DOCTORAL THESIS OF SORBONNE UNIVERSITÉ

Specialty: Physics | Doctoral School No. 564: Physics in Île-de-France
Carried out at the Kastler–Brossel Laboratory

presented by
Kyrylo GERASHCHENKO

Thesis title:

Probing the quantum motion of a macroscopic mechanical oscillator with a radio-frequency superconducting qubit

Defended on 24/09/2025
Before the jury composed of:

Albert	SCHLIESSER	Reviewer (Rapporteur)
Max	HOFHEINZ	Reviewer (Rapporteur)
Denis	VION	Examiner (Examinateur)
Philippe	CAMPAGNE-IBARCQ	Examiner (Examinateur)
Agnes	MAÎTRE	Examiner, Jury President (Examinatrice, Président du jury)
Antoine	HEIDMANN	Thesis Supervisor (Directeur de thèse)
Thibaut	JACQMIN	Thesis Co-supervisor (Co-directeur de thèse)
Samuel	DELÉGLISE	Thesis Co-supervisor (Co-directeur de thèse)

Last corrected: May 18, 2026

"I think," said Dirk, "you will be impressed. Consider this. An intractable problem. In trying to find the solution to it I was going round and round in little circles in my mind, over and over the same maddening things. Clearly I wasn't going to be able to think of anything else until I had the answer, but equally clearly I would have to think of something else if I was ever going to get the answer. How to break this circle? Ask me how."

"How?" said Miss Pearce obediently, but without enthusiasm.

"By writing down what the answer is!" exclaimed Dirk. "And here it is!" He slapped the piece of paper triumphantly and sat back with a satisfied smile.

Miss Pearce looked at it dumbly.

"With the result," continued Dirk, "that I am now able to turn my mind to fresh and intriguing problems, like, for instance . . . "

He took the piece of paper covered with its aimless squiggles and doodlings, and held it up to her.

"What language," he said in a low, dark voice, "is this written in?"

— Douglas Adams, *Dirk Gently's Holistic Detective Agency*

Acknowledgments

This thesis exists thanks to many people who supported me and kept me moving forward.

It all started when, for bureaucratic reasons, my initially planned M2 internship was cancelled, and I found out that Thibaut Jacqmin (my M2 teacher at the time) was looking for a PhD student. At that time, I was not sure which direction in quantum computing I wanted to pursue, but in retrospect I am very lucky that things happened this way, since it led me to discover the field of superconducting qubits.

I am grateful to my supervisors. Samuel Deléglise, your energy throughout this work was one of the main reasons behind the success of this project. It is a pleasure to work with you, and I am grateful for all your support, even though I know how hard it was for you to keep things positive. I learned a lot from you, and my work with you definitely changed my personality. Thibaut Jacqmin, despite your heavy teaching load, you always found time to keep the project moving. I also appreciate your calm and relaxed feedback. Special thanks for your help with the clean room and membrane simulations at the beginning of my PhD. Antoine Heidmann, I am grateful for your help with the bureaucracy and for being lab director all these years.

I am also grateful to our collaborators. Emmanuel Flurin, I appreciate all your phone calls to Samuel and the motivation you added to this project. Zaki Leghtas and Philippe Campagne-Ibarcq, thank you for the discussions and advice that helped this project to be possible. Antoine Tilloy, thank you for your help with the quantum gravity calculations. I am also grateful for the discussions and comments I received during different conferences and schools, which helped refine the project and my understanding of the field. And finally, I am grateful to the members of my thesis jury, who took the time to read my thesis and made my defense more engaging. Thank you for all your questions during the defense—especially those related to quantum gravity—and for your feedback on the thesis afterward.

This project would not have been possible without our teamwork. Léo Balembois, I appreciate your help with fabrication and your positive mindset, even when fabrication did not go as expected. Your organizational skills always inspire me. Rémi Rousseau, I am grateful that you brought your experience from A&B to our lab, and for your participation in design, theory, and experiments. You made me understand how easy it can be to simulate circuits. Himanshu Patange, I appreciate your work developing and testing the flip-chip assembly; it made this project much easier. B.-L. Najera-Santos, thank you for your initial work establishing the heavy fluxonium fabrication process in our group. I also want to acknowledge the new generation of PhD students. Paul Manset, your positivity changed the atmosphere in the group and motivated me to be more positive in my feedback. I also appreciate your scientific contributions and the way you push your projects forward. Alice Desbois, even though I did not have the chance to work with you, I am glad you joined our group to continue with the interferometer and to fabricate new types of membranes that had never been done in our group.

I am grateful to the other PhD students of my generation who brought a friendly atmosphere to the lab. Pierre-Edouard Jacquet, your interest in new physics led to many interesting discussions, even though we did not work on the same topic. Emile Ferreux, I appreciate your positivity and our sports activities. I am also grateful to other members of our team who, even outside my project, tried to help. Tristan Briant, I appreciate your engineering skills and your offer of teaching hours. Pierre-François Cohadon, thank you for your recommendations.

I am grateful to all LKB administrative staff for your help with contracts, missions, and other bureaucratic matters. I also thank the electrical workshop for letting me use your equipment and for your advice, and the mechanical workshop for doing your best and helping us within reasonable timeframes. I am grateful to the INSP and ENS clean room staff—especially José Palomo, Aurélie Pierret, and Michael Rosticher—for your help with our fabrication recipes.

In this rare opportunity to express my gratitude, I would like to acknowledge my high school teachers, who gave me the freedom to explore my own interests. I am also grateful to my alma mater, Karazin Kharkiv National University, and specifically the School of Physics and Technology. It taught me how to learn and how to understand Landau's books, which I never thought would be possible. I am very grateful to all my physics professors at X for your patience and explanations.

I am very grateful to my friends outside the lab. Oryna, it is always a pleasure to see you. Andrii and Anton, even though we cannot see each other as often, I am very happy that I met you during my bachelor days. It's always a pleasure to see you. Darina and Vsevolod, I appreciate your optimism and joy for life. Yegor, your openness always delighted me, and I am always surprised by how you combine sport with your food habits. Oleksii, thanks for motivating me to start kayaking again. Anastasia, thank you for your New Year's cards. And I am grateful to many more who supported me during this journey.

Of course, I cannot forget that during the second year of my PhD, the full-scale invasion of Ukraine by russia began. It was a tough time. I am grateful to my lab for understanding and for letting me spend more time volunteering. I am also grateful to all the volunteers I worked with in different centers; it is a pity that we did not have the chance to get to know each other better, but your strength to keep things going was inspiring. And most importantly, I am grateful to the Ukrainian army. I would not be able to discover the microworld without you defending the macroworld.

I finish by thanking my family. I am grateful to my parents and grandparents for your unconditional support and love. My sister Jane, I appreciate that you are always ready to help me. My partner Daria, I am grateful for your emotional support, your help, and simply for being there with me.

I am grateful to all of you.

Contents

List of Figures

1 Introduction

1.1 Context of this work

Probing quantum gravity with electromechanical systems

Quantum Mechanics (QM) and General Relativity (GR) have each achieved remarkable experimental success, yet reconciling them remains an open challenge in modern physics. A central tension arises when considering relativistic time dilation in quantum superpositions: GR predicts that proper time flows differently at distinct locations within a massive object, whereas QM assumes a single, universal time parameter for coherent evolution [OZ16]. The physical consequences of an appreciable differential proper-time phase shift remain debated. Diosi and Penrose [Dió89; Pen96] gravity-induced decoherence models attribute the quantum-to-classical transition solely to gravitational effects.

[OZ16]: Oosterkamp et al. (2016), *A clock containing a massive object in a superposition of states; what makes Penrosian wavefunction collapse tick?*, arXiv

[Dió89]: Diósi (1989), 'Models for universal reduction of macroscopic quantum fluctuations', Physical Review A

[Pen96]: Penrose (1996), 'On Gravity's role in Quantum State Reduction', General Relativity and Gravitation

In the Diósi–Penrose (DP) framework, a spatial superposition of a massive object acquires a finite coherence time due to the difference in its gravitational self-energy. For a resonator of mass m prepared in the "Schrödinger-cat" superposition

$$|\Psi\rangle = \frac{|\alpha\rangle + |-\alpha\rangle}{\sqrt{2}}, \tag{1.1}$$

one defines the gravitational self-energy difference [Kle+08; Che13]

$$\begin{aligned} \Delta E &= 2E_{12} - E_1 - E_2, \\ E_{ij} &= -G \iint \frac{\rho_i(\mathbf{r})\rho_j(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} d^3r\, d^3r', \end{aligned} \tag{1.2}$$

[Kle+08]: Kleckner et al. (2008), 'Creating and verifying a quantum superposition in a micro-optomechanical system', New Journal of Physics

[Che13]: Chen (2013), 'Macroscopic Quantum Mechanics', Journal of Physics B: Atomic, Molecular and Optical Physics

where ρ_i and ρ_j are the mass-density distributions of the two components of the superposition and G the gravitational constant. The collapse time is then

$$\tau_G = \frac{\hbar}{\Delta E}. \tag{1.3}$$

Observing such a mass-dependent quantum-to-classical crossover requires two conditions to be satisfied:

1. The superposition's spatial separation ΔX exceeds the typical width of the mass distribution;
2. The environmental decoherence time τ_{dec} is much longer than τ_G.

Both conditions depend crucially on the object's mass distribution, and there is considerable debate about its appropriate form for a macroscopic body [Dió07; Kle+08]. However, a plausible assumption is that the mass is arranged in a regular array of spheres—each of mass m_i and radius r_{n}—representing the nuclei of the material. Under this assumption, the gravitational self-energy of each nucleus is $\Delta E \approx 48\pi G\, m_i^2/5\, r_{\mathrm{n}}$, leading

to a total decoherence time

$$\boxed{\tau_G \approx \frac{5\hbar\, r_\mathrm{n}}{48\pi\, G\, m\, m_i} = \frac{1}{m} \times 3 \times 10^{-15}\,\mathrm{kg}\cdot\mathrm{s},} \tag{1.4}$$

where m is the total mass, G the gravitational constant, and for silicon (used in this work) $m_i = 4.6 \times 10^{-26}$ kg and $r_\mathrm{n} = 2.7$ fm.

One can envision an experiment in which a vibrational mode of a nanofabricated mechanical oscillator is prepared in a Schrödinger-cat state, i.e. a superposition of two coherent states of complex amplitude α with opposite phases, as illustrated in Fig. 1.1. The spatial separation between the two components of the superposition is

$$\Delta X = 2|\alpha|\, X_\mathrm{zpf},$$

where the zero-point fluctuation—i.e., the characteristic size of the ground state gaussian wavefunction—is $X_\mathrm{zpf} = \sqrt{\hbar/(2m\omega_m)}$. For a mechanical mode with energy damping rate Γ_m, thermal occupation n_th and prepared in the Fock state $|n\rangle$, the decoherence time is [HR06; GS21]

$$\boxed{\tau_\mathrm{dec} \approx \frac{1}{\Gamma_m}\frac{1}{n_\mathrm{th}}} \tag{1.5}$$

In practice, increasing the resonator mass m shortens the collapse time τ_G (see Eq. 1.4), but at fixed cat amplitude $|\alpha|$ the spatial separation $\Delta X \propto |\alpha|(m\omega_m)^{-1/2}$ decreases, forcing a larger phonon number $n = |\alpha|^2$ to satisfy $\Delta X \geq 2r_\mathrm{n}$. Increasing n then reduces the environmental coherence time τ_dec (see Eq. 1.5), since for $n, n_\mathrm{th} \gg 1$ one finds $\tau_\mathrm{dec} \propto 1/n$. Consequently, optimizing a platform demands a careful trade-off among m, ω_m, and $|\alpha|^2$ to ensure both $\Delta X \geq 2r_\mathrm{n}$ and $\tau_\mathrm{dec} \gg \tau_G$.

Several platforms, interfacing diverse physical domains have been proposed for preparing and detecting mechanical states. Fig. 1.2 summarizes the key mechanical parameters of these platforms, comparing their thermal and gravitational decoherence times. Among these platforms, *softly-clamped* silicon-nitride membranes [Tsa+17; Iva20] are particularly promising: thanks to high tensile stress and phononic patterning, they reach quality factors $Q \sim 10^9$ at cryogenic temperatures, yielding thermal decoherence times τ_dec on the order of 100 ms. Crucially, these membranes also exhibit zero-point motion x_zpf comparable to the nuclear scale, and satisfy $\tau_\mathrm{dec} \gg \tau_G$, making them the only system known to simultaneously fulfill both requirements for observing gravitational decoherence.

In order to engineer a generic quantum state of a mechanical oscillator, one can resonantly couple it to a fully controlled qubit [Hof+09][Chu+18]. These considerations motivate the central theme of this work: coupling the mechanical resonator to a superconducting qubit operating in the MHz frequency range, such that their frequencies are closely matched.

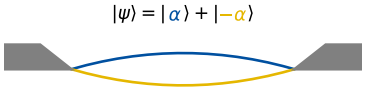

Fig. 1.1: Illustration of the membrane in a Schrödinger-cat state. Experimental test with a mechanical oscillator initialized in a Schrödinger-cat state being a superposition of two coherent states with opposite phases (color-code consistent with the one of Fig. **??**).

[HR06]: Haroche et al. (2006), *Exploring the Quantum*, Oxford University Press
[GS21]: Gely et al. (2021), 'Superconducting electro-mechanics to test Diósi–Penrose effects of general relativity in massive superpositions', AVS Quantum Science

[Tsa+17]: Tsaturyan et al. (2017), 'Ultracoherent nanomechanical resonators via soft clamping and dissipation dilution', Nature Nanotechnology
[Iva20]: Ivanov (2020), 'Fabrication of patterned silicon nitride nanomembranes at the LKB', Zenodo

[Hof+09]: Hofheinz et al. (2009), 'Synthesizing arbitrary quantum states in a superconducting resonator', Nature

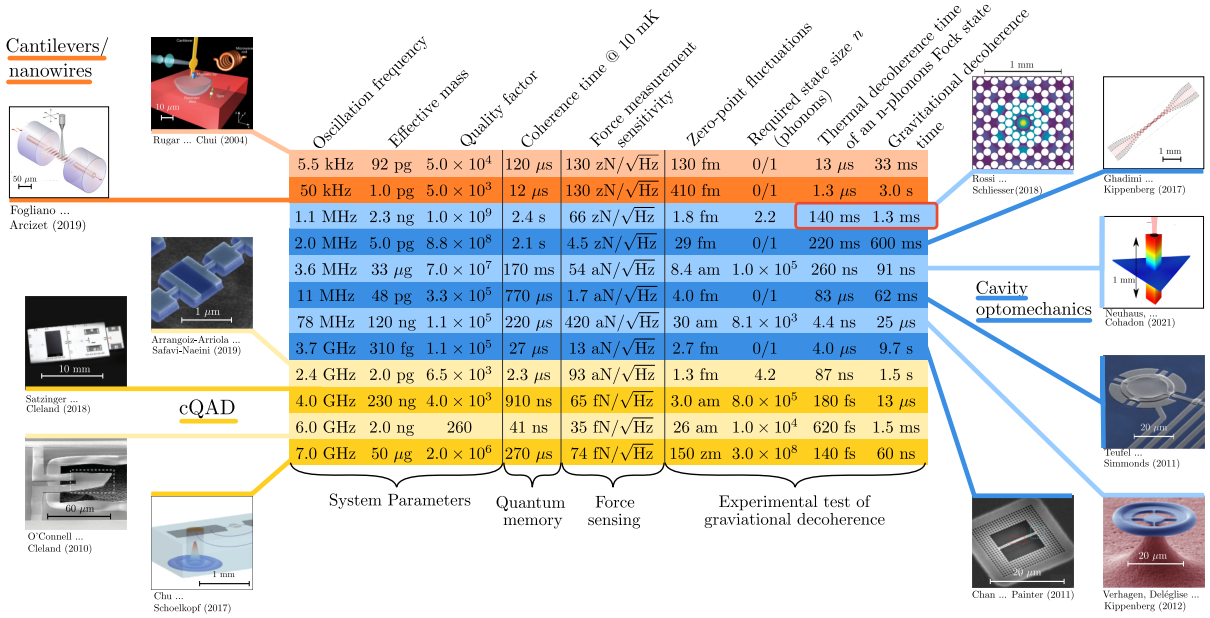

Fig. 1.2: Comparison of various mechanical systems for quantum memory, sensing, and quantum-gravity tests, focusing on three key parameters: state size, zero-point fluctuation amplitude (x_{zpf}), and thermal and gravitational decoherence times. The force sensitivity is calculated as $F_{\text{min}} = \sqrt{\hbar\, m\, \omega_m^2 (2n_{\text{th}} + 1)/Q}$ (see Ref. [Sid+95]). The required cat-state size for quantum-gravity experiments is $n = \left(2.7\,\text{fm}/x_{\text{zpf}}\right)^2$. The thermal decoherence time is obtained from Eq. 1.5, and the gravitational decoherence time from Eq. 1.4. For systems where small cat states ($n < 1$) or the ground state ($n = 0$) suffice, we denote $n = 0/1$ and compute τ_{dec} for the $|1\rangle$ state. The system proposed in this work is highlighted in red. Data and images are from Rossi et al. [Ros+18], Neuhaus et al. [Neu+21], Fogliano et al. [Fog+21], Teufel et al. [Teu+11], Rugar et al. [Rug+04], Arrangoiz et al. [Arr+19], O'Connell et al. [OCo+10], Satzinger et al. [Sat+18], Chan et al. [Cha+11], Chu et al. [Chu+17], Verhagen et al. [Ver+12], and Ghadimi et al. [Gha+18]. Adapted from Ref. [GS21].

Interfacing low-frequency mechanical systems with superconducting qubits

Near-resonant coupling between a mechanical resonator and a superconducting circuit enables the coherent transfer and manipulation of single phonons using circuit Quantum Electrodynamics (cQED) techniques [OCo+10; Chu+18; Sat+18; Bie+19]. While GHz-frequency implementations have achieved macroscopically distinguishable mechanical cats [Bil+23]. However, probing gravitational decoherence requires operation in the *megahertz* regime, which in turn demands a superconducting circuit element with a comparably low resonant frequency.

Indirect and dispersive interfaces Early attempts to couple the two systems were conducted in a regime where the qubit frequency significantly exceeded that of the mechanical resonator. Steele *et al.* coupled a 200 MHz mechanical mode to a transmon via an RF cavity [Gel+19], while Lehnert's group detected 20 MHz drum phonons by dispersive coupling to a Cooper-pair box at 4.6 GHz [VML18; Ma+21]. However, the CPB, coherence time ∼40 ns-limited by low-frequency charge noise—precludes coherent phonon operations.

The heavy fluxonium solution A direct, resonant interface becomes possible with the heavy fluxonium qubit [Ear+18; Zha+21]. As depicted in Fig. 1.3, the circuit consists of a Josephson junction (energy E_J) shunted

[OCo+10]: O'Connell et al. (2010), 'Quantum ground state and single-phonon control of a mechanical resonator', Nature
[Chu+18]: Chu et al. (2018), 'Creation and control of multi-phonon Fock states in a bulk acoustic-wave resonator', Nature
[Sat+18]: Satzinger et al. (2018), 'Quantum control of surface acoustic-wave phonons', Nature
[Bie+19]: Bienfait et al. (2019), 'Phonon-mediated quantum state transfer and remote qubit entanglement', Science
[Bil+23]: Bild et al. (2023), 'Schrödinger cat states of a 16-microgram mechanical oscillator', Science
[Gel+19]: Gely et al. (2019), 'Observation and stabilization of photonic Fock states in a hot radio-frequency resonator', Science
[VML18]: Viennot et al. (2018), 'Phonon number sensitive electromechanics', Physical Review Letters
[Ma+21]: Ma et al. (2021), 'Non-classical energy squeezing of a macroscopic mechanical oscillator', Nature Physics

by a large capacitance (energy E_C) and a superinductor (energy E_L), with $E_L < E_C \ll E_J$. Its Hamiltonian, expressed in terms of the superconducting phase operator $\boldsymbol{\varphi}$ across the inductor and the dimensionless charge operator $\boldsymbol{n}$ on the capacitor, is

$$\boldsymbol{H} = 4\,E_C\,\boldsymbol{n}^2 - E_J\,\cos\!\left(\boldsymbol{\varphi} - \varphi_{\text{ext}}\right) + \frac{E_L}{2}\,\boldsymbol{\varphi}^2.$$

Here φ_{ext} is the externally applied flux bias. At half-flux bias ($\varphi_{\text{ext}}/2\pi = 0.5$), $\boldsymbol{\varphi}$ and $\boldsymbol{n}$ experience a double-well potential; the two lowest circuit eigenstates are the symmetric and antisymmetric combinations of the local ground states in each well. The tunneling rate varies exponentially with the circuit parameters, such that the qubit splitting ω_q can be experimentally reduced to a few megahertz. Previously, we have shown that heavy fluxonium qubits are highly sensitive to charge modulations at MHz frequencies, enabling the detection of charge fluctuations as small as $33\,\mu\text{e}/\sqrt{\text{Hz}}$ [Naj+24]. Thus, they are excellent candidates for capacitive resonant coupling to an MHz-frequency mechanical membrane.

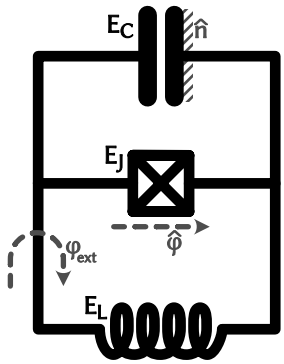

Fig. 1.3: Lumped-element equivalent circuit of the fluxonium qubit. It consists of a capacitor (C), an inductor (L), and a Josephson junction (J) connected in parallel. The phase $\hat{\varphi}$ is the superconducting phase difference between the left and right nodes. Depending on the gauge choice, an additional phase offset φ_{ext} may be included in the Josephson-junction or inductive terms.

Resonant electromechanical coupling. In the project of this thesis (*Mecaflux*), we have chosen to couple the fluxonium qubit to the first antisymmetric mode of the mechanical membrane, as illustrated in Fig. 1.4. The capacitive coupling is realized via a flip-chip assembly with a small separation $d \approx 2.5\,\mu$m and by applying a DC bias voltage V_b to the membrane. This coupling appears as an additional charge offset in the fluxonium circuit, modulated by the mechanical displacement operator $\boldsymbol{X}$. One can show that, to first order in X/X_{zpf}, the coupling term is

$$H_{\text{coupling}} = 8\,E_C\,\frac{\partial C_m^{+}}{\partial X}\,\frac{V_b}{2e}\,\boldsymbol{n}\,\boldsymbol{X}.$$

This linear coupling implies that the interaction between the fluxonium qubit and the membrane resonator (with ladder operators $\boldsymbol{a}$ and $\boldsymbol{a}^\dagger$) is described by the Jaynes–Cummings Hamiltonian,

$$\boldsymbol{H}_{JC}/\hbar = \omega_m\,\boldsymbol{a}^\dagger\boldsymbol{a} + \frac{\omega_q}{2}\,\boldsymbol{\sigma}_z + i\,\frac{\Omega}{2}\left(\boldsymbol{\sigma}_-\,\boldsymbol{a}^\dagger - \boldsymbol{\sigma}_+\,\boldsymbol{a}\right),$$

where the effective vacuum Rabi frequency Ω increases with the bias voltage V_b and decreases with the separation d, scaling as $\Omega \propto \omega_q\,V_b\,X_{\text{zpf}}/d$.

Earlier results in the group

A 1.8 MHz fluxonium qubit. The first fluxonium qubit from our group exhibited a qubit transition at 1.8 MHz with a relaxation time $T_1 = 34\,\mu$s and a dephasing time $T_2^* = 39\,\mu$s [Naj+24]. It was fabricated by a former PhD student, B.-L. Najera-Santos. In this work, in which I was fully involved during my second year of the PhD, we demonstrated the high charge sensitivity of our qubit, $33\,\mu\text{e}/\sqrt{\text{Hz}}$. At that time, we demonstrated that, in a hybrid electromechanical system where the membrane motion would manifest as charge modulation on the qubit, the measured charge sensitivity would, in principle, enable probing the membrane dynamics with the qubit.

[Ear+18]: Earnest et al. (2018), 'Realization of a Λ System with Metastable States of a Capacitively Shunted Fluxonium', Physical Review Letters

[Zha+21]: Zhang et al. (2021), 'Universal fast flux control of a coherent, low-frequency qubit', Physical Review X

[Naj+24]: Najera-Santos et al. (2024), 'High-Sensitivity ac-Charge Detection with a MHz-Frequency Fluxonium Qubit', Physical Review X

Softly-clamped mechanical membranes. Previously, Edouard Ivanov and Thibault Capelle, in their PhD thesis [Iva21; Cap20], developed a technique to fabricate silicon-nitride membranes in our group. Soft clamping (to minimize the effect of the substrate) is achieved with phononic patterning. It was demonstrated that these membranes can achieve quality factors of around $Q \sim 10^7$ at room temperature, given thermal decoherence time sufficiently higher than the gravitational decoherence time. During the first year of my PhD, I continued to develop this technique: we tested different membrane designs inspired by the work of Seis *et al.* [Sei+22], trying to maximize phononic shielding of the central mode, and rebuilt the room-temperature measurement setup.

To validate the electrostatic bias-coupling scheme, another PhD student, Himanshu Patange [Pat25], who began his PhD the same year, investigated a simpler system: a 3 MHz mechanical resonator coupled to a 7 GHz LC circuit. This study validated key technological components of the *Mecaflux* project—a flip-chip assembly process was established, and a bias-line filtering was characterized to confirm the achievable noise performance—informing the refined design presented in Chap. 5. Although we initially focused on a softly clamped membrane to maximize mechanical quality factors, Patange's results showed that reaching this goal would require another iteration of the phononic-crystal membrane design[1]. Consequently, to simplify the microwave design, we chose to focus on square membranes, deferring softly-clamped membrane coupling to future enhancements aimed at reaching the regime where thermal coherence exceeds gravitational coherence.

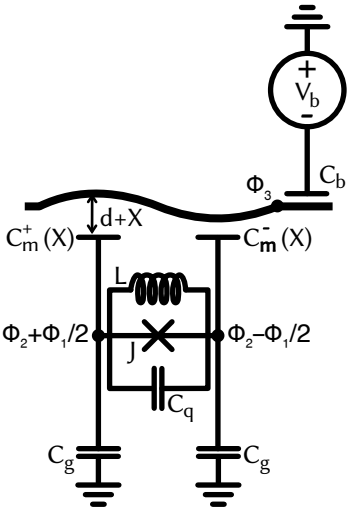

Fig. 1.4: Antisymmetric qubit-membrane coupling circuit. Schematic of a fluxonium qubit formed by a Josephson junction J in parallel with an inductance L and capacitance C_q, bridging two nodes held at flux biases $\phi_2 \pm \frac{\phi_1}{2}$. Each node is grounded through a capacitance C_g. An antisymmetric mechanical membrane mode—displaced by $\pm X$ from its equilibrium separation d—serves as an electrode and couples to both qubit nodes via capacitances $C_m^{\pm}(X)$. The membrane node ϕ_3 is biased by a DC voltage V_b through a coupling capacitor C_b.

Objectives of this work

Thus, during the first two years of my PhD, I worked on various aspects of the *Mecaflux* project. Given all the successes and challenges briefly described above, at the beginning of my third year, the main goal of my PhD was established: to achieve resonant coupling between a mechanical resonator operating at MHz frequencies (a square membrane) and a superconducting qubit (a heavy fluxonium). This goal can be divided into the following sub-objectives:

- Design a platform that couples a low-frequency heavy fluxonium qubit with a mechanical membrane.
- Demonstrate, for the first time, resonant coupling between a MHz-frequency mechanical resonator and a superconducting qubit.
- Probe the quantum motion of the mechanical oscillator.
- Identify the main technical bottlenecks and propose key improvements to achieve stronger coupling in future iterations.

1.2 Outline of the thesis

This thesis is organized as follows. Chapter 2 introduces the theoretical formalism used throughout the work, including the harmonic oscillator

[Iva21]: Ivanov (2021), 'Optimization of silicon nitride membranes for hybrid superconducting-mechanical circuits', Thesis
[Cap20]: Capelle (2020), 'Electromechanical cooling and parametric amplification of an ultrahigh-Q mechanical oscillator', Thesis
[Sei+22]: Seis et al. (2022), 'Ground state cooling of an ultracoherent electromechanical system', Nature Communications
[Pat25]: Patange (2025), 'RF Transduction in a DC-Biased Superconducting Electromechanical System', Thesis

1: I initially designed a softly-clamped membrane with a first vibrational mode at 5 MHz, but Patange's experiment showed that coupling to the symmetric mechanical mode was not feasible due to induced noise. Thus, to achieve coupling to the antisymmetric mode, one must redesign the phononic shield with two defects, necessitating an additional cycle of phononic crystal membrane fabrication and testing which could take another year of development.

notation, nonlinear circuit elements, qubit description, and basic quantum control techniques. Chapter 3 presents the fluxonium qubit—our primary platform—providing physical intuition for its key parameters and analysing its coupling to the environment. In Chapter 4, we describe the design and properties of the square membrane that will later interact with the qubit. Chapter 5 details the fluxonium chip layout and design considerations, while Chapter 6 covers fabrication procedures and the experimental setup.

The core experimental results begin in Chapter 7, where we report the basic characterization of the qubit. Chapter 8 explores the qubit–membrane interaction in both the dispersive regime and the resonant regime with a large ($\sim 10^5$ phonons) membrane coherent state. Chapter 9 presents the main findings of this thesis: the quantum motion of the mechanical resonator probed via its absorption and emission spectra. Finally, Chapter 10 summarizes our conclusions, discusses the impact of our work, and outlines potential improvements and future directions.

1.3 Publications

- B.-L. Najera-Santos, R. Rousseau, K. Gerashchenko, H. Patange, A. Riva, M. Villiers, T. Briant, P. -F. Cohadon, A. Heidmann, J. Palomo, M. Rosticher, H. le Sueur, A. Sarlette, W. C. Smith, Z. Leghtas, E. Flurin, T. Jacqmin, S. Deléglise, *High-Sensitivity ac-Charge Detection with a MHz-Frequency Fluxonium Qubit*, Phys. Rev. X 14, 011007 (2024)
- P. Manset, J. Palomo, A. Schmitt, K. Gerashchenko, R. Rousseau, H. Patange, P. Abgrall, E. Flurin, S. Deléglise, T. Jacqmin, L. Balembois, *Hyperinductance based on stacked Josephson junctions*, preprint, *arXiv:2505.02764* (2025)
- K. Gerashchenko*, R. Rousseau*, L. Balembois*, H. Patange, P. Manset, W. C. Smith, Z. Leghtas, E. Flurin, T. Jacqmin, S. Deléglise, *Probing the quantum motion of a macroscopic mechanical oscillator with a radio-frequency superconducting qubit*, preprint, *arXiv:2505.21481* (2025)
- H. Patange, K. Gerashchenko, R. Rousseau, P. Manset, L. Balembois, T. Capelle, S. Deléglise, T. Jacqmin, *Cryogenic RF-to-Microwave Transducer based on a DC-Biased Electromechanical System*, preprint, *arXiv:2508.01066* (2025)

Quantum harmonic and non-harmonic oscillators

2

The main goal of this thesis is to achieve quantum control of a mechanical resonator. This chapter introduces the theoretical foundations of the harmonic oscillator and the superconducting circuits used in this work, as well as the formalism for describing the qubit.

Outline. This chapter begins with the classical harmonic oscillator and its quantization, starting from the Lagrangian of the system and moving to the Hamiltonian description—a framework used throughout this thesis to derive system Hamiltonians. We then discuss the equally spaced energy levels of the harmonic oscillator, illustrating the fundamental asymmetry introduced by the ladder operators – a key feature that leads to the asymmetric spectrum measured in this work. Next, we introduce the LC resonator and the Josephson junction, which introduce nonlinearity and lift the degeneracy of the harmonic oscillator. We build intuition for superconducting circuits using the transmon qubit as an example, explaining its main features and limitations. We then introduce the two-level system framework, which provides the main formalism for describing the qubit throughout this thesis, including the Bloch sphere representation of its state, the fundamentals of qubit control and readout, decoherence times (T_1, T_2^*), and basic experimental characterization sequences. Finally, this chapter concludes with an introduction to the quantum electrodynamics framework, the primary formalism for describing interactions between the qubit and harmonic oscillators (such as LC resonators or mechanical membranes).

2.1 Quantum Harmonic Oscillator

In this section, we introduce the quantum harmonic oscillator by deriving and quantizing the Hamiltonian of a mass–spring system, revealing its equally spaced energy levels and zero–point fluctuations. We then map this framework onto an LC circuit—identifying flux and charge as conjugate variables—to prepare for the addition of nonlinearity in later sections.

2.1.1 Hamiltonian of a Mechanical Oscillator

We begin with the simplest mechanical resonator: a mass m attached to a spring of stiffness k (see Fig. 2.1). When displaced by X from its equilibrium position, the spring exerts a restoring force $-k\,X$ on the mass. By definition, the kinetic energy T is the energy of motion and depends on the velocity $\dot{X}$, $T = \frac{1}{2}\,m\,\dot{X}^2$, and the potential energy V is stored in the spring, proportional to the displacement, $V = \frac{1}{2}\,k\,X^2$.

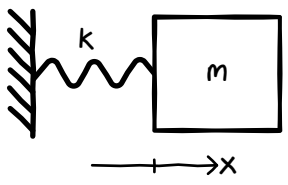

Fig. 2.1: Mechanical oscillator. A mass m attached to a spring of stiffness k.

Identifying these two energies, the Lagrangian $L = T - V$ becomes

$$\boxed{L = \tfrac{1}{2}\, m\, \dot{X}^2 - \tfrac{1}{2}\, k\, X^2 = \tfrac{1}{2}\, m\, \dot{X}^2 - \tfrac{1}{2}\, m\, \omega_m^2\, X^2, \quad \omega_m^2 \equiv \frac{k}{m},} \tag{2.1}$$

where ω_m denotes the mechanical oscillator's natural angular frequency.

The momentum conjugate to X is defined by

$$P \equiv \frac{\partial L}{\partial \dot{X}} = m\, \dot{X}. \tag{2.2}$$

achieving the common definition of momentum.

In classical mechanics, the equations of motion follow from the Euler–Lagrange equation:

$$\frac{d}{dt}\big(P\big) - \frac{\partial L}{\partial X} = 0 \implies m\, \ddot{X} + k\, X = 0,$$

which simplifies to

$$\ddot{X} + \omega_m^2\, X = 0.$$

The general solution is a simple harmonic motion,

$$X(t) = X_0\, \sin\big(\omega_m t + \theta\big),$$

with arbitrary amplitude X_0 and phase θ.

To prepare for quantization, we switch to the Hamiltonian formalism, treating X and P as independent canonical variables:

$$H = P\, \dot{X} - L = \frac{P^2}{2m} + \tfrac{1}{2}\, m\, \omega_m^2\, X^2. \tag{2.3}$$

Classically, H represents the total energy of the closed system. In quantum mechanics, we promote X and P to operators $\boldsymbol{X}$ and $\boldsymbol{P}$ satisfying:

$$[\boldsymbol{X}, \boldsymbol{P}] = i\hbar, \tag{2.4}$$

encoding the Heisenberg uncertainty principle: $\Delta X\, \Delta P \geq \frac{1}{2}\hbar$. Since $\boldsymbol{X}$ and $\boldsymbol{P}$ no longer commute, $P = m\dot{X}$ cannot be maintained and we cannot define a classical trajectory using Euler-Lagrange equations. Instead, $\boldsymbol{X}$ and $\boldsymbol{P}$ are treated as independent canonical variables, and the dynamics are encoded in the Hamiltonian

$$\boldsymbol{H} = \frac{\boldsymbol{P}^2}{2m} + \tfrac{1}{2}\, m\, \omega_m^2\, \boldsymbol{X}^2. \tag{2.5}$$

To diagonalize $\boldsymbol{H}$, we define dimensionless ladder operators $\boldsymbol{a}$, $\boldsymbol{a}^\dagger$:

$$\boldsymbol{a} = \sqrt{\frac{m\omega_m}{2\hbar}}\, \boldsymbol{X} + i \sqrt{\frac{1}{2m\hbar\omega_m}}\, \boldsymbol{P}, \quad \boldsymbol{a}^\dagger = \sqrt{\frac{m\omega_m}{2\hbar}}\, \boldsymbol{X} - i \sqrt{\frac{1}{2m\hbar\omega_m}}\, \boldsymbol{P}. \tag{2.6}$$

They satisfy

$$\boxed{[\boldsymbol{a}, \boldsymbol{a}^\dagger] = 1.} \tag{2.7}$$

Inverting gives

$$\boldsymbol{X} = \sqrt{\frac{\hbar}{2m\omega_m}}\,(\boldsymbol{a} + \boldsymbol{a}^\dagger), \quad \boldsymbol{P} = -\,i\sqrt{\frac{\hbar m\omega_m}{2}}\,(\boldsymbol{a} - \boldsymbol{a}^\dagger). \tag{2.8}$$

Substituting into Eq. 2.3 yields the quantum Hamiltonian

$$\boxed{\boldsymbol{H} = \hbar\,\omega_m\left(\boldsymbol{a}^\dagger\boldsymbol{a} + \tfrac{1}{2}\right).} \tag{2.9}$$

This finding deserves particular attention, as any system whose Hamiltonian is quadratic in $\boldsymbol{X}$ and $\boldsymbol{P}$ can be mapped onto such harmonic oscillator Hamiltonian. In particular, this form applies directly to the LC circuit (with flux playing the role of X and charge playing the role of P), as will be shown later.

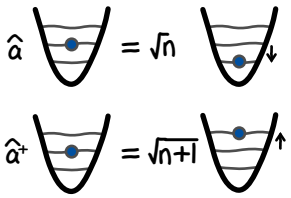

Fig. 2.2: Illustration of the ladder operators $\boldsymbol{a}$ and $\boldsymbol{a}^\dagger$ acting on the state $|n\rangle$.

2.1.2 Energy Spectrum and Zero–Point Fluctuations

The diagonal form of the harmonic oscillator Hamiltonian (Eq. 2.9) can be used to calculate the Harmonic oscillator spectrum. Indeed, following Dirac's elegant argument [1], one shows that the number operator $\boldsymbol{N} = \boldsymbol{a}^\dagger\boldsymbol{a}$ must have integer eigenvalues purely from algebraic considerations based on the commutation relation $[\boldsymbol{a}, \boldsymbol{a}^\dagger] = 1$. Let's denote by $\{|n\rangle\}$ the orthonormal eigenbasis of $\boldsymbol{N}$, with $\boldsymbol{N}\,|n\rangle = n\,|n\rangle$, $n = 0, 1, 2, \ldots$. The ladder operators then satisfy [2]

$$\boxed{\boldsymbol{a}\,|n\rangle = \sqrt{n}\;|n-1\rangle\,, \quad \boldsymbol{a}^\dagger|n\rangle = \sqrt{n+1}\;|n+1\rangle} \tag{2.10}$$

Applying Eq. 2.9 to $|n\rangle$ gives

$$\boldsymbol{H}\,|n\rangle = \hbar\omega_m\left(n + \frac{1}{2}\right)|n\rangle\,, \tag{2.11}$$

so the energy eigenvalues are

$$\boxed{E_n = \hbar\omega_m\left(n + \frac{1}{2}\right),} \quad \text{with} \quad n \in \mathbb{N}. \tag{2.12}$$

These energy eigenvalues form an evenly spaced "ladder" with spacing $\hbar\omega_m$, as is illustrated in Fig. 2.3. Although the lowest point of the potential parabola is at zero, the lowest energy level of the harmonic oscillator, $E_0 = 1/2\,\hbar\omega_m$, is nonzero and represents the zero-point energy. This gives rise to intrinsic *zero-point fluctuations* in position and momentum:

$$x_{\text{zpf}}^2 \equiv \langle 0|\boldsymbol{x}^2|0\rangle\,, \quad p_{\text{zpf}}^2 \equiv \langle 0|\boldsymbol{p}^2|0\rangle\,,$$

with

$$\boxed{x_{\text{zpf}} = \sqrt{\frac{\hbar}{2m\omega_m}}, \quad p_{\text{zpf}} = \sqrt{\frac{\hbar m\omega_m}{2}}.} \tag{2.13}$$

This quantization and spectrum discussion is fundamental to any harmonic oscillator. Since every transition $|n\rangle \rightarrow |n \pm 1\rangle$ occurs at the

1: $\boldsymbol{N}$ is positive definite, i.e.

$$\langle\psi|\boldsymbol{a}^\dagger\boldsymbol{a}|\phi\rangle = \|\,\boldsymbol{a}\,|\psi\rangle\,\|^2 \geq 0.$$

Hence, the eigenvalues of $\boldsymbol{N}$ are all greater than or equal to zero. Denoting by $|\lambda\rangle$ an eigenvector of $\boldsymbol{N}$ with eigenvalue $\lambda > 0$, and using the commutation relation $[\boldsymbol{N}, \boldsymbol{a}] = -\boldsymbol{a}$, we obtain

$$\boldsymbol{N}\,\boldsymbol{a}\,|\lambda\rangle = \big([\boldsymbol{N}, \boldsymbol{a}] + \boldsymbol{a}\,\boldsymbol{N}\big)\,|\lambda\rangle = (\lambda - 1)\,\boldsymbol{a}\,|\lambda\rangle\,.$$

In other words, $\boldsymbol{a}\,|\lambda\rangle$ is an eigenvector of $\boldsymbol{N}$ with eigenvalue $\lambda - 1$, unless $\lambda = 0$, in which case $\boldsymbol{a}\,|\lambda\rangle = 0$. We can then prove by contradiction that λ must be an integer: starting from a state $|\lambda\rangle$ with non-integer λ, repeated application of $\boldsymbol{a}$ lowers the eigenvalue by one each time, eventually producing a negative eigenvalue, which is impossible. The only way to terminate this descending ladder is to reach the state $|0\rangle$ with eigenvalue 0; this happens only if $\lambda \in \mathbb{N}$. Similarly, using the commutation relation $[\boldsymbol{N}, \boldsymbol{a}^\dagger] = \boldsymbol{a}^\dagger$, one shows that $\boldsymbol{a}^\dagger\,|\lambda\rangle$ is an eigenvector of $\boldsymbol{N}$ with eigenvalue $\lambda + 1$. Hence, by repeated application of $\boldsymbol{a}^\dagger$, one generates an infinite ladder of eigenstates called Fock states $\{|n\rangle\} = \boldsymbol{a}^{\dagger n}\,|0\rangle\,/\sqrt{n!}$, with positive integer eigenvalues.

2: Since $\boldsymbol{a}\,|n\rangle$ is an eigenvector of $\boldsymbol{N}$ with eigenvalue $n - 1$, $\boldsymbol{a}\,|n\rangle \propto |n-1\rangle$. To find the normalization factor, one can evaluate the norm

$$\|\,\boldsymbol{a}\,|n\rangle\,\|^2 = \langle n|\boldsymbol{N}|n\rangle = n.$$

Similarly, one can show an analogous relation for $\boldsymbol{a}^\dagger\,|n\rangle$.

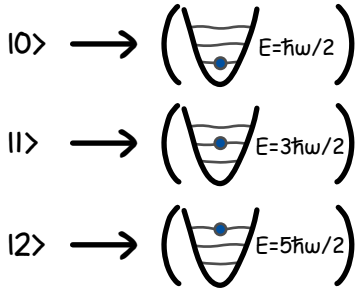

Fig. 2.3: Illustration of the first energy levels of the harmonic oscillator.

same frequency ω_m, individual transitions cannot be addressed selectively, precluding quantum control. Breaking this symmetry requires introducing a suitable nonlinearity into the system. In mechanical implementations, however, realizations of such nonlinear elements [Som+25; Bag+11; Asa+17] requires large motion amplitudes that are much greater than the zero-point fluctuations. In the following, we therefore turn to nonlinear platforms—where quantum control is well established—and show how coupling a harmonic oscillator to these nonlinear systems enables probing of its quantum dynamics. By the end of this work, we will present experimental results demonstrating the intrinsic asymmetry of the ladder operators Eq. 2.10: the emission amplitude $|n\rangle \rightarrow |n-1\rangle$ scales as $\sqrt{n}$, while the absorption amplitude $|n\rangle \rightarrow |n+1\rangle$ scales as $\sqrt{n+1}$. This asymmetry reveals the quantum character of the harmonic oscillator. To quantify it explicitly, we introduce the spectral noise density.

[Som+25]: Sommer et al. (2025), 'Spatial Mapping of Intrinsic and Readout Nonlinearities in a Strongly Driven Micromechanical Membrane', Physical Review Letters
[Bag+11]: Bagheri et al. (2011), 'Dynamic manipulation of nanomechanical resonators in the high-amplitude regime and non-volatile mechanical memory operation', Nature Nanotechnology
[Asa+17]: Asadi et al. (2017), 'Mechanism of geometric nonlinearity in a nonprismatic and heterogeneous microbeam resonator', Physical Review B

2.1.3 LC Quantum Harmonic Oscillator

As outlined in the introduction, we couple the mechanical oscillator to an electrical resonator. The simplest version of such a resonator is the LC circuit—the fundamental building block of superconducting qubits—which consists of a capacitor C_0 and an inductor L_0 in parallel (see Fig. 2.4). This single–degree–of–freedom circuit can be described either by its voltage $V(t)$ or equivalently by its current $I(t)$. In circuit quantum electrodynamics (cQED), however, it is more natural to use the magnetic flux Φ through the inductor and the electric charge Q on the capacitor as the canonical conjugate variables. These are related by Faraday's law and charge conservation:

$$V(t) = \dot{\Phi}(t), \quad I(t) = -\dot{Q}(t). \tag{2.14}$$

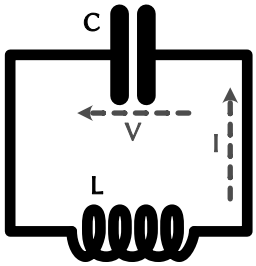

Fig. 2.4: Electrical circuit of the LC-resonator. It consists of a capacitor (C) and an inductor (L) connected in parallel.

By definition, the capacitance and inductance satisfy

$$\begin{aligned} C_0 = \frac{Q}{V} = \frac{Q}{\dot{\Phi}} \quad &\Longrightarrow \quad \dot{\Phi} = \frac{Q}{C_0}, \\ L_0 = \frac{\Phi}{I} = -\frac{\Phi}{\dot{Q}} \quad &\Longrightarrow \quad \dot{Q} = -\frac{\Phi}{L_0}. \end{aligned} \tag{2.15}$$

The Lagrangian is given by the difference between the kinetic and potential energy:

$$L = T - U = \int_0^Q V dQ' - \int_0^\Phi I d\Phi'.$$

Hence, the Lagrangian expressed in the flux variable Φ is

$$\boxed{L = \frac{1}{2} C_0 \dot{\Phi}^2 - \frac{1}{2} \frac{\Phi^2}{L_0}.} \tag{2.16}$$

From this Lagrangian, the momentum conjugate to Φ follows directly:

$$\frac{\partial L}{\partial \dot{\Phi}} = C_0 \dot{\Phi} \equiv Q, \tag{2.17}$$

confirming that Q and Φ form a canonical pair. Equivalently, one may formulate a dual Lagrangian in terms of Q. Using the Euler–Lagrange equation applied to the Lagrangian above, the circuit equations of motion admit the harmonic solutions

$$\begin{aligned} \Phi(t) &= \Phi_0 \sin(\omega_0 t + \theta), \\ Q(t) &= C_0 \dot{\Phi}(t) = \Phi_0 \, C_0 \, \omega_0 \cos(\omega_0 t + \theta) = \frac{\Phi_0}{Z_0} \cos(\omega_0 t + \theta), \end{aligned} \tag{2.18}$$

where we have introduced the resonance frequency and characteristic impedance,

$$\omega_0 \equiv \frac{1}{\sqrt{L_0 C_0}}, \qquad Z_0 \equiv \sqrt{\frac{L_0}{C_0}}. \tag{2.19}$$

Having identified these constants, we perform the Legendre transformation to obtain the classical Hamiltonian in the conjugate variables Φ and Q, and apply Dirac quantization by promoting them to operators. The resulting Hamiltonian in the flux and charge variables, which satisfies the commutation relation $[\boldsymbol{\Phi}, \boldsymbol{Q}] = i\hbar$, is

$$\boldsymbol{H} = \frac{\boldsymbol{Q}^2}{2C_0} + \frac{\boldsymbol{\Phi}^2}{2L_0} = \frac{\omega_0}{2}\left(Z_0 \, \boldsymbol{Q}^2 + \tfrac{1}{Z_0} \, \boldsymbol{\Phi}^2\right). \tag{2.20}$$

This form is directly analogous to the mechanical oscillator Hamiltonian described earlier, with Q playing the role of momentum and Φ that of position; C_0 corresponds to mass and $1/L_0$ to the spring constant.

Tab. 2.1: Mechanical-electrical harmonic oscillator analogy.

Mechanical		LC
$\boldsymbol{x}$	$\longleftrightarrow$	$\boldsymbol{\Phi}$
$\boldsymbol{p}$	$\longleftrightarrow$	$\boldsymbol{Q}$
m	$\longleftrightarrow$	C_0
k	$\longleftrightarrow$	$1/L_0$

We then define the dimensionless ladder operators

$$\boxed{\boldsymbol{c} = \sqrt{\frac{1}{2\hbar Z_0}}\boldsymbol{\Phi} + i\sqrt{\frac{Z_0}{2\hbar}}\boldsymbol{Q}, \quad \boldsymbol{c}^\dagger = \sqrt{\frac{1}{2\hbar Z_0}}\boldsymbol{\Phi} - i\sqrt{\frac{Z_0}{2\hbar}}\boldsymbol{Q}\,,} \tag{2.21}$$

which satisfy

$$\boxed{[\boldsymbol{c}, \boldsymbol{c}^\dagger] = 1\,.} \tag{2.22}$$

Inverting Eq. 2.21 gives

$$\boldsymbol{\Phi} = \Phi_{\text{zpf}}(\boldsymbol{c} + \boldsymbol{c}^\dagger), \quad \boldsymbol{Q} = -iQ_{\text{zpf}}(\boldsymbol{c} - \boldsymbol{c}^\dagger), \tag{2.23}$$

$$\text{with } \Phi_{\text{zpf}} \equiv \sqrt{\frac{\hbar Z_0}{2}} \quad Q_{\text{zpf}} \equiv \sqrt{\frac{\hbar}{2Z_0}}.$$

Substituting back into the Hamiltonian (2.20) yields the standard quantum harmonic oscillator Hamiltonian,

$$\boldsymbol{H} = \hbar\omega_0\left(\boldsymbol{c}^\dagger \boldsymbol{c} + \tfrac{1}{2}\right). \tag{2.24}$$

As discussed in Sec. 2.1.2, this Hamiltonian supports an equidistant energy spectrum and therefore does not permit selective quantum control. To break this symmetry, one must introduce a nonlinear circuit element. In the following section, we introduce the Josephson junction and show how its nonlinearity enables the construction of a qubit.

2.2 Josephson junction: The key non-linear element

In 1962, during his PhD, B. Josephson predicted that a dissipationless supercurrent of Cooper pairs can tunnel through a thin insulating barrier between two superconductors, even when the normal-state single-electron transmission is negligible (see Fig. 2.5). This supercurrent is driven by the superconducting phase difference $\varphi = \varphi_2 - \varphi_1$ between the two superconducting electrodes and is proportional to the normal-state tunnel conductance of the junction [Jos62]. Josephson's work led to the celebrated Josephson relations, which determine the electrical behavior of such junctions:

$$\begin{aligned} I &= I_c \sin(\varphi), \\ \frac{\partial \varphi}{\partial t} &= \frac{2e}{\hbar} V \;=\; \frac{V}{\phi_0}, \end{aligned} \tag{2.25}$$

where I and V are the junction current and voltage, respectively and

$$\Phi_0 = \frac{h}{2e}, \quad \phi_0 = \frac{\Phi_0}{2\pi} = \frac{\hbar}{2e}$$

is the (reduced) flux quantum. I_c is the critical current of Cooper pairs through the junction,

$$I_c = \frac{\pi \Delta}{2e\, R_N}, \tag{2.26}$$

Here, Δ is the Bardeen–Cooper–Schrieffer (BCS) gap energy of the superconducting electrodes, and R_N is the normal-state resistance of the junction. The resistance R_N varies with the oxidation-layer thickness, which can be tuned during fabrication, and is inversely proportional to the junction area, set by the design.

Since $V = \dot{\Phi}$, Eq. 2.25 allow us to define the phase across the junction as a function of the flux

$$\boxed{\varphi = \frac{\Phi}{\phi_0}\,.} \tag{2.27}$$

This dimensionless phase φ will be our preferred dynamical variable for describing the circuit. To determine the energy stored in the Josephson junction, we first differentiate the current–phase relation Eq. 2.25 with respect to time:

$$\frac{dI}{dt} = \frac{dI}{d\varphi}\frac{d\varphi}{dt} = I_c \cos(\varphi)\,\dot{\varphi} = \frac{I_c}{\phi_0}\cos(\varphi)\,V. \tag{2.28}$$

Comparing with the inductor relation $dI/dt = V/L$, we identify the phase-dependent Josephson inductance

$$L(\varphi) = \frac{\phi_0}{I_c \cos\varphi} \;\equiv\; L_J \frac{1}{\cos\varphi}, \quad L_J = \frac{\phi_0}{I_c}, \tag{2.29}$$

showing that the Josephson junction behaves as a non-linear inductor. For small phase excursions $\varphi \ll 1$, using Eq. 2.25 gives

$$L(\varphi) \approx L_J\left(1 + \frac{1}{2}\varphi^2\right) \approx L_J\left(1 + \frac{1}{2}\frac{I^2}{I_c^2}\right). \tag{2.30}$$

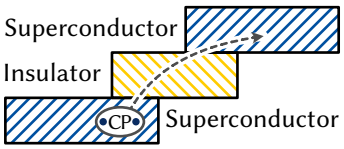

Fig. 2.5: Illustration of a Josephson junction that consists of an insulating barrier between two superconducting electrodes. Cooper pairs can tunnel through this barrier.

[Jos62]: Josephson (1962), 'Possible new effects in superconductive tunnelling', Physics Letters

Thus, in the weak-current regime $I \ll I_c$, the φ^2 correction is negligible and the junction behaves as a linear inductance L_J. The ratio I/I_c can be tuned by adjusting the junction size. In this work, we exploit both regimes of the Josephson junction: we use small junctions, for which the phase excursions are large, to add anharmonicity to the circuit—lifting the equidistant spacing of the oscillator spectrum and enabling quantum control; and we use large junctions, for which the phase excursions are small, to realize a typical linear inductance L_J.

Finally, the Josephson energy follows from

$$E = \int_{-\infty}^{t} V(t')\, I(t')\, dt' = \phi_0 I_c \int_{-\infty}^{t} \dot{\varphi}\, \sin\varphi\, dt', \tag{2.31}$$

giving the energy stored in the junction

$$\boxed{E = -\phi_0 I_c \cos\varphi = -E_J \cos\varphi, \quad E_J \equiv \phi_0 I_c = \frac{\phi_0^2}{L_J} = \frac{\Phi_0 I_c}{2\pi}.} \tag{2.32}$$

In the next section, we will embed this Josephson junction into a superconducting circuit, quantize the phase φ, derive its Hamiltonian, and thereby reveal how the junction's nonlinearity introduces anharmonicity into the energy spectrum.

2.3 The Transmon Qubit: Weakly anharmonic oscillator

As we have seen, the Josephson junction provides a dissipationless nonlinear element essential for building anharmonic oscillators. The Cooper–pair box (CPB) is the simplest such implementation. In this section, we derive the CPB Hamiltonian and show how its pronounced sensitivity to gate-charge noise leads to rapid dephasing. We then introduce the transmon regime, in which charge fluctuations are exponentially suppressed while the spectrum retains sufficient anharmonicity for qubit operation.

2.3.1 Cooper-pair box

The Cooper–pair box (CPB) is the simplest superconducting qubit, exploiting the nonlinearity of a Josephson junction. Its circuit (see Fig. 2.6) comprises a superconducting island connected to ground via a Josephson junction (with energy E_J) and a capacitance C_J, and coupled to a gate voltage V_d through a gate capacitance C_d. Using the node flux Φ (with $\varphi = \Phi/\phi_0$), the Lagrangian reads

$$L = \frac{C_J}{2}\dot{\Phi}^2 + \frac{C_d}{2}\big(\dot{\Phi} - V_d\big)^2 + E_J \cos\Big(\frac{\Phi}{\phi_0}\Big). \tag{2.33}$$

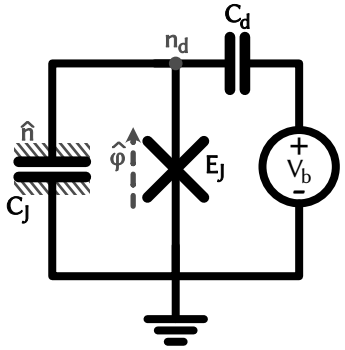

Fig. 2.6: Electrical circuit of the transmon qubit. It consists of a capacitor with energy E_C and a Josephson junction with energy E_J connected in parallel. Additionally, it is shunted by a gate capacitor C_d, biased at V_d, which tunes the offset charge n_d.

The conjugate charge is

$$Q = \frac{\partial L}{\partial \dot{\Phi}} = C_\Sigma \dot{\Phi} - C_d V_d, \quad C_\Sigma = C_J + C_d.$$

Performing the Legendre transform and discarding an overall constant gives the following Hamiltonian:

$$H = \frac{\left(Q + C_d V_d\right)^2}{2C_\Sigma} - E_J \cos(\varphi). \tag{2.34}$$

Promoting $\Phi \to \mathbf{\Phi}$, $Q \to \boldsymbol{Q}$ with $[\mathbf{\Phi}, \boldsymbol{Q}] = i\hbar$, and introducing the dimensionless operators

$$\boxed{\boldsymbol{n} \equiv \frac{\boldsymbol{Q}}{2e}, \quad n_d \equiv -\frac{C_d V_d}{2e}, \quad E_C \equiv \frac{e^2}{2C_\Sigma}, \quad [\boldsymbol{\varphi}, \boldsymbol{n}] = i,} \tag{2.35}$$

we obtain the standard CPB Hamiltonian:

$$\boxed{\boldsymbol{H}_{\text{CPB}} = 4E_C \left(\boldsymbol{n} - n_d\right)^2 - E_J \cos(\boldsymbol{\varphi}).} \tag{2.36}$$

Here $\boldsymbol{n}$ denotes the excess Cooper pairs on the junction electrodes, and n_d is the gate-induced offset charge. In the charge-qubit regime $E_J/E_C \ll 1$, the spectrum depends strongly on n_d; thus fluctuations of n_d translate directly into qubit-frequency noise and rapid dephasing (for example, pure-dephasing times T_2^* are often limited to tens of nanoseconds [NPT99]). Although one can operate at the charge "sweet spot" $n_d = 0.5$, where the qubit frequency is first-order insensitive to n_d $\left(\partial\omega_q/\partial n_d = 0\right)$, maintaining this bias is difficult. Higher-order effects of the $1/f$ noise or quasiparticle poisoning further destabilize the operating point[Ith+05; Koc+07]. For instance, quasiparticles—essentially unpaired electrons in the superconductors—tunnel stochastically across the junctions, switching the island parity and shifting the offset charge by ±0.5[Fer+06; Ris+13]. This acute charge sensitivity motivates the use of the low-impedance, or heavy, version of this qubit, in which $E_J/E_C \gg 1$ [Koc+07], greatly reducing the dependence of the spectrum on n_d and thereby enhancing coherence. This heavy version of the Cooper-pair box is the celebrated transmon qubit.

[NPT99]: Nakamura et al. (1999), 'Coherent control of macroscopic quantum states in a single-Cooper-pair box', Nature

[Ith+05]: Ithier et al. (2005), 'Decoherence in a superconducting quantum bit circuit', Physical Review B

[Koc+07]: Koch et al. (2007), 'Charge insensitive qubit design derived from the Cooper pair box', Physical Review A

[Fer+06]: Ferguson et al. (2006), 'Microsecond Resolution of Quasiparticle Tunneling in the Single-Cooper-Pair Transistor', Physical Review Letters

[Ris+13]: Ristè et al. (2013), 'Millisecond charge-parity fluctuations and induced decoherence in a superconducting transmon qubit', Nature Communications

2.3.2 Transmon Hamiltonian

In the transmon regime [Koc+07], where $E_J \gg E_C$, charge dispersion is exponentially suppressed, making the energy levels essentially independent of n_d and allowing us to neglect any static offset charge. We set $n_d = 0$ without loss of generality, leading to the transmon Hamiltonian

$$\boldsymbol{H}_{\text{transmon}} = 4E_C \boldsymbol{n}^2 - E_J \cos(\boldsymbol{\varphi}). \tag{2.37}$$

In the limit $E_J/E_C \gg 1$, the Josephson potential energy dominates the charging (kinetic) energy, localizing the phase (position-like) coordinate φ near zero with small fluctuations. This permits us to expand the potential around $\varphi = 0$. We introduce the canonical ladder operators $\boldsymbol{c}$

and $\boldsymbol{c}^\dagger$ via

$$\boxed{\boldsymbol{\varphi} = \varphi_{\mathrm{zpf}}\,(\boldsymbol{c} + \boldsymbol{c}^\dagger), \qquad \boldsymbol{n} = -\,i\,N_{\mathrm{zpf}}\,(\boldsymbol{c} - \boldsymbol{c}^\dagger), \quad [\boldsymbol{c}, \boldsymbol{c}^\dagger] = 1.} \tag{2.38}$$

Here, zero-point amplitudes

$$\varphi_{\mathrm{zpf}} = \left(\frac{2E_C}{E_J}\right)^{1/4}, \quad N_{\mathrm{zpf}} = \frac{1}{2}\left(\frac{E_J}{2E_C}\right)^{1/4}$$

are chosen so that the quadratic part of Eq. 2.37 reproduces a harmonic oscillator of frequency $\omega_q = \sqrt{8E_C E_J}/\hbar$, and the commutation relation $[\boldsymbol{\varphi}, \boldsymbol{n}] = i$ holds[1]. Substituting (2.38) into (2.37) and expanding the cosine to fourth order,

$$\begin{aligned} -E_J \cos(\boldsymbol{\varphi}) &\approx -E_J\left(1 - \tfrac{1}{2}\boldsymbol{\varphi}^2 + \tfrac{1}{24}\boldsymbol{\varphi}^4\right) \\ &= -E_J + \tfrac{1}{2}E_J\varphi_{\mathrm{zpf}}^2(\boldsymbol{c} + \boldsymbol{c}^\dagger)^2 - \tfrac{E_J\varphi_{\mathrm{zpf}}^4}{24}(\boldsymbol{c} + \boldsymbol{c}^\dagger)^4. \end{aligned}$$

1: Expanding the commutator leads to $[\varphi_{\mathrm{zpf}}(\boldsymbol{c}+\boldsymbol{c}^\dagger), -iN_{\mathrm{zpf}}(\boldsymbol{c}-\boldsymbol{c}^\dagger)] = 2i\,\varphi_{\mathrm{zpf}}N_{\mathrm{zpf}}$, so requires $\varphi_{\mathrm{zpf}}N_{\mathrm{zpf}} = \frac{1}{2}$.

Discarding the constant terms and adding the kinetic term $4E_C\boldsymbol{n}^2$ gives

$$\boldsymbol{H} = \hbar\omega_q \boldsymbol{c}^\dagger \boldsymbol{c} - \frac{E_C}{12}\left(\boldsymbol{c} + \boldsymbol{c}^\dagger\right)^4 + \cdots$$

where omitted terms are higher-order in $\varphi_{\mathrm{zpf}}^2 \sim (E_C/E_J)^{1/2}$. Applying the rotating-wave approximation (keeping only number-conserving quartic terms) yields the transmon Hamiltonian in terms of ladder operators:

$$\boxed{\boldsymbol{H}_{\mathrm{transmon}} = \underbrace{\hbar\omega_q\,\boldsymbol{c}^\dagger\boldsymbol{c}}_{\substack{\text{harmonic}\\ \text{oscillator}}} - \underbrace{\frac{E_C}{2}\,\boldsymbol{c}^{\dagger 2}\boldsymbol{c}^2}_{\substack{\text{Kerr}\\ \text{nonlinearity}}}\,.} \tag{2.39}$$

The last term represents the Kerr nonlinearity, by analogy with the optical Kerr effect discovered by John Kerr. In nonlinear optics, the Kerr effect manifests as an intensity-dependent refractive index; here, the quartic term corresponds to an effective photon–photon interaction that gradually reduces the spacing between adjacent energy levels as the excitation number increases..

2.3.3 Energy Levels and Anharmonicity of the Transmon

The quartic term $\boldsymbol{c}^{\dagger 2}\boldsymbol{c}^2$ introduces a weak anharmonicity in the spectrum. Starting from the transmon Hamiltonian in (2.39), the eigenenergies associated with the Fock basis $\{|n\rangle\}$ are[2]

$$E_n \;=\; \langle n|\boldsymbol{H}_{\mathrm{transmon}}|n\rangle = \hbar\omega_q\, n - \frac{E_C}{2}\, n(n-1), \quad n \in \mathbb{N}. \tag{2.40}$$

2: Note that, by omitting constant terms, the origin of our energy scale is shifted: the energy zero no longer coincides with the potential minimum, and the typical ground-state energy $\hbar\omega/2$ is therefore excluded.

Hence the transition frequency between levels $|n\rangle \to |n+1\rangle$ is

$$\omega_{n+1,n} = \frac{E_{n+1} - E_n}{\hbar} = \omega_q - \frac{E_C}{\hbar}\, n.$$

In particular, $\omega_{10} = \omega_q$, $\omega_{21} = \omega_q - \frac{E_C}{\hbar}$. The anharmonicity is defined as

$$\alpha = \omega_{21} - \omega_{10} = -\frac{E_C}{\hbar} < 0, \tag{2.41}$$

quantifying the non-uniform level spacing that enables selective qubit control (see Fig. 2.7).

We can intuitively understand the behavior of the system in this potential by recalling an earlier point: in LC circuits, the capacitance C is analogous to the mass of a particle confined in a potential well (Sec. 2.1.3). Similarly, here the charging energy term $4E_C\boldsymbol{n}^2$ plays the role of kinetic energy in the phase basis. Reducing E_C increases the effective mass, thereby confining the modes more tightly within the potential well.

To make this intuition quantitative, we estimate how the number of bound states scales with the circuit parameters. Ignoring the anharmonicity for a moment, we estimate the energy level spacing as $\omega_q = \sqrt{8E_C E_J}/\hbar$. Then, the number of energy levels localized in the central well can be roughly estimated by dividing the well depth ($\sim 2E_J$) by this characteristic spacing, yielding the scaling

$$N_{\text{levels}} \propto \sqrt{E_J/E_C}.$$

This implies that the number of bound states increases either by deepening the well (increasing E_J) or by increasing the effective mass (decreasing E_C). In Fig. 2.7, we illustrate this effect by varying the ratio E_J/E_C while keeping E_J constant, showing the confinement of wavefunctions in the potential well.

In typical transmons with $E_J/E_C \sim 50$ and $E_J \sim 10$ GHz, one finds

$$\omega_q/2\pi \sim 3.8 \text{ GHz}, \quad |\alpha|/2\pi \sim 230 \text{ MHz}, \quad \frac{|\alpha|}{\omega_q} \sim 6\%.$$

The transmon qubit, by operating in the regime $E_J/E_C \gg 1$, achieves a remarkable suppression of charge dispersion and correspondingly long coherence times [Koc+07; Sch+08]. However, its small anharmonicity makes it practically impossible to lower the transition frequency into the MHz range while preserving the ability to control and read out the qubit state using standard circuit QED techniques. Despite the fabrication challenges, at a transition frequency of 4 MHz with 6% anharmonicity ($|\alpha|/2\pi \approx 240$ kHz), qubit control pulses must have durations of at least 10 μs to avoid leakage into higher others levels. Also, readout via a nearly resonant $\lambda/4$ coplanar-waveguide resonator at 4 MHz entails a transmission line of about $\sim$ 15 m , which is impractically long for standard circuit QED setups.

Before proceeding, our aim is to understand the basic principles of coherent control of a harmonic resonator. In the following, we review the two-level formalism and the standard circuit QED (cQED) framework—treating qubits as two-level systems coupled to bosonic resonators—which provides the essential language and tools for coherent quantum control throughout this work.

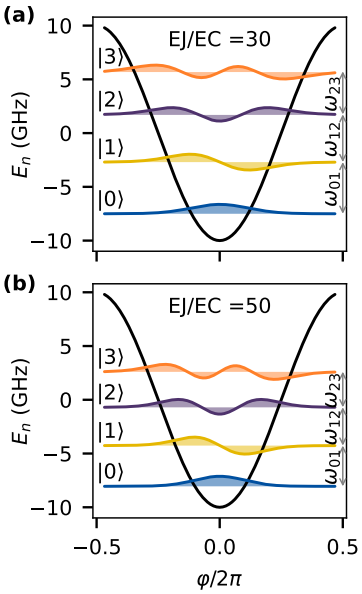

Fig. 2.7: Transmon spectrum. The first four energy levels of the transmon qubit are shown with their corresponding wavefunctions. The energy levels are labeled by the Fock states $|n\rangle$, and the transition frequencies $\omega_{n,n-1}$ between adjacent levels are indicated. In both subplots, $E_J/h = 10$ GHz is held constant.

[Koc+07]: Koch et al. (2007), 'Charge insensitive qubit design derived from the Cooper pair box', Physical Review A
[Sch+08]: Schreier et al. (2008), 'Suppressing charge noise decoherence in superconducting charge qubits', Physical Review B

2.4 Two-Level System and basic manipulations

As we have just seen in the transmon qubit example, incorporating a Josephson junction into a superconducting circuit produces an intrinsically anharmonic energy spectrum. This anharmonicity is the key enabler for defining an effective two-level system (TLS): by choosing two eigenstates whose transition frequency is well separated from all others, one can isolate those levels as a qubit basis. Consequently, a selected $|i\rangle \leftrightarrow |j\rangle$ transition can be driven with high fidelity, while all other levels stay far off-resonant and decoupled from qubit operations. In this section, we develop the minimal theoretical framework for such an anharmonically isolated two-level system and introduce the Bloch-sphere representation as our primary tool for visualizing and manipulating the qubit state.

2.4.1 Pure State

An arbitrary pure state of a two-level system can be expressed in the computational basis $\{|0\rangle, |1\rangle\}$ as

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle, \quad \alpha, \beta \in \mathbb{C}, \quad |\alpha|^2 + |\beta|^2 = 1. \tag{2.42}$$

Here, α and β are complex probability amplitudes, with $|\alpha|^2$ and $|\beta|^2$ giving the probabilities of finding the system in $|0\rangle$ or $|1\rangle$, respectively. It can be shown that any such state can be parametrized using spherical coordinates as

$$\begin{gathered} |\psi\rangle = \cos\!\left(\frac{\theta}{2}\right) e^{+i\phi/2}\,|0\rangle + \sin\!\left(\frac{\theta}{2}\right) e^{-i\phi/2}\,|1\rangle, \\ 0 \le \theta \le \pi, \quad 0 \le \phi < 2\pi. \end{gathered} \tag{2.43}$$

which represents a point on a unitary sphere, called the Bloch sphere.

Projecting the full Hamiltonian $\boldsymbol{H}$ onto the two-level subspace spanned by $\{|0\rangle, |1\rangle\}$ defines the qubit transition frequency ω_q (cf. Eq. 2.44):

$$\langle 0|\,\boldsymbol{H}\,|0\rangle = 0, \quad \langle 1|\,\boldsymbol{H}\,|1\rangle = \hbar\omega_q, \quad \langle 0|\,\boldsymbol{H}\,|1\rangle = \langle 1|\,\boldsymbol{H}\,|0\rangle = 0. \tag{2.44}$$

In matrix form, adopting the convention $|0\rangle = (1,0)^T$ and $|1\rangle = (0,1)^T$, the truncated two-level subspace inherits an SU(2) structure, so that any operator on this space can be decomposed in the Pauli basis: [1]

$$\boxed{\begin{gathered} \boldsymbol{\sigma}_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \quad \boldsymbol{\sigma}_y = \begin{pmatrix} 0 & i \\ -i & 0 \end{pmatrix}, \quad \boldsymbol{\sigma}_z = \begin{pmatrix} -1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \boldsymbol{I} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \\ \boldsymbol{\sigma}_+ = \frac{1}{2}\left(\boldsymbol{\sigma}_x + i\boldsymbol{\sigma}_y\right) = \begin{pmatrix} 0 & 0 \\ 1 & 0 \end{pmatrix}, \quad \boldsymbol{\sigma}_- = \frac{1}{2}\left(\boldsymbol{\sigma}_x - i\boldsymbol{\sigma}_y\right) = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} \end{gathered}} \tag{2.45}$$

The operators σ_x and σ_y generate bit flips about the x- and y-axes of the Bloch sphere, respectively, while σ_z induces a phase flip. The ladder operators $\boldsymbol{\sigma}_+$ and $\boldsymbol{\sigma}_-$ serve as the analogues of the creation and annihilation operators, respectively: $\boldsymbol{\sigma}_+ \leftrightarrow \boldsymbol{a}^\dagger$ and $\boldsymbol{\sigma}_- \leftrightarrow \boldsymbol{a}$. Their action

1: For an arbitrary qubit state $\alpha|0\rangle + \beta|1\rangle$, the Pauli operators act as follows:

- σ_x: Bit-flip - rotation about the x–axis on the Bloch sphere.
$$\sigma_x\big(\alpha|0\rangle + \beta|1\rangle\big) = \alpha|1\rangle + \beta|0\rangle$$
- σ_y: Bit-flip plus phase — rotation about the y–axis on the Bloch sphere.
$$\sigma_y\big(\alpha|0\rangle + \beta|1\rangle\big) = i\beta|0\rangle - i\alpha|1\rangle$$
- σ_z: Phase-flip — leaves $|0\rangle$ unchanged, flips the sign of $|1\rangle$.
$$\sigma_z\big(\alpha|0\rangle + \beta|1\rangle\big) = -\alpha|0\rangle + \beta|1\rangle$$
- σ_+: Raising/excitation — promotes $|0\rangle \to |1\rangle$.
$$\sigma_+\big(\alpha|0\rangle + \beta|1\rangle\big) = \alpha|1\rangle$$
- σ_-: Lowering/decay — demotes $|1\rangle \to |0\rangle$.
$$\sigma_-\big(\alpha|0\rangle + \beta|1\rangle\big) = \beta|0\rangle$$

on the computational basis is simply

$$\sigma_+ |0\rangle = |1\rangle , \qquad \sigma_- |1\rangle = |0\rangle . \tag{2.46}$$

Finally, projecting $\boldsymbol{H}$ into this truncated basis and discarding an inconsequential constant offset yields the compact Hamiltonian of the two-level system:

$$\boxed{\boldsymbol{H} = \frac{\hbar\omega_q}{2}\,\boldsymbol{\sigma}_z.} \tag{2.47}$$

This concise Pauli-matrix form captures the essential dynamics of the qubit whenever higher excited levels remain far off-resonant, unifying the description across different physical implementations. Its validity, however, rests on truncating the full Hilbert space to the selected pair of levels and ensuring that any intrinsic or drive-induced coupling between selected levels and all other modes remains much weaker than the corresponding energy detunings, thereby suppressing leakage and preserving faithful two-level behavior.

2.4.2 Mixed State and Density Matrix

In realistic settings, a two-level system may not occupy a single pure state but rather a statistical ensemble of states. Such a mixture cannot be described by a single wavefunction. Instead, one uses the density operator

$$\boldsymbol{\rho} = \sum_i p_i \, |\psi_i\rangle \langle\psi_i| , \quad p_i \geq 0, \quad \sum_i p_i = 1, \tag{2.48}$$

where each $|\psi_i\rangle$ is a pure state occurring with classical probability p_i. For instance, an equal-probability mixture of $|0\rangle$ and $|1\rangle$ differs fundamentally from the coherent superposition $\frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.

The density operator evolves under the system Hamiltonian $\boldsymbol{H}$ according to the von Neumann equation,

$$\boxed{\frac{d\boldsymbol{\rho}}{dt} = -\frac{i}{\hbar}\,[\,\boldsymbol{H}, \boldsymbol{\rho}\,],} \tag{2.49}$$

which generalizes Schrödinger's equation to statistical mixtures, preserving both trace and positivity of $\boldsymbol{\rho}$.

In the $\{|0\rangle , |1\rangle\}$ basis, the density matrix components are

$$\rho = \begin{pmatrix} \rho_{00} & \rho_{01} \\ \rho_{01}^* & \rho_{11} \end{pmatrix}, \quad \rho_{00} + \rho_{11} = 1, \tag{2.50}$$

where ρ_{00} and ρ_{11} denote the populations (probabilities of $|0\rangle$ and $|1\rangle$), and ρ_{01} encodes the coherence between the two levels.

Equivalently, one may expand $\boldsymbol{\rho}$ in the Pauli basis,

$$\boldsymbol{\rho} = \frac{1}{2}\big(\boldsymbol{I} + r_x\,\boldsymbol{\sigma}_x + r_y\,\boldsymbol{\sigma}_y + r_z\,\boldsymbol{\sigma}_z\big), \tag{2.51}$$

introducing the real Bloch-vector components (r_x, r_y, r_z) satisfying $r_x^2 + r_y^2 + r_z^2 \leq 1$. The bijective mapping between matrix elements and Bloch

coordinates is

$$r_z = 2\,\rho_{00} - 1, \quad r_x = \rho_{01} + \rho_{01}^*, \quad r_y = i\,(\rho_{01}^* - \rho_{01}), \tag{2.52}$$

with the inverse relations

$$\rho_{00} = \tfrac{1}{2}(1 + r_z), \quad \rho_{01} = \tfrac{1}{2}(r_x - i\,r_y), \quad \rho_{11} = \tfrac{1}{2}(1 - r_z). \tag{2.53}$$

This Bloch representation unifies populations and coherences into a single geometric picture, facilitating intuitive analysis of mixed-state dynamics.

2.4.3 Bloch-Sphere Representation

Previously, we have seen that a pure state (Eq. 2.43) can be unambiguously represented by a spherical coordinates (θ, ϕ). This can be extended to any mixed state. The density matrix $\boldsymbol{\rho}$ can be written in the Pauli basis with coordinates $(r_x, r_y, r_z) \in \{\mathbf{r} \in \mathbb{R}^3 : \|\mathbf{r}\| \leq 1\}$. This space corresponds to all points inside a unit sphere in $\mathbb{R}^3$, and thus any density matrix corresponds to a point within this sphere. We represent the state as the Bloch vector $\vec{r}$ originating from the center of the unit sphere, called the Bloch sphere. The vector's norm

$$r = \sqrt{r_x^2 + r_y^2 + r_z^2} \leq 1$$

quantifies the state's purity: $r = 1$ corresponds to a pure state whose Bloch vector terminates on the sphere's surface, whereas $r < 1$ indicates a mixed state residing in the interior.

The polar angle θ and azimuthal angle ϕ of the Bloch vector are given by classical spherical coordinates:

$$\theta = \arccos\big(r_z/r\big), \quad \phi = \arctan(r_y/r_x). \tag{2.54}$$

This provides an exact transformation between the spherical coordinates (r, θ, ϕ) and the density matrix $\boldsymbol{\rho}$ representation in the Pauli basis. Fig. 2.8 illustrates Bloch vectors for the ground state (vector pointing to the -z axis), the excited state (vector pointing to the +z axis), and an equatorial superposition (vector in the x–y plane).

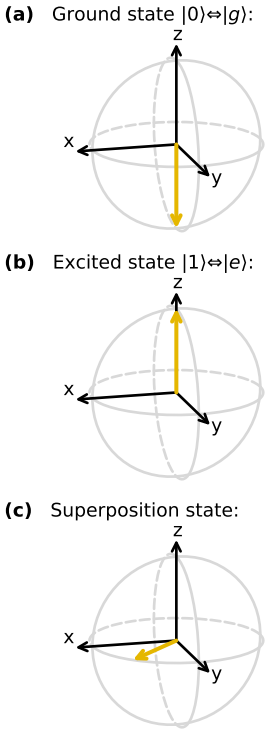

Fig. 2.8: Bloch-sphere visualization of key states. From top to bottom: the ground state $|0\rangle$, the excited state $|1\rangle$, and the equal-superposition state $|+\rangle$.

2.4.4 Qubit control and measurement

A qubit drive can be implemented by coupling the qubit to a sinusoidal, time-dependent field. By moving to a reference frame rotating at the drive frequency and applying the rotating-wave approximation (RWA) to discard rapidly oscillating terms, the drive Hamiltonian reduces to

$$\boldsymbol{H}_{\mathrm{D}} = \frac{\hbar\,\delta\omega}{2}\,\boldsymbol{\sigma}_z \;+\; \frac{\hbar\,\Omega}{2}\,\big(\boldsymbol{\sigma}_x \cos\phi + \boldsymbol{\sigma}_y \sin\phi\big), \tag{2.55}$$

where Ω is the on-resonance Rabi frequency (set by the drive amplitude and coupling), ϕ is the drive phase, and $\delta\omega$ is the detuning between the qubit transition frequency and the drive frequency.

In Sec. 3.2.1 we will derive the drive Hamiltonian for the specific case of our circuit. For now, it is crucial to develop the Bloch-sphere picture of Rabi oscillations and to see that

$$\frac{\hbar\Omega}{2}\left(\sigma_x \cos\phi + \sigma_y \sin\phi\right) \tag{2.56}$$

represents a rotation about an arbitrary axis defined by the phase ϕ and the Rabi frequency Ω.

Evolution operator. One may rewrite the drive Hamiltonian $\boldsymbol{H}_\mathrm{D}$ (Eq. 2.55) in vector form,

$$\boldsymbol{H}_\mathrm{D} = \frac{\hbar\Omega_R}{2}\left(\vec{n}\cdot\vec{\sigma}\right), \quad \Omega_R = \sqrt{\Omega^2+\delta\omega^2},$$
$$\vec{n} = \frac{1}{\Omega_R}(\Omega\cos\phi, \Omega\sin\phi, \delta\omega), \quad \vec{\sigma} = (\sigma_x, \sigma_y, \sigma_z)$$

Here, $\vec{n}$ is the unit vector along the effective drive field in Bloch-sphere space, $\vec{\sigma}$ is the Pauli-vector operator, and Ω_R is the detuned Rabi frequency. The time evolution of an arbitrary initial state $|\psi(0)\rangle$ is then described with the evolution operator[1] $\boldsymbol{U}_t$,

$$\begin{aligned} |\psi(t)\rangle &= \boldsymbol{U}(t)\,|\psi(0)\rangle\,, \\ \boldsymbol{U}(t) &= \exp\left(-i\boldsymbol{H}_\mathrm{D}t/\hbar\right) = \exp\left(-i\frac{\Omega_R}{2}\vec{n}\cdot\vec{\sigma}t\right) \\ &= \cos\Theta\mathbb{1} - i\sin\Theta\,(\vec{n}\cdot\vec{\sigma}), \end{aligned} \tag{2.57}$$

with rotation angle $\Theta = \Omega_R\, t/2$. Thus, the evolution operator describes a rotation of the qubit state on the Bloch sphere around the axis defined by $\vec{n}$.

π– and $\pi/2$–pulses. Consider an example when a qubit is initially prepared in the ground state, $|\psi_0\rangle = |0\rangle$. When driven on resonance ($\delta\omega = 0$), the qubit evolves as

$$|\psi\rangle = \boldsymbol{U}(t)\,|0\rangle = \cos\left(\frac{\Omega}{2}t\right)|0\rangle + i\,e^{i\phi}\,\sin\left(\frac{\Omega}{2}t\right)|1\rangle\,. \tag{2.58}$$

After a pulse of duration $\boxed{\tau_\pi = \pi/\Omega}$, the qubit has evolved exactly from the ground state $|0\rangle$ to the excited state $|1\rangle$. Such pulses are termed π-pulses, as they rotate the qubit's polar angle by π (see Fig. 2.9(a)).

In contrast, a pulse of duration $\tau_{\pi/2} = \pi/(2\Omega)$ yields

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + i\,e^{-i\phi}\,|1\rangle\right) = \begin{cases} |+\rangle_x\,, & \phi = \frac{\pi}{2} + 2\pi n, \\ |+\rangle_y\,, & \phi = 0 + 2\pi n, \end{cases} \quad n\in\mathbb{Z}. \tag{2.59}$$

These $\pi/2$-pulses transfer the qubit state between the Bloch-sphere poles and the equator. The final position in the equatorial plane defined by the phase ϕ (see Fig. 2.9(b-c)). Since we routinely employ both π and $\pi/2$ pulses in this study, it is essential to recognize that varying ϕ provides

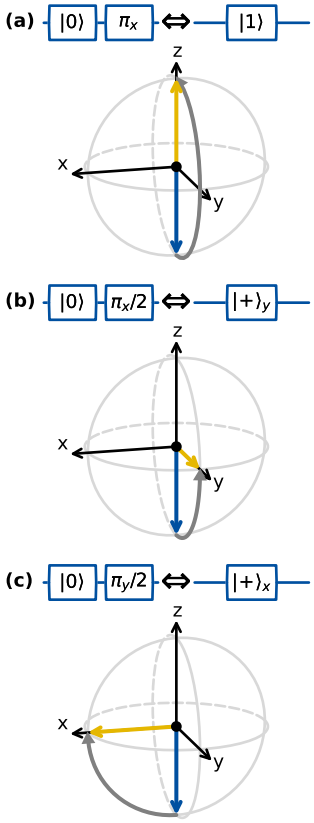

Fig. 2.9: Bloch-sphere illustration of drive pulses. (a) A π-pulse flips the qubit from $|0\rangle$ to $|1\rangle$. **(b)** A $\pi_x/2$-pulse about the x axis prepares $|+\rangle_y$ from $|0\rangle$. **(c)** A $\pi_y/2$-pulse about the y axis prepares $|+\rangle_x$ from $|0\rangle$.

1: One writes $\boldsymbol{U}(t)$ in terms of cosines and sines using the power-series expansion of the exponential, separating even and odd powers and the Pauli-algebra identity $\sigma_i\sigma_j = \delta_{ij}\,\mathbb{1} + i\varepsilon_{ijk}\sigma_k$, which implies $(\boldsymbol{n}\cdot\boldsymbol{\sigma})^2 = \mathbb{1}$.

precise control over the rotation axis and hence the direction of motion on the Bloch sphere.

To summarize, note that $\boldsymbol{\sigma}_x$ generates rotations about the x–axis, and $\boldsymbol{\sigma}_y$ about the y–axis, which illuminates the interpretation of the drive Hamiltonian $\boldsymbol{H}_{\text{Drive}}$ in Eq. 2.55. The operator combination $\boldsymbol{\sigma}_x \cos\phi + \boldsymbol{\sigma}_y \sin\phi$ thus enacts a rotation of the Bloch vector about an axis in the x–y plane at angle ϕ. Under resonant driving, these rotations proceed at the Rabi frequency Ω.

Detuned Rabi oscillations. As the detuning $\delta\omega$ increases, the effective Rabi frequency $\Omega_R = \sqrt{\Omega^2 + \delta\omega^2}$ increases, and the rotation axis tilts out of the equatorial plane of the Bloch sphere. The Bloch vector then follows a smaller circular path, leading to shorter period (faster oscillations). Fig. 2.10 shows this tilted rotation. In practice, we use the dependence of Ω_R on $\delta\omega$ (see Fig. 2.11) to find the qubit's resonance. By sweeping the drive frequency and measuring the oscillation rate, the characteristic chevron-like pattern reveals the zero-detuning point, identifying the qubit transition frequency.

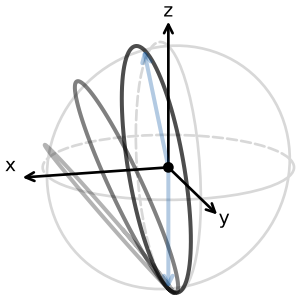

Fig. 2.10: Trajectories on the Bloch sphere under detuned drive. For a drive detuned by $\delta\omega$, the effective rotation axis tilts by $\arctan(\Omega/\delta\omega)$ from the equatorial plane, causing the Bloch vector to precess about this tilted axis. From dark- to light-gray curves, the detuning increases as $\delta\omega = -0.1\,\Omega$, $\delta\omega = -0.4\,\Omega$, and $\delta\omega = -0.8\,\Omega$.

Qubit measurement. The qubit state is read out via a projective measurement onto the computational basis, as detailed later in this work. Measurement inherently collapses the wave function onto $|0\rangle$ or $|1\rangle$, with probabilities given by the squared amplitudes of the state vector. For a general state $|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$, the outcome probabilities are

$$P(0) = |\alpha|^2, \tag{2.60}$$
$$P(1) = |\beta|^2. \tag{2.61}$$

After measurement, the qubit remains in the observed eigenstate (e.g. obtaining $|0\rangle$ leaves the qubit in $|0\rangle$).

Since each measurement yields only a binary outcome, any experiment is repeated more than thousands of times to estimate the excited-state probability $P(|e\rangle)$. Although direct readout is performed in the σ_z basis, we precede the measurement with a $\pi/2$-pulse to rotate the qubit, enabling effective readout in the σ_x or σ_y basis. In the following, we refer to this procedure simply as "readout in the $\sigma_{x/y}$ basis," bearing in mind it consists of a $\pi/2$-pulse followed by a σ_z measurement (see Fig. 2.12).

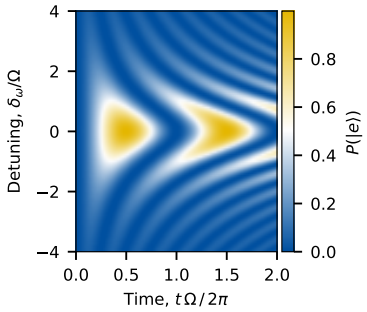

Fig. 2.11: Rabi oscillations (chevron pattern) as a function of detuning. The color scale gives the probability $P(|e\rangle)$ of finding the qubit in the excited state after a pulse of duration t, revealing the characteristic chevron structure determined by Ω and $\delta\omega$.

2.4.5 Decoherence channels of the two-level system. The Lindblad Master Equation

Though the Hamiltonian $\boldsymbol{H}$ describes an ideal, closed system, real devices interact with control lines and spurious environmental modes, leading to dissipation and dephasing. To capture both unitary and irreversible dynamics, we use the Lindblad master equation [HR06]:

$$\boxed{\dot{\rho} = -\frac{i}{\hbar}[\boldsymbol{H}, \rho] + \sum_{i,j}\left(L_{ij}\,\rho\,L_{ij}^\dagger - \tfrac{1}{2}\{L_{ij}^\dagger L_{ij},\,\rho\}\right),} \tag{2.62}$$

where the first term generates coherent evolution under H, and the sum over jump operators L_{ij} encodes distinct loss channels. In a two-level system, since Pauli matrices form a basis, any losses operator L_{ij} can be expressed in terms of the following independent operators [1]:

$$\begin{aligned} \boldsymbol{L}^{\Phi} &= \sqrt{\frac{\gamma_\Phi}{2}}\,\boldsymbol{\sigma}_z, \\ \boldsymbol{L}_- &= \sqrt{\Gamma_\downarrow}\,\boldsymbol{\sigma}_-, \quad \Gamma_\downarrow = \gamma_1(n_{\text{th}}+1), \\ \boldsymbol{L}_+ &= \sqrt{\Gamma_\uparrow}\,\boldsymbol{\sigma}_+, \quad \Gamma_\uparrow = \gamma_1\, n_{\text{th}}, \end{aligned} \tag{2.63}$$

where $\boldsymbol{L}^{\Phi}$ describes pure dephasing of the qubit state, while $\boldsymbol{L}_-$ and $\boldsymbol{L}_+$ describe energy relaxation and excitation, respectively, for a system submitted to a thermal environment of occupation n_{th}. Here, we introduce Γ_Φ as the pure dephasing rate, and γ_1 as the depolarization (energy relaxation) rate, which is the rate at which the qubit relaxes from $|1\rangle$ to $|0\rangle$ and vice versa. The thermal occupancy at frequency ω and temperature T is

$$n_{\text{th}} = \left(e^{\hbar\omega/k_B T} - 1\right)^{-1}. \tag{2.64}$$

Substituting into Eq. 2.62 together with the qubit Hamiltonian Eq. 2.47 yields the Bloch-type equations

$$\begin{aligned} \dot{\rho}_{00} &= \Gamma_\downarrow \rho_{11} - \Gamma_\uparrow \rho_{00} \\ \dot{\rho}_{11} &= \Gamma_\uparrow \rho_{00} - \Gamma_\downarrow \rho_{11} \\ \dot{\rho}_{01} &= -\left(\gamma_\Phi + \tfrac{\Gamma_\uparrow+\Gamma_\downarrow}{2} + i\omega_q\right)\rho_{01} \\ \dot{\rho}_{10} &= -\left(\gamma_\Phi + \tfrac{\Gamma_\uparrow+\Gamma_\downarrow}{2} - i\omega_q\right)\rho_{10} \end{aligned}$$

Solving these differential equations with initial conditions $\rho(0)$ and ensuring $\mathrm{Tr}[\rho] = 1$ gives

$$\begin{aligned} \rho_{00}(t) &= \rho_{00}^{\text{eq}} + \left(\rho_{00}(0) - \rho_{00}^{\text{eq}}\right) e^{-\Gamma_1 t}, \\ \rho_{01}(t) &= \rho_{01}(0) e^{-\Gamma_2 t} e^{-i\omega_0 t}, \end{aligned} \tag{2.65}$$

where we have defined the thermal equilibrium populations

$$\rho_{00}^{\text{eq}} = \frac{n_{\text{th}}+1}{2n_{\text{th}}+1},$$

and the characteristic rates and timescales[2]:

$$\boxed{\Gamma_1 \equiv \frac{1}{T_1} \equiv \gamma_1\,(2n_{\text{th}}+1), \quad \Gamma_2 \equiv \frac{1}{T_2} \equiv \gamma_\Phi + \frac{1}{2T_1}.} \tag{2.66}$$

This relation defines the two principal qubit timescales: the energy-relaxation time T_1 and the total decoherence time T_2, the latter encompassing both energy decay and pure-dephasing contributions.

Physically, T_1 characterizes how quickly the qubit exchanges energy with its environment and returns to thermal equilibrium. On the Bloch sphere, this manifests as the decay of the population difference (the z-axis projection) toward its steady-state value (see Fig. 2.13(b)). The decoherence time T_2 governs the decay of the off-diagonal density-matrix elements and the shrinking of the Bloch vector in the transverse plane.

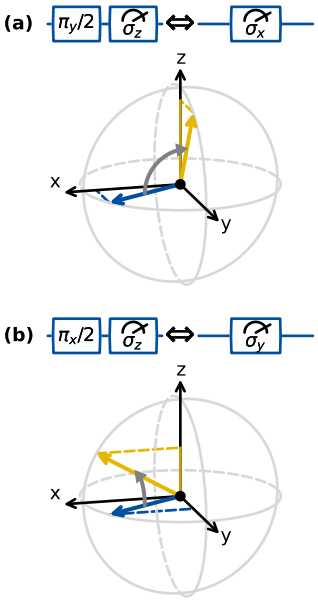

Fig. 2.12: Qubit readout in the σ_x and σ_y bases. Measurement in the σ_x basis is achieved by applying a $\pi/2$ rotation about the y–axis before the standard σ_z readout, mapping the x–axis onto z. Likewise, a $\pi/2$ rotation about the x–axis enables readout in the σ_y basis.

[HR06]: Haroche et al. (2006), *Exploring the Quantum*, Oxford University Press

1: In the full Fock space, each losses operator L_{ij} typically decomposed into pure dephasing operator $\boldsymbol{L}_{ii}^{\Phi}$ and depolarization operator $\boldsymbol{L}_{ij}^{R}$:

$$\begin{aligned} \boldsymbol{L}_{ii}^{\Phi} &= \sqrt{\frac{\gamma_\Phi}{2}}\,|i\rangle\langle i|\,; \\ \boldsymbol{L}_{ij}^{R} &= \sqrt{\tilde{\gamma}_{ij}}\,|j\rangle\langle i|\,, \quad (i \neq j). \end{aligned}$$

Here, $\tilde{\gamma}_{ij}$ is an effective depolarization rate.

2: Each γ is an angular rate (rad/s), so to express rates in hertz one divides by 2π, e.g.

$$\Gamma_1^{(\text{Hz})} = \frac{\Gamma_1^{(\text{rad/s})}}{2\pi} = \frac{1}{2\pi\, T_1}.$$

Both rates grow with the thermal photon number n_{th}, which poses an additional challenge for low-frequency qubits.

In practice, dephasing arises not only from the Markovian pure-dephasing term γ_Φ (with Lindblad operator $\boldsymbol{L}_\Phi$) but also from low-frequency noise sources (instrument drifts, magnetic-field fluctuations, etc.). These extra noise contributions shorten the observed coherence, so the free-induction decay time, including all noise sources, is denoted T_2^*. In the next subsections, we discuss measurement techniques for T_1 (relaxation time measurements), T_2^* (Ramsey experiments), and T_2 (echo experiments) in detail. For now, note that

$$\boxed{T_2^* \leq T_2 \leq 2\,T_1,}$$

so in the ideal limit $T_2^* = T_2 = 2\,T_1$.

2.4.6 Basic characterization sequence: Relaxation and Ramsey oscillations

The coherence times (T_1, T_2^*) and qubit frequency ω_q are essential for understanding the qubit's properties, and in what follows we discuss the experimental techniques used to measure them.

Relaxation. The ability of the qubit to store classical-like states (analogous to 0 and 1 in a classical computer) is characterized by the energy-relaxation time T_1. For a typical transmon qubit, $T_1 \sim 100\,\mu$s [Rig+12; Pla+21; Mam+21], and with advanced techniques this can be extended to $T_1 \sim 1$ ms [Bla+25]. For the MHz-frequency fluxonium qubit used in this work, typical values are $T_1 \sim 35\,\mu$s [Naj+24] for our devices, and up to $T_1 \sim 200\,\mu$s for the state-of-the-art heavy fluxonium qubit [Zha+21].

To measure T_1, one prepares the qubit in the excited state $|1\rangle$, allows it to decohere freely under the influence of its environment, and measures the population in $|1\rangle$ over time (i.e., the projection onto the z-axis of the Bloch sphere). This sequence is illustrated in Fig. 2.13, where the Bloch vector at specific times is shown on the sphere. From Eq. 2.65, the excited-state population decays exponentially toward the thermal equilibrium, so fitting the measured population $P(t)$ yields the value of T_1.

Ramsey oscillations. The ability of the qubit to maintain quantum coherence (superposition of $|0\rangle$ and $|1\rangle$) is characterized by the coherence time T_2^*. To measure the total decoherence time, one employs the Ramsey sequence:

1. **State preparation.** Start in $|0\rangle$ and apply a $\pi/2$ pulse about the y-axis to prepare
$$|+\rangle = \frac{1}{\sqrt{2}}\big(|0\rangle + |1\rangle\big).$$

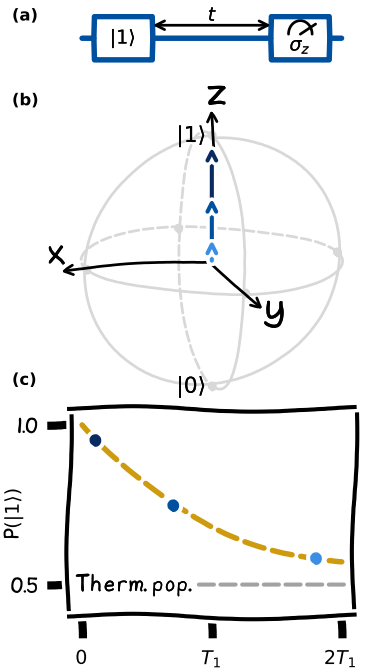

Fig. 2.13: Illustration of the qubit relaxation sequence. (a) Pulse schematic: prepare in $|1\rangle$, wait time t, then measure; **(b)** Illustration of the qubit at different times on the Bloch-sphere; **(c)** excited-state population vs. t with exponential T_1 fit. Points of different shades of blue correspond to the vectors on the Bloch sphere.

[Rig+12]: Rigetti et al. (2012), 'Superconducting qubit in waveguide cavity with coherence time approaching 0.1ms', Physical Review B

[Pla+21]: Place et al. (2021), 'New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds', Nature Communications

[Mam+21]: Mamin et al. (2021), 'Merged-Element Transmons: Design and Qubit Performance', Physical Review Applied

[Bla+25]: Bland et al. (2025), '2D transmons with lifetimes and coherence times exceeding 1 millisecond', Preprint

[Naj+24]: Najera-Santos et al. (2024), 'High-Sensitivity ac-Charge Detection with a MHz-Frequency Fluxonium Qubit', Physical Review X

[Zha+21]: Zhang et al. (2021), 'Universal fast flux control of a coherent, low-frequency qubit', Physical Review X

2. **Free evolution.** Let the system evolve for a time t. In the rotating frame of the control pulses, the state becomes

$$\rho(t) = \frac{1}{2}\begin{pmatrix} 1 & e^{-t/T_2^* + i\Delta t} \\ e^{-t/T_2^* - i\Delta t} & 1 \end{pmatrix},$$

where Δ is the detuning between the qubit's natural frequency and the rotating frame, and the factor e^{-t/T_2^*} captures coherence decay.

3. **Readout σ_x.** Apply a second $\pi/2$ pulse with phase zero (to measure σ_x) and then measure σ_z, yielding following expectation value

$$P(|+\rangle_x) = \exp\left(-t/T_2^*\right)\frac{1+\cos(\Delta t)}{2}.$$

4. **Readout σ_y.** To lift the sign ambiguity of Δ, we alternate the phase of the final $\pi/2$ pulse, thereby measuring both σ_x and σ_y. For the σ_y measurement,

$$P(|+\rangle_y) = \exp\left(-t/T_2^*\right)\frac{1+\sin(\Delta t)}{2}.$$

A combined fit of $P(|+\rangle_x)$ and $P(|+\rangle_y)$ then yields both Δ and T_2^*. In practice, one chooses Δ of order a few kilohertz to observe several oscillations before the envelope decays.

Precise frequency measurement via Ramsey oscillations is essential not only for qubit calibration. As we will see in the following section, the qubit can couple to another system, such as a mechanical mode. This coupling shifts qubit's transition frequency. Probing this shift with the Ramsey protocol enables direct characterization of the coupled system

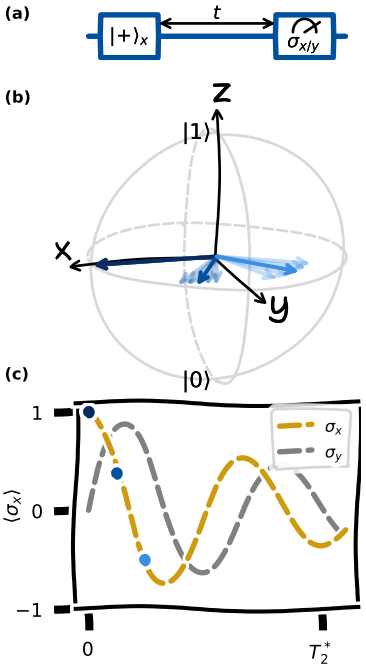

Fig. 2.14: Illustration of the Ramsey oscillations. (a) Pulse sequence: prepare the qubit in $|+\rangle$, wait for a variable delay t, then measure the projection onto σ_x or σ_y. **(b)** Bloch-sphere representation of the qubit state at different times. The spread of the Bloch vectors illustrates dephasing caused by slow fluctuations in the qubit transition frequency. **(c)** Measured $\langle\sigma_x\rangle$ as a function of t, showing an oscillating signal with exponential decay. The different shades of blue correspond to the Bloch vectors shown in panel (b).

Echo sequence. Low-frequency (quasi-static) fluctuations in the qubit transition frequency $\delta\omega(t)$ cause dephasing and limit the Ramsey coherence time T_2^*. To isolate these slow noise contributions and probe only faster decoherence processes, we use the Hahn echo sequence, as illustrated in Fig. 2.15:

1. Prepare the qubit in $|+\rangle_x = (|0\rangle + |1\rangle)/\sqrt{2}$ with a $\pi/2$ pulse about the y-axis.
2. Let the system evolve freely for a time $t/2$, during which a near-constant detuning $\delta\omega$ accumulates a phase $\phi = \delta\omega\, t/2$.
3. Apply a refocusing π pulse about the x-axis, which flips the Bloch vector across the equator and reverses the sense of any phase accrued from $\delta\omega$.
4. Let the system evolve for another $t/2$; the quasi-static phase contributions cancel, refocusing the Bloch vector.
5. Conclude with a second $\pi/2$ pulse about the y-axis and measure $\langle\sigma_x\rangle$ or with a $\pi/2$ pulse about the x-axis and measure $\langle\sigma_y\rangle$.

Because truly static or slowly varying detunings are time-reversed by the π pulse, they do not dephase the echo signal. The resulting echo envelope decays with a characteristic time T_2, which reflects only higher-frequency noise processes. In the following, we emphasize that this decoherence time was measured using the echo sequence and denote it T_2^{echo}. In an ideal system, when low-frequency noise is negligible, one finds $T_2^{\text{echo}} \approx T_2^*$.

Any deviation from this equality thus provides a quantitative metric for the contribution of low-frequency noise to qubit dephasing.

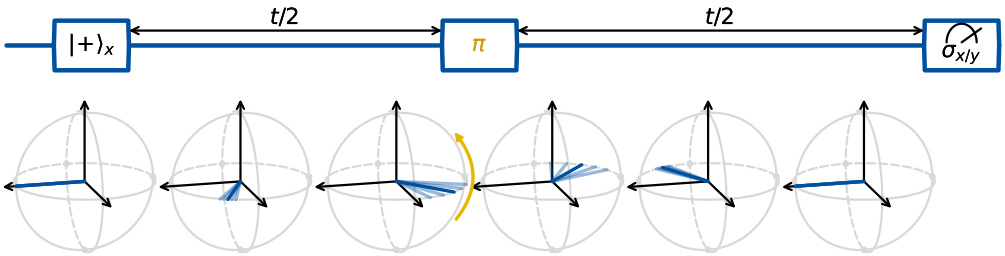

Fig. 2.15: Refocusing of the Bloch vector using the echo sequence. Slowly varying fluctuations in the qubit transition frequency can be mitigated by an echo sequence. A π pulse, applied halfway through the free evolution, reverses the accumulated phase error, effectively refocusing the Bloch vector.

2.5 The framework of circuit QED

Circuit QED exploits the coupling between a nonlinear superconducting qubit ("artificial atom"[1]) and a quantized resonator mode. In this section, we first derive the Jaynes–Cummings Hamiltonian to describe coherent excitation exchange and show how resonant Rabi oscillations enable high-fidelity control of the resonator via the qubit ancilla. We then turn to the dispersive regime of large detuning, where state-dependent frequency shifts both characterize the mechanical mode and enable quantum-non-demolition qubit readout.

1: By analogy to traditional QED—where natural atoms interact with optical cavities—the superconducting qubit plays the role of an engineered "artificial atom," while any harmonic oscillator (microwave or mechanical) serves as the cavity mode.

2.5.1 Jaynes–Cummings Hamiltonian

In circuit QED, the coherent interaction between a superconducting qubit and a single mode of a microwave resonator originates from the dipole coupling of the qubit's charge (or flux) degree of freedom to the quantized field. As will be shown in Chap. 5, starting from the circuit Lagrangian one obtains an interaction term proportional to $(\boldsymbol{a}+\boldsymbol{a}^\dagger)\,\sigma_x$ when the resonator's harmonic oscillator (described by $\boldsymbol{a}, \boldsymbol{a}^\dagger$) is capacitively coupled to the qubit. Transforming to the interaction picture and applying the rotating-wave approximation (valid for $g \ll \omega_q, \omega_r$) removes the counter-rotating terms, yielding the Jaynes–Cummings Hamiltonian[2]:

$$\boxed{\boldsymbol{H}_{\mathrm{JC}} = \hbar\omega_r\, \boldsymbol{a}^\dagger \boldsymbol{a} + \frac{\hbar\omega_q}{2}\,\sigma_z + i\frac{\hbar\Omega}{2}\big(\boldsymbol{a}^\dagger\sigma_- - \boldsymbol{a}\,\sigma_+\big).} \tag{2.67}$$

This Hamiltonian is the cornerstone of circuit QED and describes any harmonic resonator, both microwave and mechanical, coupled to a qubit. In the rotating-wave approximation (RWA), the total excitation number $\boldsymbol{N} = \boldsymbol{a}^\dagger\boldsymbol{a} + (\sigma_z + 1)/2$ is conserved, so the Hamiltonian $\boldsymbol{H}_{\mathrm{JC}}$ can be restricted to a subspace with a fixed excitation number N for coherent dynamics. Consider the subspaces spanned by $|g, n+1\rangle$ (qubit in the ground state and $n+1$ coherent photons in the resonator) and $|e, n\rangle$

2: Here we have the freedom to choose the phase of the off-diagonal terms; we fix it so to simplify the phase of the final qubit state evolution.

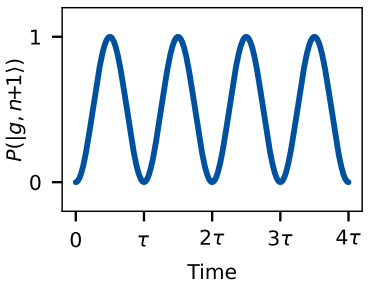

Fig. 2.16: Coherent Rabi oscillations. The population of $|g, n+1\rangle$ oscillates coherently with that of $|e, n\rangle$ with a period $\tau = 2\pi/(\Omega\sqrt{n+1})$ (Eq. 2.68).

(qubit excited and n resonator photons). Restricted to this subspace,

$$\boldsymbol{H}_{\mathrm{JC}}^{(1)} = \hbar\omega_q \left(n + \frac{1}{2}\right)\mathbb{I} + \frac{\hbar\Omega}{2}\sqrt{n+1}\begin{pmatrix} 0 & i \\ -i & 0 \end{pmatrix}.$$

so an initial $|e, n\rangle$ state, in the $\omega_q\sqrt{n+1/2}$ rotating frame, evolves as[1]

$$|\psi(t)\rangle = \cos\!\left(\frac{\Omega\sqrt{n+1}\,t}{2}\right)|e, n\rangle + \sin\!\left(\frac{\Omega\sqrt{n+1}\,t}{2}\right)|g, n+1\rangle\,, \quad (2.68)$$

exhibiting coherent swaps at frequency $\Omega\sqrt{n+1}$. This equation illustrates the system's energy-conserving behavior: each qubit de-excitation $|e\rangle \to |g\rangle$ is accompanied by creation of exactly one photon in the resonator $|0\rangle \to |1\rangle$.

Given the mature toolbox for qubit control and readout, resonant exchange permits precise manipulation of the resonator state via the qubit ancilla. High-fidelity swaps between the qubit and the resonator require the coherent coupling rate g to be much larger than both the qubit decay rate κ_q and the resonator decay rate κ_r, ensuring that coherent excitation exchange dominates over dissipative processes. In this strong-coupling regime, deterministic state-synthesis protocols become possible. Hofheinz et al. [Hof+09] demonstrated preparation of arbitrary Fock-state superpositions in a superconducting microwave resonator by interleaving qubit rotations and on-resonance swap pulses according to the Law–Eberly algorithm [LE96]. Similarly, Chu et al. [Chu+18] show multi-phonon ladder-climbing in a bulk acoustic-wave resonator, demonstrating coherent, resonant coupling of a superconducting qubit to a macroscale mechanical object, thereby extending quantum control techniques to mechanical systems. By initializing the qubit in a specific superposition $\alpha_n|g\rangle + \beta_n|e\rangle$ at each step rather than $|e\rangle$, one can extend this protocol to synthesize any state in the subspace spanned by $\{|0\rangle, |1\rangle, \ldots, |n\rangle\}$ in n steps [Hof+09].

In the weak-coupling regime $g < \kappa_q, \kappa_r$, full swaps are damped by dissipation, yet the short-time Rabi oscillations governed by Eq. 2.68 still perturb the qubit population, revealing the resonant interaction. In this work, although strong coupling was not reached with our thermally occupied mechanical resonator, repeated weak-coupling exchanges allowed us to demonstrate resonant interaction between the qubit and the mechanical oscillator in its thermal state.

2.5.2 Dispersive Regime

In the limit of large detuning $\Delta = \omega_q - \omega_r \gg g$, one can eliminate the direct exchange in (2.67) via a Schrieffer–Wolff transformation[2], yielding to second order in g/Δ the effective dispersive Hamiltonian [Gam+06; Koc+07]:

$$\boldsymbol{H}_{\mathrm{disp}} = \underbrace{\hbar\omega_r\,\boldsymbol{a}^\dagger\boldsymbol{a}}_{\text{bare resonator}} + \underbrace{\frac{\hbar\omega_q}{2}\,\boldsymbol{\sigma}_z}_{\text{bare qubit}} + \underbrace{\hbar\chi\,\boldsymbol{a}^\dagger\boldsymbol{a}\,\boldsymbol{\sigma}_z}_{\substack{\text{state-dependent}\\ \text{frequency pull}}} - \underbrace{\hbar K\,\boldsymbol{a}^{\dagger 2}\boldsymbol{a}^2}_{\substack{\text{resonator}\\ \text{self-Kerr}}} \quad (2.69)$$

1: The Hamiltonian reduces to

$$\boldsymbol{H}_{\mathrm{JC}}^{(1)} = \hbar\omega_q\left(n + \frac{1}{2}\right)\mathbb{I} + \frac{\hbar\Omega}{2}\sqrt{n+1}\,\sigma_y.$$

Therefore the evolution operator $\boldsymbol{U}(t) = e^{-i\boldsymbol{H}_{\mathrm{JC}}^{(1)}t/\hbar}$ splits into the global frame rotation and a σ_y rotation:

$$e^{-i\omega_q(n+\frac{1}{2})t}\left[\cos(\theta t)\,\mathbb{I} - i\sin(\theta t)\,\sigma_y,\right]$$

where $\theta = \Omega\sqrt{n+1}/2$. We used the fact that for any involutory Λ,

$$e^{-i\theta\Lambda} = \cos(\theta)\,\mathbb{I} - i\sin(\theta)\,\Lambda.$$

Projecting $\boldsymbol{U}(t)|e, n\rangle$ onto the computational basis and removing the global rotation at frequency $\omega_q(n+\frac{1}{2})$ yields Eq. 2.68.

[Hof+09]: Hofheinz et al. (2009), 'Synthesizing arbitrary quantum states in a superconducting resonator', Nature

[Chu+18]: Chu et al. (2018), 'Creation and control of multi-phonon Fock states in a bulk acoustic-wave resonator', Nature

2: The Jaynes–Cummings Hamiltonian can be separated into diagonal and off-diagonal parts:

$$\underbrace{\hbar\omega_r\,\boldsymbol{a}^\dagger\boldsymbol{a} + \frac{\hbar\omega_q}{2}\,\boldsymbol{\sigma}_z}_{\boldsymbol{H}_0} + \underbrace{i\frac{\hbar\Omega}{2}\left(\boldsymbol{a}^\dagger\boldsymbol{\sigma}_- - \boldsymbol{a}\,\boldsymbol{\sigma}_+\right)}_{\boldsymbol{V}}.$$

To remove $\boldsymbol{V}$ at first order, choose S such that $[S, \boldsymbol{H}_0] = -\boldsymbol{V}$, which yields

$$S = i\frac{\Omega}{2\Delta}\left(\boldsymbol{a}\,\boldsymbol{\sigma}_+ + \boldsymbol{a}^\dagger\,\boldsymbol{\sigma}_-\right),$$

with $S^\dagger = -S$, so that e^S is unitary. The transformed Hamiltonian $\tilde{H} = e^S H_{\mathrm{JC}} e^{-S}$ can then be expanded, for $\Omega/\Delta \ll 1$, via the Baker–Campbell–Hausdorff formula:

$$\tilde{H} = H_{\mathrm{JC}} + [S, H_{\mathrm{JC}}] + \tfrac{1}{2}[S, [S, H_{\mathrm{JC}}]] + \ldots$$

The leading correction is

$$\frac{1}{2}[S, H_{\mathrm{JC}}] = \hbar\frac{\Omega^2}{4\Delta}\left(\boldsymbol{a}^\dagger\boldsymbol{a}\,\boldsymbol{\sigma}_z + \tfrac{\boldsymbol{\sigma}_z+1}{2}\right),$$

which corresponds to a state-dependent frequency shift. The next nonzero term (the fourth commutator) gives a self-Kerr interaction. Together, these yield the dispersive Hamiltonian Eq. 2.69.

with

$$\chi = \frac{g^2}{\Delta}, \quad K = \frac{g^4}{\Delta^3}.$$

In the dispersive regime ($|\Delta| \gg g$), no real photons are exchanged between qubit and resonator. Instead, the dispersive term $\chi\, \boldsymbol{a}^\dagger \boldsymbol{a}\, \boldsymbol{\sigma}_z$ can be viewed in two equivalent ways. From the qubit's perspective, its transition frequency is pulled by $2\chi n$ when the resonator contains n photons.

$$\boldsymbol{H}_{\text{disp}} = \frac{\hbar}{2}\left(\omega_q + 2\,\chi\, \boldsymbol{a}^\dagger \boldsymbol{a}\right) \boldsymbol{\sigma}_z + \dots.$$

Equivalently, from the resonator's perspective, its frequency shifts by $\pm\chi$ depending on whether the qubit is in $|g\rangle$ or $|e\rangle$:

$$\boldsymbol{H}_{\text{disp}} = \hbar\,(\omega_r + \chi\, \boldsymbol{\sigma}_z)\, \boldsymbol{a}^\dagger \boldsymbol{a} - \hbar K\, \boldsymbol{a}^{\dagger 2} \boldsymbol{a}^2 + \dots$$

This state-dependent frequency shift is the basis for many applications in circuit QED. In particular, when the qubit is dispersively coupled to a mechanical harmonic mode, one can infer mechanical properties (e.g. the membrane frequency and the damping rate) by measuring the qubit's frequency shift, as demonstrated in Sec. 8.1.2. Also, this dispersive coupling forms the basis of quantum-non-demolition qubit readout: when the qubit is coupled to an LC circuit, its state can be inferred by probing the resonator's response to a weak microwave tone as described in the next subsection. Finally, the self-Kerr term K introduces a weak nonlinearity to the resonator, which become relevant at high photon numbers or smaller detunings; both regime are not considered in this work.

2.5.3 Basics of Measurements in cQED: Readout Technique

In circuit QED, the qubit is typically coupled to a readout resonator mode $\boldsymbol{c}$ that is probed by sending a weak microwave tone into the cavity (see Fig. 2.17). Because the resonator frequency shifts depending on the qubit state, we can infer the qubit state by measuring the reflected signal. In what follows, we derive the reflection coefficient S_{11} of the readout cavity and examine how it changes with the qubit state.

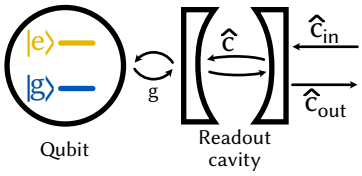

Fig. 2.17: Input–output model of the readout cavity coupled to the qubit. Single-mode cavity field $\boldsymbol{c}$ coupled to the qubit at rate g and to the external fields input $\boldsymbol{c}_{\text{in}}$ and output $\boldsymbol{c}_{\text{out}}$ modes.

Input-output formalism. In the semiclassical approach (valid for a coherent drive in a linear regime), input–output framework [GC85, eq. 2.22], the input ($\boldsymbol{c}_{\text{in}}$), output ($\boldsymbol{c}_{\text{out}}$), and cavity field ($\boldsymbol{c}$) operators satisfy

$$\boldsymbol{c}_{\text{out}} = \boldsymbol{c}_{\text{in}} + \sqrt{\kappa_c}\, \boldsymbol{c}, \tag{2.70}$$

where κ_c is the external coupling rate. Thus the reflected field is the sum of the incident wave and the fraction of the intra-cavity field that leaks out. The time evolution of the cavity operator follows [GC85, eq. 2.12]:

$$\frac{d\boldsymbol{c}}{dt} = -\frac{i}{\hbar}\,[\boldsymbol{c}, \boldsymbol{H}] - \frac{\kappa_i + \kappa_c}{2}\, \boldsymbol{c} - \sqrt{\kappa_c}\, \boldsymbol{c}_{\text{in}}, \tag{2.71}$$

[GC85]: Gardiner et al. (1985), 'Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation', Physical Review A

where κ_i is the intrinsic loss rate. For weak drives,

$$\boldsymbol{H} = \hbar\,\widetilde{\omega}_r\,\boldsymbol{c}^\dagger\boldsymbol{c} \quad\Longrightarrow\quad [\boldsymbol{c},\boldsymbol{H}] = \hbar\widetilde{\omega}_r\,\boldsymbol{c}.$$

with the effective cavity frequency $\tilde{\omega}_r = \omega_r \pm \chi$ depending on the qubit state. Defining the total decay rate $\kappa = \kappa_i + \kappa_c$, the equation becomes

$$\frac{d\boldsymbol{c}}{dt} = -\left(i\widetilde{\omega}_r + \frac{\kappa}{2}\right)\boldsymbol{c} - \sqrt{\kappa_c}\,\boldsymbol{c}_{\text{in}}. \tag{2.72}$$

Moving to the frequency domain[1], $\boldsymbol{c}(t) \to \boldsymbol{c}[\omega]$, we solve

$$-i\omega\,\boldsymbol{c}[\omega] = -\left(i\widetilde{\omega}_r + \tfrac{\kappa}{2}\right)\boldsymbol{c}[\omega] - \sqrt{\kappa_c}\,\boldsymbol{c}_{\text{in}}[\omega],$$

which yields

$$\boldsymbol{c}[\omega] = \frac{-\sqrt{\kappa_c}\,\boldsymbol{c}_{\text{in}}[\omega]}{\kappa/2 - i\Delta}, \tag{2.73}$$

where $\Delta \equiv \omega - \widetilde{\omega}_r$ is the detuning between the input signal frequency ω and the cavity's resonance frequency $\widetilde{\omega}_r$. The reflection coefficient $S_{11} = \boldsymbol{c}_{\text{out}}/\boldsymbol{c}_{\text{in}}$ then follows:

$$\boxed{\begin{aligned} S_{11} &= \frac{(\kappa_i - \kappa_c)/2 - i\Delta}{(\kappa_i + \kappa_c)/2 - i\Delta}, \\ &= \frac{2\kappa_c}{\kappa_i + \kappa_c}\,\frac{1}{1 + i\frac{\kappa_i+\kappa_c}{2\Delta}} + \frac{\kappa_i - \kappa_c}{\kappa_i + \kappa_c}. \end{aligned}} \tag{2.74}$$

This equation shows that:

- The *linewidth* (FWHM) of the resonance is set by the total loss rate $\kappa = \kappa_i + \kappa_c$.
- The *dip depth* in $|S_{11}|$ is $(\kappa_i - \kappa_c)/(\kappa_i + \kappa_c)$, reflecting the balance between intrinsic and external losses.

Dispersive shift readout. We send an input probe tone $s_{\text{in}}(t) = \tilde{s}_{\text{in}}(t)\,e^{i\omega_d t}$ into the readout resonator and record the reflected output $s_{\text{out}}(t) = S_{11}(\omega_d)\,s_{\text{in}}(t)$. Its amplitude and phase encode the qubit state via the response function $S_{11}(\omega_d)$. In Fig. 2.18, we plot $S_{11}(\omega_d)$ versus detuning $\Delta = \omega_d - \omega_r$ for the qubit prepared in $|g\rangle$ (blue) and $|e\rangle$ (yellow). Each trace is a Lorentzian centered at $\omega_r \pm \chi$ with width κ_+. This state-dependent response is the fundamental readout mechanism in circuit QED and underlies the fluxonium readout described in Sec. 3.2.3.

IQ plane readout. Practically, one don't need to sweep the drive frequency to distinguish between states, instead for a fixed probe frequency ω_d the reflected signal[2]

$$\tilde{s}_{\text{out}}(t) = S_{11}(\omega_d)\,\tilde{s}_{\text{in}}(t)$$

differs for the two qubit states. In an ideal, noiseless world, one could read out the state from $\tilde{s}_{\text{out}}(t)$ at a single instant, but thermal and amplifier noise scatter each sample. To accumulate the information, we integrate

1: In the Heisenberg picture, the cavity operator $\boldsymbol{c}(t)$ carries the full time dependence. We define its Fourier transform,

$$\boldsymbol{c}[\omega] \equiv \int_{-\infty}^{\infty} e^{\,i\omega t}\,\boldsymbol{c}(t)\,dt,$$

which yields the *frequency–domain Heisenberg operator* $\boldsymbol{c}[\omega]$.

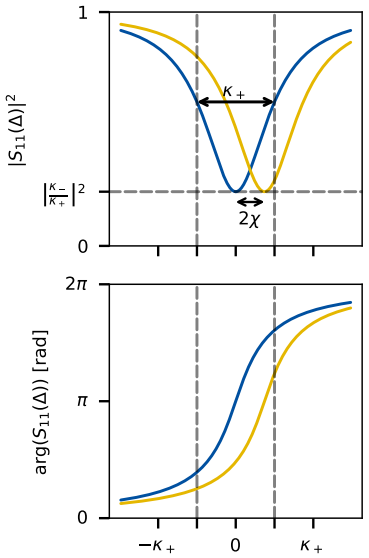

Fig. 2.18: Magnitude and phase of the readout reflection coefficient S_{11} when the qubit is in $|g\rangle$ (blue) and $|e\rangle$ (yellow). (*Top*) $|S_{11}(\Delta)|^2$ versus detuning Δ from Eq. 2.74. The full-width at half-maximum is $\kappa_+ = \kappa_i + \kappa_c$, and the minimum reflectivity is $\left(\kappa_-/\kappa_+\right)^2$ with $\kappa_- = \kappa_i - \kappa_c$. (*Bottom*) Phase $\arg S_{11}(\Delta)$ versus Δ.

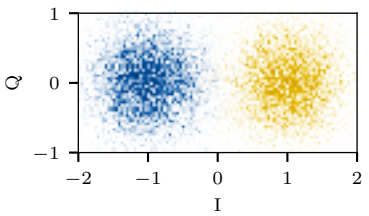

Fig. 2.19: Illustration of the readout of a qubit in the IQ plane, for the qubit prepared in $|g\rangle$ (blue) and $|e\rangle$ (yellow).

2: To obtain this complex envelope $\tilde{s}_{\text{out}}(t)$, one performs IQ demodulation by mixing the reflected signal with a local oscillator at ω_d and low-pass filtering the result. This procedure is explained in more detail in Sec. 7.5.1.

the envelope over a readout window τ (typically 3–7 μs), producing the single complex outcome

$$S = \int_0^\tau \tilde{s}_{\text{out}}(t)\,\mathrm{d}t.$$

Plotting Re$[S]$ (the in-phase I) versus Im$[S]$ (the quadrature Q) in the IQ plane yields two Gaussian clusters—one for $|g\rangle$ (left), one for $|e\rangle$ (right), as shown in Fig. 2.19. We then classify each new measurement by its nearest cluster center. When the clusters overlap minimally, as in Fig. 2.19, a single measurement is sufficient to determine the state, leading to a *single-shot readout*.

Fluxonium theory | 3

The heavy fluxonium qubit is the primary platform employed in this thesis to achieve resonant coupling with a few-MHz mechanical resonator. This chapter introduces the theory of the fluxonium qubit, its underlying physical principles, and relevant considerations.

Outline. This chapter begins with introducing the fluxonium qubit and its different regimes. We aim to build an intuition behind the main fluxonium parameters, and explain how one can achieve fluxonium with the lowest transition frequency as low as a few MHz. We then introduce the main coupling scheme in order to control the fluxonium qubit experimentally. Finally, we introduce the main decoherence mechanisms and their impact on the qubit coherence.

3.1 Fluxonium fundamentals

The fluxonium qubit was first proposed in 2009 by Manucharyan *et al.* [Man+09] as an extension of the transmon architecture, wherein a large inductor—referred to as a superinductor—is placed in parallel with the shunt capacitor, suppressing offset-charge noise. In this section, we will explore the theoretical foundations of the fluxonium qubit, focusing on its Hamiltonian, energy landscape, and transitions, with the goal of building an intuitive understanding of its characteristic parameters. At the end of this section, we will describe the heavy fluxonium regime, which is the focus of this work.

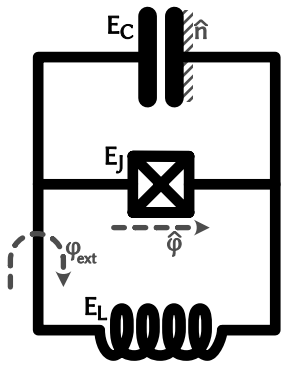

Fig. 3.1: Lumped-element equivalent circuit of the fluxonium qubit. It consists of a capacitor (C), an inductor (L), and a Josephson junction (J) connected in parallel. The phase $\hat{\varphi}$ is the superconducting phase difference between the left and right nodes. Depending on the gauge choice, an additional phase offset φ_{ext} may be included in the Josephson-junction or inductive terms.

[Man+09]: Manucharyan et al. (2009), 'Fluxonium: single Cooper pair circuit free of charge offsets', Science

3.1.1 Fluxonium Hamiltonian

The lumped-element circuit of the fluxonium (see Fig. 3.1) can be obtained by augmenting the transmon Hamiltonian (Eq. 2.37) with an inductive term (Eq. 2.20). Since this inductive contribution enters only as an additional potential energy term, it does not modify the canonical momentum operator previously derived for the transmon. Concretely, the full Hamiltonian reads[1]

$$\boldsymbol{H} = \underbrace{4E_C\boldsymbol{n}^2}_{\substack{\text{kinetic}\\ \text{energy}}} + \underbrace{\frac{E_L}{2}\boldsymbol{\varphi}^2}_{\substack{\text{harmonic}\\ \text{potential}}} \underbrace{- E_J\cos\big(\boldsymbol{\varphi} - \varphi_{\text{ext}}\big)}_{\substack{\text{anharmonic}\\ \text{potential}}}, \tag{3.1}$$

1: φ_{ext} may equivalently be introduced in the inductive term; here we include it in the Josephson cosine by choice.

where

- $E_C \equiv e^2/2C$ is the capacitive energy, associated with the capacitance C.
- E_J is the Josephson energy of the junction.
- $E_L \equiv (\Phi_0/2\pi)^2/L \equiv \phi_0^2/L$ is the inductive energy determined by the inductance L.

- $\boldsymbol{n}$ and $\boldsymbol{\varphi}$ are the dimensionless charge and phase operators, respectively, obeying $[\boldsymbol{\varphi}, \boldsymbol{n}] = i$.
- $\varphi_{\text{ext}} \equiv \Phi_{\text{ext}}/\phi_0$ is the external phase bias introduced by the magnetic flux Φ_{ext} threading the qubit loop.

Unlike the Cooper-pair box/transmon, where an isolated superconducting island renders the phase variable φ compact ($\varphi \equiv \varphi + 2\pi$) and thereby enforces an integer-valued conjugate charge operator $\hat{n} \in \mathbb{Z}$ (Cooper-pair number on the island), the fluxonium circuit is shunted by a large inductance. The resulting inductive energy $\frac{1}{2}E_L\varphi^2$ is not 2π-periodic, so in the standard node-flux formulation φ is naturally treated as a non-compact coordinate on $\mathbb{R}$. Consequently, the canonical conjugate $\hat{n}$ corresponds to a continuous capacitor charge (in units of $2e$) and has a continuous spectrum. This should not be conflated with the discreteness of Cooper-pair tunneling through the Josephson junction: tunneling events remain quantized, but $\hat{n}$ in the fluxonium Hamiltonian denotes the continuous canonical charge rather than a literal counter of transferred Cooper pairs.

At first glance, the fluxonium Hamiltonian appears similar to that of the transmon, and it is not immediately obvious why it exhibits insensitivity to static charge offsets. Intuitively, this insensitivity arises because the charge node (island) is no longer galvanically isolated: it is shunted by a large inductance. This connectivity effectively eliminates the island's charge-offset degree of freedom. Indeed, any static offset charge $N_{\text{offset}} = 2e\, n_{\text{offset}}$ can be removed by the unitary transformation $\boldsymbol{n} \to \boldsymbol{n} - n_{\text{offset}}$, implemented via $U = e^{in_{\text{offset}}\boldsymbol{\varphi}}$. This transformation leaves the fluxonium Hamiltonian invariant due to the unbounded nature of the phase operator $\boldsymbol{\varphi}$. In contrast, this procedure does not apply to the transmon Hamiltonian, as $\boldsymbol{\varphi}$ in that context is a compact variable with 2π periodicity, rendering the operator U non-single-valued. As a result, while the Cooper-pair box only achieves partial charge insensitivity in the large-E_J/E_C limit (transmon regime) and primarily for the lowest-lying eigenstates, the fluxonium Hamiltonian displays genuine charge-offset insensitivity. In its original conception [Man+09], the fluxonium was operated in a regime characterized by the energy scale hierarchy $E_J > E_C > E_L$. In this work, however, we focus on the so-called *heavy fluxonium* regime [Zha+21], where the hierarchy is $E_J \gg E_C > E_L$.

Note that the canonical commutation relation $[\boldsymbol{\varphi}, \boldsymbol{n}] = i$ implies that the matrix elements of $\boldsymbol{\varphi}$ and $\boldsymbol{n}$ are related by[1]:

$$\boxed{n_{\alpha\beta} = \frac{i\hbar\omega_{\alpha\beta}C}{(2e)^2}\varphi_{\alpha\beta} = \frac{i\hbar\omega_{\alpha\beta}}{8E_C}\varphi_{\alpha\beta},} \tag{3.2}$$

where $\boldsymbol{n}_{\alpha\beta} = \langle\alpha|\boldsymbol{n}|\beta\rangle$, $\boldsymbol{\varphi}_{\alpha\beta} = \langle\alpha|\boldsymbol{\varphi}|\beta\rangle$ are the matrix elements between the qubit states $|\alpha\rangle$ and $|\beta\rangle$, $\omega_{\alpha\beta} = \omega_\alpha - \omega_\beta$ is the transition frequency, and C is the total capacitance of the qubit[2].

In the next subsections, we will explore why this particular regime is of interest and what it reveals about the behavior of the fluxonium. We'll examine how the different parameters—E_J, E_C, E_L, and φ_{ext}—contribute to shaping the qubit's energy landscape, with the goal of building an intuitive understanding of the fluxonium qubit and its transitions.

[Zha+21]: Zhang et al. (2021), 'Universal fast flux control of a coherent, low-frequency qubit', Physical Review X

1: This expression is valid for an arbitrary nonlinear oscillator characterized by a linear capacitance C and a nonlinear inductance. Further discussion on this topic can be found in [Man12, p. 96]. The relation can be derived by evaluating the matrix elements using the commutation relationship: $\langle i|[\boldsymbol{\varphi}, \boldsymbol{H}]|j\rangle = \hbar\omega_{ji}\langle i|\boldsymbol{\varphi}|j\rangle = i8E_C\langle i|\boldsymbol{n}|j\rangle$.

[Man12]: Manucharyan (2012), 'Superinductance', Thesis

2: Note that while we often represent the capacitance as a single element, other components, such as small junction capacitance, coupling capacitors, stray capacitance in the array, and the stray capacitance of the leads, also contribute to the total capacitance, which is included in $e^2/(2E_C)$ [Mas+12, p. 47].

3.1.2 Fluxonium spectrum: Plasmon/Fluxon transitions

Let us begin by examining the fluxonium potential, which is the sum of a global parabolic confinement, $\frac{1}{2}E_L\varphi^2$, and a periodic Josephson term, $-E_J\cos(\varphi-\varphi_{\text{ext}})$. The cosine creates an infinite array of wells centered near $\varphi = 2\pi m + \varphi_{\text{ext}}$, whose depths are controlled by the ratio E_J/E_L.

Classically, each well (labeled by integer k) hosts an equilibrium at $\varphi_{\min,k}$, determined by

$$\left.\frac{\partial U}{\partial\varphi}\right|_{\varphi_{\min,k}} = E_L\,\varphi_{\min,k} + E_J\sin\left(\varphi_{\min,k}-\varphi_{\text{ext}}\right) = 0.$$

At this point, a persistent current

$$I_k = I_c\,\sin\left(\varphi_{\min,k}-\varphi_{\text{ext}}\right)$$

flows around the loop. We call these localized classical states "fluxons." In an ideal superconducting ring, such a current would remain perfectly constant, flowing clockwise or counterclockwise depending on the sign of the phase difference. However, the inclusion of the Josephson junction permits 2π phase-slips that allow the system to tunnel between adjacent wells. These phase-slip events—i.e. the tunneling of a fluxon through the junction—are central to the quantum dynamics of the fluxonium.

When $E_J \gg E_L$—the typical fluxonium regime—deep cosine wells appear within a shallow global parabolic confinement (see Fig. 3.2(a)). In contrast, if E_L becomes comparable to or larger than E_J (i.e., $E_J/E_L \lesssim 1$), the quadratic term dominates and the potential resembles a broad harmonic well with minor ripples from the cosine (see Fig. 3.2(b)). Since this latter regime behaves like a weakly perturbed harmonic oscillator, we focus on the more relevant regime $E_J \gg E_L$, where the cosine structure becomes important.

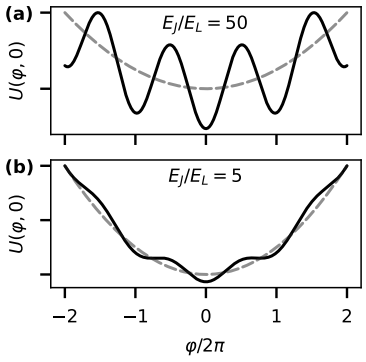

Fig. 3.2: Fluxonium potential at $\varphi_{\text{ext}} = 0$ as a function of φ for $E_J = 4.82$ GHz and $E_L = 0.128$ GHz. The harmonic potential $\frac{E_L}{2}\boldsymbol{\varphi}^2$ is plotted with dashed lines.

A defining feature of the fluxonium qubit is its tunability via the external magnetic flux φ_{ext}. As this parameter is varied continuously from 0 to 2π, the shape of the potential energy landscape $U(\boldsymbol{\varphi}, \varphi_{\text{ext}})$ changes significantly (see Fig. 3.3). At $\varphi_{\text{ext}} = 0$, the potential consists of a dominant central well due to the alignment of the cosine minima with the origin (see Fig. 3.3(a)). As φ_{ext} increases toward π, the landscape gradually transforms into a symmetric double-well potential, with two minima centered near $\boldsymbol{\varphi} \approx \pm\pi$ (see Fig. 3.3(c)).

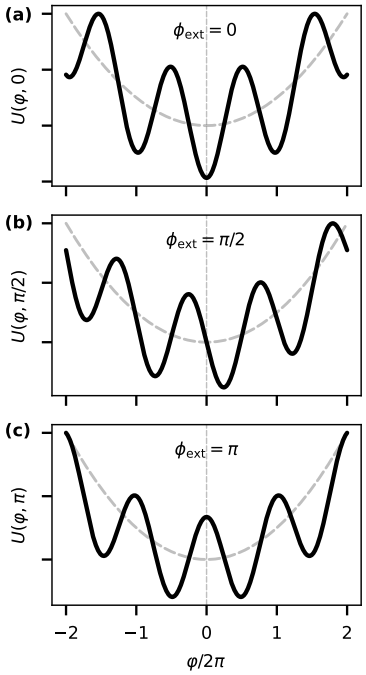

Fig. 3.3: Fluxonium potential at different φ_{ext}: $U(\boldsymbol{\varphi}, \varphi_{\text{ext}}) = \frac{E_L}{2}\boldsymbol{\varphi}^2 - E_J\cos(\boldsymbol{\varphi}-\varphi_{\text{ext}})$, for $E_L = 0.128$ GHz and $E_J = 4.82$ GHz. The potential is shown over $\boldsymbol{\varphi} \in [-20, 20]$ for several values of φ_{ext}. The dashed line shows the harmonic confinement from the inductive term.

To understand the nature of the transitions in such a potential, consider a representative intermediate value such as $\varphi_{\text{ext}} = 0.25\pi$ (see Fig. 3.3(b)). Here, the wells are generally well-separated and deep enough to support localized eigenstates (provided $E_J > E_C$). In this regime, relevant qubit transitions can be broadly classified into two types, based on the spatial localization of the wavefunctions involved.

The first class includes transitions that occur within a single potential well—e.g., from the ground state to the first excited state localized in the same well (see Fig. 3.4(a)). These transitions closely resemble those in a transmon qubit, involving no jumps in the average phase $\langle\boldsymbol{\varphi}\rangle$, and are thus referred to as *plasmon-like transitions*. Their frequencies are

primarily determined by the local curvature of the potential and are well-approximated by the plasma frequency:

$$\hbar\omega_p = \sqrt{8E_C E_J}.$$

The second class consists of transitions between states whose wavefunctions are localized in different wells (see Fig. 3.4(b)). These are known as *fluxon-like transitions* because they involve a significant change in the average phase $\langle\boldsymbol{\varphi}\rangle$, corresponding to the addition or removal of a flux quantum in the loop. In this case, the energy difference between states is governed by the relative depths of the wells. When the wavefunctions are well-localized in adjacent wells, the transition energy can be approximated by comparing the parabolic offset between wells located near $\boldsymbol{\varphi} \approx \varphi_{\text{ext}}$ and $\boldsymbol{\varphi} \approx \varphi_{\text{ext}} - 2\pi$. This gives:

$$\hbar\omega_{01} \approx \frac{E_L}{2}\left((\varphi_{\text{ext}} - 2\pi)^2 - \varphi_{\text{ext}}^2\right) = 2\pi^2 E_L\left(1 - \frac{2|\varphi_{\text{ext}}|}{\pi}\right). \tag{3.3}$$

This expression highlights the approximate linear dependence of the fluxon transition frequency on both E_L and the external phase φ_{ext}, away from the symmetry points $\varphi_{\text{ext}} = 0$ and $\varphi_{\text{ext}} = \pi$. Although this formula is a rough approximation, it already shows that as φ_{ext} increases from 0 to π, the fluxon transition frequency decreases. This makes $\varphi_{\text{ext}} = \pi$ an especially interesting operating point, where the lowest transition frequency reaches a minimum.

This behavior is confirmed by numerical diagonalization of the full Hamiltonian. In the following sections, we will explore this special case in more detail. This point—known as the **flux frustration point**—corresponds to a symmetric double-well potential, where the energy splitting is dominated by quantum tunneling between nearly degenerate wells.

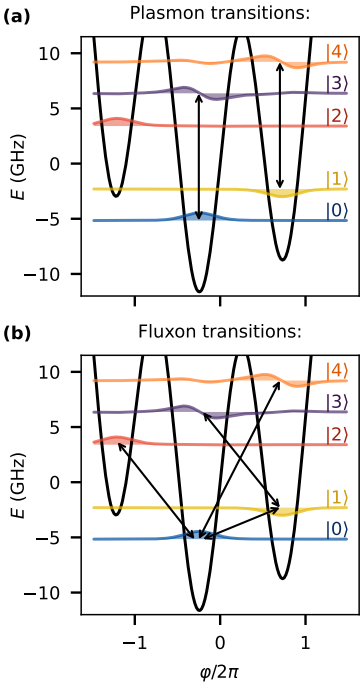

Fig. 3.4: Plasmon/Fluxon transitions in fluxonium. The first five energy levels of the fluxonium qubit and their corresponding wavefunctions, for E_J = 12 GHz, E_C = 2 GHz, E_L = 0.128 GHz, and $\varphi_{\text{ext}} = 0.3\pi$. Black arrows indicate plasmon-like transitions (a) and fluxon-like transitions (b).

State labeling convention. In what follows, we focus on the lowest four energy levels of the fluxonium. The ground and first excited states are denoted $|g\rangle$ and $|e\rangle$, respectively, and the next two levels are labeled $|f\rangle$ and $|h\rangle$ (the subsequent letters after "e" in the alphabet).

3.1.3 Flux Frustration Point: $\varphi_{\text{ext}} = \pi$ ($\Phi_{ext} = \Phi_0/2$)

Ground and first excited states. The flux frustration point, where the external flux Φ_{ext} equals half a flux quantum (i.e. $\varphi_{\text{ext}} = \pi$), is of particular operational significance for fluxonium qubits. This point is attractive because it features reduced sensitivity to flux noise (first-order insensitivity, $\partial\omega_{ge}/\partial\varphi_{\text{ext}} = 0$) and allows for very low qubit transition frequencies, as we will detail in following.

At $\varphi_{\text{ext}} = \pi$, the potential becomes:

$$U(\boldsymbol{\varphi}, \pi) = \frac{E_L}{2}\boldsymbol{\varphi}^2 - E_J\cos(\boldsymbol{\varphi} - \pi) = \frac{E_L}{2}\boldsymbol{\varphi}^2 + E_J\cos(\boldsymbol{\varphi}). \tag{3.4}$$

The cosine term has minima at $\boldsymbol{\varphi} = \pm\pi, \pm 3\pi, \ldots$, while the inductive term provides a global parabolic confinement. Together, these produce

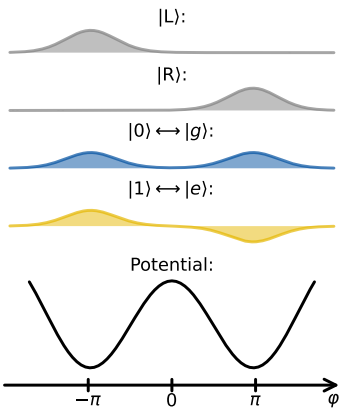

Fig. 3.5: Illustration of the fluxonium wavefunctions at $\varphi_{\text{ext}} = \pi$. Top: the localized left and right wavefunctions $|L\rangle$ and $|R\rangle$ (which are not eigenstates of the fluxonium Hamiltonian). Middle: the true eigenstates $|g\rangle$ and $|e\rangle$, which form symmetric and antisymmetric superpositions of $|L\rangle$ and $|R\rangle$, respectively. Bottom: the corresponding fluxonium potential. All wavefunctions and the potential are plotted in arbitrary units.

a symmetric double-well potential whose two lowest-energy wells lie at $\boldsymbol{\varphi}_{\pm} \approx \pm\left(1 - \frac{E_L}{2E_J}\right)\pi$, (i.e. around $\pm\pi$ but pulled slightly inward by the ratio E_L/E_J). Classically, this potential supports states localized in the left ($|L\rangle$) and right ($|R\rangle$) wells, corresponding respectively to anticlockwise ($\langle\boldsymbol{\varphi}\rangle < 0$) and clockwise ($\langle\boldsymbol{\varphi}\rangle > 0$) circulating currents. Quantum tunneling through the barrier dresses these localized states into the first two fluxonium eigenstates, which form symmetric and antisymmetric superpositions:

$$|g\rangle = \frac{1}{\sqrt{2}}\big(|L\rangle + |R\rangle\big), \quad |e\rangle = \frac{1}{\sqrt{2}}\big(|L\rangle - |R\rangle\big),$$

as illustrated in Fig. 3.5.

Parity symmetry and selection rules. Notably, at $\varphi_{\text{ext}} = \pi$, the fluxonium Hamiltonian Eq. 3.4 is invariant under the parity transformation[1]

$$\boldsymbol{P}: \quad \boldsymbol{\varphi} \mapsto -\boldsymbol{\varphi},$$

Hence, the eigenstate $|\alpha\rangle$ is characterized by its parity $p_\alpha \in \{\pm 1\}$ (+1 for symmetric, −1 for antisymmetric). This symmetry enforces selection rules: matrix elements of the odd operator $\boldsymbol{\varphi}$ vanish between same-parity states,

$$\langle\psi_\alpha|\boldsymbol{\varphi}|\psi_\beta\rangle \;\mapsto\; \langle\psi_\alpha(-\varphi)|-\boldsymbol{\varphi}|\psi_\beta(-\varphi)\rangle = -\langle\psi_\alpha|\boldsymbol{\varphi}|\psi_\beta\rangle\,, \quad \text{if } p_\alpha = p_\beta.$$

Thus, transitions between same-parity levels are forbidden. Concretely, among the first four levels only these transitions are allowed: $|g\rangle \to |e\rangle$, $|g\rangle \to |h\rangle$, $|e\rangle \to |f\rangle$, and $|f\rangle \to |h\rangle$.

First-order insensitivity to flux noise. Moreover, this reflection symmetry renders the qubit transition frequency ω_{ge} first-order insensitive to flux noise at $\varphi_{\text{ext}} = \pi$. Indeed,

$$\hbar\,\frac{\partial\omega_{ge}}{\partial\varphi_{\text{ext}}} = \langle e|\partial_{\varphi_{\text{ext}}}\boldsymbol{H}|e\rangle - \langle g|\partial_{\varphi_{\text{ext}}}\boldsymbol{H}|g\rangle\,.$$

Since $\partial_{\varphi_{\text{ext}}}\boldsymbol{H} = E_J\sin\boldsymbol{\varphi}$ anticommutes with $\boldsymbol{P}$ ($\boldsymbol{P}\sin\boldsymbol{\varphi}\,\boldsymbol{P} = -\sin\boldsymbol{\varphi}$), one has

$$\langle\alpha|\partial_{\varphi_{\text{ext}}}\boldsymbol{H}|\alpha\rangle = -\langle\alpha|\partial_{\varphi_{\text{ext}}}\boldsymbol{H}|\alpha\rangle = 0, \quad \forall\,\alpha,$$

and thus $\partial\omega_{ge}/\partial\varphi_{\text{ext}} = 0$. Consequently, the fluxonium qubit not only inherits inherent immunity to static charge offsets, but also becomes first-order insensitive to flux noise at the frustration point $\varphi_{\text{ext}} = \pi$. This operating condition is known as the *flux sweet spot*, which we adopt as our primary bias point throughout this work.

3.1.4 Heavy Fluxonium Regime

Previously, we have shown that increasing the charging energy E_C is analogous to increasing the mass of a particle in a harmonic potential: the larger E_C, the more tightly the wavefunctions confine within each well,

1: The parity operator $\boldsymbol{P}$ is unitary ($\boldsymbol{P}^2 = \mathbb{1}$) and acts as

$$\boldsymbol{P}\,|\psi_\alpha(\varphi)\rangle = |\psi_\alpha(-\varphi)\rangle = p_\alpha\;|\psi_\alpha(\varphi)\rangle\,,$$
$$\Leftrightarrow \boldsymbol{P}\,|\alpha\rangle = p_\alpha\;|\alpha\rangle\,,$$

where $p_\alpha = \pm 1$ is the state's parity (eigenvalue of $\boldsymbol{P}$). Since Hamiltonian is invariant under parity transformation

$$[\boldsymbol{H}, \boldsymbol{P}]\Big|_{\varphi_{\text{ext}}=\pi} = 0$$

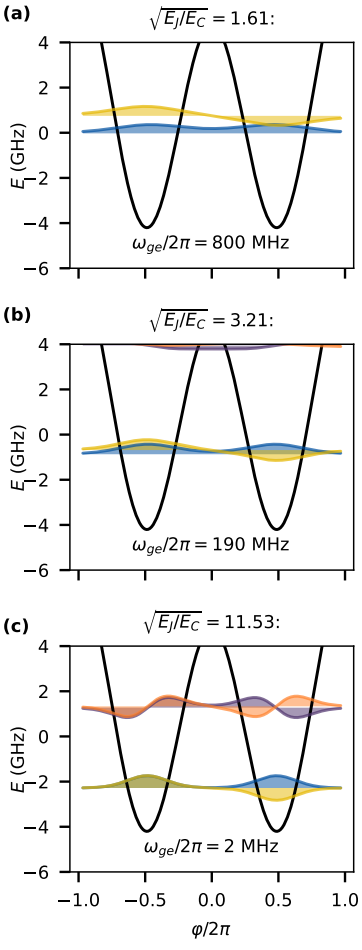

Fig. 3.6: Fluxonium spectrum at different E_J/E_C with constant E_L.

yielding smaller intrawell level spacings. At the flux-frustration point $\varphi_{\text{ext}} = \pi$, however, the lowest transition $|g\rangle \to |e\rangle$ is not an intrawell excitation but a *fluxon* transition—arising from quantum tunneling between the left and right wells. In the following, we develop a more detailed intuitive picture of the fluxon tunneling mechanism that underlies this transition.

We consider the potential near the minima at $\varphi_0 = \pm\pi$. Neglecting E_L for simplicity, we expand the cosine around the minimum:

$$E_J \cos(\varphi) \approx -E_J + \frac{E_J}{2}(\varphi - \varphi_0)^2,$$

which yields local harmonic left and right wells of the form $V_{L,R}(\varphi) \approx \frac{E_J}{2}(\varphi \mp \pi)^2$. The ground state wavefunctions in each well ($|L\rangle$ and $|R\rangle$) are approximately Gaussian:

$$\psi_{L,R}(\varphi) = \frac{1}{(\pi\varphi_{\text{zpf}}^2)^{1/4}} \exp\left(-\frac{(\varphi \mp \pi)^2}{2\varphi_{\text{zpf}}^2}\right), \quad \varphi_{\text{zpf}} = \left(\frac{2E_C}{E_J}\right)^{1/4}. \tag{3.5}$$

The qubit transition frequency $\omega_{ge} = (E_e - E_g)/\hbar$ is set by the tunnel splitting between these two states and is approximately proportional to the overlap integral $\langle L|R\rangle$:

$$\hbar\omega_{ge} \sim \langle L|R\rangle \sim \exp\left(-\sqrt{E_J/E_C}\right). \tag{3.6}$$

More accurate expressions for ω_{ge} at $\varphi_{\text{ext}} = \pi$ can be obtained via the WKB approximation. An early result (neglecting E_L) was derived by Koch *et al.* [Koc+07, Eq. 2.5], giving:

[Koc+07]: Koch et al. (2007), 'Charge insensitive qubit design derived from the Cooper pair box', Physical Review A

$$\hbar\omega_{ge} \approx \hbar\omega_p \cdot \lambda \cdot \left(\frac{E_J}{E_C}\right)^{1/4} \exp\left(-\sqrt{\frac{8E_J}{E_C}}\right), \tag{3.7}$$

where λ is a numerical prefactor. A significantly improved WKB approximation, incorporating the inductive energy E_L and showing excellent agreement with numerical simulations, was recently developed by Vakhtel *et al.* [Vak+24, Eq. 6, Fig. 6] and Ardati *et al.* [Ard+24, Eq. 6]. It reads:

[Vak+24]: Vakhtel et al. (2024), 'Tunneling of fluxons via a Josephson resonant level', Physical Review B

[Ard+24]: Ardati et al. (2024), 'Using Bifluxon Tunneling to Protect the Fluxonium Qubit', Phys. Rev. X

$$\boxed{\omega_{ge} = 4\omega_p \left(\frac{2E_J}{\pi^2 E_C}\right)^{\frac{1}{4}} \exp\left(-\sqrt{\frac{8E_J}{E_C}} + \frac{14\zeta(3)E_L}{\omega_p}\right),} \tag{3.8}$$

where $\zeta(3) \approx 1.20205$ is Apéry's constant. This expression shows that, at $\varphi_{\text{ext}} = \pi$, the qubit transition frequency ω_{ge} depends exponentially on the ratio $\sqrt{E_J/E_C}$. Consequently, even small variations in E_J or E_C produce large shifts in ω_{ge}, making the transition both extremely tunable by design and sensitive to fabrication tolerances.

By deliberately increasing the ratio E_J/E_C, one can push ω_{ge} far below the typical gigahertz range into the megahertz range. This is the essence of the *heavy fluxonium regime* [Zha+21]. While standard fluxonium already

[Zha+21]: Zhang et al. (2021), 'Universal fast flux control of a coherent, low-frequency qubit', Physical Review X

operates with $E_J \gtrsim E_C > E_L$, the heavy fluxonium pushes this hierarchy further—typically by engineering E_C to be particularly small, which is achieved by adding a large shunt capacitance ($C = e^2/(2E_C)$). The resulting parameter hierarchy is

$$E_J \gg E_C > E_L$$

with E_L still being the smallest scale. In Fig. 3.6, we hold E_J and E_L fixed and vary E_C, showing that ω_{ge} decreases rapidly as E_C is reduced.

The practical realization of such low-frequency qubits was convincingly demonstrated by Zhang *et al.* [Zha+21], who fabricated heavy fluxonium qubits with transition frequencies as low as 14 MHz while maintaining good coherence. This low-frequency operation is especially attractive for applications involving resonant coupling to systems that naturally operate in the MHz range, such as the nanomechanical resonators used in this work. In 2023, we demonstrated [Naj+24] that we could fabricate heavy fluxonium qubits with frequencies around 2 MHz, achieving a charge sensitivity of 33 $\mu e/\sqrt{\text{Hz}}$. This result served as a proof of concept that heavy fluxonium can be used to couple to nanomechanical resonators and gave us the confidence to pursue this line of research further and achieve the results presented in this thesis.

[Zha+21]: Zhang et al. (2021), 'Universal fast flux control of a coherent, low-frequency qubit', Physical Review X

[Naj+24]: Najera-Santos et al. (2024), 'High-Sensitivity ac-Charge Detection with a MHz-Frequency Fluxonium Qubit', Physical Review X

Note that the final spectrum of the heavy fluxonium qubit (see Fig. 3.6(c)) is highly anharmonic compared to a transmon: only the $|g\rangle \rightarrow |e\rangle$ transition lies in the MHz range, while all higher transitions reside in the GHz band. This pronounced separation isolates the qubit levels($\{|g\rangle, |e\rangle\}$) from thermally unpopulated levels, facilitating both fast control pulses and high-fidelity readout.

3.2 Fluxonium Control and Readout

The isolated fluxonium qubit has limited utility unless its state can be both manipulated and measured. In Sec. 2.4.4 we introduced the basic principles of qubit control, and in Sec. 2.5.3 we outlined fundamental readout techniques. In this section, we derive the mechanisms we previously took for granted. We begin by examining direct capacitive coupling to a microwave drive line for qubit state manipulation (Sec. 3.2.1). We then turn to inductive coupling via a dedicated flux-bias line to tune the external flux φ_{ext} (Sec. 3.2.2). Finally, we analyze the dispersive interaction between the fluxonium and an auxiliary LC resonator, the workhorse for nondestructive readout (Sec. 3.2.3).

3.2.1 Direct Capacitive Coupling. Fluxonium Drive

To interact with and control the qubit, we drive it by applying a time-dependent voltage source $V_d(t)$ to the fluxonium via a coupling capacitor C_d (see Fig. 3.7). This capacitive connection injects charge into the qubit node, creating an electric-dipole interaction that directly drives transitions between its energy levels.

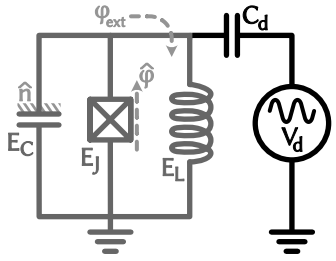

Fig. 3.7: Lumped-element equivalent circuit of the fluxonium qubit coupled capacitively to a drive line.

Hamiltonian derivation. Using the previously derived expressions for the fluxonium qubit's energy (Eq. 3.1) and the capacitor's energy (Eq. 2.17), the system's Lagrangian can be expressed as:

$$\mathcal{L} = \underbrace{\frac{C_q}{2}\dot{\Phi}^2 + E_J \cos\left(\frac{\Phi - \Phi_{\text{ext}}}{\phi_0}\right) + \frac{\Phi^2}{2L}}_{\text{Fluxonium}} + \underbrace{\frac{C_d}{2}\left(\dot{\Phi} - V_d(t)\right)^2}_{\text{Drive coupling}}. \tag{3.9}$$

By applying the Legendre transformation and introducing the canonical variable Q with $\dot{\Phi} = \frac{Q+C_d}{C_\Sigma}$, where $C_\Sigma = C_d + C_q$, we can represent the Hamiltonian of the system (discarding constant terms) as:

$$\boldsymbol{H} = \frac{(Q + C_d \cdot V_d)^2}{2C_\Sigma} - E_J \cos\left(\frac{\boldsymbol{\Phi} - \Phi_{\text{ext}}}{\phi_0}\right) + \frac{\boldsymbol{\Phi}^2}{2L}. \tag{3.10}$$

Using the dimensionless charge and flux operators, $\boldsymbol{n} = \boldsymbol{Q}/2e$ and $\boldsymbol{\varphi} = \boldsymbol{\Phi}/\phi_0$, we retain the inductive energy $E_L = \phi_0^2/L$ and redefine the capacitive energy to include the full capacitances as $E_C = e^2/2C_\Sigma$. This yields the Hamiltonian:

$$\begin{aligned} \boldsymbol{H} &= 4E_C(\boldsymbol{n} + n_d)^2 - E_J \cos(\boldsymbol{\varphi} - \varphi_{\text{ext}}) + \frac{E_L}{2}\boldsymbol{\varphi}^2 \\ &= H_{\text{fluxonium}} + \underbrace{8E_C \cdot n_d \cdot \boldsymbol{n}}_{H'}, \end{aligned} \tag{3.11}$$

where the charge induced by the capacitive drive is $n_d = C_d V_d/2e$. This Hamiltonian demonstrates that a simple capacitive drive offsets the charge.

Two-level system projection. The fluxonium is driven by a resonant pulse at the target transition frequency. By applying the rotating-wave approximation, we discard all nonresonant terms and retain only the resonant coupling. Consequently, we restrict our analysis to an effective two-level system, since the other energy levels lie far off resonance.

We begin by projecting the full system Hamiltonian onto the two-level subspace spanned by the states $|\alpha\rangle$ and $|\beta\rangle$. Defining the projector

$$\boldsymbol{P}_{\alpha\beta} = |\alpha\rangle\langle\alpha| + |\beta\rangle\langle\beta|,$$

any operator $\boldsymbol{O}$ is reduced as $\boldsymbol{P}_{\alpha\beta}\,\boldsymbol{O}\,\boldsymbol{P}_{\alpha\beta}$. In particular, for a drive voltage $V_d(t) = V_0\cos\big(\omega_d t + \phi\big)$, the projected Hamiltonian becomes

$$\boldsymbol{H} = \boldsymbol{P}_{\alpha\beta}\,\boldsymbol{H}\,\boldsymbol{P}_{\alpha\beta} = \underbrace{\frac{\hbar\,\widetilde{\omega}_{\alpha\beta}}{2}\,\boldsymbol{\sigma}_z}_{\text{Qubit}} + \underbrace{\hbar\,\Omega\cos\big(\omega_d t + \phi\big)\,\boldsymbol{\sigma}_x}_{\text{Drive coupling}}. \tag{3.12}$$

Here we have omitted any overall constants that only shift the zero of energy and defined the dressed qubit frequency $\widetilde{\omega}_{\alpha\beta}$ with the Rabi frequency (coupling strength) Ω.

$$\widetilde{\omega}_{\alpha\beta} = \omega_{\alpha\beta} + 8\,\frac{E_C}{\hbar}\,n_d\,(n_{\beta\beta} - n_{\alpha\alpha}), \qquad \Omega = \frac{8\,E_C}{\hbar}\,\frac{C_d\,V_0}{2e}. \tag{3.13}$$

The term proportional to $n_d(n_{\beta\beta} - n_{\alpha\alpha})$ describes a small, time-dependent shift of the qubit frequency (a dynamic Stark shift). Since this shift oscillates at the drive frequency or faster, we will neglect it under the rotating wave approximation (RWA) in what follows.

Rotating-wave approximation. To remove the strong time dependence at the drive frequency ω_d, we move into a frame rotating about the z-axis at ω_d. We define the unitary

$$\boldsymbol{U}(t) = e^{\,i(\omega_d t/2)\,\sigma_z},$$

under which the Hamiltonian transforms as

$$\boldsymbol{H} \;\mapsto\; \boldsymbol{H}' = \boldsymbol{U}\,\boldsymbol{H}\,\boldsymbol{U}^\dagger + i\hbar\,\frac{d\boldsymbol{U}}{dt}\,\boldsymbol{U}^\dagger.$$

Using $i\hbar\,\frac{d\boldsymbol{U}}{dt}\,\boldsymbol{U}^\dagger = -\frac{\hbar\,\omega_d}{2}\,\boldsymbol{\sigma}_z$, we find

$$\boldsymbol{U}\,\frac{\hbar\,\omega_{\alpha\beta}}{2}\,\boldsymbol{\sigma}_z\,\boldsymbol{U}^\dagger + i\hbar\,\frac{d\boldsymbol{U}}{dt}\,\boldsymbol{U}^\dagger = \frac{\hbar\,(\omega_{\alpha\beta} - \omega_d)}{2}\,\boldsymbol{\sigma}_z.$$

We next transform the coupling term,

$$\boldsymbol{H}_{\text{coupling}} = \hbar\,\Omega\cos\!\big(\omega_d t + \phi\big)\,\boldsymbol{\sigma}_x = \frac{\hbar\,\Omega}{2}\Big(e^{i(\omega_d t+\phi)} + e^{-i(\omega_d t+\phi)}\Big)\big(\boldsymbol{\sigma}_+ + \boldsymbol{\sigma}_-\big).$$

Using the identity $e^{i\theta\,\sigma_z}\,\boldsymbol{\sigma}_\pm\,e^{-i\theta\,\sigma_z} = e^{\pm 2i\theta}\,\boldsymbol{\sigma}_\pm$, and omitting the terms oscillating at $\pm 2\omega_d$, we obtain

$$\boldsymbol{U}\,\boldsymbol{H}_{\text{coupling}}\,\boldsymbol{U}^\dagger = \frac{\hbar\,\Omega}{2}\Big(\cos\phi\;\boldsymbol{\sigma}_x + \sin\phi\;\boldsymbol{\sigma}_y\Big).$$

Qubit drive Hamiltonian. Under the rotating wave approximation (RWA) we discard the rapidly oscillating components, leaving only the slowly varying drive terms. Putting these pieces together yields the final effective Hamiltonian in the rotating frame,

$$\boxed{\boldsymbol{H}_{\text{Drive}} = \frac{\hbar\,\delta\omega}{2}\,\boldsymbol{\sigma}_z + \frac{\hbar\,\Omega}{2}\big(\boldsymbol{\sigma}_x\cos\phi + \boldsymbol{\sigma}_y\sin\phi\big),} \tag{3.14}$$

where the detuning is defined as the difference between the qubit frequency and the drive frequency, $\delta\omega = \omega_{\alpha\beta} - \omega_d$. The Hamiltonian was discussed in Sec. 2.4.4, where we show that it leads to coherent qubit oscillations at the generalized Rabi frequency $\Omega_R = \sqrt{\delta\omega^2 + \Omega^2}$. Since $\Omega \propto V_0$, adjusting the drive amplitude V_0 directly controls the qubit rotation rate.

3.2.2 Inductive coupling: flux line

For a fluxonium qubit, precise control of the external flux φ_{ext} is essential. To achieve this, we inductively couple the fluxonium loop to a flux-bias line. A current I_f through the flux line threads the loop with a flux MI_f,

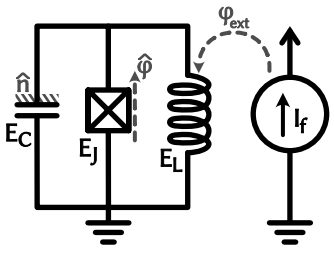

Fig. 3.8: Lumped-element equivalent circuit of the fluxonium qubit coupled inductively to a flux line.

where M is the mutual inductance. The fluxonium Hamiltonian then becomes:

$$\boldsymbol{H} = 4E_C\,\boldsymbol{n}^2 - E_J\cos\boldsymbol{\varphi} + \frac{E_L}{2}\Big(\boldsymbol{\varphi} + \underbrace{\frac{M\,I_\mathrm{f}}{\phi_0}}_{\varphi_\mathrm{ext}}\Big)^2. \tag{3.15}$$

Defining $\varphi_\mathrm{ext} = MI_\mathrm{f}/\phi_0$ makes clear that adjusting I_f directly controls the external flux bias. Practically, we calibrate M by measuring the current $I_\mathrm{f}^{2\pi}$ required to change φ_ext by 2π, yielding

$$M = \frac{\Phi_0}{I_\mathrm{f}^{2\pi}} = \frac{h}{2e\,I_\mathrm{f}^{2\pi}}. \tag{3.16}$$

3.2.3 Readout Dispersive Shift

To read out the fluxonium qubit, we capacitively couple it to a microwave resonator and infer its state from the resonator's dispersive frequency pull (see Sec. 2.5.3). We begin by deriving the full qubit–resonator Hamiltonian from the circuit Lagrangian. Next, we follow Sec. 3.2.1 of [Man12] to apply second-order perturbation theory and obtain the analytic expression for the dispersive shift $\chi_{\alpha\beta}$. Finally, we validate and illustrate these results through a numerical evaluation of χ_ge for representative heavy-fluxonium parameters.

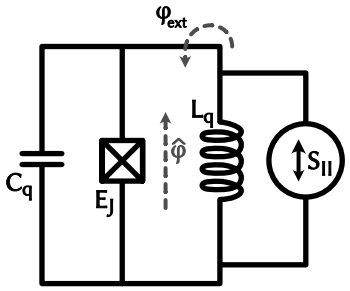

Fig. 3.9: Fluxonium-readout coupling circuit.

[Man12]: Manucharyan (2012), 'Superinductance', Thesis

Coupling Hamiltonian derivation. The fluxonium qubit and its readout resonator are coupled capacitively via the capacitor C_s (see Fig. 3.9). We introduce the generalized flux variables of the qubit (Φ_q) and readout (Φ_r) flux variable, as well as the capacitance matrix

$$C_M = \begin{pmatrix} C_q + C_s & -C_s \\ -C_s & C_r + C_s \end{pmatrix}, \quad \Phi = \begin{pmatrix} \Phi_q \\ \Phi_r \end{pmatrix}, \tag{3.17}$$

where C_q and C_r are, respectively, the bare qubit and resonator capacitances. The Lagrangian is then

$$\mathcal{L} = \underbrace{\frac{1}{2}\dot{\Phi}^T C_M \dot{\Phi}}_{\text{Kinetic energy}} - \underbrace{\frac{\Phi_r^2}{2L_r}}_{\text{Resonator potential}} - \underbrace{\frac{\Phi_q^2}{2L_q} + E_J\cos\left(\frac{\Phi_q - \Phi_\mathrm{ext}}{\phi_0}\right)}_{\text{Fluxonium potential}}. \tag{3.18}$$

The conjugate charge operators are defined by

$$Q^T = \frac{\partial\mathcal{L}}{\partial\dot{\Phi}} = \dot{\Phi}^T\,C_M, \qquad Q = \begin{pmatrix} Q_q \\ Q_r \end{pmatrix}.$$

Inverting, $\dot{\Phi} = C_M^{-1}\,Q$, where

$$C_M^{-1} = \frac{1}{C_\Sigma^2}\begin{pmatrix} C_r + C_s & C_s \\ C_s & C_q + C_s \end{pmatrix}, \quad C_\Sigma^2 \;\equiv\; C_rC_s + C_q\,(C_r + C_s).$$

The dressed mode frequencies follow from the quadratic part of $\mathcal{L}$:

$$\omega_q^2 = \frac{(C_M^{-1})_{00}}{L_q}, \qquad \omega_r^2 = \frac{(C_M^{-1})_{11}}{L_r},$$

and the corresponding mode impedances are

$$Z_q = \sqrt{L_q\,(C_M^{-1})_{00}}. \qquad Z_r = \sqrt{L_r\,(C_M^{-1})_{11}},$$

The off-diagonal entries of C_M^{-1} give the coupling coefficient

$$\tilde{g} = \frac{(C_M^{-1})_{01} + (C_M^{-1})_{10}}{4\sqrt{Z_r\,Z_q}} = \frac{C_s}{2\,C_\Sigma^2\,\sqrt{Z_r\,Z_q}}.$$

Performing the Legendre transform and promoting (Φ_i, Q_i) to operators $(\boldsymbol{\Phi}_i, \boldsymbol{Q}_i)$ with $[\boldsymbol{\Phi}_i, \boldsymbol{Q}_j] = i\hbar\,\delta_{ij}$ (with $i \in \{q, r\}$), one obtains

$$\begin{aligned} \boldsymbol{H} &= \frac{1}{2}\,L_R\,\omega_r^2\,\boldsymbol{Q}_R^2 + \frac{\boldsymbol{\Phi}_r^2}{2L_R} && \text{| Resonator} \\ &+ \frac{1}{2}\,L_J\,\omega_q^2\,\boldsymbol{Q}_q^2 + \frac{\boldsymbol{\Phi}_J^2}{2L_J} - E_J\cos\left(\frac{\boldsymbol{\Phi}_q - \Phi_{\text{ext}}}{\phi_0}\right) && \text{| Fluxonium} \\ &+ 2\,\tilde{g}\,\sqrt{Z_R\,Z_q}\,\boldsymbol{Q}_R\,\boldsymbol{Q}_q && \text{| Coupling} \end{aligned} \tag{3.19}$$

Rewriting the resonator in terms of its ladder operators $\boldsymbol{c}$ and $\boldsymbol{c}^\dagger$ (Eq. 2.23) and the fluxonium in its charge–phase basis (Eq. 2.35), the total Hamiltonian becomes

$$\boxed{\begin{aligned} \boldsymbol{H} &= \hbar\omega_r\,\boldsymbol{c}^\dagger\boldsymbol{c} && \text{| Resonator} \\ &+ 4\,E_C\,\boldsymbol{n}^2 + \frac{1}{2}\,E_L\,\boldsymbol{\varphi}^2 - E_J\cos(\varphi - \varphi_{\text{ext}}) && \text{| Fluxonium} \\ &+ i\,g\,\boldsymbol{n}\left(\boldsymbol{c}^\dagger - \boldsymbol{c}\right). && \text{| Coupling} \end{aligned}} \tag{3.20}$$

where the renormalized coupling is

$$g = 4e\,\tilde{g}\,\sqrt{\frac{\hbar\,Z_q}{2}}.$$

Dispersive shift. The interaction Hamiltonian between the fluxonium qubit and the readout resonator is

$$\boldsymbol{H}_{\text{int}} = i\,g\,\boldsymbol{n}\left(\boldsymbol{c}^\dagger - \boldsymbol{c}\right),$$

where g is small compared to the relevant detunings (the dispersive regime). We therefore treat $\boldsymbol{H}_{\text{int}}$ as a perturbation on the uncoupled eigenbasis $\{|\alpha, n\rangle\}$, with α the qubit state and n the cavity Fock state. The unperturbed energies are

$$E_{\alpha,n}^{(0)} = E_\alpha + n\,\hbar\omega_r,$$

and including corrections up to second order,

$$E_{\alpha,n} = E^{(0)}_{\alpha,n} + \delta E^{(1)}_{\alpha,n} + \delta E^{(2)}_{\alpha,n}.$$

The first-order correction vanishes,

$$\delta E^{(1)}_{\alpha,n} = \langle \alpha, n | \boldsymbol{H}_{\text{int}} | \alpha, n \rangle = 0,$$

since $\langle n|\boldsymbol{c}|n\rangle = \langle n|\boldsymbol{c}^\dagger|n\rangle = 0$. The second-order shift is

$$\delta E^{(2)}_{\alpha,n} = \sum_{\substack{\beta\neq\alpha \\ m\neq n}} \frac{\left|\langle \alpha, n|\boldsymbol{H}_{\text{int}}|\beta, m\rangle\right|^2}{E^{(0)}_{\alpha,n} - E^{(0)}_{\beta,m}} \tag{3.21}$$

Because $\boldsymbol{c} + \boldsymbol{c}^\dagger$ only connects neighboring Fock states, the sum over m reduces to nearest neighbors

$$\delta E^{(2)}_{\alpha,n} = g^2 \sum_{\beta\neq\alpha} \Big[\frac{n\,|n_{\alpha\beta}|^2}{E_\alpha - \left(E_\beta - \hbar\omega_r\right)} + \frac{(n+1)\,|n_{\alpha\beta}|^2}{E_\alpha - \left(E_\beta + \hbar\omega_r\right)} \Big],$$

where $n_{\alpha\beta}$ is the charge matrix element between the fluxonium states $|\alpha\rangle$ and $|\beta\rangle$. When the fluxonium is in state $|\alpha\rangle$, it shifts the bare cavity frequency ω_r by an amount χ_α. Equivalently, χ_α is the change in cavity frequency between n and $n+1$ photons in the cavity, given the qubit state $|\alpha\rangle$. Hence, introducing the fluxonium transition frequencies $\hbar\omega_{\alpha\beta} = E_\alpha - E_\beta$, the difference in cavity frequency between two fluxonium states, *dispersive shift*, is

$$\boxed{\begin{aligned} \chi_{\alpha\beta} &= \chi_\alpha - \chi_\beta, \\ \chi_\alpha &= g^2 \sum_{\beta\neq\alpha} \left|n_{\alpha\beta}\right|^2 \frac{2\omega_{\alpha\beta}}{\omega^2_{\alpha\beta} - \omega^2_r}. \end{aligned}} \tag{3.22}$$

This expression holds in the *dispersive regime*, defined by $\left|\omega_{\alpha\beta} - \omega_r\right| \gg g$, or equivalently $\chi_\alpha/g \ll 1$.

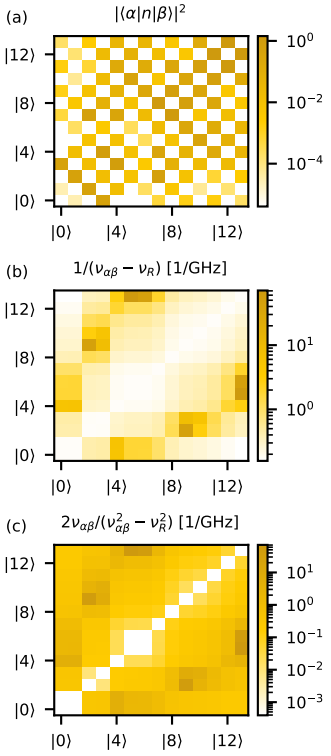

Fig. 3.10: Contributions to χ_α for a fluxonium with $E_J = 4.3\,\text{GHz}$, $E_C = 0.47\,\text{GHz}$, $E_L = 0.17\,\text{GHz}$, $\varphi_{\text{ext}} = \pi$, and $\omega_r = 5.8\,\text{GHz}$.

Numerical evaluation. Fig. 3.10 shows the matrix elements $n_{\alpha\beta}$ and their detunings from the resonator, $|\omega_{\alpha\beta} - \omega_r|$, for our fluxonium qubit coupled to a resonator at 5.8 GHz. These detunings typically exceed the coupling strength g (on the order of 100 MHz), validating the dispersive approximation. As noted in Sec. 3.1.3, transitions between states of the same parity are forbidden, producing the checkerboard-like pattern in Fig. 3.10a.

Although the fundamental $|g\rangle \to |e\rangle$ transition of our heavy fluxonium lies at a few megahertz, the dispersive frequency pull χ_{ge} remains strongly influenced by virtual excitations through higher-energy levels. To quantify this effect, we compute χ_{ge} using the standard sum-over-states expression (Eq. 3.22) while systematically increasing the number of fluxonium levels included. Fig. 3.11 shows the normalized dispersive shift as a function of the number of levels for a parameter sweep $E_J \in [3, 6]$ GHz and $E_C \in [0.3, 0.6]$ GHz. In all cases, convergence of χ_{ge} requires inclusion of at least ten excited states. Since analytic approximations for the high-lying

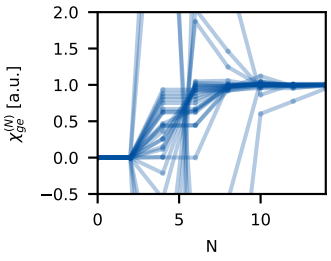

Fig. 3.11: Convergence of χ_{ge}/g^2 versus number of fluxonium levels N included in the sum-over-states expression (Eq. 3.22). Curves correspond to different (E_J, E_C) pairs, normalized so $\chi_{\text{ge}}/g^2 \to 1$ at large N.

matrix elements are not available, these results rely entirely on numerical evaluation.

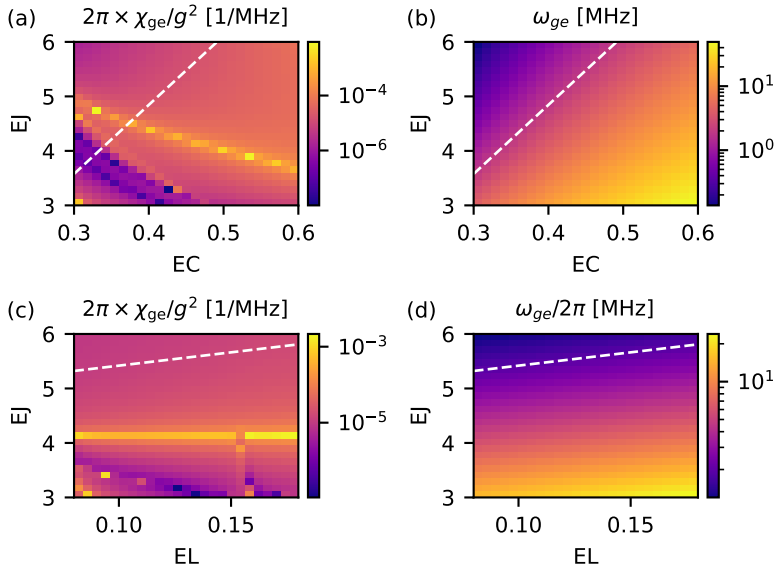

Fig. 3.12: Colormap of χ_{ge}/g^2 and ω_{ge} at $\varphi_{ext} = \pi$ for varying fluxonium parameters. The coupling strength g is expressed in units of rad s^{-1}. The dashed white line indicates a constant qubit frequency, $\omega_{ge}/2\pi = 4$ MHz.

Finally, in Fig. 3.12, we display the normalized dispersive shift χ_{ge}/g^2 at the flux sweet spot for a range of fluxonium parameters. The plot highlights parameter regions where transitions are close to resonance with the cavity, leading to a breakdown of the dispersive approximation, and regions where the dispersive shift remains stable. Outside these resonance zones, the shift remains nearly constant for a given transition frequency, enabling parameter tuning to focus primarily on achieving the desired frequency—except in cases where strong hybridization must be avoided.

3.3 Fluxonium Decoherence

Up to this point, we have treated our circuits as perfectly isolated, neglecting any interaction with the external world. In realistic superconducting circuits, however, both engineered couplings to control and readout apparatus and unavoidable interactions with environmental degrees of freedom introduce dissipation and dephasing.

In this chapter, we begin by formulating the dynamics of an open quantum system using the Lindblad master equation, which captures both energy relaxation and pure dephasing processes. We introduce the standard coherence metrics—the energy-relaxation time T_1 and the dephasing time T_2^*—which quantify a qubit's lifetime and phase memory, respectively, and establish the notation used throughout this thesis.

We then analyze the intrinsic decoherence mechanisms specific to the fluxonium qubit, detailing how dielectric loss, inductive loss, quasiparticle generation, and phase slips arise from material properties and circuit design.

3.3.1 Fermi's Golden Rule for Linear Interaction

Although we have discussed how to incorporate losses into the equations of motion (using jump operators; see Sec. 2.4.5), we have not yet introduced dissipative elements into the circuit itself. In the following, we derive the decoherence rate of a fluxonium qubit linearly coupled to its environment using the convenient form of Fermi's golden rule. The resulting "plug-and-play" expression will then be applied to compute decoherence rates for various loss channels.

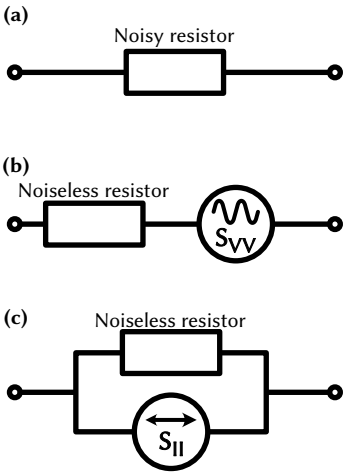

Fig. 3.13: Equivalent circuits for a noisy resistor. (a) A "noisy" resistor at a nonzero temperature with internal Johnson–Nyquist thermal noise. **(b)** An equivalent circuit representation consisting of a voltage source with spectral density $S_{VV}[\omega]$ in series with a noiseless resistor R (Thévenin equivalent circuit). **(c)** An equivalent circuit representation consisting of a current source with spectral density $S_{II}[\omega]$ in parallel with a noiseless resistor R (Norton equivalent circuit).

Fermi's Golden Rule for spectral noise density. Consider a superconducting circuit coupled to an external environment. The interaction is described by a coupling Hamiltonian, denoted $\boldsymbol{H}'$. For small perturbations, we can assume that the coupling Hamiltonian $\boldsymbol{H}'$ is linear in both the circuit and the environmental bath operators. The general form of such a coupling is then

$$\boldsymbol{H}' = g \times \boldsymbol{\Lambda}_{\text{env}} \times \boldsymbol{O}, \tag{3.23}$$

where g is a coupling constant, $\boldsymbol{O}$ is an operator associated with the superconducting circuit, and $\boldsymbol{\Lambda}_{\text{env}}$ is an operator associated with the environmental bath. For a bath in a mixed state, described by the density matrix $\sum_\mu p_\mu |\mu\rangle\langle\mu|$, Fermi's golden rule gives the transition rate between two circuit states, $|\alpha\rangle$ and $|\beta\rangle$, as:

$$\underline{\Gamma}_{\alpha\to\beta} = \frac{2\pi}{\hbar} \sum_{\mu,\nu} p_\mu \left\| \langle \alpha, \mu | \boldsymbol{H}' | \beta, \nu \rangle \right\|^2 \delta(E_{\alpha\mu} - E_{\beta\nu}) \tag{3.24}$$

Here, $|\alpha, \mu\rangle$ represents the joint state of the fluxonium in state $|\alpha\rangle$ and the bath in state $|\mu\rangle$. The total energy of this joint state is $E_{\alpha\mu}$. In the zeroth-order of perturbation theory, this energy is the sum of the individual energies of the circuit and the bath: $E_{\alpha\mu} = E_\alpha + E_\mu$. The Dirac delta function, $\delta(E)$, ensures the conservation of energy during the transition. By substituting the expression for the coupling Hamiltonian $\boldsymbol{H}'$ and using the relationship for the transition frequency, $\hbar\omega_{\alpha\beta} = E_\alpha - E_\beta$, we can rewrite the rate as:

$$\boxed{\underline{\Gamma}_{\alpha\to\beta} = \frac{g^2}{\hbar^2} \left\| \langle \alpha | \boldsymbol{O} | \beta \rangle \right\|^2 \underline{S}_{\Lambda\Lambda}[\omega_{\alpha\beta}]} \tag{3.25}$$

The term $\underline{S}_{\Lambda\Lambda}[\omega_{\alpha\beta}]$ is the spectral noise density of the bath operator $\boldsymbol{\Lambda}_{\text{env}}$ evaluated at the transition frequency. According to the Wiener-Khinchin theorem, it can be derived from the $\boldsymbol{\Lambda}_{\text{env}}$ autocorrelation function [Sch+02, Eq. 18][Cle14, Eq. 2.13], and is given by:

$$\underline{S}_{\Lambda\Lambda}[\omega_{\alpha\beta}] = 2\pi\hbar \sum_{\mu,\nu} p_\mu \left\| \langle \mu | \boldsymbol{\Lambda}_{\text{env}} | \nu \rangle \right\|^2 \delta(\hbar\omega_{\alpha\beta} - (E_\nu - E_\mu))$$

[Sch+02]: Schoelkopf et al. (2002), 'Qubits as Spectrometers of Quantum Noise', NATO Science Series

[Cle14]: Clerk (2014), 'Quantum noise and quantum measurement', Les Houches Lecture Notes

Intrinsic Losses Coupling Term. So far, we have assumed purely imaginary (nondissipative) impedances; losses arise when the real part of Z is nonzero, as in a resistor. At finite temperature T, a resistor exhibits thermal fluctuations known as Johnson–Nyquist noise. This noise can be

modeled in two equivalent ways [Cle+10, Fig. 16]: as a voltage source with spectral density $S_{VV}[\omega]$ in series with a noiseless resistor (Fig. 3.13(b)), or as a current source with spectral density $S_{II}[\omega]$ in parallel with a noiseless resistor (Fig. 3.13(c)).

To account for the intrinsic losses of the fluxonium qubit, such as those arising from a dielectric substrate at a specific bath temperature, we introduce a model where a noise current source, $\boldsymbol{I}_{\text{env}}$, is placed in parallel with the qubit circuit, as shown in Fig. 3.14. This source is characterized by a spectral density $S_{II}[\omega]$. In this framework, all dielectric losses are captured by assigning a real component to the impedance (or, equivalently, the admittance) of the circuit elements (as will be discussed in following subsections).

The environmental noise current, $\boldsymbol{I}_{\text{env}}$, induces a flux $\boldsymbol{\Phi}_{\text{env}} = L\boldsymbol{I}_{\text{env}}$ across the fluxonium inductance. Consequently, the inductive energy term transforms as:

$$\frac{E_L}{2}\boldsymbol{\varphi}^2 \quad \longrightarrow \quad \frac{E_L}{2}\left(\boldsymbol{\varphi} + \frac{\boldsymbol{\Phi}_{\text{env}}}{\phi_0}\right)^2$$

From this, we can extract the interaction Hamiltonian between the fluxonium and its environment:

$$H' = E_L \frac{L\boldsymbol{I}_{\text{env}}}{\phi_0}\boldsymbol{\varphi} = \underbrace{\phi_0}_{g} \times \underbrace{\boldsymbol{I}_{\text{env}}}_{\boldsymbol{\Lambda}_{\text{env}}} \times \underbrace{\boldsymbol{\varphi}}_{O}. \tag{3.26}$$

Intrinsic Losses Transition Rate. We identify the terms in analogy with the general form of a system-bath coupling Hamiltonian shown in Eq. 3.23. Applying Fermi's Golden Rule, as formulated in Eq. 3.25, yields the transition rate between states $|\alpha\rangle$ and $|\beta\rangle$:

$$\underline{\Gamma}^{\text{int}}_{\alpha\beta} = \frac{\phi_0^2}{\hbar^2}\left\|\langle\alpha|\boldsymbol{\varphi}|\beta\rangle\right\|^2 \underline{S}^{\text{env}}_{II}[\omega_{\alpha\beta}] = \frac{1}{(2e)^2}\left\|\langle\alpha|\boldsymbol{\varphi}|\beta\rangle\right\|^2 \underline{S}^{\text{env}}_{II}[\omega_{\alpha\beta}], \tag{3.27}$$

where $\underline{S}^{\text{env}}_{II}[\omega_{\alpha\beta}]$ is the two-sided quantum current noise spectral density of the operator $\boldsymbol{I}_{\text{env}}$. This current noise spectral density[1] is given by [Sch+02, eq. 28][2]:

$$\underline{S}^{\text{env}}_{II}[\omega] = \hbar\omega \text{Re}[Y_{\text{env}}(|\omega|)]\left(\coth\left(\frac{\hbar\omega}{2k_BT}\right) + 1\right), \tag{3.28}$$

where $\text{Re}[Y_{\text{env}}]$ is the real part of the admittance of the lossy element.

The total transition rate between two states, $\Gamma_{\alpha\beta}$, is the sum of the rates for absorption and emission. To simplify the final expression, we define the one-sided spectral density, $S^{\text{env}}_{II}[|\omega|]$, which combines contributions from positive and negative frequencies: [3]:

$$\begin{aligned} S^{\text{env}}_{II}[|\omega|] &= \underline{S}^{\text{env}}_{II}[|\omega|] + \underline{S}^{\text{env}}_{II}[-|\omega|] \\ &= 2\hbar|\omega|\ \text{Re}[Y_{\text{env}}(\omega)]\ \coth\left(\frac{\hbar|\omega|}{2k_BT}\right). \end{aligned} \tag{3.29}$$

[Cle+10]: Clerk et al. (2010), 'Introduction to Quantum Noise, Measurement and Amplification', Reviews of Modern Physics

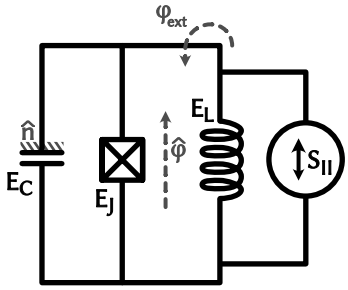

Fig. 3.14: Equivalent circuit for fluxonium intrinsic noises. The circuit consists of a standard fluxonium where every element can have a real part of its impedance, shunted by a current source $\boldsymbol{I}_{\text{env}}$ with spectral density $S_{II}[\omega]$.

Tab. 3.1: Impedances, $Z(\omega)$, and Admittances, $Y(\omega)$, of Basic Circuit Elements

Component	$Z(\omega)$	$Y(\omega)$
Resistor	R	$1/R$
Capacitor	$1/(i\omega C)$	$i\omega C$
Inductor	$i\omega L$	$1/(i\omega L)$

1: The voltage noise spectral density is related by:

$$\underline{S}^{\text{env}}_{VV}[\omega] = \text{Re}[Z_{\text{env}}(|\omega|)]^2 \underline{S}^{\text{env}}_{II}[\omega],$$

where $Z_{\text{env}} = 1/Y_{\text{env}}$ is the impedance of the environment.

[Sch+02]: Schoelkopf et al. (2002), 'Qubits as Spectrometers of Quantum Noise', NATO Science Series

2: A derivation of the spectral current noise density can be found in [Sch+02, Chap. 3] or [Wat99, Lect. 20]. However, here is a small hint to help remember it. The quantum noise from a resistor can be modeled using its quantum description as a bath of harmonic oscillators (Caldeira–Leggett model). The energy of the states around ω is given by the energy of a single state, $\hbar\omega$, multiplied by the density of states, which is $(\coth(\hbar\omega/2k_BT) + 1)/2$. The spectral density of the current must have units of amperes squared per hertz, equivalent to energy per ohm. This dimensional analysis leads to Eq. 3.28.

3: Note that

$$\coth\frac{\hbar|\omega|}{2k_BT} = 1 + 2n_{\text{BE}}(|\omega|),$$

where n_{BE} is the Bose-Einstein distribution.

This allows us to write the total transition rate in a compact form:

$$\boxed{\Gamma^{\text{int}}_{\alpha\beta} = \underline{\Gamma}^{\text{int}}_{\alpha\beta} + \underline{\Gamma}^{\text{int}}_{\beta\alpha} = \frac{1}{(2e)^2} \left\| \langle \alpha | \boldsymbol{\varphi} | \beta \rangle \right\|^2 S^{\text{env}}_{II}(|\omega_{\alpha\beta}|)} \tag{3.30}$$

Substituting specific environmental admittances $Y_{\text{env}}(\omega)$ into these expressions enables evaluation of various loss channels. For the following numerical estimates, we use the following fluxonium parameters:

$$E_J = 4.82 \text{ GHz}, \quad E_C = 0.418 \text{ GHz}, \quad E_L = 0.128 \text{ GHz},$$

which yield transition frequencies $\omega_{ge} \approx 2.1$ MHz, $\omega_{ef} \approx 3.56$ GHz and $\omega_{gh} \approx 3.62$ GHz.

3.3.2 Dielectric Losses in the Capacitor

Dielectric losses arise from the qubit's interaction via the electric field with the dielectric material on which it is fabricated. These losses can be modeled by a lossy capacitor, where the permittivity of the material is complex, $\epsilon = \epsilon_0 - i\epsilon_1$. The admittance of such a capacitor is given by:

$$Y_{\text{diel}} = i|\omega| \frac{\epsilon_0 - i\epsilon_1}{\epsilon_0} C = i|\omega|C + |\omega|C \tan\delta, \tag{3.31}$$

where C is the lossless capacitance of the qubit. The loss tangent, $\tan\delta = \epsilon_1/\epsilon_0$, characterizes the material's dielectric losses. It is often referred to as the inverse of the dielectric quality factor, $Q_{\text{diel}} = 1/\tan\delta$.

The real part of the admittance, which accounts for dielectric losses, is:

$$\text{Re}[Y_{\text{diel}}(\omega)] = |\omega|C \tan\delta. \tag{3.32}$$

Using the relationship $C = e^2/(2E_C)$, the decoherence rate due to dielectric losses is expressed as:

$$\Gamma^{\text{diel}}_{\alpha\beta} = \frac{\hbar\omega^2_{\alpha\beta}}{4E_C} \cdot \tan\delta \cdot |\langle \alpha | \boldsymbol{\varphi} | \beta \rangle|^2 \cdot \coth\left(\frac{\hbar|\omega_{\alpha\beta}|}{2k_B T}\right). \tag{3.33}$$

For an aluminium circuit fabricated on silicon, we estimate[1] $\tan\delta \approx 10^{-5}$, corresponding to $Q_{\text{diel}} = 10^5$. Note that E_C reflects the total capacitance of the qubit, which includes contributions from various interfaces, making $\tan\delta$ an effective average across different material interfaces.

We obtain the upper bounds on the decoherence rates due to dielectric losses, as illustrated in Fig. 3.15. At the flux frustration point, these bounds are given by:

$$\begin{aligned} T^{\text{diel}}_{ge} &\gtrsim 826\ \mu\text{s}, \\ T^{\text{diel}}_{ef,gh} &\gtrsim 4.34\ \mu\text{s}. \end{aligned} \tag{3.34}$$

Based on the measurement of the fluxonium decoherence (Chap. 7), we believe that dielectric losses are significant source of decoherence

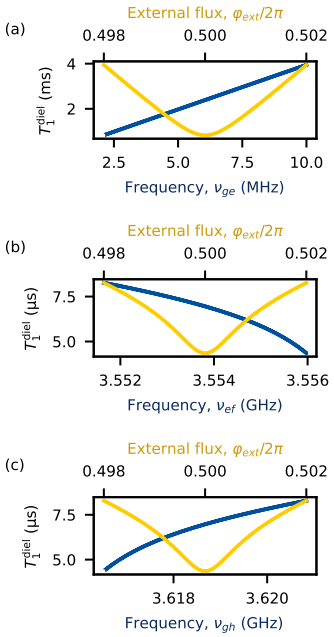

Fig. 3.15: Dielectric losses as a function of external flux, based on $\tan\delta \sim 10^{-5}$. φ_{ext}, in yellow, and as a function of the transition frequency, ν_{ij}, in blue. **(a, b, c)** show the losses for the transitions $|g\rangle \to |e\rangle$, $|e\rangle \to |f\rangle$, and $|g\rangle \to |h\rangle$, respectively.

1: The loss tangent $\tan\delta$ for certain interfaces, such as the metal-substrate interface, can be higher, with values around 3×10^{-4} [Mel+20]. However, most of the field is distributed across regions with much lower loss tangents, often on the order of 10^{-6}. For our estimation, we take a middle value, $\tan\delta \approx 10^{-5}$ ($Q_{\text{diel}} = 10^5$). More on this can be found in [Mel+20] and [Pop+14, Tab. 2].

for the transitions $|g\rangle \rightarrow |h\rangle$ and $|e\rangle \rightarrow |f\rangle$ (lower-to-upper manifold transitions).

3.3.3 Resistive Losses in the Inductor

Similarly, we can determine the losses induced by magnetic field interactions. These losses are modeled using the complex permeability of the inductance, $\mu = \mu_0 - i\mu_1$.

This leads to the admittance of the lossy inductor being:

$$Y_{\text{ind}} = \left(i|\omega|L + |\omega|L\tan\delta_\mu\right)^{-1}, \tag{3.35}$$

where L is the total inductance of the qubit, and the permeability loss tangent is defined as $\tan\delta_\mu = \mu_1/\mu_0$.

The real part of the admittance is:

$$\text{Re}[Y_{\text{ind}}(\omega)] = \frac{1}{|\omega|L}\cdot\frac{\tan\delta_\mu}{1+\frac{\tan^2\delta_\mu}{|\omega|^2L^2}} \simeq \frac{\tan\delta_\mu}{|\omega|L}. \tag{3.36}$$

The decoherence rate due to inductive losses is then:

$$\Gamma^{\text{ind}}_{\alpha\beta} = \frac{2E_L}{\hbar}\cdot\tan\delta_\mu\cdot|\langle\alpha|\boldsymbol{\varphi}|\beta\rangle|^2\cdot\coth\left(\frac{\hbar|\omega_{\alpha\beta}|}{2k_BT}\right). \tag{3.37}$$

The functional difference between inductive losses, $\text{Re}[Y_{\text{ind}}(\omega)]$, and capacitive losses, $\text{Re}[Y_{\text{diel}}(\omega)]$, lies in a factor of ω^2. Inductive losses increase more rapidly with lower frequencies, $\Gamma^{\text{ind}}_{\alpha\beta}/\Gamma^{\text{diel}}_{\alpha\beta} \propto 1/\omega^2$, making them more significant for low-frequency transitions.

The value of $\tan\delta_\mu$ is less commonly reported in the literature compared to the dielectric loss tangent $\tan\delta$. Based on available data, we estimate $\tan\delta_\mu \lesssim 2\times10^{-9}$ ($Q_{\text{ind}} = 5\times10^8$) [Pop+14, Tab. 2].

This results in the following upper bounds on the decoherence rates due to inductive losses, as illustrated in Fig. 3.16. At the flux frustration point, these bounds are given by:

$$\begin{aligned} T^{\text{ind}}_{ge} &\gtrsim 42\ \mu\text{s}, \\ T^{\text{ind}}_{gh,ef} &\gtrsim 652\ \text{ms}. \end{aligned} \tag{3.38}$$

Based on the measurement of the fluxonium decoherence (Chap. 7), we believe that the inductive losses are a significant source of decoherence of the lowest-frequency qubit transitions.

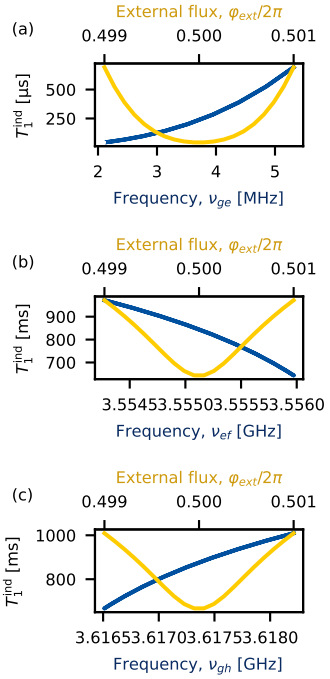

Fig. 3.16: Inductive losses as a function of external flux, based on $\tan\delta_\mu \sim 2\times10^{-9}$. φ_{ext}, in yellow, and as a function of the transition frequency, ν_{ij}, in blue. **(a, b, c)** show the losses for the transitions $|g\rangle \rightarrow |e\rangle$, $|e\rangle \rightarrow |f\rangle$, and $|g\rangle \rightarrow |h\rangle$, respectively.

[Pop+14]: Pop et al. (2014), 'Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles', Nature

3.3.4 Phase-Slip Losses in the Junction Array

Another source of decoherence arises from the fabrication approach of the superinductance, which consists of many Josephson junctions connected in series. The total phase across the array is the sum of the phases across all the junctions: $\varphi_a = \sum_i \varphi_i$. Due to the 2π periodicity of Josephson junctions, each junction can undergo a $\pm2\pi$ phase slip, referred to as

a quantum phase slip (QPS). Multiple phase slips occurring along the array lead to spatial interference effects caused by the Aharonov-Casher effect [AC84]. This phenomenon is referred to as coherent quantum phase slips (CQPS) [Man+12].

[AC84]: Aharonov et al. (1984), 'Topological Quantum Effects for Neutral Particles', Physical Review Letters

[Man+12]: Manucharyan et al. (2012), 'Evidence for coherent quantum phase-slips across a Josephson junction array', Physical Review B

[Ran+24]: Randeria et al. (2024), 'Dephasing in Fluxonium Qubits from Coherent Quantum Phase Slips', arXiv

The phase-slip energy ϵ_{ps} associated with quantum fluctuations of the phase across a single Josephson junction can be determined using the generalized WKB approximation [Man+12, ref. 24] [Ran+24, Eq. 1]:

$$\epsilon_{\text{ps}_i} = 2\sqrt{\frac{2}{\pi}}\sqrt{8E_{J_i}E_C}\sqrt[4]{\frac{8E_{J_i}}{E_C}}\exp\left(-\sqrt{\frac{8E_{J_i}}{E_C}}\right), \tag{3.39}$$

where E_{J_i} is the Josephson energy of a single junction i in the array. For a homogeneous array, E_{J_i} can be expressed as $E_{J_i} = E_L \times N_{\text{JJs}}$, where N_{JJs} is the total number of junctions in the superinductance array.

The total CQPS energy E_{CQPS} across the array is given by summing over the phase-slip energies of all junctions, including the Aharonov-Casher phase $\eta_{g,i}$:

$$E_{\text{CQPS}} = \sum_{i=1}^{N} \epsilon_{\text{ps},i} e^{-i2\pi\eta_{g,i}}. \tag{3.40}$$

For a homogeneous array, the CQPS energy can be approximated as $|\text{Re}[E_{\text{CQPS}}]| \leq N_{\text{JJs}}\epsilon_{\text{ps}}$. For large N_{JJs}, the Aharonov-Casher phases $\eta_{g,i}$ can be treated as random variables. By the central limit theorem, probability distribution of $\text{Re}[E_{\text{CQPS}}]$ is Gaussian with a standard deviation $\sigma_{\text{CQPS}} \approx \sqrt{N_{\text{JJs}}/2}\epsilon_{\text{ps}}$[Ran+24, Sec. II.D]. Thus, under quasi-static noise [Ith+05, Eq. 16], the dephasing rate is

$$\Gamma_\varphi^{\text{CQPS}} \approx \pi\sqrt{N_{\text{JJs}}}\epsilon_{\text{ps}}|\mathcal{F}_{\alpha\beta}| \lesssim \sqrt{2}\pi N_{\text{JJs}}|\mathcal{F}_{\alpha\beta}|, \tag{3.41}$$

where $\mathcal{F}_{\alpha\beta}$, the "structure factor" associated with the CQPS process, is given by:

$$\begin{aligned}\mathcal{F}_{\alpha\beta} &\approx \langle\psi_\beta|e^{-i2\pi\boldsymbol{n}}|\psi_\beta\rangle - \langle\psi_\alpha|e^{-i2\pi\boldsymbol{n}}|\psi_\alpha\rangle \\ &\approx \int_{-\infty}^{\infty} d\varphi\, \psi_\beta(\varphi)\psi_\beta(\varphi - 2\pi) - \int_{-\infty}^{\infty} d\varphi\, \psi_\alpha(\varphi)\psi_\alpha(\varphi - 2\pi),\end{aligned} \tag{3.42}$$

where $\psi_\beta(\varphi)$ represents the eigenstate β of the fluxonium Hamiltonian in phase space (Eq. 3.1).

We obtained that the dephasing rate depends exponentially on the ratio $\sqrt{E_{J_1}/E_C}$ (Eq. 3.39). Using the plasma frequency $E_p = \sqrt{8\,E_{J_1}\,E_C}$, one may write $\sqrt{E_{J_1}/E_C} \propto E_{J_1}/E_p$. Since $E_{J_1} \propto A$ and $E_C \propto 1/A$, the plasma frequency E_p is independent of the junction area A, while the ratio E_{J_1}/E_p scales linearly with A. For our fabrication parameters (an array of $N_{\text{JJs}} = 360$ junctions of area $A = 0.67\ \mu\text{m}^2$ and plasma frequency $E_p/h = 13.9$ GHz) we find $E_{J_1}/E_p \simeq 3.3$, yielding

$$T_\varphi^{\text{CQPS}} \approx 2.4\ \text{s}. \tag{3.43}$$

Thus CQPS is not a limiting factor for our qubit. However, due to the exponential dependence on $\sqrt{E_{J_1}/E_C}$, even a modest reduction of the junction area can bring T_φ^{CQPS} into the microsecond regime. For instance,

choosing $A = 0.25\ \mu\text{m}^2$ gives $E_{J_1}/E_p \simeq 1$ and, numerically, $T_\varphi^{\text{CQPS}} \approx 8\ \mu\text{s}$, which would be a significant limiting factor. The small junction area was the main issue in one of the first devices we fabricated.

Another crucial aspect is that fabrication inhomogeneity may produce a single junction smaller than the target size, which then acts as a "weak link" through which fluxons preferentially tunnel, dominating the overall dephasing rate. When, after fabrication, we observe coherence times much shorter than expected, junction non-uniformity becomes a primary suspect. To investigate this, we typically perform scans using an atomic-force microscope (AFM) or an electron microscope. An example of such an AFM scan is shown later in Fig. 6.6.

3.3.5 Quasiparticle Losses

Quasiparticle population. Another source of decoherence arises from quasiparticle tunneling across the small Josephson junction. The normalized quasiparticle density, relative to the density of Cooper pairs in thermal equilibrium, can be expressed as [Cat+11, Eq. 19]:

$$x_{\text{qp}} = \sqrt{2\pi k_B T/\Delta}\, e^{-\Delta/k_B T}, \tag{3.44}$$

where Δ is the superconducting gap, with $\Delta/k_B = 2.1\ \text{K}$ for aluminium. At cryogenic temperatures, this population is negligibly low, approximately 10^{-19} at 50 mK. However, experimental observations consistently show that the population exceeds the thermal equilibrium value, typically in the range of $(10^{-7}, 10^{-5})$ [Ste21, p. 42][Pop+14, p. 3]. This discrepancy is attributed to factors such as cosmic radiation, radioactive decay, and insufficient filtering and shielding in experimental setups.

Quasiparticle-induced losses in the junction. Quasiparticle-induced losses in the Josephson junction are nonlinear. Using the quasiparticle Hamiltonian and Fermi's Golden Rule, the decoherence rate can be approximated as [Cat+11, Eq. 32][1]:

$$\Gamma_{\alpha\beta}^{\text{qp}} = \left|\langle\alpha|\sin\left(\frac{\boldsymbol{\varphi} - \varphi_{\text{ext}}}{2}\right)|\beta\rangle\right|^2 S^{\text{qp}}(\omega_{\alpha\beta}), \tag{3.45}$$

where $S^{\text{qp}}[\omega] = \underline{S}^{\text{qp}}[\omega] + \underline{S}^{\text{qp}}[-\omega]$ is the single-sided spectral density of quasiparticles.

The spectral density $\underline{S}^{\text{qp}}$, at low temperatures $k_B T \ll \Delta$, is given by [Cat+11, Eq. 35]:

$$\underline{S}^{\text{qp}}[\omega] = \frac{16 E_J}{\pi\hbar}\exp\left(-\frac{\Delta}{k_B T}\right)\exp\left(-\frac{\hbar\omega}{2k_B T}\right)K_0\left(\frac{\hbar|\omega|}{2k_B T}\right), \tag{3.46}$$

where $K_0(x)$ is the modified Bessel function of the second kind with asymptotic forms:

$$K_0(x) \simeq \begin{cases} -\ln(x/2) - \gamma_E, & x \ll 1, \\ e^{-x}\sqrt{\pi/2x}, & x \gg 1, \end{cases} \tag{3.47}$$

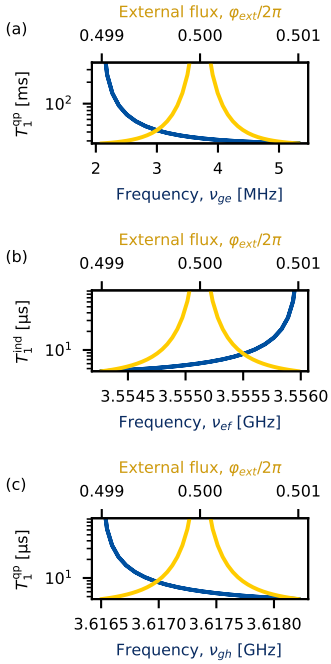

Fig. 3.17: Quasiparticle-induced losses as a function of external flux. φ_{ext} (yellow) and transition frequency ν_{ij} (blue). **(a), (b), (c)** show losses for the transitions $|g\rangle \to |e\rangle$, $|e\rangle \to |f\rangle$, and $|g\rangle \to |h\rangle$, respectively.

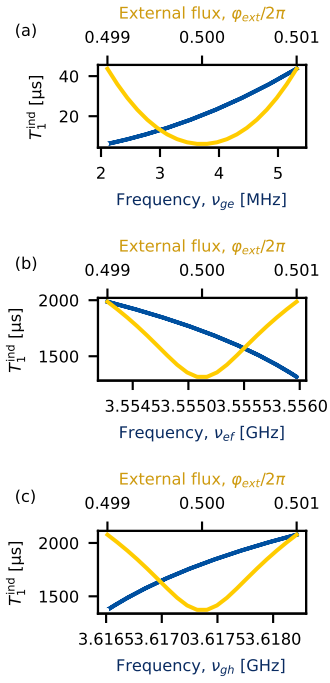

Fig. 3.18: Quasiparticle tunneling losses through the junction array. φ_{ext} (yellow) and transition frequency ν_{ij} (blue). **(a), (b), (c)** show losses for transitions $|g\rangle \to |e\rangle$, $|e\rangle \to |f\rangle$, and $|g\rangle \to |h\rangle$, respectively.

and $\gamma_E = 0.577$ is the Euler-Mascheroni constant.

The single-sided spectral noise density $S^{\text{qp}}(\omega)$ is then given by:

$$S^{\text{qp}}[\omega] = \frac{16E_J}{\pi\hbar} \exp\left(-\frac{\Delta}{k_B T}\right) K_0\left(\frac{\hbar|\omega|}{2k_B T}\right) \left(e^{-\frac{\hbar\omega}{2k_B T}} + e^{\frac{\hbar\omega}{2k_B T}}\right). \tag{3.48}$$

A more typical form of the spectral density emerges in the high-frequency limit, where $\hbar\omega \gg k_B T$. Under this condition, the modified Bessel function can be approximated as $K_0(\hbar\omega/2k_B T) \simeq \sqrt{\frac{\pi k_B T}{\hbar\omega}} e^{-\hbar\omega/2k_B T}$. Substituting this approximation, the spectral density simplifies to:

$$\begin{aligned} S^{\text{qp}}_{\omega\gg T}[\omega] &= \frac{16E_J}{\sqrt{\pi}\hbar} \exp\left(-\frac{\Delta}{k_B T}\right) \sqrt{\frac{k_B T}{\hbar\omega}} \\ &= \frac{8E_J}{\pi\hbar} x_{\text{qp}} \sqrt{\frac{2\Delta}{\hbar\omega}}. \end{aligned} \tag{3.49}$$

While the high-frequency assumption holds for typical transmon qubits and higher transitions in fluxonium qubits, it does not apply to the lowest transition of our heavy fluxonium qubit, where $\hbar\omega \ll k_B T$. In this low-frequency limit, the modified Bessel function can be approximated as $K_0(\hbar\omega/2k_B T) \simeq \ln(4k_B T/\hbar\omega) - \gamma_E$. Substituting this approximation, the spectral density becomes:

$$S^{\text{qp}}_{\omega\ll T}[\omega] = \frac{32E_J}{\pi\hbar} \exp\left(-\frac{\Delta}{k_B T}\right) \left(\ln\left(\frac{4k_B T}{\hbar\omega}\right) - \gamma_E\right), \tag{3.50}$$

Notably, at $\varphi_{\text{ext}} = \pi$, the matrix element vanishes, resulting in a zero decoherence rate. However, the decoherence rate acquires finite values near the π-flux point. Therefore, we evaluate it at $\varphi_{\text{ext}} = 0.999\pi$[1]. Furthermore, based on prior observations of superconducting qubits, we assume the quasiparticle population is higher than the equilibrium prediction, with $x_{\text{qp}} \sim 10^{-5}$. This corresponds to an effective temperature of $T^{\text{qp}}_{\text{eff}} \sim 187$ mK. Under these conditions, the decoherence rates due to quasiparticle tunneling through a single-standing junction at $\varphi_{\text{ext}} = 0.999\pi$, as shown in Fig. 3.17, are given by:

$$\begin{aligned} T^{\text{qp}}_{ge} &\gtrsim 80 \text{ ms}, \\ T^{\text{qp}}_{gh,ef} &\gtrsim 16.3\ \mu\text{s}. \end{aligned} \tag{3.51}$$

Quasiparticle-induced losses in the superinductance array. Quasiparticles can also induce losses due to tunneling through the junctions in the superinductance array. In this case, the total decoherence rate is the sum over all junctions[Cat+11, p. 13] and can be simplified under the assumption of small φ/N. Assuming identical junctions, each with energy $E_L \times N_{\text{JJA}}$, the final expression for the decoherence rate is:

$$\overline{\Gamma}^{\text{qp}}_{ij} = \left|\langle i|\frac{\boldsymbol{\varphi}}{2}|j\rangle\right|^2 S^{\text{qp}}_{JJ}(\omega_{ij}), \tag{3.52}$$

where the spectral density S^{qp}_{JJ} is the same as for a single junction (Eq. 3.48, Eq. 3.49, Eq. 3.50), but with E_J replaced by E_L [2].

[Cat+11]: Catelani et al. (2011), 'Relaxation and frequency shifts induced by quasiparticles in superconducting qubits', Physical Review B

[Ste21]: Stevens (2021), 'Effect of the Environment on Fluxonium Qubits and Thermodynamics of Quantum Measurement', Thesis

[Pop+14]: Pop et al. (2014), 'Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles', Nature

1: Compared to the reference, an additional term, φ_{ext}, is included inside the sine function based on the definition of the fluxonium Hamiltonian in Eq. 3.1.

1: This value is chosen based on typical experimental flux variations during measurements.

[Cat+11]: Catelani et al. (2011), 'Relaxation and frequency shifts induced by quasiparticles in superconducting qubits', Physical Review B

2: Not $E_L \times N_{\text{JJA}}$, as the $1/N_{\text{JJA}}$ factor is already included.

Unlike $\sin((\boldsymbol{\varphi} - \varphi_{\text{ext}})/2)$, the $\boldsymbol{\varphi}/2$ matrix element does not vanish at half flux, making it a potential limiting factor at $\varphi_{\text{ext}} = \pi$. For the chosen parameters, at $\varphi_{\text{ext}} = \pi$ and assuming $x_{\text{qp}} \sim 10^{-6}$ ($T_{\text{eff}}^{\text{qp}} \sim 156$ mK), the calculated decoherence rates are:

$$\begin{aligned} T_{ge}^{\text{qp in JJs}} &\gtrsim 172\ \mu\text{s}, \\ T_{gh,ef}^{\text{qp in JJs}} &\gtrsim 30\ \text{ms}. \end{aligned} \tag{3.53}$$

Conclusion. We conclude that quasiparticle losses, both in the small junction and in the superinductance array, do not constitute a significant source of decoherence for our qubit. However, as we will see in Chap. 7, the decoherence of the lower-to-upper manifold transitions is slightly affected by what we attribute to quasiparticle losses. Although we have not performed a detailed analysis to determine the exact quasiparticle population, a rough estimate suggests it is on the order of 10^{-5} to 5×10^{-5}.

3.3.6 Losses Due to Capacitive Coupling to a Lossy Channel

Introducing additional coupling to the environment results in increased decoherence of the qubit. To derive the decoherence rate due to capacitive coupling to the control and readout lines, we apply Fermi's Golden Rule Eq. 3.25 and use the interaction Hamiltonian from Eq. 3.11:

$$H' = \underbrace{8E_C \frac{C_d}{2e}}_{\alpha} \cdot \underbrace{V_d}_{\Lambda_{\text{env}}} \cdot \underbrace{\boldsymbol{n}}_{O}. \tag{3.54}$$

This leads to the following expression for the decoherence rate:

$$\begin{aligned} \overline{\Gamma}_{ij} &= \Gamma_{ij} + \Gamma_{ji} = \\ &= \pi \frac{C_d^2}{C_\Sigma^2} \frac{R}{R_{\text{Q}}^{\text{CP}}} \omega_{ij} \coth \frac{\hbar\omega_{ij}}{2kT} \left|\langle i|\boldsymbol{n}|j\rangle\right|^2 = \\ &= \frac{C_d^2}{4\pi} \cdot R R_{\text{Q}}^{\text{CP}} \cdot \omega_{ij}^3 \coth \frac{\hbar\omega_{ij}}{2kT} \left|\langle i|\boldsymbol{\varphi}|j\rangle\right|^2, \end{aligned} \tag{3.55}$$

where $R = 50\ \Omega$ is the effective impedance of the microwave line as seen by the qubit through the coupling capacitor C_d, $R_{\text{Q}}^{\text{CP}} = h/(2e)^2$ is the resistance quantum of a Cooper pair, and $C_\Sigma = e^2/2E_C$ is the total capacitance of the qubit.

Numerical evaluations of these losses, assuming a typical drive line coupling of $C_d = 1$ fF near the flux frustration point, are shown in Fig. 3.19 for various qubit transitions. At $\varphi_{\text{ext}} = \pi$, the resulting relaxation times are:

$$\begin{aligned} T_{ge} &\gtrsim 60\ \text{s}, \\ T_{gh,ef} &\gtrsim 1.6\ \text{ms}, \end{aligned} \tag{3.56}$$

which are far too long to act as a limiting factor.

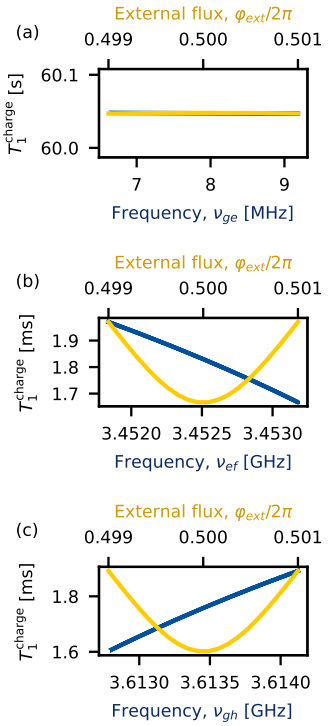

Fig. 3.19: Charge line induced decoherence. Losses due to qubit coupled capacitively to the environment via charge drive line as a function of external flux, φ_{ext}, in yellow, and as a function of the transition frequency, ν_{ij}, in blue. Panels **(a)**, **(b)**, and **(c)** show the losses for the transitions $|g\rangle \rightarrow |e\rangle$, $|e\rangle \rightarrow |f\rangle$, and $|g\rangle \rightarrow |h\rangle$, respectively.

3.3.7 Losses Due to Inductive Coupling to a Lossy Channel

The decoherence rate due to inductive coupling to the flux control line, can be derived using Fermi's Golden Rule Eq. 3.25 and the interaction Hamiltonian from Eq. 3.15:

$$H' = \underbrace{\frac{E_L}{\phi_0} M}_{g} \cdot \underbrace{I_f}_{\Lambda_{\text{env}}} \cdot \underbrace{\boldsymbol{\varphi}}_{O} . \tag{3.57}$$

This gives the following expression for the decoherence rate:

$$\Gamma_{\alpha\beta} = 4\pi \frac{R_Q^{\text{CP}}}{R} \frac{M^2}{L^2} \omega_{\alpha\beta} \coth \frac{\hbar\omega_{\alpha\beta}}{2kT} \left|\langle\alpha|\boldsymbol{\varphi}|\beta\rangle\right|^2 , \tag{3.58}$$

where $R = 50\ \Omega$ is the effective impedance of the microwave line as seen by the qubit through the inductive mutual coupling M, $R_Q^{\text{CP}} = h/(2e)^2$ is the resistance quantum of a Cooper pair, and $L = \phi_0^2/E_L$.

Numerical evaluations of these losses, assuming a typical flux line coupling of $M = \Phi_0/1.6\ \text{mA} = 1.3\ \text{pH}$ near the flux frustration point, are shown in Fig. 3.20 for various qubit transitions. At $\varphi_{\text{ext}} = \pi$, the resulting relaxation times are:

$$\begin{aligned} T_{ge} &\sim 22\ \text{ms}, \\ T_{gh,ef} &\sim 170\ \text{ms}, \end{aligned} \tag{3.59}$$

which are far too long to act as a limiting factor.

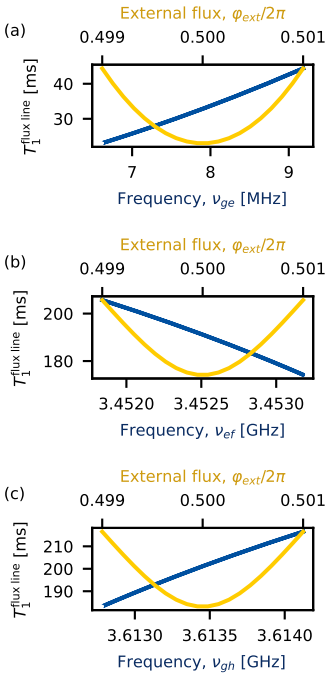

Fig. 3.20: Flux line induced decoherence. Losses due to the qubit's inductive coupling to the environment via the flux line, as a function of external flux φ_{ext} (yellow) and transition frequency ν_{ij} (blue). Panels **(a)**, **(b)**, and **(c)** show the losses for transitions $|g\rangle \to |e\rangle$, $|e\rangle \to |f\rangle$, and $|g\rangle \to |h\rangle$, respectively.

3.3.8 Purcell Losses

Additional coupling of the qubit to its environment via the readout cavity introduces an extra relaxation channel known as the Purcell effect. Originally identified by Purcell in spin–cavity systems [PTP46], this relaxation occurs through virtual resonator excitation: the qubit decays into the external continuum seen through the resonator filter, without significant real energy accumulation in the cavity. In circuit QED, the Purcell decay rate, governed by the resonator linewidth κ and the qubit–readout coupling g, has been derived by several complementary approaches: master-equation (GKSL) and dressed-state analyses for two-level systems [SGK14, Sec. II][Bla+04, Sec. II]; Fermi's Golden Rule and perturbation theory for transmon qubits [Koc+07, Sec. IV.B]; spectral-density methods [Bla+04, Sec. V][Hou+08, Fig. 1]; and Fermi's Golden Rule combined with numerical Hamiltonian diagonalization [Zha+21, App. G.5][Gro+18, App. C]. Here, we summarize the key result without reproducing these derivations.

In the dispersive regime $|\Delta| \equiv |\omega_q - \omega_R| \gg g, \kappa$, the zero-temperature decay rate is [SGK14, Eq. 13]:

$$\Gamma_P = \frac{\kappa\, g^2}{\Delta^2 + (\kappa/2)^2} \approx \frac{g^2}{\Delta^2}\kappa \quad (\Delta \gg \kappa), \tag{3.60}$$

[PTP46]: Purcell et al. (1946), 'Resonance Absorption by Nuclear Magnetic Moments in a Solid', Physical Review

[SGK14]: Sete et al. (2014), 'Purcell effect with microwave drive: Suppression of qubit relaxation rate', Physical Review B

[Bla+04]: Blais et al. (2004), 'Cavity quantum electrodynamics for superconducting electrical circuits: an architecture for quantum computation', Physical Review A

[Koc+07]: Koch et al. (2007), 'Charge insensitive qubit design derived from the Cooper pair box', Physical Review A

[Hou+08]: Houck et al. (2008), 'Controlling the Spontaneous Emission of a Superconducting Transmon Qubit', Physical Review Letters

[Zha+21]: Zhang et al. (2021), 'Universal fast flux control of a coherent, low-frequency qubit', Physical Review X

[Gro+18]: Groszkowski et al. (2018), 'Coherence properties of the 0-π qubit', New Journal of Physics

where g is the qubit–cavity coupling.

One must choose g, κ, and detuning $\Delta = \omega_q - \omega_R$ to balance readout speed and qubit protection. In the dispersive regime, the qubit induces a frequency shift

$$\chi \propto \frac{g^2}{\Delta},$$

so stronger coupling g enhances readout contrast and helps overcome the noise in the measurement line. However, larger g also increases the bare Purcell decay. One can suppress Γ_P by reducing the cavity linewidth κ, but smaller κ narrows the measurement bandwidth and prolongs the integration time, which must remain much shorter than the qubit's coherence time. While increasing Δ also reduces Purcell losses, Δ is typically already several gigahertz and cannot be made arbitrarily large. The typical coupling strength in our system is $g \sim 50$ MHz, while the cavity losses are $\kappa \sim 30$ MHz. Nevertheless, The most effective remedy is a *Purcell filter*: a bandpass network that passes the cavity frequency ω_R with minimal loss while strongly attenuating at the qubit frequency ω_q, thus preserving both readout performance and qubit lifetime. We discuss such filter designs in the fluxonium design chapter (Chap. 5).

3.3.9 Conclusion

We suspect that material losses constitute the primary source of decoherence for the fluxonium qubits used in our experiments. Unfortunately, a detailed analysis of decoherence was only perfomed on the qubits which coherence times approach the sub-microsecond scale, as then it becomes limiting factor for our experiments. In this case we found that decoherence is typically induced by phase slips due to asymmetry in the junction array or by defects in the substrate surrounding the main fluxonium elements.

Square Membrane Resonators | 4

Mechanical resonator. In this work, we employ highly stressed stoichiometric silicon nitride (Si_3N_4, or simplify "SiN") membranes as our mechanical resonators, owing to their exceptionally high quality factors [Tsa+17; Gha+18; Cup+24; Ber+22] made possible by dissipation dilution [Fed+19]. SiN is largely transparent to both optical and microwave fields, and can be fabricated in a variety of resonator geometries—squares [YCS15], circles [Ser+18], nanostrings [Fau+12; Kum+21], whispering-gallery modes [Bak+12], trampolines [Rei+16], frame membranes [Sch+16], and more. Patterned resonators such as 1D phononic crystals [Dav+13][Gha+18] or 2D phononic crystals [Tsa+17; Iva21][Ree+19] further suppress losses.

My contribution. This PhD was started with an ambitious goal of coupling phononic crystal membranes to superconducting qubits, which given the high mechanical quality factors of these structures, would allow to create long-lived membrane cat state when the strong coupling regime is achieved. In the first year of my thesis, together with an already former PhD student H. Patange, I've spend most of my time on developing phononic-crystal membranes: design, fabrication, and characterization. However, after several unsuccessful attempts to achieve a flip-chip, we understood that starting directly with phononic crystal membranes is too ambitious, so we have decided to switch to a simpler square membranes. During this project, these membranes were fabricated and primarily measured by H. Patange; it is discussed extensively in his thesis [Pat25]. Despite the time spent, I will not talk about phononic-crystal membranes in this thesis. However in this chapter I will briefly summarize the theoretical background of square membranes needed for better understanding our further decisions and overall workflow of this project.

Outline. This chapter starts with the theoretical background of square membrane resonators in Sec. 4.1, including equation of motion under uniform tension, and analytical mode shapes and frequencies; we define and compute the effective mass and the zero-point fluctuations for each mode; we then discuss the energy dissipation mechanisms and the mechanical quality factor, detailing intrinsic (material) losses, bending (anchor) losses, and radiation (acoustic) losses; We then proceed to discuss the specific design choices (Sec. 4.2) made for our SiN square membranes; we perform FEM simulation of the mechanical modes (frequencies and losses) to validate the design choices. Finally, we present the room-temperature characterization with an optical interferometer, to validate the design and fabrication.

[Tsa+17]: Tsaturyan et al. (2017), 'Ultra-coherent nanomechanical resonators via soft clamping and dissipation dilution', Nature Nanotechnology

[Gha+18]: Ghadimi et al. (2018), 'Elastic strain engineering for ultralow mechanical dissipation', Science

[Cup+24]: Cupertino et al. (2024), 'Centimeter-scale nanomechanical resonators with low dissipation', Nature Communications

[Ber+22]: Bereyhi et al. (2022), 'Perimeter Modes of Nanomechanical Resonators Exhibit Quality Factors Exceeding 10 9 at Room Temperature', Physical Review X

[Fed+19]: Fedorov et al. (2019), 'Generalized dissipation dilution in strained mechanical resonators', Physical Review B

[YCS15]: Yuan et al. (2015), 'Silicon nitride membrane resonators at millikelvin temperatures with quality factors exceeding 108', Applied Physics Letters

[Ser+18]: Serra et al. (2018), 'Silicon Nitride MOMS Oscillator for Room Temperature Quantum Optomechanics', Journal of Microelectromechanical Systems

[Fau+12]: Faust et al. (2012), 'Microwave cavity-enhanced transduction for plug and play nanomechanics at room temperature', Nature Communications

[Kum+21]: Kumar et al. (2021), *Microwave optomechanical measurement of non-metallized SiN strings at mK temperatures*, Preprint

[Bak+12]: Baker et al. (2012), 'Optical instability and self-pulsing in silicon nitride whispering gallery resonators', Optics Express

[Rei+16]: Reinhardt et al. (2016), 'Ultralow-Noise SiN Trampoline Resonators for Sensing and Optomechanics', Physical Review X

[Sch+16]: Schwarz et al. (2016), 'Deviation from the Normal Mode Expansion in a Coupled Graphene-Nanomechanical System', Physical Review Applied

[Dav+13]: Davanço et al. (2013), 'Si3N4 nanobeam optomechanical crystals', Optica Publishing Group

[Ree+19]: Reetz et al. (2019), 'Analysis of Membrane Phononic Crystals with Wide Band Gaps and Low-Mass Defects', Physical Review Applied

[Pat25]: Patange (2025), 'RF Transduction in a DC-Biased Superconducting Electromechanical System', Thesis

4.1 Mode Theory of a Rectangular SiN Membrane

We consider a thin silicon-nitride (SiN) membrane of in-plane dimensions l_y (along y) and l_z (along z), thickness $h \ll l_y, l_z$, mass density ρ, and uniform in-plane tensile stress σ. The membrane is clamped along all four edges at $y = \pm l_y/2$ and $z = \pm l_z/2$. Let's denote the out-of-plane displacement by $u(y, z, t)$. This section derives the Lagrangian, the equation of motion, the analytic mode shapes and frequencies for a purely tensioned (pinned) membrane, then incorporates finite bending rigidity, which is the main source of mechanical loss in this system. Finally we compute the expected quality factor Q of the membrane.

4.1.1 Lagrangian and Equation of Motion

A differential element of area $(\mathrm{d}y\,\mathrm{d}z)$ has mass $(\rho\, h\,\mathrm{d}y\,\mathrm{d}z)$, where ρ is the volumetric mass density and h is the thickness. In the tension-dominated regime (neglecting bending rigidity), the kinetic energy surface density is

$$\mathcal{T}(y, z, t) = \frac{1}{2}\,\rho\, h\,\left(\partial_t u\right)^2,$$

and the potential (elastic) energy surface density due to tension is

$$\mathcal{V}(y, z, t) = \frac{\sigma}{2}\left[(\partial_y u)^2 + (\partial_z u)^2\right],$$

where σ is the in-plane tensile stress. Hence the Lagrangian density is

$$\mathcal{L} = \mathcal{T} - \mathcal{V} = \frac{\rho\, h}{2}\,(\partial_t u)^2 - \frac{\sigma}{2}\left[(\partial_y u)^2 + (\partial_z u)^2\right]. \tag{4.1}$$

Integrating over the rectangular domain $\Omega = [-\frac{l_y}{2}, \frac{l_y}{2}]\times[-\frac{l_z}{2}, \frac{l_z}{2}]$ yields

$$\begin{aligned} L(u) &= \iint_\Omega \mathcal{L}(y, z, t)\,\mathrm{d}y\,\mathrm{d}z \\ &= \frac{\rho\, h}{2}\iint_\Omega (\partial_t u)^2\,\mathrm{d}y\,\mathrm{d}z - \frac{\sigma}{2}\iint_\Omega \left[(\partial_y u)^2 + (\partial_z u)^2\right]\mathrm{d}y\,\mathrm{d}z. \end{aligned} \tag{4.2}$$

Applying the Euler–Lagrange equation

$$\frac{\partial}{\partial t}\left(\rho\, h\,\partial_t u\right) - \sigma\,(\partial_{yy} + \partial_{zz})\,u = 0 \tag{4.3}$$

leads to the wave equation

$$\rho\, h\,\partial_{tt} u(y, z, t) = \sigma\left(\partial_{yy} + \partial_{zz}\right) u(y, z, t), \tag{4.4}$$

with clamped (Dirichlet) boundary conditions

$$u\left(y = \pm\tfrac{l_y}{2},\, z,\, t\right) = u\left(y,\, z = \pm\tfrac{l_z}{2},\, t\right) = 0\,, \quad \forall t. \tag{4.5}$$

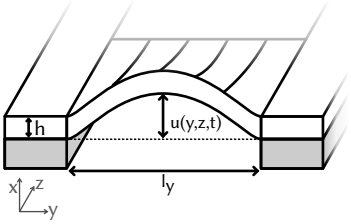

Fig. 4.1: Rectangular membrane side view illustration with the variables used in the derivation.

Define the wave speed $c = \sqrt{\sigma/(\rho\, h)}$. Then Eq. 4.4 gives the following equation of motion with pinned boundary conditions:

$$\partial_{tt} u = c^2 \left(\partial_{yy} + \partial_{zz}\right) u, \quad u\big|_{\partial\Omega} = 0.$$

4.1.2 Pinned-Membrane Mode Shapes and Frequencies

Assume a separable solution $u_{mn}(y,z,t) = A \times W_{mn}(y,z) \times X_{mn}(t)$, where A is a normalization constant. Substituting into Eq. 4.4 yields

$$\frac{\ddot{X}_{mn}(t)}{X_{mn}(t)} = c^2 \frac{\left(\partial_{yy} + \partial_{zz}\right) W_{mn}(y,z)}{W_{mn}(y,z)} = -\omega_{mn}^2,$$

where ω is the mode angular frequency. The spatial part $W(y,z)$ therefore satisfies

$$\left(\partial_{yy} + \partial_{zz}\right) W(y,z) + k^2\, W(y,z) = 0, \quad k \equiv \frac{\omega}{c},$$

subject to $W = 0$ on $\partial\Omega$. On $\Omega = [-\frac{l_y}{2}, \frac{l_y}{2}] \times [-\frac{l_z}{2}, \frac{l_z}{2}]$, unnormalized eigenfunctions are

$$W_{mn}(y,z) = \sin\left(\frac{m\,\pi}{l_y}\left(y + \frac{l_y}{2}\right)\right) \sin\left(\frac{n\,\pi}{l_z}\left(z + \frac{l_z}{2}\right)\right), \quad m,n \in \mathbb{N}. \quad (4.6)$$

The corresponding wavenumbers and eigenfrequencies are

$$k_{mn}^2 = \left(\frac{m\,\pi}{l_y}\right)^2 + \left(\frac{n\,\pi}{l_z}\right)^2, \quad \omega_{mn} = c\, k_{mn} = \pi \sqrt{\frac{\sigma}{\rho\, h}} \sqrt{\frac{m^2}{l_y^2} + \frac{n^2}{l_z^2}}. \quad (4.7)$$

The time-dependent part $X_{mn}(t)$ satisfies

$$\ddot{X}_{mn}(t) + \omega_{mn}^2\, X_{mn}(t) = 0,$$

This means that membrane mode is a simple harmonic oscillator with frequency ω_{mn}.

4.1.3 Effect of Hard Clamping (Finite Bending Rigidity)

In realistic devices, the membrane's edge is bonded to a substrate, imposing not only $u = 0$ but also zero slope at the edge (hard-clamp). Including finite bending rigidity $D = E\, h^3/[12(1-\nu^2)]$ (where E is Young's modulus and ν Poisson's ratio), the mode shape $W(y,z)$ satisfies [Lei69, Eq. 10.2]:

$$\left[\left(\partial_{yy} + \partial_{zz}\right) - \frac{D}{\sigma}\left(\partial_y^4 + \partial_z^4 + 2\,\partial_y^2\,\partial_z^2\right)\right] W(y,z) + k^2\, W(y,z) = 0\,, \quad (4.8)$$

subject to the boundary conditions at $y = \pm\frac{l_y}{2}$ or $z = \pm\frac{l_z}{2}$: $W = 0$ and $\partial_n W = 0$, where $\partial_n W$ is the normal derivative at each edge. The resulting mode profiles can be expressed approximately as a product of one-dimensional functions [Iva21][YPR12, Eq. S25-S26]:

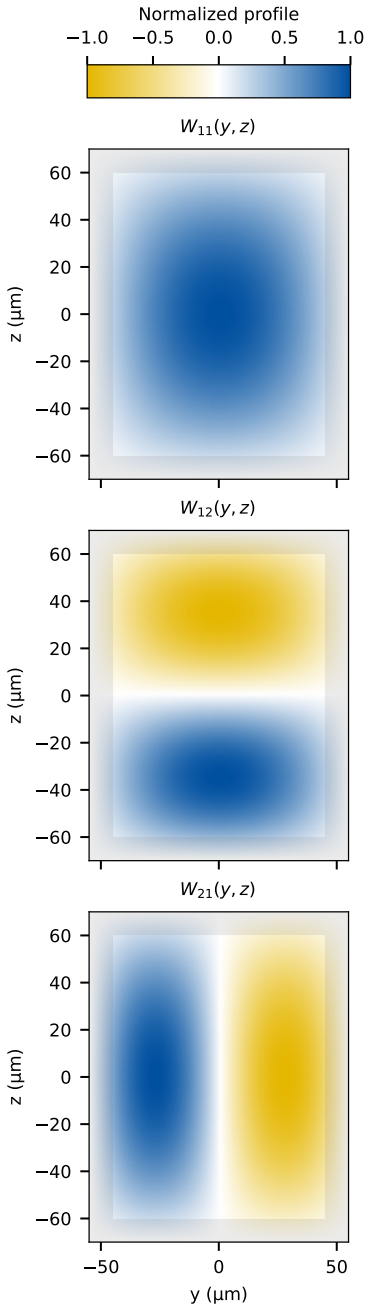

Fig. 4.2: Membrane modes illustrations. The grey areas illustrates the edges of the membrane. This part is not covered with Aluminium.

[Lei69]: Leissa (1969), *Vibration of plates*, NASA

[YPR12]: Yu et al. (2012), 'Control of Material Damping in High- Q Membrane Microresonators', Physical Review Letters

$$W_{mn}(y,z) = \gamma_y(m,y)\,\gamma_z(n,z), \tag{4.9}$$

with

$$\gamma_q(k,s) = \sin\left(\frac{k\,\pi}{l_q}\left(s+\frac{l_q}{2}\right)\right) + \epsilon_k(s)\,\frac{k\,\Lambda_q\,\pi}{2}\left[e^{-\frac{l_q-2|s|}{\Lambda_q\,l_q}} - \cos\left(\frac{k\,\pi}{l_q}(|s|-\frac{l_q}{2})\right)\right],$$

where $q \in \{y,z\}$, $k \in \{m,n\}$, $\epsilon_k(s)$ is a sign normalization factor and Λ_q is the dilution factor [Fed+19],

$$\epsilon_k(s) = \begin{cases} 1, & s<0, \\ (-1)^{k+1}, & s>0, \end{cases} \qquad \Lambda_q = \sqrt{\frac{E\,h^2}{3\,\sigma\,l_q^2\,(1-\nu^2)}},$$

This approximation holds for $\Lambda_y m \ll 1$ and $\Lambda_z n \ll 1$. In this regime, bending energy is concentrated in a narrow edge region of width $\sim \Lambda_y l_y$ and $\sim \Lambda_z l_z$, significantly altering the mode profile compared to the purely pinned case ($\Lambda \to 0$) (see Fig. 4.3). To reduce dissipation, one minimizes the dilution factor Λ: since Young's modulus E and Poisson ratio ν are fixed by the material, this entails increasing the pre-tension σ toward its fabrication limit. Importantly, hard clamping alters dissipation without appreciably shifting the mode frequency, so its effect can be neglected in total energy calculations.

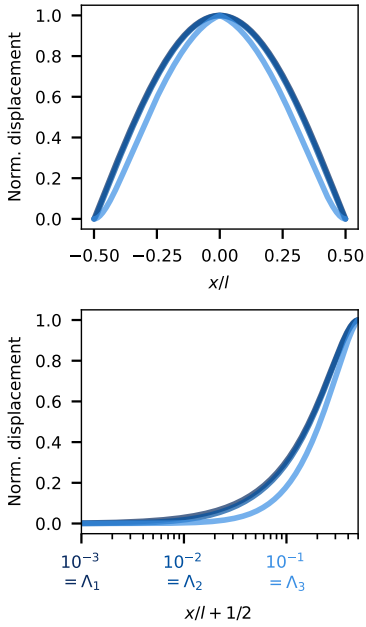

Fig. 4.3: Membrane hard clamping mode profile. Displacement profile of the fundamental mode (1, 1) normalized by the maximal displacement. The curves were calculated for $\Lambda = 10^{-3}, 10^{-2}, 10^{-1}$, from dark blue to light blue.

[Fed+19]: Fedorov et al. (2019), 'Generalized dissipation dilution in strained mechanical resonators', Physical Review B

4.1.4 Energy and Quality Factor

Given the mode effective mass (discussed in detail below), the total mechanical energy stored in the membrane (m,n) mode can be written, in direct analogy with a harmonic oscillator, as

$$\boxed{E_{mn} = \frac{1}{2}\,m_{\text{eff}}\,\dot{X}_{mn}^2 + \frac{1}{2}\,m_{\text{eff}}\,\omega_{mn}^2\,X_{mn}^2} \tag{4.10}$$

where the first term is the kinetic energy and the second term is the potential energy stored in the mode.

To quantify dissipation, we introduce the quality factor Q_{mn}, defined as the ratio of the energy stored in the mode to the energy lost per cycle ΔE:

$$\boxed{Q_{mn} = 2\pi\,\frac{E_{mn}}{\Delta E} = \frac{\omega_{mn}}{\Gamma_{mn}}} \tag{4.11}$$

where Γ_{mn} is the energy decay rate of the (m,n) mode. Physically, Q_{mn} measures how many oscillations occur before the energy decays by a factor $e^{-2\pi} \simeq 1.8 \times 10^{-3}$.

For the membranes used in this project (see Sec. 4.2), our estimates show that dissipation is overwhelmingly set by bending (anchor) losses, yielding a theoretical upper bound of

$$Q \approx 5 \times 10^5$$

as detailed in Sec. 3.1.3.4 of Patange's thesis [Pat25].

[Pat25]: Patange (2025), 'RF Transduction in a DC-Biased Superconducting Electromechanical System', Thesis

4.1.5 Mode Normalization and Effective Mass

The mode displacements $u_{mn}(y,z,t) = A \times W_{mn}(y,z) \times X_{mn}(t)$ are solutions of the linear elasticity equations (Eq. 4.4), thus there is a freedom in choosing the normalization constant A, leading to an ambiguity in the motional mass m_{eff}. To determine it, we start by calculating the kinetic energy of the membrane mode:

$$T_{mn} = \frac{1}{2}\,\dot{X}_{mn}^2\,A^2 \iint h\rho\, W_{mn}^2(y,z)\,\mathrm{d}y\,\mathrm{d}z. \tag{4.12}$$

Identifying with the harmonic oscillator kinetic energy, $T = \frac{1}{2} m_\lambda \dot{X}^2$, we define the effective mass as

$$m_{\text{eff},mn} = A^2 \iint \rho\, W_{mn}^2(y,z)\,\mathrm{d}y\,\mathrm{d}z\,. \tag{4.13}$$

For the mode of our interest $(2,1)$, this leads to the effective mass $m_{\text{eff},21} = A^2\,M/4$, where M is the total mass of the membrane. Thus the effective mass itself is determined up to a normalization factor A^2, which we can fix. There is a convenient choice that will simplify the calculations for the capacitive coupling scheme employed in this work. Imagine the membrane is coupled to a capacitor with the area of overlap $\mathcal{A}$ and mean distance d between the membrane and the capacitor plate. Then the capacitance of the capacitor is given by

$$\begin{aligned} C_m(X_{mn}(t)) &= \iint_{\mathcal{A}} \frac{\epsilon_0}{d + A\,X_{mn}(t)\,W_{mn}(y,z)}\,\mathrm{d}y\,\mathrm{d}z \\ &\approx \frac{\epsilon_0 \mathcal{A}}{d}\left(1 - \frac{X_{mn}(t)}{d\mathcal{A}} A \iint_{\mathcal{A}} W_{mn}(y,z)\,\mathrm{d}y\,\mathrm{d}z\right). \end{aligned} \tag{4.14}$$

We choose the normalization constant A such that the integral of the mode profile $W_{mn}(y,z)$ over the overlap area equals the overlap area, that is,

$$A \iint_{\mathcal{A}} W_{mn}(y,z)\,\mathrm{d}y\,\mathrm{d}z = \mathcal{A}, \tag{4.15}$$

which allows us to express the capacitance to first order in the displacement as if it was a parallel-plate capacitor,

$$C \simeq C_0\left(1 - \frac{X}{d}\right), \tag{4.16}$$

where $C_0 = \epsilon_0\,\mathcal{A}/d$ is the capacitance at the equilibrium position $X = 0$. With this normalization, the effective mass corresponds to the mass of the membrane region sensed by the fluxonium electrodes, as if the mode profile were uniform and the device formed an ideal parallel-plate capacitor. For the mode $(2,1)$, this normalization yields $A = 1.3$, and hence $m_{\text{eff},21} = 2.3$ ng for our membrane parameters provided in Sec. 4.2.

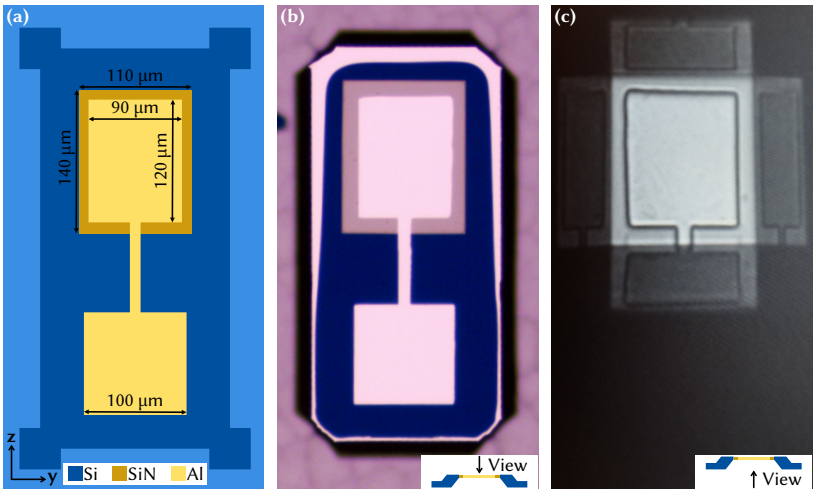

Fig. 4.4: Membrane design. (a) Top: suspended SiN rectangular membrane in orange (only the non metalized edge is visible). Aluminium layer in gold. Silicon substrate in blue. Bottom: square biasing electrode lying on the substrate. (b) Optical microscope image of the front-side of the final device. The suspended membrane and the biasing electrode sit on a silicon mesa, the typical rough etch bottom, arising from KOH etch can be seen at the edges of the picture. (c) Optical image of the same sample, seen from the back side. The released metalized SiN membrane can be seen at the top. Mirror images on the four edges of the membrane come from the smooth crystalline planes obtained after etching with KOH trough the substrate.

4.2 Membrane Design and Fabrication

4.2.1 Design and FEM simulations

In earlier works of our team [Cap20; Iva21], membrane resonance frequencies lay in the few-hundred-kilohertz range. To achieve resonant coupling to a fluxonium qubit operating at megahertz frequencies, we deliberately designed a membrane with $\omega_{21}/2\pi \approx 5\,\mathrm{MHz}$. Substituting into Eq. 4.7 yields lateral dimensions of approximately $110 \times 140\,\mu\mathrm{m}$[1]. This was checked with FEM simulations using *Comsol Multiphysics*. The resulting mode profiles with corresponding frequencies are shown in Fig. 4.5.

For capacitive coupling, the membrane was metallized with a 30 nm Al film, deliberately leaving the clamped edges metal-free to facilitate alignment during flip-chip assembly and avoid additional bending losses. The resulting metallized region measures $90 \times 120\,\mu\mathrm{m}$ (see Fig. 4.4(a)).

Capacitive biasing is achieved via matching square electrodes patterned on the membrane substrate and on the fluxonium chip (see the following section), which together form the coupling capacitance. Finally, to simplify the flip-chip assembly, all unnecessary substrate was removed by etching (see Fig. 4.4(b)). Details of this fabrication process are discussed in the thesis of H. Patange [Pat25].

[Cap20]: Capelle (2020), 'Electromechanical cooling and parametric amplification of an ultrahigh-Q mechanical oscillator', Thesis

[Iva21]: Ivanov (2021), 'Optimization of silicon nitride membranes for hybrid superconducting-mechanical circuits', Thesis

1: In this calculation, we used $\rho_{\mathrm{SiN}} = 3180\,\mathrm{kg/m^3}$ and an effective stress $\sigma \approx 1\,\mathrm{GPa}$ (the manufacturer specifies $\geq 0.8\,\mathrm{GPa}$, but fits to measured data and FEM simulations in [Pat25] indicate higher values).

[Pat25]: Patange (2025), 'RF Transduction in a DC-Biased Superconducting Electromechanical System', Thesis

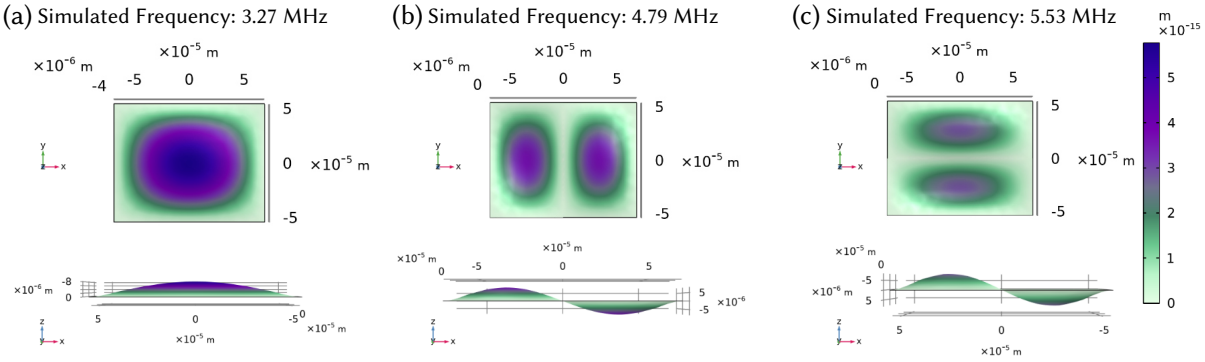

Fig. 4.5: Simulated membrane mode shapes. Upper and lower panels show top and side views, respectively, of the three lowest-frequency mode shapes of the membrane. The membrane parameters are $l_x = 140\,\mu$m, $l_y = 110\,\mu$m, $h = 0.9\,\mu$m, $\sigma = 0.8\,$GPa, $\nu = 0.27$, and $E = 270\,$GPa. All plots share the same colorbar and arbitrary amplitude shown here. These simulations are originally presented in [Pat25].

4.2.2 Membrane fabrication

A membrane fabrication process was developed previously in our group by E. Ivanov [Iva20] and was further updated for the flip-chip by H. Patange [Pat25]. In the following, we give a short overview of how freestanding SiN membranes are fabricated.

- **Starting material:** Silicon wafers (280 μm thick) coated on both sides with 100 nm LPCVD SiN.
- **Wafer thickness measurement:** Measures silicon thickness to determine etch time and achieve the target membrane size.
- **Backside lithography:** Defines square windows in the backside SiN by applying and patterning a resist mask aligned to the crystal orientation.
- **Backside RIE:** Uses reactive-ion etching to remove the exposed SiN, creating openings down to the silicon substrate.
- **KOH etching:** Immerses the frontside-protected wafer in 30% KOH at 85°C, anisotropically etching silicon to form V-shaped cavities and release the SiN membrane.
- **Membrane size adjustment:** Optionally extends the KOH etch to widen any undersized membranes along the (111) sidewalls.
- **HF cleaning:** Dips the wafer in dilute HF to remove residual etch by-products and contaminants with 10% SiN thinning.
- **Frontside lithography:** Patterns the frontside by laser lithography and KOH to remove 25 μm of Si from the frontside, increasing the further gap after the flip-chip outside of the membrane and pillars.
- **Aluminium deposition:** Patterns a bi-layer resist, evaporates aluminium to form metal pads, and lifts off the resist together with unnecessary metal parts.
- **Dicing:** Cuts the wafer into 6 mm×6 mm chips under a protective resist layer, then strips the resist in solvent.

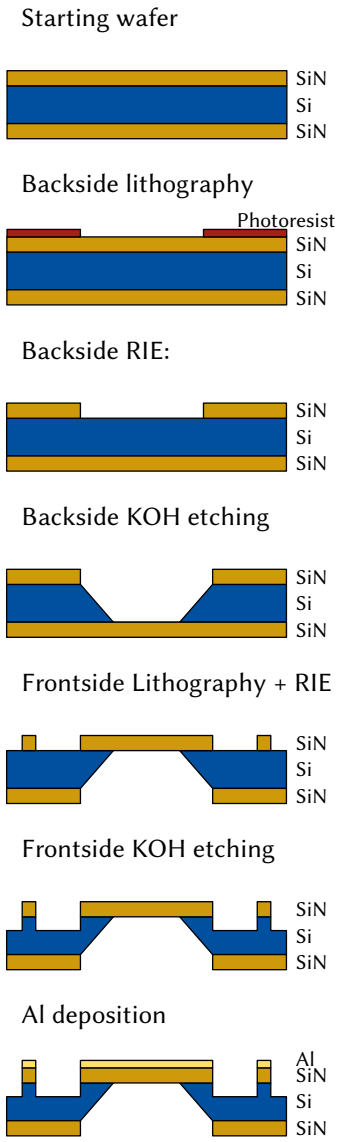

Fig. 4.6: Illustration of the membrane fabricating steps. See the main text for details.

4.3 Room Temperature Characterization

Prior to flip-chip integration, each membrane is characterized at room temperature to determine its mode shape, resonance frequency, and quality factor using a high-sensitivity optical interferometer. In the first year of this project, I fully rebuilt the interferometer originally built by Ivanov [Iva21], simplifying the optical paths by eliminating redundant mirrors and beam splitters, and replacing standard lenses with aspheric optics. These changes reduced the beam waist on the membrane from 23 µm to 6 µm, substantially improving spatial resolution and signal strength. We also rewrote the data acquisition software, implementing automated beam-scanning routines and synchronized detection, which transformed a previously manual, multi-day measurement into an unattended overnight experiment.

In what follows, we first describe the interferometer hardware and phase-sensitive detection scheme in Sec. 4.3.1, which I implemented during the first year of my PhD, and then present the room-temperature measurement results for our square membranes in Sec. 4.3.2.

4.3.1 Optical Interferometer Setup

The interferometer is designed to measure the local displacement of a SiN membrane with shot-noise-limited sensitivity, enabling observation of its thermal motion and reconstruction of its vibrational mode shapes by scanning the beam focus across the membrane.

The optical layout is shown in Fig. 4.7. A 1064 nm Nd:YAG laser provides the input beam. Ignore for now the amplitude modulator, which generally is set to constant transmission. The beam goes through a $\lambda/2$ waveplate and a polarizing beam splitter PBS_0, which controls the total intensity and sets the polarization. The beam is then split by the PBS_1 into the signal arm, with field amplitude α_s, and the local-oscillator (LO) arm, with amplitude α_{LO}. The relative intensity between the two arms is controlled by a $\lambda/2$ waveplate before the PBS_1.

The signal beam α_s passes through a dichroic mirror and focusing lens before entering a vacuum chamber, where it reflects off the SiN membrane, with a reflectivity of 30%. The dichroic mirror transmits the 1064 nm beam while reflecting visible light, allowing simultaneous illumination of the sample and imaging with a camera for alignment and focus. Upon reflection, the signal beam acquires a phase shift $\phi_s(X(t))$ that encodes the membrane's displacement $X(t)$ at specific position.

The LO beam α_{LO} reflects from a mirror mounted on a piezo actuator, acquiring a tunable phase $\phi_{LO}(t)$. Quarter-waveplates ($\lambda/4$) in each arm ensure that, on their return to the beamsplitter, the two beams exit toward the detection port rather than back toward the laser.

After recombination it goes through a final $\lambda/2$ to balance amplitude and to couple the horizontally polarized α_s with vertically polarized α_{LO}. Finally, the beam is split equally by PBS_2 onto two photodiodes PD_1 and

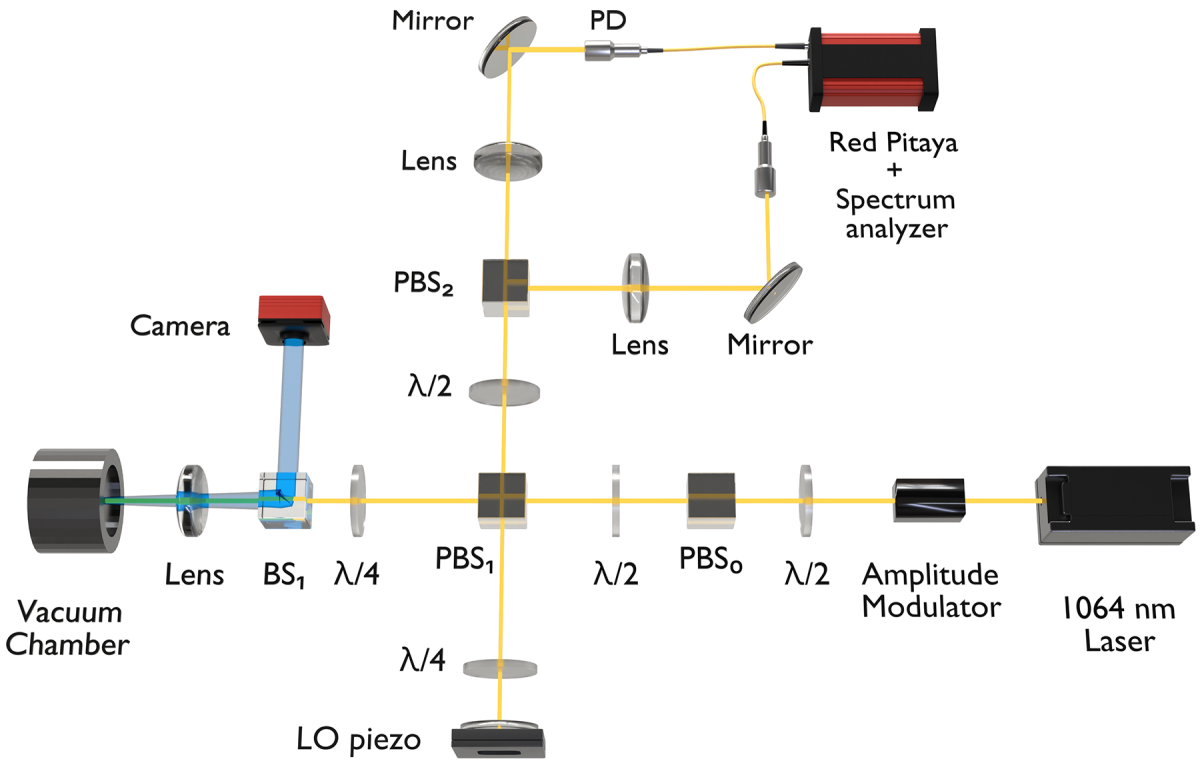

Fig. 4.7: Interferometer setup for room-temperature characterization of membrane resonators. The working principle is described in the main text. In the experiment, light first passes through the AM inside an optical fiber, then enters the interferometer. This configuration allows all interferometer components (except the vacuum chamber, AM, and laser) to be mounted on a motorized micro-positioning stage for scanning the laser spot across the membrane surface. The local oscillator (LO) mirror, mounted on a piezo actuator, is driven by a RedPitaya board.

PD_2. A differential amplifier subtracts the two photocurrents to produce the difference signal:

$$\begin{aligned} \Delta I &= I_1 - I_2 \\ &\propto \mathrm{Re}\left[\alpha_\mathrm{s}\,\alpha_\mathrm{LO}\,e^{i\left(\phi_\mathrm{s}(X(t))-\phi_\mathrm{LO}(t)\right)}\right] \\ &\propto \alpha_\mathrm{s}\,\alpha_\mathrm{LO}\,\cos\left(\phi_\mathrm{s}(X(t)) - \phi_\mathrm{LO}(t)\right). \end{aligned} \tag{4.17}$$

Because the membrane oscillates about its equilibrium, we write the signal phase as $\phi_\mathrm{s}(X(t)) = \phi_\mathrm{s}^0(t) + \delta\phi_\mathrm{s}(X(t))$, where $\phi_\mathrm{s}^0(t)$ captures slow experimental drifts (e.g., macroscopic vibrations) and the phase fluctuation is directly related to the membrane displacement $\delta\phi_\mathrm{s}(X(t)) = 4\pi X(t)/\lambda$, where λ is the wavelength of the laser. For $\delta\phi_\mathrm{s}(X(t)) \ll 1$, expanding Eq. 4.17 gives

$$\Delta I \propto \alpha_\mathrm{s}\,\alpha_\mathrm{LO}\Big(\cos\Delta\phi(t) - \sin\Delta\phi(t)\cdot\delta\phi_\mathrm{s}(X(t)) + O(\delta\phi_\mathrm{s}^2)\Big). \tag{4.18}$$

where $\Delta\phi(t) = \phi_\mathrm{s}^0(t) - \phi_\mathrm{LO}(t)$. To eliminate the constant offset and be linearly sensitive to the membrane's displacement, we lock the phases such that

$$\phi_\mathrm{s}^0 - \phi_\mathrm{LO} = \frac{\pi}{2}.$$

This condition is enforced by sending the low-frequency component of ΔI into a RedPitaya board and running a PID lock in *PyRPL* [Neu+24], which drives the piezo actuator. The high-frequency component of ΔI is routed

[Neu+24]: Neuhaus et al. (2024), 'Python Red Pitaya Lockbox (PyRPL)', Review of Scientific Instruments

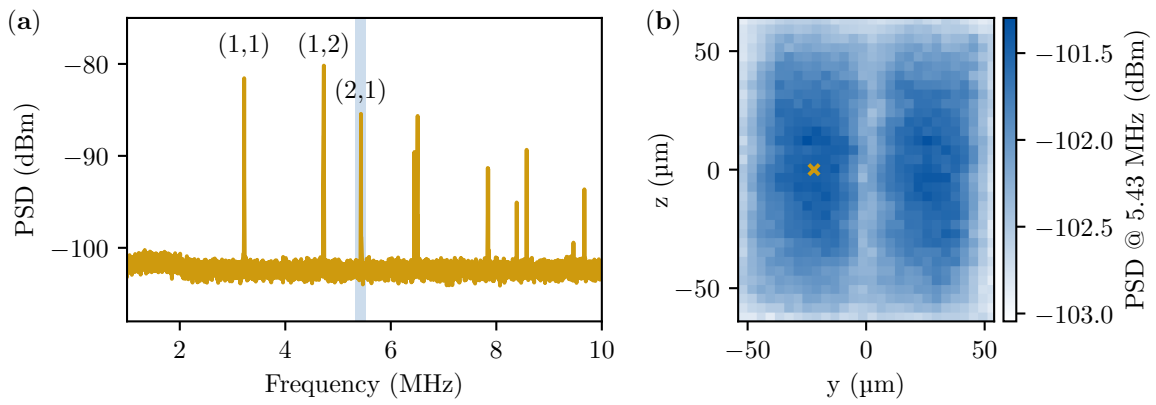

Fig. 4.8: Spectrum and measured $(2,1)$ **mode profile of the square membrane. (a)** Power spectral density (PSD) of the thermal displacement fluctuations of the membrane, measured with the room-temperature interferometer at the position marked by the yellow cross in (b). First three modes are labeled. **(b)** Spatial mode profile reconstructed by taking the peak of PSD within a narrow window around $\omega_{21}/2\pi = 5.43\,\mathrm{MHz}$ (blue band in (a)) for each (y, z) position.

directly to a spectrum analyzer to record the membrane's displacement noise spectrum S_{XX}.

4.3.2 Mode profile measurements

The interferometer is mounted on a motorized micro-positioning stage with 1 μm resolution, enabling a raster scan of the laser spot across the entire membrane surface and thus reconstruction of the full mode shapes. To facilitate coarse alignment, a camera views the membrane through a dichroic mirror that reflects visible light while transmitting the 1064 nm beam.

The membrane with the same design that is later used in the flip-chip integration was measured at room temperature [Pat25], yielding mode frequencies

$$\omega_{11}/2\pi = 3.22\,\mathrm{MHz}, \quad \omega_{12}/2\pi = 4.73\,\mathrm{MHz}, \quad \omega_{21}/2\pi = 5.43\,\mathrm{MHz}.$$

Fig. 4.8 also shows the measured spatial profile of the antisymmetric $(2,1)$ mode.

4.3.3 Quality factor measurements

To determine the mechanical quality factor Q, we perform ring-down measurements by resonant radiation-pressure excitation of a selected mode via intensity modulation of the laser (using the fibered amplitude modulator).

Once the drive is turned off, the mechanical amplitude decays exponentially, and fitting decay of PSD yields $Q = \omega_m\,\tau$, where τ is the measured decay time constant. Measured quality factors are

$$Q_{11} \approx 5.2 \times 10^4, \quad Q_{12} \approx 6.6 \times 10^4, \quad Q_{21} \approx 6.0 \times 10^4.$$

In Fig. 4.9, we show an example of the measured ring-down signal for the $(2,1)$ mode.

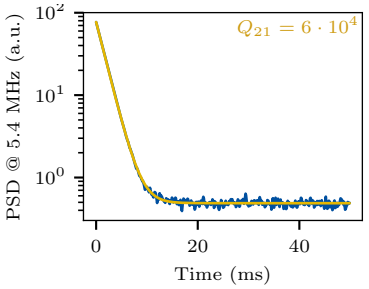

Fig. 4.9: Ring-down measurement of the $(2,1)$ **vibrational mode.** The plot shows the power spectral density (PSD) of the membrane's displacement at the mode frequency $\omega_{21}/2\pi = 5.43\,\mathrm{MHz}$ (blue) overlaid with the fitted exponential decay envelope (yellow). The data were smoothed using a Savitzky–Golay filter. The dataset was digitized from the original rasterized figure in [Pat25, Fig. 3.9], preserving its statistical properties.

Mecaflux Chip Design 5

The coupling between the qubit and the membrane-the main theme of this thesis-is achieved by a *flip-chip* assembly where the membrane chip is placed on top of the fluxonium chip. This chapter aims to detail the design of the fluxonium chip and its couplings.

My contribution. At the start of my PhD, I conducted experiments with early generations of heavy fluxonium chips designed by senior PhD student B.-L. Najera-Santos. It was a grounded heavy fluxonium chip [Naj+24] that was designed to be coupled to the symmetric mode of the membrane. However as we will show in the begin of this chapter, coupling the qubit to the symmetric mode of the membrane introduced significant common-mode noise to the qubit. To address this issue, Dr. L. Balembois, current PhD student R. Rousseau, and I redesigned the qubit to couple to an antisymmetric mode of the membrane.

Outline. This chapter starts with the simplest scheme: a single vacuum-gap capacitor biased at voltage V_d that linearly converts membrane displacement of the first symmetric mode into charge on the qubit island. We derive its interaction Hamiltonian to illustrate an additional voltage-bias–induced noise on the qubit, which makes this idea impractical. We then consider the antisymmetric coupling configuration, where two opposing, equally biased capacitors transduce motion while canceling common-mode noise. The Hamiltonian of this antisymmetric mode constitutes the primary Hamiltonian for our qubit–membrane flip-chip system.

We then introduce our approach to designing and simulating the QED circuit. The physical fluxonium chip is inherently a two-dimensional platform, with all circuit elements arranged in a planar layout. We present a detailed overview of the layout employed in this work as well as examine the environmental couplings.

5.1 Coupling to a Symmetric Membrane Mode

We begin with the simplest coupling mechanism. As before, we consider a grounded fluxonium qubit, where one side of the junction is at ground potential. The qubit is coupled to the membrane via an additional capacitance C_m created between the conducting pad of the qubit and the metalized membrane. The membrane is biased at a voltage V_b through a direct connection to a DC source.

This configuration represents our first attempt at qubit-membrane coupling. Understanding its shortcomings will help clarify the motivation for the improved design presented in section Sec. 5.2.

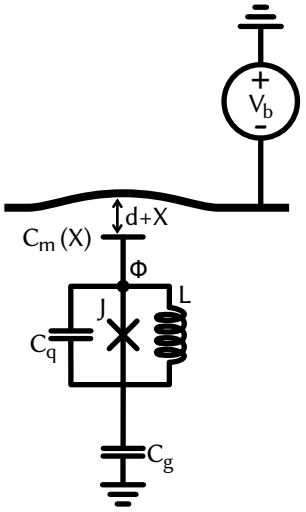

Fig. 5.1: Symmetric qubit-membrane coupling circuit. Schematic of a fluxonium qubit formed by a Josephson junction J in parallel with an inductance L and capacitance C_q. One terminal is grounded through C_g, while the other is coupled to a mechanical membrane via the position-dependent capacitance $C_m(X)$. The distance between the qubit and the membrane is $d + X$, where d is the equilibrium separation and X is the membrane's displacement. The membrane is biased by a voltage source V_b.

The potential part of the Lagrangian remains unaffected by the capacitive coupling and is not involved in the Legendre transformation. We write the kinetic part of the qubit-membrane system as:

$$T = \frac{C_g}{2}\dot{\Phi}^2 + \frac{C_m(x)}{2}\left(V_b - \dot{\Phi}\right)^2. \tag{5.1}$$

Using the Legendre transformation, $\partial T/\partial\dot{\Phi} = Q$, quantization, the dimensionless phase operator $\boldsymbol{\varphi} = \boldsymbol{\Phi}/\phi_0$, and the conjugate charge operator $\boldsymbol{n} = \boldsymbol{Q}/2e$, we can write the Hamiltonian of the system as:

$$\boldsymbol{H} = 4E_C(\boldsymbol{x})\left(\boldsymbol{n} + n_g(\boldsymbol{x})\right)^2 + \frac{E_L}{2}\boldsymbol{\varphi}^2 - E_J\cos\left(\boldsymbol{\varphi} - \varphi_{\text{ext}}\right) - \frac{C_m(\boldsymbol{x})\,V_b^2}{2}, \tag{5.2}$$

where the potential energy is added from the uncoupled fluxonium Hamiltonian, and:

$$E_C(\boldsymbol{x}) \equiv \frac{e^2}{2(C_g + C_m(\boldsymbol{x}))}, \qquad n_g(\boldsymbol{x}) \equiv \frac{C_m(\boldsymbol{x})}{2e}\cdot V_b. \tag{5.3}$$

For small displacements, the membrane capacitance can be approximated as:

$$C_m(\boldsymbol{X}) = \frac{C_m}{1 + \boldsymbol{X}/d} \approx C_m\left(1 - \frac{2x_{\text{zpf}}}{d}\boldsymbol{x}\right) \tag{5.4}$$

where C_m is the equilibrium capacitance of the membrane, x_{zpf} is the zero-point fluctuation of the membrane, and $\boldsymbol{x} \equiv \boldsymbol{X}/2x_{\text{zpf}} \equiv (\boldsymbol{a} + \boldsymbol{a}^\dagger)/2$ is the dimensionless displacement operator of the membrane. Thus, the coupling term between the qubit and the membrane takes the form

$$\boldsymbol{H}_{\text{coupling}} = \frac{4\,C_m\,E_C}{e}\left(1 + \frac{2\,C_g}{C_g + C_m}\,\frac{x_{\text{zpf}}}{d}\,\boldsymbol{x}\right)V_b\,\boldsymbol{n}, \tag{5.5}$$

where $E_C = E_C(x = 0)$. To include bias-line noise, we set $V_b \to V_b + \delta V_b$, which yields

$$\boldsymbol{H}_{\text{coupling}} = \frac{4\,E_C\,C_m}{e}\left(V_b + \delta V_b - \frac{2\,C_g}{C_g + C_m}\,\frac{x_{\text{zpf}}}{d}\,(V_b + \delta V_b)\boldsymbol{x}\right)\boldsymbol{n}. \tag{5.6}$$

Here, the first term is the static charge offset from V_b; the second term represents common-mode noise; the third term is the desired qubit–mechanics coupling; and the final term is second-order noise. In this grounded fluxonium configuration, fluctuations δV_b are directly transduced into charge noise on the fluxonium, so there is no rejection of the common-mode component. Moreover, since $x_{\text{zpf}}/d \sim 10^{-9}$, the noise contribution easily overwhelms the coupling.

Furthermore, the $V_b^2\,\boldsymbol{x}$ term in Eq. 5.2 gives rise to a mechanical drive,

$$-\frac{2\,E_C\,C_g^2\,C_m}{e^2\,(C_g + C_m)}\,\frac{x_{\text{zpf}}}{d}\left(V_b^2 + 2\,V_b\,\delta V_b + \delta V_b^2\right)\boldsymbol{x},$$

In particular, the cross term $\propto V_b\,\delta V_b\,\boldsymbol{x}$ pumps the membrane motion. Its prefactor is large enough that, experimentally, even a fraction of a volt of offset bias drove the membrane above its thermal noise floor [Pat25],

precluding both sensitive thermal-motion detection and further ground-state cooling.

At the outset of my PhD, we set out to realize this grounded-mode coupling scheme. After several attempts and extensive filtering efforts, it became clear that suppressing environmental δV_b noise to acceptable levels was not feasible. We therefore adopted a new architecture in which the qubit couples to a membrane's antisymmetric mode, inherently canceling common-mode voltage fluctuations, as explained in the following section.

5.2 Coupling to an Antisymmetric Membrane Mode

To mitigate noise introduced by the membrane bias, we consider the configuration shown in Fig. 5.2. Qubit–membrane coupling is achieved via the two capacitors C_m^+ and C_m^- [VML18], formed between the qubit's conducting pad and the metallized membrane. The membrane is biased at a voltage V_b through a coupling capacitor C_b. This avoids a galvanic connection between to bias the membrane, which cannot be guaranteed in fabrication from our experience [Pat25]. The fluxonium qubit consists of a floating superconducting loop that is symmetrically coupled to ground through two identical capacitors C_g. The electrical schematic is shown in Fig. 5.2, where any additional parasitic capacitances have been neglected. The dynamics are described by the membrane displacement coordinate X and the vector of effective flux-node coordinates

$$\vec{\Phi}^T = \big(\Phi_1,\ \Phi_2,\ \Phi_3\big), \tag{5.7}$$

where Φ_1 is the flux across the qubit, and $\Phi_{2,3}$ denote the flux coordinates at the coupling nodes in the absence of the fluxonium circuit (see Fig. 5.2).

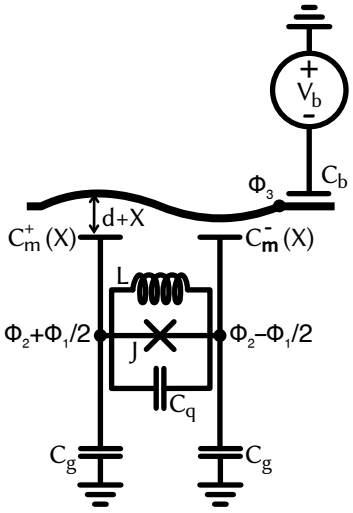

Fig. 5.2: Antisymmetric qubit-membrane coupling circuit. Schematic of a fluxonium qubit formed by a Josephson junction J in parallel with an inductance L and capacitance C_q, bridging two nodes held at flux biases $\phi_2 \pm \frac{\phi_1}{2}$. Each node is grounded through a capacitance C_g. An antisymmetric mechanical membrane mode—displaced by $\pm X$ from its equilibrium separation d—serves as an electrode and couples to both qubit nodes via capacitances $C_m^{\pm}(X)$. The membrane node ϕ_3 is biased by a DC voltage V_b through a coupling capacitor C_b.

[Pat25]: Patange (2025), 'RF Transduction in a DC-Biased Superconducting Electromechanical System', Thesis

5.2.1 Hamiltonian derivation

We express the total Lagrangian as the sum of mechanical and electrical contributions:

$$\mathcal{L} = \mathcal{L}_{\text{mec}}(X, \dot{X}) + \mathcal{L}_{\text{elec}}(\Phi_1, \dot{\vec{\Phi}}, X)\,, \tag{5.8}$$

where overdots indicate time derivatives. The membrane Lagrangian is that of a harmonic oscillator:

$$\mathcal{L}_{\text{mec}} = \frac{1}{2}\, m\, \dot{X}^2 - \frac{1}{2}\, m\, \omega_m^2\, X^2, \tag{5.9}$$

with m the effective mass of the membrane and ω_m its resonance frequency. The electrical Lagrangian is split into kinetic and potential parts,

$\mathcal{L}_{\text{elec}} = T_{\text{elec}} - V_{\text{elec}}$:

$$\begin{aligned} T_{\text{elec}} = {} & \frac{C_g}{2}\left(\dot{\Phi}_2 - \tfrac{1}{2}\dot{\Phi}_1\right)^2 + \frac{C_g}{2}\left(\dot{\Phi}_2 + \tfrac{1}{2}\dot{\Phi}_1\right)^2 \\ & + \frac{C_m^+}{2}\left(\dot{\Phi}_3 - \dot{\Phi}_2 - \tfrac{1}{2}\dot{\Phi}_1\right)^2 + \frac{C_m^-}{2}\left(\dot{\Phi}_3 - \dot{\Phi}_2 + \tfrac{1}{2}\dot{\Phi}_1\right)^2 \\ & + \frac{C_b}{2}\left(V_b - \dot{\Phi}_3\right)^2 + \frac{C_q}{2}\,\dot{\Phi}_1^2, \end{aligned} \tag{5.10}$$

$$V_{\text{elec}} = -E_J \cos\!\left(2\pi\Phi_q/\Phi_0\right) + \frac{\Phi_1^2}{2L}. \tag{5.11}$$

Here, C_g is the ground capacitance, C_b the membrane bias capacitance, C_q the qubit shunt capacitance, V_b the bias voltage, E_J the Josephson energy, $\Phi_0 = h/2e$ the flux quantum, and L the inductance of the fluxonium.

For the Legendre transform we rewrite T_{elec} in matrix form:

$$T = \frac{1}{2}\sum_{i,j=1}^{3} \dot{\Phi}_i\, M_{ij}(X)\, \dot{\Phi}_j - C_b\, V_b\, \dot{\Phi}_3, \tag{5.12}$$

with the position-dependent capacitance matrix

$$M(x) = \begin{pmatrix} \frac{1}{2}\left(C_g + C_m^\Sigma + 2C_q\right) & C_m^\Delta & -C_m^\Delta \\ C_m^\Delta & 2\left(C_g + C_m^\Sigma\right) & -2\,C_m^\Sigma \\ -C_m^\Delta & -2\,C_m^\Sigma & C_b + 2\,C_m^\Sigma \end{pmatrix},$$

with $C_m^\Sigma \equiv (C_m^+ + C_m^-)/2$ and $C_m^\Delta \equiv (C_m^+ - C_m^-)/2$ are symmetric and antisymmetric combinations of the membrane-qubit capacitances. The canonical flux-momenta are defined by $Q_k = \partial T/\partial\dot{\Phi}_k$, where $k = 1, 2, 3$. This yields

$$\dot{\Phi}_k = \sum M_{ki}^{-1}\,\widetilde{Q}_i\,, \text{ where } \widetilde{Q}_i = Q_i + C_b\, V_b\, \delta_{i3} \tag{5.13}$$

Here δ_{ij} is the Kronecker delta. The membrane momentum is simply $P_X = m\,\dot{X}$. Performing the Legendre transform, $H = \sum_{i=1}^{3} Q_i\,\dot{\Phi}_i + P_X\,\dot{X} - \mathcal{L}$, yields the Hamiltonian

$$\begin{aligned} H = {} & \frac{1}{2}\sum_{i,j=1}^{3} \widetilde{Q}_i\, M_{ij}^{-1}\, \widetilde{Q}_j - E_J \cos\!\left(2\pi\Phi_1/\Phi_0\right) + \frac{\Phi_1^2}{2L} \\ & + \frac{P_X^2}{2m} + \frac{1}{2}\, m\, \omega_m^2\, X^2. \end{aligned} \tag{5.14}$$

Since $\mathcal{L}$ (Eq. (5.8)) does not depend on Φ_2 or Φ_3, these coordinates are cyclic and their conjugate charges are strictly conserved by Noether's theorem. Physically, these conserved quantities correspond to the static charges on the fluxonium and membrane islands, and may therefore be treated as fixed scalar parameters in the subsequent analysis.

We quantize the system by introducing the canonical variables for the qubit phase φ and Cooper pair number n, along with the dimensionless canonical variables for the membrane position x, momentum p and $\boldsymbol{a}$

($\boldsymbol{a}^\dagger$) the mechanical annihilation (creation) operators:

$$\begin{aligned} \boldsymbol{\varphi} &= 2\pi\frac{\boldsymbol{\Phi}_1}{\Phi_0}, \qquad \boldsymbol{x} = \frac{\boldsymbol{X}}{2\,\mathrm{X}_{\mathrm{zpf}}} = \frac{\boldsymbol{a}+\boldsymbol{a}^\dagger}{2}, \\ \boldsymbol{n} &= \frac{\boldsymbol{Q}_1}{2e}, \qquad \boldsymbol{p} = \frac{\boldsymbol{P}}{2P_{\mathrm{zpf}}} = i\frac{\boldsymbol{a}^\dagger-\boldsymbol{a}}{2}, \end{aligned} \tag{5.15}$$

where $X_{\mathrm{zpf}} = \sqrt{\hbar/(2m\omega_m)}$ and $P_{\mathrm{zpf}} = \sqrt{\hbar m\omega_m/2}$ are the membrane zero-point fluctuations of position and momentum, respectively. These variables satisfy the commutation relations $[\boldsymbol{\varphi}, \boldsymbol{n}] = i$ and $[\boldsymbol{x}, \boldsymbol{p}] = i/2$, and operators from different sets commute. The capacitance matrix depends on the membrane position, $\boldsymbol{M}_{ij}^{-1} = \boldsymbol{M}_{ij}^{-1}(\boldsymbol{x})$, and must therefore be treated as an operator. In this basis the Hamiltonian assumes the familiar form for a fluxonium coupled to a harmonic mode:

$$\begin{aligned} \boldsymbol{H} &= 4\,\boldsymbol{E}_C(\boldsymbol{x})\big[\boldsymbol{n}+\boldsymbol{n}_g(\boldsymbol{x})\big]^2 - E_J\cos\boldsymbol{\varphi}_q + \tfrac{1}{2}\,E_L\,\boldsymbol{\varphi}_q^2 \\ &\quad + \hbar\omega_m\big(\boldsymbol{p}^2+\boldsymbol{x}^2\big) + \boldsymbol{F}_m(\boldsymbol{x}), \end{aligned} \tag{5.16}$$

where

$$E_L = \left(\frac{\Phi_0}{2\pi}\right)^2\frac{1}{L}, \quad \boldsymbol{E}_C(\boldsymbol{x}) = \frac{e^2}{2}\,\boldsymbol{M}_{11}^{-1}(\boldsymbol{x}),$$
$$\boldsymbol{n}_g(\boldsymbol{x}) = \frac{e}{4\,\boldsymbol{E}_C(\boldsymbol{x})}\sum_{i=2}^{3}\widetilde{Q}_i\,\boldsymbol{M}_{1i}^{-1}(\boldsymbol{x}).$$

This form of the Hamiltonian shows that the coupling between the fluxonium and the membrane induces an additional charge $\boldsymbol{n}_g(\boldsymbol{x})$ on the fluxonium. The membrane also experiences an additional electrostatic membrane shift potential $\boldsymbol{H}_{ms}(\boldsymbol{x})$, which, as will be shown later, leads to a reduction of the membrane frequency. The general form of $\boldsymbol{H}_{ms}$ is given by

$$\boldsymbol{H}_{ms} = 4\boldsymbol{E}_C\boldsymbol{n}_g^2 + \boldsymbol{F}_m = \frac{1}{2}\sum_{i,j=2}^{3}\widetilde{Q}_i\,\boldsymbol{M}_{ij}^{-1}\,\widetilde{Q}_j. \tag{5.17}$$

5.2.2 Qubit–membrane coupling

For small membrane displacements we perform a first-order expansion of the coupling Hamiltonian in the dimensionless coordinate $\boldsymbol{x}$, yielding

$$\begin{aligned} H_{\mathrm{coupling}} &= 8\,\boldsymbol{E}_C(\boldsymbol{x})\,\boldsymbol{n}_g(\boldsymbol{x}) \\ &\approx\; 4E_C\,\frac{\partial C_m^+}{\partial x}\,\frac{V_{\mathrm{eff}}(V_b)}{e}\,\boldsymbol{n}\,\boldsymbol{x} \;\equiv\; g\,\boldsymbol{n}\,\boldsymbol{x}, \end{aligned} \tag{5.18}$$

Here,

$$E_C = \frac{e^2}{C_g + C_m + 2C_q} \tag{5.19}$$

is the unperturbed capacitive energy of the fluxonium (neglecting all terms beyond linear order in x), and

$$\frac{\partial C_m^\pm}{\partial x} = 2\,\mathrm{X}_{\mathrm{zpf}}\,\frac{\partial C_m^\pm}{\partial X}$$

relates the derivative with respect to the dimensionless coordinate $\boldsymbol{x}$ back to the physical membrane displacement X. The effective bias, $V_{\text{eff}}(V_b)$, is given by:

$$\begin{aligned} V_{\text{eff}}(V_b) &= \frac{C_g\, V_b - Q_2/2 + (C_g/C_b)\, Q_3}{C_g + C_m + 2\,(C_g\, C_m/C_b)} \\ &\longrightarrow \frac{C_g\, V_b - Q_2/2}{C_g + C_m}\,, \end{aligned} \tag{5.20}$$

where the last expression indicates the limit $C_g/C_b \ll 1$. This indicates that trapped charges, Q_2 and Q_3, act as an effective bias on the qubit. Introducing the dilution factor $\beta = C_g/(C_g + C_m)$ and the offset voltage $V_{\text{offset}} = Q_2/2C_g$, we have

$$V_{\text{eff}}(V_b) = \beta(V_b - V_{\text{offset}}) \tag{5.21}$$

Projecting onto the lowest two levels of the fluxonium and invoking the rotating-wave approximation yields the Jaynes–Cummings Hamiltonian,

$$\boldsymbol{H}_{JC}/\hbar = \omega_m\, \boldsymbol{a}^\dagger \boldsymbol{a} + \frac{\omega_q}{2}\, \boldsymbol{\sigma}_z + i\, \frac{\Omega}{2}\, \left(\boldsymbol{a}^\dagger\, \boldsymbol{\sigma}_- - \boldsymbol{\sigma}_+\, \boldsymbol{a}\right), \tag{5.22}$$

where the effective vacuum Rabi frequency is

$$\Omega = \omega_q \left|\langle g|\, \boldsymbol{\varphi} |e\rangle\right| \frac{dC_m^+}{dx}\, \frac{V_{\text{eff}}}{2e}\,. \tag{5.23}$$

Here, ω_q is the qubit transition frequency and $\boldsymbol{\sigma}_\pm$ are the qubit raising and lowering operators; the particular phase convention for $\boldsymbol{\varphi}_q$ has been chosen to simplify later expressions.

5.2.3 Common-mode noise rejection

The antisymmetric bias-coupling scheme intrinsically suppresses voltage noise from the bias line by converting any fluctuation δV_b into an identical perturbation on both gap capacitors C_m^+ and C_m^-. Because the qubit's flux variable Φ_1 corresponds to the differential flux between its two pads, a uniform shift of both electrode potentials does not affect the qubit, whereas the mechanical antisymmetric mode involves opposite potentials on the two electrodes and therefore couples to the qubit. This can be seen from the symmetry of the system. Under left–right exchange the common-mode noise is invariant (even), while both $\boldsymbol{n}$ and $\boldsymbol{x}$ change sign (odd), so an even perturbation cannot drive these odd operators; thus neither the qubit nor the mechanical mode experiences first-order dephasing or excitation from δV_b. Consequently, the coupling Hamiltonian retains only the differential component and can be written as

$$H_{\text{coupling}} \;\propto\; \left(V_b + \delta V_b\right) \boldsymbol{n}\, \boldsymbol{x},$$

with no first-order common-mode noise term. Crucially, the Hamiltonian also lacks the linear noise–driven mechanical drive term $\propto x\, V_b\, \delta V_b$ that plagued the symmetric-mode scheme (see Sec. 5.1).

If the left and right ground capacitors differ by a small amount ΔC_g, a

residual common-mode noise term appears, linear in this asymmetry:

$$-\frac{4\,C_m\,E_C}{e\,(C_m+C_g)}\,\Delta C_g\,(V_b+\delta V_b)\,\boldsymbol{n},\quad \Delta C_g \equiv \frac{C_{g,L}-C_{g,R}}{2}.$$

5.3 Design Workflow for a QED Circuit

Before delving into the specifics of the design, it is worth providing an overview of our design workflow, highlighting the tools and methods used.

The design is created using the Python library *HFSSdrawpy*, which enables script-based circuit generation. This allows parameterization of the circuit, simplifying parameter sweeps and automatic redrawing. The library outputs either a *.gds* file for fabrication or a design in Ansys Multiphysics software.

We begin by simulating the capacitance of various components using Ansys Q3D Extractor to evaluate E_C. Subsequently, we perform eigenmode simulations using Ansys HFSS (High Frequency Structure Simulator) to determine coupling values and losses. The qubit junctions are not optimized in these simulations, since their properties depend not only on geometry but also on the fabrication process. Therefore, they are excluded from Q3D simulations and modeled as linear inductances in HFSS simulations. The coupling rates between the resonators (qubit, readout resonator, and Purcell filter) are extracted from an eigenmode simulation in which the frequency of the qubit plasmon resonance is swept across that of the readout resonator, and the coupling strength is obtained by fitting the anticrossing behavior of the eigenfrequencies.

The final fluxonium chip design is shown in Fig. 5.4. It consists of the fluxonium qubit in the center, coupled differently to control and readout lines. In the following, we will describe this in more details. It also includes (i) test junctions to enable room-temperature junctions measurements (this technique described in Sec. 6.2.5); (ii) 600 nm pillars to ensure separation between the fluxonium and the membrane; (iii) "plasma test" measurement line which consits of the microwave line capacitively coupled to the both-size-opened 150 junctions chain, allowing to measure the plasma frequency of junction (purly fabrication property) ω_p independently. The fluxonium chip is coupled to the membrane via a flip-chip assembly, by putting the membrane at the pillars and creating the device illustrated in Fig. 5.3.

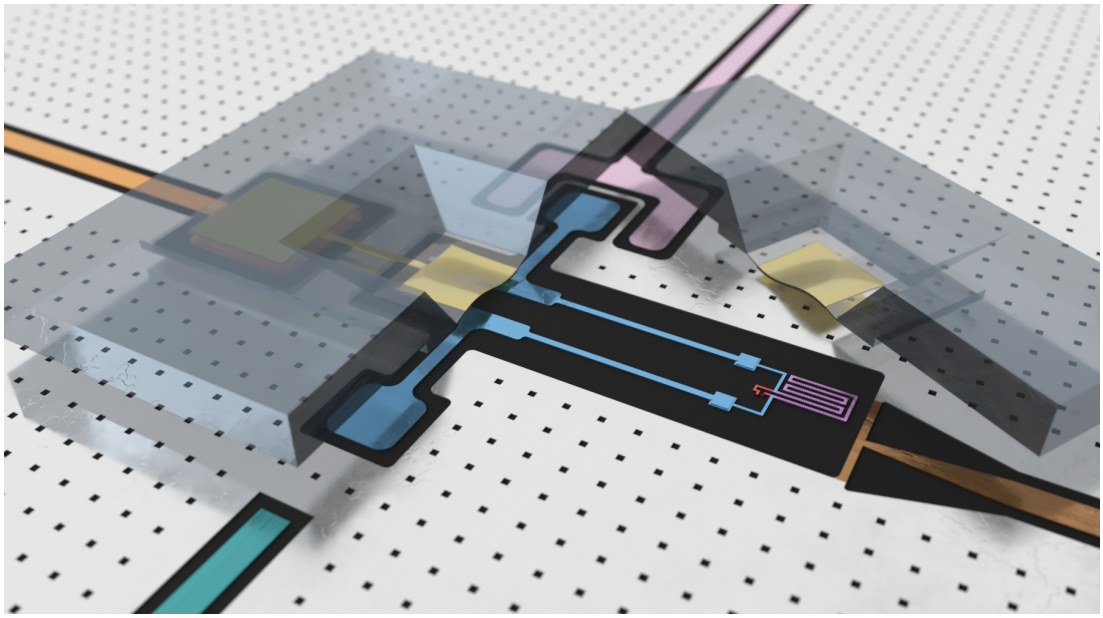

Fig. 5.3: 3D schematic of the mechanical-fluxonium device. A superconducting loop composed of a small Josephson junction (red) and a superinductor (purple), is threaded by an external flux φ_{ext} (brown). The small junction is shunted by a capacitor (blue). One capacitor electrode couples to a readout resonator (light pink) and the other to a control waveguide (cyan). The suspended SiN membrane, metallized with Al (yellow), and the two fluxonium electrodes form a vacuum-gap capacitor. A charge gate (orange) provides DC bias and drives to the membrane. For illustration purposes, the membrane substrate size has been shrunk, and the aspect ratio is not to scale. Additionally, the pillars are not shown as they are far away from the qubit.

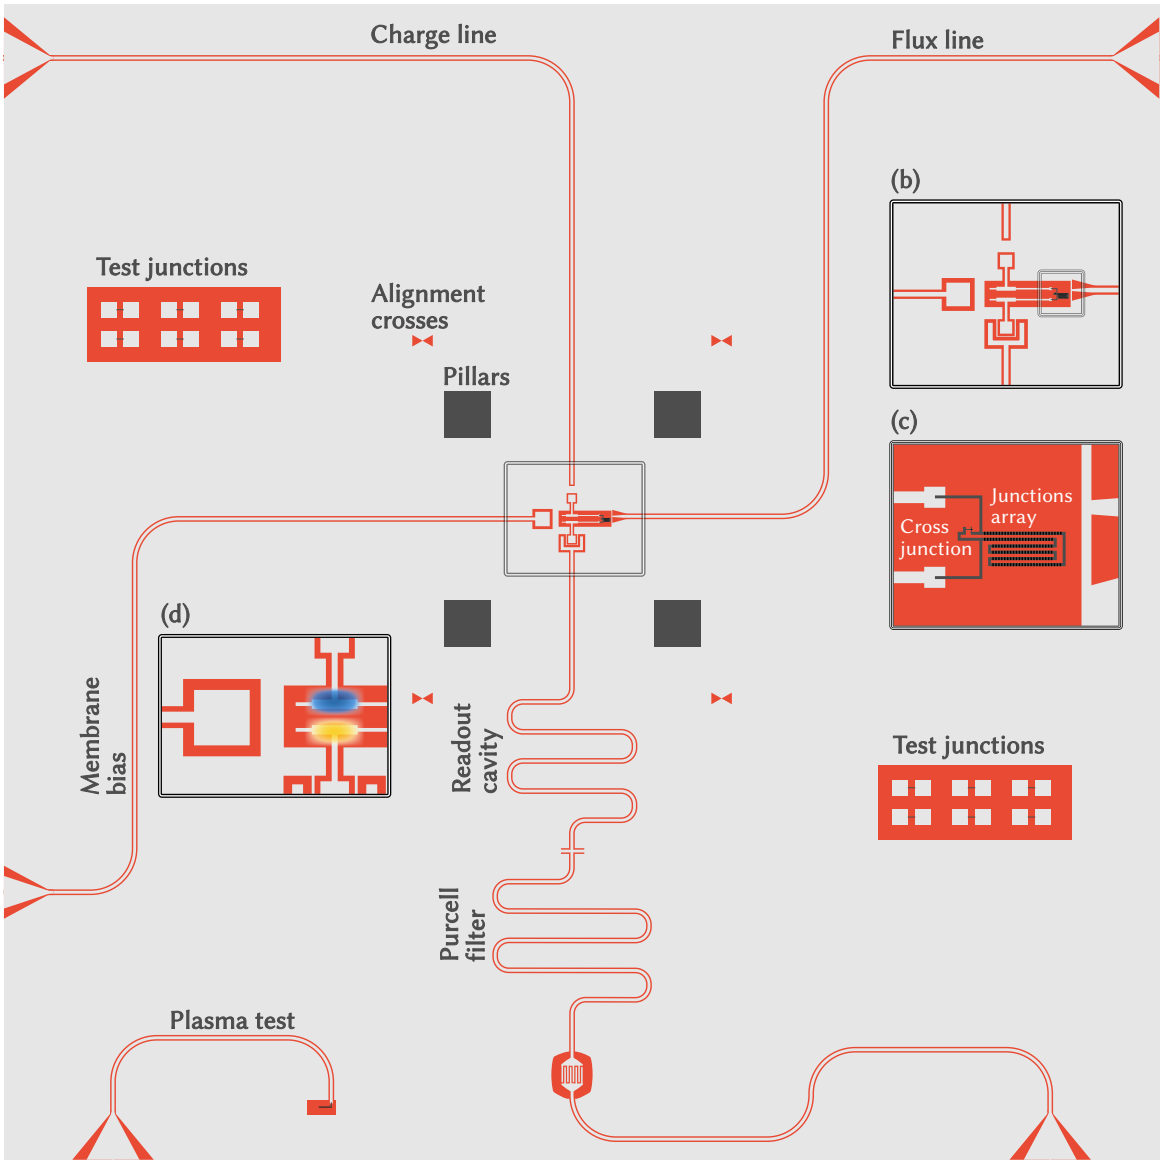

Fig. 5.4: Overview of the chip design including the fluxonium qubit (b); Josephson junction array with cross junction (c); qubit-membrane coupling alignment, where the membrane profile is illustrated with a blue-yellow gradient (d); charge and flux lines, membrane bias line, Purcell filter, readout cavity, and test structures (test junctions, alignment crosses, and plasma test point). Pillars limit the flip-chip distance.

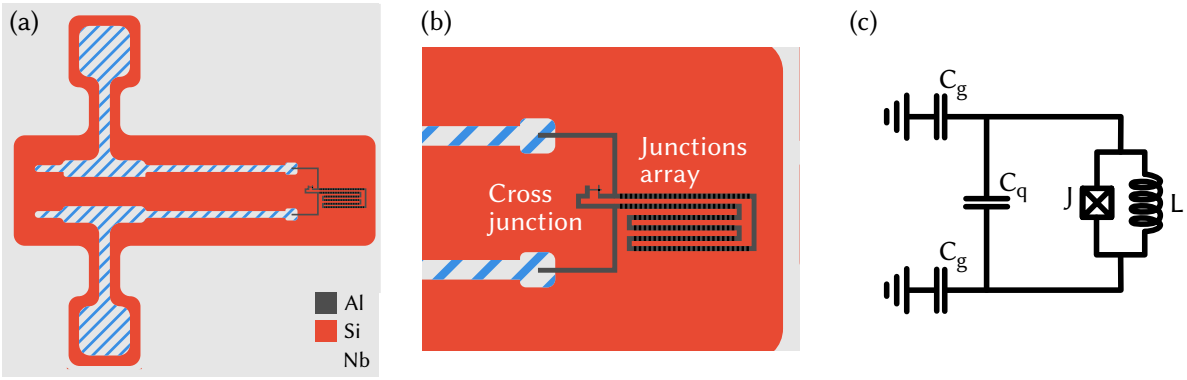

Fig. 5.5: Illustration of the isolated fluxonium qubit design. (a) Overall layout. **(b)** Zoomed-in view of the cross-junction and chain of junctions. **(c)** Lumped-element equivalent circuit.

5.4 Isolated Fluxonium Design Overview

The isolated fluxonium qubit consists of a single Josephson junction (J), an inductance (L), and a qubit capacitance (C_q), as shown in Fig. 5.5(c). The inductance is implemented using an array of Josephson junctions, depicted in Fig. 5.5(b). With a sufficiently large number of identical junctions, the flux distributes evenly, allowing each junction to operate in a regime where it behaves linearly, effectively acting as an inductor.

To enable coupling with the antisymmetric vibrational mode of the membrane, the qubit design incorporates a floating ground configuration. This introduces an additional capacitance to ground, denoted as C_g. The qubit capacitance, C_q, is formed by two large pads on opposite sides of the Josephson junction array, as illustrated in Fig. 5.5(a) with diagonal blue hatching. These pads feature extended sections that establish mutual capacitance between them, while also protruding into the ground plane, thereby contributing to C_g. Consequently, the qubit's total capacitance includes both C_q (between the pads) and C_g (qubit to ground) (Eq. 5.19)

To achieve the desired qubit–membrane interaction, the qubit pads are aligned with the antinodes of the membrane's first antisymmetric mode, as shown in Fig. 5.6(a). The membrane is suspended above the chip and biased via a dedicated DC pad. The top view in Fig. 5.6(c) illustrates the membrane's placement above the chip, while Fig. 5.6(b) highlights the metallized region of the membrane in close proximity to the qubit pads.

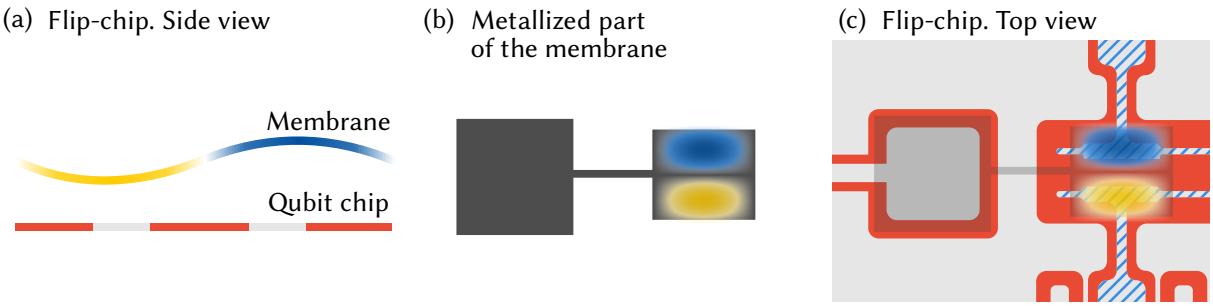

Fig. 5.6: Coupling scheme between the membrane and the qubit. (a) Side view of the flip-chip assembly, showing the membrane above the qubit chip. Hatched areas illustrate the capacitor electrodes of the qubit. (b) Metallized section of the membrane, consisting of the SiN region where vibrations occur and a bias pad where the Si substrate remains. (c) Top view of the pad-to-membrane coupling arrangement.

In our floating-fluxonium design (Fig. 5.7), the voltage across the vacuum-gap capacitors $C_m^{\pm}$ is not simply the applied bias V_b. Instead, V_b is divided across the series combination of the bias-pad capacitance C_b, the pair of membrane capacitors $2C_m$, and the two ground-reference capacitors $2C_g$. The resulting dilution factor is

$$\beta = \frac{C_g}{C_g + C_m + \frac{2C_g C_m}{C_b}},$$

so that $V_{\text{eff}} = \beta V_b$ (Eq. 5.20). Maximizing C_b therefore rigidly ties the membrane potential to the bias line and is always beneficial. By similar logic, one might attempt to increase C_g to hold the fluxonium island at ground. However, increasing C_g also reduces the capacitive energy of the fluxonium, given by Eq. 5.19. Imposing $\partial g/\partial C_m = 0$ at fixed E_C shows that the optimum occurs when C_g and C_m contribute equally to the total capacitance ($C_g = C_m$), while minimizing C_q increases C_g and C_m for a fixed E_C.

We used Ansys Q3D simulations to optimize the geometry of the qubit capacitance pads. We fixed the vacuum-gap $d = 600$ nm in our design and aimed to tune the qubit frequency just below the membrane's mode frequency of approximately 4 MHz. However, in the experiment the distance d was found to vary between 1 and 3 μm [Pat25], so the final qubit–membrane chip did not reach the optimal coupling. We attribute this larger-than-expected separation to thermal bending of the fluxonium chip and mechanical bending of the membrane after chip release. Improved control of d during fabrication therefore provides a straightforward path to enhance coupling in future experiments.

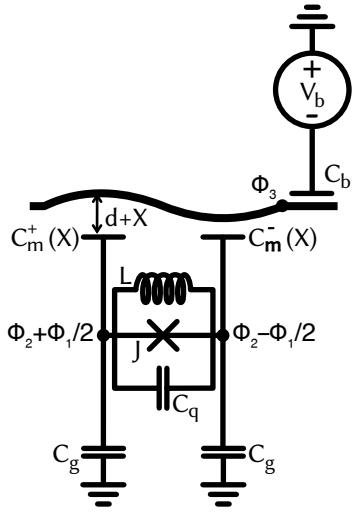

Fig. 5.7: Antisymmetric qubit-membrane coupling circuit. Duplicate of Fig. 5.2 for the sake of clarity.

[Pat25]: Patange (2025), 'RF Transduction in a DC-Biased Superconducting Electromechanical System', Thesis

5.5 Qubit Control and Coupling

5.5.1 Charge drive line

To control the qubit via microwave signals, the qubit is capacitively coupled to a drive line that connects directly to the external control electronics, as theoretically described in Sec. 3.2.1. The coupling strength

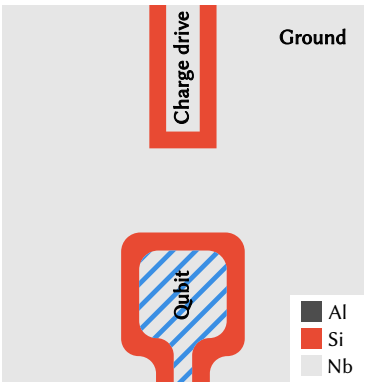

Fig. 5.8: Zoomed view of the charge coupling. The charge drive line is capacitively coupled to the qubit.

of this line must exceed the qubit's decay rate to ensure efficient manipulation, while excessive coupling can introduce additional losses. Using Q3D simulations, we engineered the drive-line capacitance to $C_d = 0.2$ fF, which leads to a plasmon-like mode decay rate of $\kappa = 2\pi \times (926\ \text{Hz})$, as determined by eigenmode HFSS simulations. This decay channel does not become the dominant loss mechanism for the qubit, as discussed in Sec. 3.3.6.

5.5.2 Flux line

Slow variations of the external flux φ_{ext} are applied using an external coil placed near the sample, as shown in the following Fig. 6.27. Rapid flux changes, however, are achieved via an on-chip flux line inductively coupled to the qubit, as described in Sec. 3.2.2. The flux line terminates close to the Josephson junction array, with its current splitting into two opposite directions to ground, as illustrated in Fig. 5.9. To generate a net magnetic flux through the array, the flux-line termination is made asymmetric with respect to the array.

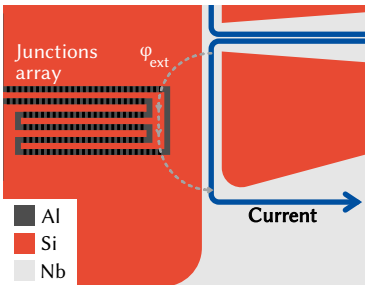

Fig. 5.9: Zoomed view of the flux coupling. The flux line is inductively coupled to the qubit.

As with the charge drive, a trade-off arises: weaker coupling requires higher currents, which can cause chip heating, while stronger coupling increases qubit losses. Eigenmode HFSS simulations for this design yield a plasmon-like mode decay rate of $\kappa = 2\pi \times (729\ \text{Hz})$, which is far from being a dominant loss mechanism.

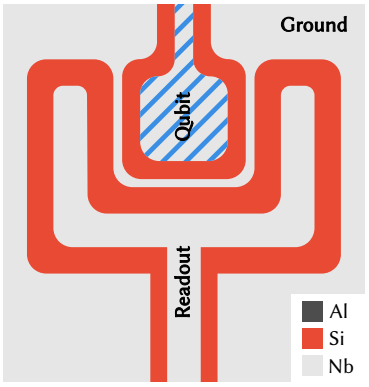

Fig. 5.10: Zoomed view of the readout coupling. The readout line is capacitively coupled to the qubit.

5.5.3 Readout line

Readout resonator. The qubit state is measured using an LC resonator, whose frequency shifts in response to the qubit state, as described in Sec. 7.5. The readout resonator is implemented as a $\lambda/4$ coplanar-waveguide resonator capacitively coupled to the qubit.

Unlike typical designs, where such long lines are wire-bonded along their length, wire bonds near the qubit are not feasible in our case due to the presence of the membrane above the chip. To prevent the division of the ground plane near the qubit and minimize potential differences, the design includes a ground-plane connection between the qubit's capacitor pad and the readout CPW resonator. The readout is designed for strong coupling via a capacitance of 4.6 fF (simulated using Q3D), resulting in a coupling strength of $g_r/2\pi = 43$ MHz, obtained by fitting the avoided-crossing of the eigenfrequencies[1] from HFSS eigenmode simulations.

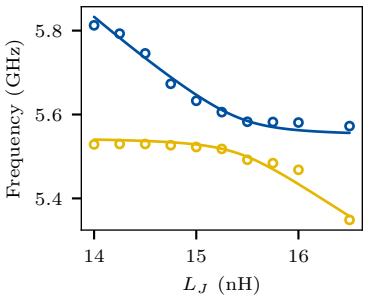

Fig. 5.11: Avoided crossing of the qubit and the readout resonator. The hybridized eigenfrequencies $\omega_\pm$ exhibit level repulsion as the effective linear inductance of the junction is varied. As this inductance is tuned, the plasmon-like mode frequency varies and an avoided crossing with the readout resonator mode is observed, characterized by a minimum frequency splitting of $2g$.

Purcell Filter. To achieve both fast, high-contrast readout and strong qubit protection, we insert a *Purcell filter*—a secondary resonator with frequency ω_f—between the readout resonator and the transmission line. We choose $\omega_f \approx \omega_r$ but ensure it is far detuned from the qubit frequency ω_q. As a result, the readout mode decays rapidly through the filter at ω_r, whereas decay at ω_q is strongly suppressed.

We model the readout mode α and the filter mode β as coupled harmonic oscillators. In a frame rotating at the drive frequency ω_d, their equations of motion are [SMK15, Sec. III]:

$$\begin{aligned}\dot{\alpha} &= -i\,\Delta_{rd}\,\alpha - i\,G\,\beta - i\,\epsilon_r,\\ \dot{\beta} &= -i\,\Delta_{fd}\,\beta - i\,G^*\,\alpha - \frac{\kappa_f}{2}\,\beta,\end{aligned} \tag{5.24}$$

where

$$\Delta_{rd} = \omega_r - \omega_d, \quad \Delta_{fd} = \omega_f - \omega_d,$$

κ_f is the linewidth of the filter resonator, G is the readout-filter coupling rate, and ϵ_r is the readout drive amplitude. In the adiabatic limit $\dot{\beta} = 0$, i.e. when the filter mode damps (and dephases) much faster than it is driven by the readout mode ($\kappa_f \gg |G|, |\Delta_{fd}|$), the filter amplitude follows

$$\beta = \frac{-i\,G^*}{\kappa_f/2 + i\,\Delta_{fd}}\,\alpha.$$

Substituting into Eq. 5.24 yields an effective single-mode equation for α:

$$\dot{\alpha} = -i\big(\Delta_{rd} + \delta\omega_r\big)\,\alpha - \frac{\kappa_r}{2}\,\alpha - i\,\epsilon_r,$$

which describes the readout-resonator decoherence modified by the filter. Here the filter induces both a frequency shift $\delta\omega_r$ and an additional, frequency-dependent decay κ_r of the readout mode:

$$\kappa_r(\omega_d) = \frac{4\,|G|^2}{\kappa_f}\,\frac{1}{1 + \big(2\,\Delta_{fd}/\kappa_f\big)^2}, \tag{5.25}$$

The Fig. 5.12 shows κ_r as a function of the drive frequency ω_d. Physically, κ_r is the rate at which the readout mode leaks through the filter into the line. It peaks when $\omega_d \approx \omega_f$, enabling fast readout, but is strongly suppressed when ω_d is far detuned from ω_f, thereby protecting the qubit. We quantify the filter's protection of the qubit by the *Purcell factor*:

$$F = \frac{\kappa_r(\omega_q)}{\kappa_r(\omega_r)} = \frac{1 + \big[2\,(\omega_r - \omega_f)/\kappa_f\big]^2}{1 + \big[2\,(\omega_q - \omega_f)/\kappa_f\big]^2} \ll 1,$$

demonstrating that decay at the qubit frequency is strongly attenuated while decay at the readout frequency remains high. The reflection coefficient S_{11} through the filter-resonator is given in [Wal+17, Eq. A.1].

The Purcell filter is implemented as a $\lambda/4$ coplanar-waveguide resonator that is inductively coupled to the readout resonator with coupling rate $g_{rp} = 2\pi \times (14\ \text{MHz})$ (simulated using HFSS simulations, see Sec. 7.5) and capacitively coupled to the environment via a 50 fF capacitor (simulated using Q3D). HFSS simulations yield a Purcell-filter mode decay rate of $\kappa_f = 2\pi \times (29\ \text{MHz})$, a readout-resonator mode decay rate of $\kappa_r(\omega_r) = 2\pi \times (23\ \text{MHz})$, and a direct plasmon-qubit mode decay rate of $\kappa_r(\omega_q) = 2\pi \times (11\ \text{Hz})$. Thus, the Purcell filter protects the qubit from direct Purcell losses, while the readout resonator still enables fast qubit-state readout.

1: When two resonant modes a_1, a_2 with bare frequencies ω_1, ω_2 are linearly coupled with coupling strength g, they hybridize and avoid crossing. The Hamiltonian

$$\begin{aligned}\boldsymbol{H} &= \hbar\omega_1\big(a_1^\dagger a_1 + \tfrac{1}{2}\big) + \hbar\omega_2\big(a_2^\dagger a_2 + \tfrac{1}{2}\big)\\ &\quad + \hbar g\,(a_1^\dagger a_2 + a_1 a_2^\dagger)\end{aligned}$$

reduces in the single-excitation basis to eigenfrequencies

$$\omega_\pm = \frac{\omega_1 + \omega_2}{2} \pm \sqrt{\Big(\frac{\omega_1 - \omega_2}{2}\Big)^2 + g^2}.$$

Similarly to classical physics, at resonance ($\omega_1 = \omega_2$), the minimum splitting $2g$ is read off from the avoided-crossing plot versus detuning to extract the coupling strength (see Fig. 5.11).

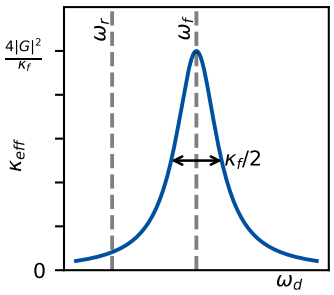

Fig. 5.12: Illustration of the effective leakage rate via the Purcell filter. $\kappa_r(\omega_d)$ vs. drive frequency, showing a Lorentzian centered at ω_f.

[SMK15]: Sete et al. (2015), 'Quantum theory of a bandpass Purcell filter for qubit readout', Physical Review A

[Wal+17]: Walter et al. (2017), 'Rapid High-Fidelity Single-Shot Dispersive Readout of Superconducting Qubits', Physical Review Applied

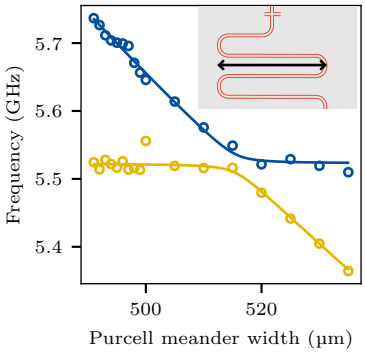

Fig. 5.13: Avoided crossing of the readout resonator and the Purcell filter. To vary the detuning, the width of the meander-shaped Purcell filter is adjusted, which changes its frequency inversely proportional to the width.

Fabrication and experimental setup

6

The fluxonium qubit is fabricated on an intrinsic silicon substrate using standard microfabrication techniques and coupled to a membrane via flip-chip assembly. The resulting mecaflux chip is mounted in a cryostat and controlled by a Field-Programmable Gate Array (FPGA)–based unit. This chapter provides both a general overview and a detailed description of the fabrication process and experimental setup.

My contribution. During my PhD, qubit chips were fabricated first by senior PhD student B.-L. Najera-Santos and later by Dr. L. Balembois. My contribution to the fabrication of the qubit chips mainly involved stepping in to parallelize and speed up the process. The cryostat experimental setup was my primary responsibility from the very beginning to the end of my PhD.

Outline. The first part of this chapter provides foundational insights into junction fabrication, introducing the process for a general reader. We then present a detailed fabrication recipe, outlining specific steps to replicate our approach, and conclude with initial room-temperature measurements of the qubit and flip-chip assembly.

The second part focuses on our experimental setup. First, we overview the essential components, explaining how each element shields the qubit from environmental disturbances and motivates our experimental setup design. Finally, we give a detailed description of the complete setup.

6.1 Josephson junction fabrication

In our qubit circuit, the Josephson junctions serve as a non-linear element E_J and as a building block in the chain that provides large inductance E_L. We begin our explanation with the fabrication process for a single Josephson junction representing E_J and then extend this concept to describe the chain of junctions constituting E_L.

6.1.1 Niemeyer–Dolan bridge technique. Junction geometry

The core of a Josephson junction consists of a three-layer structure, with two superconducting layers separated by an insulating oxide layer (see Sec. 2.2). The energy of the junction is determined by the thickness of the oxide layer and the overlap area between the layers.

Josephson junctions were fabricated using the shadow evaporation technique, specifically the Dolan bridge method [Dol77], in which a suspended bilayer resist "bridge" defines a nanoscale aperture for two angled Al deposits that overlap to form the tunnel barrier. Although

[Dol77]: Dolan (1977), 'Offset masks for lift-off photoprocessing', Applied Physics Letters

the resist overhang is delicate, a well-tuned recipe produces uniform, high-quality junctions. An alternative is the Manhattan overlap technique [Pot+01], which eliminates bridges—trading mechanical robustness for larger junction area. A current PhD student in our group, Paul Manset, is developing a stacked-junction Manhattan variant [Man+25] to shrink the footprint, but for this work we adhere to the established Dolan-bridge process.

[Pot+01]: Potts et al. (2001), 'CMOS compatible fabrication methods for submicron Josephson junction qubits', IEE Proceedings - Science, Measurement and Technology

[Man+25]: Manset et al. (2025), 'Hyperinductance based on stacked Josephson junctions', Preprint

The basic concept of the Dolan bridge is illustrated in Fig. 6.1(a). We create a suspended structure at a certain height above the wafer that serves as a mask. Then, we evaporate aluminium at a 22-degree angle relative to the wafer's normal. This coats the wafer with the first aluminium layer, as illustrated in Fig. 6.1(c, d). This layer is then oxidized by introducing oxygen into the vacuum chamber for several minutes. Next, another aluminium layer is evaporated at a -22-degree angle, as shown in Fig. 6.1(e, f). At the intersection, we obtain the Josephson junction defined by the overlap area and oxide layer.

To create the suspended structure, we use a bilayer resist technique, where the bottom layer is more sensitive to the electron beam than the top layer. This allows the top layer to remain after development, while the bottom layer is removed. To obtain the desired structure shown in Fig. 6.1(a), we can categorize it into the following areas, illustrated in Fig. 6.1(b):

- **Window**: Areas where both layers are removed, allowing evaporation on the wafer. Illustrated with inclined blue lines.
- **Undercut**: Areas where only the bottom layer is removed, with the top layer serving as a mask for subsequent evaporation. Illustrated with inclined white lines.
- **Full protection**: Areas where both layers of resist remain intact, fully protecting the wafer. Illustrated in red.

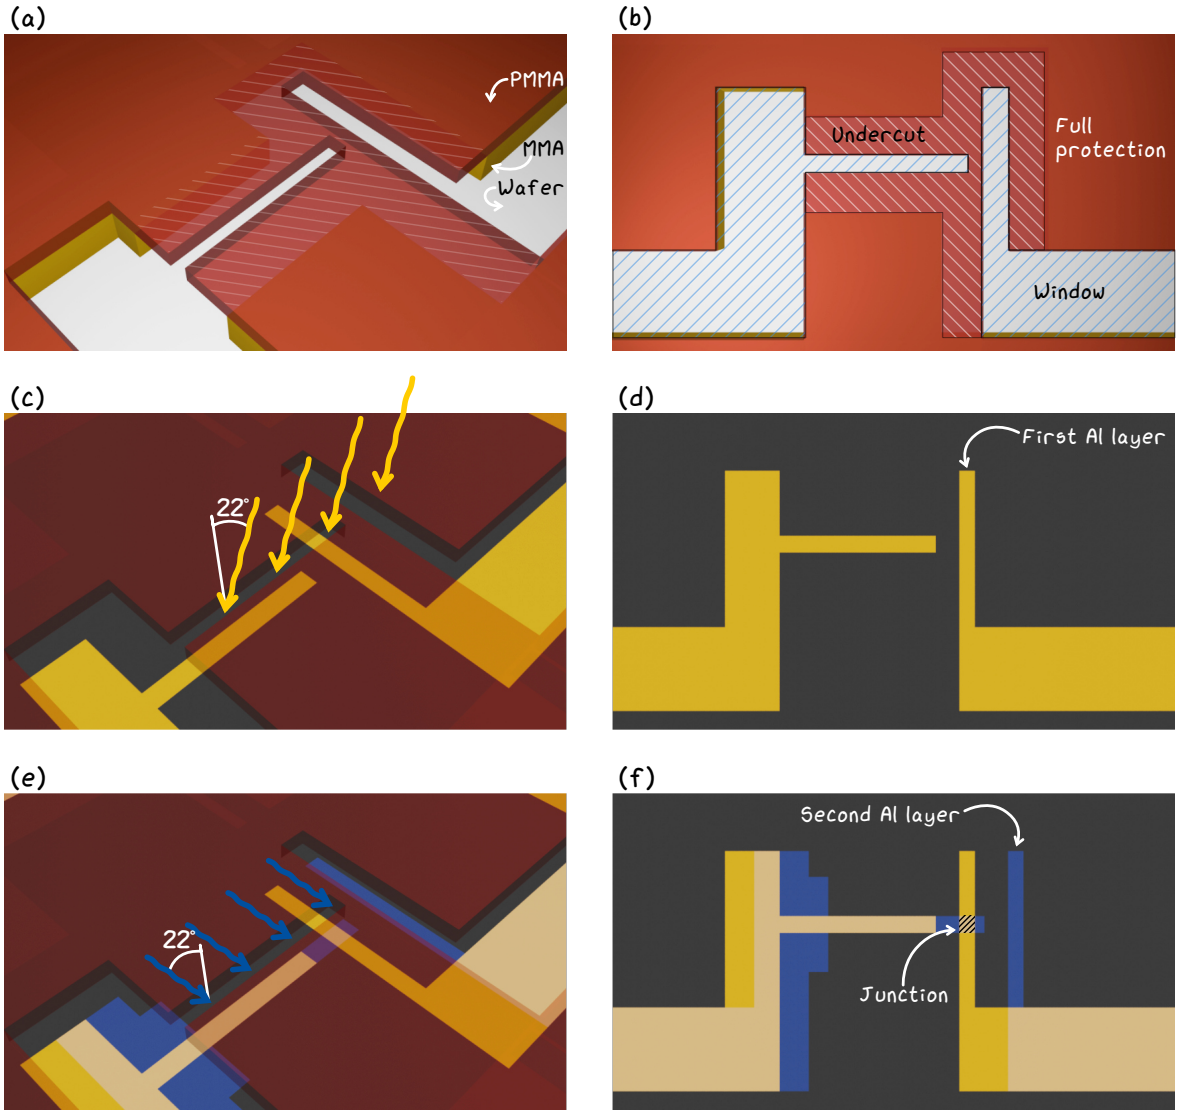

Fig. 6.1: Illustration of a single "cross" Josephson junction fabrication process. (a) presents a 3D view of the mask composed of two resist layers (MMA and PMMA). (b) provides a top view of the mask, with three types of areas highlighted by inclined lines; refer to the main text for further explanation. (c) and (d) show the first layer of evaporated aluminium. (e) and (f) illustrate the second layer of evaporated aluminium and the overlap region with the first oxidised layer, which defines the junction.

This technique allows us to create the main non-linear element of the circuit—a junction characterized by a Josephson energy E_J. It is sometimes referred to as the "cross" or "black sheep" junction. In Fig. 6.2, we provide an example of a typical junction observed under an electron microscope. The two layers of aluminium are shown in yellow and blue. In this image, one can clearly distinguish the two layers, with higher contrast at their intersection.

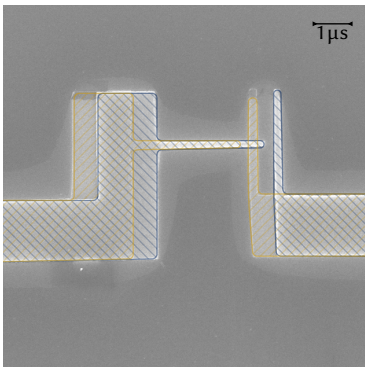

Fig. 6.2: Photo of a typical cross junction taken using an electron microscope. In post-processing, inclined lines were added to illustrate the two layers of evaporated aluminium. The first layer is highlighted in yellow, and the second layer in blue.

For the bilayer technique to work, the exposure dose should be well-calibrated, typically achieved through a dose test chip that is repeated after any fabrication pause. This involves creating multiple junctions with a sweep of exposure doses for each layer, followed by inspection using an electron microscope. For the underexposed junction, the overlap area lacks a well-defined rectangular shape, resembling blots. In the overexposed junction, all angles appear overly smooth. The correctly exposed junction, however, displays a well-defined rectangular shape with sharp, straight edges and right angles.

To ensure that the two "fingers" of each evaporation cross correctly, we require the horizontal finger of the second evaporation to cover and extend beyond the vertical finger of the first evaporation, as illustrated in Fig. 6.3(b). Using basic geometrical reasoning[1], the condition on the mask can be written as:

$$(2h_{\mathrm{MMA}} + h_{\mathrm{PMMA}}) \tan\theta - b - s > 0, \tag{6.1}$$

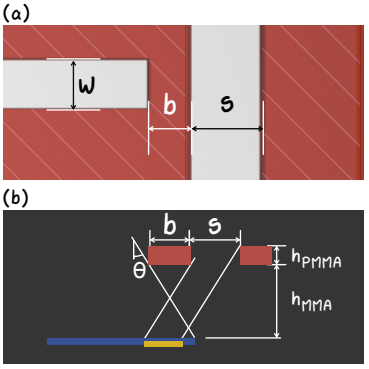

Fig. 6.3: Illustration of the geometrical parameters of a single cross junction. (a) shows the top view, while (b) provides the front view of a cross-section through the middle of subfigure (a). The top PMMA layer is depicted in red, with the first and second evaporation layers shown in yellow and blue, respectively.

where h_{MMA} and h_{PMMA} are the heights of the MMA and PMMA resist layers, respectively, θ is the angle of evaporation, b is the width of the bridge, and s is the width of the vertical "finger", as illustrated in Fig. 6.3(a). From this, we can also derive the expected area of the junction:

$$A = w * (s - h_{\mathrm{PMMA}} \tan\theta), \tag{6.2}$$

where w is the width of the horizontal finger.

1: We also neglect mask cluttering by the first evaporation and edge effects.

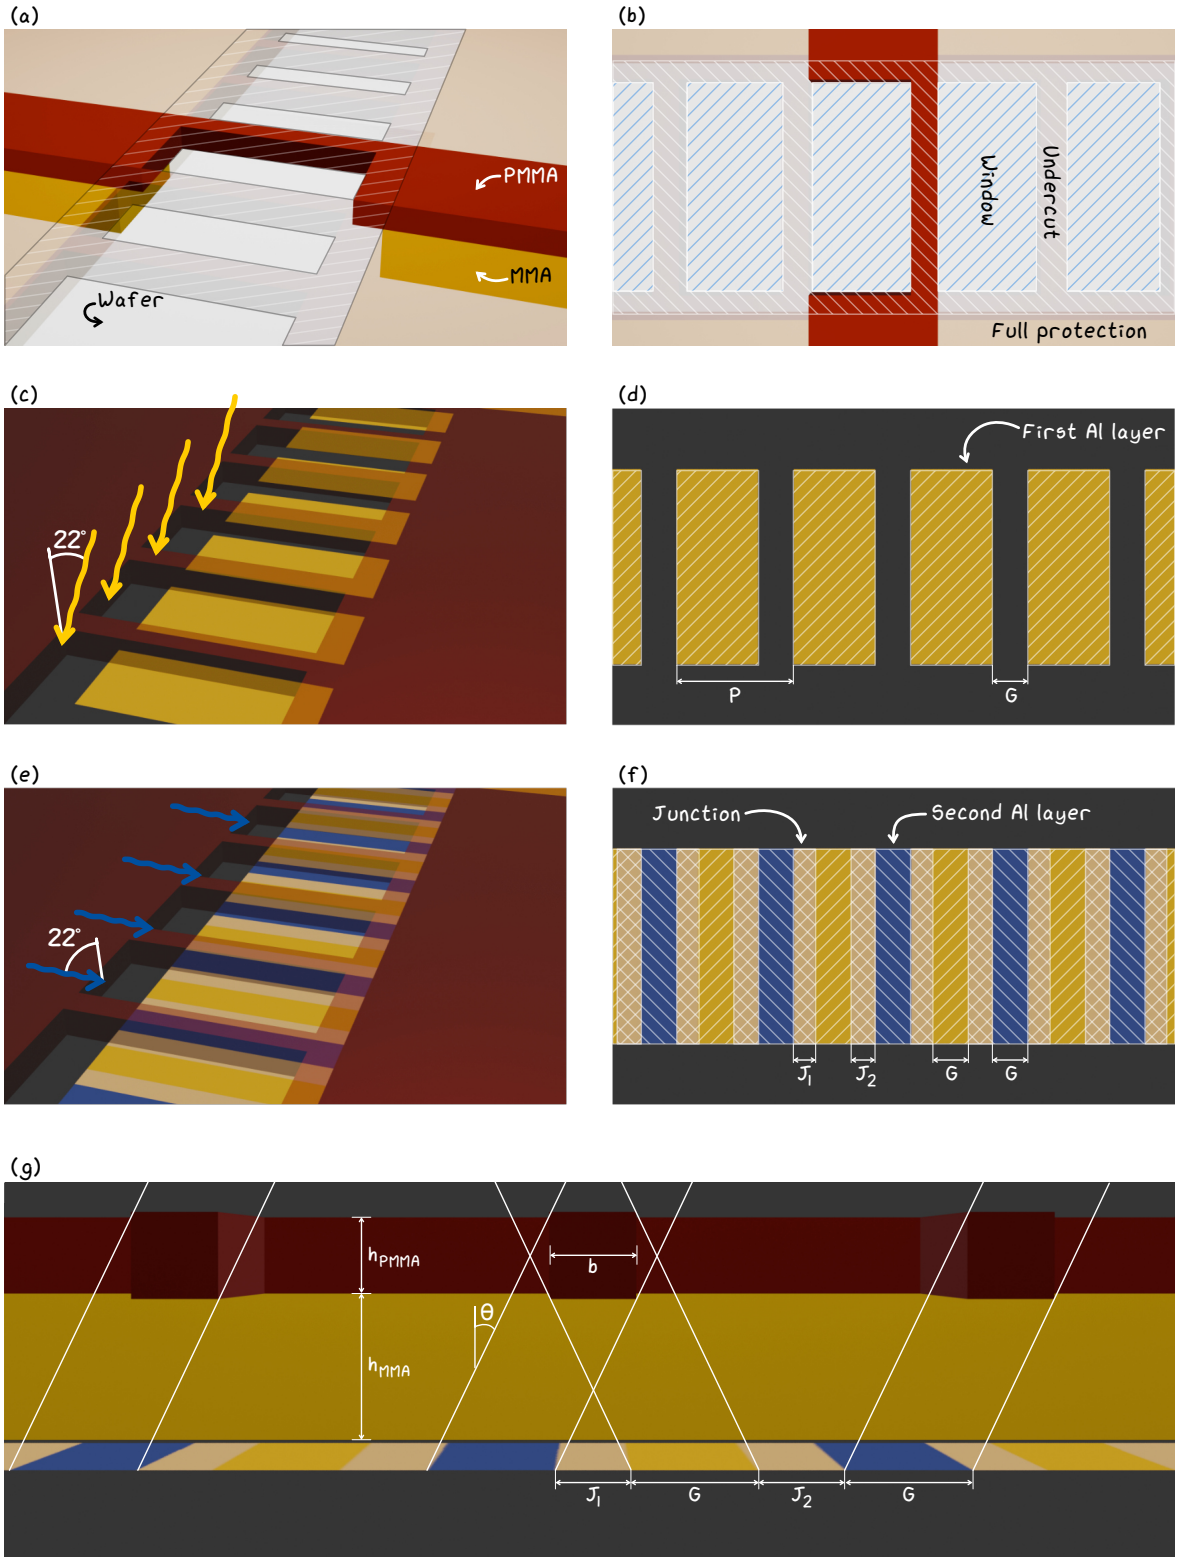

Fig. 6.4: Illustration of a chain of Josephson junction fabrication process. (a) presents a 3D view of the mask composed of two resist layers (MMA and PMMA). (b) provides a top view of the mask, with three types of areas highlighted by inclined lines; refer to the main text for further explanation. (c) and (d) show the first layer of evaporated aluminium. (e) and (f) illustrate the second layer of evaporated aluminium and the overlap region with the first oxidised layer, which defines the junction. (g) displays a front view, highlighting the geometrical parameters of the junction and the mask.

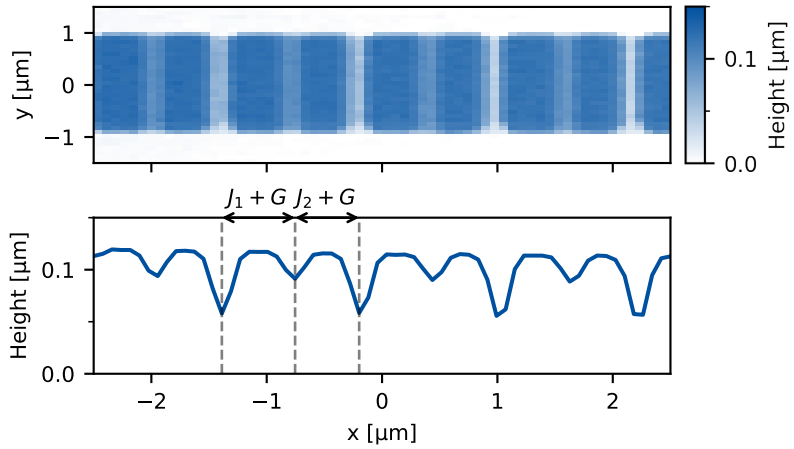

Fig. 6.6: An AFM measurement of a chain of junctions. In the top subfigure, the profile of the surface is plotted, where the junctions are clearly visible. In the bottom subfigure, the mean along the vertical axis is plotted. The wells in the junction profile curve correspond to the gaps in evaporation.

6.1.2 Junction chain fabrication. Large inductance

The superinductance E_L is implemented using N_{bridges} Dolan bridges in the same two-angle shadow-evaporation step as the small junctions described in section 6.1. A mask with pitch P and opening width b defines these bridges. The first Al evaporation at $+22°$ produces a series of aluminium stripes offset from the bridge openings (Fig. 6.5c,d), and after oxidation, the second evaporation at $-22°$ deposits complementary stripes shifted in the opposite direction (Fig. 6.5e,f). Each bridge thus gives rise to two junctions—one centered under the bridge and one between adjacent bridges—yielding $2N_{\text{bridges}} - 1$ junctions in series. It is critical to choose P and b so that all overlap regions have identical areas, preventing any weak link where phase slips could preferentially occur and thereby preserving qubit coherence.

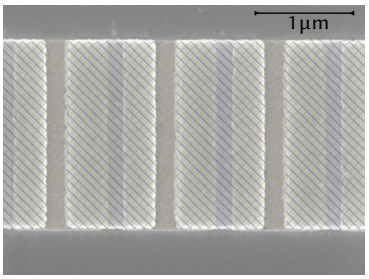

Fig. 6.5: Photo of a typical chain of junctions taken using an electron microscope. In post-processing, inclined lines were added to illustrate the two layers of evaporated Aluminium. The first layer is shown in yellow, and the second in blue.

The final geometrical size of the junctions is determined by the mask itself and the evaporation angle θ. All these parameters are illustrated in Fig. 6.4(g). From this figure, we can derive the following relations for the junction lengths:

$$\begin{aligned} J_1 &= 2h_{\text{MMA}} \tan\theta - b \\ J_2 &= P - b - 2(h_{\text{MMA}} + h_{\text{PMMA}}) \tan\theta \end{aligned} \tag{6.3}$$

where J_1 and J_2 are the junction lengths in the chain, P is the periodicity of the chain, b is the width of the bridge, h_{MMA} and h_{PMMA} are the heights of the MMA and PMMA resist layers, respectively. These relations can be rewritten as:

$$\begin{aligned} \delta J \equiv J_2 - J_1 &= P - 2(h_{\text{PMMA}} + 2h_{\text{MMA}}) \tan\theta \\ J \equiv \frac{J_2 + J_1}{2} &= P/2 - (b + h_{\text{PMMA}} \tan\theta) \end{aligned} \tag{6.4}$$

Therefore, to ensure equal junctions in the chain, we require $\delta J = 0$. This links fabrication parameters, such as h_{MMA} and h_{PMMA}, to design parameters like b and P.

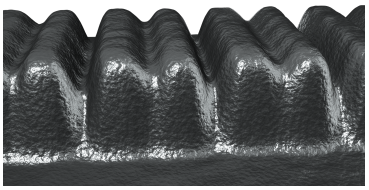

Fig. 6.7: 3D model of a chain of junctions reconstructed from an AFM measurement. The model has been smoothed and visualized with an aluminium texture for illustrative purposes.

To verify that the junctions are uniform, we use plan-view SEM to capture two-dimensional images (Fig. 6.5) and/or AFM topography to measure the surface height profile (Fig. 6.6 and Fig. 6.7). The AFM scan of a junction chain shows alternating wells and raised plateaus: each plateau marks the overlap region of the two angled aluminium depositions, while the wells correspond to single-layer regions from either the first- or second-angle evaporation. Since each evaporation deposits a different metal thickness, these single-layer regions appear at two different heights, providing clear landmarks for measuring the lateral overlap. Comparing the plateau widths confirms that all junction areas match within our fabrication tolerance.

AFM measurements can be preferable when measuring a small number of junctions, as unlike SEM, there is no need to pump the chamber before measurement, and alignment is easier. However, AFM measurements take much longer per measurement if one neglects this preparation time. AFM also provides a non-destructive way to verify the geometry of the junction, in contrast to SEM, where exposure to the electron beam can modify the junctions' properties.

In the next section, we will provide a detailed recipe for the fabrication of the junctions and less complex structures.

6.2 Qubit chip fabrication. Workflow and recipe

The fabrication of a qubit circuit is a multi-step process; however, the general approach for each step can be summarized as follows: we protect selected areas with resist, perform a process over the entire wafer, and then remove the resist. The technical details are typically defined by material deposition limitations.

The qubit circuit fabrication process can be broadly divided into three main stages, as illustrated in Fig. 6.8:

1. **Fabrication of larger Niobium-deposited structures**: larger structures like capacitor pads and coplanar waveguide resonators are made with Niobium. Niobium (Nb) is deposited across the entire wafer. Then, using resist and UV-laser lithography, we create a mask to protect areas where Nb should remain. Nb is then removed from non-target areas using reactive ion etching (RIE).
2. **Fabrication of larger Aluminium-deposited structures**: pillars for a membrane. The wafer is coated with a bilayer of resist to minimise edge effect during the subsequent lift-off process. Then, using UV-laser lithography, a mask is created to define non-target areas. Aluminium is evaporated across the wafer and subsequently removed from non-target areas by a lift-off process, thanks to the mask.
3. **Fabrication of smaller Aluminium structures**: Josephson junctions and their interconnections. This follows the process described above in Sec. 6.1. A suspended structure is fabricated using a bilayer resist. Aluminum is then evaporated at different angles with controlled oxidation between deposition steps, and finally the resist and excess aluminum are removed.

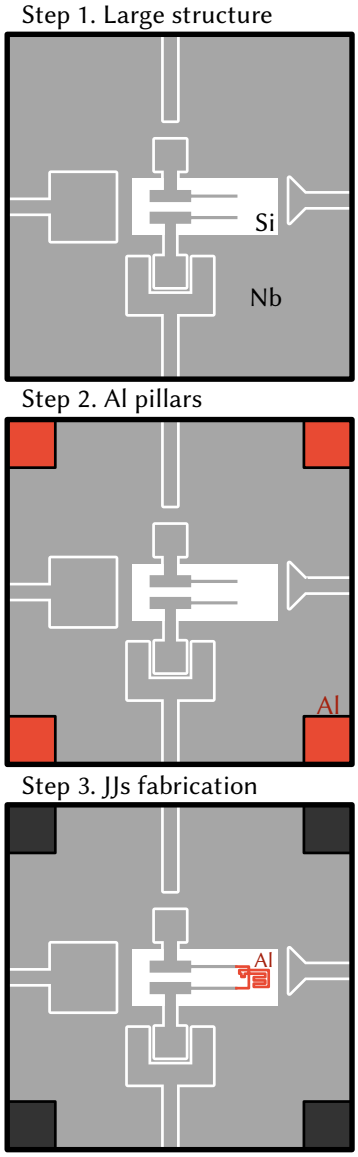

Fig. 6.8: Overview of the qubit chip fabrication process. The mask geometry highlights the main features, which have been simplified and rescaled for better visibility.

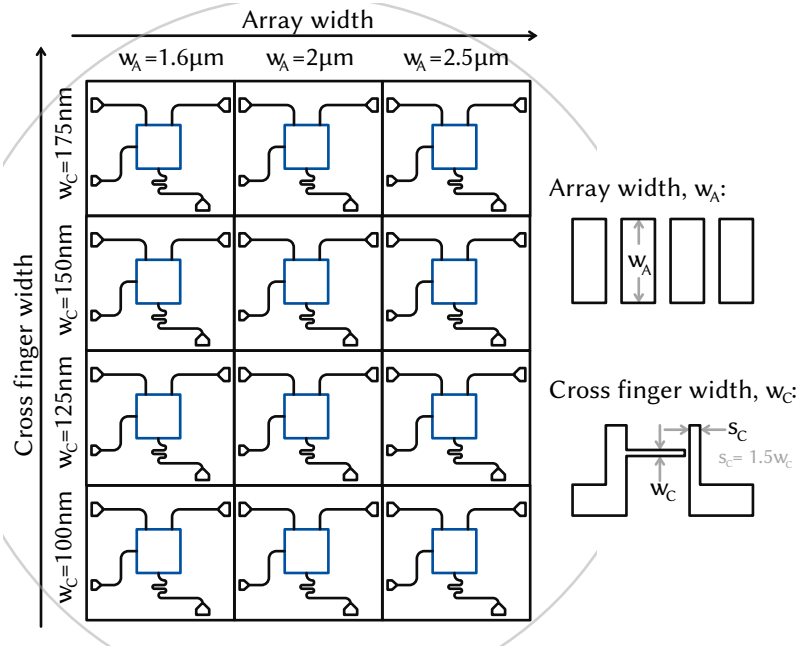

Fig. 6.9: Illustration of the chip arrangement on the wafer. The left subfigure shows 12 chips distributed across the wafer. In the vertical direction, the finger size of the cross junction w_c is varied, while in the horizontal direction, the array width w_a is varied. The grey circle represents the wafer's edge. On the right, a reminder of the parameters w_c and w_a is provided for reference.

For details on the original debugging and development of the fabrication process, see Najera-Santos's thesis [San24, Chap. 4]. The procedure has since been updated—most notably with a revised pillar design—and below we outline our current workflow, including all relevant parameters.

6.2.1 Wafer scale fabrication

To reduce the influence of the fabrication uncertainties and increase the yield of the fabrication process, we fabricate 12 (10.5 μ×9.5 µm) chips on a single 2-inch wafer at once. Each chip contains a slightly different junction geometry, allowing us to select the chip with the desired parameters after the fabrication process.

In Fig. 6.9, we illustrate the arrangement of the chips on the wafer. In the horizontal direction, E_L is varied by changing the array width w_A, as the final energy is proportional to the junction area. w_A is varied from 1.6 μm to 2.5 μm. In the vertical direction, E_J is varied by changing both finger sizes of the cross junction w_C, s_C. The horizontal finger size, w_C, is varied from 100 nm to 175 nm, while the vertical finger size, s_C, is kept at $s_C = 1.5 \times w_C$. This leads to around 25% variation in E_J and E_L, which is sufficient to compensate for our fabrication uncertainties

In Fig. 6.10, we provide an example of the energy distribution of the 12 chips estimated at room temperature by measuring the resistance of the junctions (method described in Sec. 6.2.5). Due to the uniformity limitations in fabrication, the energy of junctions with the same nominal geometrical parameters can vary. In this example of our recent wafer, there is approximately a 2% variation in the nominally identical E_L and a 5% variation in the nominally identical E_J.

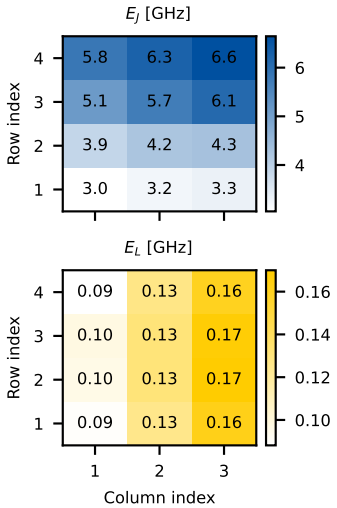

Fig. 6.10: E_J and E_L energy distribution of the 12 chips on the wafer estimated by resistance measurements. The text inside each square corresponds to the corresponding energy. Energy is in GHz units. Fabrication uncertainties cause nominally identical design parameters to yield different energies, resulting in variations of approximately 5% in E_J and 3% in E_L. Name of the wafer: *MF5*.

6.2.2 Stage 1. Fabrication of structures over a micrometer: UV-Laser Lithography

In this stage, we deposit Niobium (Nb) across the entire wafer, and then use UV-laser lithography to create a mask and remove Nb from non-target areas. By the end of this step, we have a silicon wafer with large circuit parts made of Nb. This procedure closely resembles the process used for SiN membrane fabrication.

The steps of this stage are illustrated in Fig. 6.11.

0. **Starting wafer**: The process begins with a 280-μm-thick silicon wafer with (100)-orientation and resistivity greater than 20 kΩ cm^2. 10% Hydrofluoric acid (HF) cleaning is performed for 10 minutes to remove natural silicon dioxide layer. It reduces[Mel+20, Table. 1] the silicon surface loss tangent, tan δ (defined in Eq. 3.32), thereby decreasing dielectric losses.
1. **Niobium sputtering**: A uniform layer of niobium (150 nm thick) is deposited on top of the silicon wafer using sputtering.
2. **Wafer cleaning**: The wafer is then cleaned in a warm acetone bath at 50°C, subjected to ultrasonic vibrations for approximately 10 minutes. It is then transferred to an isopropyl alcohol (IPA) solution for 1 minute of ultrasonic bath. The wafer is dried with pressurized nitrogen. If the wafer's cleanliness is unsatisfactory, brief oxygen plasma cleaning can be performed for a couple of minutes, followed by repeating the bath cleaning steps.
3. **Photoresist spin coating and baking**: A layer of UV-reactive negative resist (S1805) is uniformly spread across the entire niobium-coated wafer using spin coating at 4000 RPM for 1 minute to achieve an expected thickness of 500 nm. The resist is baked at 115°C for 1 minute.
4. **Photoresist exposure: UV-laser lithography**: The areas where Niobium (Nb) should remain for the final circuit are exposed to UV lithography machine (*Microtech Laserwriter LW405D*) at a dose of 100 mJ/cm^2.
5. **Photoresist development**: The exposed resist is developed by immersing the wafer in the MIF-319 developer for 1 minute with gentle swirling. The wafer is then rinsed in water for 1 minute.
6. **Dry-etching of Nb**: The niobium not covered by photoresist is removed via reactive ion etching (RIE) using SF_6 plasma, with an additional 10% of over-etching to ensure complete removal of small Nb parts.
7. **Cleaning remaining resist**: The remaining resist is cleaned using O_2 plasma for a few minutes. The wafer is then subjected to a double bath cleaning using fresh acetone and IPA to ensure it is clean for the next stage.

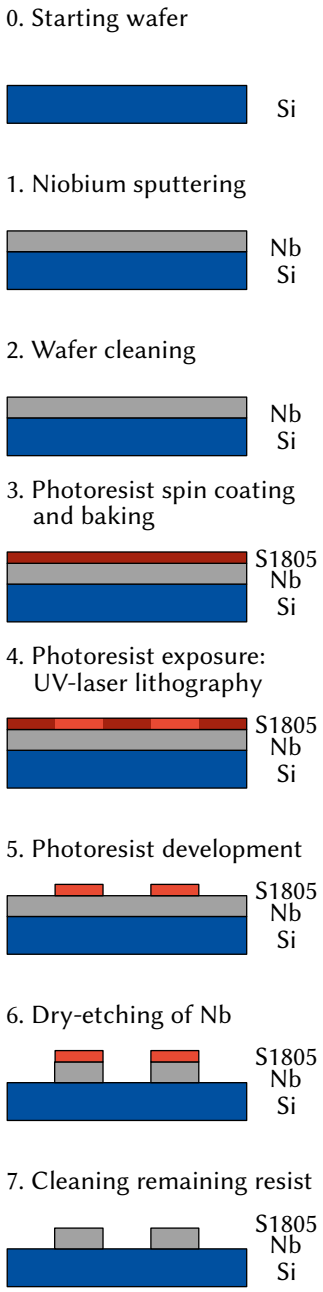

Fig. 6.11: Outline of the UV-laser lithography process for the qubit chip. See the main text for details.

6.2.3 Stage 2. Aluminium pillars deposition

In this step, we deposit the aluminium layer for the pillars.

The steps of this stage are illustrated in Fig. 6.12.

0. **Starting point**: The process begins with a clean wafer from the previous stage.
1. **Bilayer photoresist deposition**: First, a positive photoresist *LOR20B* is deposited using spin coating at a 5500 RPM for 60 seconds (including 5 seconds of acceleration) in order to achieve an expected height of 1.2 μm. It's then baked for 5 minutes at 165 °C A second layer of positive photoresist *S1805* is deposited on top using the same spin coating technique at a speed of 4000 RPM for 60 seconds in order to achieve an expected height of 500 nm. It's baked for 1 minutes at 115 °C.
2. **Photoresist exposure: UV-laser lithography**: Since the goal is to deposit Aluminium (Al) in specific areas to form pillars, these areas are exposed to a ultraviolet laser, specifically the *Laserwriter LW405D from Microtech*, delivering a dose of 100 mJ/cm^2.
3. **Photoresist development**: The exposed resist is developed by immersing the wafer in the *MF-319* (*TMAH* based) developer and gently swirling the beaker for one minute. The wafer is then rinsed in water for one minute.
4. **Aluminium evaporation**: A 600 nm layer of aluminium is evaporated using a *Plassys MEB550 e-beam evaporator*. A pressure under 10^{-6} mbar is acceptable. Ion milling is not required. The evaporation occurs perpendicular to the surface.
5. **Removal of unwanted aluminium**: The unwanted aluminium is removed by dissolving the remaining resist. For this, the wafer is submerged into NMP solution at room temperature overnight, or, less preferably, by placing it in NMP at 80°C for 30 minutes.
6. **Cleaning**: The cleaning process is the same as step 7 in Stage 1.

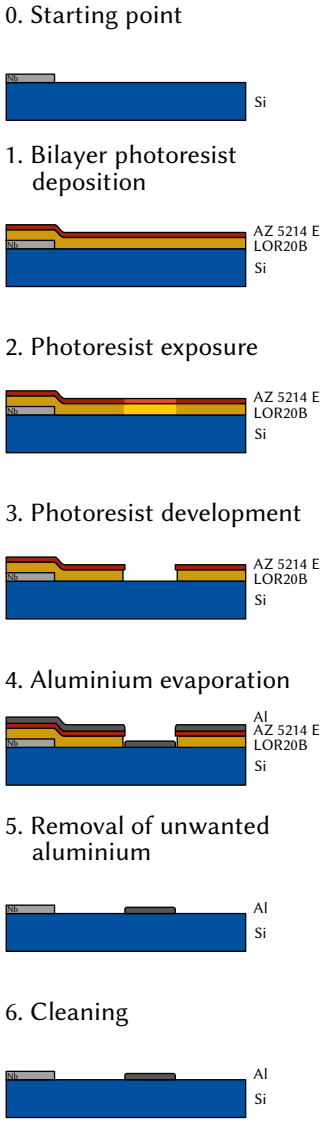

Fig. 6.12: Outline of the pillars fabrication process. The starting point is a wafer with Niobium (Nb) structures (light grey rectangle) fabricated in the previous stage. For more details, refer to the main text.

6.2.4 Stage 3. Josephson Junction Fabrication: Electron beam lithography

In the final fabrication step, we fabricate the Josephson junctions. The principle of the fabrication is explained in Sec. 6.1. Since these structures are small, the corresponding mask is patterned by electron-beam lithography.

0. **Starting point**: The process begins with a clean wafer from the previous stage.
1. **Bilayer electro-resist deposition**: A bilayer deposition is a crucial step in creating Dolan bridges, which is one of the methods of Josephson junction fabrication (Sec. 6.1.1). Initially, a thick layer of the electro-sensitive resist MMA EL13 is deposited using spin coating at a speed of 3000 RPM for 1 minute to achieve an expected height of 700 nm. It is then baked for 1 minute at 195°C. Following this, a thin layer of the electro-sensitive resist PMMA A3 is deposited using the same spin coating method, but at a speed of 4000 RPM for 1 minute to achieve an expected height of 100 nm. It is then baked for 30 minutes at 195°C.
2. **Resist exposure: E-beam lithography**: The necessary areas are exposed using electron beam lithography at a working voltage of 20 kV and pressure below 2×10^{-5} mbar. The base exposure dose is 280 μC/cm^2 with a d-step of 10nm by 10nm. However, not all parts receive the same dose. Dose tests should be conducted to determine the exact dose for different design parts (Sec. 6.2.5).
3. **Resist development**: The exposed resist is developed in a mixture of deionized water (di-H_2O) and IPA at a 1:3 ratio, at 6°C for 90 seconds. Immediately afterward, the wafer is transferred into room temperature IPA and dried with pressurized nitrogen.
4. **Aluminium evaporation**: The two layers of aluminium are evaporated using the *Plassys MEB550* machine. Pumping begins overnight to achieve pressures under 2×10^{-7} mbar. Next, ion milling is performed to remove the native oxide layer on the wafer. Ions are accelerated at 100 V with 35 mA current for 6 seconds. Too strong ion milling can lead to increased decoherence due to dielectric losses [Qui+14; Dun+17]. The first aluminium layer of 35 nm is evaporated at an angle of -22° relative to the wafer's normal vector. This layer is then oxidized by bringing the vacuum chamber to 200 mbar of pure oxygen and waiting for 10 minutes. The chamber is then pumped back to 5×10^{-6} mbar, and a second aluminium layer of 100 nm is evaporated at an angle of 22°. The same oxidation process is performed, in order to have a controlled oxide layer.
5. **Removal of unwanted aluminium**: The unwanted aluminium is removed by dissolving the remaining resist. For this, the wafer is immersed into NMP solution at room temperature overnight, or, less preferably, placed in NMP at 80°C for 30 minutes.

6.2.5 Room temperature junctions test

As soon as the wafer with different chips is fabricated, it is possible to test the junctions and estimate E_J and E_L before putting the chip into the cryostat.

We estimate these energies using the Ambegaokar-Baratoff relation [AB63], which links the normal state junction resistance $R_J^{(T_c)}$ at the critical temperature T_c, the superconducting gap Δ at $T = 0K$, and the resistance quantum for Cooper pair $R_Q^{\mathrm{CP}} = 4R_K = \frac{h}{(2e)^2}$[1]:

$$E_J = \frac{R_Q \Delta}{2R_J^N}$$

When using the room temperature resistance, an empirical factor should be applied, which is calibrated for each fabrication process and clean room. With the process outlined above and the equipment available in 2024 at the Paris ENS clean room, the calibration factor used is $R_J^{(T_c)} = 1.45\, R_J^{(T_{room})}$. This calibration factor captures two main effects: (i) the temperature dependence of the tunnel resistance, since our room-temperature measurement must be scaled to the normal-state resistance at T_c; (ii) the difference between the superconducting gap Δ_{film} of our thin aluminium films and the bulk gap Δ used in the Ambegaokar–Baratoff relation.

Measuring the resistance of the junctions involved in the qubit circuit itself is impractical and could lead to large uncertainties, due to not negligible resistance of the substrate connected in parallel. Therefore, around each chip, we fabricate 6 single-junctions and 6 junction chains (see Fig. 6.13), which we measure using a probe station.

We apply a constant current through the junction/chain at $I_0 = 100\,\mathrm{nA}$ with a compliance of 0.1 V to avoid any possible voltage spikes. We measure the voltage drop V_m across the junction and calculate the resistance $R_m = V_m/I_0$. However, since the Si substrate is also conductive at room temperature, the actual resistance of the junctions must be adjusted for this by measuring the resistance in the absence of junctions (when all current flows through the substrate) and subtracting the resistance of the connecting cables R_C. The equivalent circuit of the experimental setup is illustrated in Fig. 6.14.

Thus, with the junction and substrate R_{sub} connected in parallel, the resistance of the junctions can be found from:

$$R_J = \frac{R_{\mathrm{sub}}(R_m - R_c)}{R_{sub} - (R_m - R_C)}$$

where R_{sub} and R_m are the resistances of the substrate and the total resistance that was measured, respectively.

Thus, even though we are not measuring the actual junctions that participate in the circuit, this method provides us with good estimations of the qubit's single junction energy E_J and the junction chain energy E_L.

[AB63]: Ambegaokar et al. (1963), 'Tunneling Between Superconductors', Physical Review Letters

1: Here are numerical values of the constants needed in the formula:

- $R_Q = 6.454\,\mathrm{k\Omega}$,
- $\Delta(\mathrm{Al30nm}) = 50.3\,\mathrm{GHz} \cdot h$.

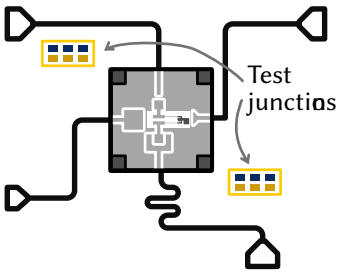

Fig. 6.13: Illustration showing the placement of the test junctions that are fabricated around the qubit. It consists of 6 junction chains illustrated in blue and 6 single junctions illustrated in yellow.

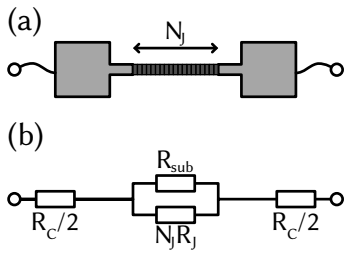

Fig. 6.14: Illustration of the junction probing connection and its corresponding circuit. (a) shows an illustration of a N_J test junctions with pads at the ends that allow contact using pins from a probe station connection. (b) shows an equivalent circuit with substrate resistance R_{sub} in parallel and cable resistance R_C in series.

6.2.6 Flip-Chip Assembly

The membrane and fluxonium chips are assembled by a flip-chip procedure: the membrane chip is placed on top of the fluxonium chip and seated onto its supporting pillars. The flip-chip assembly workflow, established by H. Patange [Pat25, Chap. 5], employs an *MJB4* mask aligner for micropositioning and a custom-built holder. The chips are aligned with a horizontal accuracy of 5 μm before contact. Bonding is achieved by applying drops of *Dymax OP-67-LC* UV-curing glue at the corners of the membrane chip. The inter-chip gap is nominally defined by 600 nm aluminium spacers on the fluxonium chip. Although room-temperature Fizeau interferometry on unmetalized test samples confirms this spacing, cryogenic qubit spectroscopy consistently reveals a gap of 1–3 μm [Pat25], which we attribute to differential bending under thermo-mechanical stress.

6.3 Practices for Cryogenic Microwave Experimental Setups

As shown in the previous fluxonium chapter (Sec. 3.3), different noise levels are crucial factors that can ruin the coherence of our qubit [Kri+19]. Environmental noise, present due to the connection of the qubit to electronics at room temperature or via environmental electromagnetic fields, leads to qubit dephasing [Ith+05; Yeh+17; Yan+16; Koc+07] as well as the creation of quasi-particles, causing dissipation and reduced relaxation times [Cat+11; Bar+11; Cor+11; Pop+14; Kre+16; Yan+16]. Hence, it becomes necessary to protect the qubit with various well-thermalized attenuators, filters, and shields placed at different temperature stages of the cryostat.

[Kri+19]: Krinner et al. (2019), 'Engineering cryogenic setups for 100-qubit scale superconducting circuit systems', EPJ Quantum Technology

In our experimental setup, we used a dilution cryostat provided by Bluefors, which has stages with different temperature levels. The stages fully accessible to the user are the 4K, Still, and MXC (Mixing Chamber) stages, corresponding to temperatures of 4K, 800mK, and 10mK, respectively. This is illustrated in Fig. 6.15. The sample is thermalized to the MXC stage.

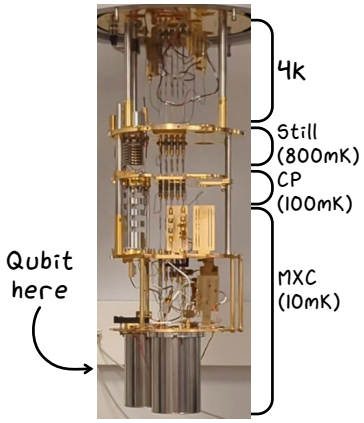

Fig. 6.15: Photo of the Bluefors cryostat with its stages and corresponding temperatures indicated. From top to bottom: 4 K, Still, Cold Plate (CP), and Mixing Chamber (MXC).

Fig. 6.16 shows the overall setup , comprising four descending RF lines, one ascending RF line, and two DC lines, with RF pulse generation and measurement handled by a *Quantum Machines QOP* FPGA. In this section, we describe the key practices underlying our cryogenic microwave experimental setup, motivating the setup design choices; the complete experimental setup will be discussed in full detail in the next section.

We begin by examining strategies for reducing environmental electrical noise and filtering the signal delivered to the qubit. We then outline the main considerations for component thermalization and the amplification chain, demonstrating that amplifier stages should be ordered by increasing noise contribution. Next, we explain the configuration of the mixers used to upconvert and downconvert signals to the qubit frequency, highlighting how imperfections in the mixing process can introduce spurious tones. Finally, we discuss microwave packaging and shielding techniques that ensure the qubit is properly thermalized and isolated from external disturbances.

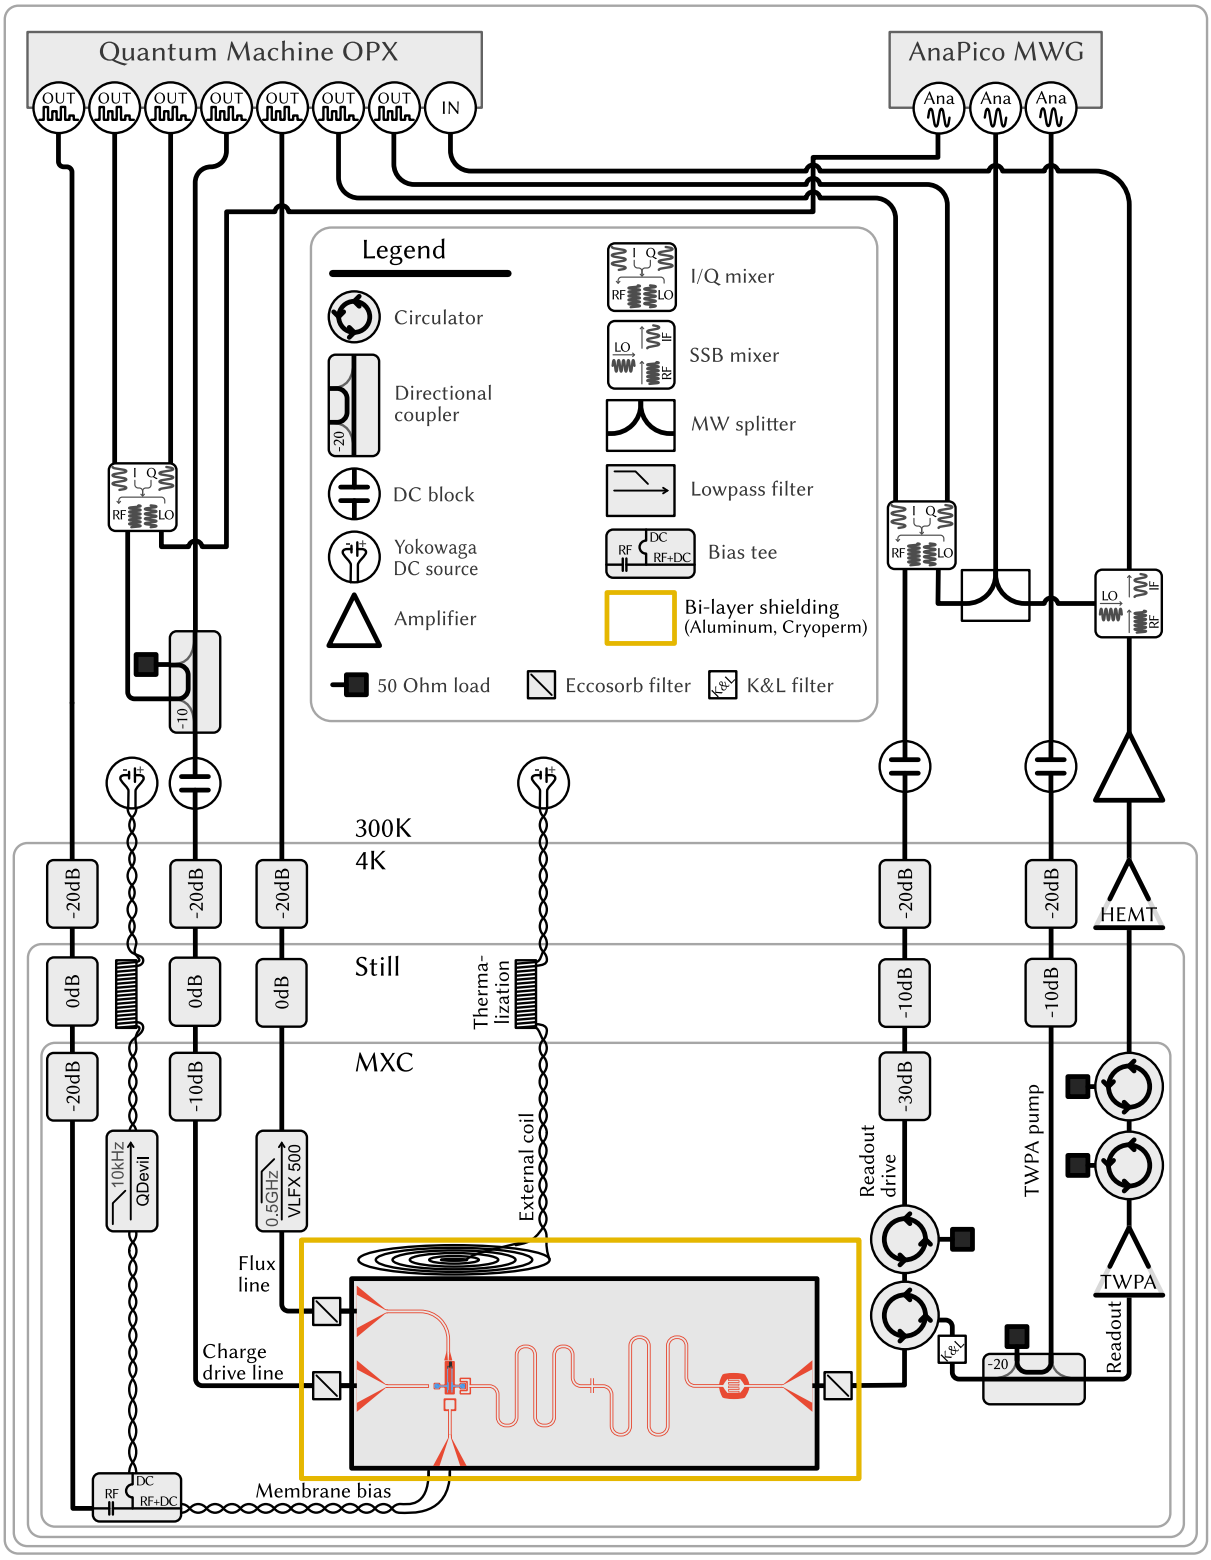

Fig. 6.16: Wiring schematic of the experimental setup. The diagram shows the cabling and connections between the cryostat, FPGA control unit, and signal generators. The qubit chip is placed at the bottom of the cryostat, where it is thermally linked to the MXC stage.

6.3.1 Noise attenuation

Electrical noise due to one-dimensional blackbody radiation in cables, known as Johnson-Nyquist noise, or $1/f$ noise, which can originate from electrical instruments themselves, decreases the coherence time of a qubit. For now, we will not delve into the details of the noise spectrum density, as we aim to reduce it in general. For any numerical examples given, a frequency of 1 GHz is chosen.

The typical signal power required to manipulate a qubit is quite low. This allows us to generate a higher-power signal and attenuate it as it descends to cryogenic temperatures. In the following discussion, we will see how such attenuation can reduce the electrical noise coming from room temperature to the qubit.

Consider a single attenuator placed on a descending line at stage s, thermalized at temperature T_s, as illustrated in Fig. 6.17. Denote the ladder operators of the incoming and outgoing signals as $\boldsymbol{a}_{s-1}$ and $\boldsymbol{a}_s$, respectively. Also denote the operators of the thermal noise due to the attenuator as $\boldsymbol{b}_{in,s}$. Then, the input-output equation can be written as follows:

$$\boldsymbol{a}_s = \sqrt{\Lambda_s} \cdot \boldsymbol{a}_{s-1} + \sqrt{1-\Lambda_s} \cdot \boldsymbol{b}_{in,s} \tag{6.5}$$

Here, we used the power transparency $\Lambda_i \in [0,1]$, which can be linked to the conventional logarithmic attenuation A_i with:

$$\Lambda_s = 10^{-|A_s|/10} \quad \rightarrow \quad |A_s| = -10 \cdot \log_{10} \Lambda_s \tag{6.6}$$

The input–output relation can be rewritten in terms of the photon occupation number of an itinerant mode at frequency ω:

$$n_s(\omega) = \Lambda_s \cdot n_{s-1}(\omega) + (1-\Lambda_s) \cdot n_{BE}(T_s, \omega) \tag{6.7}$$

Here, we used the fact that thermal noise can be found from the Bose–Einstein statistic, or $\langle \boldsymbol{b}^\dagger_{in,s} \boldsymbol{b}_{in,s} \rangle = n_{BE}(T_s) = \frac{1}{e^{\hbar\omega/k_B T_s}-1}$, where T_s is the temperature of the attenuator.

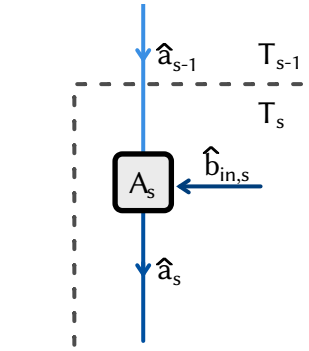

Fig. 6.17: Attenuation connection schematic. Illustration of an attenuator on a descending line used to reduce thermal noise. The input signal a_{n-1}, from the previous stage, enters the attenuator thermalized at temperature T_n, yielding an output signal denoted a_n. In this and subsequent illustrations, the temperature decreases from top to bottom, i.e., $T_{i-1} > T_i$.

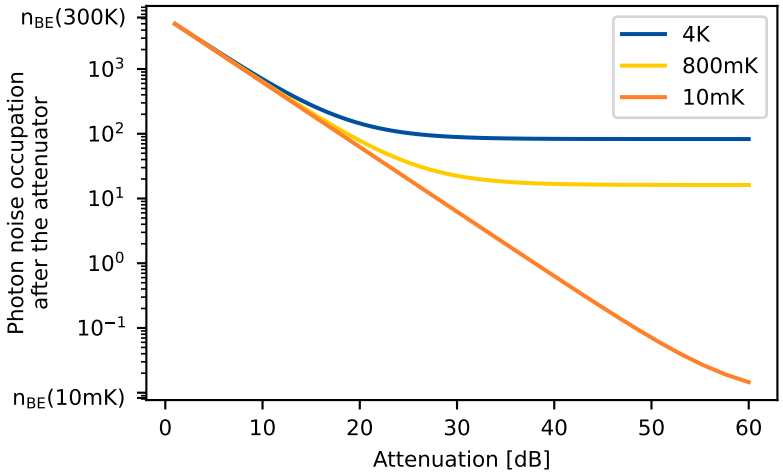

Fig. 6.18: Reduced photon occupation due to various attenuations. Photon occupation n_i as a function of attenuation A thermalized at different temperatures: 4 K in blue, 800 mK in yellow, and 10 mK in orange. The input signal for every level corresponds to the room temperature noise $n_{BE}(300\text{K})$.

Photon occupation equation (Eq. 6.7) is illustrated in Fig. 6.18, where the photon noise occupation n_i is shown as a function of the logarithmic attenuation $|A| = -10 \cdot \log_{10} \Lambda$, for the attenuator thermalized to different temperatures. The input signal corresponds to the thermal population at room temperature, $n_{BE}(300\text{K})$.

The figure demonstrates that for each stage, further increasing attenuation beyond a certain point does not lead to significant improvement. It also shows that photon noise occupation reduces significantly when attenuation is applied at lower temperatures. However, as we'll see in the following section, due to its limited cooling power at lower levels, it is not feasible to place all attenuation at the lowest temperature stage, and thus attenuation must be distributed accordingly across all stages.

6.3.2 Cooling Power

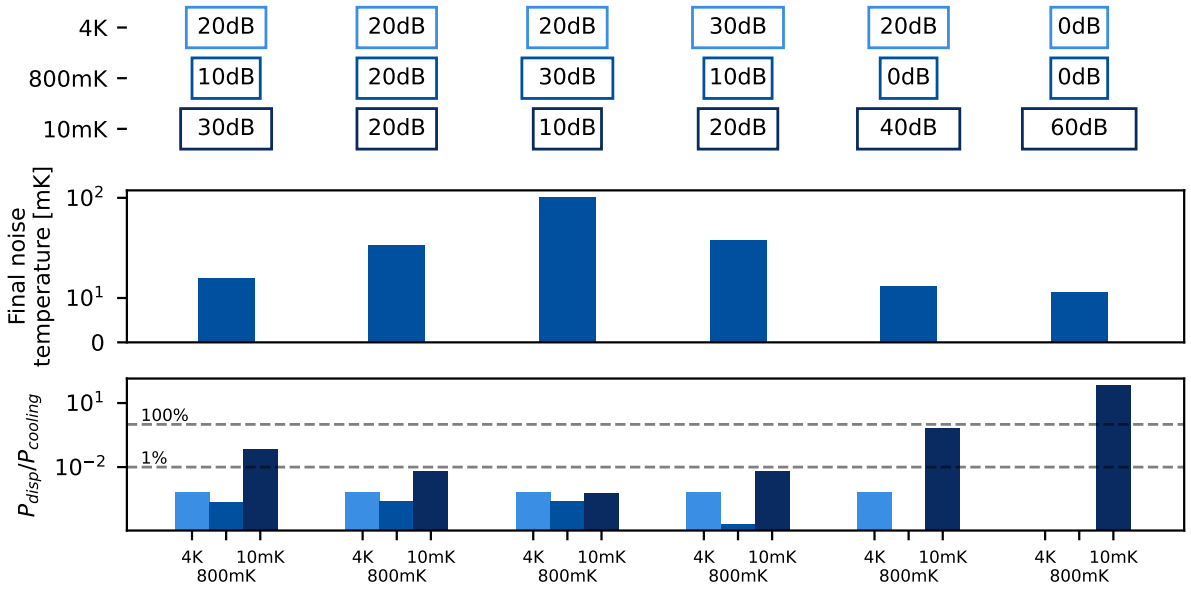

Fig. 6.19: Effective noise temperature and dissipated energy for different distributions of a total attenuation of 60 dB with an input power of 0 dBm at the top stage. . This plot shows (i) the final effective noise temperature—calculated from the photon occupation n_i—and (ii) the energy dissipated at each cryogenic stage for various ways of distributing the attenuation.

When a signal is sent into the cryostat, the power dissipated in the attenuator, $P_{\text{disp}} = (1 - \Lambda_i)P_{i-1}$, is not negligible, since it must be transferred out from the cryostat, which has its cooling limits. These limits, defined by the manufacturer, are listed in Tab. 6.1. As a result, this setup imposes restrictions on the maximum attenuation that can be applied at each stage.

Tab. 6.1: Cooling power limits at different stages of the Bluefors *BF-LD250* cryostat. The temperature in parentheses indicates the steady-state temperature of the corresponding stage when the specified power is absorbed.

Stage name	T	$P_{cooling}$
4K	2.9K	1.5 W (4.2K)
Still	830mK	40 mW (1.2mK)
MXC	9mK	15 μW (20mK)

The final attenuation for each line is defined by the following factors: the total power needed to interact with the qubit, the cooling power limits of the cryostat, and practical limitations on the output power of the source. Thus, different lines have different total attenuations, which will be explained in more detail in the following.

For now, to provide more intuition and demonstrate the importance of placing attenuation at lower-temperature stages, consider, for example,

a total maximum attenuation of −60 dB and examine how it should be distributed across the different stages. In Fig. 6.19, the final noise temperature is plotted alongside the energy dissipated at each stage for various attenuator distributions. The energy dissipated at each stage is normalized to the maximum available cooling power to illustrate how close we operate to this limit. The final effective noise temperature is given by

$$T_{\text{eff}} = \frac{\hbar\omega}{k_B \ln\left(1 + 1/n_s\right)},$$

where n_s is the photon noise occupation. The dissipated energy is calculated assuming an input power of 0 dBm at the top stage, which corresponds approximately to the maximum output power of our pulse controller (FPGA).

This figure shows that placing more attenuation at the warmer stages results in an increased final noise temperature. When attenuation is placed at the MXC stage, the dissipation energy is closer or exceeds the cooling power limit, which, with many lines, can influence the final steady-state temperature. While here we are not limited by the cooling power, adding more attenuation beyond a certain point does not lead to significant improvement in the final noise temperature. This explains the common cryostat configuration: 20 dB at 4K, 10 dB at Still, and the remaining at MXC.

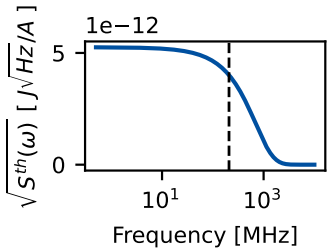

Fig. 6.20: Example of a one-sided Johnson-Nyquist noise spectrum at T = 10mK. Black dashed line indicates the thermal noise cutoff frequency, $k_B T/h$.

6.3.3 Filtering

Up to now, the frequency dependency of the noise wasn't considered. However, when the final power of the signal dictates the global attenuation, we can additionally attenuate frequencies outside the required pulse range.

Firstly, the one-sided voltage noise spectral density of the Johnson-Nyquist noise emitted by a resistor R can be written as follows:

$$S^{th}(\omega, T_i) = 4\,R\,\hbar\,|\omega|\,\frac{1}{e^{\hbar|\omega|/k_B T_i} - 1} \tag{6.8}$$

This decays rapidly with higher frequencies ($\hbar|\omega| > k_B T$), but it is not the only noise source in the system. For example, there are additional $1/f$ noise sources with different cut-off frequencies [Ith+05; Koc+07]. For further reasoning, there is no need to go into the details of their nature and reconstruct an exact spectral density profile. Instead, we can imagine a flat initial noise level and explore how to filter it out.

The required operating range of our qubit was globally attenuated. However, the bandwidth of the XMA cryogenic attenuators used (2082-6418-(Attenuation)-CRYO) ranges from DC to ~ 18 GHz. To attenuate any higher frequency noises [Ser+19; Cor+11], each line is connected through Eccosorb filters, which are placed as close as possible to the sample. These filters are homemade, composed of Eccosorb CR-110 epoxy (ranged 26+GHz) and encased in a copper box for proper thermalization and shielding. The frequency performance of this filter up to 18GHz is shown in Fig. 6.21. To reduce QP-induced dissipation, it is crucial to place

[Ith+05]: Ithier et al. (2005), 'Decoherence in a superconducting quantum bit circuit', Physical Review B
[Koc+07]: Koch et al. (2007), 'Charge-insensitive qubit design derived from the Cooper pair box', Physical Review A

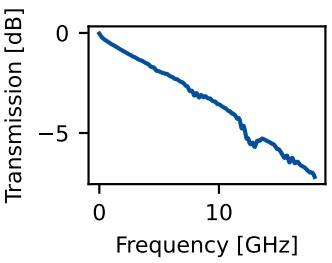

Fig. 6.21: Transmission (S_{21}) of the homemade Eccosorb absorber. Measured transmission as a function of frequency from 0 to 15 GHz at room temperature. The most relevant frequency range for Eccosorb filters lies above 15 GHz; however, our measurements were limited by the available experimental equipment.

[Ser+19]: Serniak et al. (2019), 'Direct Dispersive Monitoring of Charge Parity in Offset-Charge-Sensitive Transmons', Physical Review Applied
[Cor+11]: Corcoles et al. (2011), 'Protecting superconducting qubits from external sources of loss and heat', Applied Physics Letters

these filters within any cryoperm and aluminium shields, which will be discussed later. To further suppress thermal radiation, infra-red blocking filters can be used [Kre+16; Gee23].

Each port has its own operational range, so filters are used to suppress noise at frequencies outside the operating range.

[Kre+16]: Kreikebaum et al. (2016), 'Optimization of infrared and magnetic shielding of superconducting TiN and Al coplanar microwave resonators', Superconductor Science and Technology
[Gee23]: Geerlings (2023), 'Improving Coherence of Superconducting Qubits and Resonators', Thesis

6.3.4 Thermalization bridges

Up to this point, we have assumed that the microwave components placed at specific stages are at the same temperature as the stage itself. However, this is not always the case due to the limited thermal conductivity between the stage and the component. In practice, proper thermalization of each element must be done to improve thermal conductivity. This can be achieved by either screwing components directly to the cryostat platforms or using thermalization braids to connect components to the platforms (see an illustration in Fig. 6.22).

Here, we will not delve into precise calculations for thermalization but instead provide basic justifications for the decisions made.

The heat flux, denoted by q, through the thermalization braids can be calculated using the thermal conductivity formula:

$$q = -k\frac{\Delta T}{L} \tag{6.9}$$

where ΔT is the temperature difference between the two ends of the thermalization braid, L is its length, and k is its thermal conductivity. Therefore, the cooling power provided by a braid can be determined from:

$$P_{cooling} = q \cdot A = -kA\frac{\Delta T}{L} \tag{6.10}$$

where A represents the cross-sectional area of the connection. From this basic estimation, we can conclude that:

- The thermalization braids should be relatively short.
- The cross-sectional area should be large enough. This also implies that washers should be of comparable size and everything must be tightly screwed to maximize contact.
- Materials with a high k value, such as OFHC CU with high RRR (Residual Resistance Ratio), should be used.

A table of different materials with their thermal conductivity values is provided in Tab. 6.2.

In practice, we initially made a mistake by underestimating the material choice and the cross-sections of the braids. When we increased the thickness of the thermalization braids and switched from ETP copper to OFHC copper, we successfully reduced the temperature of our qubit holder from 30-40 mK to 15-20 mK. See Fig. 6.23 for an example of a thermalization braid used in our setup.

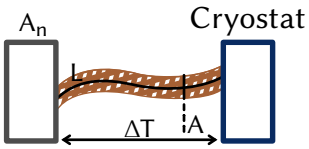

Fig. 6.22: Thermalization braid schematic. Schematic showing an attenuator (or any filter) thermally connected to the cryostat via thermalization braids of length L and cross-sectional area A. The temperature drop between the cryostat stage and the attenuator is represented by ΔT.

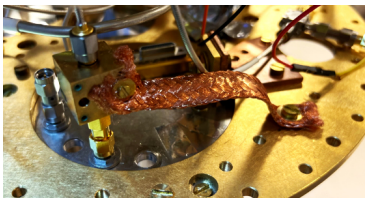

Fig. 6.23: Photo of a thermalization braid used. Here, it connects the Eccosorb filter to the mixing-chamber (MXC) stage of the cryostat.

Tab. 6.2: Thermal conductivity at 4K of typical materials for thermalization braids. Here, ETP CU stands for Electro Tough-Pitch Copper and OFHC CU for Oxygen-free high thermal conductivity copper.

Name	k [W/(m · K)]
ETP CU (50 RRR)	$2 \cdot 10^2$
OFHC CU (100 RRR)	$6.5 \cdot 10^2$
OFHC CU (300 RRR)	$2 \cdot 10^3$
OFHC CU (2000 RRR)	10^4

Another improvement was the use of bulkhead attenuators (*4880-5523-(Attenuation)-CRYO, XMA Corp*) that can be directly screwed onto the cryostat platforms. This ensures better thermalization without the need for additional thermalization braids.

6.3.5 Signal amplification

The only line in our setup that cannot be protected by attenuation and filtering is the output readout line, which brings the signal from the readout cavity out of the cryostat to room temperature. This signal is too weak compared to the noise level at room temperature and must be amplified. In the following, we will examine how the chain of amplifiers influences the signal-to-noise ratio (SNR) of the signal. Here, we only consider the signal traveling in the upward direction: from the lowest temperature levels to room temperature. All entering noise traveling in the opposite direction is suppressed by circulators as will be discussed later.

Consider a single amplifier with gain G_1 illustrated in Fig. 6.24. The total input P_0 can generally be written as the sum of the signal P_0^{S} and noise already in the system P_0^{N}:

$$P_0 = P_0^{\mathrm{S}} + P_0^{\mathrm{N}} \tag{6.11}$$

Thus, the input signal-to-noise is $\mathrm{SNR}_0 = P_0^{\mathrm{S}}/P_0^{\mathrm{N}}$. The amplifier, thermalized at temperature T_1, contributes its own noise to the output, $G_1\,P_1^{\mathrm{N}}$, where, by convention, the amplifier noise P_1^{N} is input-referred, i.e. as if injected directly at the input. Consequently, the total output power P_1 is

$$P_1 = G_1 \cdot (\underbrace{P_0^{\mathrm{S}}}_{\text{Information}} + \underbrace{P_0^{\mathrm{N}} + P_1^{\mathrm{N}}}_{\text{Noise}}) \tag{6.12}$$

The SNR of the output is then:

$$\mathrm{SNR}_1 \equiv \frac{P_0^{\mathrm{S}}}{P_0^{\mathrm{N}} + P_1^{\mathrm{N}}} \tag{6.13}$$

This reveals that the SNR after the amplifier decreases due to the noise added by the amplifier itself. This result might seem counterintuitive at first glance and motivates examining what happens when the signal passes through an additional amplifier, as shown in Fig. 6.24. Here, P_1 is the input to the second amplifier, which has gain G_2 and input-referred noise P_2^{N}. The output of the second amplifier is then:

$$\begin{aligned} P_2 &= G_2 \cdot (P_1 + P_2^{\mathrm{N}}) = \\ &= G_2 G_1 \cdot (\underbrace{P_0^{\mathrm{S}}}_{\text{Information}} + \underbrace{P_0^{\mathrm{N}} + P_1^{\mathrm{N}} + P_2^{\mathrm{N}}/G_1}_{\text{Total noise}}) \end{aligned} \tag{6.14}$$

Therefore, the SNR at the output of the second amplifier is:

$$\mathrm{SNR}_2 = \frac{P_0^{\mathrm{S}}}{P_0^{\mathrm{N}} + P_1^{\mathrm{N}} + \frac{P_2^{\mathrm{N}}}{G_1}} \tag{6.15}$$

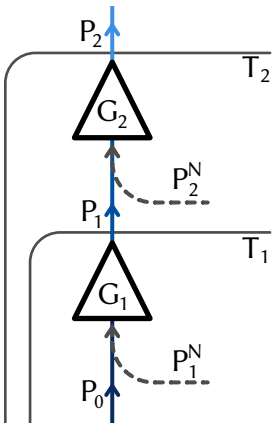

Fig. 6.24: Amplification chain schematic. Each amplifier is thermalized at a different cryostat stage (dashed lines). Each amplifier, with gain G_i, amplifies its input signal P_{i-1} and adds its own noise P_i^{N}. The output of the stage is then $P_i = G_i(P_{i-1} + P_i^{\mathrm{N}})$.

This highlights the importance of the amplification chain, as the noise contribution from the second amplifier is reduced by a factor of G_1. In other words, while the SNR after an amplifier does decrease due to added noise, the amplifier's primary purpose is to boost the signal enough to surpass subsequent noise contributions in the chain. This principle leads to a practical rule for designing an amplification chain: place high-performance, low-noise amplifiers with high gain at the start of the chain, after which more affordable amplifiers can be used.

In our setup, we use a traveling-wave parametric amplifier (TWPA) as the first amplifier in the chain. The TWPA is one of the best amplifiers currently available, providing 20 dB of gain with only a few photons of added noise. This ensures that the noise contributed by the second amplifier—a high-electron-mobility transistor (HEMT) operated at 4 K with approximately 10 quanta of added noise—is only a marginal fraction of the total noise.

6.3.6 Mixer Configuration

Our FPGA controller handles signals ω_{IF} in the intermediate-frequency range ($-350\,\mathrm{MHz} < \omega_{\mathrm{IF}} < 350\,\mathrm{MHz}$). To operate in the 3–6 GHz band, we use an external local oscillator at frequency ω_{LO} and IQ mixers to upconvert IF signals into the GHz range $\omega_{\mathrm{RF}} = \omega_{LO} + \omega_{\mathrm{IF}}$, and to downconvert RF outputs back to IF. However, real mixers introduce unwanted spurious tones if not properly balanced.

An IQ mixer multiplies two baseband inputs $I(t)$ and $Q(t)$ by in-phase and quadrature LO signals, producing an output

$$V_{\mathrm{out}}(t) = I(t)\cos(\omega_{\mathrm{LO}}t) - Q(t)\sin(\omega_{\mathrm{LO}}t)$$

For ideal single-sideband up-conversion, one chooses

$$I(t) = A\cos(\omega_{\mathrm{IF}}t), \quad Q(t) = A\sin(\omega_{\mathrm{IF}}t),$$

which yields

$$V_{\mathrm{out}}(t) = A\cos\left[(\omega_{\mathrm{LO}} + \omega_{\mathrm{IF}})t\right]$$

and suppresses both the LO (ω_{LO}) and the lower sideband ($\omega_{\mathrm{LO}} - \omega_{\mathrm{IF}}$).

Imperfections and Spur Generation. In practice, the DACs and mixer introduce [SL08, p. 144][Sab98, p. 2][Mic15]:

- DC offsets I_0, Q_0 in the baseband paths (i.e., the low-frequency I and Q channels before upconversion),
- Amplitude imbalance δ between I and Q,
- Phase error ε from deviation from the exact 90° phase shift.

To quantify the spurious tones generated by these imperfections, we substitute the imperfect low-frequency I and Q signals (the baseband signals) into the mixer's output expression and collect terms at ω_{LO} (carrier leakage) and $\omega_{\mathrm{LO}} - \omega_{\mathrm{IF}}$ (image tone). One models the imperfections as

$$I(t) = A\cos(\omega_{\mathrm{IF}}t) + I_0, \quad Q(t) = A(1+\delta)\sin(\omega_{\mathrm{IF}}t + \varepsilon) + Q_0.$$

[SL08]: Schenk et al. (2008), *RF imperfections in high-rate wireless systems: impact and digital compensation*, Springer

[Sab98]: Sabah (1998), 'Design and Calibration of IQ-Mixers', CERN

[Mic15]: Microwave (2015), *5 Ways to Compensate for Passive IQ Mixer Imbalance*, Documentation

This yields carrier leakage at ω_{LO}:

$$V_{\mathrm{LO}}(t) = I_0 \cos(\omega_{\mathrm{LO}} t) - Q_0 \sin(\omega_{\mathrm{LO}} t),$$

which depends only on the DC offsets (I_0, Q_0). The image tone at $\omega_{\mathrm{LO}} - \omega_{\mathrm{IF}}$ appears approximately as

$$V_{\mathrm{img}}(t) \approx -\frac{A}{2}\left[\delta \cos\big((\omega_{\mathrm{LO}} - \omega_{\mathrm{IF}})t\big) + \varepsilon \sin\big((\omega_{\mathrm{LO}} - \omega_{\mathrm{IF}})t\big)\right],$$

whose amplitude scales as $A\sqrt{\delta^2 + \varepsilon^2}/2$.

Cable Matching To minimize additional phase and amplitude skew, the I/Q coaxial cables should have identical length and loss. At ~ 350 MHz, the wavelength in typical coaxial cable is $\lambda \approx 0.6$ m.

Automated Two-Step Calibration Uncalibrated IQ mixers typically achieve only ~ 20 dB suppression of the carrier and image. The Quantum Machine Interface (QMI) applies per-channel DC offsets and a 2×2 complex correction matrix to compensate for amplitude imbalance δ and phase error ε. We automate the calibration by driving a continuous-wave IQ tone and using a spectrum analyzer to monitor spurs:

1. **Offset nulling:** Adjust (I_0, Q_0) via gradient descent to minimize the LO peak at ω_{LO}.
2. **Image suppression:** With DC offsets fixed, adjust the amplitude and phase entries of the correction matrix to minimize the lower sideband at $\omega_{\mathrm{LO}} - \omega_{\mathrm{IF}}$.

After convergence, both carrier leakage and image are suppressed by > 40 dB, yielding a clean single-sideband mixer response.

6.3.7 Microwave Packaging and Shielding

The qubit chip is glued and wire-bonded onto a printed circuit board (PCB), which is then integrated into a copper sample holder known as JAWS (Joint Assembly for the Wiring of Superconducting circuits). This holder was designed by Marius Villier from the LPENS group [Vil23]. It protects the qubit from spurious modes. Once the JAWS is closed, the formed cavity has box modes with frequencies above 10 GHz. This is higher than both the qubit and readout modes, providing protection.

Shielding the qubit from stray electric and magnetic fields is essential to preserve coherence [Kre+16]. We employ a nested two-layer enclosure (Fig. 6.25): it consists of an outer *Cryoperm* shield that passively attenuates ambient magnetic fields, and an inner superconducting Al shield that repels residual magnetic flux via the Meissner effect [NP14; Son+09]. The *Al* shield resides inside the *Cryoperm* to ensure that it transitions to the superconducting state in a nearly zero magnetic field. These shields consist of two parts: a top cover that is rigidly attached to the cryostat mixing chamber, and a can (an open-top cylinder) that can be easily removed to access the sample. The sample itself is placed on a copper frame that is thermally anchored to the lowest stage of the cryostat.

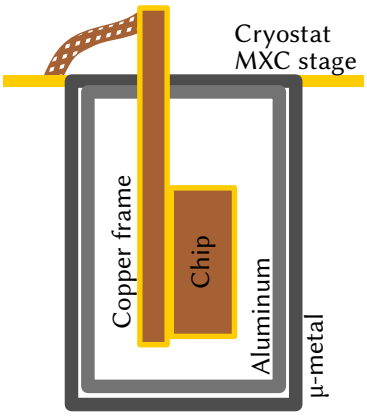

Fig. 6.25: Chip shielding assembly. Schematic of the shielding screens surrounding the chip holder (labelled "chip"). The holder is mounted on a copper frame, which is gold-plated to prevent oxidation and enhance thermal contact, and is thermalized to the mixing-chamber (MXC) stage via thermalization braids. Nested μ-metal and aluminium screens enclose the chip to block external magnetic and electromagnetic interference.

[Vil23]: Villiers (2023), 'Dynamically Enhancing Qubit-Photon Interactions with Anti-Squeezing', Thesis

[Kre+16]: Kreikebaum et al. (2016), 'Optimization of infrared and magnetic shielding of superconducting TiN and Al coplanar microwave resonators', Superconductor Science and Technology

[NP14]: Nsanzineza et al. (2014), 'Trapping a single vortex and reducing quasiparticles in a superconducting resonator', Physical Review Letters

[Son+09]: Song et al. (2009), 'Microwave response of vortices in superconducting thin films of Re and Al', Physical Review B

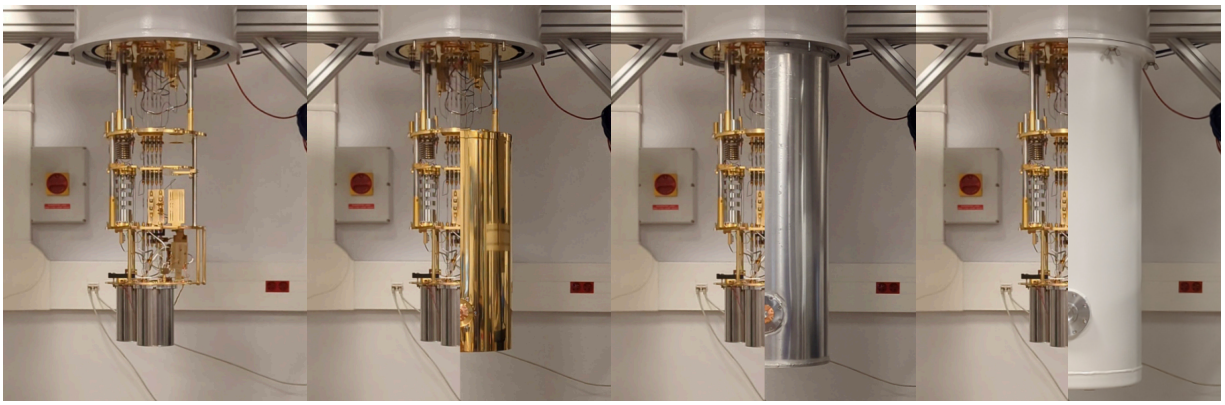

Fig. 6.26: Illustration of the cryostat sealing process. Each Each step adds a shielding screen. From left to right, the cans are thermalized at 800 mK, 4 K, and 50 K, before sealing the vacuum enclosure.

The mixing-chamber stage is shielded from blackbody radiation from the upper stages by successive metallic enclosures (see Fig. 6.26). The first shield is a copper can attached to the 800 mK still plate. Since the radiative power scales as T^4, the radiation emitted by this shield toward the CP and MXC plates is negligible in terms of heat load. Next are two aluminium shields thermalized at the 4 K and 50 K plates. As discussed in Sec. 6.3.1, distributing the heat load across successive stages prevents exceeding the cryogenic capacity of the lowest stages. Finally, the cryostat vacuum enclosure serves as a Faraday cage, isolating the interior from external electromagnetic noise.

6.4 Assembled Microwave Setup

Our qubit chip has four ports: RF charge drive, RF readout, RF+DC flux, DC membrane bias. A complete illustration of our microwave setup is shown in Fig. 6.16.

To eliminate any DC offset, we use DC-blocks on every RF line at room temperature. To generate the RF pulses required for controlling or reading the qubit state, we utilize a Quantum Orchestration Platform (QOP) called OpX from Quantum Machine. This platform allows us to create any type of signal with its resolution and conduct analog-to-digital conversion to read the reflected signal within a 350 MHz bandwidth.

DC lines

DC lines are used to connect the coil around the qubit and to bias the membrane. The coil, made of 300 turns of NbTi superconducting wire (see Fig. 6.27), is placed next to the sample to create the magnetic flux and control φ_{ext}. We used superconducting cables to connect it. This choice is critical, as any cables with considerable resistance would heat up and emit radiation.

For each connection, we use two DC lines for voltage and ground that descend in twisted pairs. They are thermalized by the manufacturer at the MXC stage. We use a Yokogawa[1] voltage source to generate a constant

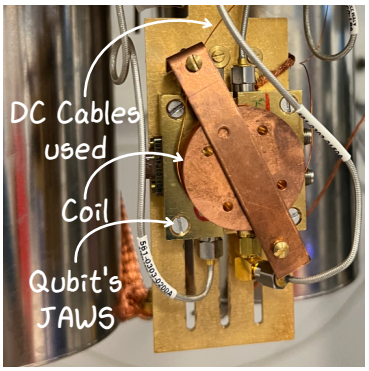

Fig. 6.27: Photo of the external coil used to set the external magnetic field and control φ_{ext}. The coil is made of 300 turns of NbTi superconducting wire, with a final approximate inductance of $L \approx 4$ mH.

1: We used the *Yokogawa 7651* model to adjust the flux offset and the *Yokogawa GS200* model for the membrane DC bias.

DC voltage or current. Minimizing noise on the DC lines is essential: flux noise couples directly to qubit fluctuations, while noise on the membrane-bias line drives unwanted membrane motion and accelerates qubit relaxation. The DC flux line connected to the coil is filtered by a room-temperature low-pass filter with a cutoff frequency of 1.4 Hz and by the coil itself, $\omega_{\text{cutoff}}/2\pi = 50\,\Omega/2\pi L \approx 2\,\text{kHz}$. The membrane bias line is filtered with a QDevil RF low-pass filter with a cutoff frequency of 10 kHz.

Flux Line

This line connects the Quantum Machines (QM) system to the qubit flux port, enabling fast control of the external flux φ_{ext} without being limited by the long response time of the external coil, $\tau = L/(50\ \Omega) \approx 80\ \mu\text{s}$. However, as a backup to this external coil, we also required the on-chip flux line to provide a tuning range spanning several periods of the device response (up to $4V_{2\pi}$). To avoid excessive heating when delivering such drive powers, we implemented no attenuation at the MXC and still stages, while adding 20 dB of attenuation at the 4 K stage. In retrospect, retaining full-period tuning capability via this line is unnecessary, and the isolation of this line can be improved in future iterations.

To additionallyreduce noise, we added a low-pass filter with a 500 MHz cutoff (VLFX-500). This sets a minimum ramp timescale of $1/f \approx 2$ ns, which is not an experimental limitation. In practice, the ramp duration is typically of order 100 ns to avoid non-adiabatic effects.

The characterization of the flux line transmission at room temperature is shown in Fig. 6.29.

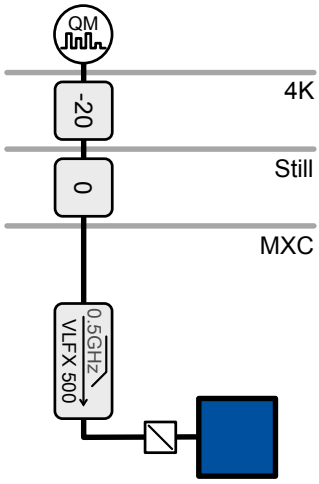

Fig. 6.28: Flux line connection schematic. This line enables adjustment of the qubit's external flux, φ_{ext}.

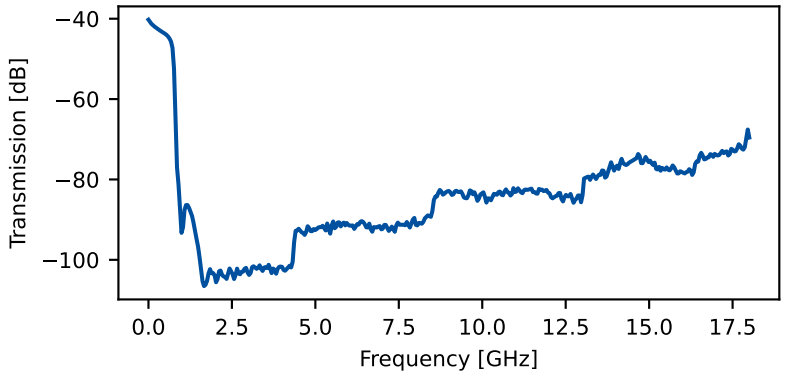

Fig. 6.29: Characterization of the flux line transmission at room temperature, measured with a spectrum analyzer.

Charge Line

This line is used to manipulate the qubit. It is connected to the qubit charge port and is directly coupled to the capacitor of the fluxonium on the chip. Therefore, shielding this line from thermal noise is crucial.

From Eq. 3.13, we can estimate the power needed to flip the qubit at a Rabi frequency of 2 MHz to $\sim$ (−60 dBm), which, given the maximal power of QM of $\sim$ (−20 dBm), results in a total attenuation on this line of 40 dB. To retain headroom, we distribute this attenuation across different stages as follows: 20 dB at 4K, 0 dB at Still, and 10 dB at MXC. No additional filters are added to this line, except for the Eccosorb. This allows us to drive:

- the $|g\rangle \rightarrow |e\rangle$ transition, which lies in the MHz frequency range. To do so, we use the QOP as the sole signal source; the generated tone passes through the through port of the directional coupler before entering the cryostat.
- the upper transitions, which lie in the GHz frequency range. To do so, we use two output channels of the QOP, providing the in-phase (I) and quadrature (Q) signals at frequency ω_{IF}. These are upconverted via IQ mixing with the local oscillator at ω_{LO}, generated by the Anapico RF source, producing a final tone at $\omega = \omega_{\mathrm{LO}} + \omega_{\mathrm{IF}}$ whose phase is set by the relative phase between the I and Q components. This upconverted signal is then routed to the coupled port of the directional coupler.

The characterization of the charge line transmission at room temperature is shown in Fig. 6.31.

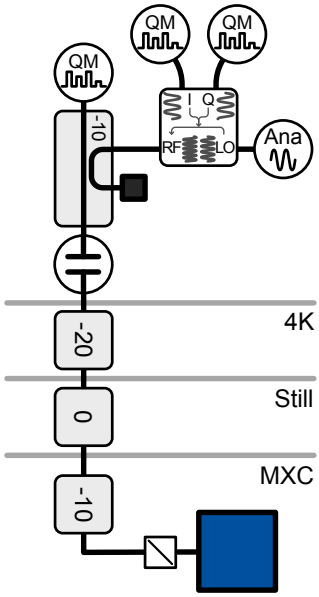

Fig. 6.30: Charge line connection schematic. This line drives the qubit through a capacitively coupled port.

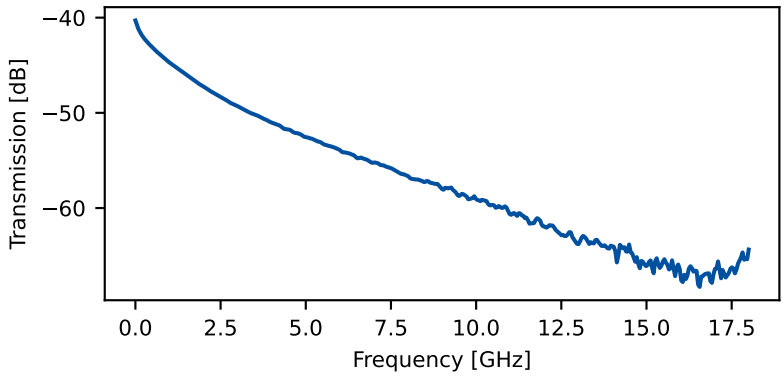

Fig. 6.31: Characterization of the charge line transmission at room temperature, measured with a spectrum analyzer.

Readout Input Line

As previously demonstrated, we need to attenuate the input signal. Therefore, any reflection measurements should include at least two lines (input and output) that are combined near the sample at lower temperatures. This line sends a probe signal to the sample's readout resonator. It is attenuated to the maximum extent possible, which, with our coupling, totals 60 dB of attenuation: 20 dB at 4K, 10 dB at Still, and 30 dB at MXC.

To ensure that the input and output lines are well separated and none of the reflected signal is sent back through the input line, we use a double circulator. The input signal passes through the two circulators, while the reflected signal from the sample, with the aid of the bottom circulator, goes to the readout output line. If there is any small leakage or reflection inside the readout output line, the top circulator directs any remaining reflected signal into a 50 Ω termination.

At room temperature, we use an IQ mixer to generate a signal in the GHz frequency range that matches the resonance frequency of the readout cavity. The QOP generates the IQ quadratures of the signal, which serve as input to a mixer, while the Anapico RF source acts as the local oscillator.

The characterization of the redout drive line transmission at room temperature is shown in Fig. 6.33.

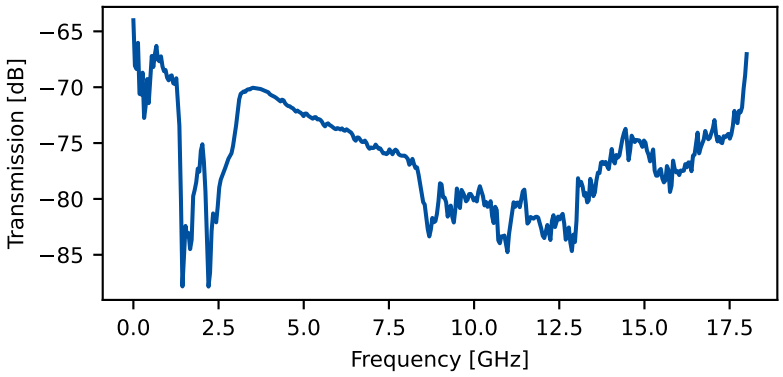

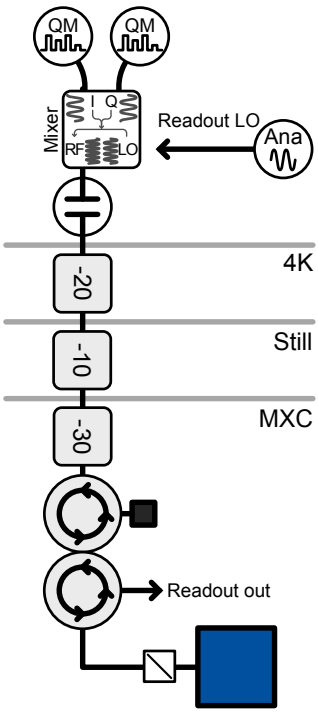

Fig. 6.32: Readout drive line connection schematic. This line is used to excite the readout cavity.

Fig. 6.33: Characterization of the readout drive line transmission at room temperature, measured with a spectrum analyzer.

TWPA Pump Line

This line powers the traveling-wave parametric amplifier (TWPA), which amplifies the signal reflected from the qubit's readout port. It has a total attenuation of 50 dB: 20 dB at 4K, 10 dB at Still, and 20 dB at MXC. The line is connected to the readout output line through a directional coupler placed just before the TWPA. The readout signal passes through the through port of the directional coupler, while the TWPA pump enters via the coupled port.

At room temperature, this line is connected to an Anapico RF source that generates the signal with the correct frequency and amplitude.

Readout Output Line

This line is the only output line in the setup. The qubit's reflected readout signal is directed toward the readout output line using a circulator. To amplify this signal, we use a traveling-wave parametric amplifier (TWPA) at the MXC stage. The pump for this amplifier is introduced through a directional coupler placed just before the TWPA.

There is no attenuation on this line, so to protect it from external thermal noise, we use double circulators with a 50 Ω termination. To minimize losses on this line, electrical connections are made using NbTi material instead of the usual CuNi. After the circulators, the signal passes through a HEMT (high-electron mobility transistor) at 4K and then through a room-temperature microwave amplifier.

At room temperature, we perform heterodyne demodulation of the signal using an image-reject (IR) mixer placed after the DC-block and room-temperature amplifier. We use as the local oscillator the same tone that was used to upconvert the readout input signal, obtained by splitting the output of the Anapico RF source, which ensures full phase correlation. This yields a demodulated low-frequency signal that is fed into the digitizing input of the QOP, allowing us to access the in-phase and quadrature components as a function of time.

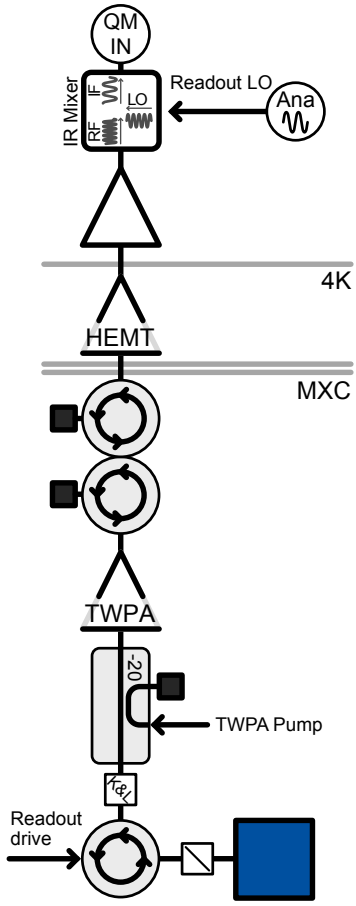

Fig. 6.34: Readout ascending line connection schematic. This line brings the signal from the readout cavity out of the cryostat for measurement.

7 Characterization of the qubit

Having completed qubit fabrication and assembled the experimental setup, we now proceed to the main experimental results of this thesis. This chapter presents the basic characterization of the qubit itself; its interaction with the membrane will be presented in the following chapters.

The results presented in this chapter are measurements of two chips from the same wafer. We first thoroughly characterized **CHIP**$-\boldsymbol{\alpha}$, the first chip of this design. A flip-chip assembly attempt with **CHIP**$-\boldsymbol{\alpha}$ failed, physically breaking the device. Next, we performed preliminary measurements on **CHIP**$-\boldsymbol{\beta}$, the second chip, but due to time constraints these were limited. A successful flip-chip assembly on **CHIP**$-\boldsymbol{\beta}$ produced **CHIP**$-\boldsymbol{\beta}-$**FC**, which served as the basis for the main results of this thesis and the corresponding paper [Ger+25]. To illustrate all the important techniques used, we present results from both **CHIP**$-\boldsymbol{\alpha}$ and **CHIP**$-\boldsymbol{\beta}$.

My contribution. Throughout my PhD, qubit characterization experiments—when feasible—were my primary occupation. I tested numerous qubit prototypes and developed Python libraries to optimize the experimental workflow, aiming to provide clean, reusable code. Obtaining the experimental results presented in this and the following chapters has been my primary occupation since wafer fabrication.

Outline. This chapter begins with readout spectroscopy of the qubit, characterizing the readout resonator. Next, we perform two-tone spectroscopy in the GHz range to extract the main fluxonium parameters E_L, E_C, and E_J, and describe how hybridization with a collective mode of the junction-chain superinductance dresses the qubit spectrum. Using these parameters, we estimate the gap between the qubit and the membrane. We then calibrate basic qubit control—determining transition frequencies and π-pulse amplitudes—and discuss reset procedures required by the high thermal occupation. We follow with an analysis of qubit readout optimization, and conclude by presenting decoherence times and the ac-Stark shift, which tunes the qubit frequency for resonant coupling to the membrane.

7.1 Readout Resonator Spectroscopy

To read out the qubit, we exploit the dispersive shift of the readout resonator frequency to infer the qubit state (Sec. 2.5.3). Therefore, before performing any qubit measurements, we first determine the resonator frequency.

7.1.1 Readout Resonator Identification

We measure the reflection coefficient S_{11} of the cavity by driving the readout resonator and recording its output. As shown previously in Sec. 2.5.3, the complex reflection coefficient can be described by a Lorentzian resonance (Eq. 2.74). The Purcell filter modifies this response, but we assume that its transmission is approximately constant in a narrow band around the readout frequency. In a realistic measurement setup, impedance mismatches and finite cable lengths introduce a fixed phase shift ζ and a drive frequency-dependent phase $\omega_d \tau_d$ due to the propagation delay τ. We capture these effects by allowing

$$\kappa_c \longrightarrow \kappa_c + i\,\kappa_c', \quad S_{11} \longrightarrow a\, e^{i\zeta}\, e^{i\omega_d \tau_d}\, S_{11}.$$

The full reflection coefficient as a function of the drive frequency ω_d then becomes

$$S_{11}(\omega_d) = a\, e^{i\zeta}\, e^{i\omega_d \tau_d} \frac{(\kappa_i - \kappa_c)/2 - i(\Delta + \kappa_c')}{(\kappa_i + \kappa_c)/2 - i\Delta}\,, \tag{7.1}$$

with $\Delta = \omega_d - \omega_R$ the detuning from the readout frequency and a is a complex amplitude.

Practically, fitting this equation requires robust initial guesses, especially for the delay τ. We found the `abcd_fit_rf` library by Ulysse Reglade [Reg24] to be an effective plug-and-play solution for complex-Lorentzian fitting.

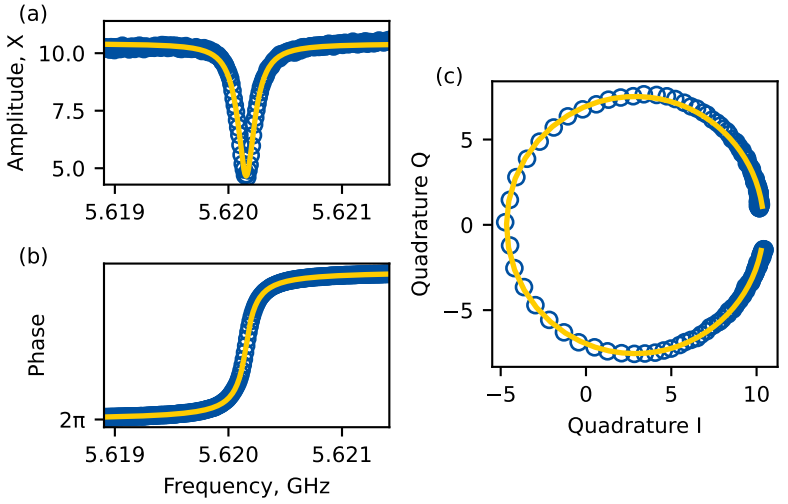

Fig. 7.1: Measured readout cavity response. (a) Amplitude vs. frequency, **(b)** phase vs. frequency, and **(c)** complex-plane trace. Data are shown in blue; the yellow curve is a fit to Eq. 7.1. All amplifiers, including the TWPA and room-temperature amplifier, are turned on. (**CHIP−*α***)

An example response is shown in Fig. 7.1. The drive frequency sweep exhibits a pronounced dip in amplitude and a smooth 2π phase roll-off centered at the readout frequency [1] $\omega_R/2\pi$. This measurement establishes the nominal readout frequency used in subsequent pulse sequences.

1: Formally, this is the dressed resonator frequency $\tilde{\omega}_R$, shifted by the readout-qubit coupling. We drop the tilde for clarity.

Regarding the amplitude of the readout pulse, we discussed that the dispersive shift depends on the photon number in the readout cavity (Sec. 2.5.2) and thus on the readout amplitude. However, with sufficiently large power, non-linear effects occur, and we exit the QND (Quantum Non-Demolition) regime. Our interest is to set the amplitude at the limit (with a reasonable gap) of the linear regime. In practice, we stay within the linear regime and optimize the exact amplitude later in readout optimization (Sec. 7.5).

7.1.2 Flux Dependence of the Readout Frequency: Single-Tone Spectroscopy

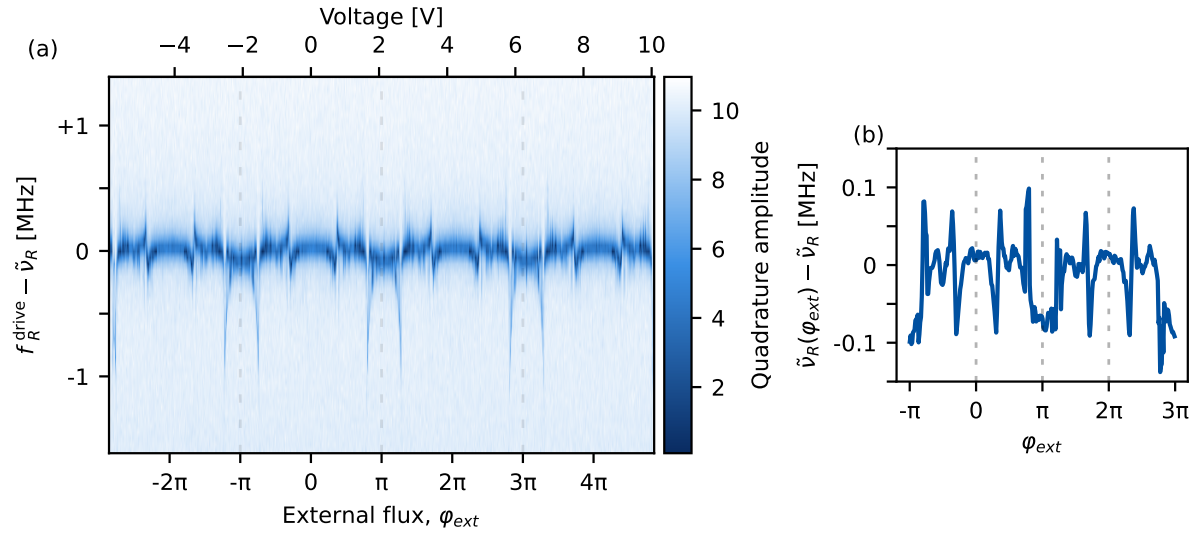

Fig. 7.2: Single-tone spectroscopy of the readout resonator as a function of external flux φ_{ext}. **(a)** Color map of the quadrature amplitude $|S_{11}|$ of the reflected signal. The horizontal axes show the voltage applied by our source (top) and the corresponding external flux φ_{ext} (bottom), calibrated using two symmetric flux points. The vertical axis shows a sweep of the drive frequency detuning $(\omega_d - \omega_R)/2\pi$. **(b)** Flux-dependent shift of the resonator frequency $\tilde{\nu}_R(\varphi_{\text{ext}}) - \tilde{\nu}_R$, obtained from Lorentzian fits of S_{11}, highlighting its periodic modulation with φ_{ext}.

The fluxonium Hamiltonian (Eq. 3.1) is symmetric about $\varphi_{\text{ext}} = 0$ and $\varphi_{\text{ext}} = \pi$ and 2π-periodic in φ_{ext}. We therefore perform a two-dimensional sweep of the drive frequency ω_d and external flux φ_{ext}. For each flux bias, we record the reflection coefficient $S_{11}(\omega_d)$ and fit it to the complex Lorentzian (Eq. 7.1). This yields the resonator frequency $\omega_R(\varphi_{\text{ext}})$ and linewidth $\kappa(\varphi_{\text{ext}})$. The resulting two-dimensional map is shown in Fig. 7.2(a), and the extracted $\omega_R(\varphi_{\text{ext}})$ appears in Fig. 7.2(b).

As φ_{ext} varies, avoided crossings appear whenever a qubit transition crosses the resonator frequency ω_R. Although detailed modeling of these anticrossings is beyond our scope (see [Zhu+13, Fig. 7]), we observe them as sharp features that unambiguously locate the flux periodicity and the symmetry points $\varphi_{\text{ext}} = 0$ and $\varphi_{\text{ext}} = \pi$, marked by dashed vertical lines in Fig. 7.2(a). One practical way to distinguish $\varphi_{\text{ext}} = 0$ from $\varphi_{\text{ext}} = \pi$ is via two-tone spectroscopy described in the next subsection. Mapping the voltage source output to these symmetry points calibrates the external flux phase. Moreover, a featureless (flat) $\omega_R(\varphi_{\text{ext}})$ would immediately signal a non-responsive or "dead" qubit.

[Zhu+13]: Zhu et al. (2013), 'Circuit QED with fluxonium qubits: Theory of the dispersive regime', Physical Review B

Although a full two-dimensional sweep of drive frequency and flux provides complete calibration, it is impractical for pure φ_{ext} calibration due to the large number of measurements required. To accelerate symmetry-point determination, we fix the readout tone at the unperturbed resonator frequency ω_R (see Sec. 7.1.1) and sweep only the external flux φ_{ext}, while recording the IQ quadratures of the reflected signal S_{11}. The qubit–resonator anticrossings then appear as sharp anomalies in amplitude and phase (see Fig. 7.3). This single-frequency line scan reduces the measurement time by an order of magnitude and still locates the symmetry points $\varphi_{\text{ext}} = 0, \pi$ with high precision.

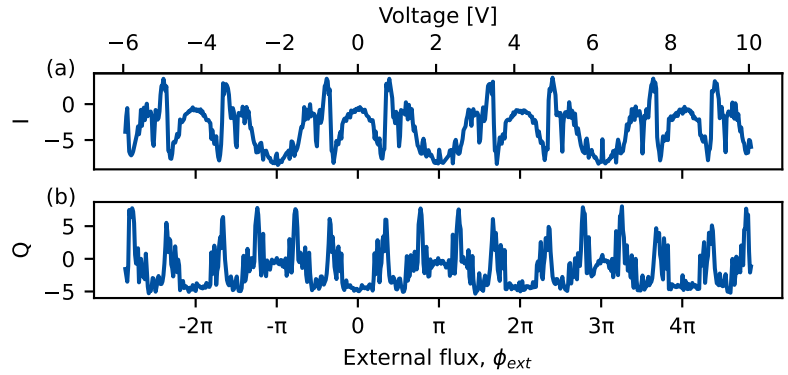

Fig. 7.3: Flux-dependent resonator response at fixed drive frequency ω_R. (a) Real and (**b**) imaginary quadrature components of the transmitted signal plotted against external flux Φ_{ext}.

7.1.3 Flux stabilization

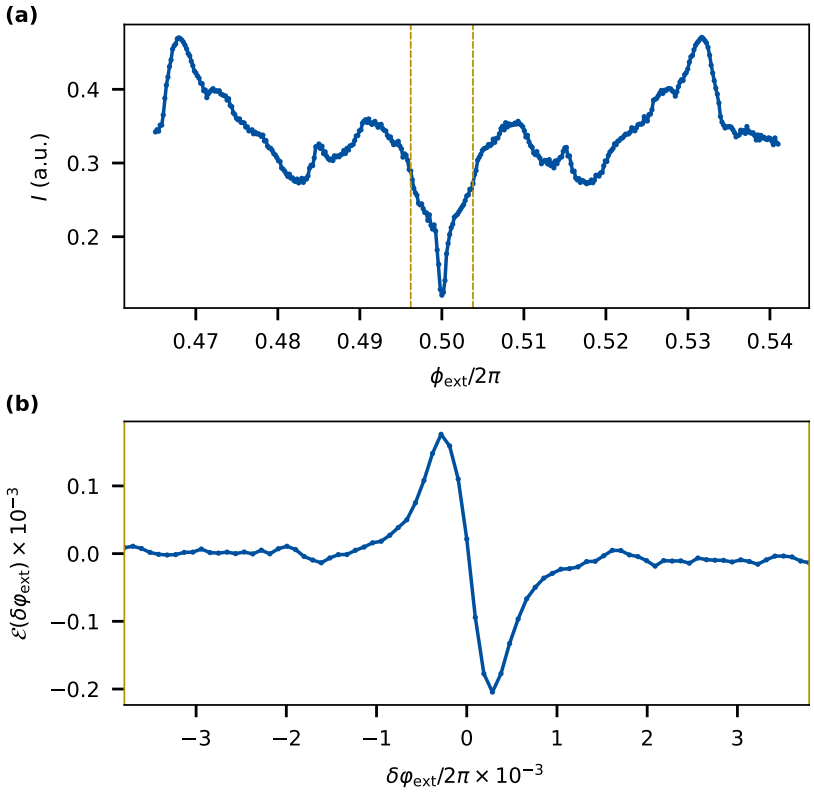

Fig. 7.4: Flux stabilization. (a) I-quadrature of the cavity response (real-part of the reflection coefficient) as a function of the external flux near $\varphi_{\text{ext}} = \pi$. Yellow lines indicate the range of $\delta\varphi_{\text{ext}}$ shown in panel (b). **(b)** Dispersive error signal $\mathcal{E}$ generated by dithering the flux around the current value plotted as a function of flux offset $\delta\varphi_{\text{ext}} = \varphi_{\text{ext}} - \pi$.

Fluxonium qubits exhibit enhanced coherence at $\varphi_{\text{ext}} = \pi$, where first-order sensitivity to external flux noise vanishes. In our setup, the external flux φ_{ext} is set by both an external coil and an on-chip flux-bias line, but environmental drifts induce slow variations in the required bias current. To maintain $\varphi_{\text{ext}} = \pi$, we implement a closed-loop feedback system based on the flux-dependent readout cavity response, which shows a pronounced dip at $\varphi_{\text{ext}} = \pi$ (see Fig. 7.4(a)). At each repetition of the experimental sequence (typically every $\sim$ 1 ms), we apply symmetric flux perturbations of $\pm 2\pi \times 1.9 \times 10^{-4}$ and record the corresponding cavity reflections. The difference between these two signals defines the dispersive error signal $\mathcal{E}$. In the range $|\varphi_{\text{ext}} - \pi| < 2\pi \times 5.7 \times 10^{-4}$, the dispersive error signal $\mathcal{E}$ is linear in the flux deviation $\delta\varphi_{\text{ext}} = \varphi_{\text{ext}} - \pi$ (see Fig. 7.4(b)). Only two measurements of $\mathcal{E}$ at symmetric flux offsets

are therefore required to determine and correct the error. Hence, at each correction step, we adjust the current in the on-chip flux line by an amount proportional to $\mathcal{E}$, realizing an integral controller that accumulates corrections over time.

7.2 Two-Tone Spectroscopy

After identifying the flux-symmetry points via single-tone spectroscopy, we probe the qubit's higher-level transitions in the GHz range using continuous-wave two-tone spectroscopy. In this experiment, we simultaneously apply (i) a drive tone ω_d through the charge-drive port to excite the qubit, and (ii) a readout tone to drive and measure the reflection from the readout cavity.

7.2.1 Plasmon/Fluxon Transitions

We apply a fixed readout tone[1] at the resonator frequency ω_R. Simultaneously, at each flux bias φ_{ext}, we sweep a weak drive tone[2] frequency ω_d. The reflected readout signal is continuously demodulated into its I and Q quadratures. Whenever the drive tone matches a transition frequency ω_{ij}, it drives sustained Rabi oscillations between states $|i\rangle$ and $|j\rangle$. These oscillations shift the resonator's dispersive response and produce a clear anomaly in the amplitude or phase of S_{11}.

A representative wide-range flux scan is shown in Fig. 7.5(a), plotting the real part of the reflection coefficient, $\text{Re}S_{11}$, versus drive frequency ω_d and external flux φ_{ext}. To assign the observed features, we numerically diagonalize the fluxonium Hamiltonian $\boldsymbol{H}(\varphi_{\text{ext}})$ [GK21] and overlay the computed transition frequencies from $\{|g\rangle, |e\rangle\}$ to $\{|f\rangle, |h\rangle\}$ as dashed curves in Fig. 7.5(b–c). The characteristic diamond-shaped feature around $\varphi_{\text{ext}} = \pi$ is clearly visible. Fitting these transitions yields the parameters $E_J = 4.30$ GHz, $E_C = 0.47$ GHz, and $E_L = 0.168$ GHz.

As discussed in Sec. 3.1.2, plasmon transitions $|g\rangle \to |f\rangle$ and $|e\rangle \to |h\rangle$ show only weak flux dependence away from the symmetry point, whereas fluxon transitions $|g\rangle \to |h\rangle$ and $|e\rangle \to |f\rangle$ vary linearly near $\varphi_{\text{ext}} = \pi$. Exactly at $\varphi_{\text{ext}} = \pi$, the Hamiltonian is invariant under $\boldsymbol{\varphi} \mapsto -\boldsymbol{\varphi}$, causing matrix elements between same-parity states to vanish and the plasmon lines to disappear (see Sec. 3.1.3). Experimentally, a full two-tone sequence can take up to a day. To accelerate measurements, we restrict the flux range to a narrow window around the target bias (e.g. $\varphi_{\text{ext}} = \pi$), lower the frequency resolution, and apply digital filters (row- and column-wise baseline subtraction, edge detection) to quickly identify transition peaks. This optimized workflow reduces the measurement time to a few hours while still extracting all relevant qubit transition frequencies.

1: To improve accuracy, one could fit or interpolate the single-tone spectroscopy data to adjust the readout pulse frequency as a function of φ_{ext}. In practice, however, we found that it is not essential, and a fixed readout frequency provides sufficient coverage for a broad flux scan.

2: In practice, we choose the drive-tone power to be just high enough to yield a clear contrast in S_{11}, but low enough to avoid excessive power broadening, AC-Stark shifts of the qubit lines or multi-photon transitions.

7.2.2 Hybridization with the Chain Mode

When measuring **CHIP−$\boldsymbol{\beta}$−FC**, we observed two overlapping diamond-shaped features in the two-tone map Fig. 7.6. The bare fluxonium Hamilto-

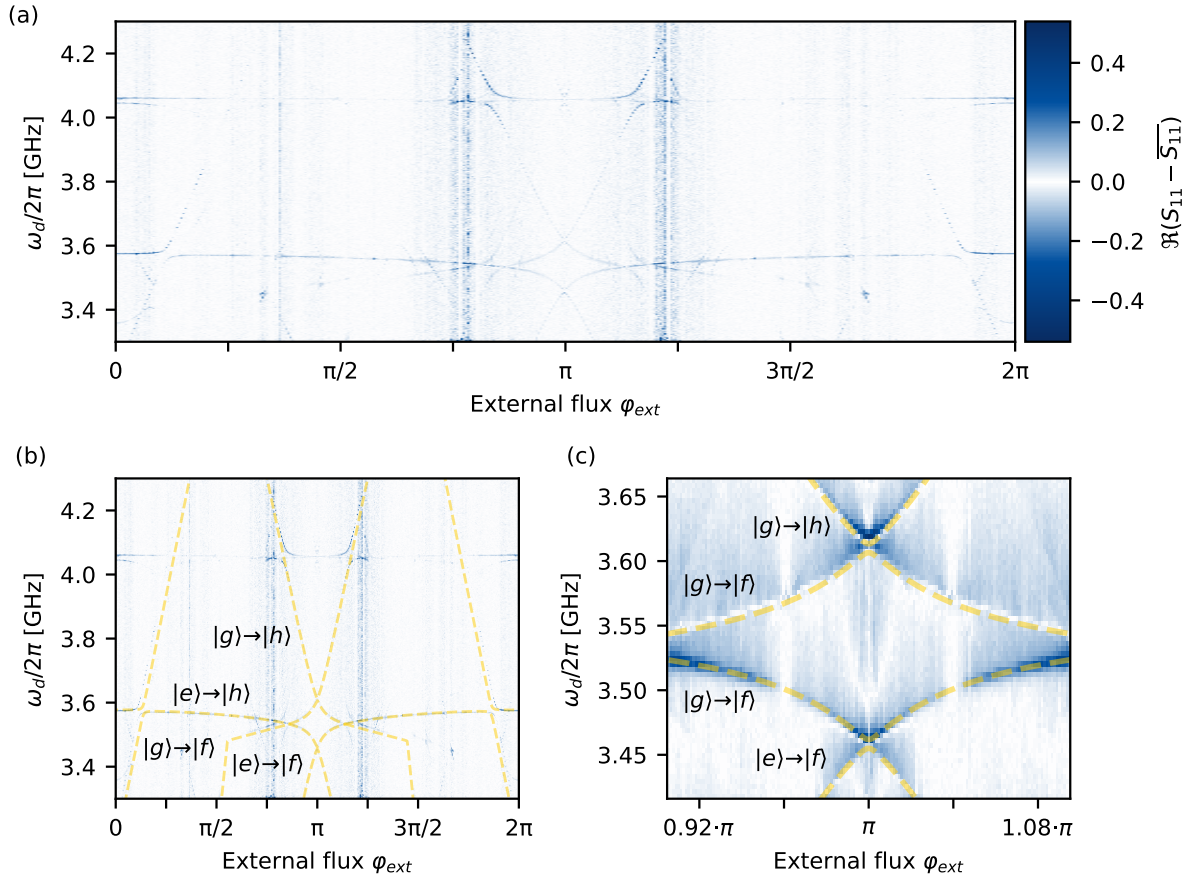

Fig. 7.5: Two-tone spectroscopy of CHIP−α. A continuous readout tone at ω_R is applied simultaneously with a swept pump tone ω_d. Color represents the real part of the reflected signal S_{11} on a common scale. **(a)** Wide-range scan over $0 \leq \varphi_{\text{ext}} \leq 2\pi$. **(b)** Same data with the theoretical fluxonium spectrum (dashed line) overlaid. (c) Zoom around the π-symmetry point, showing multiple excited-state transitions.

nian predicts only a single diamond. We attribute the secondary diamond to coupling between the fluxonium and a collective bosonic mode of the junction chain superinductance [Mas+12][1].

The Hamiltonian of the coupled system, written in terms of the fluxonium flux and charge operators $(\boldsymbol{\varphi}, \boldsymbol{n})$ and the chain mode ladder operators $(\boldsymbol{b}, \boldsymbol{b}^\dagger)$, is [VC15]

$$\begin{aligned} \boldsymbol{H} &= 4E_C\, \boldsymbol{n}^2 - E_J \cos\left(\boldsymbol{\varphi} - \varphi_{\text{ext}}\right) + \frac{1}{2} E_L\, \boldsymbol{\varphi}^2 \\ &\quad + \hbar\omega_{\text{chain}}\, \boldsymbol{b}^\dagger \boldsymbol{b} + \hbar g_{qc}\, (\boldsymbol{b}^\dagger + \boldsymbol{b})\, \boldsymbol{n}. \end{aligned} \tag{7.2}$$

We diagonalize this Hamiltonian numerically to obtain its eigenenergies and fit them to the experimental spectrum. Beyond reproducing the extra diamond, the chain-mode coupling also shifts the primary fluxonium transitions $\{|g\rangle, |e\rangle\} \to \{|f\rangle, |h\rangle\}$.

To improve fit robustness, one can include higher-level transitions as additional constraints. However, significant attenuation of the drive line above 4.5 GHz[2] prevented clear identification of these transitions. To further constrain the model, we use the qubit frequency independently measured via Ramsey techniques (Sec. 2.4.6), $\omega_q/2\pi = 2.35$ MHz at

[Mas+12]: Masluk et al. (2012), 'Microwave Characterization of Josephson Junction Arrays: Implementing a Low Loss Superinductance', Physical review letters

1: Unfortunately, we did not perform a wide-band two-tone sweep on **CHIP−β** before the flip-chip assembly, so its chain modes were not directly measured. In the analysis below, we assume that **CHIP−β** prior to flip-chip shares the same chain-mode frequencies and coupling strengths as **CHIP−β−FC** after flip-chip. However, the flip-chip process can modify the junction chain's boundary conditions and thus shift its mode frequencies and couplings.

2: We attribute this to filtering effects on the drive line.

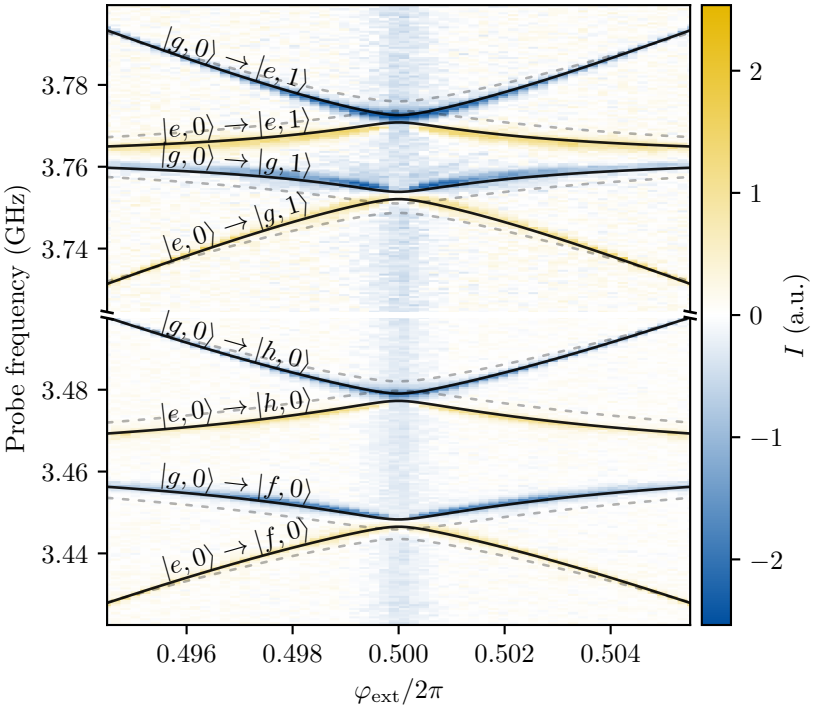

Fig. 7.6: Two-tone spectroscopy of CHIP–β–FC: reflected readout amplitude versus probe frequency and external flux φ_{ext}. The primary "diamond" at ∼ 3.46 GHz corresponds to the $\{|g\rangle, |e\rangle\} \to \{|f\rangle, |h\rangle\}$ transitions. A secondary diamond near ∼ 3.65 GHz, labeled by the transitions $\{|g,0\rangle, |e,0\rangle\} \to \{|g,1\rangle, |e,1\rangle\}$ (chain-mode Fock number in the ket), reveals coupling to a single chain mode. Solid lines are the best fit including that mode with parameters E_J/h = 4.886 GHz, E_C/h = 0.408 GHz, E_L/h = 0.135 GHz, $\omega_{\text{chain}}/2\pi$ = 3.650 GHz, and $g_{qc}/2\pi$ = 197 MHz. Dashed lines show an alternative fit fixing $\omega_{ge}/2\pi$ = 2.32 MHz, yielding E_J/h = 4.757 GHz, E_C/h = 0.427 GHz, E_L/h = 0.121 GHz, $\omega_{\text{chain}}/2\pi$ = 3.641 GHz, and $g_{qc}/2\pi$ = 203 MHz.

$\varphi_{\text{ext}} = \pi$. Fitting only the primary transitions [1] with the bare fluxonium model (no chain mode) yields $\omega_q/2\pi = 0.9$ MHz at $\varphi_{\text{ext}} = \pi$. Including the chain mode and fitting all transitions presented in Fig. 7.6 (solid-line fit) raises $\omega_q/2\pi = 1.7$ MHz, closer to the measured value.

Although adding more bosonic modes could improve theory–experiment agreement, each additional degree of freedom dramatically increases the Hilbert-space dimension and diagonalization cost. We therefore retain only the single lowest chain mode[2]. We perform two fits: one using the primary transitions with a cost function computed uniformly across each transition, and a second constraining the qubit gap $\omega_{ge}/2\pi \approx 2.35$ MHz via the penalty method (dashed-line fit in Fig. 7.6). Taking the average of these two fits as our best estimate and half their difference as the uncertainty yields

$$E_J/h = 4.82(6)\,\text{GHz}, \quad E_C/h = 0.418(9)\,\text{GHz}, \quad E_L/h = 0.128(7)\,\text{GHz},$$
$$\omega_{\text{chain}}/2\pi = 3.641(0)\,\text{GHz}, \quad g_{qc}/2\pi = 203\,\text{MHz}. \tag{7.3}$$

7.2.3 Qubit-membrane distance estimation

The vacuum-gap capacitors $C_m^{\pm}$ are highly sensitive to the separation distance d between the fluxonium and the membrane. To determine d precisely, we measured the same fluxonium device both before and after the flip-chip assembly. Plasmon transitions ($\omega_{ef}, \omega_{gh} \propto \sqrt{E_C}$) depend only weakly on E_C, whereas the low-frequency qubit transition ω_q depends exponentially on E_C, making it a much more sensitive probe of the capacitor gap. At half a flux-quantum bias, ω_q is well approximated

1: More specifically, transitions $\{|g\rangle \to |e\rangle\}$, $\{|g\rangle \to |f\rangle\}$, $\{|g\rangle \to |h\rangle\}$, $\{|e\rangle \to |f\rangle\}$, $\{|e\rangle \to |h\rangle\}$, together with the chain mode transitions $\{|g,0\rangle \to |g,1\rangle\}$, $\{|e,0\rangle \to |e,1\rangle\}$, at $\varphi_{\text{ext}} = \pi$ and $\varphi_{\text{ext}} = 0.495\,\pi$.

2: We clearly have a design and fabrication oversight in this version of the chip. In our previous design, the first chain mode was at 7 GHz [Naj+24, Fig. 4.19]. Here, the smaller junctions and reduced plasma frequency pushed the mode much lower, which not only perturbs the fluxonium transitions but may also introduce additional decoherence channels.

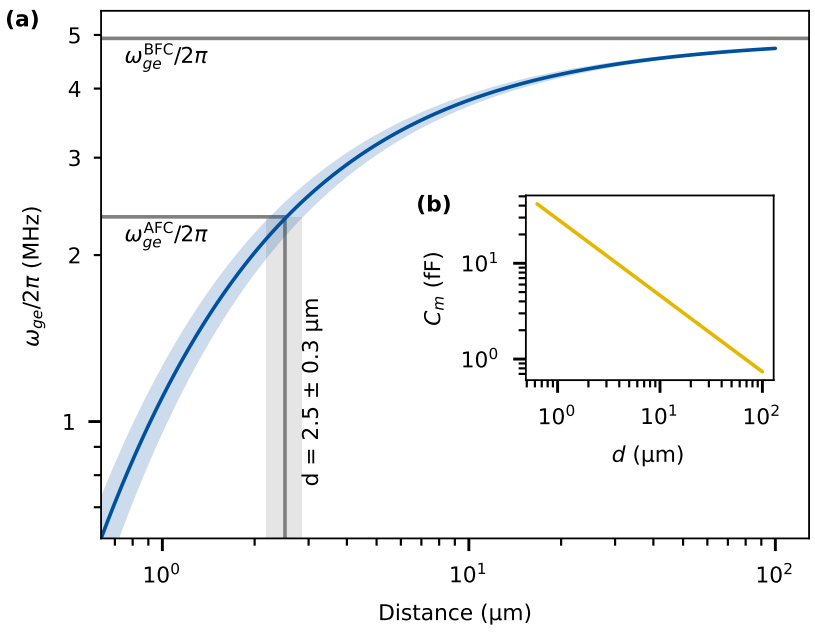

Fig. 7.7: Qubit frequency and mutual capacitance versus flip-chip separation. (a) The qubit frequency is plotted as a function of the flip-chip separation d. The frequency before flip-chip assembly (ω_q^{BFC}) and after flip-chip assembly (ω_q^{AFC}) are shown as a horizontal line. The shaded band represents the uncertainty in the qubit frequency due to the uncertainty in the fluxonium parameters. This translates to an uncertainty in the flip-chip separation of $d = 2.5(3)\,\mu\mathrm{m}$. **(b)** The mutual capacitance is plotted as a function of the flip-chip separation d. The interpolated capacitance model from Ansys Q3D simulations is used to compute the mutual capacitance $C_m(d) = \epsilon_0 \times 54 \times 10^{-9}/d^{0.8}$, where both the vacuum permittivity, ϵ_0, and the flip-chip separation, d, are in SI units.

by [Vak+24; Ard+24]

$$\omega_q = \frac{8\,2^{3/4}}{\sqrt{\pi}}\,E_J^{3/4}\,E_C(d)^{1/4}\exp\Big[-\sqrt{\frac{8E_J}{E_C(d)}}+\frac{14\,\zeta(3)\,E_L}{\sqrt{8\,E_J\,E_C(d)}}\Big]. \tag{7.4}$$

Here, $\zeta(3) \approx 1.2$ is the Riemann zeta function.

Using Ramsey techniques (Sec. 2.4.6), we measure the qubit frequency before flip-chip assembly as $\omega_q^{\mathrm{BFC}} = 4.93$ MHz. Inverting Eq. 7.4 for the known E_J and E_L yields the bare capacitive energy $E_C(\infty)$. Next, we simulate the qubit–membrane capacitance $C_m(d)$ versus the gap d using Ansys Q3D (see Fig. 7.7(b)). From the definition (Eq. 5.19), we compute

$$E_C(d) = \frac{e^2}{\frac{e^2}{E_C(\infty)} - C_m(d)}.$$

Substituting $E_C(d)$ into Eq. 7.4 then gives the predicted qubit frequency $\omega_q(d)$ for fixed E_J, E_L, and ω_q^{BFC}. To account for parameter uncertainty, we assign errors to E_J and E_L (from our two-tone spectroscopy in Sec. 7.2) and propagate these via a Monte Carlo analysis. After flip-chip assembly, the qubit frequency drops to $\omega_q^{\mathrm{AFC}} = 2.35$ MHz, which corresponds to a capacitor gap $d = 2.5(3)\,\mu\mathrm{m}$. The resulting shaded confidence band appears in Fig. 7.7(a).

7.3 Basic qubit control

Up to now, our experiments have been limited to spectroscopy with long, continuous-wave drives, without precise control of the qubit state. Having characterized the fluxonium transition frequencies, we can now perform basic time-domain qubit control experiments. Later in this chapter (Sec. 7.5), we describe in detail the qubit readout procedure used to determine the qubit state. For now, we simply note that we

employ dispersive readout, as introduced in Sec. 2.5.3, and that we can distinguish between the first four qubit states.

7.3.1 Rabi Spectroscopy

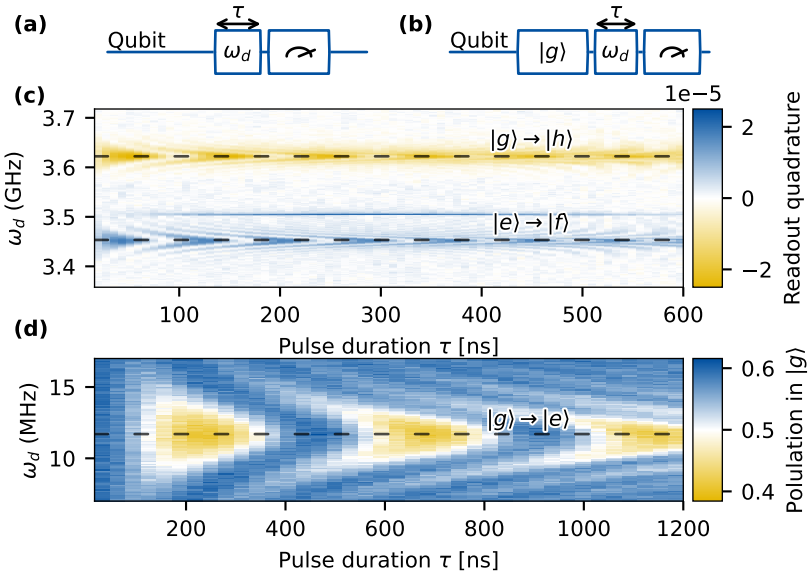

Fig. 7.8: Rabi spectroscopy. (a,b) Pulse sequences for higher-level spectroscopy and for $|g\rangle \rightarrow |e\rangle$ Rabi oscillations. **(c)** Mean readout quadrature versus drive frequency ω_d and pulse duration τ, revealing the $|g\rangle \rightarrow |h\rangle$ and $|e\rangle \rightarrow |f\rangle$ resonances (dashed lines). **(d)** Ground-state population P_g versus drive amplitude and pulse duration τ for the $|g\rangle \rightarrow |e\rangle$ transition.

Higher-level transitions. After identifying the $|g\rangle \rightarrow |f\rangle$ and $|e\rangle \rightarrow |h\rangle$ lines via two-tone spectroscopy, we perform Rabi spectroscopy at $\varphi_{\text{ext}} = \pi$ (pulse sequence in Fig. 7.8(a)). We drive the qubit at frequency ω_d for a variable duration τ, then immediately measure the reflected quadrature. We expect to observe behaviour similar to that shown in Fig. 2.11. Fig. 7.8(c) plots the mean quadrature as a function of ω_d and τ. When the drive is detuned by $\delta = \omega_{ij} - \omega_d$, we observe Rabi oscillations at frequency

$$\omega_r = \sqrt{\Omega^2 + \delta^2},$$

where Ω is the on-resonance Rabi rate (Sec. 2.4.4). The slowest oscillations (zero detuning) occur at $\omega_{ef}/2\pi = 3.453\,\text{GHz}$ and $\omega_{gh}/2\pi = 3.622\,\text{GHz}$ (dashed lines), corresponding to coherent control on $\sim 50\,\text{ns}$ timescales.

Lower-level transition. To probe the fundamental $|g\rangle \rightarrow |e\rangle$ transition, we first prepare the qubit in $|g\rangle$ using the active reset protocol described later in Sec. 7.4.1. Then, we drive at frequency ω_d for duration τ and measure the ground-state population P_g. Fig. 7.8(d) displays $P_g(\omega_d, \tau)$, showing detuning-dependent Rabi oscillations. The slowest oscillations (horizontal dashed line) indicate zero detuning and directly identify the qubit frequency ω_{ge}.

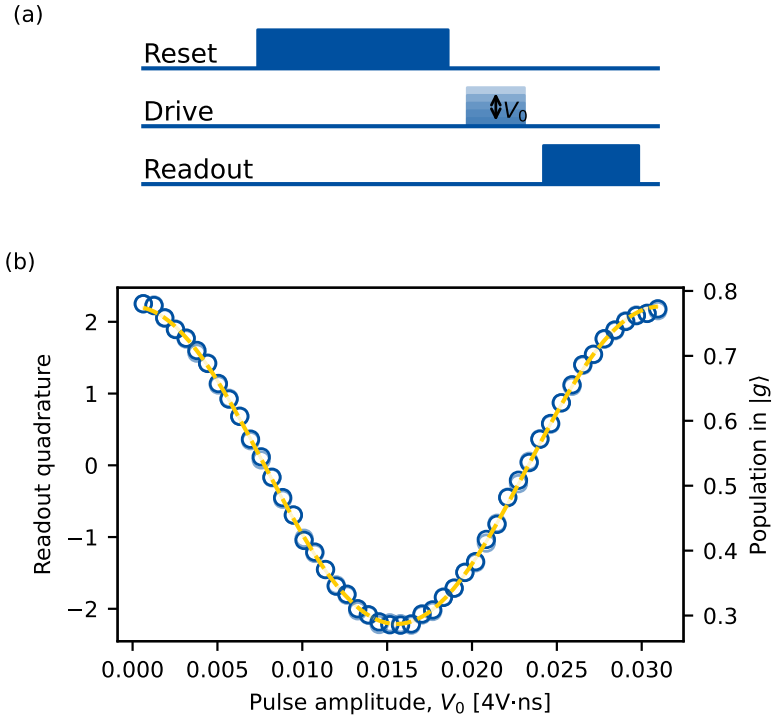

Fig. 7.9: Calibration of the $|g\rangle \rightarrow |h\rangle$ π-pulse. (a) Pulse sequence: reset to $|g\rangle$, apply a 48 ns pulse at ν_{gh} with variable amplitude, then read out. **(b)** Measured I quadrature (solid) and ground-state population $P(|g\rangle)$ (shaded) versus drive amplitude; the dashed yellow line is a sinusoidal fit that determines the π-pulse amplitude.

7.3.2 π-Pulse

A π-pulse implements a full flip of the qubit population on resonance. From Eq. 2.58, at zero detuning ($\delta_\omega = 0$) the rotation angle reduces to

$$\theta = \tau\,\Omega \propto \tau\,V_0.$$

To achieve $\theta = \pi$, one adjusts either the pulse duration τ or the drive amplitude V_0. In practice, we choose $\tau \ll T_{1,2}$ to increase fidelity, while avoiding excessively short pulses whose broad spectra distort in the control lines and drive unwanted transitions. Thus, quantum gate pulses must balance temporal speed against spectral purity. For a carrier at f_0 and pulse duration τ, the Fourier-transform-limited full width at half maximum (FWHM) is[1]

$$\Delta f_{\mathrm{FWHM}} \approx \begin{cases} 0.886/\tau, & \text{rectangular envelope,} \\ 0.441/\tau, & \text{Gaussian envelope.} \end{cases}$$

The pulse envelope must vary slowly compared to the carrier ($1/\tau \ll f_0$) to avoid distorting the carrier frequency. In Fig. 7.10 we plot the Fourier magnitude of a 40 ns Gaussian pulse at $\omega_d/2\pi = 1$ GHz, which produces a narrow spectral peak at the carrier. However, for a qubit frequency as low as 2.35 MHz, the envelope must be much longer. Fig. 7.11 shows that 50 ns and 400 ns pulses cannot resolve the 2 MHz carrier, whereas a 700 ns pulse yields a clear peak at ω_d and a correspondingly narrower bandwidth.

Accordingly, we use $\tau \approx 48$ ns π-pulses for the $\{|g\rangle, |e\rangle\} \rightarrow \{|f\rangle, |h\rangle\}$ transitions and $\tau \approx 700$ ns for the $|g\rangle \rightarrow |e\rangle$ transition. Since our Quantum

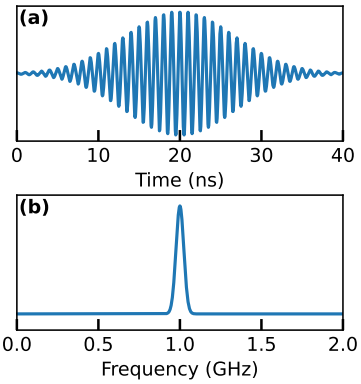

Fig. 7.10: High-frequency control pulse. (a) Gaussian envelope of a 40 ns pulse at 1 GHz. **(b)** Fourier magnitude of the pulse, showing a narrow spectral peak of width $\sim 0.44/(40\,\mathrm{ns})$.

1: The factor 0.886 for a rectangular pulse follows from solving $\mathrm{sinc}^2(\pi f \tau) = 1/2$, and 0.441 for a Gaussian pulse follows from $\tau\,\Delta f = 4\ln 2/\pi \approx 0.882$.

Machine OpX provides voltage resolution of order $10^{-5}\,[4\,\mathrm{V}\cdot\mathrm{ns}]$ but timing resolution of only $\sim 4\,\mathrm{ns}$, we fine-tune the pulse area by adjusting the drive amplitude V_0.

Fig. 7.9(b) shows the measured I quadrature and the ground-state population[1] $P(|g\rangle)$ versus V_0 for $\tau = 48\,\mathrm{ns}$. We fit these data to

$$A\cos\left(\pi\,\frac{V_0}{V_0^{\pi}} + \phi\right) + \text{offset},$$

extracting the amplitude V_0^{π}. We implement $\pi/2$-pulses by setting amplitude to $V_0^{\pi}/2$.

7.4 State Preparation

At our operating temperature of 20 mK, the thermal energy corresponds to

$$k_B T/h \approx 417\,\mathrm{MHz}.$$

Since the qubit transition frequency $\omega_{ge}/2\pi \lesssim 10$ MHz is much lower than $k_B T/h$, the $|g\rangle$ and $|e\rangle$ states are thermally populated with nearly equal probability. In contrast, the higher-manifold levels ($\sim$ 3 GHz) remain effectively unpopulated. Therefore, we must actively prepare the qubit in pure $|g\rangle$ or $|e\rangle$ states. Below, we describe our preparation methods and quantify fidelity of the main preparation methods by mapping populations onto the unoccupied auxiliary levels $|f\rangle$ and $|h\rangle$.

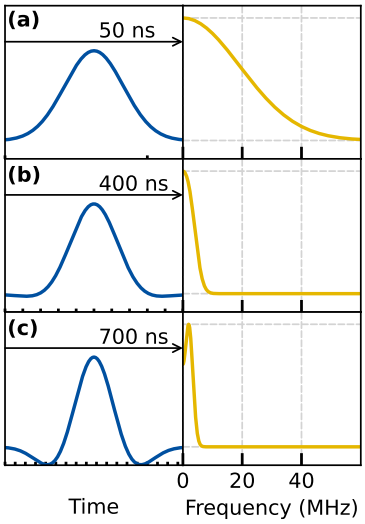

Fig. 7.11: Low-frequency control pulses. (a–c) Left: time-domain waveforms of a 2 MHz carrier modulated by Gaussian envelopes of duration (a) 50 ns, (b) 400 ns, and (c) 700 ns. Right: corresponding Fourier magnitudes (yellow), showing how longer pulses (τ) produce narrower spectral bandwidths.

1: Note that I and P_g are formally distinct observables and not bijectively related; with large ensemble averages they become proportional.

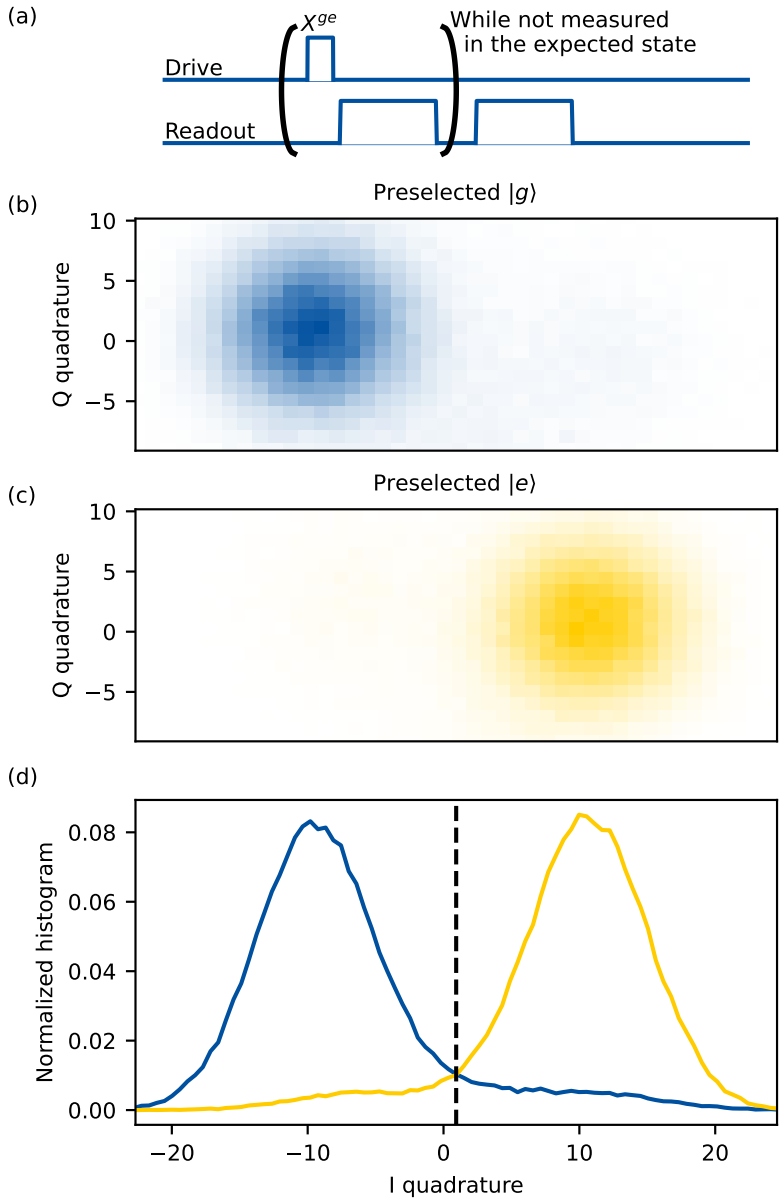

Fig. 7.12: Preparation by selection for **CHIP−α**. (a) Pulse sequence. (b-c) Single-shot histograms obtained after preselecting on measurement outcomes with a quadrature signal below or above the threshold, thereby preparing the $|g\rangle$ or $|e\rangle$ state, respectively, with limited fidelity. (d) Normalized I-quadrature histograms with the separation threshold (dashed line). This measuments gives us the following correlation matrix:

7.4.1 Qubit Preparation Methods

Post-Selection Reset. Our simplest reset protocol uses single-shot readout and post-selection. As illustrated in Fig. 7.12(a), we apply a strong, frequency-selective measurement pulse to project the qubit onto $|g\rangle$ or $|e\rangle$. If the measurement outcome matches the target state, we proceed; otherwise, we apply a corrective π-pulse and repeat until the desired state is achieved. This method requires only measurement and at most one control pulse, introduces no additional microwave power into the cryostat, and—given fast readout—has a turnaround time comparable to active reset techniques described in the following. Post-selection relies on well-separated single-shot readout distributions, which in optimal devices can exceed $\sim 99\%$ fidelity [Wal+17]. Unfortunately, for the main chip (**CHIP**$-\boldsymbol{\beta}-$**FC**) used in this work the separation was lower and this method was not used. Figure 7.12(b–c) shows the single-shot histograms for **CHIP**$-\boldsymbol{\alpha}$ obtained after preselecting on measurement outcomes with a quadrature signal below or above the threshold, thereby preparing the $|g\rangle$ or $|e\rangle$ state, respectively, with limited fidelity.

By integrating the histogram on either side of the threshold (Fig. 7.12(d)), we obtain the correlation matrix

$$\begin{pmatrix} P(I_2 < T_I \mid I_1 < T_I) & P(I_2 > T_I \mid I_1 < T_I) \\ P(I_2 < T_I \mid I_1 > T_I) & P(I_2 > T_I \mid I_1 > T_I) \end{pmatrix} = \begin{pmatrix} 90.3\% & 9.7\% \\ 7.7\% & 92.3\% \end{pmatrix},$$

where I_1 and I_2 are the quadratures of the first and second measurements, respectively, and T_I is the threshold indicated by the dashed line. We cannot directly infer the preparation and readout fidelity from this matrix because decoherence during the readout is not negligible[1]. Nevertheless, in the idealized limit of negligible intrinsic and measurement-induced decoherence, this matrix implies a preparation fidelity of $\approx 95\%$ for this method. This highlights that higher-fidelity state discrimination would enable an efficient post-selection-based reset.

1: Later, in Sec. 7.5.4, we show that the intrinsic readout fidelity can be determined by independently measuring the state-preparation fidelity.

Reset via Readout Resonator. A more active reset method exploits the qubit–readout-resonator coupling by driving a sideband transition. We apply a strong microwave drive to the readout resonator, resonant with the $|g,0\rangle \rightarrow |e,1\rangle$ or $|e,0\rangle \rightarrow |g,1\rangle$ sideband transition, where the first index denotes the qubit state and the second index the resonator photon number. Due to fast decay of the readout resonator, this method swaps one qubit state while leaving the other unchanged, thereby resetting the qubit. This process is illustrated in Fig. 7.13. Since these sideband transitions are forbidden at the sweet spot ($\varphi_{\text{ext}} = \pi$), we first ramp the external flux to $\varphi_{\text{ext}} = \pi \pm 0.002\pi$ over $\sim 4\,\mu$s, to maximize preparation contrast without inducing non-adiabatic transitions. The full pulse sequence is shown in Fig. 7.14(a). This method is independent of measurement fidelity but requires reduced attenuation on the readout-pump line, which can introduce additional noise and heating. Using our first heavy-fluxonium chip, we demonstrated a reset fidelity of 97.7% [Naj+24].

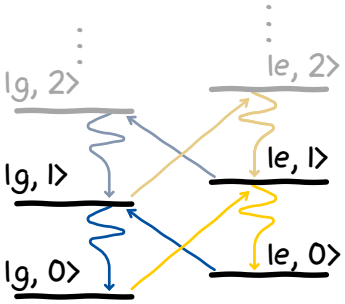

Fig. 7.13: Sideband preparation illustration.

Figure 7.14(b) shows the population of **CHIP−$\boldsymbol{\alpha}$** as a function of external flux phase φ_{ext} and sideband drive frequency. The calculated sideband resonances $\omega_R \pm \omega_q(\varphi_{\text{ext}})$ are shown as dark violet dashed lines, with their linear approximation $\omega_R \pm 2\pi E_L|\varphi_{\text{ext}} - \pi|$ in violet. To probe the reset dynamics, we fix $\varphi_{\text{ext}} = \pi + 0.029$ and vary the sideband-pulse duration τ. Figure 7.14(c) shows the state purity rising exponentially to saturation, with time constants of 1.3 μs for $|g\rangle$ and 1.0 μs for $|e\rangle$. The sideband drive may be applied via the readout line (typical) or via the charge line, as shown in Fig. 7.14(d). We observed no significant difference between the two methods.

Reset via Auxiliary Chain Mode. Another method of an active reset is to exploit the coupling between the qubit and the collective mode of the junction chain. To initialize the qubit **CHIP−$\boldsymbol{\beta}$−FC** in a pure state, we exploit a two-photon transition to the second chain mode at 6.42 GHz.

First, to prepare the $|e\rangle$ state, we drive the $|g,0\rangle \rightarrow |e,1\rangle$ transition, where first index denotes the qubit state and second index denotes the second chain mode state, with a continuous-wave tone at $(\omega_{\text{chain}} + \omega_q)/2 = 3.2118$ GHz. Owing to the short lifetime of the chain-mode, this quickly transfers all qubit population to $|e\rangle$ via optical pumping. We also investigated preparing $|g\rangle$ via the reverse two-photon process $|e,0\rangle \rightarrow |g,1\rangle$, but found its fidelity to be significantly lower. Instead, we reset the qubit to $|g\rangle$ by first preparing $|e\rangle$ as above and then applying a resonant π-pulse on the $|e,0\rangle \rightarrow |g,0\rangle$ transition.

This is the main method we use for reset of **CHIP−$\boldsymbol{\beta}$−FC**. For the main experiments of this work (spectrum analyser using qubit), it's necessary to maintain a equal initialization time for both $|g\rangle$ and $|e\rangle$ states. To ensure this, we pad the faster $|e\rangle$ preparation with an additional delay equal to the π-pulse duration, ensuring that both $|g\rangle$ and $|e\rangle$ initializations occupy the same total time.

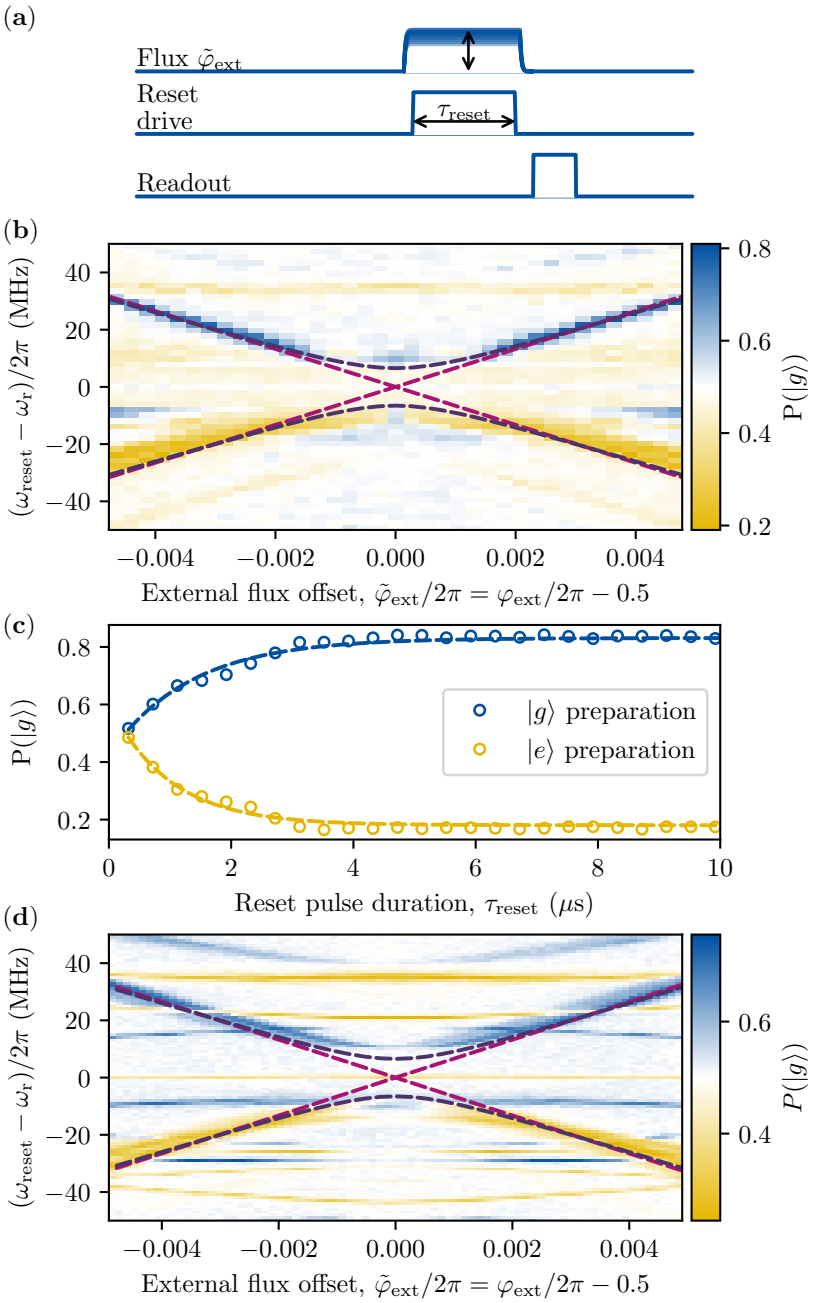

Fig. 7.14: Sideband preparation. (a) Pulse sequence. **(b)** 2D map of the population in the $|g\rangle$ state after sideband preparation for a fixed drive duration of 20 μs using the readout drive. The map is shown as a function of the external flux phase φ_{ext} and the sideband drive frequency Dark violet lines are the calculated sideband resonances $\omega_R \pm \omega_q(\varphi_{\text{ext}})$ while, violet lines are their linear approximation $\omega_R \pm 2\pi E_L|\varphi_{\text{ext}} - \pi|$. **(c)** Time evolution of the state purity for $|g\rangle$ and $|e\rangle$ preparation protocols. **(d)** 2D map of the population in the $|g\rangle$ state after sideband preparation for a fixed drive duration of 20 μs using the drive through the charge line. The map is shown as a function of the external flux phase φ_{ext} and the sideband drive frequency. Dark violet lines are the calculated sideband resonances $\omega_R \pm \omega_q(\varphi_{\text{ext}})$ while, violet lines are their linear approximation $\omega_R \pm 2\pi E_L|\varphi_{\text{ext}} - \pi|$.

7.4.2 Qubit Preparation Fidelity

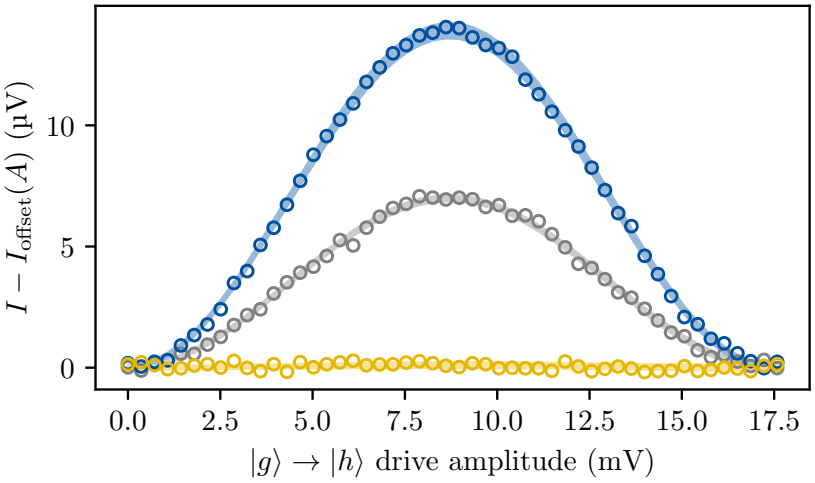

Fig. 7.15: Qubit preparation fidelity. Rabi fringes obtained by driving the $|g\rangle \rightarrow |h\rangle$ transition after three different initialization sequences: $|g\rangle$ (blue), $|e\rangle$ (red), and thermal state (gray). Solid lines are simultaneous fits to Eq. 7.7, yielding the preparation fidelities η_g and η_e.

To evaluate the fidelity of qubit preparation, we perform a three-step protocol and explore an auxiliary level $|h\rangle$ that is unpopulated in thermal equilibrium at 10 mK.

Step 1: state preparation. The system is initialized in one of three target states: ground $|g\rangle$, excited $|e\rangle$, or thermal equilibrium. The true post-initialization reduced density operator in the qubit subspace is

$$\rho_0^i = p_i \, |g\rangle\langle g| + (1 - p_i) \, |e\rangle\langle e| \,, \tag{7.5}$$

where $i \in \{g, e, \text{thermal}\}$ labels the preparation and p_i is the population in $|g\rangle$. In particular,

$$p_e = 1 - \eta_e, \quad p_g = \eta_g, \quad p_{\text{thermal}} = \frac{1}{2}, \tag{7.6}$$

where η_e and η_g denote the preparation fidelities of $|e\rangle$ and $|g\rangle$, respectively, and $p_{\text{thermal}} = 1/2$ corresponds to an equiprobable mixture of $|g\rangle$ and $|e\rangle$, i.e. a baseline fidelity of 50%.

Step 2: transfer to auxiliary state $|h\rangle$. In order to measure η_e and η_g, we apply a 152-ns pulse resonant with the $|g\rangle \rightarrow |h\rangle$ transition, varying its amplitude to induce a Rabi rotation by angle θ between $|g\rangle$ and $|h\rangle$. Restricting attention to the diagonal elements, this pulse transforms

$$|g\rangle\langle g| \longrightarrow \frac{1 + \cos\theta}{2} \, |g\rangle\langle g| + \frac{1 - \cos\theta}{2} \, |h\rangle\langle h| \,,$$

while $|e\rangle$ remains uncoupled. The full state after the pulse is

$$\begin{aligned} \rho_0^i \longrightarrow (1 - p_i) \, |e\rangle\langle e| + p_i \frac{1 + \cos\theta}{2} \, |g\rangle\langle g| \\ + p_i \frac{1 - \cos\theta}{2} \, |h\rangle\langle h| \,. \end{aligned} \tag{7.7}$$

Step 3: dispersive readout. Finally, for each initial preparation we perform a 2-μs dispersive readout optimized to distinguish the $\{|g\rangle, |e\rangle\}$

manifold from $|h\rangle$. We obtain three Rabi oscillation contrasts C^i (Fig. 7.15), from which we extract

$$C^i = v\, p_i,$$

where v is a global visibility factor. Since $p_{\text{th}} = \frac{1}{2}$, $p_g = \eta_g$, and $p_e = 1-\eta_e$, we find

$$\eta_g = \frac{C^g}{2\,C^{\text{th}}} = 99.5(1.7)\%, \quad \eta_e = 1 - \frac{C^e}{2\,C^{\text{th}}} = 99.2(5)\%.$$

Uncertainties are estimated via bootstrapping, where fits are repeated using subsets of the amplitude points and the standard deviation of the extracted parameters is taken as the error.

7.5 Qubit Readout

The dispersive shift of the readout resonator is used to discriminate between the qubit states. In Sec. 2.5.3, we described that readout outcomes are represented in the IQ plane, where I and Q correspond to the in-phase and quadrature components of the demodulated reflected signal. Up to this point, we implicitly used an optimized readout configuration, with the separation axis aligned along the I component. In this section, we delve into the optimization of the readout process.

The pulse reflected from the cavity does not contain the same amount of information throughout its duration. At times short compared to the cavity lifetime, the intracavity field has not had time to build up, and there is essentially no information about the qubit state. Conversely, if the pulse duration is comparable to the qubit lifetime, the qubit may flip during the measurement, so the end of the pulse carries less information about the initial qubit state. The measurement noise, dominated by amplifier noise, remains constant. To maximize the signal-to-noise ratio, one should (i) shape the incoming pulse waveform to populate the cavity as quickly as possible and extract information about the qubit state before it relaxes, and (ii) adapt the demodulation weights to favor times of maximum informational content.

We begin by outlining the algorithm that computes the IQ quadratures of the reflected signal and derives the optimal integration weights to maximize the separation between the two qubit states. Next, we introduce the criteria used to evaluate readout performance and describe the method for selecting the decision threshold. We then discuss various readout protocols for distinguishing $\{|g\rangle, |e\rangle\}$, including strategies for cases where the dispersive shift alone is insufficient for single-shot discrimination. Finally, we present the measured readout fidelities for **CHIP−β−FC**.

7.5.1 Readout Optimization: Integration Weights

In our experiments, we probe the qubit state by sending a microwave pulse at the readout resonator frequency with the envelope illustrated in Fig. 7.16(a) and measuring the reflected signal. The reflected wave is

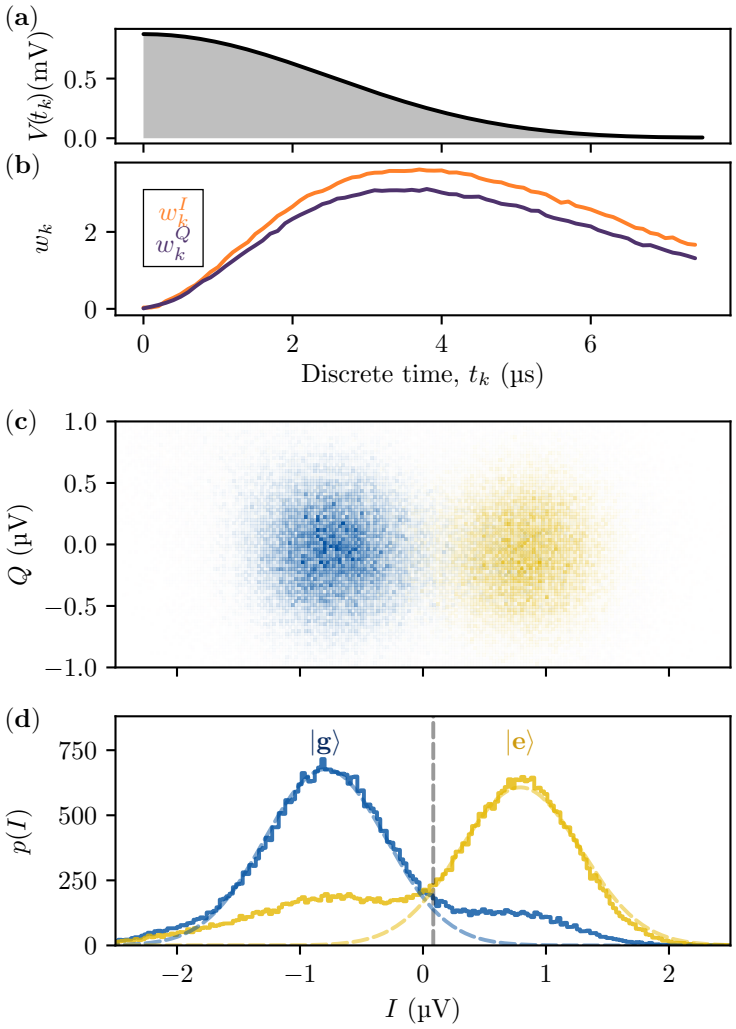

Fig. 7.16: Qubit readout fidelity for **CHIP–β–FC**. **(a)** Amplitude envelope of the 7.5 µs readout pulse. **(b)** Time-domain integration weights w_k^I and w_k^Q used to compute the I–Q quadratures: these weights are chosen as the optimal linear discriminant (matched filter) between the average $|g\rangle$ and $|e\rangle$ responses. **(c)** Histogram of single-shot outcomes in the I–Q plane, collected after preparing $|g\rangle$ (blue) or $|e\rangle$ (yellow). **(d)** Histograms of the I quadrature for the same data. Colored dashed lines show the fitted Gaussian components (equal variance) that represent the ideal distributions in the absence of preparation and readout errors. The yellow vertical dashed line indicates the discrimination threshold chosen to equalize the contrast of the spectra.

first combined in an image-rejection mixer with a local-oscillator tone, down-converting it to an intermediate frequency (IF). This IF waveform, $s(t)$, is digitized in the FPGA at times $t_k = k\,\Delta t$, with $\Delta t = 4$ ns, yielding samples $s(t_k)$. Each sample is demodulated into in-phase and quadrature components via

$$I_k = s(t_k)\cos(\omega_{\mathrm{IF}} t_k), \qquad Q_k = s(t_k)\sin(\omega_{\mathrm{IF}} t_k).$$

To correct any residual phase offset and concentrate all state-dependent information onto the I quadrature, we apply at each time step a rotation mixed with weighting. Defining two sets of weights w_k^c, w_k^s, the final integrated quadratures become

$$\begin{pmatrix} \tilde{I} \\ \tilde{Q} \end{pmatrix} = \sum_{k=1}^{N} \begin{pmatrix} w_k^c & w_k^s \\ -w_k^s & w_k^c \end{pmatrix} \begin{pmatrix} I_k \\ Q_k \end{pmatrix}.$$

This form both rotates the instantaneous (I_k, Q_k) into the optimal basis and applies time-dependent weights that emphasize intervals of maximum state separation. To determine the best set of weights $\{w_i^c, w_i^s\}$, we measure the set of readout components when the qubit is prepared in state $|u\rangle$: $\{I_i^{|u\rangle}, Q_i^{|u\rangle}\}$ as well prepared in state $|v\rangle$: $\{I_i^{|v\rangle}, Q_i^{|v\rangle}\}$. Maximizing information on the I component, leads to maximizing the difference

between the two states $\tilde{I}^{|u\rangle} - \tilde{I}^{|v\rangle}$, which leads to the following expression for each weights normalized by $\sum_i w_i = 1$:

$$w_i^c = \frac{I_i^{|u\rangle} - I_i^{|v\rangle}}{\sum_j \left(I_j^{|u\rangle} - I_j^{|v\rangle}\right)}, \quad w_i^s = \frac{Q_i^{|u\rangle} - Q_i^{|v\rangle}}{\sum_j \left(Q_j^{|u\rangle} - Q_j^{|v\rangle}\right)}. \tag{7.8}$$

Fig. 7.16(b) illustrates the resulting weights for our qubit **CHIP–β–FC**, showing that most information arrives in the beginning of the cavity ring-down and indicating the importance of using non-uniform integration weights. Fig. 7.16(c) shows the histogram of the resulting single-shot outcomes in the I–Q plane, where the two Gaussian blobs corresponding to the $|g\rangle$ and $|e\rangle$ states are clearly separated.

7.5.2 Readout Optimization: Criteria, Threshold

To optimize readout, we maximize the separation of two[1] Gaussian response blobs in the (I, Q) plane for states $|u\rangle$, $|v\rangle$. Each blob satisfies

$$p\left(z \mid |u\rangle\right) = \frac{1}{\sqrt{2\pi}\,\sigma_u} \exp\left(-\frac{|z-\mu_u|^2}{2\sigma_u^2}\right).$$

Assuming equiprobable states and equal variances[2], the likelihood ratio becomes

$$\frac{p(|u\rangle \mid z)}{p(|v\rangle \mid z)} = \exp\left(|z-\mu_v|^2 - |z-\mu_u|^2\right).$$

Since there is no information in the Q component, we consider z real. Thus, the decision rule $|z-\mu_v|^2 - |z-\mu_u|^2 \gtrless 0$ simplifies via

$$|z-\mu_v|^2 - |z-\mu_u|^2 = 2(\mu_u - \mu_v)\left(z - \frac{\mu_u + \mu_v}{2}\right).$$

The sign of this criterion indicates the most probable state. Thus, the decision rule is simply to compare the measured quadrature z with the threshold:

$$\gamma = \frac{\mu_u + \mu_v}{2},$$

which lies midway between the two state means. We quantify readout fidelity by the misclassification probability. Without loss of generality, assume $\mu_u < \mu_v$. Thus, if $z > \gamma$, we infer $|v\rangle$; otherwise, $|u\rangle$. The probability of classifying $|u\rangle$ as $|v\rangle$ is

$$p(z > \gamma \mid |u\rangle) = \int_\gamma^\infty p(z \mid |u\rangle)\, dz = \frac{1}{2}\left[1 - \mathrm{erf}\left(\frac{\mu_v - \mu_u}{\sqrt{2}\,\sigma}\right)\right], \tag{7.9}$$

so the optimization metric becomes $\frac{\mu_v - \mu_u}{\sigma}$.

We optimize the readout by adjusting the pulse amplitude and frequency via a two-dimensional sweep, while the waveform shape and duration are tuned manually. Fitting histograms during this process proved impractical, as fits become unstable when the signal-to-noise ratio drops or the distributions distort. Instead, we compute statistical properties directly from the acquired samples.

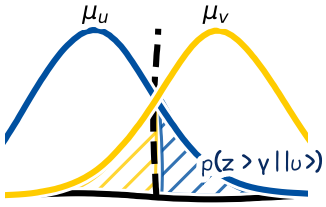

Fig. 7.17: Misclassification due to Gaussian tail overlap. The tails of the two Gaussian distributions overlap beyond the decision threshold γ, causing events from one distribution to be incorrectly assigned to the other.

1: In general, fluxonium readout may exhibit multiple blobs, especially when higher states (e.g., $|f\rangle$ or $|h\rangle$) are populated. We tune the readout for each targeted transition to separate responses along Re(z) only. Theoretically, this may not always be achievable, however, in practice, it was never a limitation.

2: Using Bayes' theorem, we invert this relation to obtain the probability that the qubit was in $|u\rangle$ given the quadrature measurement z. The likelihood ratio becomes

$$\frac{p(|u\rangle \mid z)}{p(|v\rangle \mid z)} = \frac{p(z \mid |u\rangle)\, p(|u\rangle)}{p(z \mid |v\rangle)\, p(|v\rangle)}$$
$$= \frac{\sigma_v}{\sigma_u} \exp\left[-\frac{|z-\mu_u|^2}{2\sigma_u^2} + \frac{|z-\mu_v|^2}{2\sigma_v^2}\right].$$

Then set $\sigma_u \approx \sigma_v$ and equal priors.

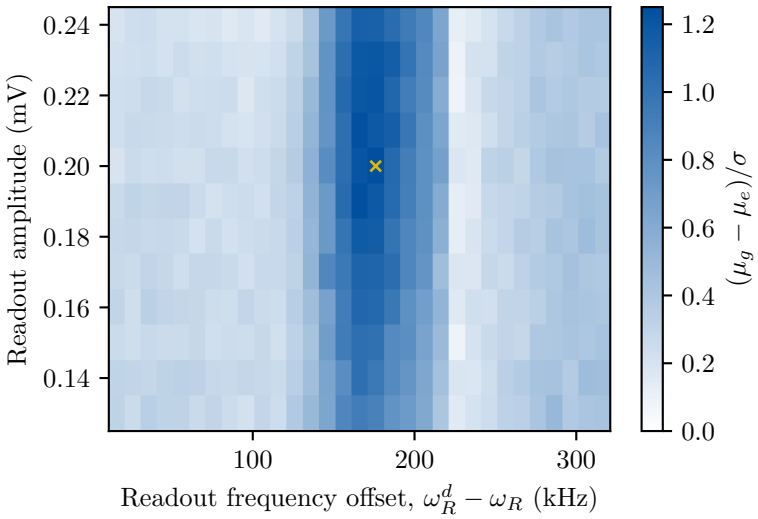

Fig. 7.18: Readout optimization. Colormap of the readout efficiency as a function of the readout frequency offset and amplitude. The optimal parameters are indicated by the yellow cross.

To optimize the g–e readout, we prepare the qubit in either the $|g\rangle$ or $|e\rangle$ state and record the measurement. For each combination of pulse amplitude and frequency, we repeat this procedure more than 10^4 times. From each amplitude-frequency point, we derive the optimal integration weights and evaluate the mean and standard deviation of the measured signal for both states. The resulting colormap is shown in Fig. 7.18, indicating the optimal readout parameters for discriminating $|g\rangle$ and $|e\rangle$.

7.5.3 Readout Protocols in the $\{|g\rangle, |e\rangle\}$ Manifold

In order to discriminate $\{|g\rangle, |e\rangle\}$ states, we use one of two methods depending on the dispersive shift:

- **Direct readout in the lower manifold.** When the dispersive shift χ_{ge} is sufficiently large, the protocol above discriminates $|g\rangle$ and $|e\rangle$ in a single shot. However, fabrication uncertainties sometimes reduce χ_{ge} below this threshold, preventing reliable lower-manifold readout.
- **Readout via the upper manifold.** If χ_{ge} is insufficient, we apply a calibrated π-pulse to transfer population into an upper level (either $|h\rangle$ or $|f\rangle$), then perform the same readout protocol. Thus we map determination between states in the lower manifold to determination between the lower- and upper-manifolds, which typically exhibits larger dispersive shifts. Although this two-step protocol adds error from the π-pulse, our first sample achieved overall readout fidelities up to 94% [Naj+24].

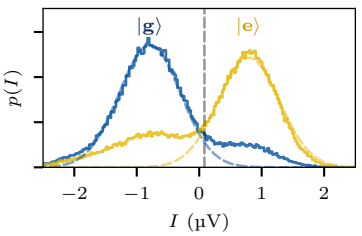

Fig. 7.19: Readout quadrature histogram. Histograms of the I quadrature of single-shot outcomes after preparation in $|g\rangle$ (blue) or $|e\rangle$ (yellow). Colored dashed lines show the fitted Gaussian components that represent the ideal distributions in the absence of preparation and readout errors (see text for more details). The vertical dashed line indicates the discrimination threshold chosen to equalize the contrast of the Rabi oscillations.

7.5.4 Qubit readout fidelity

One could imagine that an ideal readout would yield two Gaussian distributions[1] corresponding to the $|g\rangle$ and $|e\rangle$ states (dashed lines in Fig. 7.19). However, finite state-preparation fidelity and decoherence during readout cause the measured histograms to deviate substantially from these ideal Gaussians. Because the state-preparation fidelities are

1: This is shown for illustration by fitting each histogram (obtained after preparation in $|g\rangle$ or $|e\rangle$) to a sum of two Gaussians with equal variances, and then plotting only the Gaussian associated with the initially prepared state.

measured independently, we correct for preparation errors to extract the intrinsic readout fidelity.

Readout fidelities are defined as $F_g = P(I < T_I \mid |g\rangle\langle g|)$ and $F_e = P(I > T_I \mid |e\rangle\langle e|)$, where T_I is the state-discrimination threshold (dashed gray line in Fig. 7.16(d)), chosen to maximize the contrast of Rabi oscillations. Experimentally, the prepared state are affected by state preparation imperfections, according to Eq. 7.6. Hence, the measured readout fidelities for **CHIP−β−FC** are

$$\begin{aligned} P(I<T_I \,|\, \boldsymbol{\rho}_0^g) &= p_g F_g + (1-p_g)(1-F_e) = 83.9(0.9)\%, \\ P(I>T_I \,|\, \boldsymbol{\rho}_0^e) &= p_e F_e + (1-p_e)(1-F_g) = 68.3(1.0)\%, \end{aligned} \tag{7.10}$$

where $\boldsymbol{\rho}_0^i$ is defined in Eq. 7.5. From these relations, we extract intrinsic readout fidelities of

$$F_g = 84.8(1.8)\%, \quad F_e = 69.1(1.2)\%. \tag{7.11}$$

7.6 Qubit Frequency Tuning: AC-Stark Shift

The AC-Stark shift provides a controllable mechanism for tuning the qubit transition frequency, at the expense of an enhanced dephasing rate. It arises from an oscillating microwave field that is detuned from a given transition. Because the drive is off-resonant, it does not induce real excitations, but it does weakly mix (dress) the energy levels through virtual absorption and emission processes, resulting in a shift of the transition frequency. In the following, we derive the AC-Stark shift for a heavy-fluxonium qubit and examine its impact on the qubit's decoherence rate.

7.6.1 Amplitude-dependent frequency shift

The lowest $|g\rangle \rightarrow |e\rangle$ transition of our fluxonium (**CHIP−β−FC**) lies at $\omega_q/2\pi = 2.35$ MHz, while the mechanical mode sits at $\omega_m/2\pi = 4.4$ MHz (see the next chapter Chap. 9). We can overcome this detuning by applying a continuous drive at $\omega_d/2\pi = 3.463$ GHz—exactly midway between the ω_{ef} and ω_{gh} lines. This drive induces an AC–Stark shift that pulls the dressed qubit into resonance with the mechanics, at the cost of an increased dephasing rate (see the following Sec. 7.6.2). To describe this driven system we therefore add a time-dependent charge-drive term to the fluxonium Hamiltonian:

$$\boldsymbol{H} = \boldsymbol{H}_{\mathrm{fluxonium}} + \hbar\Omega_d \cos(\omega_d t)\boldsymbol{n}, \tag{7.12}$$

where $\boldsymbol{H}_{\mathrm{fluxonium}}$ is the undriven fluxonium Hamiltonian and Ω_d is the drive amplitude. We restrict our attention to transitions near ω_d and retain only the lowest four fluxonium levels, $\{|g\rangle, |h\rangle, |e\rangle, |f\rangle\}$. Projecting the driving term of the Hamiltoninan into this four-level subspace, and

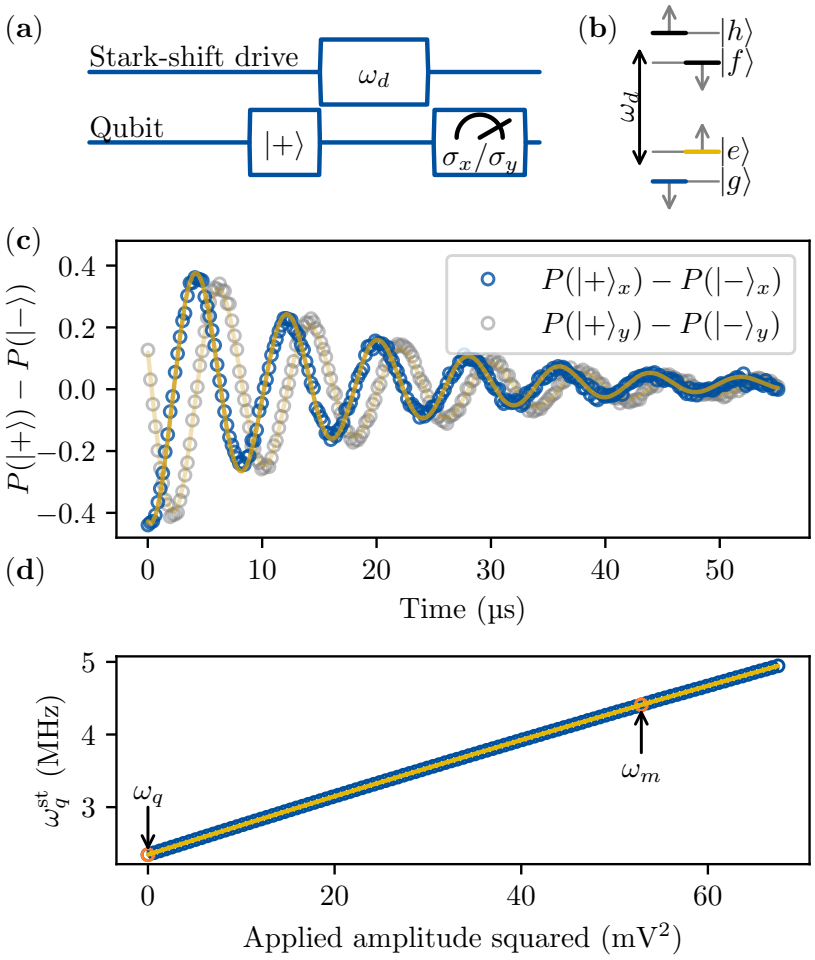

Fig. 7.20: Qubit AC Stark shift. (a) Ramsey pulse sequence with a continuous drive applied during the idle interval. **(b)** Energy-level diagram of the four lowest fluxonium eigenstates. The drive frequency ω_d is chosen midway between the $|e\rangle \to |f\rangle$ and $|g\rangle \to |h\rangle$ transitions, inducing level repulsion and opposite AC–Stark shifts on the $\{|g\rangle, |h\rangle\}$ and $\{|e\rangle, |f\rangle\}$ manifolds. **(c)** Difference between the measurement outcome probabilities for $|+\rangle$ and $|-\rangle$ along the σ_x (blue) and σ_y (gray) axes, measured at the bare qubit frequency. From these data, we reconstruct the complex transverse Bloch vector and fit it to exponentially decaying oscillations, $\propto \exp\left(-\Gamma_2^* t + i\,\omega_q^{\mathrm{st}} t\right)$. The real and imaginary components of the fit are overlaid in yellow. **(d)** Qubit transition frequency extracted from **(c)** type of experiment as a function of the applied drive amplitude squared (proportional to Ω_d^2). The yellow line shows a fit to Eq. 7.18. The bare qubit frequency ω_q and membrane frequency ω_m are indicated by orange circles.

considering only non-zero matrix elements at $\varphi_{\mathrm{ext}} = \pi$ yields

$$H_{\mathrm{Stark}} = \hbar\Omega_d \cos\omega_d t \begin{pmatrix} 0 & n_{\mathrm{gh}} & n_{\mathrm{ge}} & 0 \\ n_{\mathrm{gh}}^* & 0 & 0 & 0 \\ n_{\mathrm{ge}}^* & 0 & 0 & n_{\mathrm{ef}} \\ 0 & 0 & n_{\mathrm{ef}}^* & 0 \end{pmatrix} \tag{7.13}$$

where $n_{ij} = \langle i|\boldsymbol{n}|j\rangle$ is the charge matrix element. We then transform into the frame rotating at ω_d via the unitary transformation

$$U = \mathrm{diag}\left(1 \quad e^{i\omega_d t} \quad e^{i\omega_{\mathrm{ge}} t} \quad e^{i(\omega_d+\omega_{\mathrm{ge}})t}\right). \tag{7.14}$$

In the rotated frame, dropping all fast rotating terms (beyond or close to ω_d), the Hamiltonian writes

$$\widetilde{H}_{\mathrm{Stark}} = \hbar \begin{pmatrix} 0 & \Omega_d n_{\mathrm{gh}}/2 & 0 & 0 \\ \Omega_d n_{\mathrm{gh}}^*/2 & \Delta_{\mathrm{gh}} & 0 & 0 \\ 0 & 0 & 0 & \Omega_d n_{\mathrm{ef}}/2 \\ 0 & 0 & \Omega_d n_{\mathrm{ef}}^*/2 & \Delta_{\mathrm{ef}} \end{pmatrix}, \tag{7.15}$$

with $\Delta_{\mathrm{gh}} = \omega_{\mathrm{gh}} - \omega_d$ and $\Delta_{\mathrm{ef}} = \omega_{\mathrm{ef}} - \omega_d$. This Hamiltonian is block-diagonal, such that it can be diagonalized separately in the two subspaces $\{|g\rangle, |h\rangle\}$ and $\{|e\rangle, |f\rangle\}$. The two lowest energy dressed eigenstates

are

$$\begin{aligned} |\tilde{g}\rangle &= \cos\frac{\Theta_{\mathrm{gh}}}{2}|g\rangle - \sin\frac{\Theta_{\mathrm{gh}}}{2}|h\rangle\ , \\ |\tilde{e}\rangle &= \cos\frac{\Theta_{\mathrm{ef}}}{2}|e\rangle - \sin\frac{\Theta_{\mathrm{ef}}}{2}|f\rangle\ , \end{aligned} \tag{7.16}$$

with mixing angles defined by $\tan\Theta_{\mathrm{gh}} = \Omega_d n_{\mathrm{gh}}/|\Delta_{\mathrm{gh}}|$ and $\tan\Theta_{\mathrm{ef}} = \Omega_d n_{\mathrm{ef}}/|\Delta_{\mathrm{ef}}|$. The eigenerergies are given by

$$\begin{aligned} \tilde{E}_g/\hbar &= \frac{1}{2}\,\mathrm{sign}(\Delta_{\mathrm{gh}})\left(|\Delta_{\mathrm{gh}}| - \sqrt{\Delta_{\mathrm{gh}}^2 + \Omega_d^2 n_{\mathrm{gh}}^2}\right), \\ \tilde{E}_e/\hbar &= \frac{1}{2}\,\mathrm{sign}(\Delta_{\mathrm{ef}})\left(|\Delta_{\mathrm{ef}}| - \sqrt{\Delta_{\mathrm{ef}}^2 + \Omega_d^2 n_{\mathrm{ef}}^2}\right). \end{aligned} \tag{7.17}$$

For the circuit at half-flux, $\varphi_{\mathrm{ext}} = \pi$, one finds $n_{\mathrm{gh}} \simeq n_{\mathrm{ef}} \equiv n_{\mathrm{up}}$, furthermore, we have chosen a drive frequency precisely midway between ω_{gh} and ω_{ef} yielding $\Delta_{\mathrm{gh}} = -\Delta_{\mathrm{ef}} \equiv \Delta_{\mathrm{st}}$, and $\Theta_{\mathrm{gh}} \simeq \Theta_{\mathrm{ef}} \equiv \Theta$. In this approximation, the AC-Stark shift is given by

$$\begin{aligned} \delta\omega_q &\equiv \tilde{E}_e - \tilde{E}_g \\ &= \sqrt{\Delta_{\mathrm{st}}^2 + \Omega_d^2 n_{\mathrm{up}}^2} - \Delta_{\mathrm{st}} \approx \frac{\Omega_d^2 n_{\mathrm{up}}^2}{2\Delta_{\mathrm{st}}}\left(1 - \frac{1}{4}\frac{\Omega_d^2 n_{\mathrm{up}}^2}{\Delta_{\mathrm{st}}^2}\right). \end{aligned} \tag{7.18}$$

In the last line, we expanded the expression to fourth order in $\Omega_d/\Delta_{\mathrm{st}}$.

Since the absolute attenuation between source and sample is unknown, we calibrate the AC–Stark shift by Ramsey interferometry Fig. 7.20a. Beginning at the fluxonium sweet-spot frequency $\omega_q/2\pi \approx 2.2$ MHz, we apply a continuous tone at $\omega_d/2\pi = 3.463$ GHz with varying source power P. A $\pi/2$–Ramsey sequence maps the drive-induced phase into σ_z, from which we extract the dressed qubit frequency. The sequence is shown in Fig. 7.20a. Plotting frequencies versus P reveals a primarily linear Stark pull plus a small quadratic correction (Fig. 7.20c). Because Eq. 7.18 fixes the ratio of the linear and quadratic terms, only one free parameter—the scaling factor between P and the on-chip Rabi rate Ω_d—is required to fit all the data. Over the accessible power range the qubit frequency tunes continuously from 2.2 MHz up to ~ 5 MHz (more than one octave), enabling resonance with the 4.4 MHz mechanical mode. The resulting single-parameter fit (solid line, Fig. 7.20c) quantitatively reproduces both slope and curvature of the Stark shift.

7.6.2 Induced decoherence

Unfortunately, the AC–Stark shift comes at the cost of increased qubit decoherence, which we measured for **CHIP−$\boldsymbol{\beta}$−FC**. We recall that, the energy-relaxation rate, $\Gamma_1 = 1/T_1$, is extracted from the exponential decay of the excited-state population; The free-induction-decay rate, $\Gamma_2^* = 1/T_2^*$, is obtained from a Ramsey experiment (Fig. 7.20a). The resulting rates are plotted in Fig. 7.21. At the bare qubit frequency we observe $\Gamma_1/\Gamma_2^* \simeq 1.12 < 2$. As the AC–Stark drive increases, the pure dephasing rate $\Gamma_\phi = \Gamma_2^* - \Gamma_1/2$ rises sharply, causing T_2^* to degrade much faster than T_1. Both rates exhibit a pronounced peak when ω_q^{st} approaches ω_m, due to the qubit transitions induced by the thermally excited mechanical mode.

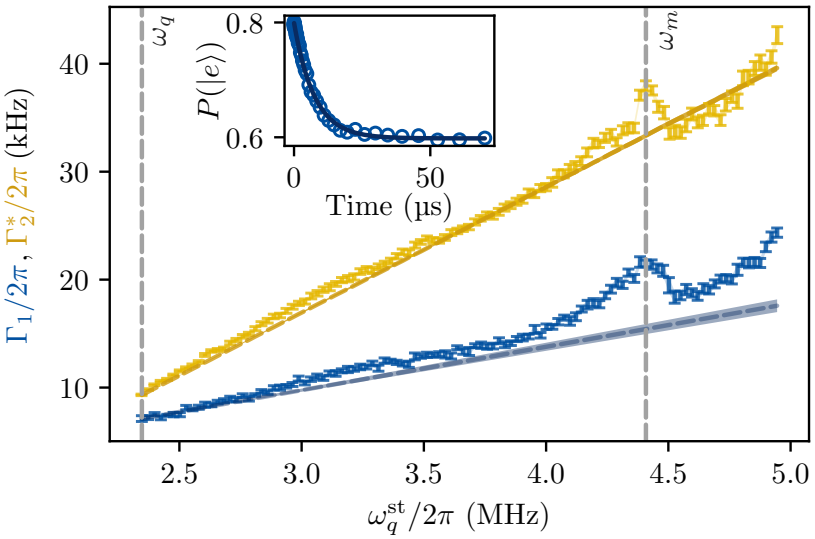

Fig. 7.21: Qubit decoherence and relaxation rates vs. AC–Stark–shifted qubit frequency. Energy-relaxation rate $\Gamma_1 = 1/T_1$ (blue markers) and free-induction-decay rate $\Gamma_2^* = 1/T_2^*$ (yellow markers) are plotted versus the AC-Stark–shifted qubit frequency ω_q^{st}. Γ_1 is extracted from the exponential relaxation of the qubit population toward its thermal equilibrium. Γ_2^* is obtained from Ramsey-fringe decay. The dashed blue line shows $\Gamma_1(\omega_q^{\mathrm{st}})$ from Eq. (7.19), with Γ_1^{ge} as the sole fit parameter. The dashed yellow line represents $\Gamma_2^* = \Gamma_1/2 + \Gamma_\phi$, using $\Gamma_\phi(\omega_q^{\mathrm{st}})$ from Eq. (7.20) and fitting $\delta\Omega_d/\Omega_d$. Vertical black dashed lines indicate the bare qubit frequency ω_q and the mechanical mode frequency ω_m. Inset: Qubit relaxation at ω_m following preparation in the state $|e\rangle$. Experimental data (blue) and corresponding exponential fit (dark blue) are shown.

We attribute the approximately linear rise of Γ_1 with detuned qubit frequency ω_q^{st} to the growing admixture of the short-lived $\{|f\rangle, |h\rangle\}$ states into the lower dressed states—an effect analogous to dressed-state spontaneous emission in optical tweezers (see Eq. 7.16). If the bare-manifold decay rates are Γ_1^{ge} and Γ_1^{fh}, then the dressed-state relaxation rate is

$$\Gamma_1(\delta\omega_q) = \cos^2\!\left(\tfrac{\Theta}{2}\right)\Gamma_1^{ge} + \sin^2\!\left(\tfrac{\Theta}{2}\right)\Gamma_1^{fh}\,.$$

For small mixing angle $\Theta \approx \delta\omega_q/\Delta_{\mathrm{st}}$, expanding to first order gives

$$\Gamma_1(\delta\omega_q) \approx \Gamma_1^{ge} + \frac{\Gamma_1^{fh}}{2\,\Delta_{\mathrm{st}}}\,\delta\omega_q\,. \tag{7.19}$$

The dashed blue line in Fig. 7.21 plots Eq. 7.19 using the independently measured value $\Gamma_1^{fh}/2\pi = 139$ kHz, leaving only Γ_1^{ge} as a fit parameter.

We attribute the enhanced dephasing rate $\Gamma_2^*(\omega_q^{\mathrm{st}})$ to amplitude fluctuations of the Stark drive Ω_d, yielding

$$\Gamma_\phi = \frac{d(\delta\omega_q)}{d\Omega_d}\,\delta\Omega_d \simeq 2\,\delta\omega_q\,\frac{\delta\Omega_d}{\Omega_d}\,, \tag{7.20}$$

where the approximation holds for $\Omega_d \ll \Delta_{\mathrm{st}}$. Assuming classical amplitude noise (i.e. constant $\delta\Omega_d/\Omega_d$), fitting this expression to the data (yellow dashed line in Fig. 7.21) gives $\delta\Omega_d/\Omega_d \approx 0.6\%$, in good qualitative agreement with the $\sim 0.2\%$ specification of our *Anapico* source.

Basic membrane-qubit interaction

8

In this chapter, we step into the main theme of this thesis: the interaction between qubit and membrane. As was shown previously, the qubit-membrane interaction is governed by the Jaynes-Cummings Hamiltonian Sec. 5.2.2:

$$\boldsymbol{H}_{\mathrm{JC}} = \hbar\omega_r\, \boldsymbol{a}^\dagger \boldsymbol{a} + \frac{\hbar\omega_q}{2}\, \boldsymbol{\sigma}_z + i\frac{\hbar\Omega}{2}\left(\boldsymbol{a}^\dagger \boldsymbol{\sigma}_- - \boldsymbol{a}\, \boldsymbol{\sigma}_+\right)$$

Outline. This chapter starts by examining the interaction in the dispersive regime, when the qubit and the membrane are out-of-resonance. Using Ramsey experiment we perform the membrane spectroscopy, extracting the membrane resonance ω_m as a function of bias V_b, and measure the mechanical relaxation time T_1^m to characterize damping.

We then turn to the resonant interaction with the membrane in a large coherent state, using the AC–Stark effect to tune ω_q into resonance with ω_m, allowing coherent Rabi oscillations. We introduce the semi-classical limit, where membrane operator $\boldsymbol{a}$ is replaced by a mean field α and show that the membrane can be seen as a classical drive on the qubit with rotation rate $\Omega|\alpha|$. We describe the experimental sequence used for every interaction experiment, including the spectrum–analyzer experiments presented in next chapter.

8.1 Dispersive interaction regime

In absence of AC-Stark effect, the qubit and the membrane modes are out-of-resonance: the large detuning $|\Delta = \omega_q - \omega_m| \gg \Omega$ suppresses direct energy exchange between the two systems. We have previosly discussed (Sec. 2.5.2) that under dispersive regime, the qubit frequency shift is proportional to the membrane's phonon population $2\chi \boldsymbol{a}^\dagger \boldsymbol{a} \boldsymbol{\sigma}_z$, where $\chi = \Omega^2/4\Delta$.

$$\boldsymbol{H}_{\mathrm{disp}} = \underbrace{\hbar\omega_m\, \boldsymbol{a}^\dagger \boldsymbol{a}}_{\text{bare resonator}} + \underbrace{\frac{\hbar}{2}\left(\omega_q + 2\chi \boldsymbol{a}^\dagger \boldsymbol{a}\right) \boldsymbol{\sigma}_z}_{\text{Bare qubit + frequency pull}}, \tag{8.1}$$

In this section, we will use this effect to characterize the qubit-membrane interaction. We will first measure the membrane frequency as a function of the applied DC bias voltage V_b to extract the membrane frequency ω_m and the electrostatic spring softening effect. Then, we will measure the membrane relaxation time T_1^m to characterize the membrane damping.

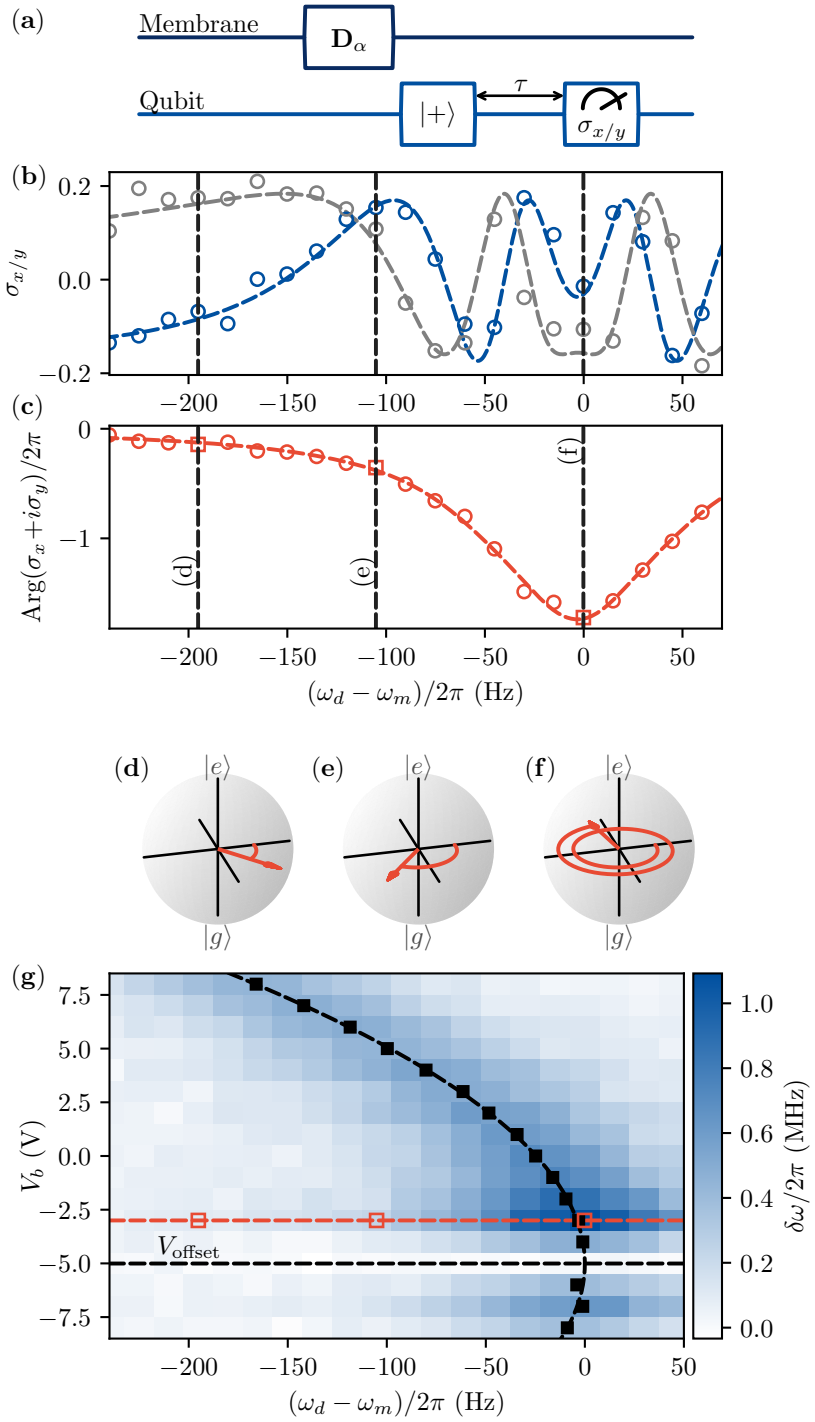

Fig. 8.1: Membrane spectroscopy using a qubit sensing. (a) Pulse sequence: membrane drive D_α, qubit prepared in $|+\rangle$, interaction for a time τ, then measurement of σ_x or σ_y (see Sec. 8.1.1). See Sec. 8.1.1 for details. **(b)** Qubit signals $\langle\sigma_x\rangle$ (blue) and $\langle\sigma_y\rangle$ (gray) as a function of drive detuning $\omega_d - \omega_m$ at $V_b = -3$ V. Dashed lines show the real and imaginary parts of a Lorentzian-phase fit to $\langle\sigma_x\rangle + i\,\langle\sigma_y\rangle$. **(c)** Unwrapped phase $\arg(\langle\sigma_x\rangle + i\,\langle\sigma_y\rangle)$ with Lorentzian fit (dashed). **(d–f)** Bloch-sphere trajectories for the detunings marked by dashed lines in (b). **(g)** Colormap of the qubit frequency shift $\delta\omega_q/2\pi$ versus V_b and detuning. Black squares indicate resonance peaks; the black dashed curve is its quadratic fit as predicted by Eq. 8.4. The drive amplitude is adjusted at each V_b to keep the peak response constant, except at $V_b = V_{\mathrm{offset}}$ (the black dashed horizontal line) where the coupling vanishes.

8.1.1 Experimental protocol

To measure the qubit frequency shift, we perform Ramsey spectroscopy (Sec. 2.4.6). The simplest protocol used in the membrane–qubit dispersive regime is illustrated in Fig. 8.2. We drive the membrane through a line connected to the RF bias-tee port and routed to the sample membrane-bias port, thereby preparing the membrane in a coherent state $|\alpha\rangle$ with large amplitude $|\alpha| \gg 1$ (at resonance, $|\alpha|^2 \approx 10^7$). This large occupation ensures that the qubit–membrane interaction produces a measurable signal. The qubit is initialized in $|+\rangle = (|g\rangle + |e\rangle)/\sqrt{2}$ by first preparing $|e\rangle$ and then applying a $\pi/2$ pulse (Fig. 8.3(a)). It then evolves for a duration $\tau = 10\ \mu$s while the membrane remains in the large coherent state, accumulating an equatorial phase $\delta\omega\,\tau = 2\chi\,\overline{n}\,\tau$ that is proportional to the mean phonon number $\overline{n}$ (Fig. 8.3(b)).

The experiment alternates between measuring the σ_x and σ_y components of the superposition across repeated realizations by applying an additional $\pi/2$ pulse before the final σ_z readout (Fig. 8.3(c)). We then extract the phase of $\langle\sigma_x\rangle + i\,\langle\sigma_y\rangle$. For large membrane coherent states, the qubit may undergo multiple rotations during the interval τ. To enable phase unwrapping, we sweep the control parameter with sufficiently fine resolution to prevent large phase discontinuities between adjacent points.

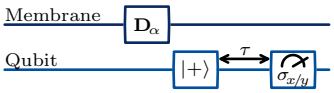

Fig. 8.2: Experimental protocol of measurements in qubit-membrane dispersive regime: membrane drive D_α, qubit prepared in $|+\rangle$, interaction for a time τ, then measurement of σ_x or σ_y (see Sec. 8.1.1). See the main text for details.

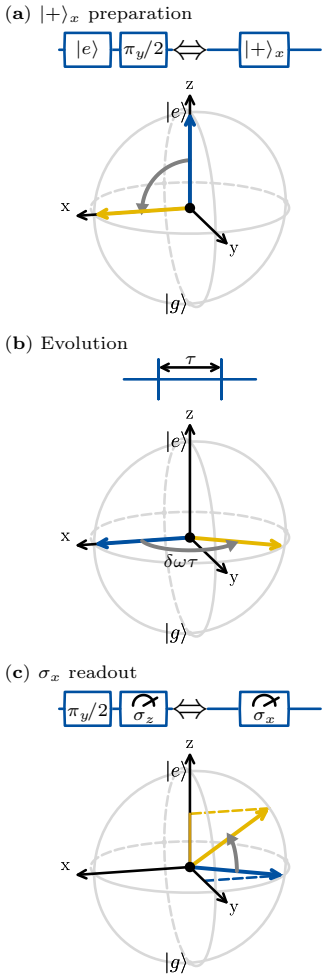

Fig. 8.3: Illustrations of difference parts of the dispersive regime experiment protocol See the main text for details.

8.1.2 Membrane Spectroscopy

The first step in the characterization of the qubit-membrane interaction is to determine the membrane frequency as a function of the applied DC bias voltage V_b. We sweep the frequency of the initial membrane drive and measure the qubit frequency shift for different bias voltages V_b. The data is shown in Fig. 8.1.

When a harmonic oscillator is driven near its resonance frequency, the cavity population follows a Lorentzian profile (Sec. 2.5.3). However, because the membrane motion is probed indirectly via the qubit phase shift, the measured quadrature does not directly follow a Lorentzian lineshape (Fig. 8.1(b)), as the phonon-induced phase shift causes the final qubit state to rotate by more than 2π. After phase unwrapping the data (Fig. 8.1(c)), a Lorentzian profile is recovered and fitted (black square markers), allowing us to extract the membrane frequency for each bias voltage V_b.

Since dispersive shift scales quadratically with bias voltage, $\chi \propto \Omega^2 \propto V_b^2$, this measurement is performed with adaptative power to ensure that the qubit exibits a detectable accumulated phase (a few turns in the Bloch sphere, i.e. $\delta\omega\tau \lesssim 6\pi$). A sweep in V_b reveals a quadratic decrease in the mechanical frequency with $V_b - V_{\text{offset}}$, explained by the electrostatic spring-softening effect that effectively softens the mechanical spring and lowers the resonant frequency.

8.1.3 Electrostatic Spring Softening

We derive the electrostatic spring softening effect [UWK09] from the qubit-membrane coupling Hamiltonian (Eq. 5.16) derived previously. The terms that are independent of the qubit state are:

$$H_{\text{membrane}} = \hbar\omega_m \left(\boldsymbol{p}^2 + \boldsymbol{x}^2\right) + \frac{1}{2}\sum_{i,j=2}^{3} \widetilde{Q}_i \, \boldsymbol{M}_{ij}^{-1} \, \widetilde{Q}_j. \tag{8.2}$$

We expand H_{membrane} to second order in the dimensionless displacement $\boldsymbol{x}$, express $\boldsymbol{x} = (\boldsymbol{a} + \boldsymbol{a}^\dagger)/2$ and $\boldsymbol{p} = i(\boldsymbol{a}^\dagger - \boldsymbol{a})/2$, and retain only energy-conserving terms. The membrane Hamiltonian then becomes

$$\begin{aligned} H_{\text{membrane}} &= \hbar\,\omega_m \left(\boldsymbol{a}^\dagger\,\boldsymbol{a} + \tfrac{1}{2}\right) \;+\; \boldsymbol{H}_{ms} \\ &= \hbar\left(\omega_m + \delta\omega_m\right)\left(\boldsymbol{a}^\dagger\,\boldsymbol{a} + \tfrac{1}{2}\right), \end{aligned} \tag{8.3}$$

where the frequency shift is

$$\delta\omega_m = -\zeta(V_b - V_{\text{offset}})^2 \tag{8.4}$$

with

$$\zeta = 2\,C_m\,(C_g + 2C_q)\,\frac{E_C}{\hbar e^2}\,\frac{X_{\text{zpf}}^2}{d^2}\,\beta^2. \tag{8.5}$$

Here we have used the parallel-plate model $C_m^{\pm}(X) = C_m/(1 \pm X/d)$ to calculate the derivatives $dC_m^{\pm}/dX$ and $d^2C_m^{\pm}/dX^2$. Since Eq. 8.4 is always negative, the electrostatic force acts as a "negative spring"[1], lowering the membrane's resonance frequency.

8.1.4 Offset voltage discussion

In our measurements, we consistently observed a residual static charge on the membrane, which offsets the voltage bias required to cancel the electrostatic force by an amount V_{offset} (see Fig. 8.1a). This offset arises from trapped charges Q_2 and Q_3 on the islands between the capacitors illustrated in Fig. 8.5 and given by Eq. 5.20. Although the qubit remains operational for a DC-bias voltage up to $|V_b| \lesssim 9\,\text{V}$ [2], the effective coupling strength depends on the difference $V_b - V_{\text{offset}}$. Therefore, the presence of this offset can be used to increase the coupling strength.

We found that V_{offset} slowly drifts toward the applied bias over a timescale of weeks. To restore the maximal coupling range, we deliberately "reset" the offset by applying a large negative bias of $V_b = -20\,\text{V}$ to the membrane for several days. While this high voltage temporarily disables measurements, it reliably shifts V_{offset} closer to $-9\,\text{V}$. After this procedure, we verify the new offset value using membrane spectroscopy and ensure that it lies within the desired measurement range.

8.1.5 Membrane Ringdown

The second step in the characterization of the qubit-membrane interaction is to measure the membrane relaxation time T_1^m. It characterizes the ability

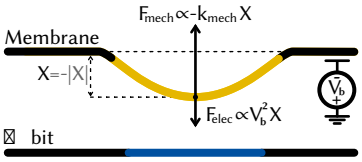

Fig. 8.4: Forces acting on a voltage-biased membrane capacitively coupled to a qubit. Yellow and blue regions mark metallized parts of the membrane and qubit that form the coupling capacitor. The mechanical restoring force $F_{\text{mech}} = -k_{\text{mech}}X$ (upward arrow) opposes the electrostatic pulling force $F_{\text{elec}} \propto V_b^2 X$ (downward arrow), both projected along the same displacement coordinate X.

1: The more classical derivation of the electrostatic spring softening effect can be done by examing the forces acting on the membrane. Since a simplified potential energy acting on the membrane mode is $U(X) = \frac{1}{2}k_m(X-d)^2 - \frac{1}{2}C_m(X)\,V_b^2$, where k_m is the spring constant of the membrane in absence of bias and $C_m(X) = C_m/(1+X/d)$ is the capacitance of the membrane, the force on the membrane is

$$F(X) = -\frac{\partial U}{\partial X} = -k_m(X-d) + \frac{V_b^2}{2}\frac{dC_m(X)}{dX}.$$

It consist of the mechanical restoring force and the electrostatic force illustrated in Fig. 8.4. It shifts the equilibrium position of the membrane from $X = d$ to X_0, However, for our parameters, X_0 can be safely identified to d. The effective spring constant is thus reduced by the electrostatic force:

$$k_{\text{eff}} = \left.\frac{\partial^2 U}{\partial X^2}\right|_{X_0\simeq d} = k_m - \frac{V_b^2}{2}\left.\frac{d^2C_m(X)}{dX^2}\right|_{X_0\simeq d}.$$

Since the $\omega_m(V_b) = \sqrt{\frac{k_{\text{eff}}}{m}}$, it leads to the decrease of the mechanical frequency, with shift that, to the first order in X/d, is proportional to $-V_b^2$. Comparing with the derivation in main text, this derivation is based on the simplified coupling version suppossing a single capacitance $C(x)$ biased at voltage V_b, which is not the case in our system. However, it provides a crucial intuition that the electrostatic force acts as a negative spring.

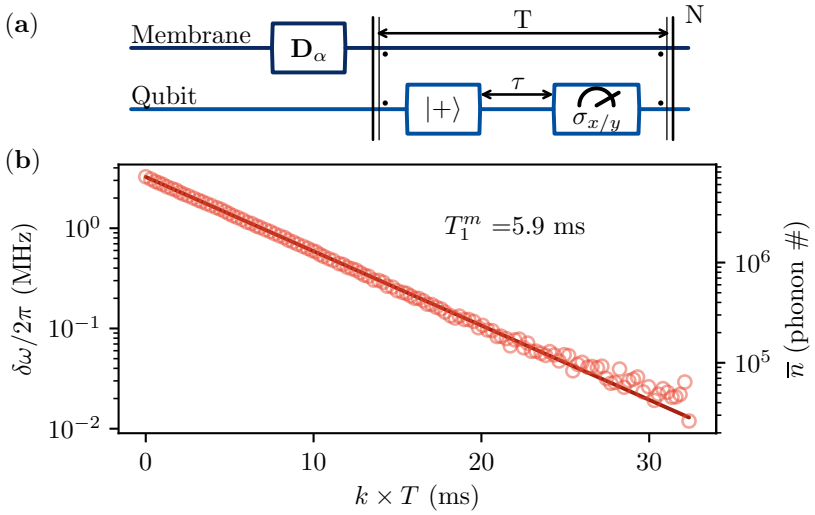

Fig. 8.6: Mechanical lifetime. (a) Pulse sequence: after resonant membrane displacement, a series of $N = 121$ Ramsey measurements are performed, sampling the instantaneous qubit frequency at discrete times $k \times T$. **(b)** Averaged qubit frequency shift $\delta\omega$ plotted versus delay $k \times T$ (for $k = 1, \ldots, N$), where data at each delay are ensemble-averaged. An exponential fit yields a mechanical lifetime $T_1^m = 5.9$ ms ($Q = 1.62 \times 10^5$). The right-hand axis translates $\delta\omega$ into the equivalent mechanical occupation in units of phonon number.

of the membrane to store energy.

We extract T_1^m by first preparing the membrane in a large coherent state $|\alpha|^2 \sim 10^7$ and then measuring the exponential decay of its phonon population with qubit at a time interval T. This decay manifests as a time variation of the qubit frequency, tracked through repeated Ramsey measurements at discrete times $k \times T$ within a single membrane decay. The unwrapped phase of the qubit frequency exhibits an exponential decay (see Fig. 8.6(b)), and an exponential fit yields $T_1^m = 5.9$ ms. The corresponding quality factor, $Q_m = \omega_m T_1^m = 1.62 \cdot 10^5$, is consistent with similar studies on silicon nitride membranes [YPR12]. It is closed to the quality factor obtained at room temperature without any metalization of the membrane ($Q_m^{\mathrm{RT}} = 6 \cdot 10^4$), thus we observed that additional losses due to the membrane metalization are over-compensated by the cryogenic boost of the quality factor (due to lowing temperature)[Pat25].

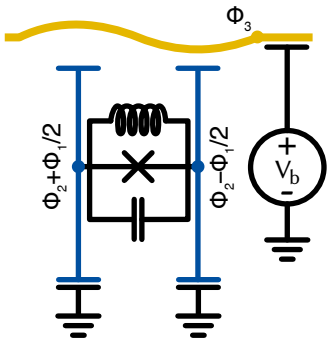

Fig. 8.5: Charge islands in the hybrid circuit. Colours indicate the two islands: the qubit islands (blue) carrying charge Q_2 at flux nodes $\Phi_2 \pm \Phi_1/2$, and the membrane electrode island (yellow) carrying charge Q_3 at flux node Φ_3, biased by voltage V_b.

8.2 Resonant Qubit–Membrane Dynamics

The tuning of the qubit frequency, thanks to the AC-Stark effect (Sec. 7.6), allows us to enter the resonant interaction regime between the qubit and the membrane. In this regime, system is governed by the resonant Jaynes-Cummings Hamiltonian (Eq. 2.67):

$$H_{\mathrm{JC}} = \underbrace{\hbar\omega_r\, \boldsymbol{a}^\dagger \boldsymbol{a} + \frac{\hbar\omega_q}{2}\,\sigma_z}_{\text{Free evolution}} + \underbrace{i\frac{\hbar\Omega}{2}\left(\boldsymbol{a}^\dagger\sigma_- - \boldsymbol{a}\,\sigma_+\right)}_{\text{Coupling}}.$$

In this section, we discuss the experimental protocol used to enter the resonant interaction regime. Then, we present the theoretical description of the resonant interaction in the semi-classical limit, i.e. when the membrabe mode operator $\boldsymbol{a}$ is approximated by a classical coherent state α. This provides the basis for understanding the spectrum analyzer experiment presented in the next chapter. And finally, we present the Rabi spectroscopy experiment that shows the coherent dynamics between the qubit and the membrane.

2: We attribute this behavior to avalanche ionization in the substrate [SLL12], likely enhanced by the tension across the 4 μm isolated gap. The exact mechanism is not fully investigated: test samples reproduced the gap dependence of the maximum voltage but not the same maximum tension (30 V), possibly due to substrate impurities or differences in cryostat setups.

[Pat25]: Patange (2025), 'RF Transduction in a DC-Biased Superconducting Electromechanical System', Thesis

8.2.1 Experimental protocol

With the AC-stark shift, we tune the qubit frequency $\omega_q^{\text{st}} = \omega_m + \Delta$ close to the membrane one, and enter the resonant interaction regime. However, since the coherence time is much longer at the qubit bare frequency, we get better fidelity preparing and reading out the qubit while it sits at its bare frequency. Thus, throughout resonant interaction experiments presented in this thesis, we used the following experimental protocol (Fig. 8.7):

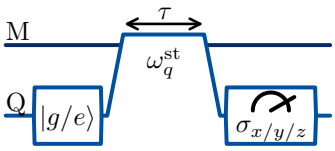

Fig. 8.7: Resonant interaction pulse sequence. M and Q are the membrane and qubit lines, respectively.

1. Prepare the qubit in the $|g\rangle$ or $|e\rangle$ state at the bare qubit frequency ω_q.
2. Tune the qubit frequency to $\omega_q^{\text{st}} = \omega_m + \Delta$ close to the membrane one during a time τ.
3. Measure $\sigma_{x/y/z}$ projection of the qubit state at the bare qubit frequency.

The evolution of the membrane–qubit state during the interaction is governed by the evolution operator

$$\boldsymbol{U}_I = e^{-i\boldsymbol{H}_{\text{JC}}\tau/\hbar}.$$

Outside this interval, the membrane exhibits free evolution under the harmonic oscillator Hamiltonian $\boldsymbol{H} = \hbar\omega_m\, \boldsymbol{a}^\dagger \boldsymbol{a}$.

8.2.2 Resonant interaction in the semi-classical limit

When the membrane is initialized in a large enough coherent state α, we can displace the mode by α and represent $\boldsymbol{a} \longrightarrow \alpha + \delta\boldsymbol{a}$, where $\delta\boldsymbol{a}$ is a small fluctuation around the coherent state α. Since $|\alpha| \gg 1$, we neglect the fluctuation operator $\delta\boldsymbol{a}$, moving to the semi-classical approximation. This yields the coupling term:

$$\begin{aligned} H_{\text{CC}} &= i\frac{\hbar\Omega}{2}\left(\alpha^*\boldsymbol{\sigma}_- - \alpha\,\boldsymbol{\sigma}_+\right) \\ &= \frac{\hbar\Omega}{2}\left(\text{Im}(\alpha)\boldsymbol{\sigma}_x + \text{Re}(\alpha)\boldsymbol{\sigma}_y\right) \end{aligned} \tag{8.6}$$

where now α is a complex number that represents the classical-like trajectory of the membrane. Additionally to the free evolution under the Hamiltonian $\boldsymbol{H}_{JC}^0$, the evolution during interaction time τ is covered by the unitary evolution operator $\boldsymbol{U}_I$ given by:

$$\boldsymbol{U}_I = \exp\left(-\frac{i}{\hbar}\boldsymbol{H}_c\tau\right) = \exp\left(-i\frac{\Omega}{2}|\alpha|\vec{n}\cdot\vec{\boldsymbol{\sigma}}\tau\right) \tag{8.7}$$

where

$$\vec{n} = \frac{1}{|\alpha|}(\text{Im}(\alpha), \text{Re}(\alpha), 0) \quad \text{and} \quad \vec{\boldsymbol{\sigma}} = (\boldsymbol{\sigma}_x, \boldsymbol{\sigma}_y, \boldsymbol{\sigma}_z).$$

Comparing with the standard $SU(2)$ rotation matrix $\exp(-i\frac{\theta}{2}\vec{n}\cdot\vec{\boldsymbol{\sigma}})$ (Eq. 2.57) shows that the qubit undergoes a rotation by $\theta = \Omega|\alpha|\tau$ about the axis $\vec{n}$ in the Bloch-sphere equatorial plane which complex direction is $i\alpha^*$.

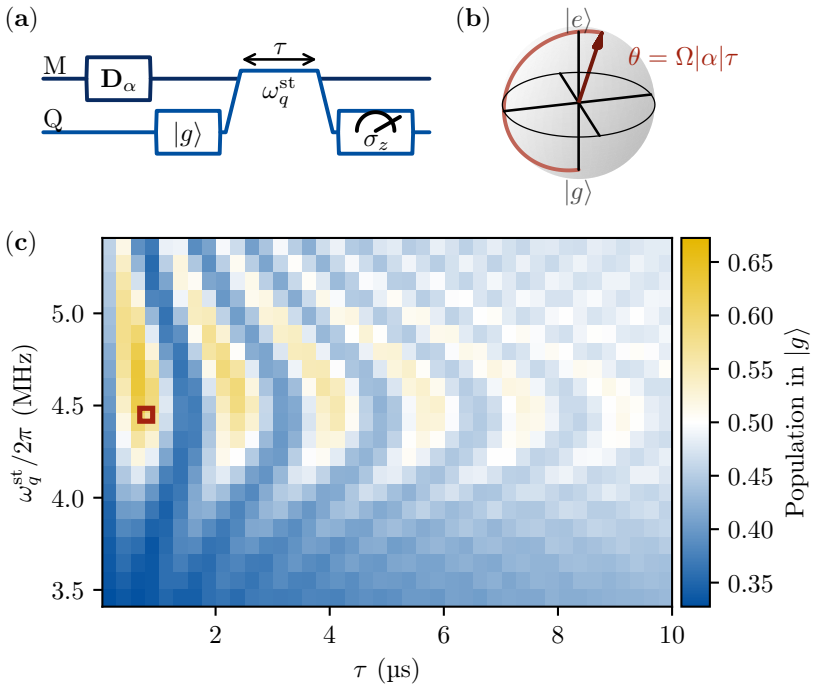

Fig. 8.8: Coherent membrane–qubit Rabi exchange. (a) Pulse sequence: the membrane is prepared in the coherent state $|\alpha\rangle$ with $|\alpha|^2 \approx 1.6 \times 10^5$; the qubit is Stark-shifted to ω_q^{st} for an interaction time τ, then read out in the σ_z basis. **(b)** Bloch-sphere illustration of the qubit rotation after a resonant exchange with the membrane. **(c)** Ground-state probability P_g as a function of τ and Stark-shifted frequency $\omega_q^{\mathrm{st}}/2\pi$. The red square marks the operating point shown in (b).

8.2.3 Rabi spectroscopy

In order to validate the physics described above, we perform a Rabi spectroscopy where the detuning Δ is varied by tuning the qubit frequency. The membrane is prepared in a large coherent state $|\alpha\rangle$ with amplitude[1] $|\alpha|^2 \simeq 1.6 \cdot 10^5$ to ensure a coherent dynamics within the qubit coherence time, $\Omega|\alpha| \gg 1/T_1$. We prepare the qubit in the $|g\rangle$ state at the bare qubit frequency ω_q, then we tune the qubit frequency to $\omega_q^{\mathrm{st}} = \omega_m + \Delta$ and after coherent interaction of duration τ, we bring back the qubit to its bare frequency and perform state detection. This sequence is illustrated in Fig. 8.8(a). By varying the detuning Δ and τ we obtain a 2D map with characteristic Rabi spectroscopy pattern with a minimum frequency when the $\omega_q^{\mathrm{st}} \simeq \omega_m$, as shown in Fig. 8.8(b).

Since the membrane is in large coherent state, these oscillations are identical to those induced by a purely classical external field. Given our current coupling strength, this classical Rabi oscillations is the only regime accessible. Indeed, the vacuum Rabi rate $\Omega/2\pi \approx 1.5$ kHz is much smaller than the qubit decay rate ($\Gamma_1/2\pi \approx 20$ kHz), placing the system outside the strong-coupling regime. As a future perspective, we expect the coupling–coherence product ΩT_1 to improve by (i) mitigating leakage current—enabled by sapphire-substrate fabrication, which allows higher bias voltages; (ii) reducing the membrane–qubit separation to sub-μm scales; and (iii) applying standard surface passivation to extend the qubit lifetime.

When the membrane is in a thermal state ($n_{\mathrm{th}} \gtrsim 50$), the resonant qubit–mechanical interaction induces a qubit precession of typical angle $\theta_1\sqrt{n_{\mathrm{th}}}$, where

$$\theta_1 = \Omega\,\tau \sim 47 \text{ mrad}$$

is the rotation angle per phonon for an interaction time $\tau = 5\,\mu$s (limited by T_1 and T_2). In the next chapter, we demonstrate that, even outside

1: This number is estimated given the Rabi frequency $\Omega|\alpha|/2\pi \approx 0.6$ MHz in this experiment and the coupling strength $\Omega/2\pi \approx 1.5$ kHz that will be estimated in the next chapter.

the strong-coupling regime, this small rotation can be resolved through sequential weak qubit measurements within a single membrane coherence time [Cuj+19; Naj+24], unambiguously revealing the mechanical oscillator's quantum spectrum.

[Cuj+19]: Cujia et al. (2019), 'Tracking the precession of single nuclear spins by weak measurements', Nature
[Naj+24]: Najera-Santos et al. (2024), 'High-Sensitivity ac-Charge Detection with a MHz-Frequency Fluxonium Qubit', Physical Review X

9 Spectrum analyzer

Qubit as a spectrum analyzer. By performing a series of weak measurements, the qubit probes the mechanical motion. The Fourier transform of the qubit measurement record reveals the mechanical mode spectrum. The fundamental commutation relation $[\boldsymbol{a}, \boldsymbol{a}^\dagger] = 1$ imposes a built-in asymmetry between the absorption and emission spectra. When the qubit is in state $|g\rangle$ (resp. $|e\rangle$), it absorbs (resp. emits) energy from the membrane (back-action), thus sensing the absorption (resp. emission) spectrum of the membrane. This chapter demonstrates the qubit's operation as a spectrum analyzer and its ability to probe the quantum-induced asymmetry of the mechanical mode spectrum.

My contribution. As mentioned in Chap. 1 and Chap. 7, the experiments described throughout this thesis were my primary focus and responsibility during this project. This chapter presents the most significant results of my thesis that led to the publication [Ger+25]. The theoretical framework for the sideband asymmetry presented here was formalized by R. Rousseau and E. Flurin. Further details on the respective contributions can be found in [Ger+25].

[Ger+25]: Gerashchenko et al. (2025), 'Probing the quantum motion of a macroscopic mechanical oscillator with a radio-frequency superconducting qubit', Preprint

Outline. We begin with a semi-classical description of the membrane to build physical intuition for the qubit spectrum analyzer. Next, we demonstrate how the qubit exerts dynamical back-action on the membrane, leading to cooling or heating of the membrane, analogous to dynamical back-action in cavity optomechanics. However, a semi-classical treatment cannot capture the intrinsic spectral asymmetry. To reveal this asymmetry, we introduce a fully quantum framework based on measurement operators. We show that initializing the qubit in $|e\rangle$ or $|g\rangle$ selectively probes the membrane's emission or absorption spectrum, respectively. This protocol enables a direct measurement of the spectral asymmetry and thus probes the quantum nature of the mechanical oscillator.

9.1 Qubit as a spectrum analyzer in the semi-classical framework

9.1.1 Probing membrane's position with the qubit

Thermal fluctuations. To build physical intuition, we first treat the degrees of freedom of the membrane—its position $\boldsymbol{x}$ and momentum $\boldsymbol{p}$—classically. In the qubit drive frame, rotating at ω_q^{st}, the complex mechanical amplitude $\alpha(t) = \tilde{\alpha}(t)e^{i\Delta t}$ appears as a signal oscillating at frequency Δ, modulated by a slowly varying envelope $\tilde{\alpha}(t)$ that diffuses in phase space at a characteristic rate κ_m. We illustrate this in Fig. 9.1(c), where we plot the classical trajectory of the membrane dimensionless position $x(t) = \mathrm{Re}(\alpha(t))$ in a frame rotating detuned by Δ from the

1: Note that the time window was chosen specifically to illustrate the "modulation of the envelope" effect, which is most evident when the random walk of $\tilde{\alpha}(t)$ in phase space passes through the origin. However, generally speaking, the loss of correlation on the timescale $1/\kappa_m$ does not necessarily coincide with such a passage.

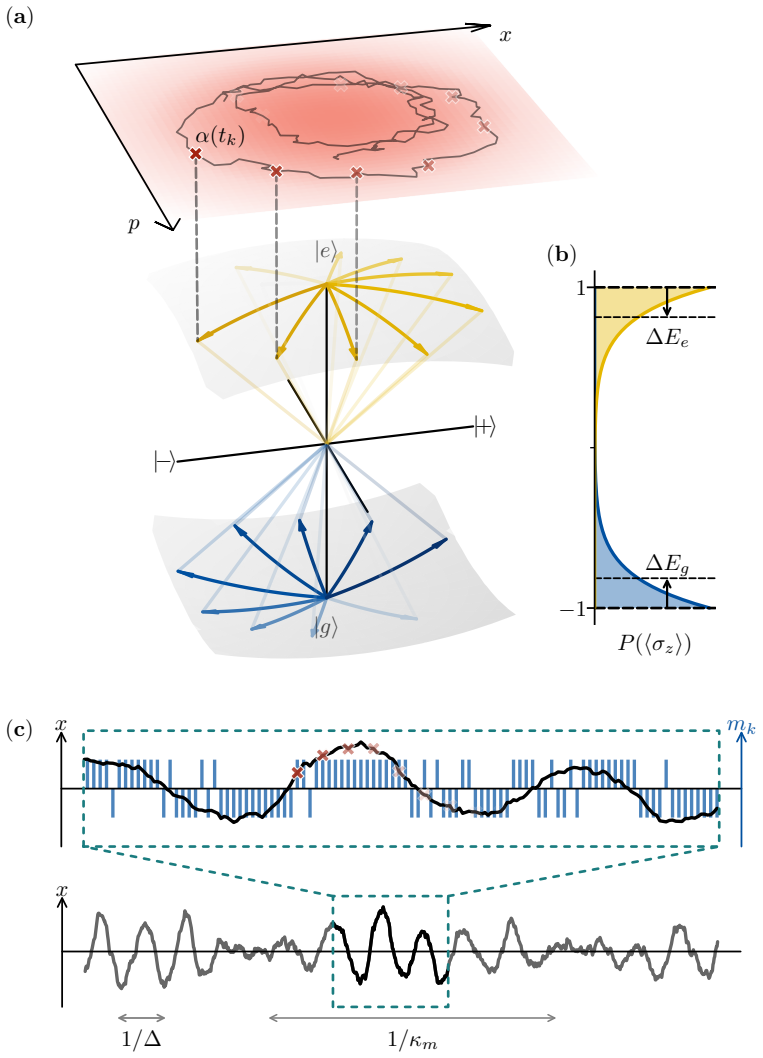

Fig. 9.1: Semi-classical explanation of the spectrum analyzer experiment. (a) Phase-space distribution of the mechanical mode amplitude, with an overlaid classical trajectory (black curve), and the qubit Bloch sphere. After initialization in $|g\rangle$ (blue) or $|e\rangle$ (yellow), the membrane mode acts as a coherent drive on the qubit. The complex mechanical mode amplitude is mapped onto the (σ_x, σ_y) plane of the Bloch sphere. **(b)** The average energy transferred from membrane to qubit during the interaction $\Delta E_{g/e}$ is displayed on the $\langle\sigma_z\rangle$ histogram on the right-hand inset. **(c)** Classical position $x(t)$ of the mechanical resonator, shown in a frame rotating at the qubit–mechanical-mode detuning Δ. The full trajectory is shown in gray, and the zoomed-in segment is highlighted in black. $x(t)$ oscillates at frequency Δ, with amplitude and phase undergoing random fluctuations at rate $\kappa_m = 1/T_1^m$. Blue bars represent qubit σ_x measurement outcomes after interaction with the mechanical mode. The data shown here is simulated for illustration purposes.

membrane frequency. One can see[1] that the global envelope of the oscillation is modulated at a time scale $1/\kappa_m$.

Measurement sequence. To detect this small qubit precession ($\theta_1\sqrt{n_{\text{th}}} \gtrsim 0.35$ rad) we implement the protocol described in Sec. 8.2.1: prepare $|g\rangle$ or $|e\rangle$, turn qubit frequency to the membrane frequency detuned by Δ during $\tau = 5\ \mu$s. Evolution during interaction is governed by

$$\boldsymbol{U}_I = \exp\left(-\frac{i}{\hbar}\boldsymbol{H}_{JC}\,\tau\right). \tag{9.1}$$

with the interaction part of the Jaynes-Cummings Hamiltonian in semi-classical approximation

$$\boldsymbol{H}_{\text{CC}} = i\frac{\Omega}{2}(\alpha^*\boldsymbol{\sigma}_- - \alpha\boldsymbol{\sigma}_+)$$

Since this interaction results in a qubit precession of a small angle $\theta_1\sqrt{n_{\text{th}}}$, we project the qubit onto σ_x. It is done by applying a $\sqrt{X}$ rotation in the $\omega_m - \Delta$ frame, followed by a σ_z measurement. This sequence is repeated N times within a single mechanical lifetime, $T \times N < T_1^m$, where $T = 15.5\ \mu$s is the total cycle time. This discrete sampling of the

mechanical amplitude $\alpha(t)$ at discrete times $t_k = k \times T$ provides a sequence of qubit measurements $\sigma_{x,k}$ after the interaction. Note that through all experiments in this chapter, the qubit is prepared and measured at its bare resonance frequency ω_q, and only tuned for the interaction with the membrane.

Mapping the mechanical mode to the qubit. In the small angle limit $|\alpha|\theta_1 \ll 1$, the rotation $\boldsymbol{U}_I$ maps the phase space coordinates $(\mathrm{Re}[\alpha(t_k)], \mathrm{Im}[\alpha(t_k)])$ of the mechanical oscillator onto the qubit's transverse components $(\sigma_{x,k}, \sigma_{y,k})$. Thereby, the expectation value of the σ_x qubit measurement is then given by

$$\langle \sigma_x(t_k) \rangle = \theta_1 x(t_k). \tag{9.2}$$

This mapping is illustrated in Fig. 9.1(a), where a complex trajectory in the phase space $(x = \mathrm{Re}(\alpha(t_k)), p = \mathrm{Im}(\alpha(t_k)))$ of the mechanical mode is mapped to the qubit's Bloch sphere; the evolution angle and phase are mapped to the amplitude and phase of the mechanical mode. By operating the qubit in a frame detuned by $\Delta > \kappa_m$, the free evolution between interaction pulses continually rotates the measurement axis, as illustrated by arrows going in the different directions. Consequently, successive $\sigma_{x,k}$ measurements sweep through all mechanical quadrature angles—probing both $\mathrm{Re}(\tilde{\alpha}(t_k))$ and $\mathrm{Im}(\tilde{\alpha}(t_k))$—in direct analogy with a heterodyne detection.

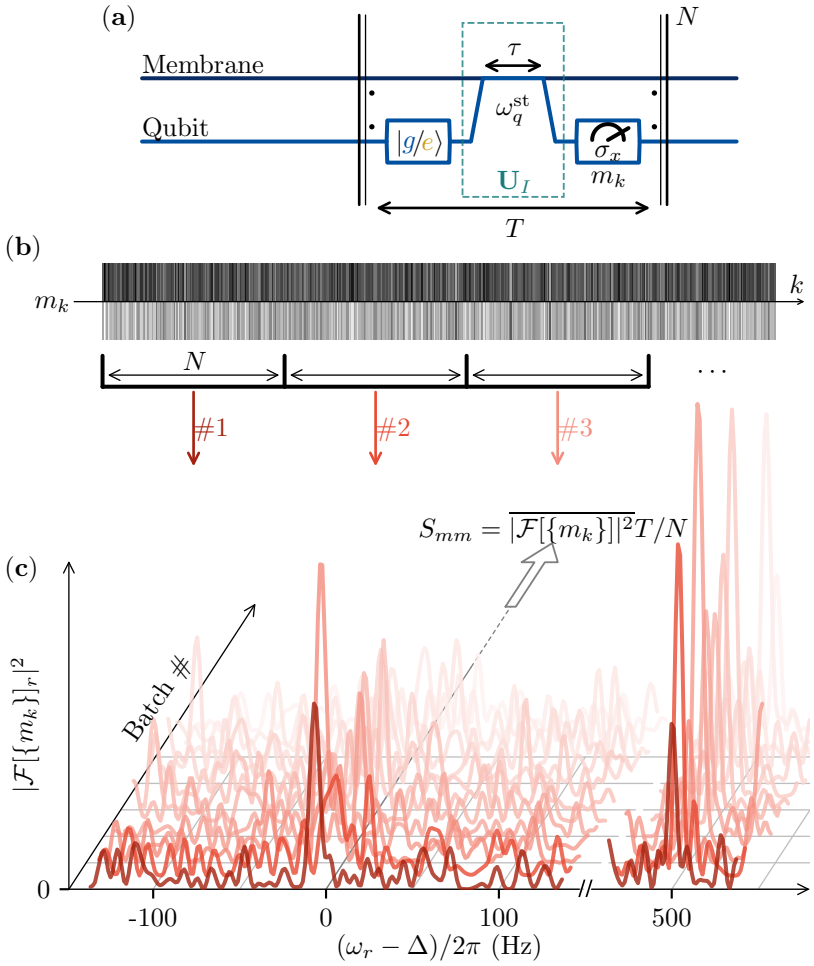

Fig. 9.2: Spectrum analyzer experimental protocol and data processing. (a) Experimental pulse sequence. The qubit, prepared in $|g\rangle$ or $|e\rangle$, is Stark-shifted to a detuning Δ from the membrane frequency for a duration τ, followed by a σ_x measurement. **(b)** Single-shot outcomes $m_k \in \{-1, 1\}$ (for initial state $|e\rangle$) are recorded, grouped into batches of length N, and each batch is discrete-Fourier-transformed. **(c)** Fourier transform of the first 10 batches represented as a waterfall plot.

9.1.2 Noise spectral density

Measuring the noise spectral density. Our goal is to characterize the steady-state thermal fluctuations of the membrane. For this, we directly compute the noise spectral density $S_{mm}^{g/e}$ of the telegraphic signal, where the superscript g/e denotes the initial qubit state used in the sequence. To do so, we first group the measurement outcomes $\{m_k\}$ into batches of $N = 10\,000$ samples (see Fig. 9.2(b)). For each batch, we compute the discrete Fourier transform (see Fig. 9.2(c)), $\mathcal{F}[m_k]_r$, where the frequency bins are defined as

$$\omega_r = \frac{2\pi r}{NT} \quad \text{for} \quad 0 \leq r < N.$$

We then average the squared magnitudes over all batches to obtain the noise spectral density

$$S_{mm}(\omega_r) = \overline{|\mathcal{F}[m_k]_r|^2} T/N, \tag{9.3}$$

with the overline indicating an average over batches corresponding to distinct classical trajectories $\{x(t_k)\}_{0\leq k<N}$. At our sampling rate, the aliasing-free frequency span is $\omega_{N/2}/2\pi \approx 32$ kHz. We therefore choose a small qubit-membrane detuning $\Delta/2\pi = 5$ kHz to ensure the oscillations of $x(t)$ remain entirely within the Nyquist zone.

Wiener-Khintchine theorem. In a semiclassical approximation where one neglects qubit backaction on the membrane, the trajectory $x(t)$ acts as an independent external drive. Using Eq. 9.2, this allow us to factorize the two-point correlator:

$$\mathcal{C}_k \equiv \overline{\langle \sigma_{x,k}\, \sigma_{x,0} \rangle} = \overline{\langle \sigma_{x,k} \rangle}\, \overline{\langle \sigma_{x,0} \rangle} = \theta_1^2 \overline{x(t_k)}\, \overline{x(0)}, \tag{9.4}$$

for $k \neq 0$, with $\mathcal{C}_0 = 1$. The Wiener–Khinchin theorem links the measured spectrum $S_{mm}^{g/e}(\omega)$ to the Fourier-transform of $\mathcal{C}_k$, giving

$$S_{mm}^{g/e}(\omega) = \frac{T}{N} \mathcal{F}[\mathcal{C}_k] = \frac{\theta_1^2}{4} S_{xx}(\omega) + T. \tag{9.5}$$

where S_{xx} is the position noise spectrum and T is the constant background arising from the sampling noise due to the binary nature of the qubit measurements. Thus, by measuring the spectrum $S_{mm}^{g/e}$ we can extract the position noise spectrum S_{xx}.

Spectrum normalization. Eq. 9.5 predicts a signal-to-noise ratio of $\theta_1^2/4T$, only dependent on the "single-phonon Rabi angle" θ_1 and periodicity T of the qubit measurements. However, due to imperfect qubit preparation and measurement, the signal gets reduced compared to this ideal case. In particular, if the preparation fidelities differs between g and e preparation, one expects systematic differences between the transduction of the mechanical spectrum between S_{mm}^{g} and S_{mm}^{e}. To compensate for those systematic effects, the spectrum is normalized using a calibration peak generated by a weak continuous tone applied on the qubit charge drive port (see Fig. 9.3(a)). It produces a coherent

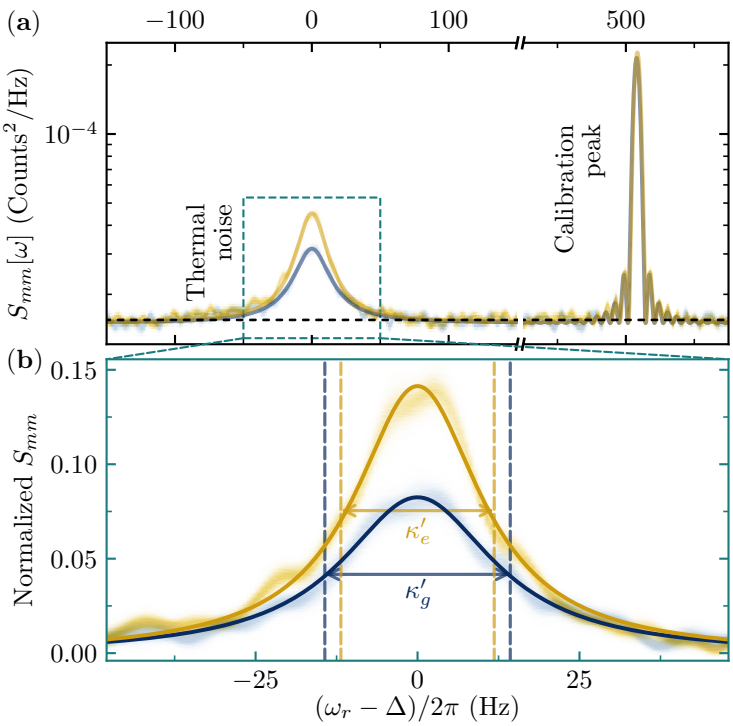

Fig. 9.3: Qubit-induced back-action. (a) Position noise spectra obtained for the qubit prepared in $|g\rangle$ (blue) and $|e\rangle$ (yellow), by averaging the squared magnitudes of the Fourier transforms shown in (e). The two Lorentzian peaks on the left are centered at the membrane frequency and correspond to the membrane thermal noise. Additionally, a weak calibration tone produces a sinc-shaped peak offset by 500 Hz towards higher frequencies. **(b)** Zoom on the Lorentzian peaks, with each spectrum normalized to the calibration-tone amplitude and flat background subtracted. The broader, lower-amplitude blue curve ($|g\rangle$) indicates qubit-induced cooling, while the narrower, higher-amplitude yellow curve ($|e\rangle$) indicates heating of the mechanical mode.

classical drive, which following the same classical reasoning (Eq. 9.5) results in a sinc-shaped (due to the finite sampling rate) peak in the spectrum purposely offset by 500 Hz towards higher frequencies. This normalization allows to compare results from different experiments with different preparation.

9.1.3 Membrane cooling and heating

Qubit-induced back-action. The spectra $S_{mm}^{g/e}$, measured with the qubit prepared in g or e, are shown in Fig. 9.3(a). Both spectra display Lorentzian peaks centered at ω_m that reflect the thermal fluctuations of the membrane, with the sampling noise background T indicated by a dashed line. Although Eq. 9.5 predicts identical spectra regardless of the initial qubit state, we observe significant differences in both the area and the linewidth of the thermal peaks, even once the spectra have been normalized with the calibration peak Fig. 9.3(b)

The discrepancies between measurements with the qubit prepared in $|g\rangle$ and $|e\rangle$ arise because our simple model neglects the energy exchange between the membrane and the qubit during their interaction. In particular, when the qubit is prepared in $|g\rangle$ it absorbs energy from the membrane, whereas when prepared in $|e\rangle$, it emits energy to the membrane (see Fig. 9.1(b)). On average, the qubit energy is lowered or raised on the Bloch sphere by

$$\frac{\Delta E_{g,e}}{\hbar\omega_q^{\mathrm{st}}} = \frac{\theta_1^2}{4} n_{\mathrm{th}}$$

due to the interaction. Due to energy conservation, the membrane must experience a qubit-induced relaxation or amplification at a rate $\kappa = \theta_1^2/4T$. The mechanical linewidth is thus modified (broadened or narrowed)

according to $\kappa'_{g/e} = \kappa_m \pm \kappa$. Furthermore the qubit being reinitialized in a pure state before each interaction, it provides a 0-occupation bath for the membrane, leading to an occupation $n'_{\mathrm{th},i} = \kappa_m n_{\mathrm{th}}/\kappa'_i$. This qubit-induced cooling (or heating) mechanism is analogous to the dynamical back-action effects observed in cavity optomechanics [Arc+06; Sch+06].

[Arc+06]: Arcizet et al. (2006), 'Radiation-pressure cooling and optomechanical instability of a micro-mirror', Nature
[Sch+06]: Schliesser et al. (2006), 'Radiation Pressure Cooling of a Micromechanical Oscillator Using Dynamical Backaction', Physical Review Letters

Membrane spectrum. Lorentzian fits to the thermal peaks yield linewidths of $\kappa'_g = 2\pi \times (28.5 \pm 1.4\ \mathrm{Hz})$ for $|g\rangle$ preparations and $\kappa'_e = 2\pi \times (23.7 \pm 0.7\ \mathrm{Hz})$ for $|e\rangle$ preparations. The inferred linewidth $\kappa_m = (\kappa'_e + \kappa'_g)/2 = 2\pi \times (26.1 \pm 0.8\ \mathrm{Hz})$ is in agreement with the previously measured $1/T_1^m = 2\pi \times (27.16 \pm 2.6\ \mathrm{Hz})$, confirming that the membrane exhibits no additional dephasing beyond its Fourier-limited linewidth. Moreover, the dynamical back-action rate $\kappa = 2\pi \times (2.4 \pm 0.8\ \mathrm{Hz})$ agrees with a full experimental model that includes finite qubit preparation and readout fidelities. The normalized, background-subtracted spectra are displayed in Fig. 9.3(b). The experimentally obtained area ratio $A_e/A_g = 1.42 \pm 0.11$ is in a reasonable agreement with the expected value $n'_{\mathrm{th},e}/n'_{\mathrm{th},g} = \kappa'_g/\kappa'_e = 1.22 \pm 0.07$. Here the dynamical backaction is moderate due to the relatively low cooperativity $\kappa/\kappa_m \sim 1/10$. Nevertheless, we expect that ground-state cooling will be within reach in future experiments—either by further increasing the mechanical lifetime [Tsa+17; Iva20] or by boosting the qubit–membrane coupling strength via a larger DC bias.

[Tsa+17]: Tsaturyan et al. (2017), 'Ultracoherent nanomechanical resonators via soft clamping and dissipation dilution', Nature Nanotechnology
[Iva20]: Ivanov (2020), 'Fabrication of patterned silicon nitride nanomembranes at the LKB', Zenodo

Error treatment. To ensure that observed differences in the spectra are not due to statistical fluctuations, we use a bootstrap approach to estimate the error. The spectra displayed in this section are obtained from $N_{\mathrm{batches}} \sim 1000$ Fourier-transformed batches. The averaging described by Eq. 9.3 is repeated over 1000 different datasets. Each dataset is obtained from the original data by drawing a resample with replacement of the N_{batches} Fourier-transformed batches. The 1000 different spectra are then accumulated in histograms, represented by the shaded regions in the spectra figures throughout this chapter (see Fig. 9.3(a-b) and Fig. 9.4(c)). The quoted errors on the fitted linewidths κ'_g, κ'_e and peak areas A_g, A_e are the standard deviations of those 1000 fits.

9.2 Spectrum analyzer in the fully quantum framework.

To fully account for the qubit-membrane interaction, we now consider the quantum evolution of the system.

9.2.1 Measurement operators

Evolution between measurements. The qubit being initialized in the pure state $|g\rangle$ or $|e\rangle$ before each interaction, we can safely assume that the initial joint qubit-membrane state is described by the factorized density operator $\boldsymbol{\rho_m} \otimes |i\rangle\langle i|$ (with $i \in \{g, e\}$). In the rotating frame, neglecting

the small qubit-membrane detuning Δ and decoherence during the short interaction time τ, the evolution operator $\boldsymbol{U}_I = e^{-i\theta_1/2(\boldsymbol{a}^\dagger\sigma - \boldsymbol{a}\sigma^\dagger)}$ writes [Rou22]

[Rou22]: Rouchon (2022), 'A tutorial introduction to quantum stochastic master equations based on the qubit/photon system', Annual Reviews in Control

$$\boldsymbol{U_I} = \begin{pmatrix} \cos\left(\frac{\theta_1}{2}\sqrt{\boldsymbol{a}^\dagger\boldsymbol{a}+1}\right) & -\boldsymbol{a}\frac{\sin\left(\frac{\theta_1}{2}\sqrt{\boldsymbol{a}^\dagger\boldsymbol{a}}\right)}{\sqrt{\boldsymbol{a}^\dagger\boldsymbol{a}}} \\ \boldsymbol{a}^\dagger\frac{\sin\left(\frac{\theta_1}{2}\sqrt{\boldsymbol{a}^\dagger\boldsymbol{a}+1}\right)}{\sqrt{\boldsymbol{a}^\dagger\boldsymbol{a}+1}} & \cos\left(\frac{\theta_1}{2}\sqrt{\boldsymbol{a}^\dagger\boldsymbol{a}}\right) \end{pmatrix}, \tag{9.6}$$

where the matrix is written in the qubit basis $\{|e\rangle, |g\rangle\}$. During the interaction, the system state evolves according to:

$$\boldsymbol{\rho}(t_k) \longrightarrow \boldsymbol{\rho}(t_{k+1}) = \boldsymbol{U}_I\left(\boldsymbol{\rho}_m(t_k)\otimes|i\rangle\langle i|\right)\boldsymbol{U}_I^\dagger. \tag{9.7}$$

Measurement operators. In the final state $\boldsymbol{\rho}(t_{k+1})$, we perform a projective measurement of the qubit in the $|\pm\rangle$ basis. The probability to measure the qubit in state $|+\rangle$ or $|-\rangle$ is then $p_{\pm,i}(t_{k+1}) = \mathrm{Tr}[\boldsymbol{\rho}(t_{k+1})|\pm\rangle\langle\pm|] = \mathrm{Tr}[\boldsymbol{M}_{\pm,i}\boldsymbol{\rho}_m(t_k)\boldsymbol{M}_{\pm,i}]$, where we defined the measurement operators

$$\begin{aligned} \boldsymbol{M}_{\pm,i} &= \langle\pm|\,\boldsymbol{U}_I\,|i\rangle \\ &\simeq \frac{1}{\sqrt{2}}\begin{cases} \mathbb{1} \mp \frac{\theta_1}{2}\boldsymbol{a} - \frac{\theta_1^2}{8}\boldsymbol{a}^\dagger\boldsymbol{a} & \text{if } i = g, \\ \mathbb{1} \pm \frac{\theta_1}{2}\boldsymbol{a}^\dagger - \frac{\theta_1^2}{8}\boldsymbol{a}\boldsymbol{a}^\dagger & \text{if } i = e. \end{cases} \end{aligned} \tag{9.8}$$

Evolution between measurements using Kraus maps. The final expression is calculated up to second order in θ_1 and we removed an irrelevant global phase term. The set of measurement operators describes a Positive Operator-Valued Measure which fulfills $\boldsymbol{M}_{+,i}^\dagger\boldsymbol{M}_{+,i} + \boldsymbol{M}_{-,i}^\dagger\boldsymbol{M}_{-,i} = \mathbb{1}$ up to second order in θ_1. Moreover, since $\boldsymbol{M}_{\pm,i}$ are close to the identity for $\theta_1 \ll 1$, this process describes more specifically a weak measurement. Upon qubit measurement, the membrane density matrix is projected to

$$\boldsymbol{\rho}_m^{\pm,i}(t_{k+1}) = \frac{\boldsymbol{M}_{\pm,i}\boldsymbol{\rho}_m(t_k)\boldsymbol{M}_{\pm,i}^\dagger}{p_{\pm,i}} \equiv \frac{\boldsymbol{\mathcal{K}}_{\pm,i}\boldsymbol{\rho}_m(t_k)}{p_{\pm,i}}. \tag{9.9}$$

Here, we defined the measurement Kraus map $\boldsymbol{\mathcal{K}}_{\pm,i}$. Introducing the dissipation and measurement superoperators as

$$\boldsymbol{\mathcal{D}}[\boldsymbol{L}]\boldsymbol{\rho} = \boldsymbol{L}\boldsymbol{\rho}\boldsymbol{L}^\dagger - \tfrac{1}{2}\{\boldsymbol{L}^\dagger\boldsymbol{L}, \boldsymbol{\rho}\}, \qquad \boldsymbol{\mathcal{M}}[\boldsymbol{L}]\boldsymbol{\rho} = \boldsymbol{L}\boldsymbol{\rho} + \boldsymbol{\rho}\boldsymbol{L}^\dagger,$$

the Kraus map can be expressed, up to second order in θ_1, as

$$\boldsymbol{\mathcal{K}}_{\pm,i} = \frac{1}{2}\left[\mathbb{1} \pm \frac{\theta_1}{2}\boldsymbol{\mathcal{M}}[\boldsymbol{L}_i] + \frac{\theta_1^2}{4}\boldsymbol{\mathcal{D}}[\boldsymbol{L}_i]\right], \tag{9.10}$$

where the weakly measured membrane operators, for each qubit state preparation are

$$\begin{aligned} \boldsymbol{L}_g &= -\boldsymbol{a}, \\ \boldsymbol{L}_e &= \boldsymbol{a}^\dagger. \end{aligned} \tag{9.11}$$

Backaction interpretation. The term proportional to $\boldsymbol{\mathcal{M}}[L]$ in Eq. 9.10 is the outcome-dependent "informational" (or quantum) backaction, which steers the mechanical state based on the measurement result, while the

$\mathcal{D}[L]$ contribution describes a measurement-independent, qubit-induced damping channel, directly analogous to dynamical backaction in cavity optomechanics [AKM14].

[AKM14]: Aspelmeyer et al. (2014), 'Cavity optomechanics', Reviews of Modern Physics

9.2.2 Dynamical back-action

Mean membrane evolution. When the measurement outcomes are disregarded, the membrane evolves under the average Kraus map

$$\overline{\mathcal{K}}_i = \mathcal{K}_{+,i} + \mathcal{K}_{-,i} = \mathbb{1} + \frac{\theta_1^2}{4}\,\mathcal{D}[\boldsymbol{L}_i]. \tag{9.12}$$

During the interval T between two successive qubit measurements, the membrane undergoes free-evolution and dissipation to the thermal bath, described by the Lindbladian

$$\mathcal{L}_m = -i[\Delta \boldsymbol{a}^\dagger \boldsymbol{a}, \cdot] + \kappa_m n_{\mathrm{th}}\,\mathcal{D}[\boldsymbol{a}] + \kappa_m (n_{\mathrm{th}} + 1)\,\mathcal{D}[\boldsymbol{a}^\dagger]. \tag{9.13}$$

Interleaving the average Kraus map (9.12) with the membrane's free evolution $e^{\mathcal{L}_m T}$ yields

$$\boldsymbol{\rho}(t_k + T) = e^{\mathcal{L}_m T}\,\overline{\mathcal{K}}_i\,\boldsymbol{\rho}(t_k) \approx e^{\mathcal{L}_{m,i} T}, \tag{9.14}$$

Coarse-grained approximation. In a coarse-grained approximation, the effective Lindbladian becomes $\mathcal{L}_{m,i} = \mathcal{L}_m + \kappa\,\mathcal{D}[\boldsymbol{L}_i]$, with the dynamical-backaction rate $\kappa = \theta_1^2/4T$. This coarse-grained evolution is identical to that induced by the coupling at a rate $\kappa'_{m,i}$ to an effective environment with thermal occupation $n'_{\mathrm{th,i}}$, where

$$\kappa'_{m,i} = \begin{cases} \kappa_m + \kappa, & i = g, \\ \kappa_m - \kappa, & i = e, \end{cases} \qquad n'_{\mathrm{th,i}} = \frac{\kappa_m n_{\mathrm{th}} + \delta_{e,i}\,\kappa}{\kappa'_{m,i}}.$$

Here $\delta_{e,i} = 1$ if $i = e$ (zero otherwise). In other words, dynamical back-action modifies both the damping rate and the apparent thermal occupation by factors $1 \pm C$, where $C = \kappa/\kappa_m$ is the cooperativity, in line with conventions commonly used in optomechanics. The last additive term in the occupation $\kappa/\kappa'_m = C/(1 - C)$ appearing when the qubit is prepared in state $|e\rangle$ corresponds to the additional fluctuations resulting from quantum backaction. In our experiment where the cooperativity is on the order of 10 %, we can safely neglect this contribution in the area of the spectrum $S^e_{mm}(\omega)$.

Classical backaction and its mitigation. In Fig. 9.3, the difference between $S^g_{xx}(\omega_m)$ and $S^e_{xx}(\omega_m)$ is largely due to classical backaction–cooling for $|g\rangle$ and amplification for $|e\rangle$. In other words, the equilibrium occupation $\langle \boldsymbol{a}^\dagger \boldsymbol{a}\rangle$ is different in the two experimental sequences. To mitigate this effect, and isolate the small quantum asymmetry, we implemented an experimental sequence where the qubit is alternatively prepared in $|g\rangle$ and $|e\rangle$ (see Fig. 9.4b). In this manner, the net evolution over a full

cycle is generated by

$$\mathcal{L}_{m,ge} = \mathcal{L}_m + \kappa\,\mathcal{D}[a] + \kappa\,\mathcal{D}[a^\dagger], \tag{9.15}$$

so that

$$\kappa'_{m,ge} = \kappa_m, \qquad n'_{\mathrm{th},ge} = n_{\mathrm{th}} + C.$$

Thus the alternating protocol cancels the first-order dynamical back-action, leaving only a negligible residual heating of order $C \ll n_{\mathrm{th}}$. Finally, note that the interval between two identical preparation steps T is effectively doubled (see Fig. 9.4b) compared to the previous protocol, which halves both the back-action rate κ and the cooperativity C.

9.2.3 Emission and absorption spectra

Full Kraus map. The full Kraus operators factor into the mean-evolution map and a pure backaction part:

$$\mathcal{K}_{\pm,i} = \overline{\mathcal{K}}_i \mathcal{K}'_{\pm,i}\ , \text{ with } \mathcal{K}'_{\pm,i} = \frac{1}{2}\left(\mathbb{1} \pm \frac{\theta_1}{2}\mathcal{M}[L_i]\right). \tag{9.16}$$

One can then calculate the joint probability to obtain outcomes m_0 and m_k separated by k full cycles becomes

$$\mathbb{P}[m_k, m_0|\rho_m, i] \simeq \mathrm{Tr}[\mathcal{K}'_{m_k,i} e^{\widetilde{\mathcal{L}}_m kT} \mathcal{K}'_{m_0,i}\rho_m], \tag{9.17}$$

where $\widetilde{\mathcal{L}}_m$ denotes the effective Lindbladian for the chosen preparation sequence—either $\mathcal{L}_{m,g}$ or $\mathcal{L}_{m,e}$ for identical preparations, or $\mathcal{L}_{m,ge}$ for alternating ground/excited protocols.

Autocorrelation. We then express the discrete-time autocorrelation of the measurement outcomes as

$$\mathcal{C}_k^i = \sum_{m_0,m_k=\pm1} m_k\, m_0\, \mathbb{P}[m_k, m_0 \mid \rho_m, i],$$

so that for $k = 0$ trivially $\mathcal{C}_0^i = 1$. For $k \neq 0$, inserting the expression (9.17), and using $\sum_{m=\pm1} m\,\mathcal{K}'_{m,i} = \frac{\theta_1}{2}\,\mathcal{M}[L_i]$ yields

$$\mathcal{C}_k^i = \frac{\theta_1^2}{4}\,\mathrm{Tr}\big[\mathcal{M}[L_i]\, e^{\tilde{\mathcal{L}}_m\, kT}\, \mathcal{M}[L_i]\,\rho_m\big].$$

Defining the Heisenberg picture operator $O(t)$ as $O\rho(t) = Oe^{\widetilde{\mathcal{L}}_m t}\rho = O(t)\rho$, the measurement correlator reads

$$\mathcal{C}_k^i\big|_{k\neq0} = \frac{\theta_1^2}{4}\,\mathrm{Re}\big\langle L_i^\dagger(t_k)\, L_i(0)\big\rangle.$$

Measured spectra. The discrete-time spectrum follows by Fourier transform.

$$\begin{aligned} \mathcal{S}^{g}_{mm}(\omega) &= T + \frac{\theta_1^2}{4}\overline{\mathcal{S}}_{a^\dagger a}(\omega), \\ \mathcal{S}^{e}_{mm}(\omega_k) &= T + \frac{\theta_1^2}{4}\overline{\mathcal{S}}_{a a^\dagger}(\omega), \end{aligned} \tag{9.18}$$

where $\overline{\mathcal{S}}_{a^\dagger a}(\omega)$ and $\overline{\mathcal{S}}_{a a^\dagger}(\omega)$ are the one-sided emission and absorption mechanical spectra, in the actual protocol, and are given by

$$\boxed{\overline{\mathcal{S}}_{L_i^\dagger L_i}(\omega) = \frac{(\tilde{n}_{\mathrm{th}} + \delta_{ei})\tilde{\kappa}_m}{(\omega - \Delta)^2 + (\tilde{\kappa}_m/2)^2},} \tag{9.19}$$

where $i \in \{g, e\}$ and $(\tilde{n}_{\mathrm{th}}, \tilde{\kappa}_m) = (n'_{\mathrm{th},p}, \kappa'_{m,p})$ with $p \in \{g, e, ge\}$ based on the used protocol. For the alternating preparation protocol, the denominator is the same for both initial states. Thus, the asymmetry in the amplitudes of the spectra for the two initial states, S^g_{mm} and S^e_{mm}, originates solely from the factor $\tilde{n}_{\mathrm{th}} + \delta_{ei}$. Experimentally, this factor can be directly accessed, and in the ideal case we expect the resulting asymmetry to be equal to 1.

9.2.4 Sideband asymmetry

The experimental spectra S^g_{mm} and S^e_{mm}, constructed from even-indexed and odd-indexed measurements corresponding to the pre-interaction ground state $|g\rangle$ and excited state $|e\rangle$, respectively, are shown in Fig. 9.4c. They exhibit a small but statistically significant difference in their areas. To confirm that this offset has a quantum origin, we vary the cryostat temperature, T_{cryo}, within the interval [10 mK, 30 mK][1]. The spectra are normalized using a calibration peak and corrected by subtracting the sampling noise background. The thermalization of the mechanical mode with the cryostat is demonstrated by the linear dependence of the mean peak area A_{ge}, such that

$$T_{\mathrm{m}} = \lambda\, A_{ge} = T_{\mathrm{cryo}} + T_{\mathrm{offset}}.$$

The fit of A_{ge} versus T_{cryo} allows recalibration of the mode area in units of quanta $\overline{n}$ (right axis of Fig. 9.4d) and yields an offset $T_{\mathrm{offset}} = 2.2$ mK. The constant offset can be attributed to a miscalibration of the generic Bluefors temperature sensor, additional acoustic noise transduced into the mechanical mode, or a slight temperature gradient between the sensor near the mixing chamber and the sample holder. Most notably, the calibrated bijection between the peak area and $\overline{n}$ allows us to calculate the difference between the peak areas in S^g_{mm} and S^e_{mm}, which shows a statistically significant offset from zero, and on average is 1.25(25) phonons. The ideal quantum asymmetry is exactly 1 phonon (Eq. 9.19). However, experimental imperfections—such as qubit decoherence and state-preparation infidelity—lead to an increased asymmetry, given by [Ger+25]:

$$\boxed{\overline{n}_e - \overline{n}_g = \frac{\tau_2}{\tau_\Sigma}\,\frac{\eta_g + \eta_e}{2\,\eta_g\,\eta_e}.} \tag{9.20}$$

1: Since the mechanical peak area scales as V^2_{offset}, any slow drift in V_{offset} (Sec. 8.1.4) during a monotonic temperature sweep would introduce a spurious temperature dependence in the sideband-asymmetry data of Fig. 9.4c-d. To avoid this, we interleaved the measurements by cycling the cryostat through each setpoint (2 h per point) for 16 complete cycles, so that any residual offset drift acts uniformly across all temperatures.

[Ger+25]: Gerashchenko et al. (2025), 'Probing the quantum motion of a macroscopic mechanical oscillator with a radio-frequency superconducting qubit', Preprint

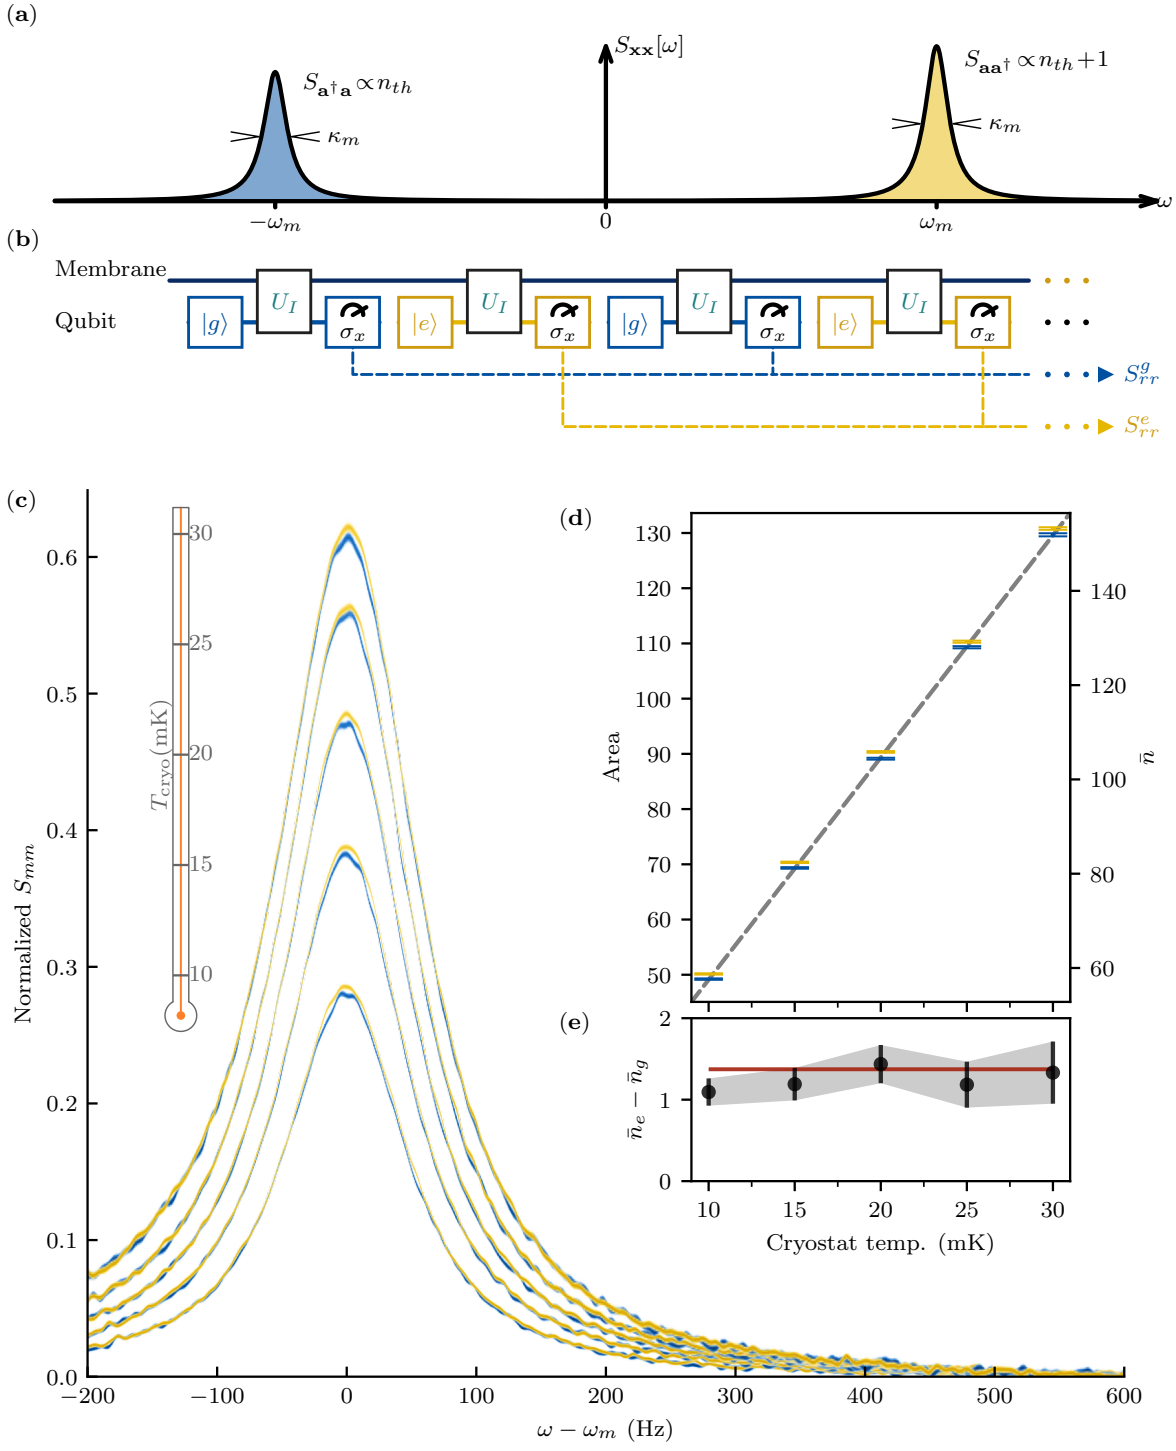

Fig. 9.4: Quantum position-spectrum asymmetry. (a) Position-noise spectrum of a harmonic oscillator at thermal equilibrium, showing absorption (blue peak at $-\omega_m$, area $\propto \overline{n}_{\mathrm{th}}$) and emission (yellow peak at $+\omega_m$, area $\propto \overline{n}_{\mathrm{th}} + 1$). **(b)** Experimental pulse sequence: to eliminate the average dynamical backaction, we alternately initialize the qubit in the ground state $|g\rangle$ and the excited state $|e\rangle$. From the resulting measurement records, we then reconstruct the emission spectrum S^g_{xx} (when starting in $|g\rangle$) and the absorption spectrum S^e_{xx} (when starting in $|e\rangle$). **(c)** Reconstructed spectra of 4×10^5 Fourier-transformed batches for $|g\rangle$ (blue) and $|e\rangle$ (yellow) preparations at varying cryostat temperatures at 10, 15, 20, 25, and 30 mK, as illustrated by the thermometer. **(d)** Lorentzian-peak areas at ω_m versus cryostat temperature: blue and yellow markers indicate S^g_{xx} and S^e_{xx} areas for each qubit preparation. A linear fit to the mean-area of the 5 temperature series is shown (dashed line), and the right-hand axis converts the fitted peak area into mean phonon number using this calibration. **(e)** Measured phonon-number difference between $|e\rangle$ and $|g\rangle$ spectra versus cryostat temperature. The horizontal red line at 1.37 quanta is the theoretical prediction, which differs from 1 due to qubit preparation infidelity and finite lifetime. All points include error bars representing the standard deviation from bootstrap fits to random FFT-batch sub-ensembles.

with η_g and η_e the state-preparation fidelities,

$$\tau_1(t) = \frac{1 - e^{-\kappa_1 t}}{\kappa_1}, \quad \tau_2(t) = \frac{1 - e^{-\kappa_2 t}}{\kappa_2}, \quad \tau_\Sigma(t) = \frac{\kappa_1\,\tau_1(t) - \kappa_2\,\tau_2(t)}{\kappa_1 - \kappa_2},$$

where κ_1 and κ_2 are the relaxation and total (relaxation + dephasing) rates, respectively. This gives a theoretical prediction of 1.37 phonons for the sideband asymmetry, in good agreement with the experimental result.

The measured spectra S^g_{mm} and S^e_{mm} in case of finite preparation fidelities (η_g and η_e), finite readout errors ($\varepsilon_g = 1 - F_g$ and $\varepsilon_e = 1 - F_e$) and finite T_1 and T_2 are given by [Ger+25]:

$$\begin{aligned}
\mathcal{S}^g_{mm}(\omega_k) =& (1 - \varepsilon_g - \varepsilon_e)^2 \eta_g \frac{\Omega^2 \tau_\Sigma^2}{4} \Big(\frac{\eta_g + \tau_2/\tau_\Sigma}{2} \bar{\mathcal{S}}_{\boldsymbol{a}^\dagger \boldsymbol{a}}(\omega_k) + \\
& \frac{\eta_g - \tau_2/\tau_\Sigma}{2} \bar{\mathcal{S}}_{\boldsymbol{a}\boldsymbol{a}^\dagger}(\omega_k) \Big) + T\,(1 - (\varepsilon_g - \varepsilon_e)^2), \\
\mathcal{S}^e_{mm}(\omega_k) =& (1 - \varepsilon_g - \varepsilon_e)^2 \eta_e \frac{\Omega^2 \tau_\Sigma^2}{4} \Big(\frac{\eta_e - \tau_2/\tau_\Sigma}{2} \bar{\mathcal{S}}_{\boldsymbol{a}^\dagger \boldsymbol{a}}(\omega_k) + \\
& \frac{\eta_e + \tau_2/\tau_\Sigma}{2} \bar{\mathcal{S}}_{\boldsymbol{a}\boldsymbol{a}^\dagger}(\omega_k) \Big) + T\,(1 - (\varepsilon_g - \varepsilon_e)^2).
\end{aligned} \tag{9.21}$$

Since all terms in the expression are determined from independent measurements except for the coupling strength Ω, fitting the amplitude of the thermal peak for preparation in $|g\rangle$ at 10 mK gives $\Omega = 1.50(8)$ kHz.

Chapter summary

In conclusion, we have reported the first observation of mechanical spectrum asymmetry via coupling to a two-level system, as originally predicted by Clerk et al. [Cle+10]. Here, we applied noise thermometry—relying on the proportionality between spectral area and temperature—to verify that the quantum asymmetry is indeed close to phonon unity. We then exploited the offset between S^g_{mm} and S^e_{mm} to calibrate the spectra in absolute quanta. In other words, this experiment can serve as a calibration-free thermometer operating at cryostat temperatures.

Most importantly, these results furnish an unambiguous demonstration of the non-commutation of the creation and annihilation operators, $[a, a^\dagger] = 1$, proving that the qubit interacts with a genuine quantum harmonic oscillator rather than with a classical noise source. In the next chapter (Conclusion and Perspectives), we will discuss the broader implications of this work and outline future directions.

Summary and Perspective 10

Summary of the experimental findings

The central achievement of this thesis is the design, fabrication, and characterization of qubit–membrane coupling, demonstrating, to our knowledge, the first resonant interaction at 4 MHz between a truly macroscopic (5.3 ng) silicon-nitride membrane and a heavy fluxonium qubit.

We demonstrated the qubit–membrane interaction in both the dispersive and resonant regimes. In the dispersive regime, we measured the membrane's resonant frequency and quality factor and demonstrated electrostatic softening. In the resonant regime, despite the weak coupling, we demonstrated classical Rabi oscillations of the membrane in a large coherent state and observed dynamical backaction cooling from thermal equilibrium with cooperativity $\kappa/\kappa_m \approx 1/10$.

Most notably, in direct analogy to Haroche's cavity-QED experiments [Gle+07; Del+08], we performed over 300 projective qubit measurements within the membrane's 6 ms lifetime. These binary outcomes yield bit-by-bit information on the membrane position, allowing reconstruction of the full position-noise spectrum and—by preparing the qubit in $|g\rangle$ or $|e\rangle$—separate extraction of its absorption (negative-frequency) and emission (positive-frequency) spectra. The observed phonon emission–absorption asymmetry provides an unambiguous signature of the membrane's quantum nature with a genuine two-level system, as proposed theoretically [Cle+10].

[Gle+07]: Gleyzes et al. (2007), 'Quantum jumps of light recording the birth and death of a photon in a cavity', Nature
[Del+08]: Deléglise et al. (2008), 'Reconstruction of non-classical cavity field states with snapshots of their decoherence', Nature

[Cle+10]: Clerk et al. (2010), 'Introduction to Quantum Noise, Measurement and Amplification', Reviews of Modern Physics

By extending coherent control into the 1–10 MHz band—over two orders of magnitude below prior platforms—we open the first experimental window onto gravity-induced collapse models, which predict spatial-superposition decoherence on millisecond timescales. The Diósi–Penrose estimate of about 0.5 ms [GS21] nearly coincides with our measured mechanical decoherence of $\sim$ 0.3 ms, providing us the first direct experimental pathway to test whether gravity alone can destroy quantum coherence in a truly macroscopic object.

[GS21]: Gely et al. (2021), 'Superconducting electro-mechanics to test Diósi–Penrose effects of general relativity in massive superpositions', AVS Quantum Science

Looking ahead, this platform naturally lends itself to quantum and adaptive feedback protocols—using membrane drives or conditional qubit preparations based on real-time readout—to achieve active stabilization, state steering, and generation of non-classical mechanical states [Say+11].

[Say+11]: Sayrin et al. (2011), 'Real-time quantum feedback prepares and stabilizes photon number states', Nature

Possible improvement of the experiment and prospects

The primary objective of *Mecaflux* project remains to reach a regime in which predicted gravity-induced decoherence of the mechanical resonator exceeds its intrinsic thermal decoherence. Several optimizations could improve the device performance of the current iteration.

First, phononic-crystal engineering is expected to boost the electromechanical cooperativity by three orders of magnitude [Sei+22], enabling ground-state cooling and a corresponding reduction in thermal decoherence.

[Sei+22]: Seis et al. (2022), 'Ground state cooling of an ultracoherent electromechanical system', Nature Communications

Second, entering the strong-coupling regime—necessary for preparing Schrödinger-cat states—requires increasing the DC bias voltage and reducing the qubit–membrane gap. A new iteration of the fluxonium chip, fabricated on a sapphire substrate by L. Balembois and P. Manset, mitigates avalanche ionization and has demonstrated preliminary voltages up to 100 V without degrading qubit performance (on an uncoupled device).

Third, standard surface passivation could extend the qubit coherence time T_2, and by more precisely tuning the fluxonium parameters to bring the qubit frequency into resonance with the membrane, we reduce the required AC–Stark shift and thereby enhance on-resonance coherence.

Together, these improvements should maximize the coupling–coherence product ΩT_2, closing the current order-of-magnitude gap and enabling operation in the strong-coupling regime—thereby opening the door to direct experimental tests of gravitational collapse models.

Bibliography

Here are the references in citation order.

[OZ16] Tjerk H. Oosterkamp and Jan Zaanen. *A clock containing a massive object in a superposition of states; what makes Penrosian wavefunction collapse tick?* arXiv:1401.0176 [quant-ph]. Oct. 2016. DOI: 10.48550/arXiv.1401.0176. URL: http://arxiv.org/abs/1401.0176 (visited on 03/05/2025) (cited on page 1).

[Dió89] L. Diósi. 'Models for universal reduction of macroscopic quantum fluctuations'. In: *Physical Review A* 40.3 (Aug. 1989). Publisher: American Physical Society, pp. 1165–1174. DOI: 10.1103/PhysRevA.40.1165. (Visited on 03/05/2025) (cited on page 1).

[Pen96] Roger Penrose. 'On Gravity's role in Quantum State Reduction'. en. In: *General Relativity and Gravitation* 28.5 (May 1996), pp. 581–600. DOI: 10.1007/BF02105068. (Visited on 03/05/2025) (cited on page 1).

[Kle+08] Dustin Kleckner et al. 'Creating and verifying a quantum superposition in a micro-optomechanical system'. en. In: *New Journal of Physics* 10.9 (Sept. 2008), p. 095020. DOI: 10.1088/1367-2630/10/9/095020. (Visited on 03/05/2025) (cited on page 1).

[Che13] Yanbei Chen. 'Macroscopic Quantum Mechanics: Theory and Experimental Concepts of Optomechanics'. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 46.10 (May 2013). arXiv:1302.1924 [quant-ph], p. 104001. DOI: 10.1088/0953-4075/46/10/104001. (Visited on 03/05/2025) (cited on page 1).

[Dió07] Lajos Diósi. 'Notes on certain Newton gravity mechanisms of wavefunction localization and decoherence'. In: *Journal of Physics A: Mathematical and Theoretical* 40.12 (Mar. 2007), pp. 2989–2995. DOI: 10.1088/1751-8113/40/12/S07 (cited on page 1).

[HR06] Serge Haroche and Jean-Michel Raimond. *Exploring the Quantum*. DOI: 10.1093/acprof:oso/9780198509141.001.0001. Oxford University Press, Aug. 2006 (cited on pages 2, 21, 22).

[GS21] Mario F. Gely and Gary A. Steele. 'Superconducting electro-mechanics to test Diósi–Penrose effects of general relativity in massive superpositions'. In: *AVS Quantum Science* 3.3 (Sept. 2021), p. 035601. DOI: 10.1116/5.0050988. (Visited on 03/05/2025) (cited on pages 2, 3, 148).

[Sid+95] J. A. Sidles et al. 'Magnetic resonance force microscopy'. In: *Reviews of Modern Physics* 67.1 (Jan. 1995). Publisher: American Physical Society, pp. 249–265. DOI: 10.1103/RevModPhys.67.249. (Visited on 03/05/2025) (cited on page 3).

[Ros+18] Massimiliano Rossi et al. 'Measurement-based quantum control of mechanical motion'. en. In: *Nature* 563.7729 (Nov. 2018). Publisher: Nature Publishing Group, pp. 53–58. DOI: 10.1038/s41586-018-0643-8. (Visited on 03/05/2025) (cited on page 3).

[Neu+21] Leonhard Neuhaus et al. *Laser cooling of a Planck mass object close to the quantum ground state*. arXiv:2104.11648 [quant-ph]. Apr. 2021. DOI: 10.48550/arXiv.2104.11648. URL: http://arxiv.org/abs/2104.11648 (visited on 03/05/2025) (cited on page 3).

[Fog+21] Francesco Fogliano et al. 'Mapping the Cavity Optomechanical Interaction with Subwavelength-Sized Ultrasensitive Nanomechanical Force Sensors'. In: *Physical Review X* 11.2 (Apr. 2021). Publisher: American Physical Society, p. 021009. DOI: 10.1103/PhysRevX.11.021009. (Visited on 03/05/2025) (cited on page 3).

[Teu+11] J. D. Teufel et al. 'Circuit cavity electromechanics in the strong-coupling regime'. en. In: *Nature* 471.7337 (Mar. 2011). Publisher: Nature Publishing Group, pp. 204–208. DOI: 10.1038/nature09898. (Visited on 03/05/2025) (cited on page 3).

[Rug+04] D. Rugar et al. 'Single spin detection by magnetic resonance force microscopy'. en. In: *Nature* 430.6997 (July 2004). Publisher: Nature Publishing Group, pp. 329–332. DOI: 10.1038/nature02658. (Visited on 03/05/2025) (cited on page 3).

[Arr+19] Patricio Arrangoiz-Arriola et al. 'Resolving the energy levels of a nanomechanical oscillator'. en. In: *Nature* 571.7766 (July 2019). Publisher: Nature Publishing Group, pp. 537–540. DOI: 10.1038/s41586-019-1386-x. (Visited on 03/05/2025) (cited on page 3).

[OCo+10] A. D. O'Connell et al. 'Quantum ground state and single-phonon control of a mechanical resonator'. en. In: *Nature* 464.7289 (Apr. 2010). Publisher: Nature Publishing Group, pp. 697–703. DOI: 10.1038/nature08967. (Visited on 03/05/2025) (cited on page 3).

[Sat+18] K. J. Satzinger et al. 'Quantum control of surface acoustic-wave phonons'. en. In: *Nature* 563.7733 (Nov. 2018). Publisher: Nature Publishing Group, pp. 661–665. DOI: 10.1038/s41586-018-0719-5. (Visited on 03/05/2025) (cited on page 3).

[Cha+11] Jasper Chan et al. 'Laser cooling of a nanomechanical oscillator into its quantum ground state'. en. In: *Nature* 478.7367 (Oct. 2011). Publisher: Nature Publishing Group, pp. 89–92. DOI: 10.1038/nature10461. (Visited on 03/05/2025) (cited on page 3).

[Chu+17] Yiwen Chu et al. 'Quantum acoustics with superconducting qubits'. In: *Science* 358.6360 (Oct. 2017). Publisher: American Association for the Advancement of Science, pp. 199–202. DOI: 10.1126/science.aao1511. (Visited on 03/05/2025) (cited on page 3).

[Ver+12] E. Verhagen et al. 'Quantum-coherent coupling of a mechanical oscillator to an optical cavity mode'. en. In: *Nature* 482.7383 (Feb. 2012). Publisher: Nature Publishing Group, pp. 63–67. DOI: 10.1038/nature10787. (Visited on 03/05/2025) (cited on page 3).

[Gha+18] A. H. Ghadimi et al. 'Elastic strain engineering for ultralow mechanical dissipation'. In: *Science* 360.6390 (May 2018). Publisher: American Association for the Advancement of Science, pp. 764–768. DOI: 10.1126/science.aar6939. (Visited on 09/04/2024) (cited on pages 3, 53).

[Tsa+17] Y. Tsaturyan et al. 'Ultracoherent nanomechanical resonators via soft clamping and dissipation dilution'. en. In: *Nature Nanotechnology* 12.8 (Aug. 2017). Publisher: Nature Publishing Group, pp. 776–783. DOI: 10.1038/nnano.2017.101. (Visited on 09/04/2024) (cited on pages 2, 53, 141).

[Iva20] Edouard Ivanov. 'Fabrication of patterned silicon nitride nanomembranes at the LKB'. In: *Zenodo* (Oct. 2020). DOI: 10.5281/ZENODO.4120417 (cited on pages 2, 59, 141).

[Hof+09] Max Hofheinz et al. 'Synthesizing arbitrary quantum states in a superconducting resonator'. en. In: *Nature* 459.7246 (May 2009), pp. 546–549. DOI: 10.1038/nature08005 (cited on pages 2, 26).

[Chu+18] Yiwen Chu et al. 'Creation and control of multi-phonon Fock states in a bulk acoustic-wave resonator'. en. In: *Nature* 563.7733 (Nov. 2018), pp. 666–670. DOI: 10.1038/s41586-018-0717-7 (cited on pages 2, 3, 26).

[Bie+19] A. Bienfait et al. 'Phonon-mediated quantum state transfer and remote qubit entanglement'. In: *Science* (2019). DOI: 10.1126/science.aaw8415 (cited on page 3).

[Bil+23] Marius Bild et al. 'Schrödinger cat states of a 16-microgram mechanical oscillator'. In: *Science* 380.6642 (Apr. 2023). Publisher: American Association for the Advancement of Science, pp. 274–278. DOI: 10.1126/science.adf7553. (Visited on 03/05/2025) (cited on page 3).

[Gel+19] Mario F. Gely et al. 'Observation and stabilization of photonic Fock states in a hot radio-frequency resonator'. en. In: *Science* 363.6431 (Mar. 2019), pp. 1072–1075. DOI: 10.1126/science.aaw3101 (cited on page 3).

[VML18] Jeremie J. Viennot, Xizheng Ma, and Konrad W. Lehnert. 'Phonon number sensitive electromechanics'. en. In: *Physical Review Letters* 121.18 (Oct. 2018). arXiv:1808.02673 [cond-mat, physics:quant-ph], p. 183601. DOI: 10.1103/PhysRevLett.121.183601 (cited on pages 3, 65).

[Ma+21] X. Ma et al. 'Non-classical energy squeezing of a macroscopic mechanical oscillator'. en. In: *Nature Physics* 17.3 (Mar. 2021), pp. 322–326. DOI: 10.1038/s41567-020-01102-1 (cited on page 3).

[Ear+18] N. Earnest et al. 'Realization of a Λ System with Metastable States of a Capacitively Shunted Fluxonium'. In: *Physical Review Letters* 120.15 (Apr. 2018). Publisher: American Physical Society, p. 150504. DOI: 10.1103/PhysRevLett.120.150504. (Visited on 03/05/2025) (cited on pages 3, 4).

[Zha+21] Helin Zhang et al. 'Universal fast flux control of a coherent, low-frequency qubit'. en. In: *Physical Review X* 11.1 (Jan. 2021). arXiv:2002.10653 [cond-mat, physics:quant-ph], p. 011010. DOI: 10.1103/PhysRevX.11.011010 (cited on pages 3, 4, 23, 31, 35, 36, 51).

[Naj+24] B.-L. Najera-Santos et al. 'High-Sensitivity ac-Charge Detection with a MHz-Frequency Fluxonium Qubit'. en. In: *Physical Review X* 14.1 (Jan. 2024), p. 011007. DOI: 10.1103/PhysRevX.14.011007 (cited on pages 4, 23, 36, 63, 109, 116, 122, 134).

[Iva21] Edouard Ivanov. 'Optimization of silicon nitride membranes for hybrid superconducting-mechanical circuits'. These de doctorat. Sorbonne université, Feb. 2021 (cited on pages 5, 53, 55, 58, 60).

[Cap20] Thibault Capelle. 'Electromechanical cooling and parametric amplification of an ultrahigh-Q mechanical oscillator'. These de doctorat. Sorbonne université, Sept. 2020. (Visited on 02/05/2025) (cited on pages 5, 58).

[Sei+22] Yannick Seis et al. 'Ground state cooling of an ultracoherent electromechanical system'. en. In: *Nature Communications* 13.1 (Mar. 2022), p. 1507. DOI: 10.1038/s41467-022-29115-9 (cited on pages 5, 149).

[Pat25] Himanshu Patange. 'RF Transduction in a DC-Biased Superconducting Electromechanical System'. PhD thesis. Sorbonne University, 2025 (cited on pages 5, 53, 57–59, 62, 64, 65, 73, 88, 131).

[Som+25] Timo Sommer et al. 'Spatial Mapping of Intrinsic and Readout Nonlinearities in a Strongly Driven Micromechanical Membrane'. In: *Physical Review Letters* 134.20 (May 2025), p. 203604. DOI: 10.1103/PhysRevLett.134.203604 (cited on page 10).

[Bag+11] Mahmood Bagheri et al. 'Dynamic manipulation of nanomechanical resonators in the high-amplitude regime and non-volatile mechanical memory operation'. In: *Nature Nanotechnology* (Nov. 2011). DOI: 10.1038/nnano.2011.180 (cited on page 10).

[Asa+17] Keivan Asadi et al. 'Mechanism of geometric nonlinearity in a nonprismatic and heterogeneous microbeam resonator'. In: *Physical Review B* 96.11 (Sept. 2017), p. 115306. DOI: 10.1103/PhysRevB.96.115306 (cited on page 10).

[Jos62] B.D. Josephson. 'Possible new effects in superconductive tunnelling'. en. In: *Physics Letters* 1.7 (July 1962), pp. 251–253. DOI: 10.1016/0031-9163(62)91369-0 (cited on page 12).

[NPT99] Y. Nakamura, Yu. A. Pashkin, and J. S. Tsai. 'Coherent control of macroscopic quantum states in a single-Cooper-pair box'. In: *Nature* 398.6730 (Apr. 1999), pp. 786–788. DOI: 10.1038/19718 (cited on page 14).

[Ith+05] G. Ithier et al. 'Decoherence in a superconducting quantum bit circuit'. en. In: *Physical Review B* 72.13 (Oct. 2005), p. 134519. DOI: 10.1103/PhysRevB.72.134519 (cited on pages 14, 47, 89, 93).

[Koc+07] Jens Koch et al. 'Charge insensitive qubit design derived from the Cooper pair box'. en. In: *Physical Review A* 76.4 (Oct. 2007). arXiv:cond-mat/0703002, p. 042319. DOI: 10.1103/PhysRevA.76.042319 (cited on pages 14, 16, 26, 35, 51, 89, 93).

[Fer+06] A. J. Ferguson et al. 'Microsecond Resolution of Quasiparticle Tunneling in the Single-Cooper-Pair Transistor'. en. In: *Physical Review Letters* 97.10 (Sept. 2006), p. 106603. DOI: 10.1103/PhysRevLett.97.106603 (cited on page 14).

[Ris+13] D. Ristè et al. 'Millisecond charge-parity fluctuations and induced decoherence in a superconducting transmon qubit'. en. In: *Nature Communications* 4.1 (May 2013), p. 1913. DOI: 10.1038/ncomms2936 (cited on page 14).

[Sch+08] J. A. Schreier et al. 'Suppressing charge noise decoherence in superconducting charge qubits'. en. In: *Physical Review B* 77.18 (May 2008), p. 180502. DOI: 10.1103/PhysRevB.77.180502 (cited on page 16).

[Rig+12] Chad Rigetti et al. 'Superconducting qubit in waveguide cavity with coherence time approaching 0.1ms'. en. In: *Physical Review B* 86.10 (Sept. 2012). arXiv:1202.5533 [quant-ph]. DOI: 10.1103/PhysRevB.86.100506 (cited on page 23).

[Pla+21] Alexander P. M. Place et al. 'New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds'. en. In: *Nature Communications* 12.1 (Mar. 2021). DOI: 10.1038/s41467-021-22030-5 (cited on page 23).

[Mam+21] H. J. Mamin et al. 'Merged-Element Transmons: Design and Qubit Performance'. In: *Physical Review Applied* 16.2 (Aug. 2021). arXiv:2103.09163 [quant-ph]. DOI: 10.1103/PhysRevApplied.16.024023 (cited on page 23).

[Bla+25] Matthew P. Bland et al. '2D transmons with lifetimes and coherence times exceeding 1 millisecond'. en. In: *Preprint* arXiv:2503.14798 (Mar. 2025). arXiv:2503.14798 [quant-ph]. DOI: 10.48550/arXiv.2503.14798 (cited on page 23).

[LE96] C. K. Law and J. H. Eberly. 'Arbitrary Control of a Quantum Electromagnetic Field'. en. In: *Physical Review Letters* 76.7 (Feb. 1996), pp. 1055–1058. DOI: 10.1103/physrevlett.76.1055 (cited on page 26).

[Gam+06] Jay Gambetta et al. 'Qubit-photon interactions in a cavity: Measurement-induced dephasing and number splitting'. en. In: *Physical Review A* 74.4 (Oct. 2006), p. 042318. DOI: 10.1103/PhysRevA.74.042318 (cited on page 26).

[GC85] C. W. Gardiner and M. J. Collett. 'Input and output in damped quantum systems: Quantum stochastic differential equations and the master equation'. en. In: *Physical Review A* 31.6 (June 1985), pp. 3761–3774. DOI: 10.1103/PhysRevA.31.3761 (cited on page 27).

[Man+09] Vladimir E. Manucharyan et al. 'Fluxonium: single Cooper pair circuit free of charge offsets'. en. In: *Science* 326.5949 (Oct. 2009). arXiv:0906.0831 [cond-mat], pp. 113–116. DOI: 10.1126/science.1175552 (cited on pages 30, 31).

[Man12] Vladimir Eduardovich Manucharyan. 'Superinductance'. PhD thesis. Yale University, 2012 (cited on pages 31, 39).

[Mas+12] Nicholas A Masluk et al. 'Microwave Characterization of Josephson Junction Arrays: Implementing a Low Loss Superinductance'. In: *Physical review letters* 109.13 (2012), p. 137002 (cited on pages 31, 108).

[Vak+24] T. Vakhtel et al. 'Tunneling of fluxons via a Josephson resonant level'. en. In: *Physical Review B* 110.4 (July 2024), p. 045404. DOI: 10.1103/PhysRevB.110.045404 (cited on pages 35, 110).

[Ard+24] Waël Ardati et al. 'Using Bifluxon Tunneling to Protect the Fluxonium Qubit'. In: *Phys. Rev. X* 14 (4 Oct. 2024), p. 041014. DOI: 10.1103/PhysRevX.14.041014 (cited on pages 35, 110).

[Sch+02] R. J. Schoelkopf et al. 'Qubits as Spectrometers of Quantum Noise'. en. In: *NATO Science Series* arXiv:cond-mat/0210247 (Oct. 2002). arXiv:cond-mat/0210247 (cited on pages 43, 44).

[Cle14] Aashish A. Clerk. 'Quantum noise and quantum measurement'. In: *Quantum noise and quantum measurement*. 2014 (cited on page 43).

[Cle+10] A. A. Clerk et al. 'Introduction to Quantum Noise, Measurement and Amplification'. en. In: *Reviews of Modern Physics* 82.2 (Apr. 2010). arXiv:0810.4729 [cond-mat, physics:quant-ph], pp. 1155–1208. DOI: 10.1103/RevModPhys.82.1155 (cited on pages 44, 147, 148).

[Wat99] Dan M. Watson. *Astronomy 203/403. Lecture Notes*. University of Rochester, 1999 (cited on page 44).

[Mel+20] A. Melville et al. 'Comparison of dielectric loss in titanium nitride and aluminum superconducting resonators'. en. In: *Applied Physics Letters* 117.12 (Sept. 2020), p. 124004. DOI: 10.1063/5.0021950 (cited on pages 45, 84).

[Pop+14] Ioan M. Pop et al. 'Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles'. en. In: *Nature* 508.7496 (Apr. 2014), pp. 369–372. DOI: 10.1038/nature13017 (cited on pages 45, 46, 48, 49, 89).

[AC84] Y. Aharonov and A. Casher. 'Topological Quantum Effects for Neutral Particles'. en. In: *Physical Review Letters* 53.4 (July 1984), pp. 319–321. DOI: 10.1103/PhysRevLett.53.319 (cited on page 47).

[Man+12] Vladimir E. Manucharyan et al. 'Evidence for coherent quantum phase-slips across a Josephson junction array'. en. In: *Physical Review B* 85.2 (Jan. 2012). arXiv:1012.1928 [cond-mat, physics:quant-ph], p. 024521. DOI: 10.1103/PhysRevB.85.024521 (cited on page 47).

[Ran+24] Mallika T. Randeria et al. 'Dephasing in Fluxonium Qubits from Coherent Quantum Phase Slips'. en. In: *arXiv* arXiv:2404.02989 (Apr. 2024). arXiv:2404.02989 [cond-mat, physics:quant-ph] (cited on page 47).

[Cat+11] G. Catelani et al. 'Relaxation and frequency shifts induced by quasiparticles in superconducting qubits'. en. In: *Physical Review B* 84.6 (Aug. 2011). arXiv:1106.0829 [cond-mat, physics:quant-ph], p. 064517. DOI: 10.1103/PhysRevB.84.064517 (cited on pages 48, 49, 89).

[Ste21] Jeremy Stevens. 'Effect of the Environment on Fluxonium Qubits and Thermodynamics of Quantum Measurement'. PhD thesis. Ecole Normale Supérieure de Lyon, 2021 (cited on pages 48, 49).

[PTP46] E. M. Purcell, H. C. Torrey, and R. V. Pound. 'Resonance Absorption by Nuclear Magnetic Moments in a Solid'. en. In: *Physical Review* 69.1–2 (Jan. 1946), pp. 37–38. DOI: 10.1103/PhysRev.69.37 (cited on page 51).

[SGK14] Eyob A. Sete, Jay M. Gambetta, and Alexander N. Korotkov. 'Purcell effect with microwave drive: Suppression of qubit relaxation rate'. In: *Physical Review B* 89.10 (Mar. 2014). arXiv:1401.5545 [quant-ph], p. 104516. DOI: 10.1103/PhysRevB.89.104516 (cited on page 51).

[Bla+04] Alexandre Blais et al. 'Cavity quantum electrodynamics for superconducting electrical circuits: an architecture for quantum computation'. In: *Physical Review A* 69.6 (June 2004). arXiv:cond-mat/0402216, p. 062320. DOI: 10.1103/PhysRevA.69.062320 (cited on page 51).

[Hou+08] A. A. Houck et al. 'Controlling the Spontaneous Emission of a Superconducting Transmon Qubit'. en. In: *Physical Review Letters* 101.8 (Aug. 2008), p. 080502. DOI: 10.1103/PhysRevLett.101.080502 (cited on page 51).

[Gro+18] Peter Groszkowski et al. 'Coherence properties of the 0-π qubit'. In: *New Journal of Physics* 20.4 (Apr. 2018), p. 043053. DOI: 10.1088/1367-2630/aab7cd (cited on page 51).

[Cup+24] Andrea Cupertino et al. 'Centimeter-scale nanomechanical resonators with low dissipation'. en. In: *Nature Communications* 15.1 (May 2024). arXiv:2308.00611 [cond-mat], p. 4255. DOI: 10.1038/s41467-024-48183-7 (cited on page 53).

[Ber+22] Mohammad J. Bereyhi et al. 'Perimeter Modes of Nanomechanical Resonators Exhibit Quality Factors Exceeding 10 9 at Room Temperature'. In: *Physical Review X* 12.2 (May 2022), p. 021036. DOI: 10.1103/PhysRevX.12.021036 (cited on page 53).

[Fed+19] S. A. Fedorov et al. 'Generalized dissipation dilution in strained mechanical resonators'. In: *Physical Review B* 99.5 (Feb. 2019). Publisher: American Physical Society, p. 054107. DOI: 10.1103/PhysRevB.99.054107. (Visited on 02/05/2025) (cited on pages 53, 56).

[YCS15] Mingyun Yuan, Martijn A. Cohen, and Gary A. Steele. 'Silicon nitride membrane resonators at millikelvin temperatures with quality factors exceeding 108'. In: *Applied Physics Letters* 107.26 (Dec. 2015), p. 263501. DOI: 10.1063/1.4938747. (Visited on 01/30/2025) (cited on page 53).

[Ser+18] Enrico Serra et al. 'Silicon Nitride MOMS Oscillator for Room Temperature Quantum Optomechanics'. In: *Journal of Microelectromechanical Systems* 27.6 (2018), pp. 1193–1203. DOI: 10.1109/JMEMS.2018.2876593. (Visited on 02/05/2025) (cited on page 53).

[Fau+12] T. Faust et al. 'Microwave cavity-enhanced transduction for plug and play nanomechanics at room temperature'. en. In: *Nature Communications* 3.1 (Mar. 2012). Publisher: Nature Publishing Group, p. 728. DOI: 10.1038/ncomms1723. (Visited on 02/05/2025) (cited on page 53).

[Kum+21] Sumit Kumar et al. *Microwave optomechanical measurement of non-metallized SiN strings at mK temperatures*. arXiv:2110.00228 [cond-mat]. Oct. 2021. DOI: 10.48550/arXiv.2110.00228. URL: http://arxiv.org/abs/2110.00228 (visited on 02/12/2025) (cited on page 53).

[Bak+12] Christophe Baker et al. 'Optical instability and self-pulsing in silicon nitride whispering gallery resonators'. EN. In: *Optics Express* 20.27 (Dec. 2012). Publisher: Optica Publishing Group, pp. 29076–29089. DOI: 10.1364/OE.20.029076. (Visited on 02/05/2025) (cited on page 53).

[Rei+16] Christoph Reinhardt et al. 'Ultralow-Noise SiN Trampoline Resonators for Sensing and Optomechanics'. In: *Physical Review X* 6.2 (Apr. 2016). Publisher: American Physical Society, p. 021001. DOI: 10.1103/PhysRevX.6.021001. (Visited on 02/05/2025) (cited on page 53).

[Sch+16] Cornelia Schwarz et al. 'Deviation from the Normal Mode Expansion in a Coupled Graphene-Nanomechanical System'. In: *Physical Review Applied* 6.6 (Dec. 2016). Publisher: American Physical Society, p. 064021. DOI: 10.1103/PhysRevApplied.6.064021. (Visited on 02/05/2025) (cited on page 53).

[Dav+13] Marcelo Davanço et al. 'Si3N4 nanobeam optomechanical crystals'. EN. In: *Optica Publishing Group* (June 2013), CW3F.2. DOI: 10.1364/CLEO_SI.2013.CW3F.2. (Visited on 02/05/2025) (cited on page 53).

[Ree+19] C. Reetz et al. 'Analysis of Membrane Phononic Crystals with Wide Band Gaps and Low-Mass Defects'. In: *Physical Review Applied* 12.4 (Oct. 2019). Publisher: American Physical Society, p. 044027. DOI: 10.1103/PhysRevApplied.12.044027. (Visited on 02/05/2025) (cited on page 53).

[Lei69] A. W. Leissa. *Vibration of plates*. NTRS Author Affiliations: NASA Headquarters NTRS Report/Patent Number: NASA-SP-160 NTRS Document ID: 19700009156 NTRS Research Center: Legacy CDMS (CDMS). Jan. 1969. URL: https://ntrs.nasa.gov/citations/19700009156 (visited on 09/04/2024) (cited on page 55).

[YPR12] P.-L. Yu, T. P. Purdy, and C. A. Regal. 'Control of Material Damping in High- Q Membrane Microresonators'. In: *Physical Review Letters* 108.8 (Feb. 2012), p. 083603. DOI: 10.1103/PhysRevLett.108.083603 (cited on pages 55, 131).

[Neu+24] Leonhard Neuhaus et al. 'Python Red Pitaya Lockbox (PyRPL): An open source software package for digital feedback control in quantum optics experiments'. In: *Review of Scientific Instruments* 95.3 (Mar. 2024), p. 033003. DOI: 10.1063/5.0178481. (Visited on 02/05/2025) (cited on page 61).

[SMK15] Eyob A. Sete, John M. Martinis, and Alexander N. Korotkov. 'Quantum theory of a bandpass Purcell filter for qubit readout'. en. In: *Physical Review A* 92.1 (July 2015). arXiv:1504.06030 [quant-ph], p. 012325. DOI: 10.1103/PhysRevA.92.012325 (cited on page 75).

[Wal+17] T. Walter et al. 'Rapid High-Fidelity Single-Shot Dispersive Readout of Superconducting Qubits'. en. In: *Physical Review Applied* 7.5 (May 2017), p. 054020. DOI: 10.1103/PhysRevApplied.7.054020 (cited on pages 75, 115).

[Dol77] G. J. Dolan. 'Offset masks for lift-off photoprocessing'. en. In: *Applied Physics Letters* 31.5 (Sept. 1977), pp. 337–339. DOI: 10.1063/1.89690 (cited on page 76).

[Pot+01] A. Potts et al. 'CMOS compatible fabrication methods for submicron Josephson junction qubits'. In: *IEE Proceedings - Science, Measurement and Technology* 148.5 (Sept. 2001), pp. 225–228. DOI: 10.1049/ip-smt:20010395 (cited on page 77).

[Man+25] Paul Manset et al. 'Hyperinductance based on stacked Josephson junctions'. In: *Preprint* (2025). DOI: 10.48550/ARXIV.2505.02764 (cited on page 77).

[San24] Baldo Luis Najera Santos. 'Radio-frequency fluxonium superconducting qubit for AC-charge sensing applications'. PhD thesis. Sorbonne University. Laboratoire Kastler Brossel, 2024 (cited on page 83).

[Qui+14] C. M. Quintana et al. 'Characterization and reduction of microfabrication-induced decoherence in superconducting quantum circuits'. en. In: *Applied Physics Letters* 105.6 (Aug. 2014). arXiv:1407.4769 [cond-mat, physics:quant-ph], p. 062601. DOI: 10.1063/1.4893297 (cited on page 86).

[Dun+17] A. Dunsworth et al. 'Characterization and Reduction of Capacitive Loss Induced by Sub-Micron Josephson Junction Fabrication in Superconducting Qubits'. In: *Applied Physics Letters* 111.2 (July 2017). arXiv:1706.00879 [quant-ph], p. 022601. DOI: 10.1063/1.4993577 (cited on page 86).

[AB63] Vinay Ambegaokar and Alexis Baratoff. 'Tunneling Between Superconductors'. en. In: *Physical Review Letters* 10.11 (June 1963), pp. 486–489. DOI: 10.1103/PhysRevLett.10.486 (cited on page 87).

[Kri+19] Sebastian Krinner et al. 'Engineering cryogenic setups for 100-qubit scale superconducting circuit systems'. en. In: *EPJ Quantum Technology* 6.1 (Dec. 2019). arXiv:1806.07862 [cond-mat, physics:quant-ph], p. 2. DOI: 10.1140/epjqt/s40507-019-0072-0 (cited on page 89).

[Yeh+17] Jen-Hao Yeh et al. 'Microwave attenuators for use with quantum devices below 100 mK'. en. In: *Journal of Applied Physics* 121.22 (June 2017), p. 224501. DOI: 10.1063/1.4984894 (cited on page 89).

[Yan+16] F. Yan et al. 'The Flux Qubit Revisited to Enhance Coherence and Reproducibility'. en. In: *Nature Communications* 7.1 (Nov. 2016). arXiv:1508.06299 [quant-ph], p. 12964. DOI: 10.1038/ncomms12964 (cited on page 89).

[Bar+11] R. Barends et al. 'Minimizing quasiparticle generation from stray infrared light in superconducting quantum circuits'. In: *Applied Physics Letters* 99.11 (Sept. 2011), p. 113507. DOI: 10.1063/1.3638063 (cited on page 89).

[Cor+11] Antonio D. Corcoles et al. 'Protecting superconducting qubits from external sources of loss and heat'. en. In: *Applied Physics Letters* 99.18 (Oct. 2011). arXiv:1108.1383 [cond-mat, physics:quant-ph], p. 181906. DOI: 10.1063/1.3658630 (cited on pages 89, 93).

[Kre+16] J M Kreikebaum et al. 'Optimization of infrared and magnetic shielding of superconducting TiN and Al coplanar microwave resonators'. en. In: *Superconductor Science and Technology* 29.10 (Oct. 2016), p. 104002. DOI: 10.1088/0953-2048/29/10/104002 (cited on pages 89, 94, 97).

[Ser+19] K. Serniak et al. 'Direct Dispersive Monitoring of Charge Parity in Offset-Charge-Sensitive Transmons'. en. In: *Physical Review Applied* 12.1 (July 2019). arXiv:1903.00113 [cond-mat, physics:quant-ph], p. 014052. DOI: 10.1103/PhysRevApplied.12.014052 (cited on page 93).

[Gee23] Kurtis Lee Geerlings. 'Improving Coherence of Superconducting Qubits and Resonators'. PhD thesis. Yale University, 2023 (cited on page 94).

[SL08] Tim Schenk and Jean-Paul Linnartz. *RF imperfections in high-rate wireless systems: impact and digital compensation*. eng. Dordrecht: Springer, 2008 (cited on page 96).

[Sab98] S. Sabah. 'Design and Calibration of IQ-Mixers'. In: 1998 (cited on page 96).

[Mic15] Marki Microwave. *5 Ways to Compensate for Passive IQ Mixer Imbalance*. Online: accessed 23. Jul. 2015. 2015. URL: https://markimicrowave.com/technical-resources/application-notes/5-ways-to-compensate-for-passive-iq-mixer-imbalance/ (cited on page 96).

[Vil23] Marius Villiers. 'Dynamically Enhancing Qubit-Photon Interactions with Anti-Squeezing'. PhD thesis. Sorbonne University, 2023 (cited on page 97).

[NP14] I. Nsanzineza and B. L. T. Plourde. 'Trapping a single vortex and reducing quasiparticles in a superconducting resonator'. en. In: *Physical Review Letters* 113.11 (Sept. 2014). arXiv:1405.0256 [cond-mat, physics:quant-ph], p. 117002. DOI: 10.1103/PhysRevLett.113.117002 (cited on page 97).

[Son+09] C. Song et al. 'Microwave response of vortices in superconducting thin films of Re and Al'. en. In: *Physical Review B* 79.17 (May 2009), p. 174512. DOI: 10.1103/PhysRevB.79.174512 (cited on page 97).

[Ger+25] Kyrylo Gerashchenko et al. 'Probing the quantum motion of a macroscopic mechanical oscillator with a radio-frequency superconducting qubit'. In: *Preprint* (2025). DOI: 10.48550/ARXIV.2505.21481 (cited on pages 103, 135, 145, 147).

[Reg24] Ulysse Reglade. 'Quantum control of a cat-qubit with macroscopic bit-flip times'. PhD thesis. 2024 (cited on page 104).

[Zhu+13] Guanyu Zhu et al. 'Circuit QED with fluxonium qubits: Theory of the dispersive regime'. en. In: *Physical Review B* 87.2 (Jan. 2013), p. 024510. DOI: 10.1103/PhysRevB.87.024510 (cited on page 105).

[GK21] Peter Groszkowski and Jens Koch. 'Scqubits: a Python package for superconducting qubits'. In: *Quantum* 5 (Nov. 2021). arXiv:2107.08552 [quant-ph], p. 583. DOI: 10.22331/q-2021-11-17-583 (cited on page 107).

[VC15] Giovanni Viola and Gianluigi Catelani. 'Collective modes in the fluxonium qubit'. In: *Physical Review B* 92.22 (2015), p. 224511. DOI: 10.1103/PhysRevB.92.224511 (cited on page 108).

[UWK09] Quirin P Unterreithmeier, Eva M Weig, and Jörg P Kotthaus. 'Universal transduction scheme for nanomechanical systems based on dielectric forces'. In: *Nature* 458.7241 (2009), pp. 1001–1004. DOI: 10.1038/nature07932 (cited on page 130).

[SLL12] S. M. Sze, Ming-Kwei Lee, and M. K. Lee. *Semiconductor devices, physics and technology*. 3. ed. Wiley, 2012 (cited on page 131).

[Cuj+19] KS Cujia et al. 'Tracking the precession of single nuclear spins by weak measurements'. In: *Nature* 571.7764 (2019), pp. 230–233. DOI: 10.1038/s41586-019-1334-9 (cited on page 134).

[Arc+06] O. Arcizet et al. 'Radiation-pressure cooling and optomechanical instability of a micro-mirror'. In: *Nature* 444.7115 (2006), pp. 71–74 (cited on page 141).

[Sch+06] A. Schliesser et al. 'Radiation Pressure Cooling of a Micromechanical Oscillator Using Dynamical Backaction'. In: *Physical Review Letters* 97.24 (Dec. 2006), p. 243905 (cited on page 141).

[Rou22] Pierre Rouchon. 'A tutorial introduction to quantum stochastic master equations based on the qubit/photon system'. In: *Annual Reviews in Control* 54 (2022), pp. 252–261 (cited on page 142).

[AKM14] Markus Aspelmeyer, Tobias J Kippenberg, and Florian Marquardt. 'Cavity optomechanics'. In: *Reviews of Modern Physics* 86.4 (2014), pp. 1391–1452 (cited on page 143).

[Gle+07] Sébastien Gleyzes et al. 'Quantum jumps of light recording the birth and death of a photon in a cavity'. en. In: *Nature* 446.7133 (Mar. 2007), pp. 297–300. DOI: 10.1038/nature05589 (cited on page 148).

[Del+08] Samuel Deléglise et al. 'Reconstruction of non-classical cavity field states with snapshots of their decoherence'. en. In: *Nature* 455.7212 (Sept. 2008), pp. 510–514. DOI: 10.1038/nature07288 (cited on page 148).

[Say+11] Clément Sayrin et al. 'Real-time quantum feedback prepares and stabilizes photon number states'. In: *Nature* 477.7362 (2011), pp. 73–77 (cited on page 148).