrahimi (2019), Phys. Rev. X **9**, 041053.

Hann, C. (2021), *Practicality of Quantum Random Access Memory*, Ph.D. thesis (Yale University).

Hann, C. T., G. Lee, S. Girvin, and L. Jiang (2021), PRX Quantum **2**, 020311.

Harrow, A. W., A. Hassidim, and S. Lloyd (2009), Phys. Rev. Lett. **103**, 150502.

Heeres, R. W., B. Vlastakis, E. Holland, S. Krastanov, V. V. Albert, L. Frunzio, L. Jiang, and R. J. Schoelkopf (2015), Phys. Rev. Lett. **115**, 137002.

Heisenberg, W. (1927), Zeitschrift für Physik **43** (3), 172.

Hofheinz, M., E. M. Weig, M. Ansmann, R. C. Bialczak, E. Lucero, M. Neeley, A. D. O'Connell, H. Wang, J. M. Martinis, and A. N. Cleland (2008), Nature **454** (7202), 310.

Houck, A. A., J. A. Schreier, B. R. Johnson, J. M. Chow, J. Koch, J. M. Gambetta, D. I. Schuster, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf (2008), Phys. Rev. Lett. **101**, 080502.

Hu, L., Y.-C. Ma, Y. Xu, W.-T. Wang, Y.-W. Ma, K. Liu, H.-Y. Wang, Y.-P. Song, M.-H. Yung, and L.-Y. Sun (2018), Science Bulletin **63** (5), 293.

Johansson, J., P. Nation, and F. Nori (2013), Computer Physics Communications **184** (4), 1234.

Johnson, B. R., M. D. Reed, A. A. Houck, D. I. Schuster, L. S. Bishop, E. Ginossar, J. M. Gambetta, L. DiCarlo, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf (2010), Nature Physics **6** (9), 663.

Jordan, P. (1935), Zeitschrift für Physik **94**, 531.

Josephson, B. (1962), Physics Letters **1** (7), 251.

Joshi, A., K. Noh, and Y. Y. Gao (2021), Quantum Science and Technology **6** (3), 033001.

Josza, R., and N. Linden (2003), Proc. R. Soc. Lond. A. **459**, 2011.

Kang, M., W. C. Campbell, and K. R. Brown (2023), PRX Quantum **4**, 020358.

Kapit, E. (2018), Phys. Rev. Lett. **120**, 050503.

Katz, O., and C. Monroe (2023), Phys. Rev. Lett. **131**, 033604.

Kerr, J. (1875), The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **50** (332), 337.

Khezri, M., A. Opremcak, Z. Chen, K. C. Miao, M. McEwen, A. Bengtsson, T. White, O. Naaman, D. Sank, A. N. Korotkov, Y. Chen, and V. Smelyanskiy (2023), Phys. Rev. Appl. **20**, 054008.

Kim, Y., A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala (2023), Nature **618** (7965), 500.

Kirchmair, G., B. Vlastakis, Z. Leghtas, S. E. Nigg, H. Paik, E. Ginossar, M. Mirrahimi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf (2013), Nature **495** (7440), 205.

Kitaev, A. (2003), Annals of Physics **303** (1), 2.

Kivlichan, I. D., J. McClean, N. Wiebe, C. Gidney, A. Aspuru-Guzik, G. K.-L. Chan, and R. Babbush (2018), Phys. Rev. Lett. **120**, 110501.

Knill, E., R. Laflamme, and G. J. Milburn (2001), Nature **409** (6816), 46.

Knill, E., R. Laflamme, and L. Viola (2000), Phys. Rev. Lett. **84**, 2525.

Knill, E., R. Laflamme, and W. H. Zurek (1998), Science **279** (5349), 342, https://www.science.org/doi/pdf/10.1126/science.279.5349.342 .

Koch, J., T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf (2007), Phys. Rev. A **76**, 042319.

Kochetov, B. A., and A. Fedorov (2015), Phys. Rev. B **92**, 224304.

Kok, P., W. J. Munro, K. Nemoto, T. C. Ralph, J. P. Dowling, and G. J. Milburn (2007), Reviews of Modern Physics **79** (1), 135–174.

Koottandavida, A., I. Tsioutsios, A. Kargioti, C. R. Smith, V. R. Joshi, W. Dai, J. D. Teoh, J. C. Curtis, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret (2024), Phys. Rev. Lett. **132**, 180601.

Krantz, P., M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver (2019), Applied Physics Reviews **6** (2), 021318, https://pubs.aip.org/aip/apr/article-pdf/doi/10.1063/1.5089550/16667201/021318_1_online.pdf .

Krastanov, S., V. V. Albert, C. Shen, C.-L. Zou, R. W. Heeres, B. Vlastakis, R. J. Schoelkopf, and L. Jiang (2015), Phys. Rev. A **92**, 040303.

Krinner, S., S. Storz, P. Kurpiers, P. Magnard, J. Heinsoo, R. Keller, J. Lütolf, C. Eichler, and A. Wallraff (2019), EPJ Quantum Technology **6** (1), 2.

Kubica, A., A. Haim, Y. Vaknin, H. Levine, F. Brandão, and A. Retzker (2023), Phys. Rev. X **13**, 041022.

Kumar, K. S., A. Vepsäläinen, S. Danilin, and G. S. Paraoanu (2016), Nature Communications **7** (1), 10628.

Lau, H.-K., and M. B. Plenio (2016), Phys. Rev. Lett. **117**, 100501.

Leibfried, D., D. M. Meekhof, B. E. King, C. Monroe, W. M. Itano, and D. J. Wineland (1996), Phys. Rev. Lett. **77**, 4281.

de Leon, N. P., K. M. Itoh, D. Kim, K. K. Mehta, T. E. Northup, H. Paik, B. S. Palmer, N. Samarth, S. Sangtawesin, and D. W. Steuerman (2021), Science **372** (6539), eabb2823, https://www.science.org/doi/pdf/10.1126/science.abb2823 .

Leonhardt, U. (1997), *Measuring the quantum state of light*, Vol. 22 (Cambridge University Press).

Lescanne, R., M. Villiers, T. Peronnin, A. Sarlette, M. Delbecq, B. Huard, T. Kontos, M. Mirrahimi, and Z. Leghtas (2020), Nature Physics **16** (5), 509.

Levine et al., (2024), Phys. Rev. X **14**, 011051.

Lisenfeld, J., A. Bilmes, S. Matityahu, S. Zanker, M. Marthaler, M. Schechter, G. Schön, A. Shnirman, G. Weiss, and A. V. Ustinov (2016), Scientific Reports **6** (1), 23786.

Liu, C., M. Wang, S. A. Stein, Y. Ding, and A. Li (2023), "Quantum memory: A missing piece in quantum computing units," arXiv:2309.14432 [quant-ph] .

Liu, G., T.-C. Chien, X. Cao, O. Lanes, E. Alpern, D. Pekker, and M. Hatridge (2017), Applied Physics Letters **111** (20), 202603.

Liu, Y., S. Singh, K. C. Smith, E. Crane, J. M. Martyn, A. Eickbusch, A. Schuckert, R. D. Li, J. Sinanan-Singh, M. B. Soley, T. Tsunoda, I. L. Chuang, N. Wiebe, and S. M. Girvin (2024), "Hybrid oscillator-qubit quantum processors: Instruction set architectures, abstract machine models, and applications," arXiv:2407.10381 [quant-ph] .

Lledó, C., R. Dassonneville, A. Moulinas, J. Cohen, R. Shillito, A. Bienfait, B. Huard, and A. Blais (2023), Nature Communications **14** (1), 6313.

Lloyd, S. (1992), Physics Letters A **167** (3), 255.

Lloyd, S., and S. L. Braunstein (1999), Phys. Rev. Lett. **82**, 1784.

Lu, Y., A. Maiti, J. W. O. Garmon, S. Ganjam, Y. Zhang, J. Claes, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf (2023), Nature Communications **14** (1), 5767.

Ma, W.-L., M. Zhang, Y. Wong, K. Noh, S. Rosenblum, P. Reinhold, R. J. Schoelkopf, and L. Jiang (2020), Phys. Rev. Lett. **125**, 110503.

Mach, L. (1892), Zeitschrift für Instrumentenkunde **12**, 89.

Maiti, A., J. Garmon, Y. Lu, Y. Wang, P. Winkel, S. Ganjam, L. Frunzio, and R. Schoelkopf (2024), "Ancilla-fault tolerant control of linear quantum memories with dynamic dispersive coupling," Bulletin of the American Physical Society.

Mallet, F., F. R. Ong, A. Palacios-Laloy, F. Nguyen, P. Bertet, D. Vion, and D. Esteve (2009), Nature Physics **5** (11), 791.

Martinis, J. M., K. B. Cooper, R. McDermott, M. Steffen, M. Ansmann, K. D. Osborn, K. Cicak, S. Oh, D. P. Pappas, R. W. Simmonds, and C. C. Yu (2005), Phys. Rev. Lett. **95**, 210503.

Miano, A., V. Joshi, G. Liu, W. Dai, P. Parakh, L. Frunzio, and M. Devoret (2023), PRX Quantum **4**, 030324.

Miano, A., G. Liu, V. V. Sivak, N. E. Frattini, V. R. Joshi, W. Dai, L. Frunzio, and M. H. Devoret (2022), Applied Physics Letters **120** (18), 184002.

Michael, M. H., M. Silveri, R. T. Brierley, V. V. Albert, J. Salmilehto, L. Jiang, and S. M. Girvin (2016), Phys. Rev. X **6**, 031006.

Milburn, G. J. (1989), Phys. Rev. Lett. **62**, 2124.

Milul, O., B. Guttel, U. Goldblatt, S. Hazanov, L. M. Joshi, D. Chausovsky, N. Kahn, E. Çiftyürek, F. Lafont, and S. Rosenblum (2023), PRX Quantum **4**, 030336.

Minev, Z. K., Z. Leghtas, S. O. Mundhada, L. Christakis, I. M. Pop, and M. H. Devoret (2021), npj Quantum Information **7** (1), 131.

Mirrahimi, M., Z. Leghtas, V. V. Albert, S. Touzard, R. J. Schoelkopf, L. Jiang, and M. H. Devoret (2014), New Journal of Physics **16** (4), 045014.

Mundhada, S. (2019), *Hardware-Efficient Autonomous Quantum Error Correction*, Ph.D. thesis (Yale University).

Müller, C., J. H. Cole, and J. Lisenfeld (2019), Reports on Progress in Physics **82** (12), 124501.

Naik, R. K., N. Leung, S. Chakram, P. Groszkowski, Y. Lu, N. Earnest, D. C. McKay, J. Koch, and D. I. Schuster (2017), Nature Communications **8** (1), 1904.

Ni, Z., S. Li, X. Deng, Y. Cai, L. Zhang, W. Wang, Z.-B. Yang, H. Yu, F. Yan, S. Liu, C.-L. Zou, L. Sun, S.-B. Zheng, Y. Xu, and D. Yu (2023), Nature **616** (7955), 56.

Nielsen, M. A., and I. L. Chuang (2010), *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press).

Noh, K., C. Chamberland, and F. G. Brandão (2022), PRX Quantum **3**, 010315.

O'Brien, J. L. (2007), Science **318** (5856), 1567.

O'Brien et al., (2023), Nature Physics .

Ofek, N., A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf (2016), "Demonstrating quantum error correction that extends the lifetime of quantum information," arXiv:1602.04768 [quant-ph] .

Oriani, A. E., F. Zhao, T. Roy, A. Anferov, K. He, A. Agrawal, R. Banerjee, S. Chakram, and D. I. Schuster (2024), "Niobium coaxial cavities with internal quality factors exceeding 1.5 billion for circuit quantum electrodynamics," arXiv:2403.00286 [quant-ph] .

Ou, Z., C. Hong, and L. Mandel (1987), Optics Communications **63** (2), 118.

Owens, C., A. LaChapelle, B. Saxberg, B. M. Anderson, R. Ma, J. Simon, and D. I. Schuster (2018), Phys. Rev. A **97**, 013818.

Paik, H., D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf (2011), Phys. Rev. Lett. **107**, 240501.

Pfaff, W., C. J. Axline, L. D. Burkhart, U. Vool, P. Reinhold, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf (2017), Nature Physics **13** (9), 882.

Place, A. P. M., L. V. H. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman, G. Cheng, T. Madhavan, H. K. Babla, X. H. Le, Y. Gang, B. Jäck, A. Gyenis, N. Yao, R. J. Cava, N. P. de Leon, and A. A. Houck (2021), Nature Communications **12** (1), 1779.

Planat, L., R. Dassonneville, J. P. Martínez, F. Foroughi, O. Buisson, W. Hasch-Guichard, C. Naud, R. Vijay, K. Murch, and N. Roch (2019), Phys. Rev. Appl. **11**, 034014.

Pozar, D. M. (2012), *Microwave Engineering* (Wiley, Hoboken, NJ).

Purcell, E. M. (1946), in *Proceedings of the American Physical Society*, Vol. 69 (American Physical Society) pp. 11–12.

Puri, S., S. Boutin, and A. Blais (2017), npj Quantum Information **3** (1), 18.

Putterman, H., J. Iverson, Q. Xu, L. Jiang, O. Painter, F. G. S. L. Brandão, and K. Noh (2022), Phys. Rev. Lett. **128**, 110502.

Putterman et al., (2024), "Hardware-efficient quantum error correction using concatenated bosonic qubits," arXiv:2409.13025 [quant-ph] .

Reagor, M., W. Pfaff, C. Axline, R. W. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, K. Chou, M. J. Hatridge, L. Frunzio, M. H. Devoret, L. Jiang, and R. J. Schoelkopf (2016), Phys. Rev. B **94**, 014506.

Reed, M. D., B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf (2010), Applied Physics Letters **96** (20), 203110.

Reinhold, P. (2019), *Controlling Error-Correctable Bosonic Qubits*, Ph.D. thesis (Yale University).

Reinhold, P., S. Rosenblum, W.-L. Ma, L. Frunzio, L. Jiang, and R. J. Schoelkopf (2020), Nature Physics **16** (8), 822.

Richardson, R., and E. Smith (1988), *Experimental Techniques In Condensed Matter Physics At Low Temperatures* (CRC Press).

Romanenko, A., R. Pilipenko, S. Zorzetti, D. Frolov, M. Awida, S. Belomestnykh, S. Posen, and A. Grassellino (2020), Phys. Rev. Appl. **13**, 034032.

Rosenblum, S., P. Reinhold, M. Mirrahimi, L. Jiang, L. Frunzio, and R. J. Schoelkopf (2018), Science **361** (6399), 266.

Sahay, K., J. Jin, J. Claes, J. D. Thompson, and S. Puri (2023), Phys. Rev. X **13**, 041013.

Sank, D., Z. Chen, M. Khezri, J. Kelly, R. Barends, B. Campbell, Y. Chen, B. Chiaro, A. Dunsworth, A. Fowler, E. Jeffrey, E. Lucero, A. Megrant, J. Mutus, M. Neeley, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, T. White, J. Wenner, A. N. Korotkov, and J. M. Martinis (2016), Phys. Rev. Lett. **117**, 190503.

Schoelkopf, R. J., A. A. Clerk, S. M. Girvin, K. W. Lehnert, and M. H. Devoret (2003), "Qubits as spectrometers of quantum noise," in *Quantum Noise in Mesoscopic Physics*, edited by Y. V. Nazarov (Springer Netherlands, Dordrecht) pp. 175–203.

Schreier, J. A., A. A. Houck, J. Koch, D. I. Schuster, B. R. Johnson, J. M. Chow, J. M. Gambetta, J. Majer, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf (2008), Phys. Rev. B **77**, 180502.

Schuster, D. I., A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, J. Majer, B. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf (2007), Nature **445** (7127), 515.

Schwinger, J. (1952), *On Angular Momentum*, Tech. Rep. NYO-3071 (United States Atomic Energy Commission).

Shapiro, A. E. (1989), Notes Rec. R. Soc. Lond. **43**, 223.

Shillito, R., A. Petrescu, J. Cohen, J. Beall, M. Hauru, M. Ganahl, A. G. Lewis, G. Vidal, and A. Blais (2022), Phys. Rev. Appl. **18**, 034031.

Shor, P. (1994), in *Proceedings 35th Annual Symposium on Foundations of Computer Science*, pp. 124–134.

da Silva, M. P., O. Landon-Cardinal, and D. Poulin (2011), Phys. Rev. Lett. **107**, 210404.

Sirois, A. J., M. A. Castellanos-Beltran, M. P. DeFeo, L. Ranzani, F. Lecocq, R. W. Simmonds, J. D. Teufel, and J. Aumentado (2015), Applied Physics Letters **106** (17), 172603, https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/1.4919759/14314706/172603_1_online.pdf .

Sivak, V., N. Frattini, V. Joshi, A. Lingenfelter, S. Shankar, and M. Devoret (2019), Phys. Rev. Appl. **11**, 054060.

Sivak, V. V., A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. L. Brock, A. Z. Ding, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret (2023), Nature **616** (7955), 50.

Sivak, V. V., S. Shankar, G. Liu, J. Aumentado, and M. H. Devoret (2020), Phys. Rev. Appl. **13**, 024014.

Stace, T. M., S. D. Barrett, and A. C. Doherty (2009), Phys. Rev. Lett. **102**, 200501.

Stehlik, J., D. M. Zajac, D. L. Underwood, T. Phung, J. Blair, S. Carnevale, D. Klaus, G. A. Keefe, A. Carniol, M. Kumph, M. Steffen, and O. E. Dial (2021), Phys. Rev. Lett. **127**, 080505.

Sun, L., A. Petrenko, Z. Leghtas, B. Vlastakis, G. Kirchmair, K. M. Sliwa, A. Narla, M. Hatridge, S. Shankar, J. Blumoff, L. Frunzio, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf (2014), Nature **511** (7510), 444.

Teoh, J. (2023), *Error Detection in Bosonic Circuit Quantum Electrodynamics*, Ph.D. thesis (Yale University).

Teoh, J. D., P. Winkel, H. K. Babla, B. J. Chapman, J. Claes, S. J. de Graaf, J. W. Garmon, W. D. Kalfus, Y. Lu, A. Maiti, *et al.* (2023), Proceedings of the National Academy of Sciences **120** (41), e2221736120.

Touzard, S., A. Grimm, Z. Leghtas, S. O. Mundhada, P. Reinhold, C. Axline, M. Reagor, K. Chou, J. Blumoff, K. M. Sliwa, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret (2018), Phys. Rev. X **8**, 021005.

Tripathi, V., H. Chen, M. Khezri, K.-W. Yip, E. Levenson-Falk, and D. A. Lidar (2022), Phys. Rev. Appl. **18**, 024068.

Tsunoda, T., J. D. Teoh, W. D. Kalfus, S. J. de Graaf, B. J. Chapman, J. C. Curtis, N. Thakur, S. M. Girvin, and R. J. Schoelkopf (2023), PRX Quantum **4**, 020354.

Tuckett, D. K., S. D. Bartlett, and S. T. Flammia (2018), Phys. Rev. Lett. **120**, 050505.

Tuckett, D. K., S. D. Bartlett, S. T. Flammia, and B. J. Brown (2020), Phys. Rev. Lett. **124**, 130501.

Tuckett, D. K., A. S. Darmawan, C. T. Chubb, S. Bravyi, S. D. Bartlett, and S. T. Flammia (2019), Phys. Rev. X **9**, 041031.

Van Harlingen, D. J., T. L. Robertson, B. L. T. Plourde, P. A. Reichardt, T. A. Crane, and J. Clarke (2004), Phys. Rev. B **70**, 064517.

Varbanov, B. M., F. Battistel, B. M. Tarasinski, V. P. Ostroukh, T. E. O'Brien, L. DiCarlo, and B. M. Terhal (2020), npj Quantum Information **6** (1), 102.

Venkatraman, J., X. Xiao, R. G. Cortiñas, and M. H. Devoret (2022), "On the static effective lindbladian of the squeezed kerr oscillator," arXiv:2209.11193 [quant-ph] .

Vlastakis, B. (2015), *Controlling Coherent State Superpositions with Superconducting Circuits*, Ph.D. thesis (Yale University).

Vlastakis, B., A. Petrenko, N. Ofek, L. Sun, Z. Leghtas, K. Sliwa, Y. Liu, M. Hatridge, J. Blumoff, L. Frunzio, M. Mirrahimi, L. Jiang, M. H. Devoret, and R. J. Schoelkopf (2015), Nature Communications **6** (1), 8970.

Vool, U., and M. Devoret (2017), International Journal of Circuit Theory and Applications **45** (7), 897, https://onlinelibrary.wiley.com/doi/pdf/10.1002/cta.2359 .

Walls, D., and G. J. Milburn (2008), "Interaction of radiation with atoms," in *Quantum Optics*, edited by D. Walls and G. J. Milburn (Springer Berlin Heidelberg, Berlin, Heidelberg) pp. 197–211.

Wang, C., C. Axline, Y. Y. Gao, T. Brecht, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf (2015), Applied Physics Letters **107** (16), 162601, https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/1.4934486/14470271/162601_1_online.pdf .

Wang, C., Y. Y. Gao, P. Reinhold, R. W. Heeres, N. Ofek, K. Chou, C. Axline, M. Reagor, J. Blumoff, K. M. Sliwa, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf (2016), Science **352** (6289), 1087.

Wang, C. S., J. C. Curtis, B. J. Lester, Y. Zhang, Y. Y. Gao, J. Freeze, V. S. Batista, P. H. Vaccaro, I. L. Chuang, L. Frunzio, L. Jiang, S. M. Girvin, and R. J. Schoelkopf (2020), Phys. Rev. X **10**, 021060.

Wang, D. S., A. G. Fowler, A. M. Stephens, and L. C. L. Hollenberg (2010), Quantum Info. Comput. **10** (5), 456–469.

Weedbrook, C., S. Pirandola, R. García-Patrón, N. J. Cerf, T. C. Ralph, J. H. Shapiro, and S. Lloyd (2012), Rev. Mod. Phys. **84**, 621.

Weiss, D. K., S. Puri, and S. M. Girvin (2024), "Quantum random access memory architectures using superconducting cavities," arXiv:2310.08288 [quant-ph] .

Wigner, E. (1931), *Gruppentheorie und ihre Anwendung auf die Quantenmechanik der Atomspektren* (Vieweg+Teubner Verlag Wiesbaden).

Wigner, E. (1932), Phys. Rev. **40**, 749.

Winkel, P., I. Takmakov, D. Rieger, L. Planat, W. Hasch-Guichard, L. Grünhaupt, N. Maleeva, F. Foroughi, F. Henriques, K. Borisov, J. Ferrero, A. V. Ustinov, W. Wernsdorfer, N. Roch, and I. M. Pop (2020), Phys. Rev. Appl. **13**, 024015.

Wu, Y., S. Kolkowitz, S. Puri, and J. D. Thompson (2022), Nature Communications **13** (1), 4657.

Xiao, X., J. Venkatraman, R. G. Cortiñas, S. Chowdhury, and M. H. Devoret (2023), "A diagrammatic method to compute the effective hamiltonian of driven nonlinear oscillators," arXiv:2304.13656 [quant-ph] .

Xu, Q., N. Mannucci, A. Seif, A. Kubica, S. T. Flammia, and L. Jiang (2022), "Tailored xzzx codes for biased noise," arXiv:2203.16486 [quant-ph] .

Xu, Y., Y. Ma, W. Cai, X. Mu, W. Dai, W. Wang, L. Hu, X. Li, J. Han, H. Wang, Y. P. Song, Z.-B. Yang, S.-B. Zheng, and L. Sun (2020), Phys. Rev. Lett. **124**, 120501.

Yamamoto, T., K. Inomata, M. Watanabe, K. Matsuba, T. Miyazaki, W. D. Oliver, Y. Nakamura, and J. S. Tsai (2008), Applied Physics Letters **93** (4), 042510.

Yan, F., S. Gustavsson, J. Bylander, X. Jin, F. Yoshihara, D. G. Cory, Y. Nakamura, T. P. Orlando, and W. D. Oliver (2013), Nature Communications **4** (1), 2337.

Yan, F., S. Gustavsson, A. Kamal, J. Birenbaum, A. P. Sears, D. Hover, T. J. Gudmundsen, D. Rosenberg, G. Samach, S. Weber, J. L. Yoder, T. P. Orlando, J. Clarke, A. J. Kerman, and W. D. Oliver (2016), Nature Communications **7** (1), 12964.

Yan, F., P. Krantz, Y. Sung, M. Kjaergaard, D. L. Campbell, T. P. Orlando, S. Gustavsson, and W. D. Oliver (2018), Phys. Rev. Appl. **10**, 054062.

Young, T. (1804), Phil. Trans. R. Soc. **94**, 1.

Yurke, B., L. R. Corruccini, P. G. Kaminsky, L. W. Rupp, A. D. Smith, A. H. Silver, R. W. Simon, and E. A. Whittaker (1989), Phys. Rev. A **39**, 2519.

Yurke, B., S. L. McCall, and J. R. Klauder (1986), Phys. Rev. A **33**, 4033.

Zakka-Bajjani, E., F. Nguyen, M. Lee, L. R. Vale, R. W. Simmonds, and J. Aumentado (2011), Nature Physics **7** (8), 599.

Zehnder, L. (1891), Zeitschrift für Instrumentenkunde **11**, 275.

Zhang, Y., B. J. Lester, Y. Y. Gao, L. Jiang, R. J. Schoelkopf, and S. M. Girvin (2019), Phys. Rev. A **99**, 012314.

Zhou, C. (2024), *Superconducting Quantum Routers, Modules, Gates, and Measurements Based on Charge-pumped Parametric Interactions*, Ph.D. thesis (University of Pittsburgh).

Zhou, C., P. Lu, M. Praquin, T.-C. Chien, R. Kaufman, X. Cao, M. Xia, R. S. K. Mong, W. Pfaff, D. Pekker, and M. Hatridge (2023), npj Quantum Information **9** (1), 54.

Zimmerman, J. E. (1971), Journal of Applied Physics **42** (11), 4483.

Appendix A

Experimental system parameters

Here I list the relevant system parameters for the different experiments in this thesis, along with explanations where parameters have changed.

Beamsplitter and SNAIL characterization

The parameters measured contemporaneously with the experiments in Chapter 4 are shown in Table A.1, when the SNAIL is at a flux point $\varphi_{\text{ext}}/2\pi = 0.35$. At this point, the SNAIL excited state population was measured to be $P_e = 2.42 \pm 0.47\%$.

		Alice	**Bob**
Transmon $e \to g$ relaxation time	T_{1q}	$127.2 \pm 1.9\ \mu s$	$57.1 \pm 0.6\ \mu s$
Transmon T_2	T_{2q}	$114.4 \pm 2.9\ \mu s$	$56.8 \pm 1.5\ \mu s$
Transmon dephasing time	$T_{\phi q}$	$208 \pm 10\ \mu s$	$113 \pm 6\ \mu s$
Transmon thermal $\lvert e\rangle$ population	P_e	$0.70 \pm 0.14\ \%$	$1.02 \pm 0.20\ \%$
Transmon readout duration	τ_{RO}	$2.1\ \mu$s	$2.1\ \mu$s
Transmon anharmonicity	$\alpha_t/2\pi$	-181.3 MHz	-184.3 MHz
Transmon-oscillator dispersive shift	$\chi/2\pi$	-0.766 MHz	-1.104 MHz
Oscillator relaxation time	T_{1c}	$482 \pm 16\ \mu s$	$91 \pm 4\ \mu s$
Oscillator dephasing time	$T_{\phi c}$	$2010 \pm 220\ \mu s$	$840 \pm 200\ \mu s$
Oscillator thermal occupation	n_{thermal}	$0.96 \pm 0.19\ \%$	$0.11 \pm 0.02\ \%$

Table A.1: **Measured experimental parameters for beamsplitter and SNAIL characterization.** Values as in Chapman *et al.* (2023).

cSWAP characterization

The characterization of the single cSWAP in Figures 5.3, 5.4 and 5.5 was performed with the same parameters as in Table A.1. Between those experiments and the repeated cSWAP experiment in Fig. 5.6, some of the coherences drifted. On re-measuring, the coherence times which changed by more than one standard deviation were $T_{1q}^{(\text{Alice})} = 112 \pm 8$ μs, $T_{1q}^{(\text{Bob})} = 45.6 \pm 9$ μs, $T_{2q}^{(\text{Alice})} = 97 \pm 10$ μs and $T_{\phi c}^{(\text{Alice})} = 1510 \pm 90$ μs.

Erasure check

The parameters measured simultaneously with the erasure check data are shown in Table A.2, at a slightly different external flux of $\varphi_{\text{ext}}/2\pi = 0.334$. The measured SNAIL parameters at this point were $\omega_c/2\pi = 5.192$ GHz, $T_1 = 69 \pm 2$ μs, $T_{2R} = 3.0 \pm 0.2$ μs, $T_{2E} = 14.4 \pm 0.6$ μs and $P_e = 4.1 \pm 0.6$ %.

In between the cSWAP experiment and the erasure check experiments, one change made was an increase in the coupling pin length to increase the linewidth of Bob's readout resonator from $\kappa_{\text{r}}/2\pi = 0.89 \pm 0.05$ MHz to $\kappa_{\text{r}}/2\pi = 1.77 \pm 0.04$ MHz. This was done to achieve faster single-shot transmon readout, with the goal of fewer false negative and idling errors. The most notable change as a result of this was a reduction in *Alice's* oscillator T_{1c}, a mode that finite element simulations suggested should have been extremely undercoupled to this port.

		Alice	**Bob**
Transmon $e \to g$ relaxation time	T_{1q}	123.4 ± 0.9 μs	42.4 ± 0.4 μs
Transmon T_2 (Ramsey)	T_{2q}	31.6 ± 0.9 μs	33.1 ± 0.9 μs
Transmon T_2 (Echo)	T_{2Eq}	40.5 ± 0.7 μs	63.2 ± 0.8 μs
Transmon thermal $\lvert e\rangle$ population	P_e	0.23 ± 0.06 %	0.52 ± 0.07 %
Transmon readout duration	τ_{RO}	1.648 μs	1.648 μs
Transmon-oscillator dispersive shift	$\chi/2\pi$	-0.777 MHz	-1.066 MHz
Oscillator relaxation time	T_{1c}	347 ± 2 μs	108.5 ± 0.6 μs
Oscillator dephasing time	$T_{\phi c}$	1210 ± 170 μs	690 ± 120 μs
Oscillator thermal occupation	n_{thermal}	3.82 ± 0.10 %	0.53 ± 0.03 %

Table A.2: **Measured experimental parameters for erasure check characterization.** Values as in de Graaf *et al.* (2025).

Appendix B

Extracting normalized drive amplitude $|\xi|$ from the impedance of the coupler embedding matrix $\mathbf{Z}$

The derivation is reproduced with permission from Chapman *et al.* (2023).

The normalized pump amplitude ξ that we can deliver for a given input power can be calculated from the impedance matrix of the embedding network of the SNAIL. Fig. B.1(a) shows a simplified schematic of the microwave drive on the SNAIL coupler. A microwave generator sourcing a voltage V_S generates a voltage V_1 at the coupling pin inside the experimental package, which in turn generates a voltage V_2 across the SNAIL. Their ratio is determined by the properties of the embedding network of the SNAIL.

If we consider just the linear part of the circuit, we can treat the entire package including the SNAIL itself as an impedance matrix $\mathbf{Z}$, as shown in Fig. B.1(b). This can be used to relate the voltage across the SNAIL to the current delivered from the source,

$$Z_{21} = \frac{V_2}{I_1}. \tag{B.1}$$

Meanwhile, the impedance presented to the source by the package is Z_{11} (Fig. B.1(c)).

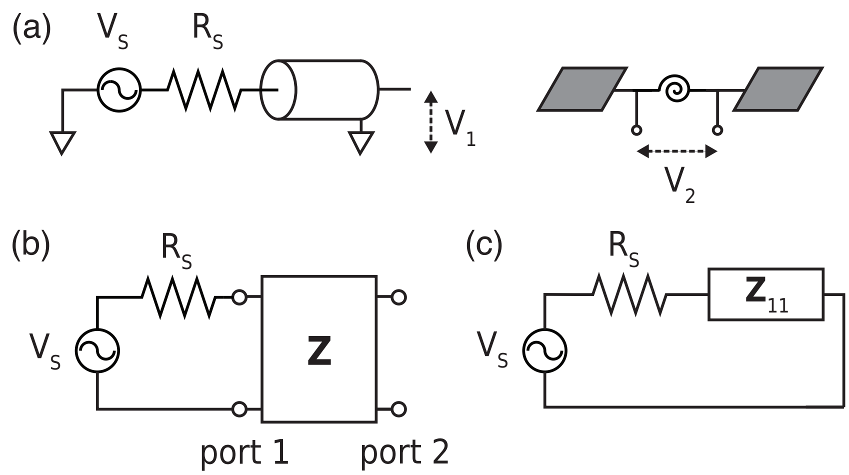

Figure B.1: **Network picture of SNAIL embedding network.** (a) A microwave source with a voltage V_S and a resistance R_S connected via a transmission line to the drive pin of the package. It generates a voltage V_1 at the drive pin (port 1) and a voltage V_2 across the SNAIL (port 2). (b) The same circuit with the package, including the SNAIL itself, expressed as a two-port network with an impedance matrix $\mathbf{Z}$. (c) When port 2 is open-circuited, we can treat the package as an impedance Z_{11} in series with the microwave source.. Figure modified with permission from Chapman *et al.* (2023).

Current conservation indicates that the source voltage V_S is related to the current I_1 by

$$V_S = (R_S + Z_{11})\, I_1. \tag{B.2}$$

We can eliminate I_1 from the equations to obtain

$$\frac{V_2}{V_S} = \frac{Z_{21}}{R_S + Z_{11}}. \tag{B.3}$$

Microwave generators typically display the RMS power they would deliver to a matched load,

$$P = \frac{|V_{\mathsf{RMS}}|^2}{2R_S} = \frac{|V_S|^2}{4R_S}, \tag{B.4}$$

allowing us to reach an expression for V_2 in terms of source power

$$|V_2| = \left|\frac{Z_{21}}{Z_{11} + R_S}\right| \sqrt{4PR_S}. \tag{B.5}$$

Assuming a monochromatic source at ω_p, the current passing through the SNAIL can be found by dividing V_2 by the impedance of the SNAIL at the pump frequency, $i\omega_p L_s$,

$$I(t) = \mathsf{Re}\left(\frac{V_2}{i\omega_p L_s} e^{i\omega_p t}\right). \tag{B.6}$$

This couples to the flux of the coupler mode in the Hamiltonian

$$\hat{H}_d = I(t)\hat{\Phi} = \mathsf{Re}\left(\frac{V_2}{i\omega_p L_s} e^{i\omega_p t}\right) \Phi_c^{\mathsf{ZPF}} \left(\hat{c} + \hat{c}^\dagger\right). \tag{B.7}$$

By comparing this to Eq., we can write

$$|\epsilon| = \frac{\Phi_c^{\mathsf{ZPF}} |V_2|}{2\hbar\omega_p L_s}. \tag{B.8}$$

From here, we can use the expression for ξ from Eq. to finally get

$$|\xi| = \left|\frac{Z_{21}}{Z_{11} + R_S}\right| \frac{\sqrt{P R_S}\,\Phi_c^{\mathsf{ZPF}}}{\hbar\omega_p L_s\left(\omega_p - \omega_c\right)}. \tag{B.9}$$

Appendix C

Logical state tomography of cat states with finite α

In Ch. 5.3.2, we discussed a method (initially described in (Vlastakis *et al.*, 2015; Wang *et al.*, 2016)) for performing logical state tomography of encoded cat states by probing only a handful of values in the Wigner tomogram of the oscillator mode(s). The relationships between the two, provided in Eq. 5.24, are exact in the limit $|\alpha| \to \infty$, however in practice, the use of a finite $|\alpha|$ will result in some corrections. Here, I provide the exact expectation values we would obtain when acting on single-mode cat states and two-mode cat-in-two-boxes states.

C.1 Single-mode cat states

For an ideal finite-α single-mode cat state

$$\left|C_\alpha^+\right\rangle = \frac{|\alpha\rangle + |-\alpha\rangle}{\sqrt{2\left(1+e^{-2|\alpha|^2}\right)}}, \tag{C.1}$$

the values extracted from Wigner tomography would yield

$$\langle I\rangle \approx W(+\alpha)+W(-\alpha)=\frac{1+2e^{-2|\alpha|^2}+e^{-8|\alpha|^2}}{1+e^{-2|\alpha|^2}} \tag{C.2}$$

$$\langle X\rangle \approx W(0)=1 \tag{C.3}$$

$$\langle Y\rangle \approx W\left(\frac{j\pi}{8\alpha}\right)=\frac{e^{-2|\alpha|^2}e^{-\frac{\pi^2}{32|\alpha|^2}}}{1+e^{-2|\alpha|^2}} \tag{C.4}$$

$$\langle Z\rangle \approx W(+\alpha)-W(-\alpha)=0 \tag{C.5}$$

C.2 Two-mode cat states

For an ideal finite-α cat-in-two-boxes state

$$\left|\Psi^+\right\rangle=\frac{|\alpha,-\alpha\rangle+|-\alpha,\alpha\rangle}{\sqrt{2\left(1+e^{-4|\alpha|^2}\right)}}, \tag{C.6}$$

the values extracted from Wigner tomography yield

$$\begin{aligned}\langle II\rangle &\approx W(+\alpha,+\alpha)+W(+\alpha,-\alpha)+W(-\alpha,+\alpha)+W(-\alpha,-\alpha)\\ &=\frac{1+4e^{-4|\alpha|^2}+2e^{-8|\alpha|^2}+e^{-16|\alpha|^2}}{1+e^{-4|\alpha|^2}}\end{aligned} \tag{C.7}$$

$$\langle IX\rangle \approx W(+\alpha,0)+W(-\alpha,0)=\frac{e^{-2|\alpha|^2}\left(3+e^{-8|\alpha|^2}\right)}{1+e^{-4|\alpha|^2}} \tag{C.8}$$

$$\langle IY\rangle \approx W\left(+\alpha,\frac{j\pi}{8\alpha}\right)+W\left(-\alpha,\frac{j\pi}{8\alpha}\right)=\frac{e^{-2|\alpha|^2}e^{-\frac{\pi^2}{32|\alpha|^2}}\left(1+e^{-8|\alpha|^2}\right)}{1+e^{-4|\alpha|^2}} \tag{C.9}$$

$$\langle IZ\rangle \approx W(+\alpha,+\alpha)-W(+\alpha,-\alpha)+W(-\alpha,+\alpha)-W(-\alpha,-\alpha)=0 \tag{C.10}$$

$$\langle XI\rangle \approx W(0,+\alpha)+W(0,-\alpha)=\frac{e^{-2|\alpha|^2}\left(3+e^{-8|\alpha|^2}\right)}{1+e^{-4|\alpha|^2}} \tag{C.11}$$

$$\langle ZI\rangle \approx W(+\alpha,+\alpha)+W(+\alpha,-\alpha)-W(-\alpha,+\alpha),-W(-\alpha,-\alpha)=0 \tag{C.12}$$

$$\langle ZX\rangle \approx W(+\alpha,0)-W(-\alpha,0)=0 \tag{C.13}$$

$$\langle ZY\rangle \approx W(+\alpha,\frac{j\pi}{8\alpha})-W(-\alpha,\frac{j\pi}{8\alpha})=0 \tag{C.14}$$

$$\langle ZZ\rangle \approx W(+\alpha,+\alpha)-W(+\alpha,-\alpha)-W(-\alpha,+\alpha)+W(-\alpha,-\alpha) \tag{C.15}$$

$$=\frac{2e^{-8|\alpha|^2}-e^{-16|\alpha|^2}-1}{1+e^{-4|\alpha|^2}} \tag{C.16}$$

For $|\alpha|\gg 1$, the measured values for II, XX, YY and ZZ tend to $+1$, $+1$, $+1$ and -1, respectively, while all the off-diagonal elements tend to 0. These are the values one expects for an ideal Bell state.

In Chapter 5, we consider a cat-in-two-boxes state with $|\alpha|=\sqrt{2}$. For an ideal (i.e. without any distortion or decoherence) cat-in-two-boxes state of this size, the diagonal elements would be measured to be $+1.0010$, $+1.0000$, $+0.7346$ and -0.9997, respectively. The YY measurement (unlike the others) suffers markedly from the finite-α effects at this size, and so we must correct for this. Indeed, for the error between the measured YY value and $+1$ to drop below ϵ requires a cat with $|\alpha|^2\approx\pi^2/16\epsilon$ photons.