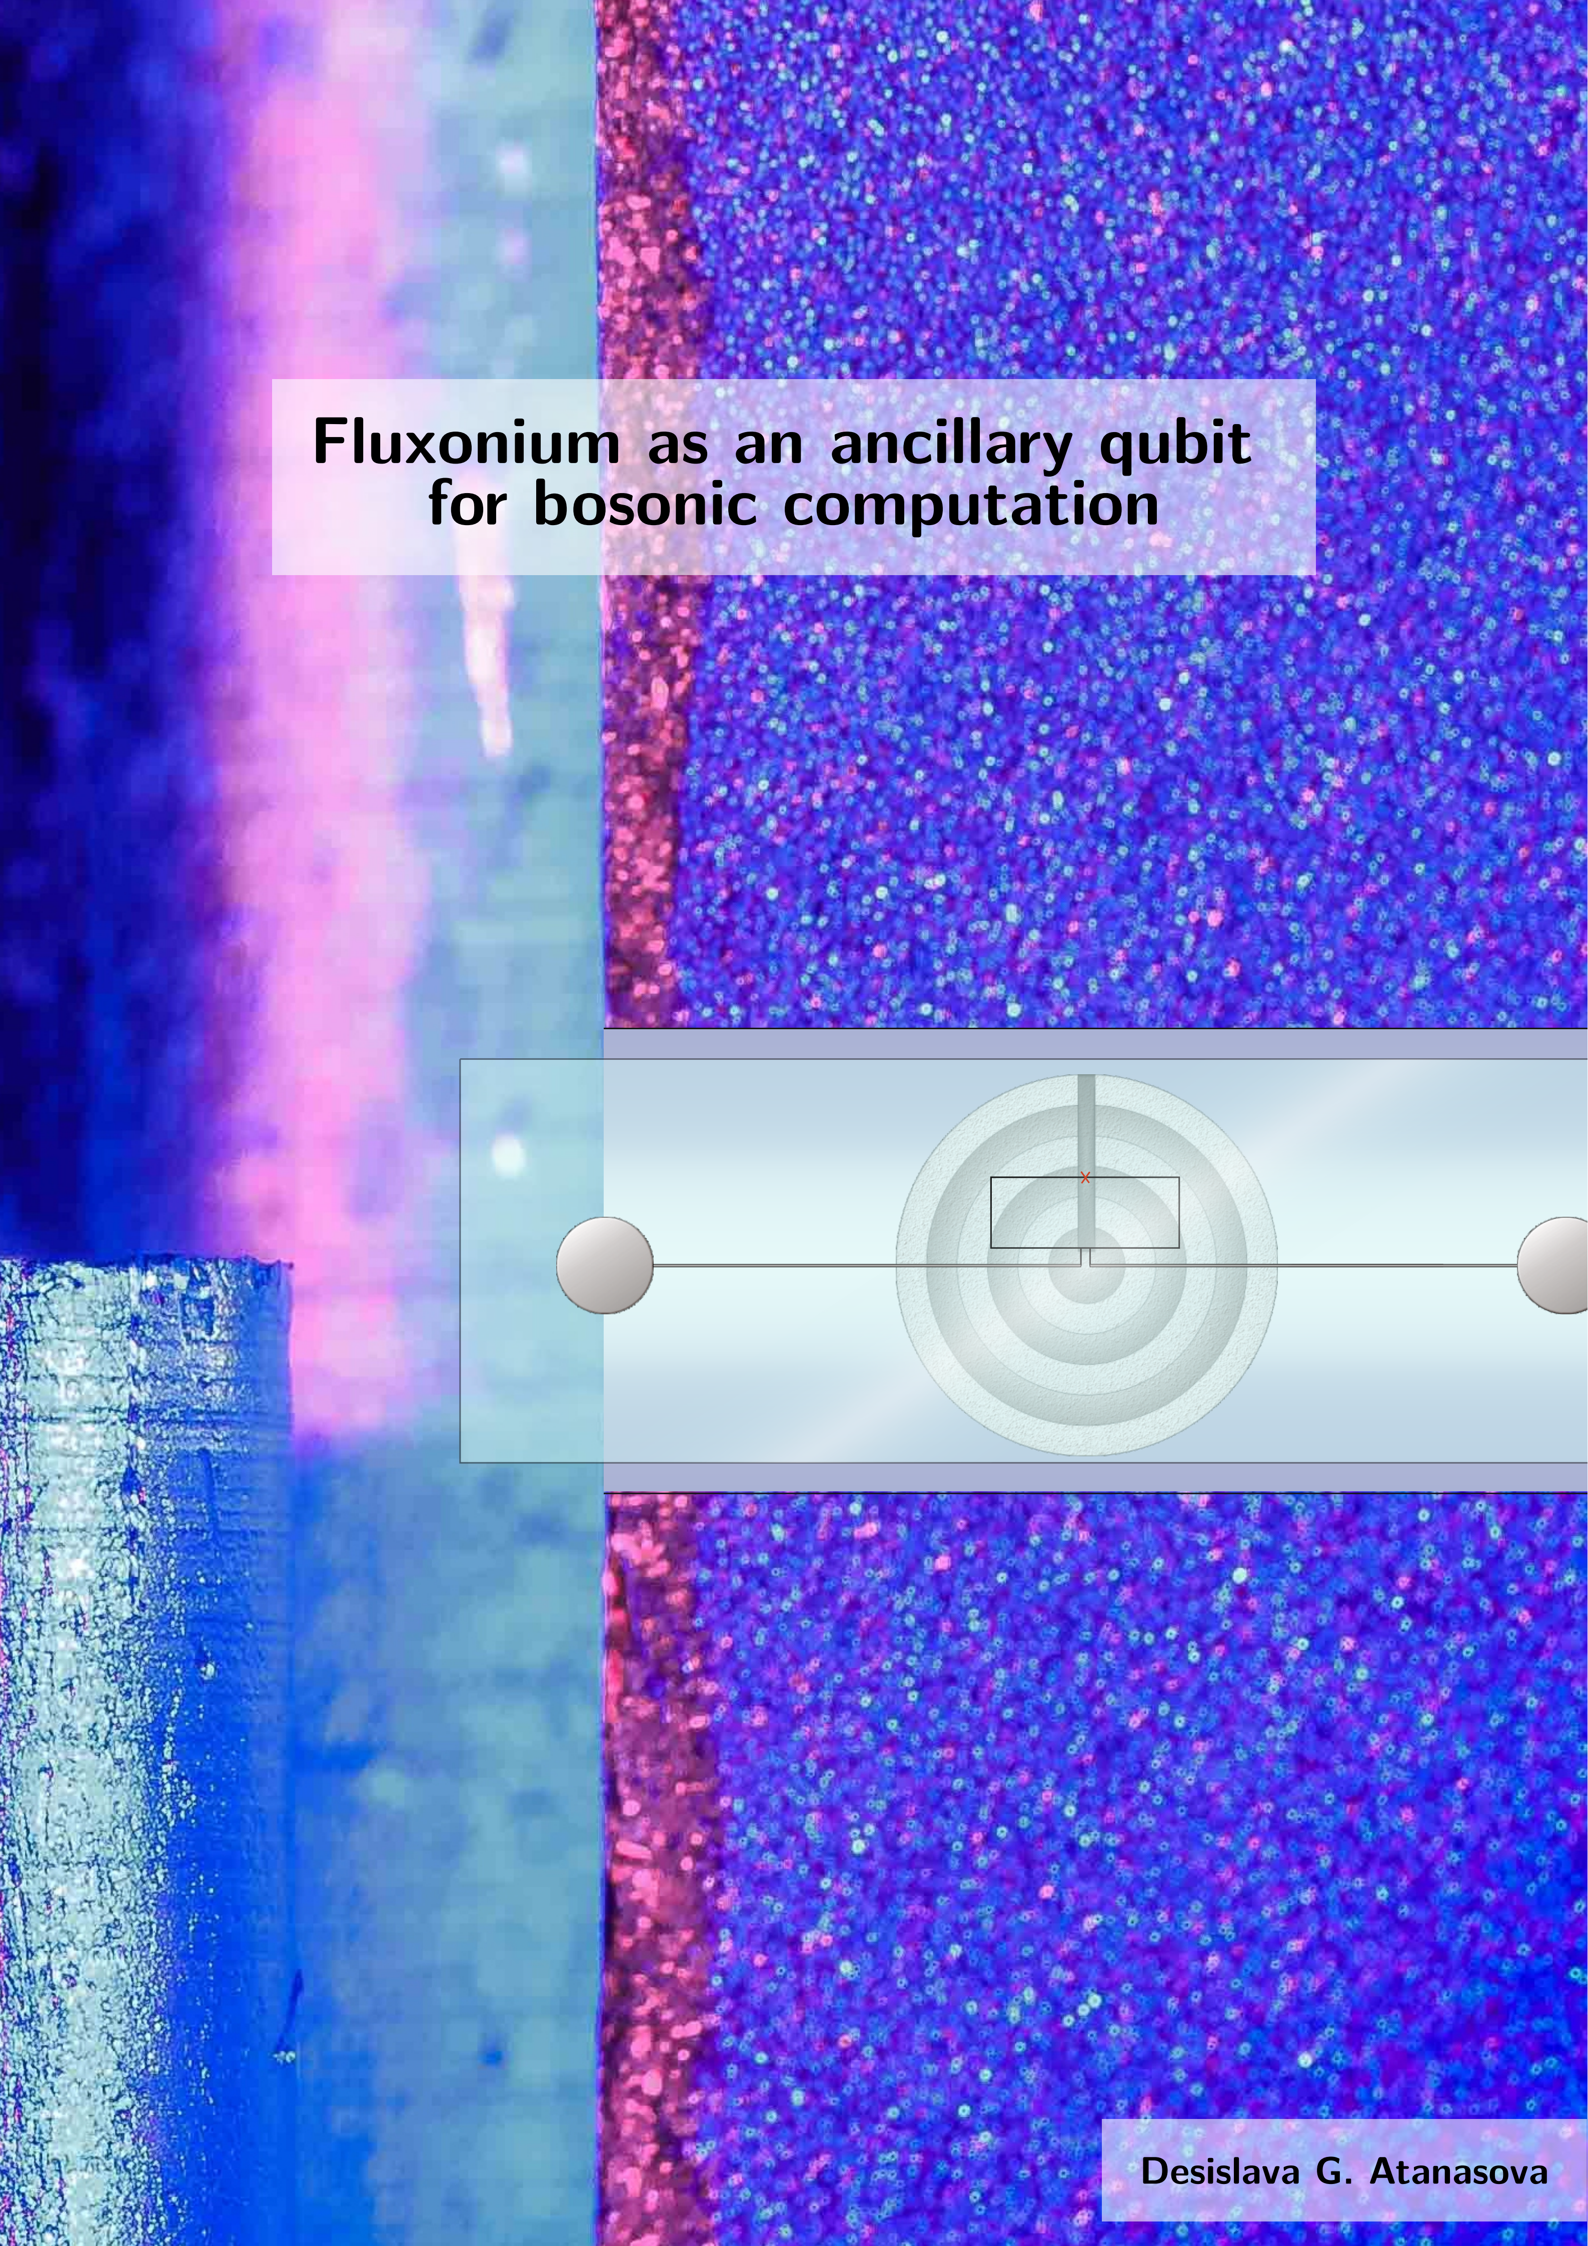

Fluxonium as an ancillary qubit for bosonic computation

Desislava G. Atanasova

Fluxonium as an ancillary qubit for bosonic computation

Desislava Gancheva Atanasova, 11705475

Innsbruck, January 2026

Dissertation

submitted to the Faculty of Mathematics, Computer Science and Physics of the University of Innsbruck in partial fulfillment of the requirements for the academic degree

Doctor of Philosophy

Doctor of Philosophy - Doctoral Program in Physics

Carried out at the Institute of Experimental Physics under the supervision of

Univ.-Prof. Dr. Gerhard Kirchmair

with the co-supervisors:

Univ.-Prof. Dr. Helmut Ritsch

Prof. Dr. Ioan Pop

Abstract

Bosonic modes provide a hardware-efficient route toward quantum information processing by exploiting the large Hilbert space of harmonic oscillators. However, in most bosonic architectures, the coherence of the encoded information is fundamentally limited by dephasing induced by an ancillary qubit used for control and readout. In this thesis, I present a novel superconducting platform that mitigates this limitation by combining a high-coherence three-dimensional post cavity with a fluxonium qubit whose interaction with the cavity can be tuned in situ.

I develop and validate a hybrid theoretical framework for modeling highly anharmonic superconducting circuits beyond the lumped-element approximation, incorporating flux quantization while retaining compatibility with distributed electromagnetic modes. This approach is implemented in the *scQubits* simulation package and enables flexible modeling of complex circuit architectures. Experimentally, I realize and characterize a fluxonium-cavity system in which the dispersive coupling can be continuously tuned to positive, negative, or vanishing values. This capability allows strong interaction during gate operations while effectively decoupling the bosonic memory from the ancilla during idle periods, thereby suppressing qubit-induced dephasing. A quantitative threshold for negligible ancilla-induced dephasing is established based on the relative magnitude of the dispersive shift and the qubit decay rate.

I further investigate the prospects of this platform for bosonic state engineering and error-protected encodings, including dissipative cat qubits and Kerr-induced cat-state generation. While fast flux control enables rapid tuning of the interaction strengths, practical challenges, most notably flux noise and reduced qubit coherence under AC flux biasing, limit the achievable performance in the present implementation. Finally, I outline concrete pathways for future improvements. Overall, this work establishes a scalable and versatile platform for bosonic quantum information processing and highlights the critical role of tunable interactions and decoupling schemes in advancing fault-tolerant quantum technologies.

Kurzfassung

Bosonische Moden bieten einen hardwareeffizienten Ansatz für die Quanteninformationsverarbeitung, indem sie den großen Hilbertraum harmonischer Oszillatoren ausnutzen. In den meisten bosonischen Architekturen ist die Kohärenz der kodierten Information jedoch grundsätzlich durch Dephasierung begrenzt, die von einem für Kontrolle und Auslese verwendeten Hilfsqubit induziert wird. In dieser Dissertation präsentiere ich eine neuar-

tige supraleitende Plattform, die diese Einschränkung adressiert, indem sie einen hochkohärenten dreidimensionalen Mikrowellenresonator mit einem Fluxonium-Qubit kombiniert, dessen Wechselwirkung mit dem Resonator *in situ* einstellbar ist.

Ich entwickle und validiere ein hybrides theoretisches Modell zum stark anharmonischen supraleitenden Schaltungen jenseits der Gültigkeit der Lumped-Element-Approximation, das die Flussquantisierung berücksichtigt und gleichzeitig mit verteilten elektromagnetischen Moden kompatibel bleibt. Dieser Ansatz wird mit dem Simulationspaket *scQubits* implementiert und ermöglicht eine flexible Modellierung komplexer Schaltungsarchitekturen. Experimentell realisiere und charakterisiere ich ein Fluxonium-Resonator-System, in dem die dispersive Kopplung kontinuierlich auf positive, negative oder verschwindende Werte eingestellt werden kann. Diese Fähigkeit erlaubt eine starke Wechselwirkung während Gatter-Operationen, der bosonische Speicher in Leerlaufphasen aber effektiv vom Hilfsqubit entkoppelt werden kann, wodurch qubitinduzierte Dephasierung unterdrückt wird. Auf Grundlage des Verhältnisses zwischen dispersiver Verschiebung und Qubit-Zerfallsrate wird eine quantitative Schwelle für vernachlässigbare Ancilla-induzierte Dephasierung bestimmt.

Darüber hinaus untersuche ich die Eignung dieser Plattform für die Erzeugung bosonischer Zustände und fehlertoleranter Kodierungen, einschließlich dissipativer "Cat-Qubits" und Kerr-induzierter Katzen-Zustände. Obwohl schnelle Flusskontrolle eine rasche Anpassung der Wechselwirkungsstärken ermöglicht, begrenzen praktische Herausforderungen, insbesondere Flussrauschen und eine reduzierte Qubit-Kohärenz unter AC-Flussbias, die erreichbare Leistung in der aktuellen Implementierung. Abschließend skizziere ich konkrete Ansatzpunkte für zukünftige Verbesserungen. Insgesamt etabliert diese Arbeit eine skalierbare und vielseitige Plattform für die bosonische Quanteninformationsverarbeitung und unterstreicht die zentrale Bedeutung einstellbarer Wechselwirkungen und effektiver Entkopplungsschemata für den Fortschritt fehlertoleranter Quantentechnologien.

Contents

CHAPTER 1

Introduction

Quantum mechanics arose in the early twentieth century as a response to a series of experimental discoveries that defied classical explanation. Phenomena such as blackbody radiation, the photoelectric effect, and atomic line spectra forced physicists to reconsider the foundations of energy, light, and matter. Attempting to explain the spectrum of thermal radiation, Planck proposed that energy is exchanged in discrete packets, or *"quanta"*, rather than continuously[1]. This idea marked the beginning of a conceptual revolution and is considered the birth of quantum theory. Only a few years later, Albert Einstein extended this notion to light itself, showing that the photoelectric effect could be understood if light were composed of particles–photons–with energy proportional to their frequency[2]. Niels Bohr then applied quantization to the structure of the hydrogen atom, attributing the stability of atomic orbits and the discreteness of spectral lines to the existence of quantized energy levels[3].

These early developments established the foundations of the "old quantum theory", a collection of rules that successfully described certain phenomena but lacked internal consistency. The full theoretical framework emerged in the mid-1920s, when Werner Heisenberg's matrix mechanics[4] and Erwin Schrödinger's wave mechanics[5] provided a consistent mathematical foundation for describing systems whose behavior could no longer be reconciled with classical intuition.

Over the following decades, quantum mechanics became a cornerstone of modern science. It enabled the development of quantum electrodynamics, the understanding of superconductivity[6–8] and semiconductor physics[9], and ultimately the digital revolution of the late twentieth century, which appeared after the transistor's invention. Around the same time, John Clarke, Michel Devoret, and John M. Martinis demonstrated quantized energy levels in a Josephson junction[10], for which they were awarded the 2025 Nobel Prize in Physics.

Today, the same principles underpin the rapidly advancing fields of quantum information science and quantum technologies. Concepts, such as entanglement, superposition, and measurement backaction, have become practical resources for communication, computation, and sensing. These technologies now hold the promise of tackling problems that are classically intractable. A major milestone came with Peter Shor's factoring algorithm[11], the first concrete demonstration that a quantum computer could surpass classical capabilities on a problem of real practical importance, sparking global interest in quantum computation. In the years that followed, the first small-scale quantum computers began to emerge[12, 13]. While trapped ions and cold atoms were amongst the first platforms used for running quantum operations, the development of the transmon qubit brought unprecedented stability and control to superconducting platforms, transforming quantum processors from theoretical constructs into scalable experimental systems. This progress drew major technology companies, university spin-offs, and even national governments[14] into an international effort to build a fully fault-tolerant quantum computer. These advances eventually paved the way for experimental demonstrations of "quantum advantage"[15] and for the first real-world applications powered by multi-qubit processors[16–21].

Bosonic modes have recently emerged as a compelling alternative to conventional multi-qubit architectures. By harnessing the infinite-dimensional Hilbert space of a single harmonic or weakly anharmonic oscillator, bosonic encodings offer a hardware-efficient route to realizing logical qubits[22–24] and performing quantum simulations[25] with reduced overhead. Unlocking this potential, however, requires both a highly coherent memory mode and a suitable nonlinear ancilla to enable universal control and readout. In most existing implementations, the ancilla becomes the primary bottleneck: even infrequent thermal excitations or relaxation events can imprint uncontrolled phase shifts onto the bosonic mode through dispersive coupling, leading to rapid dephasing of the encoded state.

This challenge motivates the central question of this thesis: Can one engineer a bosonic quantum information platform in which the ancilla-induced dephasing is minimized or even eliminated, without compromising control capabilities? Addressing this question requires not only the development of appropriate theoretical tools but also the design, fabrication, and characterization of a new experimental system capable of supporting such a regime.

1.1 Overview of thesis

This thesis begins with a brief introduction to the fields of quantum information, quantum error correction, and circuit Quantum ElectroDynamics (QED). Later in Chapter 2, I bridge these topics by describing some of the most popular encodings used in quantum error correction with superconducting circuits. Chapter 3 describes the common techniques in circuit QED for modeling a setup, including their applicability to a toy model used in this thesis. It also suggests a new numerical diagonalization method to overcome the limitations of the other methods.

The following chapter 4 introduces the design and fabrication used to build the setup from this thesis, which consists of a post cavity, fluxonium, and a flux hose. Specific attention is paid to the improvements made during my PhD, including a more stable cavity etching process, a simplified magnetic flux hose design, and improved shielding.

Chapter 5 presents the full setup characterization, addresses the challenge of ancilla-induced dephasing, and offers a solution to this challenge. Additionally, it elaborates on the possibility of combining this setup with a Superconducting Nonlinear Asymmetric Inductive eLement (SNAIL) to engineer a two-photon dissipation, suitable for dissipative cat qubits. In Chapter 7, I integrate fast flux control and utilize the flux-tunable self-Kerr shift of the cavity to create a nonclassical state. Therefore, I demonstrate universal control of the system at any flux bias point.

CHAPTER 2

Theory

This chapter presents the theoretical framework underlying quantum information and quantum error correction, with a particular focus on their implementation in superconducting circuits. The concepts introduced here provide the foundation for the experimental work discussed in subsequent chapters. Parts of this chapter are also covered in numerous textbooks, review articles, and theses (e.g. [26–32]). I begin by outlining the fundamental principles of quantum information, describing the physical carriers of quantum states and their unwanted coupling to the environment, which motivates the need for quantum error correction.

Next, in Section 2.2, I introduce superconducting circuits as engineered artificial atoms governed by the laws of quantum mechanics. By quantizing basic electrical elements (inductors, capacitors, and Josephson junctions), we can build nonlinear oscillators capable of encoding and manipulating quantum information. Building on this foundation, I discuss the effective Hamiltonians that describe these systems, with particular emphasis on fluxonium qubits. The interaction between qubits and resonators is then explored, focusing on the dispersive regime that enables high-fidelity quantum control and readout.

Section 2.3 is devoted to the theoretical principles of quantum error correction (QEC) in the context of superconducting hardware. We distinguish between discrete-variable codes, which rely on multi-qubit redundancy, and continuous-variable (bosonic) codes, which encode information in the Hilbert space of harmonic oscillators. Together, these discussions establish the theoretical basis for understanding how superconducting circuits can function as a scalable platform for fault-tolerant quantum computation.

2.1 Quantum information processing

2.1.1 Information carriers

2.1.1.1 Two-level systems: Qubits

The *qubit* (quantum bit) is a fundamental part of quantum information, analogous to the classical bit, but exhibiting uniquely quantum-mechanical properties. It is realized with a physical two-level system with two orthogonal states, $|0\rangle$ and $|1\rangle$. Unlike a classical bit, which can take one of two discrete states, a qubit can exist in a coherent superposition of both basis states simultaneously. In general, its state is described by a normalized vector $|\psi\rangle$ in a two-dimensional complex Hilbert space $\mathscr{H}_2$

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle\,, \tag{2.1}$$

where $\alpha, \beta \in \mathbb{C}$ are complex probability amplitudes satisfying the normalization condition $|\alpha|^2 + |\beta|^2 = 1$. The computational basis states are represented as

$$|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad |1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}.$$

Due to the global phase invariance of quantum states, it is sufficient to describe the qubit using two real parameters, ϕ and θ

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\phi}\sin\left(\frac{\theta}{2}\right)|1\rangle\,, \tag{2.2}$$

where $\theta \in [0, \pi]$ and $\phi \in [0, 2\pi)$. This leads to a geometric representation of the qubit on the surface of the *Bloch sphere*, a unit sphere in three-dimensional space (see Fig. 2.1). The north and south poles of the Bloch sphere correspond to the computational basis states $|0\rangle$ and $|1\rangle$, while coherent superpositions populate the equator and the surface.

The Hamiltonian of an ideal qubit reads

$$\hat{H}_q = \frac{\hbar\omega_q}{2}\hat{\sigma}_z \tag{2.3}$$

where ω_q is the energy difference between the computational states[1].

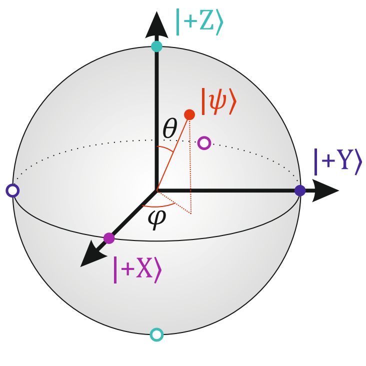

Figure 2.1: Bloch sphere representation of a qubit. The cardinal states are $|+Z\rangle = |0\rangle\,, |-Z\rangle = |1\rangle\,, |\pm X\rangle = \frac{|0\rangle \pm |1\rangle}{\sqrt{2}}, |\pm Y\rangle = \frac{|0\rangle \pm i|1\rangle}{\sqrt{2}}$. Any state vector $|\psi\rangle$ pointing on the sphere can be represented as a linear combination of $|\pm Z\rangle$ (see Eq. 2.2).

[1]Here, $\hat{\sigma}_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ is the Pauli-Z matrix and represents the qubit states, and $\hbar = h/2\pi$ is the reduced Planck constant.

The realization of qubits depends strongly on the physical platform chosen. While all qubits must obey the fundamental principles of quantum mechanics, their physical embodiments can vary widely. Each platform encodes qubit states using different degrees of freedom, with corresponding advantages and challenges. Some examples are

- **Trapped ions**[33, 34]: electronic or hyperfine spin states of the trapped ion;
- **Photonic qubits**[35]: polarization (horizontal $|H\rangle$ vs. vertical $|V\rangle$), time-bin (encoded in the time of arrival of the photon);
- **Semiconducting qubits**[36]: electron and nuclear spins ($|\uparrow\rangle$ vs. $|\downarrow\rangle$);
- **Superconducting circuits**[37, 38]: superconducting condensate phase oscillations (transmon), circulating supercurrents (fluxonium).

2.1.1.2 Multi-level systems

While binary systems are traditionally employed for quantum computation, information carriers exhibiting multi-level structures are naturally more abundant. Those include *qudits*—quantum systems with $d > 2$ orthogonal basis states, and harmonic oscillators. In contrast to the two-level systems, the Hilbert space associated with either multi-level system grows from two to d and infinite dimensions, respectively.

The use of multi-level systems in quantum computing offers potential advantages, such as an increased *information density* per physical carrier. The Hilbert space of a single harmonic oscillator truncated at $|N\rangle$ can store an equivalent amount of information to a processor of $\log_2 N$ qubits or of $\log_d N$ qudits (see scaling in Fig. 2.2). Consequently, the number of physical systems required to realize a given quantum algorithm may be reduced compared to a purely qubit-based architecture[39], and the circuit complexity can be lowered[40, 41]. Moreover, some error correction schemes based on qudits have demonstrated enhanced error thresholds for magic state distillation as compared to their qubit counterparts[42].

Qudits Notable examples of qudits in some of the previously mentioned platforms for quantum computing include:

- **Trapped ions** can serve as multi-level systems because they possess rich internal structures (e.g. Zeeman levels)[39];
- **Spin systems**[43, 44][2] with spin-$s > 1/2$ particles offer $(2s + 1)$-dimensional state spaces[45];
- **Photonic platforms**, where different spatial modes can be used to define a qudit basis[46];

[2] Not to be confused with the qubit platform based on electron spins.

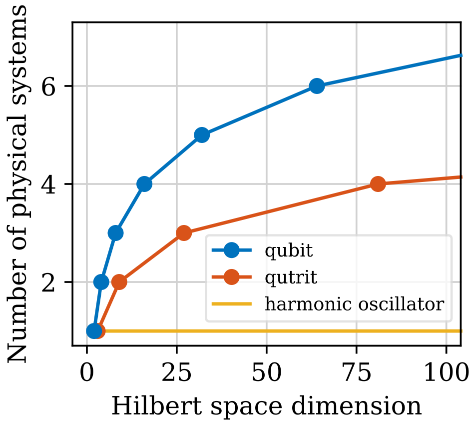

Figure 2.2: Number of physical systems required for a given Hilbert space dimension.

- **Superconducting circuits**, where higher excited states of a device (e.g., a transmon[47, 48]) can be coherently accessed and manipulated.

Control of qudits requires generalized unitary operations acting within a d-dimensional Hilbert space $\mathscr{H}_d$. Similarly, measurement schemes must be capable of resolving all d basis states, which can introduce additional experimental challenges compared to binary readout[39].

Quantum computing with qudits is an active area of research, driven by the potential to optimize quantum algorithms[49, 50], and maximize the computational capabilities of existing physical platforms. Nevertheless, all advantages come at the cost of experimental challenges, such as qudit readout, more stringent requirements for qudit coherence[51], and increased complexity of the control electronics.

Harmonic oscillators The quantum harmonic oscillator plays a central role in quantum information processing, particularly in the study and implementation of bosonic modes for quantum computation and communication. Unlike qudit systems, the harmonic oscillator possesses an infinite-dimensional Hilbert space, enabling the encoding of quantum information into continuous-variable or bosonic degrees of freedom. Its Hamiltonian

$$\hat{H} = \hbar\omega\left(\hat{a}^\dagger\hat{a} + \frac{1}{2}\right) \tag{2.4}$$

produces a spectrum of equally spaced energy levels, separated by frequency ω. Here $\hat{a}(\hat{a}^\dagger)$ represents the annihilation(creation) operator.

Several experimental platforms have demonstrated high-fidelity bosonic quantum control, examples of which are:

- **Trapped ions**, where bosonic modes arise naturally from the collective motion of the ions in the potential well or the motion of individual ions[52]. These phonon modes are conventionally used as a bus to create all-to-all connectivity between the ions in the trap. Still, there has been increasing interest in directly applying these modes in quantum computation[53, 54] and simulation[55].

- **Superconducting circuits**, where quantum information is stored in the quantized electromagnetic field of a superconducting cavity mode. These cavities can be either three-dimensional (3D) or integrated as planar resonators (2D). While mostly used for readout, there has been rapid progress in bosonic computing with superconducting microwave cavities[38, 56, 57], making superconducting circuits the most advanced platforms for bosonic quantum information processing.

Bosonic quantum computation represents a fundamentally distinct paradigm from qubit- or qudit-based approaches. Leveraging the infinite-dimensional Hilbert space of bosonic modes introduces unique challenges, particularly in benchmarking and characterizing continuous-variable quantum processors. Nevertheless, bosonic systems have emerged as strong contenders for realizing quantum algorithms, with significantly reduced hardware overhead compared to conventional qubit platforms.

2.1.2 Environment coupling

A more convenient way of representing the state of a system is the density matrix $\hat{\rho}$, a positive semi-definite (i.e., all eigenvalues are real and non-negative), Hermitian operator with unit trace. It is especially preferred for open systems, i.e., systems interacting with the environment, as it can represent both coherent and incoherent superpositions, in contrast to the state vector, which can only represent the former. Without any assumptions, a system with a set of states $|\psi_n\rangle$ has a density matrix

$$\hat{\rho} = \sum_n p_n |\psi_n\rangle\langle\psi_n| \tag{2.5}$$

where p_n corresponds to the probability that the system is in the state $|\psi_n\rangle$ and it obeys $\sum_n p_n = 1$. The system is said to be in a pure state if $\rho = |\psi\rangle\langle\psi|$ and the end of the state vector lies on the surface of the Bloch sphere; otherwise, it is mixed and the state vector points somewhere inside the Bloch sphere.[3] The diagonal terms represent basis state population, while the off-diagonal terms give information about the phase difference between the states.

For a closed system comprised of a qubit coupled to its environment, i.e, with a state vector $|\psi\rangle = \sum_n p_n |\psi_{q,n}\rangle \otimes |\psi_{\mathrm{env},n}\rangle$, the evolution is governed by the Schrödinger-von Neumann equation

$$\frac{d\hat{\rho}}{dt} = -\frac{i}{\hbar}[\hat{H}, \hat{\rho}], \tag{2.6}$$

where the system Hamiltonian is a sum of the qubit, environment, and their interaction Hamiltonians $\hat{H} = \hat{H}_q + \hat{H}_{\mathrm{env}} + \hat{H}_{\mathrm{int}}$. To extract the dynamics of the qubit, we begin by tracing out the environment

$$\hat{\rho}_q = \mathrm{Tr}_{\mathrm{env}}[\hat{\rho}] = \sum_n \langle\psi_{\mathrm{env},n}|\hat{\rho}|\psi_{\mathrm{env},n}\rangle. \tag{2.7}$$

[3] Note that the 3D Bloch sphere representation is non-trivial for the qutrit case.

The evolution of such a system is generally described by the *Lindblad master equation*

$$\frac{d\hat{\rho}_q}{dt} = -\frac{i}{\hbar}[\hat{H}, \hat{\rho}_q] + \sum_k \hat{L}_k \hat{\rho}_q \hat{L}_k^\dagger - \frac{1}{2}\hat{\rho}_q \hat{L}_k^\dagger \hat{L}_k - \frac{1}{2}\hat{L}_k^\dagger \hat{L}_k \hat{\rho}_q \tag{2.8}$$

where $\hat{L}_k = \sqrt{\gamma_k}\hat{A}_k$ are operators, each representing a specific type of coupling of the qubit to the environment via the operator $\hat{A}_k$ with a rate γ_k. To derive the Lindblad master equation, several approximations are made:

- **Separability**: there are no correlations between the qubit and the environment at time $t = 0$, i.e, $\hat{\rho}(0) = \hat{\rho}_q(0) \otimes \hat{\rho}_{\text{env}}(0)$.

- **Born approximation**: the interaction is weak and the environment is much larger than the system, and therefore not significantly influenced by it, i.e., $\hat{\rho}(t) = \hat{\rho}_q(t) \otimes \hat{\rho}_{\text{env}}$, where $\hat{\rho}_{\text{env}}$ is the stationary state of the environment;

- **Markov approximation**: the environment has no memory, i.e. the environmental correlations decay much faster than the timescale over which the qubit evolves (τ_q), so the past interactions between the qubit and the environment can be neglected.

- **Secular approximation**: the transition frequencies are much bigger than $(\min(\tau_q))^{-1}$, allowing us to neglect all fast rotating terms.

Energy Relaxation and Excitation Energy relaxation, or *amplitude damping*, refers to the irreversible loss of excitation from a higher-energy state (e.g., $|1\rangle$) to a lower-energy state (e.g., $|0\rangle$) due to coupling to a dissipative bath. This process is captured by the jump operator

$$\hat{L} = \sqrt{\gamma}\hat{\sigma}_-, \tag{2.9}$$

where γ is the decay rate and $\hat{\sigma}_- = |0\rangle\langle 1|$ is the lowering operator. This coupling is often referred to as *transverse*, as the noise couples to $\hat{\sigma}_x$. If the bath has a finite temperature, the qubit can both get excited by it and decay into it, resulting in the master equation:

$$\frac{d\hat{\rho}_q}{dt} = -\frac{i}{\hbar}[\hat{H}, \hat{\rho}_q] + \frac{\gamma}{2} n_{\text{th}} \hat{D}[\hat{\sigma}_+]\hat{\rho}_q + \frac{\gamma}{2}(n_{\text{th}} + 1)\hat{D}[\hat{\sigma}_-]\hat{\rho}_q, \tag{2.10}$$

where $\hat{\sigma}_+ = |1\rangle\langle 0|$ is the raising operator, $n_{\text{th}} = (e^{\hbar\omega_q/k_B T} - 1)^{-1}$ is the mean photon number of the bath with temperature T at the frequency ω_q.[4] We have introduced the dissipative Lindblad superoperator $\hat{D}[\hat{\mathcal{O}}]\rho = 2\hat{\mathcal{O}}\hat{\rho}\hat{\mathcal{O}}^\dagger - \hat{\mathcal{O}}^\dagger\hat{\mathcal{O}}\hat{\rho} - \hat{\rho}\hat{\mathcal{O}}^\dagger\mathcal{O}$. Equation 2.10 describes the exponential decay of the ground- and excited-state populations as well as a decay of off-diagonal terms.

The total rate

$$\Gamma_1 = \frac{1}{T_1} = \Gamma_1^\uparrow + \Gamma_1^\downarrow = \gamma(2n_{\text{th}} + 1) \tag{2.11}$$

defines the relaxation time of the qubit T_1. In thermal equilibrium, the ratio between excitation and relaxation rates is

$$\frac{\Gamma_1^\uparrow}{\Gamma_1^\downarrow} = e^{-\hbar\omega_q/k_B T} \tag{2.12}$$

[4] k_B is the Boltzmann constant.

determined by Boltzmann statistics. Note that only the thermal population of the bath at the qubit frequency[5] affects the qubit relaxation time.

Dephasing Dephasing refers to the decay of off-diagonal elements of the density matrix without a change in energy populations, or in other words, randomization of the phase without energy exchange. In contrast to the energy decay, the noise couples to the qubit along the z-axis; such a coupling is called *longitudinal*. The corresponding jump operator is

$$\hat{L} = \sqrt{\Gamma_\phi}\,\hat{\sigma}_z, \tag{2.13}$$

where $\Gamma_\phi = 1/T_\phi$ is the pure dephasing rate and $\sigma_z = |0\rangle\langle 0| - |1\rangle\langle 1|$. The master equation becomes

$$\frac{d\hat{\rho}_q}{dt} = -\frac{i}{\hbar}[\hat{H}, \hat{\rho}_q] + \frac{\Gamma_\phi}{2}\hat{D}[\hat{\sigma}_z]\hat{\rho}_q, \tag{2.14}$$

which damps the off-diagonal coherence terms without affecting the populations. Pure dephasing is not a resonant phenomenon, in contrast to relaxation, and up to a certain extent can be reversed by using decoupling sequences[58].

Decoherence In realistic systems, both energy relaxation and pure dephasing are present. Their combined effect is characterized by the total coherence time T_2

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi}, \tag{2.15}$$

which governs the decay of the off-diagonal elements of $\hat{\rho}$. In addition to the pure dephasing, relaxation or excitation causes a complete loss of the phase of any initial superposition state. In matrix form, the evolution of a qubit's density matrix under decay and decoherence is

$$\hat{\rho}_q(t) = \begin{pmatrix} 1 + (\rho_{00}(0) - 1)e^{-t/T_1} & \rho_{01}(0)e^{-t/T_2} \\ \rho_{10}(0)e^{-t/T_2} & \rho_{11}(0)e^{-t/T_1} \end{pmatrix}, \tag{2.16}$$

with populations $\rho_{00}(t)$ and $\rho_{11}(t)$ governed by T_1, and coherence $\rho_{01}(t)$ and $\rho_{10}(t)$ by T_2, where $\rho_{ij} = \langle i|\hat{\rho}_q|j\rangle$. Combined with energy relaxation, they degrade quantum information and motivate the need for quantum error correction in all scalable quantum computing architectures. The matrix from Eq. 2.16 is often referred to as the Bloch-Redfield density matrix[31].

2.1.3 Quantum error correction

In quantum mechanics, the evolution of a system's state must always preserve the fundamental properties of the density matrix: it must remain Hermitian, positive semi-definite, and have unit trace. These properties must be maintained not only during ideal unitary evolution under a Hamiltonian but also when the system undergoes noise, dissipation, or other imperfections over a time interval Δt.

[5] Realistically, the qubit linewidth plays a role as well.

The most general class of physical operations that map a valid density matrix to another valid density matrix is known as *completely positive trace-preserving (CPTP)* maps. These maps can be expressed in the so-called *Kraus form*

$$\hat{\rho}(t+\Delta t) = \mathcal{E}(\hat{\rho}) = \sum_{\mu} \hat{M}_{\mu}\hat{\rho}\hat{M}_{\mu}^{\dagger}, \tag{2.17}$$

where the set $\{\hat{M}_{\mu}\}$ is known as *Kraus operators*, which represent different possible outcomes or types of errors during the evolution. Even though individual $\hat{M}_{\mu}$ may be non-unitary, the map as a whole must preserve the trace of ρ, requiring the completeness condition

$$\sum_{\mu} \hat{M}_{\mu}^{\dagger}\hat{M}_{\mu} = \hat{\mathbb{I}}, \tag{2.18}$$

with $\hat{\mathbb{I}}$ denoting the identity operator.

Importantly, the Kraus decomposition of a quantum channel is not unique. If $\hat{U}$ is a unitary matrix acting on the space of Kraus operators, a new set defined by

$$\hat{N}_{\mu} = \sum_{\nu} \hat{U}_{\mu\nu}\hat{M}_{\nu}$$

will describe the same quantum operation

$$\sum_{\mu} \hat{N}_{\mu}\hat{\rho}\hat{N}_{\mu}^{\dagger} = \sum_{\mu} \hat{M}_{\mu}\hat{\rho}\hat{M}_{\mu}^{\dagger}.$$

If a quantum error is correctable, then there exists a *recovery operation* $\mathcal{R}$ such that the composition of the error and recovery maps returns the state to its original form:

$$(\mathcal{R} \circ \mathcal{E})(\hat{\rho}) = \hat{\rho}, \quad \forall \hat{\rho} \in \mathscr{C}, \tag{2.19}$$

where $\mathscr{C}$ is the subspace (code space) used to encode the logical quantum information.

To determine whether a particular error model $\{\hat{M}_{\mu}\}$ is correctable for a code, one uses the *Knill-Laflamme condition*[59]. Suppose the code encodes a single logical qubit in two orthonormal states $\{|\psi_0\rangle, |\psi_1\rangle\}$. The necessary and sufficient condition for the errors $\{\hat{M}_{\mu}\}$ to be correctable is

$$\langle\psi_l|\hat{M}_{\mu}^{\dagger}\hat{M}_{\nu}|\psi_k\rangle = c_{\mu\nu}\delta_{lk}, \tag{2.20}$$

for all $k, l \in \{0, 1\}$, and where $c_{\mu\nu}$ are the elements of a Hermitian matrix[6]. This equation ensures two things:

- The different error operators act in distinguishable ways on the code space, i.e., each error maps the code into an orthogonal error subspace.

- The logical basis states remain orthogonal even after errors, ensuring that no information about the logical state is leaked to the environment.

[6] δ_{lk} is the Kronecker delta.

When these conditions are satisfied, one can identify the error that occurred by measuring a suitable set of observables (syndrome measurements) and applying a corresponding recovery operation, typically a unitary inverse, to restore the original encoded state.

If a set of errors $\{\hat{M}_\mu\}$ is correctable by a recovery operation $\mathcal{R}$, then any other set of error operators $\{\hat{N}_\mu\}$, formed as linear combinations of the $\hat{M}_\mu$, will also be correctable by the same recovery map.

In the case of a logical qubit, arbitrary single-qubit errors can be expressed as linear combinations of the Pauli operators $\hat{\sigma}_x$, $\hat{\sigma}_z$, and $\hat{\sigma}_y$, where $\hat{\sigma}_x$ corresponds to bit-flip errors ($|0\rangle \leftrightarrow |1\rangle$), and $\hat{\sigma}_z$ to phase-flip errors ($|+\rangle \leftrightarrow |-\rangle$). Since $\hat{\sigma}_y = i\hat{\sigma}_x\hat{\sigma}_z$, it follows that any single-qubit error can be decomposed into a combination of bit and phase flips. Therefore, a code that can correct both $\hat{\sigma}_x$ and $\hat{\sigma}_z$ errors is, in principle, capable of correcting any arbitrary single-qubit error.

2.2 Superconducting circuits

While there are numerous platforms for encoding quantum information, superconducting circuits have emerged as one of the leaders for realizing scalable quantum computation[60]. These circuits harness macroscopic quantum coherence by integrating superconductors into engineered electrical networks. Their compatibility with well-established nanofabrication techniques enables the rapid prototyping and integration of multiple qubits on a single chip, making them particularly well-suited for building large-scale quantum processors. Moreover, the circuits can be engineered to create arbitrary interactions, i.e., longitudinal or transversal, and vary their strengths.

2.2.1 Building blocks

2.2.1.1 Quantum LC circuit

One of the simplest circuits one can build is an LC resonator. Whether classical or quantum, it consists of an inductor L and a capacitor C connected in parallel. The field within the circuit oscillates between the magnetic field stored in the inductor and the electric field in the capacitor.
The circuit has the following Hamiltonian

$$\hat{H} = \frac{\hat{Q}^2}{2C} + \frac{\hat{\Phi}^2}{2L} \tag{2.21}$$

where $\hat{\Phi} = \int_{-\infty}^{t} \hat{V}(\tau)d\tau$ is the branch flux and $\hat{Q} = C\dot{\hat{\Phi}}$ is the conjugate variable, the branch charge. A detailed derivation of this Hamiltonian can be found in Ref. [61]. The flux and charge operators satisfy the commutation relation $[\hat{\Phi}, \hat{Q}] = i\hbar$. In the second quantization representation, they can be expressed as

$$\hat{\Phi} = \Phi_{\mathrm{zpf}}(\hat{a} + \hat{a}^\dagger) \tag{2.22}$$

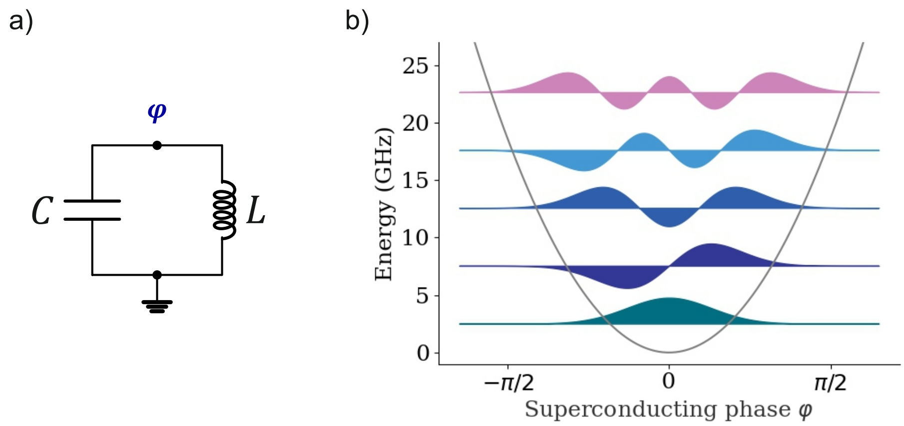

Figure 2.3: Quantum LC oscillator. a) Circuit; b) Energy spectrum as a function of superconducting phase. The parabolic potential is plotted in gray, and the wavefunctions are scaled for better visibility.

$$\hat{Q} = -iQ_{\text{zpf}}(\hat{a} - \hat{a}^\dagger) \tag{2.23}$$

where $\Phi_{\text{zpf}} = \sqrt{\frac{\hbar Z}{2}}$ and $Q_{\text{zpf}} = \sqrt{\frac{\hbar}{2Z}}$ are the zero-point fluctuations of the flux and charge variables, respectively, which depend on the impedance of the harmonic oscillator $Z = \sqrt{\frac{L}{C}}$.

The Hamiltonian 2.21 can be rewritten in terms of the phase across the inductor $\hat{\varphi} = \frac{\hat{\Phi}}{\phi_0}$[7] and the charge number operator $\hat{n} = \frac{\hat{Q}}{2e}$[8] denoting the excess Cooper pairs on the capacitor island

$$\hat{H} = 4E_C\hat{n}^2 + \frac{1}{2}E_L\hat{\varphi}^2. \tag{2.24}$$

Here, we have defined the charging energy to add a Cooper pair to the island $E_C = \frac{e^2}{2C}$, and the inductive energy $E_L = \frac{\phi_0^2}{L}$. The unitless operators obey the commutation relation $[\hat{\varphi}, \hat{n}] = i$. They can also be expressed via the creation and annihilation operators $\hat{n} = -in_{\text{zpf}}(\hat{a} - \hat{a}^\dagger)$ and $\hat{\varphi} = \varphi_{\text{zpf}}(\hat{a} + \hat{a}^\dagger)$, with the zero-point fluctuations $n_{\text{zpf}} = \sqrt[4]{\frac{E_L}{32E_C}}$ and $\varphi_{\text{zpf}} = \sqrt[4]{\frac{2E_C}{E_L}}$, respectively. The Hamiltonian can be diagonalized analytically, yielding an infinite series of eigenstates, equidistantly spaced in energy by $\hbar\omega = \hbar\sqrt{8E_CE_L} = \hbar\frac{1}{\sqrt{LC}}$ as illustrated in Fig. 2.3b.

Harmonic oscillators in practice

There are several possibilities for building a harmonic oscillator with superconducting circuits in both planar (2D) and three-dimensional (3D) architectures. Still, they all obey the same physical laws (see Sec. 2.1.1.2). Their total lifetime is governed by the coupling to the environment's degrees of freedom, both controlled and uncontrolled. We can quantify these effects with their corresponding coupling rates κ_{ext} and κ_{int}, referred to as external

[7] $\phi_0 = \frac{\Phi_0}{2\pi}$ is the reduced flux quantum, while $\Phi_0 = \frac{h}{2e}$ is the magnetic flux quantum.
[8] e is the electron charge.

and internal coupling rates. They are, in turn, connected to the external and internal quality factors[9], i.e.

$$\kappa_{\text{ext(int)}} = \frac{\omega}{Q_{\text{ext(int)}}}, \tag{2.25}$$

where ω is the resonance frequency of the harmonic oscillator. The total linewidth of the resonance, also referred to as the loaded coupling rate, is $\kappa_L = \kappa_{\text{int}} + \kappa_{\text{ext}}$, provided that there is no dephasing present[10]. These quantities can be related to loss operators:

- Single-photon loss: $\hat{L}_{1,\downarrow} = \sqrt{\kappa_1^{\downarrow}}\hat{a}$,
- Single-photon excitation: $\hat{L}_{1,\uparrow} = \sqrt{\kappa_1^{\uparrow}}\hat{a}^{\dagger}$,
- Dephasing: $\hat{L}_{\phi} = \sqrt{\kappa_{\phi}}\hat{a}^{\dagger}\hat{a}$.

Here, processes $\hat{L}_{1,\downarrow}$ and $\hat{L}_{1,\uparrow}$ depend on the number of thermal photons in the cavity n_{th}, similar to Eq. 2.11, and their ratio is the same as in Eq. 2.12. The first two loss operators can be generalized for the case of n photons: $\hat{L}_{n,\downarrow(\uparrow)} = \sqrt{\kappa_n^{\downarrow(\uparrow)}}\hat{a}^{n(\dagger n)}$. The total single-photon coupling rate κ_1 and the loaded coupling rate κ_L are correlated, i.e. the resonance linewidth will be partially determined by κ_1.

There has been an immense effort to reduce the single-photon loss rate in superconducting resonators. Various architectures have been explored (see Fig. 2.4 for examples). Generally, a realistic resonator is based on transmission lines. Infinite transmission line can support an infinite continuum of modes[62]. Introducing a discontinuity creates boundary conditions that restrict which modes can live in the line. Depending on the boundary conditions, the type of resonators can be divided into two categories: half-wave ($\lambda/2$) resonators, which have voltage antinodes at each end, and quarter-wave ($\lambda/4$) resonators, which have one voltage node and one voltage antinode. Starting from the planar implementations, some examples include microstrip[63], coplanar-waveguide (CPW)[64, 65], or stripline-based resonators[66]. Moreover, the latter are often bent into various shapes, e.g., forming a hairpin resonator[66, 67].

Three-dimensional architectures can also be used as microwave resonators. The shape determining the resonance is carved into a solid piece of a superconductor, and depending on the geometry, the cavity can be rectangular[68], post[69, 70], flute[71], mushroom[72], etc.. Generally, 3D cavities have a longer lifetime, as the mode volume is bigger than in the 2D case and is much less influenced by any lossy surfaces. The rectangular cavity normally consists of two separate pieces; their joint, the seam, is lossy[73] due to the natural oxides that grow on the surfaces. The post and flute cavity qualify as "seamless" as they don't suffer from seam losses, while the seam losses in the mushroom cavity are avoided by welding the parts together, leading to the best recorded coherence time in the field. Finally, another promising approach is the micromachined cavities[74, 75], combining the small footprint of the planar architectures with the reduced losses of the 3D cavities.

[9]The quality factor reflects the number of times an excitation can "bounce around" in the oscillator, i.e. the number of field oscillations, before escaping.

[10]Generally, $\kappa_L \propto 1/T_2$.

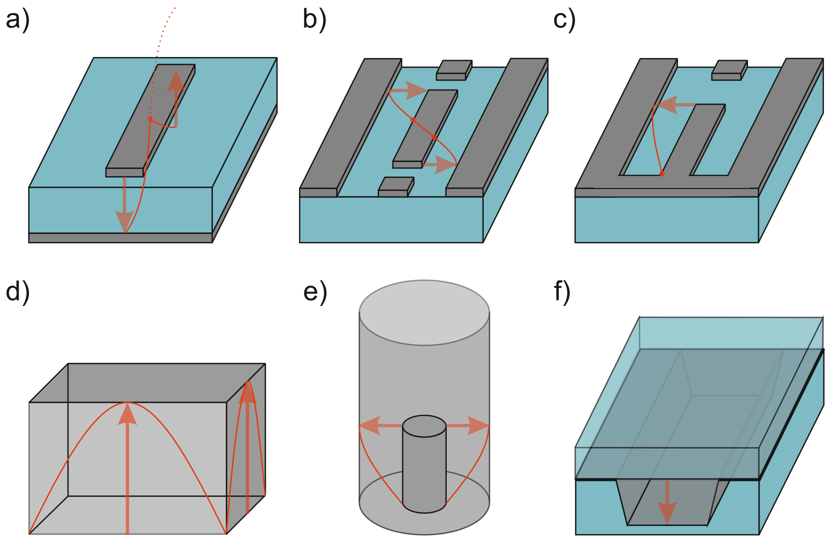

Figure 2.4: Microwave resonators. a)-c) planar (2D) architectures, d),e) 3D architectures, f) 2.5D architecture. The turquoise color represents the substrate: either silicon or sapphire, while the metalization is in gray. The field envelope is in red with an arrow pointing towards its maxima, and a red dot denoting zero or no field. The exact types are as follows: a) Microstrip $\lambda/2$ resonator. b) CPW $\lambda/2$ resonator. c) CPW $\lambda/4$ resonator. d) Rectangular cavity. e) Post cavity. f) Micromachined cavity. For the microstrip, most of the field resides in the substrate, illustrated by both maxima in the sapphire, while the dashed line show the extension of the $\lambda/2$ field envelope outside the substrate. The layers in the 2D and 2.5D cases are not to scale.

While harmonic oscillators are suitable for storage, the equidistant energy levels of the harmonic oscillator pose a limitation to their usage. Upon a resonant drive, the harmonic oscillator would be displaced to a coherent state with phase and amplitude proportional to that of the drive tone. To be able to define a computational subspace and use an LC circuit as a qubit, a non-linearity is required to lift the degeneracy of the transitions. The most commonly used non-linear element for superconducting circuits is the Josephson junction.

2.2.1.2 Junctions

A Josephson junction, by definition, is a local weakening of the superconducting condensate. It can be created in several ways, including placing a thin insulating layer between two superconductors (SIS tunnel junction) or reducing the size of the superconductor to inhibit the flow of Cooper pairs (constriction junction). The SIS junctions, shown in Fig. 2.5, have a relatively straightforward fabrication, which includes evaporating a thin superconducting film, oxidizing its surface at a fixed pressure for a certain time, and evaporating the second layer on top. Due to that, they are the most common type and, therefore, we focus on them from this point on.

Cooper pairs tunnel through the insulating barrier, resulting in a non-linear phase-current relation

$$I = I_c \sin\varphi. \tag{2.26}$$

Here $\varphi = \varphi_1 - \varphi_2$ is the phase difference of the two superconducting condensates on each island. Equation 2.26 is called the first Josephson relation, and I_c is the critical current, which depends on the superconductor and the barrier properties, as well as the junction geometry. This relation tells us that a supercurrent can flow through the junction even in the absence of a voltage, purely due to a phase difference between the superconducting condensates. The second Josephson equation

$$\frac{d\varphi}{dt} = \frac{2eV}{\hbar} \tag{2.27}$$

relates the rate of change of the phase to the voltage V across the junction. It describes the AC Josephson effect, where a constant voltage is applied, causing the phase to increase linearly over time, and the resulting current to become oscillatory.

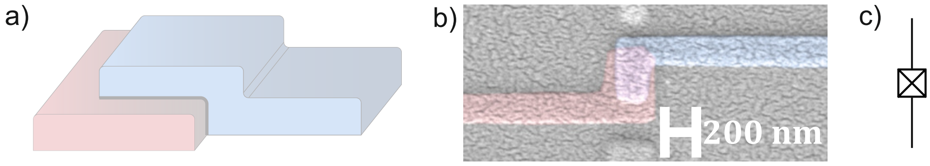

Figure 2.5: SIS junction. a) Schematic. First and second superconducting layers in blue and pink, respectively, and a thin insulating layer (oxide) in gray. b) False-colored Scanning-Electron-Microscope (SEM) image. First and second superconducting layers in blue and pink, respectively. c) Electric symbol of a Josephson junction.

From the two Josephson relations, the Hamiltonian for a single junction can be derived

$$\hat{H} = -E_J \cos\hat{\varphi} + 4E_{C_J}\hat{n}^2 \tag{2.28}$$

where we have introduced the Josephson energy $E_J = \phi_0 I_c = \frac{\phi_0^2}{L_J}$, which also defines the junction inductance L_J. The capacitive energy stems from the plate capacitor C_J formed between the two superconducting plates. For small zero-point fluctuations of the phase, we can expand the cosine term in Eq. 2.28 and see that a Josephson junction behaves as a nonlinear inductor.

2.2.1.3 Kinetic inductance materials

To introduce nonlinearity into a linear superconducting circuit, a single Josephson junction is often sufficient. However, many advanced applications, such as protected qubits[76], enhanced coupling schemes[77], parametric conversion[78], and amplification[79], require stronger nonlinearity or larger zero-point phase fluctuations, compared to those of a single junction. These properties are closely linked to the circuit's impedance, which must be high. Achieving this typically involves maximizing inductance while minimizing capacitance.

All conducting materials possess geometric inductance to oppose any change in the electric current flowing through them. This property depends purely on the physical layout of the circuit—the shape, size, and spacing of the conductors—and not on the material properties. Creating a large geometric inductor is associated with an increased footprint and advanced fabrication techniques[80]. Kinetic inductance in superconductors, on the other hand, arises in superconductors due to the inertia of Cooper pairs. Unlike conventional geometric inductance, which impedes the charge carriers simply by dimensional constraints, kinetic inductance stems from the energy stored in the motion of these charge carriers. This effect becomes particularly significant in superconducting materials with low carrier density.

As previously mentioned, Josephson junctions act as nonlinear inductors with a high kinetic contribution. One straightforward approach to creating large inductors is to place many junctions in series to form an array, as their inductance is easily tunable during the fabrication process. However, this method faces practical limitations: the junction capacitance can degrade performance, as the arrays' parasitic modes may compromise device performance[81]. An alternative solution is to use disordered superconductors, which naturally exhibit high kinetic inductance, much greater than that of conventional superconducting materials, while maintaining a compact footprint and minimizing their self-capacitance.

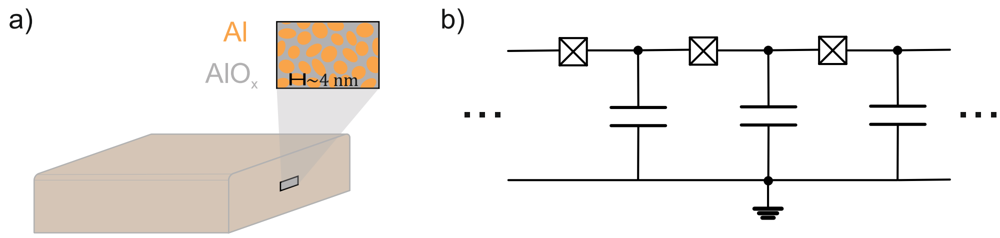

Figure 2.6: Granular aluminum. a) Schematic. Granular aluminum consists of aluminum grains with a diameter of a few nanometers[82, 83], separated by a layer of aluminum oxide. b) Equivalent circuit. The 3D network of superconducting islands separated by an oxide barrier can be mapped onto a 1D chain of Josephson junctions and the self-capacitance of the superconducting islands. Figure adapted from Ref. [84].

A prominent example of a high-kinetic inductance material is granular aluminum (grAl)[84]. This is a disordered superconductor consisting of nanometer-scale aluminum grains embedded in an amorphous aluminum oxide matrix. Its fabrication is compatible with standard Josephson junction processes, typically requiring only aluminum evaporation in a controlled oxygen environment. Granular aluminum exhibits several advantageous properties:

- High kinetic inductance per square, tunable within orders of magnitude via deposition parameters like oxygen pressure and evaporation rate;
- Larger superconducting gap, critical magnetic field, and critical temperature compared to that of pure aluminum;
- Low microwave loss, making it suitable for building circuits (e.g., resonators, qubits, kinetic inductance detectors).

Owing to this combination of high kinetic inductance, low loss, and fabrication compatibility, granular aluminum has emerged as a leading material for building compact, high-impedance superconducting circuits.

2.2.2 Qubits

Having all the necessary ingredients, more complex circuits can be built, such as qubits. The simplest way to create a qubit is by shunting a junction with a large capacitance, which suppresses the charge sensitivity of this circuit, schematically illustrated in Fig. 2.7a. Equivalently, the same circuit is built by replacing the linear inductor of the LC resonator with a junction, which lifts the level degeneracy and makes it possible to isolate two levels for a computational subspace. This qubit is called a *transmon*[85] and has the following Hamiltonian

$$\hat{H}_{transmon} = -E_J \cos\hat{\varphi} + 4E_{C_\Sigma}\hat{n}^2 \tag{2.29}$$

where $E_{C_\Sigma} = E_C + E_{C_J}$. A requirement for the charge noise suppression is that $\frac{E_J}{E_{C_\Sigma}} \gg 1$, which suppresses the phase fluctuations and allows for expanding and truncating the cosine potential in Eq. 2.29. Then the qubit transition energy is $\hbar\omega_{01} = \sqrt{8E_{C_\Sigma}E_J} - E_{C_\Sigma}$ and all higher transitions are detuned by multiples of E_{C_Σ} with respect to the fundamental transition. Due to its simplicity, this qubit is the most widely used superconducting qubit. However, there is a trade-off between having sufficient anharmonicity[11] and sufficient protection against noise, which places a constraint on the duration of the pulses that can be applied. This can be bypassed in more intricate circuits, such as the fluxonium.

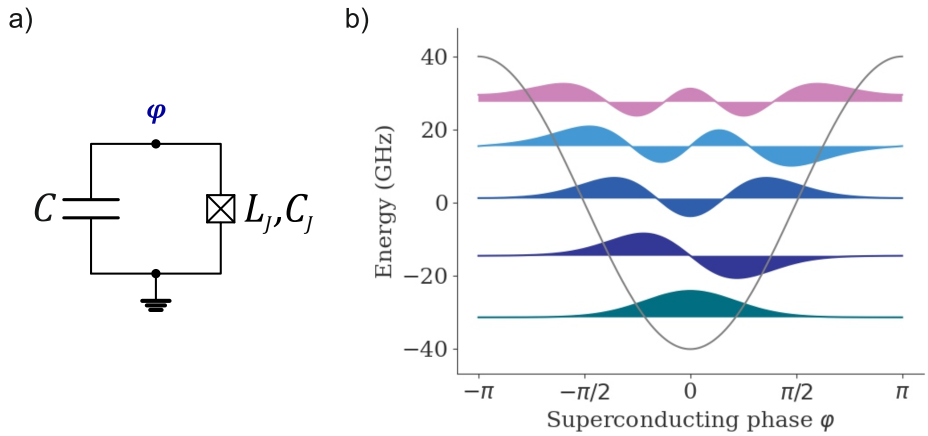

Figure 2.7: Transmon qubit. a) The electrical circuit consists of a junction in parallel with a capacitor. b) The cosine potential (gray) forms a weakly anharmonic spectrum, with the first five energy levels and their wavefunctions plotted. The lowest lying levels form the computational space. The potential is plotted with the help of *scQubits* for $E_J/h = 40\,\mathrm{GHz}$ and $E_C/h = 1\,\mathrm{GHz}$.

[11]The anharmonicity is defined as $\alpha = |\omega_{01} - \omega_{12}|$ and equals $-E_C/\hbar$ for a transmon.

2.2.2.1 Fluxonium

The *fluxonium*[86] is a superconducting qubit formed by shunting a Josephson junction with a superinductance. A superinductor has a characteristic impedance exceeding the resistance quantum at the operating frequencies typical of superconducting qubits, i.e., $\Im(Z(\omega)) = \omega L > R_Q = \frac{h}{(2e)^2} \approx 6.45\,k\Omega$. The superinductor protects the qubit against low-frequency charge noise by suppressing the charge fluctuations below $2e$ at the expense of larger flux fluctuations. The Hamiltonian of this qubit is

$$\hat{H}_{fluxonium} = 4E_C\hat{n}^2 + \frac{E_L}{2}\hat{\varphi}^2 - E_J\cos\left(\hat{\varphi} - \varphi_{\text{ext}}\right), \tag{2.30}$$

where $\varphi_{\text{ext}} = \Phi_{\text{ext}}/\phi_0$ is the normalized external flux through the loop, and the Josephson junction sets the charging energy E_C.

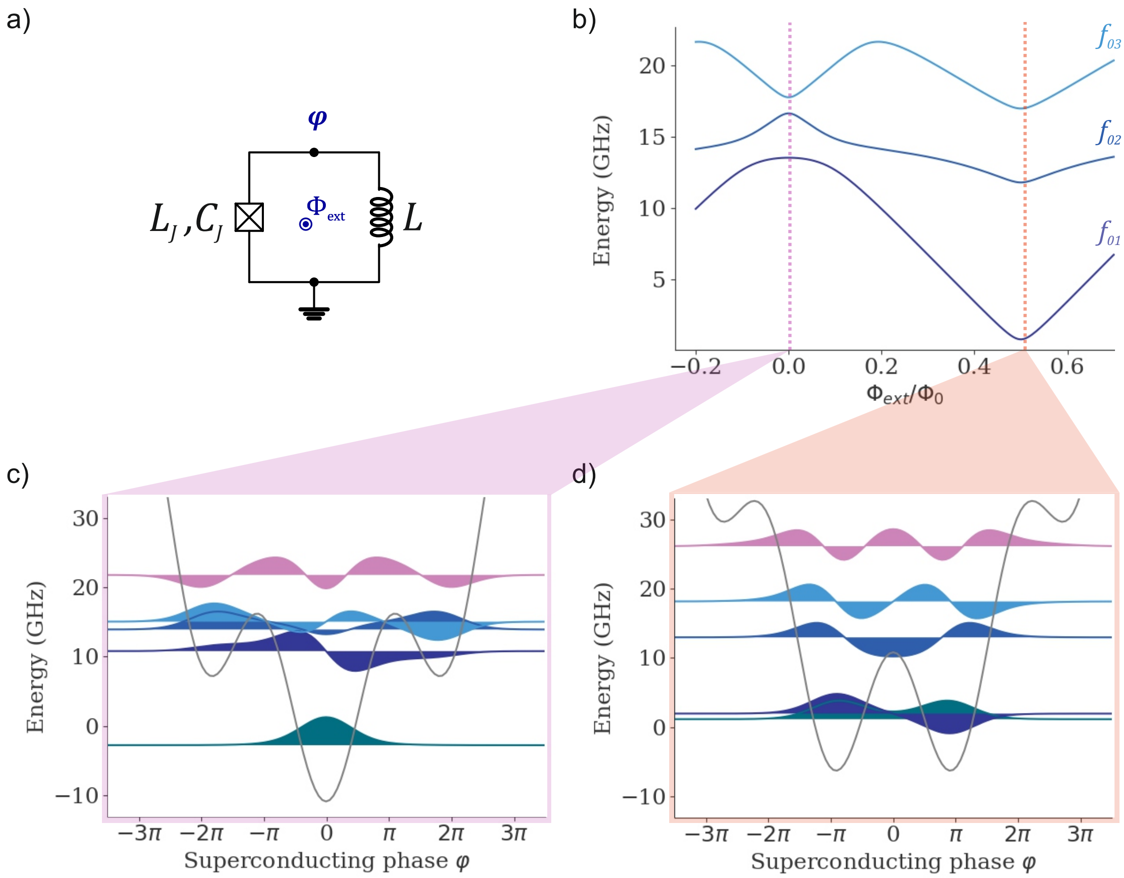

Figure 2.8: Fluxonium qubit. a) The electrical circuit consists of a junction shunted by a superconductor. The leads form a loop, which is threaded by an external magnetic field. b) Energy transitions as a function of the external flux. The frequencies are calculated as $f_{ij} = (E_j - E_i)/h$. c) and d) Potentials (gray) with the first five energy levels and their wavefunctions at $\Phi_{\text{ext}} = 0$ and $\Phi_{\text{ext}} = 0.5\Phi_0$, respectively. The two lowest lying levels form the computational space. The potential and the transitions are plotted with the help of *scQubits* for $E_J/h = 10.8\,\text{GHz}$, $E_C/h = 3.4\,\text{GHz}$, and $E_L/h = 1\,\text{GHz}$.

The fluxonium has gathered significant attention in the past few years due to its rich flux-tunable spectrum, as shown in Fig. 2.8b and in Fig. 2.9c and f. The potential consists of a parabola with a cosine superimposed on top (see Fig. 2.8c and d and in Fig. 2.9a, b, d, and e). Adjusting the flux trapped in the loop smoothly shifts the cosine term with respect to the parabola, creating a different potential and, therefore, a different energy spectrum. Two sweet spots exist at $\Phi_{\text{ext}} = 0$ and at $\Phi_{\text{ext}} = 0.5\Phi_0$, at which the transition energies are insensitive to flux to first order, denoted with dotted lines in Fig. 2.8b. The former is characterized by the lowest energy state being localized in a single well, similar to the transmon (see Fig. 2.7b), and the first energy transition resembles a *plasmon*[87, 88] excitation. This intra-well transition does not change the fluxoid number[89, 90]; it only excites collective motion of the superconducting condensate. At the latter flux bias point, the two lowest lying levels are spread out in a double-well potential. The energy splitting is given by the tunneling rate between the two wells; this transition is called a *fluxon* excitation. Such an inter-well transition corresponds to a change of the fluxoid number by one fluxon, resulting in a swap of the persistent current direction.

The charging, inductive, and Josephson energies fully determine the potential landscape. If we divide the Hamiltonian from Eq. 2.30 by the charging energy, we see that the parabola is set by the ratio $\frac{E_L}{E_C}$, and the cosine amplitude by $\frac{E_J}{E_C}$. This effect is depicted in Fig. 2.9, where a) and b) show a less steep parabola compared to the reference case from Fig. 2.8c and d, due to the former having a ratio $\frac{E_L}{E_C}$ half of that of the latter, while $\frac{E_J}{E_C}$ stays the same. Instead, in 2.9d and e $\frac{E_L}{E_C}$ was kept constant, while the ratio $\frac{E_J}{E_C}$ was reduced to half, compared to the reference in Fig. 2.8, resulting in a shallower cosine potential.

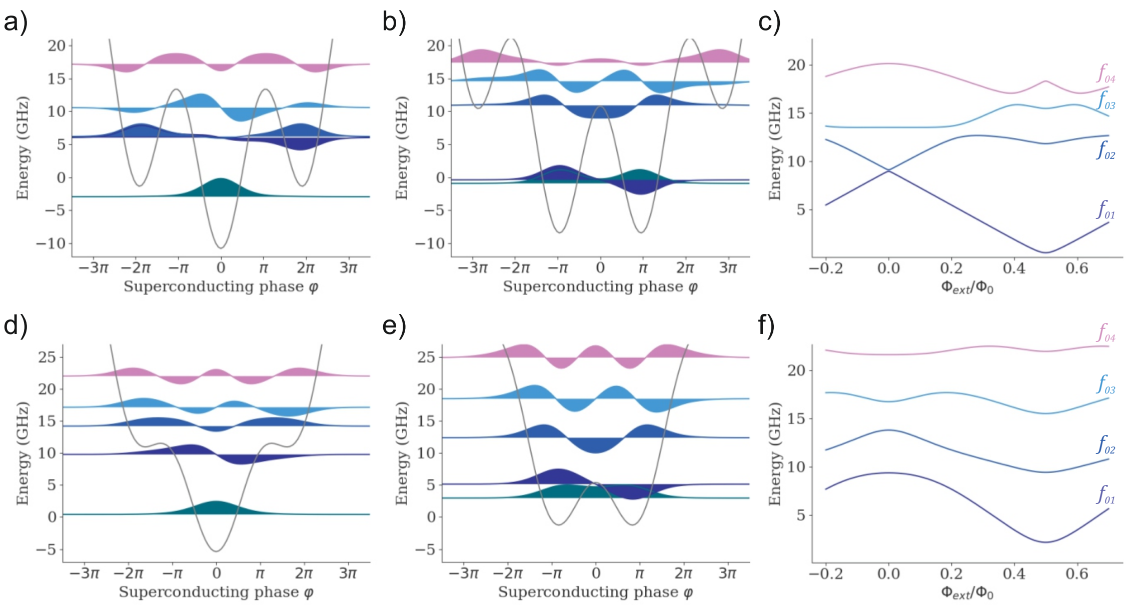

Figure 2.9: Potentials and energy spectra for different E_L/E_C and E_J/E_C ratios. First row: E_L/E_C is half the ratio of Fig. 2.8, E_J/E_C - same. Second row: E_J/E_C is half the ratio of Fig. 2.8, E_L/E_C - same. a) and d) Potentials and the energy levels at $\Phi_{\text{ext}} = 0$. b) and e) Potentials and the energy levels at $\Phi_{\text{ext}} = 0.5\Phi_0$. c) and f) First four energy transitions as a function of external flux.

2.2.3 Light-matter interaction

With the main components introduced, we turn to their interactions. We can engineer an interaction Hamiltonian that couples the degrees of freedom of the individual components, enabling entanglement between them. Physically, this coupling is realized by introducing capacitive[91] or inductive[92] elements shared between the subsystems[12]. For instance, coupling a superconducting qubit to a resonator, capacitively or inductively, allows us to read the qubit state indirectly without destroying the encoded quantum information. The following paragraphs describe the dispersive readout, the most common readout scheme used in superconducting circuits. More detailed derivations can be found in Ref. [37, 93].

Dispersive readout has its origins in cavity quantum electrodynamics (cQED), where it arises from the transverse interaction between an atom and a cavity. In superconducting circuit implementations, the bosonic mode is hosted by an on-chip 2D resonator or a 3D cavity coupled to an artificial atom (e.g., fluxonium) by a linear capacitor or inductor. Due to this resemblance, the field is known as *circuit QED*.

A qubit coupled to a harmonic oscillator, neglecting any damping in the system, is described by the Jaynes-Cummings Hamiltonian

$$\hat{H} = \hbar\omega_r \hat{a}^\dagger \hat{a} + \frac{\hbar\omega_q}{2}\hat{\sigma}_z + \hbar g(\hat{a}^\dagger + \hat{a})\hat{\sigma}_x. \tag{2.31}$$

where the first two terms are the free Hamiltonians from Eq. 2.4 and Eq. 2.3, and the third term is the interaction Hamiltonian.

The transverse interaction is inherently non-Quantum-Non-Demolition (non-QND) as it does not commute with the free qubit Hamiltonian, resulting in state evolution during the measurement. However, in the *dispersive regime*, where the detuning $\Delta = \omega_q - \omega_c$ between the atom and the cavity is much greater than the coupling strength $\Delta \gg g$, the off-resonant transverse interaction becomes QND to first order. This enables indirect measurement of the qubit state via its effect on the resonator frequency.

Under the assumption that the coupling strength is much smaller than the frequencies of the individual elements, $g \ll \omega_r, \omega_q$, the fast-rotating terms from the interaction can be neglected, which is referred to as the rotating wave approximation (RWA). The Hamiltonian in Eq. 2.31 is simplified to

$$\hat{H} = \hbar\omega_r \hat{a}^\dagger \hat{a} + \frac{\hbar\omega_q}{2}\hat{\sigma}_z + \hbar g(\hat{a}^\dagger\hat{\sigma}_- + \hat{a}\hat{\sigma}_+). \tag{2.32}$$

Under the unitary transformation $\hat{U} = \exp\left[\frac{g}{\Delta}(\hat{a}\hat{\sigma}_+ - \hat{a}^\dagger\hat{\sigma}_-)\right]$ and expanding up to second order in g, we obtain the dispersive Hamiltonian

$$\hat{H}' = \hat{U}\hat{H}\hat{U}^\dagger \approx \hbar\left(\omega_r + \frac{\chi}{2}\hat{\sigma}_z\right)\hat{a}^\dagger\hat{a} + \frac{\hbar\left(\omega_q + \frac{\chi}{2}\right)}{2}\hat{\sigma}_z. \tag{2.33}$$

[12] Mutual capacitance is introduced by placing two conductors in close proximity. To realize a mutual inductance, either two looped circuits are brought close to or on top of one another, or they share the same Josephson junction or a kinetic inductor.

The qubit experiences a Lamb shift due to its coupling to the cavity; this shift is used to define the *dispersive shift* as $\chi = 2g^2/\Delta$. The resonator, on the other hand, acquires a qubit state-dependent frequency shift. Conversely, each resonator photon would further pull the qubit frequency by χ. For a more detailed derivation, refer to Ref. [94, 95].

2.2.3.1 Dispersive shift of a fluxonium inductively coupled to a resonator

The model of a qubit coupled to a harmonic oscillator is idealized and serves to illustrate the dispersive readout qualitatively. For the realistic scenario in which we couple a superconducting qubit to a readout resonator, a full treatment of all higher-lying energy states is required for a quantitative agreement with the experiment. The frequency shift, induced by the coupling of those elements, can be found with second-order perturbation theory (more on the derivation can be found in Ref. [29]).

The inductive coupling of a fluxonium and a readout resonator is described by the Hamiltonian

$$H_{\text{int}} = -\hbar g \hat{\varphi}(\hat{a} - \hat{a}^\dagger) \tag{2.34}$$

where the flux across the shared inductance couples with strength g to the current through the resonator. The bare states of the system are $E^0_{\alpha,l} = E_\alpha + l\hbar\omega_r$, where E_α is the qubit energy of state $|\alpha\rangle$ and l is the l-th energy level $|l\rangle$ of the resonator spectrum. The dressed energy states are shifted in energy by

$$\begin{aligned} \delta E_{\alpha,l} &= \sum_{\substack{\beta\neq\alpha \\ m\neq l}} \frac{|\langle \beta, m|\hbar g\hat{\varphi}(\hat{a} - \hat{a}^\dagger)|\alpha, l\rangle|^2}{E^0_{\alpha,l} - E^0_{\beta,m}} \\ &= \sum_{\beta\neq\alpha} \hbar^2 g^2 |\varphi_{\alpha\beta}|^2 \frac{(2l+1)E_{\alpha\beta} + \hbar\omega_r}{E^2_{\alpha\beta} - (\hbar\omega_r)^2} \end{aligned} \tag{2.35}$$

where $\varphi_{\alpha\beta} = \langle\beta|\hat{\varphi}|\alpha\rangle$ is the phase matrix element between qubit states $|\alpha\rangle$ and $|\beta\rangle$ and $E_{\alpha\beta} = E_\alpha - E_\beta$ is the transition energy between the same states. The frequency shift imposed on the resonator due to the state of the qubit is given by

$$\chi_\alpha = \frac{1}{\hbar}(E_{\alpha,l+1} - E_{\alpha,l}) - \omega_r = \hbar g^2 \sum_{\beta\neq\alpha} |\varphi_{\alpha\beta}|^2 \frac{2\hbar\omega_r}{E^2_{\alpha\beta} - (\hbar\omega_r)^2}. \tag{2.36}$$

The observed dispersive shift in the resonator due to a transition between two qubit states is

$$\chi_{\alpha\beta} = \chi_\alpha - \chi_\beta. \tag{2.37}$$

The result is identical for any other superconducting qubit coupled inductively to the readout resonator. For capacitive coupling, the interaction in Eq. 2.34 would change, leading to a change in the matrix elements in equations 2.35 and 2.36, but the derivation remains the same.

2.3 Quantum error correction with superconducting circuits

Superconducting qubits, like all other platforms, require coupling to the environment for control and readout. Moreover, they couple to uncontrollable degrees of freedom. Both controllable and uncontrollable coupling to the environment greatly reduces their lifetime and coherence, rendering them inadequate for scaling and running useful quantum algorithms.

While controllable coupling is necessary, reducing the number of bath modes to which the qubit could couple uncontrollably or the strength of the coupling itself would greatly improve its computational capabilities. An effective approach is to reduce the number of baths by both improving the qubit fabrication and carefully picking the wiring and shielding of the setup. The significant efforts in this direction have led to an improvement in coherence time of several orders of magnitude. Further discussion can be found in Ref. [96–98]. A different strategy involves engineering the Hamiltonian in a way that the qubit is topologically protected from noise. Examples of such superconducting qubits are the $0-\pi$ qubit[99] and the $\cos 2\varphi$ qubit[100]. Another promising direction is error mitigation[101], a suite of techniques aimed at suppressing, estimating, or compensating for errors at the software or control level. While such methods can substantially enhance the effective fidelity of quantum operations in Noisy Intermediate-Scale Quantum (NISQ) devices, they scale unfavorably with circuit depth (number of gate operations), limiting their practicality to relatively shallow quantum algorithms.

An alternative approach for reducing the decoherence relies on quantum error correction (QEC)[102]. By delocalizing the information, it becomes less susceptible to local noise, hence the total error rate is suppressed. The coherence of the resulting logical qubit is prolonged by detecting the occurrence of errors locally due to their undesired coupling with the environment and correcting for them. The process is depicted in Fig. 2.10. Generally, the logical qubit is encoded into $|\psi\rangle_L$, and continuous measurements map the leakage out of the computational space onto ancillary qubits. Such measurements are called *stabilizers* as they leave the logical state $|\psi\rangle_L$ unchanged. The stabilizers are designed in a way that they commute with the logical qubit operators but anti-commute with the errors. The ancillary qubits are measured, and the outcome S–*the error syndrome*–is fed to a decoder D. In the event of an error, a recovery operation $\mathcal{R}$ is applied, as determined by the decoder output.

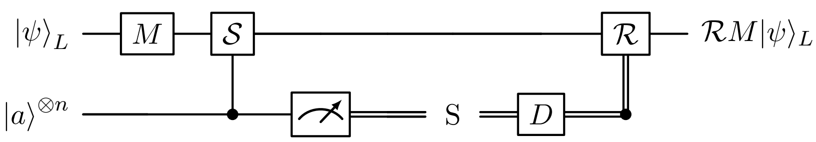

Figure 2.10: General quantum error correction circuit. The logical qubit is subjected to an error, represented by the gate M, stabilizer measurements $\mathcal{S}$ are executed with the help of ancillary qubits to extract the syndrome S. The latter is fed into a decoder D, and the outcome determines the respective recovery operation $\mathcal{R}$ applied. Figure adapted from Ref. [102].

Depending on the types of encoding, the two general subdivisions of QEC are discrete variable and continuous variable, which will be discussed with examples in the following subsections.

2.3.1 Discrete variable

Discrete variable encoding leverages two-level (qubits) or d-level (qudits) systems as quantum information carriers. Each physical system is described by a set of orthonormal basis states (e.g., $|0\rangle$ and $|1\rangle$ for qubits), and operations correspond to unitary transformations within this space. Discrete variable QEC relies on encoding a logical qubit into a larger number of physical qubits in a manner that allows errors to be detected and corrected without directly measuring and collapsing the quantum information. The amount of error suppression scales favorably with the number of qubits involved, provided that the single-qubit error rate is below a certain threshold[28, 102], determined by the type of encoding.

2.3.1.1 Repetition code

The repetition code is one of the simplest and most intuitive quantum error-correcting codes, as it most closely resembles classical error correction. It protects against bit-flip errors ($\hat{\sigma}_x$) by redundantly encoding the information of a single qubit across multiple physical qubits.

In its n-qubit form, the logical states are encoded as

$$|0\rangle_L \equiv |0\rangle^{\otimes n}, \quad |1\rangle_L \equiv |1\rangle^{\otimes n} \tag{2.38}$$

n is often referred to as the *code distance*. A majority vote across the n qubits allows for the detection and correction of up to $(n-1)/2$ bit-flip errors. For instance, the 3-qubit repetition code encodes a single logical qubit as

$$|0\rangle_L = |000\rangle\,, \quad |1\rangle_L = |111\rangle\,.$$

This code can detect and correct a single bit-flip error among the three qubits. To identify which qubit, if any, has flipped, one measures parity checks using ancillary qubits without collapsing the encoded quantum information. The stabilizers are

$$\hat{S}_1 = \hat{\sigma}_{z_1}\hat{\sigma}_{z_2}, \quad \hat{S}_2 = \hat{\sigma}_{z_2}\hat{\sigma}_{z_3}$$

and they are measured by performing controlled-X gates[13] with the data qubits as the control and ancillary qubits as the target (see Fig. 2.11).

The measurement outcomes of ancilla qubits a_1 and a_2 give a two-bit syndrome that identifies a single bit-flip error

[13] Controlled-X, or controlled-NOT (CNOT), gate flips the state of the target qubit only if the control qubit is in state $|1\rangle$.

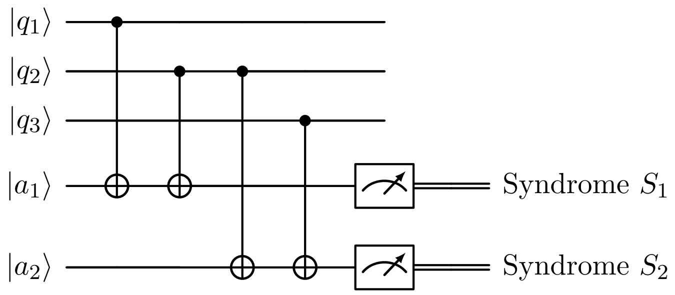

Figure 2.11: Repetition code circuit. The data qubits (q_i, $i = \{1, 2, 3\}$) store quantum information, while ancilla (a_j, $j = \{1, 2\}$) qubits are used to measure local stabilizers that detect errors. The controlled-X gate is represented by a vertical line connecting the control qubit (full circle) to the target qubit (empty circle).

Syndrome (a_1, a_2)	**Likely Error**
(0, 0)	No error
(1, 0)	Bit-flip on q_1
(1, 1)	Bit-flip on q_2
(0, 1)	Bit-flip on q_3

While the repetition code effectively corrects bit-flip errors, as demonstrated already in 2015[103] and later in 2021[104], it does not protect against phase-flip ($\hat{\sigma}_z$) errors. More advanced codes, such as the surface code, build upon the repetition code to offer full quantum error correction.

2.3.1.2 Surface code

The surface code is one of the most promising quantum error-correcting codes for scalable, fault-tolerant quantum computation due to its high error threshold[14]. The surface code operates on a 2D lattice of physical qubits ideally connected only by nearest-neighbor interactions. There are two types of stabilizers:

- X-type stabilizers, corresponding to phase-flip checks of neighboring qubits (Fig. 2.12c),
- Z-type stabilizers, corresponding to bit flip checks (Fig. 2.12b).

These stabilizers are measured repeatedly in time to detect both errors. Errors manifest as changes in the syndromes over time. A decoding algorithm (such as minimum-weight perfect matching[105]) is used to infer the most likely error configuration and apply corrections accordingly.

[14] In QEC, the error threshold refers to the single physical qubit error probability below which adding more physical qubits increases the error protection of the logical qubit.

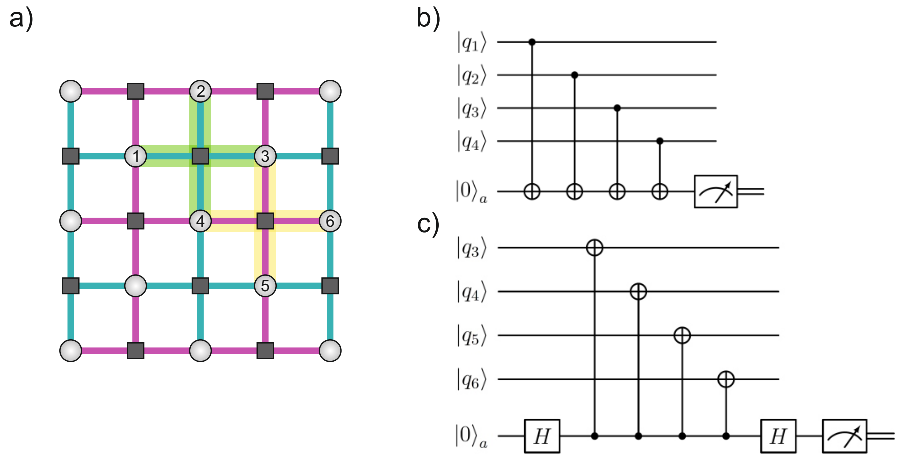

Figure 2.12: Surface code of distance three. a) Grid layout: circles represent data qubits, squares - ancillary qubits. The cyan and magenta lines depict controlled-X and controlled-Z operations, respectively. b) Z-stabilizer measurement, corresponding to the green highlighted region in a). c) X-stabilizer measurement, corresponding to the yellow highlighted region in a).

The logical X operator of a surface code acts on the data qubits along the boundary of the code on which the Z-stabilizers are applied (i.e., magenta lines in Fig. 2.12a). Similarly, the logical Z operator is defined as a chain of Z-operators on the qubits along a boundary on which the X-type stabilizers are applied (i.e., cyan lines in Fig. 2.12a). Because of how logical operations are executed, errors must form continuous chains to affect the logical information; this is the source of the code's topological protection.

Surface code with superconducting circuits operating below its error threshold has recently been achieved[106] after significant efforts from both academic groups[107] and industrial companies[104, 108]. This milestone is referred to as *break-even*, a point at which the logical qubit becomes better than any available physical qubit in the system. Despite achieving break-even, the surface code has an inherent drawback: it requires hundreds to thousands of physical qubits per logical qubit, which need to be controlled, and an elaborate decoder, which is computationally expensive. To run any useful algorithm, more than a thousand protected qubits are required[109], leading to a large hardware overhead.

2.3.2 Continuous variable

Continuous-variable (CV) systems offer an alternative encoding strategy that harnesses the infinite-dimensional Hilbert spaces of harmonic oscillators. The information is stored in a state of a bosonic mode, such as that of a superconducting LC resonator, hence the name *bosonic codes*. CV quantum error correction leverages these continuous degrees of freedom to protect quantum information against noise in a hardware-efficient manner, as the noise suppression increases by accessing higher photon-number states. CV codes are usually integrated into hybrid architectures where logical qubits are encoded in bosonic modes but controlled using DV systems[110].

2.3.2.1 Binomial Code

The *binomial code*[111] is a type of bosonic quantum error-correcting code that encodes a logical qubit in a finite superposition of Fock states of a harmonic oscillator, such as a superconducting cavity. This code is designed to correct for the most common error in bosonic systems–single photon loss.

Generally, a binomial code is defined by a distribution of Fock states with binomial weighting and fixed photon-number spacing, allowing for structured error tracking and syndrome recovery. The simplest binomial code, often referred to as *kitten code*, that corrects one photon loss, gain, and dephasing error, encodes logical states as

$$|0\rangle_L = \frac{1}{\sqrt{2}} \left(|0\rangle + |4\rangle \right), \quad |1\rangle_L = |2\rangle .$$

The Wigner function[112] representation of these states is depicted in Fig. 2.13. A photon loss would change the parity $\hat{\Pi} = \exp(i\pi\hat{a}^\dagger\hat{a})$ of the logical state, i.e., $\alpha |0\rangle_L + \beta |1\rangle_L \to \alpha' |3\rangle + \beta' |1\rangle$. The error space is orthogonal to the code space, allowing for the measurement of the syndrome via parity measurement with the help of an ancilla without disturbing the logical state. A recovery operation can then be applied, provided that the mean photon number of the logical states is the same. Otherwise, the state with a higher photon number is more likely to decay, resulting in altered weights after the recovery operation, compared to the initial state before the photon loss occurred.

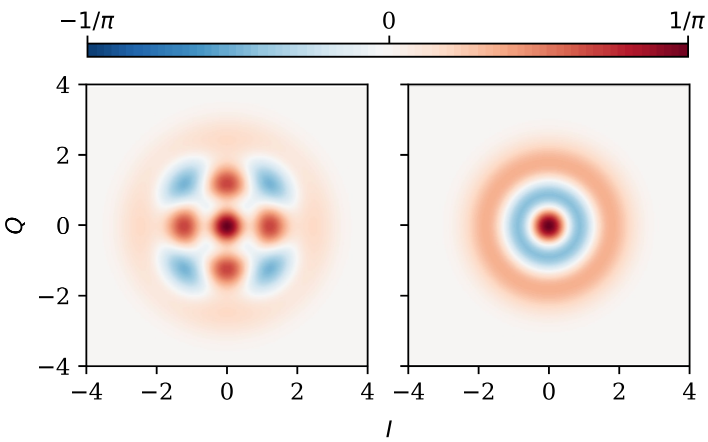

Figure 2.13: Wigner function representation of the logical states $|0\rangle_L$ (left) and $|1\rangle_L$ (right) for a binomial kitten code. The kitten code corrects for single-photon loss, an error that would dramatically alter the Wigner function representation (see Fig. 4 in Ref. [113]).

The first implementation of binomial code in superconducting circuits[114] fell short of reaching break-even compared to a single photon Fock state encoding, mainly limited by the ancilla. A more recent experiment[56] then managed to prolong the cavity lifetime by 16%, successfully reaching break-even, but was limited by two-photon loss, which the kitten code cannot correct for.

2.3.2.2 Cat code

The cat code encodes logical qubits into superpositions of coherent states (i.e., "cat states") of a single harmonic oscillator mode. The superposition state that consists of $2N$ coherent states is protected against the loss of $N-1$ photons. Here, we focus on the simplest cat code–the two-legged case, where $N=2$.

The cat states are a superposition of coherent states and are defined as

$$|C^{\pm}\rangle = \mathcal{N}_{\pm}(|\alpha\rangle \pm |-\alpha\rangle) \tag{2.39}$$

where the normalization factor $\mathcal{N}_{\pm} = [2 \pm 2\exp(-2|\alpha|^2)]^{-0.5}$ equals approximately $1/\sqrt{2}$ for large $|\alpha|^2$. These two states have opposite parity, even and odd, which is determined by the superposition of Fock states that make up the cat state. The two logical states are a superposition of the two cat states from Eq. 2.39:

$$|0\rangle_L = \frac{1}{\sqrt{2}}(|C^+\rangle + |C^-\rangle) = |\alpha\rangle + \mathcal{O}(e^{-2|\alpha|^2}), \quad |1\rangle_L = \frac{1}{\sqrt{2}}(|C^+\rangle - |C^-\rangle) = |-\alpha\rangle + \mathcal{O}(e^{-2|\alpha|^2}). \tag{2.40}$$

These states are located at the poles of the Bloch sphere, while the cat states with different parity are situated on the equator (see Fig. 2.14a). The loss of a photon flips the parity of the cat states. The parity is the code stabilizer. The syndrome is continuously observed via an ancillary qubit. After an error occurs, a parity flip is detected; however, the state remains within the code space and therefore cannot be corrected.

Increasing the amplitude $|\alpha|^2$ improves the distinguishability of the two logical states, as their overlap scales as $e^{-4|\alpha|^2}$. On the other hand, the photon loss rate is enhanced by the mean photon number $\bar{n} = |\alpha|^2$, i.e., $\kappa_1^{\text{eff}} = \kappa_1 \bar{n}$, and thus increases dephasing. Favoring one type of error creates a *noise bias*, which is used to develop autonomous protection against bit flips in the case of a cat code. In contrast, the phase flips need to be corrected by another layer of encoding, such as the repetition code.

To maintain the bit-flip protection at high mean photon numbers, the two-legged cat codes are usually stabilized; otherwise, the state amplitude would decay, and the logical states would become indistinguishable. The stabilization can be either active (measurement-based)[115] or passive (Hamiltonian or dissipation engineering)[38, 116]. Stabilized cat codes have been successfully demonstrated in several circuit QED platforms, with the first experiment dating back to 2016[115]. In this work, four-legged cat states were actively stabilized using real-time feedback, with pulses carefully selected to counteract the evolution of the quantum system. The two flavors of passive stabilization will be discussed in more detail.

Kerr cat code This stabilization scheme involves a parametric pump tone that helps engineer the Hamiltonian in a way that stabilizes the two lowest energy eigenstates of the system. A system suited for this type of stabilization is a superconducting cavity, coupled to a Superconducting Nonlinear Asymmetric Inductive eLement (SNAIL) with a

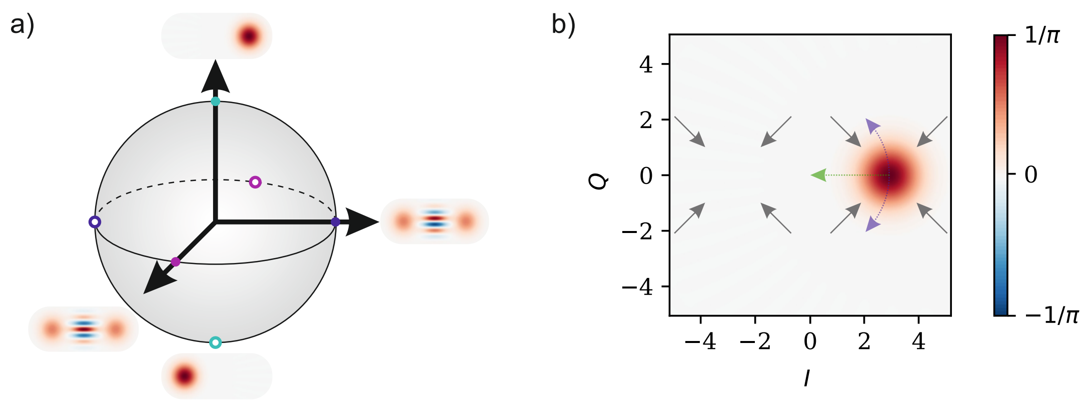

Figure 2.14: Cat code. a) Bloch sphere of the cat code. The states on the poles correspond to the logical states from Eq. 2.40, while on the equator $|\pm X\rangle$ are the states from Eq. 2.39, and $|\pm Y\rangle = \mathcal{N}_{\pm}(|\alpha\rangle \pm i\,|-\alpha\rangle)$. Only the Wigner function representations of $|\pm Z\rangle, |+X\rangle$ and $|+Y\rangle$ are shown. b) Passively stabilized coherent state. To compensate for the energy decay (green dotted arrow) and avoid dephasing (purple dotted arrow) due to, e.g. self-Kerr effect, the state needs to be stabilized (gray arrows). For a two-legged cat the manifold spanned by $\{|\alpha\rangle, |-\alpha\rangle\}$ is stabilized, where $\pm\alpha = \pm\sqrt{\epsilon_2/K_c}$ for a Kerr cat type of stabilization, or $\pm\alpha = \pm\sqrt{-\epsilon_d/g_2^*}$ for a dissipative cat type of stabilization.

squeezing drive applied to it[116]. The system Hamiltonian at the frame rotating at the cavity frequency, and after RWA in the limit of small cavity nonlinearity, is

$$\frac{\hat{H}_{Kerr}}{\hbar} = -K_c(\hat{a}^\dagger)^2\hat{a}^2 + (\epsilon_2(\hat{a}^\dagger)^2 + \epsilon_2^*\hat{a}^2) \tag{2.41}$$

where ϵ_2 is the squeezing drive amplitude and K_c is the cavity non-linearity, inherited from the SNAIL, also known as the self-Kerr coefficient. The eigenstates of this system are the coherent states $|\pm\alpha\rangle$, where $\pm\alpha = \pm\sqrt{\epsilon_2/K_c}$. These states are degenerate and separated from the rest of the spectrum by $4K_c\bar{n}$, which determines the gate and readout speed. Increasing the pump amplitude results in higher protection against bit flips. However, this behavior is yet to be observed experimentally. Instead, a staircase dependence of the bit-flip time on the photon number[117] indicates that the residual thermal population causes excess bit-flip errors. A thermal excitation to higher states enables tunneling between the potential wells, causing a bit flip[118]. It should be noted that recent efforts[119] combine Hamiltonian with dissipation engineering to mitigate leakage out of the computational manifold.

Dissipative cat code An alternative way of stabilization is to engineer the environment, i.e., create a dissipation operator that stabilizes the computational space[120]. A lossy mode, often referred to as *buffer*, is introduced to the long-lived *memory* mode. Their interaction is given by

$$\frac{\hat{H}_{\text{int}}}{\hbar} = g_2\hat{m}^2\hat{b}^\dagger + g_2^*(\hat{m}^\dagger)^2\hat{b}. \tag{2.42}$$

Here, $\hat{m}, \hat{b}$ are the annihilation operators for the memory and buffer, respectively, and g_2 is the rate at which two memory photons are converted to one photon in the buffer. The buffer is exposed to single photon loss at a rate κ_b, so an additional drive is applied to it to compensate for the energy loss

$$\frac{\hat{H}_d}{\hbar} = \epsilon_d \hat{b}^\dagger + \epsilon_d^* \hat{b}. \tag{2.43}$$

This drive also repopulates the memory mode in pairs of photons due to the coupling with the buffer. Depending on the physical implementation, it is either combined with a strong pump[121] that matches the frequency difference $\omega_p = 2\omega_m - \omega_b$ or it acts alone if the condition $2\omega_m = \omega_b$ is satisfied[122].

In the limit $\kappa_b \gg g_2$, we can perform adiabatic elimination (see Appendix B), and the effective dissipation of the memory becomes

$$\hat{L}_2 = \sqrt{\kappa_2}(\hat{m}^2 - \alpha^2) \tag{2.44}$$

with an effective two-photon loss rate $\kappa_2 = 2|g_2|^2/\kappa_b$. This dissipator stabilizes the manifold spanned by $\{|\alpha\rangle, |-\alpha\rangle\}$ with $\pm\alpha = \pm\sqrt{-\epsilon_d/g_2^*}$.

The first proposal for a dissipative cat code[38] includes a transmon coupled to both the memory and buffer modes. Because the Josephson junction by itself is a four-wave mixing element, another strong pump is applied to compensate for the energy difference. The first experimental demonstration of two-photon stabilization[120] followed shortly after the theory proposal, but it was limited by the short lifetime of the memory and by the low ratio κ_2/κ_1. In the following generation of dissipative cat qubits[121], roles of tomography and stabilization were split between two distinct elements: the transmon and the Asymmetrically Threaded SQUID (ATS). In addition to the resonant microwave drive on the buffer, a flux parametric pump is applied to the ATS for energy matching of the two-photon conversion process, i.e. $\omega_p = 2\omega_m - \omega_b$, making the conversion rate dependent on the drive amplitude.

In contrast to the Kerr cat variation, this realization of the dissipative cat code has demonstrated exponential suppression of the bit-flip rate in 2020[121]. The main limitation was the heating of the transmon ancilla, resulting in saturation of the exponential suppression after a certain cat size. In a more recent experiment[24], removing the ancilla allowed for measuring bit flip times above 10 seconds. The caveat is that, due to the weaker higher-order nonlinearities used for state tomography, the measurement times were increased drastically compared to other cat code experiments. A way out of this situation is to utilize resonant two-photon conversion in a three-wave mixing setup, thereby eliminating the need for additional pump tones without compromising tomography speed. Such a platform[123] showed bit flip times up to 0.3 seconds.

2.3.2.3 Gottesman-Kitaev-Preskill (GKP) Code

The GKP code provides intrinsic protection against the dominant noise channels in bosonic systems by discretizing continuous phase-space displacements caused by photon loss or dephasing into correctable shift errors. This code encodes a qubit into a harmonic oscillator using grid states[124]: superpositions of Dirac delta functions[15] on a lattice in phase space. Depending on the lattice cell, the code can be square, rectangular, or hexagonal; here, we discuss only the square variation. The ideal Gottesman-Kitaev-Preskill square code states are given by

$$|0\rangle_L \propto \sum_{n\in\mathbb{Z}} |x = n\sqrt{2\pi}\rangle$$
$$|1\rangle_L \propto \sum_{n\in\mathbb{Z}} \left|x = \left(n + \frac{1}{2}\right)\sqrt{2\pi}\right\rangle$$

and they also form a grid in momentum space

$$|0\rangle_L \propto \sum_{n\in\mathbb{Z}} \left|p = \frac{n}{2}\sqrt{2\pi}\right\rangle$$
$$|1\rangle_L \propto \sum_{n\in\mathbb{Z}} (-1)^n \left|p = \frac{n}{2}\sqrt{2\pi}\right\rangle$$

The ground state is centered at $(x, p) = (0, 0)$, and the excited at $(\sqrt{\pi/2}, 0)$. They remain identical under displacements of $\sqrt{2\pi}$ in both position and momentum, meaning that the X and Z stabilizers are $\hat{D}(\sqrt{2\pi})$ and $\hat{D}(i\sqrt{2\pi})$, respectively. The logical Pauli operations are displacement operators by half the amplitude of the stabilizers. They commute with the stabilizers up to a phase.

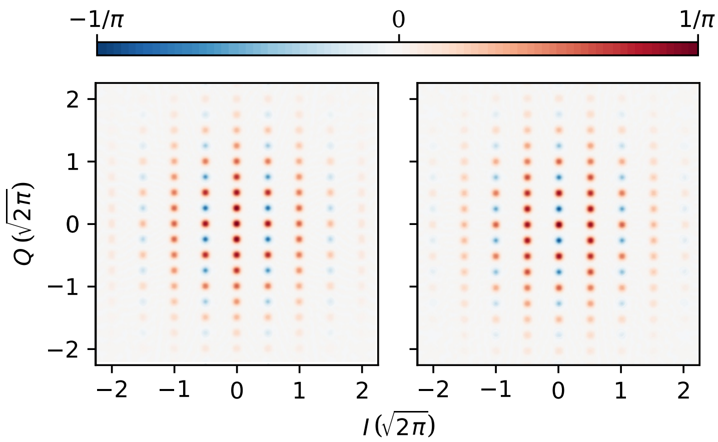

Figure 2.15: Wigner function representation of the logical states $|0\rangle_L$ (left) and $|1\rangle_L$ (right) for a finite energy GKP code. At the same positions where the ideal GKP code has Dirac delta functions, here there are finitely squeezed states. The GKP code corrects for small displacements in I and Q, meaning that the error space would have similar Wigner representations to the computational space.

GKP states are unphysical as they require infinite energy, but they provide the theoretical foundation for the approximate GKP code used in practice. Approximate versions of the ideal states can be constructed with finite squeezing of the individual states and by applying a Gaussian envelope to the ideal codewords[125, 126], as shown in Fig. 2.15. In contrast to the binomial and cat codes, which are rotationally symmetric, the GKP code is

[15] By definition $\delta(x) = \begin{cases} 0, & x \neq 0 \\ \infty, & x = 0 \end{cases}$

translationally symmetric and can correct small displacement errors in both quadratures. In the ideal case, any displacement smaller than $\sqrt{\pi/2}$ can be corrected for. Still, in reality, the correction is limited by the squeezing: the higher the squeezing of the states, the smaller the displacements that can be corrected for. The single photon loss sets a limit on the possible squeezing of the states and, therefore, the code's correction capabilities.

GKP states have been created in the motional modes of a trapped ion chain[127], preceding the first implementation with superconducting circuits[126]. The latter experiment succeeded in breaking even after greatly enhancing the lifetime and coherence of the transmon ancilla[22]. More recently, the break-even has been reached with GKP encoding in a qudit[128]. Despite reaching such a milestone, the platforms were still limited by the ancilla qubit.

2.3.3 Withstanding challenges

Significant experimental progress has been made in both DV and CV QEC with superconducting circuits, including reaching break-even. Despite these advances, DV QEC remains resource-intensive. Realistic surface code implementations require on the order of hundreds to thousands of physical qubits per logical qubit. Furthermore, the process of syndrome extraction introduces additional error channels, which, combined with the need for precise calibration, qubit connectivity, and cross-talk mitigation, complicates scaling up towards a quantum computer.

Unlike DV codes, CV QEC can offer protection using a single physical mode, drastically reducing overhead. Nevertheless, CV codes are not without challenges. The main limitation of most experimental platforms remains–the ancilla, which is required for bosonic control. Additionally, photon loss and Kerr nonlinearity (except for the case of a Kerr-cat code) pose obstacles to the scaling. In practice, hybrid schemes, where bosonic codes serve as building blocks for DV codes, are also under active investigation. Examples of such are the repetition code on top of dissipative cat code, as well as the surface-GKP code[129].

Continued advances in coherence, control, and measurement are essential for scalable quantum error correction in superconducting platforms. This thesis suggests an approach to relieve some of the aforementioned limitations. By carefully picking the bosonic mode and the ancilla, the detrimental effects of either the Kerr effect or the finite lifetime of both the bosonic mode or the ancilla can be alleviated.

CHAPTER 3

Circuit design

A basic system for bosonic quantum information processing consists of a microwave harmonic oscillator, a superconducting qubit as an ancilla, and a readout resonator. Determining the full Hamiltonian of such a composite system requires different approaches depending on the physical dimensions of its constituents compared to the wavelengths of interest. Superconducting circuits normally operate in the range of a few gigahertz, as their thermal occupation is negligible once cooled down to cryogenic temperatures in the range of a few millikelvin. This chapter focuses on the different ways of describing a microwave system and focuses on a specific use case of the setup in this work. The setup consists of a microwave cavity, which serves the purpose of a memory mode, coupled capacitively to the readout resonator of a fluxonium, which both share an inductor. The final goal of this chapter is to find a model for describing the system that allows us to predict its behavior.

3.1 Lumped element simulation

The most common technique involves a simplified representation of a physical system that assumes each component to be treated as a single, idealized element with all of its properties concentrated ("*lumped*") into one point, rather than spread out in space. This approximation is valid as long as the physical size of the circuit is much smaller than the wavelength of the signals involved, so it ignores wave propagation and treats voltage and current as being uniform across each component. In this scenario, one can draw the equivalent electrical circuit of the system (e.g., Fig. 2.3a, Fig. 2.7a, Fig. 2.8a). If, however, the size of the component is comparable to or larger than the wavelength of the signals it carries, it is referred to as *distributed*.

In the lumped element approximation, we combine the circuits from Fig. 2.3a and Fig. 2.7a to the corresponding composite circuit as depicted in Fig. 3.1. The following derivations were done by our collaborator Dr. Daria Gusenkova[1], during her PhD at Karlsruhe Institute of Technology (KIT).

[1]Current address: IQM, Munich, Germany.

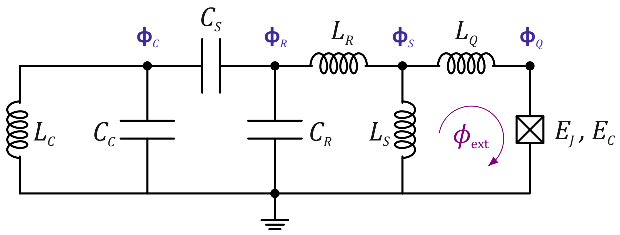

Figure 3.1: Lumped element circuit representation of the system with the reduced flux at each node $\phi_i = \Phi_i/2\pi$.

The task starts with deriving the classical Lagrangian[2]

$$\begin{aligned}\mathcal{L} =& \frac{C_C}{2}\dot{\Phi}_C^2 + \frac{C_R}{2}\dot{\Phi}_R^2 + \frac{C_Q}{2}\dot{\Phi}_Q^2 + \frac{C_s}{2}(\dot{\Phi}_R - \dot{\Phi}_C)^2 \\ &- \frac{1}{2L_C}\Phi_C^2 - \frac{1}{2L_R}(\Phi_R - \Phi_s)^2 - \frac{1}{2L_Q}(\Phi_Q - \Phi_s)^2 - \frac{1}{2L_s}\Phi_s^2 + E_J\cos(\varphi_Q - \varphi_{\text{ext}})\end{aligned} \tag{3.1}$$

where Φ_i, $i = \{C, R, Q, s\}$ are the fluxes at each node close to a certain component (cavity, resonator, qubit, coupling element), as defined previously in the thesis (see Sec. 2.2.1.1), φ_Q is the phase difference across the junction, and φ_{ext} is the phase offset caused by the externally applied field. To derive this Lagrangian, the method of branches was used[61].

The Euler-Lagrange equation of motion for the variable Φ_s is

$$\frac{d}{dt}\frac{\partial\mathcal{L}}{\partial\dot{\Phi}_s} = \frac{\partial\mathcal{L}}{\partial\Phi_s}. \tag{3.2}$$

Since $\frac{\partial\mathcal{L}}{\partial\dot{\Phi}_s} = 0$, the system dynamics does not depend on this variable, which allows us to reduce the number of variables by making the following substitution

$$\Phi_s = \frac{L_R L_Q L_s}{L_\Sigma^2}\left(\frac{\Phi_R}{L_R} + \frac{\Phi_Q}{L_Q}\right) \tag{3.3}$$

where we have defined $L_\Sigma^2 = L_R L_Q + L_R L_s + L_Q L_s$. With the conjugate charge $Q_i = \frac{\partial\mathcal{L}}{\partial\dot{\phi}_i}$ for each element, we can derive the Hamiltonian in the following way

$$H = \sum_i \dot{\Phi}_i \frac{\partial\mathcal{L}}{\partial\dot{\Phi}_i} - \mathcal{L} \tag{3.4}$$

which finally gives us the system Hamiltonian

$$\begin{aligned}H =& \frac{C_R + C_s}{2C_\Sigma^2}Q_C^2 + \frac{C_C + C_s}{2C_\Sigma^2}Q_R^2 + \frac{1}{2C_Q}Q_Q^2 \\ &+ \frac{1}{2L_C}\Phi_C^2 + \frac{L_Q + L_s}{2L_\Sigma}\Phi_R^2 + \frac{L_R + L_s}{2L_\Sigma}\Phi_Q^2 \\ &+ \frac{C_s}{C_\Sigma^2}Q_C Q_R - \frac{L_s}{L_\Sigma^2}\Phi_R\Phi_Q - E_J\cos\varphi_Q\end{aligned} \tag{3.5}$$

with a similar normalization factor for the capacitive elements $C_\Sigma^2 = C_R C_s + C_C C_R + C_C C_s$. Note that we have neglected the contribution from the external field for now.

[2]From this point on, all hats of the operators are omitted.

From this point on, one can proceed with the diagonalization in two ways: in the normal mode (dressed) basis or in the bare harmonic oscillator (Fock) basis. For the former one, following the method from Ref. [130], we rewrite the linear part of the Hamiltonian in matrix form

$$H_{\text{lin}} = \frac{1}{2}\mathbf{Q}^T||C||^{-1}\mathbf{Q} + \frac{1}{2}\mathbf{\Phi}^T||L||^{-1}\mathbf{\Phi} \tag{3.6}$$

where $\mathbf{Q} = (Q_C, Q_R, Q_Q)^T$ and $\mathbf{\Phi} = (\Phi_C, \Phi_R, \Phi_Q)^T$ are the charge and flux vectors. The capacitance and inductance matrices are defined as

$$||C||^{-1} = \begin{pmatrix} (C_R + C_s)/C_\Sigma^2 & C_s/C_\Sigma^2 & 0 \\ C_s/C_\Sigma^2 & (C_C + C_s)/C_\Sigma^2 & 0 \\ 0 & 0 & 1/C_Q \end{pmatrix} \tag{3.7a}$$

$$||L||^{-1} = \begin{pmatrix} 1/L_C & 0 & 0 \\ 0 & (L_Q + L_s)/L_\Sigma^2 & -L_s/L_\Sigma^2 \\ 0 & -L_s/L_\Sigma^2 & (L_R + L_s)/L_\Sigma^2 \end{pmatrix} \tag{3.7b}$$

respectively. Their units correspond to [F^{-1}] and [H^{-1}], hence the power of -1.

Diagonalizing this matrix yields three normal modes — cavity-like, readout-like, and qubit-like — rescaled due to the coupling elements L_s and C_s. Those modes are coupled only by the non-linear cosine term, which is a function of the qubit phase $\varphi_Q = \sum_{i=\{q,r,c\}} \lambda_i \varphi_i$, now a linear combination of all normal phases. In the normal mode harmonic oscillator basis $\{|qcr\rangle\}$, where q, c and r correspond to the number of excitations in the qubit-like, cavity-like, and readout-like modes, the Hamiltonian can be expressed with an infinite-dimensional matrix

$$\begin{aligned} H/\hbar &= H_{\text{lin}} - E_J \cos(\varphi - \varphi_{\text{ext}}) \\ &= \sum_{q,c,r} (\tilde{\omega}_q q + \tilde{\omega}_c c + \tilde{\omega}_r r)|qcr\rangle\langle qcr| \\ &- E_J \Big[\cos(\varphi_{\text{ext}}) \left(c_{qq'}c_{cc'}c_{rr'} - s_{qq'}s_{cc'}c_{rr'} - s_{qq'}c_{cc'}s_{rr'} - c_{qq'}s_{cc'}s_{rr'} \right) \\ &+ \sin(\varphi_{\text{ext}}) \left(s_{qq'}c_{cc'}c_{rr'} + c_{qq'}s_{cc'}c_{rr'} + c_{qq'}c_{cc'}s_{rr'} - s_{qq'}s_{cc'}s_{rr'} \right) \Big] |qcr\rangle\langle q'c'r'| \end{aligned} \tag{3.8}$$

where $\tilde{\omega}_i$, $i = \{q, c, r\}$ are the Lamb-shifted mode frequencies extracted from the diagonalization of Eq. 3.6. The cosine $c_{ii'}$ and sine $s_{ii'}$ matrix elements can be computed analytically with the Laguerre polynomials

$$\begin{aligned} c_{kl}(\lambda_i\varphi_i) &= \langle k|\cos(\lambda_i\varphi_i)|l\rangle \\ &= \begin{cases} (-1)^{\frac{l-k}{2}} \sqrt{\frac{k!}{l!}} (\lambda_i\varphi_{i,zpf})^{l-k} e^{-(\lambda_i\varphi_{i,zpf})^2/2} \mathscr{L}_k^{l-k} \left[(\lambda_i\varphi_{i,zpf})^2 \right] & \text{for } k+l \text{ even} \\ 0 & \text{for } k+l \text{ odd} \end{cases} \end{aligned} \tag{3.9a}$$

$$\begin{aligned} s_{kl}(\lambda_i\varphi_i) &= \langle k|\sin(\lambda_i\varphi_i)|l\rangle \\ &= \begin{cases} 0 & \text{for } k+l \text{ even} \\ (-1)^{\frac{l-k+1}{2}} \sqrt{\frac{k!}{l!}} (\lambda_i\varphi_{i,zpf})^{l-k} e^{-(\lambda_i\varphi_{i,zpf})^2/2} \mathscr{L}_k^{l-k} \left[(\lambda_i\varphi_{i,zpf})^2 \right] & \text{for } k+l \text{ odd.} \end{cases} \end{aligned} \tag{3.9b}$$

For clarity, I have omitted the arguments of the matrix elements, but I summarize them here: $y_{qq'}(\lambda_q\varphi_q), y_{rr'}(\lambda_r\varphi_r), y_{cc'}(\lambda_c\varphi_c)$, where $y = \{s, c\}$.

Usually, the diagonalization and matrix element evaluation are done numerically with the truncation of the Hilbert space for all available modes. As calculating the matrix elements with Laguerre polynomials is computationally intensive, another method is presented, following Ref. [30]. In the bare basis, $|QCR\rangle$ with phase variables $\{\varphi_Q, \varphi_R, \varphi_C\}$, we use Eq. 3.5 to create a Hamiltonian matrix. Its diagonal terms contain the rescaled bare frequencies[3], as well as the matrix elements corresponding to $Q = Q', C = C', R = R'$.[4] On the other hand, the off-diagonal terms of the Hamiltonian are given by both the inductive and capacitive coupling elements, L_s and C_s,[5] as well as the matrix elements corresponding to $Q \neq Q', C = C', R = R'$. In this case, all the Laguerre polynomials are calculated only for the phase of the fluxonium φ_Q, which substantially reduces the computational power required for the diagonalization, allowing for the simulation of high-power readout[131]. This method is used later in this thesis in Sec. 6.1.

3.2 Distributed element simulation of weakly anharmonic systems

In the regime of weak nonlinearity, quantum electrical circuits can be effectively diagonalized using the *black-box quantization* (BBQ) framework. Within this method, the Josephson junction is modeled by explicitly including its shunt capacitance C_J, and the cosine potential is expanded to isolate the leading-order linear inductive term L_J. The residual nonlinear contributions—arising from the higher-order terms in the expansion—are collectively captured by an abstract component known as the *"spider" element*. By assigning a

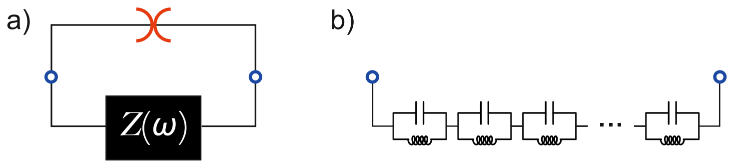

Figure 3.2: Black box quantization working principle. a) The spider element is depicted in red, while the rest of the circuit is grouped in a black box. The open blue circles represent the port. b) Equivalent circuit of the black box from a).

port at the spider element, the impedance of the surrounding linear network can be evaluated (see Fig. 3.2). For a lossless system without voltage or current sources, the impedance function of an arbitrary network can be expressed using *Foster's reactance theorem*[132] as follows

[3]For instance, the rescaled cavity frequency is $\bar{\omega}_C = \sqrt{8\bar{E}_C\bar{E}_L}$, where $\bar{E}_C/h = e^2\left(\frac{C_R+C_s}{2C_\Sigma^2}\right)/2h$ and $\bar{E}_L/h = h/16\pi^2e^2L_C$.

[4]i.e. $-E_J\left[\cos\varphi_{\text{ext}}c_{QQ}(\varphi_Q) - \sin\varphi_{\text{ext}}s_{qq}(\varphi_Q)\right]$.

[5]Or, equivalently, from the terms $\sim Q_CQ_R$ and $\sim \Phi_R\Phi_Q$ in Eq. 3.5.

$$Z(\omega) = \sum_k \left(j\omega C_k + \frac{1}{j\omega L_k} \right)^{-1} = \sum_k \frac{j\omega / C_k}{\omega_k^2 - \omega^2}, \tag{3.10}$$

where k indexes the normal modes of the circuit, and $j = i$. The resonance frequencies $\omega_k = 1/\sqrt{L_k C_k}$ correspond to the zeros of the admittance $Y(\omega) = 1/Z(\omega)$, satisfying $Y(\omega_k) = 0$. The mode capacitances are obtained from the frequency derivative of the admittance

$$\left. \frac{\partial Y(\omega)}{\partial \omega} \right|_{\omega=\omega_k} = j2C_k. \tag{3.11}$$

Given the impedance from Eq. 3.10, the linearized Hamiltonian of the circuit can be written as a sum of independent harmonic oscillators

$$H_{\text{lin}} = \sum_k \left(\frac{Q_k^2}{2C_k} + \frac{\Phi_k^2}{2L_k} \right) = \sum_k \hbar \omega_k a_k^\dagger a_k, \tag{3.12}$$

where Q_k and Φ_k are the canonical conjugate charge and flux operators, respectively, and $a_k^\dagger$, a_k are the bosonic creation and annihilation operators for each mode. The corresponding zero-point flux fluctuations are given by $\phi_k^{\text{zpf}} = \sqrt{\hbar/(2\omega_k C_k)}$.

According to Kirchhoff's voltage law, the total node flux across the nonlinear element (i.e., the spider element) is equal to the sum of flux contributions from each mode, up to an arbitrary reference phase (which is set to zero). Applying the second Josephson relation, the total phase difference across the junction can thus be expressed as

$$\varphi = \sum_k \Phi_k = \sum_k \phi_k^{\text{zpf}} (a_k^\dagger + a_k). \tag{3.13}$$

Substituting this expression into the quartic (and higher-order) terms of the cosine potential (see Eq. 2.28) and using the Rotating-Wave Approximation (RWA) yields the nonlinear interaction terms in the Hamiltonian. These include the dispersive shift as well as self- and cross-Kerr interactions between modes, all of which are essential for capturing the system's anharmonic dynamics.

In practice, this theoretical framework is typically used in conjunction with finite-element simulations, such as those performed in *Ansys HFSS*. These simulations provide accurate evaluations of the circuit admittance, regardless of whether the system comprises lumped elements, distributed elements, or a combination thereof. The extracted modal parameters are then employed to compute quantities of interest, including anharmonicities, mode couplings, dispersive shifts, and Kerr coefficients, thereby enabling precise modeling of complex superconducting quantum devices. Due to the high zero-point phase fluctuation inherent to the fluxonium, RWA is invalid[133, 134], deeming this method incompatible with the system of interest. Moreover, in contrast to the lumped element model, this method is incapable of accounting for flux quantization and, therefore, does not allow for loops in the circuit under consideration.

3.3 Distributed element simulation of highly anharmonic systems

The high nonlinearity of the fluxonium, in combination with distributed elements coupled to it, makes the system hard to simulate. Finite element simulators, such as *Ansys HFSS*, in conjunction with either BBQ or the *Qiskit Metal* extension: *pyEPR*[135], have until recently[136] failed to simulate highly anharmonic qubits. *scQubits*[137] provides a possible solution, but a few subtleties must be addressed.

The full Hamiltonian of the system used in this thesis reads

$$H = H_Q + H_R + H_C + H_{QR} + H_{RC}, \tag{3.14}$$

where H_Q is the fluxonium Hamiltonian from Eq. 2.30. We represent the resonator and cavity parts in their normal mode bases, i.e. $H_R = \hbar\omega_R r^\dagger r$ and $H_C = \hbar\omega_C c^\dagger c$ with ω_R, ω_C being the fundamental mode frequencies of the resonator and the cavity, respectively, and r, c - their annihilation operators respectively. The inductive coupling between the qubit and the resonator is modeled as $H_{QR} = -\hbar g_{QR}\varphi(r^\dagger + r)$, where g_{QR} is the interaction strength. Additionally, the capacitive coupling between the resonator and the cavity is $H_{RC} = \hbar g_{RC}(r^\dagger - r)(c^\dagger - c)$ with the interaction strength g_{RC}. This model does not assume any of the approximations used in the previously mentioned methods. To diagonalize the system Hamiltonian, *scQubits* takes as inputs the fluxonium energies E_C, E_L and E_J, the fundamental frequencies of the harmonic oscillators $\omega_C/2\pi$ and $\omega_R/2\pi$, and the coupling strengths between the elements g_{RC} and g_{QR}.

With the help of *scQubits*, all system parameters, such as dispersive shifts, Kerr nonlinearities, matrix elements, etc., can be directly calculated after experimentally extracting the couplings g_{RC} and g_{QR}. This method was successfully employed in this thesis in Chap. 5.

CHAPTER 4

Experimental setup

In the experimental setup, shown in Fig. 4.1, an aluminum post cavity serves as the memory mode. The control is provided by a fluxonium, consisting of an aluminum-aluminum oxide-aluminum (Al-AlO_x-Al) Josephson junction, shunted by a granular aluminum superinductance. The fluxonium is inductively coupled to a readout resonator via a shared grAl section. The readout resonator shares a capacitance with the cavity. The additional control knob for the magnetic bias of the fluxonium is provided by a small superconducting coil sitting on top of a home-built magnetic flux hose, which guides the field into the superconducting enclosure. This chapter focuses on each of the described elements and their improvements over time.

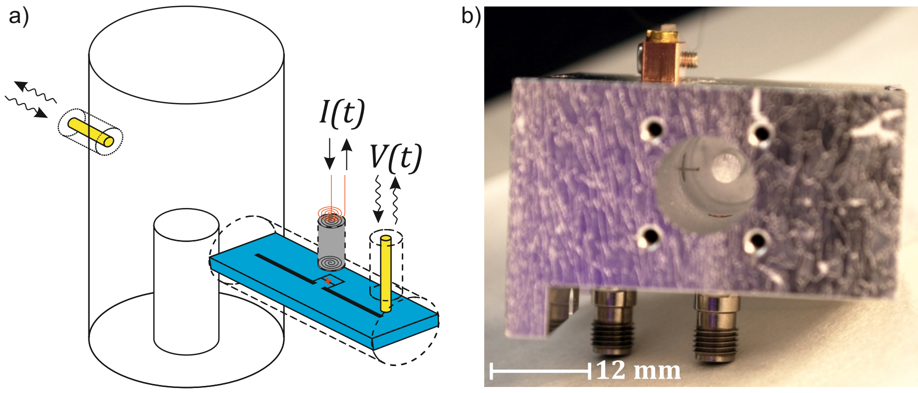

Figure 4.1: Experimental setup. a) The device consists of a high-Q cavity, fluxonium (the junction is marked with a red cross), a readout resonator on a sapphire chip (blue), and a flux hose (gray) and superconducting coil (red), and coupling pins (yellow). b) A photograph of the actual setup used. The chip is visibly poking into the cavity.

4.1 Post cavity

The workhorse of any bosonic qubit is the memory mode. The main source of error in microwave cavities is the single-photon loss. As mentioned in Sec. 2.2.1.1, 3D architectures are longer-lived due to the reduced energy participation in the lossy surfaces[135]. In particular, the lifetime of the seamless cavities and their ease of manufacturing make them the preferred choice for bosonic information processing.

In this thesis, a post cavity was used. As depicted in Fig. 4.1a, it consists of a coaxial part, short-circuited at one end and open at the other, as well as a circular waveguide part. The former forms a $\lambda/4$ cavity, where the height of the post determines the frequency of the mode $l \approx \lambda/4$[1]. As the cut-off frequency of the waveguide section is much lower than the fundamental frequency of the cavity, the mode can only propagate evanescently, making the seam at the top virtually invisible for the mode. This effect is enhanced by elongating the waveguide section[2].

4.1.1 Design, machining and etching

To design a high-Q cavity, several key parameters must be taken into account. Here, the length of the stub is 14.8 mm and has been chosen to achieve the frequency of the fundamental mode, 4.5 GHz. This length has been determined after eigenmode simulations in Ansys HFSS. The outer radius of the coaxial cavity is 5.8 mm, while the inner is 2.2 mm. They were picked in a way that, after etching, their ratio is three to one to minimize the losses from surface resistance[138].

The waveguide section continues 39.2 mm after the top of the stub. The waveguide supports several modes with different propagation constants: TE and TM modes. The former one have propagation constant[62] $\beta_{nm} = \sqrt{\left(\frac{2\pi}{\lambda}\right)^2 - \left(\frac{p'_{nm}}{b}\right)^2}$, with p'_{nm} being the m-th root of the derivative of Bessel function of first kind J'_n with respect to its argument. The latter has the same propagation constant except p'_{nm} is now replaced by p_{nm}, which is the m-th root Bessel function of the first kind J_n. The mode with the smallest propagation constant propagates the fastest, and the electric field in the cavity decays as $|\vec{E}| \propto e^{-2i\beta z}$, meaning the TE_{11} mode would be the first to carry away any excitation from the cavity. For the quoted dimension, the electric field at the end of the waveguide section is suppressed to approximately $1.6 \cdot 10^{-10}$ of its original strength.

The tunnel, where the qubit chip is hosted, has a radius of 2 mm, a value picked as a trade-off between ease of manufacture and cavity leakage into this tunnel. Its axial position is aligned with the end of the post. The chip orientation can be used to reduce the coupling's susceptibility to variations in the chip position, and, therefore, to chip vibrations, compared to the setup in Ref. [139]. Moreover, in this configuration, the coupling is less susceptible to vibrations that might occur due to the pulse tube of the cryostat or to any imperfections in the chip positioning. The chip itself dielectrically loads the cavity, lowering its frequency.

The cavity is machined into a block of high-purity aluminum. The waveguide and other tunnels are milled away, while the coaxial part is manufactured using a plunge erosion electric discharge machining (EDM) process, with the electrode specifically designed to match the cavity dimensions. The ready piece is later cleaned with acetone in an ultrasonic bath, rinsed with isopropanol, and blow-dried with nitrogen.

[1] The waveguide capacitively loads the cavity, leading to a reduced mode frequency compared to $l = \lambda/4$.
[2] Note that the waveguide section length needs to be adjusted for a different post length to achieve the same seam participation.

Next, the aluminum piece is etched to remove scratches and residuals from the machining process. The etchant used is a commercial aluminum etchant, Transene Type A[140]. The screw threads are protected with aluminum, later with polytetrafluoroethylene (PTFE or Teflon) screws. The process is exothermic, and the etch rate highly depends on the temperature of the etchant. The etching process runs for a total of four hours, stabilized at 50° C by a PID-temperature-controlled water system running between the walls of the beaker, with feedback from a temperature probe inserted into the etchant. This process is interrupted after two hours to replenish the etchant, as it gets depleted over time. The aluminum piece must be positioned so that no air pockets form inside. This process concludes with thoroughly rinsing the cavity with deionized (DI) water and blow-drying it with nitrogen[3]. In total, about 200 μm of material is removed from each surface, less for surfaces with reduced etchant flow, such as the qubit tunnel.

4.2 Fluxonium fabrication

The qubit design was provided by Dr. Daria Gusenkova and Prof. Ioan Pop from Karlsruhe Institute of Technology, similar to the design from Ref. [131, 141]. Originally, the fluxonium is inductively coupled to a readout resonator via a piece of granular aluminum. Additionally, the resonator itself is loaded with granular aluminum sections, allowing for a reduction in total footprint while maintaining the readout frequency. This design was adapted for the current dissertation application. To bypass additional coupling elements between the qubit and the 3D cavity, the readout resonator was extended to 7 mm total length, and the coupling pads at the end were changed from rectangles to circles to avoid field concentration at the sharp edges. Since the fluxonium couples to the cavity indirectly via the resonator, the mutual qubit-readout resonator inductance was increased by thinning the mutual grAl connection to attain reasonable qubit-cavity coupling. The design was finalized and fabricated in the Quanten-Nano-Zentrum Tirol (QNZT) cleanroom in the University of Innsbruck. A detailed fabrication recipe can be found in Appendix I.

The fluxonium is fabricated using a single-step electron-beam lithography process with a Raith eLINE Plus 30kV. The design has two length scales: one for the readout resonator and another for the qubit itself. This dictates the need for different writefield sizes: 1 mm and 200 μm, respectively. As the resonator does not fit within a single writefield, stitching errors may occur depending on the machine settings at the time of exposure. To ensure overlap of the structures in adjacent writefields, the zoom factor is increased by 0.2%. For the smaller length scale, the design is proximity-corrected in Beamer GenISys software[142]. The aperture sizes used are 60 μm and 10 μm for the large and small length scales, respectively.

After development[4], the wafer is loaded into an electron beam evaporator Plassys MEB550S, where we evaporate the aluminium and granular aluminium at three different angles to create a Dolan-bridge junction[143] and the superinductance of the fluxonium. The wafer is left in N-Methyl-2-pyrrolidone (NMP) for the lift-off, and later rinsed in isopropanol. After probing, the wafer is diced using a laser cutting system, Coherent ExactCut 430[144].

[3] The cavity is also rinsed between etching rounds.
[4] The development solution is an isopropanol-DI water mixture chilled in a 6° C water bath.

4.2.1 Parameters extraction

Each wafer contains 16 qubit chips, with three different resonator lengths, $l = \{6, 6.5, 7\}$ mm, and four different granular aluminum sections, $w = \{130, 145, 160, 180\}$ nm. Moreover, 108 test structures are placed between the qubit chips to determine the junction and grAl parameters (see Fig. 4.2). The Ambegaokar-Baratoff relation[145] gives an estimate for the critical current of the junction from the normal state resistance R_n

$$I_c(T) = \frac{\pi\Delta(T)}{2eR_n} \tanh\frac{\Delta(T)}{2k_BT} \underset{T\to 0}{=} \frac{\pi\Delta(T)}{2eR_n} \tag{4.1}$$

where $\Delta(T)$ is the temperature-dependent energy gap of the superconductor[5]. From the extracted critical current, we can estimate the Josephson energy of the junction E_J, as pointed out in Sec. 2.2.1.2.

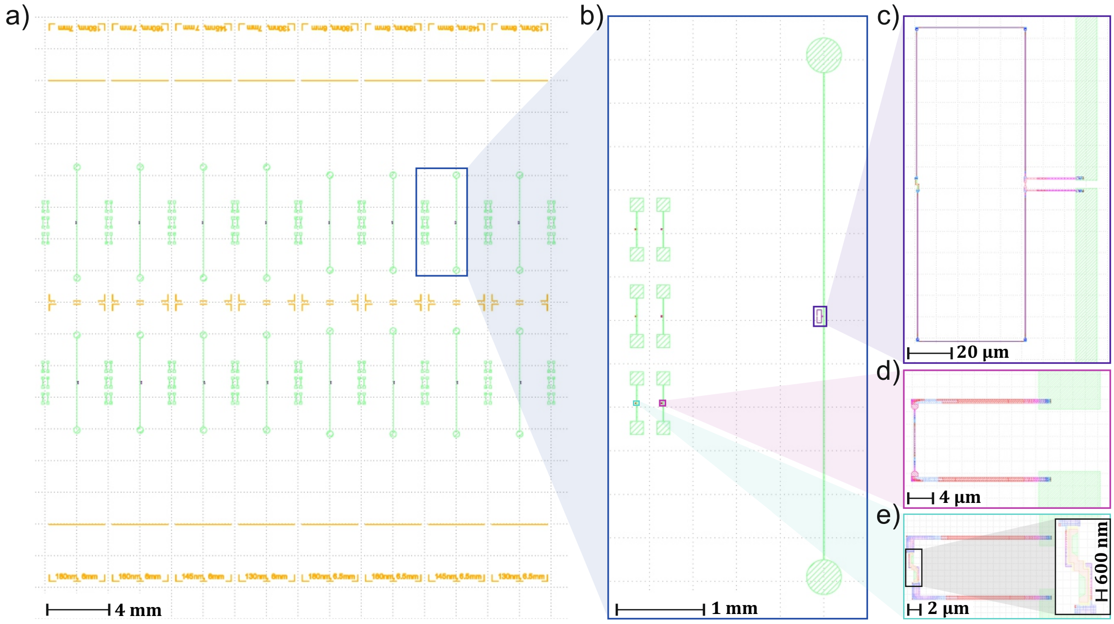

Figure 4.2: GDS design file. Different colors represent different layers: the orange layer includes qubit information, dicing lines, and a line for clamp aligning; the green layer includes all pads and antennas for the structures; the remaining layers are dedicated to the qubits and the test structures. a) Full wafer design, b) fluxonium with its readout resonator and six test structures with their pads and antennas, c) fluxonium with the shared inductance, d) grAl test strip, e) test junction with a zoom-in (inset). All pictures are screenshots from KLayout.

The relevant parameter for the granular aluminum is the film resistivity ρ. For a grAl strips with width W, length L and thickness t, the measured resistance is

$$R = \rho\frac{L}{Wt}. \tag{4.2}$$

[5]For aluminum $\Delta_{\mathrm{Al}}(0) \approx 180\,\mu eV$.

With the extracted resistivity, we can calculate the kinetic sheet inductance[6]

$$L_{\mathrm{kin}}^{\square} = \frac{\hbar}{\Delta_{\mathrm{grAl}}(0)\pi} \frac{\rho}{t}. \tag{4.3}$$

Here, $\Delta_{\mathrm{grAl}} = 1.76 k_B T_c$, where T_c is the critical temperature of the granular aluminum[7]. From Eq. 4.3, we can estimate the mutual inductance between the fluxonium and the readout resonator L_s, as well as the superinductance of the qubit L_Q. The inductive energy E_L of the qubit depends on their sum.

Lastly, in contrast to the transmon, here the charging energy is dominated by the junction. In its essence, a junction is a plate capacitor with dielectric constant $\varepsilon_{r,\mathrm{AlOx}} \approx 10$[146, 147]. Therefore, we can calculate its capacitance $C_J = \varepsilon_0 \varepsilon_{r,\mathrm{AlOx}} \frac{A_J}{t}$ with the area A_J being determined from SEM images and assuming the oxide barrier has a thickness $t \approx 2\,\mathrm{nm}$[148], which gives approximately $3\,\mathrm{fF}$, or equivalently an estimated charge energy $E_C^{\mathrm{est}}/h \approx 6\,\mathrm{GHz}$.

A qubit, similar to the one used in this dissertation, is depicted in Fig. 4.3. The junction size, measured with an SEM, is approximately $150 \times 350\,\mathrm{nm}$. The grAl width on the qubit superinductance is in the range of $200 - 220\,\mathrm{nm}$. The shared inductance section width varies depending on the qubit. From the test structures, the estimated energies are $E_J^{\mathrm{est}}/h \approx 14.8\,\mathrm{GHz}$, $E_L^{\mathrm{est}}/h \approx 0.65\,\mathrm{GHz}$. When cooled down and measured, however, these values were always different, i.e., $E_J^{\mathrm{meas}}/h = 10.8\,\mathrm{GHz}$ and $E_L^{\mathrm{meas}}/h = 1.014\,\mathrm{GHz}$, partially explainable by the $10 - 20\%$ difference in size between the test structures and the qubit. The measured capacitive energy was always about a factor of two smaller than the predictions for the measured junction size, i.e., $E_C^{\mathrm{meas}}/h = 3.5\,\mathrm{GHz}$, which raises the suspicion of a different dielectric constant or thickness of the oxide layer.

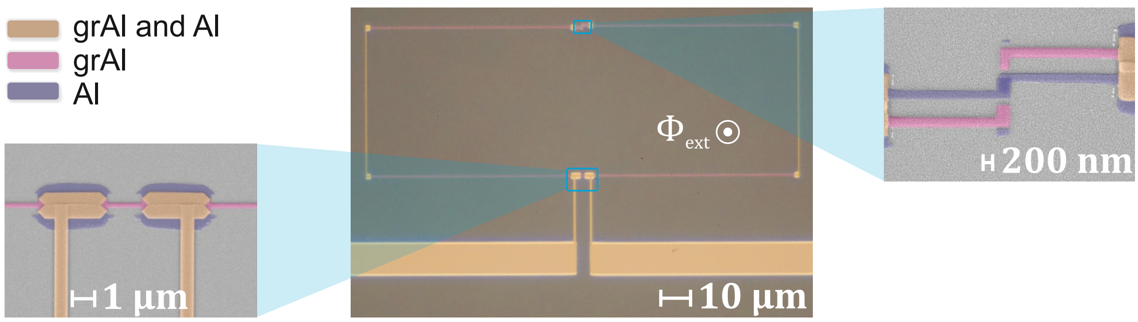

Figure 4.3: False-colored images of a fluxonium. In the center, the full fluxonium is visible with a small part of the readout resonator. The left inset depicts the granular aluminum section that serves as a mutual inductance L_s. The inset on the right shows the junction; the additional leads are formed due to the angle of granular aluminum evaporation. Insets: SEM images; center: optical microscope image.

[6] The sheet inductance is the inductance of a square sheet of metal. The total inductance is $L = NL^{\square}$, where $N = L/W$ is the number of squares.

[7] For the estimation, we use $T_c = 2\,\mathrm{K}$[84], but this value generally varies with the film resistivity.

4.3 Evolution of the magnetic flux hose

Biasing flux-tunable qubits presents significant technical challenges. In 2D architectures, this is typically achieved using on-chip flux bias lines, which function like coplanar waveguides shorted to ground near the qubit. However, these lines can capacitively couple to the qubit, introducing unwanted losses. In copper 3D enclosures, such as waveguides or cavities, tunability can be realized by winding a coil around the device. By contrast, in superconducting enclosures, the Meissner effect[6] shields qubits from both stray and control magnetic fields, preventing direct flux biasing. To overcome this limitation, several approaches have been developed: introducing a coil with an orifice to guide the field, and using on-chip pickup loops to direct the field toward the qubit[149], or integrating flux bias lines directly into 3D enclosures[150]. Both approaches, however, present a significant engineering challenge due to the additional elements involved. An alternative approach is leading the magnetic field in through a so-called magnetic flux hose[151].

The first generation consists of alternating layers of $\mu-$metal and aluminum alloy shells with a $\mu-$metal rod in the center. The magnetic sheets guarantee the transport of the magnetic field created by a coil at the end of the structure, while the superconductor diminishes the radial component of the field. The transport is maximized for a short magnetic hose with many layers[152]. The superconducting shells protrude 1 mm further towards the qubit, compared to the $\mu-$metal, maximizing the field transport while avoiding the qubit from coupling to the lossy magnetic material. The cylindrical structure has a slit to ensure no Eddy currents are created in the $\mu-$metal, and prevents the flux quantization in the superconducting shells, so the field can be smoothly ramped. Moreover, the majority of the magnetic field lines close through this slit, with the leftover closing through the $\mu-$metal layers.

The coil body is fabricated from a Stycast 1266 rod, machined on a lathe. A superconducting wire is wound around the body and guided through a groove to preserve its insulation, while a thin Stycast 1266 coating provides additional protection against scratching. The coil has a diameter of approximately 1.8 mm, smaller than the outer superconducting shell in which it is inserted. The aluminum shells are wire-eroded into the required cylindrical shape with an initial thickness of 200 μm, then further etched down to 100 μm. The $\mu-$metal sheets are also wire-eroded and subsequently hand-formed into the desired geometry. The entire hose assembly, including the coil, is held together with a copper clamp that must be well thermalized (e.g., with a copper braid). For fast flux pulsing, an SMA connector should be soldered as close as possible to the assembly to minimize parasitic effects, which would lead to slow operation due to stray inductance. Since the connector contains normal metal that can heat due to resistive losses, its clamp should be kept separate from that of the hose assembly. A technical drawing and a photograph of the completed assembly are shown in Fig. 4.4a and b.

To overcome the laborious assembly process of the first-generation hoses, a new design was introduced, consisting of a single cylindrical piece of superconductor. A coil is wound around the cylinder with a slit added to prevent flux quantization (see Fig. 4.4c). When the coil applies a magnetic field, the superconductor generates screening currents. These currents, in turn, create a field suitable for tuning a qubit[155]. Beyond simplifying fabrica-

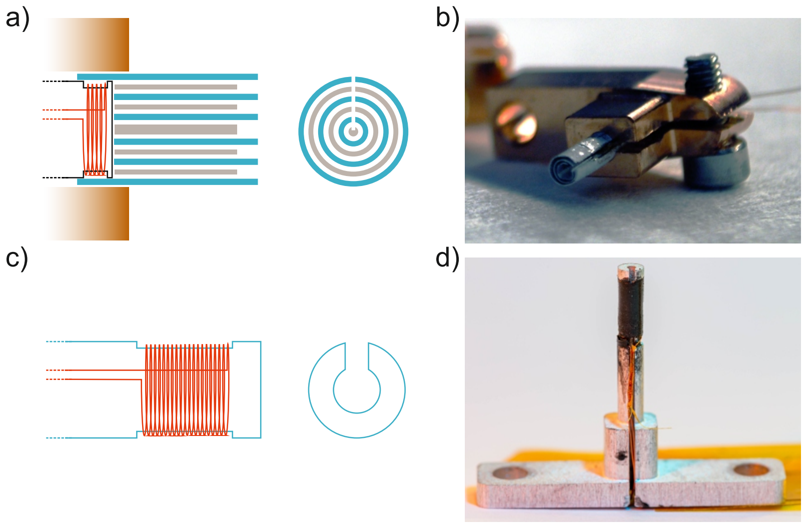

Figure 4.4: Magnetic hose generations. The first and second rows correspond to the first and second generations, respectively. The first column contains the schematic of the axial and radial cross sections, while the second column shows a picture of an assembled magnetic hose. Cyan represents aluminum, gray - $\mu-$metal, red - superconducting coil, and orange - copper. The picture in b) shows the actual device used in this dissertation, while the one in d) is taken from Ref. [153], showing a device manufactured for a different project[154]. Both hoses have a diameter of about 2 mm.

tion, this design also allows the hose to be positioned closer to the qubit, since it eliminates the need for lossy magnetic materials.

The device in Fig. 4.4d was the first proof of principle as a part of a master's project[154]. The flux hose was machined with EDM, which allows for a 200 μm slit width. It then followed the same etching procedure as the cavities, but the duration was reduced to one hour. The coil is wound directly on the superconductor and then protected by Stycast 2850. In the case of fast flux applications, there is no limit to the speed besides the inductance of the coil itself. However, while a transmon qubit has been tuned with this design, the speed is yet to be tested.

The second-generation hose supports a microwave mode because, when mounted, it effectively forms a $\lambda/4$ post cavity.[8] In principle, the different field symmetries between the qubit and hose modes should prevent coupling; however, small imperfections in chip alignment lead to residual interactions. The microwave properties of this mode were previously investigated in a master's thesis[156] that I supervised. That study showed that, for the design shown in Fig. 4.4d, the mode quality factor is primarily limited by dielectric losses originating from the coil insulation and the Stycast. To mitigate these losses, the geometry was subsequently optimized to reduce the electromagnetic participation of these lossy ma-

[8]Note that also the first-generation hose forms a microwave mode, but its shape is not an ideal cylinder, which makes this mode harder to predict. Moreover, it might get shorted by an accidental touch between the hose and the walls of the enclosure.

terials. With the improvements made in Ref. [156], the mode can potentially be employed both for flux biasing and for qubit readout[157].

4.4 Wiring

A schematic of the cryogenic and room temperature setup is shown in Fig. 4.5. The measurements were conducted in an Oxford Triton DU7-200 cryofree dilution refrigerator system. The sample is mounted on the base plate of the fridge, thermalized to 20 mK. Thermalization to a temperature much lower than the critical temperature of the superconductor is necessary to reduce the thermal population in the microwave modes, allowing them to be initialized in their ground state.

4.4.1 Room temperature wiring

A Quantum Machine Operator X+ (OPX+) was used together with a Quantum Machine Octave to generate and digitize the microwave signals. The OPX+ has arbitrary waveform generators that produce the pulses used in the experiment, which are then mixed up to the relevant frequencies and amplified in the Octave. For readout, the Octave down-mixes the resonator and cavity pulses, which are digitized by the OPX+. In the low-frequency range $f \in [1, 6.8]$ GHz, the fluxonium pulses were up-mixed via a double-super-heterodyne[158] (DSH) scheme that employs two local oscillators and two single-sideband mixers to produce a wide range of frequencies without the need for calibration. For higher qubit frequencies, the pulses are directly mixed by the Octave. Some details on the IQ and DSH mixing are outlined in Appendix D. The pulse generation setup includes other amplifiers, filters, fast microwave switches, and DC blocks to produce a high amplitude of the required microwave signal, reduce leakage of unwanted signals, and prevent the formation of ground loops.

4.4.2 Cryogenic wiring

The radio frequency (RF) input coaxial lines were attenuated by 20 dB at the 4K plate and 10 dB at the still (1K) plate. On the base plate, the input signal for the qubit line was attenuated further by 30 dB (a 20 dB attenuator and a Quantum Microwave 10 dB thermalized attenuator) and then filtered with a Mini-Circuits DC-14 GHz low-pass filter. The experiment was in a reflection configuration, and a Low Noise Factory single junction 4–12 GHz circulator was used. The signal was finally routed through a home-built infrared filter. The input signal for the cavity has the same configuration up to the 100 mK stage. A low-pass K&L filter filters the cavity signal, which is then attenuated by a Quantum Microwave 20 dB thermalized directional coupler and 10 dB attenuator from the same brand. A narrow-band Mini-Circuits 4–4.8 GHz band-pass filter and an infrared filter are placed just outside the shield. A Quinstar 4–8 GHz circulator allows for direct reflection measurements of the cavity.

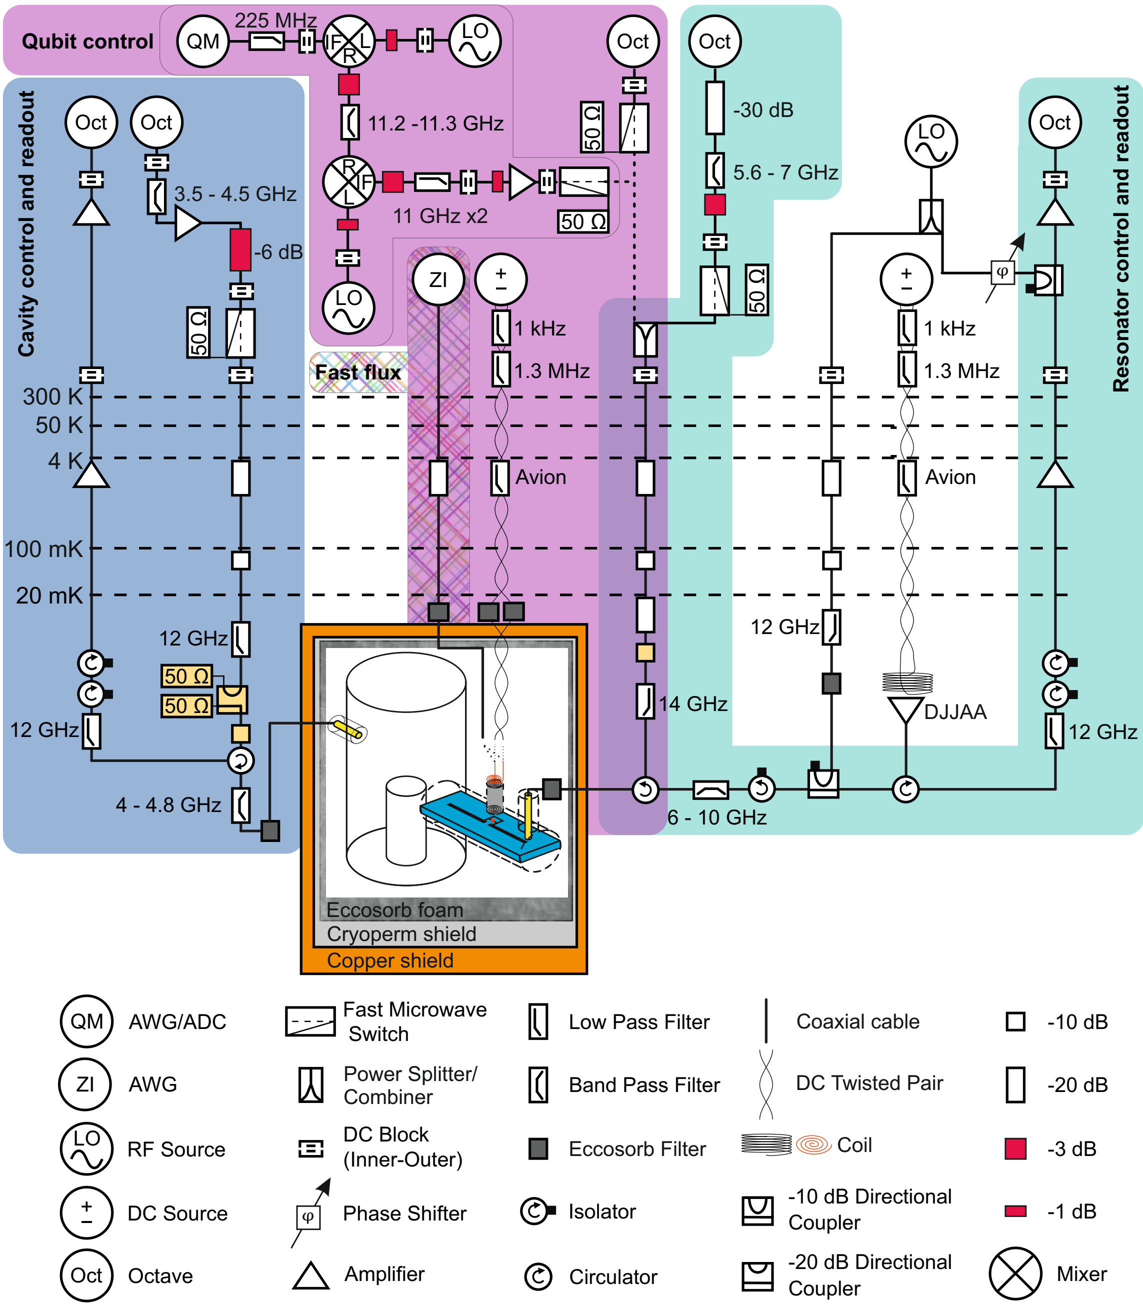

Figure 4.5: Wiring diagram of the experiment. The lined part of the qubit control corresponds to the low-frequency microwave setup. The dashed lines refer to the different plates of the cryostat. Depending on the measurement, either the low-, or high-frequency branch, or both branches are combined with the readout tone. For the hose coil, either DC or fast flux setup is used. The elements highlighted in yellow are specially thermalized to the baseplate. The components related only to the parametric amplifier are not highlighted.

The output signal from the resonator was filtered by a 6–10 GHz Microtronics band-pass filter and then routed through an isolator, a Krytar 10 dB directional coupler, and a Low

Noise Factory circulator before reaching a parametric DJJAA (a Dimer Josephson Junction Array Amplifier[159, 160]). The 10 dB directional coupler is used to couple the amplifier pump line. After the parametric amplifier, the output signal goes through two Quinstar isolators before being amplified at the 4K stage by a high-electron-mobility transistor (HEMT) and again at room temperature. To avoid amplifier saturation, an additional 10 dB directional coupler is added before the room-temperature amplifier to cancel the DJJAA pump tone from the fridge.

The flux biasing is achieved by leading DC signals through twisted pairs and commercial low-pass filters. Two home-built high-attenuation filters[9] further refine the qubit bias. Details on their fabrication can be found in Appendix J. The flux biasing in Chapter 5 was done with a current source, while for Chapter 7, the HDAWG of Zurich Instruments provided the bias offset and generated flux pulses once triggered by the OPX+. The signal in the latter case was carried through a high-attenuation filter to the sample via a coaxial line.

4.4.2.1 Shielding

The sample was thermalized to the base plate via copper clamps. The experiment was shielded by a μ-metal shield and enclosed by a copper can, covered with eccosorb paint on the inside. In the inner can, the sample was surrounded by eccosorb foam.

The concerns of cross-talk between two different setups in the same shield gave grounds for a new shielding design, which hosts a single setup. The requirements include:

- The main body with good thermalization to the base plate;
- Well thermalized superconducting layer with a high critical temperature, which expels the magnetic field due to the Meissner effect[6];
- Well thermalized cryogenic μ-metal shield to attenuate any residual stray magnetic fields.

According to these requirements, the main body was constructed from copper and electroplated with tin. The choice of tin was motivated by its critical temperature being higher than that of both aluminum and granular aluminum, which would prevent flux trapping on the sample. Electroplating guarantees good thermalization of the superconducting layer. Before plating the can, small samples of copper were electroplated with tin and cycled several times to cryogenic temperatures to ensure no tin pest appears.[10] The SolidWorks drawings of the new shielding are shown in Fig. 4.6.

[9]The filters for the DC and fast flux biasing are made out of Eccosorb CR-124, while for the microwave filters, Eccosorb CR-112 was used.
[10]The electroplating was done externally by Ulrich Bauer from CCI Eurolam GmbH, Dreieich, Germany.

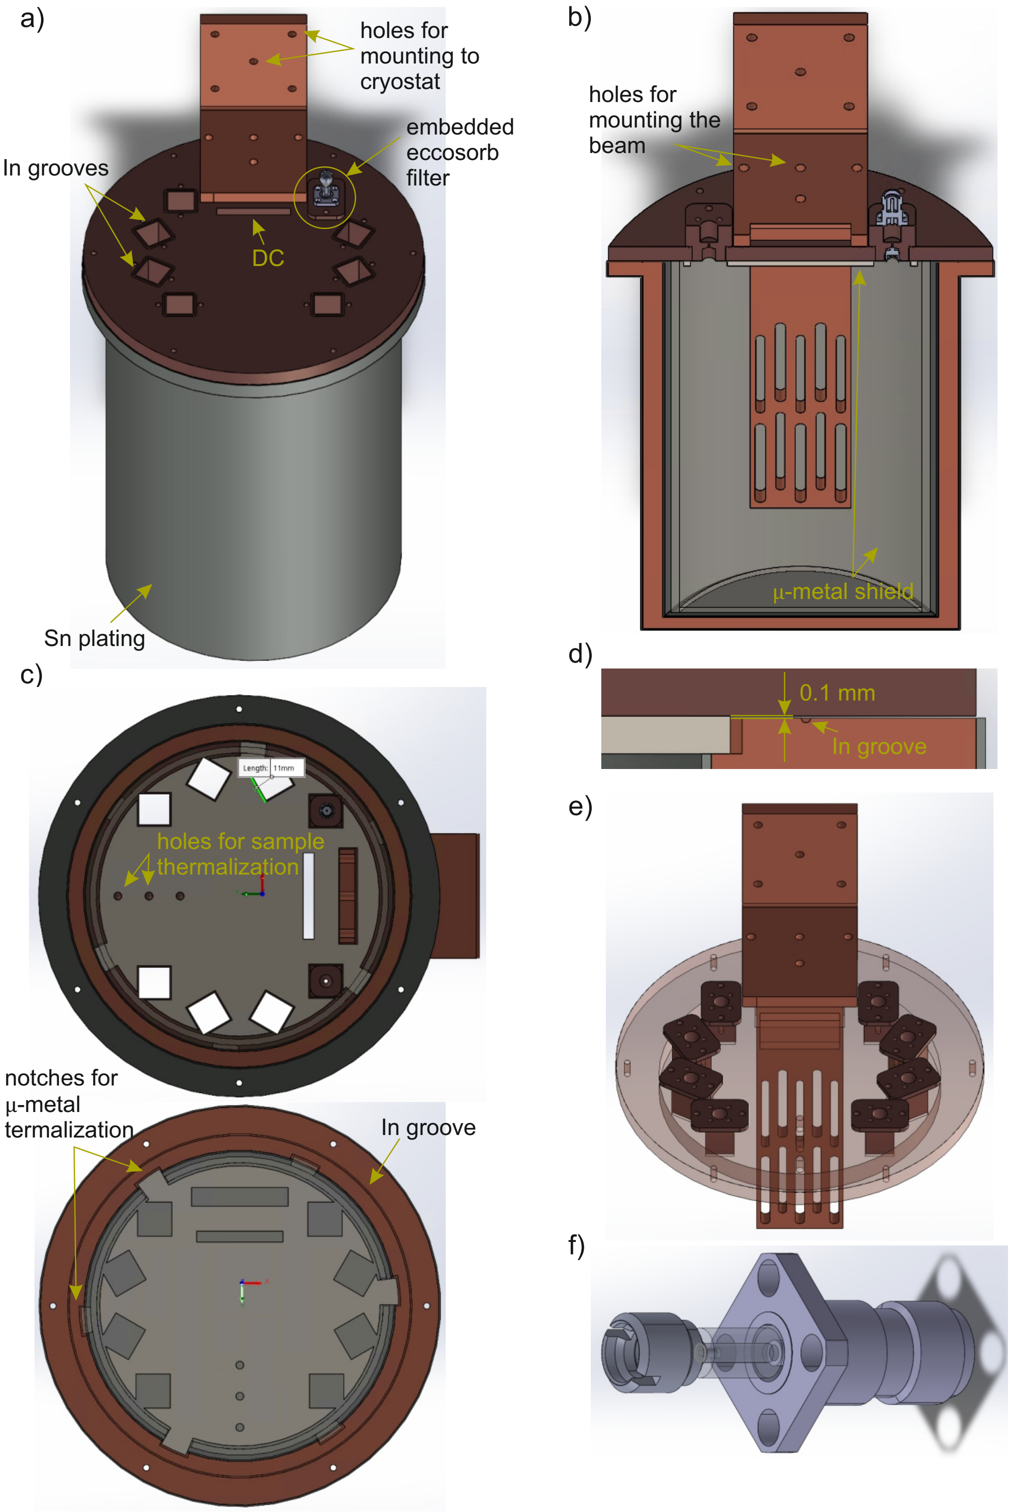

Figure 4.6: SolidWorks drawings of the new shielding can. a) Full assembly, the gray represents the tin plating; b) Cut assembly so the interior is visible; the width of the beam is 30 mm. c) Upper panel: view from the bottom of the assembly cut in the middle; Lower panel: view from the top of the assembly containing only the copper can, the $\mu-$metal can, and the $\mu-$metal lid; d) Side zoom-in where the copper can and lid touch the $\mu-$metal parts. The distance of 0.1 mm between the top of the $\mu-$metal parts and the copper can, and the indium put in the groove would guarantee good thermalization. e) Lid with eccosorb filter bodies; f) Integrated eccosorb filter assembly: SMP connector is on the left, SMA - on the right.

To avoid photon-assisted quasiparticle tunneling through the junctions[161], the interior of the can is protected from stray photons by cryogenically sealing all distinct parts with indium. The μ-metal shield, located inside the copper, is comprised of two parts as well. Those are directly thermalized to the copper can by making them extrude 0.1 mm above the copper can (see Fig. 4.6d). In this way, adding indium wire of diameter 1 mm in a triangular groove of depth 0.5 mm would ensure both good thermalization and light-tightness.

Furthermore, the can design needs to be compact and ideally stackable to fit several setups on the base plate. For this reason, the SMA connectors on the inside are replaced with SMP connectors, which have a smaller footprint. Finally, according to the requirements for light-tightness, the IR eccosorb filters are integrated into the lid[11]. Those filters have separate bodies, which makes them easily exchangeable in case they break.

The DC wiring was planned to be integrated via six pairs of 2.54 mm connectors on each side of the can, connected by a flexible PCB. The gap between the lid of the can and the PCB would be filled with Eccosorb.

DC measurements A silicon wafer was first sputtered with a 300 nm layer of gold to improve conductivity and subsequently electroplated with tin. A long and narrow structure ($l \sim$cm and $w \sim \mu$m) was patterned using a laser writer and a negative photoresist. However, both wet and dry etching attempts proved unsuccessful due to the substantial roughness and thickness of the tin layer ($t_{\mathrm{Sn}} = 5\,\mu$m), as well as the low melting temperature of tin, which complicated standard processing techniques.[12]

As an alternative approach, a narrow strip of the wafer ($l \sim$cm and $w \sim$mm) was mechanically cleaved and glued onto a thin copper substrate using GE varnish to facilitate mounting on the cryostat base plate (see Fig. 4.7a). Four superconducting wires were attached to the tin sample using a low-temperature Surface-Mount-Device (SMD) solder paste, followed by annealing in an oven to ensure reliable electrical contact. The sample was then mounted on the base plate of the cryostat, and its critical temperature and critical magnetic field were investigated using a four-terminal sensing configuration[162].

In the critical temperature measurement, the resistance was monitored during cooldown using a Stanford Research Systems SR830 lock-in amplifier. A resistor $R_0 = 1\,\mathrm{k}\Omega$ was placed before the cryostat to transform the voltage source into a current source. A sharp drop in resistance was observed at $T \approx 9\,$K, which is inconsistent with the expected critical temperature of tin thin films, $T_{c,\mathrm{Sn}} \approx 4\,$K[163]. This discrepancy is likely attributable to the superconducting leads, which contain niobium and therefore exhibit a higher critical temperature. A plausible explanation is that the leads formed an unintended electrical shortcut, leading to imperfect subtraction of their contribution to the measured resistance. Additional complications were encountered: the resistance did not approach zero in the superconducting state, and the measured signal exhibited a pronounced frequency depen-

[11] More information on their manufacturing, impedance matching, and characteristics can be found in Appendix J.
[12] The etching tests were executed by Christian Schlager (dry) and Julian Daser (wet).

dence, suggesting the presence of capacitive or inductive loading, likely originating from the DC filtering in the cryostat.

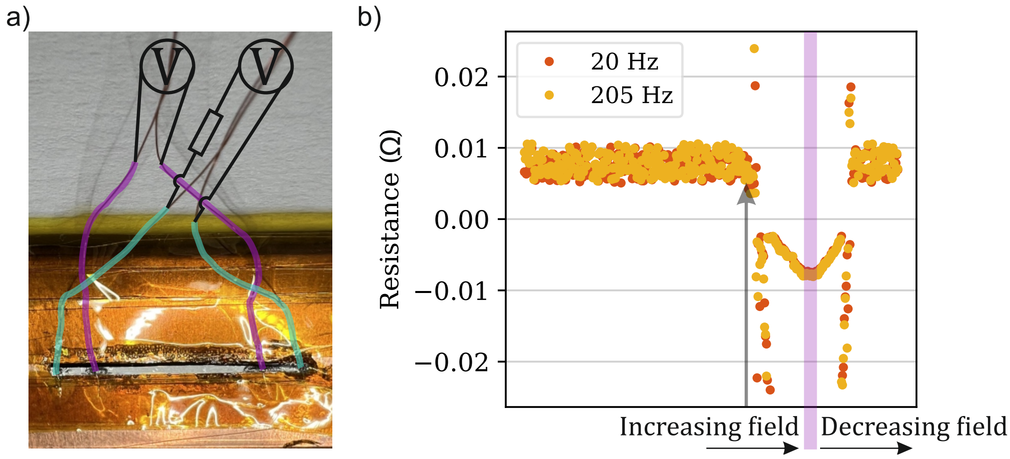

Figure 4.7: DC measurements. a) Setup. The Sn sample is mounted on a copper piece, and four wires are soldered to it. The cyan color highlights the current path, while the violet indicates where the voltage drop is being measured. b) Critical field measurement. The gray arrow indicates where the ramp to 200 mT magnetic field in the magnet started. The magenta region highlights a constant field after the last ramp-up. Afterwards, the field was ramped down. The negative values are an error coming from the lock-in amplifier when a resistance overload is reached. The x-axis is nonuniform in the magnetic field and uniform in time. These and the critical temperature measurements of the tin samples were executed together with Christian Schlager.

Rather than modifying the cryostat wiring, a new sample was prepared and transferred to a different cryostat (Qinu L Pro), which is equipped with an electromagnet but lacks DC filtering. The sample was positioned beneath the electromagnet, while another experiment occupied the primary sample space. The magnetic field was ramped to various values while continuously monitoring the resistance. The resulting data are shown in Fig. 4.7b. The resistance sharply changes at a field indicated with a gray arrow. At this timestamp was the onset of a field ramp to 200 mT in the electromagnet, corresponding to an estimated field of approximately 50 mT at the sample location, after accounting for its radial and axial displacement from the magnet center. This estimated critical magnetic field is in good agreement with previously reported values[164].

CHAPTER 5

Characterization

In this chapter, I describe the measurement techniques used to characterize the Hamiltonian of the experiment. The measurements used for extracting the Hamiltonian parameters can be split into two categories: frequency- and time-domain measurements. In the former, one or numerous tones are applied, and their frequency is swept to observe the system response. These tones can be either used for probing or pumping the system. The latter records the system's temporal evolution.

The first two sections present the measurements used to characterize the readout resonator, the qubit, and the cavity. Sections 5.3 and 5.4 focus on their interactions. Some of the findings of this chapter can be found in Ref. [165].

5.1 On-chip resonator and fluxonium

5.1.1 Spectroscopy

For the initial characterization, a Vector Network Analyzer (VNA) was used to measure the scattering parameters S_{11} and S_{22}[62], corresponding to reflections from the readout and cavity ports, respectively. A weak probe tone is applied to each port and is mostly reflected, except when the probe frequency matches a resonance. At resonance, the reflected signal exhibits a reduction in amplitude and a distinct phase shift, indicating the presence of a mode at that frequency. We repeat this procedure for different flux bias points to acquire the flux maps of the cavity and the readout resonator. Each trace is then fitted to a Lorentzian curve to extract the resonant frequencies, which are plotted in Fig. 5.1 left and right. The tunability of both elements is inherited from the fluxonium.

The qubit frequency is measured with two-tone spectroscopy. We monitor the resonator resonance while sweeping the frequency of a second tone with a signal generator. At the qubit frequency, the resonator shifts, and our readout signal changes. We repeat the procedure for different applied magnetic fields and fit each data to a Lorentzian curve to extract the qubit resonance frequency. The higher transition f_{12} is extracted in the same way around $\Phi_{\mathrm{ext}}/\Phi_0 = -0.5$, where the non-zero excited state population of the qubit allows direct excitation to the next higher level. For the f_{12} data points around the high sweet spot at $\Phi_{\mathrm{ext}}/\Phi_0 = 0$, the qubit is first excited by applying a π-pulse to

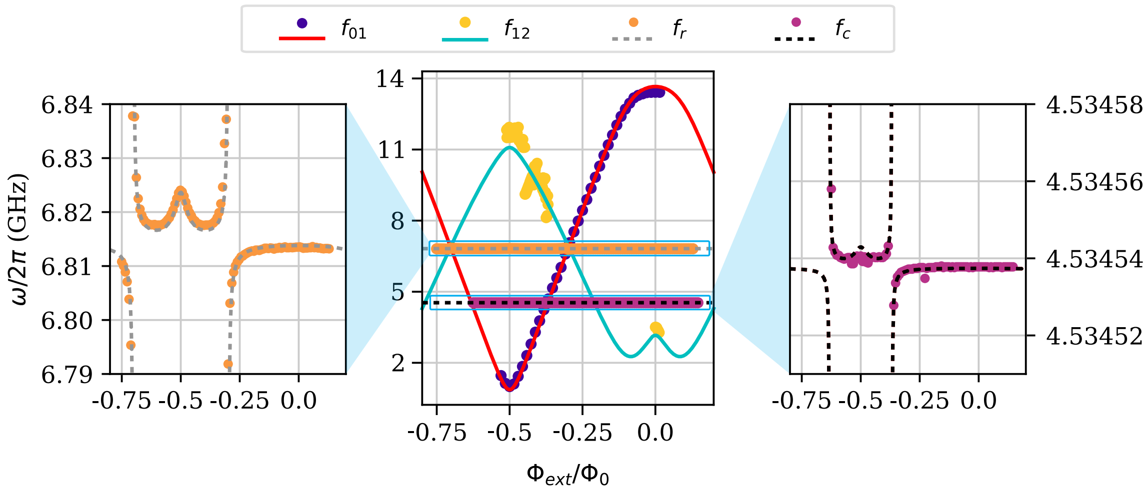

Figure 5.1: System spectrum. (center) The qubit frequency is measured as a function of flux bias. By fitting this spectrum, we obtain the qubit energies $E_C/h = 3.5\,\mathrm{GHz}$, $E_L/h = 1.014\,\mathrm{GHz}$, and $E_J/h = 10.8\,\mathrm{GHz}$. The measured data for the fundamental transition are plotted in purple, together with its fit in red. The higher transition f_{12} is in yellow, with the *scQubits* simulation for it in cyan. The readout resonator (left) and high-Q cavity spectra (right) are measured as a function of flux, respectively. The qubit is tuned through both the readout resonator and high-Q cavity, resulting in avoided crossings with coupling strengths $g_{QR}/2\pi \approx 25.2\,\mathrm{MHz}$ and $g_{QC}/2\pi \approx 1.5\,\mathrm{MHz}$. The dashed lines correspond to their fits.

the fundamental transition, followed by a time-domain two-tone spectroscopy around the frequencies, where we expect the transition to appear.

We fit the f_{01} data with scQubits (Fig. 5.1 center) and obtain the qubit energies $E_C/h = 3.5\,\mathrm{GHz}$, $E_L/h = 1.014\,\mathrm{GHz}$, and $E_J/h = 10.8\,\mathrm{GHz}$. The Hamiltonian is then constructed by including the harmonic oscillators and the interactions, as described in Sec. 3.3. The numerical results of the diagonalization are manually fitted to the cavity and resonator spectroscopy (Fig. 5.1, right and left, respectively) as a function of the external magnetic field. The avoided crossings between the resonator and the qubit, and between the cavity and the qubit, give us the interaction strengths $g_{QR}/2\pi \approx 25.2\,\mathrm{MHz}$ with an accuracy of about 10% and $g_{RC}/2\pi \approx 8$ MHz with 15%. The splitting of $2g_{QC}/2\pi \approx 3$ MHz we extract from the avoided crossing in Fig. 5.1 (right) is a measure of the effective coupling rate between the qubit and the cavity, which is mediated by the resonator and determined by both g_{QR} and g_{RC}.

Near the high sweet spot, there is a discrepancy between the model and the measured data. A possible explanation is the cavity's higher modes, which we expect around 13.37 GHz, 13.73 GHz, and 13.78 GHz, according to HFSS simulations. Each of those modes will couple to the qubit via the readout resonator. In addition, the symmetry of the second additional mode will allow it to leak into the tunnel of the qubit and couple directly with the fluxonium. Owing to all the additional interactions at those frequencies, the model fails to fully capture the qubit behavior at the high frequencies.

The higher cavity modes cannot be neglected and will influence the bare frequencies of the harmonic oscillators, needed for the fit. They will also push the qubit resonance to lower frequencies than anticipated. In Fig. 5.2, we show the influence of higher cavity modes on the qubit fundamental frequency. Three higher modes of the cavity are included with direct coupling to the resonator; the highest-frequency mode is also coupled directly to the qubit, as the simulations show that this mode leaks further into the qubit tunnel. The coupling strengths were based on expectation values from the scaling with $\sqrt{\omega}$ and then further increased to amplify their effect.[1] As the frequencies of those modes were never measured[2] but extracted from simulations, the actual cavity modes could coincide with the qubit levels, which would cause a further shift in the transition f_{01}. Such a coincidence could explain why, around $\Phi_{\mathrm{ext}}/\Phi_0 = 0$, the measured qubit frequencies overlaps with the cavity mode at f_c^1.

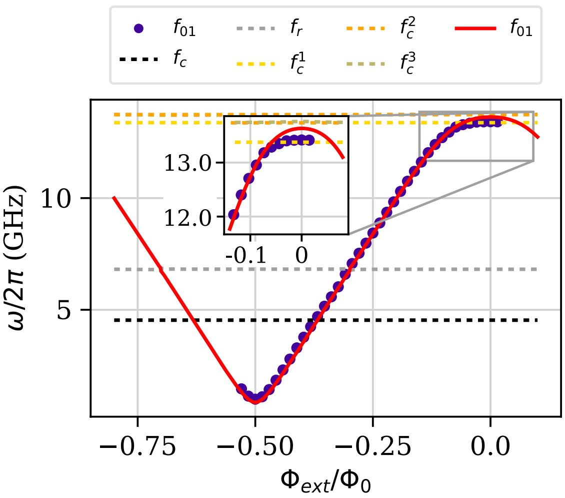

Figure 5.2: Fluxonium spectroscopy with simulation curves including higher cavity modes. In the model, apart from the fundamental modes of the cavity f_c and the resonator f_r, three higher cavity modes are included: $f_c^1 = 13.375\,\mathrm{GHz}$, $f_c^2 = 13.735\,\mathrm{GHz}$, and $f_c^3 = 13.734\,\mathrm{GHz}$. They are capacitively coupled to the resonator with rates $g_{RC}^{1,2,3}/2\pi = 16\,\mathrm{MHz}$. Moreover, the highest frequency mode at f_c^3 is directly coupled to the qubit, as it leaks into the tunnel of the qubit, while the modes f_c^1 and f_c^2 do not. This mode does not significantly affect the f_{01} transition frequency, even with a direct coupling strength of $g_{QC}^3/2\pi = 16\,\mathrm{MHz}$. Inset: zoom-in of the region around $\Phi_{\mathrm{ext}}/\Phi_0 = 0$.

5.1.2 Temperature measurements

The effective temperature of a mode in the circuit is typically higher than the physical temperature of the dilution refrigerator's base plate. This discrepancy arises from imperfect thermalization and residual noise coupling through the control and measurement lines. Thermal photons originating from higher-temperature stages leak into the device despite the use of attenuators and filters, as these components cannot completely suppress

[1]This scaling becomes evident from the derivations in Appendix E.
[2]They are out of the $4-8\,\mathrm{GHz}$ frequency range of the HEMT.

noise transmission. Additional contributions stem from back-action noise of amplifiers, especially HEMTs at 4 K, as well as insufficient filtering of infrared and high-frequency radiation, which can generate quasiparticles[161] or directly excite the modes[166]. Additional noise from room-temperature electronics further populates the qubit and cavity with thermal photons. Consequently, even at base temperatures of 10-20 mK, the effective mode temperature is typically in the range of 40-100 mK.

To characterize the mode temperature, one must measure the population distribution of the mode. This can be achieved using either spectroscopy or selective Rabi oscillations of the higher transition[167]. Those methods are suited to different temperature regimes. For low temperatures, the qubit mode temperature can be extracted by measuring the Rabi oscillations between the energy levels $|0\rangle_q - |1\rangle_q$ and $|1\rangle_q - |2\rangle_q$ and comparing their amplitudes.

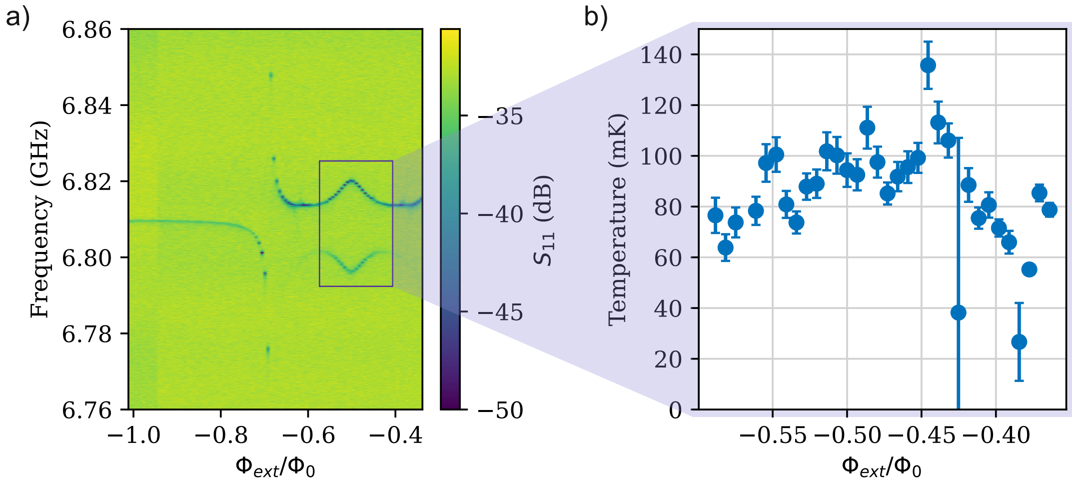

Figure 5.3: Qubit temperature extracted from readout resonator spectroscopy. a) The flux map of the resonator exhibits a double resonance around $\Phi_{\text{ext}} = -0.5\Phi_0$, due to the non-zero excited state qubit population. b) In the region around the low sweetspot, the qubit temperature can be directly extracted, using the qubit frequencies from Fig. 5.1 and Eq. 5.1.

At higher mode temperatures, the excited-state population becomes directly observable in the resonator spectroscopy, as seen on the flux map in Fig. 5.3a. The two resonances, corresponding to the qubit being in $|0\rangle_q$ and $|1\rangle_q$, are separated by the dispersive shift, and the ratio of their amplitudes provides a measure of the qubit temperature as follows

$$\frac{A_{|1\rangle_q}}{A_{|0\rangle_q}} = e^{-\frac{hf_{01}}{k_bT}} \tag{5.1}$$

where $A_{|0/1\rangle_q}$ is the amplitude of the readout resonance for the qubit in ground/excited state. The qubit temperatures, as extracted from the resonator flux map, are summarized in Fig. 5.3b.

The procedure for extracting the resonator temperature is similar, but $A_{|0/1\rangle_r}$ would now be the amplitude of the qubit peak corresponding to 0/1 photons in the resonator, and the frequency will be replaced with f_r. The extracted temperature of the resonator at the high sweet spot is $112 \pm 29\,\mathrm{mK}$, extracted from the qubit spectroscopy in Fig. 5.4. Nevertheless, the extracted temperature may exceed the true mode temperature, as the measurement procedure involves opening the switch on the readout input before the qubit spectroscopy. This process can lead to an increased photon population within the resonator due to excitation from the amplified noise[139] and the local oscillator leakage while the switch is open[3], which can artificially elevate the inferred thermal occupation relative to its equilibrium value.

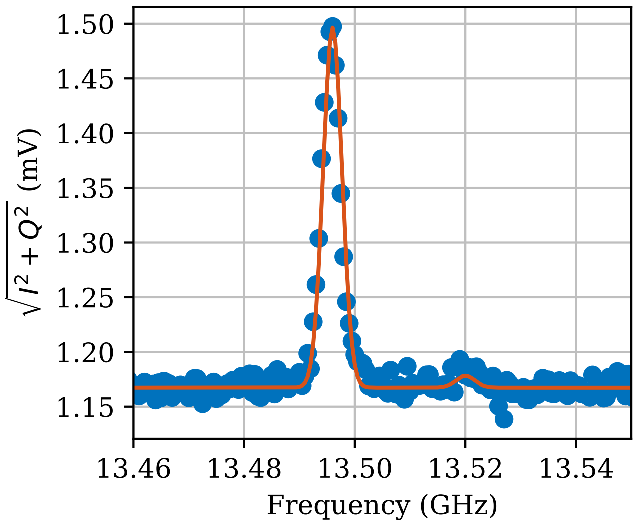

Figure 5.4: Qubit spectroscopy at the high sweet spot. From the amplitude ratio of the two peaks in the spectroscopy, one can extract the resonator temperature. Due to the positive dispersive shift, the qubit resonance corresponding to one photon in the readout resonator is at a higher frequency.

It is worth noting that the mode temperature changes over time. Over the span of a six-month cooldown, two different mode temperatures have been measured for the qubit, $T_Q = \{112, 45\}\,\mathrm{mK}$. The lowest qubit temperature was observed after the helium refill of the cryostat, which reduced the base plate temperature from $\gtrsim 40\,\mathrm{mK}$ to $< 26\,\mathrm{mK}$[4].

5.1.3 Fluxonium lifetime

For this system to be suitable as a bosonic qubit, further understanding of the fluxonium losses is required. We measure the lifetime of the fluxonium qubit as a function of the flux bias and estimate the T_1^Q limit from the dominant decay channels[168–170]: dielectric loss, inductive loss, Purcell limit[171], and quasiparticle loss[172].

[3]The LOs are constantly on during the measurements.

[4]The exact temperature is unknown, because even the sensor needed to be recalibrated as this temperature was out of its scope.

5.1.3.1 Loss model

In this section, we discuss in detail the limitations on the fluxonium lifetime via the various qubit decay channels[168, 169, 173]. The decay rate of the qubit is given by Fermi's Golden Rule[170] of the type

$$\Gamma_{01} = \frac{1}{\hbar^2}|\langle 0|\mathscr{O}|1\rangle|^2[S(\omega_{01}) + S(-\omega_{01})]. \tag{5.2}$$

The operator $\mathscr{O}$ is the qubit operator which couples to the lossy bath, and $S(\omega_{01})$ is the noise spectral density at the qubit frequency. The first term corresponds to emission into the bath, and the second to absorption.

We consider the most common loss mechanisms. For dielectric loss, the flux operator $\mathscr{O} = \phi_0\varphi$ couples to the voltage of the bath with density

$$S_{VV}(\omega_{01}) + S_{VV}(-\omega_{01}) = \frac{2\hbar\omega_{01}^2 C_J}{Q_{\text{diel}}(\omega)}\coth\left(\frac{\hbar|\omega|}{2k_BT}\right) \tag{5.3}$$

where $C_J = e^2/2E_C$ is the junction capacitance, and T is the qubit mode temperature. The quality factor of the dielectric loss is weakly frequency-dependent[168, 169, 174]

$$Q_{\text{diel}}(\omega) = Q_{\text{diel}}\left(\frac{2\pi\, 6\,\text{GHz}}{|\omega|}\right)^{0.7}. \tag{5.4}$$

Here, the Q_{diel} is the nominal quality factor quoted for 6 GHz.

In the case of inductive losses, the flux operator $\mathscr{O} = \phi_0\varphi$ couples to the current of the bath with density[168, 169]

$$S_{II}(\omega_{01}) + S_{II}(-\omega_{01}) = \frac{\hbar}{LQ_{\text{ind}}}\coth\left(\frac{\hbar|\omega|}{2k_BT}\right) \tag{5.5}$$

where $L = \phi_0^2/E_L$ is the fluxonium superinductance and Q_{ind} is the inductive quality factor.

The quasiparticle loss would result from the qubit operator $\mathscr{O} = 2\phi_0\sin\frac{\varphi}{2}$ coupling to a noise variable with density

$$S_{\text{qp}}(\omega) + S_{\text{qp}}(-\omega) = 2\hbar\omega\,\Re\{Y_{\text{qp}}(\omega)\}\coth\left(\frac{\hbar|\omega|}{2k_BT}\right) \tag{5.6}$$

with the dissipative part of the Josephson junction admittance having the form[169, 172]

$$\begin{aligned}\Re\{Y_{\text{qp}}(\omega)\} = \quad & \sqrt{\frac{2}{\pi}}\frac{8E_J}{R_Q\Delta_{\text{Al}}}\left(\frac{2\Delta_{\text{Al}}}{\hbar\omega}\right)^{1.5} x_{\text{qp}}\sqrt{\frac{\hbar\omega}{2k_BT}} \\ & \times K_0\left(\frac{\hbar|\omega|}{2k_BT}\right)\sinh\left(\frac{\hbar|\omega|}{2k_BT}\right).\end{aligned} \tag{5.7}$$

Here, R_Q is the resistance quantum, and Δ_{Al} the Al superconducting gap with critical temperature T_c^{Al}. The fitted value here is the quasiparticle density x_{qp} relative to the number of Cooper pairs. For Fig. 5.5, the qubit operator has been transformed to $\mathscr{O} = \phi_0\sin(\frac{\varphi}{2} + \frac{I}{2})$[137], where I is the identity operator.

5.1.3.2 Purcell limit

The Purcell loss denotes the qubit relaxation arising from the coupling to the resonator. It is modeled according to the supplementary material in Ref. [171]. We assume that the resonator is coupled to a bath of harmonic oscillators $H/\hbar = \sum_k \omega_k b_k^\dagger b_k$. The interaction Hamiltonian is given by

$$H_{\text{int}}/\hbar = \sum_k g_k(r b_k^\dagger + r^\dagger b_k). \tag{5.8}$$

This interaction is treated perturbatively, causing energy transitions at a rate

$$\gamma_{if} = \frac{2\pi}{\hbar}\delta(E_i - E_f)|\langle\psi_f| H_{\text{int}} |\psi_i\rangle|^2 \tag{5.9}$$

from the dressed state $|\psi_i\rangle$ with energy E_i to the dressed state $|\psi_f\rangle$ with energy E_f. Here, δ is the delta function. The total rate is a sum over all initial and final state configurations of the bath, taking into account the thermal probability of occupying a given initial configuration.

The absorption rate from the bath of oscillators through the readout resonator[171] is

$$\Gamma^{\text{Purcell}\uparrow}_{|0\rangle_q \to |1\rangle_q} = P_{\text{res}}(0)\kappa n_{\text{th}}(\omega_{01})|\langle 01|r^\dagger|00\rangle|^2 \tag{5.10}$$

while the emission rate into the bath is

$$\Gamma^{\text{Purcell}\downarrow}_{|1\rangle_q \to |0\rangle_q} = P_{\text{res}}(0)\kappa[n_{\text{th}}(-\omega_{01}) + 1]|\langle 00|r|01\rangle|^2. \tag{5.11}$$

Here, $P(n) = [1 - e^{-\hbar\omega_R/k_B T_R}]e^{-n\hbar\omega_R/k_B T_R}$ is the probability to have n photons in a resonator with mode temperature T_R, κ is the total loss rate of the resonator[5], and $n_{\text{th}}(\omega) = \frac{1}{e^{\hbar\omega/k_B T_R - 1}}$. The state $|00\rangle$ corresponds to the dressed state of the coupled qubit-resonator system with the largest overlap to $|0\rangle_r|0\rangle_q$ in the bare basis. Equations 5.10 and 5.11 are valid in case the thermal occupation of the resonator is negligible, i.e. $P(n > 0) \approx 0$. Otherwise, additional terms are introduced to account for other thermally occupied initial states of the resonator.

5.1.3.3 Summary

Figure 5.5 summarizes the fluxonium lifetime results. From the fitted bounds for the quality factors of each dissipation channel, we deduce that the lifetime of the fluxonium at the zero field spot is limited by dielectric losses. Besides the range where the qubit is close to resonant with the readout resonator, it is evident that the qubit is also Purcell limited[173, 175] at $\Phi_{\text{ext}}/\Phi_0 = 0.5$. These limits are further reduced at this flux bias when including the high thermal population of both the qubit and the resonator.

For the fit in Fig. 5.5, a qubit mode temperature of $T_Q = 45$ mK, a nominal resonator mode temperature $T_R = 50$ mK, and the total loss rate at the low sweet spot $\kappa_R/2\pi = 922$ kHz were used. However, the qubit mode temperature obtained before the sample was

[5] The loss rate can be expressed in terms of the density of states of the bath $\rho(\omega)$, i.e. $\kappa = 2\pi\hbar\rho(\omega)|g_k|^2$.

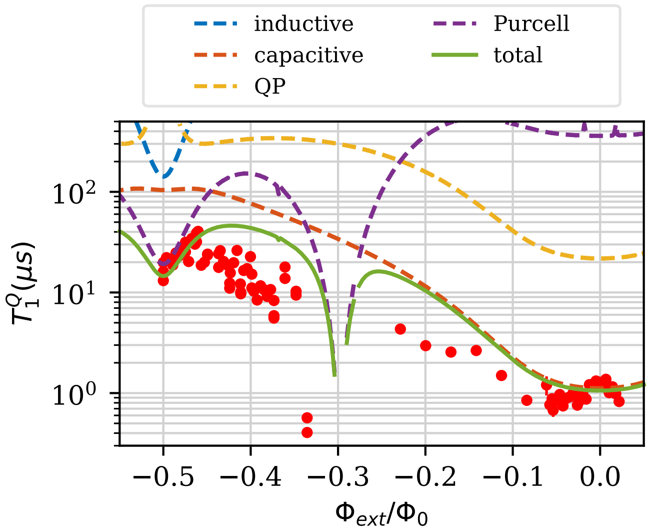

Figure 5.5: Fluxonium lifetime and loss mechanisms. The measured fluxonium T_1^Q is plotted over the lifetime limits set by different loss channels: dielectric, inductive, non-equilibrium quasiparticles (QP), and Purcell. The losses are calculated for qubit mode temperature of $T_Q = 45$ mK and resonator mode temperature of $T_R = 50$ mK. The following quality factors were used for the fit: $Q_{\mathrm{diel}} = 0.15 \cdot 10^6$, $Q_{\mathrm{ind}} = 30 \cdot 10^6$, and $x_{\mathrm{qp}} = 1.25 \cdot 10^{-6}$.

fully thermalized to the base plate of the cryostat is $T_Q = 112 \pm 2$ mK, measured similarly to Fig. 5.3. The minimum measured qubit temperature is 45 mK five months after the cooldown. In Fig. 5.6, the effect of qubit mode temperature on the quasiparticle and dielectric losses is shown. Depending on the qubit mode temperature, the quasiparticle density (Fig. 5.6a) and the dielectric quality factors (Fig. 5.6b) need adjustment to fit the same qubit lifetimes.

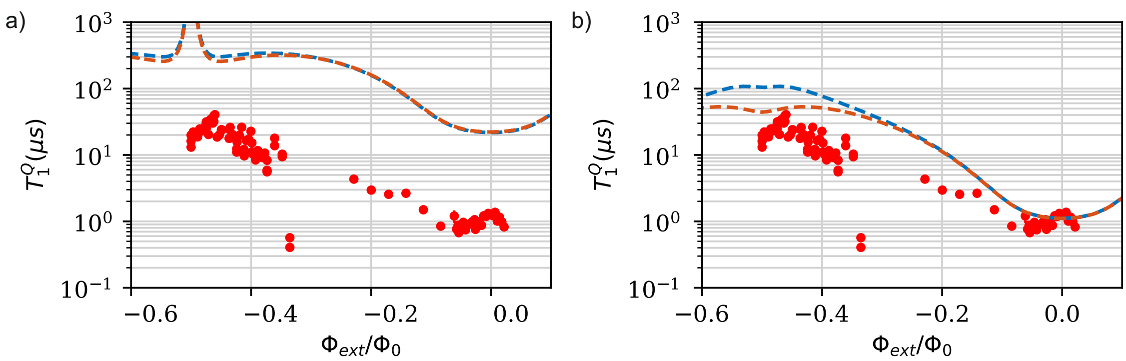

Figure 5.6: Temperature dependence of the quasiparticle and dielectric losses for 45 (blue dashed line) and 112 mK (red dashed line) qubit mode temperature. a) Quasiparticle losses for quasiparticle density $x_{\mathrm{qp}} = 1.25 \cdot 10^{-6}$. b) Dielectric losses for quality factor $Q_{\mathrm{diel}} = 1.5 \cdot 10^5$. The biggest influence the temperature has is around $\Phi_{\mathrm{ext}}/\Phi_0 = -0.5$ as at this bias the qubit frequency is low, making the qubit more susceptible to thermal excitation (increased $\Gamma_1^{\uparrow}$). Only the respective loss channel for each plot is considered, and all other loss contributions are neglected.

From a qubit spectroscopy measurement (see Fig. 5.4), we set the upper limit on the resonator mode temperature to $T_R = 112 \pm 29$ mK. For this temperature, the Purcell limit set on the qubit lifetime lies below the measured data. As the biggest contributors to the Purcell decay around $\Phi_{\text{ext}}/\Phi_0 = 0.5$ are the coupling strength g_{QR} and the resonator temperature, we vary both, as shown in Fig. 5.7. The lower temperature of 50 mK is chosen to be slightly above the minimum measured qubit mode temperature, while the variation of the coupling strength corresponds to 10% of its nominal value (refer to Table A.1). Even with reduced interaction strength, the Purcell limit at $T_R = 112$ mK lies below the measured fluxonium lifetimes, which dictates our choice to use $T_R = 50$ mK instead. For the lower resonator mode temperature and reduced coupling strength, the qubit lifetime is fully Purcell-limited around the low sweet spot.

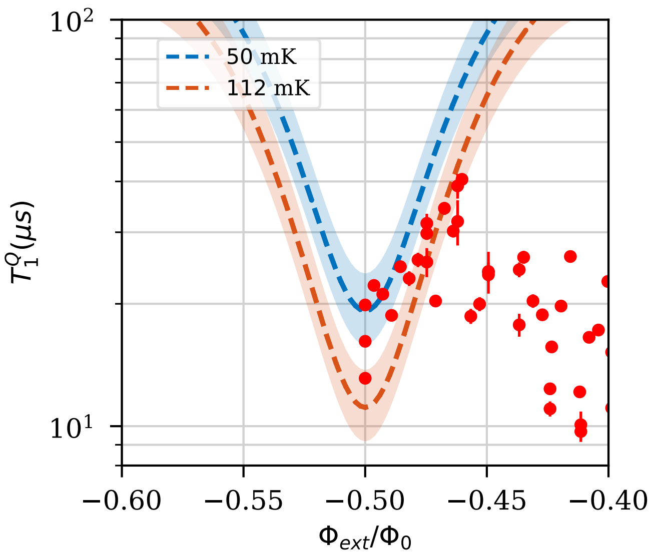

Figure 5.7: Adjusted Purcell limit. The temperature of the resonator mode is adjusted from 112 to 50 mK to show the temperature change during the measurements. The shaded regions correspond to a change in the coupling rate g_{QR} of $\pm 10\%$ for each resonator temperature. All other loss mechanisms are neglected.

To justify the limits we have set on the inductive losses, we plot the bound set by both Purcell decay and inductive loss with varying quality factor Q_{ind} (Fig. 5.8). The biggest impact of these dissipation channels is around the low sweet spot. When dielectric loss is included, the total bound is lowered further, imposing a high inductive quality factor to fit the qubit lifetimes around $\Phi_{\text{ext}}/\Phi_0 = -0.5$. It is worth noting that another model for the Purcell loss exists[176], according to which the lifetime of this qubit at the low sweet spot would be limited by inductive loss rather than Purcell decay. This model, however, assumes that lumped element approximation is valid, which is not true for this system.

The system would greatly benefit from a Purcell filter and improved fabrication techniques. A possible improvement in the fabrication would be to treat the sample with hydrofluoric (HF) acid after fabrication, which has been shown to reduce dielectric losses due to two-level system defects in resonators on silicon[177]. Another possibility is to adapt the fabrication

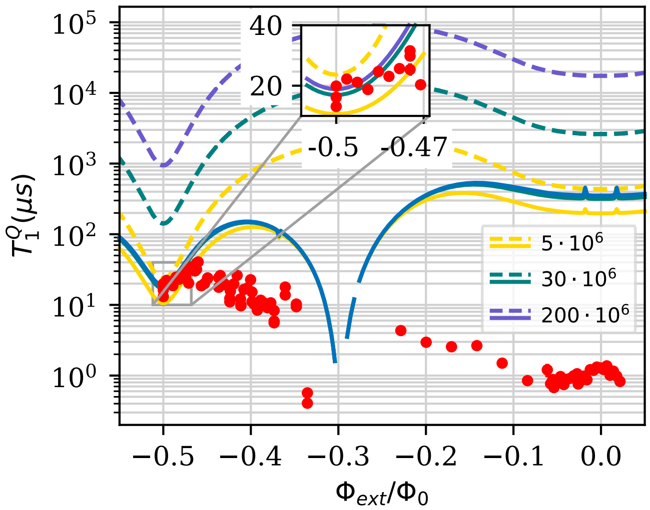

Figure 5.8: Limit for Purcell decay and inductive losses with varying inductive quality factor. The solid line represents the sum of inductive and Purcell loss, while the dashed line is only the limit set by inductive loss. The inset focuses on the illustrated losses around the low sweet spot.

process so it is compatible with more aggressive surface cleaning by swapping the Al with either tantalum[178, 179] or niobium capped by another metal[180].

5.1.4 Fluxonium decoherence

Ramsey T_2^R and echo T_2^E times provide valuable insight into the frequency components of the environment noise to which the qubit couples. Ramsey measurements capture the inhomogeneous dephasing time caused by both environmental flux noise and other sources of low-frequency noise. Echo T_2^E measurements provide complementary information about higher-frequency components of the noise spectrum. We address the flux stability in our system by measuring both T_2^R and T_2^E times of the fluxonium as a function of flux bias (see Fig. 5.9). At the flux sweet spot $\Phi_{\text{ext}} = -0.5\Phi_0$, the qubit frequency is first-order insensitive to flux variations. At this flux bias, the Ramsey time approaches the echo time but decreases rapidly away from the sweet spot. This behavior indicates the primary source of dephasing originates from low-frequency $1/f$ flux noise[181]. Further filtering and shielding of the setup, as well as the flux bias line, are required to reduce the susceptibility of the fluxonium to flux noise. The decoherence at the high sweet spot, although not shown here, is tenfold worse.

5.2 Cavity measurements

This section presents the measurements employed to characterize the cavity. Since the coupling between the qubit and the cavity is relatively weak compared to other decoherence rates, the system operates in the so-called weak dispersive regime. While this regime

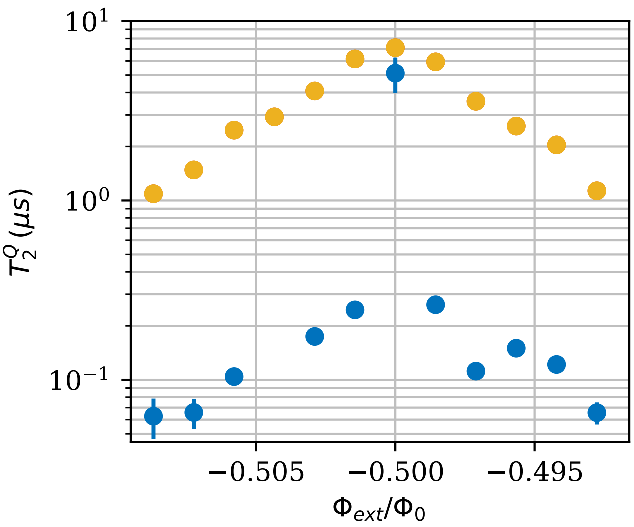

Figure 5.9: Fluxonium coherence times from Ramsey T_2^R (blue) and echo T_2^E (yellow) measurements as a function of flux bias. The sharp drop of T_2^R away from the sweet spot, compared to T_2^E, indicates that the main coherence limitation is low-frequency flux noise.

offers the advantage of reduced sensitivity to ancilla-induced noise, it also introduces challenges, as achieving universal control requires a more unconventional set of measurement techniques.

5.2.1 Echoed conditional displacements

When a resonant drive is applied to the cavity, the qubit-state-dependent frequency shift leads to different displacement trajectories in phase space, effectively realizing a conditional displacement[182]. The rate at which the states $|0,\alpha\rangle$ and $|1,\alpha\rangle$ would separate in phase space depends on the dispersive shift. To avoid ancilla-induced dephasing, a small dispersive shift is desirable, which would considerably increase the gate and the state tomography time.

To maintain a low dispersive shift, i.e., $\chi \lesssim \Gamma_1, \Gamma_2, \kappa$, while not compromising on gate times, a new technique was developed, known as Echoed Conditional Displacement (ECD)[126, 183]. It combines large displacements on the cavity with a cleverly engineered sequence of pulses to achieve fast state separation (i.e. faster than $\sim \chi^{-1}$). Therefore, ECD enables the extraction of the Hamiltonian parameters even when the dispersive shift is smaller than the qubit decoherence rate.

The working principle of the protocol is illustrated in Fig. 5.10. Consider an ideal two-level system dispersively coupled to a cavity mode. The system is initialized in a superposition of the qubit ground and excited states, with the cavity in the vacuum state. A cavity drive at frequency $(\omega_c^{|0\rangle_q} + \omega_c^{|1\rangle_q})/2$ is applied, displacing the cavity into a coherent state $|\alpha\rangle$. Due to the dispersive interaction, the two states $|0,\alpha\rangle$ and $|1,\alpha\rangle$ acquire opposite phase rotations during a wait time t_w. The cavity is then displaced back towards the

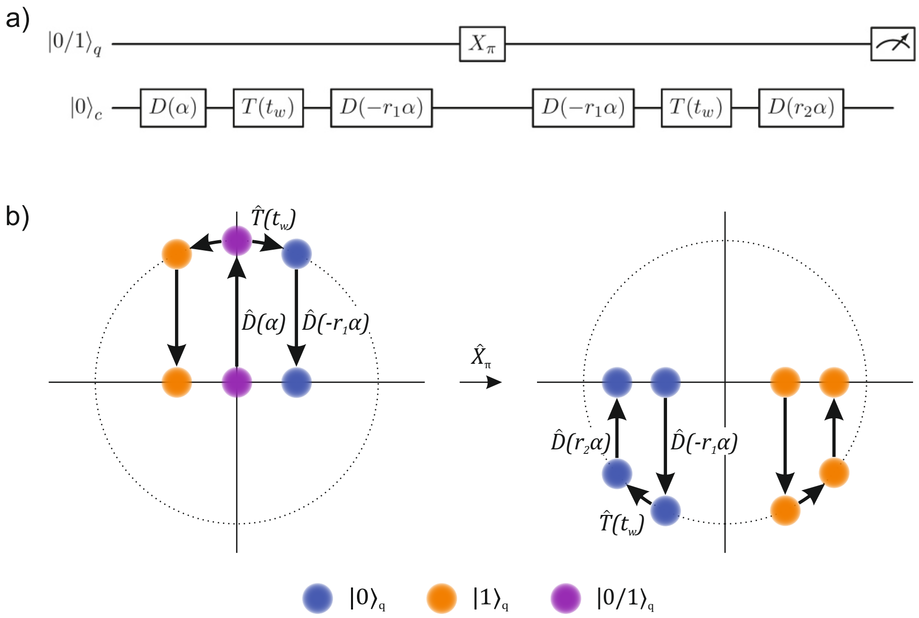

Figure 5.10: Echoed conditional displacement: Sequence and phase-space trajectories. a) For the sequence, the initial state could be either ground or excited, while the cavity is said to be in vacuum, but can be in an arbitrary state. b) The sequence $D(r_2\alpha)T(t_w)D(-r_1\alpha)X_\pi D(-r_1\alpha)T(t_w)D(\alpha)$ is split in two for visual clarity. The coefficients $r_1 = \cos(\chi t_w/2)$ and $r_2 = \cos(\chi t_w)$ ensure the states are displaced only to an amplitude $|\alpha|$ and back to the I-axis. The illustration assumes instantaneous pulses.

$Q = 0$ axis by an amount $-\alpha\cos(\chi t_w/2)$; the cosine term compensates for the different displacement amplitude required due to this evolution. A non-selective π-pulse flips the qubit population, after which a second displacement $D(-\alpha\cos(\chi t_w/2))$ is applied. During the subsequent wait time t_w, the two components of the state separate further in phase space before being displaced back to the I-axis, $D(\alpha\cos(\chi t_w))$. Due to the symmetric sequence and the intermediate echo pulse, this protocol achieves a conditional displacement while simultaneously canceling unwanted dynamical phases arising from dispersive and ac Stark shifts. The ECD gate is the workhorse of all cavity characterization methods used in this manuscript.

5.2.2 Semi-classical trajectories

To develop an understanding of the working principle underlying the following measurements, we first consider the theory of semiclassical trajectories. A detailed derivation of the formalism can be found in the Supplementary Information of Ref. [183].

We restrict our analysis to the qubit-cavity subsystem, described by the Hamiltonian

$$H/\hbar = \Delta c^\dagger c + \chi q^\dagger q c^\dagger c + \chi'(q^\dagger q)^2 c^\dagger c + K_c(c^\dagger c)^2 + K_q(q^\dagger q)^2 + \epsilon^*(t)c + \Omega(t)q + \text{h.c.}, \tag{5.12}$$

expressed in the frame rotating at the respective drive frequencies after applying the rotating-wave approximation. Here, $\epsilon(t)$ and $\Omega(t)$ denote complex-valued, time-dependent drives applied to the cavity and qubit, respectively. In addition to the dispersive shift χ, we include a higher-order cross-Kerr term χ', while K_c and K_q represent the self-Kerr nonlinearities of the cavity and the qubit.

If the qubit is not driven, one can extract the cavity phase-space trajectories conditioned on the qubit state. Choosing an appropriate time-dependent displaced frame $U = D^\dagger(\alpha(t))$ removes the terms linear in α in the transformed Hamiltonian. This condition leads to the following coupled differential equations:

$$\partial_t \alpha_{|0\rangle_q}(t) = i\Delta\alpha_{|0\rangle_q}(t) - 2iK_c|\alpha_{|0\rangle_q}(t)|^2\alpha_{|0\rangle_q}(t) - \frac{\kappa}{2}\alpha_{|0\rangle_q}(t) - i\epsilon(t) \tag{5.13a}$$

$$\partial_t \alpha_{|1\rangle_q}(t) = i\Delta\alpha_{|1\rangle_q}(t) - 2iK_c|\alpha_{|1\rangle_q}(t)|^2\alpha_{|1\rangle_q}(t) - \frac{\kappa}{2}\alpha_{|1\rangle_q}(t) - i\epsilon(t) - i(\chi + 2\chi'|\alpha_{|1\rangle_q}(t)|^2)\alpha_{|1\rangle_q}(t). \tag{5.13b}$$

In the limit of small displacements, the terms proportional to the cavity Kerr coefficient and the non-linear dispersive shift can be neglected, and for the initial condition $\alpha(t = 0) = 0$, these equations can be solved to obtain the semiclassical cavity trajectories

$$\alpha_{|0/1\rangle_q}(t) = e^{-\frac{1}{2}(\pm i\chi+\kappa)t}\left(e^{\frac{1}{2}(\pm i\chi+\kappa)t_0}\alpha_{|0/1\rangle_q}(t_0) - i\int_{t_0}^{t} e^{\frac{1}{2}(\pm i\chi+\kappa)\tau}\epsilon(\tau)d\tau\right). \tag{5.14}$$

Here, we have picked $\Delta = \chi/2$. In the case of an initial coherent state at $t = t_0$ and no drive $\epsilon(t) = 0$, this state will acquire a phase $\pm\chi t/2$ with the sign depending on the qubit state $|0/1\rangle_q$. Moreover, it will decay towards the vacuum at a rate $\kappa/2$. An applied drive would change the position of this state in phase space; for non-instantaneous drives on the time scales of χ^{-1}, the free evolution would curve the trajectories during the drive.

5.2.3 Out-and-back measurement

In systems operating in the strong dispersive regime, the Hamiltonian can be characterized by displacing the cavity and performing spectroscopy on the qubit. The resulting photon-number-split peaks in the qubit spectrum can then be fitted to extract the effective dispersive shift. However, in the weak dispersive regime, these number-split peaks can be difficult or impossible to resolve. The sequence *out-and-back* was developed in Ref. [183] to deal with this problem and extract Hamiltonian parameters, which are small compared to the qubit decoherence.

The measurement sequence, illustrated in Fig. 5.11a, begins with the preparation of the qubit in either $|0\rangle_q$ or $|1\rangle_q$ by applying a selective $1.2\,\mu$s Gaussian pulse. The preparation is executed through active reset, enabled by the parametric amplifier (see Appendix G for details on active reset). The cavity is then displaced with a 40 ns square pulse and allowed to evolve for a duration $t_w = 1\,\mu$s before being displaced back toward vacuum by a second pulse with a variable phase ϕ. Large displacements on the cavity enable the extraction of the Hamiltonian parameters even when the dispersive shift is smaller than the qubit decoherence rate. At the end of the protocol, we flip the qubit with a narrow-bandwidth π-pulse and then read it out with a $2\,\mu$s square pulse. The π-pulse is only successful if the phase for the displacement back to the cavity ground state is correct, i.e., when the phase of the second displacement pulse matches the accumulated phase of the cavity evolution. The measurement sensitivity is enhanced by repeating the cavity sequence five times.

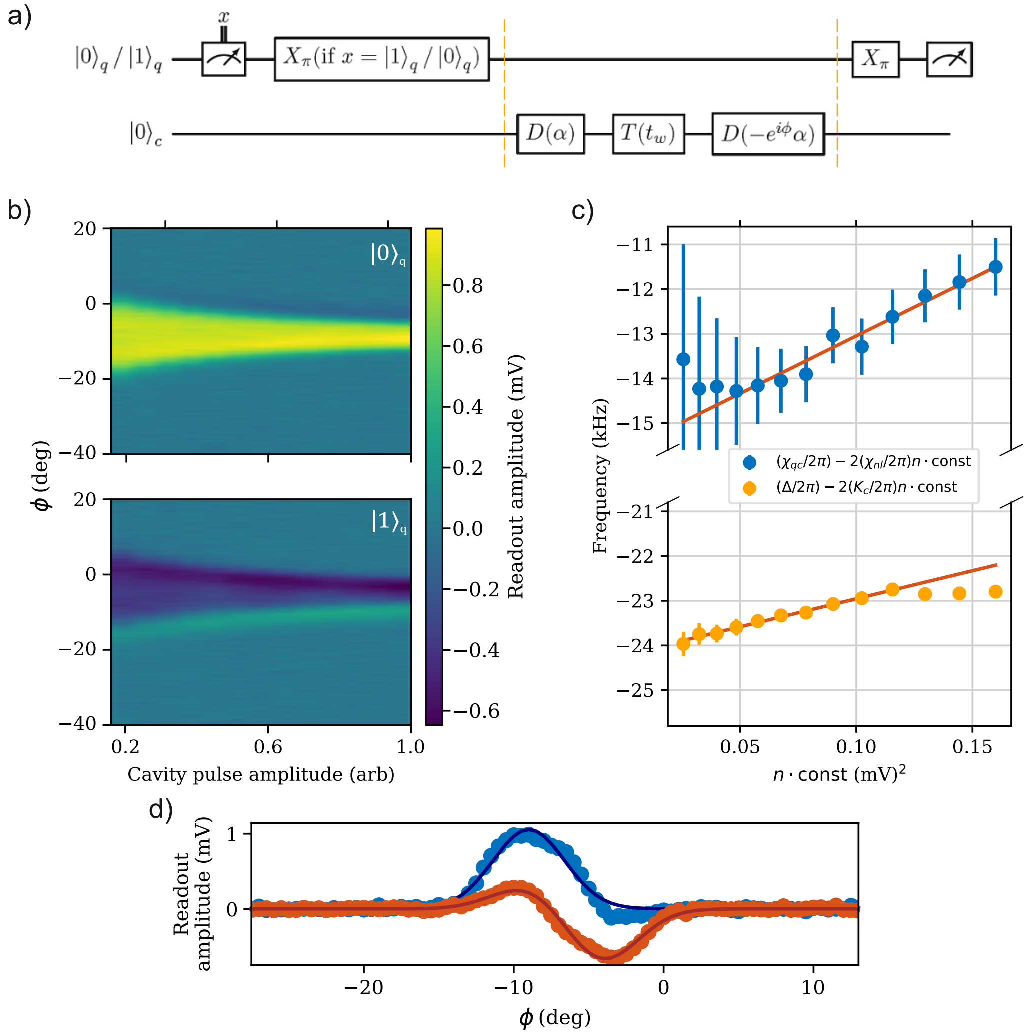

Figure 5.11: Out-and-back measurement. a) Sequence. The first part includes an active reset conditioned on the measurement result x. The cavity sequence $D(-e^{i\phi}\alpha)T(t_w)D(\alpha)$ (enclosed between the two yellow dashed lines) is repeated five times to enhance the effect. b) Raw data from the measurement at $\Phi_{\text{ext}} = 0.5\Phi_0$ for $|0\rangle_q$ and $|1\rangle_q$. The dataset corresponding to each cavity pulse amplitude is fitted with either a single Gaussian for $|0\rangle_q$ or two Gaussians for $|1\rangle_q$, accounting for decay. c) Frequencies extracted from the fit of b) as a function of the pulse voltage squared. The rotation frequency $\Delta_{|0\rangle_q}/2\pi$ is plotted in yellow, while the difference $(\Delta_{|1\rangle_q} - \Delta_{|0\rangle_q})/2\pi$ is plotted in blue. The data sets are fitted by linear regression (red). From the fit, we extract $\Delta/2\pi = -24.21 \pm 0.03$ kHz and $\chi_{QC}/2\pi = -15.6 \pm 0.2$ kHz, the latter of which we use to find the unknown constant from the x-axis and thus calibrate the photon number n. d) Data sets and fits from the last linecut in b), i.e. cavity pulse amplitude is one. The blue corresponds to $|0\rangle_q$, while the red to $|1\rangle_q$.

The raw measurement data are shown in Fig. 5.11b, where the upper panel corresponds to the qubit initially prepared in $|0\rangle_q$ and the lower panel to preparation in $|1\rangle_q$. For the $|0\rangle_q$ case, each line cut is fitted with a Gaussian function, whose mean position and standard deviation are then used to constrain the fit for the $|1\rangle_q$ data, which exhibits a sum of two Gaussian components. This double-Gaussian behavior deviates from the single-peaked response observed in the initial implementation reported in Ref. [183]. The origin of this difference lies in the relatively short qubit relaxation time: during the sequence, the qubit can spontaneously decay or become thermally excited, and this stochastic behavior averages out over multiple repetitions, producing a composite signal that reflects contributions from both the intended $|1\rangle_q$ state and the decay to $|0\rangle_q$.

Finally, the mean phase accumulation obtained from the Gaussian fits is converted from degrees to frequency by dividing it by the waiting time t_w. The resulting data are shown in Fig. 5.11c. The displacement amplitude is scaled by the voltage range of the AWG (0.4 V), and the squared value of this quantity is proportional to the average photon number in the cavity, up to a constant factor. Using the semiclassical trajectories derived in Eq. 5.13, the cavity rotation frequencies for the qubit in its ground and excited states can be expressed as

$$\Delta_{|0\rangle_q} = \Delta - K_c n \tag{5.15a}$$

$$\Delta_{|1\rangle_q} = \Delta - \chi - (2K_c + 2\chi_{nl})n \tag{5.15b}$$

where n denotes the mean photon number in the cavity. Here, χ_{nl} is the same as χ' from Eq. 5.13. The experimental data are fitted with a linear regression to extract the intercepts, corresponding to the cavity detuning and the qubit-cavity dispersive shift. For this flux bias point, the extracted parameters are $\Delta/2\pi = -24.21 \pm 0.03$ kHz and $\chi_{QC}/2\pi = -15.6 \pm 0.2$ kHz. To determine the slopes accurately, the calibration between the AWG pulse amplitude and the cavity photon number must be established.

5.2.4 Photon number calibration

An essential part of extracting the Hamiltonian parameters is calibrating the conversion coefficient between the output of the AWG and the number of photons in the cavity. In the strong dispersive regime, one can use conditional π-pulses on $\omega_q^{|0\rangle_c}, \omega_q^{|1\rangle_c}, \ldots, \omega_q^{|n\rangle_c}$. By simultaneously fitting the traces corresponding to each qubit frequency for varying AWG pulse amplitudes, the conversion factor can be extracted[184]. In the weak dispersive regime, the qubit frequencies for different numbers of photons in the cavity are not resolved, no matter how selective the pulse driving the qubit is.

Calibrating the displacements in the weak dispersive regime exploits the geometric phase acquired by the qubit during an ECD-like sequence, depicted in Fig. 5.12. We prepare the qubit in a superposition $(|0\rangle_q + |1\rangle_q)/\sqrt{2}$ and apply a sequence described by the drive

$$\epsilon(t) = \epsilon_0[g(t) - rg(t - (t_p + t_w)) - rg(t - (2t_p + t_w)) + g(t - (2t_p + 2t_w))] \tag{5.16}$$

where ϵ_0 is its amplitude and $g(t)$ can be an arbitrary pulse shape, but I have used a simple square pulse of duration $t_p = 40$ ns for both simulations and measurements. This sequence is designed to create the trajectories, depicted in Fig. 5.12b, which enclose an area in phase

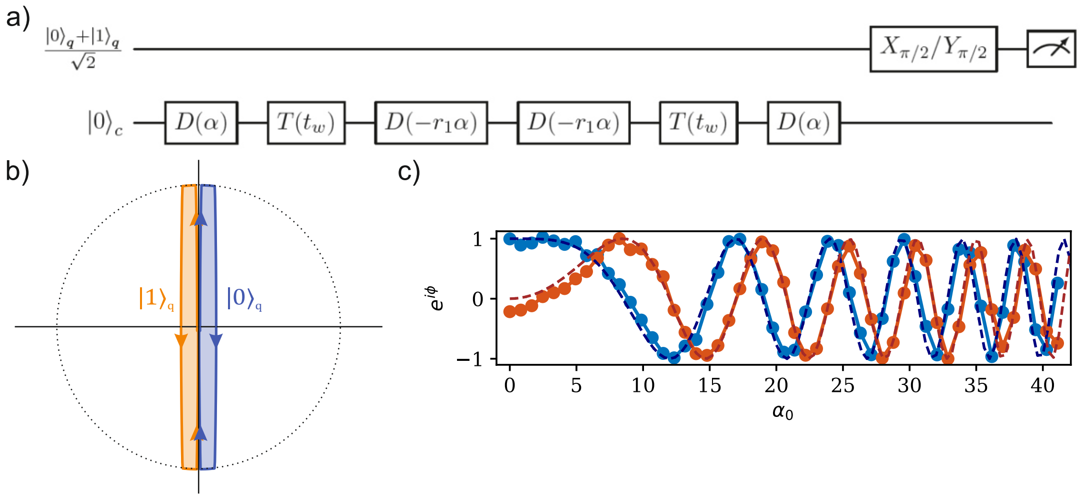

Figure 5.12: Geometric phase. a) The sequence $D(\alpha)T(t_w)D(-r_1\alpha)D(-r_1\alpha)T(t_w)D(\alpha)$ is used to measure the geometric phase acquired by the qubit. b) Qubit trajectories during the ECD-like sequence, depicted in a). The qubit will acquire a geometric phase proportional to the enclosed area, and therefore to the displacement α. c) Measured (dots) and fitted (lines) data for the phase acquired during the sequence for varying displacements. The red color corresponds to $X_{\pi/2}$ at the end of the sequence, and blue to $Y_{\pi/2}$. The extracted conversion coefficient is about 41.1. The measurements were run for $t_w = 200\,$ns, while the simulation in a) was chosen to have a ten-fold higher dispersive shift to make the enclosed area more visible.

space for both ground and excited qubit states. These trajectories have been numerically calculated by solving Eq. 5.14. At the end of the sequence, the qubit and the cavity are disentangled, i.e. the cavity is back to $n = 0$ no matter the qubit state, and the qubit state is $(|0\rangle_q + e^{i\phi}|1\rangle_q)/\sqrt{2}$, where ϕ is the enclosed area. To extract this phase, we need to measure both expectation values, $\langle\sigma_x\rangle$ and $\langle\sigma_y\rangle$, by applying either $X_{\pi/2}$ or $Y_{\pi/2}$ before the readout. This procedure is executed at the low sweet spot $\Phi_{\text{ext}} = 0.5\Phi_0$, and the results are plotted in Fig. 5.12c. From here on, the displacement amplitude for a 40 ns square pulse is calculated as $|\alpha| = 41.1A_{\text{AWG}}$, where A_{AWG} is the amplitude of the pulse at the AWG. With this scaling and by using Eq. 5.15, we can extract the self-Kerr coefficient of the cavity $K_c/2\pi = 1.22\pm0.13\,$Hz, and the higher-order nonlinearities $\chi_{nl}/2\pi = 0.4\pm0.1\,$Hz.

5.2.5 Characteristic function

To complete the toolbox of cavity measurements, a method for state tomography is needed. While the conventional method involves the use of Wigner[112, 185] or Husimi Q[186] functions, these methods both benefit from a high dispersive shift. The time it takes for a Wigner function measurement is $t = \pi/\chi$, plus the duration of the pulses on both the cavity and the qubit. If this dispersive shift is on the order of a hundred kilohertz, the time for this measurement would exceed 60 μs, requiring coherence times of the ancilla, which are inaccessible for many labs. The Husimi Q function measurement also requires a strong dispersive regime to play qubit pulses selective on the number of photons in the cavity.

The characteristic function $\mathcal{C}(\xi)$ provides an alternative and powerful way to represent quantum states in phase space. It encodes the complete information of the density operator, and is connected to the other probability function via the Fourier transform. A full description of the characteristic function can be found in App. C. By measuring the system's response to controlled sequences of displacements and phase rotations, one can experimentally reconstruct $\mathcal{C}(\xi)$.

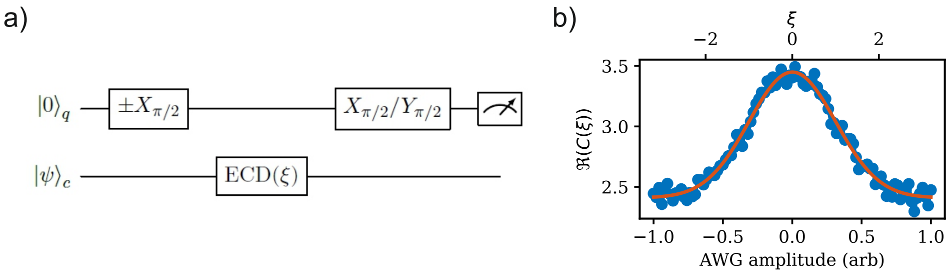

Figure 5.13: Characteristic function. a) Sequence for measuring the characteristic functions. The $\pm X_{\pi/2}$ pulse compensates for both decay and excitation of the ancilla during the readout. The two measurements are subtracted. The pulse $X_{\pi/2}$ or $Y_{\pi/2}$ at the end corresponds to measuring the real or the imaginary part of $\mathcal{C}(\xi)$. b) Scaling. A linecut of $\mathcal{C}(\xi)$ is used for calibrating the scaling in phase space. Here, $t_w = 600\,\text{ns}$, $t_p = 40\,\text{ns}$, giving the scaling $|\xi| \approx A_{\text{AWG}}/0.312$.

To measure the characteristic function, the pulse sequence from Fig. 5.13a is played on the system. It consists of an ECD gate embedded in a qubit Ramsey sequence, which maps the qubit expectation values $\langle \sigma_x \rangle$ and $\langle \sigma_y \rangle$ to the real and imaginary parts of the characteristic function. This state tomography method is more forgiving in terms of qubit coherence times, as the dispersive shift does not limit the time it takes to measure $\mathcal{C}(\xi)$. The phase space sampling depends both on t_w and t_p from the ECD sequence. For a fixed set of parameters, we can calibrate the phase space displacement by measuring a vacuum state in the cavity. As the characteristic function of this state is a 2D Gaussian function with $\sigma = 1$, a simple line cut is sufficient to extract the scaling. Fig. 5.13b depicts $\Re(\mathcal{C}(\xi))$ for $t_w = 600\,\text{ns}$ and $t_p = 40\,\text{ns}$. From the fit, we extract the phase space scaling $|\xi| \approx A_{\text{AWG}}/0.312$.

5.2.5.1 Cavity lifetime

The lifetime of the cavity mode is a key parameter characterizing the overall performance of a system as a bosonic qubit. There are several methods to measure this quantity, depending on whether the cavity is directly accessible or coupled to an ancillary qubit. In the absence of a qubit, the most direct approach involves populating the cavity with a weak coherent pulse and monitoring the reflected or transmitted signal using heterodyne detection[70]. The exponential decay of the signal amplitude yields the photon lifetime. When the cavity is coupled dispersively to a qubit, its lifetime can also be inferred indirectly through the qubit. One method is to use a displacement pulse on the cavity with a variable delay time afterward, and probe the state similar to a Husimi Q measurement[153]. This, however, requires the ability to play number-selective pulses on the qubit, conditioned on no photons in the cavity. The alternative approach is to measure the characteristic function of a decaying coherent state, as done in Ref. [126].

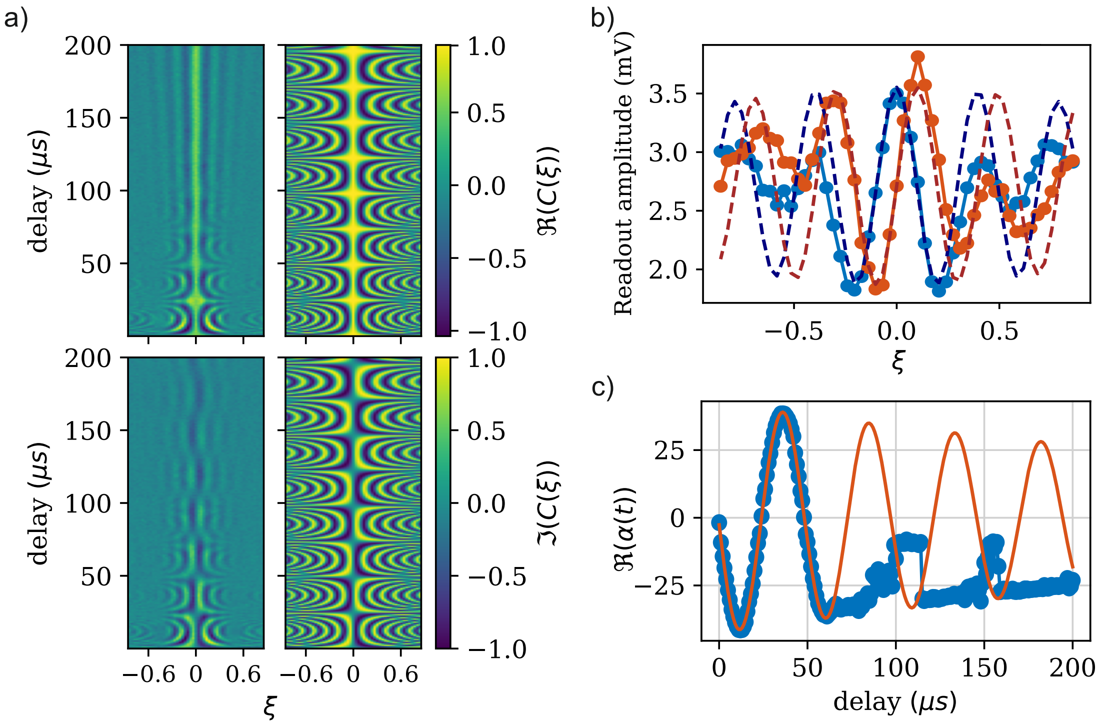

Figure 5.14: Cavity T_1 measured via characteristic function. a) Real (top) and imaginary (bottom) parts of the characteristic function, measured at the axis $Q = 0$, or equivalently $\Im(\xi) = 0$. The left column presents the normalized data, while the right column presents the fit to the data. b) In order to get the fit parameters in a), we first fit $\mathcal{C}(\xi)$ for each time stamp $t_i \in [0, 200]\,\mu$s to extract the coherent state at the time $\alpha(t_i)$. Here, the blue and red dots represent the line cuts from a) at $t_d = 2.016\,\mu$s with their respective fits. The fit is run for $\xi \in [-0.5, 0.5]$. c) Once we have extracted $\alpha(t)$, we fit its real part to $Ae^{-t/2T_1}\cos(2\pi\Delta f + \phi) + B$ to extract the cavity lifetime $T_1 = 210 \pm 40\,\mu$s. The fit is restricted to times smaller than $66\,\mu$s, as the data quality does not allow the extraction of $\alpha(t > 66\,\mu s)$.

We use the same protocol for the ECD gate from Fig. 5.10a with $t_w = 200\,$ns[6], but preceded by a pulse, creating the coherent state $\alpha_0 = \alpha(t = 0) = i(42.4 \pm 0.3)$[7]. Then the characteristic function is measured after a varying delay, allowing the state to decay. The process is depicted in Fig. 5.14a. The characteristic function of a coherent state is

$$\mathcal{C}(\alpha(t)) = e^{-\frac{|\alpha(t)|^2}{2}} e^{2i\Im(\alpha(t)\xi^*)} \tag{5.17}$$

[6] The different wait time results in a different scaling between AWG amplitude and ξ, which is $|\xi| \approx A_{\mathrm{AWG}}/1.166$.

[7] This value is extracted from the amplitude of the fit in Fig. 5.14c. It is purely imaginary, as the real part is zero at $t = 0$.

where $\alpha(t)$ is the decaying coherent state

$$\alpha(t) = \alpha_0 e^{-i\Delta t} e^{-\frac{t}{2T_1^c}}. \tag{5.18}$$

There are two competing effects in the data plots of Fig. 5.14a. Firstly, the imperfect $\pi-$ and $\pi/2-$pulses on the qubit result in stray oscillation on top of the data (see Fig. 5.14a left for $t_d \gtrsim 70\,\mu$s), which becomes more prominent as the delay increases. Secondly, the width of the coherent state is ideally the same as that of the vacuum state; however, since cavity decay is enhanced by the number of photons present, the measured characteristic function appears to have a smaller width compared to the simulations plotted on the right side of each map. Nonetheless, this data quality is sufficient to extract the cavity lifetime.

For the extraction of $\alpha(t)$, each dataset, consisting of $\Re(\mathcal{C}(t))$ and $\Im(\mathcal{C}(t))$, is fitted to Eq. 5.17, an example of which is shown in Fig. 5.14b. The characteristic function $\mathcal{C}(\xi)$ is measured only for $\xi \in \mathbb{R}$, which means that we only have access to the real part of $\alpha(t)$. To measure the imaginary part, one would need the same dataset for $\xi \in i\mathbb{R}$. As both the real and imaginary part decay with the same constant, and due to the lengthy measurements required to extract this additional information (>24 hours), we extract the cavity lifetime $T_1 = 210 \pm 40\,\mu$s only from the fit of $\Re(\alpha(t))$, as shown in Fig. 5.14c. This procedure was only feasible at $\Phi_{\text{ext}} = -0.5\Phi_0$ as the fluxonium coherence time was insufficient at other flux bias points. More simulations and considerations on these measurements can be found in Appendix C.

The cavity mode is continuously coupled to the readout resonator, allowing excitations to leak out through Purcell decay potentially. To rule out the possibility that the cavity lifetime is limited by this coupling, we investigate the Purcell effect on the cavity. The simulations follow the method outlined in Section 5.1.3.2, with the matrix element adjusted to account for excitation loss from the cavity. In this analysis, we neglect internal losses of the fluxonium; only the direct loss via the readout resonator is included. The measured cavity lifetime lies below the calculated Purcell limit of $T^{\text{Purcell}}_{|1\rangle_C \to |0\rangle_C} \approx 2.8$ ms. This limit remains nearly constant across the entire flux bias range, except near the point where the fluxonium's first excited state intersects with the fundamental cavity mode.

5.3 Dispersive shifts

Having determined all Hamiltonian parameters, the dispersive shift between the qubit and the resonator χ_{QR} and between the qubit and the cavity χ_{QC} can be calculated as depicted in Fig. 5.15. We extract the dispersive shifts χ_{QR} and χ_{QC} using various measurements, which will be described in detail below. The different methods are required as χ_{QR} is higher than the qubit linewidth at any external flux, and χ_{QC} is lower.

The extracted data is plotted together with the numerical model predictions in Fig. 5.15. The dispersive shift between the resonator and the fluxonium is much higher than both the linewidth of the resonator and the qubit, allowing for easy extraction of the quantity as a function of flux bias. For Fig. 5.15a, several approaches were used. Firstly, because of the high mode temperature of the qubit, the excited state population is non-negligible

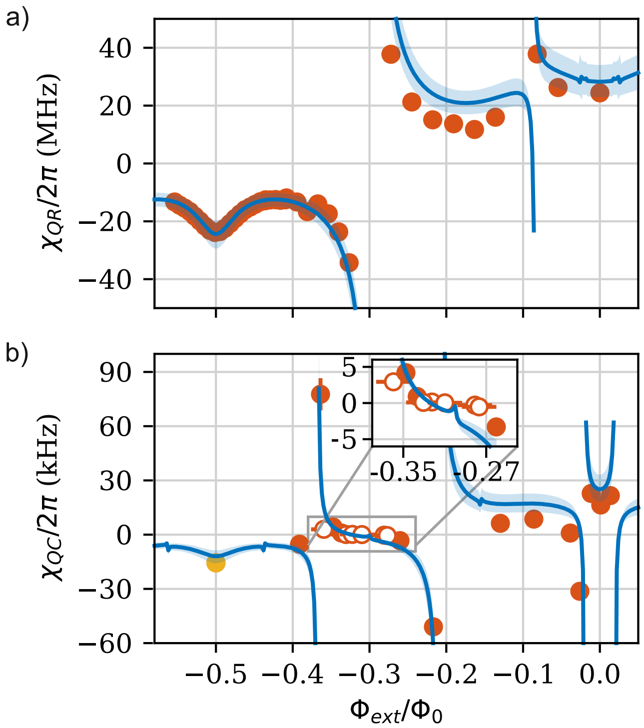

Figure 5.15: System dispersive shifts. a) The dispersive shift between the qubit and readout resonator as a function of flux bias. The shaded region corresponds to a 10% variation in the coupling strength g_{QR}. b) The dispersive shift between the qubit and high-Q cavity as a function of flux bias. The yellow point is measured with an echoed conditional displacement sequence, while the rest (red) are extracted from qubit spectroscopy as a function of the number of photons in the cavity, measured in the same (full) or consecutive (empty) cooldown. The shaded region is a confidence interval, which corresponds to 15% uncertainty of the coupling strength g_{RC}. Note the smooth transition between positive and negative values of the interaction around $\Phi_{\mathrm{ext}}/\Phi_0 = -0.35$.

around $\Phi_{\mathrm{ext}}/\Phi_0 = 0.5$, causing the resonator response to split by an amount equal to the dispersive shift (see Fig. 5.4a). At higher qubit frequencies, i.e., $f_{01} \geq 5\,\mathrm{GHz}$, the qubit cannot be thermally excited by the bath, and the resonator has a single resonance. At these flux bias points, a two-tone measurement was used, in which an extra tone at the qubit frequency was applied, while the resonator spectroscopy was monitored. At the high sweet spot $\Phi_{\mathrm{ext}}/\Phi_0 = 0$, χ_{QR} was extracted from the qubit time-domain spectroscopy measured after populating the resonator by applying a pulse.

The yellow point in Fig. 5.15b was measured with the out-and-back method[8]. As this method is time-consuming, only the dispersive shift at $\Phi_{\mathrm{ext}}/\Phi_0 = 0.5$ was measured this way. As the cavity frequency does not change much as a function of flux threading the fluxonium loop, we use the same photon number calibration from the measured dispersive shift for the rest of the points in Fig. 5.15b. These points were extracted by monitoring the qubit frequency after the cavity had been displaced to a different coherent state. We

[8]The raw data for this measurement can be found in Fig. 5.11.

fit the qubit resonance shift observed as a function of the photon number and extract the dispersive shift $\chi_{QC}/2\pi$. These subsequent measurements allowed for fast characterization of the dispersive shift across the entire flux map.

We notice that in Fig. 5.15b, the dispersive shift χ_{QC} is smoothly tunable between positive and negative values through zero, meaning we can switch off the interaction between the fluxonium and the cavity. The feature at $\Phi_{\text{ext}} \approx -0.3\Phi_0$, stems from the crossing between the fluxonium and the readout resonator (see Fig. 5.1 left). This feature is not detrimental as the required qubit frequency change to go from $\chi_{QC}/2\pi = 0\,\text{Hz}$ to this crossing is about a hundred times larger than the qubit linewidth, which is limited by flux noise at this flux bias point, and it would require a large change in Φ_{ext}. The largest negative dispersive shift we measured is $\min(\chi_{QC})/2\pi = -51 \pm 2\,\text{kHz}$, which is comparable to other recent works [56, 57, 183]. Also, the closest to zero we measured was $\min(|\chi_{QC}|)/2\pi = 30 \pm 94$ Hz. While such measured results can also be achieved with a tunable transmon as an ancilla[187], in this setup, it is possible to tune the interaction strictly to zero without any additional elements or pump tones. Furthermore, we can tune to the opposite sign of χ_{QC}, with our largest measured value at 78 ± 9 kHz, similar to tunable transmon implementations using additional elements like an inductive coupler[188].

The last remark concerns the fitting routine from Fig. 5.1 center. Initially, the non-linear optimization and curve fitting of the qubit resonance as a function of flux were executed with a Python solver *lmfit* with the model given by the *scQubits* simulation for f_{01}. However, the energies from the initial fit were adapted from $E_J/h = 10.34 \pm 0.04\,\text{GHz}$, $E_C/h = 3.36 \pm 0.04\,\text{GHz}$, and $E_L = 1.038 \pm 0.003\,\text{GHz}$ to the values from the caption of Fig. 5.1 to accommodate the avoided crossing in Fig. 5.15b, observed around the high sweet spot. Upon further investigation, the cause of this crossing is the two-photon qubit transition f_{13} crossing the cavity resonance. The exact location of the higher energy levels plays an important role in how χ behaves. Even small changes in the qubit energies E_J, E_C, E_L result in new features, such as the aforementioned avoided crossing. For the most accurate predictions of the dispersive shift, the fit must include the higher fluxonium transitions.

5.3.1 Qubit-induced dephasing of the cavity

When the qubit relaxation rate Γ_1^Q is much smaller than the dispersive shift $|\chi_{QC}|$, the cavity dephasing time depends solely on the thermal population of the qubit and its lifetime[189, 190]. If this relationship swaps, i.e., $|\chi_{QC}| \ll \Gamma_1^Q$ (around the zero-crossing flux bias), the cavity dephasing becomes proportional to χ_{QC}^{-2},

$$T_\phi^C = \frac{\Gamma_1^Q}{\chi_{QC}^2 n_{\text{th}}^Q}, \tag{5.19}$$

where n_{th}^Q is the thermal population of the fluxonium. If we take our best qubit lifetime $T_1^Q \approx 40\,\mu\text{s}$, this equation becomes valid for a dispersive shift $|\chi_{QC}|/2\pi \ll 4$ kHz. Compared to the limit set by the cavity lifetime of $210\,\mu\text{s}$ and assuming the same thermal excitation of the fluxonium as measured at $\Phi_{\text{ext}} = -0.5\Phi_0$, the qubit-induced decoherence becomes negligible for $|\chi_{QC}|/2\pi \ll 1.8\,\mathrm{kHz}$. In fact, this limit is further raised by the smaller thermal occupation of the qubit at higher frequencies.

As a result, in the regions where the qubit lifetime is compromised by Purcell decay through the readout resonator $\Phi_{\mathrm{ext}} \approx -0.3\Phi_0$, the cavity coherence is preserved as the dispersive shift is much lower than the qubit relaxation rate (see Fig. 5.15 and Fig. 5.5). To measure this effect on the cavity dephasing times, fast flux control is required as qubit coherence is only sufficient at $\Phi_{\mathrm{ext}} = 0.5\Phi_0$.

5.4 Kerr shifts

In the coupled fluxonium-cavity-resonator system, nonlinearities originating from the Josephson potential of the fluxonium give rise to both *self-Kerr* and *cross-Kerr* frequency shifts. These nonlinear interactions determine how the resonance frequencies of the cavity, resonator, and qubit depend on their photon populations and on the excitation states of one another. Depending on the context, Kerr shifts can either limit bosonic state fidelities[153, 191] or be harnessed as a resource for generating[192] and stabilizing[117] nonclassical states.

The *self-Kerr shift* K_i of a mode manifests as a photon-number-dependent correction to its resonance frequency, such that the effective frequency becomes

$$\omega_i(n_i) = \omega_i + K_i n_i, \tag{5.20}$$

where n_i is the average photon number in the mode with bare frequency ω_i.

The *cross-Kerr shift* quantifies the photon-number-dependent coupling between two modes, leading to a frequency shift of one subsystem depending on the occupation of the other[9]. This interaction can be described by

$$H_{\text{cross-Kerr}} = \hbar K_{cr} n_c n_r, \tag{5.21}$$

where n_c and n_r denote the photon-number operators of the cavity and readout resonator. While Eq. 5.12 does not explicitly include the resonator-related nonlinear terms K_r and K_{rc}, these contributions are essential for a complete Hamiltonian description of the system.

Experimentally, the self-Kerr coefficient can be extracted by measuring the cavity resonance frequency as a function of the average photon number, typically calibrated through the drive amplitude. Alternatively, the *out-and-back* protocol (see Sec. 5.2.3) offers a more sensitive method, capable of resolving frequency shifts on the order of a few hertz. The data presented in Fig. 5.16a are obtained following the same analysis procedure as in Fig. 5.11. Notably, the cavity self-Kerr coefficient changes sign at the avoided crossing with the qubit. Although its magnitude remains on the order of a few hertz across most

[9]The dispersive shifts χ_{QR} and χ_{QC} are often referred to as the cross-Kerr coefficients between the qubit and the resonator, and the qubit and the cavity, respectively.

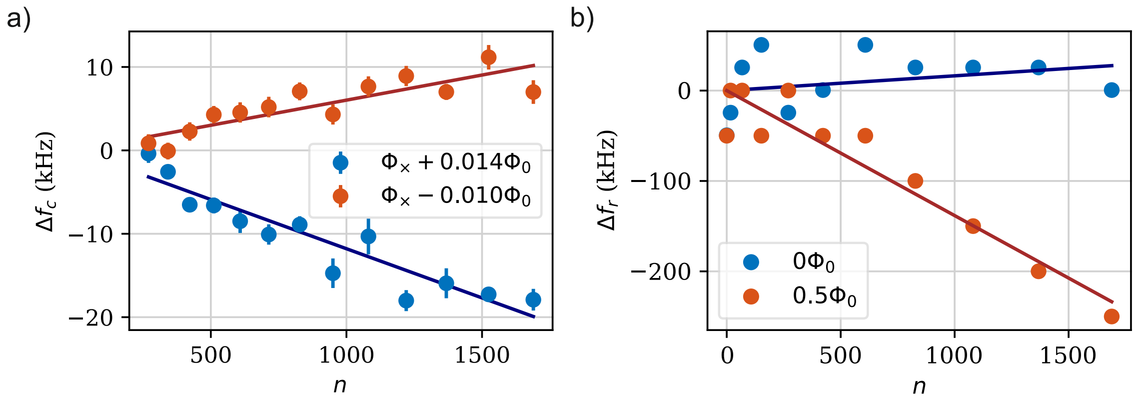

Figure 5.16: Kerr shifts. a) Cavity self-Kerr K_c is extracted with the help of out-and-back protocol. This is done for both above and below the bias point $\Phi_\times \approx -0.372\Phi_0$, where the cavity and the qubit cross. The Kerr shift changes its sign, i.e. $K_c/2\pi(\Phi_\times - 0.010\Phi_0) = -6 \pm 1\,\mathrm{Hz}$, and $K_c/2\pi(\Phi_\times + 0.014\Phi_0) = 3 \pm 1\,\mathrm{Hz}$. b) The cross-Kerr coefficient between the cavity and the resonator is measured via the shift of the resonance frequency of the latter. The extracted shifts for the low and high sweet spots are respectively $K_{rc}/2\pi(0.5\Phi_0) = -138 \pm 15\,\mathrm{Hz}$ and $K_{rc}/2\pi(0\Phi_0) = 16 \pm 18\,\mathrm{Hz}$. All data has been shifted by the offset $\Delta/2\pi$, extracted from the linear fit to Eq. 5.15.

of the flux bias range, it diverges near the avoided crossing[10], highlighting a regime where the enhanced Kerr nonlinearity could be harnessed for generating cat states. The self-Kerr of the readout resonator was not measured.

The cross-Kerr shift is experimentally measured by monitoring the readout resonator frequency shift after applying a displacement pulse of varying amplitude on the cavity. The frequencies are extracted by taking the minimum of the spectroscopy transitions, which is why no error bars are present on the plot. The results for both flux sweet spots are summarized in Fig. 5.16b. This concludes the characterization of the system, with the results summarized in Table A.1 from App. A.

[10]The cavity self-Kerr diverges also around $\Phi_{\mathrm{ext}} \approx -0.23\Phi_0$.

CHAPTER 6

Dissipative cat code with SNAIL in 3D cavities

The initial goal of this project was to realize a logical dissipative cat qubit. In retrospect, this objective was overly ambitious for the scope of a PhD project, though this became clear only during my research stay at LPNS in Paris. As my understanding of dissipation engineering deepened, so did my awareness of its practical limitations, and the prospect of achieving a 3D dissipative cat—or even a logical qubit—became increasingly remote. The following analysis ultimately led to the decision to abandon this direction.

To advance beyond the current state-of-the-art architectures, we considered a 3D seamless cavity resonantly coupled to a two-photon dissipation element. Several approaches to engineering such an interaction were evaluated, including the Asymmetrically Threaded SQUID[121] (ATS), the Superconducting Nonlinear Asymmetric Inductive eLement[193] (SNAIL), and the rf-SQUID[194]. Their respective advantages and drawbacks in the context of a 3D implementation are summarized in Table 6.1. An honorable mention goes to the Josephson Ring Modulator[195] (JRM), which also supports three-wave mixing but involves three modes. Even when tuning the memory and buffer to two of these modes, the Wheatstone-bridge geometry cancels mixing terms of the form $\varphi_X\varphi_Y^2$ (or equivalently $\varphi_b\varphi_m^2$)[1], which are essential for implementing two-photon dissipation.

The ATS represents the most widely used method for realizing two-photon dissipation. While most implementations rely on 2D resonators as memory modes, Ref. [150] reports the first integration of an ATS within a 3D cavity. In that setup, flux bias lines wire-bonded to a printed circuit board (PCB) are employed to tune the device, and an on-chip Purcell filter is added to prevent memory leakage through the buffer. As summarized in Table 6.1, the main challenge in this design is achieving differential biasing. To circumvent this, we propose using two offset flux hoses in place of the bias lines, together with a modular Purcell filter[153] instead of the on-chip version.

The ATS creates a two-photon exchange with rate $g_2 = \frac{E_J}{2}\epsilon_p\varphi_m^2\varphi_b$, where E_J is the Josephson energy of the junctions and ϵ_p is the amplitude of the pump tone, used to match the frequencies, i.e., $\omega_p = |2\omega_m - \omega_b|$. A typical parameter regime for a 3D implementation is $\epsilon_p < 1$, $\varphi_m \approx 0.01$ and $\varphi_b \approx 0.1$.

[1]Here, φ_i, short for $\varphi_{zpf,i}$, denotes the phase zero-point fluctuations of mode i; X, Y are two of the modes that JRM supports, while the indices b, m correspond to buffer, memory.

Table 6.1: Advantages and disadvantages of different two-photon dissipation elements in a 3D architecture.

Element	Advantages	Disadvantages	Biasing
ATS	1. flux bias for desired operation at the sweet spot 2. high g_2	differential flux bias in 3D	1. PCB integrated[150] 2. two offset flux hoses
rf-SQUID	1. easy to implement in 3D 2. high g_2	flux bias for desired operation not at the sweet spot	flux hose
SNAIL	easy to implement in 3D	flux bias for desired operation 1. not at the sweet spot 2. coincides with $\omega_b = 2\omega_m$	flux hose

An alternative approach employs an rf-SQUID as the mixing element. Its main advantage lies in its simplicity of fabrication, requiring only a single Josephson junction shunted by a superconducting loop. This circuit naturally exhibits a Kerr-free point at $\varphi_{\text{ext}} = \pi/2$. In this configuration, the two-photon exchange rate is given by $g_2 = \frac{E_J}{6}\varphi_m^2\varphi_b$, which, for comparable junction parameters, can reach values similar to those of the ATS (depending on ϵ_p). However, this design also presents drawbacks: the frequency exhibits a steep gradient at the Kerr-free bias point, and to avoid hysteresis, the junction inductance must exceed that of the loop, placing constraints on the Josephson energy. Furthermore, the loop inductance cannot be reduced arbitrarily, as it is limited by practical considerations on the strength of the applied magnetic field.

The final candidate for a dissipation element is the SNAIL. While it circumvents the biasing difficulty, it possesses other disadvantages compared to alternative wave-mixing elements. Those will be discussed in the following subsection. However, at the time, it seemed like the most straightforward option for engineering the two-photon dissipation of our memory mode.

6.1 Diagonalization

The SNAIL consists of a superconducting loop interrupted by several Josephson junctions arranged asymmetrically: typically one small junction in parallel with n larger ones. The Hamiltonian of a SNAIL embedded in an LC resonator (see the left part of the circuit in Fig. 6.1) reads

$$H_{SNAIL} = 4E_C n^2 + \frac{E_L}{2}(\varphi - \varphi_\alpha)^2 + U_S(\varphi_\alpha). \tag{6.1}$$

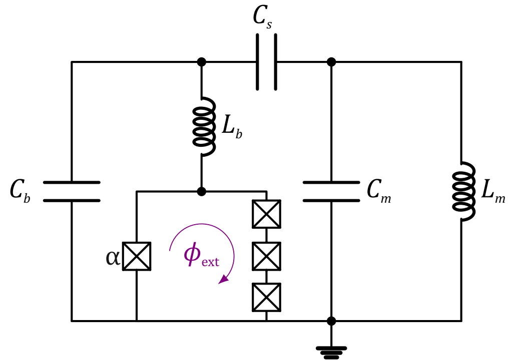

Figure 6.1: Lumped element circuit representation of a memory and a buffer mode with a SNAIL. For the sake of visual clarity, the junction notations are missing, except for α, denoting the junction with energy αE_J.

where we have introduced φ as the total phase across the capacitive portion of the circuit. The non-linear potential energy is

$$U_S = -\alpha E_J \cos\varphi_\alpha - nE_J \cos\left(\frac{\varphi_{\text{ext}} - \varphi_\alpha}{n}\right) \tag{6.2}$$

where φ_α is the superconducting phase across the small junction, α is the junction asymmetry ratio, and φ_{ext} is the reduced external flux threading the loop. The SNAIL has the following current-phase relation

$$I_S(\varphi_s) = \frac{1}{\phi_0}\frac{\partial U_S}{\partial \varphi_s} = \frac{E_J}{\phi_0}\left[\alpha \sin(\varphi_s) + \sin\left(\frac{\varphi_s - \varphi_{\text{ext}}}{n}\right)\right]. \tag{6.3}$$

The asymmetry between the junctions breaks the inversion symmetry of the current-phase relation, enabling odd-order nonlinearities (e.g., three-wave mixing) that are forbidden in symmetric circuits like the SQUID. In the limit of small phase fluctuations, the potential U_S can be Taylor expanded around its minimum, which is determined by the condition that no DC current flows across the SNAIL, i.e., $I_S(\varphi_{\alpha,min}) = 0$. We get the effective potential in terms of $\tilde{\varphi}_\alpha = \varphi_\alpha - \varphi_{\alpha,min}$

$$U_{S,eff}(\tilde{\varphi}_s)/E_J = \frac{c_2}{2!}\tilde{\varphi}_\alpha^2 + \frac{c_3}{3!}\tilde{\varphi}_\alpha^3 + \frac{c_4}{4!}\tilde{\varphi}_\alpha^4 + ... \tag{6.4}$$

The coefficients c_3 and c_4 are flux-tunable and correspond to the three- and four-wave mixing capabilities of the SNAIL.

For a memory mode, capacitively coupled to a buffer mode with embedded SNAIL (see circuit in Fig. 6.1, the total Lagrangian is

$$\begin{aligned}\mathcal{L} = \frac{C_m}{2}\dot{\Phi}_m^2 - \frac{1}{2L_m}\Phi_m^2 + \frac{C_s}{2}(\dot{\Phi}_m - \dot{\Phi}_b)^2 - \frac{1}{2L_b}(\Phi_\alpha - \Phi_b)^2 + \frac{C_b}{2}\dot{\Phi}_b^2 \\ + \frac{\alpha C_J}{2}\dot{\Phi}_\alpha + \frac{C_J}{2n}\left(\frac{\dot{\Phi}_\alpha}{n}\right)^2 + \alpha E_J \cos\varphi_\alpha + nE_J\cos\left(\frac{\varphi_{\text{ext}} - \varphi_\alpha}{n}\right)\end{aligned} \tag{6.5}$$

Note that, due to the added junction capacitance, the SNAIL possesses a mode with an independent phase variable φ_α, otherwise $\varphi_\alpha = \varphi_\alpha(\varphi_b)$. Rewriting the Lagrangian to group all flux variables, we get

$$\begin{aligned}\mathcal{L} = &\frac{C_b + C_s}{2}\dot{\Phi}_b^2 + \frac{C_m + C_s}{2}\dot{\Phi}_m^2 + \frac{C_{J,eff}}{2}\dot{\Phi}_\alpha^2 - C_s\dot{\Phi}_b\dot{\Phi}_m \\ &- \frac{1}{2L_b}\Phi_b^2 - \frac{1}{2L_m}\Phi_m^2 - \frac{1}{2L_b}\Phi_\alpha^2 + \frac{1}{L_b}\Phi_\alpha\Phi_b \\ &+ \alpha E_J \cos\varphi_\alpha + nE_J\cos\left(\frac{\varphi_{\text{ext}} - \varphi_\alpha}{n}\right)\end{aligned} \tag{6.6}$$

where we have introduced an effective junction capacitance $C_{J,eff} = C_J\left(\alpha + \frac{1}{n^3}\right)$.

We use Legendre transform to define the conjugate charges $Q_i = \partial\mathcal{L}/\partial\dot{\Phi}_i$

$$(C_b + C_s)\dot{\Phi}_b - C_s\dot{\Phi}_m = Q_b \tag{6.7a}$$

$$(C_m + C_s)\dot{\Phi}_m - C_s\dot{\Phi}_b = Q_m \tag{6.7b}$$

$$C_{J,eff}\dot{\Phi}_\alpha = Q_\alpha. \tag{6.7c}$$

We then calculate the Hamiltonian using $H = \sum_i \dot{\Phi}\frac{\partial\mathcal{L}}{\partial\dot{\Phi}_i} - \mathcal{L}$

$$\begin{aligned}H = &\frac{C_b + C_s}{2}\left(\frac{(C_m + C_s)Q_b + C_sQ_m}{C_{norm}}\right)^2 + \frac{C_m + C_s}{2}\left(\frac{(C_b + C_s)Q_m + C_sQ_b}{C_{norm}}\right)^2 \\ &+ \frac{1}{2C_{J,eff}}Q_J + \frac{C_s}{C_{norm}^2}\left((C_m + C_s)Q_b + C_sQ_m\right)\left((C_b + C_s)Q_m + C_sQ_b\right)\right) + U(E_J)\end{aligned} \tag{6.8}$$

with the normalization factor defined as $C_{norm} = C_mC_b + C_sC_b + C_sC_m$.

We introduce renormalized capacitances for each mode

$$C_{b,n}^{-1} = ((C_b + C_s)(C_m + C_s)^2 - C_s^2(C_m + C_s))/C_{norm}^2 \tag{6.9a}$$

$$C_{s,n}^{-1} = ((C_m + C_s)(C_b + C_s)C_s - C_s^3)/C_{norm}^2 \tag{6.9b}$$

$$C_{m,n}^{-1} = ((C_m + C_s)(C_b + C_s)^2 - C_s^2(C_b + C_s))/C_{norm}^2 \tag{6.9c}$$

and rewrite the Hamiltonian as

$$H = \frac{Q_b^2}{2C_{b,n}} + \frac{Q_m^2}{2C_{m,n}} + \frac{Q_bq_m}{C_{s,n}} + \frac{Q_J^2}{2C_{J,eff}} + \frac{1}{2L_b}\Phi_b^2 + \frac{1}{2L_m}\Phi_m^2 + \frac{1}{2L_b}\Phi_\alpha^2 - \frac{1}{L_b}\Phi_b\Phi_\alpha + U(E_J). \tag{6.10}$$

Similarly to the fluxonium, we use the Fock basis $\{|bma\rangle\}$ for numerically diagonalizing this circuit, where b, m and a represent the number of excitations in the buffer-like, memory-like, and SNAIL-like modes, respectively.

Defining the capacitive $E_{C,i} = \frac{e^2}{2C_{i,n}}$ and inductive energies $E_{L,i} = \frac{\phi_0^2}{L_i}$ for each mode i, we get the following expression

$$\begin{aligned} H = & 4E_{C,b}\langle bma|n_b^2|b'm'a'\rangle + 4E_{C,m}\langle bma|n_m^2|b'm'a'\rangle + 8E_{C,s}\langle bma|n_m n_b|b'm'a'\rangle \\ & + 4E_{C,\alpha}\langle bma|n_\alpha^2|b'm'a'\rangle + \frac{E_{L,b}}{2}\langle bma|\varphi_b^2|b'm'a'\rangle + \frac{E_{L,b}}{2}\langle bma|\varphi_\alpha^2|b'm'a'\rangle \, . \\ & - E_{L,b}\langle bma|\varphi_b\varphi_\alpha|b'm'a'\rangle + \frac{E_{L,m}}{2}\langle bma|\varphi_m^2|b'm'a'\rangle + \langle U(E_J)\rangle \end{aligned} \tag{6.11}$$

The dressed frequencies of each element are $\omega_i = \sqrt{8E_{C,i}E_{L,i}}$, and the zero-point fluctuations are rescaled with the new energies, i.e., $\varphi_{zpf,i} = \sqrt[4]{\frac{2E_{C,i}}{E_{L,i}}}$ for the flux and $n_{zpf,i} = \sqrt[4]{\frac{E_{L,i}}{32E_{C,i}}}$ for the charge number variable. Moreover, we can introduce the coupling rates between the buffer and the SNAIL modes $g_i = E_{L,b}\varphi_{zpf,b}\varphi_{zpf,\alpha}$, as well as between the buffer and the memory modes $g_c = 8E_{C,s}n_{zpf,b}n_{zpf,m}$. The final Hamiltonian that goes into the numerical simulation is the following

$$\begin{aligned} H = \hbar & \sum_{b,m,a\in\mathbb{N}} (\omega_b b + \omega_m m + \omega_\alpha a)|bma\rangle\langle bma| \\ & - \hbar g_c \sum_{b,m,a,b',m'\in\mathbb{N}} \langle bma|mb + m^\dagger b^\dagger - m^\dagger b - mb^\dagger|b'm'a\rangle \\ & - \hbar g_i \sum_{b,m,a,b',a'\in\mathbb{N}} \langle bma|ab + a^\dagger b^\dagger + a^\dagger b + ab^\dagger|bm'a'\rangle \\ & -\alpha E_J \sum_{b,m,a,a'\in\mathbb{N}} c_{aa'}(\varphi_\alpha)|bma\rangle\langle bma'| \\ & -nE_J \Bigg[\cos(\varphi_{\text{ext}}) \sum_{b,m,a,a'\in\mathbb{N}} c_{aa'}\left(\frac{\varphi_\alpha}{n}\right)|bma\rangle\langle bma'| \\ & \quad - \sin(\varphi_{\text{ext}}) \sum_{b,m,a,a'\in\mathbb{N}} s_{aa'}\left(\frac{\varphi_\alpha}{n}\right)|bma\rangle\langle bma'|\Bigg] \end{aligned} \tag{6.12}$$

where the matrix elements have been defined in Eq. 3.9.[2] In this Hamiltonian, all terms contribute to the diagonal elements, except for the ones in the second and third lines of the equation. The frequencies of the dressed modes, together with their tunability, are derived by numerically diagonalizing Eq. 6.12.

[2]For ease of reading, I repeat Eq. 3.9 here

$$\begin{aligned} c_{kl}(\varphi) &= \langle k|\cos(\varphi)|l\rangle \\ &= \begin{cases} (-1)^{\frac{l-k}{2}}\sqrt{\frac{k!}{l!}}(\varphi_{zpf})^{l-k}e^{-(\varphi_{zpf})^2/2}\mathscr{L}_k^{l-k}\left[(\varphi_{zpf})^2\right] & \text{for } k+l \text{ even} \\ 0 & \text{for } k+l \text{ odd} \end{cases} \end{aligned} \tag{6.13a}$$

$$\begin{aligned} s_{kl}(\varphi) &= \langle k|\sin(\varphi)|l\rangle \\ &= \begin{cases} 0 & \text{for } k+l \text{ even} \\ (-1)^{\frac{l-k+1}{2}}\sqrt{\frac{k!}{l!}}(\varphi_{zpf})^{l-k}e^{-(\varphi_{zpf})^2/2}\mathscr{L}_k^{l-k}\left[(\varphi_{zpf})^2\right] & \text{for } k+l \text{ odd.} \end{cases} \end{aligned} \tag{6.13b}$$

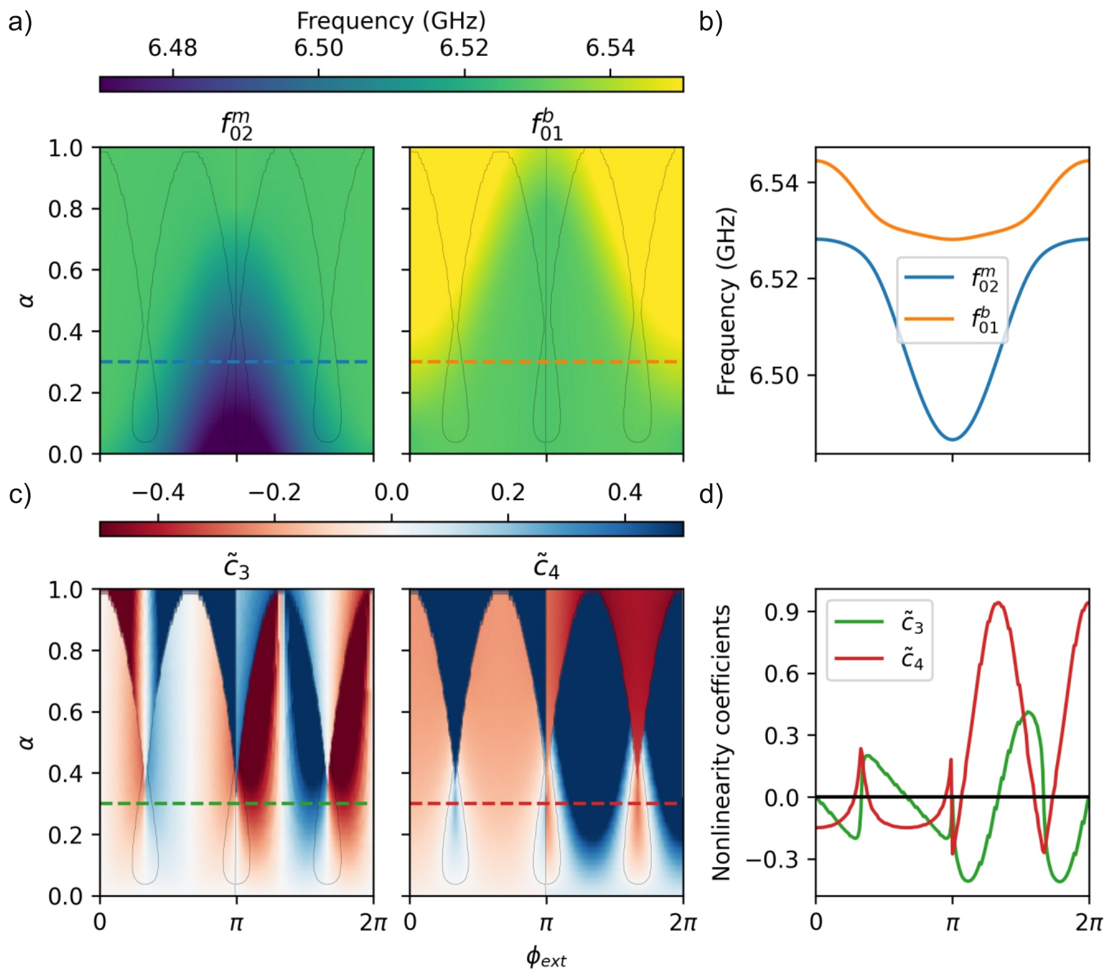

Figure 6.2: Dressed frequencies and nonlinearity coefficients of a SNAIL, embedded in a buffer mode, coupled to a memory mode. a) Frequencies of the two-photon memory transition f_{02}^m and the single photon buffer transition f_{01}^b as a function of both the external bias and the junction imbalance α. b) Line cuts from a) for $\alpha = 0.3$. c) Nonlinear coefficients for three- and four-wave mixing, $\tilde{c}_3$ and $\tilde{c}_4$, respectively. d) Line cuts from c) for $\alpha = 0.3$. Parameters: $E_J = 20\,\text{GHz}$, $Z_b = 226$, $Z_m = 66$, $f_m = 4.484\,\text{GHz}$ $f_b = 7.833\,\text{GHz}$, $\alpha = 0.3$, $n = 3$, $C_{mb} = 108\,fF$, $C_J = 10\,fF$. The parameters are picked in a way that the zero crossing of $\tilde{c}_4$ is at the same flux bias as the avoided crossing between f_{01}^b and f_{02}^m, i.e. $\phi_{\text{ext}} = 0.38\pi$. The thin black line in a) and c) shows where the coefficients $\tilde{c}_3$ and $\tilde{c}_4$ change sign.

Figure 6.2 presents the results of the diagonalization. The first row shows the dressed frequencies corresponding to the two-photon transition of the memory mode and the one-photon transition of the buffer mode. As seen in Fig. 6.2b, these modes exhibit an avoided crossing around $\phi_{\text{ext}} = 0.38\pi$, indicating a resonant two-to-one-photon exchange. At this same bias point, the four-wave mixing coefficient vanishes, resulting in a purely three-wave mixing process—that is, a Kerr-free interaction. One of the main challenges lies in designing and fabricating such a dissipative element with sufficient precision to achieve this delicate balance required for a resonant, Kerr-free operation.

6.2 Dissipation engineering

Once we have diagonalized the Hamiltonian from Eq. 6.12, we can express it in the second quantization notation

$$H/\hbar = \tilde{\omega}_m m^\dagger m + \tilde{\omega}_b b^\dagger b + \tilde{\omega}_\alpha a^\dagger a + U(E_J) \tag{6.14}$$

where $\tilde{\omega}_i$ are Lamb-shifted frequencies of each mode.

In the rotating frame $U = e^{i\omega_d b^\dagger b} e^{i\frac{\omega_d}{2} m^\dagger m} e^{i\tilde{\omega}_\alpha a^\dagger a}$, where ω_d is the drive frequency, the Hamiltonian is transformed to

$$H_{rot}/\hbar = \left(\tilde{\omega}_m - \frac{\omega_d}{2}\right) m^\dagger m + (\tilde{\omega}_b - \omega_d) b^\dagger b - \alpha E_J \cos\varphi + 3E_J \cos\left(\frac{\varphi_{\text{ext}} - \varphi}{3}\right) \tag{6.15}$$

where the phase operators take the form $\varphi = \sum_{i=\{b,m,a\}} \varphi_i(\tilde{i}^\dagger + \tilde{i})$,[3] with the newly defined rotated annihilation operators $\tilde{a} = e^{-i\tilde{\omega}_\alpha t} a$, $\tilde{b} = e^{-i\omega_d t} b$, and $\tilde{m} = e^{-i\omega_d t/2} m$.
For $n = 3$ and low zero-point fluctuations, we can expand the cosine terms in the potential

$$U(E_J) \approx -\alpha E_J \left(1 - \frac{\varphi^2}{2!} + \frac{\varphi^4}{4!}\right) - 3E_J \left[\cos\varphi_{\text{ext}} \left(1 - \frac{\varphi^2}{3^2 2!} + \frac{\varphi^4}{3^4 4!}\right) + \sin\varphi_{\text{ext}} \left(\frac{\varphi}{3} - \frac{\varphi^3}{3^3 3!}\right)\right]. \tag{6.16}$$

By replacing the operators with their expressions as a function of the rotated annihilation and creation operators, we get a plethora of interactions of different orders. After Rotating-Wave Approximation (RWA), the terms $\tilde{i}^\dagger \tilde{i}$ would shift the respective mode frequency, the terms $\tilde{i}^\dagger \tilde{j}$ ($\tilde{i}\tilde{j}^\dagger$) would alter the coupling between the mode with creation (annihilation) operator $\tilde{i}^\dagger$ ($\tilde{i}$) and the mode with annihilation (creation) operator $\tilde{j}$ ($\tilde{j}^\dagger$). Third-order terms correspond to three-wave mixing processes. In particular, what we are interested in is the mixing terms converting two photons in the memory to one photon in the buffer mode, i.e., terms of the type $\varphi_m^2 \varphi_b$

$$H_{3WM}/\hbar = \frac{1}{54} E_J \sin\varphi_{\text{ext}} \varphi_m^2 \varphi_b ((\tilde{m}^\dagger)^2 \tilde{b} + \tilde{m}^2 \tilde{b}^\dagger) = g_2 (\tilde{m}^\dagger)^2 \tilde{b} + g^* \tilde{m}^2 \tilde{b}^\dagger. \tag{6.17}$$

Upon driving the buffer on resonance $\omega_d = \omega_b = 2\omega_m$, photons are injected into the system and immediately converted into pairs of memory photons. The process is described by the Hamiltonian

$$H_{3WM}/\hbar = g_2((\tilde{m}^\dagger)^2 \tilde{b} + \tilde{m}^2 \tilde{b}^\dagger) + \epsilon_d \tilde{b}^\dagger + \epsilon_d^* \tilde{b} = g_2 \left((\tilde{m}^\dagger)^2 + \frac{\epsilon_d^*}{g_2}\right) \tilde{b} + g_2^* \left(\tilde{m}^2 + \frac{\epsilon_d}{g_2^*}\right) \tilde{b}^\dagger. \tag{6.18}$$

After adiabatically eliminating the buffer mode, the memory dynamics are described by the dissipator from Eq. 2.44, which stabilizes the manifold $\{|\alpha\rangle, |-\alpha\rangle\}$ with $\alpha = \sqrt{\epsilon_d / g_2^*}$.

Any other mixing process left after the RWA is parasitic. The most pronounced parasitic process is the four-wave mixing, governed by the Hamiltonian

$$\begin{aligned} H_{4WM}/\hbar = &-\frac{K_\alpha}{2} (\tilde{a}^\dagger)^2 \tilde{a}^2 - \frac{K_m}{2} (\tilde{m}^\dagger)^2 \tilde{m}^2 - \frac{K_b}{2} (\tilde{b}^\dagger)^2 \tilde{b}^2 \\ &- K_{\alpha b} \tilde{a}^\dagger \tilde{a} \tilde{b}^\dagger \tilde{b} - K_{mb} \tilde{m}^\dagger \tilde{m} \tilde{b}^\dagger \tilde{b} - K_{\alpha m} \tilde{a}^\dagger \tilde{a} \tilde{m}^\dagger \tilde{m}. \end{aligned} \tag{6.19}$$

[3]Here, we use the shortened notation $\varphi_i \equiv \varphi_{zpf,i}$.

Here, we have introduced both self- and cross-Kerr coefficients, i.e,

$$K_i = E_J \varphi_i^4 \left(\frac{\alpha}{2} + \frac{1}{54} \cos \varphi_{\text{ext}} \right) \tag{6.20}$$

and

$$K_{ij} = E_J \varphi_i^2 \varphi_j^2 \left(\frac{\alpha}{6} + \frac{1}{162} \cos \varphi_{\text{ext}} \right), \tag{6.21}$$

respectively, which can be canceled by tuning φ_{ext} such that the expression in the brackets for either K_i or K_{ij} nulls.

6.3 Discussion

One of the most important figures of merit for a dissipative cat is the ratio between the two-photon dissipation and the single-photon loss rates. This ratio determines the protection of the dissipative cat qubit against phase flips. Therefore, a ratio of at least $\kappa_2/\kappa_1 > 10^2$ is desirable for the operation. So far, the highest ratio reported is in Ref. [122], which implements a resonant two-photon exchange, similar to the concept described here.

To boost this ratio, one can take two different paths: decreasing the single-photon loss of the memory or increasing the two-photon dissipation rate. The former can be achieved by using 3D seamless cavities instead of 2D resonators, as they generally have higher lifetimes. However, the zero-point fluctuations reduce significantly as these structures have much higher self-capacitance. This, in turn, harms the two-photon exchange g_2. As a reminder, for adiabatic elimination we require that $|g_2| \ll \kappa_b$, and the two-photon dissipation rate is $\kappa_2 = 4|g_2|/\kappa_b$. A delicate balance between these parameters needs to be achieved, as simply increasing κ_b would decrease the two-photon dissipation rate and compromise the cat qubit operation.

In 3D cavities, the zero-point phase fluctuations φ_{zpf} are typically on the order of ~ 0.01, and since they enter the two-photon exchange rate quadratically, their impact is further amplified. These fluctuations can be increased by enhancing E_J, g_i, or g_c. However, it should be noted that modifying E_J also strongly affects the mode frequencies, which complicates the task of matching the Kerr-free bias point to the avoided crossing. Considering state-of-the-art experiments, such as the ones from Ref. [122] and [24], and the time it takes to engineer such interaction by a one-man team, after my stay in LPNS, a decision was taken: the project will not proceed in this direction.

CHAPTER 7

Fast flux implementation

The ultimate objective of this work is to utilize the system for generating quantum states. As discussed earlier, the dissipative approach to realizing cat states is no longer considered feasible within the current experimental constraints. Alternative strategies were therefore explored, including the coherent exchange of excitations between the fluxonium and the cavity for Fock state preparation, the generation of resonant cat states[196], the application of variational or machine-learning-based optimization techniques for arbitrary state synthesis[57, 183], and the exploitation of the self-Kerr nonlinearity to engineer multi-legged cat states[192]. For the majority of these protocols, the state preparation and subsequent tomography are performed at distinct flux bias points, thereby requiring fast flux control of the device. Precise control of the magnetic flux threading a superconducting qubit loop is essential for achieving high-fidelity gate operations, especially when fast flux pulses are used to dynamically tune the couplings.

Implementing fast flux control in superconducting circuits presents several experimental challenges. The primary difficulty lies in balancing the need for high bandwidth with the requirement for low noise. Flux lines must be heavily filtered and attenuated to suppress thermal and flux noise, but these components inherently limit the achievable pulse speed. Moreover, fast flux pulses can suffer from distortions and dispersion due to impedance mismatches, reflections, bandwidth limitations, and filtering effects along the wiring, necessitating precise calibration and pulse pre-compensation. Overcoming these challenges requires careful line design, shielding, and calibration to enable high-fidelity, fast flux operations without compromising qubit performance.

This chapter outlines the strategies developed for fast flux tuning and pulse pre-compensation. It explores the possible approaches for generating non-classical states within the system and discusses the associated experimental challenges. Finally, the chapter presents some efforts towards the experimental realization of a Kerr cat state.

7.1 Filtering and precompensation

Flux control waveforms generated by an arbitrary waveform generator (AWG) are subject to distortions along the flux control line. These distortions arise from various sources, including the finite bandwidth of the AWG, the wiring and components between the room-temperature electronics and the cryogenic device, as well as the device packaging. Filters placed along the signal path each introduce their own time constant, which is determined by their design parameters, such as type and cutoff frequency. As a result, compensating for these distortions requires addressing effects across a wide range of time scales, from nanoseconds to tens of microseconds.

The step response of the full signal path contains all the information to model the response to an arbitrary pulse[1]. From that information, we can extract the time constants of the path response and apply digital filters at the AWG level that adjust the pulse in a way that compensates for the distortions. Depending on the time scales and the response, different filters can be used. All digital filter information is summarized in Table 7.1.

Table 7.1: Types of digital filters[197] applied for correcting a distorted signal. The first three filters are of the type Infinite Impulse Response (IIR) filters, which respond indefinitely, while the last is a Finite Impulse Response (FIR) filter, which has a finite duration. Here, τ is the filter time constant, $k(t)$ is the filter kernel, A is the filter amplitude, $*$ denotes convolution, and $x(t)$ is the input signal. This table is rewritten from the Zurich Instruments (ZI) HDAWG manual[198].

Compensation	Model response $y(t)$	Applications	Time scales
High-pass	$e^{-t/\tau}$	DC-block, Bias-tee	few hundred ps to hundreds of ms
Exponential	$1 + A_1 e^{-t/\tau}$	LCR elements, i.e. low-pass filters	few ns to a few ms
Bounce	$x(t) + A_2 y(t - \tau)$	reflections due to impedance mismatch	up to a hundred ns
FIR	$k(t) * x(t)$	short timescale distortions	few tens of ns

While the room-temperature part of the signal path can be directly characterized by measuring the response to a square pulse with an oscilloscope, the section inside the cryostat remains inaccessible. Due to the components on the line, the signal behaves differently under cryogenic conditions, even if measured at room temperature. This necessitates specialized methods to characterize and correct the full response of the flux control line. In the following section, I discuss the most common strategies for measuring and compensating distortions in flux bias lines.

7.1.1 Cryoscope

The Cryoscope technique, developed in Ref. [199], uses the qubit itself as a fast, high-sensitivity flux-to-frequency transducer. The principle relies on a modified Ramsey interferometry sequence, schematically shown in Fig. 7.1. In this protocol, a flux pulse of controlled amplitude and duration t_p is applied to the qubit between two $\pi/2$ pulses. The qubit accumulates a phase $\phi(t_p)$ that depends on the instantaneous qubit frequency $\omega_q(t)$,

[1]Ideally, the impulse response to a Dirac delta pulse should be measured, but such a pulse is inaccessible with an AWG.

which in turn depends on the applied flux $\Phi(t)$. The accumulated phase can be expressed as

$$\phi(t_p) = \int_0^{t_p} \delta\omega_q(\Phi(t))\, dt, \tag{7.1}$$

where $\delta\omega_q(\Phi)$ is the deviation of the qubit frequency from its idle value. By measuring the expectation values $\langle\sigma_x\rangle$ and $\langle\sigma_y\rangle$ as a function of t_p (see Fig. 7.1b and c), one can extract $\phi(t_p)$ as shown in Fig. 7.1d and thus reconstruct $\delta\omega_q(t) = \frac{\phi_{t+\delta t}-\phi_t}{2\pi\delta t}$, providing a direct measurement of the flux waveform experienced by the qubit. By comparing the measured

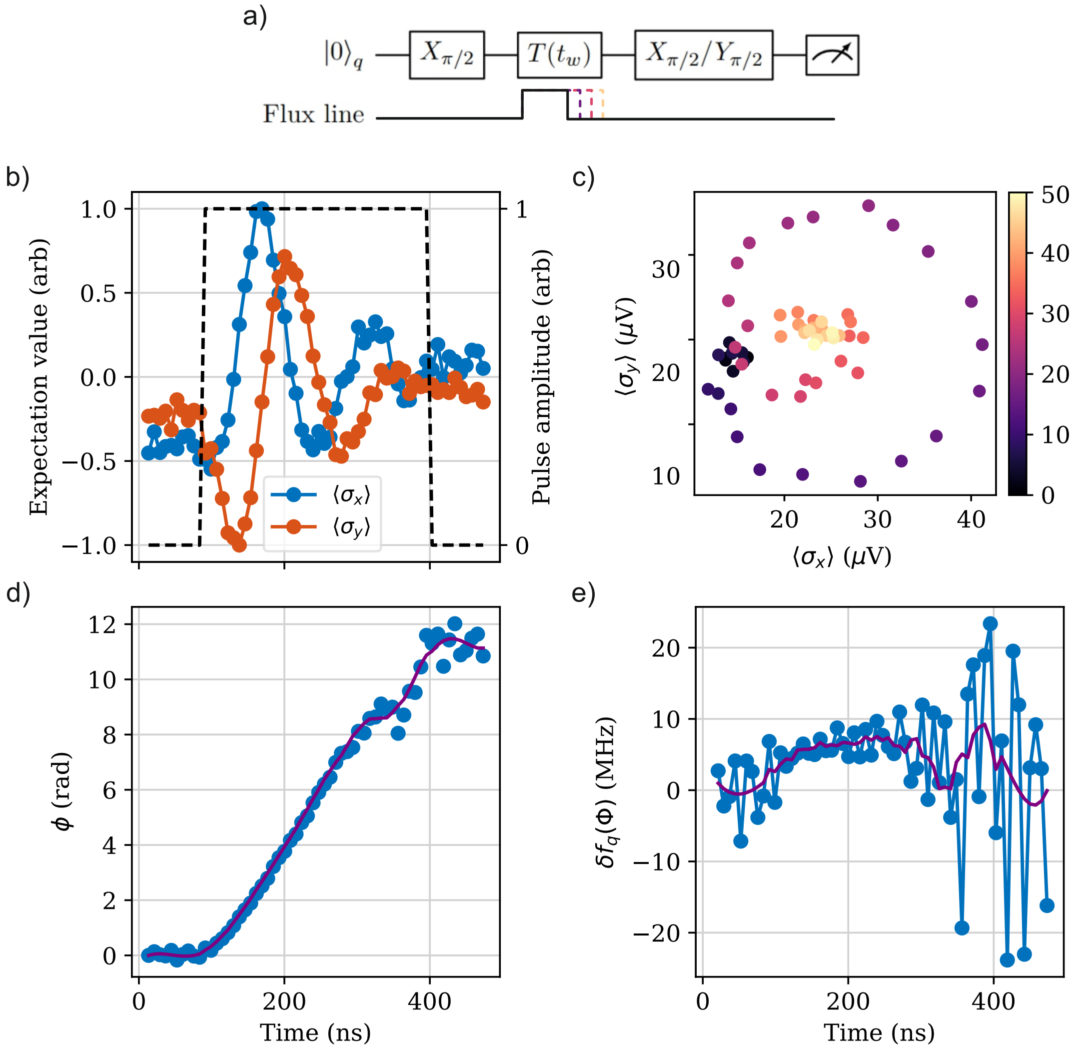

Figure 7.1: Cryoscope. a) Sequence. The flux pulse duration $t_p \in [0, 340]$ ns is always smaller than the fixed time $t_w = 528$ ns between the two pulses and starts 67 ns after the start of $T(t_w)$. b) Measured $\langle\sigma_x\rangle$ and $\langle\sigma_y\rangle$ as a function of time, with the longest flux pulse amplitude (right y-axis) plotted in black. c) Qubit trajectory in phase space as a function of time, indicated by the colorbar. d) Accumulated phase. e) Extracted change in frequency due to the applied flux pulse. In both d) and e), the data is plotted in blue, while the solid purple line is the data smoothed by applying a fourth-order Savitzky-Golay filter with a window size of 20 samples.

frequency response with the target waveform, one can extract the impulse response function (or equivalently, the transfer function) of the flux line. Once the line's transfer function is known, pre-compensation can be applied to future control pulses–typically by filtering the waveform through the inverse transfer function–so that the flux at the qubit follows the desired temporal shape with high accuracy.

The Cryoscope technique enables precise calibration of flux control lines over a wide dynamic range and across multiple time scales, without requiring direct access to the cryogenic circuitry. However, the accessible time window is ultimately limited by the qubit's Ramsey coherence time T_2^R. In this setup, T_2^R was observed to decrease from approximately 10 μs to 2 μs after replacing the DC wiring with a coaxial line. The formation of ground loops most likely causes this degradation. Since the coaxial cable is directly connected to the refrigerator without a DC block, ground noise is coupled into the flux line, effectively translating into additional flux noise at the qubit. Moreover, the accuracy of this method is directly correlated to the amplitude of the pulse: the bigger the displacement, the better the contrast, but also the higher the sensitivity to flux noise. Due to all these constraints, the extracted change in frequency has an insufficient signal-to-noise ratio even after smoothing the data with a fourth-order Savitzky-Golay filter (see Fig. 7.1e). Therefore, alternative methods must be used to calibrate the flux control line.

7.1.2 Spectroscopy

A less accurate method involves measuring the spectroscopy of the qubit as the flux pulse is applied. This method depends on the qubit lifetime T_1 instead of the Ramsey decoherence time T_2^R. This method is used in this manuscript due to the limitations of the qubit decoherence time. The protocol is presented in Fig. 7.2a. A $\pi-$pulse is played after some varying delay $t_{w,1}$. For the readout contrast to be consistent, the time after the $\pi-$pulse $t_{w,2}$ is varied such that the sum $t_{w,1} + t_{w,2}$ remains constant. For an uncorrected pulse, the response corresponds to a low-pass filter (see Fig. 7.2b top left), which is consistent with other flux hose devices of the type[151]. The frequency is extracted by fitting each resonance to a Lorentzian curve. Due to the non-linear voltage-frequency relation, the frequencies need to be converted to bias currents and then fitted to the filter responses from Table 7.1 and corrected for.

Using a set of five exponential filters, we compensate for most of the distortions arising from the low-pass response of the flux control line, as illustrated in Fig. 7.2b (bottom left). However, the accuracy of this correction is inherently limited by the $\pi-$pulse duration. Achieving compensation on shorter time scales requires the application of shorter π-pulses. The minimum achievable pulse length is constrained by the sampling rate of the arbitrary waveform generator: in our setup, the Quantum Machines hardware allows for a minimum pulse duration of 16 ns, while the Zurich Instruments AWG enables pulses as short as 13.33 ns. Such short pulses exhibit a broad spectral width, which in turn broadens the qubit resonance, making the precise determination of its frequency more challenging (see right column of Fig. 7.2b). A zoomed-in view of the response in the bottom left panel is shown in the top right panel of Fig. 7.2b. By including four additional filters, we observe

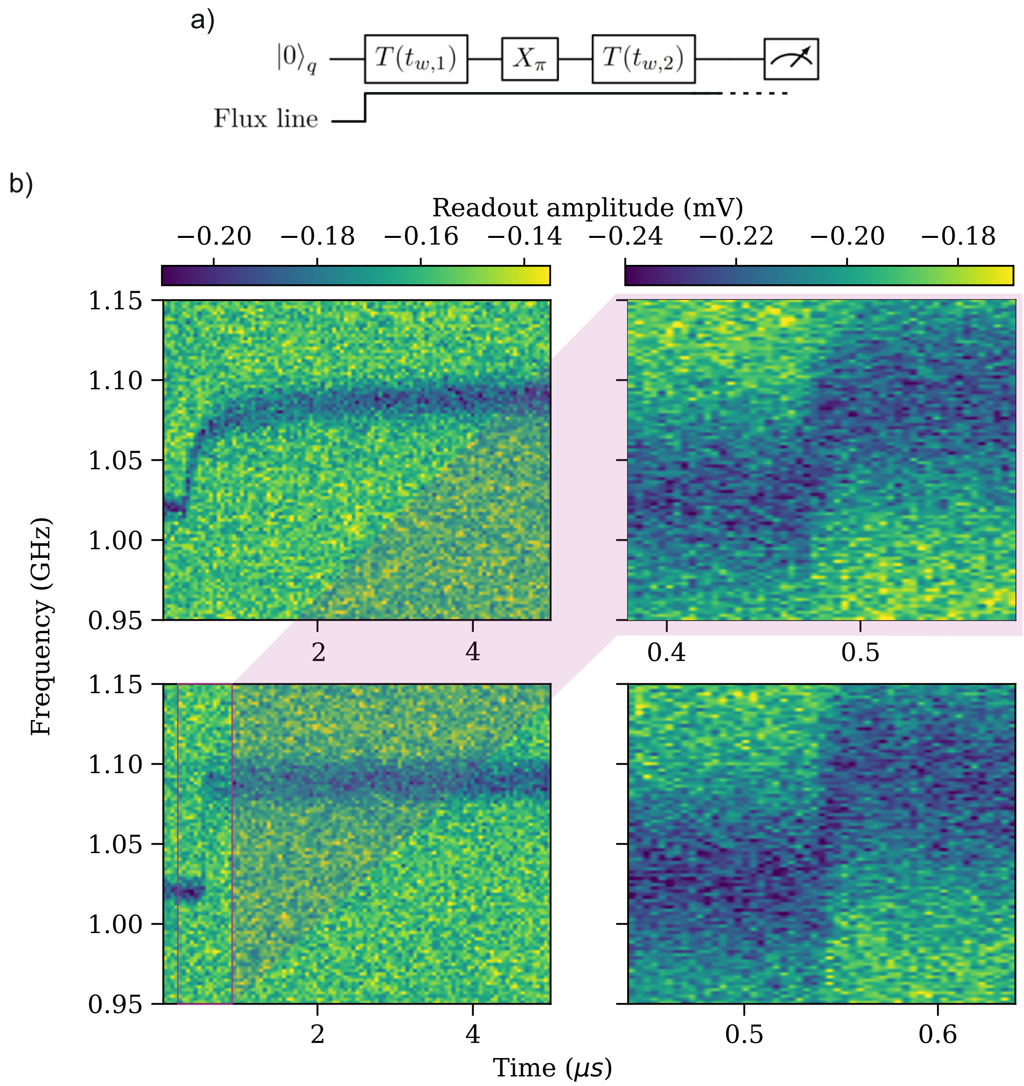

Figure 7.2: Flux line calibration via spectroscopy. a) Pulse sequence. The sum $t_{w,1}+t_{w,2}$ remains constant while both wait times are varied. b) Compensation. The π-pulse length for the left column is a Gaussian with $\sigma = 25.33\,\text{ns}$, and for the right column with $\sigma = 2.67\,\text{ns}$. The top left panel is the raw response to a square flux pulse. The bottom left panel is corrected with five exponential filters with the following parameters: $A_1 = -0.0407$, $\tau_1 = 7.119\,\mu s$, $A_2 = -0.4969$, $\tau_2 = 66.5\,ns$, $A_3 = -0.1058$, $\tau_3 = 999\,ns$, $A_4 = -0.4839$, $\tau_4 = 12.51\,ns$, $A_5 = -0.157$, $\tau_5 = 188.8\,\text{ns}$. The top right panel is a zoom of the area indicated with a magenta square. The bottom right panel is corrected with three additional exponential filters with the parameters: $A_6 = -0.4839$, $\tau_6 = 12.51\,ns$, $A_7 = -0.4839$, $\tau_7 = 23\,ns$, $A_8 = -0.4865$, $\tau_8 = 15.58\,\text{ns}$. All the filters are multiplied.

a further improvement in compensation performance and response speed, as shown in Fig. 7.2b (bottom right).

The corrections were confirmed by measuring the Ramsey time as a function of qubit frequency 70 ns after applying the flux pulse. This measurement is more sensitive to pulse distortions compared to a simple spectroscopy measurement, but less sensitive compared to the cryoscope protocol. As long as the Ramsey time is stable over the time it takes to measure the characteristic function, the state tomography would be possible. The results of such a measurement are presented in Fig. 7.3. The only distortion which is uncompensated for is a weak oscillation in the first 200 ns, which is correctable with a bounce filter.

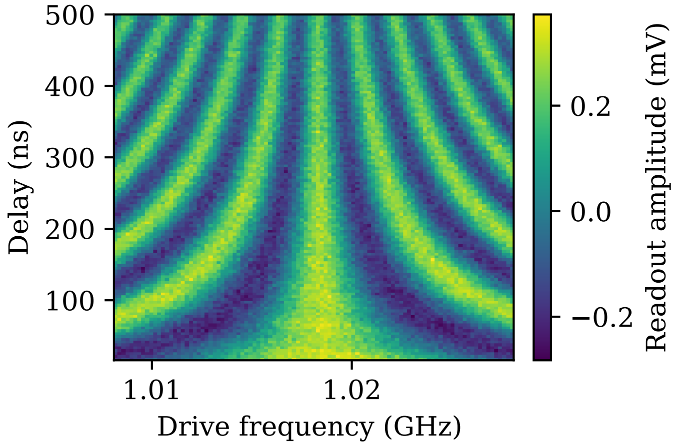

Figure 7.3: Ramsey decoherence time measured 70 ns after a flux jump from the avoided crossing between the qubit and the cavity to the low sweet spot of the qubit.

7.2 State creation

The preparation of non-classical bosonic states in superconducting circuits is a key ingredient for exploring quantum information processing with continuous variables and for implementing hardware-efficient error correction schemes. Without two-photon stabilization, we explore different techniques to generate a non-classical state in the system. Depending on the strength of the qubit-cavity coupling and the accessible control techniques, various approaches can be employed to generate and manipulate bosonic states. In this section, we focus on three conceptually distinct protocols relevant to the fluxonium-cavity architecture: vacuum Rabi state transfer, resonant cat state creation[196], and Kerr-cat state generation[192].

In the dispersive regime, the qubit-cavity system is well described by the dispersive Hamiltonian from Eq. 2.33[2]. However, when the detuning between the two subsystems is reduced

[2]For ease of the reader, I include the dispersive Hamiltonian here: $H_{\text{disp}} = \hbar(\omega_r + \chi\sigma_z)a^\dagger a + \frac{\hbar(\omega_q+\chi)}{2}\sigma_z$.

via a fast flux pulse, the system temporarily enters the resonant regime, governed by the Jaynes-Cummings Hamiltonian,

$$H_{\mathrm{JC}}/\hbar = \omega_c a^\dagger a + \frac{\omega_q}{2}\sigma_z + g_{QC}(a^\dagger \sigma_- + a\sigma_+), \tag{7.2}$$

where g_{QC} is the coupling rate. In this regime, the excitation coherently oscillates between the qubit and cavity at the vacuum Rabi frequency $\Omega_R = 2g_{QC}$, enabling deterministic preparation of single-photon Fock states through timed excitation swaps. Higher-order Fock states can be synthesized by repeating the sequence of excitation and swap operations, provided fast flux control with nanosecond precision is available[200, 201]. Additionally, this protocol requires the lifetime of both the qubit and the cavity to be longer than the time required for swapping the excitation $t_{\mathrm{swap}} = \pi/2g_{QC}$.

Beyond Fock-state generation, resonant cat state creation exploits the evolution of the system under the JC Hamiltonian after preparing the cavity in a coherent state and the qubit in $|\pm Y\rangle$ or $|\pm Z\rangle$. After a time $t_{\mathrm{collapse}} = \sqrt{2}/g_{QC}$, the vacuum Rabi oscillations collapse and a cat state is formed, as demonstrated in Ref. [196]. Achieving this regime requires precise control of both the flux trajectory and pulse timing, as well as high coherence times of the qubit.

Alternatively, cat states can be generated by exploiting the cavity's intrinsic self-Kerr nonlinearity. In this regime, the phase-space evolution of an initial coherent state undergoes self-interference due to the photon-number-dependent phase accumulation, leading to the formation of multi-component cat states[192], i.e., n-legged cat state after a waiting time $t_{\mathrm{cat},n} = 2\pi/nK$. In this system, the Kerr nonlinearity can be tuned via the flux bias, opening a path toward Kerr-cat creation.

Each of these state-preparation techniques presents distinct experimental challenges. Among them, Kerr-cat creation is the most practical choice, as it requires optimizing only a single step of the square flux pulse rather than both steps needed for the vacuum Rabi protocol. Moreover, its performance depends primarily on the qubit relaxation time T_1^Q,[3] rather than the coherence time T_2^R, which limits the resonant cat protocol. For these reasons, we focus on generating a four-legged Kerr cat state using this approach. The following sections describe the implementation of the method, the associated calibration procedures, and the experimental results obtained in the fluxonium-cavity system.

7.2.1 Geometric phase calibration

We use characteristic function measurement for the bosonic state tomography. For this purpose, the qubit is prepared in a superposition $|\psi_i\rangle = |+\rangle = (|0\rangle_q + |1\rangle_q)/\sqrt{2}$. After the ECD gate, the qubit acquires an additional phase, i.e., the final qubit state is $|\psi_f\rangle = (|0\rangle_q + e^{i\theta}|1\rangle_q)/\sqrt{2}$. This phase can be analytically calculated by solving the Schrödinger equation for the system during the gate (see Supplementary material of Ref. [183] for

[3] Due to the stronger cavity-qubit hybridization close to their avoided crossing.

detailed derivations). We only present the results of this calculation. The phase is given by

$$\theta(t) = -2\int_0^t \Re[\epsilon^*(\tau)\delta(\tau)]d\tau. \tag{7.3}$$

Here, $\delta(t)$ is the conditional displacement

$$\delta(t) = -\int_0^t \sin[\phi(\tau) - \phi(t)]\epsilon(\tau)d\tau. \tag{7.4}$$

$\epsilon(t)$ is the pulse sequence from Eq. 5.16 and $\phi(t) = -\frac{\chi_{QC}}{2}\int_0^t z(\tau)d\tau$ represents a qubit-state-dependent rotation of the cavity with $z(t) = \pm 1$ being the sign of σ_z. This accumulated phase can either be calculated and compensated for analytically or it can be directly extracted from experimental measurements.

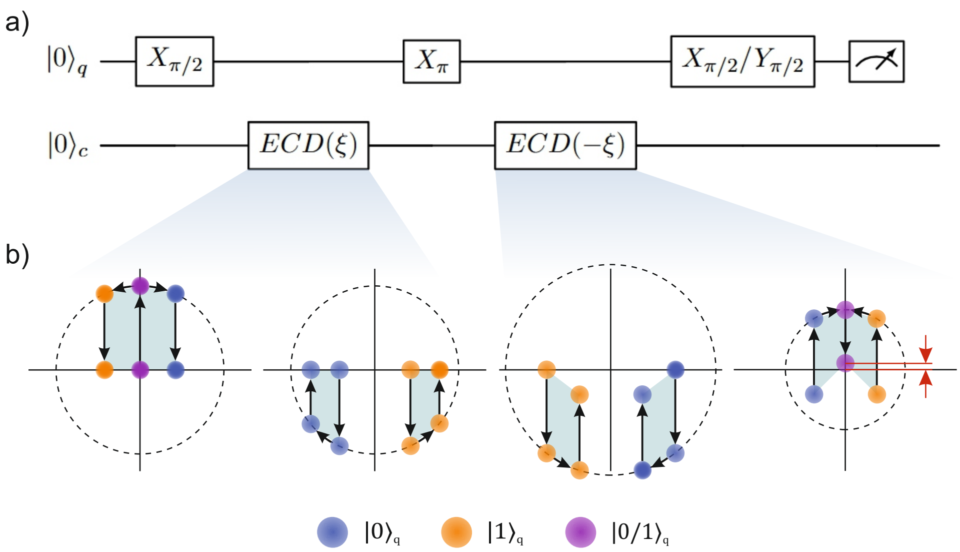

Figure 7.4: Phase accumulation during ECD. a) Protocol for measuring the accumulated phase. b) Phase space trajectories. The first two panels are identical to those in Fig. 5.10b and represent the first ECD gate. The other two panels represent the second ECD gate after the $\pi-$pulse. The enclosed areas, which correlate to the accumulated phase, are highlighted for each panel. The red lines and arrows highlight the imperfection between the final state and the vacuum.

The employed protocol is shown schematically in Fig. 7.4a. The qubit is initialized in a superposition state, followed by the application of an ECD gate. To extract the geometric phase, the phase-space trajectories must form closed loops; hence, a second ECD gate with the opposite displacement is applied after a $\pi-$pulse on the qubit. The expectation values $\langle\sigma_x\rangle$ and $\langle\sigma_y\rangle$ are then measured to determine the acquired phase θ. Note that the measured accumulated phase is 2θ, due to the phase space trajectories encompassing twice the area of the ECD during the protocol.

As illustrated in Fig. 7.4b, the trajectories do not close perfectly after applying $ECD(-\xi)$. A more accurate implementation would replace $ECD(-\xi)$ with the inverse operation

$ECD^{-1}(\xi)$, achieved by reversing the gate order[4]. If the correction factors r_1 and r_2 are sufficiently close to unity, the residual error due to imperfect closure is negligible.

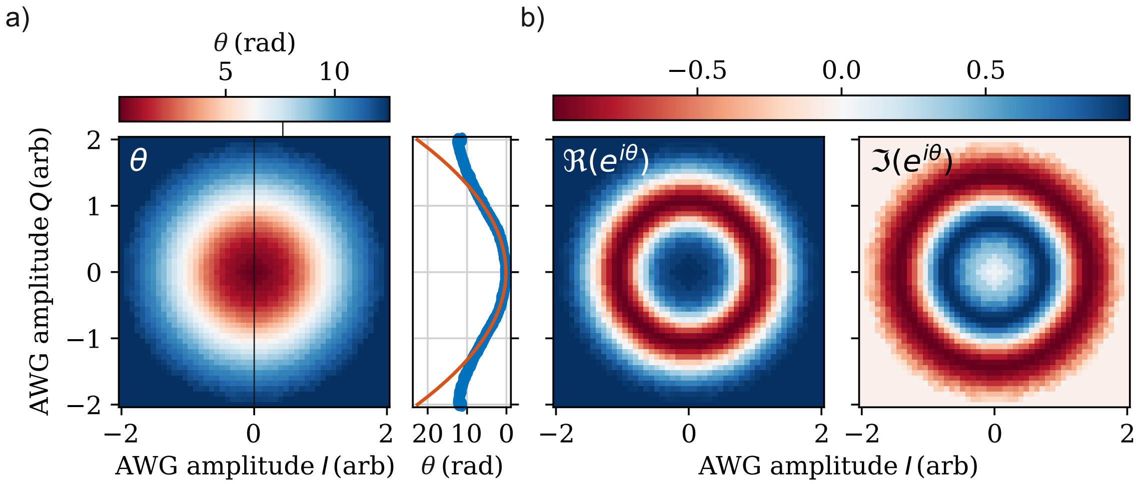

Figure 7.5: Reconstructed accumulated phase and corrections. a) Reconstructed phase (left) from the measurement (right). The black line on the left plot denotes where the data was sampled. Blue dots in the right plot are the extracted phase from the measurements, while the red line represents the ideal case of a parabolic dependence. b) From the reconstructed accumulated phase θ, the correction $e^{i\theta}$ is calculated, with its real and imaginary parts plotted here.

The results of the phase calibration measurements are presented in Fig. 7.5. The phase was measured along the $I = 0$ axis with small amplitude steps. Since the correction depends only on $|\xi|$, the complete phase correction map can be reconstructed for the specific phase-space grid later used in the characteristic function measurements. This is done for each point in phase space with coordinates (A_I, A_Q)[5], calculating $\sqrt{A_I^2 + A_Q^2}$ and taking the measured θ corresponding to the closest value A_Q measured in phase space. For the phase space points, where $\sqrt{A_I^2 + A_Q^2} > 2$, the last measured point $\sqrt{A_I^2 + A_Q^2} = 2$ was taken for the reconstruction.

To ensure sufficient qubit coherence during tomography, the waiting time in the ECD gate was decreased to $t_w = 200\,\text{ns}$. This modification necessitates a recalibration of the conversion factor between the AWG output amplitude and the corresponding displacement in phase space; therefore, the axes in Fig. 7.5 are expressed in units of AWG amplitude.

The nominal output voltage recommended for the Octave module is limited to 0.125 V, which was initially used for the pulse amplitudes. The reason for this is most likely to avoid saturation of the built-in amplifiers. However, to sample regions of phase space farther away from the vacuum, the maximum pulse voltage was intentionally exceeded up to 0.3 V. As a result, the measured phase deviates from the expected parabolic dependence on ξ, as illustrated by the discrepancy between the data and the red curve in Fig. 7.5a (right). Consequently, the phase-space response becomes increasingly nonlinear once the normalized amplitude exceeds unity.

[4] i.e. $D(-\alpha)T(t_w)D(r_1\alpha)X_\pi D(r_1\alpha)T(t_w)D(-r_2\alpha)$
[5] Here, $A_{I/Q}$ are the real and imaginary parts of the AWG amplitudes.

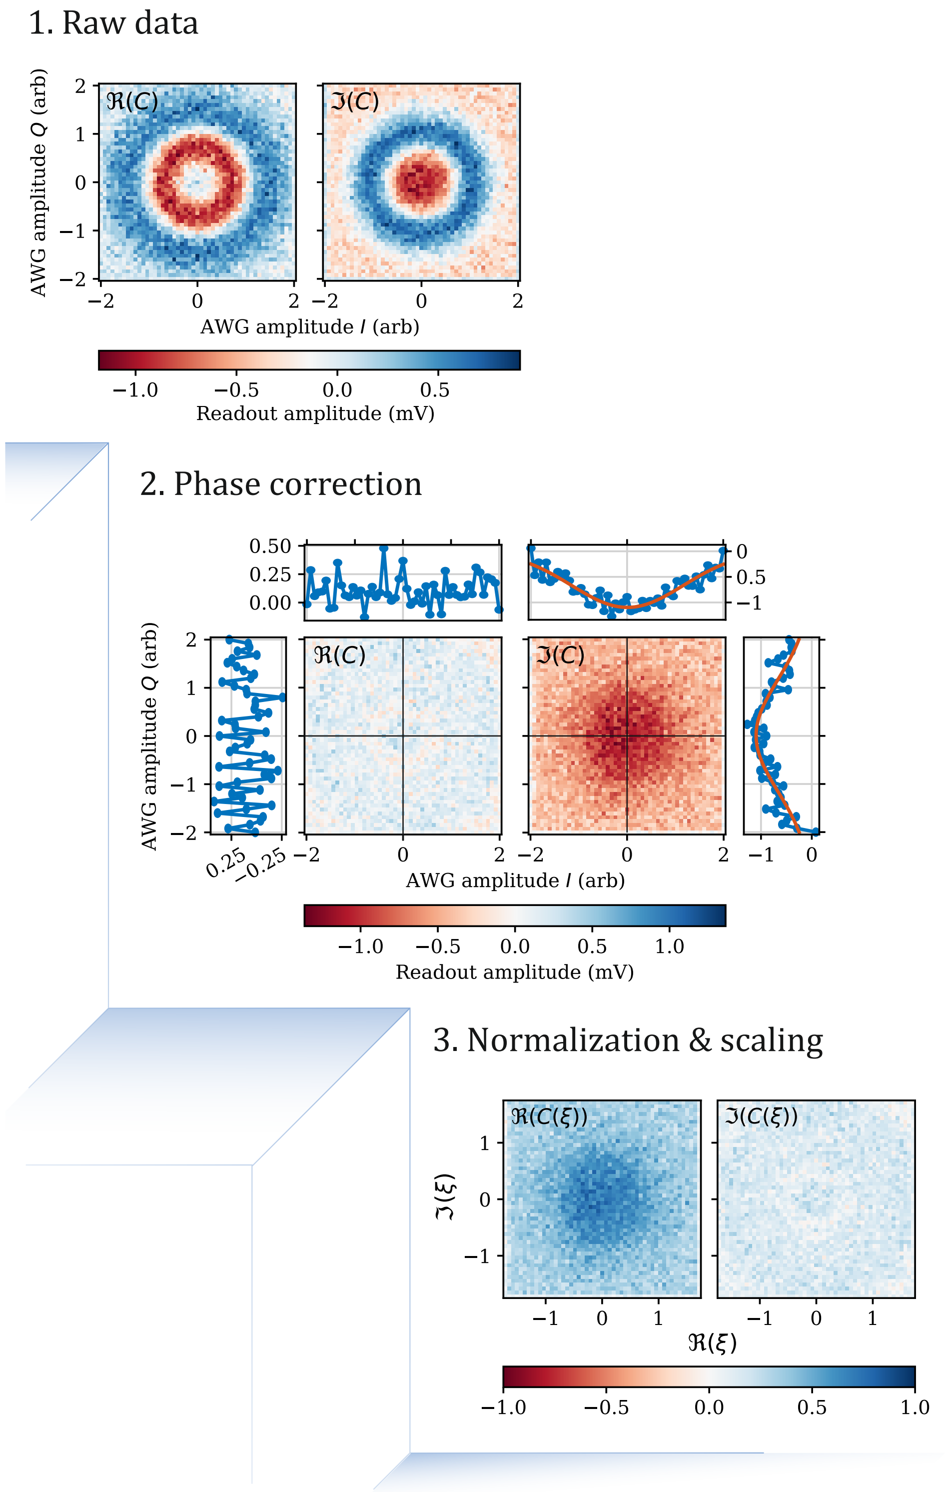

Figure 7.6: Steps for correcting the characteristic function of a vacuum state. 1. Shows the raw data. 2. The same state after applying the correction from Fig. 7.5b. The black lines denote the locations from which the linecuts were taken for each colormap. The red lines represent a Gaussian fit to the data (the same for both linecuts). 3. From the σ of the fit, the ξ can be extracted. The characteristic function is limited in the range $[-1, 1]$, so the data is normalized and multiplied by i (the imaginary unit).

The simplest benchmark state for validating the phase correction is the vacuum state. As previously discussed, the real part of its characteristic function should form a two-dimensional Gaussian centered at the origin, while the imaginary part should ideally vanish. However, the measured data exhibit nonzero values in both components, as shown in Fig. 7.6 (top). To compensate for the acquired phase, we apply the transformation $\mathcal{C} \cdot e^{i\theta}$, with the resulting characteristic function displayed in Fig. 7.6 (middle). Interestingly, the real and imaginary components appear swapped relative to the expected behavior. Multiplying the corrected data by an additional factor of i restores the correct orientation. The precise cause of this anomaly remains unclear, though it is likely related to hardware implementation details, such as the OPX-Octave interface (see App. F).

Even after applying the phase correction, residual features remain visible in the real part of the characteristic function. These appear either near the vacuum state or at larger displacements, specifically for points where $\sqrt{A_I^2 + A_Q^2} > 2$. The latter artifacts originate from the extrapolated correction beyond the measured region and could be mitigated by performing the phase calibration along one of the diagonal axes, $I = Q$ or $I = -Q$. The residual signal near the vacuum requires further investigation, as its origin is not immediately clear. The linecuts, however, show that the residual information is negligible compared to the signal in the imaginary part.

Finally, linecuts of $\Im(\mathcal{C})$ along the $I = 0$ and $Q = 0$ axes are fitted to extract an updated conversion factor between the AWG pulse amplitude and the corresponding phase-space coordinate ξ. Moreover, the characteristic function $\mathcal{C}(\xi)$ is multiplied by the imaginary unit i and subsequently normalized to recover the expected form of the vacuum state's characteristic function, as shown in Fig. 7.6 (bottom).

7.2.2 Kerr cat

As a final demonstration of universal control, we prepare a four-legged cat by using the Kerr-cat protocol. For this purpose, we bias the system close to a cavity–qubit avoided crossing, where the self-Kerr coefficient of the cavity diverges. Care must be taken that this bias is not identical to that of the avoided crossing, which would cause a Rabi swap, or the cavity lifetime would be compromised by the qubit.

The protocol begins by displacing the cavity and waiting for a variable time for the cavity to evolve. The waiting time is varied as the exact Kerr coefficient could not be measured so close to the avoided crossing. Then, a flux pulse is applied such that the qubit bias is shifted to the low sweet spot, where the state tomography is executed. At the right waiting time $t_{\mathrm{cat},4} = \pi/2K$, a four-legged cat state is formed in the cavity.

To get an idea of how the characteristic function of a four-legged cat would look, we simulate it with *QuTip*[202] where the initial displacement is $\alpha = 2$ and the evolution time is exactly $t_{\mathrm{cat},4} = \pi/2K$. The Hamiltonian used for the evolution is $H/\hbar = \chi_{QC} c^\dagger c q^\dagger q + K_c (c^\dagger c)^2$. Figure 7.7 shows the resulting characteristic and Wigner functions for both the ideal and realistic scenarios. In the realistic case, which incorporates the finite lifetimes of the cavity and qubit, as well as qubit decoherence, the contrast is significantly reduced, and the state acquires a small additional rotation.

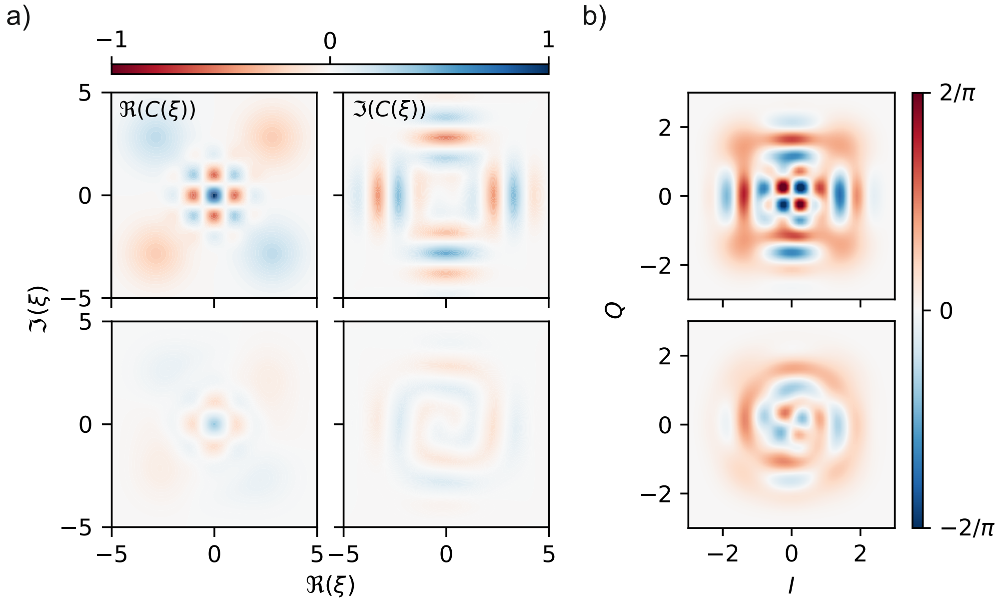

Figure 7.7: Four-legged cat $\alpha = 2$ simulation. a) Characteristic function. The first column corresponds to the real part, the second column to the imaginary part. b) Wigner function for comparison. The first row always shows the idealized case, while the second has been simulated for $T_1^C = 100\,\mu\text{s}$, $T_1^Q = 10\,\mu\text{s}$, and $T_2^Q = 100\,\text{ns}$.

We track the evolution of a coherent state with amplitude $\alpha = 3$ over a total duration of $60,\mu$s, recording snapshots every $2,\mu$s. The characteristic function is measured using the same parameters described in Sec. 7.2.1. Figure 7.8 shows a selected subset of the data, displaying snapshots spaced by $10\,\mu$s for clarity. For comparison, we simulate the evolution using the same Hamiltonian employed in Fig. 7.7. For a cavity self-Kerr coefficient of $K_c/2\pi = 10.8\,\text{kHz}$, simulated snapshots are generated every $1.05\,\mu$s. To match the visual density of the experimental data, only every third snapshot is shown, corresponding to an effective step of $3.15\,\mu$s between plotted frames, except for the final snapshot, where the interval to the previous frame is $4.2\,\mu$s.

We can see that after $t_d = 10\,\mu$s, there is barely any Kerr curving of the coherent state fringes. Instead, the rotation of the state allows us to set a limit on the detuning between the cavity and the drive $|\Delta/2\pi| < 250\,\text{kHz}$. After $20\,\mu$s, the curvature becomes more pronounced and can be compared to the second snapshot of the simulated evolution. For a four-legged cat, the imaginary part almost vanishes. However, in the last measured dataset, the imaginary part of the evolution has a similar amplitude to the real part. This behavior matches the simulated state from either the last frame of the first row or the first frame of the second row in Fig. 7.7. This allows us to put an estimate on the cavity self-Kerr coefficient at this bias $K_c/2\pi \in [1.7, 2.3]\,\text{kHz}$. We can therefore conclude that a cat state has not yet formed. Based on the comparison, approximately twice the current evolution time would be required to generate a four-legged cat state.

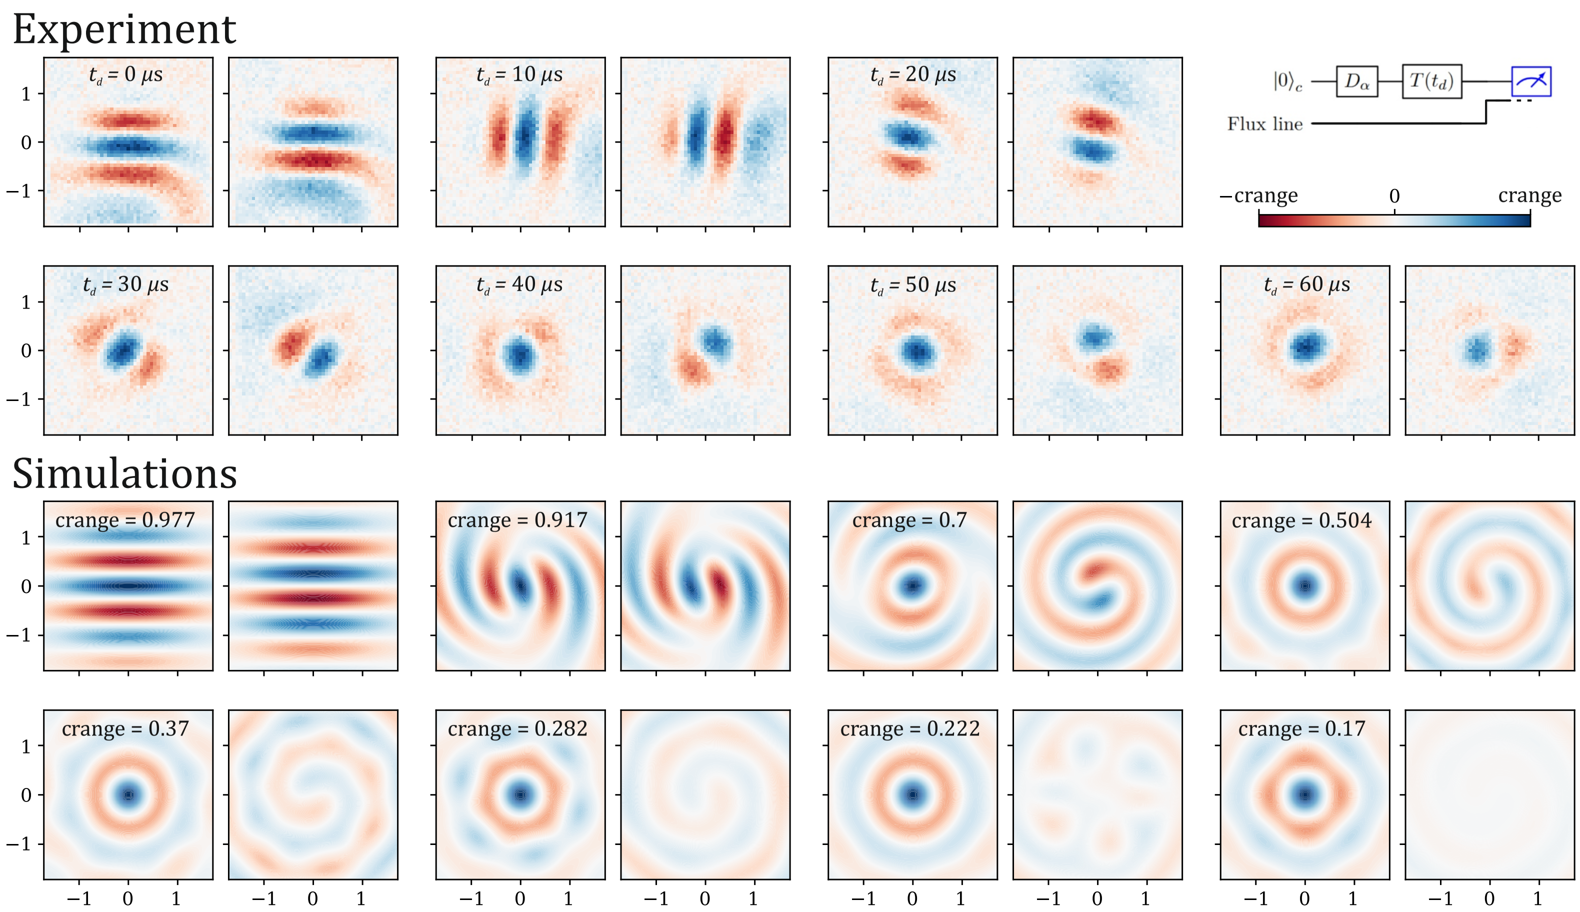

Figure 7.8: Evolution of a coherent state. The top right corner shows the protocol used, where the blue meter denotes the whole sequence from Fig. 7.2a. For each data pair, the left plot represents $\Re(\mathcal{C}(\xi))$, and the right one - $\Im(\mathcal{C}(\xi))$. The axes labels, $\Re(\xi)$ for x and $\Im(\xi)$ for y, are omitted for better visibility. The colorbar for the experimental data was normalized from -1 to 1 (crange $= 1$) from the measurement without any delay ($t_d = 0\,\mu s$). The same scaling was used for the rest of the experimental data. The simulation data was normalized to the maximum of the absolute value of both the real and imaginary parts of the characteristic function, i.e., crange $= \max(|\Re(\mathcal{C}(\xi))|, |\Im(\mathcal{C}(\xi))|$. From this measurement, we estimate $K_c/2\pi \in [1.7, 2.3]\,$kHz.

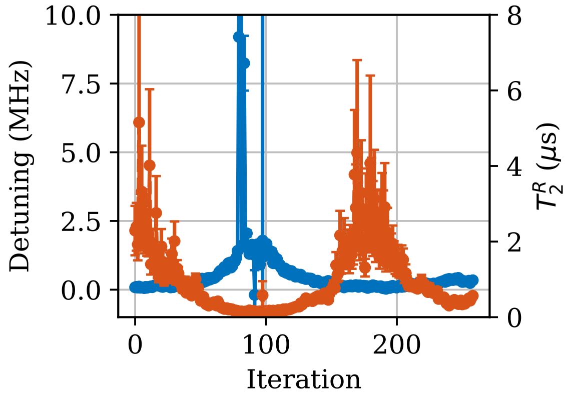

Figure 7.9: Qubit stability. The data is extracted from a detuned Ramsey measurement after applying a flux pulse. The blue data is the qubit detuning from its sweet spot, while the red points correspond to the measured T_2^R.

When attempting to repeat the experiment with longer delay times, the evolution became unstable and nonuniform, e.g., consecutive snapshots occasionally appeared identical or even a bit rotated in the opposite direction. Such behavior may arise from flux noise, which can couple either through the ground, as the flux line is grounded, or through the microwave control used for the pulses. To monitor this effect, the qubit frequency as a function of time was extracted from Ramsey measurements, as shown in Fig. 7.3. The time picked for this experiment was identical to the one required for reconstructing the full characteristic function (i.e., ~ 3 hours). Each measurement takes 18 s and is dominated by the initialization of the OPX+. An example of such a frequency stability measurement is shown in Fig. 7.9. The two peaks in decoherence time are spaced 45 min apart. The increase in frequency appears at the same time as the T_2^R time plummets. Despite trying to investigate external culprits for this instability, as other labs nearby work with high magnetic fields (e.g., Zeeman slowers, magneto-optical traps), as well as power supplies that can induce ground noise[6], we could not find the cause of this noise.

[6]Many thanks to Dr. Peng Yin for sharing his ground noise measurements with me at the time.

CHAPTER 8

Conclusions and Outlook

In this thesis, I present a novel platform for quantum information processing based on a bosonic mode. It combines the high coherence properties of a post cavity and the rich energy spectrum of a fluxonium. This approach addresses a longstanding limitation common to most bosonic quantum computing architectures–dephasing induced by the ancillary qubit–and proposes an alternative solution. The following paragraphs summarize the goals and results of this work and highlight open questions for future research.

Chapter 2 is written with a pedagogy purpose in mind, targeting a broad audience of physics students and researchers with basic knowledge of quantum mechanics. It provides a general overview of quantum information processing and illustrates key concepts with examples drawn from various physical platforms. The theoretical framework for information carriers is extended to encompass open quantum systems, introducing the mechanisms of energy relaxation and dephasing. This chapter also offers a comprehensive introduction to circuit quantum electrodynamics, covering the most common superconducting circuits and their interactions. In addition, it includes a concise overview of quantum error correction, with particular emphasis on its implementation in superconducting architectures. While not intended to be an exhaustive reference, the chapter provides ample pointers to the relevant literature for interested readers.

The first scientific question addressed in this thesis is how to model highly anharmonic circuits that lie beyond the validity of the lumped-element approximation, such as the circuit used in this work. This challenge is compounded by the need to incorporate flux quantization. I discuss several modeling approaches, including lumped-element circuit quantization and black-box quantization, highlighting both their applicability and their limitations. The lumped-element method requires all components to satisfy the lumped-element approximation, but it can accurately capture the flux tunability of a fluxonium circuit. In contrast, black-box quantization is compatible with distributed elements but relies on the rotating-wave approximation to extract key system parameters and cannot account for flux quantization.

To overcome these limitations, I propose a hybrid modeling strategy in which the fluxonium is treated within the lumped-element approximation, while harmonic modes and their couplings are incorporated in a manner similar to *QuTiP*. This approach is implemented in *scQubits*, a Python package for simulating superconducting circuits. It enables the introduction of arbitrary interactions without imposing constraints on the physical size of circuit elements relative to the relevant wavelengths. A related methodology has recently been introduced in Ref. [136] and implemented in the *pyEPR* package. Although both the *scQubits*- and *pyEPR*-based approaches provide powerful tools for modeling highly anharmonic circuits, they have yet to be tested on larger-scale systems, for instance, those incorporating multiple flux-tunable or distributed components.

In Chapter 4, I describe the experimental setup and its fabrication. Several key components of the platform were improved through my contributions. These include establishing a temperature-controlled cavity-etching procedure, developing a new design of the flux hose–the superconducting flux hose–and later repurposing it for both readout and flux tuning of a transmon qubit[157]. Additionally, I propose an improved shielding design for cryogenic operation. While the benefits of temperature-controlled cavity etching and the simplified fabrication of flux hoses are clear, the performance of the new shielding concept remains to be experimentally tested.

Chapter 5 validates the model introduced in Sec. 3.3. The system is fully characterized with the help of recently developed techniques for probing a harmonic oscillator weakly coupled to a qubit[183]. Due to this weak interaction, the cavity lifetime of $T_1^C = 210 \pm 40\,\mu\mathrm{s}$ is not substantially compromised by the lossy qubit. The key result is the successful realization of a platform in which the dispersive shift between the cavity and the fluxonium can be tuned *in situ* to positive, negative, or vanishing values[165]. This tunability enables precise control during gate operations while allowing the interaction to be effectively switched off during idle periods, thereby suppressing qubit-induced dephasing of the bosonic mode. A similar behavior was recently observed with a 2D platform[203] developed in parallel with the setup from this thesis.

This tunability has important implications for protecting the encoded bosonic information from ancilla-induced dephasing. When the qubit decay rate is much lower than the dispersive shift, a single thermal excitation in the qubit persists long enough to impose a full phase rotation on the cavity state, thereby completely dephasing the encoded information. In this regime, the dephasing rate is set solely by the qubit's lifetime and thermal population and is essentially independent of the magnitude of χ_{QC}.

However, when the hierarchy is reversed, i.e., the qubit decays at a much faster rate than the dispersive interaction, a thermally excited ancilla induces only a small, random phase rotation on the bosonic mode before relaxing. In this case, reducing the dispersive shift directly suppresses cavity dephasing. For our system parameters, even in the most pessimistic scenario, the ancilla-induced dephasing becomes negligible when the dispersive shift satisfies $|\chi_{QC}|/2\pi \ll 1.8\,\mathrm{kHz}$. This establishes a quantitative threshold below which the cavity can effectively be decoupled from qubit thermal fluctuations.

The potential of this platform to operate as a dissipative cat qubit was evaluated by coupling the cavity to a SNAIL, used as a buffer mode. By diagonalizing the Hamiltonian of the composite system, we extracted the parameters required for resonant operation of a dissipative cat qubit, specifically the condition $\omega_b = 2\omega_m$. Although the increased lifetime of the post cavity would, in principle, improve the dephasing properties of the cat qubit, the intrinsically small zero-point fluctuations of the 3D architecture suppress the two-photon exchange rate, ultimately limiting efficient qubit operation. In addition, incorporating an extra nonlinear element would have introduced substantial experimental overhead, well beyond the scope and timeline of this project. For these reasons, implementing two-photon dissipation was deferred and remains an interesting direction for future work, building on the results presented here.

Fast flux control is the final ingredient needed to achieve universal control of this platform. It enables coherent vacuum Rabi swaps between the qubit and the cavity—an essential capability for preparing coherent superpositions of Fock states. Implementing fast flux, however, introduces significant challenges, including the need for a clean, loop-free ground and careful pre-compensation of pulse distortions arising from impedance mismatches and filtering along the flux-control line.

Transitioning the flux-bias control from DC to AC unexpectedly reduced the fluxonium coherence to $T_2^R \approx 2\,\mu$s,[1] adding substantial operational complexity. Protocols relying on long qubit coherence times, such as the one demonstrated in Ref. [204], became impractical. Consequently, the focus shifted toward generating a four-legged cat state by exploiting the cavity's self-Kerr nonlinearity. Although the initial evolution of the coherent state appeared promising, the formation of a cat state was ultimately prevented by the onset of sudden flux noise in the system—noise whose origin could not be identified despite extensive investigation.

For a future generation of this setup, incorporating a Purcell filter will be essential, as decay through the readout resonator currently sets the dominant limit on the fluxonium lifetime at the low sweet spot. Beyond extending coherence, a Purcell filter could also enable single-shot readout, substantially reducing measurement time. The system would further benefit from improved qubit fabrication, e.g., introducing a hydrofluoric acid dip after liftoff to reduce dielectric loss, which is the main limitation at the high sweet spot. Enhancing the qubit lifetime in this way would also make the bosonic mode more resilient to ancilla-induced dephasing at flux bias points where the dephasing rate is insensitive to the dispersive shift.

The cat-state formation experiments highlight the need for improved wiring to avoid ground loops, as well as the overall magnetic shielding, partially achievable with the shielding design proposed in this thesis. Inspired by the setup at LPNS in Paris, an external $\mu-$metal shield was added around the cryostat at room temperature, but further refinements are necessary. These include eliminating magnetic components within the shielded region at cryogenic temperatures, regularly demagnetizing $\mu-$metal shields, and exploring the use of *finemet* films[205] as part of a multilayer magnetic shielding strategy. An alternative strategy to the flux noise problem is implementing a direct stabilization to counteract offsets in real time, such as the Pound-Drever-Hall locking technique[206].

As a final remark The infinite Hilbert space of bosonic modes continues to expand the landscape of quantum exploration, enabling advances in simulation[25, 207, 208], precision measurements[209–212], and the pursuit of fault-tolerant quantum information processing[189, 213–215]. As experiments become increasingly complex, the importance of effective decoupling schemes grows accordingly. Novel interactions do not always require

[1] The 50 Hz peak of the flux line, extracted via fast Fourier transform of the signal from a fast oscilloscope, had a height of 107 μV.

complex coupling elements; they can be achieved by combining already existing components.

Bibliography

[1] M. Planck, "Über das gesetz der energieverteilung im normalspectrum", Annalen der Physik **309**, 553 (1901).

[2] A. Einstein, "Über einen die erzeugung und verwandlung des lichtes betreffenden heuristischen gesichtspunkt", Annalen der Physik **322**, 132 (1905).

[3] N. Bohr, "The quantum postulate and the recent development of atomic theory", Nature **121**, 580 (1928).

[4] W. Heisenberg, "Über quantentheoretische umdeutung kinematischer und mechanischer beziehungen", Zeitschrift für Physik **33**, 879 (1925).

[5] E. Schröodinger, "An undulatory theory of the mechanics of atoms and molecules", Phys. Rev. **28**, 1049 (1926).

[6] W. Meissner and R. Ochsenfeld, "Ein neuer Effekt bei Eintritt der Supraleitfähigkeit", Naturwiss. **21**, 787 (1933).

[7] J. Bardeen, L. N. Cooper, and J. R. Schrieffer, "Microscopic theory of superconductivity", Phys. Rev. **106**, 162 (1957).

[8] B. Josephson, "Possible new effects in superconductive tunnelling", Physics Letters **1**, 251 (1962).

[9] M. Riordan, L. Hoddeson, and C. Herring, "The invention of the transistor", Rev. Mod. Phys. **71**, S336 (1999).

[10] 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", Phys. Rev. Lett. **55**, 1908 (1985).

[11] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer", SIAM Journal on Computing **26**, 1484 (1997).

[12] I. L. Chuang, N. Gershenfeld, and M. Kubinec, "Experimental implementation of fast quantum searching", Phys. Rev. Lett. **80**, 3408 (1998).

[13] F. Schmidt-Kaler, H. Häffner, M. Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt, "Realization of the Cirac–Zoller controlled-not quantum gate", Nature **422**, 408 (2003).

[14] "The quantum state of affairs", Nature Physics **19**, 605 (2023).

[15] Google Quantum AI, R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, D. Bacon, J. C. Bardin, J. Basso, A. Bengtsson, S. Boixo, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, D. M. Debroy, A. Del Toro Barba, S. Demura, A. Dunsworth, D. Eppens, C. Erickson, L. Faoro, E. Farhi, R. Fatemi, L. Flores Burgos, E. Forati, A. G. Fowler, B. Foxen, W. Giang, C. Gidney, D. Gilboa, M. Giustina, A. Grajales Dau, J. A. Gross, S. Habegger, M. C. Hamilton, M. P. Harrigan, S. D. Harrington, O. Higgott, J. Hilton, M. Hoffmann, S. Hong, T. Huang, A. Huff, W. J. Huggins, L. B. Ioffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, P. Juhas, D. Kafri, K. Kechedzhi, J. Kelly, T. Khattar, M. Khezri, M. Kieferová, S. Kim, A. Kitaev, P. V. Klimov, A. R. Klots, A. N. Korotkov, F. Kostritsa, J. M. Kreikebaum, D. Landhuis, P. Laptev, K.-M. Lau, L. Laws, J. Lee, K. Lee, B. J. Lester, A. Lill, W. Liu, A. Locharla, E. Lucero, F. D. Malone, J. Marshall, O. Martin, J. R. McClean, T. McCourt, M. McEwen, A. Megrant, B. Meurer Costa, X. Mi, K. C. Miao, M. Mohseni, S. Montazeri, A. Morvan, E. Mount, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, H. Neven, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, M. Y. Niu, T. E. O'Brien, A. Opremcak, J. Platt, A. Petukhov, R. Potter, L. P. Pryadko, C. Quintana, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, M. J. Shearn, A. Shorter, V. Shvarts, J. Skruzny, V. Smelyanskiy, W. C. Smith, G. Sterling, D. Strain, M. Szalay, A. Torres, G. Vidal, B. Villalonga, C. Vollgraff Heidweiller, T. White, C. Xing, Z. J. Yao, P. Yeh, J. Yoo, G. Young, A. Zalcman, Y. Zhang, and N. Zhu, "Suppressing quantum errors by scaling a surface code logical qubit", Nature **614**, 676 (2023).

[16] J. Robledo-Moreno, M. Motta, H. Haas, A. Javadi-Abhari, P. Jurcevic, W. Kirby, S. Martiel, K. Sharma, S. Sharma, T. Shirakawa, I. Sitdikov, R.-Y. Sun, K. J. Sung, M. Takita, M. C. Tran, S. Yunoki, and A. Mezzacapo, "Chemistry beyond the scale of exact diagonalization on a quantum-centric supercomputer", Science Advances **11**, eadu9991 (2025).

[17] J. Cobos, J. Fraxanet, C. Benito, F. d. Marcantonio, P. Rivero, K. Kapás, M. A. Werner, O. Legeza, A. Bermudez, and E. Rico, "Real-Time Dynamics in a (2+1)-D Gauge Theory: The Stringy Nature on a Superconducting Quantum Simulator" (2025).

[18] A. Kotil, E. Pelofske, S. Riedmüller, D. J. Egger, S. Eidenbenz, T. Koch, and S. Woerner, "Quantum approximate multi-objective optimization", Nature Computational Science , 1 (2025).

[19] A. Ameri, E. Ye, P. Cappellaro, H. Krovi, and N. F. Loureiro, "Quantum algorithm for the linear vlasov equation with collisions", Phys. Rev. A **107**, 062412 (2023).

[20] S. Bravyi, R. Manson-Sawko, M. Zayats, and S. Zhuk, "Quantum simulation of a noisy classical nonlinear dynamics" (2025).

[21] N. Gomes, G. Sharma, and J. Pathak, "Hamiltonian Simulation for Advection-Diffusion Equation with arbitrary transport field" (2025).

[22] V. V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. L. Brock, A. Z. Ding, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, "Real-time quantum error correction beyond break-even", Nature **616**, 50 (2023).

[23] Z. Ni, S. Li, X. Deng, Y. Cai, L. Zhang, W. Wang, Z.-B. Yang, H. Yu, F. Yan, S. Liu, C.-L. Zou, L. Sun, S.-B. Zheng, Y. Xu, and D. Yu, "Beating the break-even point with a discrete-variable-encoded logical qubit", Nature **616**, 56 (2023).

[24] U. Réglade, A. Bocquet, R. Gautier, J. Cohen, A. Marquet, E. Albertinale, N. Pankratova, M. Hallén, F. Rautschke, L.-A. Sellem, P. Rouchon, A. Sarlette, M. Mirrahimi, P. Campagne-Ibarcq, R. Lescanne, S. Jezouin, and Z. Leghtas, "Quantum control of a cat qubit with bit-flip times exceeding ten seconds", Nature **629**, 778 (2024).

[25] C. S. Wang, J. C. Curtis, B. J. Lester, Y. Zhang, Y. Y. Gao, J. Freeze, V. S. Batista, P. H. Vaccaro, I. L. Chuang, L. Frunzio, L. Jiang, S. M. Girvin, and R. J. Schoelkopf, "Efficient multiphoton sampling of molecular vibronic spectra on a superconducting bosonic processor", Phys. Rev. X **10**, 021060 (2020).

[26] S. Haroche and J.-M. Raimond, "Exploring the quantum: Atoms, cavities, and photons" (Oxford University Press, 2006).

[27] D. Walls and G. Milburn, "Quantum Optics" (Springer Berlin Heidelberg, 2008).

[28] M. A. Nielsen and I. L. Chuang, "Quantum computation and quantum information", 10th ed. (Cambridge University Press, Cambridge ; New York, 2010).

[29] N. Masluk, "Reducing the losses of the fluxonium artificial atom", Ph.D. thesis, Yale University (2012).

[30] D. Gusenkova, "Quantum non-demolition readout of superconducting artificial atoms", Ph.D. thesis, Karlsruher Institut für Technologie (KIT) (2021), 47.12.01; LK 01.

[31] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, "A Quantum Engineer's Guide to Superconducting Qubits", Applied Physics Reviews **6**, 021318 (2019).

[32] A. Joshi, K. Noh, and Y. Y. Gao, "Quantum information processing with bosonic qubits in circuit QED", Quantum Science and Technology **6**, 033001 (2021).

[33] J. I. Cirac and P. Zoller, "Quantum computations with cold trapped ions", Phys. Rev. Lett. **74**, 4091 (1995).

[34] B. B. Blinov, D. Leibfried, C. Monroe, and D. J. Wineland, "Quantum Computing with Trapped Ion Hyperfine Qubits", Quantum Information Processing **3**, 45 (2004).

[35] C. Adami and N. J. Cerf, in *Quantum Computing and Quantum Communications*, edited by C. P. Williams (Springer, Berlin, Heidelberg, 1999) pp. 391–401.

[36] D. Loss and D. P. DiVincenzo, "Quantum computation with quantum dots", Phys. Rev. A **57**, 120 (1998).

[37] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, "Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation", Phys. Rev. A **69**, 062320 (2004).

[38] M. Mirrahimi, Z. Leghtas, V. V. Albert, S. Touzard, R. J. Schoelkopf, L. Jiang, and M. H. Devoret, "Dynamically protected cat-qubits: a new paradigm for universal quantum computation", New Journal of Physics **16**, 045014 (2014).

[39] M. Ringbauer, M. Meth, L. Postler, R. Stricker, R. Blatt, P. Schindler, and T. Monz, "A universal qudit quantum processor with trapped ions", Nature Physics **18**, 1053 (2022).

[40] M. Luo and X. Wang, "Universal quantum computation with qudits", Science China Physics, Mechanics & Astronomy **57**, 1712 (2014).

[41] Y. Wang, Z. Hu, B. C. Sanders, and S. Kais, "Qudits and High-Dimensional Quantum Computing", Frontiers in Physics **8**, 10.3389/fphy.2020.589504 (2020).

[42] H. Anwar, E. T. Campbell, and D. E. Browne, "Qutrit magic state distillation", New Journal of Physics **14**, 063006 (2012).

[43] N. Gershenfeld and I. L. Chuang, "Quantum Computing with Molecules", Scientific American **278**, 66 (1998).

[44] B. E. Kane, "A silicon-based nuclear spin quantum computer", Nature **393**, 133 (1998).

[45] X. Yu, B. Wilhelm, D. Holmes, A. Vaartjes, D. Schwienbacher, M. Nurizzo, A. Kringhøj, M. R. v. Blankenstein, A. M. Jakob, P. Gupta, F. E. Hudson, K. M. Itoh, R. J. Murray, R. Blume-Kohout, T. D. Ladd, N. Anand, A. S. Dzurak, B. C. Sanders, D. N. Jamieson, and A. Morello, "Schrödinger cat states of a nuclear spin qudit in silicon", Nature Physics **21**, 362 (2025).

[46] M. Karácsony, L. Oroszlány, and Z. Zimboras, "Efficient qudit based scheme for photonic quantum computing", SciPost Phys. Core **7**, 032 (2024).

[47] V. Tripathi, N. Goss, A. Vezvaee, L. B. Nguyen, I. Siddiqi, and D. A. Lidar, "Qudit dynamical decoupling on a superconducting quantum processor", Phys. Rev. Lett. **134**, 050601 (2025).

[48] L. B. Nguyen, N. Goss, K. Siva, Y. Kim, E. Younis, B. Qing, A. Hashim, D. I. Santiago, and I. Siddiqi, “Empowering a qudit-based quantum processor by traversing the dual bosonic ladder”, Nature Communications **15**, 10.1038/s41467-024-51434-2 (2024).

[49] P. Liu, R. Wang, J.-N. Zhang, Y. Zhang, X. Cai, H. Xu, Z. Li, J. Han, X. Li, G. Xue, W. Liu, L. You, Y. Jin, and H. Yu, “Performing SU(d) operations and rudimentary algorithms in a superconducting transmon qudit for $d = 3$ and $d = 4$”, Phys. Rev. X **13**, 021028 (2023).

[50] A. S. Nikolaeva, E. O. Kiktenko, and A. K. Fedorov, “Efficient realization of quantum algorithms with qudits”, EPJ Quantum Technology **11**, 1 (2024).

[51] D. Janković, J.-G. Hartmann, M. Ruben, and P.-A. Hervieux, “Noisy qudit vs multiple qubits: conditions on gate efficiency for enhancing fidelity”, npj Quantum Information **10**, 59 (2024).

[52] V. G. Matsos, C. H. Valahu, T. Navickas, A. D. Rao, M. J. Millican, X. C. Kolesnikow, M. J. Biercuk, and T. R. Tan, “Robust and deterministic preparation of bosonic logical states in a trapped ion”, Phys. Rev. Lett. **133**, 050602 (2024).

[53] C. Flühmann and J. P. Home, “Direct characteristic-function tomography of quantum states of the trapped-ion motional oscillator”, Phys. Rev. Lett. **125**, 043602 (2020).

[54] B. de Neeve, T.-L. Nguyen, T. Behrle, and J. P. Home, “Error correction of a logical grid state qubit by dissipative pumping”, Nature Physics **18**, 296 (2022).

[55] M. Kang, H. Nuomin, S. N. Chowdhury, J. L. Yuly, K. Sun, J. Whitlow, J. Valdiviezo, Z. Zhang, P. Zhang, D. N. Beratan, and K. R. Brown, “Seeking a quantum advantage with trapped-ion quantum simulations of condensed-phase chemical dynamics”, Nature Reviews Chemistry **8**, 340 (2024).

[56] Z. Ni, S. Li, X. Deng, Y. Cai, L. Zhang, W. Wang, Z.-B. Yang, H. Yu, F. Yan, S. Liu, C.-L. Zou, L. Sun, S.-B. Zheng, Y. Xu, and D. Yu, “Beating the break-even point with a discrete-variable-encoded logical qubit”, Nature **616**, 56 (2023).

[57] V. Sivak, A. Eickbusch, B. Royer, S. Singh, I. Tsioutsios, S. Ganjam, A. Miano, B. L. Brock, A. Z. Ding, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, “Real-time quantum error correction beyond break-even”, Nature **584**, 368 (2023).

[58] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. Cory, Y. Nakamura, J. Tsai, and W. Oliver, “Noise spectroscopy through dynamical decoupling with a superconducting flux qubit”, Nature Physics **7**, 565 (2011).

[59] E. Knill, R. Laflamme, and L. Viola, “Theory of quantum error correction for general noise”, Phys. Rev. Lett. **84**, 2525 (2000).

[60] D. Castelvecchi, "Quantum computers ready to leap out of the lab in 2017", Nature **541**, 9 (2017).

[61] U. Vool and M. Devoret, "Introduction to quantum electromagnetic circuits", International Journal of Circuit Theory and Applications **45**, 897 (2017).

[62] D. M. Pozar, "Microwave engineering", 4th ed. (Wiley, 2012).

[63] S. Hähnle, K. Kouwenhoven, B. Buijtendorp, A. Endo, K. Karatsu, D. Thoen, V. Murugesan, and J. Baselmans, "Superconducting microstrip losses at microwave and submillimeter wavelengths", Phys. Rev. Appl. **16**, 014019 (2021).

[64] A. Megrant, C. Neill, R. Barends, B. Chiaro, Y. Chen, L. Feigl, J. Kelly, E. Lucero, M. Mariantoni, P. O'Malley, D. Sank, A. Vainsencher, J. Wenner, T. White, Y. Yin, J. Zhao, C. Palmstrøm, J. Martinis, and A. Cleland, "Planar superconducting resonators with internal quality factors above one million", Applied Physics Letters **100** (2012).

[65] G. Calusine, A. Melville, W. Woods, R. Das, C. Stull, V. Bolkhovsky, D. Braje, D. Hover, D. Kim, X. Miloshi, D. Rosenberg, A. Sevi, J. Yoder, E. Dauler, and W. Oliver, "Analysis and mitigation of interface losses in trenched superconducting coplanar waveguide resonators", Applied Physics Letters **112** (2017).

[66] S. Ganjam, Y. Wang, Y. Lu, A. Banerjee, C. Lei, L. Krayzman, K. Kisslinger, C. Zhou, R. Li, Y. Jia, M. Liu, L. Frunzio, and R. Schoelkopf, "Surpassing millisecond coherence in on chip superconducting quantum memories by optimizing materials and circuit design", Nature Communications **15** (2024).

[67] D. Zöpfl, *Characterisation of stripline resonators in a waveguide*, Master's thesis, University of Innsbruck (2017).

[68] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, "Observation of high coherence in Josephson junction qubits measured in a three-dimensional circuit QED architecture", Physical Review Letters **107** (2011).

[69] M. Reagor, H. Paik, G. Catelani, L. Sun, C. Axline, E. Holland, I. M. Pop, N. A. Masluk, T. Brecht, L. Frunzio, M. H. Devoret, L. Glazman, and R. J. Schoelkopf, "Reaching 10 ms single photon lifetimes for superconducting aluminum cavities", Appl. Phys. Lett. **102**, 192604 (2013).

[70] P. Heidler, C. M. F. Schneider, K. Kustura, C. Gonzalez-Ballestero, O. Romero-Isart, and G. Kirchmair, "Non-Markovian Effects of Two-Level Systems in a Niobium Coaxial Resonator with a Single-Photon Lifetime of 10 milliseconds", Physical Review Applied **16**, 034024 (2021).

[71] S. Chakram, A. E. Oriani, R. K. Naik, A. V. Dixit, K. He, A. Agrawal, H. Kwon, and D. I. Schuster, “Seamless high-q microwave cavities for multimode circuit quantum electrodynamics”, Phys. Rev. Lett. **127**, 107701 (2021).

[72] O. Milul, B. Guttel, U. Goldblatt, S. Hazanov, L. M. Joshi, D. Chausovsky, N. Kahn, E. Çiftyürek, F. Lafont, and S. Rosenblum, “Superconducting cavity qubit with tens of milliseconds single-photon coherence time”, PRX Quantum **4**, 030336 (2023).

[73] G. Catelani, K. Li, C. Axline, T. Brecht, L. Frunzio, R. Schoelkopf, and L. Glazman, “Ac losses in field-cooled type i superconducting cavities”, Superconductor Science and Technology **35** (2022).

[74] T. Brecht, W. Pfaff, C. Wang, Y. Chu, L. Frunzio, M. Devoret, and R. Schoelkopf, “Multilayer microwave integrated quantum circuits for scalable quantum computing”, npj Quantum Information **2** (2015).

[75] T. Brecht, Y. Chu, C. Axline, W. Pfaff, J. Z. Blumoff, K. Chou, L. Krayzman, L. Frunzio, and R. J. Schoelkopf, “Micromachined integrated quantum circuit containing a superconducting qubit”, Phys. Rev. Appl. **7**, 044018 (2017).

[76] F. Hassani, M. Peruzzo, L. N. Kapoor, A. Trioni, M. Zemlicka, and J. M. Fink, “Inductively shunted transmons exhibit noise insensitive plasmon states and a fluxon decay exceeding 3 hours”, Nature Communications **14**, 3968 (2023).

[77] A. Torras-Coloma, L. Cozzolino, A. Gómez-del Pulgar-Martínez, E. Bertoldo, and P. Forn-Díaz, “Superinductor-based ultrastrong coupling in a superconducting circuit” (2025).

[78] M. Khalifa, P. Feldmann, and J. Salfi, “Kinetic inductance parametric converter”, Phys. Rev. Appl. **22**, 024025 (2024).

[79] N. Zapata, I. Takmakov, S. Günzler, S. Geisert, S. Ihssen, M. Field, A. Nambisan, D. Rieger, T. Reisinger, W. Wernsdorfer, and I. M. Pop, “Granular aluminum parametric amplifier for low-noise measurements in tesla fields”, Phys. Rev. Lett. **133**, 260604 (2024).

[80] M. Peruzzo, A. Trioni, F. Hassani, M. Zemlicka, and J. M. Fink, “Surpassing the resistance quantum with a geometric superinductor”, Phys. Rev. Appl. **14**, 044055 (2020).

[81] S. Singh, G. Refael, A. Clerk, and E. Rosenfeld, “Impact of josephson junction array modes on fluxonium readout” (2025), arXiv:2412.14788 [cond-mat.mes-hall] .

[82] G. Deutscher, H. Fenichel, M. Gershenson, E. Grünbaum, and Z. Ovadyahu, “Transition to zero dimensionality in granular aluminum superconducting films”, Journal of Low Temperature Physics **10**, 231 (1973).

[83] I. Pietka, *Microstrip Resonators made from Granular Aluminium for Microwave Circuit Quantum Electrodynamics*, Master's thesis, Ruhr University Bochum, Bochum (2019).

[84] N. Maleeva, L. Grünhaupt, T. Klein, F. Levy-Bertrand, O. Dupre, M. Calvo, F. Valenti, P. Winkel, F. Friedrich, W. Wernsdorfer, A. V. Ustinov, H. Rotzinger, A. Monfardini, M. V. Fistul, and I. M. Pop, "Circuit quantum electrodynamics of granular aluminum resonators", Nature Communications **9**, 10.1038/s41467-018-06386-9 (2018).

[85] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Charge-insensitive qubit design derived from the cooper pair box", Phys. Rev. A **76**, 042319 (2007).

[86] V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret, "Fluxonium: Single cooper-pair circuit free of charge offsets", Science **326**, 113 (2009).

[87] D. Pines and D. Bohm, "A collective description of electron interactions: Ii. collective vs individual particle aspects of the interactions", Phys. Rev. **85**, 338 (1952).

[88] D. Bohm and D. Pines, "A collective description of electron interactions: Iii. coulomb interactions in a degenerate electron gas", Phys. Rev. **92**, 609 (1953).

[89] F. London, "On the problem of the molecular theory of superconductivity", Phys. Rev. **74**, 562 (1948).

[90] T. K. Hunt and J. E. Mercereau, "Fluxoid conservation by superconducting thin film rings", Phys. Rev. **135**, A944 (1964).

[91] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, "Approaching unit visibility for control of a superconducting qubit with dispersive readout", Phys. Rev. Lett. **95**, 060501 (2005).

[92] J. Q. You, Y. Nakamura, and F. Nori, "Fast two-bit operations in inductively coupled flux qubits", Phys. Rev. B **71**, 024532 (2005).

[93] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, "Circuit quantum electrodynamics", Rev. Mod. Phys. **93**, 025005 (2021).

[94] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, "Circuit quantum electrodynamics", Rev. Mod. Phys. **93**, 025005 (2021).

[95] A. Borgognoni, *Mode temperature analysis of a readout resonator*, Master's thesis, University of Innsbruck (2020).

[96] J. M. Kreikebaum, A. Dove, W. Livingston, E. Kim, and I. Siddiqi, "Optimization of infrared and magnetic shielding of superconducting tin and al coplanar microwave resonators", Superconductor Science and Technology **29**, 104002 (2016).

[97] J. M. Martinis, "Surface loss calculations and design of a superconducting transmon qubit with tapered wiring", npj Quantum Information **8**, 26 (2022).

[98] J. J. Tian, Y. Song, P. Liu, W. G. Zhang, H. F. Yu, and Y. R. Jin, "High-density wiring solution for 500-qubit scale superconducting quantum processors", Review of Scientific Instruments **96**, 104709 (2025).

[99] P. Brooks, A. Kitaev, and J. Preskill, "Protected gates for superconducting qubits", Phys. Rev. A **87**, 052306 (2013).

[100] W. C. Smith, A. Kou, X. Xiao, U. Vool, and M. H. Devoret, "Superconducting circuit protected by two-cooper-pair tunneling", npj Quantum Information **6**, 10.1038/s41534-019-0231-2 (2020).

[101] K. Temme, S. Bravyi, and J. M. Gambetta, "Error mitigation for short-depth quantum circuits", Phys. Rev. Lett. **119**, 180509 (2017).

[102] J. Roffe, "Quantum error correction: an introductory guide", Contemporary Physics **60**, 226–245 (2019).

[103] J. Kelly, R. Barends, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O'Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, A. N. Cleland, and J. M. Martinis, "State preservation by repetitive error detection in a superconducting quantum circuit", Nature **519**, 66 (2015).

[104] Z. Chen, K. J. Satzinger, J. Atalaya, A. N. Korotkov, A. Dunsworth, D. Sank, C. Quintana, M. McEwen, R. Barends, P. V. Klimov, S. Hong, C. Jones, A. Petukhov, D. Kafri, S. Demura, B. Burkett, C. Gidney, A. G. Fowler, A. Paler, H. Putterman, I. Aleiner, F. Arute, K. Arya, R. Babbush, J. C. Bardin, A. Bengtsson, A. Bourassa, M. Broughton, B. B. Buckley, D. A. Buell, N. Bushnell, B. Chiaro, R. Collins, W. Courtney, A. R. Derk, D. Eppens, C. Erickson, E. Farhi, B. Foxen, M. Giustina, A. Greene, J. A. Gross, M. P. Harrigan, S. D. Harrington, J. Hilton, A. Ho, T. Huang, W. J. Huggins, L. B. Ioffe, S. V. Isakov, E. Jeffrey, Z. Jiang, K. Kechedzhi, S. Kim, A. Kitaev, F. Kostritsa, D. Landhuis, P. Laptev, E. Lucero, O. Martin, J. R. McClean, T. McCourt, X. Mi, K. C. Miao, M. Mohseni, S. Montazeri, W. Mruczkiewicz, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Newman, M. Y. Niu, T. E. O'Brien, A. Opremcak, E. Ostby, B. Pató, N. Redd, P. Roushan, N. C. Rubin, V. Shvarts, D. Strain, M. Szalay, M. D. Trevithick, B. Villalonga, T. White, Z. J. Yao, P. Yeh, J. Yoo, A. Zalcman, H. Neven, S. Boixo, V. Smelyanskiy, Y. Chen, A. Megrant, J. Kelly, and Google Quantum AI, "Exponential suppression of bit or phase errors with cyclic error correction", Nature **595**, 383 (2021).

[105] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, "Topological quantum memory", Journal of Mathematical Physics **43**, 4452 (2002).

[106] R. Acharya, D. A. Abanin, L. Aghababaie-Beni, I. Aleiner, T. I. Andersen, and M. e. a. Ansmann, "Quantum error correction below the surface code threshold",

Nature **638**, 920 (2024).

[107] S. Krinner, N. Lacroix, A. Remm, A. Di Paolo, E. Genois, C. Leroux, C. Hellings, S. Lazar, F. Swiadek, J. Herrmann, G. J. Norris, C. K. Andersen, M. Müller, A. Blais, C. Eichler, and A. Wallraff, "Realizing repeated quantum error correction in a distance-three surface code", Nature **605**, 669 (2022).

[108] R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, D. Bacon, J. C. Bardin, J. Basso, A. Bengtsson, S. Boixo, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, D. M. Debroy, A. Del Toro Barba, S. Demura, A. Dunsworth, D. Eppens, C. Erickson, L. Faoro, E. Farhi, R. Fatemi, L. Flores Burgos, E. Forati, A. G. Fowler, B. Foxen, W. Giang, C. Gidney, D. Gilboa, M. Giustina, A. Grajales Dau, J. A. Gross, S. Habegger, M. C. Hamilton, M. P. Harrigan, S. D. Harrington, O. Higgott, J. Hilton, M. Hoffmann, S. Hong, T. Huang, A. Huff, W. J. Huggins, L. B. Ioffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, P. Juhas, D. Kafri, K. Kechedzhi, J. Kelly, T. Khattar, M. Khezri, M. Kieferová, S. Kim, A. Kitaev, P. V. Klimov, A. R. Klots, A. N. Korotkov, F. Kostritsa, J. M. Kreikebaum, D. Landhuis, P. Laptev, K.-M. Lau, L. Laws, J. Lee, K. Lee, B. J. Lester, A. Lill, W. Liu, A. Locharla, E. Lucero, F. D. Malone, J. Marshall, O. Martin, J. R. McClean, T. McCourt, M. McEwen, A. Megrant, B. Meurer Costa, X. Mi, K. C. Miao, M. Mohseni, S. Montazeri, A. Morvan, E. Mount, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, H. Neven, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, M. Y. Niu, T. E. O'Brien, A. Opremcak, J. Platt, A. Petukhov, R. Potter, L. P. Pryadko, C. Quintana, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, M. J. Shearn, A. Shorter, V. Shvarts, J. Skruzny, V. Smelyanskiy, W. C. Smith, G. Sterling, D. Strain, M. Szalay, A. Torres, G. Vidal, B. Villalonga, C. Vollgraff Heidweiller, T. White, C. Xing, Z. J. Yao, P. Yeh, J. Yoo, G. Young, A. Zalcman, Y. Zhang, N. Zhu, and Google Quantum AI, "Suppressing quantum errors by scaling a surface code logical qubit", Nature **614**, 676 (2023).

[109] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: Towards practical large-scale quantum computation", Phys. Rev. A **86**, 032324 (2012).

[110] Y. Liu, S. Singh, K. C. Smith, E. Crane, J. M. Martyn, A. Eickbusch, A. Schuckert, R. D. Li, J. Sinanan-Singh, M. B. Soley, T. Tsunoda, I. L. Chuang, N. Wiebe, and S. M. Girvin, "Hybrid oscillator-qubit quantum processors: Instruction set architectures, abstract machine models, and applications", PRX Quantum **7**, 010201 (2026).

[111] M. H. Michael, M. Silveri, R. T. Brierley, V. V. Albert, J. Salmilehto, L. Jiang, and S. M. Girvin, "New class of quantum error-correcting codes for a bosonic mode", Phys. Rev. X **6**, 031006 (2016).

[112] E. Wigner, "On the quantum correction for thermodynamic equilibrium", Phys. Rev. **40**, 749 (1932).

[113] W. Cai, M. Yuwei, W. Wang, C.-L. Zou, and L. Sun, "Bosonic quantum error correction codes in superconducting quantum circuits", Fundamental Research **1** (2021).

[114] L. Hu, Y. Ma, W. Cai, X. Mu, Y. Xu, W. Wang, Y. Wu, H. Wang, Y. P. Song, C.-L. Zou, S. M. Girvin, L.-M. Duan, and L. Sun, "Quantum error correction and universal gate set operation on a binomial bosonic logical qubit", Nature Physics **15**, 503 (2019).

[115] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, "Extending the lifetime of a quantum bit with error correction in superconducting circuits", Nature **536**, 441 (2016).

[116] A. Grimm, N. E. Frattini, S. Puri, S. O. Mundhada, S. Touzard, M. Mirrahimi, S. M. Girvin, S. Shankar, and M. H. Devoret, "Stabilization and operation of a kerr-cat qubit", Nature **584**, 205 (2020).

[117] A. Hajr, B. Qing, K. Wang, G. Koolstra, Z. Pedramrazi, Z. Kang, L. Chen, L. B. Nguyen, C. Jünger, N. Goss, I. Huang, B. Bhandari, N. E. Frattini, S. Puri, J. Dressel, A. N. Jordan, D. I. Santiago, and I. Siddiqi, "High-coherence kerr-cat qubit in 2d architecture", Phys. Rev. X **14**, 041049 (2024).

[118] N. E. Frattini, R. G. Cortiñas, J. Venkatraman, X. Xiao, Q. Su, C. U. Lei, B. J. Chapman, V. R. Joshi, S. Girvin, R. J. Schoelkopf, S. Puri, and M. H. Devoret, "Observation of pairwise level degeneracies and the quantum regime of the arrhenius law in a double-well parametric oscillator", Physical Review X **14**, 10.1103/physrevx.14.031040 (2024).

[119] F. Adinolfi, D. Z. Haxell, A. Bruno, L. Michaud, V. H. Kamrul, P. Pandey, and A. Grimm, "Enhancing kerr-cat qubit coherence with controlled dissipation" (2025), arXiv:2511.01027 [quant-ph] .

[120] Z. Leghtas, S. Touzard, I. M. Pop, A. Kou, B. Vlastakis, A. Petrenko, K. M. Sliwa, A. Narla, S. Shankar, M. J. Hatridge, M. Reagor, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret, "Confining the state of light to a quantum manifold by engineered two-photon loss", Science **347**, 853.

[121] R. Lescanne, M. Villiers, T. Peronnin, A. Sarlette, M. Delbecq, B. Huard, T. Kontos, M. Mirrahimi, and Z. Leghtas, "Exponential suppression of bit-flips in a qubit encoded in an oscillator", Nature Physics **16**, 509 (2020).

[122] A. Marquet, A. Essig, J. Cohen, N. Cottet, A. Murani, E. Albertinale, S. Dupouy, A. Bienfait, T. Peronnin, S. Jezouin, R. Lescanne, and B. Huard, "Autoparametric resonance extending the bit-flip time of a cat qubit up to 0.3 s", Phys. Rev. X **14**, 021019 (2024).

[123] A. Marquet, A. Essig, J. Cohen, N. Cottet, A. Murani, E. Albertinale, S. Dupouy, A. Bienfait, T. Peronnin, S. Jezouin, R. Lescanne, and B. Huard, "Autoparametric

resonance extending the bit-flip time of a cat qubit up to 0.3 s", Phys. Rev. X **14**, 021019 (2024).

[124] D. Gottesman, A. Kitaev, and J. Preskill, "Encoding a qubit in an oscillator", Phys. Rev. A **64**, 012310 (2001).

[125] B. Royer, S. Singh, and S. Girvin, "Stabilization of finite-energy gottesman-kitaev-preskill states", Physical Review Letters **125**, 10.1103/physrevlett.125.260509 (2020).

[126] P. Campagne-Ibarcq, A. Eickbusch, S. Touzard, E. Zalys-Geller, N. E. Frattini, V. V. Sivak, P. Reinhold, S. Puri, S. Shankar, R. J. Schoelkopf, L. Frunzio, M. Mirrahimi, and M. H. Devoret, "Quantum error correction of a qubit encoded in grid states of an oscillator", Nature **584**, 368 (2020).

[127] C. Flühmann, T. L. Nguyen, M. Marinelli, V. Negnevitsky, K. Mehta, and J. P. Home, "Encoding a qubit in a trapped-ion mechanical oscillator", Nature **566**, 513 (2019).

[128] B. L. Brock, S. Singh, A. Eickbusch, V. V. Sivak, A. Z. Ding, L. Frunzio, S. M. Girvin, and M. H. Devoret, "Quantum error correction of qudits beyond break-even", Nature **641**, 612 (2025).

[129] K. Noh and C. Chamberland, "Fault-tolerant bosonic quantum error correction with the surface–gottesman-kitaev-preskill code", Phys. Rev. A **101**, 012316 (2020).

[130] W. C. Smith, A. Kou, U. Vool, I. M. Pop, L. Frunzio, R. J. Schoelkopf, and M. H. Devoret, "Quantization of inductively shunted superconducting circuits", Phys. Rev. B **94**, 144507 (2016).

[131] D. Gusenkova, M. Spiecker, R. Gebauer, M. Willsch, D. Willsch, F. Valenti, N. Karcher, L. Grünhaupt, I. Takmakov, P. Winkel, D. Rieger, A. V. Ustinov, N. Roch, W. Wernsdorfer, K. Michielsen, O. Sander, and I. M. Pop, "Quantum nondemolition dispersive readout of a superconducting artificial atom using large photon numbers", Phys. Rev. Applied **15**, 064030 (2021).

[132] R. M. Foster, "A reactance theorem", Bell System Technical Journal **3**, 259 (1924).

[133] D. A. Rower, L. Ding, H. Zhang, M. Hays, J. An, P. M. Harrington, I. T. Rosen, J. M. Gertler, T. M. Hazard, B. M. Niedzielski, M. E. Schwartz, S. Gustavsson, K. Serniak, J. A. Grover, and W. D. Oliver, "Suppressing counter-rotating errors for fast single-qubit gates with fluxonium", PRX Quantum **5**, 040342 (2024).

[134] M. F. S. Zwanenburg, S. Singh, E. Y. Huang, F. Yilmaz, T. V. Stefanski, J. Hu, P. Kumaravadivel, and C. K. Andersen, "Single-qubit gates beyond the rotating-wave approximation for strongly anharmonic low-frequency qubits", Phys. Rev. Res. **7**, 043290 (2025).

[135] Z. Minev, Z. Leghtas, S. Mundhada, L. Christakis, I. M. Pop, and M. Devoret, "Energy-participation quantization of josephson circuits", npj Quantum Inf **7** (2021).

[136] F. Yilmaz, S. Singh, M. F. S. Zwanenburg, J. Hu, T. V. Stefanski, and C. K. Andersen, "Energy participation ratio analysis for very anharmonic superconducting circuits" (2024), arXiv:2411.15039 [quant-ph] .

[137] P. Groszkowski and J. Koch, "Scqubits: a Python package for superconducting qubits", Quantum **5**, 583 (2021).

[138] J. Blumoff, "Multiqubit experiments in 3D circuit quantum electrodynamics", Ph.D. thesis, Yale University (2017).

[139] I. Yang, T. Agrenius, V. Usova, O. Romero-Isart, and G. Kirchmair, "Hot schrödinger cat states", Science Advances **11**, 10.1126/sciadv.adr4492 (2025).

[140] "Aluminum etchants | transene, https://transene.com/aluminum/", (accessed: 01.09.2025).

[141] L. Grünhaupt, M. Spiecker, D. Gusenkova, N. Maleeva, S. T. Skacel, I. Takmakov, F. Valenti, P. Winkel, H. Rotzinger, A. V. Ustinov, and I. M. Pop, "Granular aluminum: A superconducting material for high impedance quantum circuits", Nat. Mater. **18**, 816 (2019).

[142] "Beamer - genisys gmbh, https://www.genisys-gmbh.com/beamer.html", (accessed: 01.09.2025).

[143] G. J. Dolan, "Offset masks for lift-off photoprocessing", Appl. Phys. Lett.; (United States) **31:5**, 10.1063/1.89690 (1977).

[144] "Laser cutting and drilling machines | coherent, https://www.coherent.com/machines-systems/laser-cutting-drilling", (accessed: 03.09.2025).

[145] V. Ambegaokar and A. Baratoff, "Tunneling between superconductors", Phys. Rev. Lett. **10**, 486 (1963).

[146] C. Wang, C. Axline, Y. Y. Gao, T. Brecht, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf, "Surface participation and dielectric loss in superconducting qubits", Applied Physics Letters **107**, 10.1063/1.4934486 (2015).

[147] Y. Y. Gao, M. A. Rol, S. Touzard, and C. Wang, "Practical guide for building superconducting quantum devices", PRX Quantum **2**, 040202 (2021).

[148] L. J. Zeng, S. Nik, T. Greibe, P. Krantz, C. M. Wilson, P. Delsing, and E. Olsson, "Direct observation of the thickness distribution of ultra thin AlO_x barriers in Al/AlO_x /Al Josephson junctions", Journal of Physics D: Applied Physics **48**, 395308 (2015).

[149] B. J. Chapman, S. J. de Graaf, S. H. Xue, Y. Zhang, J. Teoh, J. C. Curtis, T. Tsunoda, A. Eickbusch, A. P. Read, A. Koottandavida, S. O. Mundhada, L. Frunzio, M. Devoret, S. Girvin, and R. Schoelkopf, "High-on-off-ratio beam-splitter interaction for gates on bosonically encoded qubits", PRX Quantum **4**, 020355 (2023).

[150] U. Reglade, "Quantum control of a cat-qubit with macroscopic bit-flip times", Ph.D. thesis, École Normale Supérieure (2024).

[151] O. Gargiulo, S. Oleschko, J. Prat-Camps, M. Zanner, and G. Kirchmair, "Fast flux control of 3D transmon qubits using a magnetic hose", Applied Physics Letters **118**, 012601 (2021).

[152] C. Navau, J. Prat-Camps, O. Romero-Isart, J. I. Cirac, and A. Sanchez, "Long-distance transfer and routing of static magnetic fields", Phys. Rev. Lett. **112**, 253901 (2014).

[153] I. Yang, "Quantum information processing with superconducting circuits and high q cavities", Ph.D. thesis, University of Innsbruck (2024).

[154] L. Québaud, *Characterization of the Purcell limitation of a transmon qubit for implementing a Purcell filter*, Master's thesis, Université Grenoble Alpes (2023).

[155] N. J. Bort-Soldevila, "Magnetic field manipulation for cutting-edge technologies", Ph.D. thesis, Universitat Autònoma de Barcelona (2025).

[156] P. Straub, *Development and Testing of a High-Q Magnetic Flux Hose*, Master's thesis, University of Innsbruck (2025).

[157] D. G. Atanasova*, P. Straub*, L. Quebaud, and G. Kirchmair, "The superconducting flux hose for qubit tuning and readout" (2026), in preparation.

[158] J. Herrmann, C. Hellings, S. Lazar, F. Pfäffli, F. Haupt, T. Thiele, D. C. Zanuz, G. J. Norris, F. Heer, C. Eichler, and A. Wallraff, "Frequency up-conversion schemes for controlling superconducting qubits" (2022), arXiv:2210.02513 [quant-ph] .

[159] P. Winkel, I. Takmakov, D. Rieger, L. Planat, W. Hasch-Guichard, L. Grünhaupt, N. Maleeva, F. Foroughi, F. Henriques, K. Borisov, J. Ferrero, A. V. Ustinov, W. Wernsdorfer, N. Roch, and I. M. Pop, "Nondegenerate parametric amplifiers based on dispersion-engineered josephson-junction arrays", Phys. Rev. Appl. **13**, 024015 (2020).

[160] D. G. Atanasova, *Fabrication and Characterization of a Dimer Josephson Junction Array Amplifier*, Master's thesis, University of Innsbruck (2020).

[161] S. Diamond, V. Fatemi, M. Hays, H. Nho, P. D. Kurilovich, T. Connolly, V. R. Joshi, K. Serniak, L. Frunzio, L. I. Glazman, and M. H. Devoret, "Distinguishing parity-switching mechanisms in a superconducting qubit", PRX Quantum **3**, 040304 (2022).

[162] M. B. Heaney, in *Electrical Measurement, Signal Processing, and Displays*, edited by J. G. Webster (CRC Press, 2003) Chap. 7, pp. 7–1–7–14.

[163] M. Sasaki, M. Ohkuma, R. Matsumoto, T. Shinmei, T. Irifune, Y. Takano, and K. Shimizu, "Enhancement of superconductivity in thin films of sn under high pressure", Phys. Rev. B **111**, 104513 (2025).

[164] R. Gross and A. Marx, "Festkörperphysik", 4th ed. (De Gruyter Oldenbourg, Berlin/Boston, 2022).

[165] D. Atanasova, I. Yang, T. Hönigl-Decrinis, D. Gusenkova, I. Pop, and G. Kirchmair, "In situ tunable interaction with an invertible sign between a fluxonium and a post cavity", PRX Quantum **6**, 020318 (2025).

[166] M. McEwen, K. C. Miao, J. Atalaya, A. Bilmes, A. Crook, J. Bovaird, J. M. Kreikebaum, N. Zobrist, E. Jeffrey, B. Ying, A. Bengtsson, H.-S. Chang, A. Dunsworth, J. Kelly, Y. Zhang, E. Forati, R. Acharya, J. Iveland, W. Liu, S. Kim, B. Burkett, A. Megrant, Y. Chen, C. Neill, D. Sank, M. Devoret, and A. Opremcak, "Resisting high-energy impact events through gap engineering in superconducting qubit arrays", Phys. Rev. Lett. **133**, 240601 (2024).

[167] X. Y. Jin, A. Kamal, A. P. Sears, T. Gudmundsen, D. Hover, J. Miloshi, R. Slattery, F. Yan, J. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver, "Thermal and residual excited-state population in a 3d transmon qubit", Phys. Rev. Lett. **114**, 240501 (2015).

[168] I. M. Pop, K. Geerlings, G. Catelani, R. J. Schoelkopf, L. I. Glazman, and M. H. Devoret, "Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles", Nature **508**, 369 (2014).

[169] W. Smith, A. Kou, X. Xiao, and M. Devoret, "Superconducting circuit protected by two-cooper-pair tunneling", npj Quantum Inf **6** (2020).

[170] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, "Introduction to quantum noise, measurement, and amplification", Rev. Mod. Phys. **82**, 1155 (2010).

[171] H. Zhang, S. Chakram, T. Roy, N. Earnest, Y. Lu, Z. Huang, D. K. Weiss, J. Koch, and D. I. Schuster, "Universal fast-flux control of a coherent, low-frequency qubit", Phys. Rev. X **11**, 011010 (2021).

[172] G. Catelani, J. Koch, L. Frunzio, R. J. Schoelkopf, M. H. Devoret, and L. I. Glazman, "Quasiparticle relaxation of superconducting qubits in the presence of flux", Phys. Rev. Lett. **106**, 077002 (2011).

[173] R. Hanai, A. McDonald, and A. Clerk, "Intrinsic mechanisms for drive-dependent purcell decay in superconducting quantum circuits", Phys. Rev. Res. **3**, 043228 (2021).

[174] V. B. Braginsky, V. S. Ilchenko, and K. S. Bagdassarov, “Experimental observation of fundamental microwave absorption in high-quality dielectric crystals”, Physics Letters A **120**, 300 (1987).

[175] E. M. Purcell, “Spontaneous Emission Probabilities at Radio Frequencies”, Physical Review **69**, 681 (1946).

[176] P. Paluch, “Dissipation and noise in granular aluminum fluxonium qubits”, Ph.D. thesis, Karlsruhe Institute of Technology (2025).

[177] L. Chayanun, J. Biznárová, L. Zeng, P. Malmberg, A. Nylander, A. Osman, M. Rommel, P. L. Tam, E. Olsson, P. Delsing, A. Yurgens, J. Bylander, and A. Fadavi Roudsari, “Characterization of process-related interfacial dielectric loss in aluminum-on-silicon by resonator microwave measurements, materials analysis, and imaging”, APL Quantum **1**, 026115 (2024).

[178] A. P. M. Place, L. V. H. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman, G. Cheng, T. Madhavan, H. K. Babla, X. H. Le, Y. Gang, B. Jäck, A. Gyenis, N. Yao, R. J. Cava, N. P. de Leon, and A. A. Houck, “New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds”, Nat Commun **12**, 1779 (2021).

[179] C. Wang, X. Li, H. Xu, Z. Li, J. Wang, Z. Yang, Z. Mi, X. Liang, T. Su, C. Yang, G. Wang, W. Wang, Y. Li, M. Chen, C. Li, K. Linghu, J. Han, Y. Zhang, Y. Feng, Y. Song, T. Ma, J. Zhang, R. Wang, P. Zhao, W. Liu, G. Xue, Y. Jin, and H. Yu, “Towards practical quantum computers: transmon qubit with a lifetime approaching 0.5 milliseconds”, npj Quantum Inf **8**, 1 (2022).

[180] M. Bal, A. A. Murthy, S. Zhu, F. Crisa, X. You, Z. Huang, T. Roy, J. Lee, D. v. Zanten, R. Pilipenko, I. Nekrashevich, A. Lunin, D. Bafia, Y. Krasnikova, C. J. Kopas, E. O. Lachman, D. Miller, J. Y. Mutus, M. J. Reagor, H. Cansizoglu, J. Marshall, D. P. Pappas, K. Vu, K. Yadavalli, J.-S. Oh, L. Zhou, M. J. Kramer, F. Lecocq, D. P. Goronzy, C. G. Torres-Castanedo, P. G. Pritchard, V. P. Dravid, J. M. Rondinelli, M. J. Bedzyk, M. C. Hersam, J. Zasadzinski, J. Koch, J. A. Sauls, A. Romanenko, and A. Grassellino, “Systematic improvements in transmon qubit coherence enabled by niobium surface encapsulation”, npj Quantum Inf **10**, 1 (2024).

[181] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J.-S. Tsai, and W. D. Oliver, “Noise spectroscopy through dynamical decoupling with a superconducting flux qubit”, Nature Physics **7**, 565 (2011).

[182] Z. Leghtas, G. Kirchmair, B. Vlastakis, M. H. Devoret, R. J. Schoelkopf, and M. Mirrahimi, “Deterministic protocol for mapping a qubit to coherent state superpositions in a cavity”, Physical Review A **87**, 042315 (2013).

[183] A. Eickbusch, V. Sivak, A. Z. Ding, S. S. Elder, S. R. Jha, J. Venkatraman, B. Royer, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, “Fast universal control of an

oscillator with weak dispersive coupling to a qubit", Nat. Phys. **18**, 1464 (2022).

[184] K. S. Chou, "Teleported operations between logical qubits in circuit", Ph.D. thesis, Yale University (2018).

[185] A. Royer, "Wigner function as the expectation value of a parity operator", Phys. Rev. A **15**, 449 (1977).

[186] K. HUSIMI, "Some formal properties of the density matrix", Proceedings of the Physico-Mathematical Society of Japan. 3rd Series **22**, 264 (1940).

[187] F. Valadares, N.-N. Huang, K. T. N. Chu, A. Dorogov, W. Chua, L. Kong, P. Song, and Y. Y. Gao, "On-demand transposition across light-matter interaction regimes in bosonic cQED", Nat Commun **15**, 5816 (2024).

[188] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O'Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, M. R. Geller, A. N. Cleland, and J. M. Martinis, "Qubit architecture with high coherence and fast tunable coupling", Phys. Rev. Lett. **113**, 220502 (2014).

[189] S. Rosenblum, P. Reinhold, M. Mirrahimi, L. Jiang, L. Frunzio, and R. J. Schoelkopf, "Fault-tolerant detection of a quantum error", Science **361**, 266 (2018).

[190] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Córcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen, "Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms", Phys. Rev. B **86**, 100506 (2012).

[191] V. Usova, "Adaptive reconstruction for closed-loop optimisation of bosonic quantum states", Ph.D. thesis, University of Innsbruck (2025).

[192] G. Kirchmair, B. Vlastakis, Z. Leghtas, S. E. Nigg, H. Paik, E. Ginossar, M. Mirrahimi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, "Observation of quantum state collapse and revival due to the single-photon Kerr effect", Nature **495**, 205 (2013).

[193] N. E. Frattini, "Three-wave mixing in superconducting circuits: Stabilizing cats with snails", Ph.D. thesis, Yale University (2021).

[194] J. T. Peltonen, P. C. J. J. Coumou, Z. H. Peng, T. M. Klapwijk, J. S. Tsai, and O. V. Astafiev, "Hybrid rf SQUID qubit based on high kinetic inductance", Scientific Reports **8**, 10033 (2018).

[195] B. Abdo, A. Kamal, and M. Devoret, "Nondegenerate three-wave mixing with the josephson ring modulator", Phys. Rev. B **87**, 014508 (2013).

[196] M. Bild, M. Fadel, Y. Yang, U. von Lüpke, P. Martin, A. Bruno, and Y. Chu, "Schrödinger cat states of a 16-microgram mechanical oscillator", Science **380**, 274 (2023).

[197] R. G. Lyons, "Understanding digital signal processing", 3rd ed. (Pearson Education, Upper Saddle River, NJ, 2010).

[198] Zurich Instruments AG, *HDAWG User Manual*, Zurich Instruments AG, Technoparkstrasse 1, 8005 Zurich, Switzerland (2025).

[199] M. A. Rol, L. Ciorciaro, F. K. Malinowski, B. M. Tarasinski, R. E. Sagastizabal, C. C. Bultink, Y. Salathe, N. Haandbaek, J. Sedivy, and L. DiCarlo, "Time-domain characterization and correction of on-chip distortion of control pulses in a quantum processor", Applied Physics Letters **116**, 10.1063/1.5133894 (2020).

[200] C. K. Law and J. H. Eberly, "Arbitrary control of a quantum electromagnetic field", Phys. Rev. Lett. **76**, 1055 (1996).

[201] M. Hofheinz, H. Wang, M. Ansmann, R. C. Bialczak, E. Lucero, M. Neeley, A. D. O'Connell, D. Sank, J. Wenner, J. M. Martinis, and A. N. Cleland, "Synthesizing arbitrary quantum states in a superconducting resonator", Nature **459**, 546 (2009).

[202] J. Johansson, P. Nation, and F. Nori, "Qutip 2: A python framework for the dynamics of open quantum systems", Computer Physics Communications **184**, 1234 (2013).

[203] K. Nie, J. N. Bradford, S. Mandal, A. Bista, W. Pfaff, and A. Kou, "Fluxonium as a control qubit for bosonic quantum information" (2025), arXiv:2505.23641 [quant-ph] .

[204] M. Bild, M. Fadel, Y. Yang, U. Von Lüpke, P. Martin, A. Bruno, and Y. Chu, "Schrödinger cat states of a 16-microgram mechanical oscillator", Science **380**, 274 (2023).

[205] P. DeVore, D. Escontrias, T. Koblesky, C. J. Lin, D. W. Liu, K. B. Luk, J. Ngan, J. C. Peng, C. Polly, J. Roloff, H. Steiner, S. Wang, J. Wong, and M. Yeh, "Light-weight flexible magnetic shields for large-aperture photomultiplier tubes", Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **737**, 222 (2014).

[206] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, "Laser phase and frequency stabilization using an optical resonator", Applied Physics B **31**, 97 (1983).

[207] C. Sparrow, E. Martín-López, N. Maraviglia, A. Neville, C. Harrold, J. Carolan, Y. N. Joglekar, T. Hashimoto, N. Matsuda, J. L. O'Brien, D. P. Tew, and A. Laing, "Simulating the vibrational quantum dynamics of molecules using photonics", Nature **557**, 660 (2018).

[208] W. R. Clements, J. J. Renema, A. Eckstein, A. A. Valido, A. Lita, T. Gerrits, S. W. Nam, W. S. Kolthammer, J. Huh, and I. A. Walmsley, "Approximating vibronic spectroscopy with imperfect quantum optics", Journal of Physics B: Atomic, Molecular and Optical Physics **51**, 245503 (2018).

[209] W. Wang, Y. Wu, Y. Ma, W. Cai, L. Hu, X. Mu, Y. Xu, Z.-J. Chen, H. Wang, Y. P. Song, H. Yuan, C.-L. Zou, L.-M. Duan, and L. Sun, "Heisenberg-limited single-mode quantum metrology in a superconducting circuit", Nature Communications **10**, 4382 (2019).

[210] Q. Zhuang, Z. Zhang, and J. H. Shapiro, "Distributed quantum sensing using continuous-variable multipartite entanglement", Phys. Rev. A **97**, 032329 (2018).

[211] Y. Xia, W. Li, W. Clark, D. Hart, Q. Zhuang, and Z. Zhang, "Demonstration of a reconfigurable entangled radio-frequency photonic sensor network", Phys. Rev. Lett. **124**, 150502 (2020).

[212] Q. Zhuang, J. Preskill, and L. Jiang, "Distributed quantum sensing enhanced by continuous-variable error correction", New Journal of Physics **22**, 022001 (2020).

[213] Y. Ma, Y. Xu, X. Mu, W. Cai, L. Hu, W. Wang, X. Pan, H. Wang, Y. P. Song, C.-L. Zou, and L. Sun, "Error-transparent operations on a logical qubit protected by quantum error correction", Nature Physics **16**, 827 (2020).

[214] P. Reinhold, S. Rosenblum, W.-L. Ma, L. Frunzio, L. Jiang, and R. J. Schoelkopf, "Error-corrected gates on an encoded qubit", Nature Physics **16**, 822 (2020).

[215] J. Guillaud and M. Mirrahimi, "Repetition cat qubits for fault-tolerant quantum computation", Phys. Rev. X **9**, 041053 (2019).

[216] A. Paquette, J. Griesmar, G. Lavoie, R. Albert, F. Blanchet, A. Grimm, U. Martel, and M. Hofheinz, "Absorptive filters for quantum circuits: Efficient fabrication and cryogenic power handling", Applied Physics Letters **121**, 10.1063/5.0114887 (2022).

APPENDIX A

System parameters

Table A.1: System Parameters. The cavity frequency $\omega_c/2\pi$ has an additional uncertainty of 400 kHz coming from the 15% variation in g_{RC}. Due to the dephasing at $\Phi_{\mathrm{ext}}/\Phi_0 = 0.5$ caused by the excited qubit population, the cavity coupling rate and linewidth were not measured at this flux bias.

Parameter	Symbol	Value
Josephson Energy	E_J/h	10.8 GHz
Inductive Energy	E_L/h	1.014 GHz
Capacitive Energy	E_C/h	3.5 GHz
Bare High Q Cavity Frequency	$\omega_{\mathrm{c}}/2\pi$	4.535854 GHz
Bare Readout Resonator Frequency	$\omega_{\mathrm{r}}/2\pi$	6.8176 GHz
Qubit-Resonator Coupling	$g_{\mathrm{QR}}/2\pi$	25.2 MHz ±10%
Resonator-High Q Cavity Coupling	$g_{\mathrm{RC}}/2\pi$	8 MHz±15%
High Q Cavity Lifetime at $\Phi_{\mathrm{ext}}/\Phi_0 = 0.5$	T_1^C	210 ± 40 µs
Qubit Residual Mode Temperature at $\Phi_{\mathrm{ext}}/\Phi_0 = 0.5$	T_{Q}	112 ± 2 mK
Qubit Excited State Population at $\Phi_{\mathrm{ext}}/\Phi_0 = 0.5$	P_e	39.1 ± 0.7 %
Resonator Residual Mode Temperature at $\Phi_{\mathrm{ext}}/\Phi_0 = 0.5$	T_{R}	110 ± 30 mK
Readout Resonator Linewidth at $\Phi_{\mathrm{ext}}/\Phi_0 = 0.5$	$\kappa_R/2\pi$	922 ± 3 kHz
Readout Resonator External Coupling Rate at $\Phi_{\mathrm{ext}}/\Phi_0 = 0.5$	$\kappa_R^{\mathrm{ext}}/2\pi$	549 ± 2 kHz
Readout Resonator Linewidth at $\Phi_{\mathrm{ext}}/\Phi_0 = 0$	$\kappa_R/2\pi$	956 ± 4 kHz
Readout Resonator External Coupling Rate at $\Phi_{\mathrm{ext}}/\Phi_0 = 0$	$\kappa_R^{\mathrm{ext}}/2\pi$	801 ± 2 kHz
Cavity Linewidth at $\Phi_{\mathrm{ext}}/\Phi_0 = 0$	$\kappa_C/2\pi$	550 ± 20 Hz
Cavity External Coupling Rate at $\Phi_{\mathrm{ext}}/\Phi_0 = 0$	$\kappa_C^{\mathrm{ext}}/2\pi$	233 ± 7 Hz
Dielectric quality factor at 6 GHz	Q_{diel}	$\gtrsim 0.15 \cdot 10^6$
Inductive quality factor	Q_{ind}	$\gtrsim 30 \cdot 10^6$
Non-equilibrium quasi-particle density	x_{qp}	$\lesssim 1.25 \cdot 10^{-6}$

APPENDIX B

Engineering dissipation

This appendix follows the lecture given by Prof. Zaki Leghtas at the Les Houches summer school in 2019. To engineer a certain type of dissipation Hamiltonian, we start with the Hamiltonian of a driven harmonic oscillator, which has some non-linearity, represented by the function $f(b^\dagger, b)$

$$\frac{H}{\hbar} = \omega_b b^\dagger b + (\epsilon_p e^{-i\omega_p t} + \epsilon_p^* e^{i\omega_p t})(b^\dagger + b) + f(b^\dagger, b). \tag{B.1}$$

where $b^\dagger(b)$ is the respective creation (annihilation) operator, ω_b is the oscillator frequency, and the drive has amplitude ϵ_p and frequency ω_p. This oscillator is coupled to the environment with a rate κ_b[1]. To reveal the dynamics of this system, we derive the quantum Langevin equation for the operator b

$$\dot{b}(t) = -\frac{i}{\hbar}[b, H] - \frac{\kappa_b}{2} b = -i\omega_b b - i(\epsilon_p e^{-i\omega_p t} + \epsilon_p^* e^{i\omega_p t}) - i[b, f(b^\dagger, b)] - \frac{\kappa_b}{2} b \tag{B.2}$$

where we have used the identity $[b, H] = \partial H(b, b^\dagger)/\partial b^\dagger$.
We undertake a couple of transformations to simplify Eq. B.2. In the displaced frame $b = \bar{b}(t) + \beta(t)$, where $\beta(t) \in \mathbb{C}$ is the classical response of the mode to the drive, Eq. B.2 changes to

$$\dot{\bar{b}} + \dot{\beta} = -i\omega_b(\bar{b} + \beta) - i(\epsilon_p e^{-i\omega_p t} + \epsilon_p^* e^{i\omega_p t}) - i[\bar{b} + \beta, f(\bar{b}^\dagger + \beta^*, \bar{b} + \beta)] - \frac{\kappa_b}{2}(\bar{b} + \beta). \tag{B.3}$$

We impose that

$$\dot{\beta} = -i\omega_b \beta - i(\epsilon_p e^{-i\omega_p t} + \epsilon_p^* e^{i\omega_p t}) - \frac{\kappa_b}{2}\beta \tag{B.4}$$

to cancel the term proportional to the drive. The response of the mode $\beta(t)$ converges, after a transient regime over a time-scale of order $1/\kappa_b$, to a solution of the type $\beta = \beta_p e^{-i\omega_p t}$, where

$$\beta_p = -\frac{i\epsilon_p}{i(\omega_b - \omega_p) + \frac{\kappa_b}{2}}. \tag{B.5}$$

This ratio provides a measure of the drive strength, its detuning from the oscillator, and the oscillator's loss.

In the rotated frame $\bar{b} = \tilde{b} e^{-i\omega_b t}$, Eq. B.3 is transformed to

$$\dot{\tilde{b}} e^{-i\omega_b t} - i\omega_b \tilde{b} e^{-i\omega_b t} = -i\omega_b \tilde{b} e^{-i\omega_b t} - i[\tilde{b} e^{-i\omega_b t}, f(\tilde{b} e^{-i\omega_b t} + \beta, \tilde{b}^\dagger e^{i\omega_b t} + \beta^*)] - \frac{\kappa_b}{2}\tilde{b} e^{-i\omega_b t}. \tag{B.6}$$

[1] Equivalent to κ_{ext}.

All the terms with exponents cancel, as well as the second term on the left side of the equal sign and the first term on the right side of the equal sign. The resulting Hamiltonian in this displaced, rotated frame is

$$\frac{\tilde{H}}{\hbar} = f(\tilde{b}e^{-i\omega_b t} + \beta, \tilde{b}^\dagger e^{i\omega_b t} + \beta^*). \tag{B.7}$$

This procedure is used to eliminate the pump term as well as the free Hamiltonian so the exact form of the two-photon exchange rate g_2 can be extracted, as done in Ref. [121]. In the case of a SNAIL coupled to a cavity, no pump tone is applied due to the resonant condition $\omega_b = 2\omega_m$.

Adiabatic elimination The following section provides a recipe for engineering an arbitrary loss operator. We introduce an additional oscillator with creation operator $m^\dagger$, which couples only to the previously introduced harmonic oscillator with a strength g, but not to the environment. For clarity, the new mode will be our memory, while the old one–the buffer. We assume their interaction Hamiltonian is

$$H/\hbar = gMb^\dagger + g^* M^\dagger b \tag{B.8}$$

where the memory operator that couples to the buffer is $M = f(m^\dagger, m)$.

We extract the dynamics of this coupled system by writing the Master equation

$$\dot{\rho} = -i[H,\rho] + \mathcal{D}(\sqrt{\kappa_b}b)\rho = -i[H,\rho] + \kappa_b\left(b\rho b^\dagger - \frac{1}{2}\{b^\dagger b, \rho\}\right) \tag{B.9}$$

where ρ is the density matrix of this system, and the second term represents the single photon loss of the buffer. In the limit, where the buffer is always close to a vacuum, i.e., $\frac{|g_2|}{\kappa_b} = \delta \ll 1$, we look for a solution of the Master equation of the type

$$\begin{aligned}\rho(t) = &\rho_{00} \otimes |0\rangle_b\langle 0|_b + \\ &\delta(\rho_{10} \otimes |1\rangle_b\langle 0|_b + \rho_{01} \otimes |0\rangle_b\langle 1|_b) + \\ &\delta^2(\rho_{11} \otimes |1\rangle_b\langle 1|_b + \rho_{20} \otimes |2\rangle_b\langle 0|_b + \rho_{02} \otimes |0\rangle_b\langle 2|_b) + \mathcal{O}(\delta^3)\end{aligned} \tag{B.10}$$

where ρ_{ij} is the memory part of ρ. By tracing out the buffer

$$\rho_m = \mathrm{Tr}_b[\rho] = \sum_{i\in\mathbb{N}} {}_b\langle i|\rho|i\rangle_b = \rho_{00} + \mathcal{O}(\delta) \tag{B.11}$$

we see that the dynamics of the memory depends only on the term ρ_{00}.

We project the Master equation into the vacuum state of the buffer

$${}_b\langle 0|\dot{\rho}(t)|0\rangle_b = \dot{\rho}_{00} = -i\delta(g^* M^\dagger \rho_{10} - g\rho_{01}M) + \delta^2\kappa_b\rho_{11} + \mathcal{O}(\delta^3) \tag{B.12}$$

and similarly extract the rest of the terms

$$\dot{\rho}_{01} = i\frac{g^*}{\delta}\rho_{00}M^\dagger - \frac{\kappa_b}{2}\rho_{01} + \mathcal{O}(\delta) \tag{B.13a}$$

$$\dot{\rho}_{10} = -i\frac{g^*}{\delta}M\rho_{00} - \frac{\kappa_b}{2}\rho_{10} + \mathcal{O}(\delta) \tag{B.13b}$$

$$\dot{\rho}_{11} = \frac{1}{\delta}i(g^*\rho_{10}M^\dagger - gM\rho_{01}) - \kappa_b\rho_{11} + \mathcal{O}(\delta^2). \tag{B.13c}$$

We simplify the equations B.12 and B.13a by substituting $\bar{M} = \frac{gM}{\delta\kappa_b}$ into them

$$\frac{1}{\kappa_b}\dot{\rho}_{00} = -i\delta(g^*\bar{M}^\dagger\rho_{10} - g\rho_{01}\bar{M}) + \delta^2\rho_{11} + \mathcal{O}(\delta^3) \tag{B.14a}$$

$$\frac{1}{\kappa_b}\dot{\rho}_{01} = i\rho_{00}\bar{M}^\dagger - \frac{1}{2}\rho_{01} + \mathcal{O}(\delta) \tag{B.14b}$$

Let us examine the equations B.14a and B.14b. The first term in Eq. B.14b can be considered as an external drive, while the second one corresponds to damping. The dynamics of the drive is much slower, i.e., proportional to δ^2, compared to the damping terms. So the system converges to the right solution exponentially fast and then follows the driving term adiabatically, meaning that ρ_{01} and ρ_{01} are continuously in their steady state, giving us the solutions of equation B.13a and B.13b. Due to the same considerations, the assumption that ρ_{11} is in its steady state is justified, solving Eq. B.13c. The solutions for ρ_{01}, ρ_{01} and ρ_{11} are all a function of ρ_{00}; we can plug the expressions back into B.12, giving us

$$\dot{\rho}_{00} = \frac{4|g|^2}{\kappa_b}\left(M\rho_{00}M^\dagger - \frac{1}{2}\{M^\dagger M, \rho_{00}\}\right) = \mathcal{D}(\sqrt{\kappa_{\text{eff}}}M)\rho_{00} \tag{B.15}$$

The dynamics of the system are fully governed by a dissipation operator $L_{\text{eff}} = \sqrt{\kappa_{\text{eff}}}M$. Ideal dissipative cat qubit operation implies working in the limit $\kappa_{\text{eff}} \gg \kappa_1$. If we substitute $M = m^2 - \alpha^2$, κ_{eff} becomes the effective two-photon dissipation rate κ_2.

APPENDIX C

Bosonic state representations

Phase-space representations provide a powerful framework for visualizing and analyzing quantum states of bosonic systems. Instead of working directly in the Hilbert space, where wavefunctions or density matrices describe states, one can map the states onto quasi-probability distributions defined over the conjugate field quadratures $\hat{x}$ and $\hat{p}$

$$\hat{x} = \frac{1}{2}(\hat{a} + \hat{a}^\dagger), \quad \hat{p} = \frac{i}{2}(\hat{a}^\dagger - \hat{a}) \tag{C.1}$$

They act similarly to the position and momentum operators of a harmonic oscillator. These observables have continuous eigenspectra, i.e.

$$\hat{x} \left| x \right\rangle = x \left| x \right\rangle, \quad \hat{p} \left| p \right\rangle = p \left| p \right\rangle \tag{C.2}$$

with continuous eigenvalues $x \in \mathbb{R}$, $p \in \mathbb{R}$. The eigenstates define two bases connected by a Fourier transform $\left| p \right\rangle = \mathscr{F}[\left| x \right\rangle]$ and inverse Fourier transform $\left| x \right\rangle = \mathscr{F}^{-1}[\left| p \right\rangle]$.

Common representations include the Wigner function, the Husimi Q-function, and the Glauber-Sudarshan P-function, each offering a different perspective on the quantum state. The Wigner function, in particular, is widely used in circuit QED and quantum optics, as it captures both classical-like features and nonclassical signatures such as negativity. All those representations can be derived from the so-called characteristic functions. Together with the Wigner function, they are usually measured using an ancillary two-level system. The characteristic function has been gaining quite some interest in the field of bosonic quantum computation[128, 183], as its measurement is faster compared to the Wigner function, which is beneficial in the case of low dispersive shifts and coherence times of the ancilla.

C.1 Characteristic function

The characteristic function is the expectation value of a displacement $D(\xi)$

$$\mathcal{C}(\xi) = \mathrm{tr}(\rho D(\xi)) = \mathrm{tr}(\rho e^{\xi \hat{a}^\dagger - \xi^* \hat{a}}). \tag{C.3}$$

The characteristic function is a complex number with $\mathcal{C}(\xi)^* = \mathcal{C}(-\xi)$. Therefore, for a density matrix with symmetric wavefunctions in x and p, e.g. of a coherent state, the imaginary part of the characteristic function vanishes.

There are three operator orderings

- symmetric $\mathcal{C}_S(\xi) = \mathrm{tr}(\rho e^{\xi \hat{a}^\dagger - \xi^* \hat{a}})$
- normal $\mathcal{C}_N(\xi) = \mathrm{tr}(\rho e^{\xi \hat{a}^\dagger} e^{-\xi^* \hat{a}})$
- anti-normal $\mathcal{C}_A(\xi) = \mathrm{tr}(\rho e^{-\xi^* \hat{a}} e^{\xi \hat{a}^\dagger})$

giving the three density distributions by taking the inverse Fourier transform

- Wigner function $W(\alpha) = \mathscr{F}^{-1}[\mathcal{C}_S(\xi)] = \frac{1}{\pi^2} \int e^{\xi^* \alpha - \xi \alpha^*} \mathcal{C}_S(\xi) d^2\xi$
- P-function $P(\alpha) = \mathscr{F}^{-1}[\mathcal{C}_N(\xi)] = \frac{1}{\pi^2} \int e^{\xi^* \alpha - \xi \alpha^*} \mathcal{C}_N(\xi) d^2\xi$
- Husimi Q-function $Q(\alpha) = \mathscr{F}^{-1}[\mathcal{C}_A(\xi)] = \frac{1}{\pi^2} \int e^{\xi^* \alpha - \xi \alpha^*} \mathcal{C}_A(\xi) d^2\xi$

To measure a characteristic function, one uses a conditional displacement of the oscillator depending on the ancilla state. In case of a weak coupling, i.e., number-unresolved regime, the conditional displacements are achieved via the Echoed Conditional Displacement (ECD) protocol. Starting with the ancilla in a superposition state $|+\rangle$, applying two ECD gates with an opposite sign and a π-pulse in between, we can recover the characteristic function. Measuring the expectation value $\langle \sigma_x \rangle$ gives the real part $\Re(\mathcal{C})$, while measuring $\langle \sigma_y \rangle$ gives its imaginary part $\Im(\mathcal{C})$.

As the characteristic function shows only correlations between the bosonic state and its displaced copy, it is sometimes harder to interpret, compared to, e.g., the Wigner function, where the negativities denote a quantum state. The vacuum state overlaps only with small displaced states, symmetrically around the zero of phase space. Therefore, its characteristic function is

$$C_S^{|0\rangle}(\xi) = e^{-\frac{|\xi|^2}{2}}. \tag{C.4}$$

In phase space, this function represents a 2D Gaussian with a standard deviation $\sigma = 1$. The width of the vacuum state is sometimes used for calibrating the AWG amplitude to x and p scaling.

Another common example is the coherent state, which has a characteristic function

$$C_S^{|\beta\rangle}(\xi) = e^{-\frac{|\xi|^2}{2}} e^{2i\Im(\xi\beta^*)}. \tag{C.5}$$

In phase space, this state has the same Gaussian envelope as a vacuum state, but its real part is amplitude-modulated in the p axis with a cosine and its imaginary part–with a sine. The modulation frequency is $\omega = 2|\beta|$.

C.2 Characteristic function: considerations

For measuring the characteristic functions, tuning precise $\pi-$ and $\pi/2-$pulses is essential. Due to the AC Stark shift, the $\pi/2-$pulse cannot simply be considered to be half the

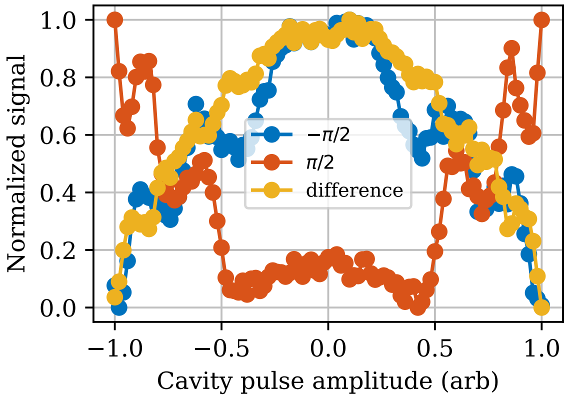

Figure C.1: Imaginary part of the characteristic function of a vacuum state with an imperfect $\pi/2-$pulse. A $\pi/2-$pulse with varying sign prepares a different superposition before the first ECD gate. An imperfect $\pi/2-$pulse leads to stray oscillations on top of the expected Gaussian for the vacuum state. If the measurements with initial $\pi/2-$ and $-\pi/2-$pulse on the qubit are combined, they partially compensate the residual oscillations.

amplitude of the $\pi-$pulse. Figure C.1 shows an example, where $\Im(\mathcal{C})$ of vacuum was measured with the $\pi/2-$pulse not properly tuned. The full Gaussian is not visible as the cavity pulse amplitude varies from -1 to $+1$, compared to -2 to $+2$ used in the main text. Preparing the qubit in a different superposition before the first ECD gate, i.e. $|\psi\rangle_i = (|0\rangle_q \pm |1\rangle_q)/\sqrt{2}$, could partially compensate for the pulse imperfection when the data is combined (see yellow data in Fig. C.1). However, this would double the measurement time. Instead, before running a long measurement, it is worth spending extra time and care on fine-tuning the pulses.

APPENDIX D

IQ and DSH mixing

High-fidelity control of superconducting quantum circuits requires precise generation of microwave signals with well-defined amplitude, phase, and frequency. In most experiments, this is typically achieved using In-phase and Quadrature (IQ) mixing, but alternative schemes such as Double-SuperHeterodyne (DSH) mixing become more and more popular. Both methods enable flexible synthesis of arbitrary microwave envelopes starting from digitally generated baseband signals.

D.1 IQ mixing

A general microwave signal at frequency ω_{LO} with a time-dependent complex envelope can be written as

$$V(t) = I(t)\cos(\omega_{\mathrm{LO}}t) + Q(t)\sin(\omega_{\mathrm{LO}}t), \tag{D.1}$$

where $I(t)$ is the in-phase component and $Q(t)$ is the quadrature component. By independently controlling $I(t)$ and $Q(t)$, one can generate microwave pulses with arbitrary amplitude and phase modulation.

Experimentally, the baseband signals, also known as intermediate-frequency (IF) signals, $I(t)$ and $Q(t)$ are produced by an AWG at low frequencies (typically a few hundred MHz). The mixing happens in an IQ mixer, seen in Fig. D.1, a device which consists of two single sideband mixers, 90°-hybrid splitter and a power combiner. The LO is applied at its respective port, where inside it is split with the hybrid splitter and applied to each of the two single sideband mixers, where it is mixed with either $I(t)$ or $Q(t)$. At the end, the resulting signals are combined. IQ mixing is used for both up-conversion before the experimental setup and down-conversion after the setup for the signal processing.

By choosing

$$I(t) = A(t)\cos(\omega_{\mathrm{IF}}t + \phi_I), \qquad Q(t) = A(t)\cos(\omega_{\mathrm{IF}}t + \phi_Q), \tag{D.2}$$

where $A(t)$ contains the envelope information, the resulting signal at the output of the IQ mixer would have components at $\omega_{\mathrm{LO}} \pm \omega_{\mathrm{IF}}$. By tuning the phases ϕ_I and ϕ_Q, one can suppress one of the two components and use this signal for quantum control. Note that the quadrature amplitudes can be picked to be different, i.e. A_I and A_Q.

In practice, imperfections in the mixer, such as amplitude imbalance between the I and Q paths, phase errors in the 90° hybrid, and DC offsets, lead to incomplete suppression

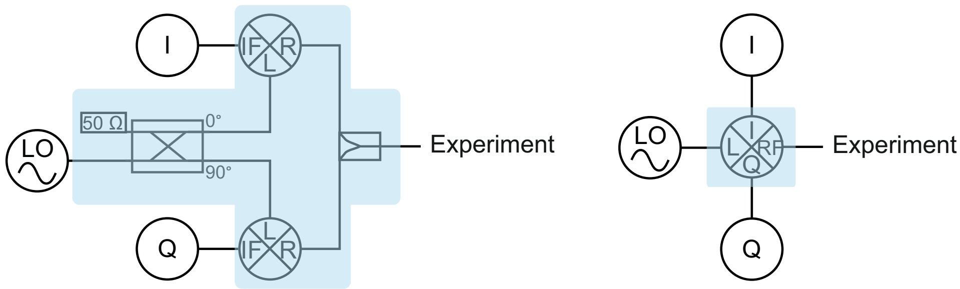

Figure D.1: Up-conversion with an IQ mixing setup: extended (left) and simplified (right). The components highlighted in blue show the IQ mixer. I and Q are the two AWG channels.

of the unwanted sideband and leakage of the LO tone. LO leakage arises from DC offsets in the I and Q channels and leads to an unwanted tone at ω_{LO}. Amplitude and phase mismatches result in imperfect sideband suppression. These errors are typically corrected by measuring the output spectrum and iteratively adjusting DC offsets and scaling factors in the AWG.

Additionally, nonlinearities in the mixer can also generate spurious tones, especially for large-amplitude pulses. This necessitates recalibration in case the input signal amplitude is changed. Moreover, recalibration is required if any of the frequencies are changed or in case the ambient temperature changes. Moreover, the higher-order mixing terms $\omega_{\mathrm{LO}} \pm 2\omega_{IF}$, $\omega_{\mathrm{LO}} \pm 3\omega_{IF}$, etc., cannot be calibrated out and can address an available transition in the system. Despite these limitations, IQ mixing remains the standard technique for microwave control in circuit QED.

D.2 Double Superheterodyne Mixing

Double superheterodyne mixing enables flexible frequency translation and improved spectral purity by performing frequency conversion in two successive stages. This approach is particularly advantageous when generating or detecting microwave signals at frequencies that are difficult to access directly, as is the case for fluxonium qubits. One of the main benefits of double superheterodyne mixing is that it does not require calibration.

The mixing scheme consists of two LOs, two single-sideband mixers and a narrow bandpass filter. Each mixing process generates both sum and difference frequency components, necessitating careful filtering after each stage to suppress image frequencies and unwanted spurious tones. In the first mixing stage, a single IF signal is mixed with the first local oscillator to produce microwave tones with frequencies $\omega_{\mathrm{LO},1}, \omega_{\mathrm{LO},1} \pm n\omega_{\mathrm{IF}}$, where n is the sideband order. A narrow band pass filter is placed after the first mixer to filter all but one tone, e.g. $\omega_{\mathrm{LO},1} + \omega_{\mathrm{IF}}$. In a second stage, this intermediate signal is mixed with another local oscillator to reach the final target frequency. The resulting tones are at frequencies $\omega_{\mathrm{LO},2}, \omega_{\mathrm{LO},2} \pm n(\omega_{\mathrm{LO},1} + \omega_{\mathrm{IF}})$. A low-pass filter removes all but the tone at frequency $\omega_{\mathrm{LO},2} \pm (\omega_{\mathrm{LO},1} + \omega_{\mathrm{IF}})$, which is used for the control. The same principle can be applied in reverse for down-conversion.

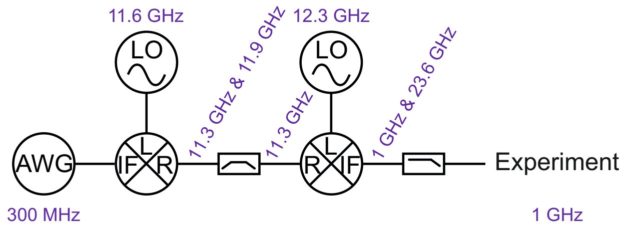

Figure D.2: Up-conversion with a DSH mixing setup with typical frequencies used in my setup. The LO frequencies are omitted in the intermediate sections.

In a typical operation, the first LO frequency is picked according to the central frequency of the filter and remains fixed, while the second LO frequency is calculated to achieve the desired frequency at the output. The IF needs to be larger than the filter bandwidth in order for the other mixing terms to be sufficiently suppressed. During large frequency sweeps, it is beneficial to first use the full bandwidth of the AWG before switching the frequency $\omega_{\mathrm{LO},2}$, as it is faster than changing the LO frequency.

The advantages of this method come at a price. The primary drawback of this approach is the increased system complexity. Frequency or phase drift in either local oscillator can directly impact the stability of the generated or detected signals. Moreover, the LOs must have high spectral purity as any higher harmonics or sub-harmonics would go into the mixing products and influence the signal at the output. Lastly, the real-time phase control is sacrificed as only one IF signal is mixed, unless the AWG can directly output a complex signal.

Due to the limitations of both mixing methods, the field has started its move towards Direct Digital Synthesis (DDS) for creating arbitrary waveforms from the reference clock.

APPENDIX E

Capacitively coupled LC oscillators

Two LC oscillators coupled via a shared capacitor C_s can be described by the following Lagrangian

$$\mathcal{L} = \frac{C_1}{2}\dot{\Phi}_1^2 - \frac{\Phi_1^2}{2L_1} + \frac{C_2}{2}\dot{\Phi}_2^2 - \frac{\Phi_2^2}{2L_2} + \frac{C_s}{2}(\dot{\Phi}_1 - \dot{\Phi}_2)^2. \tag{E.1}$$

The Hamiltonian of the system, instead, is

$$H = \frac{C_2 + C_s}{2C_\Sigma}Q_1^2 + \frac{\Phi_1^2}{2L_1} + \frac{C_1 + C_s}{2C_\Sigma}Q_2^2 + \frac{\Phi_2^2}{2L_2} + \frac{C_s}{C_1 C_2}Q_1 Q_2 \tag{E.2}$$

where here $C_\Sigma = C_1C_2 + C_1C_s + C_2C_s$.

In the second quantization language, the interaction Hamiltonian (last term of Eq. E.2) becomes

$$H_{\text{int}} = -\frac{\hbar}{2}\frac{C_s}{C_1C_2}\sqrt[4]{\frac{C_1C_2}{L_1L_2}}(a_1^\dagger - a_1)(a_2^\dagger - a_2) \tag{E.3}$$

where we have already replaced $Q_i = i\sqrt{\frac{\hbar}{2Z_i}}(a_i^\dagger - a_i)$. We can rewrite this equation in terms of the oscillators' frequencies $\omega_i = (L_iC_i)^{-1/2}$ to get

$$H_{\text{int}} = -\frac{\hbar}{2}C_s\sqrt{\frac{\omega_1\omega_2}{C_1C_2}}(a_1^\dagger - a_1)(a_2^\dagger - a_2) \equiv -\hbar g_{12}(a_1^\dagger - a_1)(a_2^\dagger - a_2) \tag{E.4}$$

from which the scaling $g_{ij} \sim \sqrt{\omega_1\omega_2}$ becomes evident.

APPENDIX F

Lessons learned

This list originated as part of the mixing chapter but was eventually promoted to a chapter of its own, as it may be the most useful part of this thesis. It is an incomplete list of things I did wrong, perhaps even multiple times, or I wish I knew at the beginning of my PhD.

- Most IQ mixing routines optimize only for the LO and the first-order products from the mixing. Higher-order mixing terms exist and will affect the spectral purity of the pulse unless filtered, i.e., with a YiG tunable filter.

- If you get a warning while calibrating in mixers in Octave, don't trust this calibration.

- If the frequency is swept, Octave + OPX looks for the calibration of the first applied frequency. If it cannot find a calibration matrix for this LO-IF frequency pair, it is unclear which, if any, calibration is applied.

- Be cautious when updating the firmware of an instrument. Known old bugs may be preferable to unknown new ones.

- The second harmonic of the LO in Octave is not filtered. This is a *well-known secret.*

- Use LOs with high spectral purity for the DSH mixing. While Valon's might be a tempting choice for IQ mixing, because they are cheap, they do have higher harmonics, as well as subharmonics, which could be detrimental for the qubit, especially if it is a fluxonium.

- Never trust your hardware[1]! (connected to the previous points) If something looks wrong, the chances are you are not discovering new physics.

- Avoid instruments powered via USB if you value your ground hygiene.

- Don't trust other people connecting your experiment. Two out of two times I did that, my wiring was messed up.

- Always note the wiring before cooling down or when changing it outside the cryostat. You never know when you might want to revert a change.

[1] Except for (maybe) the fast oscilloscope and the signal analyzer.

- The better the bookkeeping, the easier it is to find the data when needed.

- Try to enjoy the process. No matter how hard you try, there will always be hiccups on the way. Take your holidays, and don't take your laptop with you.

APPENDIX G

Active reset

The presence of a parametric amplifier allows us to do an active reset on the qubit by measuring its excited state population and playing a $\pi-$pulse depending on the result. For this purpose, we first measure the qubit via the readout resonator, multiple times without averaging, and build a 2D histogram in phase space, similar to the ones shown in Fig. G.1. We observe two coherent states corresponding to $|0\rangle_q$ and $|1\rangle_q$ due to the finite temperature, and thus the finite excited state population of the qubit.[1] We fit a sum of 2D Gaussian distributions to determine their crossing point and the threshold voltage, which is the feedback for the AWG. The 2D fits are convenient for determining whether there is a residual rotation in phase space that needs to be compensated for before setting the threshold. Note that there must be some delay before playing the conditional $\pi-$pulse, as there might still be photons in the readout resonator, which would influence the state preparation fidelity. This delay needs to be tuned for optimal operation so the qubit maintains the same state we have already measured.

From Fig. G.1, we extract that the excited state population in steady state is around 30.3%; after preparation in the ground state, it is reduced to 25.6%. The preparation of the qubit in an excited state with active reset gives us 79.5% excited state population, compared to 72.4% without feedback from the parametric amplifier. These numbers can be improved by repeating the reset sequence multiple times. The preparation fidelity is limited by the qubit lifetime and the external coupling rate of the readout resonator. To enable single-shot readout without compromising the qubit lifetime, a Purcell filter is necessary.

[1]If the qubit mode is cold, a $\pi/2-$pulse would be necessary prior to the resonator measurement.

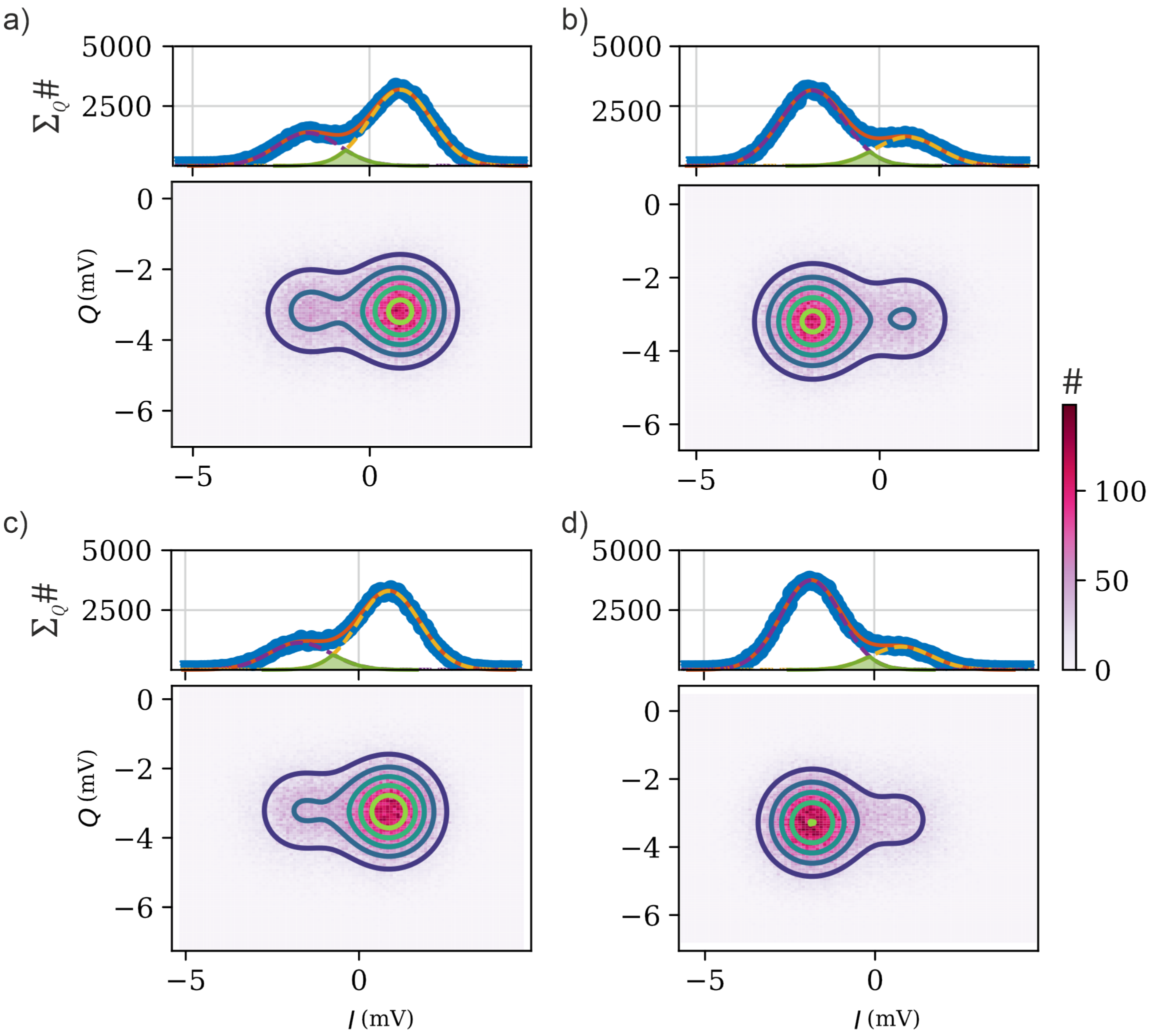

Figure G.1: Comparison of state preparation with (upper row) and without (lower row) active reset. The right column corresponds to the ground state, while the left column corresponds to the excited state preparation. In each subfigure, the top plot shows a sum over all bins in Q from the map under it. The sum is plotted in blue, with the two-Gaussian fit represented by a solid red line. Single Gaussian fits are shown in yellow and purple, and the overlap of the fits is indicated in green. The 2D two-Gaussian fit is plotted on top of the colormap with contours of changing color to signify the z-axis of the fit. # is the number of datapoints in each bin. The data was taken with the following parameters: square readout pulse of duration $2.4\,\mu$s, 100 ns delay, Gaussian qubit pulse of duration 20 ns, truncated at 6σ, threshold at -0.7415 mV.

APPENDIX H

SNAIL: Wave-mixing

This Appendix follows the derivations in Ref. [193]. The mixing capabilities of the SNAIL can be investigated by Taylor expanding the potential from Eq. 6.2 about one of its minima φ_{min}, resulting in the different order mixing coefficients

$$c_k = \frac{1}{E_J}\left(\frac{d^k U}{d\varphi_s^k}\right)_{\varphi_\alpha=\varphi_{\alpha,min}}. \tag{H.1}$$

We determine this minimum by enforcing the condition that no DC current flows across the SNAIL, but only circulating current within it in response to the applied flux. This is equivalent to numerically solving the following equation

$$c_1 = \alpha\sin(\varphi_{\alpha,min}) + \sin\left(\frac{\varphi_{\alpha,min}-\varphi_{\text{ext}}}{n}\right) = 0. \tag{H.2}$$

The effective potential as a function of the new phase variable $\tilde{\varphi}_\alpha = \varphi_\alpha - \varphi_{\alpha,min}$ is

$$U_{S,eff}(\tilde{\varphi}_s)/E_J = \frac{c_2}{2!}\tilde{\varphi}_\alpha^2 + \frac{c_3}{3!}\tilde{\varphi}_\alpha^3 + \frac{c_4}{4!}\tilde{\varphi}_\alpha^4 + ... \tag{H.3}$$

The coefficients c_k are flux-tunable, allowing us to tune the wave-mixing processes in situ.

The same procedure can be extended by including the harmonic contribution of the buffer mode (see Eq. 6.1). Instead of the Hamiltonian from Eq. 6.1, we work with the Lagrangian

$$\mathcal{L} = \frac{C_J\phi_0^2}{2}\dot{\varphi}^2 - \frac{E_L}{2}(\varphi-\varphi_\alpha)^2 + U_S(\varphi_\alpha) = \frac{C_J\phi_0^2}{2}\dot{\varphi}^2 - U_{\text{pot}}(\varphi,\varphi_\alpha) \tag{H.4}$$

where we combine the total potential energy of the system into the term U_{pot}. Finding the local minimum is equivalent to numerically solving the following equation

$$0 = \frac{\partial U_{\text{pot}}}{\partial\varphi_\alpha} = E_L(\varphi_\alpha-\varphi) + \frac{dU_S}{d\varphi_\alpha} = x_j(\varphi_\alpha-\varphi) + \alpha\sin(\varphi_\alpha) + \sin\left(\frac{\varphi_\alpha-\varphi_{\text{ext}}}{n}\right) \tag{H.5}$$

where we have introduced the junction-to-linear inductance ratio $x_j = L_J/L_b = E_{L,b}/E_J$.

In the regime of small phase fluctuations, we use the Taylor expansion of U_{pot} about its minimum $\tilde{\varphi}_{min}$ to reveal the mixing coefficients. The minimum is determined from the condition

$$\tilde{c}_1 = x_j(\tilde{\varphi}_{min} - \varphi_\alpha(\tilde{\varphi}_{min})) = 0. \tag{H.6}$$

The third and fourth order nonlinearity coefficients are

$$\tilde{c}_3 = -x_j \frac{d^2\varphi_\alpha}{d\varphi^2}(\varphi_{min}) \tag{H.7}$$

and

$$\tilde{c}_4 = -x_j \frac{d^3\varphi_\alpha}{d\varphi^3}(\varphi_{min}), \tag{H.8}$$

respectively. The explicit form of the first, second, and third derivatives are

$$\frac{d\varphi_\alpha}{d\varphi} = \left(1 + \frac{1}{E_L}\frac{d^2U_S}{d\varphi_\alpha^2}\right)^{-1} = \frac{x_j}{(x_j + \alpha\cos(\tilde{\varphi}_{min}) + \frac{1}{n}\cos(\tilde{\varphi}_{min}/n - \varphi_{\text{ext}}))}, \tag{H.9}$$

$$\frac{d^2\varphi_\alpha}{d\varphi^2} = \frac{-1}{E_L}\frac{d^3U_S}{d\varphi_\alpha^3}\left(\frac{d\varphi_\alpha}{d\varphi}\right)^3 = \frac{1}{x_j}\left(\frac{d\varphi_\alpha}{d\varphi}\right)^3\left(\alpha\sin(\tilde{\varphi}_{min}) + \frac{1}{n^2}\sin\left(\frac{\tilde{\varphi}_{min}}{n} - \varphi_{\text{ext}}\right)\right), \tag{H.10}$$

$$\begin{aligned}\frac{d^3\varphi_\alpha}{d\varphi^3} &= \frac{-1}{E_L}\left[\frac{d^4U_S}{d\varphi_\alpha^4}\left(\frac{d\varphi_\alpha}{d\varphi}\right)^4 - \frac{3}{E_L}\left(\frac{d^3U_S}{d\varphi_\alpha^3}\right)^2\left(\frac{d\varphi_\alpha}{d\varphi}\right)^5\right] \\ &= \frac{-1}{x_j}\Bigg[(-\alpha\cos(\tilde{\varphi}_{min}) - \frac{1}{n^3}\cos(\tilde{\varphi}_{min}/n - \varphi_{\text{ext}}))\left(\frac{d\varphi_\alpha}{d\varphi}\right)^4 - \\ &\qquad \frac{3}{x_j^2}(-\alpha\sin(\tilde{\varphi}_{min}) - \frac{1}{n^2}\sin(\tilde{\varphi}_{min}/n - \varphi_{\text{ext}}))^2\left(\frac{d\varphi_\alpha}{d\varphi}\right)^5\Bigg].\end{aligned} \tag{H.11}$$

The presence of the series linear inductor does not change the location of each SNAIL's potential minimum, i.e. $\tilde{\varphi}_{min} = \varphi_{s,min}$. These equations are used for the calculations presented in Fig. 6.2b.

Nanofabrication recipe

This appendix contains a detailed recipe that we developed, together with Dr. Teresa Hönigl-Decrinis, which I used for the fluxonium fabrication. The lithography was done in a Raith eLINE Plus 30kV, and the evaporation in a Plassys Bestek MEB 550S with the structures facing away from the laser and without a protective resist. After the fabrication, the wafer was diced in a Coherent 430 laser dicer. Some details on the dicing parameters can be found in Appendix D of Ref. [153].

Table I.1: Fluxonium fabrication recipe

<table>
<tr><th>Step</th><th colspan="2">Description</th></tr>
<tr><td>Piranha Cleaning</td><td colspan="2">Solution $H_2SO_4 : H_2O_2 = 3 : 1$
Clean for 5 min, rinse in DI water, blow-dry with N_2</td></tr>
<tr><td>Resist spinning</td><td colspan="2">MMA (8.5) EL13
Speed: 1500 rpm
Acceleration: 2000 rpm
Time: 100 s
Resist quantity: 1.05 ml
Bake: 200° C for 5 min
Thickness: 910 nm

950 PMMA A4
Speed: 2000 rpm
Acceleration: 2000 rpm
Time: 100 s
Resist quantity: 0.55 ml
Bake: 200° C for 5 min
Thickness: 250 nm</td></tr>
<tr><td rowspan="2">Discharge layer</td><td>Ti gettering</td><td>2 min at a rate 0.2 nm/s</td></tr>
<tr><td>Al evaporation</td><td>8 nm at a rate 1 nm/s, planar rotation at 5 rpm</td></tr>
<tr><td>Electron beam lithography</td><td colspan="2">Aperature size: 10 μm, Writefield size: 200 μm for the small structures
Aperature size: 120 μm, Writefield size: 1000 μm for the big structures
Zoom factor: 0.2%, Exposure time for a 2-inch wafer: about 3 hours</td></tr>
<tr><td>Etching</td><td colspan="2">TMAH-based developer for 1 min 30 s,
rinsed in DI water, blow-dried with N_2</td></tr>
<tr><td>Development</td><td colspan="2">Solution Isopropanol : $H_2O = 3 : 1$ at 6° C
Develop for 1 min 30 s, stop in DI water, blow-dry with N_2</td></tr>
</table>

Evaporation	**Loading and pumping overnight** Loadlock pressure: $9.9 \cdot 10^{-8}$ mbar Chamber pressure: $1.8 \cdot 10^{-8}$ mbar **Descum** Ar flow: Start at 10 sccm, ramped down to 5 sccm O_2 flow: 5 sccm Duration: 2 min V_{beam}: 200 V V_{acc}: 40 V I_{beam}: 5 mA Planetary rotation **Ti gettering** Duration: 2 min Rate: 2 nm/s **First Al layer** Angle: 28.5° Thickness: 20 nm Rate: 0.5 nm/s **Static oxidation** Pressure: 29.6 mbar Duration: 5 min Pressure ramp up: 1 min 26 s Pressure ramp down: 3 min 23 s **Second Al layer** Angle: -28.5° Thickness: 30 nm Rate: 0.5 nm/s **grAl layer** Angle: 0° Thickness: 20 nm Rate: 0.5 nm/s Flow: 2.2 sccm Chamber pressure: $3.8 \cdot 10^{-5}$ mbar Planetary rotation **O_2 capping** Pressure: 10 mbar Duration: 10 min Duration: approximately 1 hour 15 min
Lift-off	NMP at 80° C for 10+ min Clean in Isopropanol in an ultrasonic bath at 135 kHz, 35% power blow-dry with N_2

Additional notes on the fabrication:

- The thickness of each resist layer was measured after baking at three different points of the wafer with a Horiba SmartSE Liquid-crystal Spectroscopic Ellipsometer.

- A sputtered gold layer was used previously for the discharge layer, but its removal in a 5% Lugold solution left visible residue after removal.

- For the discharge layer, I did not note the pressures in the chamber. The recipe was started immediately after sufficient pressure for the machine operation was reached.

- TMAH does not attack the resist used; therefore, any TMAH-based developer available works for etching the Al discharge layer.

- The plasma in the descum process is ignited with the help of the argon. To ensure the process stays uninterrupted, we increase the argon flow at the beginning of the process and reduce it before opening the shutter to ensure the resist thickness does not change significantly.[1]

- For the qubit evaporation, the angles for the first two layers are calculated depending on the measured resist thickness.

[1] Argon milling is more aggressive and is used for etching metal by transfer of momentum.

APPENDIX J

Filters

J.1 Fast flux filters

This recipe for fabricating filters was shown to me by Dr. Romain Albert[1] and has been conceptualized and realized in Ref. [216]. This appendix simply details their fabrication.

Things you need:

- wax
- holders (2 per filter)
- Eccosorb CR-124
- SMA connector toolkit
- RG-402 semi-rigid copper coax cable
- 0.5 mm wire (best out of stainless steel covered in tin)
- SMA connectors from Radiall
  - R125.055.002W
  - R125.052.002W
  - R125.225.000W
  - R125.222.000W

Preparations:

The fabrication requires two holders (see Fig. J.1) per filter. If using the old holders, first they should be thoroughly cleaned with acetone and isopropanol. Ensure the inner

[1]Current address: Silent Waves, Grenoble, France.

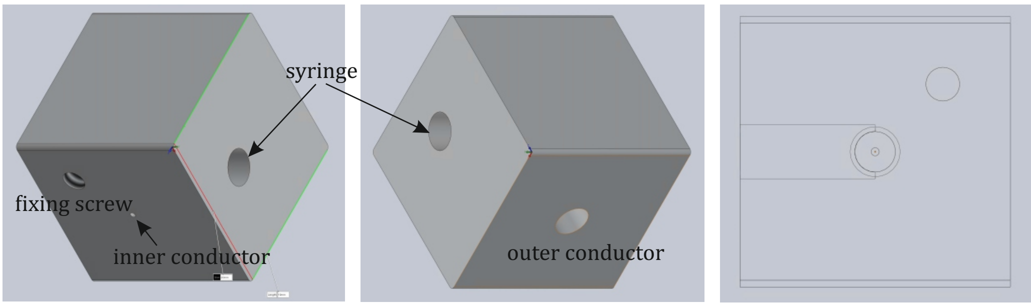

Figure J.1: Fast flux filter holders. The size of the cube (20 mm) sets the scale for all figures. The holders are usually machined out of aluminum alloy.

conductor hole is not clogged. Once clean, the holders have to be submerged in hot, melted wax; it creates a thin layer that prevents the eccosorb from sticking to the metal.

Cut a piece of RG-402 semi-rigid copper coax cable with a length corresponding to the desired cut-off frequency (10 cm $\leftrightarrow$ 300 MHz approximately). Create a clean cut with a saw, then use sandpaper to smooth the cut further. Use the cable stripping tool from the SMA toolbox to cut and strip the cable out of the dielectric and the outer conductor (see Fig. J.2a). Secure the coaxial cable with a wire and dip it into liquid nitrogen (LN2) for a few seconds. Pull the inner conductor, so you take it out together with the dielectric, leaving you with only the copper tube.

Clean the tube in deionized water and RBS or citric acid solution. Rinse thoroughly with deionized water and finish with isopropanol. Cut a piece of wire for the inner connector, longer than the tube, as you would have to fix it, and wipe it with isopropanol. Insert the wire into the tube, then into the holders together with the copper tube. Fix the wire with the screws on the holders. Do that for both sides, ensuring the wire is straight by looking into the holes for the syringe. The syringe holes should face up.

Heat up the eccosorb container (Part A) up to 60°C so it mixes easier[2]. Then mix with Part B[3] (check ratios on the datasheet) and put the mixture in a desiccator (not the same as the one for the samples) to remove any bubbles that formed during the mixing. There is a caveat: if you leave it in for too long, the solvent will evaporate; a few minutes should be enough. Once satisfied, pour it into a small syringe (10 ml or smaller so it is easier to push it in later)[4]. Put the end of the syringe into the hole on top of one of the holders and push the eccosorb inside. Be careful not to push the syringe too hard inside and break the inner connector. Stop when you see some eccosorb peeking out of the copper tube into the holder on the other side (Fig. J.2b). Cure the holders with the filter attached for 12 hours at 80°C.

Release the wire from the screws. Clean the holders from the eccosorb by drilling into the syringe holes (drill size 4 mm) until you first hear the inner conductor breaking. The metal should be visible at the bottom of the hole. To take out the broken wire, first push it a bit

[2] Pro tip: ask somebody who is regularly climbing to mix it.
[3] The ratios and the curing conditions can be found on the datasheet.
[4] Pro tip: pour it under such an angle that the outlet has all the air so you can easily push it out later without creating bubbles.

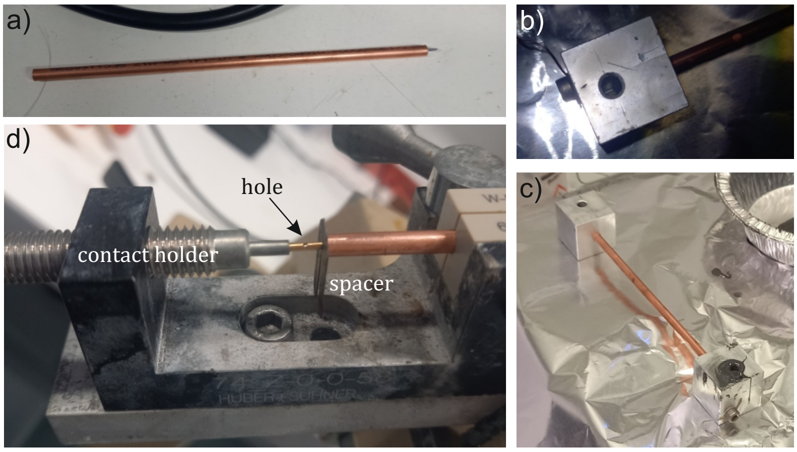

Figure J.2: Filter preparation. a) A piece of RG-402 semi-rigid copper coax cable, prepared for extracting the inner conductor with its dielectric. b) The holder that has not been used for the filling, the eccosorb poking in is visible. c) Filter after the curing. d) Fixing the filter to solder the connectors.

in and then out to prevent it from breaking inside the holder. Remove the holders from the filter; heating them up to melt the wax should help. Use the cable stripping tool from the SMA toolbox again to cut only the copper, not the eccosorb. Create a groove in the eccosorb with a scalpel; you shouldn't need too much pressure. Break the eccosorb and take it out while pulling straight to prevent the inner conductor from bending or breaking. Check that you have only eccosorb and no wax left inside the filter.

Check the length of the inner conductor needed for soldering (it should be on the datasheet of the connectors) and cut it accordingly. For each side, you would need a set of two connectors as they fit either RG-405 or RG-402 coaxial cables, and we use a combination; for example, use the inner conductor piece and dielectric from R125.052.002W and the outer conductor piece from R125.055.002W. Before you continue, check that the socket fits the outer conductor of the filter; otherwise, file the copper until it does.

Screw in the contact holder to the soldering fixture from the SMA toolbox as shown in Fig. J.2d. Fix the filter on the fixture. Use the thinnest spacer/soldering gauge between the pin and the end of the filter. Apply flux on the conductor, put solder on the iron beforehand for good thermal contact, and fill the hole on the pin with solder. The recommended temperature is about 260 degrees. If you use too much, remove the excess with the copper wick.

Remove the contact holder and screw in the corresponding locator tool, which should fix the position of the socket with respect to the pin. Fit the tool to the soldered pin and fix everything. Put solder on the copper. Slide in the socket. Use the resistance soldering tweezers for this part. If the tips are too oxidized, file them before usage. Prepare the amount of solder you think you will use in one hand, and with the other, grab the bottom of the socket, and use the pedal to heat it up while adding the solder between the copper and the socket. Cool the connector with isopropanol immediately. Rinse off the flux with water and isopropanol afterwards.

Use the corresponding dielectric insertion tool to put the Teflon piece into the connector. The dielectric should be either flush with the pin or with the step of the pin, depending on the connector type. If that's not the case, use the locator tool to push it further. Then your filter is ready to be tested.

J.2 Microwave filters

The filter was conceptualized by Dr. Romain Albert; my contribution is in the time-domain reflectometry simulations used for the impedance matching of the filter, as well as finalizing the design.

Time-domain reflectometry (TDR) is a powerful technique for identifying impedance discontinuities along an electrical path by analyzing the time-resolved response to a voltage pulse. The time at which a reflected signal appears indicates the location of a discontinuity, while the detailed shape of the reflection provides insight into its physical origin. Here, TDR is used to impedance-match the filter at both the SMP and SMA interfaces.

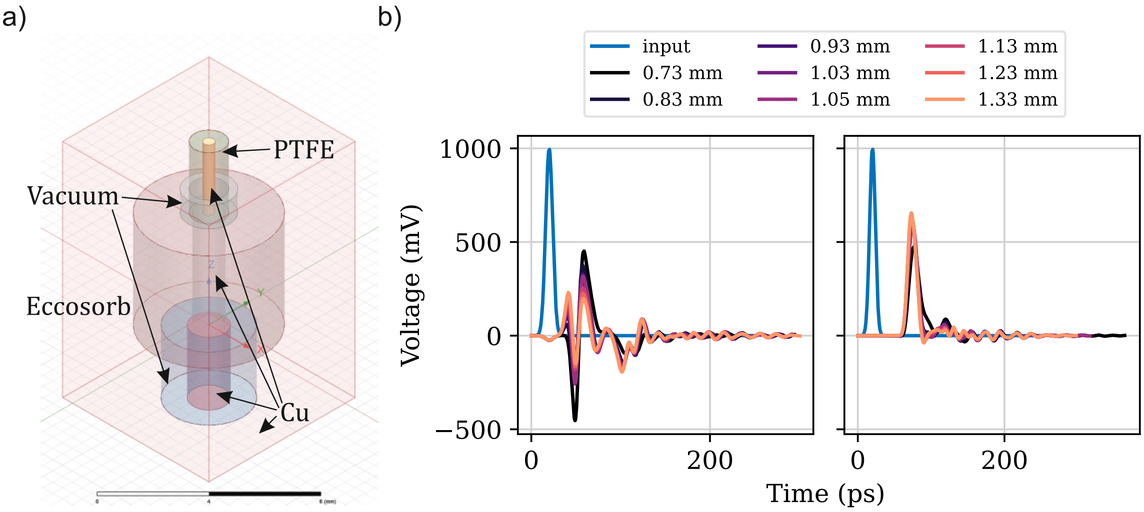

Figure J.3: Time-domain reflectometry simulations. a) Model in *Ansys HFSS*. The solution type used for the simulations in *Transient Network*. b) Reflection at the SMP side (left) and transmission from the SMP to the SMA (right)for different radii of the vacuum section. The input pulse is always plotted in blue, while the remaining curves show the output signal.

The filter is modeled as shown in Fig. J.3a, with all dimensions taken directly from the connectors used. The filter body and all inner conductors are treated as copper. The SMP connector (top) contains a PTFE dielectric between its inner and outer conductors, whereas the SMA connector (bottom) features a vacuum dielectric. The Eccosorb[5] is placed between the two connectors.

[5] A custom material was defined in *Ansys HFSS* to reproduce the microwave properties of Eccosorb CR-112 accurately.

Optimal impedance matching is achieved by simultaneously minimizing the reflection amplitude and maximizing the transmission. To this end, the radii of both the inner conductor in the middle section $r_{\text{pin}} = 0.58\,\text{mm}$ and the surrounding Eccosorb region $r_{\text{ecco}} = 2.7\,\text{mm}$ are selected based on a parameter sweep in the TDR simulations. These prior simulations identified the SMP–Eccosorb interface as the dominant source of impedance mismatch. To compensate for this, a short vacuum section is introduced at the boundary between the Eccosorb and the SMP connector, and its radius is varied in the simulation to achieve the best possible matching. The resulting reflection at the SMP port and the transmission from the SMP to the SMA side are shown in Fig. J.3b. In the present design, this condition is satisfied for radii in the range $r \in [1.23, 1.33]\,\text{mm}$.

Things you need:

- filter body out of copper
- 23_PC35-50-0-51/199_U SMA connector
- R222.561.301B SMP connector
- Eccosorb CR-112
- 4 customized screws
- low-temperature (SMD) solder
- inner conductor

Preparations:

Initially, the two sides of the copper part are not connected (see Fig. J.4a left), so the eccosorb can be poured. Working with Eccosorb CR-112 is substantially easier as the viscosity is lower and therefore does not require heating of Part A in order to be mixed. Similarly, the mixed eccosorb is put in a desiccator before pouring and before curing to avoid the formation of bubbles. Once the eccosorb is cured, the top part is evened out by milling so the eccosorb is flush with the body, the screw holes for the SMA connector are drilled, the whole eccosorb is drilled axially through to fit the inner conductor, and the SMP part together with the inner conductor is connected to the eccosorb (see Fig. J.4a right).

The inner conductor is made out of a beryllium copper tube, which has a radius that would match the eccosorb section to $50\,\Omega$, i.e. $r_{\text{pin}} = 0.58\,\text{mm}$ as determined by the TDR simulations. At the SMA side, it has a reduced radius in order to fit the connector. The inner conductor is soldered to the SMP connector by filling the interior of the tube with some low-temperature solder. The solder should not come out of the tube once pressed to the SMP connector. The two pieces are soldered together by heating them from the top with a heat gun. Care must be taken that the inner conductor is not bent with respect to the connector during the soldering.

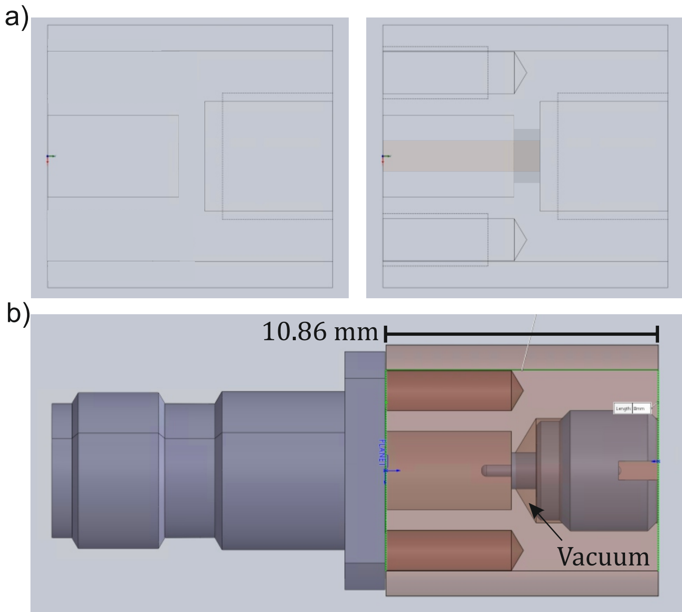

Figure J.4: Microwave eccosorb filters. a) Filter body before (left) and after (right) the eccosorb pouring and curing. After curing, the eccosorb and the copper are drilled through to accommodate the inner conductor (orange). Additionally, the vacuum section is created (darker gray). b) Full assembly. The section between the SMP filter and the eccosorb has a cylindrical shape, instead of the depicted conical one. The length scale of 10.86 mm for all subfigures is set by the length of the filter body.

As previously mentioned, time-domain reflectometry simulations were used to determine the approximate radius of the section between the SMP connector and the eccosorb. The correct volume ratio of the vacuum part and the dielectric part of the SMP connector results in impedance matching. Several different diameters were tested, each time incrementally increasing the diameter of the vacuum section by 0.1 mm. For the final design, a diameter of 2.6 mm was used, as a deterioration in the reflection was observed for the next size. This is consistent with the expectations from the TDR simulations. The scattering matrix elements are plotted in Fig. J.5. The filter is suitable for the typical working range $[0.5, 17]$ GHz, with its bandwidth limited by the SMP side.

To assemble the filter, first, the SMP connector with the inner conductor soldered to it is screwed into the body. Then the SMA connector is pushed into the top of the inner conductor. Finally, we screw the SMA connector with the four customized[6] screws.

[6]The screw heads were too big to fit the connector, so we milled them down.

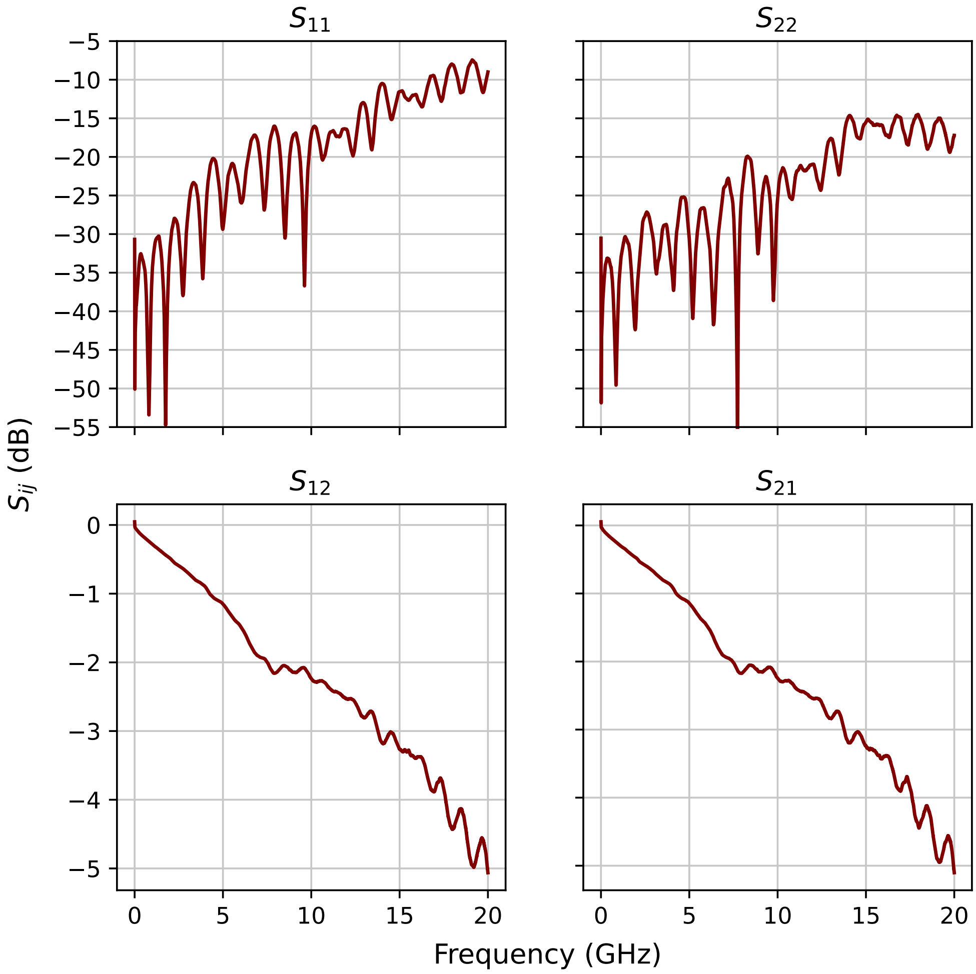

Figure J.5: Scattering matrix elements of the final successful iteration of the microwave Eccosorb filter. The SMP connector is on side 1, while the SMA connector is on side 2.

Acknowledgements

These will be my first acknowledgements, so brace yourself, they might be long.

First and foremost, I would like to thank Prof. Gerhard Kirchmair for believing in me. I would have never imagined myself having a PhD in physics if it wasn't for you, and frankly, I probably wouldn't have done a PhD with any other supervisor. I will forever cherish the conversations with you, whether scientific or not-so-scientific, the great barbecues, which I hope I still can receive invitations for, and the freedom you gave me to establish my own scientific approach. I still have a lot to learn from you, and hopefully I will have the opportunity to do so in the future. In the meantime, I hope you find yourself more often in the lab and less buried in bureaucracy, because otherwise another student might never find out how to properly use a fast oscilloscope.

I express my deepest gratitude to Prof. Ioan Pop. I appreciate all our productive discussions, whether it is about the fluxonium, its Purcell limits, or the Balkan ways of doing things. Thank you also for the week-long research stay in Karlsruhe and for inviting me to your HiQCuing group retreat, it is one of the greatest weeks of my whole PhD. I would also like to thank all the group members, as well as non-group members, for making this HiQCuing experience so memorable and for sharing their knowledge with me during my stay in Karlsruhe. And while I'm on the people from KIT, I want to thank Dr. Daria Gusenkova for writing out the lumped element diagonalization of my system, as well as the code for it, which I later used as a reference for how to diagonalize the SNAIL coupled to the cavity. I thank Dr. Patrick Winkel, who, in the end, had to spend two days with me in the cleanroom, for me to graduate with my master's degree.

I would like to thank Prof. Zaki Leghtas for allowing me to do the research stay at LPNS. I appreciated all the discussions with you and your team during which I learned everything I know about two-photon dissipation, especially in 3D cavities. I greatly enjoyed the INRIA day and having the opportunity to meet so many amazing physicists and get an overview of the state-of-the-art dissipative cat experiments and theory. I especially want to thank Dr. Ulysse Reglade and Adrien Bocquet not only for sharing their scientific knowledge, but also for helping me get into the Alice & Bob parties at the March meeting. I have to mention also Dr. Alvise Borgognoni, with whom I shared an office in Innsbruck, where we cried about our master's projects together. Thank you for keeping me caffeinated during my stay, and thank you for your continuing friendship.

I want to express my gratitude to Dr. Ian Yang for introducing me to the first cavity measurements and always sharing his secret stash of microwave components with me. I admire your dedication to science and your perseverance. I see a lot of success in front of you! I have told you before, but I am so grateful for introducing the team spirit in this lab, one of your greatest contributions besides the scientific ones. Also, thank you for the great food!

I'm grateful that Dr. Teresa Hönigl-Decrinis helped me immensely during the first fluxonium samples with both their fabrication and measurements. It made the process a bit less painful. Moreover, thank you for accompanying me to Manchester and allowing me to spin fire, which I will always remember. I know your future group will be in good hands.

I thank Philip Straub for all the work he put into the flux hose. I hope we successfully finish this project (now from three different countries) with a nice paper. I will be keeping a close eye on your future career, which I'm sure will be bright.

I would like to thank Dr. Romain Albert, who was the mastermind of a big part of my side project. You came into a group and immediately saved me by helping with the amplifier measurements, when nobody else had any idea about them. I learned so much from you regarding wet chemistry, impedance matching, good food, and good wine, for which I am grateful. I am sure that everyone at Silent Waves is more than happy to have you there.

I thank Christian Schlager for all the help with the tin measurements, as well as for the mindless chats we had in the office when I didn't feel like working that much. I believe efforts will pay off, and the lab will soon have its very own TWPA.

I thank Stefan Oleschko for teaching me how to craft a flux hose. It is a skill I hope I will never have to use again. I'm looking forward to your defense!

I extend my thanks to Vasilisa Usova, who wrote most of the code for the QM sequences. I know you will not agree, but I believe you are a great physicist! I am certain your stubbornness and knowledge would allow you to achieve whatever you set your mind to. I know the PhD was painful for both of us, but your friendship is one of the most valuable things I would take out of this PhD.

I need to mention Dr. Christian Schneider (the OG Christian in the group). When I started my PhD, I said to Gerhard, "I don't know, I can't imagine myself becoming as good as Christian". You helped me so much during my master's project, when I didn't know anything about the data module, and I'm so grateful for it. Don't forget to take care of yourself; everything else you can already do.

I want to thank the last high-Q team member, Lucien Québaud, who was the first tester of the superconducting flux hose. Although our efforts towards a modular Purcell filter have not yet been successful, it was a pleasure to share this side project. I know I am leaving the can project in good hands. I am happy your PhD is on the right track, and I'm looking forward to your future success.

I want to thank Lukas Deeg for the time we shared crafting input lines... three times, and for keeping my plants alive while I was away. I will cherish all the summer schools together and the complaining sessions we had about the PhD. I have a great trust in your scientific capabilities, and I am certain you can achieve wherever else you set your mind to.

I thank Dr. Bhargava Thyagarajan for proofreading part of this thesis. I thank Dr. Aleksei Sharafiev for showing me the lab for the first time. I want to thank Prof. Mathieu L. Juan for teaching me perseverance while trying to wind my DJJAA coil a week before moving from Innsbruck to start his own group in Canada. I thank Dr. Markus Weiss for the introduction to the cleanroom in Innsbruck.

I want to thank Julian Daser for taking care of the things nobody wants to (computers and networks), for all the liquid nitrogen ice cream, and for all the drinks and parties we shared on two different continents. I thank Dr. David Zöpfl for organizing the sunset tea. Thanks to Raamamurthy Sathyanarayanan, who successfully ran the first D&D campaigns I participated in, I was introduced to a whole new world of entertainment. I thank Dr. Oscar Gargiulo for teaching me the art of making coffee with the Vesuvio analogy. I thank Dr. Phani R. Muppalla for giving me one of the best curry powders. I thank specifically the pasta crew (Teresa, David, Lukas, and Julian) for always keeping me satiated during Wednesday lunch. And a huge thanks to the rest of the group members: Dr. Maximillian Zanner, Filip Hudek, Dr. Tanmoy Bera, Christian Dejaco, Niklas Nocker, Deniz E. Molavali, Emily Kunz, Isabel Pietka, Markus Sifft, Stephanie Wögerer, Paul Heidler, and all our summer/winter/spring/fall students for all the great discussions, all the group activities, and all the coffee breaks. You all made this whole experience much nicer in those couple of years.

I want to thank the mechanical and electrical workshop of IQOQI: Andreas Strasser, Bernhard Öttl, Tanja Waidacher, Gerhard Hendl, and David Arnold, who always had to do something for me, and the deadline for it was yesterday. I owe you partially my thesis and a huge amount of beer and cookies.

I would like to thank the IT department of IQOQI: Valentin Staubmann and, especially, David Jordan, for always fixing my laptop issues and keeping the systems running. I want to thank the whole administration staff: Elke Wölflmaier, Elisabeth Huck, Nikolaus Falschlunger, Claudia Mavenkamp, Katrin Nagele, Birgit Weihs-Dopfer, and Markus Knabl for making my life so much easier by dealing with all the documentation for traveling and ordering. Your work is much appreciated! I want to thank Dr. Mustapha Chouiki for keeping the cleanroom running. And I would like to thank our gender officers, Dr. Karoline Irschara and Angélique Sanchez, for always organizing great events and for advocating for inclusivity in the STEM field.

I want to thank the whole community of the DK-ALM, especially the administrative staff: Dr. Heide Streicher and Dr. Eric Endres, who organized endless events, summer schools, etc., and whose doors were always open for us, the students. It was a pleasure to be a student representative of this program for a couple of years. I express my gratitude towards the whole Quantum Machines team, who did their best to solve any issues that occurred as soon as possible. I extend my thanks to Ulrich Bauer for doing a splendid job of electroplating all the parts that I provided so far.

A big thanks to all the friends I made in Innsbruck, who are too many to list here, but I appreciate you all. I express my gratitude to Bastiaan, who was always there for me. I wouldn't be where I am if it weren't for you.

To Luigi, whom I sometimes admire for putting up with my behavior, I am grateful to have met you during my PhD. I know my coding skills can be frustrating to you sometimes, but if it works, it works. I am thankful that you helped me with the Linux machine crashing and other coding adventures, and thank you for always listening to my complaints (and I am a bit less thankful for the unwanted advice you gave me sometimes). And while I can imagine my life without you, I simply don't want to. No matter how overwhelmed or stressed I am about moving to another country, I'm ecstatic to do this with you. Ti amo.

И, естествено, кратко обращение на български, в случай че някой отвори дипломната ми. Огромно благодаря на семейството ми, най-вече родителите ми, Виолета и Ганчо, и брат ми ,Галин, които въпреки че не винаги разбираха какво правя или защо го правя, винаги ме подкрепяха в начинанието. Обичам ви безкрайно. Искам да благодаря и на приятелите ми в България, които след седем години скитане в Австрия още не са ме изхвърлили от животите си и са винаги насреща. Благодаря на дипломния ми ръководител през бакалавъра, доц. д-р Петър Александров Иванов, работата с теб беше винаги удоволствие. И последно благодаря на учителките ми по физика от гимназията, Антонина Димчева Христова и Милена Боева, които възпламениха тази страст към физиката.