Controlling and Protecting Quantum Information in Superconducting Oscillators

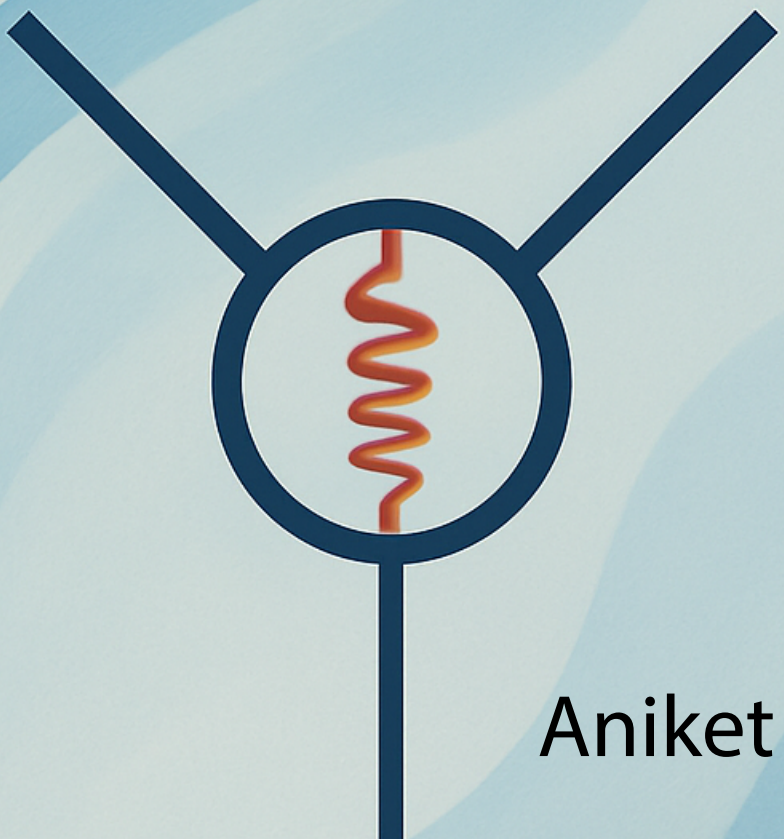

Aniket Maiti

Abstract

Controlling and Protecting Quantum Information in Superconducting Oscillators

Aniket Maiti

2025

Modern quantum experiments allow the precise manipulation and measurement of many-body quantum states, pushing quantum mechanics from a testable theory to a utilizable technology. The central promise of these experiments is to process quantum information for exponential advantages in computing, sensing, and communication. An interesting way to achieve such a processor is to manipulate quantum information stored in the continuous-variable (bosonic) phase space of electromagnetic radiation. Since photons in free space do not interact, such an approach necessarily requires the introduction of nonlinearity through strong light-matter couplings. However, since all matter is lossy, this inevitably introduces a trade-off between the speed of control and the inherited decoherence of the 'light.'

This thesis explores the control of microwave radiation trapped in superconducting oscillators through interactions with Josephson junction-based nonlinearities. I first demonstrate novel ways to exchange single photons between two detuned oscillators through carefully constructed driven nonlinearities, achieving orders of magnitude higher fidelity than previously possible. Using such protected driven interactions, I then implement a bosonic control architecture and that is protected from any nonlinearity when idle, and implements clean photon-exchanges when driven. Finally, I introduce ways to utilize such photon-exchanges to dynamically hybridize light and matter, in a way that regains universal control without reintroducing the inherited decoherence. Together, this thesis provides a promising path toward error-resilient bosonic quantum processors.

Controlling and Protecting Quantum Information in Superconducting Oscillators

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

by
Aniket Maiti

Dissertation Director: Robert J. Schoelkopf

May 2025

Acknowledgments

Graduate school is an interesting experiment – it is of course, a curiosity-driven degree, where one chooses to forego what perhaps may have been more financially lucrative opportunities, to extend human knowledge in a relatively niche topic and learn a little more about how the universe works. Yet it is precisely this environment that brings some extremely interesting and like-minded people together. Practically every person I have worked with at the Yale Quantum Institute has brought to me a unique perspective on life and why they find what we do interesting. This has been true even in the broader context of the Yale Physics department, where I have, over the last six years, found community, friends in various sub-fields of Physics, and been part of efforts that aim to protect that community beyond just their involvement in physics research.

Perhaps the most impact I have felt within these past six years has been from my two closest colleagues: **Yao Lu** and **John Garmon**. Yao, who now leads his own group at Fermilab, was my primary mentor – teaching me why parametric operations are interesting, how to think about driven circuits, and the innate requirement for being confident in your analysis and numerical techniques. Practically every idea in this thesis was started through a discussion with Yao, in particular, providing significant design input on every successful device that you will find in these pages. When we first started getting useful data on the differentially driven SQUID, the primary device on which many of our later concepts and devices were based, we were joined by John. John, while initially playing the role of a mentee, quickly became an equal collaborator, with a particular enthusiasm for diving into FPGA code and making experiments work. He served as an extremely important check to the many ideas I would constantly try to come up with – asking me to slow down, critically think through things over a blackboard, and often proving me wrong or significantly changing how I think about certain concepts, like balanced mixers. These projects could not have succeeded without constant effort from John's end, both scientific

(including learning fabrication and developing a process for the LINC from scratch), and through his friendship and humor. I have no doubt that the expanse of his knowledge and his ability to think carefully and consider concepts through multiple lenses will make him an excellent scientist.

I had the fortunate experience during my time at YQI of directly collaborating and working with so many others, all of whom taught me so many new things. **Chao Zhou** joined RSL towards the end of my PhD, and quite fortunately, was placed in the same office as me. I was soon able to convince him that some of my crazy ideas might work, and Chao became a bedrock for translating the LINC project into reality. He brought with him an immense experience in designing and running superconducting experiments, and a quick and eager desire to learn. Soon after I had described the problem to him, proposing an outlandish combination of implementing a new coupler in a new architecture using a new control scheme with control electronics that were still under development, he got to work. About six months later, we had our first data on a working LINC device implementing dynamic light-matter coupling using an RF-SoC. I have no doubt he will apply the LINC and other ideas in the most wonderful ways, and I would be most fortunate to work with him again. There were many others who the LINC project could not have succeeded without – in particular **Gautham Umasankar**, who joined us right before I left the group and soon helped fabricate and take some of the first measurements on the LINC devices, **Yanhao Wang**, who taught us the basics of fabrication, **Suhas**, who provided support through his precise and extensive design and fabrication intuition, and **Billy Kalfus** who was a one-man production team that designed our first RF-SoC firmware, and an absolute joy as an office-mate. Parallel to, and perhaps initially motivated by, the LINC project, was an effort to create high-quality compact inductances using disordered superconductors (in particular, grAl). This effort was led by **Patrick Winkel**, who in addition to helping me understand all things dual-rail and LINC, was one of the closest friends I had in grad

school. I spent many coffee breaks and weekend hikes with Patrick, talking about physics, life as a physicist, and the world in general. The grAl project was taken up by **Vishakha Gupta**, **Neel Thakur**, and **Jenny Giampalmo**, all of whom were both close friends, and collaborators – I am excited to see where this platform goes!

Surprisingly, despite being in an experimental group, a majority of my projects and papers were theoretical. This is a reflection of the beautiful learning environment that YQI provided to me; we did not simply categorize ourselves as experimentalists with theory collaborators, we dove deep and worked out both the theory and the everyday experimental issues ourselves. The average person in YQI could equally well hold a conversation about measurement theory and dissipatively protected manifold and debug a ground loop on a fridge, with about equal ease. One of the theory projects I found myself working on quite early on, was to ask the question "given my ability to universally control an oscillator, what pulse should I play to stabilize any particular code"? This idea initially sprang out of conversations with **Nathaniel Cottet**, but soon became the ground rock of a multi-year collaboration with **Jahan Claes**, whose astounding theoretical ability made this project concrete and provided real answers. Through our weekly conversations, we went from two people who barely knew anything about bosonic codes, to coming up with a universal framework that protects them and a fairly deep understanding of what made each popular code work. We were helped by **Alec Eickbusch**, a bosonic code expert who inspired much of my interest in bosonic control and error correction, and later **Kaavya Sahay** and **Harsh Babla**, who provided the much-needed support required to bring this project to completion and translate it into actual optimized pulses. I am hopeful that one day, this framework can be tested on a bosonic device (perhaps the LINC), and provide the first universal demonstration of stabilizing *arbitrary* codes beyond break-even. My other main theory collaboration was in developing the framework of dynamic light-matter interactions. Here, **Takahiro Tsunoda** helped answer critical questions about what this dynamic dressing

actually looked like and how fast it could occur (in addition to providing excellent soccer opposition), and **Danny Weiss** provided critical theory support and helped optimize novel protected pulses using these control techniques. There were also those that I did not ever officially collaborate with, but learnt so much from – this list includes **Nick Frattini**, **Chris Wang**, **Jaya Venkatraman**, **Akshay Koottandavida** (also a close friend and ex-roommate), **Vidul Joshi**, **Shraddha Singh**, **Alex Read**, and many others.

An interesting mix of things is required to make a place like YQI, and prime among them is the multi-decade close collaboration among the research groups of three brilliant scientists – Rob, Michel and Steve. I was lucky enough to have all three be a part of my graduate school experience and my dissertation committee. In particular, of course, I am indebted to **Rob Schoelkopf**, who, despite being one of the largest figures in the field, made time for all of us young graduate students and personally taught me so much as an advisor. I still remember the first day I walked into Becton, and was pulled into an immediate deep dive into debugging ground loops on a dilution fridge, one of Rob's favorite topics in lab. He eventually went on to teach me both concepts within physics – such as what matters in a high-Q experimental device, and a phenomenal practical intuition for microwave design – but also broader things like a general approach to research, being critical and precise while simultaneously learning how to prioritize and balance curiosity and practicality. As an older graduate student I also eventually became friends with Rob outside of our discussions of physics, no doubt through the many conferences we attended together, and the many pool parties he so graciously hosted. **Steve Girvin** perhaps has the most astounding amount of intuition for quantum mechanics out of all the people I have known. The most proud and happy that I have been with the physics I have conducted is when I have been able to impress Steve Girvin – a feat that even Rob admits is difficult. He was instrumental in my primary theory project, that of finding a universal theory for bosonic codes, and I will considerably miss being able to walk into the office of some-

one of his capability and just clarify every physics concept. **Michel Devoret**, one of the undoubted fathers of this field, was another such figure, who I had extremely interesting conversations with, and learned so much from. I'm very lucky to continue working and learning from him during my next phase at Google Quantum, and look forward to the many wild ideas he will surely get me thinking about. Finally, **Shruti Puri** has always provided for me the piercing questions that I have often asked others, and that has helped keep my (and many other's) physics practical and precise. Her group has come an incredible way in a very short time, and I have been lucky to watch it grow from scratch and am excited to see where it will go in the future!

While the people above shaped me scientifically, a scientist is not just the research they do. Much of my goals, my life in America, and my evolution from a naive fresh-grad to a person with strong interests and opinions occurred in graduate school. My primary support through all of this was my family, in particular my parents and Bumba, and my partner Hannah. It is very difficult to express the absolute bedrock a family can provide, and I am lucky to have one that has remained so to this day. Everything I do, I do for Bumba. This section is not nearly long enough to contain my love for them, but I hope I can continue to support them as I grow into a more capable adult. **Hannah Bossi** was one of the first people I met in the physics program. We started out as close friends going on frequent adventures, and slowly that has evolved into a close companionship that I cannot imagine my life without. She has taught me many of the things specific to adult life in America, as has her family, who have many a time hosted me and provided a home away from home to rest and celebrate holidays in.

Finally, grad school came with many close friends, who made the last seven years feel like a breeze. My memories at Yale will always be characterized by the frequent pizza nights at **Maggie Pavlovich** and **Patrick Hulin's** house, usually attended by Patrick, Stijn and Jacob. **Stijn de Graaf** in particular was one of my closest friends at Yale, and for

good reason – we started this experience together, in the same lab, with extremely similar projects and interests. I'll always cherish our many lunches, coffee breaks, bike rides, soccer games and conversations together – and hopefully make some more memories, in this next phase at Google! Maggie was another incredibly close friend. It is difficult to describe the depth of our friendship, but it was always incredibly comfortable – even when we argued through practically every political debate we could lay our hands on. **Jacob Curtis**, and previously **Yanni Dahmani**, brought much-needed life to these parties, and always some delicious food. I will miss this group dearly. Similarly, before my life in Applied Physics, I made some close friends in the physics department. Prime among them was **Ridge Liu**, who remains one of my closest friends, and continues be the most interesting combination of flaky yet incredibly dependable when you need him. This friend group included **James Chow** and **Giacomo Scanavini**, and my many shenanigans with them defined my early years of graduate school. I hope we keep in touch, and I will always be missing all of you!

Thank you all, and most importantly to you, dear reader, who finds interest in my work!

Contents

List of Figures

List of Tables

Chapter 1

Modern Quantum Experiments

What crosses your mind when you hear the word *quantum*? If you asked me this question when I was starting out at Yale, I would probably tell you about classically non-intuitive results like wave-particle duality in the double-slit experiment, or the discrete quantization of spins in the Stern-Gerlach experiment. These concepts are still how quantum mechanics is introduced in many undergraduate textbooks [1], but they are motivated by an era where quantum phenomena were just being discovered and seen as unintuitive. Much of quantum physics in that era involved thought experiments that explored how 'weird' entanglement was, like the EPR paradox or Bell tests, or analytic solutions to tractable quantum systems like perfect harmonic oscillators or quantum spins. However, even if these theories were difficult to believe, they were undeniably accurate – particularly those describing the interactions of electrons and light, or Quantum Electrodynamics (QED) [2, 3]. QED provided some of the most accurate predictions of Nature that humans have ever made, like the magnetic moment of an electron, which matches experimental measurements to 12 decimal places [4]. It also describes the behavior of many extremely successful technologies that require quantum mechanics to predict their parameters, even though their dynamics and effects in the real world are completely classical. These include the laser and the Global Positioning System (GPS) based on the atomic clock, both of which rely on

the fine optical transitions of atomic states that are set by their quantum interactions with an electromagnetic environment. It also describes the transport of electrons in semiconductors, forming the basis of transistors and diodes that define modern integrated circuits, including the laptop computer on which this thesis was written.

Today's quantum science has moved leaps and bounds beyond these initial demonstrations and their effectively classical outputs. Where once Bell tests were seen as a test of quantum mechanics, modern experiments use Bell measurements to benchmark their experimental setup [5]. Instead of observing interference through a double-slit setup, we now precisely manipulate the quantum state of hundreds of atoms and interfere them to simulate complex phenomena like long-range topological effects [6, 7], or lattice gauge theories [8, 9]. This leap in experimental ability has been driven by the promise of significant advantages that might come from directly controlling information stored in quantum states, a realization that is often termed the second quantum revolution. Such 'quantum information' is encoded in the precise amplitude and phases of quantum superpositions, making it simultaneously fragile and powerful. Its fragility arises because any quantum system that can be manipulated and measured by a user, can also have its quantum state collapsed by the environment at approximately the same rate. Yet this same sensitivity enables novel abilities beyond those of any classical device, for example, sensing beyond the standard limits set by a thermal environment and communication channels whose security is guaranteed by the laws of physics.

Perhaps the most important promise of being able to precisely manipulate and measure the state of a quantum many-body ('qubit') system is to utilize it for computation [10]. This is made concrete by quantum algorithms, which take either a classical or quantum input, perform a dance of destructive and constructive interference, and result in either a quantum output or a probability distribution of classical outputs [11]. In general, these computations are reversible (since Hamiltonian evolution is unitary) and so far are believed

to provide an exponential advantage in a select few classically hard problems, potentially breaking the Extended Church-Turing hypothesis [12]. This forms an important statement about the computational ability of the universe, and in particular, it means that certain problems that we think are hard are only so because of the way we have implemented our computer. Importantly, this statement is not strictly proven [13, 14], and proving a verifiable advantage (for example, on an NP problem) would have profound implications on computational complexity theory (e.g., proving $P \subset BQP \subseteq NP \implies P \neq NP$ [15]). Fig. 1.1a draws out this hierarchical Venn diagram explicitly. Practically, the potential existence of quantum algorithmic primitives with an advantage over their known classical counterparts updates the balance between cryptography (hard problems with verifiable solutions) and computability (ideally, easy problems once you have quantum). Perhaps the most famous example of this trade-off is Shor's set of algorithms, which finds the prime factors (or discrete logarithm) of a large integer (or modular subgroup size) N in time $\sim \mathcal{O}\left((\log \mathrm{N})^2\right)$. This is exponentially faster than the best-known classical algorithm, the general number sieve, which is polynomial in N. Building a sufficiently powerful quantum computer would thus make a significant fraction of past and current encryption protocols that use public key schemes based on factorization or discrete-logarithms, like the RSA and elliptic-curve cryptography, redundant.

An equally interesting application of controllable many-body quantum systems is the simulation of various inherently quantum phenomena [18]. Studying the natural world purely analytically is nigh impossible – in the presence of a large number of interacting particles, such as in a bio-molecule, a solid state material, or even in more fundamental problems like quark confinement in a gauge theory, it is imperative to perform large and precise calculations of the emergent physical phenomena. This is provably exponentially hard with only classical computation, in particular because the propagation of entanglement through such systems can grow quickly for strong coupling strengths, which in turn

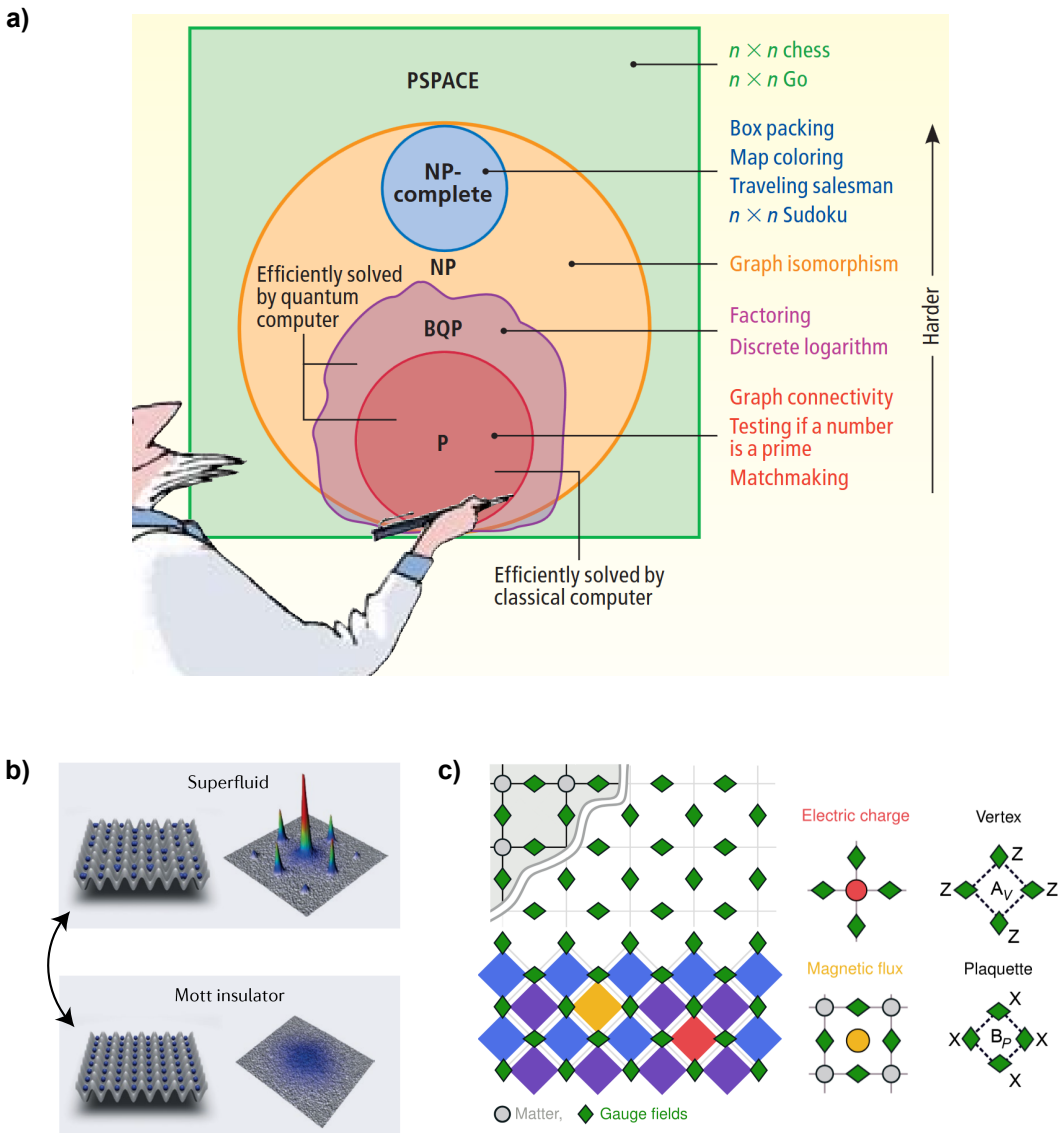

Figure 1.1: **Applications of quantum information processing a)** Current best estimate of the complexity classification for classical computational problems. PSPACE represents all problems that can be stored in polynomial space. BQP represents problems easy to solve with bounded error on a quantum system, which ideally includes problems outside P (classically deterministically easy) and perhaps even NP (classically deterministically hard). See [16] for more details. All rights for the illustration lie with the artist, Dušan Petričić. **b)** Example of quantum simulation on an atomic lattice for a superfluid to Mott insulator transition [17]. **c)** Example of quantum simulation on a superconducting device for $(2+1)$D lattice-gauge theories of QED [9].

requires the simulation of the exponentially large Hilbert space of the system. This is easy to see in an example problem like the dynamics of protein folding. The production and assembly of proteins is a fundamental function of disease-carrying viruses, and blocking this assembly through external molecules is an important way that drugs function. Simulating the assembly and production of these proteins thus directly enables novel drug discovery. Importantly, it is not enough to predict the final state of the protein, which can be achieved by well-trained machine learning models [19]. The dynamics of protein folding, which rely on quantum or semi-classical interactions between a large number of atoms and the solution they are submerged in, are important to understand and simulate. This would be made significantly easier on a quantum system that mimics the Hamiltonian of the protein, or even assists a semi-classical computation by providing accurate strengths for the interaction fields [20]. A less obvious example lies in simulating systems with strongly interacting fermions. Here, the sign problem [21] makes numerical solutions to the dynamics of the fermions being described by the fine-tuned cancellation of rapidly oscillating terms, which is provably NP-hard. However, a quantum system that natively contains fermions in the simulation could make such simulations relatively trivial [22].

While these useful simulations may remain a while away, many current experiments have already demonstrated interesting analog simulations with lattices of quantum modes or bits that can interact with each other in a nearest-neighbor fashion. These experiments can probe correlations that evolve as a function of system parameters to show interesting many-body phenomena that are somewhat resilient to decoherence, like phase transitions and scaling laws. Relevant examples of this include the Mott-insulator to superfluid transition in a Bose-Hubbard model (Fig. 1.1b), the insulator-metal transition in the Fermi-Hubbard model [23], or interacting models with long-range topological order [6]. Instead of engineering highly tailored Hamiltonians for analog simulation, it is also possible to Trotterize a many-body Hamiltonian evolution and decompose it into simpler discrete op-

erations on individual sites, ie, perform unitary computations through quantum gates [24]. This often scales very poorly in the number of operations required to approximate small evolutions (rotations) well, which inevitably leads to accumulated noise 'washing out' relevant quantum features at long simulation times. Some success has still been shown in current experiments by utilizing classical post-processing to mitigate errors [25]. Alternatively, one can demonstrate interesting phenomena that don't require a lot of gates, such as the braiding of effective anyons [7] and small quantum electrodynamic simulations that can be computed through lattice-based propagation of electric and magnetic fields (Fig. 1.1c). It is clear that many interesting emergent phenomena in physics can be probed with these highly controlled quantum systems, but performing practically useful simulations will inevitably require dealing with their fragility to noise.

Overall, this promise of exponential advantage for problems important to society (like better encryption protocols) and to physicists (like simulating complex quantum systems) has sparked enormous interest in the quest to build a working quantum computer (Fig. 1.2). In my time in graduate school, 'spooky action at a distance' became 'high-fidelity remote entanglement', which in turn became a necessary component for scalable quantum processors [26, 27]. Yet the quest for building a quantum system that actually performs useful computation remains a while away, since the fragility of quantum information grows exponentially with larger systems. In fact, I will try to convince you in this thesis that, at its core, useful quantum phenomena arise simply from a system's ability to navigate its available Hilbert space before interacting with its environment (which includes you, the observer). In particular, this thesis describes quantum systems composed purely of bosonic excitations, where visible quantum behavior is simply a trade-off between nonlinearity and dissipation. Most of the foundational concepts required to understand this thesis can be found in "Exploring the Quantum" by Serge Haroche and Jean-Michel Raimond [28], an introduction to quantum noise and measurement by A. Clerk et al [29], and

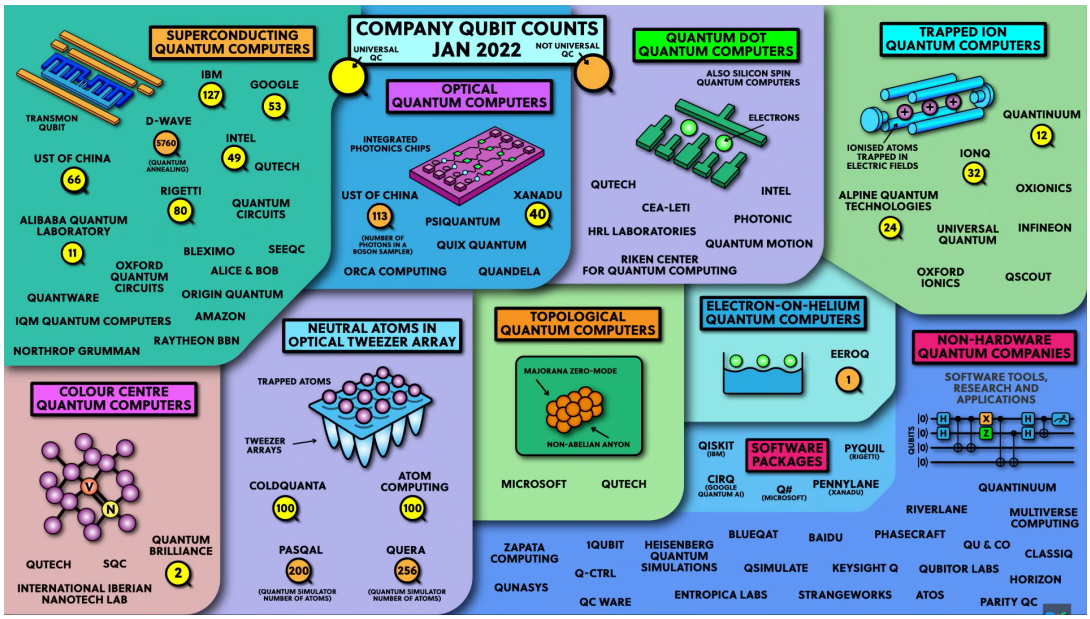

Figure 1.2: **Quantum industrial presence in January 2022**, Potential promises of significant advantage in both broadly relevant tasks (like breaking encryption and better machine learning models) and scientifically important demonstrations (like simulating many-body quantum phenomena) have led to a burst in the industrial pursuit of quantum processors. Illustration by Dominic Walliman for Domain of Science.

Steve Girvin's Les Houches notes on circuit-QED [30], all of which are excellent reads for those looking to learn more.

1.1 More control means more noise

Let us pause briefly and make the trade-off between the fragility and computational power of quantum systems concrete. A typical undergraduate quantum physics problem deals with 'closed' quantum systems - where the Hilbert space of a small system with explicit boundary conditions is entirely set by a time-independent Hamiltonian. However, performing useful computation requires executing a sequence of Unitaries and measuring the final state, which necessarily requires external input for time-dependent control, and output of information through a measurement apparatus. This means all realistic experiments

need to instead be performed on an open quantum system, where the environment has the opportunity to act on the system, both during the system's manipulation and while it is idle. For a tangible example, consider a simple two-level system (a qubit) coupled to a control 'port'. Realistically, this could be a Fabry-Perot cavity with a partially silvered mirror, an isolated spin coupled to an external magnetic field, or a microwave resonator weakly coupled to a transmission line. All of these systems allow the user to talk to and change the effective Hamiltonian of the system through their ports. Specifically for the systems mentioned above, this means that some external (classical) field is linearly coupled to an observable of the system (e.g., the cavity field or a spin Pauli operator). Now, any change to the value of this field automatically changes the system - and the system has no idea whether that change was caused by the user, or by the environment! In fact, any control knob the user has access to is also necessarily accessible by the environment, since the user and the control apparatus are part of the environment. This means that increasing control strength **always** introduces more capacity for errors.

Mathematically, this trade-off is given by general linear response theory and the Fermi Golden Rule [29]. Assume a control field $F(t)$ with intrinsic stochastic fluctuations (noise) that is linearly coupled to a system operator, for e.g., to the Pauli σ_x operator of a qubit. Then, the interaction is given by:

$$\hat{V} = \hbar g F(t) \sigma_x. \tag{1.1}$$

Here, driving $F(t)$ enacts a Rabi oscillation on the qubit, allowing a continuous set of control pulses. At the same time, assume $F(t)$ has environment-induced fluctuations with a noise spectral density of $S_F[\omega]$ (which is related to the Fourier transform of the auto-correlation of F, $\langle F(t)F(0)\rangle$ [31]). This then enacts an incoherent transition on the qubit,

given by Fermi's Golden Rule:

$$\Gamma_{\downarrow,\uparrow} = g^2 S_F[\pm\omega_q], \tag{1.2}$$

where $\Gamma_{\downarrow,\uparrow}$ are the transition rates between the excited and ground states of the qubit $|e\rangle \leftrightarrow |g\rangle$, and ω_q is its transition frequency. When the fluctuations in F are thermally distributed with some effective thermal population $\bar{n}$, then in equilibrium, these rates differ by the ratio $\Gamma_\uparrow/\Gamma_\downarrow = \bar{n}/(\bar{n}+1)$. We will see later how keeping these (effective) thermal populations low will be incredibly important in high-fidelity quantum control.

A few comments are in order. First, while we have assumed a linear coupling between a classical control and the qubit operator here, in general, the control field may be quantum, and its influence on the system's operator may be more complicated. Correctly treating this stochastic interaction to first order involves a general susceptibility function $\chi(t)$, and the Kubo formula [31]. We will see that in the presence of non-linearity and periodic drives, even this general susceptibility approach will fail – the system's dissipation will depend on the noise spectrum at many different frequencies, which will require the Floquet formalism to properly describe. Second, the qubit operator σ_x makes no assumption on how the qubit subspace is realized, and it may correspond to a fancier qubit encoded in a larger Hilbert space. Correspondingly, $F(t)$ may be composed of many control drives at different frequencies required to enact a gate on this encoded qubit. As a consequence, this trade-off between control strength and noise also exactly holds for novel protected qubits, like the fluxonium, zero-pi, or Kerr-cat superconducting qubits [32, 33, 34] – their gate infidelities are necessarily second-order sensitive to noise in any of their drives, exactly like a Rabi drive on a simple physical qubit. As an example, consider a virtual Raman process that drives a transition between two-qubit states in the fluxonium that otherwise have a vanishing matrix element [35]. The fluctuations in those drives can then be decomposed

as follows:

$$\begin{aligned} \hat{V}_{\text{gate}} &= \hbar g \; F_{\omega_1} F_{\omega_2} \; \sigma_x \\ \omega_1 - \omega_2 &= \omega_q \\ \implies S_F[\omega_q] &\approx F_{\omega_1}^2 S_{F_{\omega_2}}[\omega_2] + F_{\omega_2}^2 S_{F_{\omega_1}}[\omega_1]. \end{aligned} \tag{1.3}$$

It is clear that the dissipation rates during the gate, and hence the gate infidelities, are as sensitive to fluctuations in either drive F_1 or F_2 as our prototype qubit was to F. Note that this does not make (Hamiltonian) protected qubits redundant, since they may still have significant gains in their idle coherence or noise bias.

Finally, it is useful to interpret this dissipation in a different manner, that of the environment gaining information on the system. Instead of a classical stochastic variable that causes the system to decay or decohere, we can consider a quantum field that entangles itself with the qubit through the interaction $\hat{V} = \hbar g \; \hat{F} \otimes \sigma_x$, which the environment then measures. Every measurement of $\hat{F}$ also effectively applies some σ_x rotation to the system, causing transitions in $\langle \sigma_z \rangle$, which we, the observer, interpret as decoherence. Alternatively, the environment gaining information on σ_x necessarily reduces information about its conjugate variables through the Heisenberg uncertainty principle! This interpretation of the environment as a witness to the lost information in a dissipative process is also useful beyond pedagogical purposes. We show in Appendix A how, by utilizing an ancilla qubit as the external degree of freedom $\hat{F}$, we can effectively engineer the bath that our system sees and use it to our advantage.

In general, for a thermal environment, the effect of frequent measurements by the environment is to extract information about the system, but then lose that information through quick re-thermalization. This creates a classical probabilistic mixture of the 'pointer' states related to the environment's measurements, returning completely classical behavior when

these dissipation rates are much quicker than all other system evolution [36]. It is convenient to represent this general loss of information about the system from the user's point of view through the system's density matrix $\rho(t)$, with its precise evolution given by the Lindblad master equation [37]:

$$\dot{\rho}(t) = \mathcal{D}[\sqrt{\gamma}\hat{S}]\rho = \gamma\left(\hat{S}\rho\hat{S}^{\dagger} - \frac{1}{2}\{\hat{S}^{\dagger}\hat{S}, \rho\}\right) \tag{1.4}$$

Here, γ is the transition or dissipation rate (here $\Gamma_{\uparrow,\downarrow}$), and $\hat{S}$ is the system operator that the environment acts on (here $\sigma^{+,-}$). The derivation and interpretation of this equation are treated well in a number of texts, including [28, 38]. Note that a common scenario is the measurement of the qubit logical state by the user, which results in the dissipator $\mathcal{D}[\sqrt{\kappa_m}\sigma_z]$, frequently called measurement-induced dephasing. We will use this dissipator notation throughout the thesis, in particular when we design engineered dissipators in Appendix A to protect quantum information by purposely measuring carefully designed $\hat{S}$ operators.

Now that the trade-off between the power and fragility of quantum systems is clear, how does one actually build a powerful quantum system while retaining the ability to coherently manipulate it? Does the universe allow any protocol at all to refine quantum coherence, i.e., to take a bad quantum system limited by the noise in its controls, but somehow play the right control pulses to purify its coherence? The answer to this question has led to some of the most interesting research in modern quantum information theory and, as the reader may well know, lies in quantum error correction (QEC). Broadly, QEC provides a way to divide up our quantum system into a part where the important information is contained (logical space), and another (error or gauge space) within which the environmental decoherence acts [39]. The act of retaining logical information then simply becomes the act of clarifying or evacuating entropy from this space before it is completely

spoiled. Importantly, simply modulating $F(t)$ cleverly is not enough, since evacuating entropy is a non-time-reversible process - one *must* include either the measurement or reset of sub-parts of the quantum system. Regardless of how it is achieved and on what physical platform, the ability to successfully error-correct arbitrarily large quantum systems will mark a significant landmark in experimental physics, and in the human ability to control and understand Nature.

1.2 From atoms to circuits

Let us now shift focus slightly and introduce the quantum platform for the experiments described in this thesis. Fundamentally, the universe is built out of quantum parts, yet most of these parts decohere at the scale and temperatures at which we interact with them. It is interesting to consider for a moment why this is the case. Practically all platforms of quantum information are based on the electromagnetic (EM) interaction. As such, photons (the gauge bosons for the interaction) are always involved in mediating quantum information, or even directly storing it. Unfortunately, photons in a vacuum are identical non-interacting particles, so their quantum information is purely stored in the number of photons in any mode (e.g., of polarization), and single-photon manipulation is practically impossible. This necessitates the interactions of photons with electrons in matter, either through their charge or their spin, to perform any useful preparation, control, or tomography.

This is where the primary problem of visible quantum phenomena comes in – matter is lossy. This is especially true when the matter is derived from macroscopic objects or atomic clouds that have a large number of degrees of freedom. Not only do these degrees of freedom all interact with each other (e.g., through phonons or plasmons), but they also generally interact with any environmental electromagnetic modes that overlap

with them. For a large number of interacting degrees of freedom, the same matter that allows the introduction of quantum control also necessarily introduces an environment that is thermalized to some effective temperature T. Note that the emergence of an effective thermal (Boltzmann-like) distribution in even a closed quantum system is a necessity for the classical physics we observe at a macroscopic scale, and explaining this emergence is the subject of active research (see the eigenstate thermalization hypothesis [40, 41]). The presence of this effective environment means that any controllable degree of freedom also eventually thermalizes and completely decoheres to this environment. Thus, to store and manipulate information at the quantum level, a primary requirement is the reduction and isolation of the available degrees of freedom.

Historically, controllable quantum phenomena were thus restricted to very tiny objects that could be confined to a clean environment, like nuclear spins and atoms (ions) in traps. It is interesting to introduce both these types of systems, as a comparison to the experiments described in this thesis. Nuclear spins embedded in solid-state crystals are intrinsically qubit-like (fermionic two-level) systems with long coherence times. Their control and measurement are enabled through external magnetic fields, oscillating at microwave frequencies, that couple to these spins' Larmor transition. This platform, called nuclear magnetic resonance (NMR), gave rise to many significant initial results in experimental quantum control, including dynamical decoupling techniques for suppressing slow variations across ensembles of particles (called inhomogeneous broadening), and optimal pulse shaping for constructing precise Unitaries [42, 43]. Both of these techniques are extensively used today, including in this thesis. However, getting nuclear spins to interact with each other, an important prerequisite for quantum information processing, is exceptionally difficult. This is in particular because in most feasible implementations, their direct spin-spin coupling is negligible, and the solid-state platforms that embed these nuclei do not easily allow them to increase this coupling (e.g., by moving them closer together).

Thus, modern systems primarily utilize NMR for sensing, most prominently in Magnetic Resonance Imaging.

Atoms and light (AMO) were the dominant platform for quantum experiments throughout the late 20th century and have emerged as one of the most promising platforms for implementing quantum information processors. In a sense, describing the interactions of individual atoms and photons is precisely why quantum mechanics (in particular QED) was formulated. Atoms of any particular species are 'perfect' identical quantum systems designed by Nature. This means that the response of an atom to a well-defined incoming optical field needs to only be characterized once, and then Nature guarantees that any atom that sees exactly the same control field will respond in exactly the same way. Controlling atoms thus amounts to achieving two feats of experimental physics - being able to design the right control fields, and being able to isolate atoms from their environment by 'trapping them'. Both of these have decades of hard science and multiple Nobel Prizes behind them. For the former, modern experimental setups utilize frequency combs, electro-optic modulators, and spatial light modulators to design arbitrary wavefronts of incredibly stable frequencies of light (Fig. 1.3a). This leaves the latter task of isolating atomic transitions from their environment, while still getting them to talk to each other. Current experiments achieve this isolation by injecting an atomic species in an ultra-high vacuum, and then suppressing the other degrees of freedom of the atom through 'traps' that suppress thermal fluctuations in the motional states of the atom. The dominant examples of such atomic traps today come in one of three flavors. The first uses ions like Calcium or Ytterbium trapped in electric fields [44, 45], which interact with each other through shared vibrational modes of the trap [46], once they are brought close together by physical movement along the trap. The second, with significant recent advances, uses high-intensity light beams with a narrow beam waist that form local potential energy minimas due to the AC stark effect, a technique known as optical tweezers [47]. These tweezers are loaded

stochastically from an initial magneto-optical trap, reconfigured into a desired lattice, and then made to interact through physical movement of the tweezers bringing atoms together, and excitations of the atoms to (Rydberg) states with higher dipole moment (Fig. 1.3) [48]. Finally, some AMO experiments explore an enhancement of the atom-atom interaction by trapping atoms within an optical cavity and making them interact through the standing mode of the cavity field.

Overall, these atoms obey many of the usual intuitions of tiny natural systems exhibiting quantum effects, and are an attractive candidate for processing quantum information. However, the identical nature of atoms also means they are difficult to engineer to the user's delight, leading to physical limitations set by the atomic species and the fundamental constants. A significant example is the small dipole moment present in most atomic states, which makes the atom's interactions with both optical fields and other atoms (and in turn, readout and control speeds) fairly slow. This is made concrete through the metric of Cooperativity, $C = 4g^2/\kappa_1\kappa_2$, where g is the coupling strength of light and matter, and $\kappa_{1,2}$ are their individual dissipation rates. It is incredibly difficult to achieve $C > 10^2$ in an atomic system [50]. As a consequence, it is also difficult to measure specific quantum information without permanently affecting the atom (i.e., a non-demolition measurement via traveling light), which is fundamental to demonstrating many quantum experiments.

Is it possible to create quantum systems with engineerable Hamiltonians and natively strong interactions? Importantly, such a quantum system cannot rely purely on an individual atom, electron, or photonic mode as the carrier of information, since their Hamiltonians and interaction strengths are pre-decided by Nature. Instead, such a quantum platform would require macroscopic visible quantum phenomena resulting from an engineered ensemble of atoms that can take on specific Hamiltonians. Importantly, since macroscopicity always comes with a large number of degrees of freedom, the observer must be able to isolate and address a small subset of them in such a system. It turns out that a great way to

a)

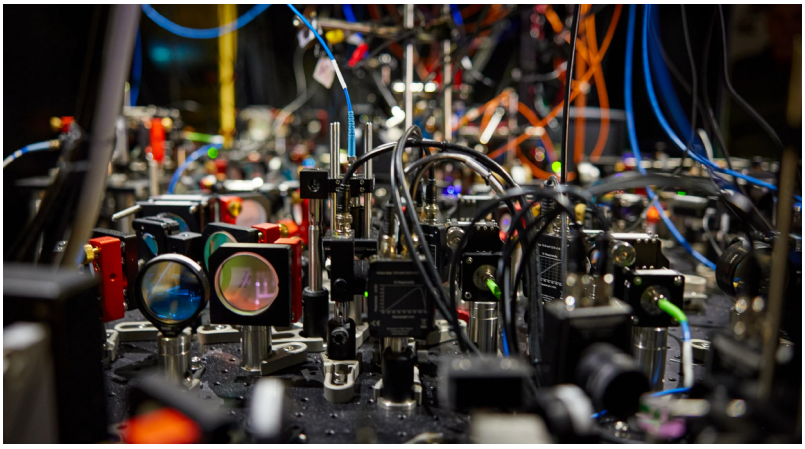

Example optical setup, Lukin lab (Quanta Magazine)

b)

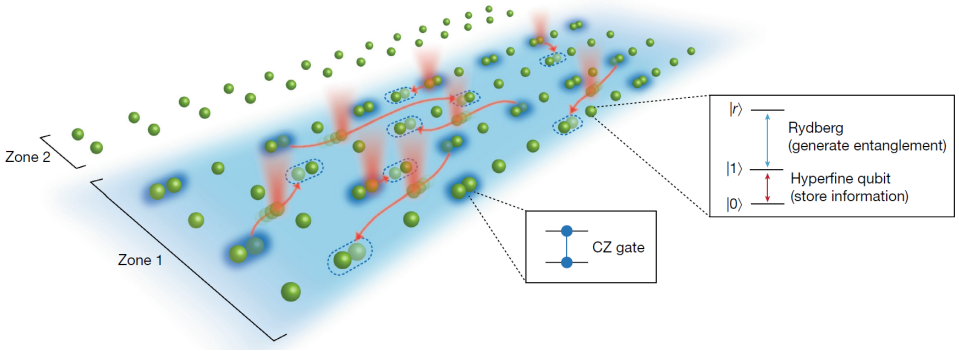

Bluvstein et al, 2022

Figure 1.3: **Example atomic optics experiment a)**, Since atomic species are identical by Nature, controlling them requires achieving well-established control optical fields with extremely high precision, through free-space and fiber optical setups. **b)**, Modern neutral-atom experiments control hundreds of atoms in optical tweezer arrays, and perform parallel gates through global optical fields. Information is usually stored in the hyperfine levels, except when entangling gates are required, when they must be excited to Rydberg levels with larger interaction dipole strength [49].

perform this isolation is to create a condensed phase within the matter, such as with a Bose-Einstein Condensate or a superconducting material. We will primarily focus on the latter, where the oscillations of Cooper pairs in the superconductor form addressable modes that are protected from the rest of the material by the superconducting gap. Quantum infor-

mation can then be stored in light interacting with this superconductor, specifically in photonic excitations with significantly lower energy than the superconducting gap, which are 'dressed' by their participation in the superconductor. The engineer-ability of these superconductors can then be introduced through Josephson junctions, which interrupt the superconductor and provide the nonlinearity required for control [51]. As an additional consideration, typical superconductors that form good Josephson junctions are Aluminum and Niobium, whose superconducting gap (80-200 GHz) implies that controllable modes must be in the microwave frequency regime ($\omega \sim 2\pi \times 0.1 - 100$ GHz). Thus, these devices can only show quantum behavior at particularly low temperatures $T \ll \hbar\omega/k_B \sim 0.5$ K (at 10 GHz), which are much smaller than the critical temperature required for superconductivity, but can be relatively easily achieved in Helium-based dilution refrigerators.

Quantum phenomena in such superconducting devices, broadly proposed by Anthony Leggett [52, 53], and made concrete through experiments at Berkeley (Clarke group [54, 55, 56]), NEC Japan (Tsai group [57]), Yale (Devoret, Schoelkopf and Girvin groups [58, 59, 60, 61, 62]), and Santa Barbara (Martinis group [63]), form the central platform for all experiments in this thesis. These experiments were phenomenal at the time of their introduction, producing true quantum behavior in macroscopic systems when many physicists still treated quantum and microscopicity as synonymous. Notice, for example, that both coherent microwave light and bulk superconductivity, while having origins in quantum phenomena, are still phenomenologically classical. The signal from a standard microwave generator often relies on an externally sourced, extremely precise frequency standard based on atomic transitions in a Rubidium clock, but the signal itself is entirely classical. Similarly, as Richard Feynman pointed out in his classic series of lectures at Caltech [64], the phenomenon of superconductivity relies on the superconducting phase, which is effectively the Schrödinger wave function for many particles (Cooper pairs) in the same state. This phase, in the absence of any junctions, can mostly be described by

classical wave mechanics. This includes fluxoid quantization in a superconducting loop, where the loop only supports discrete excitations when periodic boundary conditions are imposed, similar to waves on a plucked string.

To understand the behavior of light and matter in a superconducting device, it is useful to study the microwave response of the superconductor, and compare it to similar examples in the optical regime. Consider, for example, the case of microwave light injected into a superconducting cavity, where it then proceeds to reflect and form a standing wave. There are a few distinguishing characteristics that make this especially different from the optical equivalent, where visible light is injected into a Fabry-Perot cavity, with mirrors made of normal metal. The first is the wavelength of photons - features inside the microwave cavity, like a cylindrical stub, can easily be at a similar size to the wavelength injected microwave light. On the other hand, a typical Fabry-Perot is orders of magnitude bigger than the wavelength of optical light. These wavelength characteristics result in even classical microwave physics (e.g., in integrated circuits and antennae) being fairly different from both optics and DC electromagnetism.

The second effect to consider is how light at these two different frequencies interacts with the metal itself. Electromagnetic radiation penetrates a normal metal up to a skin depth, which is $\propto \sqrt{\rho/\omega}$, where ρ is the resistivity and ω is the frequency of the drive. For optical frequencies, this skin depth is extremely small ($\sim$ 1 nm), and individual metallic atoms within that depth do not have sufficient dipole moments or resonant transitions that can interact with the light (assuming $\omega \ll \omega_{\text{plasma}}$). This makes the light primarily reflect off the metal, populating a series of higher harmonics of the Fabry-Perot cavity resonance, and forming a wave packet that travels back and forth according to the rules of geometrical optics. In a superconductor, where $\rho = 0$, this skin effect is instead replaced by the London penetration depth, which is purely a function of the intrinsic properties of the charge carriers. When either the penetration depth is smaller than the thickness of the superconductor,

a)

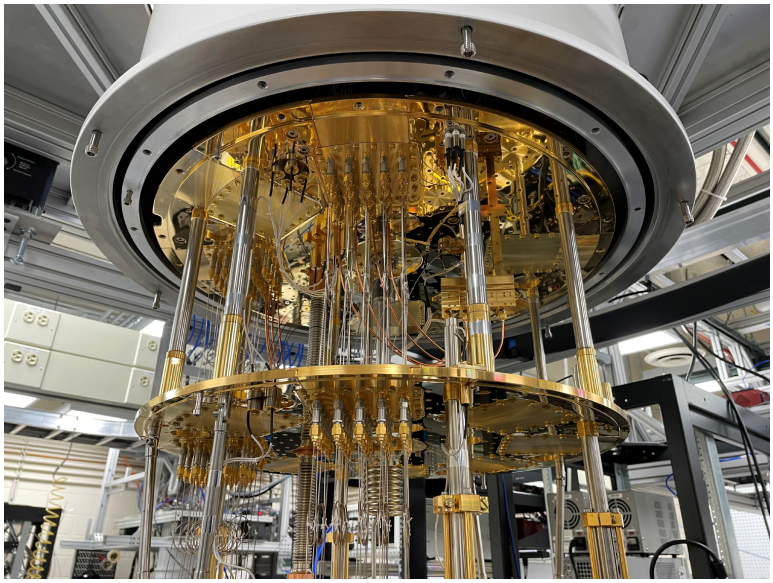

Wired-up dilution cryostat (Neo), Schoelkopf Lab

b)

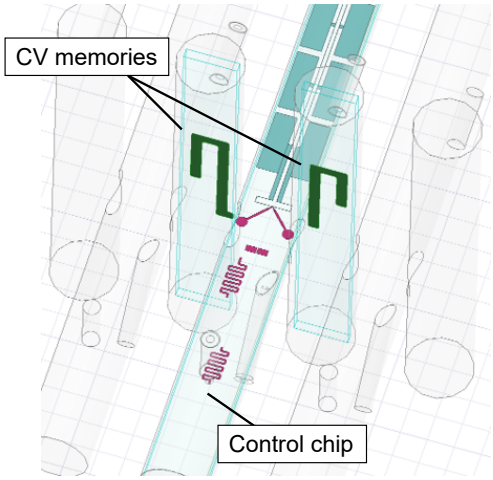

This thesis, 2024

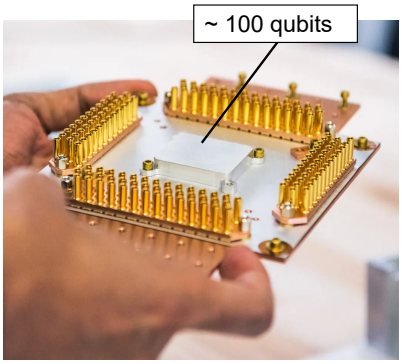

Google Willow experiment, 2024

Figure 1.4: **Example superconducting circuit experiment a)**, Superconducting devices must be operated at temperatures close to 10 mK, which can be achieved in Helium dilution cryostat like Neo, a fridge I set up for this thesis with Y. Lu and J. Garmon. **b)**, superconducting devices are operated at the base stage of the dilution cryostat, with external packaging and microwave control lines. The device architecture shown on the left (from this thesis) controls and measures information stored in CV microwave resonators. Industrial superconducting experiments (example packaged device from Google Quantum shown on right) contain more densely packed modes, with $O(100)$ qubits and control lines.

or the field is sufficiently weak, the superconductor behaves equivalently to a perfect conductor, with Cooper pairs driven by the input light field. While these Cooper pairs face truly zero impedance at DC, at non-zero frequencies, they face an inductive impedance that has both a geometric contribution, and a kinetic contribution from the presence of impurities (disorder). This inductance, in combination with a geometric self-capacitance, results in a set of harmonic LC-oscillator modes in the superconducting 'circuit' that can have (plasma) frequencies that are comparable to the incoming radiation, and dipole moments comparable to the total size of the superconductor [62]. Any microwave photons interacting with this superconductor then participate in the superconductor's modes by driving these Cooper pairs through an effective distributed LC circuit, with their exact coupling set by the frequency and impedance of the light and superconducting modes.

The language of modes, participations, and impedance matching serves as a very useful way to interpret the behavior of superconducting circuits, and harnesses extensively developed RF-engineering theory, with its rich history from the second world war [65, 66]. It also makes clear that the interaction of a pure junction-free superconductor and microwave photons (with homodyne or heterodyne measurements) is indeed classical, since the system simply behaves like a set of driven coupled harmonic oscillators. The key to observing quantum phenomena is then to introduce non-linearity and number-resolved measurement, which one does through the Josephson junction. Crucially, while in the absence of any junctions, the superconducting phase behaves classically, it still does possess quantum-mechanical phase coherence. It is simply impossible to utilize or measure this coherence in the absence of non-linearity.

A Josephson junction acts as a non-linear inductor with a sinusoidal current-phase relationship [67], which separates the usually equally-spaced levels of a superconducting oscillator. The simplest circuit that demonstrates this is, therefore, a superconducting LC oscillator, where the 'L' has been replaced by a junction with inductance L_J, commonly

known as the transmon [61, 68]. Its Hamiltonian is represented by:

$$H_q = 4E_C\hat{n}^2 - E_J \cos\hat{\theta}, \tag{1.5}$$

where $E_C = e^2/2C$ is the capacitive 'charging' energy, and $E_J = (\Phi_0/2\pi)^2/L_J$ is the junction's inductive energy. Importantly, the natural conjugate variable to the superconducting phase is the charge on the capacitor, which we turn dimensionless by instead quantifying the number of Cooper pairs $\hat{n}$. For such an ideal junction, the potential energy is the periodic $\cos\hat{\theta}$ function. This means that the potential energy is bounded, with the existence of free-running (ionized) states outside the 'cosine well', and that its conjugate variable $\hat{n}$ is necessarily discretely quantized. The existence of the nonlinearity splits up the oscillator's spectrum into several discrete coherent transitions, with the possibility of ionization – thus exhibiting exactly the same behavior as an artificial atom. Selecting a particular transition from this effective atom then forms a two-level system that can be used to store information, commonly called a qubit. For example, it turns out that the ground state $|g\rangle$ and first-excited state $|e\rangle$ of this circuit are separated from the rest of the transitions in the circuit by a gap $\sim E_C$, when $E_J \gg E_C$. Thus RF radiation that drives the Cooper pairs in the circuit near the $|g\rangle \leftrightarrow |e\rangle$ transition can be analyzed much more simply through the language of a two-level system [61, 68]:

$$\begin{aligned} H_q &= 4E_C\hat{n}^2 - E_J\cos\hat{\theta} + \hbar F(t)\hat{n} \\ &\approx \frac{\hbar\omega_q}{2}\hat{\sigma}_z + \hbar\tilde{F}(t)\,\hat{\sigma}_x, \end{aligned} \tag{1.6}$$

where $F(t)$ ($\tilde{F}(t)$) is a time-dependent drive on the Cooper pairs (qubit). When this drive is tuned near resonance, with drive frequency $\sim \omega_q$, the qubit exhibits Rabi oscillations similar to an atomic transition or a spin. This ability to simplify the complex macroscopic circuit to two-level quantum mechanics is also why it is tractable to understand and control

modern chips with thousands or even millions of such circuits – and correspondingly, the failure of this two-level approximation is also what leads to many of the unsolved problems in superconducting quantum processors today.

Now that we have our engineer-able 'atom', what does a quantum light-matter interaction look like? Of course, there is the previously discussed interaction between a strong RF drive and the Cooper pairs, but the drive here is purely classical. In fact, this classicality is well-defined: the drive exists in a coupling regime where its exchanges with the superconducting device are negligible compared to the total energy in the field, i.e., it is 'stiff'. This means the drive's behavior can be simply modeled as the expectation value of its quantum mechanical field – or equivalently, any entanglement with the atom can be ignored. This 'stiff pump approximation' will be especially important when we engineer drives for more complicated driven processes later in this thesis. To define a quantum-mechanical light-matter interaction, we need observable quantum phenomena near the single-photon level. As a result, we will find it useful to first cast the Josephson circuit we introduced in the language of photonic excitations.

A helpful result for the junction-based oscillator presented above, which in the regime of $E_J \gg E_C$ is called the transmon, is that it largely behaves like a harmonic oscillator with only a slight anharmonicity. This means that the usual basis for a harmonic oscillator, that of the phase-space and ladder operators, can be successfully used to describe the circuit. In particular, we can quantize this circuit by taking the conjugate variables of charge $\hat{n}$ and phase $\hat{\theta}$ to be the momentum and position operators:

$$\begin{aligned} \hat{n} &= n_{\text{zpf}}\, i(q - q^\dagger), \quad n_{\text{zpf}} = \left(\frac{E_J}{32E_C}\right)^{1/4} \\ \hat{\theta} &= \theta_{\text{zpf}}(q + q^\dagger), \quad \theta_{\text{zpf}} = \left(\frac{2E_C}{E_J}\right)^{1/4}, \end{aligned} \tag{1.7}$$

where $q, q^\dagger$ annihilate and create photons in the circuit. This casting of charge and phase as

an oscillator phase space is not random – in fact, it precisely fits the description of an electromagnetic oscillation like light, with the electric and magnetic fields being the conjugate momentum and position variables. What's more interesting in this transformation is that for the periodic Josephson potential, the superconducting phase and charge operators are periodic and discrete variables, respectively, due to the boundary conditions they satisfy. Casting them as continuous phase space operators requires their deviations from the origin to be small, which is equivalent to approximating the cosine potential as a quadratic one. Alternatively, introducing a (parasitic) linear inductor makes the potential non-periodic, regaining usual phase-space behavior [69]. The approximately harmonic behavior in the regime $E_J \gg E_C$ also exponentially suppresses the sensitivity of the transmon's spectrum to noise in its charge degree of freedom, which is discussed in detail in [61].

Sticking to small deviations, one can Taylor expand the cosine potential to get the effective transmon Hamiltonian, which behaves like a quantum mechanical Duffing oscillator:

$$\begin{aligned} &H_q/\hbar \approx \omega_q \hat{q}^\dagger \hat{q} + \frac{K_q}{2} \hat{q}^{\dagger 2} \hat{q}^2, \text{ where} \\ &\hbar\omega_q = \sqrt{8E_J E_C} - E_C, \;\; \hbar K_q = -E_C. \end{aligned} \tag{1.8}$$

This anharmonic oscillator, with the gap between subsequent levels of the oscillator being given by K_q (called the 'Kerr nonlinearity'), then analytically describes most of the transmon's relevant physics. If a resonant drive tone interacts with the transmon, at sufficiently low strength compared to the anharmonicity, it behaves like the qubit in Eq. 1.6, with the transformation $\sigma_z = 2\hat{q}^\dagger \hat{q} - 1$ and $\hat{\sigma}_x = (\hat{q} + \hat{q}^\dagger)$. Importantly, to display quantum behavior, this Kerr nonlinearity must be stronger than the transmon's decoherence; otherwise, the transmon simply behaves like a classical (Duffing) oscillator.

We are now sufficiently prepared to introduce the quantum light-matter coupling. Let us assume a second (standing) electromagnetic mode, which is a perfect harmonic oscilla-

tor described by the ladder operators $\hat{a}, \hat{a}^\dagger$. If it exists in the same geometric space as the transmon, its field has some overlap with the circuit geometry and can drive the Cooper pairs in it. In general, since the junction's inductance is extremely localized, but its capacitors are not, we will assume this to be an interaction through its electric field. In the language of linear optics, this then looks like a coupling along the momentum axis:

$$H_{qA}/\hbar = \omega_q \hat{q}^\dagger \hat{q} + \frac{K_q}{2} \hat{q}^{\dagger 2} \hat{q}^2 + \omega_A \hat{a}^\dagger \hat{a} + g_{qA}(\hat{q} + \hat{q}^\dagger)(\hat{a} + \hat{a}^\dagger), \tag{1.9}$$

where ω_A is the frequency of the EM mode and g_{qA} is its linear coupling to the anharmonic circuit, sometimes referred to as its vacuum Rabi strength. There exist perturbatively accurate methods that diagonalize this Hamiltonian, but at the lowest order, we can simply utilize the fact that both the transmon and oscillator modes here are approximately linear. This means the coupling produces two new normal modes, a transmon-like mode $(\tilde{q}, \tilde{q}^\dagger)$ and an oscillator-like mode $(\tilde{a}, \tilde{a}^\dagger)$. These are given by the transformation:

$$\begin{aligned} \tilde{q} &\approx \hat{q} + \left(\frac{g_{qA}}{\Delta_{qA}}\right) \hat{a} \\ \tilde{a} &\approx \hat{a} - \left(\frac{g_{qA}}{\Delta_{qA}}\right) \hat{q}, \end{aligned} \tag{1.10}$$

where Δ_{qA} is the detuning of the two modes, $\omega_q - \omega_A$. Note that this assumes a strict hierarchy of rates, specifically $g, K \ll \Delta_{qA} \ll \omega_q + \omega_A$. The effective participation ratio $\beta_{qA} = (g_{qA}/\Delta_{qA})$ is of critical importance; it tells us how much these two modes hybridize, and governs most of the physics that results from this interaction.

The Hamiltonian for the 'dressed' light and matter modes of the above system can be

written down as a generalized version of the Jaynes-Cummings Hamiltonian [70]:

$$H_{qA}/\hbar = \tilde{\omega}_q \tilde{q}^\dagger \tilde{q} + \frac{\tilde{K}_q}{2} \tilde{q}^{\dagger 2} \tilde{q}^2 + \tilde{\omega}_A \tilde{a}^\dagger \tilde{a} + \chi_{qA} \tilde{q}^\dagger \tilde{q}\, \tilde{a}^\dagger \tilde{a} + \frac{K_A}{2} \tilde{a}^{\dagger 2} \tilde{a}^2, \quad \text{with } \chi_{qA} \approx \beta_{qA}^2\, K_q, \quad K_A \approx \beta_{qA}^4\, K_q \tag{1.11}$$

The primary effect of the hybridization, other than some perturbations to the original transmon and oscillator Hamiltonians, is the addition of two new terms, χ_{qA} and K_A. The latter is an effective anharmonicity inherited by the oscillator due to its participation in the transmon. The former is a new term entirely, and it provides much of the quantum functionality that is unique to superconducting circuits. χ_{qA} is a dispersion of the oscillator's (transmon's) frequency according to the state of the transmon (oscillator). This 'dispersive' shift can be significantly stronger than any decoherence in the system, which means that this shift is observable even at the single excitation level (a strong-dispersive coupling regime [59]). This means that the frequency of any photons leaving the oscillator also carries information about the qubit's excitation, but they do not change the qubit's state [71]. If these photons are monitored, for example, by a careful measurement setup that purposely probes the oscillator's frequency with a coherent drive tone, then one gains a quantum non-demolition (QND) measurement of the qubit's state [72]. Alternatively, if the oscillator is actually a second anharmonic mode, i.e., a qubit, this directly enables two-qubit interactions [73]. In combination, this simple driven circuit enables a macroscopic engineerable Hamiltonian whose quantum information can be measured and manipulated with high fidelity.

The above effects of hybridization are only those at the Hamiltonian level. When a superconducting circuit is coupled to another electromagnetic mode, it also shares some of its loss and dephasing. Assuming a Markovian environment (unstructured noise), this again can be derived simply from the Eq. 1.10, with the inherited decay rate of the trans-

mon given by:

$$\gamma_1 = \beta_{qA}^2 \, \kappa_A, \tag{1.12}$$

where κ_A is the dissipation rate of the coupled mode. Additionally, even a weak dispersive interaction (χ_{qA}) with any coupled modes allows them to measure the transmon's state, which means any thermal fluctuations of photons in these modes also result in a dephasing of the transmon:

$$\gamma_\varphi \approx \bar{n}_{\text{th},A} \, \kappa_A \left(\frac{\chi_{qA}^2}{\chi_{qA}^2 + \kappa_A^2} \right). \tag{1.13}$$

In a general electromagnetic environment, the 'Purcell' decay rate of the transmon through each coupled mode adds up, as does the thermal dephasing, and both can be very significant when uncontrolled. Isolating this quantum system thus requires the introduction of a literal electromagnetic box, which usually takes the form of an outer superconducting 'package. ' Ideally, this box then forms a nearly closed system, with access allowed only through well-defined electromagnetic ports, practically formed by microwave transmission lines. In a typical experiment, this package, with an embedded circuit, sits at the base stage of a dilution fridge (see Fig. 1.4). Both control and measurement of the circuit occur through microwave radiation traveling through coaxial transmission lines that carry information through the various stages of the refrigerator, with carefully chosen filters, attenuators, and amplifiers at each stage [74].

Now that we have introduced the core of the platform, it is useful to take a step back and put it into the context of other quantum systems and experiments. Superconducting circuits permit highly engineerable Hamiltonians that introduce strong nonlinearity into electromagnetic modes in the microwave regime. This nonlinearity is intrinsic to the Josephson Junction and is significantly stronger than the nonlinearities available in other bosonic platforms, like optical systems. The large dipole moments available in macroscopic circuits also allow strong linear couplings between circuits, and to external photons.

As a reference, typical coupling strengths in this thesis are around $g/2\pi \sim 100$ MHz, and typical dissipation rates are $\kappa/2\pi \sim 1$ KHz. The resulting cooperativity of two such modes is $C = 4g^2/\kappa^2 \sim 10^{10}$, which is approximately eight orders of magnitude stronger than the best atomic systems. Combining the strong nonlinearity and coupling allows a strong nonlinear interaction (the dispersive shift), which in the quantum information context allows fast entangling gates and high-fidelity measurement.

The platform of superconducting circuits lies at the intersection of macroscopic quantum phenomena in microwave photons and solid-state systems. This places it in the interesting regime of being equally well described by the language of quantum optics, and by that of RF engineering. As an example, the non-linearity in these circuits is often easily cast in the basis of photonic excitations, with the presence of higher-order nonlinear interactions 'dressing' its effective spectrum. This dressing, which perturbatively adds the contributions of various (possibly virtual) interactions both within the circuit and with its electromagnetic environment, is well described by Feynman-diagram-like rules that enumerate the possible mixing processes. Such a dressing and energy renormalization is fundamental to the field of quantum electrodynamics – for example, the self-energy (and thus mass) of the electron is precisely calculated through a very similar process, accounting for decays into other photons and particles. Correspondingly, the interaction of light and matter in these circuits takes on extremely analogous behavior to true quantum optic experiments, and cavity quantum-electrodynamics. On the other hand, many of the phenomena that affect these circuits – like their hybridization with other modes and their dissipation through electromagnetic ports, are well captured by just the linear part of their Hamiltonians. As a result, one can utilize classical linear electromagnetic simulations to directly calculate much of the circuit's behavior, and then perturbatively add in the nonlinear quantum effects [75, 76]. At the same time, the circuits' solid-state nature allows their experimental realization through lithographic techniques derived from integrated cir-

cuits, where tools like electron-beam lithography and scanning-electron microscopes give extreme precision over the design and characterization of each device. Together, the combination of fundamentally quantum behavior and classically precise engineerability make the field of 'circuit quantum electrodynamics' (circuit-QED) an extremely rich platform for both academic and industrial research.

As a small aside, another very important aspect of the quantum behavior in circuit-QED is the ability to send and receive precise microwave radiation. For context, in a typical setup, one uses an extremely stable atomic clock reference to produce room-temperature signals at close to 10 milli-watts, which are eventually used to control and measure about three-millionths of an atto-Joule of energy (a single photon at 5 GHz). This process requires extremely carefully engineered drive lines that filter and attenuate the input drives while adding minimal noise to the control signal as it travels down the cryostat.

Overall, circuit-QED enables a combination of high performance and high engineerability in relatively tabletop quantum information experiments. Yet, the challenge of actually building a fault-tolerant quantum system is extremely grand – and perhaps only achievable with very large and extremely precise systems. As such, what interesting quantum phenomena can one show that both take advantage of circuit-QED's special abilities and are achievable in a near-term small-scale experiment? A promising candidate is to use the strong light-matter couplings available in circuit-QED, not to manipulate the circuits as quantum bits, but to instead manipulate and measure quantum phenomena in the light. Such 'bosonic' quantum experiments, described in detail in the next Chapter, underlie all of the work introduced in this thesis.

As an overall outline, we will first discuss bosonic quantum information, including the natural ways to describe bosonic states and the errors that they can incur. We will show that natural bosonic control and measurement techniques available through the strong-dispersive coupling in circuit-QED can often propagate errors and spoil the bosonic modes.

We will thus propose an alternative form of control, that will involve driven non-linear circuits with special symmetries. These circuits, called 'parametric couplers', will activate a coupling between external modes whenever desired, while themselves remaining in the ground state. Such a coupler forms the central element of a novel bosonic control architecture, where light-matter interactions can be introduced dynamically. We will show that this allows non-linearity to be injected into the system in a precise manner only when desired, while retaining a degree of native tolerance to decoherence in the matter. If such an architecture can be demonstrated, we can utilize it to autonomously protect any appropriately encoded quantum information in a bosonic system. Together, the theories and experiments in this thesis provide a platform for near-term quantum information systems that are well-suited to simulate bosonic physics and demonstrate bosonic error correction.

Chapter 2

Bosonic Quantum Information

Modern quantum information theory primarily deals with quantum bits of information, typically encoded into physical two-level systems, like spins or atomic transitions. Circuit-QED is no stranger to this - most superconducting systems encode information into the ground and excited states of transmons, or its fancier cousins like the 'fluxonium' [32] and 'zero-pi' [77] circuits. Yet the intrinsically high cooperativity present in circuit-QED can be utilized to store information in a fundamentally different way - we can store it in the phase-space of a linear oscillator. Such an implementation is interesting because it directly utilizes the bosonic nature of the microwave photons we manipulate.

Circuits that behave like a Kerr oscillator (like the transmon) can in general smoothly interpolate between bosonic and fermionic behavior. The Kerr nonlinearity sets how anharmonic the oscillator's level spacing is. If the Kerr is much larger than the dissipation rate, the difference between subsequent levels is not washed out, and a sufficiently weak probe tone can pick out and manipulate a single pair of levels. A large Kerr implies that photons interacting with one level face a large barrier in populating subsequent levels, which is effectively an exclusion principle (but unlike Pauli's, it does not arise from exchange statistics). This means that the system can emulate fermionic behavior, despite arising from bosons, and this is made mathematically concrete through the Jordan-Wigner

transform. On the other end of the spectrum, if an oscillator's Kerr is far smaller than its dissipation linewidth, any photons interacting with the oscillator can continuously populate levels, having access to an effectively infinite Hilbert space. This then is described by a fundamentally different class of quantum information, that of the 'bosonic' kind, which comes with intrinsically different physics. A usual superconducting system constructed from a combination of transmon-like and linear oscillator-like modes constitutes a hybrid boson-fermion system, and can display both kinds of physics.

There are two main advantages of utilizing bosonic quantum information in circuit-QED. The first is the nature of bosonic information - bosonic phase space is described by a continuous variable, and can support highly non-classical states that require effectively infinite levels to perfectly describe. In fact, given the formally infinite nature of an oscillator's Hilbert space, one could technically simulate arbitrarily large quantum problems in a single oscillator. This is of course, only true for a perfectly dissipation-less and Kerr-free oscillator - for any realistic loss or Kerr, higher levels in the oscillator always have higher dissipation or energy gap difference, eventually making it intractable to store bosonic information in them. However, many continuous variable states exist that have an exponentially decaying overlap with higher oscillator levels, yet still show interesting physics that is difficult to replicate in qubit-based systems given available native controls. There even exist fundamental no-go theorems in qubit-based quantum information [78], that may be possible to break in continuous variable systems. More practically, systems with bosonic modes that couple to each other, even with access to only a part of their total Hilbert space, offer very different connectivity than a grid of qubits. This makes the propagation of information and entanglement in such systems much faster, assuming one has natively bosonic controls, which makes it much harder to classically simulate such systems. As an example, estimating the ground state energy of a sufficiently complex bosonic Hamiltonian is provably NP-complete [79], but of course trivial with the right bosonic

system. It is in fact an open problem whether the complexity class of continuous-variable quantum algorithms exactly overlaps with qubit-based ones (BQP $\subseteq$ CV-BQP). This is especially important because much of the universe is indeed bosonic - from the carriers of all fundamental forces, to complex bosonic molecules, and collective excitations in condensed matter systems - and simulating these phenomena might be most tractable on an intrinsically bosonic 'processor'.

The second major advantage lies in the specific implementation of bosonic systems in circuit-QED. Bosons in circuit-QED are excitations of a linear electromagnetic mode, generally contained in a standing wave of a superconducting resonator. As a result, they require no junctions or special elements. A simple superconducting stub or stripline of the appropriate length (some fraction of the desired photon's wavelength) can host a bosonic mode. If the resonator is a section of superconductor (tantalum, aluminum, or niobium) deposited on a (sapphire or silicon) substrate and surrounded by a ground plane, it is referred to as a 2D resonator, with its electromagnetic field confined to near the stripline. If it instead exists in a physical cavity, within standing modes of a rectangular or cylindrical box (often with a stub), it is referred to as a 3D resonator. In general, the participation of the oscillator mode in the superconductor and in the vacuum around it can be tuned over a wide range, with its 'mode volume' setting its sensitivity to lossy matter, and hence its dissipation rate. For either configuration, these pure resonators generally have a higher quality factor (defined as $Q = \omega/\kappa$) and longer lifetimes than junction based circuits. However, that is not all – an oscillator's frequency, and hence phase evolution, is extremely stable and set primarily by their macroscopic geometry. This means they do not natively incur any dephasing of quantum information, and the only error channel that affects them is the occasional loss of a photon to external modes. As such, they have a highly structured error channel, which is advantageous to any bosonic implementations that attempt to detect or correct the resonator's errors. Some of these, notably the dual-rail encoding described

later in this chapter, even use resonators as two-level systems, simply benefiting from their inherent noise bias and not their large Hilbert space.

The rest of this chapter will describe continuous variable (CV) quantum information in more detail, including how to control and encode information in bosonic modes. We will then outline typical bosonic architectures in cQED, what errors in those architectures look like, and introduce parametric controls as a potentially better solution.

2.1 Quantum mechanics in an oscillator

The Hilbert space of an oscillator can be described in two interesting ways. The first uses the Fock basis and describes the oscillator's behavior in terms of the number of photons in the oscillator. This is precisely the basis we have been using so far, with the 'levels' of an oscillator corresponding to different occupation numbers. This discrete photon number basis is conjugate to the periodic oscillator phase, which is distinct from the superconducting phase that makes up the oscillations. In the absence of nonlinear controls, this oscillator phase φ determines the phase on each Fock state $|n\rangle$ in any superposition state as $\varphi_n = n\varphi$. This number-phase conjugate basis is extremely useful as it provides a simple mapping from continuous variable phenomena to discrete variables through truncation in the Fock space - making it the default choice for classically simulating these systems. It is also the natural basis to describe natively available nonlinearities like the Kerr and the dispersive shift. Controlling the oscillator in this language then has two requirements. The first is that one must be able to change the amplitude of Fock states and introduce photons into the system. The second is that one must be able to change the phase associated with each Fock state in a superposition. Despite the oscillator's infinite nature, one can prove that having these two controls is sufficient for enacting any Unitary operation (ie universal control) on the oscillator's state. As we will see, in a dispersively coupled oscillator-qubit

system in circuit-QED, these two controls can be provided by just the complex-valued linear drives on an oscillator and qubit [80].

The second way to describe oscillators is much more directly suited to their continuous variable nature - it is by directly using states and operators that are best described in the language of it phase space. The position and momentum of an oscillator, in normalized coordinates, is given by:

$$\begin{aligned} q &= \hat{x}/\sqrt{2}\ x_{\text{zpf}} = (\hat{a} + \hat{a}^\dagger)/\sqrt{2} \\ p &= \hat{p}/\sqrt{2}\ p_{\text{zpf}} = i(\hat{a} - \hat{a}^\dagger)\sqrt{2}, \end{aligned} \tag{2.1}$$

where $x_{\text{zpf}} = \sqrt{\hbar Z/2}$ and $p_{\text{zpf}} = \sqrt{\hbar/2Z}$ are the uncertainties (std. deviation) of the vacuum state $|0\rangle$ in the $\hat{x}$ and $\hat{p}$ directions, with $Z = \sqrt{L/C}$ as the impedance of the oscillator. The two uncertainties minimize the Heisenberg uncertainty principle, with $x_{\text{zpf}} p_{\text{zpf}} = \hbar/2$. The normalized coordinate operators further have these uncertainties symmetrized and equal to 1, with their commutator $[q, p] = i$, and their Heisenberg evolution given by $q^2 + p^2 =$ const., greatly simplifying the analysis of the oscillator. When the phase-space operators are driven by a resonant drive, the oscillator undergoes a displacement transform, entering a coherent state:

$$D_\alpha = \exp\{i(\alpha\hat{a}^\dagger - \alpha^*\hat{a})\} \tag{2.2}$$

$$D_\alpha|0\rangle = |\alpha\rangle. \tag{2.3}$$

This coherent state $|\alpha\rangle$ also has the same symmetric minimal-uncertainty in $\{q, p\}$ as the vacuum state, and α can take continuous complex values – driving along the q and p directions changes the value of α along the real and imaginary axes respectively. In fact, the drives precisely change the mean value of the wavefunction in $\{q, p\}$ by $\{\text{Re}(\alpha), \text{Im}(\alpha)\}$, which justifies naming this operation a 'displacement'. Thus the set of all coherent states

$\{|\alpha\rangle\}$ forms an overcomplete basis for describing phase space, with a clear one-to-one mapping to $\{q, p\}$, and is complementary to the number-phase basis for describing oscillator dynamics.

The coherent state basis provides a much more natural intuition for intrinsically bosonic controls and has some important properties with respect to dissipation that only perfectly hold in the truly continuous (infinite) Hilbert-space limit. Intrinsic bosonic controls are controls that don't select out particular oscillator levels to interact with - they have non-zero action on the entire Hilbert space. As such, in the Fock basis, they can only be described by a sum of an infinite series, but they are much more simply described in terms of polynomials of the ladder operators $\hat{a}, \hat{a}^\dagger$. The displacement transform above is the unitary formed by such a control, specifically the linear drive $\epsilon_a(t)\hat{a}^\dagger + \epsilon_a^*(t)\hat{a}$ (with $\alpha = \int \epsilon_a(t)dt$). We can then list all relevant control Hamiltonians up to second-order in the ladder operators, which we will see later are directly achievable in circuit-QED. Here we list them for two oscillators (with ladder operators $\hat{a}$ and $\hat{b}$), but in general, these can be extended to any number of oscillators (with phase-space dimension 2^n):

$$H_{\text{contr.}}/\hbar = \delta\omega_a(t)\ \hat{a}^\dagger\hat{a} \qquad \text{..'Detuning'} \tag{2.4a}$$

$$= g_{sq}(t)\ \hat{a}^2 + g_{sq}^*(t)\ \hat{a}^{\dagger 2} \qquad \text{..'Squeezing'} \tag{2.4b}$$

$$= g_{bs}(t)\ \hat{a}^\dagger\hat{b} + g_{bs}^*(t)\ \hat{a}\hat{b}^\dagger \qquad \text{..'Beamsplitting'} \tag{2.4c}$$

$$= g_{tms}(t)\ \hat{a}^\dagger\hat{b}^\dagger + g_{tms}^*(t)\ \hat{a}\hat{b} \qquad \text{..'Two-mode sq.'} \tag{2.4d}$$

$$= |g_{\text{q-q}}(t)|\ (e^{-i\phi_a(t)}\hat{a} + h.c.)(e^{-\phi_b(t)}\hat{b} + h.c.). \qquad \text{..'Quadrature coupl.'} \tag{2.4e}$$

Note that by Eq. 2.1, these also describe possible controls based on polynomials of $\{q, p\}$ up to second order. Each control Hamiltonian realizes a unitary transform $U_{\text{contr.}} = \exp\left[-\frac{i}{\hbar}\int H_{\text{contr.}}(t)dt\right]$ on the oscillator state, assuming no other controls are simultaneously acting (otherwise there will be additional effects due to the commutators of those

controls, following the rules of time-ordering).

The first of these controls is a simple temporary shift in the oscillator frequency, effectively imparting a phase rotation $U_{\text{rot}} = \exp\left[i\,\delta\phi_a\,\hat{a}^\dagger\hat{a}\right]$, where $\delta\phi_a = \int \delta\omega_a(t)dt$. The second of these, commonly known as 'single-mode squeezing' or just 'squeezing', performs the unitary $U_{\text{sq}} = \exp\left[\frac{1}{2}(za^2 - z^*a^{\dagger 2})\right]$, where $z = \int g_{sq}(t)dt$. As its name suggests, U_{sq} squeezes the uncertainty of a coherent state along an axis given by $\arg[z]$, by an amount $e^{-2|z|}$ (and correspondingly elongates it along the orthogonal axis). This squeezing of uncertainty has fundamental usefulness in sensing and quantum measurement, allowing a system to bypass the limits set by standard fluctuations of a vacuum or thermal state. Intuitively, by pre-squeezing the input measurement field, one can exponentially (in $|z|$) increase sensitivity to displacements along the squeezed direction. Similarly, squeezing the output field from the oscillator in the orthogonal direction can protect it from fluctuations that occur further along the measurement chain, allowing amplification and measurement limited by quantum fluctuations (see [81, 29, 82]). The action of $U_{sq}(z)$ is perhaps even more simply stated by its action on the ladder operators of the oscillator:

$$U_{sq}^\dagger(z)\;\hat{a}\;U_{sq}(z) = \hat{a}\cosh|z| - e^{i\arg[z]}\;\hat{a}^\dagger\sinh|z|, \tag{2.5}$$

and the action on $\hat{a}^\dagger$ is given by $\hat{a} \to \hat{a}^\dagger$ and $\arg(z) \to -\arg(z)$.

The next two control Hamiltonians (Eqs. 2.4c and 2.4d), commonly referred to as beamsplitting and two-mode squeezing in quantum optics, provide a way to add or subtract correlated photons in two oscillators. As a result, these interactions provide a way to couple oscillators for multi-mode control and create entanglement from any initial Fock state (including vacuum [83]). The beamsplitting Hamiltonian Eq. 2.4c in particular, with a controllable exchange interaction between the two modes, will play a central part in this thesis and be analyzed in detail later. For now, we simply list the effect of these controls

on the oscillators' ladder operators, as we did for U_{sq}:

$$U_{bs}^{\dagger}\ \hat{a}\ U_{bs} = \cos\theta_{bs}\ \hat{a} - i\sin\theta_{bs}\ e^{i\varphi_{bs}}\ \hat{b} \tag{2.6a}$$

$$U_{tms}^{\dagger}\ \hat{a}\ U_{tms} = \cosh r_{tms}\ \hat{a} + \sinh r_{tms}\ e^{i\varphi_{tms}}\ \hat{b}^{\dagger}. \tag{2.6b}$$

Here, $U_{bs}(\theta_{bs}, \varphi_{bs}) = \exp\left[i\theta_{bs}(e^{i\varphi_{bs}}\hat{a}^{\dagger}\hat{b} + h.c.)\right]$ represents the beamsplitter Unitary, characterized by the variables $\theta_{bs} = \int 2|g_{bs}(t)|dt$ and $\varphi_{bs} = \arg[g_{bs}]$. The exchange coupling resonantly hybridizes the two modes $\hat{a}$ and $\hat{b}$, with the modes fully exchanging their behavior at time $\frac{\pi}{2g_{bs}}$, which we will refer to as a SWAP gate. In fact, a simple representation of U_{bs} can be intuitively understood as a rotation on a Heisenberg-picture Bloch sphere, with poles corresponding to $\hat{a}$ and $\hat{b}$ – a representation that makes computing the effect of beamsplitter-based gates significantly easier (see [84, 85]). The second operation, $U_{tms}(z_{tms})$, represents a two-mode squeezing unitary, with $z_{tms} = r_{tms}e^{i\varphi_{tms}} = \int g_{tms}dt$. As an example, on an initial state with vacuum in both oscillators, this enacts a correlated squeezing of phase space, creating the entangled state $\cosh^{-1}(r_{tms})\sum_n (\tanh(r_{tms}))^n |n_a\rangle|n_b\rangle$.

The final control, g_{q-q}, is simply a sum of the beamsplitting and two-mode squeezing interactions with equal strength, ie with $|g_{q-q}| = |g_{bs}(t)| = |g_{tms}(t)|$, and phases $\phi_{a,b} = (\phi_{tms} \pm \phi_{bs})/2$. While this makes its representation through ladder operators clear, its structure actually allows a more intuitive description – it precisely enacts a coupling of the phase space quadratures of the two oscillators. This means it can be re-written as $H_{\text{contr.}}/\hbar = |g_{q-q}(t)|\ (\cos\phi_a p_a + \sin\phi_a q_a)\ (\cos\phi_b p_b + \sin\phi_b q_b)$, which emulates a resonant inductive or capacitive coupling at a $\phi_{a,b}$ of zero and π respectively.

The full set of control Hamiltonians in Eq. 2.4 has some interesting properties. The first is that they all preserve the (joint) photon-number parity in the oscillator(s), which means utilizing any combination of them for a quantum operation preserves parity as a syndrome

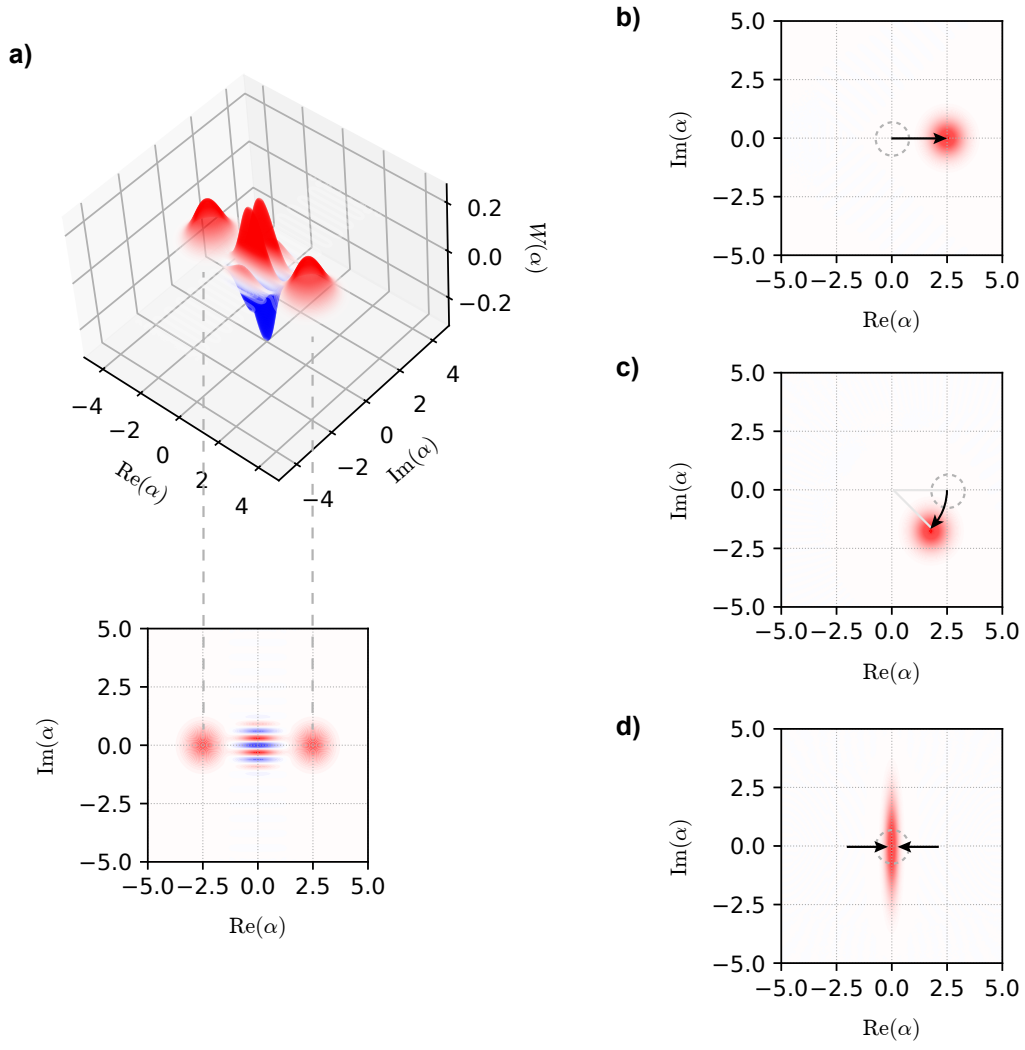

Figure 2.1: **Wigner functions and CV controls a)**, The Wigner function for complex bosonic (cat) state, visualized by a three-dimensional contour plot as a function of $\alpha \sim q + ip$, and its corresponding projection onto a two-dimensional color plot. **b, c, d)**, Examples of Gaussian control on coherent states, with each operation as a direct simulation of the drives in Eq. 2.4. The displacement of $\alpha = 2.5$ is enacted by $\epsilon_a = 2\pi \times 1$ MHz for $t = \alpha/\epsilon_a = 398$ ns. The rotation by $\phi_a = \pi/4$ is a result of a frequency detuning by $\delta\omega = 2\pi \times 1$ MHz that acts for 125 ns. Finally, the squeezing reduces the variance in the x-quadrature by 60%, through $g_{sq} = e^{i\pi/2} \times 2\pi$ MHz, enacted for 75 ns.

for whether an error has occurred. This is only broken by the linear displacement drive, which allows access to the entire Hilbert space. However, even with the ability to create entanglement and access the entire Hilbert space, one cannot always use these controls

to enact arbitrary quantum controls (Unitaries) – ie this set of controls is not necessarily universal. To test the construction of arbitrary Unitaries, it is useful to first understand how primitive controls can combine to take on more complex forms. Consider playing one Hamiltonian for time δt, followed by another. Their combined action is given by the Baker-Campbell-Haussdorf formula:

$$e^{iH_1 t}.e^{iH_2 t} = e^{i(H_1+H_2)t+\frac{1}{2}[H_1,H_2]t^2+\ldots}. \tag{2.7}$$

This means that a set of Hamiltonians can produce new effects if their commutators are non-trivial. Unfortunately, the controls in Eq. 2.4 can only produce other Hamiltonians within the same set (they form a closed Lie group). In fact, it can be shown that any initial state that is a Gaussian in phase space (like vacuum or a thermal state) can only transform into other Gaussians under these operations - i.e., these only perform symplectic transformations to phase space. This makes their non-universal nature obvious, but it also means that their effect on Gaussian states is efficient to classically simulate, since one only needs to keep track of the changes to the mean and variance of the Gaussian under any of these transformations. Thankfully, introducing a non-Gaussian initial state or measurement (often called 'resource'), which is trivial in a dispersively coupled cQED system, is sufficient to make these operations universal. More broadly, by a curious coincidence of mathematics, **any** control which is $O(a^3)$ or higher is sufficient to build arbitrarily high-order controls through commutators – the dispersive interaction $\sim \hat{a}^\dagger\hat{a}\hat{b}^\dagger\hat{b}$ simply provides one such control [86]. For the full set of circuit-QED bosonic controls available prior to this thesis, we refer the reader to [87] and A. Eickbusch's thesis [88].

Now that it's clear that there exist quantum states and controls that are best described using the phase-space or ladder operators, one may ask what the best way might be to visualize these states and their evolutions. It turns out that this question has a very concrete

answer, and leads to a parallel formulation of quantum mechanics that entirely bypasses the notion of the Hilbert space and state vectors [89]. Central to this theory is the representation of any quantum density matrix ($\hat{\rho}$) as a distribution in phase space:

$$W(x,p) = \frac{1}{\pi\hbar}\int_{-\infty}^{\infty}\langle x-y|\hat{\rho}|x+y\rangle e^{-2ipy/\hbar}dy. \tag{2.8}$$

This distribution, known as the Wigner function, conveys an effective density or occupation of phase space, except that the function can take on negative values. In fact, Wigner negativity is a signature of quantum behavior, and it disappears when the distribution is blurred (averaged) with a filter that has an area larger than a few $\hbar$, therefore returning classical behavior in the limit $\hbar \to 0$. One can produce a more directly interpretable probability density when the Wigner function is integrated over one quadrature:

$$\begin{aligned}\langle x|\hat{\rho}|x\rangle &= \int dp\; W(x,p)\\ \langle p|\hat{\rho}|p\rangle &= \int dx\; W(x,p),\end{aligned} \tag{2.9}$$

which leads to $W(x,p)$ being called a 'quasi-probability' distribution. The existence of the Wigner representation is highly non-trivial, as it represents a quantum state in two conjugate variables that don't commute and must satisfy the uncertainty relation with respect to each other. Yet, in combination with a special construction called the $\star$-product, it reproduces standard quantum theory in a manner that's much more directly comparable to the Poisson-bracket formulation of classical Hamiltonian mechanics. A Wigner function's time evolution is given by:

$$\frac{\partial W}{\partial t} = \frac{i}{\hbar}(W\star H - H\star W) = \{\{H,W\}\}, \tag{2.10}$$

where $\{\{A,B\}\}$ is the quantum analog of the Poisson bracket, named the Moyal bracket.

Readers interested in the full phase-space formulation of quantum mechanics can read [90] as an excellent reference on both the physics and the history of this subject.

In this thesis, the Wigner function will primarily be used as a way to visually represent and experimentally perform tomography of states. As such, the function has a much more direct representation in terms of available controls in circuit-QED; it is equivalent to a displaced parity measurement:

$$\begin{aligned} W(\alpha) &= \mathrm{Tr}\left[D_\alpha e^{i\pi \hat{a}^\dagger \hat{a}} D_\alpha^\dagger,\ \hat{\rho}\right] \\ &= \langle D_{-\alpha}\psi | e^{i\pi \hat{a}^\dagger \hat{a}} | D_{-\alpha}\psi \rangle \end{aligned} \tag{2.11}$$

for a pure state $|\psi\rangle$ that is displaced by a linear drive ($D_{-\alpha}$). Here, Re(α) and Im(α) are used interchangeably with p and q, and all the previous intuitions of phase space controls carry over. For example, Gaussian states like coherent or thermal states appear as Gaussian 'blobs' in Wigner space, and the displacement and squeezing transform precisely shift the mean and squeeze the variance of their Wigner functions. It is also useful to gain more intuition on the link between the Fock basis and phase-space. Each Fock state appears as a circularly symmetric Wigner function, since it does not have a preference for a particular quadrature, and its value at the origin is +1 (red) or -1 (blue) depending on whether the Fock state is even or odd. Higher Fock states have alternating positive and negative rings around the central blob, with the number of additional rings corresponding to the number of photons. This means that more complicated Wigner functions with finer details, where the function changes rapidly over a small portion of phase-space (sometimes called sub-Planckian features), require larger Fock states. These finer features set the sensitivity of the state to errors (eg, small displacements), which then translates to its error detection and quantum sensing capabilities – in general, these grow with photon number. Example Wigner functions showing some continuous variable states and controls are illustrated

in Fig. 2.1.

2.2 Bosonic errors and error correction

Now that we understand how to represent and manipulate oscillator-based information, let us take a look at the type of errors they can encounter. We will also introduce how information may be resiliently encoded in these oscillators, but details on how the errors may actually be corrected will be left to Appendix A.

If a linear superconducting resonator ($\hat{a}$) is coupled to a linear electromagnetic environment, the only possible noisy variables that could affect it are of the form:

$$H_{\text{noise}}/\hbar = F(t)\hat{a} + F^*(t)\hat{a}^\dagger. \tag{2.12}$$

This translates to a noisy fluctuation in both its quadratures, $\{q, p\}$, which for a symmetric noise spectral density $S_F[\omega]$ would result in small displacement errors. However, if the fluctuations in the EM modes are purely thermal (which is generally a good approximation), the noise spectral density is actually highly asymmetric, with $S_F[-\omega_a] = \bar{n}_{th}\ S_F[\omega_a]$. Here, $\bar{n}_{th}$ is the thermal population of the resonator's environment, which can generally be cooled down to $O(10^{-2} - 10^{-3})$. This instead results in a single dominant error channel, that of amplitude damping, given by:

$$\dot{\rho} = \mathcal{D}[\sqrt{\kappa_1}a]\hat{\rho} = \kappa\left(a\hat{\rho}a^\dagger - \frac{1}{2}\{a^\dagger a, \hat{\rho}\}\right). \tag{2.13}$$

The action of the damping channel can be more intuitively understood through separating the two parts of the master equation as follows. The first part, $a\hat{\rho}a^\dagger$, enacts 'jumps' in the quantum evolution of $\hat{\rho}$ where the system loses a single photon. These are discrete transitions that occur on the thermalization timescale of the environment, which for

a general environment is effectively instantaneous (much faster than any control or measurement rates). The time at which they occur cannot be determined deterministically, but the probability that they occur within a time δt is $p_{\text{jump}} = \text{Tr}\left[\hat{\rho}, a^\dagger a\right]\ \kappa_1 \delta t$. However, given that these jumps are uncorrelated, any finite time evolution can have more than one jump. In general, the probability of n jumps is given by $(p_{\text{jump}})^n$, and error-correction strategies must either protect against a large enough number of jumps such that the remaining errors become negligible, or use clever encodings that are only sensitive to a finite number of jumps.

Interestingly, when no jumps are occurring, the system still undergoes a loss of energy to preserve the rules of probability (keep the channel trace-preserving). This takes the form of a Bayesian update to the current state, resulting in the evolution $\hat{\rho}(t) = \mathbb{1} - e^{-\frac{\kappa_1}{2} a^\dagger a \delta t} \hat{\rho}(t) e^{-\frac{\kappa_1}{2} a^\dagger a \delta t}$. Note that the periods (or trajectories) when no jump is occurring are also probabilistically distributed, with probability $\sim 1 - p_{\text{jump}}$. However, during the time that it *is* evolving under this process, its evolution is highly smooth and continuous. In fact, it can be alternatively described exactly as a Schrödinger-like evolution under an effective Hamiltonian:

$$\dot{\rho}(t) = \left[H_{\text{eff}}, \hat{\rho}\right], \quad H_{\text{eff}}/\hbar = \left(\omega_0 + \frac{i\kappa_1}{2}\right) a^\dagger a. \tag{2.14}$$

This representation provides a simple way to calculate quantities like the inherited (Purcell) loss through hybridization, and engineered dissipation rates. More importantly, it makes clear that Hamiltonian evolution and the no-jump evolution are on similar footing, which means that the former can be used to protect against the latter, as we show explicitly in Appendix A.

The above approximation for linear oscillators and linear couplings holds reasonably well for isolated superconducting resonators, but of course does not allow universal quan-

tum computing. Superconducting devices must purposely introduce nonlinearity for control and measurement, through circuits like the transmon. Even without a purposeful nonlinear mode, given that these devices are realized in solid-state systems, they may have impurities like spins that couple to the resonator's magnetic field, and contribute nonlinear effects. Interacting with either the nonlinear ancillary circuit or a nonlinear environment can then introduce additional error mechanisms to the oscillator. These nonlinear errors can be understood by writing down the general Hamiltonian of the oscillator's interaction with a nonlinear mode:

$$H_{\text{nl}}/\hbar = \chi s^\dagger s \; a^\dagger a + g(s^\dagger + s) \; a^\dagger a + \frac{K_a}{2} {a^\dagger}^2 a^2, \tag{2.15}$$

where $(s + s^\dagger)$ is either the field of a quantum of classical mode, or the σ_x operator for an environmental spin, both of which are assumed to be noisy (have random fluctuations). The first term in this Hamiltonian is the dispersive coupling, which is a frequency shift of the resonator dependent on the state of s. If this state fluctuates, the resonator automatically inherits some frequency fluctuations, or dephasing noise. The noise spectrum for this dephasing technically has a bandwidth that is a function of χ and the fluctuation rate, but for any system with a strongly coupled nonlinear ancilla, the resonator dephasing is essentially Markovian, given by $\mathcal{D}[\sqrt{\kappa_\varphi} a^\dagger a]$. The second term also has a similar effect – here the resonator frequency depends on the value of a field given by $\phi_s = \langle s + s^\dagger \rangle$, and any fluctuations in this field impart some resonator dephasing. A common example of this is when the resonator is coupled to a flux-tunable device with a loop, where fluctuations in the magnetic field inside the loop can change the frequency of the resonator. Such flux-noise dephasing can be very non-Markovian and has been measured to generally have a noise spectrum that is dominantly low frequency, with $S_{\phi_s}[\omega] \propto 1/\omega$. In either form of dephasing, the Wigner function for the cavity state undergoes random rotations, which

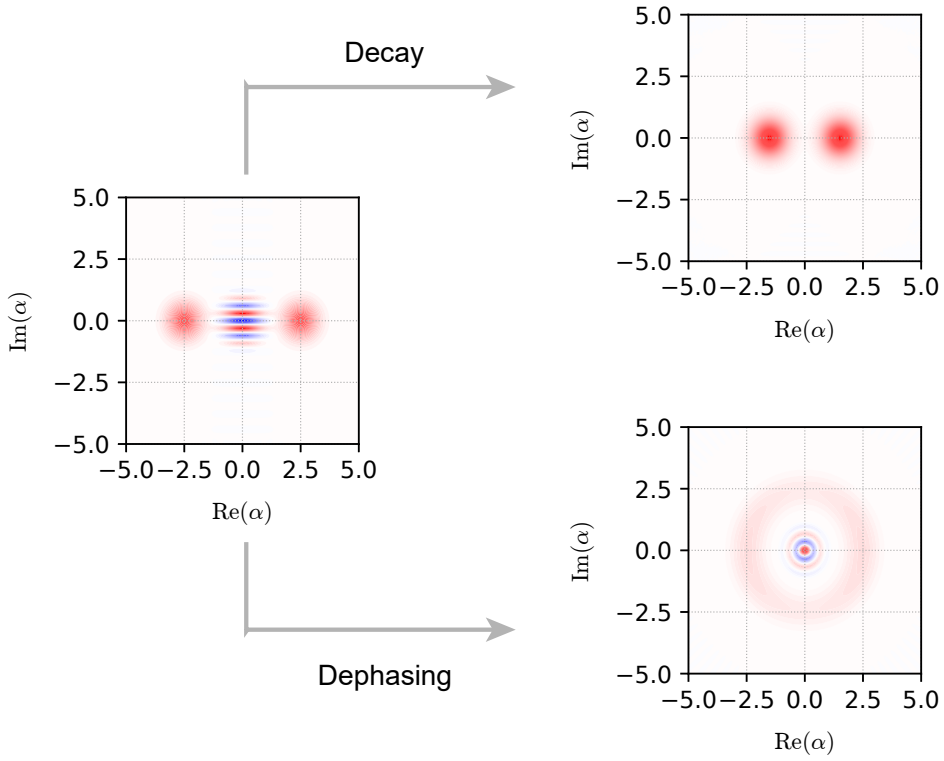

Figure 2.2: **Errors in an oscillator**, An example of photon-loss and dephasing acting on the Schrödinger cat state in an oscillator. The photon-loss rate and dephasing rate are set to $(1\text{ ms})^{-1}$ for each individual simulation, and the state is plotted at time $1/\kappa$. Decay makes the cat state quickly lose its fringes due to jumps, and shrink slightly due to no-jump evolution. Dephasing makes the cat spread out in phase, but does not spoil its fringes or affect its parity.

need to either be corrected or (for low-frequency noise) prevented through stabilization. Examples of both the decay and dephasing processes on a continuous variable (cat) state have been shown in Fig. 2.2. Finally, the third term in H_{nl} is an inherited resonator Kerr. While this Kerr is not a directly noisy variable, it can still be a source of error when it is not explicitly taken into account. As an example, most oscillator-based encodings are completely stationary in an appropriate rotating frame when the oscillator is linear, but have a slow evolution that decreases fidelity when they are subjected to such an inherited Kerr.

Now that we have a sense of the errors that can occur in an oscillator, let us look at some ways to encode oscillator information that allow recovery from these errors – i.e., bosonic error correcting codes. We start by zooming out and understanding a few statements about correcting errors in general. Any system's evolution over time, in the presence of noise, can be cast into the language of one person (Alice) sending a message to another person (Bob) through a noisy communication 'channel'. If the channel is not too noisy, one can technically send the same message through the channel multiple times, or some other redundantly encoded version of the message, and be able to recover the original information completely. The rate at which a noisy channel can send information such that it is still recoverable to arbitrary precision is well-defined, and known as the Shannon bound on channel capacity. A good encoding and decoding strategy should make full use of this channel capacity (saturate the Shannon bound). However, the bounds on a quantum channel are much harder to define, especially given that one is not allowed to create direct copies of an arbitrary quantum state (the no-cloning theorem). Quantum channels also have interesting properties; for example, a channel's capacity to transmit recoverable information is also its capacity to generate entanglement between Alice and Bob, and even channels with individually zero capacity can have non-zero capacity when combined. Overall, quantum codes must still build in some form of redundancy, but these must not violate the no-cloning theorem. It turns out is indeed possible to do this, as long as the encoding satisfies special conditions with respect to the channel's errors, known as the Knill-Laflamme conditions [91]. For special types of error channels, like a channel composed of uncorrelated Pauli errors, one can find the channel (Hashing) bound by considering random encoding strategies. For a more general error channel, like the amplitude damping + dephasing channels common in oscillators, finding this bound is much more difficult [92]. However, one can still efficiently numerically evaluate an optimum recovery channel given these errors, which while generally impractical to achieve using available

control, prove the existence of 'good codes ' [93].

Let us now look at some actual error-correcting codes. We first note that all the errors listed earlier in this section also apply when these circuits are used as physical qubits (two-level systems), like in the transmon. This is important because qubit-based error-correcting codes are primarily tailored to an error model where the errors are Pauli jumps, which have no no-jump evolution ($\sigma_i^\dagger \sigma_i = \mathbb{1} \forall i$). These encodings generally use a grid of qubits, where the logical information is spread over a many-qubit wavefunction, and local errors can be detected through measuring strings of Pauli operators, known as stabilizers. To preserve logical information in this simplified model, one does not even need to necessarily fix the errors once they happen, since there is no way to jump to a non-correctable state. Instead, one can simply keep track of frame changes that result from the Pauli errors, and apply the frame changes to subsequent gates in the computation. However, larger grids of qubits always introduce more errors from the additional qubits – which raises the important question, could such an approach ever work? One of the most important theorems in quantum error correction is that there indeed exists a threshold probability for errors that occur within a single stabilizer measurement cycle, below which the error-correcting code always performs *exponentially* better with more physical qubits [94]. Given that realistic errors look like the amplitude damping channel mentioned above, or leakage due to finite Kerr effects, special changes need to be made in physical implementations of such codes (see Pauli twirling [95] and leakage reduction units [96]) to get them to show realistic improvements [97].

Bosonic error-correcting codes are fairly different – their potentially infinite Hilbert space makes them significantly more susceptible to leakage to uncontrolled error states. As such, it is currently unknown whether a threshold theorem for encoding logical qubits in oscillators even exists, unless it is effectively concatenated with a qubit-based encoding that has a threshold! However, utilizing continuous variables and higher-dimensional

spaces also allows unique encodings that can directly combat complex errors like amplitude damping, even within a single oscillator. The most direct example of this is in how coherent states interacts with photon loss. A coherent state α, prepared simply by a linear drive on a linear oscillator, is an eigenstate of the annihilation operator: $a|\alpha\rangle = \alpha|\alpha\rangle$. This is very un-intuitive to someone who does not routinely work with such states, since the coherent state has a well defined average energy, $E = |\alpha|^2\hbar\omega$. Yet every time it loses a photon, its average energy remains unchanged. There is no discrete variable state that has this property – indeed, one requires non-zero (but exponentially decaying) overlap with infinite Fock states to exactly replicate this behavior. The coherent state still loses energy through the amplitude damping channel (as, necessarily, does any state that is not vacuum), but it is only through the channel's no-jump evolution: $e^{-\kappa_1 t a^\dagger a}|\alpha\rangle = |\alpha e^{-\kappa_1 t}\rangle$. This is again very unintuitive since it is precisely when you don't observe a photon leave the system that you can be confident that the system has lower average energy.

Practically, this means that coherent states are enormously useful when trying to encode quantum information. Since a single state does not contain any information, consider instead an encoded qubit described by a Bloch sphere with the $|0\rangle, 1\rangle$ states given by the pair of coherent states $|\pm\alpha\rangle$. Under photon loss jump errors, each coherent state is individually preserved, but their relative phase incurs jumps: $a(x|\alpha\rangle+y|-\alpha\rangle) = \alpha(x|\alpha\rangle-y|-\alpha\rangle$. This is easier to interpret by noticing that both $|\alpha\rangle \pm |-\alpha\rangle$ have a well-defined photon number parity of ± 1 (as can be seen from their Wigner functions), and losing a photon always exactly flips this parity. Photon jumps on this Bloch sphere then cause practically no error in states near the poles, but a full logical error near the equator – thus providing an encoded qubit with intrinsically biased jump errors (bit flips $\ll$ phase flips). The only requirement to realize this qubit is then to stabilize it against the no-jump evolution, which is possible because of its smooth and continuous nature. Dominant methods use either an engineered dissipation or a Hamiltonian whose steady state is this encoding [34, 98],

and Appendix A explains precisely how these work. Note that $|\alpha\rangle$ and $|-\alpha\rangle$ are not truly orthogonal states, which means there is still a finite probability for a bit-flip transition for any finite $|\alpha|$. However, the overlap of the states, and therefore their bit-flip rates, grows exponentially smaller with larger $|\alpha|$ ($\propto e^{-2|\alpha|^2}$), while the additional photon loss rate (and therefore phase-flips) only grow linearly. Readers interested in more details can read N. Frattini's thesis [99].

The above encoding is part of a family of encodings known as Schrödinger cat-codes, with this encoding in particular named the 2-component cat (or 2-cat, for short). To understand how this encoding can be extended to actually perform error correction, consider two separate 2-cat codes themselves as the two logical states [100, 93]:

$$\begin{aligned} |0_C\rangle &= |C_\alpha^+\rangle = |\alpha\rangle + |-\alpha\rangle, \\ |1_C\rangle &= |C_{i\alpha}^+\rangle = |i\alpha\rangle + |-i\alpha\rangle. \end{aligned} \tag{2.16}$$

These two states are exactly rotated versions of each other in phase space, but for high enough $|\alpha|$, can be orthogonal enough to use as code words in an encoding named the 4-cat code. Importantly, these states now both have positive parity, which means a photon loss event from any of these states moves them precisely to their negative parity counterparts $\{|C_\alpha^-\rangle, |C_{i\alpha}^-\rangle\}$ (see Fig. 2.3). In fact, the states on the equator of the 4-cat Bloch sphere actually have well-defined 4-parity ($n \mod 4$), with the states denoted $|\pm_C\rangle = |0_C\rangle \pm |1_C\rangle = |C_\alpha^{\{0,2\}}\rangle$. Each subsequent photon loss changes the 4-parity of these states, moving in a cyclic manner through $\{|C_\alpha^0\rangle, |C_\alpha^3\rangle, |C_\alpha^2\rangle, |C_\alpha^1\rangle\}$ – a phenomenon only possible in a continuous variable encoding. This means that the photon loss jumps are correctable, but one can equivalently just keep track of the parity jumps without correcting them, as in the case of qubit Pauli errors. Importantly, this assumes that the jumps are caught in time – if two jumps occur before the detection step, then the error syndrome (2-parity) remains unchanged, but a logical error occurs. In other words, the 4-cat is said

to be correctable (satisfy the Knill Laflamme conditions) against a single photon-loss error [101]. Note that similar to the 2-cat, the 4-cat will also require stabilization against the no-jump evolution [100].

Observing at the phase-space representation of the 2-cat and 4-cat codes reveals an interesting structure, where both their Wigner functions have rotational symmetry. This is, of course, no coincidence, since we constructed them using coherent states of equal magnitude and equally spaced out phase. However, it turns out that there exists a different way to generalize such rotation-symmetric codes [102, 103], by instead focusing on their number-phase representation. Since number and phase are conjugate bases, if one defines an encoding where a rotation by some fraction of π enacts a Z-rotation, ie $Z_N = e^{i\frac{\pi}{N}a^\dagger a}$ (with $Z_N^2 = \mathbb{1}$), then this necessarily implies that the logical codewords always have a photon-number parity N. In equations, this means that:

$$\begin{aligned} |\pm_C\rangle &= \frac{1}{\mathcal{N}_\pm}\sum_{m=0}^{2N-1}(\pm 1)^m e^{i\frac{m\pi}{N}a^\dagger a}|\Theta\rangle \\ &= \frac{1}{\mathcal{N}_\pm}\sum_{k=0}^{\infty} f_n|n_\pm\rangle, \quad n_\pm = \{2kN, (2k+1)N\}. \end{aligned} \tag{2.17}$$

Here $|\Theta\rangle$ (and correspondingly f_n), represents (Fock space components of) the state that is being rotated, which in the cat codes is the coherent state $|\alpha\rangle$. This representation makes the code's resilience to photon loss errors clear – the code moves to an orthogonal error state with different N-parity every time it incurs a photon loss, until it has N loss errors, which then cause a logical X_L error. Similarly, depending on the angular spread of $|\Theta\rangle$ (which is determined by its average photon number), the state is resilient to phase errors up to approximately a rotation of $\pi/2N$.

This idea of symmetry and error resilience can also be translated to the position-momentum basis, forming translational symmetric bosonic codes. The most popular among these is an encoding where the logical codewords form a grid in phase-space, named the

'GKP' code after its inventors Gottesman, Kitaev and Preskill [104]. Similar to the cat codes, where rotated copies of a coherent state are summed over to form the codewords, one can define a (square) GKP code by summing over translated copies of a coherent state:

$$\begin{aligned} |\mu_C\rangle &= \sum_{n_1,n_2} |\alpha^{\mu}_{n_1,n_2}\rangle \\ |\alpha^{\mu}_{n_1,n_2}\rangle &= D_{\sqrt{\frac{\pi}{2}}(2n_1+\mu)}\, D_{i\sqrt{\frac{\pi}{2}}n_2}\, |0\rangle, \end{aligned} \tag{2.18}$$

where $\mu = \{0, 1\}$ represent the Z_L eigenstates. This definition also makes it immediately clear that enacting a logical X_L operation corresponds to displacing the lattice along the real axis by $\sqrt{\frac{\pi}{2}}$ – and it turns out, a Z_L operation is correspondingly a displacement along the imaginary axis by the same amount. Combined with a $CNOT$ gate that can be realized by the quadrature coupling, this means that all Clifford gates for this code are composed of Gaussian operations. However, these ideal phase space lattices are non-physical. Since total probability in phase space is conserved, spreading the lattice out over all of phase space ($n_1, n_2 \in \mathbb{Z}^2$) reduces the wavefunction at each lattice point to a delta function. The average photon number of this ideal encoding is infinite, as is the rate at which it incurs errors, making it non-correctable. A more realistic encoding is then the same lattice with a Gaussian envelope, given by:

$$|\mu^{\Delta}_C\rangle = \sum_{n_1,n_2} e^{-\frac{\pi}{2}\Delta^2\left[(2n_1+\mu)^2+n_2^2\right]}\, |\alpha^{\mu}_{n_1,n_2}\rangle. \tag{2.19}$$

In effect this leads to a finite squeezing of each blob that makes up the lattice, and an overall Gaussian envelope that confines the code in phase-space. Any logical gate on this finite-energy code then necessarily requires non-Gaussian resources, since the envelope does not commute with Gaussian operations.

Similar to the rotation-symmetric codes, the translation symmetry of the GKP code

allows it to correct a finite number of small (displacement) jumps in both conjugate variables of position and momentum. Since these displacements are related to the logical Pauli operations, both the errors and their correction follow the same stabilizer-like description that holds for qubit-based codes [105]. However, as discussed previously, displacement errors are in general not a good description of the errors in a superconducting oscillator – in any environment with finite temperature, they primarily face photon loss and dephasing errors. It turns out that GKP codes are surprisingly resilient to photon loss errors. In some sense, being able to correct both $(a + a^\dagger)$ and $(a - a^\dagger)$ jumps also lets it correct errors that are just a and $a^\dagger$, as can be seen from a density matrix that has suffered half of each error:

$$\begin{aligned} \rho_\epsilon &= \frac{\kappa}{2}\left(\frac{a+a^\dagger}{\sqrt{2}}\,\rho\,\frac{a+a^\dagger}{\sqrt{2}} + \frac{a-a^\dagger}{\sqrt{2}}\,\rho\,\frac{a-a^\dagger}{\sqrt{2}}\right) \\ &= \frac{\kappa}{2}\left(a\rho a + a^\dagger \rho a^\dagger\right), \end{aligned} \tag{2.20}$$

where ρ_ϵ is the part of the density matrix that has incurred a jump error at rate κ. A more intuitive way to understand the resilience of this code to photon loss is to realize that it has even parity (the Wigner function at the origin is always $+1$, see Fig. 2.3). Each subsequent photon jump then flips the code's parity, but does not take both codewords to perfectly orthogonal states if the code has finite energy. In fact, the orthogonality of its error spaces grows exponentially with the squeezing parameter Δ, but so does its average photon number and photon-loss rate, resulting in an optimum Δ for any given κ [106]. Surprisingly, this holds for not just single jump errors, but for any number of photon jumps, which leads to the GKP code being 'optimal' against photon loss errors, assuming an ideal recovery channel [93, 92]. This has sparked a large amount of interest in GKP error correction, and despite this ideal recovery channel being a numerical construction that may be impractical to achieve with available controls, has led to many successful near-term demonstrations [107, 108]. However, despite the GKP code's resilience to amplitude

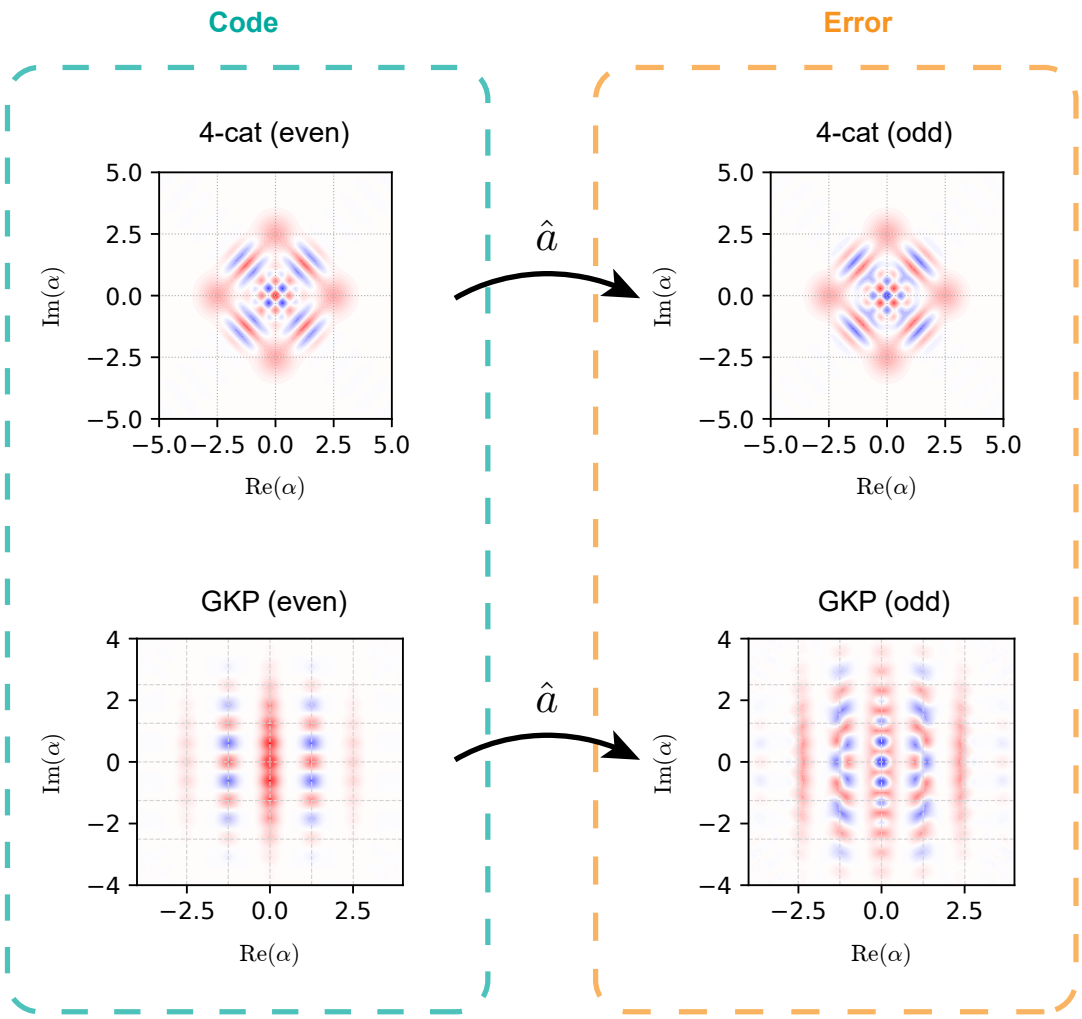

Figure 2.3: **Encodings in an oscillator**, Examples of a rotationally symmetric (4-cat, $\alpha = 2.5$) code state and a translational symmetric (GKP, $\Delta^{-1} = 2.5$) code state. Both have well-defined photon-number parity (even), which flips under a photon-loss error, which allows both the detection and correction of such errors.

damping, it has so far faced limited success since it is fairly fragile against dephasing and Kerr errors, which are difficult to avoid in circuit-QED experiments.

The above continuous variable codes have rich physics and many interesting properties worth academic study. As an interesting example, consider that the resilience of the above 'coherent-state constellation' type codes to certain types of errors grows exponentially with photon number, and therefore with the available Hilbert space size (effective distance

$d \propto \bar{n} \sim \dim[\mathcal{H}]$). This is in stark contrast to qubit-based encodings like the surface code, where resilience to Pauli errors grows only linearly with Hilbert space size, with distance $d \propto \sqrt{N_{\text{qubits}}} \sim \log(\dim[\mathcal{H}])$. This means that for the same Hilbert space size, these bosonic codes provide exponentially more protection than qubit-based codes. However, this does not give any immediate advantage in a practical setting, since achieving a Hilbert space size of ~ 1000 only takes about 10 qubits, but is practically intractable in any single oscillator. Since protecting and controlling these continuous variable codes for high photon numbers inevitably leads to difficulties, one can instead ask the question – what happens if we try to encode information in a truncated portion of the Hilbert space instead? The answer to this question lies in discrete variable encodings, of which we will discuss the binomial and dual-rail codes below.

The minimum requirements (KL conditions) on any logical encoding that protects against photon loss, are that the codewords move to orthogonal error spaces under photon jumps, and have equal average photon number. The former ensures that the jump errors can be detected and corrected, and the latter that any evolution under the effective Hamiltonian $\left(\omega + \frac{i\kappa}{2}\right) a^\dagger a$ does not shear the logical information. It turns out that these conditions are perfectly achievable using just a few photons in one (or a few) oscillator(s). Perhaps the simplest way to realize this is to directly numerically optimize the coefficients of available Fock states for these conditions, but these can lead to encodings that are much harder to control and interpret [109]. A more intuitive class of error correction codes that satisfy these conditions exactly are the binomial codes, which were derived at Yale shortly before my arrival [102]. These codes are essentially the rotation-symmetric codes introduced above with truncated photon number ($f_{n>N} = 0$ in Eq. 2.17), and with well-defined number parity that changes with photon loss errors. The simplest of these is a truncated

version of the 4-cat code, called the kitten code:

$$|0_C\rangle = \left(|0\rangle + |4\rangle\right)/\sqrt{2}, \quad |1_C\rangle = |2\rangle. \tag{2.21}$$

The behavior of this code is very simple to understand; it clearly stores information in even 4-parity codewords, with equal average photon number, which jump to odd 4-parity states ($|3\rangle$ and $|1\rangle$) under a photon loss error. Detecting these jumps simply requires being able to detect 2-parity for states with up to 4 photons. Being able to selectively change the phase of Fock $|2\rangle$ is sufficient for enacting arbitrary Z_L rotations, and X_L rotations require more complicated number-changing pulses. The no-jump evolution of this state moves the entire Bloch sphere towards the vacuum state $|0\rangle$, but this evolution can be combated through the general strategies discussed in Appendix A.

Overall, the above oscillator encodings have all been confined to a single oscillator, and have directly provided resilience to either errors along one axis or on the entire Bloch sphere. However, for manageable photon numbers, these encodings usually can only correct one (or a few) photon loss and dephasing errors. Even if they were to perfectly correct one error with no additional overhead, this would at most square the effective error rate; ie $p_{\text{log.}} \sim (\kappa\delta t)^2$. For typical bosonic setups, the speed of control limits $\delta t \sim 1\mu$s, and typical superconducting resonators with attached ancillae have $\kappa \sim 1$ms, which means that $p_{\text{log.}} \sim 10^{-6}$. This would be a huge improvement over typical physical qubits like transmons – yet there are a number of difficulties that need to be taken account. The first is that this is only the error rate for a quantum *memory*, and actual computation would require error-corrected quantum gates, which might have a much lower fidelity and be difficult to implement. The second is that while these error rates are low, they are still many orders of magnitude higher than typical estimates for the required error rates for useful quantum algorithms ($p_{\text{log.}} < 10^{-10}$ to decode RSA-2048 [110]). This means that these bosonic en-

codings would still need to be concatenated with higher-level qubit-based encodings that have a threshold and can be scaled to reach these required error rates. However, bosonic qubits come with a large amount of overhead in qubit complexity, since each of these encodings requires active error correction and complex controls to activate gates and logical measurements. This then begs the question – can one still have oscillator-based encodings that provide similar gains when concatenated with a higher-level code, yet retain controls and measurements that have minimal complexity? Recent work has answered this question in the affirmative, with a broader class of error correction strategies that utilize **erasure errors** [111, 112]. In these, the errors in the physical system only need to raise a detectable flag, that can then be used by a higher-level surface code to provide a very similar advantage to codes that completely correct a single error at the physical level. This thesis contains the first high-fidelity demonstration of such an 'erasure qubit' in circuit-QED – the dual-rail cavity qubit, which is the final encoding we will introduce here.

The dual-rail cavity qubit [113] breaks the norms for oscillator encodings discussed so far, as it is neither in a single oscillator nor is it intrinsically error-correctable. It consists of a single photon living in two oscillators (denoted a and b), with the simple logical encoding:

$$|0_C\rangle = |0_a, 1_b\rangle, \quad |1_C\rangle = |1_a, 0_b\rangle. \tag{2.22}$$

It is clear that the average photon number of both code states are equal. However, a single photon loss event does not take the these codewords to orthogonal states. Instead, it takes any state on the effective Bloch sphere precisely to vacuum, where both oscillators remain until they are purposefully brought back through a reset protocol. This means that the qubit can raise a clear flag through vacuum checks by an ancillary observer (like a dispersively coupled transmon), and once an error occurs, it stays benign and does not propagate errors to the rest of the concatenated code. The single-qubit gates on this encod-

ing are also surprisingly straightforward – they simply result from turning on an exchange (swap) interaction between the two oscillators, ie through beamsplitter-based gates, and two-qubit gates are enabled by controlled-swap or controlled-phase gates. The dual-rail qubit is by no means a concept specific to circuit-QED, instead being first invented for linear optical quantum computing to enable qubit operations through physical beamsplitters. However, superconducting cavities provide an excellent combination of noise bias (dephasing $\ll$ photon loss $\Rightarrow$ Pauli errors $\ll$ erasure errors), high fidelity gates (demonstrated in this thesis) and error detection fidelity, which allows this qubit to particularly shine in its circuit-QED implementation [113]. A detailed study of the dual-rail cavity qubit, including proposals for its implementation in a fault-tolerant architecture, can be found in J. Teoh's thesis [114].

Overall, every bosonic encoding discussed here requires a particular error structure, and particular forms of oscillator control and measurement that don't spoil that error structure. The next section explores whether such 'error-preserving' control and measurement is possible, and goes over typical circuit-QED implementations of bosonic architectures.

2.3 Controlling an oscillator without dephasing it

Let us take a look at available bosonic control architectures in circuit-QED, and consider what knobs are available and what drawbacks they may have. A minimal bosonic device (see Fig. 2.4a) is composed of a tiling of high-Q superconducting resonators (often 3D stub cavities), that are coupled to a control chip containing a non-linear ancilla (typically a transmon qubit) and a low-Q readout mode (typically a stripline resonator). In a 3D architecture, these high-Q cavities will generally be monolithic with the outer superconducting package, and have their fields confined around a $\lambda/4$ stub, where λ is the wavelength of its fundamental mode. The field then mostly lives in vacuum, but has some participation in

the control chip, and therefore in the ancilla transmon. The transmon also has participation from the readout mode, which is strongly coupled to an external port with characteristic impedance 50Ω. Here, strongly coupled implies that the readout photons that are lost through the port dominate the readout mode's overall loss rate, which is critical for the maximal extraction of information from the readout mode.

As before, the hybridizations of the oscillator and readout modes in the qubit can be written down in terms of their effective linear coupling strengths g_{qA}, g_{qr} and frequency detunings Δ_{qA}, Δ_{qr}. These linear couplings, along with the effective energy decay rate κ_r of the readout, can be derived from an analysis of their linear equivalent lumped-element circuit similar to Fig. 2.4b:

$$\begin{aligned} g_{12} &\approx \frac{C_g}{C_1 C_2}\, q_{zpf,1}\, q_{zpf,2} \\ \kappa_{1,2} &\approx \left(\frac{C_g}{C_{1,2}+C_g}\right)^2 \frac{2R}{L_{1,2}}, \end{aligned} \tag{2.23}$$

for any two modes described by $L_{1,2}$ and $C_{1,2}$ that are coupled via a capacitance C_g either to each other, or to a resistance R, and $q_{zpf,i} = 2e\ n_{zpf,i}$. However, this circuit representation is only useful as a visual representation and for some design intuition – practically none of these devices actually obey lumped element physics, and analytic calculation of their participations and decay rates is in general infeasible. Instead, in practical experiment design, these can be extracted much more directly from fully numerical 3D EM simulations performed in High-frequency Simulation Software (HFSS). Specifically, these simulations provide the normal modes of the system, which include their field distributions and decay rates through coupling to ports with well-defined impedance. Exciting one of these modes and measuring the corresponding voltage drop across the transmon junction directly provides the mode's participation [76]. Additionally, a well-defined voltage drive at a port enacts a corresponding phase drop across the junction, providing the

frequency-dependent linear driven response of the circuit. As discussed previously, this response provides both the effective drive strength for any driven process, and the ability for external noise at the port to couple into the system.

How exactly does this architecture store, manipulate, and measure quantum information? We will see that while the operations on the qubit and oscillator and their evolution are best described in a quantum optics language, the readout and control drives on the circuit are much easier to formulate through classical microwave physics, clearly displaying the intersectional regime of this experimental platform. Let us start with the readout port, which while necessary for both drives and measurement, also allows the transmon and the oscillator to spontaneously decay through the port. This spontaneous decay is fundamental to quantum physics, leading to formulations like the Fermi-Golden rule and exotic behavior like super-radiance – but here we can simply treat (and cure) it through the language of microwave design. The readout port primarily requires access to only photons at the readout frequency. This means that the density of states available through that port can be narrowed to a small frequency window that contains the readout frequency. Specifically, one can add a bandpass Purcell filter at the port that prevents transmon and oscillator photons from leaking out through it, significantly increasing their lifetimes [115]. One can also use alternative strategies for mitigating this decay, like utilizing the geometric variation in the transmon's (readout's) electric fields to create a node (anti-node) at the port [116]. In fact, breaking this trade-off between imposed (Purcell-limited) lifetimes and the coupling to the port is critical for effective reset and error correction. To still be able to control these two elements, one often adds custom-filtered under-coupled drive lines to both the oscillator and the ancilla transmon. This then allows a direct linear displacement on the oscillator, and Rabi oscillations on the transmon. It turns out that these two (complex-valued time-dependent) controls are sufficient to universally control and read out the joint oscillator-ancilla system [117].

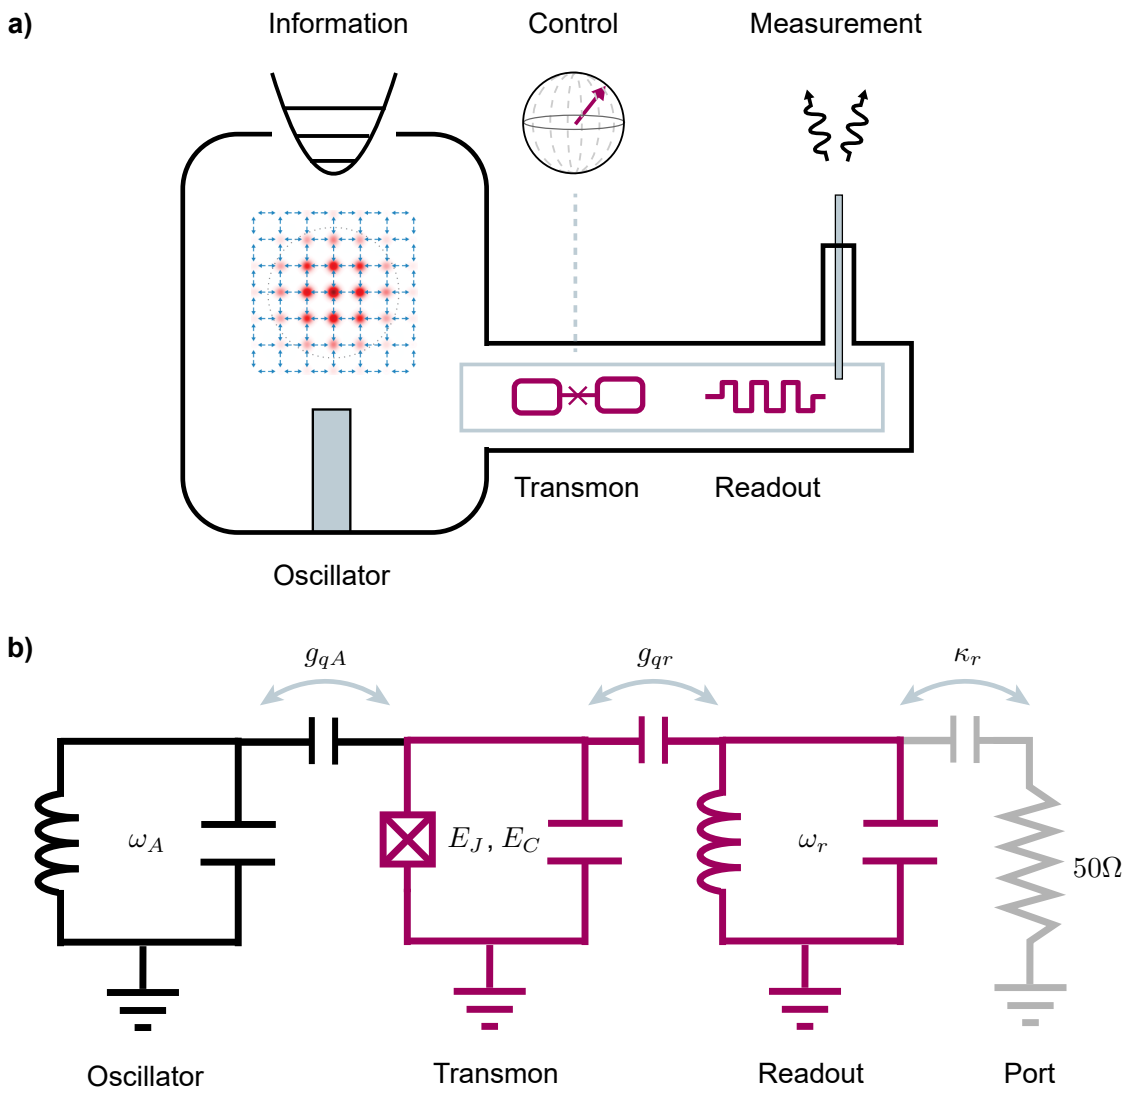

Figure 2.4: **Bosonic control in circuit-QED. a)**, a minimal setup for controlling CV information (inset Wigner from [108]) in a superconducting oscillator (here a 3D $\lambda/4$ stub cavity), using a control chip that contains an ancillary qubit (here a transmon) and an on-chip readout stripline mode that is coupled to a measurement port. Any additional drive ports are not shown. **b)**, lumped-element circuit representation for the architecture shown above. Both the 3D oscillator and the readout mode participate in the qubit and correspondingly have a dispersive shift to it. The measurement port effectively acts as a resistive (50Ω) termination. The decay rate through this resistor κ_r has purposely been shown as two-sided since external drives and noise can also couple in through this port at the same rate.

Central to manipulating quantum information in this system is the dispersive shift of both the oscillator and the readout with respect to the ancilla's state [72]. As before, these are approximately given by $\chi_{qA,\,qr} \approx \beta^2_{qA,\,qr}\, K_q$, where K_q is the Kerr-induced shift of the ancilla's $|e\rangle \leftrightarrow |f\rangle$ transition with respect to its $|g\rangle \leftrightarrow |e\rangle$ transition. Both these dispersive shifts are typically on the order of $2\pi \times 0.5 - 5$ MHz, which places some bounds on the time for an average control or readout operation on the system to be $\sim$ 100 ns-1 μs. A full formulation of χ and other non-linearities that result from the resonator-ancilla hybridization will be worked out in Chapter 5.

Given the drives and the non-linear coupling, how does one actually perform control and readout of both the ancilla and the oscillator? First, let us consider just the ancilla qubit. A single dedicated drive that is at the ancilla's transition frequency can enact arbitrary control on it through Rabi oscillations. The primary caution one must take here is to not have too high a drive amplitude, which may cause additional phase shifts and leakage due to the higher levels of the ancilla [118]. The more difficult task is then measuring its state, which occurs through the dedicated readout mode, whose frequency dispersively shifts depends on the ancilla's state. If a probe tone is sent into the readout port, it can reflect off of the port with a phase that depends on its detuning to the readout frequency, and the readout's decay rate. Specifically, this takes the form of a 'phase roll' [119, 120]:

$$\theta_r = \tan^{-1}\left(\frac{2\Delta_{pr}}{\kappa_r}\right), \tag{2.24}$$

where Δ_{pr} is the detuning between the probe tone (ω_p) and readout frequency (ω_r), with θ_r covering a full π phase shift as one sweeps the probe from $\Delta_{pr} \ll -\kappa_r$ to $\Delta_{pr} \gg \kappa_r$. Due to the dispersive shift, Δ_{pr} will have a difference of χ between the $|g\rangle$ and $|e\rangle$ states, allowing a sufficiently averaged measurement (for eg, on a Vector Network Analyzer, or 'VNA') to always distinguish between these states if χ is comparable or bigger than κ.

Specifically, for a fixed probe frequency halfway between the two frequencies, the reflected voltage will pick up a phase difference of $2\tan^{-1}(\chi_{qr}/\kappa_r)$ between the two ancilla qubit states. However, to study quantum dynamics, one tries to maximize the rate of information extraction about the ancilla state, especially since its inherent decoherence may change its state during the measurement. This can be shown to depend on the steady-state photon population in the readout mode and the maximum information gain per photon, which in steady state is exactly the measurement-induced dephasing rate of Eq. 1.13 [72]. Note that realistic qubit readout is much more complicated, since the exact frequency configuration of the system and environment can set hard limits on the measurement drive due to spurious driven transitions [121], and optimal readout pulse shapes may spend very little time in steady state. Studying and optimizing qubit readout (and reset) is still an open problem, and this includes the amplification chain that follows the initial dispersive measurement. This is in particular interesting because the high 'single-shot' non-demolition readout fidelity achievable in circuit-QED allows one to probe fundamental quantum behavior, like the transition from weak to strong measurements, and measurement-induced phenomena like the Zeno and anti-Zeno effects. For an overview of measurements in circuit-QED and their implications, the interested reader may try the theses of D. Sank [122], Z. Wang [123], and V. Joshi.

The above combination of ancilla drives and readout enables the ancilla to be used as an effective qubit for the primary goal of controlling and measuring the oscillator. To understand this, it is easiest to study some examples of measurement and control protocols. First, consider the driven response of the qubit in the presence of the dispersive shift:

$$H_q/\hbar = \left(\omega_q^0 + \chi_{qA}\, a^\dagger a\right)\sigma_z + \epsilon_{qx}(t)\sigma_x + \epsilon_{qy}(t)\sigma_y, \tag{2.25}$$

where $\epsilon_{qx,qy}$ represent the two quadratures (or real and imaginary parts) of a Rabi drive

on the qubit. Clearly, if the oscillator is in a well-defined Fock state $|n\rangle$, the qubit drive is only resonant at $\omega_q^0 - n\chi$. If one calibrates the amplitude of the drive $|\epsilon_{qx}|$ such that in some fixed pulse time t_p it exactly flips the state of the qubit (a π_x pulse), then one can use such a pulse to detect which state the oscillator is in [124]. Specifically, if the pulse is a Gaussian with a standard deviation of t_p, then the qubit will flip its state only within a Gaussian centered at $\omega_q^0 - n\chi$, with a standard deviation $2\pi/t_p$, which corresponds to the bandwidth of the pulse. If this bandwidth and the qubit's intrinsic decoherence rate γ_2 are both sufficiently smaller than χ, then the pulse can directly measure (and hence prepare) the oscillator's state in the Fock basis. We will call such a (narrow-bandwidth π_x) pulse a 'selective' π pulse, and the regime for such qubit-oscillator control the number-splitting regime.

For measurements like the parity syndromes in the bosonic codes discussed above, one would ideally like to not fully measure the oscillator state in the Fock basis, which would collapse any logical information. Instead, it would be much more convenient to directly measure just the photon number parity [125] – which would then, combined with an oscillator displacement, also provide a direct measurement of the oscillator's Wigner function! It turns out this is indeed possible, through a conveniently simple pulse construction. Let us work in a frame rotating at the nominal qubit frequency ω_q^0, and assume all drives and dynamics are described in this frame. If the qubit is prepared in the $+Y$ state with a calibrated $\pi/2$ pulse, then it will start to rotate around the equator at the qubit's (number-dependent) frequency, picking up phase $\varphi_q(t) = \chi_{qA}\ a^\dagger a\ t$. This means that at a time given by $t_p = \pi/\chi_{qA}$, the Bloch vector for all the even Fock states in the cavity line up back at $|+Y\rangle$ ($\varphi_q = 2m\pi$), while the odd Fock states line up at the flipped state $|-Y\rangle$ ($\varphi_q = (2m+1)\pi$). A final $-\pi/2$ pulse then puts the qubit back on the measurement axis, landing in $g\rangle$ or $|e\rangle$ states for even and odd parity respectively, providing a photon-number parity measurement through the qubit. As an important caveat, both the initial and final

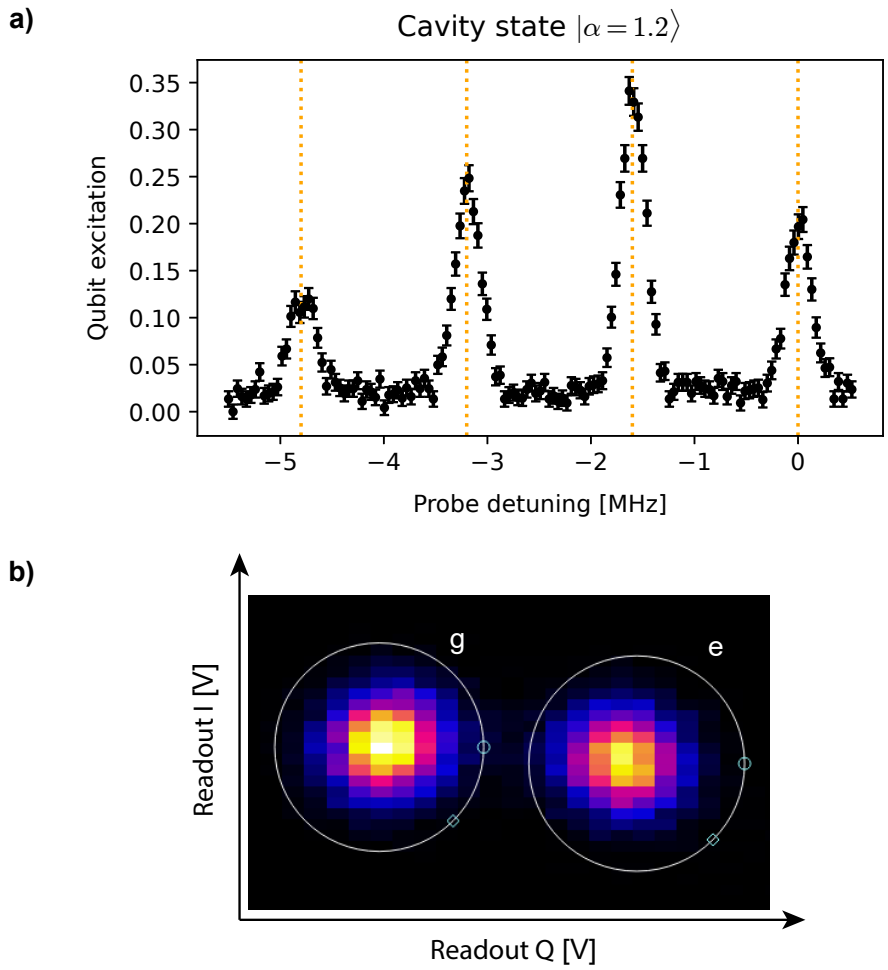

Figure 2.5: **Dispersive shifts for control and measurement a)**, The number splitting regime in circuit-QED, measured in an experimental setup with calibrated selective π pulses ($\sim$ 1000 averages, SQUID experiment). The qubit is dispersively coupled to a 3D post-cavity that is displaced to $\alpha = 1.2$. Data (black) is normalized to range between 0 and $|\langle\alpha|1\rangle|^2$, and with dotted orange lines corresponding to exact multiples of 1.6 MHz. **b)**, Histograms from a reflection readout homodyne measurement, with and without a calibrated π pulse on a qubit (LINC experiment). The x and y axes in such a plot will generally be in voltage units from an analog to digital converter, with exact values depending on the amplification chain. The separation of the 'blobs' for the $|g\rangle$ and $|e\rangle$ qubit states allow high-fidelity single-shot readout.

$\pi/2$ pulses in this protocol must be 'unselective' to the oscillator's state, which can only be achieved with a short, high amplitude pulse whose bandwidth is significantly larger than $n\chi$.

While the above pulse sequences enable a variety of measurements of the oscillator in the Fock basis, they do not directly provide a knob for Unitary control. This can also be achieved in the number splitting regime using selective pulses, through a protocol called Selective Number Arbitrary Phase (SNAP) [117]. This requires the utilization of an additional degree of freedom that we have ignored until now, that of the direction of the qubit Rabi drive in the XY plane. Specifically, if the qubit drive is modified such that $|\epsilon_x|^2 + |\epsilon_y|^2$ is kept constant but $\varphi_d = \tan^{-1}(\epsilon_y/\epsilon_x)$ is changed, the pulse always enacts a π flip on the qubit, but follows a different longitudinal circle given by φ_d. Now, if this phase of this drive is varied between the selective pulses on each Fock state $|n\rangle$, then at the end of the pulse, the qubit in all cases reaches the same state $|e\rangle$, but the information of which trajectory it took to get there imprints a number-dependent phase φ_n on each $|n\rangle$. A final un-selective π pulse then resets the qubit state to $|g\rangle$, applying the overall Unitary transformation $\sum_n f_n|n\rangle \rightarrow \sum_n f_n e^{i\varphi_n}|n\rangle$ on the oscillator state. As we discussed previously, this number-dependent phase control, along with the number-changing displacement drive, is sufficient to enact arbitrary Unitaries on the oscillator [80]. In fact, it turns out this 'Displacement+SNAP' gate set is by no means the only way to achieve universal control in the dispersively coupled oscillator-qubit system. Most modern implementations simply feed the time-dependent complex controls $\epsilon_a(t)$ and $\epsilon_q(t)$ to a numerical optimizer, which calculates the resulting final state $|\psi_f\rangle$ from a set of initial states $|\psi_i\rangle$ under their driven Schrödinger evolution, and then tunes each drive to maximize the fidelity $|\langle\psi_f|U_{\text{targ.}}|\psi_i\rangle|^2$ for a target Unitary $U_{targ.}$. For a broader introduction to parity measurements, SNAP, and such gradient-ascent 'Optimal Control Theory' (OCT) techniques, one can refer to P. Reinhold's thesis [126].

Overall, the strong dispersive shift available in circuit-QED provides a natural number-selectivity, which in turn enables number-selective gates and measurements. However, being restricted to this form of control comes with a number of caveats. First, remember that

the number-phase conjugate bases are only one way of representing the oscillator Hilbert space – many natural continuous variable states and controls are only well-represented in the phase-space bases or through ladder operators. Implementing such continuous-variable control or error correction with number-selective gates is impractical, since one would have to apply selective pulses on an arbitrarily high number of Fock states to fully capture the intended evolution. The second caveat is much more dire – an always-on strong dispersive interaction is nefarious for propagating errors from the ancilla to the oscillator. Even when lying idle, any natural thermal fluctuations in the strongly coupled ancilla can cause oscillator dephasing at rate $\kappa_\varphi \approx \bar{n}_{q,th}$, which is derived from the previously discussed measurement-induced dephasing with $\chi \gg \kappa$. This dephasing immediately spoils the natural error hierarchy present in the oscillator mode, making phase-sensitive encodings like the GKP or erasure-dominated strategies like the dual-rail encoding fail. In addition to incoherent idle effects, the coupled ancilla also adds parasitic coherent nonlinearity like oscillator Kerr, a non-negligible oscillator-readout dispersive shift, and a nonlinearity in the dispersive shift itself, making most protocols fail at high photon numbers. The propagated error is even worse when the ancilla qubit is actively entangled with the oscillator state, for example during a SNAP or OCT pulse. Here, any part of the joint state with the ancilla in $|e\rangle$ can decay or dephase at an unknown time during the pulse, drastically changing the final Unitary enacted on the oscillator. Such an ancilla error can irreversibly spoil any intended computational trajectory for the oscillator state, or directly cause logical errors in an error-corrected encoding. Finally, while the above dispersive control and measurement techniques work well for a single oscillator, they are difficult to scale to a multi-oscillator system. One possible trick is to have both oscillators simultaneously participate in the same ancilla, with their joint dispersive shifts providing means for entanglement [127]. However, this quickly becomes intractable for high photon numbers and leads to correlated errors, necessitating a better, more scalable strategy for coupling

oscillators.

Numerous attempts have been made to alleviate these issues in through novel hardware and control techniques. Central to most of these strategies is the addition of one or more RF drives that modify the effect of the dispersive nonlinearity, allowing more continuous variable control, more tolerance to ancilla error-propagation, and better ways to control and couple oscillators. Let us study a few dominant examples. Consider a resonant drive on the oscillator that displaces its phase space for the entire duration of the control sequence. The dispersive shift can than simply be treated in the displaced frame:

$$\begin{aligned} a &\rightarrow a + \alpha \\ H_{\text{disp}}/\hbar &\rightarrow \left(\omega_q^0 + \chi_{qA}\ (a^\dagger + \alpha)(a + \alpha)\right)\ \sigma_z \\ &\approx H_{\text{disp}}/\hbar + \chi_{qA}|\alpha|^2\sigma_z + \alpha\chi_{qA}\ (a + a^\dagger)\sigma_z. \end{aligned} \tag{2.26}$$

This displacement transformation gives rise to two additional terms. The first of these, $\chi_{qA}|\alpha|^2\sigma_z$, is known as the A.C. Stark shift, and leads to an effective re-definition of the qubit frequency $\omega_q^0 \rightarrow \omega_q^0 + \chi_{qA}|\alpha|^2$ when driven. This shift will generally arise in any driven qubit-like system, even in more complicated control techniques introduced later in this thesis, and it will be important to track and correct it. The second term, $\alpha\chi_{qA}\ (a + a^\dagger)\sigma_z$, is much more interesting – it is a displacement of the oscillator conditioned on the state of the qubit (or a 'cross-resonance' gate between two qubits [128, 129]). This control is intrinsically non-number-selective and amenable to CV control, and has immediate use in measuring the qubit state through a homodyne detection of the conditional displacement of a readout mode [130, 131]. Importantly, this also serves as a non-Gaussian resource that is sufficient for universal oscillator control, so the residual dispersive shift in H_{disp} is no longer necessary. This means the static χ_{qA} can be strongly suppressed, but the conditional displacement can still be made fast through a large displacement strength

$|\alpha|$ [132]. In particular, this strategy is highly amenable for the GKP encoding, where these conditional displacements directly map translation errors to the ancilla, and the low static-coupling prevents idle dephasing or Kerr errors that the encoding is sensitive to. Importantly however, this does not solve the issue of error propagation *during* the ancilla-entangling control, nor does it allow one to regain number-selective control when desired.

A more generalizable set of strategies is to instead apply *off-resonant* drives that modify the nonlinearity of the ancilla at various orders. We will understand them here through some historic examples of control based on drives on the transmon, and then treat these strategies more generally in the next chapter. First, we must break away from the simplification of treating the transmon ancilla as a two-level system – let us instead consider the full driven Hamiltonian to fourth order with no rotating wave approximations. We assume we know the participation of both the oscillator mode and an external EM drive in the superconducting phase across the junction – as a reminder, these can be directly calculated from an HFSS simulation. The driven Hamiltonian, for a drive $\epsilon_d(t) = |\epsilon_d(t)|e^{-i\omega_d t}$, is given by:

$$H_d(\epsilon_d)/\hbar = \omega_A\ a^\dagger a + \omega_q\ q^\dagger q + \frac{K}{12}\left(q + \beta_{qA}\ a + \xi_d(t) + h.c.\right)^4 \tag{2.27}$$

where $\xi_d(t) = \frac{\beta_{qd}}{\theta_{\text{zpf}}}\ \epsilon_d(t)$ is the off-resonant displacement of the transmon at the drive frequency, and the factor of 12 is a result of combinatoric coefficients. Now, by simply choosing the appropriate drive frequency for $\epsilon_d(t)$, one can directly create arbitrary Gaussian controls on the transmon-oscillator system! As an example, consider a drive at half the *detuning* of the transmon and the oscillator, ie with $\omega_d \approx (\omega_a - \omega_q)/2$ – we will show that it creates a driven exchange (beamsplitting) interaction. In the undriven Hamiltonian, the only terms in the fourth-order nonlinearity that are energy conserving and contribute significantly to the system's dynamics are the transmon and oscillator Kerrs and dispersive

shift. However, the drive creates two additional energy conserving terms:

$$\begin{aligned} H_d(\epsilon_d)/\hbar - H_d(0)/\hbar =& K|\xi_d|^2(q^\dagger q + \beta_{qA}^2\ a^\dagger a) + \\ & K\beta_{qA}\ \left(\xi_d^2(t)\ a^\dagger q + h.c.\right). \end{aligned} \tag{2.28}$$

The first represents the previous discussed A.C. Stark shift, acting now on both the oscillator and transmon and shifting the frequency of any driven processes involving the two modes. The second is exactly the beamsplitting Hamiltonian from Eq. 2.4c, with $g_{bs}(t) = K\beta_{qA}\ \xi_d^2(t)$. More complicated nonlinear elements and drives can even provide non-Gaussian interactions, potentially eliminating the need for transmon-based control [133, 134].

In effect, the above driven treatment is an example of utilizing the strong light-matter interactions and nonlinearities in circuit-QED to produce new effective interactions in a driven frame. The possible interactions that can be created in such a manner can be enumerated through Feynman-like diagrams [135] of purely photonic 'four-wave mixing' interactions (a scalar ϕ^4 theory). These interactions are generally dependent on the drive like in the beamsplitter above, and therefore can be controlled with extremely high precision through room-temperature classical control of the microwave drives. We will see that these drives can often be viewed as modulating a parameter in the Hamiltonian, which leads to the resulting driven interactions being called 'parametric processes'. Specifically, the parametric beamsplitter interaction above enables any manipulations and measurements of the transmon to be directly mapped onto the oscillator. In fact, a purely resonant exchange interaction actually allows universal control of the oscillator, through a technique frequently called 'sideband control' [136]. It also allows interesting modifications to the dispersive interactions between the oscillator and the transmon. For example, driving an off-resonant beamsplitter to the transmon $e-f$ transition allows one to change the dispersive shift of the oscillator with respect to the $g-e$ and $e-f$ levels [137]. Clever manipulations

can then either null out the dispersive shift entirely, or create precise symmetries in the interaction such that utilizing the first three levels of the transmon effectively prevents the propagation of transmon errors to the oscillator [138, 139]. Finally, if instead of a single oscillator, the transmon was participating in two separate oscillators (a and b), it could also directly turn on a beamsplitting interaction between these two oscillators, allowing a scalable multi-mode oscillator control [140].

These parametric interactions, and the symmetries and multi-mode controls that they allow, are enormously useful. Indeed, unlocking the full range of parametric controls that are tractable in circuit-QED would provide direct control knobs on all Gaussian and dispersive interactions, enabling a powerful mix of both continuous-variable and number-selective controls. However, the transmon as a driven quantum 'mixer' suffers multiple non-idealities. The first one is its 'always on' anharmonicity, which inevitably leads to the oscillator mode inheriting some Kerr, and to driven stark shifts. More importantly, the same anharmonicity also leads to cavity-state dependent (dispersive) shifts on the resonance conditions for these parametric processes, which then become impractical to turn on in an unselective manner. The second non-ideality is much more complicated. The transmon's nonlinearity stems from the cosine potential of the Josephson junction, which formally allows mixing processes at all orders. This means for every desired process one may want to turn on, there exist uncountable parasitic processes that may be turned on instead [141]. The frequency selectivity of specific parasitic processes having specific resonance conditions is somewhat helpful, but due to frequency shifts and the high density of driven processes, one inevitably collides with a parasitic process at high enough drive strength. Additionally, the environment provides a continuous density of states that can also participate in these driven processes which can be difficult to predict, providing another dimension to the ways a driven process may be spoiled. As an example, in previous experiments for oscillator-oscillator beamsplitting using a transmon, such effects have

limited the process fidelity to $\sim 98\%$, primarily due to undesired drive-induced excitation of the transmon [140].

The primary work in this dissertation is to understand in detail how these nonlinear open quantum systems behave in the presence of drives, and to suppress these parasitic drive-induced excitations by about three orders of magnitude. We will go through the general framework and results for optimizing such driven quantum processes in the next chapter, focusing on a frequency-tunable variant of the transmon, the SQUID. Using the general strategies we develop in that chapter, we will construct a novel Josephson-based circuit that has zero anharmonicity when idle and can turn on parametric processes cleanly when driven (Chapter 4). This will then unlock a new architecture for bosonic quantum experiments that allows direct drive-activated continuous-variable and number-selective controls, which we will describe in detail in Chapter 5.

Chapter 3

Optimizing parametric control

How does one understand the interactions that arise from a driven quantum nonlinearity? Indeed, the theory required to properly describe the 'parametric processes' teased at the end of the previous chapter is very general – and it relates to a complete description of a driven open quantum system. A full treatment of this problem would involve predicting a practically infinite number of driven (resonant) processes, understanding how the system's interaction with the environment changes in the presence of the drives, and then generalizing the theory to arbitrary time-dependent control. To make this problem more tractable, we will divide the problem up into several levels of complexity. Our primary goal will be to understand the workflow for designing driven nonlinear elements within the circuit-QED setting. To do this, we will assume we have a set of desired interactions we would like to turn on by utilizing appropriate RF drives. At the minimal level of complexity, we will want some analytical intuition for predicting how these processes may be turned on, how they would depend on the drive, and what leading order undesired effects we might encounter. We will then want to make each of these predictions more concrete. First, we will learn how to engineer such a nonlinear element and its drive delivery, and extract the relevant driven Hamiltonian and noise directly from a full electromagnetic simulation. Then, we will study a way to more accurately calculate the

strengths and resonance conditions for coherent driven processes through a numerically exact (Floquet) treatment of the extracted driven Hamiltonian. Finally, we will introduce the effects of noise through a (structured) environment that can also participate in these processes, which we will describe through the Floquet-Markov theory. We will tie all of this together with experimental results of a precisely engineered driven nonlinearity, the 'Differentially-Driven SQUID' (DDS), which performs several orders of magnitude better than the transmon described in the previous chapter.

3.1 Understanding parametric processes

Let us start by studying the nature of parametric processes perturbatively, to gain some design intuition for them. To make our study more concrete, we will specifically study and compare two nonlinear elements that have different 'parameters' that are driven, and therefore different driven behavior, but can both activate the same desired processes. The first element is the charge-driven transmon Hamiltonian, which we introduced in Eq. 1.6:

$$H_{\text{trans}} = 4E_C \left(\hat{n} - n_g(t)\right)^2 - E_J \cos\hat{\theta}. \tag{3.1}$$

Here, $n_g(t) = \sum \epsilon_j \sin(\omega_j t)$ is a time-dependent gate voltage applied across the transmon's capacitor, which couples to its charge degree of freedom and acts as the driven parameter. It is significantly easier to analyze this circuit by performing a displacement transformation such that the drive instead displaces the junction phase:

$$\begin{aligned} \tilde{H}_{\text{trans}} &\rightarrow 4E_C\hat{n}^2 + \frac{E_J}{2}\hat{\theta^2} - E_J \cos_{NL}\left(\hat{\theta} + \phi_c(t)\right) \\ \phi_c(t) &\approx 8E_C\, n_{\text{zpf}}\theta_{\text{zpf}} \sum \epsilon_j \left[\frac{1}{\omega_j - \omega_q} + \frac{1}{\omega_j + \omega_q}\right] \cos(\omega_j t), \end{aligned} \tag{3.2}$$

Here, $n_{\text{zpf}}, \theta_{\text{zpf}}$ are the fluctuations defined in Eq. 1.7, and the driven displacement of the transmon phase is given by $\phi_c(t) = \xi(t)\theta_{\text{zpf}}$ (in radians), in terms of the usual oscillator-like displacement $\xi(t)$. As before, note that the equation above for $\phi_c(t)$ is purely for formal completeness, and is not useful in an actual experiments since one generally cannot measure the actual AC voltage applied at the transmon pads – instead, we will describe how to calculate this driven phase directly through numerical simulations later in this chapter. Particularly important is the term $\cos_{NL}(x) = \cos(x) + x^2/2$, which will give rise to all the interesting nonlinear effects that come from the Josephson junction in the displaced frame (Fig. 3.1b).

In contrast, the alternative Hamiltonian we will study is a simple extension of the single-junction transmon circuit to two junctions, sharing a superconducting loop. This circuit, named the Superconducting Quantum Interference Device (SQUID), was invented long before circuit-QED, and has been used for a broad range of experiments ranging from tests of superconductivity to extremely sensitive magnetic sensors. In particular, this is because the superconducting loop containing the two junctions can be threaded by magnetic flux, providing an additional control knob for the system. In the most general case, these junctions (E_{J1}, E_{J2}) each have an independent superconducting phase drop across them, which have opposite contributions from an external (uniform) flux threading the loop (Φ_{ext}):

$$H_{\text{SQUID}} = 4E_C\hat{n}^2 - E_{J_1}\cos\left(\hat{\theta}_1 + \frac{\Phi_{\text{ext}}}{2\phi_0}\right) - E_{J_2}\cos\left(\hat{\theta}_2 - \frac{\Phi_{\text{ext}}}{2\phi_0}\right), \tag{3.3}$$

where $\phi_0 = \hbar/2e$ is the reduced magnetic flux quantum [142]. However, if the two outer junctions are perfectly symmetric ($E_{J1} = E_{J2} = E_J/2$), an interesting interference effect leads to them being described by a simple Hamiltonian that is close to the regular

transmon:

$$H_{\text{DDS}} = 4E_C\hat{n}^2 - E_J\cos(\phi_d)\cos\left(\hat{\theta}\right). \tag{3.4}$$

Here, $\hat{\theta}$ is the same variable as in the regular transmon, being conjugate to the charge on the capacitors ($\hat{n}$), and is given by the common superconducting phase across both junctions $(\hat{\theta}_1 + \hat{\theta}_2)/2$. The second variable, ϕ_d, corresponds to the orthogonal *differential* phase across the two junctions, ie $\hat{\phi}_d = (\hat{\theta}_1 - \hat{\theta}_2)/2$. Why is this degree of freedom treated as classical? It is essentially because this differential degree of freedom corresponds to an extremely high-frequency oscillator, whose frequency is determined by the total loop inductance and the loop's self-capacitance, with the latter being tiny. Any changes in the external flux, by comparison, 'drive' this degree of freedom effectively adiabatically, creating essentially negligible (virtual) excitation in the actual quantum mode and simply redefining its mean value. Intuitively, this is like moving the pivot of a pendulum by hand, but doing it so slowly that the pendulum never oscillates, and only shifts its mean position. More specifically for this circuit, any disagreement between the flux threaded through the loop and the differential phase across the junction faces a huge energy penalty because of the negligible loop inductance. This leads to the mean value of this differential phase following the external flux very closely, with any variations being corrected on timescales much faster than any dynamics of the common phase (at the detuning of the drive and the differential mode). This means that $\phi_d = \langle\hat{\phi}_d\rangle$ is a 'frozen' degree of freedom, and exactly equal to the the external flux up to a unit conversion, $\phi_d = \Phi_{\text{ext}}/2\phi_0$. This can then precisely function as our driven parameter, $\phi_d(t)$.

There are two important things to note in the above description of the parameter ϕ_d. The first is that it forms an independent degree of freedom to the common mode phase θ, and does not displace it [143, 144]. This means that, barring any higher order nonlinear processes, the SQUID entirely stays in its *undriven* ground state (instead of a displaced

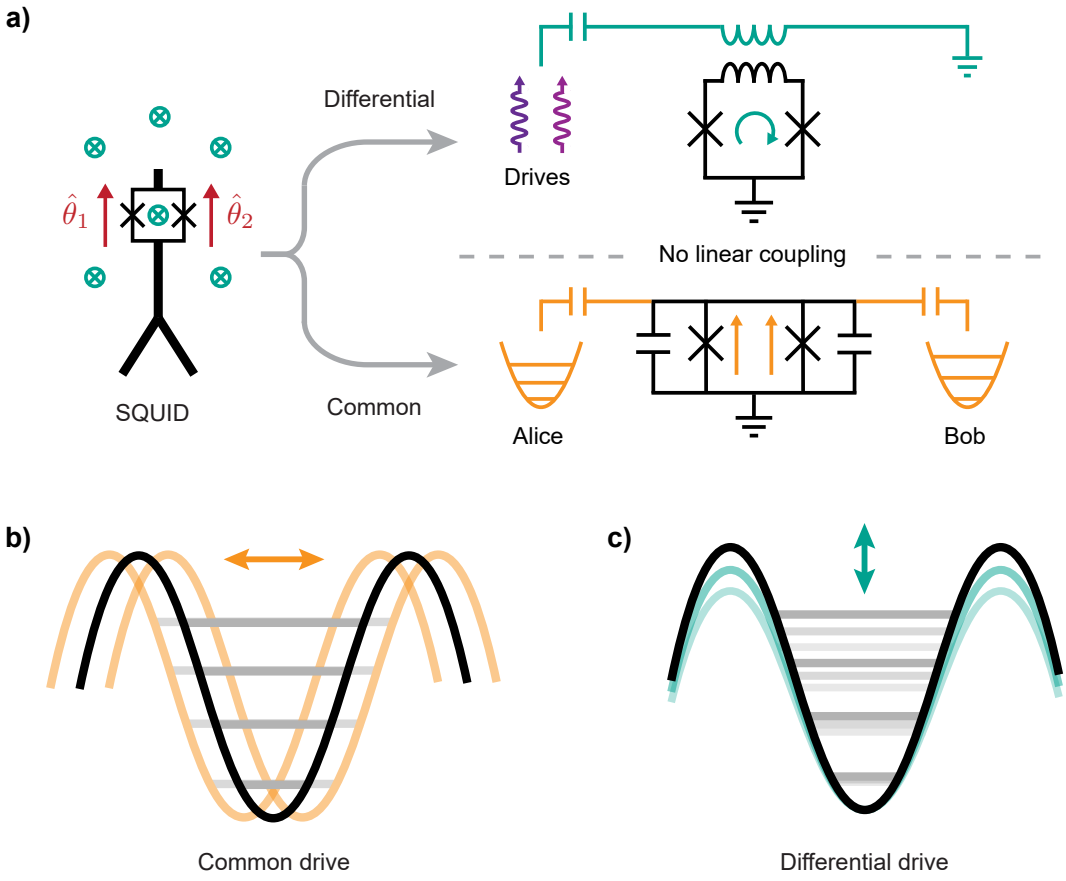

Figure 3.1: **Diifferent parametric drives on a SQUID a)**, A mixer with at least two degrees of freedom like the SQUID provides orthogonality between the driven parameter ($\phi_d = \langle\hat{\theta}_1 + \hat{\theta}_2\rangle/2$) and the dynamical quantum variable ($\hat{\theta} = (\hat{\theta}_1 + \hat{\theta}_2)/2$). Utilizing this orthogonality requires selectively driving ϕ_d (ie a differential drive), and coupling the information storing modes purely to $\hat{n},\ \hat{\theta}$ (a capacitive coupling). **b)**, If one were to drive the common-mode of the SQUID, it would displace θ and thus the potential energy, and behave exactly like a charge-driven transmon. **c)**, If one instead selectively drives the differential degree of freedom, it is equivalent to a pure modulation of $E_J(t)$, and hence the height of the potential. These two types of drive activate different parametric processes.

state) even when driven. In fact, one can understand this parametric drive even more simply – its effect is to simply modulate the junction energy, with $E_J(t) = E_J \cos(\phi_d(t))$. The effect of this modulation, at first order, is to change the height of the potential energy function and therefore its eigen-energies, but mostly not affect its actual wavefunction (Fig. 3.1c). We will see later, in a more complicated nonlinear element, that modulating this E_J does have some effect on the spread of each wave-packet (which is determined

by its impedance $\propto \sqrt{E_C E_J}$), but the effect of this modulation is largely invisible in the SQUID. The lowest order driven effect is then a frequency modulation of the SQUID, given by $\omega(t) = \omega_0 + \frac{1}{2}\frac{d^2\omega}{d\phi_d^2}\ \phi_d^2(t)$, assuming $\phi_d = 0$ when un-driven, with ω_0 being its idle frequency. We will see shortly that this latter term is effectively a 'four-wave mixing' process, and that most driven nonlinear effects can be derived by Taylor-expanding in a similar manner. The second important note is that because the parameter ϕ_d does not linearly couple to (is orthogonal to) the SQUID mode described by $\hat{n}, \hat{\theta}$, it also has no special constraints on the efficiency of drive delivery at different drive detunings from the SQUID frequency (Fig. 3.1a). This is in stark contrast to a charge-driven element like the transmon above, where there is only one degree of freedom available, and displacing it gets significantly more difficult at larger drive detunings (see Eq. 3.2). This allows an independent optimization of drive frequencies for an orthogonal parameter like in the SQUID, easing the constraints on designing high-fidelity parametric processes. Importantly, this also raises questions about whether the charge drive can even be considered a parameter if it is not orthogonal to the dynamical degree of freedom, but such a discussion is primarily about semantics.

Let us now provide analytic intuition for the processes that these parametric drives can turn on. Consider a general circuit-QED element where the sources of nonlinearity are Josephson-junctions (as opposed to phase-slip elements [145]). Also assume, for simplicity, that the nonlinear quantum 'mixer' has a single dynamical degree of freedom described by $(\hat{n}, \hat{\theta})$ – we can generalize this to an arbitrary number of modes later. Similarly, we will for now assume a single driven parameter $\phi(t) = \sum_j \phi_j \cos\big(\omega_{d_j} t + \varphi_j\big)$, with carrier frequency components $\{\omega_{d_j}\}$, phases φ_j, and amplitude envelopes ϕ_j. Note that these amplitudes ϕ_j will in general be time-dependent pulses, but for now we are simply interested in what happens when they are on at some constant value (i.e., steady-state behavior). The bare nonlinear circuit can then be simply represented by an effective parameter-dependent

nonlinear potential energy:

$$H_{\text{mixer}} = 4E_C\hat{n}^2 + U_J\left(\hat{\theta},\ \phi(t)\right) \tag{3.5}$$

The primary goal for such a parametric 'mixer' will be to turn on desired interactions, like beamsplitting or squeezing on external (bosonic) modes, without themselves participating in or contributing to the process. As such, we will assume that each of these modes of interest somehow participate in the junction, ie:

$$\hat{\theta} = \theta_{\text{zpf}}\left(\tilde{c} + \tilde{c}^\dagger\right) \approx \theta_{\text{zpf}}\left((c + c^\dagger) + \sum \beta_i(a_i + a_i^\dagger)\right), \tag{3.6}$$

where $\tilde{c}$ is the bare junction phase ladder operator, c is the corresponding dressed operator, and $\{a_i\}$ are the external dressed mode operators. Importantly, note that this decomposition already makes a couple of assumptions that are usually reasonable for transmon-like circuits. The first is that it is okay to represent a potentially compact (periodic) variable like $\hat{\theta}$ in terms of explicitly non-compact ladder operators from the phase-space of the external modes. This assumption is reasonable if the total phase fluctuations $\sum \beta_i\theta_{\text{zpf}}$ are significantly smaller than the periodicity of the potential. The second is that the participation of these external modes are simply represented by such a linear hybridization model. This assumes that all modes involved in the system are primarily harmonic, with their anharmonicity much smaller than their individual frequencies and frequency detunings. This is explicitly untrue for more complicated circuits like the fluxonium [32], whose lowest like states can be near degenerate ($\sim$zero frequency), but can still have a large gap (anharmonicity) to the higher excited states.

So what does one need to be able to turn on parametric processes? Every desired parametric process has an explicit energy gap that disallows it from always being on. This energy gap is extremely simple to calculate – one needs to simply enter the rotating

frame for each of the participating modes at their respective frequencies, and calculate the effective rotation of that process. As an example, each desired process has a form that can be explicitly written out in terms of its ladder operators:

$$\begin{aligned} H_{\text{contr.}}/\hbar &= g(\phi) \prod_j (a_j^\dagger)^{m_j} a_j^{n_j} \\ a_j e^{-i\omega_j t} &\Rightarrow H_{\text{contr.}} \times \exp\left[i \sum_j (m_j - n_j)\omega_j t \right]. \end{aligned} \tag{3.7}$$

This net rotation frequency of the process, $\sum(m_j - n_j)\omega_j$ sets the energy required to activate it, which must come from the drive appearing in the function $g(\phi(t))$. Specifically, the non-linear potential energy $U_J\left(\hat{\theta},\ \phi(t)\right)$ must allow a process that activates this control using k drive photons, such that:

$$\begin{aligned} U_J\left(\hat{\theta},\ \phi(t)\right) &\to g_{\text{eff}}\ \phi^k(t)\ \prod_j (a_j^\dagger)^{m_j} a_j^{n_j} \\ \sum k_i\ \omega_{d_i} &= \sum (m_j - n_j)\omega_j, \quad \text{s.t. } \sum k_i = k, \end{aligned} \tag{3.8}$$

where multiple drives at frequencies ω_{d_i} conspire to make the process resonant. This equation is much simpler to interpret when using a single drive frequency ω_d, whose resonance condition is given by $\omega_d = \omega_{\text{gap}}$, with $\omega_{\text{gap}} = \sum(m_j - n_j)\omega_j$. Finally, finishing this general treatment, we can derive how to get the above process from a general nonlinearity U_J. This is actually surprisingly simple – our potential energy is a function of two variables, $\hat{\theta}$ and ϕ, and we can simply perform a bi-variate Taylor expansion:

$$U_J\left(\hat{\theta},\ \phi(t)\right) = \sum_{l,k} \frac{\partial^k}{\partial \phi^k} \left(\frac{\partial^l\ U_J}{\partial \hat{\theta}^l}\bigg|_{\hat{\theta}_{\text{min}}} \frac{\left(\hat{\theta} - \hat{\theta}_{\text{min}}\right)^l}{l!} \right) \bigg|_{\phi_{\text{DC}}} \frac{\phi_{\text{AC}}^k}{k!}, \tag{3.9}$$

where $\phi(t) = \phi_{\text{DC}} + \phi_{\text{AC}}(t)$ has been split into its static and driven components.

Let us pause and analyze this Taylor expansion in a little more detail. First, notice that we must the evaluate these Taylor expansions at the bi-variate minima for the potential energy $\left(\hat{\theta}_{\text{min}},\ \phi_{\text{DC}}\right)$. Here, θ_{min} is defined as (classical) value of the $\hat{\theta}$,at which $\frac{\partial U_J}{\partial \hat{\theta}}\Big|_{\hat{\theta}_{\text{min}},\ \phi_{\text{DC}}} = 0$. For either a linear inductor, or any potential that is effectively a single junction, this minimum is simply at $\theta_{\text{min}} = 0$, but it may be nonzero for more complicated mixers. Second, the Taylor expansion above is still a step away from the desired process in Eq. 3.8. Specifically, we must expand $\hat{\theta}$ in terms of the ladder operators of various modes Eq. 3.6, which gives:

$$\begin{aligned} g_{\text{eff}} &= \frac{\partial^k}{\partial \phi^k}\left(\frac{\partial^l U_J}{\partial \hat{\theta}^l}\Big|_{\theta_{\text{min}}} \theta_{\text{zpf}}^l\right)\Big|_{\phi_{\text{DC}}} \times \prod_j \beta_j^{(m_j+n_j)} \\ &:= g_{l,k}[\phi_{\text{DC}}] \times \beta_{\text{eff}}^l. \end{aligned} \tag{3.10}$$

Here, $l = \sum_j (m_j + n_j)$ denotes the order of the parametric process, and $\beta_{\text{eff}} = \prod \beta_j^{\frac{m_j+n_j}{l}}$ is the geometric mean of the mode participations. In general, the hybridizations β_j are kept small (~ 0.1) in order to suppress effects like Purcell-induced decay and inherited Kerr in the bosonic modes, which means that multi-oscillator processes can quickly grow very weak. Additionally, notice that the above expansion means that we have included θ_{zpf} inside the derivative with respect to ϕ. This is a little subtle to understand – the Hamiltonian by definition must track the parameter $\phi(t)$ closely, ie it's potential energy takes on a new value every time ϕ changes, effectively instantaneously. This means the effective impedance, and spread of the wavefunction, can also change with the parameter instantaneously. Specifically,

$$\theta_{\text{zpf}} = \left(\frac{2E_C}{\partial^2 U_J/\partial \hat{\theta}^2\Big|_{\hat{\theta}_{\text{min}}}}\right)^{1/4}, \tag{3.11}$$

which can still be a function of $\phi(t)$! Previous literature has often intuitively taken this into account, by writing frequency-modulated processes as a Taylor expansion with derivatives of frequency [146, 147], for eg:

$$g_{4wm} = \frac{1}{2}\frac{d^2\omega}{d\phi^2}\phi_{\text{AC}}^2\, c^\dagger c = \frac{\partial^2}{\partial\phi^2}\left(\frac{\partial^2 U}{\partial\hat{\theta}^2}\theta_{\text{zpf}}^2\right)\frac{(c+c^\dagger)^2}{2}. \tag{3.12}$$

The correction due to this modulation might not be large in the transmon or the SQUID, but we will see that it will have very visible effects in protected mixers where one would naively have expected certain terms to be zero (see Zeeman shift in Chapter 4).

Now that we have a general form for perturbatively expressing parametric processes, let us specifically analyze some processes of interest activated by the driven transmon and SQUID. In particular, we can consider these mixers coupled to two bosonic modes, Alice (a) and Bob (b), on which we want to enact the Gaussian controls listed in Eq. 2.4. In any other charge-driven mixer, including the transmon, the description of parametric processes given in Eq. 3.9 can be simplified. The modulated parameter in the displaced frame appears as a shift in $\hat{\theta} \to \hat{\theta} + \phi_c$. This means that taking a derivative with respect to ϕ_c is equivalent to taking one with respect to $\hat{\theta}$! This translates to its driven processes simply being enumerated as:

$$U_{\text{charge}} = \sum_m \frac{\partial^m U_J}{\partial\hat{\theta}^m}\,\theta_{\text{zpf}}^m\,\frac{(\tilde{c}+\xi_c(t)+h.c.)^m}{m!}, \tag{3.13}$$

where $\xi_c(t) = \phi_c(t)/\theta_{\text{zpf}}$. For the transmon, this Taylor expansion is just that of the cosine function:

$$U_{\text{trans}} = \sum_{m\in\text{even}} E_J\left(\frac{2E_C}{E_J}\right)^{m/4}(-1)^m\frac{(\tilde{c}+\xi_c(t)+h.c.)^m}{m!}. \tag{3.14}$$

It is immediately clear that the transmon can only turn on processes that have a *total* number of excitations participating that are even. The lowest order of these is the fourth-

order nonlinearity, which allows four photons (or 'waves') to mix:

$$U^{(4)}_{\text{trans}} = -\frac{E_C}{12}\left(\tilde{c} + \xi_c(t) + h.c.\right)^4 . \tag{3.15}$$

As an example, for a single drive frequency ω_d, we can decompose this nonlinearity into each parametric process that it activates:

$$\begin{aligned}
U^{(4)}_{\text{trans}} = &- E_C|\xi(t)|^2\left(\beta_a^2 a^\dagger a + \beta_b^2 b^\dagger b\right) && \text{Stark shift} && (3.16)\\
&- E_C\beta_a\beta_b\left(\xi^2(t)a^\dagger b + h.c.\right) && \text{Beamsplitting} && (3.17)\\
&- E_C\xi^2(t)\left(\frac{\beta_a^2}{2}a^{\dagger 2} + \frac{\beta_b^2}{2}b^{\dagger 2} + \beta_a\beta_b a^\dagger b^\dagger\right) + h.c. && \text{Squeezing} && (3.18)\\
&- \frac{E_C}{3}\xi(t)\left(\beta_a^3 a^{\dagger 3} + 3\beta_a^2\beta_b a^{\dagger 2} b + \dots\right) + h.c. && \text{Extra} && (3.19)\\
&- \frac{E_C}{3}\xi^3(t)\left(\beta_a a^\dagger + \beta_b b^\dagger + \dots\right) + h.c. && \text{Extra, sub-harmonic.} && (3.20)
\end{aligned}$$

The stark shift occurs regardless of drive frequency, and can cause driven frequency shifts even when they are undesired. The beamsplitting, squeezing, and two-mode squeezing terms specifically occur at resonance conditions given by $\omega_d = (\omega_a - \omega_b)/2,\ \omega_a,\ \omega_b,\ (\omega_a + \omega_b)/2$ respectively, and allow complete Gaussian control. Finally, there exist some extra processes that are undesired in this scenario and could lead to process degradation.

The SQUID potential is much more directly expanded by Eq. 3.9, since $\phi = \phi_d(t)$ is a

true orthogonal parameter (Fig. 3.1a). Specifically, at $\phi_{\text{DC}} = 0$, this gives:

$$\begin{aligned} U_{\text{DDS}} &= -E_J \cos\phi_d \cos\hat{\theta} \\ &= -E_J \sum_{\substack{l\in\text{even} \\ k\in\text{even}}} \frac{\partial^k}{\partial\phi_d^k}\left(\frac{\partial^l U_{\text{DDS}}}{\partial\hat{\theta}^l}\frac{\theta_{\text{zpf}}^l}{l!\ k!}\right)\Bigg|_{\hat{\theta},\ \phi_{\text{DC}}=0} \phi_{\text{AC}}^k\ (\tilde{c}+\tilde{c}^\dagger)^l, \\ &\approx -E_J \sum_{k\in\text{even}} \frac{(-\phi_d)^k}{k!} \sum_{l\in\text{even}} \frac{(-\theta_{\text{zpf}})^l}{l!}\left(\tilde{c}+\tilde{c}^\dagger\right)^l, \end{aligned} \tag{3.21}$$

where the last line assumes that θ_{zpf} does not vary with the flux ϕ_d, which is at least true at first order corrections ($\partial\theta_{\text{zpf}}/\partial\phi_d|_{\phi_{\text{DC}}=0} = 0$). This expression is incredibly simple, and imposes an additional symmetry on top of the terms in the driven transmon: all non-zero terms must be even in *both* the number of drive excitations *and* the number of mode excitations (ladder operators). At lowest order, this again gives a four-wave mixing term, but one which has this symmetry rule explicitly imposed:

$$U_{\text{DDS}}^{(4)} \approx -\frac{\hbar\omega_q}{2}\phi_d^2\left(\tilde{c}+\tilde{c}^\dagger\right)^2, \tag{3.22}$$

where we have used $\omega_q \approx \sqrt{8E_JE_C}$ to absorb the factor of θ_{zpf}^2. It is then easy to derive that this potential gives the same Gaussian processes in Eq. 3.20, yet it automatically nulls out all of the extra processes. In fact, this protection against undesired terms actually occurs at all orders of non-linearity. This is significantly easier to observe if one re-writes the transmon potential using a simple trignometric expansion:

$$U_{\text{trans}} = -E_J \underbrace{\cos\phi_c \cos\hat{\theta}}_{\text{Even Parity}} - E_J \underbrace{\left(\sin\phi_c \sin\hat{\theta} - \phi_c\ \hat{\theta}\right)}_{\text{Odd Parity}}. \tag{3.23}$$

The first term, which is of even parity in both ϕ_d and $\hat{\theta}$ (drive and mode operators), is precisely U_{DDS}. The second, with odd-parity, represents all the extra undesired terms in the

single junction potential. Importantly, the SQUID circuit can also function exactly like the transmon Hamiltonian (with tunable $E_J(\phi_d)$), if one instead drives it through its charge degree of freedom (Fig. 3.1b). This would then activate a displacement of the *common* mode phase across the two SQUID junctions, making them behave equivalently to a single junction transmon – hence the subscript 'c' in ϕ_c. Enforcing the additional selection rule in the DDS requires three conditions overall; its junctions must be symmetric, it must be *differentially* driven (Fig. 3.1a,c), and it must sit at the flux sweet-spot of $\phi_{\text{DC}} = 0$. Importantly, a slight violation of any of these conditions only brings it closer to the regular transmon, and does not completely spoil its driven performance.

The suppression of approximately half the driven terms in the DDS compared to the transmon is of fundamental importance to designing clean parametric processes. One can in fact generalize this principle, by realizing that it arises from simply a quantum analog of mixer-balancing in classical electrical engineering. Since potential functions composed of the Josephson cosine nonlinearity are highly amenable to engineering parity-type selection rules, we call this broad strategy **parity protection**. To understand why this suppression is important, lets first put a parametric mixer in an actual experimental context. A typical operation one may need to perform, either between qubits or between bosonic modes, is a two-mode entangling gate. Performing this operation with a dedicated parametric mixer involves involves at least three-modes: a mixer that drives beamsplitting (which we will call a 'coupler'), and the information storing modes, Alice and Bob. Even with this minimal system, and just considering the fourth-order nonlinearity in Eq. 3.20, it is clear that there are a significant number of parasitic processes one could run into – remember, in this context, *every* process except the beamsplitter itself is undesirable. Even if one were to avoid these resonant processes at low drive powers through careful engineering of the Hamiltonian, the AC Stark shift causes all of these processes to shift significantly with drive power, which can lead to eventual collisions. Moreover, the cosine nonlinearity in

a typical Josephson junction provides nonlinear terms at all orders, and therefore a nearly infinite forest of driven processes that one might need to avoid, all of which can shift with drive power. Thankfully, higher order processes are generally exponentially suppressed in the drive amplitude and participations (since they go as $\sim \phi^k \beta_j^l$), but colliding with any process on resonance could completely spoil fidelity.

The above analysis is just for coherent processes within the necessary modes of the system. A general experimental device, even at the scale of an academic experiment, contains numerous distributed modes with varying coherence, all of which may participate in the mixer. In general, this forms a highly structured environment that is coupled to the driven nonlinear element, and the drives can turn on new decoherence mechanisms on both the mixer and the information storing modes. This is exactly the same argument as the linear response relation detailed in Chapter 1 – any degree of freedom (parametrically driven process) of the mixer that the desired modes or the user has access to, the environment also has access to. Together, these parasitic resonances within the Hamiltonian of the desired modes, and any interactions with the general environment created by the 'extra' modes, are practically infeasible to predict or intuit analytically. Instead, the intuitions above only provide guiding principles for designing new driven nonlinear Hamiltonians from a quantum-optics or control perspective. When converting such an idea into experiment, further analysis is required that can numerically predict the additional modes and environmental structure in the device, and accurately calculate all driven coherent and incoherent processes that can arise from the known system and environment. Even with a parity-protected Hamiltonian like the DDS that contains a smaller fraction of resonant terms, optimal performance still requires such a careful numerical analysis. Our primary tool to perform this analysis and optimization will be **Floquet engineering** – which is composed of coherent Floquet and incoherent Floquet-Markov simulations, aided by information from numerical EM simulations of the package. We will outline this broad

strategy below, but specific details for the Floquet analysis can be found in [148, 149, 141], and methods for extracting the driven Hamiltonian and environment in [150].

3.2 Floquet engineering

Let us now study methods to fully numerically analyze a driven open quantum system in circuit-QED. Overall, the workflow for analyzing such a system is as follows. First, we create a linear 3D EM model for the full experiment that matches the engineered device to the best of our knowledge. In general, this model is constrained to a (usually superconducting) metal package that acts as a boundary condition for the simulation, and contains one or more substrates with quasi-2D 'on-chip' circuits that contain Josephson junctions approximated as linear inductances. The simulation will also contain (usually coaxial transmission-line) ports with well defined characteristic impedances. We will choose a tractable number of modes within this device that we will treat as a full quantum-mechanical Hilbert space (e.g., a coupler and two cavity modes). The rest of the modes in the system will be treated classically, but their effects will be captured through the structure of the environmental density of states, and the spectral filtering of the drives. Specifically, we will be interested in the modes, drives and environment from the perspectives of the junctions involved in our simulation, which will then provide all the nonlinear driven effects. Finding the drives and mode participations in each junction will provide the full driven Hamiltonian, from which we can calculate the Floquet modes of the system, which form an accurate basis to describe our all coherent dynamics of the driven time-periodic Hamiltonian. Studying the effective eigen-energies and overlaps of various Floquet modes will directly provide important driven effects, like drive-induced anharmonicity and frequency shifts, and will specify the drive conditions for every nonlinear resonance within the system. Finally, we will include the effect of the other modes and the ports through

the previously calculated mode and noise spectrum, and perform Floquet-Markov simulations to calculate any drive-induced decoherence on the modes of interest. This will also provide the steady-state of the system in the presence of noise, which we will see provides an important proxy for the 'clean drive-space' available to any parametric mixer. Once we can accurately analyze the full system, we can optimize the geometry, layout, and frequency stack of our circuits to optimize their performance. The rest of this section will describe this workflow in more detail, including providing the necessary ingredients for understanding Floquet theory.

3.2.1 Extracting Hamiltonians, drives and noise

We start with extracting the driven Hamiltonian and noise from an EM model of our device, using a high-frequency simulation software (HFSS) [151]. The full-driven treatment of a superconducting circuit in a manner that is compatible with quantization is non-trivial. A general 3D device does not usually allow a simple lumped-element description of confined modes with well-defined couplings. Additionally, the drive can take a geometrically distributed form of time-varying electric and magnetic fields that do not interact with the driven circuit in simple terms, like a "voltage at a capacitive pad" or a "flux in a loop". As an example, even if one explicitly calculates the total time-dependent flux threading a SQUID loop, it does not fully capture the effect of the drives on the SQUID's individual junctions. In fact, there exists a unique way to allocate the total phase drop due to an AC flux to each branch in the loop, and it depends on not just the loop itself, but also the distributed capacitance matrix that shunts the loop [152, 153]! One way to understand this is to consider each junction of the SQUID to be shunted by a different (asymmetric) capacitance. These capacitors form additional loops around the junctions that are not galvanically connected, but can still be driven by AC flux and contribute a different phase drop to each junction depending on their capacitances. The total phase drop across each

junction is then determined by all such loops that the junction participates in, which makes calculations significantly complicated for distributed capacitances. Additionally, any realistic drive-line that delivers a flux drive will not purely drive geometrically uniform flux – the drive will inevitably create a flux gradient, and will also likely have some stray capacitive coupling to the SQUID's common mode. While the combined asymmetry due to these effects may be complicated, each effect is completely described by classical linear EM theory. This, instead of calculating their effects individually, we can simply calculate their total driven effect on each SQUID junction in a classical EM simulator.

This is the primary strategy employed to calculate the driven Hamiltonian of an arbitrary circuit – we utilize the fact that each junction is a lumped element inductor, and any nonlinear effects of the fields in the circuit must result from some driven current through the junction, whose resulting phase drop can be directly calculated. Specifically, one can define a line across the linear equivalent lumped inductor for each junction, and then integrate the electric field (given by HFSS) along that line to get the voltage drop V_0 across the junction. Since any modulation that we will consider will be sinusoidal, the phase drop across the junction can be directly calculated using:

$$\phi_j(t) = \frac{1}{\phi_0}\int dt\ V_j e^{i\omega_j t} = \frac{V_j}{\omega\phi_0} e^{i\omega_j t}. \tag{3.24}$$

We can utilize this phase drop to get both the idle participation of the other modes in the system, and the effect of the drives on each junction. For the former, we perform an 'eigenmode' simulation, which calculates the normal modes of the system and their electric and magnetic field distributions. Importantly, the same simulation also takes into account the effects of any impedances at ports, which appear as imaginary parts of the mode frequency, ie $\omega + i\kappa/2$. This provides both the coupling of each mode to the ports, and the idle lifetime (or Q) of each mode if they were only limited by these ports. Now,

to calculate the participation of each normal mode in a junction, we need to simply excite that mode with a known amount of energy (ideally $|\alpha = 1\rangle$), and measure the phase drop across the junction at that mode's frequency. This is equivalent to the well-established method of computing the energy participation ratio [76], and directly provides the junction phases fluctuation amplitude due to the presence of each mode, $\beta_j \theta_{\text{zpf}}$. The general idea of integrating field participation can also be extended within the same simulation to calculate other idle effects, like the losses inherited from a seam in the package or a substrate with well-defined $\tan\delta$, and the interested reader should refer to S. Ganjam's thesis [154] for details of this analysis.

What remains is to calculate the strength of the drives and noise on these junctions. To do this, one essentially uses the linear response theory detailed in Chapter 1 explicitly. We simply find the conversion factor from a voltage drop at each port to the voltage drop across each junction, through a 'driven modal' or 'driven terminal' simulation. Specifically, this factor is a function of frequency, given by:

$$A_{k,s}(t) = \chi_{k,s}[\omega_d] \, \frac{V_s}{\omega_d \phi_0} \, e^{i(\omega_d t + \phi)}, \tag{3.25}$$

where $A_{k,s}(t)$ is the phase drop at the kth junction due to a drive at the sth port, and $\chi_{k,s}[\omega_d]$ is a complex-valued function of frequency that captures the effect of any drive delivery filters and all the modes in the system. Note that $\chi[\omega]$ provides a simplistic view of the effect of the other normal modes – which is that they simply act as a distributed linear filter network and have no dynamics of their own. Importantly, if there are multiple modes participating in the junction, then the drive can displace each of these individual modes, which can in turn enact a phase drop across the junction through their participation in it. Finding a direct voltage conversion ratio ignores the decomposition of the drive into these individual components, and does not contain any information about which

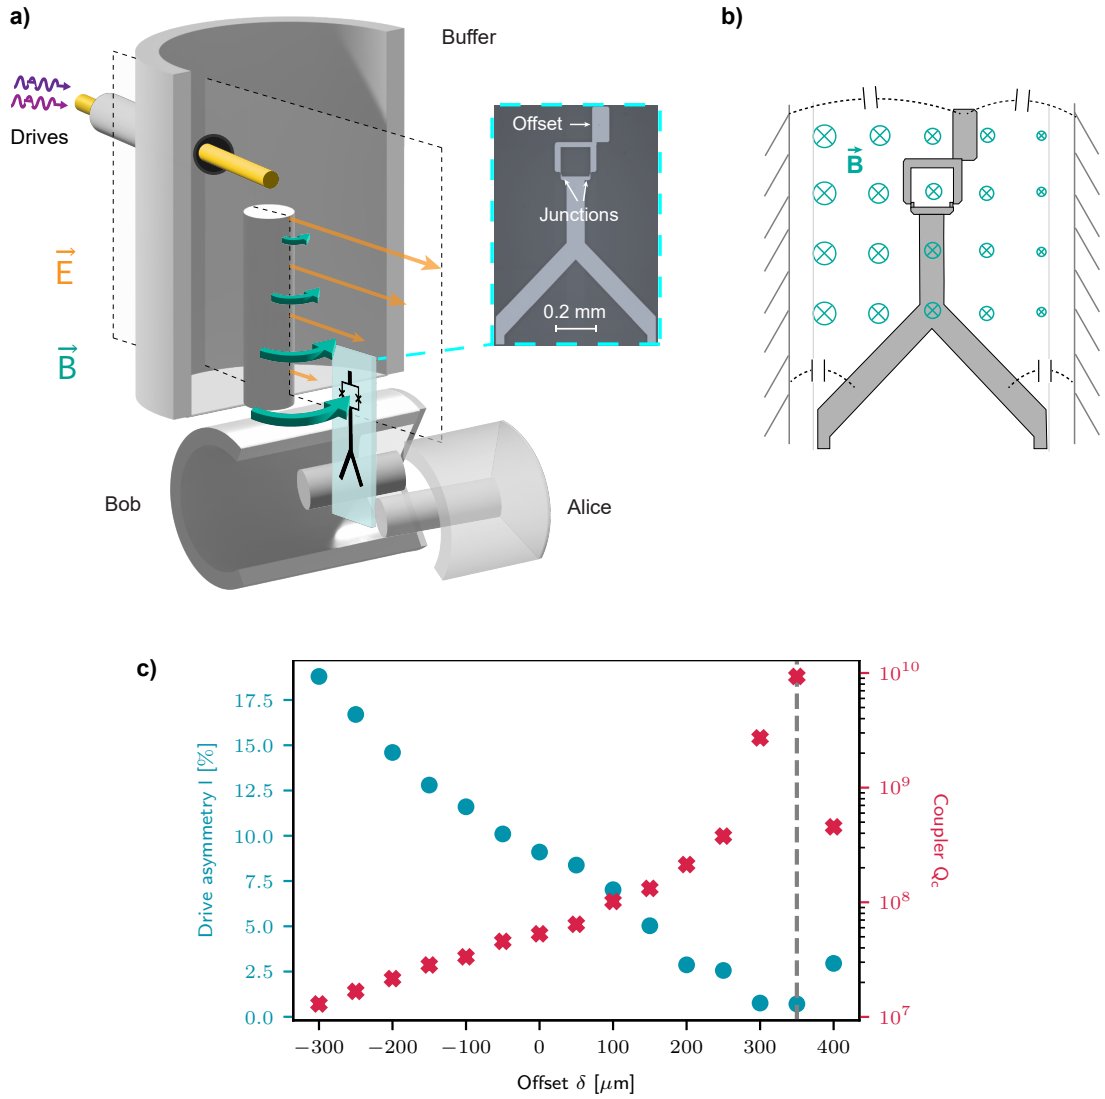

Figure 3.2: **Example of drive optimization for the DDS a**, The DDS in a 3D cavity package, acting as a high-fidelity coupler between bosonic modes Alice and Bob. The drive is delivered through a $\lambda/4$ buffer filter cavity, with the SQUID at the simultaneous B-field anti-node and E-field node. Inset is an optical micrograph of actual device. **b**, The package's cylindrical geometry means that the B-field has a non-uniform distribution along its radial direction (from left to right in figure). Additionally, capacitance between the antenna pads of the SQUID and the wall of the package (dashed lines) may also allow some parasitic charge drives. **c**, By changing the offset δ of one of the SQUID's capacitive pads, we can fine tune the ratio of common mode to differential mode coupling (blue circles), extracted directly from integrated phase drops across the two junctions. At $\delta \approx 350\mu$m, this common mode coupling is minimized and coupler Q (red crosses) is maximized.

modes are displaced. Fully taking these displacements into account requires computing the overlap between the driven field and each mode, integrated over the full 3D space, and separately calculating any residual phase drop from the drive that is not associated with these modes [150]. Additionally, both these methods still assume that linear hybridization is a good model for the coupled Hamiltonian of the system, which relies on any effects due to the anharmonicity being perturbative, but this is okay for transmon-like circuits.

Since circuit-QED allows a high degree of engineerability through such microwave analysis, how specifically should one design the coupling to the drive port? The primary requirements on the drive coupling are that it:

1. Delivers sufficient phase drop across the circuit's junctions at at all required frequencies (over a wide enough bandwidth) with reasonable drive powers,

2. Can be treated as effectively classical, ie has low direct coupling and large mean field value when driven such that any effects due to quantum fluctuations, like shot noise dephasing, are negligible,

3. Does not Purcell-limit or introduce lossy elements into any important modes in the system, and

4. Obeys any special symmetry requirements for correlated phase drops across different junctions.

As an example, in the drive-delivery for the DDS experiment, we require a two-tone beamsplitting drive around ~ 3 GHz that specifically drives the SQUID junctions differentially. To achieve this, we utilize a 3D $\lambda/4$ stub cavity filter at 3 GHz, dubbed the 'buffer mode', that is over-coupled to a capacitive drive pin to provide a wide enough bandwidth for both tones, and provides a geometric separation between driven electric and magnetic fields (Fig. 3.2). Placing the SQUID at the simultaneous B-field antinode and E-field node

mostly takes care of the differential driving, but flux gradients and parasitic capacitances prevent this condition from being met perfectly. We then utilize the above methods to directly calculate and optimize the 'differential-ness' of the driven phase drops across the two junctions. Specifically, we tune the drive asymmetry by changing the SQUID's capacitive matrix through an offset in one of its capacitive pads. When a fully differential coupling to the drive is achieved, this also makes the Purcell limit on the coupler due to the drive port negligible, since the common mode cannot decay through this port. Finally, by tuning the loop size and SQUID position, we can suppress the flux-coupling to the buffer mode and ensure that it can be treated as classical (stiff). This is possible because in the limit of small dispersive shift to a filter ($\chi_f \ll \kappa_f$), any shot noise dephasing due to the filter mode scales as χ_f^2, but the four-wave mixing process strength only scales as $\sim \chi_f$, so the former can be suppressed without significantly reducing the latter.

The general treatment described above provides a full driven Hamiltonian in the displaced frame of each junction, for a M-junction circuit with N modes of interest:

$$H_{\text{disp}} = \sum_{i=1}^{N} \hbar\omega_i a_i^\dagger a_i - \sum_{k=1}^{M} E_{J_k} \cos_{NL}\left(\sum_{i=1}^{N} \beta_{ki}(a_i + a_i^\dagger) + \sum_s A_{k,s}(t)\right), \tag{3.26}$$

where each port s must also be associated with some noise spectrum that is added to the parametric processes. This noise can often be simply approximated as the Johnson noise on the voltage $V_s(t)$ through the characteristic impedance of the port, at the base temperature of the fridge. However, a general control line will contain a filter and attenuation stack that will determine the non-zero contribution of noise propagated down from room temperature controls and higher stages. The total noise spectrum can then be explicitly calculated, including any self-heating of the cryogenic attenuators that dissipate excess drive power [155, 156], and added to the Johnson noise at the port. In general, this noise will enter into the junction's nonlinearity and can cause highly non-trivial effects in combi-

nation with the drive, especially since Johnson noise has spectral content at all frequencies and can therefore contribute 'noise photons' to all processes. This is in addition to the idle decay of each mode through the ports, which can be added into the simulation by hand, along with any estimates for sources of noise that do not come from the ports, like on-chip flux noise [157].

3.2.2 Floquet-analysis of driven coherent and incoherent processes

What does one do with the time-dependent Hamiltonian and noise derived in the previous section? Let us first consider just the driven Hamiltonian, and specifically in the presence of a single drive tone. This Hamiltonian is then explicitly periodic at the periodicity of the drive, ie $H_d(t+T) = H_d(t)$, where $T = 2\pi/\omega_d$. It turns out that such a periodic Hamiltonian has an exact eigenbasis that is also similarly periodic, in the time analogue of Bloch's theorem for periodic lattices. Specifically, the Schrödinger equation for $H_d(t)$ has solutions called 'Floquet states', given by:

$$\begin{aligned} \left(H_d(t) - i\frac{\partial}{\partial t}\right)|\psi_{\alpha,k}\rangle &= \hbar(\epsilon_\alpha + k\ \omega_d)|\psi_{\alpha,k}\rangle \\ |\psi_{\alpha,k}\rangle &= \exp\left[i(\epsilon_\alpha + k\omega_d)t\right]\ |\varphi_\alpha\rangle \\ |\varphi_\alpha(t)\rangle &= |\varphi_\alpha(t+T)\rangle, \end{aligned} \tag{3.27}$$

for eigenstate indices α, k and eigen-frequencies ϵ_α. How do we interpret these solutions? Consider that if the Hamiltonian is periodic, and if the system is in an eigenstate of the evolution, we expect probabilities to repeat with at least the same period. As such, we could naively hypothesize that $|\langle\psi_{\alpha,k}(t)|\psi_{\alpha,k}(t+T)\rangle|^2 = 1$, ie that the wave-function perfectly overlaps with itself one time period away, up to an overall phase. This automatically means that $|\psi_{\alpha,k}(t+T)\rangle = e^{i\omega_\alpha T}|\psi_{\alpha,k}(t)\rangle$, for some constant value ω_α. However, $\omega_\alpha T$ is just a phase and so it is only defined to up to an addition of $2\pi\ k$, where $k \in \mathcal{Z}$ is any

integer. This implies that the eigen-energy, ω_α, is perfectly valid even when multiplies of $2\pi/T = \omega_d$ are added, ie:

$$\begin{aligned} \omega_\alpha^{(k)} &= \epsilon_\alpha + k\omega_d \\ \Rightarrow \epsilon_\alpha &= \omega_\alpha^{(k)} \mod \omega_d. \end{aligned} \tag{3.28}$$

Here ϵ_α is only defined up to multiples of ω_d, and hence is called a *quasi-energy*. Finding the true energy of the system is then equivalent to assigning a 'Brillouin zone' [158] to the quasi-energy, and we will learn how to do this later in this section. The part of the wavefunction that is not this phase $e^{i\omega_\alpha^{(k)}t}$, is then explicitly periodic in T, and is defined as the the *Floquet mode* $|\varphi_\alpha(t)\rangle$. Importantly, this Floquet mode is also a 'quasi-state' – in that it is not the true eigenstate of the system, but is used to generate the true eigenstates $\{|\psi_{\alpha,k}\rangle\}$, and it describes the system's dynamics *modulo* ω_d. Solving the Schrödinger's equation for this system then simply requires finding $|\varphi_\alpha(t)\rangle$ and ϵ_α given $H_d(t)$ and ω_d. This is relatively easy, since the periodicity of $|\varphi_\alpha(t)\rangle$ means that it is an eigenstate of the Hamiltonian propagator for one time period, and one can just solve the eigensystem equation:

$$\begin{aligned} U_d(t, t+T)\ |\varphi_\alpha(t)\rangle &= \epsilon_\alpha\ |\varphi_\alpha(t)\rangle \\ U_d(t_1, t_2) &= \mathcal{T}\left[e^{-i\int_{t_1}^{t_2} H_d(t)\ dt}\right], \end{aligned} \tag{3.29}$$

where $\mathcal{T}$ is the usual time-ordering operator. Any standard quantum simulation package, like `qutip` [159] or `dynamiqs` [160], will have functions defined that solve this problem.

Now that we have a mathematical framework to describe the system, lets gain some physical intuition for it and put it to use. Our eigenstates of the system in the presence of the drive are given by $\{|\psi_{\alpha,k}\rangle\}$. However, the number of such eigenstates is of course

infinite due to the index k – but this can't possibly be true, since the size of our Hilbert space cannot be affected by the drive. Instead, the true eigenstates of the system are simply $|\psi_\alpha\rangle$ (for any k), each of which can be traced back to the original *undriven* eigenstates at $\omega_d = 0$ or zero drive amplitude. The presence of the index k instead denotes that each eigenstate has its wavefunction's behavior distributed over frequency space. Specifically, each eigen-energy of the system now has copies, or *sidebands*, spaced by ω_d, in a very similar manner to classical frequency modulation. This is simply another way to state that the transition frequency between any two states can be aided by the gain or loss of k drive photons in the presence of the drive! Any driven transitions between two states can then simply be analyzed by asking the questions:

1. Do the two states have any sidebands that cross?

2. Do the states have a non-zero matrix element in the interaction Hamiltonian?

If the answer to both these questions are yes, then the drive will activate a resonant transition between those states, aided by $k = k_1 - k_2$ drive photons for state sideband indices k_1, k_2 (and the process will scale at least as $|\phi_d|^k$ with the drive amplitude). Note that these sidebands are not simply mathematical tools for our description of the system, they are very physical and can be directly measured in an experiment, as visible in Fig. 3.3a. In the presence of a strong drive, any weak probe that interacts with a transition will also interact with it at frequencies spaced by ω_d, providing direct spectroscopic evidence for the existence and the relative amplitude of sidebands to the main transition 'peak'. Both the experimental data and the Floquet simulation also clearly show that certain sideband collisions will not show an avoided crossing when prevented by Hamiltonian selection rules, like the coupler $k = 0$ transition and Bob $k = 3$ sideband. To gain a full understanding of such Floquet physics, the reader is highly encouraged to read [148] (Chapters 2 and 9), or [149] (Appendix B) for a concise summary.

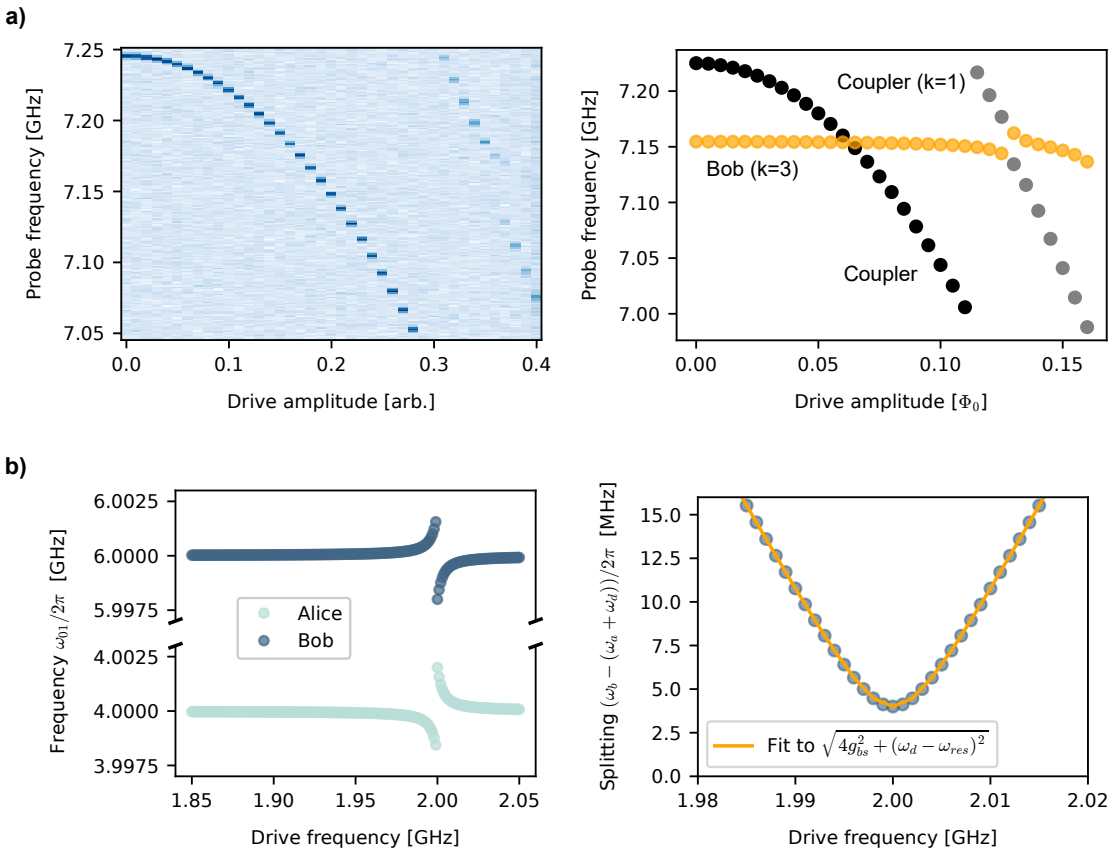

Figure 3.3: **Coherent Floquet analysis a**, Experimental data (left) and corresponding Floquet simulation (right) of coupler sidebands in the presence of a drive. The sideband is 'real', in the sense that it is visible in the spectroscopy, and it shows a resonant collision with a different sideband of a coupled bosonic mode (Bob). **b**, Floquet simulations of a three-wave mixing beamsplitter between two modes, Alice and Bob, with $g_{bs}/2\pi = 2$ MHz. An avoided crossing is visible (left) around the beamsplitter resonance condition ($\omega_d/2\pi = 2$ GHz), which can then be extracted and fitted (right) to confirm the beamsplitting strength and resonance frequencies.

While diagonalizing a static Hamiltonian automatically provides real energies and eigenstates ordered by these energies, diagonalizing the propagator in Eq. 3.29 only provides quasi-energies and corresponding Floquet modes in an arbitrary order. This results in two important but simple tasks that are required to extract the experimentally relevant physics of the system. The first is an assignment of each Floquet mode to a known undriven state, often called *quantum number assignment*. This is relatively simple – each undriven state's quantum number (which can be multi-indexed for multiple modes) must

be assigned to the Floquet mode that it has most overlap with at low drive amplitude. In fact, at zero drive amplitude, this overlap will always be complete (1). However, at high enough drive amplitude, this overlap may decrease significantly, especially when close to a resonance or in a chaotic regime (in the latter case, the mode overlap with every undriven state becomes small). Resolving this issue sometimes requires tracking the assigned quantum number as one increases drive amplitude, and assigning the states at each subsequent amplitude $|\phi_d|_i$ by their overlap with states at the previous amplitude $|\phi_d|_{i-1}$. This amplitude-tracking has been taken care of in the Floquet packages used in this thesis, and later was implemented and published by the Blais group in [161]. Once the Floquet modes have been successfully assigned to their respective eigenstates, we can return to the problem of converting their quasi-energies into real energies. This simply consists of adding an appropriate number of ω_d's to the quasi-energy to make it similar to the expected real energy (for e.g., the energy of that state at zero drive amplitude), a process we call *Brillouin zone assignment*. If the amplitude tracking for state-assignment has been implemented, one can calculate the Brillouin zone at each amplitude with respect to real energies calculated at the previous amplitude. We will refer to these driven real energies as $E_\beta\left[\phi_d\right] = k_\beta\,\omega_d + \epsilon_\beta$, with β indexing their assigned quantum state.

Once state and Brillouin zone assignments are complete, we formally have access to all the coherent physics in the system in an easily interpretable manner. Non-resonant processes, like driven frequency (Stark or Zeeman) shifts or driven self and cross-Kerr, can simply be calculated by taking appropriate differences of the driven energies:

$$\begin{aligned}
\Delta\omega\left[\phi_d\right] &= \left(E_1\left[\phi_d\right] - E_0\left[\phi_d\right]\right) - \left(E_1[0] - E_0\left[0\right]\right) \\
K[\phi_d] &= \left(E_2\left[\phi_d\right] - E_1\left[\phi_d\right]\right) - \left(E_1\left[\phi_d\right] - E_0\left[\phi_d\right]\right) \\
\chi[\phi_d] &= E_{11}\left[\phi_d\right] - E_{01}\left[\phi_d\right] - E_{10}\left[\phi_d\right] + E_{00}\left[\phi_d\right],
\end{aligned} \tag{3.30}$$

where the indices label Fock states of a one or two mode system. For desired resonant processes, one is generally interested in the resonance condition and strength of the process as a function of drive amplitude. Both of these can be extracted by tracking the energies of the two states that make up the transition, which will always have an avoided crossing around the resonance. As an example, for beamsplitting, one could look at the avoided crossing in between E_{10} and E_{01}, which would have a Lorentzian lineshape as a function of ω_d:

$$E_{10}[\phi_d] - E_{01}[\phi_d] + k\hbar\omega_d \approx \hbar\sqrt{4g_{bs}^2 + k^2\left(\omega_d - \omega_{res}\right)^2}, \tag{3.31}$$

where k is the number of drive photons required to make the process resonant ($k = \pm 1, \pm 2$ for three and four-wave mixing respectively), g_{bs} is the strength of the beamsplitting process, and $\omega_{res} \approx \left(E_{01}[0] - E_{10}[0]\right)/\hbar k$ is the resonant frequency. An example, Fig. 3.3b shows an explicit Floquet simulation of a three-wave mixing beamsplitting process after state and Brillouin-zone assignment, along with the extracted splitting that agrees with Eq. 3.31. Note that while the extracted energies from Floquet are exact, the fit formula assumes contribution from only one sideband – a general formula for multiple tones and sidebands can be found in Eq. 24 of [141].

Finally, to find undesired resonances, one can simply look at either state overlaps or quasi-energies as a function of drive frequencies and amplitudes. The quasi-energies for any system will always show avoided crossings at strong resonances, and along with state assignment and Brillouin zone assignment, provides a clear picture of what resonances are occurring and how many drive photons are required to activate it. However, often it is convenient to have a single number as a metric for whether a certain combination of drive frequency and amplitude could lead to clean desired processes. For a single-mode mixer coupled to external bosonic modes, it is possible to get such a number from mode overlaps by making the assumption that any infidelity in the driven process will primarily

arise from the mixer leaving its driven ground state, since it is the only nonlinear element in this system. One can then calculate the ideal driven ground state of the mixer, e.g., a displaced state for a charge-driven mixer or the undriven ground state for a flux-driven mixer, and calculate the true driven state's overlap with this ideal state:

$$\mathcal{F} \approx |\langle \varphi_\alpha | U_{ideal} | g \rangle|^2. \tag{3.32}$$

Any regions where this number is significantly less than 1 must have undesired driven (coherent) transitions. This is also the method used for applying Floquet theory to the optimization of transmon readout [162, 161], where both $|g\rangle$ and $|e\rangle$ are considered as initial states.

The above analysis provides a complete picture of coherent driven dynamics. However, we still need to account for the noise in our system, which we had previously extracted from our driven EM simulation. To analyze the effect of this noise, we will utilize Floquet-Markov theory, which we will broadly cover here with more details found in [149, 141]. The central intuition for studying the driven effects of noise is to calculate the decay and decoherence of specific elements in the system's Floquet-mode density matrix:

$$\rho_{\alpha\beta} = |\varphi_\alpha\rangle\langle\varphi_\beta|. \tag{3.33}$$

Overall, external noise will cause two types of processes – it will either cause incoherent hopping between states (eg $\rho_{\alpha\alpha} \rightarrow \rho_{\beta\beta}$, or it will cause fluctuations in their energy and hence their phase ($\rho_{\alpha\beta} \rightarrow e^{-\kappa_\varphi t}\rho_{\alpha\beta}$). Similar to Fermi's golden rule, these processes will depend on the matrix elements of the noisy operator in the Floquet basis and the noise-spectral density of the operator. However, importantly, both the matrix elements and the noise spectral density will be sampled at the nominal frequency of the process plus all possible multiples of ω_d, i.e., have contributions from all possible sidebands.

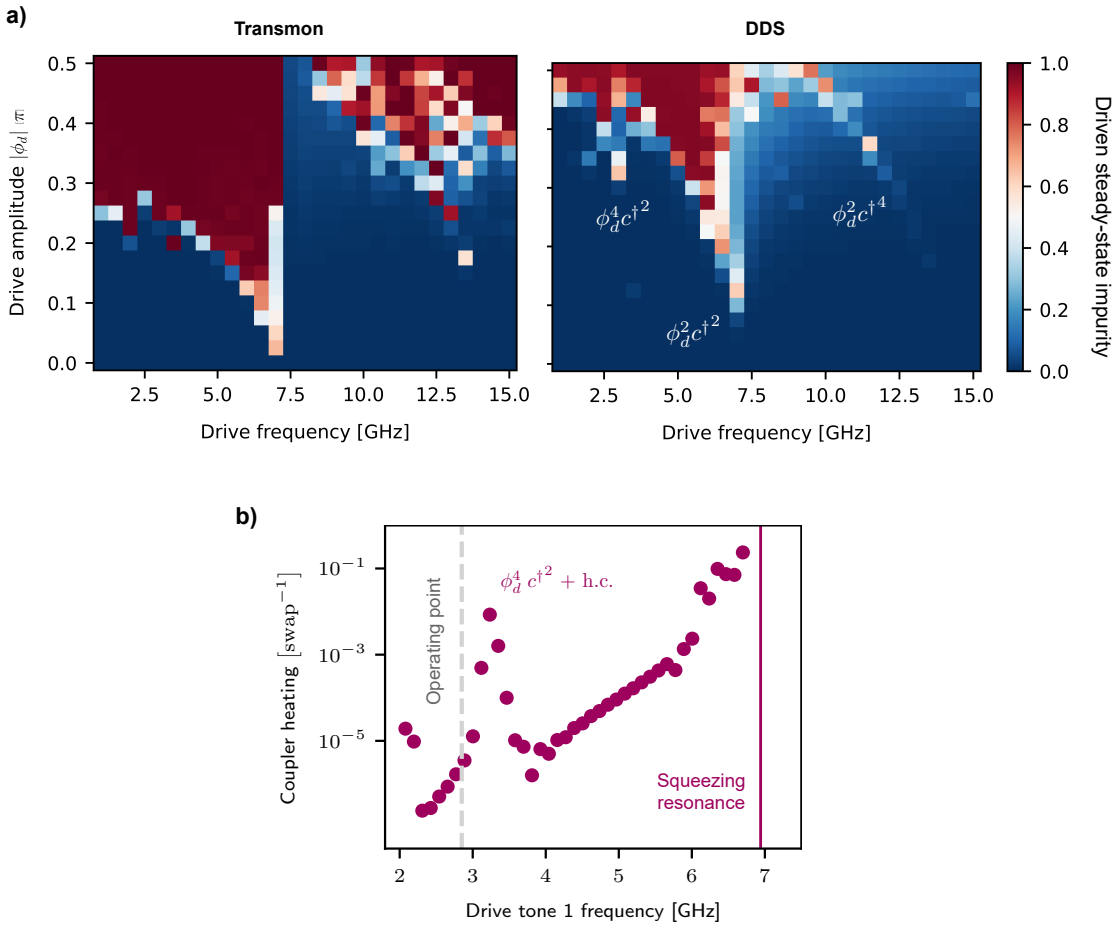

Figure 3.4: **Incoherent Floquet-Markov analysis a**, Driven steady-state impurity $(1 - \mathcal{F}_{\text{incoh.}})$ with a single tone for the transmon and DDS mixers, with fixed coupler decay ($\kappa_c = (100\ \mu\text{s})^{-1}$). The parity protection in the latter shows the clear improvement in the available drive space within which the coupler remains pure. Important parasitic processes are labelled in white. Adding more coupled modes or more drive tones will always add fewer parasitic processes to the DDS, and will either preserve or enhance this improvement in drive-space. **b**, Full two-tone Floquet-Markov impurity simulation for the DDS at a fixed drive amplitude ($|\phi_d| = 0.2\pi$), as a function of drive frequency. The two tones are swept simultaneously, while maintaining their difference at the Alice-Bob beamsplitting resonance condition ($\omega_{d_2} - \omega_{d_1} = \omega_b - \omega_a$). The experimental operating point is chosen just below the $\phi_d^4 c^{\dagger 2}$ sub-harmonic squeezing transition.

Overall the changes in the density matrix in the presence of noise are captured through

the following set of formulae:

$$\begin{aligned}\dot{\rho}_{\alpha\alpha} &= \sum_{\beta} V_{\alpha\beta}\ \rho_{\beta\beta} - V_{\beta\alpha}\ \rho_{\alpha\alpha} \\ \dot{\rho}_{\alpha\beta} &= -\frac{1}{2}\rho_{\alpha\beta}\left(\sum_{\nu\neq\alpha} V_{\nu\alpha} + \sum_{\nu\neq\beta} V_{\nu\beta} + W_{\alpha\beta}\right),\end{aligned} \tag{3.34}$$

where V_{ij} are direct hopping rates from state $|i\rangle \rightarrow |j\rangle$ and W_{ij} are dephasing rates due to energy fluctuations. To calculate these rates, we must find the noise spectral density at the frequency at which we are sampling, the noise matrix element demodulated at that frequency, and the resulting 'nominal' hopping rate, defined respectively as:

$$\begin{aligned} S_{AA}[\Delta_{\alpha\beta,k} = \epsilon_\alpha - \epsilon_\beta + k\omega_d], \\ A_{\alpha\beta,k} = \frac{i}{\kappa}\int_0^T e^{-ik\omega_d t}\langle\varphi_\alpha|A|\varphi_\beta\rangle \text{ , and} \\ \gamma_{\alpha\beta,k} = 2\pi\Theta(\Delta_{\alpha\beta,k})\ S_{AA}[\Delta_{\alpha\beta,k}]\ |A_{\alpha\beta,k}|^2, \end{aligned} \tag{3.35}$$

where $\Theta(\Delta)$ is the Heavyside step function that helps separate positive and negative frequency contributions for thermal noise sources. In the previously described method of finding the noise due to a port by calculating the effective phase drop across the junction due to port voltage fluctuations, $A = i(\tilde{c} - \tilde{c}^\dagger)$ is the charge operator. Finally, the hopping rates $V_{\alpha\beta}$ and $W_{\alpha\beta}$ at some well-defined bath temperature (denoted by the thermal distribution $n_{th}[\omega]$), are given by:

$$\begin{aligned} V_{\alpha\beta} &= \sum_k \gamma_{\alpha\beta,k}(1 + n_{th}[|\Delta_{\alpha\beta,k}|]) + \gamma_{\beta\alpha,-k}\ n_{th}[|\Delta_{\alpha\beta,k}|] \\ W_{\alpha\beta} &= 2\pi\sum_k \Theta(\Delta_k)|A_{\alpha\alpha,k} - A_{\beta\beta,k}|^2\ S_{AA}[\Delta_k]\ (1 + n_{th}[|\Delta_k|]). \end{aligned} \tag{3.36}$$

These equations are complicated, but they essentially say that the bath at the port is thermal and can contribute photons to any sideband according to this thermal distribution.

However, this thermal bath is filtered by the non-white spectral density of the coupling between the system operator (which for us is the junction phase) and the port. The rest is simply an extension of Fermi's golden rule to account for the different frequencies and matrix elements at each sideband.

We have now found our coherent dynamics, described by Floquet modes and their energies, and any incoherent hopping or dephasing within these modes due to environmental noise. How do we use this information? For a clean parametric control on bosonic modes, we want:

1. No parasitic coherent resonances that spoil information in the bosonic modes
2. No driven non-linear effects in the bosonic modes, which includes self-Kerr, cross-Kerr and un-tracked stark-shifts
3. No dephasing of information in the bosonic modes
4. The driven decay rates of the bosonic modes should remain similar to their undriven decay rates
5. The parametric mixer stays in a pure state, ideally the driven ground state, throughout the process.

We have already explored the first two of these, through our analysis of the Floquet modes and energies. The driven dephasing and the decay of the bosonic modes, or a relevant subspace like an encoded dual-rail Bloch sphere, can be directly calculated through Eq. 3.34.

The last requirement is a particularly interesting one. Since it is the mixer that activates the parametric processes, the quasi-energies corresponding to each process may have different values and amplitude or frequency dependence based on the state of the mixer. This means that both the strength and resonance condition of these processes, which are given by the avoided crossing of the quasi-energies, can have an explicit dispersion with

respect to the mixer's state. Hence if the mixer starts in the ground state but incoherently hops to other states, a mechanism often called drive-induced excitation or 'Floquet heating', it would automatically dephase the process. This was a primary limiting factor in transmon-based beamsplitter experiments [163, 141], and will be important to avoid in high-fidelity parametric mixers. At high enough drive amplitudes, such drive-induced incoherent excitation could even lead to chaotic behavior, and cause leakage to Floquet modes that are extremely difficult to recover from, such as those corresponding to energies greater than E_J in the transmon Hamiltonian. Such an event would completely spoil process fidelity, and is also hard to capture in Floquet simulations due to a breakdown of effective state-assignment. Thankfully, one can come up with a simple metric as a proxy for Floquet heating that does not even require state assignment – which is the driven steady state impurity of the coupler density matrix:

$$\mathcal{F}_{\text{incoh.}} = \text{Tr}(\rho_{ss}^2). \tag{3.37}$$

Here, ρ_{ss} can be calculated by setting $\dot{\rho} = 0$ in Eq. 3.34 for the full system, and then tracing out (entanglement to) all other modes, leaving just the coupler. Performing such an analysis for the DDS shows the clear advantage that its parity-protection provides over the regular charge-driven transmon (Fig. 3.4a), even when we simulate just the coupler in the presence of decay and a single drive-tone. Extending this simulation to more tones (Fig. 3.4b) and more modes can be costly, but the parity protection present in the DDS implies that it will always be protected against ∼half the additional processes that might be introduced, preserving its advantage over the transmon. Note that simulating multiple drive tones in Floquet requires using the lowest common multiple of both the periodicity of each tone as the periodicity T, and extending the sideband index k to one index for each drive frequency.

Overall, the above framework provides a way to analyze any nonlinear mixer-based system for any driven process, in the presence of noise. In general, once one has a workflow to extract the metrics for good performance that are listed above, one can optimize the frequency stack within the Floquet analysis, and then re-optimize the EM design to achieve that frequency stack, often through simply tuning junction inductances and drive frequencies. Finally, one must remember that this framework only holds for the continuously driven steady-state performance of the system. Real pulses are more complicated, and include finite ramps. Offering a clean drive space in steady state is generally a good goal for system optimization, but ultimately the exact dynamics must be checked in a full time-domain driven simulation with actual pulse sequences.

3.3 High-fidelity beamsplitting with the DDS

Our primary goal remains bosonic control – as we learnt in this chapter, parametric processes are special drive activated control knobs that could potentially let us achieve clean bosonic control if engineered correctly. To experimentally test this hypothesis, it is important to implement such a protected parametric mixer that remains in its ground state despite being driven strongly, and activates fast control without propagating its own errors to attached bosonic modes. Specifically, we will demonstrate a parametric beamsplitter interaction between two bosonic modes with the DDS. However, as a reminder, the DDS still has Kerr, and therefore can still spoil the linearity of the bosonic modes. We will thus primarily utilize this example to demonstrate that one can indeed be protected from spurious drive-induced excitation, and carefully characterize the coupler's behavior within just the single-photon subspace of two bosonic modes. Even this subspace already implements the useful dual-rail cavity encoding described in Chapter 2, and allows single-qubit gates within this encoding that are at least two orders of magnitude better than any previous

demonstration. We will later extend such (parity) protections to a more linear coupler that is compatible with full bosonic control in the next chapter.

The rest of this chapter will carry over and discuss the beamsplitter results presented in [147]. As a reminder, the protected DDS Hamiltonian is given by:

$$\hat{\mathcal{H}}_{\text{DDS}} = 4E_C\hat{n}_c^2 - E_J \underbrace{\cos(\phi_d)\cos\left(\hat{\theta}_c\right)}_{\text{Even Parity}}, \tag{3.38}$$

where we will use the slightly different notation of $\hat{n}_c, \hat{\theta}_c$ to describe the common-mode quantum operators. We will also use the convenient nomenclature of calling the SQUID's common mode the 'coupler' and its driven differential degree of freedom the 'actuator'. Implementing the protected Hamiltonian in Eq. 3.38 amounts to fabricating a SQUID with symmetric junctions, calibrating it to zero DC flux, and delivering a purely differential drive. The first two conditions are achieved through careful fabrication, and through a dedicated EM coil that tunes the flux to the zero point. The differential drive is implemented through the $\lambda/4$ buffer cavity filter, and careful calibrations of the SQUID's (capacitive) geometry as we described earlier. Importantly, we can directly measure any drive asymmetry present in the experimental device by specifically measuring the rate of 'protected' transitions, like the $\phi_d^3\ c^\dagger$ subharmonic process. We compare the rate of this process in our device to the coupler Zeeman shift at the same drive amplitude, and experimentally bound the asymmetry of the drive to $< 1\%$. A full rendition of the device, as used in HFSS simulations, is shown in Fig. 3.5.

3.3.1 Demonstrating a high-coherence beamsplitter

We now present the demonstration of a high-fidelity beamsplitter that strongly suppresses undesirable coupler heating. The storage bosonic modes, named Alice and Bob respec-

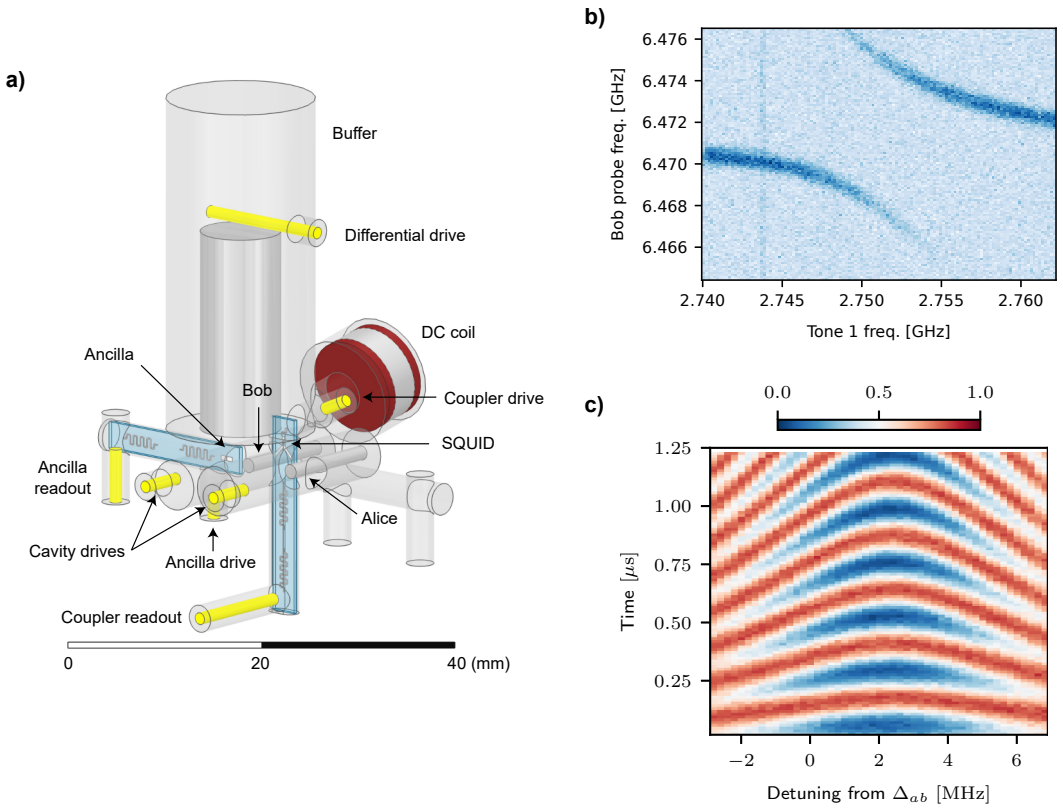

Figure 3.5: **The SQUID package a**, The SQUID experimental package with all relevant details, including dive and readout ports, and accurate size representations for the chips and cavities. **b**, Avoided crossing in the Bob mode, measured using ancilla qubit-assisted spectroscopy, in the presence of two beamsplitting drive tones on the DDS. The first drive tone is swept across this crossing (x-axis), while keeping the second tone fixed at ~ 3.026 GHz. **c**, Corresponding chevron for the time-evolution for the Alice-Bob avoided crossing, at $g_{bs}/2\pi \approx 2$ MHz. The frequency of oscillations at each drive detuning corresponds to the energy splitting in the avoided crossing found in spectroscopy. The x-axis corresponds to detuning from the resonance condition expected from Alice and Bob's static frequencies, which is changed by the drive-induced Zeeman shift.

tively, are capacitively coupled to the Y-shaped antenna of the SQUID [163] and have a negligible mutual inductance to the SQUID loop, ensuring that they exclusively participate [75, 76] in the coupler mode:

$$\hat{\theta}_c \approx \left(\frac{2E_C}{E_J}\right)^{\frac{1}{4}} \left(\frac{g_a}{\Delta_a}\,\hat{a} + \frac{g_b}{\Delta_b}\,\hat{b} + \hat{c}\right) + \text{h.c.} \tag{3.39}$$

Here $\hat{a}$, $\hat{b}$ are the ladder operators for our dressed storage modes, while $\hat{c}$ represents the dressed coupler. The coupling strengths ($g_{a,b}$) and mode detunings ($\Delta_{a,b}$) between the storage modes and the coupler are chosen to be in the dispersive regime $\left(\frac{g_a}{\Delta_a}, \frac{g_b}{\Delta_b} \sim 0.1\right)$. Additionally, we can prepare and readout Fock states in Bob through a dispersively coupled ancilla transmon [117] and a dedicated stripline readout resonator. The SQUID coupler also includes a dedicated readout resonator and drive pin, for explicit characterization of frequency shifts and drive-induced excitation. The measured device parameters are presented in the tables below.

	Coupler	Alice	Bob	Ancilla
Frequency (GHz)	7.245	6.225	6.46	5.663
T_1 (μs)	$\sim$60	375	300	120
T_2^* (μs)	25-35	450	250	5.5
T_2^E (μs)	40	N/A	N/A	25
n_{th}	0.02	$<$0.03	$<$0.03	0.06

Table 3.1: Measured frequencies, coherence times, and thermal populations for the SQUID's coupler mode, the storage modes Alice and Bob, and Bob's coupled ancilla transmon.

Kerrs	Coupler	Alice	Bob	Ancilla
Coupler	-125 MHz	-1.7 MHz	-2.6 MHz	N/A
Alice		-4.9 KHz	-11 KHz	N/A
Bob			-14.6 KHz	-1.2 MHz
Ancilla				-180 MHz

Table 3.2: Measured self-Kerrs and cross-Kerrs for the four modes. The cross-Kerr between the ancilla and the coupler or Alice modes have not been measured, but are assumed to be negligible.

We activate and control the amplitude and phase of our beamsplitter with a bi-chromatic drive on the actuator: $\phi_{d_{1,2}}(t) = |\phi_{d_{1,2}}| \cos\left(\omega_{d_{1,2}} t + \varphi_{d_{1,2}}\right)$. When the difference in our drive frequencies (Δ_d) is close to our cavity detuning (Δ_{ab}), Eq. 3.38 and Eq. 3.39 com-

bine to create a tunable beamsplitter Hamiltonian:

$$\hat{\mathcal{H}}_{\text{BS}}/\hbar = \Delta_{\text{BS}}\ \hat{a}^\dagger\hat{a} + g_{\text{BS}}\ (e^{i\varphi_{\text{BS}}}\hat{a}^\dagger\hat{b} + e^{-i\varphi_{\text{BS}}}\hat{a}\hat{b}^\dagger), \tag{3.40}$$

$$\text{with} \quad g_{\text{BS}} \approx \frac{\omega_c}{2}\frac{g_a\ g_b}{\Delta_a\Delta_b}\ J_1\left(|\phi_{d_1}|\right) J_1\left(|\phi_{d_2}|\right), \tag{3.41}$$

$$\Delta_{\text{BS}} = \Delta_{ab} - \Delta_d + \Delta_{\text{Z},ab} \tag{3.42}$$

where $J_1(|\phi_{d_{1,2}}|)$ is the first-order Bessel function of the drive amplitudes, and φ_{BS} is the beamsplitter phase, controlled by the relative phase of the drives. The corresponding avoided crossing and time-evolution as a function of drive detuning is shown in Fig. 3.5. This time-evolution pattern is a common characterization of driven transitions and and the speed of oscillations as a function of detuning precisely follow the energy splitting in its corresponding avoided crossing (Eq. 3.31 with $k = 2$). The chevron experiment can be used to calibrate any drive-induced frequency offset from the relative AC-Zeeman shift of the cavities ($\Delta_{\text{Z},ab}$), and we can experimentally find the amplitude-dependent resonance condition $\Delta_{\text{BS}} = 0$ to execute a resonant beamsplitter.

We characterize our beamsplitter interaction using the joint single-photon subspace of our storage cavities, which forms a microwave implementation of a dual-rail qubit (Fig. 3.6a). We initialize a single photon in Bob with a preparation fidelity of $\sim 94\%$, by displacing Bob to a coherent state ($\alpha_b = \sqrt{2}$) and using number-resolved measurements through the ancilla to post-select the desired Fock state. We then apply the resonant beamsplitter interaction for a range of times up to 32 μs (Fig. 3.6c), to estimate both the beamsplitting strength and the driven decoherence time of the dual-rail qubit. The evolution of the time-dependent photon population in Bob follows (see Methods in [147] for derivation):

$$P_{\text{Bob}} = \frac{1}{2}e^{-\kappa_1 t}\left(1 + e^{-\kappa_\varphi t}\ \cos\left(2g_{\text{BS}}t\right)\right) \tag{3.43}$$

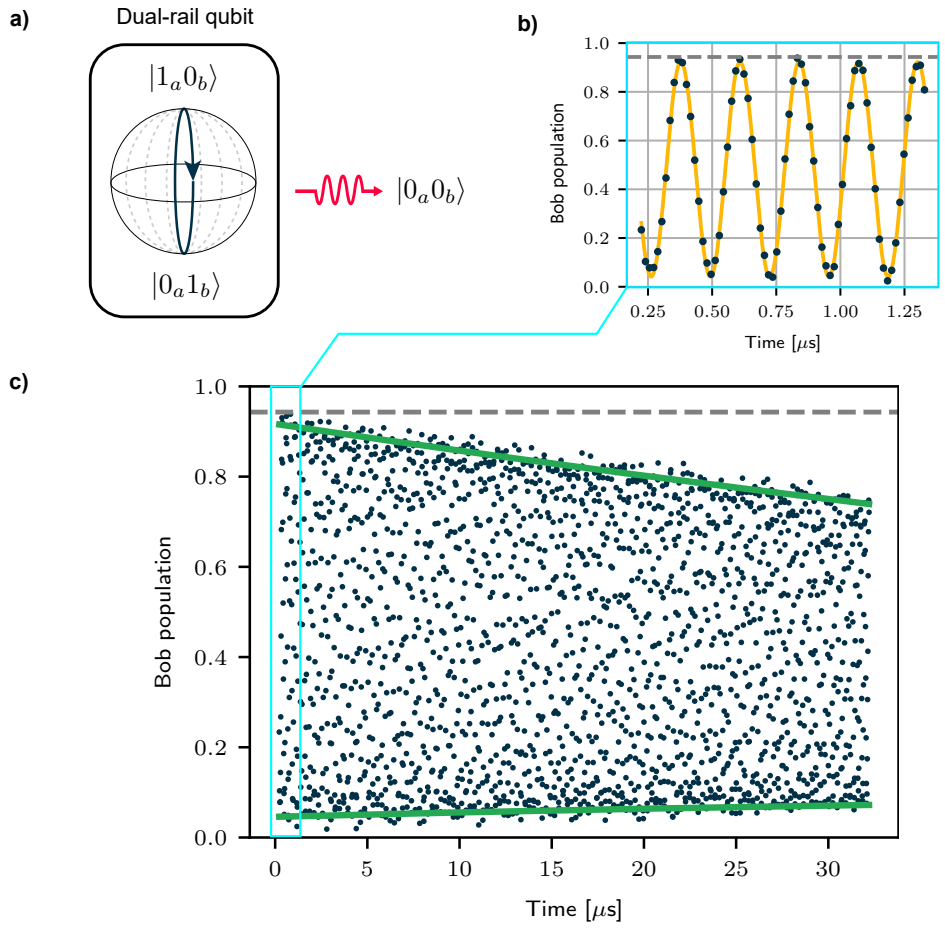

Figure 3.6: **Beamsplitting with the differentially-driven SQUID. a**, Beamsplitting implements an effective driven Rabi evolution in the Bloch sphere of the dual-rail qubit formed by the single photon subspace Alice and Bob, where decay can be detected by monitoring the vacuum state. **b**, Resonant evolution of a single-photon prepared in Bob, with resonant drive conditions found from a Chevron experiment. The data is normalized for readout infidelity, and state preparation fidelity is shown as a dashed gray line. **c** Zooming out in time makes the ratio between the rate of oscillations and the envelope decay clear. The fast coherent oscillations (black dots) between the cavities are fitted to Eq. 3.43 (green lines show envelope) to obtain both the decay and dephasing time-scales.

Here, κ_1 is the mean of the driven cavities' single-photon decay rates, and is the effective rate of population leakage out of the dual-rail subspace into vacuum $|0_a 0_b\rangle$. The dual-rail qubit may also experience dephasing at a rate κ_φ, which would drive this evolution towards an evenly mixed state within the qubit subspace. Combining these lets us place a lower

bound on the expected decoherence limit on the fidelity of a single beamsplitter operation:

$$\mathcal{F} \approx 1 - \frac{\pi}{4}\frac{\kappa_{\mathrm{BS}}}{g_{\mathrm{BS}}}, \quad \kappa_{\mathrm{BS}} = \kappa_1 + \frac{\kappa_\varphi}{2}, \tag{3.44}$$

which is a more accurate metric than the one used in [163].

We choose our operating point to maximize this expected fidelity $\mathcal{F}$ with respect to drive strength, which we characterize by analyzing sections of the resulting driven long-time evolution at various drive amplitudes (Fig. 3.7a). At each amplitude, we extract the g_{BS} and κ_{BS} by fitting two short sections of the evolution to Eq. 3.43. We are able to obtain a maximum $g_{\mathrm{BS}}/2\pi$ exceeding 5 MHz, with the effective decoherence-limited fidelity surpassing 99.9% for a wide range of beamsplitting strengths. In particular, at the operating point, we fit the evolution in Fig. 3.6b to find $g_{\mathrm{BS}}/2\pi = 2.16 \pm 0.01$MHz, $\kappa_1 = (197 \pm 8\ \mu\mathrm{s})^{-1}$ and $\kappa_\varphi = (313 \pm 40\ \mu\mathrm{s})^{-1}$. Our effective $\kappa_{\mathrm{BS}} = (150 \pm 25\ \mu\mathrm{s})^{-1}$ places a decoherence-based upper bound on the fidelity of $\mathcal{F} = 99.96 \pm 0.01\%$, which is almost two orders of magnitude better than the previous transmon-based implementation. Crucially, this fidelity is also limited primarily by photon loss in the cavities, preserving their advantageous noise bias.

To directly quantify the suppression of drive-induced coupler excitation, we measure the coupler population as a function of drive amplitude after resonantly evolving for ten swaps. With our coupler prepared in the ground state, we apply this pulse and measure the coupler's population through a protocol that is robust to readout infidelity [164]. We measure no correlated increase of its driven population as a function of drive amplitude, up to our measurement uncertainty of $\sim 0.2\%$ (Fig. 3.7b). At the operating point, we explicitly quantify the increase in coupler excitation as a function of number of swaps (Fig. 3.7c). We evolve the system up to 100 swaps and observe a total heating rate below $\sim 4 \times 10^{-5}$ excitations per swap, which is consistent with the undriven heating rate of the

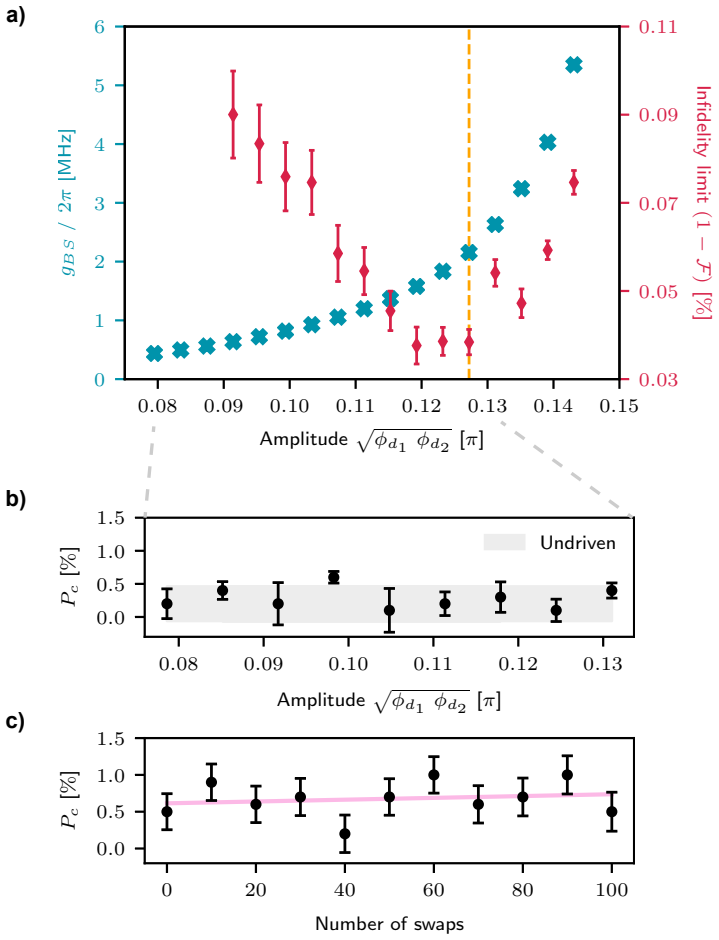

Figure 3.7: **Beamsplitting with the differentially-driven SQUID. a**, Sweeping both drive amplitudes simultaneously and repeating the resonant experiment lets us quantify g_{BS} (blue crosses), and the decoherence limit on beamsplitter infidelity (red diamonds) at various drive strengths. We choose a drive strength with simultaneously low infidelity and high beamsplitter rate as our operating point (yellow dashed line). **b**, The coupler's driven excitation (P_c) after evolving for 10 swaps is directly quantified through a dedicated on-chip readout mode. We observe no monotonic correlation with respect to drive amplitude, and driven populations mostly remain within the range of the undriven population (grey region). **c**, Coupler population as a function of number of swaps at the operational driving point. The heating rate is nearly immeasurable, with a fitted (pink line) slope of $(1.2 \pm 2.4) \times 10^{-5}$ excitation per swap, which is within expectation for our natural thermal background ($\gamma_{c,\uparrow} \sim (3.3\ \mathrm{ms})^{-1}$). The non-zero offset of the fit arises from preparation and readout infidelities. Error bars in both **b** and **c** represent fit errors from the protocol described in [164].

coupler, implying no additional drive-induced heating. This substantial suppression (three orders of magnitude better than transmon-based implementations [163]) eliminates limi-

tations placed by coupler-induced dephasing on the fidelity of the beamsplitter, allowing us to harness the long lifetimes and even longer dephasing times of the 3D cavities.

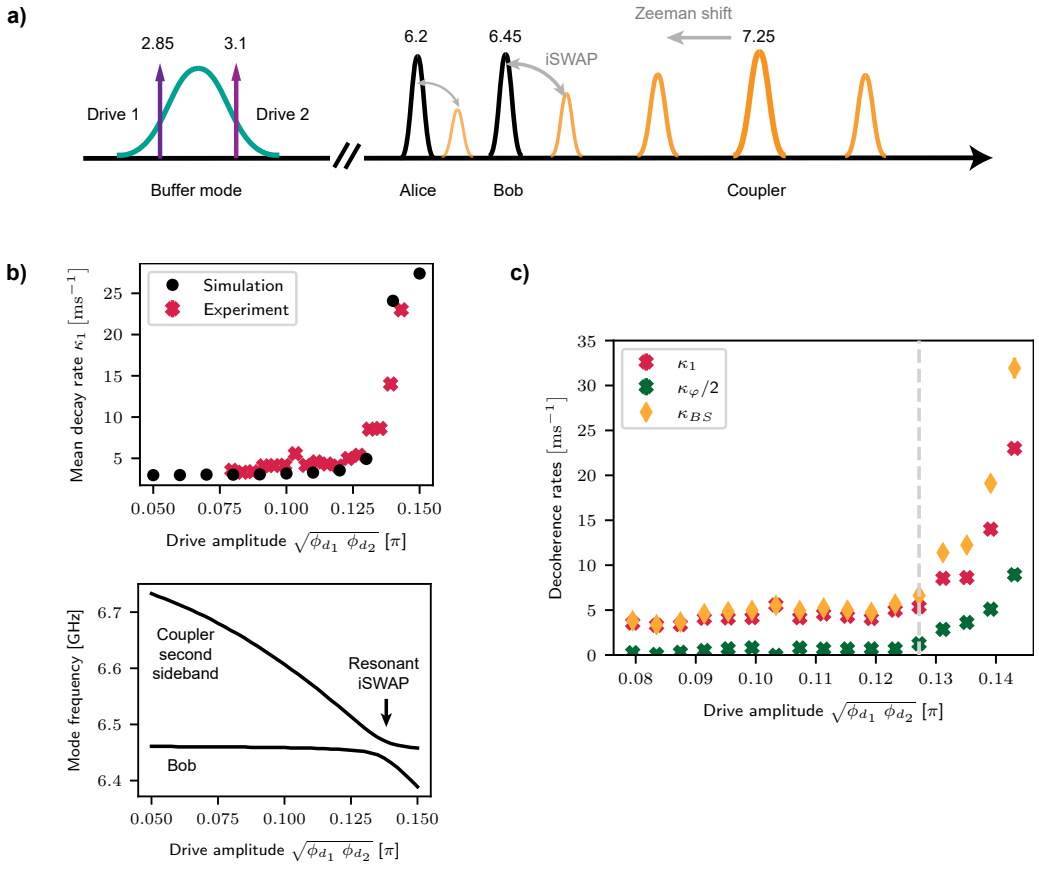

Figure 3.8: **Limits due to coupler-Bob sideband interaction.** **a**, Non-ideal frequency stack leads to sideband collisions between coupler and cavity modes. **b**, Mean decay rate of Alice and Bob as a function of drive amplitude, both theory and experiment. The decay is enhanced by the off-resonant sideband interaction (below), but is well predicted by Floquet-Markov simulations. **c**, The effect of the sideband collision on dual-rail decoherence rates, measured from long-time resonant evolutions at each drive amplitude.

Unfortunately, the fabricated coupler frequency was lower than the optimal frequency found in Floquet simulations, resulting in certain non-idealities. We observe that upon increasing the drive strength, the coupler frequency shifts closer to the cavities by ~ 200 MHz, and a direct sideband interaction between the coupler and the cavities likely increases their hybridization. This effect results in a faster-than-quadratic dependence of

the beamsplitting strength on our drive amplitudes, but also reduces the coherence of the dual-rail subspace at higher amplitudes due to the aggravated Purcell loss, creating an optimum infidelity point. Floquet-Markov simulations capture this sideband collision well, with accurate predictions for Bob's Purcell limited decay rate (see Fig. 3.8). Additionally, this also causes a drive-induced increase in the self-Kerr of the cavities, with up to 128 KHz of inherited Kerr at the operating point. This can lead to coherent errors when trying to operate the beamsplitter in higher-photon manifolds. Numerous avenues exist to minimize this driven non-linearity if desired, including fabricating a slightly higher coupler frequency, arraying multiple SQUIDs, or dynamical Kerr cancellation [137]. Fully linear bosonic control will however require using an alternative scheme like Kerr-free three-wave mixing [165, 166]. We will discuss both these avenues towards the end of this chapter.

3.3.2 Benchmarking fidelity in the single-photon subspace

We now precisely characterize the fidelity and noise bias of our beamsplitter by using it to implement universal control of the dual-rail qubit subspace, allowing techniques akin to standard randomized benchmarking (RB) protocols [39, 167, 168]. First, we identify the amplitude and frequency required for a fixed-length pulse (Fig. 3.9a) to achieve the beamsplitter unitary

$$U_{\text{BS}}\left(\varphi{=}0\right) = e^{i\pi/4(\hat{a}^\dagger\hat{b}+\hat{a}\hat{b}^\dagger)}, \tag{3.45}$$

by repeating the pulse to perform swaps between Alice and Bob, and iteratively checking its performance up to ~ 1000 such repetitions. This calibrates an effective $X_{\pi/2}$ gate for the dual-rail qubit, with a relative amplitude precision of less than 3×10^{-6} and frequency precision of less than 3×10^{-7}. Through phase control and repetition, we use this pulse to

construct a set of native gates,

$$G_{\mathrm{DR}} = \{X_{\pi/2}, Y_{\pi/2}, X_{-\pi/2}, Y_{-\pi/2}, X_{\pi}, Y_{\pi}\}, \tag{3.46}$$

that generate the Clifford group for the dual rail qubit. This allows a form of direct randomized benchmarking [169], which under uniform sampling should convert both dephasing and coherent control errors into an effective depolarization channel. The dominant but detectable error of cavity photon loss appears as a leakage to the orthogonal state $|0_a 0_b\rangle$, which is not converted to depolarization under this protocol, but can be separately quantified and selected out in post-processing.

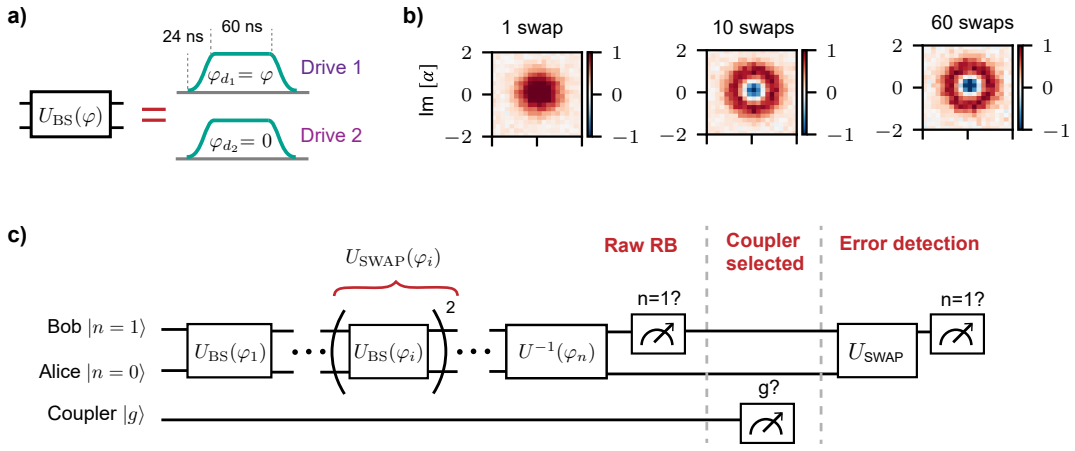

Figure 3.9: **Pulse sequences for randomized benchmarking. a)**, The gate-sets required for the above protocols are generated from calibrated beampslitter pulses with tanh-shaped ramps, where different $U_{\mathrm{BS}}(\varphi)$ are obtained by changing the relative phase of our drives. **b)**, Wigner function of Bob after preparing $|0_a 1_b\rangle$ and implementing 1, 10 and 60 calibrated swaps. **c)**, The benchmarking sequences consist of randomly generated pulses that, under ideal operation, map $|0_a 1_b\rangle$ back to itself. After each sequence, we measure whether the coupler is in its ground state, the presence of a photon in Bob, and the presence of a photon in Alice using an additional swap gate. All sequences are also conditioned on Bob's ancilla ending in its ground state, to discount first-order effects of ancilla heating.

The RB protocol consists of initializing the system in $|0_a 1_b\rangle$ with the coupler and ancilla prepared in their ground states, and running sequences of varying lengths of randomly

chosen gates from G_{DR}. Each sequence ends with an additional gate from G_{DR} that maps the state back to $|0_a 1_b\rangle$, after which the presence of an excitation in the coupler, Bob, and Alice are measured (Fig. 3.9c). We explicitly discount the effects of the ancilla in all results shown by separately measuring the ancilla and only including sequences where it ends in the ground state. The sequences range up to 8100 gates, where we choose up to 900 random gates (limited by FPGA memory), and repeat each gate nine times to fully capture the fidelity decay timescale (this is roughly equivalent to performing nine times as many random gates and provides a lower-bound for the fidelity of a non-repeated sequence). We average over $\sim 10^5$ such semi-random sequences at each sequence length.

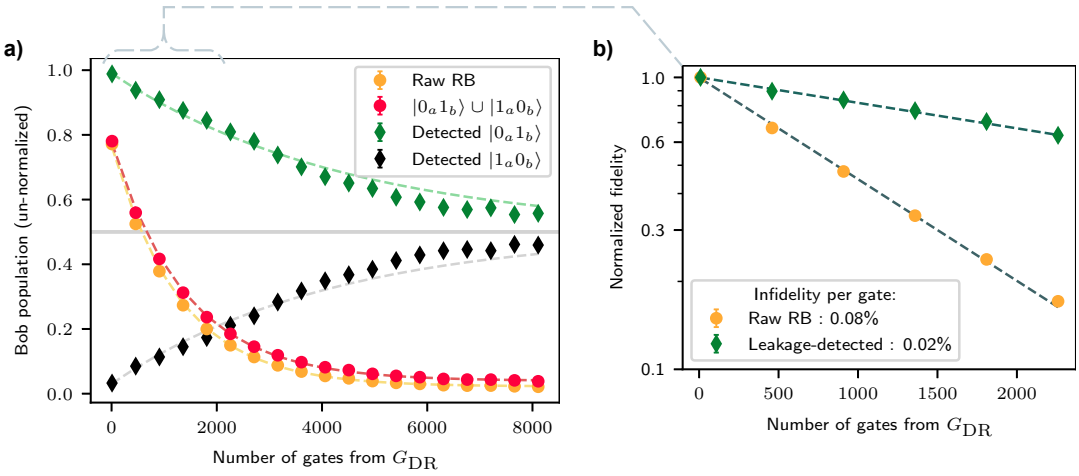

Figure 3.10: **Randomized benchmarking of a dual-rail cavity qubit. a)**, We analyze in detail a single RB experiment with random sets of 9 repeated gates, under various selection protocols, with no normalization. Dashed lines for each curve represent exponential fits. Curves without the error-detected postselection on photon-loss errors decay to zero, whereas the error-detected curves decohere to 0.5. **b)**, Focusing on the first 2250 gates, we use measurements of both cavities to post-select on sequences in which no photon loss event occurred (green diamonds). We compare these sequences to a normalized version of the raw RB in **a)** (yellow), showing an improvement in average gate infidelity from $0.078 \pm 0.001\%$ to $0.020 \pm 0.001\%$.

We observe that the average success probability of returning to Bob decays exponentially, and plot and fit these curves for various selection protocols in Fig. 3.10a. The raw randomized benchmarking (yellow) shows a decay constant of $\tau_{\mathrm{RB}} = 1271 \pm 4$ gates, with

its reduced amplitude corresponding to a SPAM error of $\sim 21\%$. We first compare this to a leakage-limited curve (red), which represents the total rate of photon loss out of the dual-rail subspace, and does not distinguish between states on the dual-rail Bloch sphere. These curves have large overlap, implying that we are limited by photon loss errors, with a leakage-limited decay constant of 1388 ± 4 gates. On detecting and selecting out these leakage events, we obtain error-detected curves where the system ended in either $|0_a 1_b\rangle$ (green) or $|1_a 0_b\rangle$ (black), but never in $|0_a 0_b\rangle$. These sequences on average represent a depolarization channel, and decay approximately towards a perfectly mixed state (0.5, grey horizontal line) with improved decay times of 4477 ± 30 and 4194 ± 32 gates respectively.

To explicitly quantify our beamsplitter's infidelity and noise bias, we focus on just the first 2250 gates, where slower effects like double jumps and reheating from vacuum are negligible (Fig. 3.10b, plotted on a log scale). Within this span, the raw decay of Bob's population suggests an un-selected single gate infidelity of $0.078 \pm 0.001\%$. For a fair comparison between the raw and error-detected RB curves, we re-scale the former's fidelity to also have a steady-state of zero. This helps illustrate the clear improvement under error-detection to an infidelity per gate of $0.020 \pm 0.001\%$, which means that discarding only one out of every ~ 1300 shots per gate can lead to a $3.9 \pm 0.2\times$ increase in gate performance, enabled by the cavity noise-bias. This reiterates the fact that the dominant errors in the system remain cavity photon loss, which is detectable or correctable with various bosonic encoding schemes, and satisfies one of the key goals for a high-performance bosonic control.

Notably, because our gate-set is crafted from nearly identical beamsplitter pulses, with a random gate from G_{DR} containing $4/3$ beamsplitters on average, we can directly convert these RB infidelities into an effective fidelity of a single beamsplitter gate. Our measurements imply an effective un-selected beamsplitter fidelity of $99.941 \pm 0.001\%$, which improves on leakage detection to $99.985 \pm 0.001\%$. The remaining errors after leakage

detection can be due to intrinsic dephasing of the cavities, drifts in our control electronics, or other effects that are not treated by the post-selection protocols, like cascaded heating and decay events of the ancilla [113].

Discussions and implications

Overall, we demonstrated that parity protected Josephson mixers can allow clean and high-fidelity control in a bosonic architecture. This took the form of a tunable cavity-cavity beamsplitter, which we characterized using a dual-rail encoding in the two cavities. We obtained a beamsplitter gate fidelity exceeding 99.94%, which corresponds to a generator for all Clifford gates on this qubit, with non-Cliffords being simple extensions to shorter pulses. Importantly, we were limited by detectable single-photon loss in the cavities, which provided a $\sim 4\times$ improvement on being post-selected out. This performance was enabled through careful symmetry and Floquet engineering that kept the coupler in its ground state even when driving a fast beamsplitter, and avoided any other parasitic driven mechanisms. In general, such high-fidelity control using a strongly driven lossy nonlinear element is a significant step forward for fast parametric operations in circuit-QED, and quantum control more broadly.

This particular system utilized 3D cavities for both the bosonic modes and the drive delivery, but the generality and versatility of the design framework far exceed this specific design. The DDS could be easily extended to performing a beamsplitter between on-chip resonators, or phononic modes in hybrid architectures. It will also be interesting to test the implementation of other Gaussian processes using the DDS, and check whether they indeed have parity protection and suppressed coupler-induced infidelity. Beyond the context of parametric interactions, this experiment also demonstrated the delivery of AC flux in a high-Q (superconducting) 3D package. This has previously proved difficult, for

example due to the Meissner effect, but the numerical simulation techniques and general high-Q hygiene utilized in this project can be easily extended towards more compact drive delivery design.

Finally, the demonstration of high-fidelity control in the dual-rail subspace motivates the hypothesis that this architecture may be used as a building block for a scaled quantum processor [170, 171, 113]. As discussed previously in Chapter 2 the error-hierarchy of detectable decay over dephasing makes the dual-rail qubit amenable to erasure conversion [172], which can potentially yield higher thresholds and effective code distances in the surface-code architecture. One can also extend the single-qubit control demonstrated here to realize a full high-fidelity gate-set [85] for multi-qubit control, which will be necessary for a general dual-rail qubit-based architecture in circuit-QED.

However, within the context of general bosonic control, the DDS had two primary pitfalls. The first is the driven frequency shift, which eventually causes the collision detailed in Fig. 3.8 and places a limit on simultaneously increasing beamsplitting speed and fidelity. As we derived earlier in this section, these driven shifts are mostly drive-frequency independent (assuming the same amplitude of driven phase across the junctions), which means they have no resonance condition that one can avoid or tune away from. This is a natural consequence of utilizing four-wave mixers, and usually causes all parasitic resonances to shift with drive amplitude, significantly complicating the navigation of available drive space. The second issue is that, again due to the fourth-order nonlinearity, there exist significant idle and driven anharmonicity in any oscillators that are coupled to DDS. This makes achieving continuous variable control in a cavity-state independent manner very difficult – for example, the resonance condition for the beamsplitter interaction in the experiment above shifts with photon number, making it impossible to swap arbitrary states.

One possibility to suppress this anharmonicity is to realize that the beamsplitting

strength of the SQUID and its Kerr nonlinearity arises from different fourth-order terms. Specifically, the former is $\propto E_J\ \phi_d^2\ \hat{\theta}_c^2$, which in the notation established in Eq. 3.10 is the $g_{2,2}$ nonlinearity, and the latter is $\propto E_J\ \hat{\theta}_c^4$ (i.e., the $g_{4,0}$ nonlinearity. Inspecting these two terms makes it clear that simultaneously suppressing the value of $\hat{\theta}_c$ and boosting E_J can preserve the beamsplitting strength while suppressing Kerr. This is precisely what arraying the SQUID element achieves. If one places N differentially-driven SQUID loops in a series, the total phase drop across this arrayed coupler is divided up into smaller phase drops across each element, ie $\hat{\theta}_c \rightarrow \hat{\theta}_c/N$. Now, if one also scales each SQUID junction area simultaneously by the same factor, which causes $E_J \rightarrow NE_J$, then the new DDS Hamiltonian is given by:

$$\begin{aligned} H_{\text{DDS}} &= 4E_C\hat{n}^2 - N^2 E_J \cos(\phi_d) \cos\left(\frac{\hat{\theta}_c}{N}\right) \\ &= 4E_C\hat{n}^2 + E_J \cos(\phi_d) \left(\frac{\hat{\theta}_c^2}{2} - \frac{1}{N^2}\frac{\hat{\theta}_c^4}{4!} + ...\right) \end{aligned} \tag{3.47}$$

It is clear that the strength of the useful four-wave mixing, and therefore all the desired parametric Gaussian processes, are preserved, while all higher order process of order $2k+2$ are suppressed by a factor of $1/N^{2k}$. This idea of arraying dipole mixing elements is fairly general, trading fabrication difficulty for resilience to higher order-nonlinearity, and we will use similar schemes in the next chapter too. Arraying has also extensively been used in the quantum-limited amplifier context, for example by varying the effective inductance of each element in the array slightly, or by matching the periodicity of the array to a traveling wave, to gain a large amplification bandwidth [173].

While arraying the DDS forms a viable solution if multiple SQUID loops can be arranged to be differentially driven, it does not completely solve the nonlinearity problem. An alternative solution would be to utilize a three-wave mixer (3WM) instead, which has its Kerr explicitly set to zero. Since every Josephson junction can only have a cosine

nonlinearity, simply combining such elements cannot create an odd-order nonlinear term. Instead, to create a 3WM, one needs to introduce an external field that explicitly breaks this symmetry – which is simply a DC external magnetic field. The most prominent three-wave mixer currently used for bosonic control is the superconducting nonlinear asymmetric inductive element (SNAIL), which is an asymmetric single-loop dipole element that is threaded by DC magnetic flux [174]. It is usually implemented as some approximation of a junction shunted by a linear inductor (which in the ideal case would be called an 'RF-SQUID'), where the inductor is instead approximated by an array of 3 to 5 junctions. This element's Hamiltonian is often expressed as:

$$H_{\text{SNAIL}} = 4E_C\hat{n}^2 - \alpha E_J \cos\left(\hat{\theta} - \phi_{\text{DC}}\right) - ME_J \cos\left(\frac{\hat{\theta}}{M}\right), \tag{3.48}$$

where α is the ratio between the small and arrayed big junctions, and M is the number of junctions in the array, with $\alpha < 1/M$ to prevent flux points where the potential minima is multi-valued. Expanding the SNAIL potential to derive its charge-driven mixing parameters follows the general formulation in Eq. 3.13 – except that unlike the transmon, its θ_{min} can be non-zero! This allows the introduction of nonlinear terms ($g_{m,0}$) at both even and odd orders, all of which depend on the flux in the SNAIL loop in different ways, and could each potentially be tuned to zero at specific flux points. Additionally, any practical realization of the SNAIL requires capacitive 'arms' that couple to the modes of interest, which adds some parasitic linear inductance in series with the mixer. We will study how to treat such parasitic inductors properly in the next chapter, but in general this requires finding the solutions to transcendental current conservation equations that change the SNAIL's nonlinearity at all orders. Finally, the actual non-linearities of interest, like the Kerr and cross-Kerr to the bosonic modes we want to control, get dressed by interactions at various orders of nonlinearity. This means simply setting the fourth order nonlinearity to zero is

not sufficient to make the SNAIL effectively linear, one must choose between either setting the Kerr, the stark shift, or the cross-Kerr to zero, each of which depend on at least $g_{3,0}$ and $g_{4,0}$. In combination, these effects make designing a SNAIL coupler a task of precise numerical optimization, and prevent simple design intuitions for its Hamiltonian. Practically, with any fabricated SNAIL device, one can find various points of interest by explicitly tuning flux and measuring parameters in the SNAIL and bosonic modes. In fact, around the same time as above beamsplitting experiment with a DDS, it was also shown that a SNAIL could successfully perform as a high-performance Kerr-free coupler, providing a significantly more linear control than the DDS [166]. More details about analyzing the SNAIL as a mixer and as a coupler for bosonic control can be found in N. Frattini and S. deGraaf's theses [99, 175].

The complicated design of the SNAIL Hamiltonian is a more minor downside – its main issue of course is that it introduces nonlinear terms at *all* orders. This makes it potentially have even more parasitic resonances than the transmon, which can only turn on processes that have total drive + mode photons as even. This begs the crucial question: since we know that we require (Kerr-free) three-wave mixers to achieve linear bosonic control, and that engineering clean mixers benefits from selection rules like parity protection, can we design a parity-protected three-wave mixer? Indeed we can, using a novel nonlinear driven element introduced in the next chapter – the Linear Inductive Coupler (LINC). Importantly, the LINC will provide a cleaner drive-space than both the DDS and the SNAIL, while simultaneously having analytically predictable driven processes, potentially allowing simple high-fidelity CV control.

Chapter 4

The Linear Inductive Coupler

Before we dive into describing the LINC, it is useful to zoom out and ask the question: What is the ideal parametric three-wave mixer? For a single-mode mixer with a well-defined frequency that can activate all possible three-wave mixing (Gaussian) processes, it would look like:

$$H_{\text{ideal}}/\hbar = \omega_c \; c^\dagger c + g_{\text{3wm}} \; \phi(t) \; (c + c^\dagger)^2. \tag{4.1}$$

This can actually be re-written in a language much closer to usual circuit analysis:

$$H_{\text{ideal}} = 4E_C \; \hat{n}^2 + \frac{E_L}{2} \left(1 + k\phi\right) \; \hat{\theta}^2, \tag{4.2}$$

which is simply an oscillator with a linearly modulated linear inductor, ie with $E_L(t) = E_L \; + \delta E_L \sin(\omega t)$ where $\delta E_L = k E_L |\phi|$ if $\phi(t) = |\phi| \sin(\omega t)$. The mode frequency is then given by $\omega_c(t) = \sqrt{8 E_L(t) E_C}$, and k sets the relative strength of the three-wave mixing process, with $g_{\text{3wm}} = k E_L / 2\theta_{zpf}^2$. This Hamiltonian has several interesting properties. The first is that when the parameter ϕ is set to zero, the mixer loses any effect of its nonlinearity – it is identical to a perfectly linear oscillator. However, the *slope* of the potential energy is actually maximum at this point. This means any noise in the parameter ϕ has maximal

effect on the mixer's Hamiltonian, ie it is at the anti-sweet spot with respect to ϕ. This is simply a direct consequence of the linear response physics described in Chapter 1, and is a natural trade-off when moving to three-wave mixing. The second is that while it is indeed true that its frequency is modulated, this modulation is *not* zero on average. Specifically, even this ideal mixer has a driven frequency shift of:

$$\begin{aligned}\hbar\Delta\omega_{\text{ideal}} &= \sqrt{8E_C E_L}\left[1 - \left\langle\sqrt{1 + k|\phi|\sin(\omega t)}\right\rangle\right] \\ &\approx -\sqrt{8E_L E_C} \times \frac{k^2|\phi|^2}{16}.\end{aligned} \tag{4.3}$$

Note that the exact shift is a complicated function of elliptic integrals, but here we have used a simply Taylor expansion in $k|\phi|$. Lastly, the ideal general purpose three-wave mixer is actually not the same as the ideal coupler that purely activates the beamsplitting interaction. The latter is given by:

$$H_{\text{ideal, BS}}/\hbar = \omega_c(1 + k\phi)\, c^\dagger c, \tag{4.4}$$

which in circuit language would be a simultaneous modulation of inductive and capacitive energies such that their ratio was constant. This ideal coupler would have no driven frequency shifts and no parasitic resonances, but might be difficult to realize in circuit-QED due to a lack of capacitive nonlinearities.

How would one realize the ideal 3WM in circuit-QED? Since Josephson tunnel junctions only allow sinusoidal nonlinearities, one could try constructing the above Hamiltonian through an arrayed nonlinear element, such that in the limit of a large enough array and small enough drives, the higher order terms from the sinusoids disappear. More precisely, by replacing any odd-order non-linear terms in Eq. 4.2 by a $\sin$ function, and even order terms by the $\cos$ function, one could convert the ideal Hamiltonian into the more

circuit-QED friendly form:

$$\begin{aligned} H_{\text{LINC}}^{(N)} &= 4E_C\hat{n}_c^2 + \frac{E_L}{2}\hat{\theta}_c^2 - N^2\alpha E_L \sin\phi_d \cos\frac{\hat{\theta}_c}{N}, \\ \Rightarrow H_{\text{LINC}}^{(\infty)} &= 4E_C\hat{n}_c^2 + \frac{E_L}{2}\left(1 + 2\alpha\ \phi_d\right)\hat{\theta}_c^2 = H_{\text{ideal}}. \end{aligned} \tag{4.5}$$

Here we have used the arraying conventions similar to Eq. 3.47, with each inductive element's energy scaled by a factor N for an array of N elements. This Hamiltonian is precisely the arrayed version of the LINC, which we will derive, study, and implement in the rest of this chapter.

To gain an intuition for how the LINC may be designed, and how to think about its symmetry protections, we can simply take the parity-protected circuit we know (the DDS) and mold it into a three-wave mixer. First, notice that the standard DDS (Eq. 3.4) already has a point, $\phi_{\text{DC}} = \pi/2$, where it appears to perform three-wave mixing:

$$H_{\text{DDS}}[\phi_{\text{DC}} = \pi/2] = 4E_C\hat{n}_c^2 - E_J \sin\phi_{AC} \cos\hat{\theta}_c. \tag{4.6}$$

However, this operating point has a glaring issue – when the drive is not on ($\phi_{AC} = 0$), the entire inductive energy disappears, and so does the mode! This can be understood as biasing the SQUID loop to a point where the junction phases exactly differ by π, which means they perfectly destructively interfere to cancel out the total inductive energy (or make the effective inductance of the element $\to \infty$). However, this is a feature, not a bug. At this special bias point, for a perfectly symmetric SQUID, all orders of nonlinearity simultaneously disappear. All we need to add to make this into a usable nonlinear element is an inductive element that does not tune with flux, and gives the circuit a well-defined mode frequency. We can simply do this with a linear inductor, placed in a precise manner such that there is no driven phase across it – for eg, in a uniform flux drive, this would

be placed exactly at the center of the loop. For convenience, we will assume the outer junctions have energy E_J each (which differs from the DDS convention by a factor of 2, but there is no single-junction equivalent to compare to here), and the inductive shunt has energy E_L. Together, this gives the LINC Hamiltonian:

$$H_{\text{LINC}} = \underbrace{4E_C\hat{n}_c^2 + E_L\frac{\hat{\theta}_c^2}{2}}_{\text{static}} + \underbrace{2E_J \sin\phi_{\text{AC}} \cos\hat{\theta}_c}_{\text{driven}}, \tag{4.7}$$

which is precisely a DDS operated at $\phi_{\text{DC}} = \pi/2$, shunted by a linear inductor. Comparing this equation to $H^{(1)}_{\text{LINC}}$ in Eq. 4.5, we see that they match, with the convention $\alpha = 2E_J/E_L$. Note that any static charge offsets are also redistributed by the shunting inductor, and hence n_g is also explicitly absent in the Hamiltonian. Finally, for the circuit's potential energy to have a single-valued solution at all values of DC flux, its total inductive energy must always be positive (its frequency must never cross zero), which constrains $E_L > 2E_J$.

It is important to state that while the simultaneous balancing of the drive and the bias of the LINC is novel and crucial to its function, the circuit topology itself is fairly common in literature [176, 177, 178, 179, 180]. The LINC is a direct sibling to the Asymmetrically Threaded SQUID (ATS), with nominally identical circuit geometries but an exactly opposing notion of DC bias point. This allows the LINC to be optimized as a three-wave mixing coupler, while the ATS is utilized for Kerr-free four-wave mixing applications, like parametric two-photon dissipation [176]. The LINC also shares its topology with the gradiometric SNAIL [177], where the flux bias in the SQUID loops is used to tune the effective junction ratio of a regular charge-driven SNAIL, and symmetry is not a strict constraint. One key physical feature where the LINC might differ from these circuits is in its shunting inductor, which for the SNAIL and ATS typically consists of an array of a few junctions. We will see that in the LINC, this center shunt sets the entire idle potential

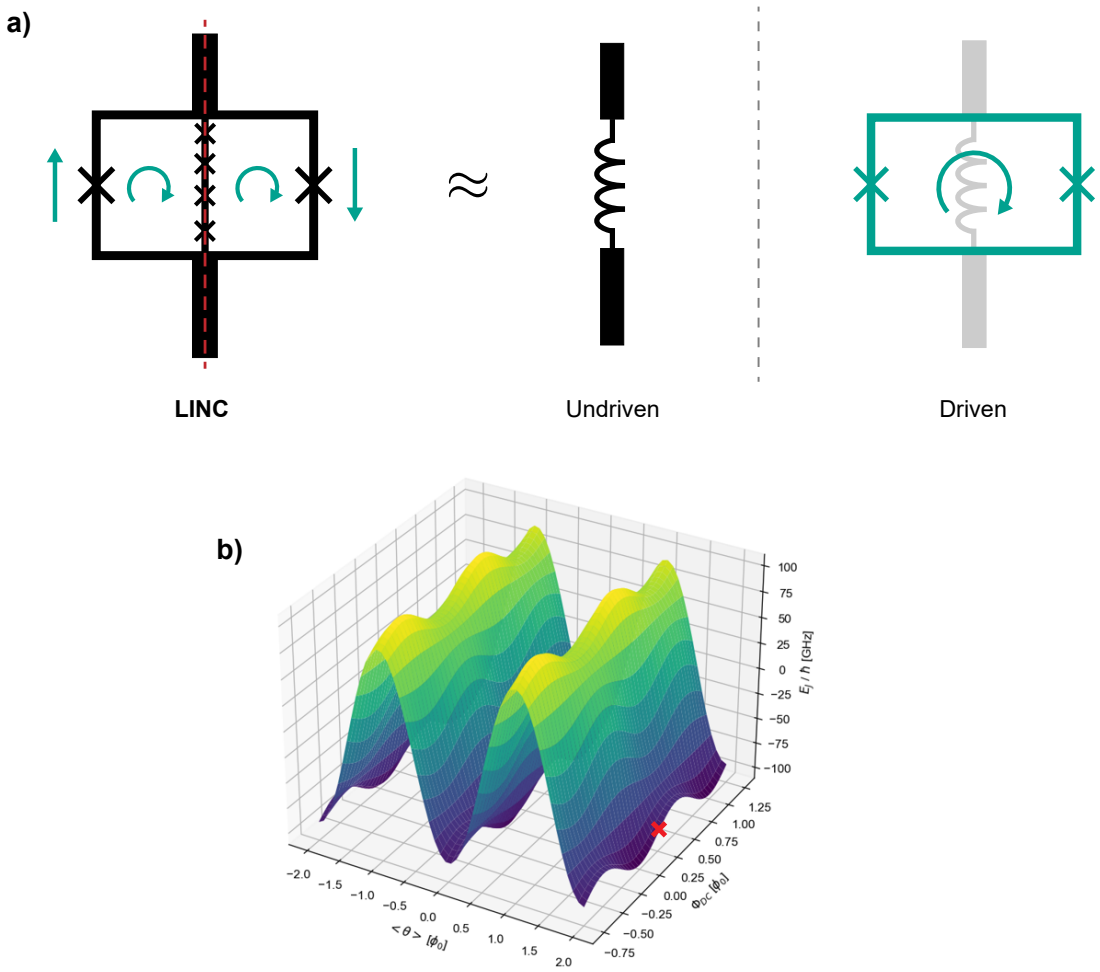

Figure 4.1: **The LINC circuit and potential a**, The LINC circuit, with its center shunt in general composed of a junction array. The flux drives and bias must exactly cancel across the shunt to maintain orthogonality. The circuit can then be split exactly into its undriven (linear oscillator) and driven (SQUID at $\phi_{\text{DC}} = \pi/2$) sub-circuits. **b**, The LINC's bi-variate potential energy $U_{\text{LINC}}(\hat{\theta}, \phi)$, as a function of the expectation value of both variables. The height and periodicity along $\hat{\theta}$ is determined by the number of shunt junctions M (here $M = 2$ for exaggeration), reaching a quadratic potential for a perfect linear inductor. Red cross denotes operating point minima.

energy of the circuit. Thus for bosonic applications that might require highly linear idle circuits, the LINC shunt could benefit from using high kinetic inductance materials such as NbN or granular Aluminum (grAl) [181, 182], or a meandering geometric inductor [183], to minimize its static nonlinearity.

4.1 Analyzing the LINC Hamiltonian

Now that we have established the LINC Hamiltonian and its relation to both the DDS and the ideal three-wave mixer, we can proceed to fully analyze its idle and driven behavior. Most of this section reproduces the work in [184].

The primary idea behind the LINC's performance is that its Hamiltonian is exactly divided into two parts, the static and the driven, as suggested in Eq. 4.7 and Fig. 4.1a. Specifically, at the special flux point of $\phi_{\text{DC}} = \frac{\pi}{2}$, the nonlinearity of the circuit completely disappears when idle, resulting in an exactly linear static Hamiltonian:

$$H_{\text{LINC}}\left[\frac{\pi}{2}\right] = 4E_C\hat{n}_c^2 + E_L\frac{\hat{\theta}_c^2}{2}. \tag{4.8}$$

This is beyond a simple 'Kerr-free' mixer like the SNAIL – at this flux point, the circuit's static nonlinearity at all orders *simultaneously* vanishes. We will show later that this property holds even in the presence of experimental imperfections like a parasitic series linear inductance, which as discussed previously, can complicate the mixer's nonlinear behavior. The mixer hence provides an extremely linear environment to any modes it mediates a coupling between in the coupler context, which is particularly useful for bosonic control. This is why we name the mixer the 'Linear INductive Coupler'. When driven, the LINC still contains beneficial symmetries that are very similar to the parity protection in the DDS. Only processes where the number of coupler and resonator photons are even can be activated. Even within these, processes that are even-order in the drive, like the drive-induced frequency shift, are suppressed. We will discuss these protections in detail later in this section.

Let us first discuss the overarching formalism behind analyzing the LINC circuit. A general LINC-like circuit contains three inductive branches, and two galvanic loops that

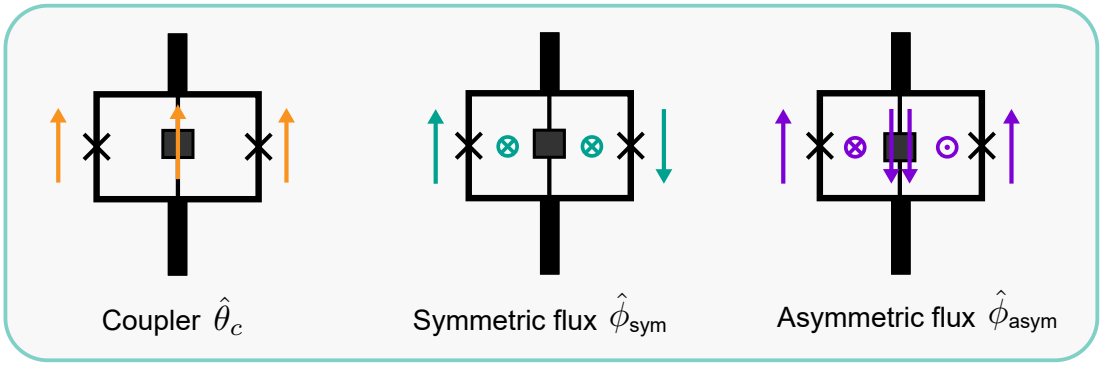

Figure 4.2: **Mode decomposition for the LINC circuit** The general LINC-like circuit, with an arbitrary inductive element as its center shunt, is defined by the three variables $\{\theta_1, \theta_2, \theta_s\}$ across the outer two symmetric branches and the shunt respectively. These can be re-written in terms of the more operationally relevant charge-dipole ($\hat{\theta}_c$), symmetric flux ($\hat{\phi}_{sym}$), and anti-symmetric flux ($\hat{\phi}_{asym}$) modes.

can be independently threaded by DC magnetic flux, setting its 'operating point'. In the presence of non-zero field, each of these three branches can incur a different voltage drop across them, with corresponding superconducting phase drops across the three inductive elements, $\hat{\theta}_1$, $\hat{\theta}_s$ and $\hat{\theta}_2$. These phase drops can be rewritten in terms of more easily interpretable independent variables (Fig. 4.2),

$$\begin{aligned} \hat{\theta}_c &= \left(\hat{\theta}_1 + \hat{\theta}_2 + \hat{\theta}_s\right)/3 \\ \hat{\phi}_{\text{sym}} &= \left(\hat{\theta}_1 - \hat{\theta}_2\right)/2 \\ \hat{\phi}_{\text{asym}} &= \left(\hat{\theta}_1 + \hat{\theta}_2 - 2\hat{\theta}_s\right)/2. \end{aligned} \tag{4.9}$$

As usual, $\hat{\theta}_c$ corresponds to the common mode of the circuit, and is the phase variable conjugate to the charge on the capacitive pads ($\hat{n}_c$). In the absence of a magnetic field, this dynamical degree of freedom fully describes the circuit, which behaves similar to an inductively shunted transmon (IST, [185]). A symmetric flux in both loops displaces $\hat{\phi}_{\text{sym}}$, and forms the bias and the differential drive for the LINC. An anti-symmetric flux on the other hand, displaces $\hat{\phi}_{\text{asym}}$ and forms the bias for the asymmetrically threaded SQUID

(ATS) mode of operating the circuit.

The full Hamiltonian, with no assumptions on symmetry, is given by:

$$\begin{aligned} H_{\text{full}} = 4E_C\hat{n}_c^2 + E_L\frac{(\hat{\theta}_c - 2\hat{\phi}_{\text{asym}}/3)^2}{2} \\ - (E_{J_1} + E_{J_2})\cos\left(\hat{\theta}_c\right)\cos\left(\hat{\phi}_{\text{asym}}/3\right)\cos\left(\hat{\phi}_{\text{sym}}\right) \\ + (E_{J_1} + E_{J_2})\sin\left(\hat{\theta}_c\right)\sin\left(\hat{\phi}_{\text{asym}}/3\right)\cos\left(\hat{\phi}_{\text{sym}}\right) \\ + (E_{J_1} - E_{J_2})\cos\left(\hat{\theta}_c\right)\sin\left(\hat{\phi}_{\text{asym}}/3\right)\sin\left(\hat{\phi}_{\text{sym}}\right) \\ + (E_{J_1} - E_{J_2})\sin\left(\hat{\theta}_c\right)\cos\left(\hat{\phi}_{\text{asym}}/3\right)\sin\left(\hat{\phi}_{\text{sym}}\right) \end{aligned} \tag{4.10}$$

To reduce this to the standard LINC Hamiltonian, we assume that all symmetry constraints in the circuit are satisfied, specifically that the outer loop is a symmetric DDS. We will explore the effects of deviating from this ideal in a later section. This constraint lets us set $\hat{\phi}_{\text{asym}} = 0$, giving:

$$\begin{aligned} (\hat{\theta}_1 + \hat{\theta}_2) &= 2\hat{\theta}_s \\ \implies \hat{\theta}_c &= \hat{\theta}_s, \\ \hat{\theta}_1 &= \hat{\theta}_c + \hat{\phi}_{\text{sym}}, \\ \hat{\theta}_2 &= \hat{\theta}_c - \hat{\phi}_{\text{sym}}. \end{aligned} \tag{4.11}$$

Finally, we must make the same assumption as in the DDS, where the symmetric flux degree of freedom $\hat{\phi}_{\text{sym}}$ is stiff and can be replaced by the classical variable $\phi_d = \langle\hat{\phi}_{\text{sym}}\rangle = \frac{\Phi_{\text{loop}}}{2\phi_0}$, where Φ_{loop} is the total flux threading the outer loop.

To understand the LINC's behavior, we can now simply find the expansion of the LINC

potential (see Fig. 4.1b) using Eq. 3.9 and Eq. 3.10, which we reproduce below:

$$\begin{aligned} U_{\text{LINC}} &= \frac{E_L}{2}\hat{\theta}_c^2 - 2E_J \cos\phi_d \cos\hat{\theta}_c \\ &= \sum_{\substack{m\in\text{even},\\ n}} \frac{\partial^n}{\partial\phi_d^n}\left(\frac{\partial^m U_{\text{LINC}}}{\partial\hat{\theta}_c^m}\bigg|_{\hat{\theta}_c=0}\frac{\hat{\theta}_c^m}{m!}\right)\bigg|_{\phi_{\text{DC}}}\frac{\phi_{AC}^n}{n!} \\ &= \sum_{\substack{m\in\text{even},\\ n}} g_{mn}[\phi_{\text{DC}}]\ \phi_{AC}^n\ \left(\hat{c}+\hat{c}^\dagger\right)^m. \end{aligned} \tag{4.12}$$

Importantly, the expansion above explicitly only contains the even terms in m, highlighting that the inherent parity-protection rule holds at arbitrary operating points. At $\phi_{\text{DC}} = \pi/2$, an additional protection is enforced which sets part of the LINC potential to zero, which makes the un-driven Hamiltonian linear:

$$g_{m>2,\,0}\big|_{\phi_{\text{DC}}=\pi/2} = 0. \tag{4.13}$$

Finally, consider the spread of the LINC's phase fluctuations, which are given by:

$$\begin{aligned} \theta_{zpf} &= \left(\frac{2E_C}{\partial^2 U_{\text{LINC}}/\partial\hat{\theta}_c^2\big|_{\hat{\theta}_c=0}}\right)^{1/4} \\ &= \left(\frac{2E_C}{E_L + 2E_J\cos\phi_d}\right)^{1/4}. \end{aligned} \tag{4.14}$$

Since $\theta_{zpf}(\phi_d)$ is itself a non-trivial function of the LINC flux point ϕ_d, modulating flux modulates not just the LINC potential, but also the effective spread of the wavefunction. While the effect of this modulation was negligible in the DDS, it results in an important re-normalization of the LINC physics, for example resulting in a non-zero driven coupler shift at $\phi_{\text{DC}} = \pi/2$.

We can now explicitly study the LINC's idle and driven behavior, both through full Floquet simulations and through the analytic formulation above. We will see that these

match surprisingly well.

4.1.1 Idle linearity and decoherence

While the well-controlled mixing properties of the LINC make it advantageous for a number of applications, we want to focus on the performance of the LINC as a mixer for high-Q bosonic quantum control. In this context, the LINC must activate a parametric coupling to one or more resonators (bosonic modes) when driven, and minimally spoil their quantum information when static. To reiterate our constraints, we desire an ideal bosonic coupler that neither limit the modes' natural decay rate (κ_{res}), nor introduces any additional nonlinearity or dephasing to the oscillators. Suppressing the decay inherited from the coupler limits the allowable energy participation [76] of the resonator in the LINC mode, which we assume for the rest of this section to be $p_{\text{res}} = \beta_{\text{res}}^2 = 0.01$. Minimizing the resonator nonlinearity and dephasing bounds the acceptable static nonlinearity of the coupler (through inherited Kerr and cross-Kerr respectively), placing bounds on the number of junctions in the center shunt (M above). However, suppressing the resonator dephasing will come with a caveat – as discussed previously, a three-wave mixing element is always linearly sensitive to any noise in the driven parameter. Specifically, since a three-wave mixer allows processes to be excited by a single drive photon (coming in through the flux port), a single noise photon at the appropriate frequency can also incoherently activate or dephase the desired process. This means that while the LINC will provide a linear environment, and have no propagation of thermal or shot-noise induced dephasing, it will incur significant dephasing due to noise in the flux, which it could propagate to the bosonic modes. However, this flux noise is explicitly low-frequency, and allows a large range of coherent control based techniques for its suppression. We analyze the effect of both the LINC's linearity and propagated flux noise below.

As motivated earlier, the ideal LINC's Kerr (α_L) goes to zero at the operating point, as

shown in Fig. 4.3a. The curious shape of this curve can be explained entirely analytically by considering the flux dependence of the fourth-order nonlinearity:

$$\begin{aligned}\alpha_L\left[\phi_{\text{DC}}\right] &= 12\ g_{4,0}[\phi_{\text{DC}}] \\ &= -2E_J\cos\phi_{\text{DC}}\ \left(\theta_{zpf}[\phi_{\text{DC}}]\right)^4 \\ &= -\frac{2E_J\cos\phi_{\text{DC}}}{E_L+2E_J\cos\phi_{\text{DC}}}E_C.\end{aligned} \tag{4.15}$$

Crucially the un-driven LINC has no third-order nonlinearity at any bias point (ϕ_{DC}), and thus has no perturbative corrections to the coupler's idle Kerr. This means that the LINC's self-Kerr and cross-Kerr to coupled modes always simultaneously go to zero at the same operating point. The LINC's frequency, including Lamb shift corrections, can also be simply calculated through:

$$\begin{aligned}\omega_L\left[\phi_{\text{DC}}\right] &= 4\ g_{2,0}[\phi_{\text{DC}}] + \alpha_L[\phi_{\text{DC}}] \\ &= \sqrt{8E_C(E_L+2E_J\cos\phi_{\text{DC}})} + \alpha_L[\phi_{\text{DC}}].\end{aligned} \tag{4.16}$$

We overlay these results with exact numerical diagonalization of a 10-junction LINC in Fig. 4.3a. The analytics describe static physics near the operating point exceptionally well, and only deviate from numerics around $\Phi_{\text{DC}} = \Phi_0$, where the truncated Taylor expansion does not accurately represent the coupler's nonlinearity.

If instead of an ideal inductor, the LINC's shunt is composed of an array of N Josephson Junctions, its static Hamiltonian at the operating point is given by:

$$\begin{aligned}H_{\text{LINC}}\left[\frac{\pi}{2}\right] &= 4E_C\hat{n}_c^2 - NE_{L,J}\cos\frac{\hat{\theta}}{N} \\ \alpha_L\big|_{\frac{\pi}{2}} &= E_C/N^2,\end{aligned} \tag{4.17}$$

where each junction in the shunting array has been scaled appropriately to preserve the

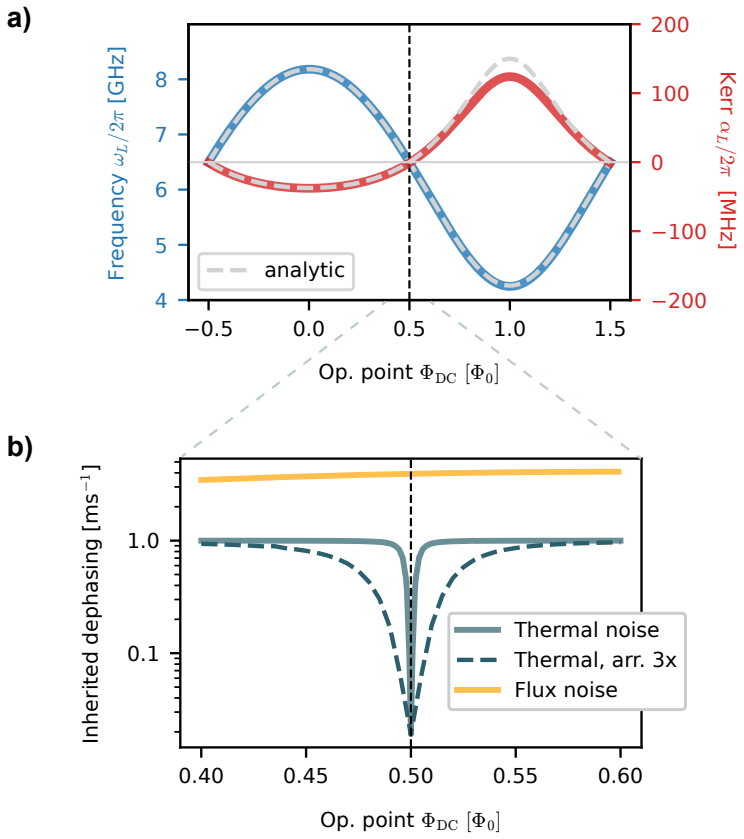

Figure 4.3: **The LINC as a function of operating point a**, The static LINC frequency and Kerr as a function of DC flux. Numerical values (solid lines) are calculated by exact diagonalization of a LINC Hamiltonian with $E_L/h = 52.8$ GHz, $E_J/h = 15.84$ GHz, $E_C/h = 100$ MHz, with the center shunt composed of an array of 10 junctions, where h is Planck's constant. Overlayed analytic curves follow Eqs. 4.15 and 4.16. At the operating point of $\phi_{DC} = \pi/2$, the coupler is linear. **b**, Inherited dephasing for a coupled quantum memory, as a function of DC flux. Thermal noise-induced dephasing is calculated for a coupler thermal population of 2% and T_1 of 20 μs, and is minimized at the operating point, where the coupler is linear. The low-frequency dephasing due to inherited flux noise is calculated through Eq. 4.19, with noise amplitude $A_\Phi = 1\mu\Phi_0/\sqrt{Hz}$. The former is flat outside a narrow region around half-flux, which is made easier to target in an array of three LINCs. The latter dominates memory coherence, but may be suppressed through dynamical decoupling.

same total inductance ($E_{L,J} = NE_L$). To keep the effect of any inherited resonator non-linearity smaller than the resonator's linewidth, the number of junctions in the LINC shunt must be greater than $(p_{res}^2 \bar{n} E_C/\kappa)^{\frac{1}{2}}$, for an intended resonator population of $\bar{n}$. However, note that the LINC behaves as a protected three-wave mixer even if it is not this linear,

including when its shunting inductor is just a single Josephson Junction [179].

As discussed previously, a major source of idle dephasing in any LINC architecture will be low-frequency flux noise. However, since the LINC itself remains in the ground state, this dephasing is primarily harmful if it propagates errors into the desired parametric process, or to the modes that are statically hybridized with the LINC. The effect of flux noise on the strength of the parametric process can be estimated by the simplified infidelity limit:

$$\begin{aligned} 1 - \mathcal{F}_p &\leq \left(\frac{2\pi}{g_{3\text{wm}}} \frac{dg_{3\text{wm}}}{d\Phi} \right)^2 \Bigg|_{\frac{\pi}{2}} \int_0^\infty S_{\Phi\Phi}[f]\ g_N[f\tau_\text{g}]\ df \\ &\approx \left(\frac{2E_J}{E_L} \right)^2 \int_{1/\tau_\text{exp}}^{1/\tau_\text{g}} \frac{1}{\Phi_0^2}\ S_{\Phi\Phi}[f]\ df. \end{aligned} \tag{4.18}$$

Here $S_{\Phi\Phi}[f]$ is the spectral density of flux noise, $g_N[f\tau_\text{g}]$ is a characteristic function defined by the pulse sequence, and the flux-noise sensitivity is analytically derived later in the next sub-section. For an upper bound on the infidelity, we approximate g_N as a rectangular window between the relevant timescales of a single gate ($\tau_\text{g} \sim 1$ ns) and the total experiment ($\tau_\text{exp} \sim 1$ s). Note that the quadratic dependence on flux-noise sensitivity assumes that the variance in the strength of $g_{3\text{wm}}$ is slow and small, and therefore appears as a coherent offset to the intended pulse evolution. For $1/f$ type noise $S_{\Phi\Phi}[f] \sim A_\Phi^2/f$, typical values of the noise amplitude $A_\Phi \sim 1\ \mu\Phi_0/\sqrt{Hz}$ at 1 Hz [157, 186, 187] lead to an estimate of $1 - \mathcal{F}_p \sim 10^{-10}$, which means this mechanism should not be a limiting factor in the mixer's performance.

A more relevant effect might be the inherited flux noise in coupled information-storing modes, which we estimate analytically in Fig. 4.3b, and compare to thermal-noise induced

dephasing around the operating point. This inherited dephasing scales $\propto p_{\text{res}}$ [188, 189]:

$$\begin{aligned}\kappa_\varphi &\propto \sqrt{\left(d\omega_{\text{res}}/d\Phi\right)^2 \int_0^\infty S_{\Phi\Phi}[f]\; g_N[f\tau]\; df} \\ &\approx p_{\text{res}} \left|\frac{d\omega_L}{d\Phi}\right|_{\phi_{\text{DC}}=\pi/2} A_\Phi\; C,\end{aligned} \tag{4.19}$$

where $C = \sqrt{2|\ln\left(2\pi\tau/\tau_{\text{exp}}\right)|} \sim 3-5$ is a slow time dependence on when κ_ϕ is evaluated (τ), and the total length of the experiment (τ_{exp}). The inherited dephasing can therefore be significant, but it is low-frequency, which means it could be mitigated with techniques like dynamical decoupling [157, 186, 187], or stabilized bosonic codes [190, 176]. This makes it potentially still beneficial to operate near the anti-sweet spot, where the coupled resonator is sensitive to inherited flux noise but not to thermal noise-induced dephasing, since the latter requires non-trivial strategies for suppression [191, 192, 193].

4.1.2 Driven performance

We now focus on the driven behavior of the LINC and its performance as a balanced quantum mixer. With ideal symmetry, the LINC's driven behavior is independent of its center shunt (up to a normalization of its impedance), and is given by:

$$H_{\text{driven}} = -2E_J \underbrace{\sin\phi_{\text{AC}}}_{\text{odd}}\; \underbrace{\cos_{\text{NL}}\hat{\theta}}_{\text{even}}. \tag{4.20}$$

There are two important points to note about this driven Hamiltonian. The first is that the drive, ϕ_{AC}, acts on an entirely orthogonal degree of freedom to the LINC mode. This means that the drive does not displace the mode, and the LINC in general remains in its undriven ground state, similar to the DDS [147]. Second, the order of the allowed parametric processes obeys a strict selection rule – the only processes allowed are of the type $\phi_{\text{AC}}^k \hat{\theta}^l$, where l is strictly even. This is a very similar 'parity protection' to that in

the DDS. Interestingly, the LINC has this parity protection at arbitrary operating points, even though it is only linear at $\phi_{\text{DC}} = \pi/2$ – at zero flux, it precisely becomes the DDS, with a slightly offset energy! However, unlike the DDS, there is only a weak parity protection in the number of drive photons k, despite appearing as an odd function. This is primarily because modulating the drive parameter ϕ_{AC} can also modulate the spread of the wave-function θ_{zpf}, causing non-trivial corrections to the strength of parametric processes. Additionally, odd-order processes can combine at higher orders in perturbation theory to form even-order processes. The simplest non-trivial effect where this is observable will be in the driven frequency shift of the LINC, which is suppressed but non-zero at the operating point. However, since the LINC is a 'true' parametric coupler, all such effects are well-predicted by measuring static properties of the LINC as a function of the parameter (ϕ_{DC}), and computing appropriate derivatives.

Let us start by studying the three-wave mixing strength of the LINC, both analytically and through Floquet simulations. Utilizing our usual notation for the bi-variate Taylor expansion, we know that this three wave mixing must contain two derivatives with respect to $\hat{\theta}_c$, and one with respect to ϕ_d, and hence be $\propto g_{2,1}[\phi_{\text{DC}}]$. This three-wave mixing strength is then given by:

$$\begin{aligned} g_{3\text{wm}}[\phi_{\text{DC}}] &= 2\, g_{21}[\phi_{\text{DC}}] \\ &= \frac{E_J}{E_L}\sqrt{\frac{2E_LE_C}{1+2\frac{E_J}{E_L}\cos\phi_{\text{DC}}}}\sin\phi_{\text{DC}} \\ \Rightarrow g_{3\text{wm}}\left[\frac{\pi}{2}\right] &= \frac{E_J}{2E_L}\,\omega_L. \end{aligned} \tag{4.21}$$

Importantly, note that because the LINC's frequency is given by $\omega_L \approx 4g_{2,0}$, the three-wave mixing strength is mathematically equal to the derivative of its frequency (up to a scaling factor), given by:

$$g_{3\text{wm}}[\phi_{\text{DC}}] = \frac{1}{2}\frac{d\omega_L}{d\phi_{\text{DC}}} \tag{4.22}$$

. The latter derivative is exactly the sensitivity of frequency to flux noise, and hence gives the flux-noise dephasing (see Eq. 4.19), reiterating the fluctuation-dissipation trade-off. As an example, intuitively, a beamsplitter interaction is essentially a frequency modulation of the mixer that two modes participate in, so this relation to noises sensitivity makes sense. In fact, experimentally, one can even directly use the relation in Eq. 4.22 to calculate the expected beamsplitting strength once they have extracted each of the two mode's frequencies as a function of the parameter ϕ. As an example, if the coupler is coupled to two modes Alice (a) and Bob (b), we have:

$$g_{\text{BS}} = \sqrt{\frac{\partial \omega_a}{\partial \phi_{\text{DC}}} \frac{\partial \omega_b}{\partial \phi_{\text{DC}}}} \left| \frac{\phi_{\text{AC}}}{2} \right|, \tag{4.23}$$

where the modes' energy participations are intrinsically accounted for through their frequency shifts. This expression has previously been utilized for other flux-driven couplers [146, 194], and in fact holds for *any* parametric 3WM beamsplitter where the parameter is a true orthogonal degree of freedom. Finally, to complete our earlier analysis of the effects of flux-noise, we can also derive how sensitive the three-wave mixing strength itself is to flux:

$$\begin{aligned} \frac{1}{g_{\text{3wm}}} \frac{dg_{\text{3wm}}}{d\phi}\bigg|_{\phi_{\text{DC}}} &= \frac{1 + \cos^2 \phi_{\text{DC}} + \frac{E_L}{E_J} \cos \phi_{\text{DC}}}{\sin \phi_{\text{DC}} \left(\frac{E_L}{E_J} + 2 \cos \phi_{\text{DC}} \right)} \\ \implies \frac{2\pi}{g_{\text{3wm}}} \frac{dg_{\text{3wm}}}{d\Phi}\bigg|_{\Phi_0/2} &\approx \frac{1}{\Phi_0} \frac{2E_J}{E_L}, \end{aligned} \tag{4.24}$$

which then provides Eq. 4.18.

Let's zoom out again – through some careful Hamiltonian analysis, we have derived the LINC's three-wave mixing strength as a function of flux. From an operational standpoint,

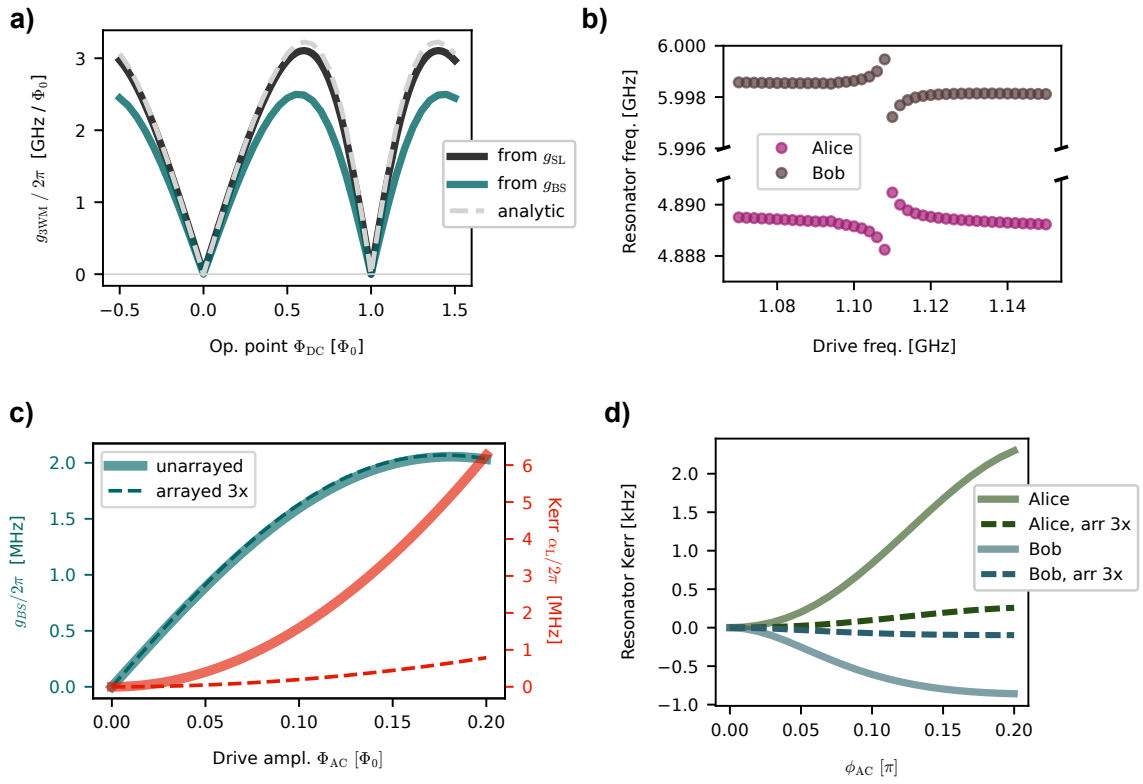

Figure 4.4: **The LINC as a beamsplitter a)**, LINC three-wave mixing strength as a function of the operating point, comparing analytic formula (grey dashed, Eq. 4.21) to exact time-domain and Floquet results from a squeezing operation within the LINC ($2g_{\mathrm{SL}}$, black) and a beamsplitter operation between two external resonators (g_{BS}, teal) respectively. The discrepancy in the beamsplitting prediction may be due to driven changes in the effective resonator-LINC participations, which are captured in the Floquet simulation but not in the analytics. **b)**, LINC beamsplitting strength (g_{BS}) and Kerr (α_L) as a function of drive strength (ϕ_{AC}), from exact Floquet simulation. While the LINC is Kerr-free when idle, higher-order nonlinearities can induce a driven Kerr. This can be suppressed by arraying multiple LINCs (arrayed 3x, dashed). **c)**, Driven avoided crossing due to the beamsplitter interaction between two storage modes, Alice (4.9 GHz) and Bob (6.0 GHz), as a function of drive detuning from beamsplitter resonance, at a fixed drive amplitude of $\phi_{AC} = 0.2\pi$. **c)**, Driven Kerr of each mode for an unarrayed (solid lines) vs arrayed (dashed lines) LINC. Arraying 3 LINCs reduces the inherited driven Kerr by a factor of ~ 9.

we can just look at overall third-order part of the LINC Hamiltonian, at the half-flux point:

$$\begin{aligned} H^{(3)}_{\mathrm{driven}}/\hbar &= \left(\frac{E_J}{2E_L}\right)\omega_L\ J_1\left(|\phi_{\mathrm{AC}}|\right)\cos\omega_d t\left(\hat{c}+\hat{c}^\dagger\right)^2 \\ &\approx g_{3\mathrm{wm}}|_{\frac{\pi}{2}}\ \phi_{\mathrm{AC}}(t)\left(\hat{c}+\hat{c}^\dagger\right)^2, \end{aligned} \tag{4.25}$$

where $\omega_L \approx \sqrt{8E_L E_C}$ is the undriven frequency of the LINC mode, $J_1(x)$ is the first-order Bessel function, and $\phi_{\text{AC}}(t) = |\phi_{\text{AC}}| \cos \omega_d t$ is the drive. These Bessel functions often appear in the strengths of parametric processes, and can be derived simply from the Jacobi-Anger expansion [195]. For two modes Alice ($\hat{a}$, ω_a) and Bob ($\hat{b}$, ω_b) that are coupled to the LINC (with energy participations p_a and p_b), driving at select frequencies can activate various desired processes, such as:

$$\begin{aligned} \omega_d = \omega_a - \omega_b &\implies g_{\text{BS}} \left(\hat{a}^\dagger \hat{b} + \hat{a} \hat{b}^\dagger \right) \\ \omega_d = \omega_a + \omega_b &\implies g_{\text{TS}} \left(\hat{a}^\dagger \hat{b}^\dagger + \hat{a} \hat{b} \right) \\ \omega_d = 2\omega_a &\implies g_{\text{Sa}} \left(\hat{a}^{\dagger 2} + \hat{a}^2 \right) \\ g_{\text{BS}} = g_{\text{TS}} = 2g_{\text{Sa}} &\approx g_{\text{3wm}} |\phi_{\text{AC}}| \sqrt{p_a p_b}, \end{aligned} \tag{4.26}$$

with $p_a = p_b$ for the single-mode squeezing process.

To verify our analysis and check the true behavior of the LINC-Alice-Bob system, we can directly analyze some of these processes through Floquet simulations, at arbitrary flux points. First, we can simulate just the coupler activating a self-squeezing process, by driving around $2\omega_L[\phi_{\text{DC}}]$. Since squeezing a perfectly harmonic oscillator (like the LINC close to half flux) spreads its wavefunctions out among all even Fock states, it is a little difficult to analyze such a process by tracking individual quasi-energies. Instead, we can simply add some artifical Kerr to the system (that does not interact with the drive), and then analyze the Rabi-like oscillations between the LINC's $|g\rangle$ and $|f\rangle$ states to extract $g_{\text{3wm}} = 2g_{SL}$. These match well with the analytic $g_{\text{3wm}}[\phi_{\text{DC}}]$ at arbitrary flux-points, as shown in Fig. 4.4a. We can then simulate a full beamsplitting process between Alice and Bob in the full 3-mode system. At any flux point and drive-amplitude, this gives an avoided crossing between two oscillators, as shown in Fig. 4.4b, from which g_{BS} can be numerically extracted. Plotting this as a function of flux provides a slightly scaled down

version of the analytic curve (Fig. 4.4a), with the mismatch likely being due to changes in the effective participation ratios (p_a, p_b) when driven. We can also numerically plot this at arbitrarily strong drive amplitudes and observe a turn-around in the beamsplitting strength due to higher-order contributions and the Bessel-function behavior. This means that there exists an optimal $|\phi_{\text{AC}}|$ after which $|g_{\text{BS}}|$ saturates, and we would like to keep our driven system 'clean' up to that drive amplitude. Finally, for bosonic control, the key improvement we were looking for over the DDS was linearity. This linearity is trivially guaranteed when undriven, but the presence of drives can introduce Kerr through higher-order effects. However, since these are indeed higher-order, we can suppress them by simply using the arraying trick in Eq. 3.47, where we have multiple LINC loops threaded by the same flux, shunted by the same capacitor. We plot the Kerr extracted for the single and arrayed LINCs, for both just the coupler and the coupled oscillators, in Fig. 4.4c,d. The latter inherited Kerr is difficult to predict analytically, as it depends on interactions through multiple sidebands of the coupler, which Floquet simulations capture exactly.

The (inherited) Kerr was only one drawback of using the DDS – another major non-ideality was its significant frequency shifts, which at small drive strengths were described by a $|\phi_d|^2 c^\dagger c$ process. If one naively considers the driven potential energy in Eq. 4.20, one might conclude that there is no such term, given that the sinusoidal modulation is an odd function that has no average driven behavior (in the absence of resonances). However, a quick Floquet simulation reveals this to be not true – there is a considerable frequency shift ($\geq g_{\text{BS}}$ even at small drive amplitudes (Fig. 4.5a). Should this be surprising? Consider that the average frequency shift of any sinusoidally driven mode is always given to

second order by:

$$\begin{aligned}\Delta\omega_L &\approx \left\langle \left.\frac{\partial\omega_L}{\partial\phi_d}\right|_{\phi_{\text{DC}}} \phi_{\text{AC}} \;+\; \frac{1}{2}\left.\frac{\partial^2\omega_L}{\partial\phi_d^2}\right|_{\phi_{\text{DC}}} \phi_{\text{AC}}^2 \right\rangle \\ &= \frac{1}{4}\left.\frac{\partial^2\omega_L}{\partial\phi_d^2}\right|_{\phi_{\text{DC}}} |\phi_{\text{AC}}|^2.\end{aligned} \tag{4.27}$$

Since the curvature of the frequency vs flux is not zero at the half flux point, neither should its Zeeman shift. In fact, this shift is exactly proportional to the flux-noise sensitivity of $g_{3\text{wm}}$ that we derived in Eq. 4.24. We can derive this explicitly by considering our usual Taylor analysis, which takes into account the non-zero flux modulations of θ_{zpf}:

$$\begin{aligned}\Delta\omega_L[\phi_{\text{DC}}, \phi_{AC}] &= 2\; g_{2,2}[\phi_{\text{DC}}]\; |\phi_{AC}|^2 \\ \hbar g_{2,2} &= -\frac{E_J}{4}\sqrt{\frac{2E_C}{E_L + 2E_J\cos\phi_{\text{DC}}}}\times \\ &\quad \frac{E_J(1+\cos^2\phi_{\text{DC}}) + E_L\cos\phi_{\text{DC}}}{E_L + 2E_J\cos\phi_{\text{DC}}} \\ \implies \Delta\omega_L[\pi/2, \phi_{\text{AC}}] &= -\left(\frac{E_J}{2E_L}\right)^2 \omega_L\; |\phi_{\text{AC}}|^2.\end{aligned} \tag{4.28}$$

This expression matches the Zeeman shift extracted from Floquet at nearly all flux points (Fig. 4.5b). Importantly, this driven shift is completely independent of the device's (driven) anharmonicity – it does *not* get suppressed by arraying the LINC, and for a large enough array simply approaches the ideal three-wave mixer's shift $\Delta\omega_{\text{ideal}}$ (Eq. 4.3). However, if we cannot calibrate it out or the experiment requires small Zeeman shifts, one can even operate at the inflection point $d^2\omega/d\phi^2\,[\phi_{\text{DC}}] = 0$, which for the perfectly symmetric LINC (with $E_J/E_L = 0.3$) occurs at $\phi_{\text{DC}} \sim 0.594\pi$ (dotted line in Fig. 4.5a).

The above analysis clearly shows that the (arrayed) LINC can perform three-wave mixing beamsplitting, with the coupled oscillators remaining reasonably linear even when

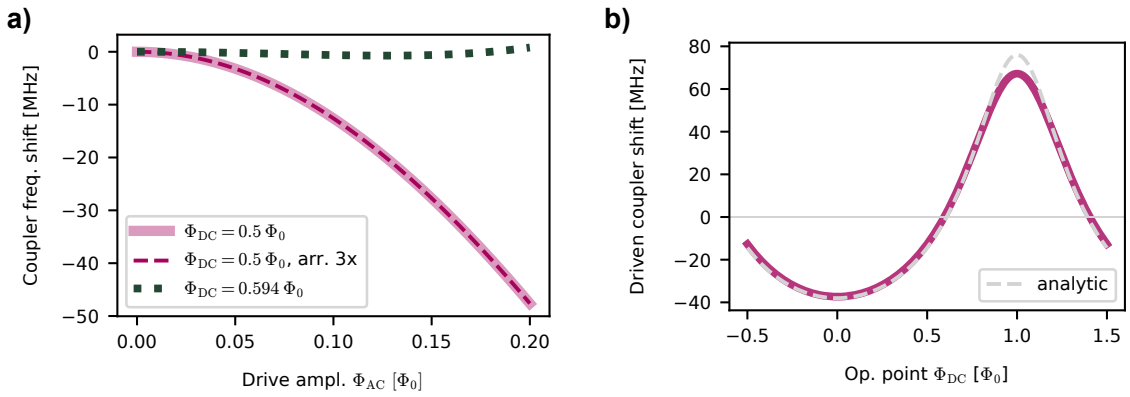

Figure 4.5: **Driven frequency shifts in the LINC a)**, Driven LINC frequency shift (solid purple) vs drive strength. Unlike in charge-driven mixers like the SNAIL, this frequency shift is independent of the coupler Kerr, and persists even if the coupler is arrayed (dashed purple). It is possible to minimize this shift by biasing to a flux point where the coupler frequency has an inflection point as a function of DC flux (dotted black). **b)**, Comparison of the analytic formulae (grey dashed, Eq. 4.28) to exact Floquet simulation for driven coupler Zeeman shift as a function of DC flux, for fixed drive amplitude $\phi_{AC} = 0.1\pi$.

driven. However, this only takes care of the leading order coherent nonidealities due to the drive – how well is the LINC protected against general parasitic processes? Remember that the LINC is parity-protected in a very similar way to the DDS. In fact, both these circuits can be exactly cast as a quantum analogues of classical mixer balancing, where two identical nonlinear elements are driven with equal and opposite drives to coherently cancel half the driven processes (Fig. 4.6a). For the DDS, these nonlinear elements were simply single junctions, but for the LINC these effective nonlinear elements are instead RF-SQUIDs that are explicitly biased to $\phi_{\text{DC}} = \pi/2$. To demonstrate how its parity protection helps suppress parasitic processes, we can examine the driven behavior of the LINC in the presence of one or more drive tones through Floquet simulations, and as a point of comparison, simulate an equivalent Kerr-free SNAIL. As a reminder, we desire both couplers to remain in their driven ground state even at strong drive amplitudes. For a realistic simulation, we can incorporate an environment that induces both coupler decay and flux-noise dephasing, which may activate into nonlinear coupler heating in the presence of the drive (see previ-

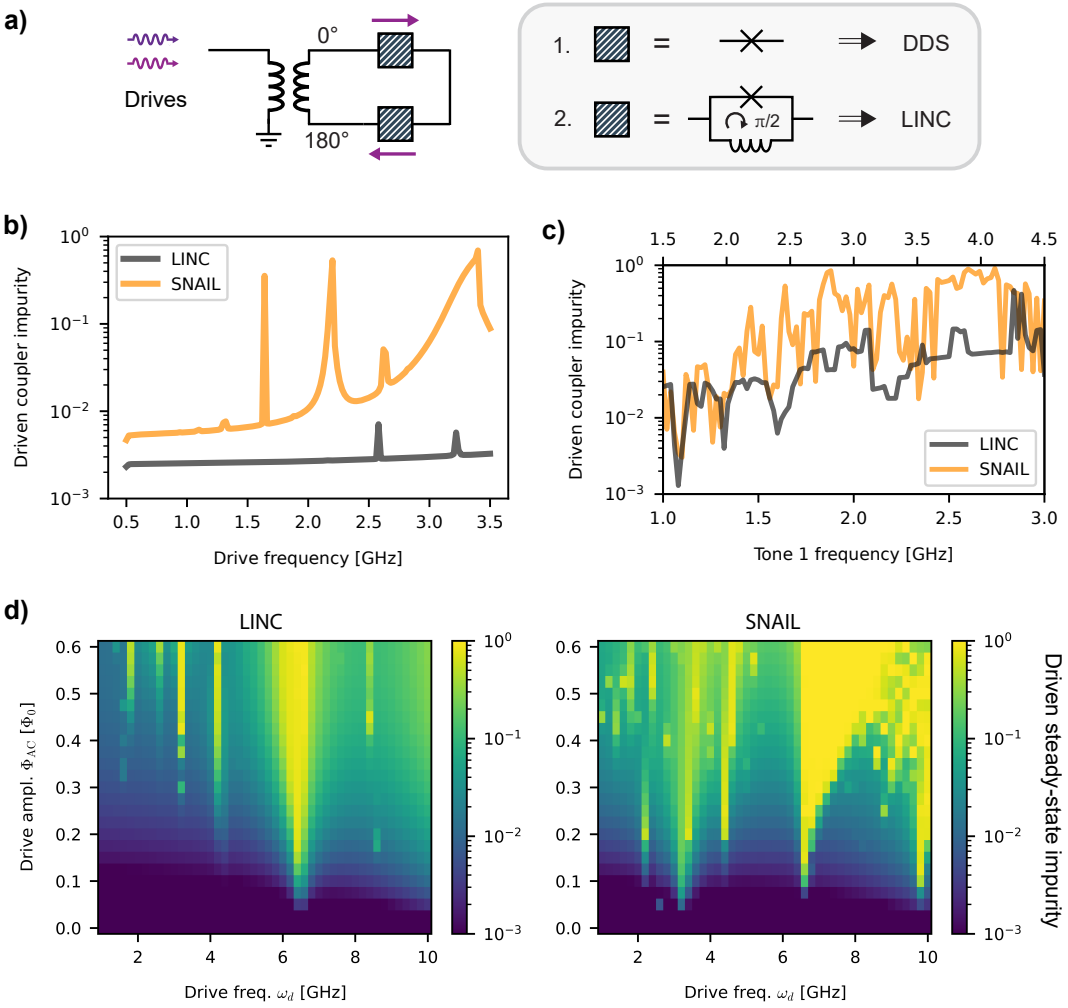

Figure 4.6: **Parasitic resonances in the LINC and SNAIL a)**, The LINC can be broadly connected to the DDS as another instance of a balanced quantum mixer, specifically acting as a balanced RF-SQUID. **b)**, Comparison of the LINC and SNAIL driven impurities at ($|\phi_{\text{AC}}| = 0.2\pi$) as a function of drive frequency, on a log scale. The simulation includes equal low-frequency flux noise ($1\mu\Phi_0/\sqrt{Hz}$), decay ($T_1 \sim 25\mu$s), and coupling to an external hot TLS (at 4.9 GHz) for both couplers. **c)**, Comparison of the LINC and SNAIL driven impurities as a function of simultaneously scaled drive frequencies of two tones (each with $\phi_{AC} = 0.1\pi$), on a log scale. The ratio of the two drive tones is maintained at $2:3$. **d)**, Full two-dimensional sweep of the driven impurity in both the LINC and SNAIL, as a function of drive amplitude and frequency, in the presence of flux noise and decay. The advantages of the protections available in the LINC are clear from the larger available drive space.

ous section, and [141, 196]). We also include a coupled, information-storing qubit mode, which must ideally remain unaffected by the coupler during any parametric process. As-

suming the coupler is not periodically reset, this coupler-qubit system will reach a driven steady state after a sufficient number of operations, and we compute the impurity of this driven steady-state as our metric for performance (Eq. 3.37).

Specifically, in Fig. 4.6b, we compute this driven purity for a 6.5 GHz LINC through Floquet-Markov simulations, with decay $\gamma_1 = (26.7\mu\text{s})^{-1}$, flux-noise dephasing $S_{\Phi\Phi}[\omega] = [1\mu\Phi_0/\sqrt{Hz}]^2/\omega$, and coupling to a qubit (4.9 GHz). Importantly, we want to compare to an *equivalent* SNAIL. We define this equivalent SNAIL as one that has identical frequency, beamsplitting strength, and decoherence rates, when operated at its Kerr-free point (and expect any advantages to hold at arbitrary flux points). We first fix drive amplitude near the beamsplitter saturation point, $\phi_{AC} = 0.2\pi$, and compare their purity o a log scale in the presence of a single drive tone. The difference is quite stunning – the LINC remains orders of magnitude more pure at *all* drive frequencies. We then subject both couplers to two simultaneous drive tones, scaling their drive frequencies together to simplify the effective periodicity of the Hamiltonian for Floquet simulations (Fig. 4.6c). Here, the LINC is still more pure at all drive frequencies - in fact the SNAIL is virtually un-operable (infidelity nearly always $> 10\%$) above $\omega_1/2\pi = 1.5$ GHz. This means the SNAIL sees substantial spurious transitions over most of the frequency range, but a large fraction of these inter-modulation products are suppressed by parity protection in the LINC. Finally, we can perform a full 2-dimensional sweep of the driven infidelity (Fig. 4.6d). This makes clear the difference in total available drive-space between the two couplers, and the significantly larger number of parasitic processes in the SNAIL, despite being Kerr-free. Importantly, unlike in the DDS, these parasitic processes are mostly straight lines that do not feel the effect of Zeeman shift on this frequency scale, implying that they are much easier to avoid. Overall, these simulations show that in a realistic environment, the LINC should offer important advantages over the SNAIL (and potentially the DDS) in high-Q and multi-tone applications.

4.2 Asymmetries and parasitic inductances

Our analysis of the LINC so far has been restricted to a perfectly symmetric circuit under a purely differential DC flux and drive. These have led to important advantages, like its perfect linearity at the half-flux point, and its parity protected performance when driven. However, practical implementations of the coupler will inevitably come with experimental imperfections. In particular, we must study how the LINC's static and driven performance changes in the presence of finite asymmetries and parasitic inductances, which are known to provide nontrivial effects in other mixing elements [165, 147].

What asymmetries and imperfections might be useful to consider? The first obvious asymmetry is that of the junctions in the outer loop, which we define by the ratio $\beta_\Delta := (E_{J1} - E_{J2})/E_L$ for junction energies $E_{J1,J2}$ respectively. From usual fabrication imperfections, we expect these to be $\lesssim 2\%$. We will also find it useful to correspondingly define $\beta_\Sigma := (E_{J1} + E_{J2})/E_L$. The second imperfection is a finite difference in the DC flux in the two LINC sub-loops, arising from a gradient in the residual magnetic field in which the experimental package cooled down. This imperfection is in principle possible to cancel in-situ with two dedicated flux lines per coupler, such that the relative currents in the flux lines can be tuned to achieve arbitrary flux biases in the two sub-loops. However, the LINC ideally only requires a single fast flux line for its drive and DC flux, which always applies symmetric flux to both sub-loops. In this scenario, any residual DC flux difference ϕ_Δ will change both the LINC's driven and undriven behavior and be uncorrectable, and so we will analyze its effect below. In the fully general case, one may also have an asymmetry in the applied drive due to imperfect drive line engineering, which might also affect performance. However, this asymmetry can be largely avoided by appropriately designing the flux distribution and capacitance matrix of the device, aided by numerical simulation techniques [147, 150]. Finally, the LINC circuit will generally have a

distributed geometry that has both parasitic capacitances and parasitic inductances. While the large capacitance between the pads of the coupler will dominate any capacitive effects, the parasitic inductances may have non-negligible corrections on the LINC's behavior. We will analyze two such inductances, that in the loop and in series with the loop (ie, in the capacitive 'arms'), below. Throughout this study, we will follow the general framework of writing out the Hamiltonian in the presence of the asymmetries and inductances, find the new potential minima θ_{min}, and utilize the general Taylor expansion (Eqs. 3.9 and 3.10) at this new minima to extract both static and driven behavior.

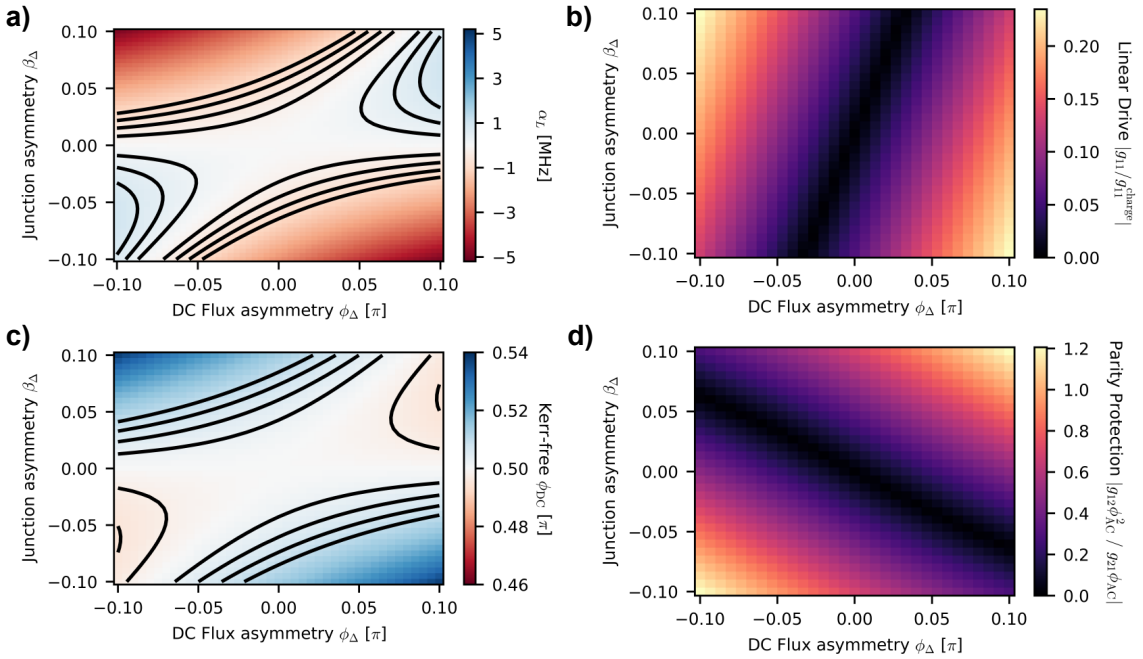

Figure 4.7: **Effect of asymmetries in the LINC a)**, Kerr of the LINC mode at $\phi_{\text{DC}} = \pi/2$, computed by direct diagonalization. Black contour lines correspond to $|\alpha_L| = 250, 500, 750,$ and 1000 KHz respectively. **b)**, Operating point ϕ_{DC} where the LINC is Kerr-free, as a function of asymmetry. Black contour lines correspond to flux values from $98\% \times \pi/2$ to $102\% \times \pi/2$ in 0.5% increments around the asymmetry-free DC operating point. **c)**, strength of g_{11}, the linear coupling of the LINC mode to a symmetric flux drive, as compared to directly charge-driving the LINC ($g_{11}^{\text{charge}} \approx E_L\theta_{zpf}$). **d)**, Parity protection within the 3rd order nonlinearity. Strength of the $g_{12}\phi_{\text{AC}}^2$ process, that would permit subharmonic-driving-like processes, as compared to the desired $g_{21}\phi_{\text{AC}}$, is shown as a function of asymmetry for the drive amplitude $|\phi_{\text{AC}}| = 0.2\pi$.

Let us first analyze the effect of junction and flux asymmetries on the idle LINC. With the definitions above, we can express the static inductive potential of the LINC in the presence of asymmetries as

$$\begin{aligned} U_{\text{LINC}}^{(\text{asym})}/E_L =& \frac{1}{2}(\hat{\theta}_c - 2\phi_\Delta/3)^2 \\ & - \beta_\Sigma \cos(\phi_d)\left[\cos(\phi_\Delta/3)\cos\left(\hat{\theta}_c\right) - \sin(\phi_\Delta/3)\sin\left(\hat{\theta}_c\right)\right] \\ & + \beta_\Delta \sin(\phi_d)\left[\cos(\phi_\Delta/3)\sin\left(\hat{\theta}_c\right) + \sin(\phi_\Delta/3)\cos\left(\hat{\theta}_c\right)\right]. \end{aligned} \tag{4.29}$$

Importantly, these asymmetries result in an effective shift in the potential minima, which can at the half-flux operating point can be found by evaluating $\left.\frac{\partial U_{\text{LINC}}^{(\text{asym})}}{\partial \hat{\theta}_c}\right|_{\theta_{min},\phi_d=\pi/2} = 0$, giving

$$\theta_c^{\text{min}} \approx \frac{2}{3}\phi_\Delta - \beta_\Delta \tag{4.30}$$

to lowest order in β_Δ and ϕ_Δ. This new minima results in a change in the undriven potential, meaning that Taylor expansions to compute parameters of interest should be expanded about this new point. In light of this shift, the static Kerr at the operating point becomes

$$\alpha_L^{\text{asym}} \approx E_c \beta_\Delta \sin\left(\phi_\Delta - (1+\sqrt{2})\beta_\Delta\right) \tag{4.31}$$

assuming the amount of asymmetry is small. Thus any static nonlinearity gained by the LINC in the presence of these asymmetries is only second-order in $\{\beta_\Delta, \phi_\Delta\}$. Notably, the sign of the individual asymmetries is important, as the two effects can either constructively or destructively interfere, and for symmetric junctions, the static Kerr is nulled for any value of ϕ_Δ. Even in the presence of asymmetry, the Kerr always has a zero crossing, and can be nulled by slightly shifting the operating point ϕ_{DC}. We plot the magnitude of Kerr and the shifted Kerr-free point as a function of asymmetry in Fig. 4.7a,b. As a reference, for a $\beta_\Delta = 2\%$ and a $\phi_\Delta = -\frac{\pi}{2} \times 1\%$, the static Kerr of the LINC is roughly -130 KHz for

$E_C = 100$ MHz. The shifted Kerr-free point can then be found by changing the symmetric flux bias by less than 0.005π.

An additional effect due to asymmetries is the presence of a non-zero third order non-linearity in the common mode:

$$\begin{aligned} g_{30} &= \frac{1}{3!}\left[\frac{\partial^3 U}{\partial \theta_c^3}\theta_{zpf}^3\right]\Bigg|_{\theta_c=\theta_c^{\min},\phi_d=\pi/2} \\ &\approx -E_L\beta_\Delta\cos(\phi_\Delta-\beta_\Delta)\theta_{zpf}^3. \end{aligned} \tag{4.32}$$

This g_{30} can change the relative position of the Kerr-free and cross-Kerr free points, which it adds different perturbative corrections to. Specifically, the cross-Kerr is given by $\chi_{\text{res,LINC}} = 24p_{\text{res}}(g_{40} + 6g_{30}^2\omega_L/(\omega_{\text{res}}^2 - 4\omega_L^2))$, while the self-Kerr is given by $\alpha_L = 12(g_4 - 5g_{30}^2/\omega_L)$ [175]. However, again, for symmetric junctions, this nonlinearity is nulled for all values of ϕ_Δ. The effect of this parasitic g_{30} will also always be significantly smaller than in dedicated charge-driven three-wave mixers like the SNAIL.

Now that we have studied the implications of asymmetry on the static behavior of the LINC, we move to examining its driven behavior. To do this, we can look at the driven Hamiltonian at the operating point in the presence of asymmetries

$$\begin{aligned} U/E_L =& \frac{1}{2}(\hat{\theta}_c - 2\phi_\Delta/3)^2 \\ &+\beta_\Sigma\sin(\phi_{AC})\left[\cos(\phi_\Delta/3)\cos\left(\hat{\theta}_c\right) - \sin(\phi_\Delta/3)\sin\left(\hat{\theta}_c\right))\right] \\ &+\beta_\Delta\cos(\phi_{AC})\left[\cos(\phi_\Delta/3)\sin\left(\hat{\theta}_c\right) + \sin(\phi_\Delta/3)\cos\left(\hat{\theta}_c\right))\right]. \end{aligned} \tag{4.33}$$

There are two primary driven processes we might care about at lowest order. The first is a parasitic linear drive that displaces the LINC, given by $H_{lin} = g_{1,1}\ \phi_{\text{AC}}\ (c + c^\dagger)$. Such a displacement would normally come from a charge drive coupled to $\hat{n}_c$, but here it can arise due to unequal driven currents in the two outer branches in the presence of junction

asymmetry. This parasitic drive scales as

$$\begin{aligned} g_{11} &= \theta_{zpf} \frac{\partial^2 U}{\partial \phi_d \partial \theta_c}\bigg|_{\theta_c=\theta_c^{\min},\phi_d=\pi/2} \\ &\approx -E_L \beta_\Sigma \theta_{zpf} \sin(\phi_\Delta - \beta_\Delta), \end{aligned} \tag{4.34}$$

which means it is linearly sensitive to small asymmetries. As a benchmark, we can compare this against the resulting driven Hamiltonian from directly displacing the LINC, which would be $g_{11}^{\text{charge}} \approx E_L \theta_{zpf}$. The ratio of these is plotted in Fig. 4.7c, and is less than 0.1 for most of this range (this ratio is always 1 for the SNAIL). The second parasitic driven term we could consider is one that breaks parity-protection within the third-order nonlinearity, given by the term $g_{12}(\hat{c} + \hat{c}^\dagger)\phi_{\text{AC}}^2$. For context, an unprotected three-wave mixer like the SNAIL would have this driven term only off by a factor of θ_{zpf} from the desired mixing process (g_{21}). For the asymmetric LINC, one can derive:

$$\begin{aligned} g_{12} &= \left[\frac{1}{2}\frac{\partial^3 U}{\partial \phi_d^2 \partial \theta_c}\theta_{zpf} + \frac{\partial^2 U}{\partial \phi_d \partial \theta_c}\frac{\partial \theta_{zpf}}{\partial \phi_d}\right]\bigg|_{\theta_c=\theta_c^{min},\phi_d=\pi/2} \\ &\approx -\frac{E_L \theta_{zpf}}{2}\left[\beta_\Delta + \frac{1}{2}\beta_\Sigma^2 \sin(\phi_\Delta - \beta_\Delta)\right], \end{aligned} \tag{4.35}$$

which is again linearly sensitive to asymmetry. Interestingly, the linear dependence in ϕ_Δ originates from a modulation of θ_{zpf}, similar to the LINC's AC Zeeman shift. We plot the ratio of the undesired and desired process strengths ($g_{12}\phi_{\text{AC}}^2$ vs. $g_{21}\phi_{\text{AC}}$) for a chosen drive strength of $\phi_{\text{AC}} = 0.2\pi$ in Fig. 4.7d.

We now turn our focus to the effect of parasitic inductances in the LINC. In general, the method for analyzing the behavior of such an inductor in series with a general nonlinear element involves looking at the total potential energy across the two elements. This total potential energy is always a function of at least two variables; the phase drop across the inductor, and the phase drop across the nonlinear element. However, because they are in

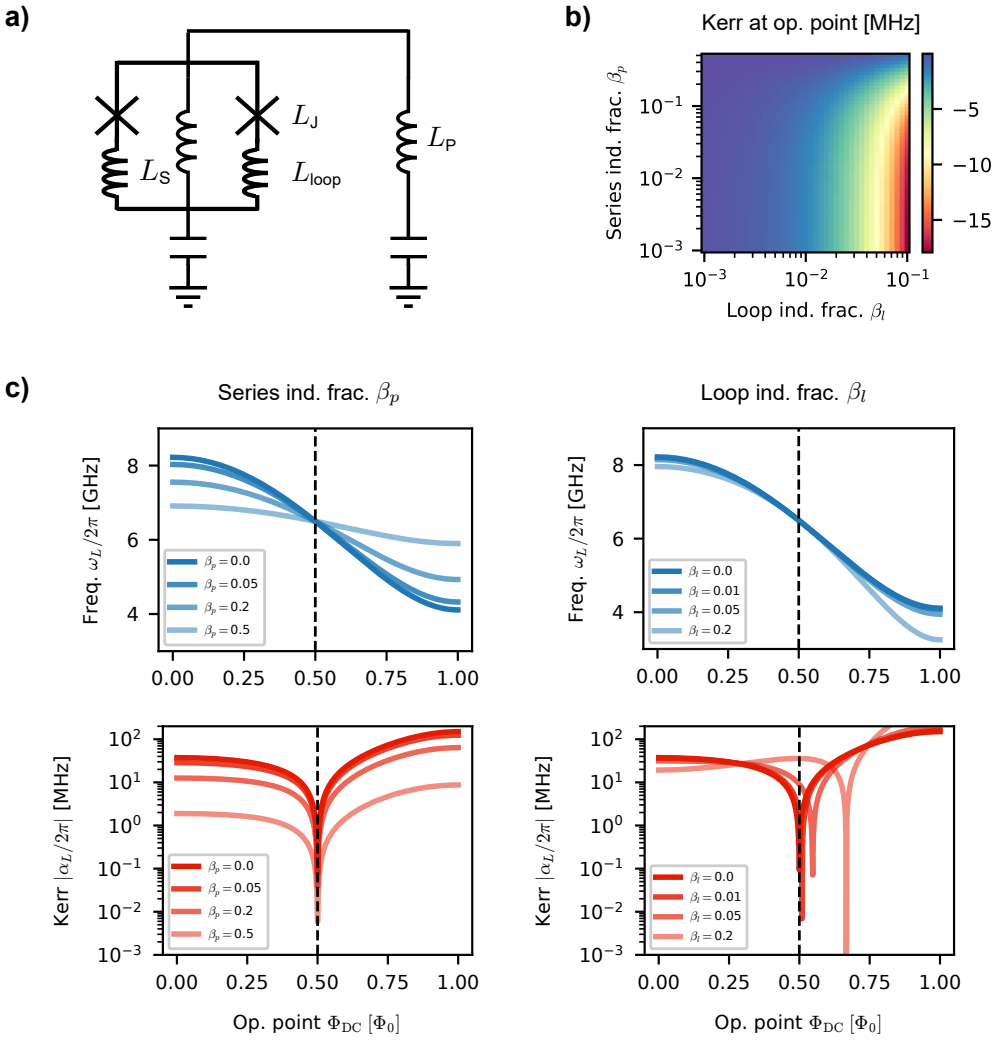

Figure 4.8: **Effect of parasitic inductance in the LINC a)**, Circuit diagram for the LINC with parasitic loop inductance (L_{loop}) and series inductance (L_{P}). We analyze their effects by quantifying their fractional branch participations $\beta_l = L_{\text{loop}}/L_{\text{J}}$ and $\beta_p = L_{\text{P}}/L_{\text{S}}$. **b)**, Kerr at the half-flux operating point as a function of inductance fractions β_l and β_p. The series inductance has a minimal effect, but a large loop inductance can introduce significant Kerr. **c)**, Frequency and Kerr vs flux for several values of series inductance (left) and loop inductance (right) fractions. A series inductor changes the participation of the loop and therefore the frequency tuning range, but does not change the Kerr-free point. A loop inductance has minimal changes to the frequency, but can change the LINC's Kerr. All plots in this figure were calculated using a symbolic engine, following [197].

series, they must satisfy an additional constraint – the current through the inductor must be equal to the current through the element [198, 165, 197]. This constraint reduces the

problem to a single variable, that of the phase drop across the total mixer. Various mixing products can then be evaluated by taking derivatives of the total potential energy with respect to this total phase drop. For a general sinusoidal element, the current conservation condition results in a transcedental equation (eg, $A\sin(\phi_1) + B\phi_2 = 0$) whose solutions are technically analytic, but can only be represented in terms of a common parameter (or solved graphically). This makes the corrections to the nonlinearity of the mixer due to this inductor nontrivial to calculate in general, and one often utilizes symbolic engines for efficient evaluation of their behavior [197].

The LINC in general will have two types of parasitic inductances; those within the loop, in series with the outer junctions (L_{loop}), and those in series with the coupler itself, arising from its arms (L_p). The general circuit diagram in the presence of these inductors is shown in Fig. 4.7a. We will study the effect of these inductances in terms of their unit-less (linear) branch participations $\beta_l = L_{\text{loop}}/L_{\text{J}}$ and $\beta_p = L_{\text{P}}/L_{\text{S}}$. Overall, we will see that the loop inductance will have a more non-trivial effect and shift the Kerr-free point. However, in usual geometries, this loop inductance is small ($\sim 1\text{pH}/\mu\text{m}$) and β_l is generally on the order of 0.01 (with $L_{\text{loop}} \sim 0.1$ nH and $L_J \sim 10$ nH). The series inductance on the other hand may be much larger, even becoming comparable to the shunt inductance ($\beta_p \sim 0.2$) in 3D implementations, where the coupler's capacitive 'arms' are elongated to enhance its dipole moment. Importantly however, the LINC's special symmetry and operating point result in the effect of a series inductor simply being a normalization factor by the linear voltage division between the shunt and parasitic inductor – ie it effectively acts as linear inductor in series with another inductor. This is primarily because the parameter ϕ_d does not drive or interact with the parasitic series inductor. The presence of the inductor normalizes all orders of nonlinearity by constant factors related to β_p, but crucially does not change any nonlinearity that was zero in the ideal LINC to a non-zero value. This means, for example, that the half-flux point always remains Kerr-free, and the effect of

any residual nonlinearity is only diluted by the inductor.

We first analyze the loop inductance in more detail. The total potential energy for the left, center and right branches of the LINC, in the presence of a symmetric flux ϕ_d, can be written as:

$$U(\theta,\phi) = U_{\mathrm{L}}\left(\theta_c - \frac{\phi_d}{2}\right) + U_{\mathrm{C}}(\theta_c) + U_{\mathrm{R}}\left(\theta_c + \frac{\phi}{2}\right), \tag{4.36}$$

where $U_{\mathrm{L,R}}$ are the potential energies of the left and right branch, respectively, while U_{C} is the potential energy of the center branch, which is nominally just the shunting inductor's energy. Fore a symmetric LINC, the functional form of U_L and U_R are the same, and can expressed in the form of a parametric equation as:

$$\begin{cases} U_{L,R} = \dfrac{E_{\text{loop}}}{2}\left(\beta_l \sin\varphi_{J_{L,R}}\right)^2 - E_J \cos\varphi_{J_{L,R}} \\ \phi_{L,R} = \varphi_{J_{L,R}} + \beta_l \sin\varphi_{J_{L,R}} \\ \phi_d = 2\left(\varphi_J + \beta_J \sin\varphi_J\right). \end{cases} \tag{4.37}$$

where E_{loop} is the energy of the stray loop inductance, β_l is the ratio between the stray loop inductance and the Josephson inductance (L_{loop}/L_J), and φ_J is the phase across the Josephson junction. It is then possible to graphically evaluate the exact expressions of the LINC Hamiltonian expansion coefficients through our usual Taylor expansions. We perform this numerically to evaluate effects on the static LINC in Fig. 4.8b, c. We expect the loop inductance of the LINC to be geometrically limited to ~ 1 pH/ μm, which for a 100μm loop side implies $L_{\text{loop}} = 0.2\text{nH} \implies \beta_l \sim 0.02$ (with $L_S \sim 3$nH, $L_J \sim 10$nH). We observe that the static frequency of the LINC as a function of flux is barely affected for in this range of loop inductance. This also implies that the parametric mixing strength $g_{3wm} \sim 0.5\ d\omega_L/d\phi_d$ remains similar. The static Kerr of the LINC is more significantly affected, with $\alpha_L = 1 - 10$ MHz at $\phi_{DC} = \pi/2$, but the Kerr still has a zero-crossing very close to half flux and can, therefore, always be nulled out for sensitive operations.

We now shift to analyzing the effect of a parasitic inductor in series with the LINC. Importantly, the effect of this inductor has an analytic solution, and so it is useful to set up the equations and framework required to understand it. We denote the potential energy of just the pure LINC as U_{LINC}, corresponding to phase $\hat{\theta}_c$. We then consider the phase across the total effective dipole (LINC + parasitic inductor) as θ, with corresponding potential energy U_{tot}. Each nonlinear process activated by the coupler is described through derivatives of this total potential energy, which can be analytically derived using the current conservation relation. In particular, our goal will be to quantify how well the LINC's parity-protection (which includes its idle linearity) is preserved in the presence of a parasitic series inductance, i.e., whether

$$\frac{\partial^{m+n} U_{\text{LINC}}(\theta_c, \phi_d)}{\partial \theta_c^m \, \partial \phi_d^n} = 0 \Rightarrow \frac{\partial^{m+n} U_{\text{tot}}(\theta, \phi_d)}{\partial \theta^m \, \partial \phi^n} \approx 0. \tag{4.38}$$

To evaluate the correspondence in Eq. (4.38), we make use of the current conservation condition $i = i_{\text{LINC}}$ to relate first-order partial derivatives of U and U_{LINC}. As $i = \frac{2\pi}{\Phi_0}\frac{\partial U}{\partial \theta}$, and $i_{\text{LINC}} = \frac{2\pi}{\Phi_0}\frac{\partial U_{\text{LINC}}}{\partial \theta_c}$, the following relation always holds:

$$\frac{\partial U_{\text{tot}}}{\partial \theta}(\theta, \phi_d) = \frac{\partial U_{\text{LINC}}}{\partial \theta_c}[\theta_c(\theta, \phi_d), \phi_d]. \tag{4.39}$$

Higher order partial derivatives can then be computed from this identity, making it possible to relate partial derivatives of U_{tot} as linear combinations of partial derivatives of U_{LINC}, weighted by partial derivatives of $\theta_c(\theta, \phi_d)$.

To explicitly calculate the relation between θ_c, θ and ϕ_d, we can write out the currents

through both the LINC and the series inductor:

$$\begin{aligned} i_{\text{LINC}}(\theta_c, \phi_d) &= \frac{\Phi_0}{2\pi L}\theta_c + 2I_J \cos\phi_d \sin\theta_c, \\ i_p &= \frac{\Phi_0}{2\pi L_p}\theta_p. \end{aligned} \tag{4.40}$$

Equating these currents and summing the phase across the two elements, we get the total phase θ as:

$$\theta = (1+\beta_p)\theta_c + 2\beta_J \cos\phi_d \sin\theta_c \tag{4.41}$$

where we have also defined $\beta_J = L_p/L_J$ for convenience. Unfortunately, to compute our desired derivatives, we actually require the partial derivatives of the inverse of this function instead. The proof for invertibility involves computing a Jacobian, but ultimately holds whenever the LINC potential is single-valued, or specifically when $1+\beta_p-2\beta_J > 0$. Computing derivatives then simply involves composing the identity $\theta = \theta\left[\theta_c(\theta, \phi_d),\ \phi_d\right]$ and taking derivatives on both sides. Details about this calculation can be found in Appendix E of [199].

The net effect of the above analysis is to produce the linear participation ratio of the LINC in the total potential, given by a simple voltage division between its total inductance and the linear inductor:

$$p_{\text{LINC}} = 1 + \beta_p + 2\beta_J \cos\phi_{\text{DC}}. \tag{4.42}$$

This inductor participation then normalizes the frequency and higher order nonlinearities of the LINC recursively. For the lowest order relevant Hamiltonian terms, we have:

$$\begin{aligned} \omega_L(\beta_p, \phi_{\text{DC}}) &= p_{10}\ \omega_L(0, \phi_{\text{DC}}) \\ \alpha_L(\beta_p, \phi_{\text{DC}}) &= p_{10}^3\ \alpha_L(0, \phi_{\text{DC}}) \\ g_{3\text{wm}}(\beta_p, \phi_{\text{DC}}) &= p_{10}^{3/2}\ g_{3\text{wm}}(0, \phi_{\text{DC}}). \end{aligned} \tag{4.43}$$

This simple normalization is clearly visible in Fig. 4.8c, where β_p has been varied such that the total inductance is the same, i.e., the frequency at half flux has been preserved. Overall, any effect due to the series inductor preserves the LINC's parity protection and is significantly simpler to analyze than in charge-driven mixers like the SNAIL.

4.3 Conclusions from LINC theory

Since the LINC is a novel driven Hamiltonian, it is useful to take a step back and consider the general takeaways from our theory analysis of its performance. Central to its benefits is the combination of the benefits of Kerr-free three-wave mixing and mixer balancing. This results in a nonlinear element that is nearly linear when idling and only activates a parity-protected nonlinearity when driven, suppressing a large fraction of drive-induced parasitic processes. The benefits of such a mixer are significant, offering possible advantages not just in bosonic and qubit control, but also general frequency conversion, and amplification.

In the bosonic context, the LINC promises to break the trade-off between fast nonlinear control and the idle errors introduced by a nonlinear ancillary mode, which has been a long-standing research topic [132, 200, 201, 202]. We will specifically explore its advantages in bosonic control through the rest of this chapter and the next. The parity protection in the LINC could also provide important advantages through its power handling and multi-tone operation. This is useful in multiple contexts, with the simplest being the activation of a simultaneous parametric coupling between multiple neighboring elements. If the LINC were used as an amplifier, an array of LINCs should perform equivalently to an array of balanced RF-SQUIDs, potentially outperforming both SNAIL and regular RF-SQUID based amplifiers. Such an implementation would potentially allow simultaneous gain at multiple frequencies, easing the constraints on multiplexed readout [203, 204]. Finally, the LINC could also simultaneously activate multiple types of parametric processes,

giving rise to new bosonic control techniques. For example, by activating a resonant beam-splitting and two-mode squeezing between two oscillators in the high-Q regime, one could realize a direct parametric quadrature-quadrature coupling between them, enabling two-qubit gates for the GKP code [205, 206].

Simulations can only take us so far – proving that this novel driven Hamiltonian indeed provides its promised advantages requires an experimental implementation. In particular, one will have to solve the engineering challenges of delivering the differential bias and drive in a compact architecture, fabricating high-Q shunting inductors or junction-array. There also exists the general design challenge of being able to isolate the advantages of the LINC's linearity on the bosonic modes, without diluting them through other nonlinear elements in the system. Thankfully, our analysis of the LINC's sensitivity to small asymmetries shows that its experimental implementation need not be impractically precise, which means that the coupler should perform well even with realistic imperfections. Our next section will explore a preliminary implementation of a LINC device, and the chapter after will introduce a general bosonic architecture based on the LINC that allows both continuous variable and number selective bosonic control.

4.4 Preliminary experiments with the LINC

To demonstrate the LINC as a high-fidelity coupler, one could hypothetically just utilize the experimental buffer-mode package from [147], replacing the SQUID with one that has a center inductive shunt. However, there are a few key design constraints that would be useful to satisfy. The first is the ability to measure the LINC to verify its performance. As a reminder, the ideal LINC is a linear oscillator at its half-flux operating point. This means that it cannot be Rabi-driven, and unlike the DDS, it would have no dispersive shift to a coupled stripline readout resonator and hence could not be directly measured. Even if the

LINC were built with an array of (~ 10) junctions instead of a compact super-inductor as its center shunt, it's dispersive shift would be much smaller than a regular transmon-like qubit, and it would still be difficult to characterize near its (shifted) Kerr-free operating point. This means in order to fully evaluate the LINC's performance, one would need to treat it as a usual bosonic mode, and incorporate a dedicated ancilla qubit (transmon) to measure its state. Such a test chip would then at least need to contain a LINC, an ancilla transmon, a readout mode and a Purcell filter. The second design constraint that would be useful to satisfy would be to make the drive and DC bias delivery more compact. Similar to the DDS, the LINC requires just a single control knob for both its drive and DC bias, one that threads symmetric flux through the outer LINC loop. In the experimental package used in [147], the drive and DC bias were separately realized using a large 3D buffer cavity and a copper electromagnetic coil respectively. Given our renewed understand of how to optimize specific driven phase drops across the junctions in EM simulations to maximize the differential-ness of the drive and bias, we can redesign the experiment to instead combine the drive and bias-lines into a single compact structure. This would make any multimode experiments using the LINC far more scalable and simpler to achieve. When combined, this total chip containing a LINC, transmon ancilla, and drive and readout apparatus can act as a single 'plug-and-play' chip for universal bosonic control and measurement.

How does one design a drive and bias line for the LINC? To be compatible with high-Q bosonic control, the drive line must:

1. Deliver AC flux in a superconducting package
2. Drive the LINC precisely differentially
3. Act as a filter that prevents the high-Q modes in the system from being limited by the drive port

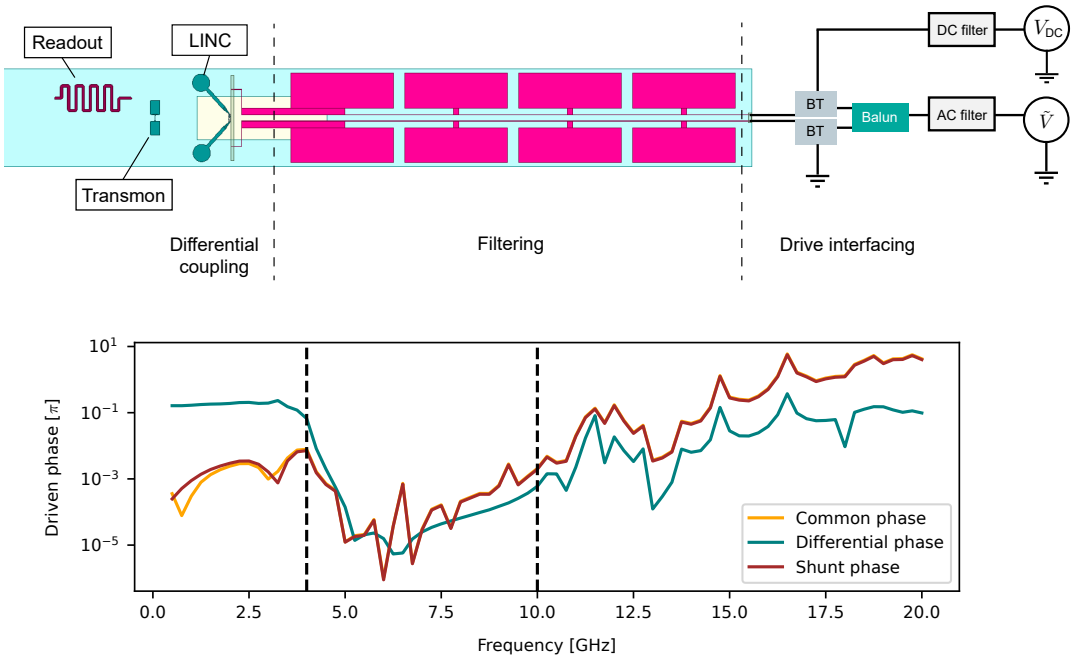

Figure 4.9: **Experimental implementation of a driven LINC Top**, Chip layout with LINC, ancilla transmon for measurement, and filtered line for differential drive delivery. The latter can be divided up into three sections for design intuition, as described in the text. BT represents Bias-Tees for combining the DC bias and AC drive, and the Balun ensures a differential excitation at the drive ports. **Bottom**, Driven phase across the outer junctions and center shunt at 10 nW of drive power for the above LINC design, on a log scale. The drive is differential up to the stopband, which extends between 4 and 10 GHz. The performance in the high-frequency band is faces convergence issues in numerical simulation, and requires more careful analysis.

4. Not contain standing modes that participate in the LINC and might therefore contribute parasitic driven processes, and

5. Be galvanically connected to allow DC-biasing of the LINC.

In particular, the LINC's differential drive has one extra constraint compared to the DDS – not only must the driven phase across each outer junction be equal and and opposite, the driven phase across the central inductive shunt must also be zero. Studying how important the satisfaction of this latter constraint might be to the performance of the coupler is an interesting subject for future research.

In general, any galvanically connected (on-chip) transmission line that carries driven current can deliver both DC and AC flux. One must then fashion this simple transmission line into a more complicated drive-line that satisfies the requirements listed above. The overall intuition for experimentally designing such a drive line is to divide it up into three effective parts, from the edge of the chip to the LINC loop, as outlined in Fig. 4.9. The first part must simply appropriately interface with external cables and connectors that deliver the drive and DC bias. This generally involves some impedance transformation from the input impedance of the on-chip drive line to the usual 50Ω. It also involves some wire-bonds that connect to the drive ports at DC, and will generally also require an external bias-Tee (or equivalent device) that combines independently filtered DC and AC control lines. The second, central section is a (multi-pole) filter that provides Purcell protection to any high-Q modes in the system. The filter can either be a low-pass filter for pure beam-splitting/ coupling operations, or a bandstop filter when both beamsplitting and squeezing like controls are required. Importantly, this filter can often contain standing modes that may have distributed fields, and one must take care to have these not couple to the LINC or the other high-Q modes in the system. Finally, the last portion is the part that couples to the LINC device, and must ensure the differential drive. This part will in general contain a current-carrying wire near the LINC loop, which can also have a stray-capacitive coupling to the LINC's dipole moment. If the LINC is arrayed, this part of the drive line must also ensure that each LINC loop receives approximately equal flux. Finally, this part of the drive line also sets the strength of coupling to the LINC loop, which can be adjusted to ensure that stiff-pump approximations hold, and one does not need to drive too much power or current at the device's drive port in order to get a reasonable drive or bias flux respectively.

To fully take advantage of high-Q bosonic modes, we will want to operate this chip in a monolithic 3D superconducting package. This comes with some additional constraints

primarily that there are no on chip ground planes, and any DC connections can only be made at one end, where the chip enters the monolithic package. This means that any DC-compatible drive line must start at that end, reach the LINC loop, and end back at the same end of the chip. This design constraints causes us to design a twin coupled transmission line for the drive and bias, instead of a single line as might be possible in architectures with on-chip ground planes. These twin transmission lines are connected at the end near the LINC, forming a short that delivers the required AC and DC flux, which we will refer to as the flux-line section. Overall, this forms a four-port network, with two ports at the driven end of the line and the other two at the LINC junctions. We can simply describe the desired S-matrix verbally – the driven phase drop across the junctions should be uniformly negligible for any common excitation of the two drive ports, and a differential excitation of the drive port must only result in a differential excitation at the junction ports. In addition, any transmission through this drive line must contain a band-stop near the LINC, transmon and oscillator frequencies. Again, due to the the absence of an on-chip ground plane, this can only be realized through distributed multi-pole filter that utilizes the tunnel as its (capacitively connected) ground plane. In particular, we implement a 4-pole twin capacitively shunted transmission lines, with a set of large capacitive pads that are intended to act as a low-pass filter up to 4 GHz (see Fig. 4.9). Due to its distributed nature, this filter automatically has a re-entrant pass-band above 10 GHz, which is in fact useful to us as this band allows the delivery of single and two-mode squeezing drives. The simulated differential, common and shunt driven phase for a LINC coupled to such a device is shown in **??**b. Note that because this implementation relies on a differential drive at the drive ports, it also needs an external balun that creates this differential drive from a single control line, and phase-matched cables inside any light-tight shielding that preserve this differential-ness. Future experiments may be able to make this filter more compact through an alternative design, utilizing high-kinetic inductance

materials like granular Aluminum, or incorporating an on-chip balun.

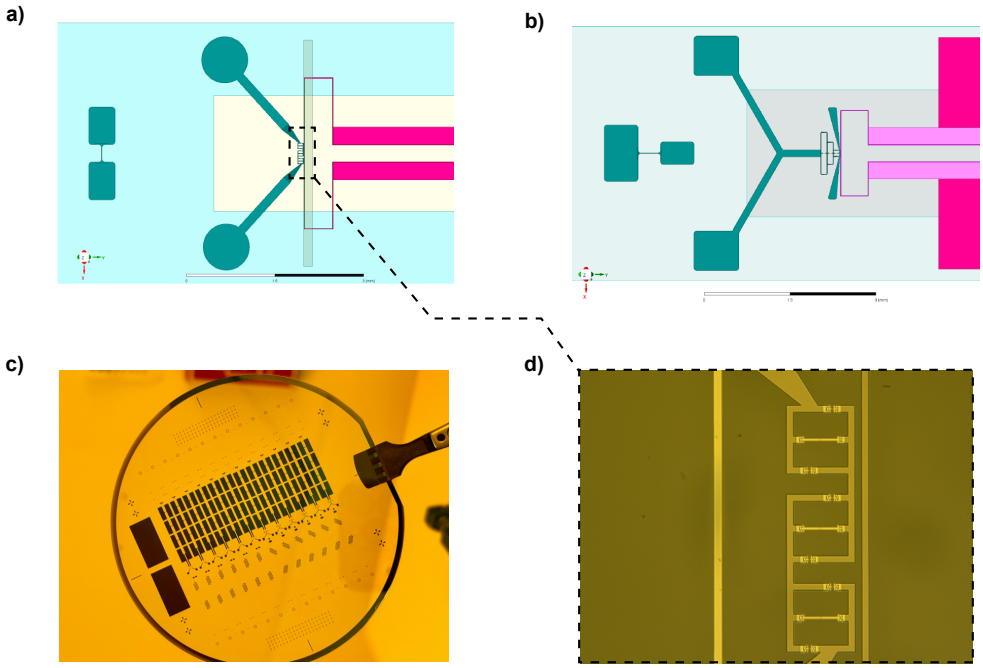

Figure 4.10: **LINC designs and fabrication a)**, LINC design with dipole moment parallel to flux-line section, which is easy to array but requires fine-tuning for drive symmetry. **b)**, LINC design with perpendicular dipole moment, which is more resilient to geometry changes and drive frequency due to its inherent symmetry. Arraying requires progressively larger loops. Both designs have transmons with dipoles oriented to maximize coupling. **c)**, Fabricated chips on a 4-inch sapphire wafer with photolithography and etched Tantalum. **d)**, Optical micrograph of arrayed parallel-dipole LINC from part a).

We test two types of LINC devices, each with an array of 3 loops, coupled the above drive lines and an ancilla transmon. The first design naively arrays the LINC along the flux-line section, to directly enforce the uniformity and symmetry constraints (Fig. 4.10a). Unfortunately, in this orientation, the voltage drop across flux-line section is exactly parallel to the LINC's large dipole moment, which can directly drive the LINC's common mode through any stray capacitive couplings. By finely tuning the angle of the LINC's capacitive arms, one can adjust the capacitance matrix across the two-junctions to achieve a differential drive. However, this adjustment is highly sensitive to slight changes in geometry, like in the orientation of the coupler arms or the size of the flux-line section, and also

does not uniformly work at all drive frequencies. In particular, this optimizes the delivery of drives in the beamsplitting band for use as a coupler, but not in the squeezing band for general Gaussian control. Alternatively, we can be slightly more clever in our design and orient the LINC dipole perpendicular to the flux line section, as shown in Fig. 4.10b. This design provides perfect symmetry that *guarantees* a differential drive at all frequencies and is resilient to any symmetric changes in design – in fact one can precisely check this symmetry by cutting the design in half with a perfect-E boundary, yet still achieving the same driven parameters. This geometry then enables universal parametric Gaussian control, assuming the mode frequencies are carefully designed to avoid obvious collisions. A caveat of this design is the difficulty in arraying LINC loops, which must be propagated perpendicular to the flux-line section. This results in loops that are further away from the section being larger in area to ensure uniformity in driven AC and DC flux, which could result in higher sensitivity to any stray magnetic fields, and higher loop inductance. We fabricate both designs on a 4-inch HEMEX sapphire wafer with major features designed in photolithography with high-purity Tantalum, and Josephson junctions designed with electron-beam lithography in a standard Al-AlO_x-Al format (Fig. 4.10c,d).

We test chips of both designs in a high-purity ($5N$ Al) superconducting package, with preliminary fabrication rounds achieving meager LINC coherences of $\lesssim 1\mu$s. Even these poor coherences let us test the basic flux biasing and beamsplitting abilities of the LINC. Specifically, we observe both the LINC and the coupled transmon to tune with flux over their expected ranges (Fig. 4.11a). The slope for these frequency tunings near half-flux gives us an estimate for the expected beamsplitting rate through Eq. 4.23. Specifically, parking at $\phi_{\text{DC}} = 0.52\pi$, we test the LINC-ancilla beamsplitting by preparing an excitation in the ancilla, sweeping the beamsplitting drive across resonance, and observing the ancilla excitation as a function of time. The corresponding chevron is shown in Fig. 4.11b, and corresponds to a Swap time of $t_{\text{SWAP}} = 24$ ns. We perform such chevron experiments as

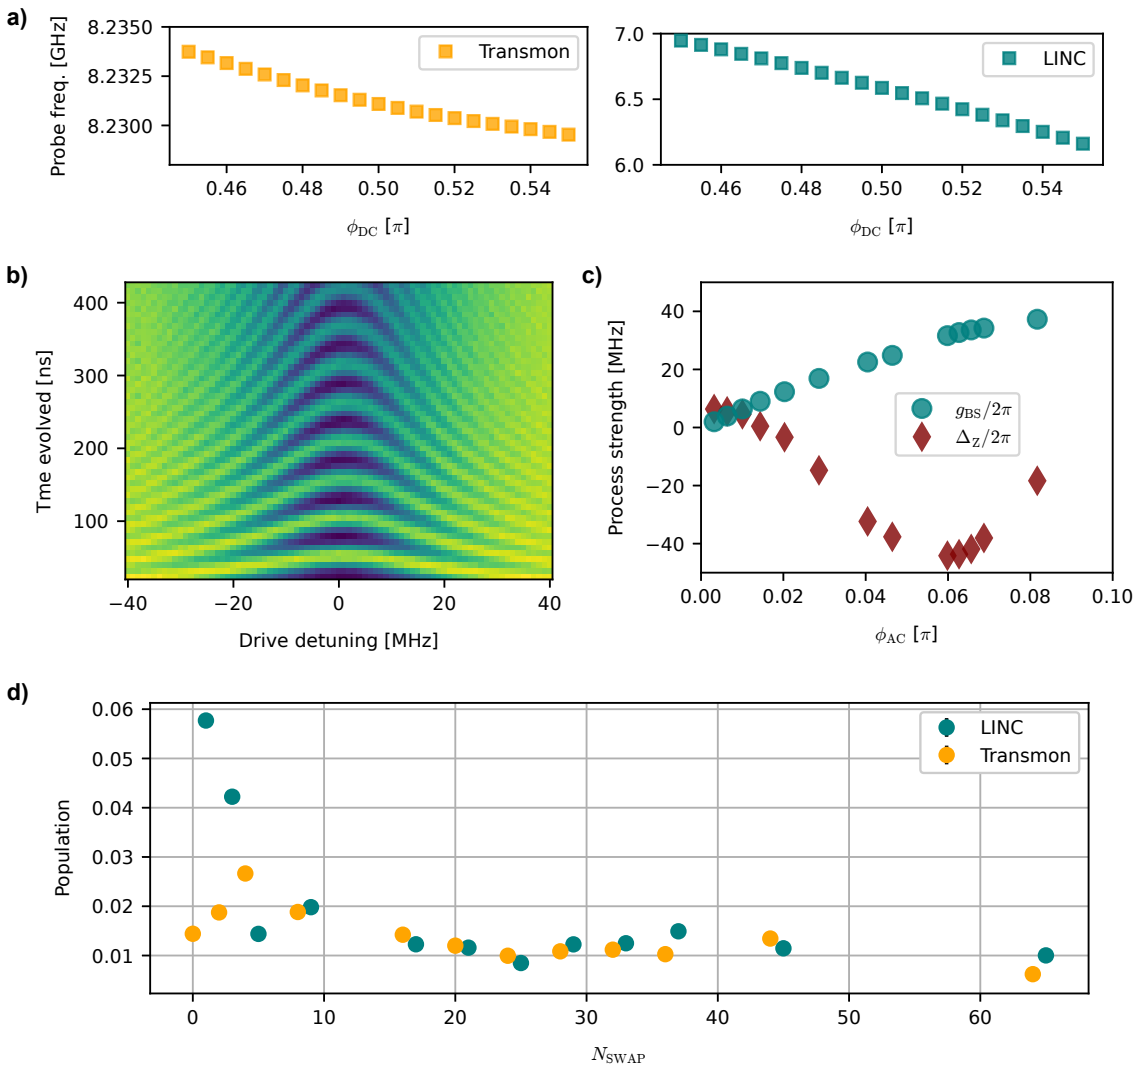

Figure 4.11: **Preliminary LINC measurements a)**, Transmon and LINC frequencies as a function of the LINC's DC bias. LINC spectroscopy is performed through the transmon, using three-tone spectroscopy (readout, transmon, LINC). **b)**, Direct chevron measurement of LINC-transmon swaps, through preparation and readout of the transmon, as a function of differential drive frequency. This roughly corresponds to a 24 ns SWAP. **c)**, Analysis of several chevron experiments extracts beamsplitting strength g_{BS} and relative Zeeman shift Δ_Z as a function of drive amplitude $|\phi_{AC}|$. These show the expected linear and quadratic variation at low drive amplitudes respectively, but the latter shows non-monotonic behavior at higher drive strengths. **d)**, Measured population in the transmon and LINC as a function of number of (calibrated) swaps. Each measurement is performed using the protocol in [164], and error-bars are smaller than the respective marker sizes.

a function of drive amplitude, and extract g_{BS} vs ϕ_{AC} up to $t_{SWAP} = 9$ ns in Fig. 4.11c.

The LINC-ancilla beamsplitting resonance also has an interesting non-monotonic Zeeman

shift, which we extract from the same chevrons and plot in Fig. 4.11d – this might be the result of non-ideal mode frequencies and complicated sideband interactions. Finally, even with these poor coherences, we can characterize any drive induced leakage or excitations by performing multiple swaps and checking the ancilla population. We see no evidence of drive-induced heating in this system. Instead, the LINC and ancilla thermal populations equilibrate after a few swaps, which effectively cools the LINC mode (Fig. 4.11d).

Overall, these measurements prove that the compact differential bias and drive delivery in these chip designs work. They also show that the LINC can perform fast three-wave mixing beamsplitting, and does not obviously cause drive-induced excitation or leakage. A proper analysis of the LINC's performance will require a high-coherence device, which is the subject of current work. Particularly interesting measurements in this minimal device will include testing the LINC's linearity near the half-flux point, the spectrum of its flux noise, and characterizing its (hopefully sparse) drive-induced transitions as a function of drive frequency and amplitude. Even with a high-coherence device, however, fully unlocking the LINC's capabilities as a bosonic control element will require the introduction of at least two high-Q modes that participate in the LINC. This will allow direct tests of the inherited Kerr, flux noise, and the performance various forms of multi-mode Gaussian control on the bosonic modes. The next chapter thus outlines a full architecture for bosonic control based on the LINC. Crucially, even this will not be enough for our final goal – in an ideal bosonic control architecture, we want both continuous variable and number-selective controls, and ideally desire 'fault-tolerant' light-matter interactions that never propagate the errors in the matter to the bosonic modes. We will thus also introduce a new form of bosonic control that is only possible with linear high-fidelity parametric couplers like the LINC, and numerically demonstrate its fault-tolerant properties.

Chapter 5

Dynamic light-matter interactions

Our overall goal remains the continuous variable control of quantum information stored in electromagnetic radiation, using the toolbox of circuit-QED. The parametric controls enabled by driven nonlinear circuits introduced in the last two chapters are an important stepping stone towards this goal. They allow direct drive-controlled activation of continuous variable (Gaussian) Hamiltonians, and also allow the propagation of entanglement across multiple (bosonic) electro-magnetic modes. In particular, a parametric mixer like the LINC executes a delicate balance – where it simultaneously presents itself as a passive linear element to the idle bosonic modes, yet activates coherent nonlinear control when desired. It also explicitly remains in its ground state even when driven, which prevents any parasitic entanglement or propagation of errors to the bosonic modes. Unfortunately, these clean parametric controls are not sufficient, as they do not allow non-Gaussian Hamiltonians, or any measurements of the bosonic mode's state. As we learned in Chapter 2, the strong dispersive shifts available in regular circuit-QED are a natural resource for such non-Gaussian (number-selective) controls and for direct measurement of important continuous variable attributes like number-parity and the Wigner function. However, naively reintroducing dispersively coupled nonlinear elements – like an ancilla transmon – to this bosonic architecture will undo much of the advantages gained by the linearity of the LINC.

Specifically, the static nonlinearity of the bosonic modes and their undesirable dephasing errors would be entirely limited by the inherited Kerr and thermal shot-noise induced dephasing from the transmon, and not by the LINC.

Since a clean Kerr-free parametric controller forms an ideal ancillary element to the bosonic modes, we would ideally like this to be the only circuit that the modes see in their environment. Practically, this means we want their nonlinear (dispersive) couplings to any other mode in the system to be much smaller than their linewidth. Since the parametric controller can turn on drive-activated couplings, we can introduce a nonlinear qubit-like ancilla, for example a transmon, as a third element that participates in this controller (Fig. 5.1a). This transmon is already present in the previously introduced LINC chip, to be able to prepare and characterize the LINC itself. Any non-gaussian control that may be desired can then be activated by turning on a coupling to this ancilla. If the ancilla has a measurement apparatus attached, this control scheme allows universal control and measurement of the oscillators purely through resonant exchanges [136, 207]. Importantly, this is purely a reconfiguration of the necessary elements required for a bosonic architecture – both the nonlinear ancilla and the coupler would have been required for multi-oscillator control anyway. Configuring them in this manner provides a number of automatic advantages. The most obvious is of course the absence of any inherited Kerr and shot-noise dephasing in the oscillators. The oscillator's coherence would then likely be limited by inherited flux-noise dephasing from the Kerr-free controller (which, as a reminder, generally does not sit at the flux sweet spot), which can be mitigated by coherent control techniques, like dynamical decoupling. The second simple advantage of this scheme is that 'unselective' control is the default. This means that any gates that are either Gaussian controls on the oscillators, or Rabi drives on the transmon, no longer need to satisfy any constraints with respect to a static dispersive shift. This also extends to measurement, where any readout of the ancilla ideally has no back-action on the oscillator state.

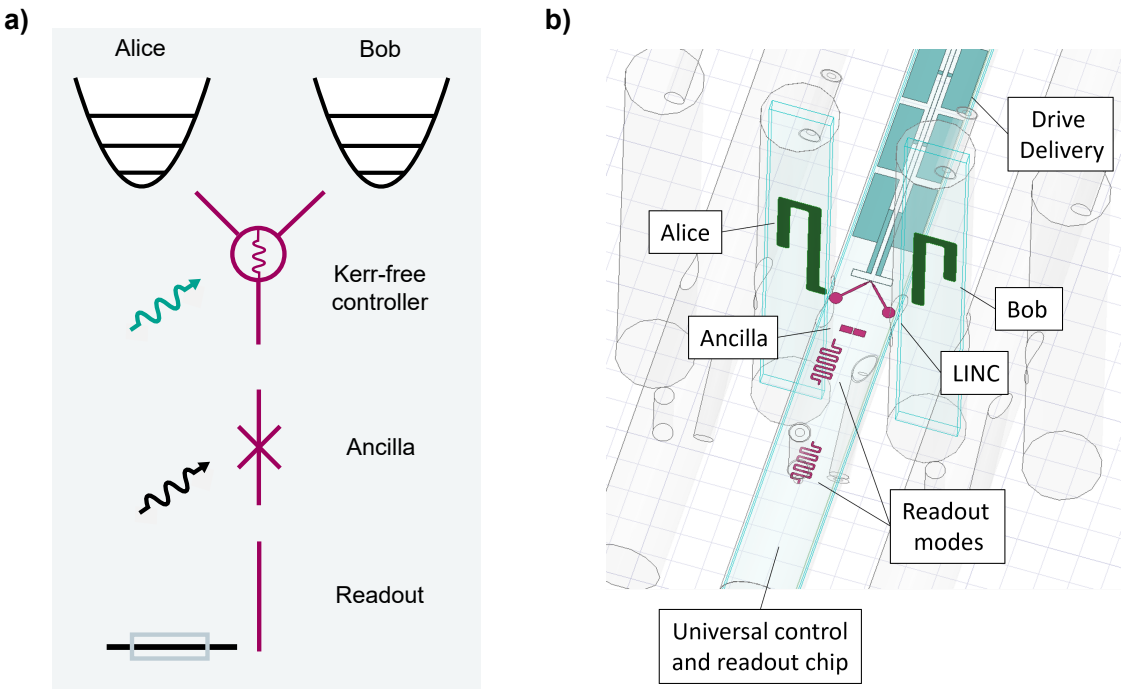

Figure 5.1: **A scalable protected bosonic architecture a)**, A general proposal for a novel bosonic control architecture with a Kerr-free parametric controller that activates direct CV control on the bosonic modes, and couples them to a nonlinear ancilla when required. **b)**, Realization of the control architecture in a) using the LINC package. The LINC chip, with its differential drive and bias delivery, ancilla transmon, and readout apparatus, forms a single chip that can be 'plugged in' to achieve clean universal control and measurement of any microwave oscillator system.

What would actual control sequences on this system look like? Since enacting Gaussian operations on non-Gaussian states provides universal computation, one could simply prepare a Fock state in an oscillator by swapping an excitation from the ancilla, perform Gaussian manipulations, then swap the state back to the ancilla when measurement is required. A simple example of this might be a coherence measurement sequence of the oscillator. To perform a direct Ramsey measurement on the oscillator, one could simply prepare a superposition state $(|g\rangle + |e\rangle)/\sqrt{2}$ in the ancilla, swap it into the oscillator and wait a varying length of time, before swapping it back to the ancilla, applying a $\pi/2$ pulse and measuring its state. This measurement sequence could also be extended to include

echoes, where one swaps the state into the ancilla, with a π pulse enacting the echo, before swapping it back. In general, these pulses essentially 'sandwich' ancilla operations between swap gates to convert them into oscillator operations. This simple trick only works well within the single-photon subspace, but it is highly amenable to the dual-rail encoding, where one would also include swaps to a second oscillator. Specifically, any state preparation, single-qubit gates, and end-of-line measurements can be directly performed through these resonant swap techniques involving the three modes.

The 'swap when nonlinearity is desired' style control is useful, and allows the bosonic modes to be protected while idling and during all Gaussian operations. However, it is still restrictive, and similar to the conditional-displacement control [132], it cannot make use of the full range of techniques that dispersive shifts allow, and may not be a natural choice to all types of bosonic states and encodings. Ideally, one would like to benefit from the clean CV controls available in this architecture while also retaining the ability to perform number-selective controls and measurements when necessary. It turns out that this is indeed possible. Since the parametric controller can couple the oscillator and ancilla on demand, it can also effectively hybridize the oscillator-ancilla system with an off-resonant coupling, regaining a dispersive interaction in the driven frame. Moreover, this dynamic dispersive interaction, which we call 'Dynamic-χ' for short, allows novel physics that could never be emulated on any statically coupled system! We will see that this unlocks a rich set of novel control techniques, from switching both the magnitude and sign of the dispersive shift mid-pulse, to creating symmetries in the dispersive interaction that allow a degree of native fault-tolerance to the ancilla's errors. The scheme we propose is a general technique in quantum control, and has no specific dependence on parametric couplers, or even on circuit-QED. We will spend the next section describing this technique, and demonstrate through simulations some simple analytic techniques that make bosonic control error-resilient, and then we will generalize the control to full numerical optimization

techniques. As a note, this project was in-progress at the time of writing this thesis, and may not contain all the details that might appear in its eventual manuscript.

5.1 Dynamic hybridization of light and matter

Our broad goal is to use controllable couplings $g_{\text{BS}}(t)$ to dress light and matter on demand, such that new nonlinear interactions emerge in the dressed frame that enable control or measurement of quantum information in the light. Central to this technique is the realization that such a coupling can exactly imitate the usual static hybridizations in standard light-matter systems, including circuit-QED. Specifically, if we start with a statically decoupled system composed of a bosonic mode (light, $\hat{a}$) and a nonlinear ancillary 'atom' (matter, $\hat{q}$), a controllable coupling can turn on:

$$H[g_{BS}]/\hbar = H_{NL}(\hat{q}, \hat{q}^\dagger)/\hbar + \Delta_s \hat{q}^\dagger \hat{q} + g_{\text{BS}}(e^{-i\omega_d t}\hat{q}^\dagger \hat{a} + e^{i\omega_d t}\hat{q}\hat{a}^\dagger), \tag{5.1}$$

where Δ_s is the static frequency difference between the atom and oscillator modes, and ω_d is the rate of modulation of the coupling, which in parametric operations is set by the drive frequency. H_{NL} is the nonlinear portion of the ancilla's Hamiltonian. Note that g_{BS} in general also has a phase, which in the context of circuit-QED can smoothly interpolate between capacitive and inductive types of coupling, but we have ignored it here for brevity. We can now enter a frame where the ancilla is rotating at the frequency of the modulation, enacting $\hat{q} \to \hat{q}e^{-i\omega_d t}$, which gives:

$$\begin{gathered} H[g_{\text{BS}}]/\hbar = H_{NL}(\hat{q}, \hat{q}^\dagger)/\hbar + \Delta_{\text{BS}}\ \hat{q}^\dagger \hat{q} + g_{\text{BS}}(\hat{q}^\dagger \hat{a} + \hat{q}\hat{a}^\dagger), \\ \Delta_{\text{BS}} = \Delta_s - \omega_d, \end{gathered} \tag{5.2}$$

where, importantly, we have assumed that the nonlinear terms in H_{NL} are all non-rotating, like Kerr. Our control schemes in the past few chapters have focused on this Hamiltonian in its resonant form ($\Delta_{\text{BS}} = 0$), where it simply enacts the usual exchanges of photons, specifically implementing an iSWAP interaction in the presence of anharmonicity. To achieve our dynamic hybridization, we will instead focus on the off-resonant case, with $\Delta_{\text{BS}} \gg g_{\text{BS}}$.

We first notice that if g_{BS} and Δ_{BS} were held constant, Eq. 5.2 is exactly the Hamiltonian of a general Jaynes-Cummings type static hybridization. However, the manner in which we have turned on these interactions provides a few important differences, which we will discuss in the specific context of circuit-QED. Consider a typical oscillator-ancilla setup where the ancilla is an anharmonic oscillator like the transmon, with anharmonicity K (and static coupling g_s). The typical regime for such an interaction is where the mode detuning is much larger than the anharmonicity, specifically $g_s \sim |K| \ll |\Delta_s|$. Operating in this regime allows a robust implementation where the resulting nonlinear interactions between the oscillator and ancilla can be effectively reduced to the off-resonant effects of a single pole in the interaction, at $\Delta_s = 0$. This gives the usual formulae for χ and inherited resonator Kerr that we have worked with so far in this thesis. However, the true oscillator-ancilla interaction has a lot of fine details that are hidden by the usual static coupling regime. In particular, the ancilla is a multi-level system, and each transition within this system has a separate Jaynes-Cummings like pole in the oscillator-ancilla hybridization. As an example, the transmon can interact with the oscillator through both its $|g\rangle \rightarrow |e\rangle$ and $|e\rangle \rightarrow |f\rangle$ transitions, and thus has at least two poles in its interaction. Any resulting dispersive shift, for example on the $|e\rangle$ level (labelled $\chi_e|e\rangle\langle e|\hat{a}^\dagger\hat{a}$), will have contributions from both these poles that can constructively or destructively interfere depending on the detuning with respect to each pole. The effective analytic formula for χ_e also clearly reflects this multi-pole behavior, which up to fourth-order in the interaction strength g is

given by (derived by J. Garmon):

$$\chi_e = \frac{2g_s^2 K}{\Delta_s(\Delta_s + K)} - \frac{4g_s^4 K(K^2 + 2\Delta_s(\Delta_s + K))}{\Delta_s^3(\Delta_s + K)^3}, \tag{5.3}$$

with clear poles at $\Delta_s = 0, -K$. Importantly, when Δ_s lies in between these two poles, the resulting dispersive shifts from each pole add up to give a significantly stronger interaction, with opposite sign! This regime, often called the 'straddling regime', would be interesting to operate in – yet it is fraught with danger, since colliding with either of the two poles of the interaction can completely spoil the oscillator's linearity and coherence.

Traditional static couplings tend to avoid the above 'straddling regime', as even a slight imprecision in fabrication can lead to enormous changes in the expected static interaction. One could slightly ameliorate this situation by introducing a frequency-tunable ancilla [208, 209, 210, 211, 212]. However, while such tunability relaxes the constraints on precise fabrication, it cannot freely explore the interaction space between the oscillator and ancilla. In particular, it can never tune *through* a pole without significantly affecting the oscillator's quantum state, and when moving between any two points of detuning, it must take a 1-dimensional route that always traverses all intermediate values of Δ_s. This can be highly detrimental in the presence of stray uncontrolled modes or two-levels systems, which can spoil the ancilla as it passes through them. This traditional approach to 'Landau-Zener' like interactions has been a long-standing cornerstone of quantum control – but parametric interactions let us fully generalize it.

If the static g_s, Δ_s are instead replaced by dynamic, parametrically activated $g_{\text{BS}}(t), \Delta_{BS}(t)$, we get three important changes to the resulting physics. The first is that, in general, parametrically activated interactions cannot be as strong as static interaction strengths ($g_{\text{BS}} \ll g_s$). While this seems like a 'bug' at first inspection, it is actually a feature – this means parametric interactions specifically operate in the regime where separate con-

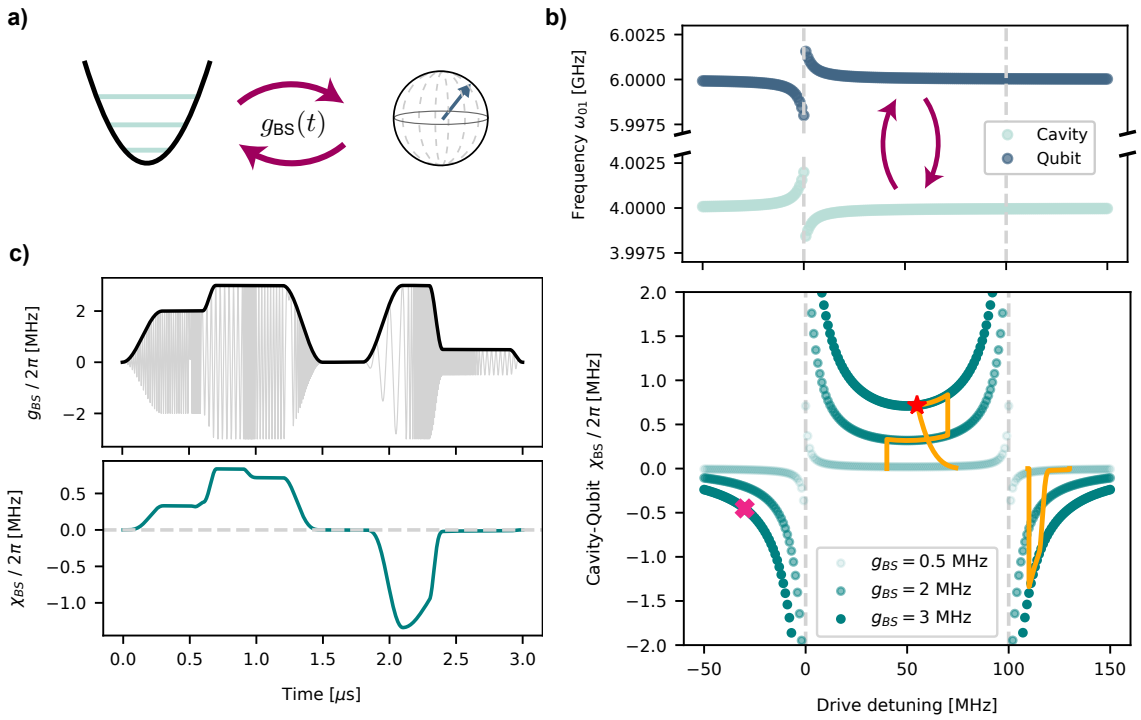

Figure 5.2: **Dynamic light-matter coupling a)**, An illustration of a general controllable coupling between a harmonic oscillator and an ancilla qubit or multi-level nonlinearity. **b)**, The controllable coupling opens up a small avoided crossing between the oscillator and ancilla. Depending on the detuning from each resonant pole in the interaction, the dynamically dressed oscillator-ancilla system can pick up large nonlinear interaction strengths. Here we plot the driven dispersive shift (teal) between an oscillator and a transmon, with $K/2\pi = -100$ MHz and $g_{\mathrm{BS}}/2\pi = 0.5, 2, 3$ MHz. We mark some example points in the interaction where one operates at the usual static hybridization (pink cross), or in the straddling regime where some higher-order nonlinearities have a zero crossing (red star). **c)**, Modulating the strength and effective detuning of the controllable interaction (top) results in a time-dependent dispersive shift (bottom), that takes a specific trajectory in the interaction space (yellow in **a**). Grey modulation in top plot shows $g_{\mathrm{BS}}(t)\cos(\Delta_{\mathrm{BS}}(t)\ t)$, with black envelope showing just $|g_{\mathrm{BS}}(t)|$.

tributions due to multiple resonances in the oscillator-atom interaction come into play, resulting in rich physics around the straddling regime. While g_{BS} itself is small, in fact being negligible with respect to the static detuning, one can still tune Δ_{BS} to be simultaneously small to recover the usual static magnitude of hybridization ($g_s/\Delta_s \approx g_{\mathrm{BS}}/\Delta_{BS}$), and hence the usual strengths of dispersive shifts (see Fig. 5.2b). This brings us to the second major feature of a parametric hybridization: both g_{BS} and Δ_{BS} are set with the

precision of the drive, which can easily be 6 orders of magnitude more precise than the spacing between interaction poles, and than usual fabrication imperfections. This means that we are no longer afraid of colliding with any unwanted interactions, since we can just set our drives to avoid them. In fact, we have this precision in a full *two-dimensional* space, since we have full control on both g_{BS} and Δ_{BS}. This means we never have to cross a pole – we can simply turn the interaction off, and change Δ_{BS} before turning it back on, effectively circumventing it through the second dimension! The overall control scheme thus fully generalizes these light matter interactions to two dimensions, avoiding Landau-Zener effects and making their exploration achievable in an experimental platform.

Finally, we arrive at the primary advantage of making our couplings controllable. If we can modulate g_{BS} as a function of time during a single control sequence, then we can utilize *dynamic* nonlinear interactions. This opens up an entirely new toolbox for bosonic control, where values of $\chi_{\text{BS}}(t)$ can vary significantly within control pulses, including taking on both positive and negative values. We show one such pulse in Fig. 5.2c, where $g_{\text{BS}}(t)$ turns on $\chi(t)$ for a desired amount of time, traversing a wide space in the interaction space including circumventing a pole. However, even this is just scratching the surface – what happens when we turn on a bichromatic parametric coupling, where we have two separate values of $g_{\text{BS}} = (g_1, g_2)$ and $\Delta_{\text{BS}} = (\Delta_1, \Delta_2)$ *simultaneously*? It turns out that there is no rotating frame in which the resulting physics is static (assuming $\Delta_1 \neq \Delta_2$), providing fundamentally new interaction physics. The resulting interaction can then be analyzed using time-dependent Schrieffer-Wolff theory [213], and results in interesting novel phenomena like oscillating dispersive shifts and Purcell effects. We will explore the use of such multi-tone couplings to engineer special symmetries into the oscillator-ancilla coupling later in this chapter.

Good ideas are often reinvented, and such is the case with these dynamic dispersive shifts too. While our approach was to discover them through the lens of perfecting para-

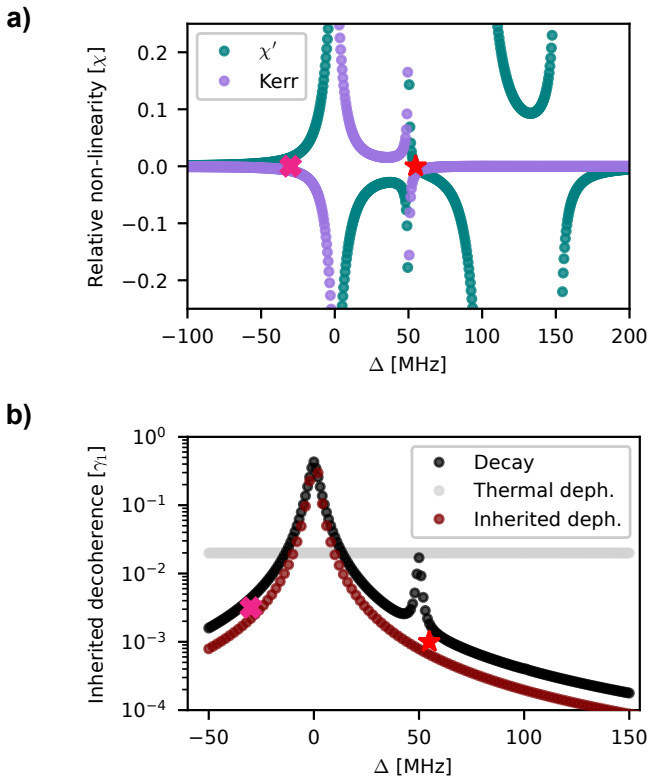

Figure 5.3: **Dynamic light-matter coupling a)**, Higher-order non-linearities in the oscillator-qubit interaction. The pink cross and red star are carried over from Fig. 5.2b, and χ'_e has a zero crossing at the latter. **b)**, Inherited Purcell decay, thermal shot noise dephasing, and direct dephasing in the oscillator from the interaction with a transmon qubit. Each is normalized by transmon coherence, with the first two being normalized by transmon decay (γ_1) and the third by transmon white-noise dephasing (γ_φ).

metric couplers and bosonic control, very similar ideas had already been proposed in the land of qubit control and measurement (see [213, 214]). In those works, the authors explored turning on (effectively static) dispersive shifts between a qubit and a readout resonator, or between two qubits, when readout or two-qubit gates are desired. We apply similar ideas to bosonic control instead, inventing new control techniques that go beyond any current bosonic instruction set architecture [87]. Of course, simply modulating $\chi(t)$ itself performs a gate that entangles the phase of the oscillator with the state of the qubit. However, the dispersive shift is just one of the many nonlinearities this dynamic dressing can turn on. Indeed, any effective hybridization between a transmon ancilla and an oscil-

lator activates nonlinearities at all orders, all of which have interesting structure resulting from the combined behavior of multiple poles. As an example, we can also write out the nonlinearity of the dispersive shift and the inherited Kerr in both oscillators to fourth order in Schreiffer-Wolff theory [215, 216], which is given by:

$$
\begin{aligned}
H_{int}/\hbar &\approx (\chi_e + \chi_e' \hat{a}^\dagger \hat{a}) \hat{a}^\dagger \hat{a} \, |e\rangle\langle e| + \frac{K_{osc}}{2} \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a} + O(|f\rangle\langle f|), \\
\chi_e' &= \frac{4g^4 K^2 (3K^3 + 11K^2\Delta + 15K\Delta^2 + 9\Delta^3)}{\Delta^3 (K + 2\Delta)(3K + 2\Delta)(K + \Delta)^3}, \\
K_{osc} &= \frac{2g^4 K}{\Delta^3 (2\Delta + K)},
\end{aligned}
\tag{5.4}
$$

where H_{int} are all the nonlinear terms involving the oscillator that arise due to the interaction, and g and Δ can be wither static or parametrically activated. The reader may notice the poles in the interactions appearing at $\Delta = 0, -K/2, -K, -3K/2$, across which these nonlinearities may change sign. The extra part $O(|f\rangle\langle f|)$ contains terms where the oscillator interacts with the $|f\rangle$ state or higher, which can be ignored in usual qubit-based ancilla techniques, but will be useful later in this chapter. Each of the terms in Eq. 5.4 have interesting dynamics and multi-pole behavior (see Fig. 5.3a), which could be useful in a general analog bosonic control, in particular for simulating complex bosonic systems. However, in control schemes based purely on dispersive shifts, these higher-order nonlinearities primarily lead to coherent errors, so we must be careful to avoid them by not operating too close to the pole. However, the same multi-pole effects that make this region complicated also create novel interesting operating points – for example, there is a point just above the $\Delta = -K/2$ pole where χ_e' has a zero crossing (red star in Fig. 5.3), and working at this point makes the dispersive shift highly linear. Note that one must also take into account the usual incoherent effects due to hybridization, like Purcell decay and inherited dephasing, while choosing an appropriate operating point (see Fig. 5.3b).

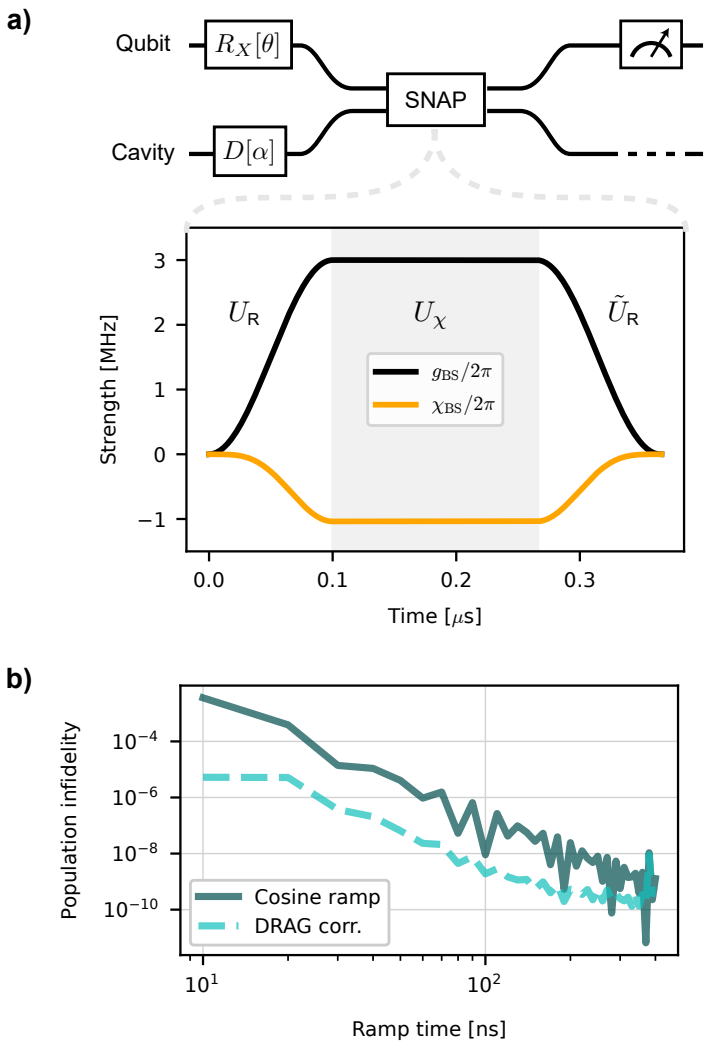

Figure 5.4: **Speed limits of dynamic dressing a)**, Example pulse sequence with initial unselective gates, an entangling operation (SNAP), and final unselective measurement. In a dynamic coupling setting, the oscillator-ancilla interaction can be selectively activated only while required, during the entangling gate. This would involve ramping up to a dressed frame (U_R), sitting in that dressed frame while dispersive control (U_χ) is applied, and then ramping back to the uncoupled frame ($\tilde{U}_R$). **b)**, Numerically calculated infidelity due to population leakage from the intended dressed state when starting in a bare eigenstate (averaged over Fock states). Faster ramp times have more errors, but using counter-diabatic techniques like DRAG can help reduce these errors.

As a reminder, each nonlinear interaction above appears in a driven dressed frame. Since one starts from a bare uncoupled oscillator-ancilla system, and moves to this effective dressed frame to utilize its nonlinear interactions, one may ask – what is the speed limit

for modulating between these frames? Such an adiabaticity speed limit must exist, since turning on the interaction instantaneously will simply result in off-resonant oscillations. As a concrete example, consider a pulse sequence like the one shown in Fig. 5.4a. Here, we would like to perform some unselective displacements and rotations on the oscillator and transmon respectively, then perform an entangling SNAP (see Chapter 2) gate, and then finally measure the state of the qubit. We would ideally like the transmon-oscillator coupling to be off during the unselective gates and measurement, and only be turned on for the duration of the entangling gate. This means we would like the oscillator and transmon to preserve their states as they are ramped to the dressed frame, where their effective dispersive shift is utilized (along with selective π pulses) to entangle them, and then they must return to their undressed frame while preserving this entanglement. One way to analytically derive the speed limit for these ramps is to analyze the inertial correction that appears in the Bogoliubov (or first-order Schrieffer-Wolff) transformation that dresses them.

$$\begin{aligned} H_{\text{inertial}}/\hbar &= i\dot{U}_{SW}U_{SW}^{\dagger} \\ U_{SW} &\approx \exp\left[-\frac{g}{\Delta_{eff}}(\hat{a}^{\dagger}\hat{q} - \hat{a}\hat{q}^{\dagger})\right], \end{aligned} \tag{5.5}$$

where Δ_{eff} is the detuning to the nearest resonant pole that primarily contributes to the dressing, which in the above equation we have approximated to be the linear resonance condition. We would like any off-diagonal contributions due to this inertial term to be much smaller than the diagonal terms, which gives:

$$\begin{aligned} \dot{\theta}_{SW} &\ll \sqrt{4g^2 + \Delta_{eff}^2}, \\ \theta_{SW} &= \frac{1}{2}\arctan\left(\frac{2g}{\Delta_{eff}}\right) \end{aligned} \tag{5.6}$$

where θ_{SW} is the mixing angle for linear hybridization. If only g is changing, and assum-

ing $g \ll \Delta_{eff}$, this speed limit is simply satisfied by $\dot{g} \ll \Delta_{eff}^2$. Crucially, this speed limit is much faster than the inherited dispersive shift χ_{BS}, which means that any ramps for the pulses that turn on the interaction contribute negligibly to the total gate time. We can numerically verify the ramp's performance and speed limit by starting in a bare eigenstate, activating this ramp, and looking at how much population ends up in the intended dressed state (Fig. 5.4b). We see that with achievable parametric couplings ($g_{\text{BS}}/2\pi = 3$ MHz, $\Delta_{\text{BS}}/2\pi = 50$ MHz) we can ramp faster than 100 ns with negligible pulse infidelity. Additionally, because the inertial errors are entirely coherent, one can also utilize counter-diabatic techniques like Derivative Removal by Adiabatic Gate (DRAG, [217, 118]) to further speed up this process. Finally, note that the above metric only captures the population errors. In general, the non-adiabatic contributions of the pulse will also enact phase errors, and fully quantifying these phase errors are part of the ongoing work for this project.

5.2 Control techniques using Dynamic dressing

Now that we understand the general intuition for modulating $g_{\text{BS}}(t)$ and the speed limits of its modulation, let us look at some instances of useful control sequences where the dynamic nature of the dispersive shift provides a measurable advantage. In particular, we will be interested in utilizing dynamic-χ to combat errors from the ancilla that can propagate into the oscillator and its control. Such propagated errors are generally the primary limitations for the performance of bosonic control, and have been shown to directly limit the available gain in bosonic error correction experiments [139, 107, 108].

We start with a technique that utilizes the combination of positive and negative dispersive shifts available in the straddling regime to combat low-frequency dephasing noise in the ancilla. Such low-frequency fluctuations result in shot-to-shot variations of the ancilla's frequency, which can appear in a wide variety of platforms, from sources such as

inhomogeneous broadening in ensembles to fluctuations in traps and control electronics. In the particular context of parametric control with three-wave mixers in circuit-QED, such noise is always present in the coupler due to $1/f$ flux noise. This noise can then propagate to both the ancilla and the oscillator. A well-established control technique to mitigate such low-frequency noise is to use dynamical decoupling, where a sequence of 'echo' pulses refocuses any slow fluctuations [42, 157, 218]. However, utilizing this technique in bosonic control runs into a key problem; any echo on the ancilla also echoes out the dispersive control itself. One way to understand this is to think about a static dispersive shift as a coherent (zero-frequency) detuning 'noise' on the ancilla, which it can decouple from. This means that in usual static dispersive control, such an echo is forbidden, and any gates enacted on the ancilla-oscillator system face the full brunt of dephasing noise. To make this more concrete, we study the particular case of an oscillator parity measurement, where the qubit is prepared on the equator and picks up an $n\pi$ phase for each Fock state n, mapping all even Fock states in the oscillator to $|+X\rangle$ and all odd states to $|-X\rangle$ (see Chapter 2). Since the qubit spends a majority of this pulse sequence on the Bloch sphere's equator, accumulating phase, it is maximally sensitive to any slow frequency fluctuations, which spoil the parity map's fidelity. An emulation of such an effectively static dispersive shift using dynamic-χ for a parity map is shown in Fig. 5.5a,c. In the corresponding Monte-Carlo simulation, the $\langle X\rangle$ trajectories of the ancilla spread out due to the dephasing, causing a clear deviation (infidelity) in where its mean trajectory ends up.

How do we use dynamic-χ to solve this? Notice that we are operating in the vicinity of the straddling regime, where χ can take on both positive and negative values. If we want to echo the ancilla qubit, but not echo out the dispersive shift we can play a neat trick – we can simply flip the qubit and flip the value of χ simultaneously, preserving the net interaction $\chi\hat{a}^\dagger\hat{a}\hat{\sigma}_z$. In particular, this would involve turning on the parametric coupling at two different drive frequencies sequentially, one inside the straddling regime and one outside

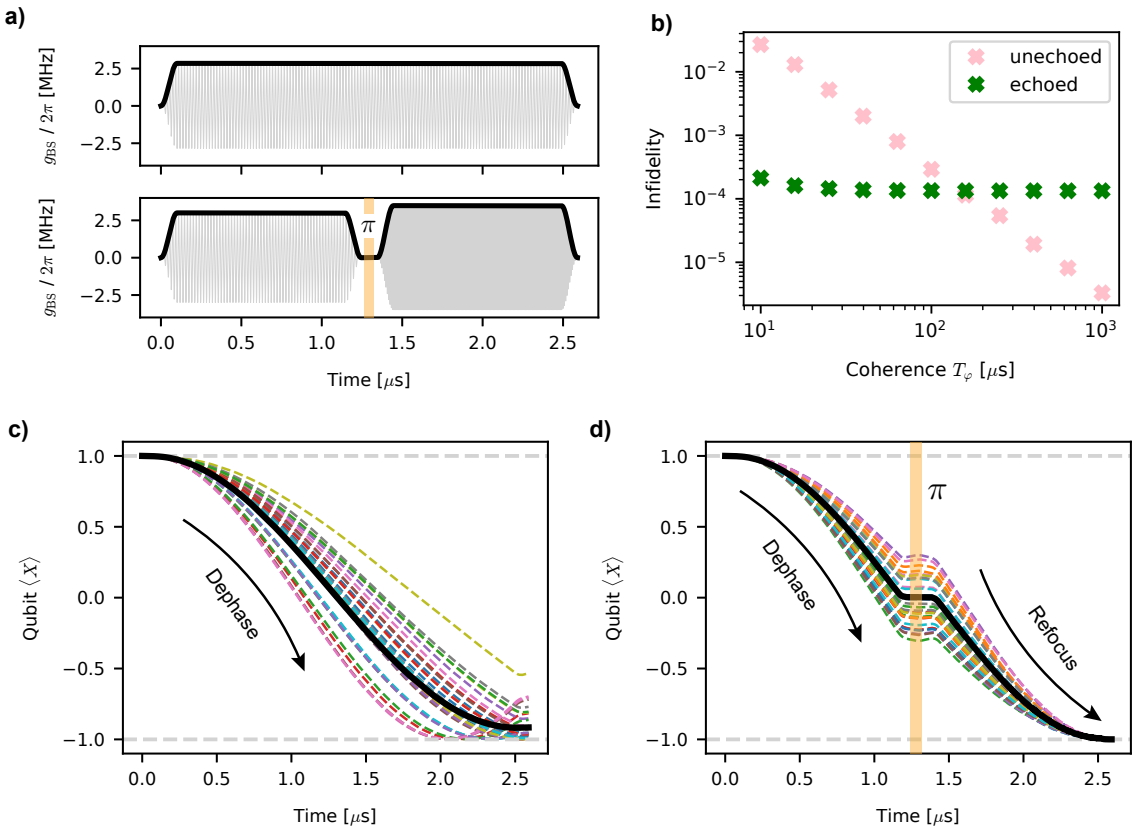

Figure 5.5: **A dynamically decoupled parity map a)**, Calibrated pulses for un-echoed and echoed parity map sequences. The former simply turns on a positive dispersive shift ($g_{\mathrm{BS}}/2\pi \sim 3$ MHz, $\Delta_{\mathrm{BS}} \sim 54.8$ MHz) for a duration such that $\int \chi(t)dt = \pi$. The latter turns on two separate pulses at $g_{\mathrm{BS}}/2\pi \sim 3$ MHz $\Delta_{\mathrm{BS}} \sim 54.8, 120$ MHz respectively, acquiring positive and negative phases such that $\int \chi_{1,2}(t)dt = \pm\pi/2$. The latter includes an echo (π rotation) on the qubit in between the two pulses. **b)**, Preliminary results for parity map infidelity as a function of strength of low-frequency dephasing (T_φ) noise. Echoed pulses are limited by coherent calibration errors and hence do not show improvement with dephasing time, but still show significantly better performance at low coherence times. **c)**, Monte Carlo simulation for individual qubit trajectories with low frequency dephasing noise. Black line represents the average trajectory, which clearly shows infidelity. **d)**, Monte Carlo trajectories with the echoed parity pulse. Low frequency dephasing appears as a shot-to-shot variation in frequency, which refocuses after the echo to provide the significant improvement in fidelity in **b**.

it. Properly executing such a pulse also requires careful calibration, since an ideal echo requires an equal amount of time spent before and after the echo such that the trajectories refocus. To perform the simulation in Fig. 5.5, we calibrate two separate $\pi/2$ rotations

with opposite values of χ. For the first, we choose $\Delta_1/2\pi = 54.8$ MHz (the operating point where $\chi'_e = 0$) and $g_1/2\pi = 3$ MHz, with 100 ns cosine ramps, and finely tune the flat-top time to calibrate a $+\pi/2$ rotation. We then fix this pulse time, and choose a second detuning $\Delta_2/2\pi = 120$ MHz outside the straddling regime, and then finely tune g_2 to achieve a $-\pi/2$ rotation. The total concatenated pulse, with an intermediate 100 ns delay where we execute a Gaussian π_x pulse on the qubit, is shown in Fig. 5.5a. When this pulse is applied to the same Monte Carlo simulation with low-frequency dephasing, the effect of the echo is clearly visible in the refocusing of trajectories (Fig. 5.5d). We fix the low-frequency cutoff for the noise at 1 KHz, sweep its overall strength, and plot the fidelities of both the un-echoed and echoed parity maps in Fig. 5.5b. The echoed pulse significantly outperforms the unechoed pulse at low coherence times, but is limited by calibration errors in this implementation and hence does not itself scale with coherence. Future numerical and experimental implementations should be able to improve this through more careful calibration.

The above sequence shows a simple way to combat pseudo-coherent errors through a coherent control technique. However, how does one create interactions that are resilient to truly incoherent errors, like Markovian dephasing or decay? The answer to this question lies in the fact that the errors, despite occurring at unknown times, are occurring on a known part of the system – the ancilla. This means that if we can engineer the oscillator-ancilla interaction to somehow have a natural symmetry with respect to these errors, they can be prevented from propagating to the oscillator. Such a technique were devised in [138], where it was shown that utilizing three levels of the ancilla and engineering special symmetries with respect to these three levels, can provide tolerance to both decay and dephasing error of the ancilla in a SNAP gate. The central intuition for this resilience is that if the ancilla is operated as a qubit within the $g - f$ manifold, then a decay error can take the ancilla to the orthogonal $|e\rangle$ state. If the oscillator cannot distin-

guish between the $|e\rangle$ and $|f\rangle$ states, such a decay error will not propagate to the oscillator (Fig. 5.6a). Specifically, this happens when the dispersive shifts to both the $|e\rangle$ and $|f\rangle$ states are equal, which is the special symmetry that needs to be engineered in this interaction. This construction also takes care of dephasing errors, which in effect project the ancilla's state onto the $|g\rangle$ or $|f\rangle$ states, which correspond to exactly performing Identity or the desired SNAP unitary on the oscillator state respectively. If the ancilla is then measured to be in the $|g\rangle$ state, the pulse can simply be repeated.

While the central idea of engineering the above symmetry is crucial, its implementation in [139, 138] was through driving parametric interactions directly through the transmon. Such driven interactions through a transmon have numerous drawbacks, covered extensively in Chapter 3, as does reliable repeated measurements of the transmon, which may directly dephase the oscillator. As a result, despite the beauty of the central idea, this scheme has not been widely utilized in bosonic architectures. However, armed with the new toolbox of dynamic-χ, we can revisit this idea and construct a simpler implementation of it. First, we must study multi-pole structure of χ_f provided to us when an interaction is turned on. Since we are interested in a point where $\chi_e = \chi_f$, we can instead study the structure of their difference $\chi_{ef} = \chi_f - \chi_e$ directly.

Unfortunately, as shown in Fig. 5.6b, there exists no point in the interaction where χ_{ef} is zero while χ_e and χ_f remain at a reasonable magnitude. However, this is where a new feature of dynamic-χ comes into play, that of multi-tone parametric couplings.

As discussed before, when two couplings (g_1, Δ_1 and g_2, Δ_2) are simultaneously activated, there is no rotating frame where the effective physics is truly static, giving rise to novel physics that is inaccessible in standard circuit-QED. Specifically, a two-tone coupling has each nonlinear term from the individual couplings, but also has physics at the beat-frequency of the two tones, $(\Delta_1 - \Delta_2)$. However, this latter physics has an amplitude that scales as $g_{1,2}^2/(\Delta_1 - \Delta_2)$, which means if these two tones are far

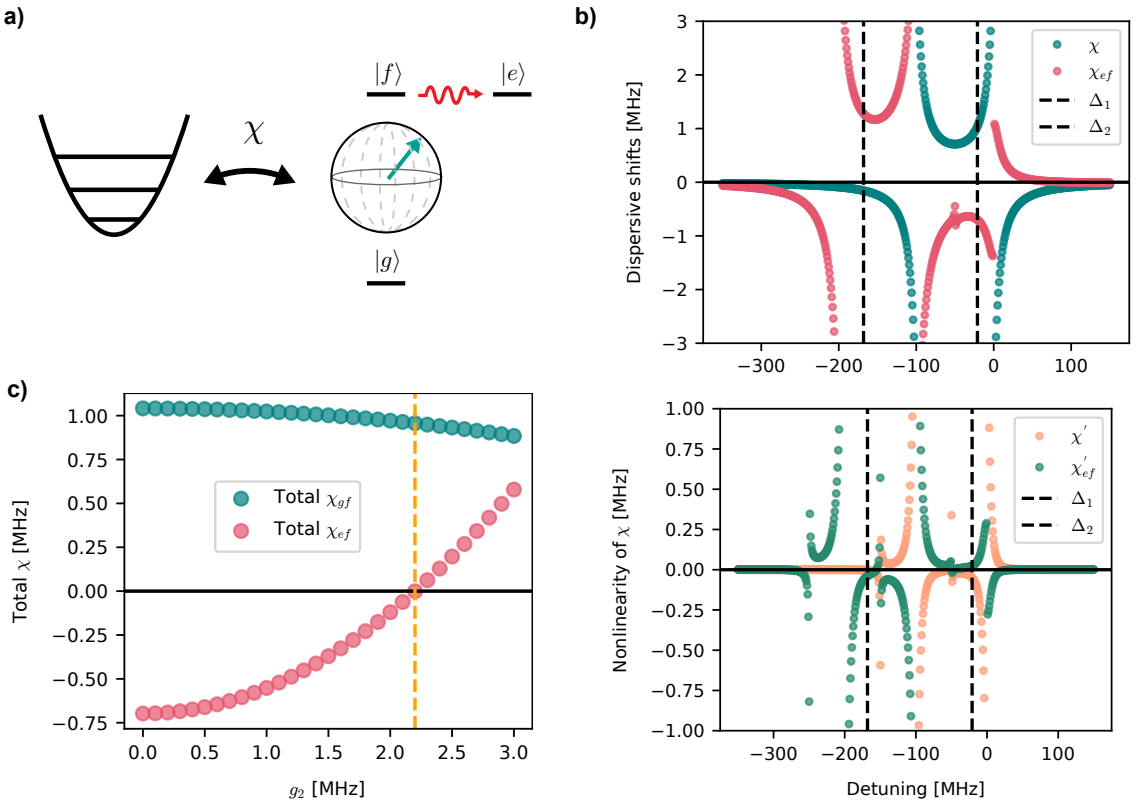

Figure 5.6: **Multi-tone dynamic-χ a)**, Schematic for engineered symmetry in the dispersive interaction to $|e\rangle$ and $|f\rangle$ levels. Creating this symmetry and operating the qubit in the $g-f$ manifold allows the oscillator to be transparent to ancilla decays. **b)**, Turning on a single-tone parametric coupling provides a rich structure in the dispersive shift to the $|e\rangle$ (teal) and $|f\rangle$ levels, but does not directly provide an operating point where their difference χ_{ef} (pink) crosses zero. However, by utilizing two separate tones, one can engineer an effective nonlinear interaction where the net χ_{ef} is nulled. Importantly, one must make sure this dispersive shift is reasonably linear, which requires ensuring that χ'_e (orange) and χ'_{ef} are small at each drive tone. **c)**, We fix the two tones at $g_1/2\pi = 3,\ \Delta_1/2\pi = 21$ MHz, and $\Delta_2/2\pi = 168$ MHz, and sweep the value of g_2 to find the chi-matching condition in a Floquet simulation. We find that at $g_2/2\pi = 2.2$ MHz, χ_{ef} has a zero crossing while $\chi_e/2\pi = \chi_f/2\pi \sim 1$ MHz remain high. This provides a possible operating point for chi-matched pulses.

apart such that $|\Delta_1 - \Delta_2| \gg g_{1,2}$, the effective non-linearities due to each tone simply add [213]! This is then precisely what we look for – we find a pair of tones, specifically $\Delta_1/2\pi, \Delta_2/2\pi = 21, 168$ MHz, where the total χ_{ef} can essentially cancel, and there are no poles in χ'_e or χ'_f either (Fig. 5.6b). We then fix one tone at a feasible strength,

eg $g_1/2\pi = 3$ MHz, and sweep the strength of the other tone g_2 until we observe a zero crossing, at which $\chi_e = \chi_f \approx 1$ MHz (Fig. 5.6c). These simulations were performed using Floquet techniques, and future implementations can calibrate actual pulses of $g_{\mathrm{BS}}(t) = g_1(t)e^{-i\Delta_1 t} + g_2 e^{-i\Delta_2 t}$ to fully test such a χ-matched SNAP gate.

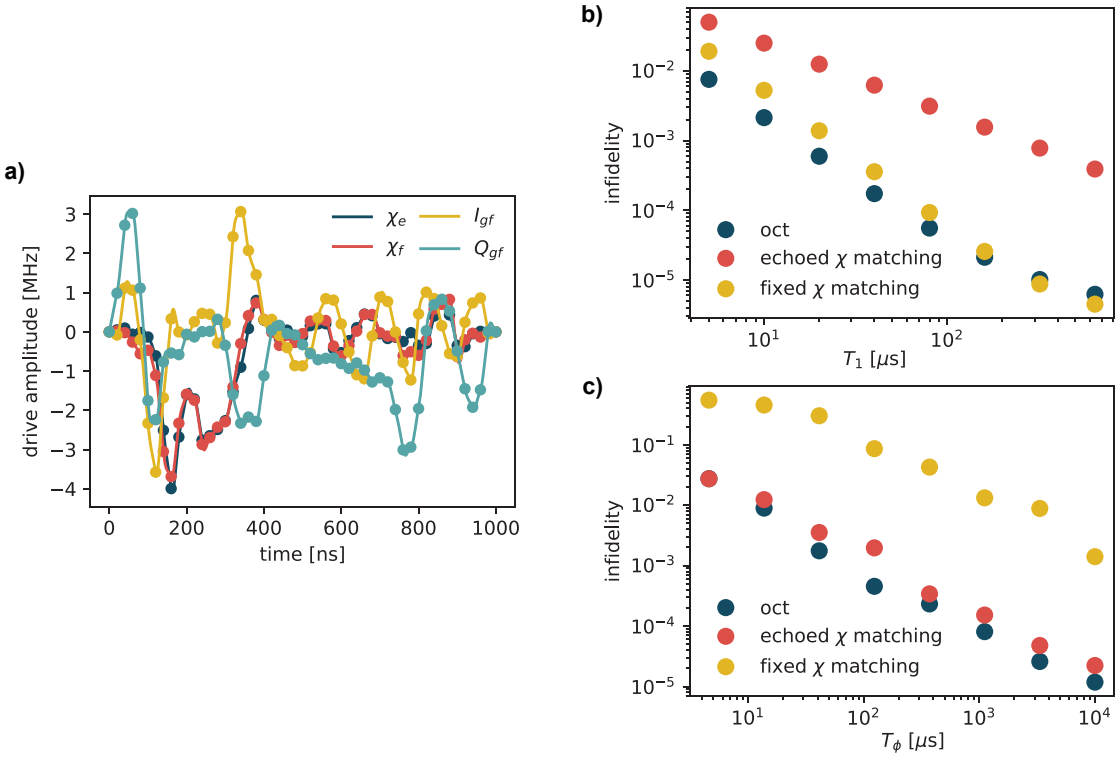

Figure 5.7: **Numerically optimized dynamic-χ a)**, Numerically optimized drive controls for a parity map unitary on the kitten code. We provide $\chi_e(t)$ and $\chi_f(t)$ controls directly to the optimizer, along with $g - f$ Rabi drive control. It automatically optimizes to a χ-matched pulse that incorporates some echoes in the latter half of the pulse. **b)**, performance of the optimized pulse in **a**, compared to perfectly chi-matched and echoed pulses, as a function of ancilla decay rate. The optimized pulse performs as well as the ideal χ-matched pulse, with its slope proving first-order resilience to ancilla decay. **c)**, The same optimized pulse also performs at least as well as the ideal echoed pulse against low-frequency dephasing.

Now that we have studied a few specific analytically derived schemes for novel control using dynamic-χ, one could ask the question – what's the most general form of control using this toolbox? Each pulse sequence outlines so far was composed of sequential sections where χ_e (or χ_f) were ramped to effectively static values, where pulses involving

the ancilla were performed for control or measurement. However, a general pulse utilizing this control technique will have no such restrictions. Indeed, $\chi(t)$ can have arbitrary pulse shapes as long as its changes are within the established speed limit, and the ancilla can be driven simultaneously with any changes in χ. To fully explore this design space, we can appeal to a numerical optimizer that has direct access to $\chi_e(t)$, $\chi_f(t)$, and ancilla drives as controls.

To demonstrate the advantages of this additional design space, we would like the optimizer to design a pulse that has capabilities beyond the two analytic pulse schemes introduced earlier. Specifically, we would like to implement a parity-map sequence that is simultaneously resilient to both ancilla decay and low-frequency dephasing. Each pulse introduced previously is only resilient to one of these two errors. Crucially, even combining these into a χ-matched, echoed pulse is only resilient to low-frequency dephasing, but not to decay errors. In fact it analytically does not seem straightforward to find a pulse that does both – in [139], for example, the authors had to introduce multiple measurements to make the parity map protocol resilient to both types of errors. We test the numerical optimizer by asking it to perform exactly this miracle, specifically to enact a decay and low-frequency dephasing tolerant parity map on the kitten code (Chapter 2). As a reminder, such a parity map is precisely the error detection required for tracking this encoding's syndromes. Specifically, we optimize the cost function:

$$C_{\text{infidelity}}^{T_1,T_\phi} = \sum_k \sum_i \left[1 - \text{Tr}(P_{gf}\rho_i^{(k)}(T)P_{gf}\rho_i^{\text{target}}) - \text{Tr}(P_e\rho_i^{(k)}(T)P_e\rho_i^{\text{jump}}) \right], \quad (5.7)$$

which sums over various instances of low-frequency dephasing trajectories (k), and separates the contributions of trajectories with jumps occurred (in which case the ancilla is projected into the $|e\rangle$ state) and no jumps occurred (in which case the ancilla is projected into the $g-f$ manifold). This Monte-Carlo OCT was implemented and performed by

D. Weiss [219]. After careful optimization, the pulse appears to prefer a matched χ, and performs a few rapid echoes in a portion of the pulse (Chapter 2a). Surprisingly, this numerically optimized pulse both performs as well as the ideal matched pulse with respect to ancilla decay errors, and as well as the ideal echoed pulse with respect to dephasing errors. However, note that this is an optimization over $\chi_{\text{BS}}(t)$ – to convert this into a more realistic simulation, we must convert this pulse into a modulated coupling $g_{\text{BS}}(t)$, and the resulting pulse may be non-unique. Future optimization techniques can explore directly optimizing $g_{\text{BS}}(t)$, converting this scheme into one that is experimentally implementable.

5.3 Initial results and summary

While this technique and the physics it unlocks are promising, our analysis so far has been purely theoretical. To give the reader a sense of its experimental feasibility, I include here some initial results from the LINC device shown in Fig. 5.1. Since this is ongoing work, any results in this section will be superseded by the eventual paper about this experiment. As a reminder, this device ideally allows the bosonic modes Alice and Bob to be statically de-coupled from each other the transmon ancilla, with any interactions in the system only being activated parametrically by the LINC. The transmon ancilla is the only mode in the system that can be directly measured, through a dedicated readout resonator. Performing any characterization of the resonators, including eventually trying to activate and measure the effects of dynamic-χ, thus depends on the strength and fidelity of the parametric beamsplitting interaction between the ancilla and the resonator. The measured chevron for this interaction, and the extracted g_{BS} as function of drive amplitude, are shown in Fig. 5.8a. These coherent swaps directly let us prepare and read out states within the single-photon manifold for Bob, where we measure $T_1 = 463\mu$s, $T_2^R = 891\mu$s and $T_2^E = 895\mu$s (Fig. 5.8b). This implies an extremely long pure dephasing time for Bob,

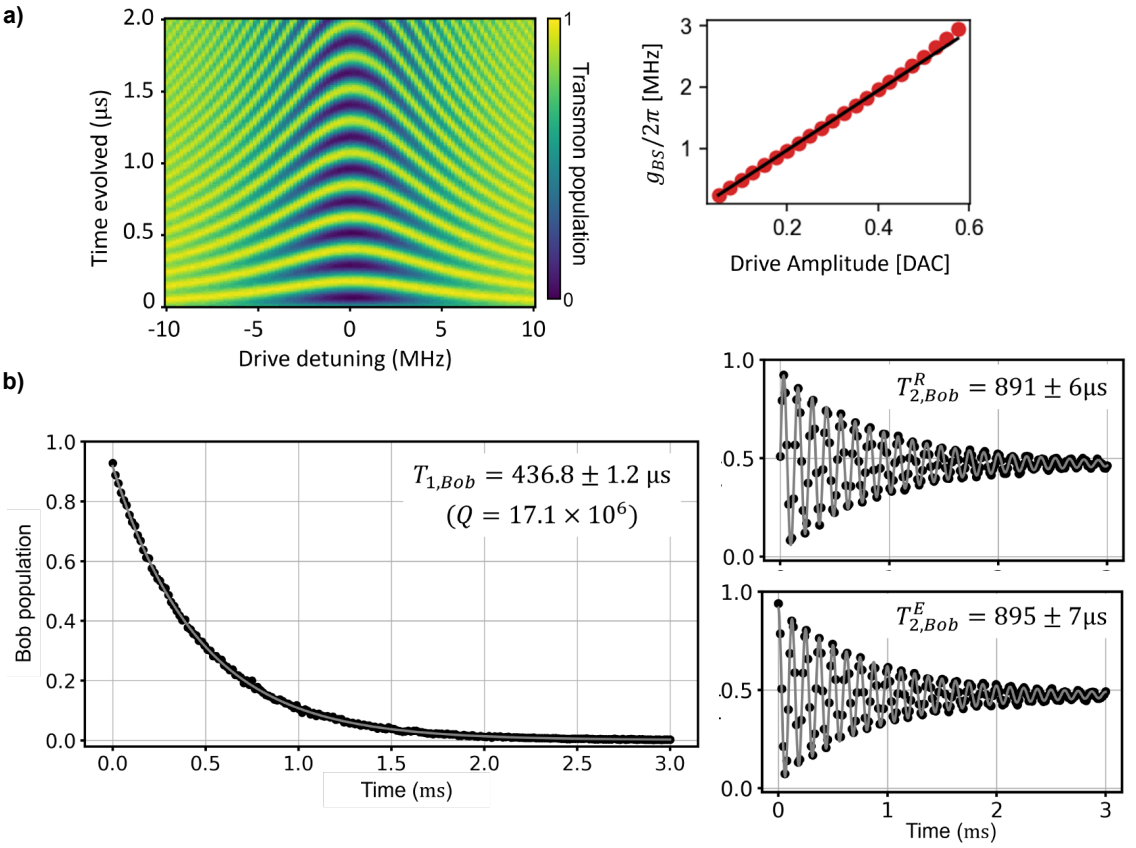

Figure 5.8: **Resonator swaps and coherence a)**, Beamsplitting between the transmon and Bob, with evolution as a function of time and rate as a function of drive amplitude. **b)**, Measured coherence for the Bob resonator, with state preparation, echoes and readout performed through swaps with the transmon.

with $T_\varphi^R \sim 26 \pm 6$ ms. Note that this for a cavity dual-rail qubit, this implies an erasure bias of > 50, which is an order of magnitude higher than that achieved in the SQUID experiment. In fact, the true bound on resonator or dual-rail dephasing is likely even higher, as this dephasing rate is consistent with a low-frequency flux noise of $3\mu\Phi_0/\sqrt{\text{Hz}}$, which can be significantly reduced through dynamical decoupling.

The strong phase coherence we measure in this device hints at bosonic modes that are not limited by cross-Kerr-mediated thermal dephasing from their ancillae. This would mean they are highly linear when idling, and be amenable to demonstrate dynamic-χ with a large on-off ratio. To directly measure the static decoupling, we perform a Ramsey

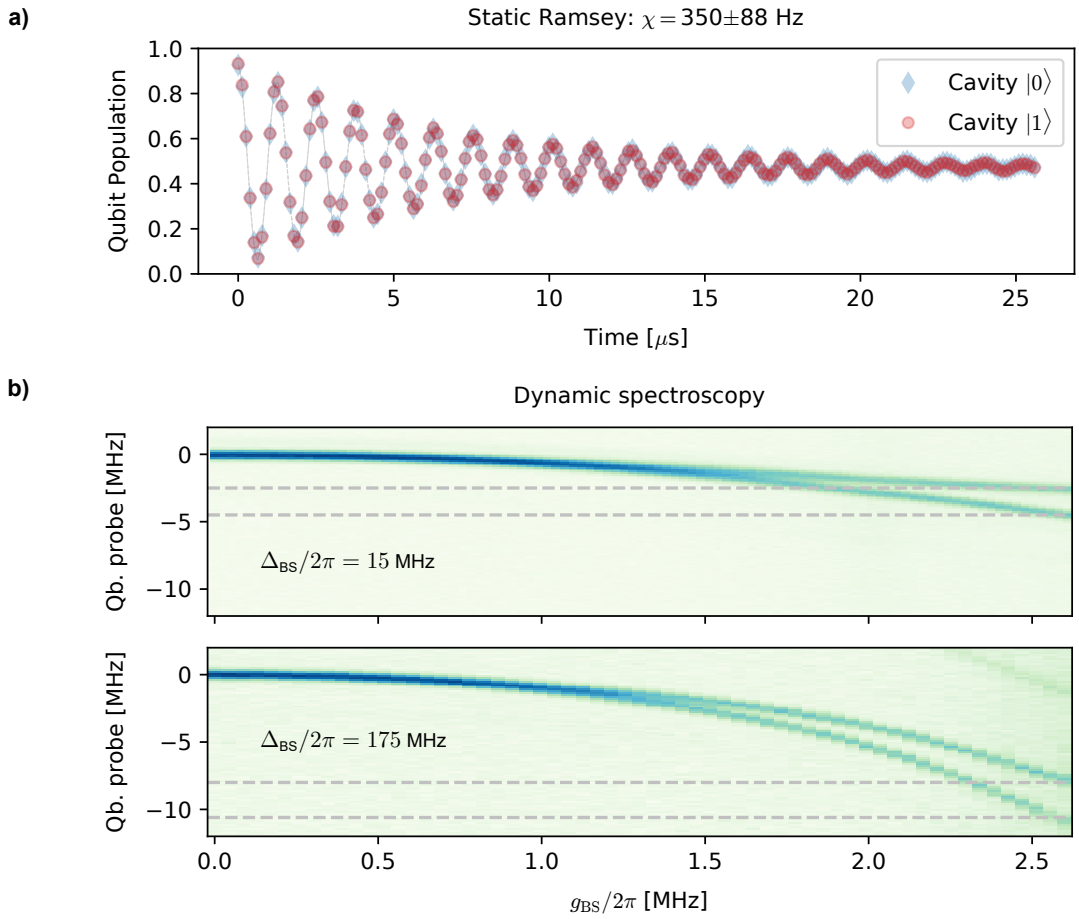

Figure 5.9: **Parametric hybridization in a bosonic system a)**, Ramsey measurements in the transmon with and without a photon in Bob, to bound the static dispersive shift. **b)**, Transmon spectroscopy in the presence of an off-resonant beamsplitter, with a superposition state in Bob, as a function of drive amplitude. The activated χ_{BS} can range up to 2.5 MHz, comfortably in the range of standard dispersive control setups. Different effective detunings produce χ_{BS} of opposite sign.

measurement on the transmon, with and without a photon in Bob. Any dispersive shift would result in a shift in the fringes between the two Ramsey curves. We see no observable shift (Fig. 5.9a), and careful fitting bounds the idle resonator-ancilla cross-Kerr to less than 350 Hz. As a reference, this is two orders of magnitude more linear than bosonic systems implementing ECD control techniques for Kerr-sensitive bosonic codes. We then utilize the techniques introduced in this chapter to turn on a resonator-ancilla dispersive interaction through an off-resonant beamsplitter. We measure this dynamic-χ, by loading a

superposition state $|0\rangle + |1\rangle$ in the oscillator, and performing spectroscopy on the transmon in the presence of the off-resonant parametric drive. As a function of the drive amplitude, we see the transmon get Zeeman-shifted, but we also see a beautiful splitting emerge due to the effective χ_{BS}, as shown in Fig. 5.9b. Crucially, at any particular drive amplitude, the speed of this dynamic dressing is given by the ramp time of the off-resonant drive. This means we are able to modulate the dispersive shift by four orders of magnitude (0.3 KHz to 2.5 MHz) in ~ 20 ns! Moreover, simply choosing a different drive frequency lets us access a different point inside the straddling regime, and obtain a dispersive shift of opposite sign instead. This experiment proves that dynamic-χ is a feasible technique in bosonic circuit-QED. It also allows a number of novel demonstrations, including (to my knowledge) the first Wigner function of a statically de-coupled ($g \sim 0$) bosonic mode, illustrated on the back of this thesis.

Overall, these last few chapters have proposed and analyzed new techniques for bosonic control that break the natural trade-off between the required nonlinearity and the undesired decoherence that stems from light-matter interactions. We now have a way to enact both continuous variable and number-selective controls on bosonic quantum information, using an architecture that may prove exceptionally resilient to errors in the ancillary matter. Central to these proposals were a novel quasi-linear driven non-linearity, the LINC, and a generalization of Landau-Zener type interactions, in the form of dynamic hybridization (or dynamic-χ). How do we take full advantage of these novel architectures and techniques once they become experimentally available? One broad application of course is either analog or digital bosonic simulation, where the rich control space and native ancilla-error tolerance can let us explore novel non-equilibrium many-body bosonic phenomena. However, this will likely require the construction of a large bosonic processor to properly demonstrate.

A more immediate application for this architecture, one that might even be achievable

in the LINC experimental package introduced earlier in this chapter, is a significant improvement in bosonic error-correction. Bosonic encodings, introduced in Chapter 2, are in general a wide variety of highly non-classical states, each with their own stabilization and error-correction frameworks. Mapping each of these frameworks to the above control techniques, in a way that properly takes advantage of their error resilience, is a daunting task. However, there exists a way to capture all bosonic error correction under a single unified framework. This framework is autonomous and can be practically translated to a sequence of repeated oscillator-ancilla unitaries, interspersed with ancilla resets, in a way that works for arbitrary bosonic codes. This is then directly amenable to our control techniques and architecture, where we construct these unitaries out of ancilla-error tolerant pulse sequences, and completely decouple the ancilla and oscillator during each reset, ensuring no back-action. This framework, and its possible implementation using the architecture introduced in this thesis, is described in detail in Appendix A.

Chapter 6

Conclusions and Outlook

So what did this thesis achieve, and what could the ideas introduced here be useful for? We address this question by enumerating a list of problems, and describing the degree to which these problems are now solved or remain unanswered. The first problem is non-scientific, but perhaps the most important of all – that of the pedagogy of future students in the field of circuit-QED. This thesis, especially its first two chapters, was written with a very broad scope in mind, ideally providing a textbook-like reference for any new graduate students reading through a copy with only a preliminary background in quantum information. This includes a general view of what modern quantum experiments allow and are useful for, and their salient roadblocks. It also includes a simple introduction to continuous variable quantum information in the circuit-QED context, which, as a younger graduate student, I had to painstakingly learn from an assortment of high-level references. Finally, it also contains a brief but comprehensive introduction to parametric interactions in circuit-QED, a form of quantum control that will undoubtedly play a dominant role in near-term quantum experiments and architectures. This thesis is by no means a fully comprehensive text, but hopefully provides references to more subject-specific works (other theses, papers proposing original ideas, and review articles), such that the interested reader can find further information whenever required.

The first major scientific problem this thesis aims to answer is – how does one understand and control a general driven Josephson circuit? This problem is of course very broad, and in general requires finding the solution to time-dependent nonlinear distributed open quantum systems, which is about as complex as quantum mechanical theory can get. However, we are helped by a few assumptions: the systems we consider are primarily harmonic with perturbative nonlinearity, their physics can mostly be confined to the dynamics of a few quantum modes, their nonlinear parts are geometrically confined, and one can gain useful information by just considering simple time-periodic drives. Within these restrictions, the work in this thesis (and in the actively developed manuscript [150]) provides a helpful construction to understand, simulate and optimize driven circuits. In particular, it describes the extraction of the relevant nonlinear Hamiltonian, drive parameters, and noise from a general distributed circuit. It then shows how to utilize Floquet and Floquet-Markov theory to extract useful information about the system in the presence of the drives, from the strength of desired processes, to the available drive space where no parasitic processes occur, to the effects of noise dressed by the drives and a structured environment. In particular, we showed that one way to bring the numerous driven processes in a general circuit under our control is to introduce symmetries in the driven Hamiltonian that only allow a certain fraction of the processes to occur, while making sure the remaining processes don't significantly shift with drive strength. These ideas gave birth to the LINC, a balanced quasi-linear driven element that can be arrayed to approach an ideal three-wave mixer. The LINC will likely have many applications from clean bosonic control to better quantum amplifiers. The general protocol for analyzing and optimizing these circuits has much broader applications, including the analysis of driven gates, measurement schemes, and general design optimization for any circuit-QED experiment. Generalizing such analyses to less harmonic elements like the fluxonium, more non-trivial drives like pulsed gates, and more distributed sources of nonlinearity in other quantum

platforms, still remain open questions.

Cleaning up the above driven processes provides a possible path towards solving another general problem – that of strong nonlinear control of continuous-variable (bosonic) quantum information. The trade-offs one faces when controlling such systems forms the central theme of this thesis. This trade-off is multi-faceted; it includes the fact that ideal bosonic systems are linear but control necessarily requires nonlinearity, and also that any nonlinear elements we introduce will inevitably incur extra errors that can propagate to the bosonic modes. The architecture presented in this thesis provides one possible path towards resolving this problem. At its core, it is based on the idea that driven nonlinear elements like the LINC can present themselves as linear when the drives are not active, and that non-linear light matter interactions can be turned on precisely when we desire them through parametric couplings. Demonstrating that this indeed works in a real experimental setup is the central achievement of this thesis. Importantly, while coupling to a general nonlinear environment inevitably propagates errors to the bosonic modes, the parametric coupling above can create special symmetries in the light-matter interaction that prevent such propagated errors at first order. There are still many open questions with regards to the architecture itself, like realizing an optimal arrayed coupler and frequency stack, more compact drive and bias delivery, and more scalable bosonic modes like in flipped-chip micromachined cavities or phononic cavities. The latter platform, that of quantum 'acousto'-dynamics, forms a particularly interesting application. Phonons can not just form compact quantum memories, but also interface with various other quantum platforms, like solid-state defects and optics, allowing the high fidelity control available in this architecture to be extended to quantum sensing and transduction. Additionally, one can only scale single packages so far without losing the ability to guarantee performance – eventual architectures may necessarily require modularity and long-range connections. This will require the introduction of high-fidelity quantum routing and long-range entan-

glement generation, all of which can be directly integrated into this architecture through LINCs connected to waveguides or bus resonators.

Even with purely resonant exchange interactions in the above architecture, one may be able to achieve significant gains in controlling quantum information in a dual-rail bosonic encoding. The linearity of the system and the modes automatically provides the large erasure bias required for this encoding, with any low-frequency flux noise induced dephasing suppressed through echoes within the dual-rail qubits. This bias is central to any advantage a dual-rail processor may provide over a simpler transmon-based architecture. Importantly, providing an idle erasure bias is not sufficient – one must also ensure that any preparation, gates and error detection on the dual rail qubits retain a similar bias. Of these, both preparation and single-qubit gates are simple resonant exchanges enacted directly by the linear coupler, which should not degrade bias. Two-qubit gates and error-detection form more interesting unanswered questions. An obvious route might be to utilize simultaneous dynamic-χ between a common ancilla transmon and both rails, the flexibility within this architecture will undoubtedly allow more clever protocols. It may be interesting to explore whether one could simply use the ancilla transmon's anharmonicity as a blockade while driving Raman-type transitions between neighboring rails, that result in different accumulated phases when the two rails have even vs odd parity. This would automatically result in joint-parity measurements, and two-qubit CZ-type gates.

Switching to off-resonant interactions, this thesis tries to solve the precarious problem of bridging two dominant regimes of modern quantum experiments. The first regime is that of always having weak non-linearities, such as in optics, where it is extremely difficult to demonstrate non-Gaussian control or achieve high cooperativities, but where idle coherence is extremely long. The second is the regime of having always on strong non-linearities, such as in circuit-QED, where control is fast and powerful but always comes with uncontrolled dephasing or leakage due to the additional nonlinearity. This thesis at-

tempts to bridge these two regimes by introducing dynamical light-matter hybridization that changes dispersive shifts by four orders of magnitude in less than a hundred nanoseconds. Such parametric hybridization fully generalizes Landau-Zener physics, and unlocks rich phenomena that require significant theory and experimental demonstrations to fully explore. As an example, it will be important to experimentally demonstrate the effects of counter-diabatic techniques during this dressing, verify that one can circumvent poles in the interaction without picking up unaccounted phases, and that one can indeed turn on multiple effective couplings to give physics beyond any single effective static coupling. Such dynamic dressing will also lead to numerous applications, from clever use of the bare and dressed bases, to engineered multi-mode symmetries in the dispersive interaction, to dynamically exploring non-dispersive nonlinearities and full numerically optimized control. I am most excited for this next generation of experiments that will utilize dynamic-χ as a tool, enabling demonstrations that are impossible with any current bosonic architecture.

Perhaps the most interesting problem that this architecture might help solve is the error correction and stabilization of bosonic codes using techniques that are resilient to ancilla errors and measurements. This idea is very general – such error correction necessarily requires controls compatible with the bosonic encoding, and the ability to evacuate entropy in a precise manner such that its logical quantum information is refined. Putting a parametric mixer between the oscillator and an ancillary qubit is one way to achieve such evacuation, where the mixer exactly transfers entropy from the bosonic encoding into the ancilla, then turns off any coupling to the ancilla and dumps its entropy by resetting it. The result is in an engineered dissipator on the bosonic mode that stabilizes the desired encoding as its steady state, with the strength of error correction (or the logical error rate) being synonymous with the strength of dissipation (or its relative strength to the rate of errors). In fact, even if the parametric mixer cannot directly turn on the interaction required

to evacuate this entropy, one can Trotterize the required interaction between the oscillator and ancilla and achieve it through any architecture with universal control. This idea forms a majority of my theoretical work during my PhD, but at the time of writing this thesis the project was still in progress, and thus it is only described in Appendix A. How well this method might eventually perform will depend on the extent of Trotterization, and the resilience of the resulting control Hamiltonian to ancilla errors. Both of these should benefit from an intrinsically protected continuous-variable and number-selective control architecture like the one introduced in this thesis. Demonstrating the above qubit-engineered dissipation on a successful LINC architecture could thus allow the stabilization of *arbitrary* bosonic codes beyond break-even in a single experimental device, which would constitute a major milestone for the general field of error correction. The primary open question that might then remain, is whether such bosonic error-correction could be resiliently scaled to a multi-oscillator system (for eg, see [220, 221, 222]) such that one establishes a bosonic threshold theorem, and I hope to work on this problem in my academic career.

In general, a scaled version of the architecture introduced in this thesis would allow an experimental platform that simulates hybrid continuous and discrete variable quantum information, or correspondingly, bosonic and fermionic physics [22]. As an example, even the package implemented in this thesis, with direct access to high-fidelity exchange and squeezing interactions, could act as a simulator of interesting chemical dynamics, from tunneling through double-well potentials [223] in multiple modes, to simulating the vibronic spectra of a general nonlinear triatomic molecule [224]. Eventually, extending this architecture to a connected two-dimensional lattice could give rise to directly studying many-body bosonic physics that is inaccessible in systems like atomic Bose-Einstein condensates and molecular gas microscopes, due to the inherent universal control and tomography present in superconducting circuits. A primary limiter to such experiments could be the finite lifetime of superconducting circuits, which are orders of magnitude

shorter lived than atomic states. However, one could compensate for these short lifetimes through faster control, and performing tricks like error-detection. In particular, the latter could involve detecting the joint many-body photon parity of these systems and only implementing controls that preserve this joint parity – which all Gaussian controls do, sans displacement, as do any dynamic dispersive shifts. Such an error-detected analog many-body CV-DV simulators could form a frontier of physics that is inaccessible by near-term qubit-based simulators, and is exponentially difficult to simulate classically due to their inherent connectivity.

Overall, this thesis provides new bosonic controls, new ways of utilizing those controls that solve many problems inherent to bosonic experiments, and new ideas for how to protect and manipulate bosonic quantum information. I look forward to the exciting times and many experiments that result from and extend this work!

Appendix A

Protecting encoded information in light

This appendix describes ongoing work on a general framework for autonomous bosonic error correction, and copies over the salient arguments from its associated manuscript.

As briefly described in Chapter 2, encoding quantum information in light takes on numerous different forms. Each such encoding comes with separate theories and experimental techniques that describe the action of errors on the logical Hilbert space, and how to stabilize and correct against those errors. Overall, this has led to efforts towards a universal language for describing bosonic codes so far not being very successful. The primary reason for this is that the continuous-variable nature of bosonic encodings means that they often do not permit a simple stabilizer-like description, and their errors can be dominated by leakage into a potentially infinite Hilbert space. Prior work on bosonic codes have thus often focused on tailor-made theories describing the behavior of individual encodings [225, 226, 106, 227]. Some attempts at unification have been made through optimal recovery maps for each type of bosonic error-channel [228, 229], but these maps are highly non-local, and in general impractical to achieve with available experimental controls. Due to experimentally available controls, realistic bosonic error correction is often a mix of continuous stabilization and discrete gate-based error correction protocols, which are challenging to present in a combined language. However, experimental bosonic

architectures in circuit-QED frequently offer universal control of oscillator-ancilla systems, which means that one can technically construct arbitrary Unitaries on the system from simpler lower-order controls. The architecture and parametric controls introduced in this thesis provide precisely one such form universal control, in a manner that is resilient to ancilla errors. When such Unitaries are combined with measurements or resets of the ancilla, they can extract entropy from the bosonic mode in a very particular manner that exactly protects against that mode's errors. Our goal therefore is to recast the protection of any bosonic encoding into a realization of such Unitaries and resets, a technique we call 'qubitized dissipation'. This technique generalizes stroboscopic strategies that have been successfully used to stabilize the GKP code [107, 225, 108], and directly relates to theoretical models of Markovian dissipation [230]. Formulating and executing qubitized dissipation then simply requires asking the following question – for a given bosonic code that we want to protect against a given set of errors, what should these Unitaries look like?

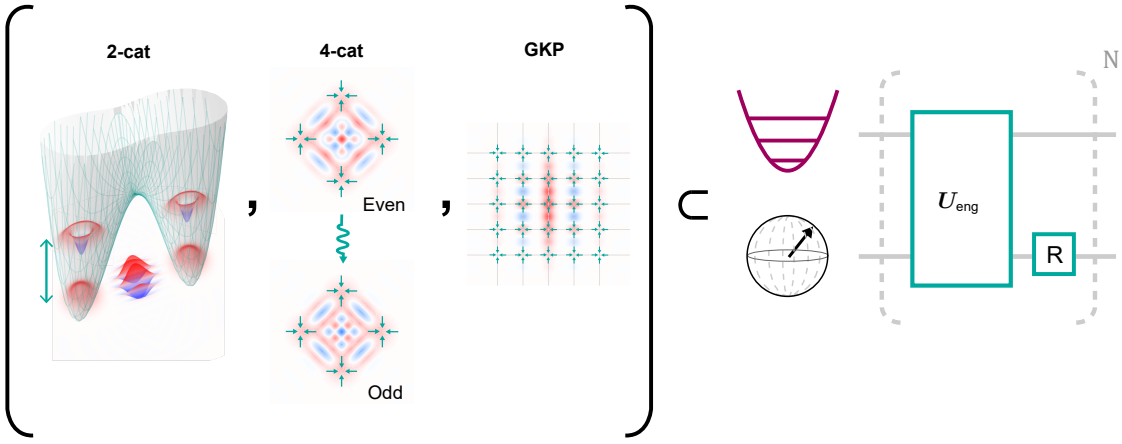

Figure A.1: **Qubitized dissipation as a universal framework** We will describe a general framework for protecting bosonic encodings using engineered Lindbladians that is universal, and captures the protection of popular bosonic codes, such as the Kerr cat (2-cat), the four-component cat (4-cat), and the Gottesman-Kitaev-Preskill (GKP) codes. For most codes, this engineered Lindbladian is not naturally available, but can always be realized in any system with universal control through periodic interactions with an ancillary qubit, followed by qubit reset.

To formulate the above Unitaries, we must first study the nature of errors on any encoded Hilbert space, and the ideal strategies one might use to combat them. Crucially, it is well known that to protect against a set of discrete errors, an encoding of quantum information must satisfy the Knill-Laflamme (KL) conditions with respect to those errors [231]. However, even though many existing systems have demonstrated such KL encodings [232, 233, 100, 102, 234, 235, 236, 237, 238, 239, 101, 108, 240] , experimental results have yet to show logical error rates that are significantly below physical error rates. The KL conditions are thus clearly not sufficient – they only mathematically guarantee some recovery map exists that corrects the system's errors. There must also be explicit conditions on the error correction protocol itself, which allows for implementation through available controls and measurements, while simultaneously not introducing uncorrectable errors into the system. These extra conditions are especially strict in real systems that interact with a Markovian environment (described by a Lindbladian evolution), where any errors on the system are independent and effectively instantaneous, and systems undergo both jump and no-jump errors.

The central result of this work is to provide a constructive framework for protecting quantum information *autonomously* against a Markovian environment, provided the encoding satisfies the KL conditions with respect to *only* the jump errors. This covers dissipative, Hamiltonian, and measurement-based strategies and generalizes previous work in [241, 242], which only proved the *existence* of autonomous dissipators for KL encodings. Our framework makes it clear that multiple error-correction strategies exist for any given encoding, and provides simple analytical predictions for the leading-order performance of each strategy at both short and long times. As an example, we show that it describes the correction of parity jumps in amplitude damping codes [243, 102], and extends both the Hamiltonian and dissipative protections of the 2-component cat code [233, 100] to arbitrary encodings. Combining these techniques with qubitized dissipation allows the

demonstration of multiple novel techniques for fully autonomous protection of any bosonic encoding, including the 4-component cat code [101] and the finite-energy GKP code [232], in standard bosonic architectures.

A.1 The nature of errors

Quantum systems are subject to a range of non-Markovian and Markovian errors, like undesired evolution due to parasitic Hamiltonian terms, low-frequency dephasing, or amplitude damping and white-noise dephasing errors. Any error with correlation times slower than the speed of control may be corrected by unitary strategies, such as dynamical decoupling and pulse-shaping [244, 218]. Hence we focus instead on errors that occur effectively instantaneously, described by Markovian error channels, which require stricter conditions on the choice of both code-space and error-correction strategies.

A Markovian error channel is well-described by a Lindbladian evolution under a set of quantum jump operators $\{E_i\}_{i=1}^{\ell}$. To be able to protect logical information (of dimension d) from this error channel, one requires a code-space $\mathcal{H}_C \subset \mathcal{H}$ that satisfies the KL conditions with respect to $\{E_i\}$. This is an important difference from previous literature [228], in that the KL conditions are not with respect to the Kraus map for a finite time evolution, nor is there any explicit separate condition required to protect the code from the no-jump back-action of $\{E_i\}$. For the context of this paper, the KL conditions simply imply that the code space $\mathcal{H}_C$ has $\bar{\ell}$ additional orthogonal copies, corresponding to independent error syndromes (see Thm 3.5 of [231]). If these errors are non-degenerate, each error syndrome corresponds to a unique correctable error, and $\bar{\ell} \geq \ell$. We call the span of these additional copies of the code space the error space, $\mathcal{H}_E$.

We provide an intuitive example through the case of encoding a single logical qubit, shown in Fig. A.2a. Here, we have multiple copies of the logical Bloch sphere, such that

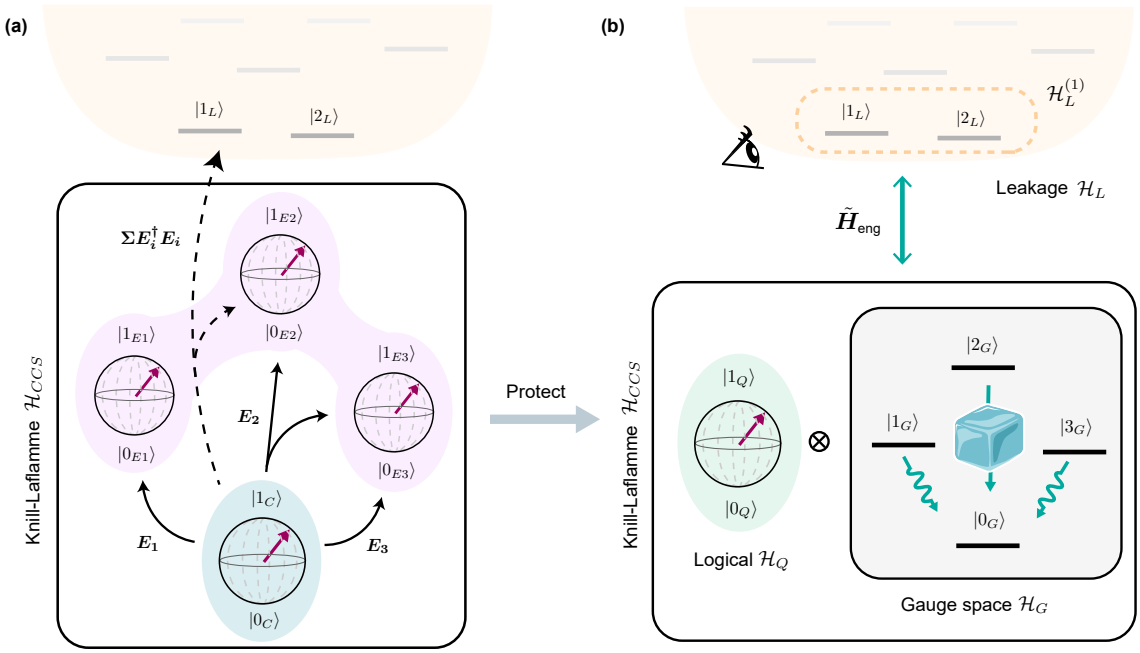

Figure A.2: **General framework for autonomous protection (a)** A logical qubit encoding that satisfies the Knill-Laflamme conditions maps the code Bloch sphere (teal) to orthogonal error Bloch spheres (pink) under environment-induced jump errors, while preserving the logical information. When there is no jump, environmental back-action may cause unrecoverable leakage (yellow). **(b)** The Knill-Laflamme conditions allow us to decompose $\mathcal{H}_{CCS}$ into a tensor product of the space containing logical information (green) and a gauge space (yellow) that indexes the copies of the Bloch spheres. To correct jump errors, we will engineer a dissipator that cools in the gauge space. Additionally, to stabilize against the no-jump backaction-induced leakage, we will engineer an effective Hamiltonian gap between $\mathcal{H}_{CCS}$ and $\mathcal{H}_L^{(1)}$, either through a Hermitian Hamiltonian, or by measuring the leakage space frequently and harnessing the quantum Zeno effect (eye).

the action of E_i always preserves the logical state vector $|\psi_C\rangle$ within each Bloch sphere. For a non-degenerate code in which every error is distinguishable, these Bloch spheres index which error has occurred (or the absence of one), and in general these indices form an additional gauge Hilbert space $\mathcal{H}_G$ of dimension $\bar{\ell}+1$. This decomposition into logical and gauge spaces (Fig. A.2b) is possible because the Knill-Laflamme conditions force E_i to only act on the gauge space (since the error is correctable):

$$E_i\left(|\psi\rangle \otimes |0\rangle\right) = |\psi\rangle \otimes \left(c_{i,0}|0_G\rangle + \sum_{m=1}^{\bar{\ell}} c_{i,m}|m_G\rangle\right), \tag{A.1}$$

where the coefficients $c_{i,m}$ are independent of the code space state $|\psi\rangle$, and the resulting state is not necessarily normalized. Thus, the action of single errors on our code-space spans a "corrupted code space" $\mathcal{H}_{CCS}$, which can be decomposed into parts containing the logical state ($\mathcal{H}_Q$), and information about which errors have occurred ($\mathcal{H}_G$):

$$\mathcal{H}_{CCS} = \mathcal{H}_C \oplus \mathcal{H}_E = \mathcal{H}_Q \otimes \mathcal{H}_G. \tag{A.2}$$

This subsystem decomposition is paramount to good error correction and fault-tolerant gates. To preserve logical information, it is sufficient to preserve the purity of just $\mathcal{H}_Q$. A mixed state in $\mathcal{H}_G$ does not destroy logical information, and a perfect decoder just extracts the information in $\mathcal{H}_Q$, tracing over $\mathcal{H}_G$. This means any operation, whether during error correction or during logical gates, is allowed to 'spoil' $\mathcal{H}_G$, which allows flexibility in engineering controls for the error-corrected system. Any error correction strategy must thus obey these broad rules: it must not act on $\mathcal{H}_Q$, and it must always evolve the system towards (have a steady state of) $\mathcal{H}_C = \mathcal{H}_Q \otimes |0_G\rangle$. We will show later that this idea can also be utilized in defining a cost function to optimize the entangling unitaries with ancilla qubits required for our proposed qubitized dissipation. Finally, note that while the subsystem decomposition in Eq. A.2 is well defined given the set $\{E_i\}$, one can generalize the problem to finding the largest possible gauge space for any given encoding $\mathcal{H}_C$. This optimal gauge space can define the full set of errors that the encoding can recover from and is an active area of research [245, 105] that we do not delve into here.

To complete the description of our Hilbert space, we can define a leakage space $\mathcal{H}_L$ as the space orthogonal to $\mathcal{H}_{CCS}$. $\mathcal{H}_L$ cannot be reached from $\mathcal{H}_C$ by any single jump error, but we will see below that it can be populated during the no-jump evolution associated with the error dissipators. In the general case, losing any population to this space leads to a permanent loss of logical information, as we have no information on what part of

the code-space the population came from. However, we will see below that the smooth and deterministic nature of the no-jump evolution still lets us prevent this leakage from happening in the first place, and therefore the code does not need to additionally obey the Knill-Laflamme conditions with respect to the no-jump evolution.

We now precisely write down the action of errors $\{E_i\}$ in this structured Hilbert space. We assume the average rate of errors to be $\propto \kappa$, with relative differences in error rates absorbed into $\{E_i\}$ for ease of notation. In the absence of any error correction, the evolution of our density matrix is given by the Lindblad equation

$$\begin{aligned}\dot{\rho} &= \kappa\left[-iH_{\text{int}}\rho + i\rho H_{\text{int}} + \sum_i \mathcal{D}[E_i](\rho)\right] \\ &= \kappa\Bigg[-iH_{\text{int}}\rho + i\rho H_{\text{int}} \\ &\qquad + \sum_i \left(E_i\rho E_i^\dagger - \frac{1}{2}E_i^\dagger E_i\rho - \frac{1}{2}\rho E_i^\dagger E_i\right)\Bigg]\end{aligned} \tag{A.3}$$

where we use the shorthand $\mathcal{D}[E](\rho) := E\rho E^\dagger - \frac{1}{2}\left(E^\dagger E\rho + \rho E^\dagger E\right)$ for the dissipator induced by the error operator E. The first two terms in Eq. A.3 are the usual Hamiltonian evolution generated by the intrinsic Hamiltonian $\mathcal{H}_{\text{int}}$, and the remaining terms are the evolution induced by the error operators $\{E_i\}$. We note that the terms involving $E_i^\dagger E_i$ can be absorbed into the Hamiltonian evolution if we define an effective non-Hermitian Hamiltonian $\tilde{H}_{\text{int}} := H_0 - \frac{i}{2}\sum_i E_i^\dagger E_i$, so that the evolution is given by

$$\dot{\rho} = \kappa\left[-i\tilde{H}_{\text{int}}\rho + i\rho\tilde{H}_{\text{int}}^\dagger + \sum_i E_i\rho E_i^\dagger\right]. \tag{A.4}$$

This splits the evolution terms into a continuous Hamiltonian-like evolution, dubbed the "no-jump" evolution, and the remaining "jump" terms $E_i\rho E_i^\dagger$. For the remainder of this work, we assume we are in a frame with $H_0 = 0$, so that the effective Hamiltonian is due

entirely to the $\{E_i\}$.

We aim to protect against the errors introduced by the Lindbladian at first order in evolution time δt. Starting in a pure code-state $\rho(0) = |\psi_C\rangle\langle\psi_C|$, we have

$$\begin{aligned}\rho(\delta t) \approx &(1 - i\kappa\delta t\tilde{H}_{\text{int}})|\psi_C\rangle\langle\psi_C|(1 + i\kappa\delta t\tilde{H}_{\text{int}}^\dagger) \\ &+ \kappa\delta t\sum_i E_i|\psi_C\rangle\langle\psi_C|E_i^\dagger.\end{aligned} \tag{A.5}$$

In a quantum trajectories picture, any state $|\psi_C\rangle \in \mathcal{H}_C$ incurs a jump E_i with probability $\kappa\delta t\langle\psi_C|E_i^\dagger E_i|\psi_C\rangle$, and in the absence of jumps it undergoes a no-jump evolution generated by $\kappa\tilde{H}_{\text{int}}$. We know that the (un-normalized) state after a jump, $E_i|\psi_C\rangle$, is in $\mathcal{H}_{CCS}$ and preserves logical information (Eq. A.1). On the other hand, the no-jump evolution causes a continuous rotation that may take a state in $\mathcal{H}_C$ to states outside of $\mathcal{H}_C$ or even outside of $\mathcal{H}_{CCS}$. To capture this leakage outside $\mathcal{H}_{CCS}$, we can define the first-order leakage space $\mathcal{H}_L^{(1)}$:

$$\mathcal{H}_L^{(1)} = \text{span}_l\left[|\psi_l\rangle\right] = \text{span}_l\left[\Pi_L\sum_i E_i^\dagger E_i\ |\psi_C^{(l)}\rangle\right], \tag{A.6}$$

for some basis $\{|\psi_C^{(l)}\rangle\}$ of the code space. Note that the size of this leakage subspace is at-most equal to the size of the code-space (d), and is independent of Kraus rank. We will see that because the no-jump evolution causes the state to always move coherently into $\mathcal{H}_\mathcal{E}$ or $\mathcal{H}_L^{(1)}$ by an amount proportional to $\kappa\delta t$, we can either exactly cancel, or prevent this evolution.

To summarize, a state in $|\psi_C\rangle \in \mathcal{H}_Q \otimes |0_G\rangle$ may occasionally jump 'up' the gauge space to $\mathcal{H}_Q \otimes \sum_m c_{i,m}|m_G\rangle$, with a probability proportional to $\kappa\delta t$. If no jump occurs, the state will evolve via a smooth Hamiltonian-like evolution into states that are not necessarily correctable (they may be in $\mathcal{H}_L$, or in $\mathcal{H}_E$ with the wrong logical information),

but the uncorrectable part will have norm $\sim \kappa\delta t$. Overall, a necessary condition for being able to protect against the Lindbladian evolution generated by the errors $\{\sqrt{\kappa}E_i\}$ is being able both to correct the effect of the jump evolution within $\mathcal{H}_G$ and stabilize against the no-jump evolution to $\mathcal{H}_L^{(1)}$. In the next two sections, we will show that the jump errors can be corrected by engineering additional jump operators that enact a cooling in the gauge space, and the no-jump errors can be prevented by engineering a Hamiltonian or dissipative (Zeno) gap that separates $\mathcal{H}_L^{(1)}$ from $\mathcal{H}_{CCS}$. Overall, the tools we will use are some combination of an engineered Hamiltonian H_{eng} and an additional set of engineered dissipators $\{\Gamma\mathcal{D}[F_i]\}_{i=1}^k$, that enact both the jump correction and the non-jump stabilization. We will treat these Hamiltonians and dissipators in a continuous manner for now, and later discretize them into 'qubitized' Lindbladians.

A.2 Correcting jump errors

We saw in the previous section that jump errors on Knill-Laflamme codes only act within the gauge space, and we will show below that their correction can also be entirely constructed within this gauge space. Starting in $\mathcal{H}_C$, jump errors take a state $|\psi_Q\rangle \otimes |0_G\rangle$ to a higher state $|\psi_Q\rangle \otimes |\psi_G\rangle$, acting as an effective "heating" of the gauge space without corrupting the logical information. However, if any population remains in the higher states of $\mathcal{H}_G$, a subsequent jump error can corrupt the logical information. We thus need our engineered evolution to "cool" the gauge space back to $|0_G\rangle$, before a second jump error occurs. Importantly, we note that it is not possible to cool the gauge space with an engineered Hermitian Hamiltonian, as a unitary evolution cannot send $|\psi_G\rangle \to |0_G\rangle$ without also acting on $|0_G\rangle$ and introducing some new error. We therefore require the gauge space heating to be corrected by engineered dissipators that break this reciprocity and allow all states in the gauge space to be cooled to the ground state. This amounts to the evacuation

of entropy from the system in a very specific manner, enforced by the structure of the gauge space.

Following this intuition, we outline below a set of necessary and sufficient conditions for a set of engineered jump operators $\{F_i\}$ to correct errors at first order, without inducing any additional errors even when $\Gamma/\kappa \to \infty$. This latter restriction is necessary for an error-correction scheme that always performs better with faster correction, or with slower physical error rates, but it is not in general necessary to achieve error-correction beyond break-even. The framing of the conditions below are purposely abstract for mathematical compactness, but we will explain them in detail and provide intuition through examples in the rest of this section.

C1 For the engineered jumps to not introduce additional errors, it is necessary that the jump operators preserve the tensored structure of $\mathcal{H}_{CCS}$ and have zero action on $\mathcal{H}_C$. Explicitly,

$$\begin{aligned} F_i\left(|\psi_Q\rangle \otimes |\psi_G\rangle\right) &= \mathbb{1}_Q|\psi_Q\rangle \otimes F_{i,G}|\psi_G\rangle \\ F_{i,G}|0_G\rangle &= 0 \end{aligned} \tag{A.7}$$

for some operator $F_{i,G}$ that acts only on the gauge degree of freedom [1].

C2 For the engineered effective Hamiltonian $\tilde{H}_{\text{eng}} = H_{\text{eng}} - \frac{i}{2}\sum_i F_i^\dagger F_i$ to not introduce additional errors, it is necessary that the effective Hamiltonian does not take $\mathcal{H}_C$ to $\mathcal{H}_E \oplus \mathcal{H}_L$, does not cause logical errors within $\mathcal{H}_{CCS}$, and has no component that takes $\mathcal{H}_{CCS}$ to $\mathcal{H}_L$. Condition **C1** already ensures $F_i^\dagger F_i$ does not act on $\mathcal{H}_C$ or cause logical errors, so we need to additionally restrict H_{eng} to only have action on the

[1] Note that we could in principle have $F_i^G|0_G\rangle \propto |0_G\rangle$. However, we are always free to add a constant to the jump operators (at the cost of modifying the Hamiltonian, see e.g. Eqs. 1.7 and 1.8 in [246]), so we fix this freedom with the convention $F_{i,G}|0_G\rangle = 0$.

gauge space within $\mathcal{H}_{CCS}$ and have zero action on $\mathcal{H}_C$ [2]:

$$\begin{aligned} H_{\text{eng}}\,\Pi_{CCS} &= \mathbb{1}_Q \otimes H_{\text{eng},G} \\ H_{\text{eng},G}|0_G\rangle &= 0. \end{aligned} \tag{A.8}$$

Additionally, we must ensure that the engineered effective Hamiltonian does not cause leakage:

$$\Pi_L \tilde{H}_{\text{eng}} \Pi_{CCS} = 0. \tag{A.9}$$

C3 Finally, it is necessary that the net action of the engineered jump operators eventually take us back to the code space. Given the form of $\{F_i\}$ and H_{eng} on $\mathcal{H}_{CCS}$ in Condition **C2**, we can consider the total engineered Lindbladian restricted to the gauge space $\mathcal{H}_G$, $\mathcal{L}_G$. To correct errors, it is necessary that $\mathcal{L}_G$ has a single zero eigenvalue (corresponding to the code space), with all other eigenvalues $\Gamma\lambda_{n,G}$ satisfying $\text{Re}(\lambda_{n,G}) \leq -\Delta_{\text{corr}}$ for some positive definite dissipative gap Δ_{corr}.

In summary, the engineered Lindbladian generated by $\{F_i\}$ must preserve $\mathcal{H}_{CCS}$, and its action on $\mathcal{H}_{CCS}$ must only be an effective cooling in the gauge space $\mathcal{H}_G$. Given this structure, the dynamics of errors and their correction can be entirely analyzed in $\mathcal{H}_G$, balancing the dynamics between the heating caused by $\{E_i\}$ and the engineered cooling. The last condition is intentionally broad – there is no requirement for each excited gauge state to be directly mapped back to the code space. This means multiple non-equivalent strategies exist to correct any code that satisfies the Knill-Laflamme conditions with respect to more than one type of error (see Fig. A.3a for illustrative examples). All such strategies enact at least first-order error-correction on the code, and have the same scaling of the rate of decay of logical information with respect to the relative strength of correction Γ/κ. In

[2] Just like for $F_{i,G}$ in Condition **C1**, in principle we could have $H_{\text{eng},G}|0_G\rangle \propto |0_G\rangle$, but since we can always add a constant to the Hamiltonian without changing the dynamics we take the convention that $H_{\text{eng},G}|0_G\rangle = 0$.

the limit of $\Gamma/\kappa \to \infty$, the engineered jump operators dominate, and the probability of two errors occurring before the state is corrected goes to zero, resulting in perfect error correction.

For any finite Γ/κ, the system's dynamics can be analyzed as a Markov chain problem within the gauge space. Regardless of the strength of correction, the system always first reaches a stationary solution at rate $\propto \kappa$. This can be understood intuitively as no correction occurring when the error population is too low, since correction doesn't act on the code space. The stationary solution can be found by setting the net rate of change of population of every gauge state to zero, at first order in κ/Γ. With populations in the gauge space denoted by $\{\pi_i\}$, and jump probabilities between states denoted $\{P_{ij}\}$, we have:

$$\dot{\pi}_i = \sum_j P_{ji}\pi_j = 0 \qquad \forall\, i \leq \bar{l} + 1.$$

In general, we have $P_{0i} \propto \kappa$ and all other jump probabilities are either 0 or $\propto \Gamma$. This means that the stationary state solution has an error population $\propto \kappa/\Gamma$, and it reaches this solution in time $\propto 1/\Gamma$. This initial equilibration always occurs for any error-correction strategy and could appear as a strict lower bound on the encoding fidelity into any error-corrected memory. However, we note that (to first order) this equilibration is entirely within the gauge space and hence preserves the logical information in $\mathcal{H}_Q$, which means that an ideal decoder that traces over the gauge space will not lose any information. Once this initial equilibrium has been reached, second-order effects kick in, and we lose logical information through subsequent jumps. These jumps occur at a rate $\propto \kappa$ from error states that have population $\propto \kappa/\Gamma$, resulting in a net loss of logical information at a rate κ^2/Γ, as expected from any first-order error correction procedure.

A.2.1 Correcting photon jumps in the CLY code

We now build intuition for these conditions using an example of a bosonic code. We choose the odd-parity Chuang-Leung-Yamamoto (CLY) code [243], under single-photon loss:

$$\begin{aligned} |0_C\rangle &= \left(|1_a, 5_b\rangle + |5_a, 1_b\rangle\right)/\sqrt{2} \\ |1_C\rangle &= |3_a, 3_b\rangle. \end{aligned} \tag{A.10}$$

This simple encoding of two oscillators (denoted a and b) balances the population in either oscillator, such that when the decay rates of the oscillators are equal, this code faces no intrinsic no-jump evolution. This allows us to independently study just the jump errors under various error-correction protocols. The reason for specifically choosing an odd-parity encoding will become clear later. The full Lindbladian on these two oscillators, in a rotating frame where each oscillator is stationary, is given by:

$$\begin{aligned} \mathcal{L}_{ab} &= \kappa\left(\mathcal{D}[\hat{a}] + \mathcal{D}[\hat{b}]\right) + \Gamma\left(\mathcal{D}[F_1] + \mathcal{D}[F_2]\right) \\ F_i &= \mathbb{1}^Q \otimes F_i^G, \end{aligned} \tag{A.11}$$

where $\hat{a}, \hat{b}$ are the annihilation operators on either cavity, and F_i^G are the engineered jump operators on the gauge space that will error-correct this code.

To engineer the appropriate F_i^Gs, we consider the gauge space structure of this code. Photon loss in either cavity changes their respective photon parity, while still preserving relative superpositions, as shown in Fig. A.3a. The gauge space $\mathcal{H}_G$ therefore has dimension 3, labeling the $|0_G\rangle$ (no error), $|1_G\rangle$ (photon loss in a) and $|2_G\rangle$ (photon loss in b) as orthogonal subspaces of $\mathcal{H}_{CCS}$ with different parities. The engineered dissipation must cool any gauge space heating that takes $|0_G\rangle$ to $|1_G\rangle, |2_G\rangle$, since any subsequent photon

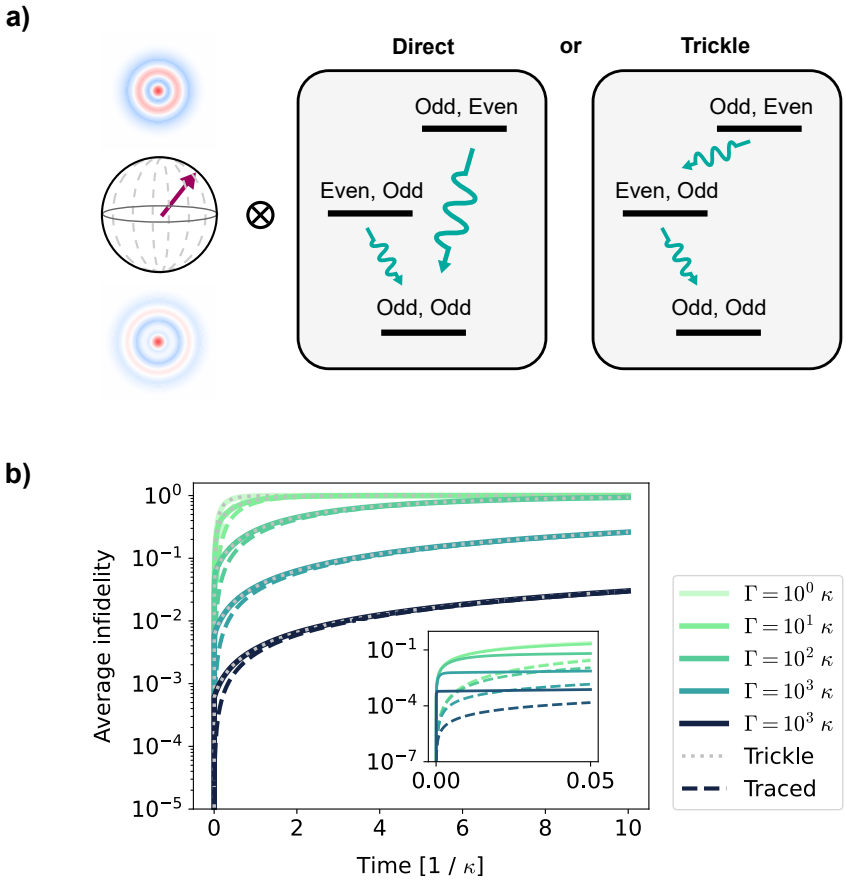

Figure A.3: **Illustrating correction strategies on the CLY code (a)** The logical Bloch sphere, with single-mode Wigner functions shown for the gauge space being in the ground state. The gauge space structure allows at least two non-equivalent error correction strategies (Direct or Trickle). **(b)** Average state fidelity, calculated as the fidelity to the initial state as a function of time, averaged over starting at all six cardinal points on the logical Bloch sphere, for a few values of Γ/κ (solid lines). The decay rates at long times decrease $\propto (\Gamma/\kappa)^{-1}$. The initial drop in fidelity from equilibration in the gauge space always occurs at rate $\propto \kappa$, and disappears once the fidelity is calculated by tracing over the gauge space (dashed lines, see inset for a zoomed-in version).

loss events will spoil logical information. We thus test the following two sets of engineered dissipators:

$$\left\{ \begin{aligned} F_1^G &= |0_G\rangle\langle 1_G| \\ F_2^G &= |0_G\rangle\langle 2_G| \end{aligned} \right\}, \text{ or} \tag{A.12}$$

$$F_3^G = \sqrt{\gamma_{rel}}\, |0_G\rangle\langle 1_G| + |1_G\rangle\langle 2_G|. \tag{A.13}$$

We refer to the former as the "Direct" strategy, mapping each error state directly to its respective code state. The latter is labeled the "Trickle" strategy, where errors in cavity b are converted into errors in cavity a, and are corrected in a trickle-down manner. Note that this is qualitatively different than the trickle-down dissipation enacted in [108], since a photon loss in b occurs at the same rate as one in a, and is not a higher-order error process.

Somewhat counter-intuitively, these two strategies work exactly the same for a finite value of $\gamma_{rel} = 2$, irrespective of Γ/κ (see Fig. A.3c). This finite value can be derived through a simple analytical calculation for the Markov-chain, by equating the stationary state populations in $|\psi_G^{(1,2)}\rangle$ for the two strategies. We see that the initial fidelity drop (from equilibration) and long-time decay rate for these strategies always get better with Γ/κ, as expected from a "true" error correcting strategy (Fig. A.3c). Importantly, note that the Direct strategy always requires as many independent dissipators as there are independent errors. This is because a single combined dissipator (like $F_{1+2}^G = F_1^G + F_2^G$) will inevitably have a dark state ($|1_G\rangle - |2_G\rangle$) under the engineered jumps, which the engineered no-jump back-action will bias the state towards at rate $\propto \Gamma$ while in the error-space. The Trickle strategy bypasses this constraint by allowing a single dissipator to act on $|1_G\rangle$ and $|2_G\rangle$, which in this example performs equivalently to a sum of two separate dissipators.

Lower-order dissipators - While both the Direct and Trickle strategies work well, they are high-order interactions that include the exchange of up to five photons. Actual experiments with similar parity-based encodings often engineer lower-order dissipators that are imperfect, but easier to implement. Since we are mapping even parity spaces to odd parity spaces, let us consider the family of dissipators that are of the form:

$$\left\{ \begin{aligned} F_{1P} &= \sum_{n_a \in \text{odd}} \gamma_n |n_a\rangle\langle n_a - 1| \otimes \mathbb{1}_b, \\ F_{2P} &= \mathbb{1}_a \otimes \sum_{n_b \in \text{odd}} \gamma_n |n_b\rangle\langle n_b - 1|, \end{aligned} \right\} \tag{A.14}$$

where $|n_{a(b)}\rangle$ is the Fock state n in cavity $a(b)$. We note that such a dissipator would not have worked if our encoding had any support on vacuum, which would be annihilated and unrecoverable through this simple mapping, and this was our reason for choosing the odd-parity CLY code as an example.

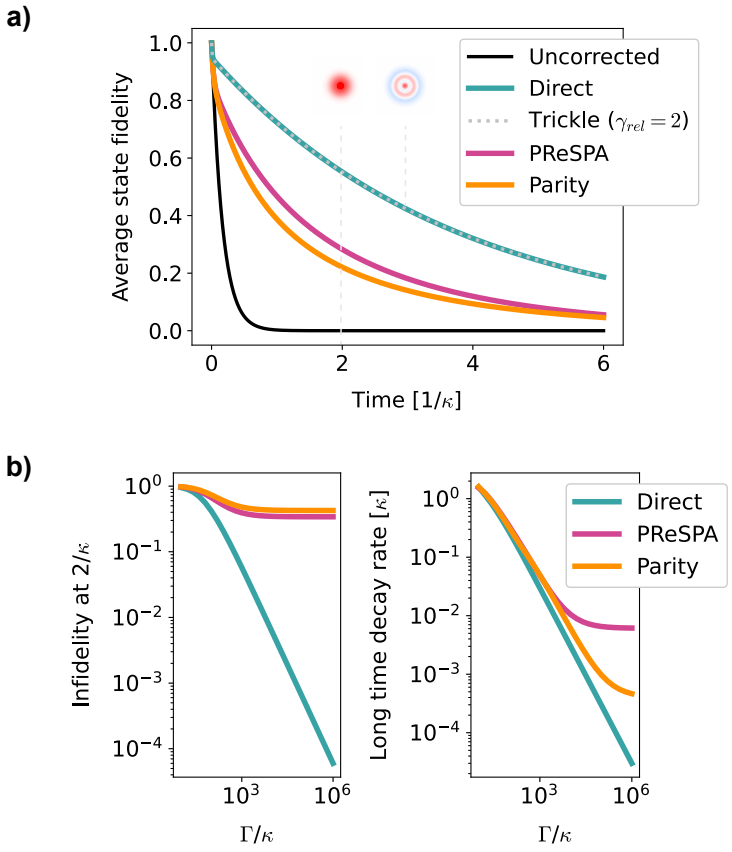

Figure A.4: **Comparing low-order correction strategies (a)** A comparison between various error correction strategies, with $\Gamma/\kappa = 100$. PReSPA and Parity refer to the engineered dissipators in Eq. A.14 with $\gamma_n = 1/\sqrt{\bar{n}}$ and $\gamma_n = 1/\sqrt{n}$ respectively. All strategies perform significantly beyond break-even. Single-mode Wigner functions are shown for the uncorrected evolution (time $\sim 2/\kappa$) and the Direct correction strategy (time $\sim 3/\kappa$) respectively, when starting in $|1_C\rangle$. **(b)** True error correcting strategies (Direct, Trickle performs identically and is not shown) are compared against imperfect strategies (PReSPA, Parity) as a function of Γ/κ. The imperfect strategies saturate in fidelity (left, plotted at time $2/\kappa$), while true correction shows persistent improvement at the expected rate $\propto \kappa^2/\Gamma$. The long-time decay rate (right, from an exponential fit near time $\sim 10/\kappa$) of the imperfect strategies are close to the true correction up to larger values of Γ/κ, but also eventually saturate.

Consider the specific case of $\gamma_n = 1$, which is nearly identical to the dissipator engineered in [247], generalized to two modes. We immediately note that this dissipator contains no information about the type of encoding it is correcting, except that 2-parity forms a good error syndrome for that code. This means that studying its behavior on the code in Eq. A.10 is equivalent to understanding its behavior on any 2-parity-based encoding. Unfortunately, this dissipation strategy does not restrict its action to just the gauge space, violating Condition **C1**, and therefore fails at large dissipation strength. We can understand this by noticing that the photon loss error weights each Fock state transition matrix differently ($\hat{a} = \sum_n \sqrt{n}|n-1\rangle\langle n|$), and any corrective dissipators that do not undo these weights will spoil logical information by entangling $\mathcal{H}_Q$ and $\mathcal{H}_G$, or cause leakage. This is numerically demonstrated in Fig. A.4b, where beyond a certain rate of Γ/κ, the engineered dissipation introduces as many errors as it corrects and therefore the error-corrected state fidelity saturates. The simulated dissipator has been normalized ($\gamma_n = 1/\sqrt{\bar{n}}$) to match the logical error rate on $|1_C\rangle$ to that of the Direct and Trickle strategies, for the sake of an even comparison.

Alternatively, we could try a different strategy that undoes these weights, where

$$\gamma_n = 1/\sqrt{n}.$$

These dissipators indeed only act on the gauge space within $\mathcal{H}_{CCS}$. However, they cause transitions from states in $\mathcal{H}_{CCS}$ to states in $\mathcal{H}_L$, violating Condition **C2** and breaking the required block-diagonal form. This is most evident in the engineered dissipation's no-jump evolution, given by:

$$\tilde{H}_{\text{eng}} = \frac{i\Gamma}{2} \sum_{r=a,b} \sum_{n_r \in \text{even}} \frac{1}{n_r + 1} |n_r\rangle\langle n_r|,$$

which causes leakage at a rate $\propto \Gamma$ from any error state $|\psi_G^{(1,2)}\rangle$. For $\Gamma \gg \kappa$, this means

that logical information is immediately lost as soon as a single error occurs, causing the error correction to fail (again, see saturation in Fig. A.4b). In general, numeric sweeps over various realizations of γ_n reveal that there is *no low-order strategy that performs true error-correction*, suggesting that our conditions for jump correction may be necessary, not just sufficient.

A.3 Stabilizing no-jump evolution

Unlike the jump errors, which obey the Knill-Laflamme conditions, the intrinsic no-jump evolution of the error dissipators might not preserve logical information. However, this no-jump evolution still possesses two important properties that let us combat it. First, the no-jump evolution has a special structure guaranteed by the Knill-Laflamme conditions – when acting on the code space, it must either do nothing, or it must take it to an orthogonal detectable space. It cannot cause mixing within the code space. Explicitly,

$$\Pi_C \sum_l E_l^\dagger E_l \Pi_C \propto \Pi_C, \tag{A.15}$$

where Π_C is the projector onto the code space $\mathcal{H}_C$. Note that because this condition is a subset of the Knill-Laflamme conditions, there may exist encodings that satisfy this condition without being fully Knill-Laflamme. Second, the no-jump evolution is a coherent (smooth) evolution that preserves the purity of the initial state. As a result, even though it is not possible to correct the leakage due to the no-jump evolution once it has occurred, it is possible to prevent it from building up in the first place.

Because of these two properties, the no-jump evolution coherently maps population from $\mathcal{H}_C$ to $\mathcal{H}_E \oplus \mathcal{H}_L$ at rate $\propto (\kappa\delta t)^2$. We wish to prevent this probability build-up in $\mathcal{H}_E \oplus \mathcal{H}_L$, since once it has occurred, we assume it is uncorrectable. In continuous-variable

codes, this probability is often more benign than jump errors because the no-jump operator often mostly preserves the code space. However, if it is not stabilized, information will eventually leak to uncorrectable portions of the oscillators' infinite Hilbert space. Let us first consider the probability build-up in $\mathcal{H}_E$ due to the no-jump evolution, which unlike the jump errors need not respect the separation of $\mathcal{H}_Q$ and $\mathcal{H}_G$, and can cause logical errors. To satisfy Condition **C3** for correcting jump errors, we already require a set of dissipators $\{\Gamma\mathcal{D}[F_i]\}$ that effectively measure the population of $\mathcal{H}_E$ every $\delta t \sim 1/\Gamma\Delta_{\text{corr}}$. If no jump errors are found, these measurements project the state back into $\mathcal{H}_C$ with probability $1-(\kappa\delta t)^2 = 1-(\kappa/\Gamma\Delta_{\text{corr}})^2$, since the probability of the no-jump evolution causing transitions from $\mathcal{H}_C$ to $\mathcal{H}_E$ is $O\left((\kappa\delta t)^2\right)$. In the limit $\Gamma/\kappa \to \infty$ the net evolution from $\mathcal{H}_C$ to $\mathcal{H}_E$ approaches zero, which is exactly the quantum Zeno effect. Equivalently, the engineered no-jump evolution due to the rapid measurements decays away any coherence between $\mathcal{H}_C$ and $\mathcal{H}_E$, preventing the build-up of probability in $\mathcal{H}_E$. Thus, the engineered dissipators $\{F_i\}$ are already guaranteed to stabilize against any intrinsic no-jump evolution to $\mathcal{H}_E$.

This leaves us with the probability build-up in $\mathcal{H}_L$, which constitutes leakage. The prevention of this leakage constitutes the "stabilization" aspect of any bosonic encoding, which restricts the dynamics to the finite-dimensional subspace $\mathcal{H}_{CCS}$ within the infinite-dimensional oscillator Hilbert space. Note that stabilization here implies that the code space, once prepared, remains within $\mathcal{H}_{CCS}$, but a subset of stabilization protocols can also prepare the code space from target initial states (like Fock states). Since we are only concerned with protection at first order, we can restrict our concerns to the subspace that the code space can leak to at first order, $\mathcal{H}_L^{(1)}$. We want to prevent any population build-up in this subspace. As we show below, this can be achieved by one of two strategies.

The first strategy creates a large effective Hamiltonian gap between the code space and $\mathcal{H}_L^{(1)}$, utilizing either the (dissipative) Zeno effect or a large (Hamiltonian) detuning to

stabilize against leakage:

C4 Each eigenvector of $\tilde{H}_{\text{eng}}$ that has any overlap with $\mathcal{H}_L^{(1)}$ must have eigenvalue $|\lambda_L^{(n)}| \geq \Delta_{\text{stab}}$, for some non-zero dimensionless gap Δ_{stab} that sets a scale for the strength of protection.

Note that the eigenvalues $\lambda_L^{(n)}$ can be complex, as long as their magnitude is large. We will call these types of strategies "strong" stabilization strategies, as they will always get better as the gap ($\Gamma\Delta_{\text{stab}}$) increases.

The second strategy uses a weak but fine-tuned effective Hamiltonian ($\kappa\tilde{H}_\kappa$) to exactly counter the evolution out of the code space:

C5 We add a combination of weak dissipators $\mathcal{D}[\sqrt{\kappa}F_\kappa]$ and a weak Hamiltonian κH_κ to the $\mathcal{L}_{\text{tot}}$ to ensure the effective engineered Hamiltonian has exactly the opposite action between the code and leakage space as the intrinsic no-jump evolution:

$$\Pi_L \tilde{H}_\kappa \Pi_C = -\Pi_L \tilde{H}_{\text{int}} \Pi_C. \tag{A.16}$$

In order for the jumps to not cause logical errors or leakage, the F_κ must obey

$$F_\kappa \left(|\psi_Q\rangle \otimes |0_G\rangle \right) = |\psi_Q\rangle \otimes F_{\kappa,G} |0_G\rangle \tag{A.17}$$

for some $F_{\kappa,G}$. In addition, we must place a restriction on H_κ to not cause additional logical errors when acting on the code space:

$$\Pi_C H_\kappa \Pi_C \propto \Pi_C. \tag{A.18}$$

Note that H_κ is allowed to otherwise mix $\mathcal{H}_C$ and $\mathcal{H}_E$, since this evolution will be suppressed by the strong engineered dissipators that are already required to correct

the jump errors. Similarly, F_κ is allowed to cause logical errors or leakage when acting on states in $\mathcal{H}_E$.

We will call these strategies "fine-tuned" strategies, as they require matching the engineered evolution to the intrinsic evolution.

While we have presented these conditions in an abstract form with full generality, we will provide more constructive forms of these conditions in the sections below and provide more intuition for them. Note that both conditions can be achieved with an engineered Hermitian Hamiltonian, engineered dissipator, or combination of the two (similar to, e.g., [248]).

To illustrate each condition, we will use the example of the stabilization of the two-component cat code (2-cat), which has been extensively researched in literature. Following standard convention, the 2-cat code space is defined as:

$$\begin{aligned} |0_C\rangle &= |C_\alpha^+\rangle + |C_\alpha^-\rangle \approx |\alpha\rangle \\ |1_C\rangle &= |C_\alpha^+\rangle - |C_\alpha^-\rangle \approx |-\alpha\rangle \\ |C_\alpha^\pm\rangle &= \frac{1}{\mathcal{N}^\pm}(|\alpha\rangle \pm |-\alpha\rangle), \end{aligned} \tag{A.19}$$

where $|\pm\alpha\rangle$ are coherent states and $\mathcal{N}^\pm = \sqrt{2(1 \pm e^{-2|\alpha|^2})}$ are normalization factors. For large α, photon loss jumps enact strongly biased errors on this encoding – the logical Z states are exponentially close to coherent states and are (nearly) impervious to the jump errors, while the logical X states $|\pm_C\rangle$ are interchanged every time a jump error occurs since the jump changes the photon number parity. Thus, even though this encoding is not Knill-Laflamme to photon loss, focusing our attention on the logical Z axis lets us ignore the jump errors and study the no-jump errors in isolation. For ease of derivation in the rest of the text, we will approximate the logical codewords as $|\pm\alpha\rangle$, but all numerical simulations in Fig. A.5 use the true logical codewords.

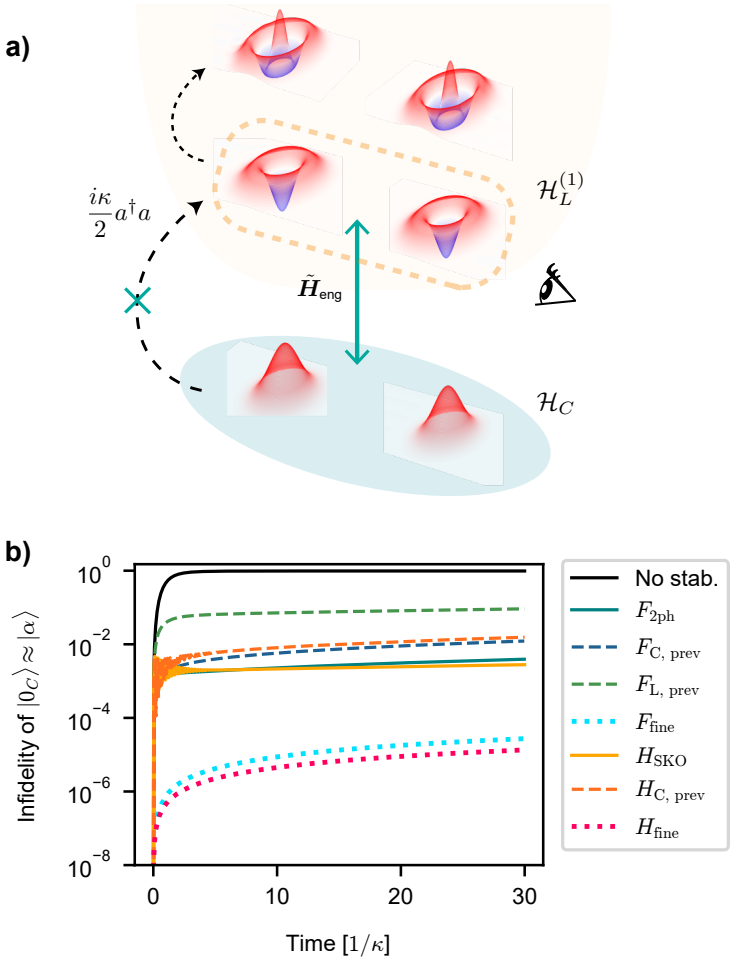

Figure A.5: **Seven ways to stabilize the 2-cat. a)** The no-jump backaction due to photon loss takes the 2-cat code states ($\sim |\pm\alpha\rangle$) to higher displaced Fock states. We present three broad strategies for stabilizating against no-jump evolution through the quantum Zeno effect (eye), an engineered Hamiltonian gap ($\tilde{H}_{\text{eng}}$), or a fine-tuned cancellation (blue cross), subdivided into seven stabilization methods. **b)** Time evolution of the Logical $|0_C\rangle$ state under seven different stabilization methods, with $\Gamma/\kappa = 60$ and $\alpha = 2$. Two of these are the well-known dissipative cat (dark green, F_{2ph}) and Kerr-cat (yellow, H_{SKO}) respectively, and the black solid line represents the state's free evolution. The other five strategies are novel and arise naturally from this paper's framework. Fine-tuned strategies (dotted) outperform the strong engineered strategies (dashed, solid) significantly.

Intuitively, the action of intrinsic no-jump evolution on the cat basis is to 'shrink' the cat basis ($\alpha \to \alpha e^{-\frac{\kappa t}{2}}$). However, for small evolution times, this shrunk cat has large overlap with the original code states. This is easy to see in when analyzing the no-jump

evolution on the logical states to first order:

$$\begin{aligned}\left(\mathbb{1} - \kappa\delta t a^\dagger a\right)|\pm\alpha\rangle \approx \left(1 - \kappa\delta t|\alpha|^2\right)|\pm\alpha\rangle \\ \mp \alpha\kappa\delta t|\pm\alpha^{(1)}\rangle,\end{aligned} \tag{A.20}$$

where $|\pm\alpha^{(1)}\rangle$ is the 'displaced Fock 1' state $D_{\pm\alpha}|1\rangle$.

While the displaced Fock states $|\alpha^{(n)}\rangle$ form the eigenbasis of a displaced harmonic well and are often used to analyze the Kerr-cat [249], they emerge naturally from the action of intrinsic photon loss on this encoding, without any assumptions on the underlying stabilization scheme (Fig. A.5a). The above equation makes it clear that the 'shrunk' cat primarily decreases the fidelity through leakage to an orthogonal space $\mathcal{H}_L^{(1)} = \text{span}\{|\alpha^{(1)}\rangle, |-\alpha^{(1)}\rangle\}$, which satisfies condition eq. (A.15), letting us apply various strategies that stabilize against this leakage. Such a stabilization will allow the (exponential) protection of $\langle Z_L\rangle$, while still being susceptible to X_L errors, thus providing a noise-biased qubit.

Note that because $|\pm\alpha\rangle$ leak to approximately orthogonal states, the 2-cat happens to be approximately Knill-Laflamme (for large $|\alpha|$) with respect to the no-jump evolution, which means that some of the leakage is recoverable even after it has been populated. This will allow some strategies to improve their performance if they send $|\pm\alpha^{(1)}\rangle \to |\pm\alpha\rangle$, which is beyond the protection one would expect from an arbitrary encoding that is only Knill-Laflamme to the jump errors. If instead of pure single-photon loss, the oscillator incurred other intrinsic errors like heating and dephasing, these displaced Fock states could form a gauge space and the above dissipation would then help with error-correction too [250].

Popular methods to stabilize the cat manifold use either a squeezed-Kerr Hamiltonian [34, 251, 252] or two-photon dissipation [100, 253, 133], such that $|\pm\alpha\rangle$ form a

degenerate ground state manifold. We show below that both of these methods can be understood as 'gapping away' the leakage space, obeying Condition **C4**, and we generalize these strategies to the stabilization of arbitrary bosonic codes. We also show that several non-equivalent dissipators and Hamiltonians exist that achieve the same function, but they have a strict hierarchy in their performance (e.g., the Kerr-cat always performs better than the dissipative cat for the same effective Hamiltonian gap). Finally, we introduce a new fine-tuned stabilization technique based on Condition **C5** that uses fine-tuned engineered Hamiltonians or dissipators to exactly cancel the leakage, which is a continuous generalization of re-pumping the cat at finite intervals [233, 100] or periodic unitary corrections in the binomial code [102, 254, 240]. The fine-tuned stabilization technique always performs better than the 'strong' Hamiltonians and dissipators, and for the 2-cat, saturates the bound on how well the code can be stabilized for finite α.

Strong dissipative stabilization

One approach to prevent probability evolution into the leakage space is to ensure the engineered dissipators $\{F_i\}$ induce the quantum Zeno effect on $\mathcal{H}_L$, just as they already do for $\mathcal{H}_E$. An example of this was introduced in [241], using dissipators of the form

$$F_{l,\text{stab}} \propto |\psi_{CCS}^{(l)}\rangle\langle\psi_l|, \tag{A.21}$$

for any set $\{|\psi_{CCS}^{(l)}\rangle\} \in \mathcal{H}_{CCS}$. However, these dissipators are fairly strict and do not sufficiently describe the dissipative stabilization of well-known codes like the cat code. Here we introduce a more general method based on the quantum Zeno intuition (similar to [102]), and suggest specific examples of dissipators that go beyond the ones proposed in [241], explaining most known strategies for dissipative stabilization of bosonic codes. Our description also helps directly compare dissipative stabilization against Hamiltonian

stabilization, outlined in the next section.

In general, we can achieve a Zeno-type protection by ensuring that the dissipators $\{F_i\}$ measure the population of $\{|\psi_l\rangle\}$ much faster than the intrinsic evolution. This is equivalent to the engineered no-jump evolution induced by these dissipators on $\mathcal{H}_L$

$$\tilde{H}_{\text{eng}}\Pi_L = \frac{i}{2}\sum_i F_i^\dagger F_i \; \Pi_L$$

satisfying Condition **C4** with all imaginary eigenvalues. For a simpler intuition for this condition, let us consider an explicit set of engineered dissipators $\{F_{j,\text{stab}}\}$ whose only purpose is to stabilize against leakage. These dissipators will stabilize against leakage at first order if:

C4a The dissipators have no action on both the code space and the error space

$$F_{j,\text{stab}}\Pi_{CCS} = 0 \; \forall j$$

and

C4b The no-jump evolution of the dissipators forms projectors onto the states in $\mathcal{H}_L$ that the code space can leak to:

$$\sum_j F_{j,\text{stab}}^\dagger F_{j,\text{stab}} = \sum_l \lambda_l |\psi_l\rangle\langle\psi_l|,$$
$$\text{where } \lambda_l \geq 2\Delta_{\text{stab}} \; \forall l$$

This engineered no-jump evolution causes the relevant states in $\mathcal{H}_L$ to rapidly decay in amplitude at a rate $\sim \Gamma\Delta_{\text{stab}}$.

Interestingly, different engineered dissipators can produce the same no-jump evolu-

tion, and hence the same Zeno effect. For example, the dissipators of the form $F_{\ell,\text{stab}} \propto |\psi_{CCS}^{(\ell)}\rangle\langle\psi_\ell|$ suggested in [241] will have no-jump evolution $\propto \sum_\ell |\psi_\ell\rangle\langle\psi_\ell|$ regardless of the states $\{|\psi_{CCS}^{(\ell)}\rangle\}$. These dissipators can be understood as measuring whether the state is in $\mathcal{H}_L$, and if it is, mapping it to $\mathcal{H}_{CCS}$, forming an attractor to $\mathcal{H}_C$. Because in the limit of strong dissipation, the Zeno effect implies we will never measure the state in $\mathcal{H}_L$, the precise state we map back to is not important for stabilizating against leakage. Understanding this Zeno intuition lets us propose an even simpler dissipator of the form

$$F_{\text{stab}} \propto \sum_\ell |\psi_\ell\rangle\langle\psi_\ell|$$

with the same no-jump evolution, which does not form an attractor to $\mathcal{H}_C$ but nonetheless stabilizes against leakage out of $\mathcal{H}_C$.

Let us now consider the example of dissipatively stabilizing the 2-cat, which has leakage levels $|\pm\alpha^{(1)}\rangle$. We can construct and study dissipative strategies to stabilize the 2-cat using Conditions C4a and C4b, with $|\psi_l\rangle = \{|\pm\alpha^{(1)}\rangle\}$. Since jump errors only cause phase flips within the code space, the 2-cat has no error space $\mathcal{H}_E$ ($\Pi_{CCS} = \Pi_C$) and thus Condition C4a implies:

$$F_{\text{stab}}|\pm\alpha\rangle = 0.$$

To satisfy Condition C4b, we can try two candidate dissipators using the leakage states $|\pm\alpha^{(1)}\rangle$:

$$\begin{aligned} F_{\text{L,prev}} &= |\alpha^{(1)}\rangle\langle\alpha^{(1)}| + |-\alpha^{(1)}\rangle\langle-\alpha^{(1)}| \\ F_{\text{C,prev}} &= |\alpha\rangle\langle\alpha^{(1)}| + |-\alpha\rangle\langle-\alpha^{(1)}|. \end{aligned} \tag{A.22}$$

The former simply measures the leakage space, while the second maps the leakage space back to the most likely code state from which the population might have been transferred.

Both of these scale similarly with κ/Γ, since at first order, both of these are stabilizating against leakage using the quantum Zeno effect (Fig. A.6a). Note that one can add arbitrary terms that only act on the rest of the leakage space, and still preserve this first-order protection:

$$\begin{aligned} F'_{\text{C, prev}} &= F_{\text{C, prev}} + F_{\text{LL}}, \\ \text{s.t. } F_{\text{LL}}|\pm\alpha\rangle &= 0 \text{ and } F_{\text{LL}}|\pm\alpha^{(1)}\rangle = 0. \end{aligned}$$

For a specific choice of F_{L}, this results in the familiar and experimentally demonstrated two-photon dissipation of the 2-cat:

$$\begin{aligned} F_{\text{2ph}} &= \hat{a}^2 - \alpha^2, \\ &\approx 2\alpha F_{\text{C, prev}} + \\ &\quad \sum_{\substack{n>2, \\ \pm\alpha}} \left(\sqrt{n(n-1)}|\alpha^{(n-2)}\rangle + 2\alpha\sqrt{n}|\alpha^{(n-1)}\rangle\right)\langle\alpha^{(n)}|, \end{aligned} \tag{A.23}$$

where the approximation for each displaced Fock state is true up to $O\left(|\langle -\alpha^{(n)}|\alpha^{(n)}\rangle|\right)$. This means that the familiar dissipative stabilization of the 2-cat code, to lowest order, simply utilizes the quantum Zeno effect.

In general, for continuous variable codes, one can often create a stabilizer-like description for dissipators that stabilize against the intrinsic no-jump evolution. To measure out the leakage space without affecting error correctability (and therefore be compatible with Conditions C4a and C4b), it is sufficient that these stabilizers act identically on the code and the error space, and distinguish them from the leakage space. This means the $+1$ eigenvalue eigenspace of the stabilizer contains the corrupted code space $\mathcal{H}_{\text{CCS}}$ and the leakage space $\mathcal{H}_L$ lies outside this eigenspace. Thus, the operator S_{CCS} can be used to

form a stabilizing dissipator F_{stab} if

$$\begin{aligned} &S_{\text{CCS}}|\psi\rangle = |\psi\rangle \forall |\psi\rangle \in \mathcal{H}_{\text{CCS}} \\ &\Rightarrow F_{\text{stab}} = S_{\text{CCS}} - \mathbb{1}. \end{aligned} \tag{A.24}$$

Note that if S_{CCS} is non-Hermitian, this dissipator, along with error correction, can also prepare the code space from an appropriate starting state. Common examples include $S_{\text{CCS}} = \hat{a}^n/\alpha^n$ for any n-component cat code, and similarly derived encodings like the squeezed cat code [255]:

$$\begin{aligned} S_{\text{CCS}} &= \hat{S}^\dagger(z)\, \frac{\hat{a}^n}{\alpha^n}\, \hat{S}(z) \\ \hat{S}(z) &= \exp\left(z^* \hat{a}^2 - z\hat{a}^{\dagger 2}\right). \end{aligned}$$

Strong Hamiltonian stabilization

In the previous section, we used a strong engineered non-Hermitian Hamiltonian generated by an engineered no-jump evolution to counteract leakage. We can achieve the same effect with a strong engineered Hermitian Hamiltonian instead, which generalizes 'Kerr-cat' like protection to arbitrary codes. In this case, rather than appealing to the quantum Zeno effect, we use the rotating wave approximation (RWA), which says that continuous evolution between states of different energy levels is suppressed when the energy difference between the levels is large compared to the magnitude of the matrix elements coupling them.

Concretely, we can satisfy Condition **C4** by adding an engineered Hamiltonian H_{eng} that acts on $\{|\psi_l\rangle\}$ with eigenvalues $\geq \Gamma\Delta_{\text{stab}}$ for some Hamiltonian gap Δ_{stab}. Explicitly,

C4c Hamiltonian H_{eng} must create a detuning $\propto \Delta_{\text{stab}}$ between the code space and the

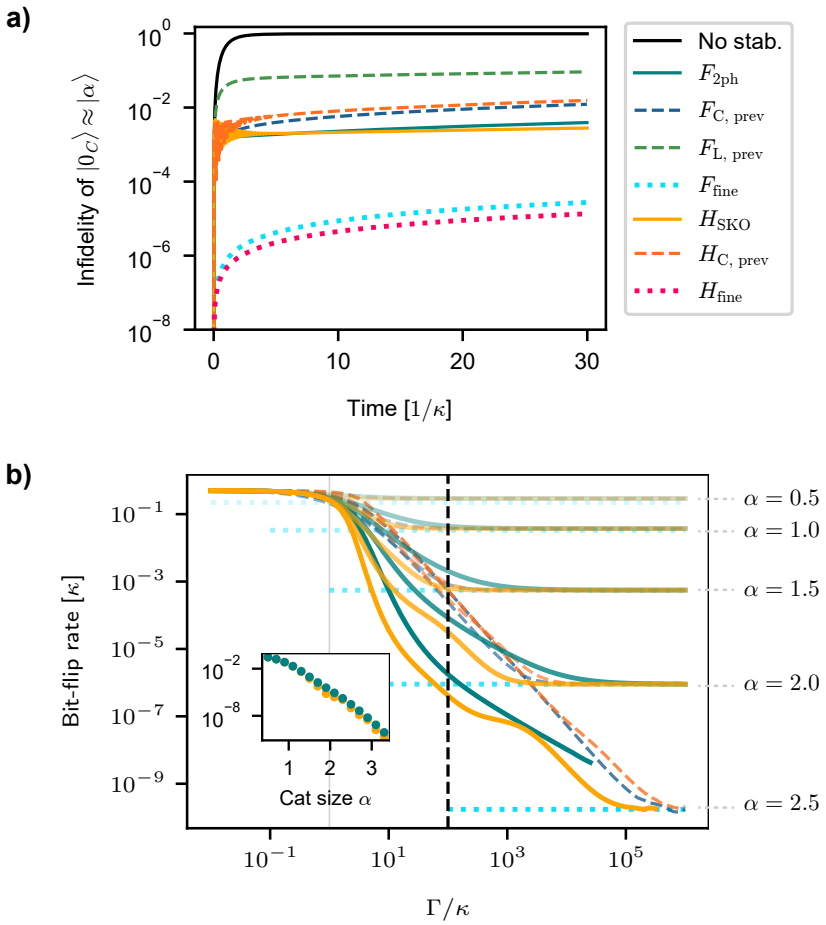

Figure A.6: **Comparing performance for 2-cat stabilization. a)**, Extracted decay rate of $\langle Z_L \rangle$ as a function of the relative strength of each stabilization method. Each strong strategy has been normalized to have an equal engineered no-jump on the first-order leakage states $|\pm\alpha^{(1)}\rangle$. The fine-tuned strategies saturate the bound provided by the exponentially small Z_L errors due to the non-zero overlap $\langle\alpha|-\alpha\rangle$. Note that both fine-tuned strategies have identical error rates, so that the curve for H_κ exactly covers F_κ **b)**, Comparison of extracted $\langle Z_L \rangle$ decay rate for the Kerr-cat and dissipative cat, compared to the strategies proposed in this frameworks, for various cat sizes $\{\alpha\}$. The proposed strong engineered strategies (F_C and H_C, blue and orange dashed lines) follow a consistent bit-flip suppression $\propto (\kappa/\Gamma)^2$, and saturate to the same bounds as the Kerr and dissipative cats. The fine-tuned dissipative (blue dotted) and Hamiltonian (not shown) strategies always saturate this bound and outperform the other strategies at all α. Within the Kerr and dissipative cat, the Hamiltonian gap always outperforms its equivalent dissipative gap for any fixed α, and shows interesting structure. At high α and Γ, the simulations for the dissipative and Kerr cat strategies become numerically unstable, so we omit the final points for $\alpha = 2.5$. At any sufficiently large Γ/κ, one still sees the exponential suppression of bit flips with the size of the 2-cat, as shown in the inset for $\Gamma/\kappa = 100$ (black dashed line).

relevant leakage subspace. Because of our convention $H_{\text{eng}}\Pi_C = 0$, this means:

$$\begin{aligned} H_{\text{eng}}\Pi_L &= \sum_l \lambda_l |\psi_l\rangle\langle\psi_l| + H_L, \\ \text{and } \lambda_l &\geq \Delta_{\text{stab}} \; \forall l \end{aligned} \tag{A.25}$$

where H_L is some Hamiltonian acting only on the leakage states outside of span$[|\psi_\ell\rangle]$

The correspondence between the Hamiltonian strategy and the Zeno strategy in Sec. A.3 goes beyond just the fact that they both protect the code space by engineering a gap in the effective Hamiltonian. For any dissipators $\{F_{j,\text{stab}}\}$ that stabilize against leakage using the Zeno strategy, we can always construct a Hamiltonian that protects the same code space:

$$H_{\text{eng}} = \frac{1}{2}\sum_j F_{j,\text{stab}}^\dagger F_{j,\text{stab}} \tag{A.26}$$

with the same effective Hamiltonian gap. We henceforth refer to this as the Hamiltonian-Zeno correspondence. We will show later that this engineered Hamiltonian always outperforms the corresponding dissipator in stabilizating against leakage, assuming the leakage is perfectly unrecoverable. However, dissipative strategies may be engineered to form an attractor to the code space, and thus can be used to prepare code states as well as stabilize them. Hamiltonian strategies, by constrast, never form an attractor to the code space and thus can only stabilize code states that have already been prepared.

For the specific example of protecting the 2-cat, we construct a Hamiltonian with a gap to the leakage state from the dissipators of appendix A.3 using the Hamiltonian-Zeno correspondence:

$$\begin{aligned} H_{\text{C, prev}} &= \frac{1}{2} F_{\text{C, prev}}^\dagger F_{\text{C, prev}} \\ &= \frac{1}{2}\left(|\alpha^{(1)}\rangle\langle\alpha^{(1)}| + |-\alpha^{(1)}\rangle\langle-\alpha^{(1)}|\right). \end{aligned} \tag{A.27}$$

Note again that the effective Hamiltonian gap here is by construction identical to that generated by $F_{L,\text{stab}}$ and $F_{C,\text{stab}}$, yet as seen in Fig. A.5b all three protections perform differently, and we derive this difference analytically in a code-independent manner in the next section. We can also derive the familiar 'squeezed Kerr Hamiltonian' [34, 256] by using the Hamiltonian-Zeno correspondence on F_{2ph}:

$$\begin{aligned} H_{2ph} &= \frac{1}{2} F_{2ph}^{\dagger} F_{2ph} = \frac{1}{2} \hat{a}^{\dagger 2} \hat{a}^2 - \left(\alpha^2 \hat{a}^{\dagger 2} + \alpha^{*2} \hat{a}^2 \right) \\ &= 4|\alpha|^2 H_{\text{C, prev}} + H_{LL} \end{aligned} \tag{A.28}$$

where H_{LL} consists of higher-order displaced Fock terms $|\alpha^{(n>1)}\rangle\langle\alpha^{(m>1)}|$. This demonstrates that the gap to the leakage space is $4|\alpha|^2$, as in previous work on the spectrum of Kerr cats [34, 249, 256]. Like the previous comparison between dissipative and Hamiltonian gaps, the Kerr-cat (H_{2ph}) and the dissipative cat (F_{2ph}) also share the same effective Hamiltonian gap, yet perform differently on stabilizing the 2-cat, as seen in Fig. A.6b. Similar Hamiltonian protections can also be formed from the stabilizer-like dissipators $F_{\text{stab}} = S_{\text{CCS}} - \mathbb{1}$ introduced in the previous subsection.

Fine-tuned leakage cancellation

The previous strategies for stabilizating against no-jump evolution relied on engineering a strong anti-Hermitian or Hermitian Hamiltonian that suppressed leakage by preventing the no-jump evolution from coherently accumulating population in the leakage space. These strategies did not require fine-tuning: they systematically improve as we increase the strength of our engineered operators ($\Gamma/\kappa \to \infty$). In contrast, our final strategy for combating leakage involves engineering a fine-tuned Lindbladian that exactly cancels the no-jump evolution out of the code space, and thus has a strength comparable to the intrinsic dissipators.

We start with a purely Hermitian cancellation of the no-jump evolution. At first glance, it seems counter-intuitive that the non-unitary no-jump evolution could be canceled by unitary evolution - and it is indeed true that we cannot simultaneously cancel the action of the no-jump evolution on both $\mathcal{H}_C$ and $\mathcal{H}_L$. However, canceling the evolution on $\mathcal{H}_C$ alone is sufficient, since a successful cancellation implies that $\mathcal{H}_L$ has no population at first order. Following Condition **C5**, it is thus sufficient for the Hermitian Hamiltonian to only cancel the intrinsic effective Hamiltonian between $\mathcal{H}_C$ and $\mathcal{H}_L$, with matrix elements between $\mathcal{H}_L$ and $\mathcal{H}_C$ following from Hermiticity:

$$\Pi_L H_\kappa \Pi_C = \frac{i\kappa}{2} \Pi_L \sum_j E_j^\dagger E_j \Pi_C$$
$$\Pi_C H_\kappa \Pi_L = (\Pi_L H_\kappa \Pi_C)^\dagger$$

We note that, unlike in the case of strong Hamiltonian stabilization, H_κ necessarily only connects leakage states to precisely the code word they leaked from.

For a more constructive approach to finding H_κ, we find an orthonormal basis for the code space, with each basis element $|\psi_C^{(j)}\rangle \in \mathcal{H}_C$ mapping to a corresponding normalized leakage state $|\psi_L^{(j)}\rangle := \frac{1}{\mathcal{N}} \Pi_L \sum_k E_k^\dagger E_k |\psi_C^{(j)}\rangle$. H_κ then has a simple Pauli-like representation in each Bloch sphere defined by $|\psi_C^{(j)}\rangle$ and $|\psi_L^{(j)}\rangle$:

$$H_\kappa = \sum_j H_\kappa^{(j)} \tag{A.29}$$
$$= \sum_j \frac{1}{2} \langle \psi_L^{(j)}| \sum_k E_k^\dagger E_k |\psi_C^{(j)}\rangle \ \sigma_{y,j}, \tag{A.30}$$
$$\sigma_{y,j} = i \left(|\psi_L^{(j)}\rangle\langle\psi_C^{(j)}| - |\psi_C^{(j)}\rangle\langle\psi_L^{(j)}| \right).$$

This simply amounts to an effective Rabi drive whose phase and strength are determined by the action of no-jump evolution $\frac{\kappa}{2} \sum_k E_k^\dagger E_k$ on each code word.

We can also enact a purely anti-Hermitian cancellation of the no-jump evolution, through an engineered fine-tuned dissipator. This dissipator is intricately linked to the Hermitian Hamiltonian we engineered above:

$$-\frac{i}{2}\,\Pi_L F_\kappa^\dagger F_\kappa \Pi_C = \Pi_L H_\kappa \Pi_C, \tag{A.31}$$

which is the fine-tuned equivalent of the Hamiltonian-Zeno correspondence [3]. The dissipator, however, also comes with its own discrete jumps and thus requires the additional condition that these jumps cause no logical error. Combining these conditions, we find the following constructive formulation for a fine-tuned dissipator:

$$F_\kappa = \gamma \Pi_{CCS} - \frac{1}{\gamma}\sum_j \langle\psi_L^{(j)}| \sum_k E_k^\dagger E_k |\psi_C^{(j)}\rangle \; |\psi_C^{(j)}\rangle\langle\psi_L^{(j)}| \tag{A.32}$$

where γ is an arbitrary constant that does not affect the strength of protection.

To complete this with our example of protecting the 2-cat, we can construct the following fine-tuned Hamiltonian or dissipator respectively:

$$\begin{aligned} H_{\text{fine, C}} &= \sum_{\pm\alpha} \frac{i\kappa\alpha}{2}\left(|\alpha^{(1)}\rangle\langle\alpha| - |\alpha\rangle\langle\alpha^{(1)}|\right) \\ F_{\text{fine, C}} &= \sum_{\pm\alpha} |\alpha\rangle\langle\alpha| + \kappa\alpha \; |\alpha\rangle\langle\alpha^{(1)}| \end{aligned} \tag{A.33}$$

Both of these strategies saturate the bound provided by the finite overlap of the code states and significantly outperform the strong engineered Hamiltonians or dissipators, as shown in Fig. A.6a.

[3] We could also have multiple dissipators $\{F_{i,\kappa}\}$ to generate this no-jump evolution, if this is easier to engineer

A.4 Discretizing protection with qubitized Lindbladians

Now that we have introduced protections against both the jump errors and no-jump back-action, our goal is to combine and implement them through qubitized dissipation. In particular, the previous sections dealt with continuous dissipators or Hamiltonians, but these are in general high-order interactions that are difficult to engineer in current circuit-QED experiments. We therefore move to a stroboscopic picture similar to previous literature, where the proposed dissipators are realized through an entangling unitary with an ancillary qubit, while preserving their autonomous nature. Crucially, we describe how to interpret this "qubitized" protection, and show that the protections from this stroboscopic recovery scale similarly to the continuous protection. In addition to experimental feasibility, moving to the qubitized picture also lets us connect to discrete gate-based realizations of bosonic error correction. It is well known that measurements of an error syndrome, with a classically conditioned feedback unitary, can be reconstructed as an autonomous protocol with a quantum conditional unitary and an unconditional reset of the qubit, as shown in Fig ... Our construction below follows precisely the same circuit for also realizing stroboscopic versions of continuous dissipators, thus being able to imitate both protocols.

Central to the qubitized dissipation are unitary interactions of the form

$$U = \exp\{-i\lambda\delta t(\sigma^+ L + \sigma^- L^\dagger)\} \tag{A.34}$$

between the qubit and the cavity, where $\sigma^\pm$ are the qubit raising/lowering operators and L is the dissipation jump operator that we desire on the cavity. If the qubit always starts in $|g\rangle$, it is immediately clear that the unitary acts as a detector, exciting the qubit when the cavity state is not in the kernel of L. To more intimately understand the action of U on the

qubit-cavity system, we can expand the unitary as (note that $(\sigma^{\pm})^2 = 0$):

$$U = \sum_n \frac{(-i\lambda\delta t)^n}{n!}(\sigma^+ L + \sigma^- L^\dagger)^n \tag{A.35}$$

$$= \sum_{n \text{ even}} \frac{(-i\lambda\delta t)^n}{n!}\left[|g\rangle\langle g|(L^\dagger L)^{\frac{n}{2}} + |e\rangle\langle e|(LL^\dagger)^{\frac{n}{2}}\right] \tag{A.36}$$

$$+ \sum_{n \text{ odd}} \frac{(-i\lambda\delta t)^n}{n!}\left[\sigma^+ L(L^\dagger L)^{\frac{n-1}{2}} + \sigma^- L^\dagger (LL^\dagger)^{\frac{n-1}{2}}\right]$$

$$= |g\rangle\langle g|\cos\left(\lambda\delta t\sqrt{L^\dagger L}\right) + |e\rangle\langle e|\cos\left(\lambda\delta t\sqrt{LL^\dagger}\right) \tag{A.37}$$

$$- i\sigma^+ L\frac{\sin\left(\lambda\delta t\sqrt{L^\dagger L}\right)}{\sqrt{L^\dagger L}} - i\sigma^- L^\dagger\frac{\sin\left(\lambda\delta t\sqrt{LL^\dagger}\right)}{\sqrt{LL^\dagger}}$$

.

We use this unitary to implement each qubitized dissipator as a repeated Kraus-rank-two quantum channel on the cavity, by entangling a qubit initialized in $|g\rangle$ with the operator U, and then autonomously resetting the qubit.

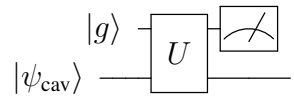

(A.38)

In general, this operation is described by the Kraus map

$$\rho \mapsto K_g \rho K_g^\dagger + K_e \rho K_e^\dagger. \tag{A.39}$$

Here, K_g is the operator associated with the qubit ending in $|g\rangle$ before reset

$$\begin{aligned} K_g|\psi_{\text{cav}}\rangle &:= (\langle g| \otimes \mathbb{1})U(|g\rangle|\psi_{\text{cav}}\rangle) \\ &= \cos\left(\lambda\delta t\sqrt{L^\dagger L}\right)|\psi_{\text{cav}}\rangle \end{aligned} \tag{A.40}$$

and K_e is the operator associated with the qubit ending in $|e\rangle$

$$\begin{aligned} K_e|\psi_{\text{cav}}\rangle &:= (\langle e| \otimes \mathbb{1})U(|g\rangle|\psi_{\text{cav}}\rangle) \\ &= L\frac{\sin\left(\lambda\delta t\sqrt{L^\dagger L}\right)}{\sqrt{L^\dagger L}}|\psi_{\text{cav}}\rangle, \end{aligned} \tag{A.41}$$

where we have neglected an irrelevant phase on K_e. The mathematical form of these Kraus operators makes them particularly amenable to physical interpretation. We have, for $\Gamma\delta t \ll 1$, $K_g \propto \left(1 - \frac{\Gamma\delta t}{2}L^\dagger L\right)$, which is precisely the no-jump evolution due to the dissipator L. This makes sense, since the qubit has been found to remain in its ground state after the interaction with the cavity, and thus the cavity has not undergone a jump under L. The action of K_e in the same limit is correspondingly a jump under L. The qubit thus exactly provides a witness to whether the engineered jump operators acted or did not, which represents the extracted entropy from the oscillator's logical encoding. Resetting the qubit then evacuates this entropy, forming the required qubitized dissipators that protect the oscillator's information.

This project is still underway, and current efforts are directed towards studying how well such a qubitized implementation protects logical information, compared to the ideal continuous Hamiltonians and dissipators. A full qubitized protection would require alternating Unitaries for jump correction and no-jump stabilization respectively. However, more importantly, each such unitary would entangle the qubit and oscillator information, which makes our logical information maximally suffer from any qubit errors. This is precisely where the novel architectures and control techniques introduced in the thesis become important. Specifically, they provide an advantage because:

1. These Unitaries in general will contain number-selective parts and continuous variable parts, and thus having natively number-selective and continuous controls ensures that they can be achieved in reasonably short Trotterized sequences.

2. Higher-order dissipators will benefit from higher order controls. As an example, constructing the $\mathcal{D}[a^4 - \alpha^4]$ dissipator from displacement and SNAP sequences is significantly more difficult than constructing it from squeezing and SNAP.

3. Ultimately, the qubit can have Markovian errors that propagate to the oscillator. Engineering symmetries and turning on interactions only when non-Gaussian gates are required, as outlined in Chapter 5, ensure the effect of such errors is minimized.

4. Qubitized dissipators necessarily require frequent qubit resets. Any back-action during such resets can then entirely spoil oscillator information, and hence being able to decouple the qubit during resets should provide a significant advantage.

Even with these resilient architectures, since each Unitary will eventually be be constructed through pulse optimization in an actual experiment, special care needs to be taken to maximally utilize the error resilience of the setup. We already know the sufficient conditions for achieving autonomous error correction – we could simply apply these to the cost function for the pulse optimization itself! Specifically, the optimizer could utilize a Monte Carlo sampling to calculate its cost function, where trajectories in which the ancilla has a jump could be separately held accountable to not propagate errors to the logical qubit, which amounts to any ancilla errors only being allowed to corrupt the gauge space. The total cost function would then look like:

$$\mathcal{F} = \left| \sum_{n \leq N} \langle n, g | U_{\text{targ}}^{\dagger} U_{NJ} | n, g \rangle \right| + \sum_{\epsilon,\, q \in g,e} \left| \sum_{\alpha,i,j} \langle \psi_{\alpha,i}, q | U_{J,\epsilon} | \psi_{\alpha,j}, g \rangle \right|, \tag{A.42}$$

where U_{targ} is the target Unitary, U_{NJ} is the control propagator under no jumps, $|n\rangle$ is the oscillator eigenstate, α and i, j are the indices corresponding to $\mathcal{H}_q$ and $\mathcal{H}_G$ respectively. Implementing this optimizer, and testing it on the LINC architecture, is the subject of actively ongoing research.

Acronyms

3WM Three-wave mixer. 119, 123, 137

AMO Atomic and Molecular Optics. 14, 15

ATS Asymmetrically Threaded SQUID. 125

BQP bounded-error Quantum Polynomial time. 3, 4, 32

CV Continuous Variable. 19, 32, 33, 38, 60, 67, 121, 167, 168

DDS Differentially Driven SQUID. 73, 84, 85, 90, 91, 100, 103–105, 117, 118, 124, 125, 127, 129, 140, 143, 144, 156

DRAG Derivative Removal by Adiabatic Gate. 176, 178

ECD Echoed Conditional Displacement. 188

EM Electro-magnetic. 12, 24, 42, 58, 68, 85–88, 99, 104, 105, 157

GKP Gottesman Kitaev Preskill. 66, 68, 156, 199

HEMEX Heat Exchanger Method (Extra pure). 162

HFSS High Frequency Simulation Software. 58, 68, 87, 88, 105

KL Knill-Laflamme. 200

LINC Linear Inductive Coupler. 64, 121, 122, 124–126, 128, 156, 165, 189, 192

NEC Nippon Energy Corporation. 17

NMR Nuclear Magnetic esonance. 13, 14

NP Nondeterministic Polynomial time. 3, 4

OCT Optimal Control Theory. 65, 66, 185

P deterministic Polynomial time. 4

PSPACE Polynomial Space complexity. 4

Q Quality factor (ω/κ). 57, 88, 90

QED Quantum Electrodynamics. 4, 14, 74, 165, 169, 199

RF Radio frequency. 120, 142

RSA Rivest–Shamir–Adleman cryptosystem. 3

SNAIL Superconducting Nonlinear Asymmetric Inductive eLement. 120, 127, 142

SNAP Selective Number Arbitrary Phase. 65, 66, 176, 177, 181, 182, 234

SQUID Superconducting Quantum Interference Device. 64, 71, 74–76, 81, 82, 87, 88, 120, 124, 142

VNA Vector Network Analyzer. 61

Bibliography

[1] J. J. Sakurai, "Modern Quantum Mechanics (Revised Edition)". Addison Wesley, 1 edition, September, 2014. ISBN 0201539292. 1

[2] R. P. Feynman, "QED: The Strange Theory of Light and Matter". Princeton University Press, 10, 2014. ISBN 978-0-691-16409-0, 978-0-691-02417-2. 1

[3] M. E. Peskin and D. V. Schroeder, "An Introduction to quantum field theory". Addison-Wesley, Reading, USA, 1995. doi:10.1201/9780429503559, ISBN 978-0-201-50397-5, 978-0-429-50355-9, 978-0-429-49417-8. 1

[4] X. Fan, T. G. Myers, B. A. D. Sukra, and G. Gabrielse, "Measurement of the electron magnetic moment", *Phys. Rev. Lett.* **130** (Feb, 2023) 071801, doi:10.1103/PhysRevLett.130.071801. 1

[5] S. Storzet al., "Complete self-testing of a system of remote superconducting qubits", 2024. https://arxiv.org/abs/2408.01299. 2

[6] G. Semeghini et al., "Probing topological spin liquids on a programmable quantum simulator", *Science* **374** (2021), no. 6572, 1242–1247, doi:10.1126/science.abi8794, arXiv:https://www.science.org/doi/pdf/10.1126/science.abi8794. 2, 5

[7] Google Quantum AI and Collaborators et al., "Non-Abelian braiding of graph vertices in a superconducting processor", *Nature* **618** (June, 2023) 264–269, doi:10.1038/s41586-023-05954-4. 2, 6

[8] M. C. Bañuls et al., "Simulating lattice gauge theories within quantum technologies", *The European Physical Journal D* **74** (August, 2020) doi:10.1140/epjd/e2020-100571-8. 2

[9] T. A. Cochranet al., "Visualizing dynamics of charges and strings in (2+1)d lattice gauge theories", 2024. https://arxiv.org/abs/2409.17142. 2, 4

[10] M. A. Nielsen and I. L. Chuang, "Quantum Computation and Quantum Information: 10th Anniversary Edition". Cambridge University Press, 2011. ISBN 9781107002173. 2

[11] A. M. Dalzellet al., "Quantum algorithms: A survey of applications and end-to-end complexities", 2023. https://arxiv.org/abs/2310.03011. 2

[12] E. Bernstein and U. Vazirani, "Quantum complexity theory", *SIAM Journal on Computing* **26** (1997), no. 5, 1411–1473, doi:10.1137/S0097539796300921. 3

[13] S. Aaronson, "Read the fine print", *Nature Physics* **11** (April, 2015) 291–293, doi:10.1038/nphys3272. 3

[14] S. Aaronson and L. Chen, "Complexity-theoretic foundations of quantum supremacy experiments", 2016. https://arxiv.org/abs/1612.05903. 3

[15] Wikipedia, "BQP — Wikipedia, the free encyclopedia", 2025. [Online; accessed 02-February-2025]. http://en.wikipedia.org/w/index.php?title=BQP&oldid=1230041542. 3

[16] S. Aaronson, "The limits of quantum", *Scientific American* **298** (2008), no. 3, 62–69. 4

[17] I. Bloch and M. Greiner, "The superfluid-to-Mott insulator transition and the birth of experimental quantum simulation", *Nature Reviews Physics* **4** (September, 2022) 739–740, doi:10.1038/s42254-022-00520-9. 4

[18] I. M. Georgescu, S. Ashhab, and F. Nori, "Quantum simulation", *Rev. Mod. Phys.* **86** (Mar, 2014) 153–185, doi:10.1103/RevModPhys.86.153. 3

[19] J. Jumper et al., "Highly accurate protein structure prediction with AlphaFold", *Nature* **596** (August, 2021) 583–589, doi:10.1038/s41586-021-03819-2. 5

[20] H. A. Scheraga, M. Khalili, and A. Liwo, "Protein-Folding Dynamics: Overview of Molecular Simulation Techniques", *Annual Review of Physical Chemistry* **58** (2007), no. Volume 58, 2007, 57–83, doi:https://doi.org/10.1146/annurev.physchem.58.032806.104614. 5

[21] Wikipedia, "Numerical sign problem — Wikipedia, The Free Encyclopedia". http://en.wikipedia.org/w/index.php?title=Numerical%20sign%20problem&oldid=1261880833, 2025. [Online; accessed 02-February-2025]. 5

[22] E. Crane et al., "Hybrid Oscillator-Qubit Quantum Processors: Simulating Fermions, Bosons, and Gauge Fields", September, 2024. arXiv:2409.03747 [quant-ph]. doi:10.48550/arXiv.2409.03747, http://arxiv.org/abs/2409.03747. 5, 196

[23] T. Hensgens et al., "Quantum simulation of a Fermi–Hubbard model using a semiconductor quantum dot array", *Nature* **548** (August, 2017) 70–73, doi:10.1038/nature23022. 5

[24] L. Clinton, J. Bausch, and T. Cubitt, "Hamiltonian simulation algorithms for near-term quantum hardware", *Nature Communications* **12** (August, 2021) 4989, doi:10.1038/s41467-021-25196-0. 6

[25] Y. Kim et al., "Evidence for the utility of quantum computing before fault tolerance", *Nature* **618** (June, 2023) 500–505, doi:10.1038/s41586-023-06096-3. 6

[26] S. Bravyi et al., "High-threshold and low-overhead fault-tolerant quantum memory", *Nature* **627** (March, 2024) 778–782, doi:10.1038/s41586-024-07107-7. 6

[27] A. Gold et al., "Entanglement across separate silicon dies in a modular superconducting qubit device", *npj Quantum Information* **7** (September, 2021) 142, doi:10.1038/s41534-021-00484-1. 6

[28] S. Haroche and J. Raimond, "Exploring the Quantum: Atoms, Cavities, and Photons". Oxford University Press, Oxford, England, 2006. 6, 11

[29] A. A. Clerk et al., "Introduction to Quantum Noise, Measurement and Amplification", *Reviews of Modern Physics* **82** (April, 2010) 1155–1208, doi:10.1103/RevModPhys.82.1155. 6, 8, 36

[30] S. M. Girvin, "Circuit QED: superconducting qubits coupled to microwave photons", in *Quantum Machines: Measurement and Control of Engineered Quantum Systems*, M. Devoret, B. Huard, R. Schoelkopf, and L. F. Cugliandolo, eds., pp. 113–256. Oxford University Press, June, 2014. doi:10.1093/acprof:oso/9780199681181.003.0003. 7

[31] Wikipedia, "Kubo formula — Wikipedia, The Free Encyclopedia". http://en.wikipedia.org/w/index.php?title=Kubo%20formula&oldid=1249042296, 2025. [Online; accessed 02-February-2025]. 8, 9

[32] V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret, "Fluxonium: Single cooper-pair circuit free of charge offsets", *Science* **326** (2009), no. 5949, 113–116, doi:10.1126/science.1175552, arXiv:https://www.science.org/doi/pdf/10.1126/science.1175552. 9, 30, 78

[33] A. Kitaev, "Protected qubit based on a superconducting current mirror", *arXiv:cond-mat/0609441* (September, 2006). 9

[34] S. Puri, S. Boutin, and A. Blais, "Engineering the quantum states of light in a Kerr-nonlinear resonator by two-photon driving", *npj Quantum Information* **3** (April, 2017) 18, doi:10.1038/s41534-017-0019-1. 9, 48, 220, 228

[35] N. Earnest et al., "Realization of a Λ system with metastable states of a capacitively shunted fluxonium", *Phys. Rev. Lett.* **120** (Apr, 2018) 150504, doi:10.1103/PhysRevLett.120.150504. 9

[36] W. H. Zurek, "Pointer basis of quantum apparatus: Into what mixture does the wave packet collapse?", *Phys. Rev. D* **24** (Sep, 1981) 1516–1525, doi:10.1103/PhysRevD.24.1516. 11

[37] D. Manzano, "A short introduction to the lindblad master equation", *AIP Advances* **10** (02, 2020) 025106, doi:10.1063/1.5115323, arXiv:https://pubs.aip.org/aip/adv/article-pdf/doi/10.1063/1.51153 11

[38] D. Steck, "Quantum and Atom Optics". 2007. 11

[39] E. Knill, R. Laflamme, and L. Viola, "Theory of quantum error correction for general noise", *Physical review letters* **84 11** (1996) 2525–8. 11, 113

[40] J. M. Deutsch, "Quantum statistical mechanics in a closed system", *Phys. Rev. A* **43** (Feb, 1991) 2046–2049, doi:10.1103/PhysRevA.43.2046. 13

[41] Wikipedia, "Eigenstate thermalization hypothesis — Wikipedia, The Free Encyclopedia". http://en.wikipedia.org/w/index.php?title=Eigenstate%20thermalization%20hypothesis&oldid=1193767723, 2025. [Online; accessed 03-February-2025]. 13

[42] E. L. Hahn, "Spin echoes", *Phys. Rev.* **80** (Nov, 1950) 580–594, doi:10.1103/PhysRev.80.580. 13, 179

[43] N. Khaneja et al., "Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms", *Journal of Magnetic Resonance* **172** (2005), no. 2, 296–305, doi:https://doi.org/10.1016/j.jmr.2004.11.004. 13

[44] Wikipedia, "Penning trap — Wikipedia, The Free Encyclopedia". http://en.wikipedia.org/w/index.php?title=Penning%20trap&oldid=1272430977, 2025. [Online; accessed 02-February-2025]. 14

[45] W. Paul, "Electromagnetic traps for charged and neutral particles", *Rev. Mod. Phys.* **62** (Jul, 1990) 531–540, doi:10.1103/RevModPhys.62.531. 14

[46] A. Sørensen and K. Mølmer, "Quantum computation with ions in thermal motion", *Phys. Rev. Lett.* **82** (Mar, 1999) 1971–1974, doi:10.1103/PhysRevLett.82.1971. 14

[47] Wikipedia, "Optical tweezers — Wikipedia, The Free Encyclopedia". http://en.wikipedia.org/w/index.php?title=Optical%20tweezers&oldid=1254946316, 2025. [Online; accessed 03-February-2025]. 14

[48] H. J. Levine, "Quantum Information Processing and Quantum Simulation with Programmable Rydberg Atom Arrays". PhD thesis, Harvard University Graduate School of Arts and Sciences, 2021. 15

[49] D. Bluvstein et al., "A quantum processor based on coherent transport of entangled atom arrays", *Nature* **604** (April, 2022) 451–456, doi:10.1038/s41586-022-04592-6. 16

[50] R. M. Kroeze et al., "High cooperativity using a confocal-cavity–qed microscope", *PRX Quantum* **4** (May, 2023) 020326, doi:10.1103/PRXQuantum.4.020326. 15

[51] J. M. Martinis and K. Osborne, "Superconducting qubits and the physics of josephson junctions", 2004. https://arxiv.org/abs/cond-mat/0402415. 17

[52] A. O. Caldeira and A. J. Leggett, "Influence of dissipation on quantum tunneling in macroscopic systems", *Phys. Rev. Lett.* **46** (Jan, 1981) 211–214, doi:10.1103/PhysRevLett.46.211. 17

[53] A. J. Leggett, "Schrödinger's cat and her laboratory cousins", *Contemporary Physics* **25** (1984), no. 6, 583–598, doi:10.1080/00107518408210731, arXiv:https://doi.org/10.1080/00107518408210731. 17

[54] M. H. Devoret, J. M. Martinis, and J. Clarke, "Measurements of Macroscopic Quantum Tunneling out of the Zero-Voltage State of a Current-Biased Josephson Junction", *Physical Review Letters* **55** (October, 1985) 1908–1911, doi:10.1103/PhysRevLett.55.1908. 17

[55] J. M. Martinis, M. H. Devoret, and J. Clarke, "Experimental tests for the quantum behavior of a macroscopic degree of freedom: The phase difference across a josephson junction", *Phys. Rev. B* **35** (Apr, 1987) 4682–4698, doi:10.1103/PhysRevB.35.4682. 17

[56] J. Clarke et al., "Quantum mechanics of a macroscopic variable: The phase difference of a josephson junction", *Science* **239** (1988), no. 4843, 992–997, doi:10.1126/science.239.4843.992, arXiv:https://www.science.org/doi/pdf/10.1126/science.239.4843.992 17

[57] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai, "Coherent control of macroscopic quantum states in a single-Cooper-pair box", *Nature* **398** (April, 1999) 786–788, doi:10.1038/19718. 17

[58] A. Blais et al., "Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation", *Physical Review A* **69** (June, 2004) 062320, doi:10.1103/PhysRevA.69.062320. 17

[59] A. Wallraff et al., "Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics",. 17, 25

[60] I. Siddiqi et al., "Rf-driven josephson bifurcation amplifier for quantum measurement", *Phys. Rev. Lett.* **93** (Nov, 2004) 207002, doi:10.1103/PhysRevLett.93.207002. 17

[61] J. Koch et al., "Charge-insensitive qubit design derived from the Cooper pair box", *Physical Review A* **76** (October, 2007) 042319, doi:10.1103/PhysRevA.76.042319. 17, 21, 23

[62] R. J. Schoelkopf and S. M. Girvin, "Wiring up quantum systems", *Nature* **451** (February, 2008) 664–669, doi:10.1038/451664a. 17, 20

[63] M. Steffen et al., "Measurement of the entanglement of two superconducting qubits via state tomography", *Science* **313** (2006), no. 5792, 1423–1425, doi:10.1126/science.1130886. 17

[64] R. P. Feynman, "Feynman Lecture 21, Caltech". https://www.feynmanlectures.caltech.edu/III_21.html. Accessed: 2025-02-04. 17

[65] Wikipedia, "MIT Radiation Laboratory — Wikipedia, The Free Encyclopedia". http://en.wikipedia.org/w/index.php?title=MIT%20Radiation%20Laboratory&oldid=1255057886, 2025. [Online; accessed 03-February-2025]. 20

[66] D. M. Pozar, "Microwave engineering; 3rd ed.". Wiley, Hoboken, NJ, 2005. 20

[67] B. D. Josephson, "Possible new effects in superconductive tunnelling", *Physics Letters* **1** (1962), no. 7, 251–253, doi:https://doi.org/10.1016/0031-9163(62)91369-0. 20

[68] J. A. Schreier et al., "Suppressing charge noise decoherence in superconducting charge qubits", *Phys. Rev. B* **77** (May, 2008) 180502, doi:10.1103/PhysRevB.77.180502. 21

[69] M. H. Devoret, "Does Brian Josephson's Gauge-Invariant Phase Difference Live on a Line or a Circle?", *Journal of Superconductivity and Novel Magnetism* **34** (June, 2021) 1633–1642, doi:10.1007/s10948-020-05784-9. 23

[70] A. D. Greentree, J. Koch, and J. Larson, "Fifty years of jaynes–cummings physics", *Journal of Physics B: Atomic, Molecular and Optical Physics* **46** (nov, 2013) 220201, doi:10.1088/0953-4075/46/22/220201. 25

[71] D. I. Schuster et al., "Resolving photon number states in a superconducting circuit", *Nature* **445** (February, 2007) 515–518, doi:10.1038/nature05461. 25

[72] J. Gambetta et al., "Qubit-photon interactions in a cavity: Measurement-induced dephasing and number splitting", *Physical Review A* **74** (October, 2006) 042318, doi:10.1103/PhysRevA.74.042318. 25, 61, 62

[73] L. DiCarlo et al., "Demonstration of two-qubit algorithms with a superconducting quantum processor", *Nature* **460** (July, 2009) 240–244, doi:10.1038/nature08121. 25

[74] S. Krinner et al., "Engineering cryogenic setups for 100-qubit scale superconducting circuit systems", *EPJ Quantum Technology* **6** (December, 2019) 2, doi:10.1140/epjqt/s40507-019-0072-0. 26

[75] S. E. Nigg et al., "Black-Box Superconducting Circuit Quantization", *Physical Review Letters* **108** (June, 2012) 240502, doi:10.1103/PhysRevLett.108.240502. 27, 106

[76] Z. K. Minev et al., "Energy-participation quantization of Josephson circuits", *npj Quantum Information* **7** (December, 2021) 131, doi:10.1038/s41534-021-00461-8. 27, 58, 89, 106, 131

[77] P. Groszkowski et al., "Coherence properties of the 0-π qubit", *New Journal of Physics* **20** (April, 2018) 043053, doi:10.1088/1367-2630/aab7cd. arXiv: 1708.02886. 30

[78] B. Eastin and E. Knill, "Restrictions on transversal encoded quantum gate sets", *Phys. Rev. Lett.* **102** (Mar, 2009) 110502, doi:10.1103/PhysRevLett.102.110502. 31

[79] U. Chabaud, M. Joseph, S. Mehraban, and A. Motamedi, "Bosonic quantum computational complexity", 2024. https://arxiv.org/abs/2410.04274. 31

[80] S. Krastanov et al., "Universal control of an oscillator with dispersive coupling to a qubit", *Physical Review A* **92** (October, 2015) 040303, doi:10.1103/PhysRevA.92.040303. 34, 65

[81] C. M. Caves, "Quantum limits on noise in linear amplifiers", *Phys. Rev. D* **26** (Oct, 1982) 1817–1839, doi:10.1103/PhysRevD.26.1817. 36

[82] J. Aumentado, "Superconducting parametric amplifiers: The state of the art in josephson parametric amplifiers", *IEEE Microwave Magazine* **21** (2020), no. 8, 45–59, doi:10.1109/MMM.2020.2993476. 36

[83] B. Abdoet al., "Teleportation and entanglement swapping of continuous quantum variables of microwave radiation", 2025. https://arxiv.org/abs/2501.05537. 36

[84] J. Schwinger, "On angular momentum", technical report, Harvard Univ., Cambridge, MA (United States); Nuclear Development Associates, Inc. (US), 01, 1952. doi:10.2172/4389568. 37

[85] T. Tsunoda et al., "Error-Detectable Bosonic Entangling Gates with a Noisy Ancilla", *PRX Quantum* **4** (June, 2023) 020354, doi:10.1103/PRXQuantum.4.020354. 37, 118

[86] S. Lloyd and S. L. Braunstein, "Quantum Computation over Continuous Variables", *PHYSICAL REVIEW LETTERS* **82** (1999), no. 8, 4. 39

[87] Y. Liuet al., "Hybrid oscillator-qubit quantum processors: Instruction set architectures, abstract machine models, and applications", 2024. https://arxiv.org/abs/2407.10381. 39, 174

[88] A. Eickbusch, "Quantum Control and Error Correction of Grid States in a Superconducting Oscillator". PhD thesis, Yale University, New Haven, USA, 2023. 39

[89] H. J. Groenewold, "Quantum mechanics and its models", *Synthese* **10** (1956), no. 1, 203–209, doi:10.1007/bf00484654. 40

[90] T. L. Curtright, D. B. Fairlie, and C. K. Zachos, "A Concise Treatise on Quantum Mechanics in Phase Space". World Scientific Publishing Co Pte Ltd, 1, 2014. doi:10.1142/8870, ISBN 978-981-4520-45-4. 41

[91] E. Knill, R. Laflamme, and L. Viola, "Theory of quantum error correction for general noise", *Phys. Rev. Lett.* **84** (Mar, 2000) 2525–2528, doi:10.1103/PhysRevLett.84.2525. 46

[92] P. Leviant, Q. Xu, L. Jiang, and S. Rosenblum, "Quantum capacity and codes for the bosonic loss-dephasing channel", *Quantum* **6** (September, 2022) 821, doi:10.22331/q-2022-09-29-821. arXiv:2205.00341 [quant-ph]. 46, 52

[93] V. V. Albert et al., "Performance and structure of single-mode bosonic codes", *Physical Review A* **97** (March, 2018) 032346, doi:10.1103/PhysRevA.97.032346. arXiv:1708.05010 [quant-ph]. 47, 49, 52

[94] S. B. Bravyi and A. Y. Kitaev, "Quantum codes on a lattice with boundary", 1998. https://arxiv.org/abs/quant-ph/9811052. 47

[95] J. J. Wallman and J. Emerson, "Noise tailoring for scalable quantum computation via randomized compiling", *Phys. Rev. A* **94** (Nov, 2016) 052325, doi:10.1103/PhysRevA.94.052325. 47

[96] P. Aliferis and B. M. Terhal, "Fault-tolerant quantum computation for local leakage faults", 2006. https://arxiv.org/abs/quant-ph/0511065. 47

[97] Google Quantum AI and Collaborators et al., "Quantum error correction below the surface code threshold", *Nature* (December, 2024) doi:10.1038/s41586-024-08449-y. 47

[98] R. Lescanne et al., "Exponential suppression of bit-flips in a qubit encoded in an oscillator", *Nature Physics* **16** (May, 2020) 509–513, doi:10.1038/s41567-020-0824-x. 48

[99] N. E. Frattini, "Three-wave Mixing in Superconducting Circuits: Stabilizing Cats with SNAILs". PhD thesis, Yale University, New Haven, USA, 2021. 49, 121

[100] M. Mirrahimi et al., "Dynamically protected cat-qubits: a new paradigm for universal quantum computation", *New Journal of Physics* **16** (April, 2014) 045014, doi:10.1088/1367-2630/16/4/045014. 49, 50, 200, 220, 221

[101] N. Ofek et al., "Extending the lifetime of a quantum bit with error correction in superconducting circuits", *Nature* **536** (August, 2016) 441–445, doi:10.1038/nature18949. 50, 200, 201

[102] M. H. Michael et al., "New Class of Quantum Error-Correcting Codes for a Bosonic Mode", *Physical Review X* **6** (July, 2016) 031006, doi:10.1103/PhysRevX.6.031006. 50, 54, 200, 221

[103] A. L. Grimsmo, J. Combes, and B. Q. Baragiola, "Quantum computing with rotation-symmetric bosonic codes", *Physical Review X* **10** (March, 2020) 011058, doi:10.1103/PhysRevX.10.011058. arXiv: 1901.08071. 50

[104] D. Gottesman, A. Kitaev, and J. Preskill, "Encoding a qubit in an oscillator", *Physical Review A* **64** (June, 2001) 012310, doi:10.1103/PhysRevA.64.012310. arXiv:quant-ph/0008040. 51

[105] M. H. Shaw, A. C. Doherty, and A. L. Grimsmo, "Stabilizer subsystem decompositions for single- and multimode gottesman-kitaev-preskill codes", *PRX Quantum* **5** (Feb, 2024) 010331, doi:10.1103/PRXQuantum.5.010331. 52, 203

[106] L.-A. Sellem, R. Robin, P. Campagne-Ibarcq, and P. Rouchon, "Stability and decoherence rates of a GKP qubit protected by dissipation", *IFAC-PapersOnLine* **56** (2023), no. 2, 1325–1332, doi:https://doi.org/10.1016/j.ifacol.2023.10.1776. 52, 198

[107] P. Campagne-Ibarcq et al., "Quantum error correction of a qubit encoded in grid states of an oscillator", *Nature* **584** (August, 2020) 368–372, doi:10.1038/s41586-020-2603-3. 52, 178, 199

[108] V. V. Sivak et al., "Real-time quantum error correction beyond break-even", *Nature* **616** (April, 2023) 50–55, doi:10.1038/s41586-023-05782-6. 52, 60, 178, 199, 200, 212

[109] V. V. Albert and P. Faist, "Numerically optimized bosonic code", in *The Error Correction Zoo*. The Error Correction Zoo, 2022. 54

[110] C. Gidney and M. Ekerå, "How to factor 2048 bit RSA integers in 8 hours using 20 million noisy qubits", *Quantum* **5** (April, 2021) 433, doi:10.22331/q-2021-04-15-433. 55

[111] A. Kubica et al., "Erasure qubits: Overcoming the T_1 limit in superconducting circuits", *Phys. Rev. X* **13** (Nov, 2023) 041022, doi:10.1103/PhysRevX.13.041022. 56

[112] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, "Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays", *Nature Communications* **13** (August, 2022) 4657, doi:10.1038/s41467-022-32094-6. 56

[113] J. D. Teoh et al., "Dual-rail encoding with superconducting cavities", *Proceedings of the National Academy of Sciences* **120** (2023), no. 41, e2221736120, doi:10.1073/pnas.2221736120, arXiv:https://www.pnas.org/doi/pdf/10.1073/pnas.2221736120. 56, 57, 117, 118

[114] J. Teoh, "Error Detection in Bosonic Circuit Quantum Electrodynamics". PhD thesis, Yale University, New Haven, USA, 2023. 57

[115] M. D. Reed et al., "Fast reset and suppressing spontaneous emission of a superconducting qubit", *Applied Physics Letters* **96** (05, 2010) 203110, doi:10.1063/1.3435463, arXiv:https://pubs.aip.org/aip/apl/article-pdf/doi/10.1063/1.34354 59

[116] Y. Sunada et al., "Fast readout and reset of a superconducting qubit coupled to a resonator with an intrinsic purcell filter", *Phys. Rev. Appl.* **17** (Apr, 2022) 044016, doi:10.1103/PhysRevApplied.17.044016. 59

[117] R. W. Heeres et al., "Cavity State Manipulation Using Photon-Number Selective Phase Gates", *Physical Review Letters* **115** (September, 2015) 137002, doi:10.1103/PhysRevLett.115.137002. 59, 65, 107

[118] J. M. Gambetta, F. Motzoi, S. T. Merkel, and F. K. Wilhelm, "Analytic control methods for high-fidelity unitary operations in a weakly nonlinear oscillator", *Phys. Rev. A* **83** (Jan, 2011) 012308, doi:10.1103/PhysRevA.83.012308. 61, 178

[119] M. S. Khalil, M. J. A. Stoutimore, F. C. Wellstood, and K. D. Osborn, "An analysis method for asymmetric resonator transmission applied to superconducting devices", *Journal of Applied Physics* **111** (03, 2012) 054510, doi:10.1063/1.3692073, arXiv:https://pubs.aip.org/aip/jap/article-pdf/doi/10.1063/1.36920 61

[120] J. Gao, "The Physics of Superconducting Microwave Resonators". PhD thesis, California Institute of Technology, Pasadena, USA, 2008. 61

[121] M. F. Dumas et al., "Measurement-induced transmon ionization", *Phys. Rev. X* **14** (Oct, 2024) 041023, doi:10.1103/PhysRevX.14.041023. 62

[122] D. Sank, "Fast, Accurate State Measurement in Superconducting Qubits". PhD thesis, University of Santa Barbara, Santa Barbara, USA, 2014. 62

[123] Z. Wang, "Quantum Microwave Engineering for the Protection and Readout of Superconducting Qubits". PhD thesis, Yale University, New Haven, USA, 2022. 62

[124] B. R. Johnson et al., "Quantum non-demolition detection of single microwave photons in a circuit", *Nature Physics* **6** (September, 2010) 663–667, doi:10.1038/nphys1710. 63

[125] L. Sun et al., "Tracking photon jumps with repeated quantum non-demolition parity measurements", *Nature* **511** (July, 2014) 444–448, doi:10.1038/nature13436. 63

[126] P. Reinhold, "Controlling Error-Correctable Bosonic Qubits". PhD thesis, Yale University, New Haven, USA, 2019. 65

[127] A. Koottandavida et al., "Erasure detection of a dual-rail qubit encoded in a double-post superconducting cavity", *Phys. Rev. Lett.* **132** (May, 2024) 180601, doi:10.1103/PhysRevLett.132.180601. 66

[128] G. S. Paraoanu, "Microwave-induced coupling of superconducting qubits", *Phys. Rev. B* **74** (Oct, 2006) 140504, doi:10.1103/PhysRevB.74.140504. 67

[129] C. Rigetti and M. Devoret, "Fully microwave-tunable universal gates in superconducting qubits with linear couplings and fixed transition frequencies", *Phys. Rev. B* **81** (Apr, 2010) 134507, doi:10.1103/PhysRevB.81.134507. 67

[130] N. Didier, J. Bourassa, and A. Blais, "Fast Quantum Nondemolition Readout by Parametric Modulation of Longitudinal Qubit-Oscillator Interaction", *Physical Review Letters* **115** (November, 2015) 203601, doi:10.1103/PhysRevLett.115.203601. 67

[131] S. Touzard et al., "Gated conditional displacement readout of superconducting qubits", *Physical Review Letters* **122** (February, 2019) 080502, doi:10.1103/PhysRevLett.122.080502. 67

[132] A. Eickbusch et al., "Fast universal control of an oscillator with weak dispersive coupling to a qubit", *Nature Physics* **18** (Dec, 2022) 1464–1469, doi:10.1038/s41567-022-01776-9. 68, 155, 168

[133] U. Réglade et al., "Quantum control of a cat qubit with bit-flip times exceeding ten seconds", *Nature* (2024) 1–6. 69, 220

[134] A. M. Eriksson et al., "Universal control of a bosonic mode via drive-activated native cubic interactions", *Nature Communications* **15** (March, 2024) 2512, doi:10.1038/s41467-024-46507-1. 69

[135] X. Xiaoet al., "A diagrammatic method to compute the effective hamiltonian of driven nonlinear oscillators", 2024. https://arxiv.org/abs/2304.13656. 69

[136] C. K. Law and J. H. Eberly, "Arbitrary control of a quantum electromagnetic field", *Phys. Rev. Lett.* **76** (Feb, 1996) 1055–1058, doi:10.1103/PhysRevLett.76.1055. 69, 166

[137] Y. Zhang et al., "Drive-induced nonlinearities of cavity modes coupled to a transmon ancilla", *Phys. Rev. A* **105** (Feb, 2022) 022423, doi:10.1103/PhysRevA.105.022423. 69, 113

[138] P. Reinhold et al., "Error-corrected gates on an encoded qubit", *Nature Physics* **16** (August, 2020) 822–826, doi:10.1038/s41567-020-0931-8. 70, 181, 182

[139] S. Rosenblum et al., "Fault-tolerant detection of a quantum error", *Science* **361** (2018), no. 6399, 266–270, doi:10.1126/science.aat3996, arXiv:https://www.science.org/doi/pdf/10.1126/science.aat3996. 70, 178, 182, 185

[140] Y. Y. Gao et al., "Entanglement of bosonic modes through an engineered exchange interaction", *Nature* **566** (February, 2019) 509–512, doi:10.1038/s41586-019-0970-4. 70, 71

[141] Y. Zhang et al., "Engineering bilinear mode coupling in circuit QED: Theory and experiment", *Physical Review A* **99** (January, 2019) 012314, doi:10.1103/PhysRevA.99.012314. 70, 86, 98, 99, 103, 143

[142] Wikipedia, "Magnetic flux quantum — Wikipedia, The Free Encyclopedia". https://en.wikipedia.org/wiki/Magnetic_flux_quantum, 2025. [Online; accessed 15-February-2025]. 74

[143] F. Lecocq, J. Claudon, O. Buisson, and P. Milman, "Nonlinear coupling between the two oscillation modes of a dc squid", *Phys. Rev. Lett.* **107** (Nov, 2011) 197002, doi:10.1103/PhysRevLett.107.197002. 75

[144] A. Kamal, J. Clarke, and M. H. Devoret, "Gain, directionality, and noise in microwave squid amplifiers: Input-output approach", *Phys. Rev. B* **86** (Oct, 2012) 144510, doi:10.1103/PhysRevB.86.144510. 75

[145] C. Purmessur, K. Chow, B. van Heck, and A. Kou, "Operation of a high-frequency, phase-slip qubit", 2025. https://arxiv.org/abs/2502.07043. 77

[146] E. Zakka-Bajjani et al., "Quantum superposition of a single microwave photon in two different 'colour'states", *Nature Physics* **7** (2011), no. 8, 599–603. 81, 137

[147] Y. Lu et al., "High-fidelity parametric beamsplitting with a parity-protected converter", *Nature Communications* **14** (September, 2023) 5767, doi:10.1038/s41467-023-41104-0. 81, 105, 108, 135, 145, 156, 157

[148] M. Grifoni and P. Hänggi, "Driven quantum tunneling", *Physics Reports* **304** (October, 1998) 229–354, doi:10.1016/S0370-1573(98)00022-2. 86, 95

[149] L. Verney et al., "Structural Instability of Driven Josephson Circuits Prevented by an Inductive Shunt", *Physical Review Applied* **11** (February, 2019) 024003, doi:10.1103/PhysRevApplied.11.024003. 86, 95, 99

[150] Y. Lu et al., "The numerical modeling of a microwave-driven josephson circuit", 2025. Manuscript in preparation. 86, 91, 145, 192

[151] Wikipedia, "Ansys HFSS — Wikipedia, The Free Encyclopedia". http://en.wikipedia.org/w/index.php?title=Ansys%20HFSS&oldid=1270088169, 2025. [Online; accessed 25-February-2025]. 87

[152] X. You, J. A. Sauls, and J. Koch, "Circuit quantization in the presence of time-dependent external flux", *Physical Review B* **99** (May, 2019) 174512, doi:10.1103/PhysRevB.99.174512. 87

[153] R.-P. Riwar and D. P. DiVincenzo, "Circuit quantization with time-dependent magnetic fields for realistic geometries", *npj Quantum Information* **8** (March, 2022) 36, doi:10.1038/s41534-022-00539-x. 87

[154] S. Ganjam, "Improving the Coherence of Superconducting Quantum Circuits through Loss Characterization and Design Optimization". PhD thesis, Yale University, New Haven, USA, 2023. 89

[155] F. C. Wellstood, C. Urbina, and J. Clarke, "Hot-electron effects in metals", *Physical Review B* **49** (March, 1994) 5942–5955, doi:10.1103/PhysRevB.49.5942. 92

[156] J.-H. Yeh et al., "Microwave attenuators for use with quantum devices below 100 mK", *Journal of Applied Physics* **121** (June, 2017) 224501, doi:10.1063/1.4984894. 92

[157] J. Bylander et al., "Noise spectroscopy through dynamical decoupling with a superconducting flux qubit", *Nature Physics* **7** (July, 2011) 565–570, doi:10.1038/nphys1994. 93, 134, 135, 179

[158] Wikipedia, "Brillouin zone — Wikipedia, The Free Encyclopedia". http://en.wikipedia.org/w/index.php?title=Brillouin%20zone&oldid=1247423038, 2025. [Online; accessed 16-February-2025]. 94

[159] N. Lambert et al., "QuTiP 5: The quantum toolbox in Python", 2024. doi:10.48550/arXiv.2412.04705, https://arxiv.org/abs/2412.04705. 94

[160] P. Guilminet al., "Dynamiqs: an open-source python library for gpu-accelerated and differentiable simulation of quantum systems", 2024. In progress. 94

[161] M. F. Dumaset al., "Unified picture of measurement-induced ionization in the transmon", February, 2024. arXiv:2402.06615 [quant-ph]. http://arxiv.org/abs/2402.06615. 97, 99

[162] Z. Xiao et al., "Perturbative Diagonalization for Time-Dependent Strong Interactions", *Physical Review Applied* **18** (August, 2022) 024009, doi:10.1103/PhysRevApplied.18.024009. 99

[163] Y. Y. Gao et al., "Programmable Interference between Two Microwave Quantum Memories", *Physical Review X* **8** (June, 2018) 021073, doi:10.1103/PhysRevX.8.021073. 103, 106, 110, 111

[164] K. Geerlings et al., "Demonstrating a driven reset protocol for a superconducting qubit", *Phys. Rev. Lett.* **110** (Mar, 2013) 120501, doi:10.1103/PhysRevLett.110.120501. 110, 111, 163

[165] N. E. Frattini et al., "Optimizing the nonlinearity and dissipation of a SNAIL Parametric Amplifier for dynamic range", *Physical Review Applied* **10** (November, 2018) 054020, doi:10.1103/PhysRevApplied.10.054020. 113, 145, 150

[166] B. J. Chapman et al., "High-On-Off-Ratio Beam-Splitter Interaction for Gates on Bosonically Encoded Qubits", *PRX Quantum* **4** (June, 2023) 020355, doi:10.1103/PRXQuantum.4.020355. 113, 121

[167] J. Emerson, R. Alicki, and K. Życzkowski, "Scalable noise estimation with random unitary operators", *Journal of Optics B: Quantum and Semiclassical Optics* **7** (sep, 2005) S347, doi:10.1088/1464-4266/7/10/021. 113

[168] K. Boone, A. Carignan-Dugas, J. J. Wallman, and J. Emerson, "Randomized benchmarking under different gate sets", *Phys. Rev. A* **99** (Mar, 2019) 032329, doi:10.1103/PhysRevA.99.032329. 113

[169] T. J. Proctor et al., "Direct randomized benchmarking for multiqubit devices", *Phys. Rev. Lett.* **123** (Jul, 2019) 030503, doi:10.1103/PhysRevLett.123.030503. 114

[170] L. Chirolli, G. Burkard, S. Kumar, and D. P. DiVincenzo, "Superconducting resonators as beam splitters for linear-optics quantum computation", *Phys. Rev. Lett.* **104** (Jun, 2010) 230502, doi:10.1103/PhysRevLett.104.230502. 118

[171] E. Zakka-Bajjani et al., "Quantum superposition of a single microwave photon in two different 'colour'states", *Nature Physics* **7** (2011), no. 8, 599–603, doi:10.1038/nphys2035. 118

[172] Y. Wu, S. Kolkowitz, S. Puri, and J. D. Thompson, "Erasure conversion for fault-tolerant quantum computing in alkaline earth rydberg atom arrays", *Nature Communications* **13** (Aug, 2022) 4657, doi:10.1038/s41467-022-32094-6. 118

[173] C. Macklin et al., "A near–quantum-limited josephson traveling-wave parametric amplifier", *Science* **350** (2015), no. 6258, 307–310, doi:10.1126/science.aaa8525, arXiv:https://www.science.org/doi/pdf/10.1126/science.aaa8525. 119

[174] N. E. Frattini et al., "3-wave mixing Josephson dipole element", *Applied Physics Letters* **110** (May, 2017) 222603, doi:10.1063/1.4984142. 120

[175] S. J. de Graaf, "Microwave beamsplitters for oscillator-based quantum information processing". Phd thesis, Yale University, Nov, 2024. Pgs. 55-56. 121, 148

[176] R. Lescanne et al., "Exponential suppression of bit-flips in a qubit encoded in an oscillator", *Nature Physics* **16** (2020), no. 5, 509–513, doi:10.1038/s41567-020-0824-x. 125, 135

[177] A. Miano et al., "Frequency-tunable Kerr-free three-wave mixing with a gradiometric SNAIL", *Applied Physics Letters* **120** (05, 2022) 184002, doi:10.1063/5.0083350. 125

[178] Y. Yeet al., "Near-ultrastrong nonlinear light-matter coupling in superconducting circuits", 2024. https://arxiv.org/abs/2404.19199. 125

[179] B. Bhandariet al., "Symmetrically threaded squids as next generation kerr-cat qubits", 2024. https://arxiv.org/abs/2405.11375. 125, 134

[180] M. Hosoya et al., "Quantum flux parametron: a single quantum flux device for josephson supercomputer", *IEEE Transactions on Applied Superconductivity* **1** (1991), no. 2, 77–89, doi:10.1109/77.84613. 125

[181] S. Kim et al., "Enhanced coherence of all-nitride superconducting qubits epitaxially grown on silicon substrate", *Communications Materials* **2** (September, 2021) 98, doi:10.1038/s43246-021-00204-4. 126

[182] V. Guptaet al., "Low loss lumped-element inductors made from granular aluminum", 2024. https://arxiv.org/abs/2411.12611. 126

[183] M. Peruzzo et al., "Geometric superinductance qubits: Controlling phase delocalization across a single josephson junction", *PRX Quantum* **2** (November, 2021) doi:10.1103/prxquantum.2.040341. 126

[184] A. Maitiet al., "A linear quantum coupler for clean bosonic control", 2025. https://arxiv.org/abs/2501.18025. 127

[185] W. C. Smith et al., "Quantization of inductively shunted superconducting circuits", *Phys. Rev. B* **94** (Oct, 2016) 144507, doi:10.1103/PhysRevB.94.144507. 128

[186] F. Yan et al., "Rotating-frame relaxation as a noise spectrum analyser of a superconducting qubit undergoing driven evolution", *Nature Communications* **4** (2013), no. 1, 2337. 134, 135

[187] J. Braumüller et al., "Characterizing and optimizing qubit coherence based on squid geometry", *Phys. Rev. Appl.* **13** (May, 2020) 054079, doi:10.1103/PhysRevApplied.13.054079. 134, 135

[188] J. M. Martinis et al., "Decoherence of a superconducting qubit due to bias noise", *Phys. Rev. B* **67** (Mar, 2003) 094510, doi:10.1103/PhysRevB.67.094510. 135

[189] P. Groszkowski et al., "Coherence properties of the 0-πqubit", *New Journal of Physics* **20** (April, 2018) 043053, doi:10.1088/1367-2630/aab7cd. 135

[190] A. Grimm et al., "Stabilization and operation of a kerr-cat qubit", *Nature* **584** (August, 2020) 205–209, doi:10.1038/s41586-020-2587-z. 135

[191] F. Valadares et al., "On-demand transposition across light-matter interaction regimes in bosonic cqed", *Nature Communications* **15** (Jul, 2024) 5816, doi:10.1038/s41467-024-50201-7. 135

[192] U. Goldblatt et al., "Recovering quantum coherence of a cavity qubit coupled to a noisy ancilla through real-time feedback", *Phys. Rev. X* **14** (Dec, 2024) 041056, doi:10.1103/PhysRevX.14.041056. 135

[193] D. G. Atanasovaet al., "In-situ tunable interaction with an invertible sign between a fluxonium and a post cavity", 2024. https://arxiv.org/abs/2409.07612. 135

[194] T. Noh et al., "Strong parametric dispersive shifts in a statically decoupled two-qubit cavity qed system", *Nature Physics* **19** (2023), no. 10, 1445–1451, doi:10.1038/s41567-023-02107-2. 137

[195] Wikipedia, "Jacobi–Anger expansion — Wikipedia, The Free Encyclopedia". http://en.wikipedia.org/w/index.php?title=Jacobi%E2%80%93Anger%20expansion&oldid=1223168846, 2025. [Online; accessed 21-February-2025]. 139

[196] M. I. Dykman, "Fluctuating nonlinear oscillators: From nanomechanics to quantum superconducting circuits", pp. 165–197. Oxford University Press, , 2012. 143

[197] A. Miano et al., "Hamiltonian extrema of an arbitrary flux-biased josephson circuit", *PRX Quantum* **4** (Aug, 2023) 030324, doi:10.1103/PRXQuantum.4.030324. 150, 151

[198] V. Sivak et al., "Kerr-free three-wave mixing in superconducting quantum circuits", *Phys. Rev. Appl.* **11** (May, 2019) 054060, doi:10.1103/PhysRevApplied.11.054060. 150

[199] A. Maiti et al., "A Linear Quantum Coupler for Clean Bosonic Control", February, 2025. arXiv:2501.18025 [quant-ph]. doi:10.48550/arXiv.2501.18025, http://arxiv.org/abs/2501.18025. 154

[200] T. Hillmann et al., "Universal gate set for continuous-variable quantum computation with microwave circuits", *Phys. Rev. Lett.* **125** (Oct, 2020) 160501, doi:10.1103/PhysRevLett.125.160501. 155

[201] A. M. Eriksson et al., "Universal control of a bosonic mode via drive-activated native cubic interactions", *Nature Communications* **15** (March, 2024) 2512, doi:10.1038/s41467-024-46507-1. 155

[202] A. Z. Dinget al., "Quantum control of an oscillator with a kerr-cat qubit", 2024. https://arxiv.org/abs/2407.10940. 155

[203] A. Remm et al., "Intermodulation distortion in a josephson traveling-wave parametric amplifier", *Phys. Rev. Appl.* **20** (Sep, 2023) 034027, doi:10.1103/PhysRevApplied.20.034027. 155

[204] P. A. Springet al., "Fast multiplexed superconducting qubit readout with intrinsic purcell filtering", 2024. https://arxiv.org/abs/2409.04967. 155

[205] D. Gottesman, A. Kitaev, and J. Preskill, "Encoding a qubit in an oscillator", *Phys. Rev. A* **64** (Jun, 2001) 012310, doi:10.1103/PhysRevA.64.012310. 156

[206] I. Rojkov et al., "Two-qubit operations for finite-energy gottesman-kitaev-preskill encodings", *Phys. Rev. Lett.* **133** (Sep, 2024) 100601, doi:10.1103/PhysRevLett.133.100601. 156

[207] I. Pietikäinen et al., "Strategies and trade-offs for controllability and memory time of ultra-high-quality microwave cavities in circuit quantum electrodynamics", *PRX Quantum* **5** (Oct, 2024) 040307, doi:10.1103/PRXQuantum.5.040307. 166

[208] L. DiCarlo et al., "Demonstration of two-qubit algorithms with a superconducting quantum processor", *Nature* **460** (2009), no. 7252, 240–244, doi:10.1038/nature08121. 171

[209] J. M. Martinis and M. R. Geller, "Fast adiabatic qubit gates using only σ_z control", *Phys. Rev. A* **90** (Aug, 2014) 022307, doi:10.1103/PhysRevA.90.022307. 171

[210] F. Yan et al., "Distinguishing coherent and thermal photon noise in a circuit quantum electrodynamical system", *Phys. Rev. Lett.* **120** (Jun, 2018) 260504, doi:10.1103/PhysRevLett.120.260504. 171

[211] F. Swiadeket al., "Enhancing dispersive readout of superconducting qubits through dynamic control of the dispersive shift: Experiment and theory", 2023. https://arxiv.org/abs/2307.07765. 171

[212] F. Valadares et al., "On-demand transposition across light-matter interaction regimes in bosonic cQED", *Nature Communications* **15** (July, 2024) 5816, doi:10.1038/s41467-024-50201-7. 171

[213] Z. Xiao et al., "Perturbative diagonalization for time-dependent strong interactions", *Phys. Rev. Appl.* **18** (Aug, 2022) 024009, doi:10.1103/PhysRevApplied.18.024009. 173, 174, 183

[214] T. Noh et al., "Strong parametric dispersive shifts in a statically decoupled two-qubit cavity QED system", *Nature Physics* **19** (October, 2023) 1445–1451, doi:10.1038/s41567-023-02107-2. 174

[215] S. Bravyi, D. P. DiVincenzo, and D. Loss, "Schrieffer–Wolff transformation for quantum many-body systems", *Annals of Physics* **326** (2011), no. 10, 2793–2826. 175

[216] S. Richer, "Perturbative analysis of two-qubit gates on transmon qubits", Master's thesis, RWTH Aachen University, 2013. 175

[217] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, "Simple pulses for elimination of leakage in weakly nonlinear qubits", *Physical Review Letters* **103** (September, 2009) 110501, doi:10.1103/PhysRevLett.103.110501. 178

[218] E. Barnes et al., "Dynamically corrected gates from geometric space curves", *Quantum Science and Technology* **7** (April, 2022) 023001, doi:10.1088/2058-9565/ac4421. 179, 201

[219] D. K. Weiss, "qontrol: Quantum optimal control based on dynamiqs, diffrax and jax", 2024. under development. 186

[220] B. Royer, S. Singh, and S. Girvin, "Encoding qubits in multimode grid states", *PRX Quantum* **3** (Mar, 2022) 010335, doi:10.1103/PRXQuantum.3.010335. 196

[221] M.-A. Lemondeet al., "Hardware-efficient fault tolerant quantum computing with bosonic grid states in superconducting circuits", 2024. https://arxiv.org/abs/2409.05813. 196

[222] Y. Xu, Y. Wang, C. Vuillot, and V. V. Albert, "Letting the tiger out of its cage: bosonic coding without concatenation", 2024. https://arxiv.org/abs/2411.09668. 196

[223] A. C. C. de Albornozet al., "Oscillatory dissipative tunneling in an asymmetric double-well potential", 2024. https://arxiv.org/abs/2409.13113. 196

[224] C. S. Wang et al., "Efficient multiphoton sampling of molecular vibronic spectra on a superconducting bosonic processor", *Phys. Rev. X* **10** (Jun, 2020) 021060, doi:10.1103/PhysRevX.10.021060. 196

[225] B. Royer, S. Singh, and S. Girvin, "Stabilization of Finite-Energy Gottesman-Kitaev-Preskill States", *Physical Review Letters* **125** (December, 2020) 260509, doi:10.1103/PhysRevLett.125.260509. 198, 199

[226] L. Gravina, F. Minganti, and V. Savona, "Critical Schrödinger Cat Qubit", *PRX Quantum* **4** (June, 2023) 020337, doi:10.1103/PRXQuantum.4.020337. 198

[227] J. Venkatraman, X. Xiao, R. G. Cortiñas, and M. H. Devoret, "On the static effective Lindbladian of the squeezed Kerr oscillator", September, 2022. arXiv:2209.11193 [cond-mat, physics:physics, physics:quant-ph]. http://arxiv.org/abs/2209.11193. 198

[228] V. V. Albert et al., "Performance and structure of single-mode bosonic codes", *Phys. Rev. A* **97** (Mar, 2018) 032346, doi:10.1103/PhysRevA.97.032346. 198, 201

[229] P. Leviant, Q. Xu, L. Jiang, and S. Rosenblum, "Quantum capacity and codes for the bosonic loss-dephasing channel", *Quantum* **6** (September, 2022) 821, doi:10.22331/q-2022-09-29-821. 198

[230] T. A. Brun, "A simple model of quantum trajectories", *American Journal of Physics* **70** (2002), no. 7, 719–737. 199

[231] E. Knill and R. Laflamme, "Theory of quantum error-correcting codes", *Physical Review A* **55** (1997), no. 2, 900. 200, 201

[232] D. Gottesman, A. Kitaev, and J. Preskill, "Encoding a qubit in an oscillator", *Phys. Rev. A* **64** (Jun, 2001) 012310, doi:10.1103/PhysRevA.64.012310. 200, 201

[233] Z. Leghtas et al., "Hardware-efficient autonomous quantum memory protection", *Phys. Rev. Lett.* **111** (Sep, 2013) 120501, doi:10.1103/PhysRevLett.111.120501. 200, 221

[234] A. L. Grimsmo, J. Combes, and B. Q. Baragiola, "Quantum computing with rotation-symmetric bosonic codes", *Phys. Rev. X* **10** (Mar, 2020) 011058, doi:10.1103/PhysRevX.10.011058. 200

[235] V. V. Albert, J. P. Covey, and J. Preskill, "Robust encoding of a qubit in a molecule", *Phys. Rev. X* **10** (Sep, 2020) 031050, doi:10.1103/PhysRevX.10.031050. 200

[236] J. A. Gross, "Designing codes around interactions: The case of a spin", *Physical Review Letters* **127** (2021), no. 1, 010504. 200

[237] S. Omanakuttan and J. A. Gross, "Multispin clifford codes for angular momentum errors in spin systems", *Phys. Rev. A* **108** (Aug, 2023) 022424, doi:10.1103/PhysRevA.108.022424. 200

[238] S. Omanakuttan et al., "Fault-tolerant quantum computation using large spin-cat codes", *PRX Quantum* **5** (Jun, 2024) 020355, doi:10.1103/PRXQuantum.5.020355. 200

[239] S. P. Jain, E. R. Hudson, W. C. Campbell, and V. V. Albert, "Æ codes", *arXiv preprint arXiv:2311.12324* (2023). 200

[240] Z. Ni et al., "Beating the break-even point with a discrete-variable-encoded logical qubit", *Nature* **616** (April, 2023) 56–60, doi:10.1038/s41586-023-05784-4. 200, 221

[241] J.-M. Lihm, K. Noh, and U. R. Fischer, "Implementation-independent sufficient condition of the Knill-Laflamme type for the autonomous protection of logical qudits by strong engineered dissipation", *Physical Review A* **98** (July, 2018) 012317, doi:10.1103/PhysRevA.98.012317. 200, 221, 223

[242] J. Lebreuillyet al., "Autonomous quantum error correction and quantum computation", *arXiv:2103.05007 [quant-ph]* (March, 2021). arXiv: 2103.05007. 200

[243] I. L. Chuang, D. W. Leung, and Y. Yamamoto, "Bosonic quantum codes for amplitude damping", *Phys. Rev. A* **56** (Aug, 1997) 1114–1125, doi:10.1103/PhysRevA.56.1114. 200, 210

[244] D. A. Lidar, "Review of Decoherence-Free Subspaces, Noiseless Subsystems, and Dynamical Decoupling", pp. 295–354. John Wiley & Sons, Ltd, 2014. doi:https://doi.org/10.1002/9781118742631.ch11. 201

[245] G. Pantaleoni, B. Q. Baragiola, and N. C. Menicucci, "Modular Bosonic Subsystem Codes", *Physical Review Letters* **125** (July, 2020) 040501, doi:10.1103/PhysRevLett.125.040501. arXiv:1907.08210 [quant-ph]. 203

[246] P. Hayden and J. Sorce, "A canonical hamiltonian for open quantum systems", *Journal of Physics A: Mathematical and Theoretical* **55** (2022), no. 22, 225302. 207

[247] J. M. Gertleret al., "Protecting a Bosonic Qubit with Autonomous Quantum Error Correction", *arXiv:2004.09322 [quant-ph]* (April, 2020). arXiv: 2004.09322. 214

[248] R. Gautier, A. Sarlette, and M. Mirrahimi, "Combined Dissipative and Hamiltonian Confinement of Cat Qubits", *arXiv:2112.05545 [quant-ph]* (April, 2022). arXiv: 2112.05545. 218

[249] S. Puri et al., "Stabilized cat in a driven nonlinear cavity: A fault-tolerant error syndrome detector", *Phys. Rev. X* **9** (Oct, 2019) 041009, doi:10.1103/PhysRevX.9.041009. 220, 228

[250] H. Putterman et al., "Stabilizing a bosonic qubit using colored dissipation", *Phys. Rev. Lett.* **128** (Mar, 2022) 110502, doi:10.1103/PhysRevLett.128.110502. 220

[251] A. Grimm et al., "Stabilization and operation of a Kerr-cat qubit", *Nature* **584** (August, 2020) 205–209, doi:10.1038/s41586-020-2587-z. 220

[252] N. E. Frattiniet al., "The squeezed kerr oscillator: spectral kissing and phase-flip robustness", 2022. https://arxiv.org/abs/2209.03934. 220

[253] Z. Leghtas et al., "Confining the state of light to a quantum manifold by engineered two-photon loss", *Science* **347** (2015), no. 6224, 853–857, doi:10.1126/science.aaa2085, arXiv:https://www.science.org/doi/pdf/10.1126/science.aaa2085. 220

[254] L. Hu et al., "Quantum error correction and universal gate set operation on a binomial bosonic logical qubit", *Nature Physics* **15** (2019), no. 5, 503–508, doi:10.1038/s41567-018-0414-3. 221

[255] Q. Xuet al., "Autonomous quantum error correction and fault-tolerant quantum computation with squeezed cat qubits", October, 2022. arXiv:2210.13406 [quant-ph]. http://arxiv.org/abs/2210.13406. 225

[256] S. Puri et al., "Bias-preserving gates with stabilized cat qubits", *Science Advances* **6** (2020), no. 34, eaay5901, doi:10.1126/sciadv.aay5901, arXiv:https://www.science.org/doi/pdf/10.1126/sciadv.aay5901. 228