



High-Fidelity Control and Stabilization of Cat Qubits

Ronan Gautier

► To cite this version:

Ronan Gautier. High-Fidelity Control and Stabilization of Cat Qubits. Quantum Physics [quant-ph]. Inria Paris, 2023. English. NNT: . tel-04362128

HAL Id: tel-04362128

<https://hal.science/tel-04362128>

Submitted on 22 Dec 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Public Domain

SORBONNE UNIVERSITÉ

DOCTORAL THESIS

High-Fidelity Control and Stabilization of Cat Qubits

Ronan GAUTIER

<i>Reviewers:</i>	Liang JIANG	- University of Chicago
	Clément PELLEGRINI	- Université Paul Sabatier
<i>Examinators:</i>	Christiane KOCH	- Freie Universität Berlin
	Jean-Michel RAIMOND	- Sorbonne Université
	Mario SIGALOTTI	- Inria Paris & Sorbonne Université
<i>Director:</i>	Alain SARLETTE	- Inria Paris & Ghent University
<i>Invited:</i>	Mazyar MIRRAHIMI	- Inria Paris
	Alexandre BLAIS	- Université de Sherbrooke

Abstract

The development of quantum computing relies on the unpleasant predicament that a quantum system should be freely controllable but also very long-lived, two often conflicting requirements. A promising path to overcome this difficulty lies in continuous-variable qubits encoded in Schrödinger cat states of a quantum harmonic oscillator. These so-called cat qubits, once stabilized with engineered two-photon dissipation or Kerr Hamiltonians, can exhibit an exponentially small sensitivity to bit-flip noise channels. This bias in noise is in turn highly beneficial for quantum error correction since it reduces experimental overheads by requiring that only phase-flip errors should be corrected for. However, towards the realization of error-corrected cat qubits, several orders of magnitudes in gate and readout fidelities are still required.

In this thesis, using precise mathematical analyses of the dynamics of cat qubits, we identify opportunities to improve their control and to achieve orders of magnitude improvements in error protection or gate fidelities. First, we study in detail the energy spectrum of Kerr cat qubits, and show how it limits the scaling of protection against bit-flip errors in the presence of thermal-induced leakage. From this study, a novel cat qubit confinement scheme is devised based on the combination of engineered dissipation and of a two-photon exchange Hamiltonian. We show how this scheme outperforms existing proposals in terms of gate fidelities while retaining the key exponential suppression of errors. Second, we focus on dissipative Zeno gates and survey the role of the buffer mode into the gate process. From this insight, we devise several designs of high-fidelity Zeno gates, either based on feeding information leaked to the environment back into the system of interest, or based on reducing the amount of leaked information in the first place. Finally, we introduce a numerical method to differentiate through a Lindblad master equation with a constant memory cost, based on the reverse time integration of an adjoint state. We then demonstrate it for the optimal control of transmon readout, and show readout fidelity improvements with interpretable pulses. The method is completely generic, has a low numerical overhead, and can be applied to a range of problems where dissipation is key to the process being optimized.

Throughout this thesis, a particular attention has been devoted to proposing schemes that are experimentally viable with the toolbox of modern superconducting circuits. In addition, more than providing concrete proposals to improve upon existing cat qubit operations and stabilization methods, we hope that this thesis can provide new insights into research on cat qubits, namely regarding their stabilization and the role of the buffer mode into the bigger picture.

Acknowledgments

Through the completion of this thesis, I have had the privilege of being guided and mentored by three outstanding individuals. I wish to extend my deepest gratitude to Alain Sarlette, Mazyar Mirrahimi, and Alexandre Blais for their supervision, patience, and unending generosity. Each in their own way, they have contributed to the depth, breadth, and richness of this scientific journey. Alain in particular has been instrumental. Beyond his wealth of research ideas, he always kept his door open for hour-long whiteboard discussions that quickly became indispensable. His humanity and unfailing tranquility are invaluable qualities. Mazyar has also been a genuine pleasure to collaborate with. More than his friendliness, he is an exceptional scientist from whom I have learned immensely in every aspect of my research, and whose advice and mentorship I truly cherish. Lastly, Alexandre, thanks to his scientific virtuosity, kindness, and distinct management acumen, has created fertile ground in Sherbrooke for top-tier research. My time in his group, though brief, has proven invaluable. I eagerly look forward to further collaborating with the three of them.

For having accepted to read and review this thesis, I would like to thank all members of the jury: Liang Jiang, Clément Pellegrini, Christiane Koch, Jean-Michel Raimond and Mario Sigalotti. Their time and effort honors me.

I would also like to thank all the colleagues I have collaborated with during this thesis. Jérémie Guillaud, both for his early guidance in following through the research he initiated and for the trust he now grants me. Lev-Arcady Sellem, for his invaluable help in solving the many mathematical dilemmas I was faced with. Diego Ruiz, for his endless curiosity and the countless ideas he was kind enough to share with me. Pierre Guilmin, for the thoroughness and enthusiasm he carries with him that made our collaboration so simple. Élie Génois, for having been a brilliant colleague but also a great friend in my time overseas. Adrien Bocquet and Ulysse Réglade, the dynamic duo, for allowing me to glimpse into experimental physics. And Ross Shillito, for the wild yet subtle ideas he brought forward that make him such a pleasure to work with. Together, they not only helped me grow as a scientist, but also made the process truly enjoyable.

Within the Quantic team, there are many others whom I wish to thank. First, all the permanent members of the group: Antoine Tilloy, Pierre Rouchon, Alexandru Petrescu, Zaki Leghtas and Philippe Campagne-Ibarcq. Together, they share an immense knowledge of physics and mathematics that I have greatly benefited from at one point or another. It is difficult not to find answers to most research questions when surrounded by such a brilliant group of scientists. Then, all the students, post-docs and interns that have been

part of the team: Michiel, Christian, François-Marie, Léon, Thiziri, Louis, Angela, Linda, Karan, Rémi, Marius, Matthieu and all the others. I will miss our lunch time discussions, evenings around a few beers and the delightful atmosphere you brought to the group. Derya, with her unmatched kindness, has also been of vital help in navigating the administration in these past few years.

Then, I would like to extend my thanks to all members of the Institut Quantique who welcomed me in Sherbrooke: Adrian, Alexander, Benjamin and Benjamin, Camille, Catherine, Christian, Cristobal, Lautaro, Manuel, Othmane, and of course Marie Frédérique. With them, my time overseas flew by all too quickly. I will keep fond memories of our time at the pub, of our climbing sessions, volleyball games and impromptu ski trips.

I also want to thank the Alice & Bob team for infusing this thesis with an exciting atmosphere and for generously sharing their fast-paced research with us. Not having them around for the March Meeting trips would also have been a lot less fun. In particular, I want to thank Raphaël and Joachim for the many stimulating discussions we had about physics.

Finally, I cannot imagine having gone through this thesis without the support of my friends and family. I already have a lifetime of memories with them, and hope for many more. From my earliest days, my parents nurtured my curiosity, allowing me to grow without boundaries but with the deepest passion they could share. For this, I will never thank them enough. Lastly, I owe a profound debt to she whose unwavering support, love, and kindness have been the foundation of this journey. Her faith in me, even in my most anxious moments, has been my guiding light. Anna, this thesis, much like the rest of my days, I dedicate to you.

Contents

List of Acronyms	ix
1 From superconducting circuits to protected cat qubits	1
1.1 Quantum computation	1
1.2 Superconducting quantum circuits	4
1.2.1 Quantum LC resonators	4
1.2.2 Artificial atoms	6
1.2.3 Light-matter interaction	9
1.2.4 The bosonic shift in paradigm	10
1.3 Cat qubits	12
1.3.1 Encoding a cat qubit in an oscillator	12
1.3.2 Shifted Fock basis	16
1.3.3 Concatenating into error-correcting codes	18
1.4 Protecting cat qubits from decoherence and leakage	19
1.4.1 Hamiltonian protection: Kerr cat qubits	20
1.4.1.1 Kerr cat Hamiltonian	20
1.4.1.2 Engineering with a SNAIL	21
1.4.1.3 Semiclassical analysis	22
1.4.2 Dissipative protection: dissipative cat qubits	24
1.4.2.1 Two-photon dissipation	24
1.4.2.2 Reservoir engineering with a buffer mode	25
1.4.2.3 Engineering with an ATS	27
1.4.2.4 Semiclassical analysis	28
1.4.3 Engineering gates on protected qubits	30
1.4.3.1 Pauli X	32
1.4.3.2 $Z(\theta)$ rotations	32
1.4.3.3 CNOT and Toffoli	34
1.4.3.4 Gate errors	35
1.5 Outline and Contributions	35
2 Combined confinement of cat qubits	39
2.1 Introduction	39
2.2 Cat confinement and bit-flip suppression	41
2.2.1 Confinement and decoherence	41
2.2.2 Bit-flip suppression	43
2.2.3 Estimating bit-flip errors for an idling qubit	46
2.2.3.1 General Framework	46
2.2.3.2 Kerr Confinement	48

2.2.4	Combined Kerr and two-photon dissipation confinement	50
2.3	Combined confinement with two-photon exchange	52
2.3.1	The two-photon exchange Hamiltonian	52
2.3.2	Combined TPE and two-photon dissipation	54
2.4	Engineering gates under combined confinement	58
2.4.1	Zeno Z gate	58
2.4.1.1	Master equation	59
2.4.1.2	Gate errors	60
2.4.1.3	Analytical derivation of gate errors	61
2.4.1.4	Numerical analysis	63
2.4.2	CNOT gate	66
2.4.2.1	Gate engineering through code deformation	66
2.4.2.2	Target stabilization-free gate protocol	68
2.4.2.3	Numerical analysis	69
2.5	Towards experimental realization	72
2.5.1	Superconducting circuit implementation	72
2.5.2	Spurious terms	75
2.5.3	Circuit Hamiltonian derivation	77
2.6	Conclusion	81
3	Designing high-fidelity Zeno gates	83
3.1	Introduction	84
3.2	Zeno gate errors	85
3.3	Summary of gate designs	88
3.3.1	Buffer photodetection with classical feedback	89
3.3.2	Cat-buffer autonomous feedback	89
3.3.3	Locally flat Hamiltonian	90
3.3.4	Discrete jump	90
3.3.5	Combined dissipation and two-photon exchange Hamil- tonian	91
3.3.6	Combining designs	91
3.4	Gate designs	91
3.4.1	Buffer photodetection with classical feedback	91
3.4.1.1	Design principle	91
3.4.1.2	Measurement strategy	93
3.4.1.3	Jump and no-jump trajectories	94
3.4.1.4	Design performance	96
3.4.1.5	Non-ideal photodetector	97
3.4.1.6	Ideal feedback action	99
3.4.2	Cat-buffer autonomous feedback	100
3.4.2.1	Design principle	100
3.4.2.2	Parity-switching dynamics	102

3.4.2.3	Design performance	102
3.4.2.4	Multi-qubit gates	104
3.4.2.5	Adiabatic elimination of the buffer mode . . .	104
3.4.3	Locally flat Hamiltonian	105
3.4.3.1	Intuition	105
3.4.3.2	Design principle	106
3.4.3.3	Results	108
3.4.3.4	Multi-qubit gates	108
3.4.3.5	Engineering with an ATS	109
3.4.3.6	Optimal drive Hamiltonians	111
3.4.4	Discrete jump	112
3.4.4.1	Design principle	112
3.4.4.2	Qutrit design	114
3.4.4.3	Results	114
3.4.4.4	Multi-qubit gates	117
3.5	Robustness to noise	119
3.5.1	On the impact of noise and spurious Hamiltonians . . .	119
3.5.2	Thermal noise in the buffer mode	121
3.5.2.1	Photodetection	121
3.5.2.2	Autonomous feedback	122
3.6	Conclusion	123
4	Optimal control of transmon readout	125
4.1	Introduction	125
4.2	Adjoint state quantum optimal control	127
4.2.1	Automatic differentiation	128
4.2.2	Adjoint state method	130
4.2.3	Derivation of adjoint state method equations	133
4.2.3.1	Adjoint state master equation	133
4.2.3.2	Gradients	134
4.3	Optimal control of transmon readout	135
4.3.1	Transmon-resonator-filter model	135
4.3.2	Dispersive readout metrics	137
4.3.3	Optimizing readout	138
4.3.4	Results	139
4.4	Conclusion	143
5	Conclusions and perspectives	145

Appendices	149
A Invariance under joint phase conjugation	151
A.1 Phase conjugation superoperator	151
A.2 Master equation invariance	152
Bibliography	155

List of Acronyms

- AC** Alternative Current. 79
- ATS** Asymmetrically Threaded SQUID. 22, 27, 28, 73, 74, 77, 78, 110, 145
- AWG** Arbitrary Waveform Generator. 125, 138
- DC** Direct Current. 6, 73, 78
- GKP** Gottesman-Kitaev-Preskill. 11, 147
- GRAPE** Gradient Ascent Pulse Engineering. 126
- QND** Quantum Non Demolition. 142
- RF** Radiofrequency. 74, 75, 79
- RWA** Rotating Wave Approximation. 20, 28, 74, 75, 79, 136
- SFB** Shifted Fock Basis. 16, 17, 20, 23, 25, 29, 33, 35, 61, 71, 85, 86, 101, 121
- SNAIL** Superconducting Nonlinear Asymmetric Inductive eLement. 21, 22
- SNR** Signal-to-Noise Ratio. 137–142
- SQUID** Superconducting Quantum Interference Device. 27, 110
- TPE** Two-Photon Exchange. 41, 52–59, 61, 63–65, 68–76, 80, 81

From superconducting circuits to protected cat qubits

In this introductory chapter, we review the invention of protected cat qubits, from the early development of quantum physics to the growth of superconducting circuits and how they gave rise to continuous-variable quantum information.

Contents

1.1	Quantum computation	1
1.2	Superconducting quantum circuits	4
1.2.1	Quantum LC resonators	4
1.2.2	Artificial atoms	6
1.2.3	Light-matter interaction	9
1.2.4	The bosonic shift in paradigm	10
1.3	Cat qubits	12
1.3.1	Encoding a cat qubit in an oscillator	12
1.3.2	Shifted Fock basis	16
1.3.3	Concatenating into error-correcting codes	18
1.4	Protecting cat qubits from decoherence and leakage	19
1.4.1	Hamiltonian protection: Kerr cat qubits	20
1.4.2	Dissipative protection: dissipative cat qubits	24
1.4.3	Engineering gates on protected qubits	30
1.5	Outline and Contributions	35

1.1 Quantum computation

The study of physics is one of the most ancient and fundamental human endeavors, dating back to the earliest civilizations. From the first attempts to

understand the movements of celestial bodies to the modern development of particle accelerators, physics strives to uncover the mysteries of the universe and to comprehend the fundamental laws that govern it. Early classical mechanics provided the first glimpses into the workings of the physical world, but as our tools for observing the universe grew more advanced, so too did our understanding of the universe itself. The development of relativity theory and quantum mechanics in the early twentieth century marked a paradigm shift in our understanding of the physical world. These theories revealed a universe that was far more complex, subtle, and counter-intuitive than previously imagined, and they challenged our fundamental assumptions about reality itself. In particular, quantum mechanics posed a significant obstacle to our intuition regarding the universe's nature. The theory revealed that at the most fundamental level, the world is inherently probabilistic and that objects may simultaneously exist in several states. Moreover, the act of observing a quantum system can fundamentally alter its state, resulting in a profound change in our perception of the observer's role in shaping reality.

But quantum mechanics is not just a theoretical construct aimed at describing the nature of the universe; its insights have also given rise to numerous applications, ranging from improved sensors and imaging technology to cryptography and quantum communication. In particular, one exciting application of quantum mechanics is the development of quantum computers, first proposed by Richard Feynman in the early 1980s [Feynman 1982]. Feynman recognized that classical computers would struggle to simulate quantum systems, and proposed a device that was specifically designed to simulate these systems much more efficiently. Feynman's work inspired a wave of research into quantum computation, and in particular the development of quantum algorithms that are exponentially faster than their deterministic classical counterparts [Deutsch 1985, Shor 1994]. While a quantum computer that can outperform any classical computer on a useful application has yet to be demonstrated, the last decade has seen tremendous progress on this endeavor [Arute *et al.* 2019].

Such speed-ups in information processing directly arise from *entanglement*, a fundamental property of quantum mechanics that allows for the creation of complex, highly correlated states between two or more quantum systems. One of the most remarkable yet simple application of entanglement is quantum teleportation [Bennett *et al.* 1993], a quantum communication protocol that relies on the creation of an entangled state between two distant quantum objects. The protocol involves a sender and a receiver, who are located at different locations and who each possess one of the two entangled objects. The sender performs a measurement on their quantum object and sends the measurement result to the receiver via a classical communication channel. Depending on this information, the receiver decides whether to perform an

operation on their quantum object or not, which can effectively teleport the initial state of the sender's system to that of the receiver. Indeed, the sender's measurement in an appropriate basis will collapse the entangled system in a given state, unknown to the receiver. Hence, classical information is required for the receiver to know which state he is left with, without himself measuring and collapsing its own quantum object. Quantum teleportation thus harnesses the fundamental principles of quantum computation, ranging from the creation of intricate entangled states, to the conservation of these states while subject to noise during classical processing steps, and to the precise readout and manipulation of quantum objects with maximum accuracy.

This protocol not only finds practical application in modern quantum information processing, but also underlines a pivotal aspect of ongoing research in quantum computing: the demand for high-fidelity control and minimum decoherence [Preskill 1998, Knill 2005]. High-fidelity control pertains to the precision with which quantum states are manipulated and measured, and is vital for ensuring accurate and dependable quantum operations. Conversely, decoherence is a fundamental limitation of quantum systems, which occurs when a quantum state becomes entangled with its environment leading the eventual destruction of the quantum state. Minimizing decoherence is therefore critical for maintaining the integrity of quantum states and preventing errors in quantum computations [Pellizzari *et al.* 1995]. It is thus fundamentally challenging to operate a quantum device without introducing noise — i.e. any interaction with the environment that may prompt decoherence —, and doing so requires careful engineering and design to achieve balance between the two conflicting requirements.

Superconducting circuits are one of a few highly promising platforms that can achieve this intricate balance [Nakamura *et al.* 1999, Vion *et al.* 2002, Koch *et al.* 2007]. By harnessing the remarkable feature of superconductivity to conduct electric current without any resistance, these circuits can explore and manipulate quantum phenomena at macroscopic scales for extended periods of time. In addition, superconducting circuits can be easily fabricated using established microfabrication techniques, enabling the integration of complex circuitry on a chip. Their compatibility with classical electronics and well-developed control and measurement techniques also facilitates seamless interfacing with conventional computational systems, thus promoting the hybridization of quantum and classical computing. These factors, combined with the rapid advancements in materials science, device fabrication, and system integration, have pushed superconducting circuits to the forefront of research in quantum computation.

In the rest of this introduction, we first present a short review of superconducting circuits from the LC resonator to the non-linear control enabled by Josephson junctions. Next, we discuss how bosonic qubits have been giving

rise to a shift in paradigm in the field of quantum information by encoding in continuous-variable systems instead of discrete ones. We then move on to the introduction of cat qubits and to the promise of low-overhead quantum computation they offer. Finally, we review how cat qubits can be protected with engineered Hamiltonians or dissipation, and how to engineer a universal gate set with either of these two approaches.

1.2 Superconducting quantum circuits

The advent of superconducting circuits has opened up new possibilities for studying and manipulating quantum states at a macroscopic level, bridging the gap between microscopic quantum phenomena and the observable macroscopic world. After quantization into lumped-elements and into normal electromagnetic modes, these circuits exhibit discrete energy levels and coherent dynamics, characteristic of quantum systems. In this section, we introduce LC resonators as a means to confine and store energy in the form of quantized electromagnetic fields, transmon artificial atoms as robust and controllable two-level quantum systems, and their light-matter-like interaction. For an in-depth review, we refer to [Blais *et al.* 2021, Vool & Devoret 2017].

1.2.1 Quantum LC resonators

The quantum LC resonator plays a crucial role in superconducting circuits, providing a means to confine and store energy in the form of quantized electromagnetic fields. With lumped elements, an LC oscillator consists of an inductor with inductance L and a capacitor with capacitance C in parallel, as represented in Figure 1.1(a). In the classical regime, the dynamics of the resonator can be described by the Hamiltonian

$$H = \frac{Q^2}{2C} + \frac{\Phi^2}{2L}, \quad (1.1)$$

where $\Phi(t) = \int_0^t V(t')dt'$ is the magnetic flux related to the integral of the voltage across the dipole, $V(t)$. Its corresponding canonical momentum is the charge, $Q(t) = C\dot{\Phi}(t) = \int_0^t I(t')dt'$ given by the integral of the current, $I(t)$. This Hamiltonian directly highlights the mechanical analogy between the LC resonator and a mass-spring oscillator of position coordinate Φ , momentum Q and mass C , and where the two Hamiltonian terms correspond to kinetic and potential energies respectively.

To quantize the system, we promote the canonical variables Φ and Q to quantum operators, $\hat{\Phi}$ and \hat{Q} , respectively. These conjugate operators satisfy




Figure 1.1: (a) Lumped-element model of a LC resonator with inductance L and capacitance C . (b) Harmonic oscillator potential with resonant frequency ω_r . Figure adapted from [Blais *et al.* 2021].

the canonical commutation relation, $[\Phi, Q] = i\hbar$, that enforces the uncertainty principle, where \hbar is the reduced Planck constant. Next, we move into the diagonal basis of the Hamiltonian by introducing the annihilation and creation operators, \mathbf{a} and \mathbf{a}^\dagger respectively, which satisfy the commutation relation $[\mathbf{a}, \mathbf{a}^\dagger] = \hbar$. By making the substitution $\Phi = \Phi_{\text{zpf}}(\mathbf{a} + \mathbf{a}^\dagger)$ and $Q = iQ_{\text{zpf}}(\mathbf{a}^\dagger - \mathbf{a})$ into the Hamiltonian, we can rewrite it in terms of ladder operators according to

$$\mathbf{H} = \hbar\omega_r \left(\mathbf{a}^\dagger \mathbf{a} + \frac{1}{2} \right) \quad (1.2)$$

where $\omega_r = 1/\sqrt{LC}$ is the resonant frequency of the LC resonator. Here, $\Phi_{\text{zpf}} = \sqrt{\hbar Z_r/2}$ and $Q_{\text{zpf}} = \sqrt{\hbar/2Z_r}$ denote the characteristic magnitude of zero-point fluctuations of the flux and charge, where $Z_r = \sqrt{L/C}$ is the characteristic impedance of the circuit. This quantized Hamiltonian reveals the discrete nature of the energy spectrum, with the ladder operator $\mathbf{a}^\dagger \mathbf{a}$ representing the number operator that counts the excitation quanta or photons of frequency ω_r in the oscillator. The eigenstates of this Hamiltonian thus satisfy $\mathbf{a}^\dagger \mathbf{a} |n\rangle = n |n\rangle$ for $n \geq 0$, and are known as Fock or number states. We can thus identify the creation and annihilation operators \mathbf{a}^\dagger and \mathbf{a} as operators that induce transitions between energy levels, and satisfy the ladder relations, $\mathbf{a} |n\rangle = \sqrt{n} |n-1\rangle$ and $\mathbf{a}^\dagger |n\rangle = \sqrt{n+1} |n+1\rangle$. Figure 1.1(b) represents the harmonic potential described by (1.2) with equidistant energy levels.

Experimentally, LC resonators can be engineered using different geometries, such as 2D stripline resonators or 3D cavities, providing flexibility in the design and performance of superconducting circuits. 2D stripline resonators consist of a planar conductor sandwiched between two ground planes a few tens of micrometers apart, forming a transmission line structure. By carefully designing the dimensions and materials of the stripline, the inductance and capacitance can be controlled, allowing for precise tuning of the resonant frequency and impedance. On the other hand, 3D resonators utilize three-dimensional structures such as hollow microwave cavities made of alu-

minum to confine the electromagnetic field. These cavities provide additional degrees of freedom in shaping the field distribution and can be tailored for specific applications. The choice between 2D stripline resonators and 3D cavities depends on factors like desired frequency range, coupling strength to other components, scalability, and fabrication constraints. Both approaches offer unique advantages and have been extensively employed in the design and implementation of superconducting circuits for diverse quantum computing and quantum information processing tasks.

1.2.2 Artificial atoms

Non-linearity plays a pivotal role in probing non-classical behavior in superconducting circuits. Indeed, the ground state of an LC resonator is classical; it is a gaussian-like distribution of number states obtained after thermalization with the environment. As such, only other gaussian states can be prepared by using linear and quadratic Hamiltonians in the ladder operators which are the standard control tools of quantum optics. Hence the introduction of Josephson junctions [Josephson 1962] was crucial for the success of superconducting circuits by enabling large non-linearities to be engineered.

A Josephson junction typically consists of two superconducting electrodes separated by a thin insulating barrier as depicted in Figure 1.2(a), forming a weak link for the flow of Cooper pairs. The behavior of Josephson junctions arises from the Josephson effect, which is described by the phase difference φ between the superconducting order parameters on either side of the junction. The Josephson effect gives rise to two distinct phenomena: the Direct Current (DC) Josephson effect, where a supercurrent flows across the junction in the absence of a voltage bias, and the AC Josephson effect, where an oscillating supercurrent is induced by an applied voltage. The dynamics of Josephson junctions are elegantly described by the Josephson equations, which connect the phase difference φ across the junction to the supercurrent according to

$$I = I_c \sin(\varphi) \quad (1.3)$$

where I_c is the critical current of the junction. Voltage across the weak link is also related to the junction phase difference through the second Josephson relation, according to

$$\frac{d\varphi}{dt} = \frac{2\pi}{\Phi_0} V(t) \quad (1.4)$$

where $\Phi_0 = h/2e$ is the magnetic flux quantum. Together, these relations describe the relation between current and voltage across a junction, and hence the introduction of Josephson junctions into regular electronic circuits becomes evident [Vion *et al.* 2002, Wallraff *et al.* 2004].




Figure 1.2: (a) Schematic of a Josephson junction, with superconducting electrodes (light blue) separated by a thin insulating barrier (grey). (b) Lumped-element model of a fixed-frequency transmon, where the crossed-out square represents a junction with Josephson energy E_J , shunted by a large gate capacitance C_S . The superconducting island is shown in light blue. (c) Cosine potential of a transmon qubit with its lower three eigenstates of decreasing transition frequency ($|g\rangle$, $|e\rangle$ and $|f\rangle$). Panels (b) and (c) are adapted from [Blais *et al.* 2021].

Certainly the most prominent circuit of modern quantum devices is that of the transmon qubit [Koch *et al.* 2007, Blais *et al.* 2007, Clarke & Wilhelm 2008], as represented in Figure 1.2(b) with lumped elements. The transmon qubit architecture involves shunting the Josephson junction with a large capacitance, resulting in a highly anharmonic energy spectrum. This anharmonicity is crucial for achieving long coherence times and suppressing certain types of errors. The transmon Hamiltonian can be derived by considering the charge degree of freedom of the Josephson junction. The charge on the junction is described by the operator \mathbf{n} , which commutes with the Hamiltonian and represents the number of Cooper pairs transferred across the junction. The transmon Hamiltonian then reads

$$\mathbf{H} = 4E_C(\mathbf{n} - n_g)^2 - E_J \cos(\varphi) \quad (1.5)$$

where E_C is the charging energy associated with the capacitance, n_g is the offset charge induced by external sources, $E_J = \Phi_0 I_c / 2\pi$ is the Josephson energy related to the junction's critical current, and φ is the phase difference across the junction. The cosine potential of this Hamiltonian is represented in Figure 1.2(c). By providing a non-quadratic potential compared to regular oscillators, transmons feature unevenly spaced energy levels (i.e. non-zero anharmonicity) which allows for the precise control of isolated two-level systems.

One of the notable advantages of the transmon qubit is its remarkable resilience against charge noise. Charge noise can cause fluctuations in the offset charge n_g and thus poses a significant challenge in superconducting qubits. However, the transmon qubit is exponentially protected against charge noise thanks to its large shunting capacitance, which significantly reduces the charging energy $E_C \propto 1/C_S$ compared to the Josephson energy E_J , effectively

placing the transmon qubit in a "soft" charging regime. In this regime, the qubit states are spread out in charge space, making them less susceptible to charge noise-induced variations in n_g . As a result, transmon qubits exhibit enhanced coherence times on the order of tens of microseconds or even longer, contributing to their suitability for quantum information processing tasks. The downside of this regime is a reduction in qubit anharmonicity, which however scales only with a weak power law with the energy ratio according to $\alpha \sim (E_J/E_C)^{-1/2}$.

The experimental realization of transmon qubits has been achieved with great success, demonstrating their feasibility and robustness against charge noise [Schreier *et al.* 2008, Paik *et al.* 2011]. Moreover, transmon qubits have shown remarkable compatibility with established superconducting circuit technology, allowing for the integration of multiple qubits and the implementation of various quantum operations. They have been employed in a wide range of applications, including the realization of multi-qubit systems, quantum gates, and quantum error correction [Arute *et al.* 2019, AI 2023, Kim *et al.* 2023]. The combination of their anharmonic energy spectrum, long coherence times, and scalability positions transmon qubits as promising candidates for the development of large-scale, fault-tolerant quantum computers. Still, reaching this degree of performance requires further improvements in the control of such qubits. Chapter 4 is thus dedicated to the design of an optimal readout protocol for transmons using a novel method of quantum optimal control.

In addition to the widely studied transmon qubit, several other superconducting qubit architectures have gained attention for their unique characteristics and potential applications in quantum information processing. One such qubit is the fluxonium qubit [Manucharyan *et al.* 2009], which utilizes a superconducting loop shunted by a large inductance made of Josephson junctions in series, also referred to as superinductance. The fluxonium qubit operates in the regime where the Josephson energy dominates, resulting in a highly anharmonic energy spectrum. This anharmonicity enables long coherence times and robustness against certain types of noise, making the fluxonium qubit promising for applications requiring high fidelity quantum operations. Another intriguing qubit design is the zero-pi qubit [Douçot & Vidal 2002, Ioffe & Feigel'man 2002, Kitaev 2006], which utilizes Josephson junctions with unconventional phase relationships. This qubit exploits the presence of a pi-phase shift in the Josephson junction to create a non-trivial energy spectrum with protected states against noise and environmental fluctuations. The zero-pi qubit offers unique properties for quantum information processing, such as tunable qubit frequency and reduced sensitivity to charge noise. These alternative superconducting qubit architectures provide a rich landscape for exploring diverse quantum phenomena and expanding the capabilities of quantum technologies.




Figure 1.3: (a) Schematic of a transmon qubit (red) capacitively coupled to a 1D transmission-line resonator (blue). (b) Lumped-element circuit of a LC resonator capacitively coupled to a transmon. Figure adapted from [Blais *et al.* 2021].

1.2.3 Light-matter interaction

In superconducting quantum circuits, the interaction between transmon qubits and LC resonators plays a crucial role in enabling efficient information transfer and manipulation. This light-matter interaction [Haroche & Raimond 2006] arises from the coupling between the quantized electromagnetic field in the resonator and the discrete energy levels of the transmon qubit. In this section, we will explore the general intuition behind this interaction and provide a short mathematical derivation of the Hamiltonians involved.

The transmon having a relatively large size (a few hundreds of microns), it can be naturally coupled to LC resonators through capacitive coupling as represented on Figure 1.3. In this case, derivation of the full lumped-element circuit Hamiltonian leads to

$$\mathbf{H} = 4E_C(\mathbf{n} + \mathbf{n}_r)^2 - E_J \cos(\boldsymbol{\varphi}) + \hbar\omega_r \mathbf{a}^\dagger \mathbf{a} \quad (1.6)$$

where \mathbf{n} and $\boldsymbol{\varphi}$ are conjugate variables for the transmon, \mathbf{a} is the annihilation operator on the resonator, and $\mathbf{n}_r \propto i(\mathbf{a}^\dagger - \mathbf{a})$ is the charge operator on the resonator. Here, we have assumed that only a single mode of the resonator couples to the transmon, i.e. that of closest frequency to the transmon. Note here that compared to the Hamiltonian of the free transmon (1.5), the role of the offset charge n_g is played by the charge bias of the resonator which leads to the charge-charge coupling of both modes.

Introducing creation and annihilation operators for the transmon, \mathbf{q}^\dagger and \mathbf{q} respectively with $[\mathbf{q}, \mathbf{q}^\dagger] = \hbar$, it is possible to diagonalize the free transmon Hamiltonian and to obtain the following reduced model of the two-mode system

$$\mathbf{H}/\hbar \approx \omega_r \mathbf{a}^\dagger \mathbf{a} + \omega_q \mathbf{q}^\dagger \mathbf{q} - \frac{1}{2\hbar} E_C \mathbf{q}^{\dagger 2} \mathbf{q}^2 - g(\mathbf{q}^\dagger - \mathbf{q})(\mathbf{a}^\dagger - \mathbf{a}) \quad (1.7)$$

where the first two terms are the harmonic components of the resonator and transmon, and the third term denotes the transmon anharmonicity (with a negative energy contribution). The last term is the coupling between both modes with typical amplitude g which is directly proportional to the capacitive coupling C_g between both modes (see Figure 1.3). In practice, this coupling is typically much smaller than the system frequencies (in the gigahertz range), $|g| \ll \omega_r, \omega_q$, such that we can assume the rotating wave approximation and arrive at a more practically-relevant Hamiltonian,

$$\mathbf{H}/\hbar \approx \omega_r \mathbf{a}^\dagger \mathbf{a} + \omega_q \mathbf{q}^\dagger \mathbf{q} - \frac{1}{2\hbar} E_C \mathbf{q}^{\dagger 2} \mathbf{q}^2 + g(\mathbf{q}^\dagger \mathbf{a} + \mathbf{q} \mathbf{a}^\dagger). \quad (1.8)$$

At this point, by restricting the transmon to its two lowest energy eigenstates, one obtains the well-known Jaynes-Cummings Hamiltonian [Jaynes & Cummings 1963] with coupling of the form $g(\sigma_+ \mathbf{a} + \sigma_- \mathbf{a}^\dagger)$ where σ_\pm are Pauli operators, that describes the light-matter interaction of electromagnetic waves (here, the resonator) with atoms (here, the transmon).

Another instructive derivation is to diagonalize the linear part of this Hamiltonian through a Bogoliubov transformation, in the regime where the detuning between both modes $\Delta \equiv \omega_q - \omega_r$ is large compared to their coupling g . This so-called dispersive regime is most commonly used with superconducting circuits because the transmon and resonator are weakly entangled, which allows the qubit to be in a well-defined quantum state. The Bogoliubov transformation reads $\mathbf{U}_{\text{disp}} = \exp[\Lambda(\mathbf{a}^\dagger \mathbf{q} - \mathbf{a} \mathbf{q}^\dagger)]$ where $\Lambda = \arctan(2\lambda)$ and $\lambda \equiv g/\Delta \ll 1$. This yields

$$\mathbf{U}_{\text{disp}}^\dagger \mathbf{H} \mathbf{U}_{\text{disp}}/\hbar \approx \tilde{\omega}_r \mathbf{a}^\dagger \mathbf{a} + \tilde{\omega}_q \mathbf{q}^\dagger \mathbf{q} - K_a \mathbf{a}^{\dagger 2} \mathbf{a}^2 - K_q \mathbf{q}^{\dagger 2} \mathbf{q}^2 - \chi_{ab} \mathbf{a}^\dagger \mathbf{a} \mathbf{q}^\dagger \mathbf{q} \quad (1.9)$$

where $\tilde{\omega}_r$ and $\tilde{\omega}_q$ are the dressed frequencies of either mode, $K_a = \mathcal{O}(\lambda^4)$ and $K_q = \mathcal{O}(1)$ are self-Kerr nonlinearities and $\chi_{ab} = \mathcal{O}(\lambda^2)$ is a cross-Kerr interaction. Here, we have also neglected several rotating terms. This expression emphasizes how both modes are now dressed by the other one, resulting in a pair of slightly hybridized modes that interact through the cross-Kerr interaction. This interaction is used for the preparation, control and readout of transmon qubits. For instance, populating the resonator would probe a transmon state-dependent frequency shift that can be measured through weak measurement of the resonator, a process hence called dispersive readout. In addition, we will also see in Section 1.4.1 how the self-Kerr term on the resonator is an essential ingredient for the Hamiltonian stabilization of so-called Kerr cat qubits.

1.2.4 The bosonic shift in paradigm

Discrete variable qubits such as the transmon are at the cornerstone of quantum computing proposals with superconducting circuits [Arute *et al.* 2019].

They are small, highly reproducible, and allow fast and standard control through their light-matter coupling to resonators as we have previously discussed. However, their discrete nature renders them susceptible to a variety of noise sources that significantly hinder their performance. The discrete noise manifests itself in the form of fluctuations in the energy levels, thermal excitations, and photon loss, to name a few. These noise sources can induce direct errors and leakage away from the two-level subspace of the qubit, posing significant challenges to error correction and fault-tolerant quantum computation. Therefore, one may want to draw inspiration from classically bistable systems, such as flip-flops in computer memories or more simply retraction mechanisms in ballpoint pens, that use noise-insensitive macroscopic variables (transistor states in a flip-flop, or the cam position of a ballpoint pen) to achieve this bistability.

Continuous variable (or bosonic) qubits thus represent a departure from the discrete nature of their predecessors. They exploit macroscopic states to improve resilience to noise and enable embedded error correction strategies. The non-linear interactions provided by superconducting circuits are then essential to create and manipulate non-trivial macroscopic states that transcend the limitations imposed by the standard gaussian Hamiltonians of quantum optics. These qubits rely on continuous-variable modes, typically defined by quadrature operators, to encode and process quantum information.

Among the various continuous variable qubits, Gottesman-Kitaev-Preskill (GKP) qubits have garnered significant attention [Gottesman *et al.* 2001]. GKP qubits leverage the principle of error correction through bosonic codes, where information is encoded in the position and momentum quadratures of the continuous-variable modes. These grid-like qubits were designed to be optimally resilient against the main source of decoherence of resonators, photon loss, allowing for error detection and correction schemes that can significantly extend coherence times. GKP qubits have been experimentally realized in superconducting circuits [Campagne-Ibarcq *et al.* 2020, Sivak *et al.* 2023] and trapped ions [De Neeve *et al.* 2022] with great success, demonstrating their promise for practical quantum computation. They are however difficult to stabilize and control, and multi-GKP qubit experiments have yet to be demonstrated.

Binomial codes are another intriguing class of continuous variable qubits [Michael *et al.* 2016]. They employ non-linearity to create non-classical states that are tailored to protect against photon loss, thermal noise or dephasing in resonators up to a given polynomial order, thanks to the structure of the encoding states in the number basis. They are also attractive candidates for error-resilient quantum information processing [Hu *et al.* 2019], but similarly have yet to demonstrate extensive stabilization and control of multi-qubit chips.

Finally, perhaps one of the simplest yet most fascinating continuous-variable codes are that of cat qubits, about which this thesis is dedicated. The cat qubit encoding [Cochrane *et al.* 1999] is characterized by superpositions of classical (i.e. coherent) states that are highly delocalized in different sides of the oscillator phase space. These states, akin to Schrödinger’s famous cat, exhibit remarkable resilience against local noise and offer the potential for low-overhead fault-tolerant quantum computation. It is usually claimed that cat qubits offer an exponential bias in one of two qubit noise channels, thus significantly reducing the experimental complexity of error correction. This thesis will discuss how this claim may or may not be valid depending on the setup, and in particular depending on the cat qubit stabilization mechanism employed.

In essence, the transition from discrete variable qubits to continuous variable qubits marks a pivotal shift in the way of engineering superconducting qubits. The light-matter coupling between discrete qubits and resonators initially developed for qubit control can now be utilized for rapid preparation and control of non-gaussian states in quantum harmonic oscillators. Resonators are then at the core of the quantum information processor and bring several advantages with them, among which longer intrinsic lifetimes.

1.3 Cat qubits

Now that we have introduced superconducting circuits and how bosonic qubits have emerged from this flourishing subfield of quantum information, this section will be dedicated to cat qubits, from their physical encoding in a quantum harmonic oscillator to how they are concatenated into error correcting codes. Some of the ideas developed in this section were first stated in [Cochrane *et al.* 1999, Ralph *et al.* 2003] and later revived with a series of recent publications, mainly [Mirrahimi *et al.* 2014, Guillaud & Mirrahimi 2019, Chamberland *et al.* 2022]. From here on out and for the rest of this thesis, we assume $\hbar \equiv 1$.

1.3.1 Encoding a cat qubit in an oscillator

A quantum harmonic oscillator is an infinite-dimensional system with equidistant energy levels. The first step in defining a bosonic qubit is to restrict the large Hilbert space to a two-dimensional subspace, which is referred to as the ‘codespace’. The codespace of cat qubits is spanned by two-component




Figure 1.4: Encoding a cat qubit in a quantum harmonic oscillator. Computational states on the Z axis of the Bloch sphere are approximately coherent states delocalized on either side of phase space, hence providing an exponentially small probability of bit-flip errors.

Schrödinger cat states [Cochrane *et al.* 1999, Ralph *et al.* 2003], according to

$$\begin{aligned} |0_L\rangle &= \frac{1}{\sqrt{2}} (|\mathcal{C}_\alpha^+\rangle + |\mathcal{C}_\alpha^-\rangle) = |\alpha\rangle + \mathcal{O}(e^{-2|\alpha|^2}) \\ |1_L\rangle &= \frac{1}{\sqrt{2}} (|\mathcal{C}_\alpha^+\rangle - |\mathcal{C}_\alpha^-\rangle) = |-\alpha\rangle + \mathcal{O}(e^{-2|\alpha|^2}) \end{aligned} \quad (1.10)$$

where $|\pm\alpha\rangle$ are coherent states of the oscillator, i.e. eigenstates of the annihilation operator, $\mathbf{a}|\pm\alpha\rangle = \pm\alpha|\pm\alpha\rangle$, characterized by their uniform gaussian distribution in phase space centered on $\pm\alpha$. Here, $|\mathcal{C}_\alpha^\pm\rangle = (|\alpha\rangle \pm |-\alpha\rangle)/\mathcal{N}_\pm$ are the even and odd parity cat states, and $\mathcal{N}_\pm = \sqrt{2(1 \pm e^{-2|\alpha|^2})}$ are normalization constants. A schematic of the encoding is represented on Figure 1.4. Note that other variants of this encoding exist, most notably four-legged cat qubits [Mirrahimi *et al.* 2014] that make use of four coherent states instead of two, or pair cat codes [Albert *et al.* 2019] encoded in two oscillators instead of one. They are however more difficult to stabilize and are not the subject of this thesis.

The logical code words $|0_L\rangle$ and $|1_L\rangle$ of cat qubits (1.10) are therefore localized on opposite sides of phase space, up to exponentially small corrections. This often leads to the claim that cat qubits are exponentially biased in noise, since local noise cannot bring one logical state to the other under a stabilization mechanism. This exponentially biased noise is particularly attractive for low-overhead quantum error correcting codes, as later discussed. Note that this is only true under the assumption of local noise — i.e. noise that may be decomposed as small oscillator displacements in the short time regime —, and most importantly, assuming that the stabilization indeed suppresses any leakage out of the codespace also in a local manner. Chapter 2 will provide a more detailed discussion of this matter.

While this definition is sufficient to describe a qubit, it is often impractical. Indeed, any local noise that the harmonic oscillator may be subject to will induce leakage in the system. At any point in time, it is then important to be able to answer the following question: what is the computational state of the qubit if the oscillator is measured? This question is definitely not trivial since a given state may be close to a computational state (in the sense of phase space distance) and yet not overlap with the codespace. As such, a complete definition of a bosonic qubit should always include a map from the infinite-dimensional oscillator to the two-dimensional codespace. Such a map is not unique, and thus should further be consistent with the readout procedure in order to capture the same dynamics between theory and experiments.

Mathematically, this qubit-defining map corresponds to a set of three observables ($\mathbf{J}_x, \mathbf{J}_y, \mathbf{J}_z$) that uniquely determine the three axes of the Bloch sphere. They are such that $\langle \sigma_i \rangle = \text{Tr}[\mathbf{J}_i \rho]$ for $i = x, y, z$ where σ_i is a Pauli operator, and ρ the density matrix of the full oscillator. They should further match the codespace definition, such that $\text{Tr}[\mathbf{J}_z |0_L\rangle\langle 0_L|] = 1$, $\text{Tr}[\mathbf{J}_z |1_L\rangle\langle 1_L|] = -1$ and similarly for the other two axes. The simplest operators that match these criteria are codespace projectors that read

$$\begin{aligned}\mathbf{J}_x^P &= |+_L\rangle\langle +_L| - |-_L\rangle\langle -_L| \\ \mathbf{J}_y^P &= |+i_L\rangle\langle +i_L| - |-i_L\rangle\langle -i_L| \\ \mathbf{J}_z^P &= |0_L\rangle\langle 0_L| - |1_L\rangle\langle 1_L|\end{aligned}\tag{1.11}$$

While this definition can be practical for numerical simulation, it may lead to large mismatches between theory and experiment due to the specific readout procedure. For instance, it may be difficult to tell two states apart if they share the same support in phase space.

For cat qubits, computational states on the x -axis of the Bloch sphere are Schrödinger cat states $|\pm_L\rangle \equiv |\mathcal{C}_\alpha^\pm\rangle$ with even and odd parity in the number of photons respectively. The extension of this definition to the entire oscillator thus naturally arises through the parity operator

$$\mathbf{J}_x = \exp(i\pi \mathbf{a}^\dagger \mathbf{a}) = \sum_{n=0}^{\infty} (-1)^n |n\rangle\langle n|\tag{1.12}$$

where \mathbf{a} is the annihilation operator on the cat qubit mode, and $|n\rangle$ is the n -th Fock state. This operator indeed verifies that $\text{Tr}[\mathbf{J}_x |\pm_L\rangle\langle \pm_L|] = \pm 1$ and that $\text{Tr}[\mathbf{J}_x \rho] \in [-1, 1]$ for any ρ . It further is an invariant of any parity-preserving dynamics, such as two-photon dissipation and Kerr cat Hamiltonians that will later be introduced.

For the other two axes, the extension of codespace projectors is not as straightforward. Since computational states along the z -axis are defined on

opposite sides of phase space, one might be tempted to define $\mathbf{J}_z \equiv \text{sign}(\mathbf{x})$ where $\mathbf{x} = (\mathbf{a} + \mathbf{a}^\dagger)/\sqrt{2}$ is the position quadrature. In this case, any population within the positive (negative) half-plane of phase space would contribute positively (negatively) to the expectation value. However, this definition does not match the cat qubit codespace as defined in (1.10) due to exponentially small populations in the opposite side of phase space (required for state normalization). For this reason, this thesis uses the invariants of two-photon dissipation as \mathbf{J}_z and \mathbf{J}_y operators. They are defined according to [Mirrahimi *et al.* 2014]

$$\mathbf{J}_z = \mathbf{J}_{+-} + \mathbf{J}_{+-}^\dagger, \quad \mathbf{J}_y = i\mathbf{J}_{+-} - i\mathbf{J}_{+-}^\dagger \quad (1.13)$$

where ($\alpha \in \mathbb{R}$ for simplicity)

$$\mathbf{J}_{+-} = \sqrt{\frac{2\alpha^2}{\sinh(2\alpha^2)}} \sum_{q=-\infty}^{\infty} \frac{(-1)^q}{2q+1} I_q(\alpha^2) \mathbf{J}_{+-}^{(q)} \quad (1.14)$$

with I_q is the modified Bessel function of the first kind, and

$$\mathbf{J}_{+-}^{(q)} = \begin{cases} \frac{(\mathbf{a}^\dagger \mathbf{a} - 1)!!}{(\mathbf{a}^\dagger \mathbf{a} + 2q)!!} \mathbf{J}_{++} \mathbf{a}^{2q+1} & q \geq 0 \\ \mathbf{J}_{++} \mathbf{a}^{\dagger(2|q|-1)} \frac{(\mathbf{a}^\dagger \mathbf{a})!!}{(\mathbf{a}^\dagger \mathbf{a} + 2|q| - 1)!!} & q < 0 \end{cases} \quad (1.15)$$

with $n!! = (n-2)!!n$ the double factorial. While the exact definition of these operators does not matter, the important feature is that \mathbf{J}_z is a very good approximation of the $\text{sign}(\mathbf{x})$ observable in the limit of large α . In addition, it continuously extends the codespace projectors, such that $\text{Tr}[\mathbf{J}_z |0_L\rangle\langle 0_L|] = 1$ and $\text{Tr}[\mathbf{J}_z |1_L\rangle\langle 1_L|] = -1$. The expectation value of \mathbf{J}_z as defined in (1.13) on a coherent state $|\beta\rangle$ is shown on Figure 1.5. As α is increased, the separation between both sides of phase space becomes sharper, showing the convergence towards the sign observable.

While the above definition arises from two-photon dissipation dynamics [Mirrahimi *et al.* 2014], we emphasize that it can still be used in other contexts for which physical readout cannot distinguish between states that are close together in phase space (e.g. beamsplitter readout of [Frattini *et al.* 2022]), and more generally any single-shot readout procedure. For higher precision and many-shot readout schemes (e.g. dispersive readout, or the buffer-based readout of [Réglaide *et al.* 2023]), definitions that take into account leakage out of the cat qubit codespace could also be brought into play, such as the one of (1.11).

As a final remark, let us highlight the following property of the \mathbf{J} operators. From (1.15), we can see that \mathbf{J}_z strictly projects the positive parity subspace to the negative one, and vice versa. As such, it swaps the parity of any state




Figure 1.5: Expectation value of the \mathbf{J}_z observable in phase space on a coherent state $|\beta\rangle$ for two values of $|\alpha|^2$. As α is increased, \mathbf{J}_z approaches $\text{sign}(\mathbf{x})$. Figure adapted from [Mirrahimi *et al.* 2014].

and thus of the $|+_L\rangle$ and $|-_L\rangle$ states. This is much alike the σ_z operator for a qubit, but would further apply to any state in the oscillator. Similarly, \mathbf{J}_x performs a rotation in phase space of angle π , thus also performing a swap of $|0_L\rangle$ with $|1_L\rangle$, i.e. an X gate with a σ_x operator. Again, this would apply to any state close to the codespace.

1.3.2 Shifted Fock basis

Now that we have introduced the cat qubit encoding, this section reviews a very handy tool for cat qubit analysis that will be used repeatedly in this thesis. The Shifted Fock Basis (SFB) was first introduced in [Chamberland *et al.* 2022] and provides an alternative basis of states for a quantum oscillator that not only can represent cat qubit dynamics particularly efficiently but also provide intuition. We refer to the original paper for a full review to the method, but this short section provides the minimum elements required for the comprehension of the derivations performed throughout the thesis.

The SFB splits the main properties of a cat qubit in two separate modes: a qubit mode that represents the logical cat qubit state, and a gauge mode that represents leakage away from the computational subspace (1.10). More precisely, let us define the non-orthonormal basis states

$$|\pm, n\rangle \equiv \frac{1}{\sqrt{2}} [\mathbf{D}(\alpha) \pm (-1)^n \mathbf{D}(-\alpha)] |n\rangle \quad (1.16)$$

where $|n\rangle$ is the n -th Fock state, and $\mathbf{D}(\alpha) = \exp(\alpha \mathbf{a}^\dagger - \alpha^* \mathbf{a})$ is the displacement operator that displaces any state in phase space without deformation,

and such that $\mathbf{D}(\alpha)|0\rangle = |\alpha\rangle$. Firstly, for $n = 0$, these states match the cat qubit computational basis states, namely $|\pm, 0\rangle \propto |\pm\rangle_L$. Second, the set of such states $\mathcal{S} = \{|\pm, n\rangle\}_{\pm, n}$ spans the full Hilbert space. They however do not define a basis of states since they are not orthonormal. Intuitively, this can be understood by noting that the number of $|\pm, n\rangle$ states is doubled compared to the regular Fock basis due to the introduction of the \pm mode¹. Although these states do not rigorously define an orthonormal basis, any two states in a separate \pm branch are exactly orthogonal due to their opposite parity in the number of photons. In addition, the first shifted Fock states are approximately orthonormal in the limit of $|\alpha| \gg 1$. To rigorously treat this basis, it can be orthonormalized for instance with a Gram-Schmidt process, but we again refer to the original publication for details.

The SFB is particularly interesting for its compactness, since it can represent any state that is locally close to the codespace with only a few basis states. Indeed, defining the truncated basis $\mathcal{S}_N = \{|\pm, n\rangle\}_{\pm, n \leq N}$, we have that the cat qubit codespace is contained in \mathcal{S}_0 , or that a codespace state that suffered leakage due to a thermal photon entering the cavity, $\mathbf{a}^\dagger |\pm\rangle_L$, is contained in \mathcal{S}_1 . More generally, it is straightforward to show that $\mathbf{a}^{\dagger N} |\pm\rangle_L$ is contained in \mathcal{S}_N such that the size of the truncated SFB grows favorably with local leakage.

Another interesting property of the SFB is how it transforms the annihilation operator \mathbf{a} . Indeed, acting on the basis states of (1.16) yields

$$\begin{aligned} \mathbf{a} |\pm, n\rangle &= \sqrt{n} |\mp, n-1\rangle + \alpha |\mp, n\rangle \\ &= \boldsymbol{\sigma}_z \otimes (\tilde{\mathbf{a}} + \alpha) |\pm, n\rangle \end{aligned} \quad (1.17)$$

where we have defined two new operators $\boldsymbol{\sigma}_z$ and $\tilde{\mathbf{a}}$ such that $\boldsymbol{\sigma}_z |\pm, n\rangle = |\mp, n\rangle$ and $\tilde{\mathbf{a}} |\pm, n\rangle = \sqrt{n} |\pm, n-1\rangle$ in analogy to a qubit and quantum oscillator mode. From (1.17) we infer the mapping of the annihilation operator from the Fock to the shifted Fock basis, which reads

$$\mathbf{a} \rightarrow \boldsymbol{\sigma}_z \otimes (\tilde{\mathbf{a}} + \alpha). \quad (1.18)$$

Again, this mapping is only valid in the limit of small shifted Fock excitation numbers because of the non-orthonormalization of the basis. It is however very helpful for understanding the intimate dynamics of cat qubit gates and for the numerical analysis of phase-flip errors.

In the following section, we dive into how the physical encoding of cat qubits we have now presented through the oscillator encoding and through the SFB can be concatenated into low-overhead error correcting codes.

¹This argument is only provided for intuition and has no mathematical ground since the Hilbert space is infinite dimensional.

1.3.3 Concatenating into error-correcting codes

In quantum information, the errors suffered by qubits can be decomposed in two types: bit-flip errors (flip along the Z axis) and phase-flip errors (flip along the X axis). Quantum error correction offers a solution to these algorithm-breaking error rates through information redundancy [Knill *et al.* 2000]. By embedding a number of physical qubits in an error-correcting code, the error rate of a reduced number of qubits (called logical qubits) is suppressed with the size of the code, but only if the physical error rates are below a given error threshold. A common error-correcting code is the surface code [Fowler *et al.* 2012] which is an optimal two-dimensional code for local and unbiased errors corrected with local measurements. It has been extensively studied in the context of transmon-based quantum computers. By encoding logical qubits in the manifold defined by stabilizer measurements, the surface code offers robustness to local noise and enables the detection and correction of errors through a series of syndrome measurements.

In this context, one particularly attractive feature of cat qubits is their noise bias. Thanks to the non-local phase space encoding, the probability of bit-flip errors is greatly reduced compared to a discrete encoding into the $|0/1\rangle$ basis states of the oscillator. With the appropriate stabilization scheme with respect to the noise channels, this bias even grows exponentially with the mean number of photons in the oscillator which is also the cat size $|\alpha|^2$. Therefore, this drastically reduces the experimental overhead of quantum error correction. This limit is particularly evident for cat qubits concatenated with a repetition code [Guillaud & Mirrahimi 2019]. In this case, logical qubits are formed only by a one-dimensional line of cat qubits instead of, for most error-correcting codes, a two-dimensional grid of qubits. This quadratic gain in qubit number could eventually be exploited for low-overhead quantum computation [Gouzien *et al.* 2023].

While the repetition code offers a straightforward and efficient method of protecting information, other error-correcting codes can also be employed with cat qubits to achieve better performance in certain scenarios. One notable example is the rectangular surface code, which can capitalize on the noise bias of the cat states to enhance the error correction capabilities further. Another interesting error-correcting code that can be forward-concatenated with cat qubits is the XZZX code [Ataides *et al.* 2021], which is specifically tailored to biased-noise qubits. The XZZX code is based on a different set of local stabilizers that feature an error-correction threshold that grows with the noise bias. By adapting the error correction strategy to the specific noise characteristics of cat qubits, the XZZX code can enhance the overall error correction performance and eventually reduce the overhead associated with error correction.

To ensure the preservation of the noise bias during operation of the device, it is crucial that any logical operation is only built with physical gates that preserve the noise bias, or bias-preserving gates. Through careful engineering, such gates are designed to suppress the introduction of certain errors that convert phase-flip errors into bit-flip errors during gate processes, thereby preserving the advantage of the noise bias. In this case, the error-correcting code can be universal (any quantum algorithm is implementable with the available gate set) and fault-tolerant (errors do not build up throughout the operation of the device).

In [Guillaud & Mirrahimi 2019], a set of bias-preserving physical operations is proposed for universal and fault-tolerant quantum computation with cat qubits. This set is composed of Pauli X , $Z(\theta)$ rotations, CNOT and Toffoli gates, in addition to the preparation and measurement of cat states along the X axis, i.e. of $|\pm_L\rangle$. Such a gate set allows for forward concatenation with various logical error-correcting codes such as ones previously mentioned. For this reason, these are the gates that will be mainly studied in this thesis. Their physical implementation then depends on the underlying qubit confinement mechanism.

In the following section, we introduce the two main variants of cat qubit confinement schemes, based on either the Kerr Hamiltonian or engineered two-photon dissipation.

1.4 Protecting cat qubits from decoherence and leakage

In the previous section, we have seen how encoding information in cat states is interesting with regard to the noise bias and to the reduction in hardware overhead. However, this promise may only hold if one can maintain the information within the cat codespace, and protect it against decoherence. The two main variants of cat qubit confinement are thus introduced in the following section. We will also discuss the standard gate designs for each scheme. The results discussed in this section are mainly developed in [Mirrahimi *et al.* 2014, Puri *et al.* 2017] for confinement mechanisms, and [Guillaud & Mirrahimi 2019, Puri *et al.* 2020] for gates.

1.4.1 Hamiltonian protection: Kerr cat qubits

1.4.1.1 Kerr cat Hamiltonian

In Section 1.2.3, we have seen how the Kerr effect is inherent to superconducting circuits with any non-linearity. While it is often employed as a source of anharmonicity to isolate lower energy levels of a qubit, it can also be at the base of a cat qubit confinement scheme [Puri *et al.* 2017]. Indeed, the Hamiltonian

$$\mathbf{H}_{\text{Kerr}} = -K(\mathbf{a}^{\dagger 2} - \alpha^{*2})(\mathbf{a}^2 - \alpha^2) \quad (1.19)$$

features a degenerate ground subspace of dimension 2, corresponding to the cat qubit codespace $\text{span}\{|\pm\alpha\rangle\}$. This degeneracy is a first essential condition that ensures no dephasing between computational states is induced by the Hamiltonian.

Studying its eigenspectrum (see Figure 1.6(a)) also shows that this ground subspace is gapped away from excited states by an energy of the order of $E_{\text{gap}} \sim 4K\alpha^2$. This second feature is required to protect the ground subspace from undesired leakage to excited eigenstates. Indeed, the adiabatic theorem guarantees that any weak and slowly-varying (compared to the energy gap) spurious Hamiltonian will only cause exponentially small leakage [Nenciu 1993]. Note that this is only true of such Hamiltonians. For instance, dissipative-like decoherence sources are not comprised in this protection. Furthermore, there is no engineered process that ensures reconvergence of the state of the oscillator to the stabilized cat qubit subspace after population leakage. As such, we favor the terminology of "confinement" rather than the more commonly used "stabilization" to refer to this protection scheme. The conditions required for exponential suppression of bit-flip errors in Kerr cat qubits are further detailed in Chapter 2.

At this point, it is instructive to make use of the SFB that we previously introduced in Section 1.3.2. Reinserting the change of basis formula $\mathbf{a} \rightarrow \sigma_z \otimes (\tilde{\mathbf{a}} + \alpha)$ into (1.19) yields

$$\mathbf{H}_{\text{Kerr}} = -K\mathbf{I} \otimes (\tilde{\mathbf{a}}^{\dagger 2}\tilde{\mathbf{a}}^2 + 2\alpha^*\tilde{\mathbf{a}}^{\dagger}\tilde{\mathbf{a}}^2 + 2\alpha\tilde{\mathbf{a}}^{\dagger 2}\tilde{\mathbf{a}} + 4|\alpha|^2\tilde{\mathbf{a}}^{\dagger}\tilde{\mathbf{a}}). \quad (1.20)$$

We can first note that this Hamiltonian acts with the identity on the qubit mode, such that no decoherence is induced by it. Second, in the limit of $|\alpha| \rightarrow \infty$, we can move into the rotating frame of the term in $\tilde{\mathbf{a}}^{\dagger}\tilde{\mathbf{a}}$ and perform a Rotating Wave Approximation (RWA) on the term $2\alpha^*\tilde{\mathbf{a}}^{\dagger}\tilde{\mathbf{a}}^2 + \text{h.c.}$ to obtain the approximate Hamiltonian

$$\mathbf{H}_{\text{Kerr}} \approx -K\tilde{\mathbf{a}}^{\dagger 2}\tilde{\mathbf{a}}^2 - 4|\alpha|^2K\tilde{\mathbf{a}}^{\dagger}\tilde{\mathbf{a}}. \quad (1.21)$$

From now on, we do not explicitly write the identity on the qubit mode. The Kerr Hamiltonian thus represents two slightly anharmonic quantum oscillators

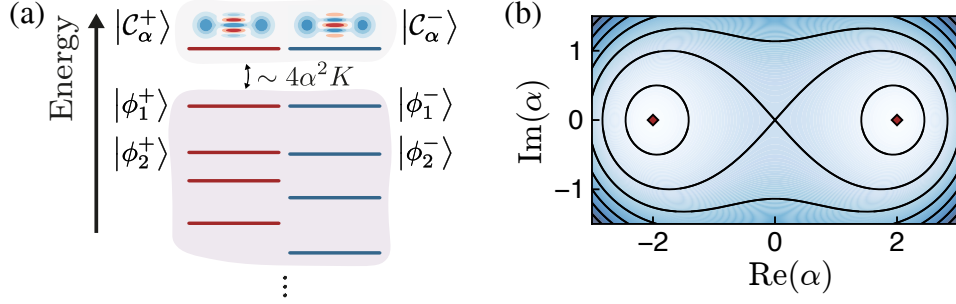


Figure 1.6: (a) Eigenspectrum of the Kerr cat Hamiltonian, separated in two branches of even and odd parity eigenstates. The degenerate ground subspace is gapped away from excited states by an energy approximately $E_{\text{gap}} \sim 4\alpha^2 K$. (b) Semi-classical potential associated to the Kerr cat Hamiltonian. Red dots represent the two stable points of the quasi-potential.

with natural frequencies (or energy gaps) $4|\alpha|^2 K$, which are centered on $\pm\alpha$. In other words, the Kerr Hamiltonian is a double well potential with two wells that are infinitely separated in the large $|\alpha|$ limit.

1.4.1.2 Engineering with a SNAIL

Remarkably, the Hamiltonian of (1.19) can be rewritten as

$$\mathbf{H}_{\text{Kerr}} = -K\mathbf{a}^{\dagger 2}\mathbf{a}^2 + \varepsilon_2\mathbf{a}^{\dagger 2} + \varepsilon_2^*\mathbf{a}^2 - |\varepsilon_2|^2/K \quad (1.22)$$

where $\varepsilon_2 = K\alpha^2$ is the amplitude of a two-photon pump. This Hamiltonian thus mainly features two terms: a self-Kerr non-linearity and a two-photon pump that exchanges pairs of photons of the oscillator with the input line. We have seen in Section 1.2 that Kerr non-linearities are inherent to Josephson junction circuits. On the contrary, the two-photon pump can be engineered with three-wave mixing for instance with a Superconducting Nonlinear Asymmetric Inductive eLement (SNAIL) [Frattini *et al.* 2017] as demonstrated in [Grimm *et al.* 2020, Frattini *et al.* 2022]. A SNAIL is made of a small Josephson junction of energy αE_J ($\alpha < 1$) in parallel with a short array of Josephson junctions ($n = 3$ junctions in the original paper) of energy E_J , as is depicted in Figure 1.7(a). This closed inductive loop can then be flux biased with an external flux φ_{ext} . The resulting inductive energy of this circuit reads

$$\mathbf{U}_{\text{SNAIL}}(\varphi) = -\alpha E_J \cos(\varphi) - n E_J \cos\left(\frac{\varphi - \varphi_{\text{ext}}}{n}\right) \quad (1.23)$$

Such a potential features a single global minimum φ_{min} for $n \geq 2$. Taylor expanding around this minimum numerically indeed yields three- and four-wave mixing terms that can be tuned through n , α and φ_{ext} .




Figure 1.7: (a) Lumped-element circuit of a SNAIL with $n = 3$, used to engineer Kerr cat qubits. Figure adapted from [Frattini *et al.* 2017]. (b) Lumped-element circuit of an ATS, used to engineer dissipative cat qubits. Figure adapted from [Lescanne *et al.* 2020b].

1.4.1.3 Semiclassical analysis

Now that we have shown how to engineer Kerr cat qubits, let us turn to their semiclassical analysis to obtain further intuition on why they can yield protected qubits. A two-photon driven Kerr nonlinear oscillator subject to single-photon loss is governed by the following master equation,

$$\frac{d\rho}{dt} = -i[\mathbf{H}_{\text{Kerr}}, \rho] + \kappa\mathcal{D}[\mathbf{a}]\rho \quad (1.24)$$

where $\mathcal{D}[\mathbf{L}]\rho = \mathbf{L}\rho\mathbf{L}^\dagger - \{\mathbf{L}^\dagger\mathbf{L}, \rho\}/2$ is the dissipation superoperator. The semiclassical output field of the oscillator, $a = \langle \mathbf{a} \rangle$, is then governed by the Heisenberg picture master equation,

$$\begin{aligned} \dot{a} &= i[H_{\text{Kerr}}, a] + \kappa\mathcal{D}^\dagger[a]a \\ &= 2iKa^*a^2 - 2i\varepsilon_2a^* - \kappa a/2 \end{aligned} \quad (1.25)$$

Further writing $a = x + iy$ and separating real and imaginary values yields the semiclassical equations of motion

$$\begin{aligned} \dot{x} &= -2Ky(x^2 + y^2) - 2\varepsilon_2y - \kappa x/2 \\ \dot{y} &= +2Kx(x^2 + y^2) - 2\varepsilon_2x - \kappa y/2 \end{aligned} \quad (1.26)$$

Finally, we get the semiclassical Hamiltonian \mathcal{H} and semiclassical dissipative potential \mathcal{V} by identifying with Hamilton's equations of motions [Hamilton 1833], $\dot{x} = -\partial_y\mathcal{H} - \partial_x\mathcal{V}$ and $\dot{y} = \partial_x\mathcal{H} - \partial_y\mathcal{V}$. They read

$$\mathcal{H}(x, y) = \frac{K}{2}(x^2 + y^2)^2 - \varepsilon_2(x^2 - y^2) \quad (1.27a)$$

$$\mathcal{V}(x, y) = \frac{\kappa}{4}(x^2 + y^2) \quad (1.27b)$$

Note that both of these equations could have been obtained directly from (1.24). Interestingly, we have thus identified two different semiclassical potentials. The first is pictured in Figure 1.6(b), and features two potential wells with degenerate global minimums in $x = \pm\alpha = \pm\sqrt{\varepsilon_2/K}$ and $y = 0$ (for $\varepsilon_2 \in \mathbb{R}$) corresponding to the cat codespace. This corroborates our previous study of the Kerr Hamiltonian in the SFB. Since it is a Hamiltonian semiclassical potential, any initial condition (x_0, y_0) will follow equipotential curves (i.e. $\mathcal{H}(x(t), y(t)) = \mathcal{H}(x_0, y_0)$ for all $t \geq 0$, as depicted in Figure 1.6(b)) and thus never fall at the bottom of either potential well. This emphasizes the non-stabilizing aspect of Kerr cat qubits that we previously discussed. The second semiclassical potential is dissipative, quadratic, and with a single global minimum in $x = y = 0$. In this case, the dynamics is contracting and will ensure that the state eventually falls in a potential minimum. The combination of these two potentials results in non-trivial dynamics described by the equations of motion of (1.26). In particular, in the limit of $\kappa \ll K$ (as desired to preserve cat qubit parity information for longer timescales), the global dynamics still features two degenerate steady states.

We now solve the three steady states of the equations of motion of (1.26) such that $\dot{x} = \dot{y} = 0$. A first trivial but unstable solution is obtained for $x = y = 0$. To simplify the derivation of the stable solutions, we introduce the radius R given by $R^2 = x^2 + y^2$, as well as the polar angle θ given by $\theta = \tan^{-1}(y/x)$. Rearranging the equations of motion, we find

$$\begin{aligned} x &= -4(\varepsilon_2 + KR^2)y/\kappa \\ y &= -4(\varepsilon_2 - KR^2)x/\kappa \end{aligned} \quad (1.28)$$

Since $x \neq 0$ or $y \neq 0$, we can always reinsert the first equation of (1.28) into the second, or the second equation into the first, and divide by y or x respectively, which yields $\kappa^2 = 16(\varepsilon_2^2 - K^2R^4)$. Further rearranging yields

$$R^2 = \frac{\varepsilon_2}{K} \sqrt{1 - (\kappa/4\varepsilon_2)^2} \quad (1.29)$$

For $\kappa \geq 4\varepsilon_2$, the only steady state solution is the trivial one at $x = y = 0$. For $\kappa < 4\varepsilon_2$ instead, there are two stable steady state solutions with a non-zero radius. Note that in the limit of $\kappa = 0$, we retrieve the known result that $R = \sqrt{\varepsilon_2/K}$. We now square both equations of (1.28) and sum them to get

$$x^2 + y^2 = \frac{16(K^2R^4 + \varepsilon_2^2)}{\kappa_1^2}(x^2 + y^2) - \frac{32KR^2\varepsilon_2}{\kappa_1^2}(x^2 - y^2) \quad (1.30)$$

Rearranging yields

$$x^2 - y^2 = R^2 \cos(2\theta) = \frac{KR^4}{\varepsilon_2} \quad (1.31)$$

Finally, we find from (1.28) that if $x > 0$, then $y < 0$ and inversely, that if $x < 0$, then $y > 0$. This yields two stable steady state solutions to the equations of motion (1.26), denoted (R, θ_0) and (R, θ_1) in polar coordinates, where the angles read respectively

$$\begin{aligned}\theta_0 &= -\frac{1}{2} \cos^{-1} \sqrt{1 - (\kappa/4\varepsilon_2)^2} \\ \theta_1 &= \pi - \frac{1}{2} \cos^{-1} \sqrt{1 - (\kappa/4\varepsilon_2)^2}\end{aligned}\tag{1.32}$$

Finally, in the limit of $\kappa \ll \varepsilon_2$, we can expand these equations to first order in $\delta = \kappa/8\varepsilon_2$ and find

$$\begin{aligned}R &= R_0 (1 - \delta^2 + \mathcal{O}(\delta^4)) \\ \theta_0 &= -\delta + \mathcal{O}(\delta^3) \\ \theta_1 &= \pi - \delta + \mathcal{O}(\delta^3)\end{aligned}\tag{1.33}$$

where $R_0 \equiv \sqrt{\varepsilon_2/K}$. Interestingly, this shows a rescaling of the cat qubit size and angle dependent on a spurious dissipative process, here the rate of single-photon loss. In Chapter 2, we will examine how purely quantum dissipative processes, such as thermal noise in the oscillator, can limit the bit-flip protection of Kerr cat qubits.

1.4.2 Dissipative protection: dissipative cat qubits

1.4.2.1 Two-photon dissipation

In the dissipative stabilization of cat qubits, a harmonic oscillator exchanges pairs of photons with its environment both through a driven dissipation process and a two-photon pump [Mirrahimi *et al.* 2014]. The Lindblad master equation that governs this oscillator is

$$\begin{aligned}\frac{d\rho}{dt} &= \kappa_2 \mathcal{D}[\mathbf{a}^2 - \alpha^2] \rho \\ &= [\varepsilon_2 \mathbf{a}^{\dagger 2} - \varepsilon_2^* \mathbf{a}^2, \rho] + \kappa_2 \mathcal{D}[\mathbf{a}^2] \rho\end{aligned}\tag{1.34}$$

where $\mathcal{D}[\mathbf{L}]\rho = \mathbf{L}\rho\mathbf{L}^\dagger - \{\mathbf{L}^\dagger\mathbf{L}, \rho\}/2$ is the dissipation superoperator, \mathbf{a} denotes the annihilation operator of the cat qubit mode, κ_2 the rate of two-photon dissipation, α the cat qubit amplitude and $\varepsilon_2 = \kappa_2\alpha^2/2$ is the two-photon pumping amplitude. With the first equality of (1.34), it is clear that $|\pm\alpha\rangle$ are steady states of the dynamics. Furthermore, since the Lindblad equation is contracting, one can show that any initial state ρ_0 converges asymptotically towards

$$\rho_\infty = c_{++} |\mathcal{C}_\alpha^+ \rangle \langle \mathcal{C}_\alpha^+| + c_{+-} |\mathcal{C}_\alpha^+ \rangle \langle \mathcal{C}_\alpha^-| + c_{-+} |\mathcal{C}_\alpha^- \rangle \langle \mathcal{C}_\alpha^+| + c_{--} |\mathcal{C}_\alpha^- \rangle \langle \mathcal{C}_\alpha^-| \tag{1.35}$$

where the coefficients $c_{\pm\pm} = \text{Tr}[\mathbf{J}_{\pm\pm}\rho_0]$ are uniquely determined by the initial condition and by the two-photon dissipation invariants previously introduced [Mirrahimi *et al.* 2014]. In particular, the final state is not necessarily pure. Because of this exponentially fast asymptotic convergence, we do allow ourselves to use the terminology of "stabilization" in this case, in the sense of dynamical systems theory [Bellman 1953].

The second equality of (1.34) emphasizes the two-way process that can engineer such a stabilization. Pairs of photons are lost the environment through two-photon dissipation while pairs of photons are pumped back into the system. This achieves an intricate equilibrium where only a subset of states are stabilized, but also highlights how purity of the initial state may not be preserved by the dissipative dynamics.

Again, it is instructive to write out (1.34) in the SFB. Performing the change of variable $\mathbf{a} \rightarrow \boldsymbol{\sigma}_z \otimes (\tilde{\mathbf{a}} + \alpha)$ yields

$$\frac{d\rho}{dt} = \kappa_2 \mathcal{D}[\tilde{\mathbf{a}}^2 + 2\alpha\tilde{\mathbf{a}}]\rho \quad (1.36)$$

which again acts with the identity on the qubit mode since it is parity preserving dynamics. Furthermore, taking the large cat limit, $|\alpha| \rightarrow \infty$, or equivalently the limit of low-leakage limit, $\langle \tilde{\mathbf{a}}^\dagger \tilde{\mathbf{a}} \rangle \ll 1$, we can neglect the second-order term in $\tilde{\mathbf{a}}^2$. This yields the simplified model

$$\frac{d\rho}{dt} = \kappa_{\text{conf}} \mathcal{D}[\tilde{\mathbf{a}}]\rho \quad (1.37)$$

where $\kappa_{\text{conf}} = 4|\alpha|^2\kappa_2$ is a local confinement rate. This is the master equation describing a lossy quantum harmonic oscillator in the rotating frame, which converges to the vacuum in a typical timescale $1/\kappa_{\text{conf}}$. Of course, here, it instead describes the exponential convergence towards both coherent states $|\pm\alpha\rangle$ of the cat codespace.

1.4.2.2 Reservoir engineering with a buffer mode

To engineer the unusual dissipation of (1.34), an ancillary buffer mode — not necessarily harmonic — is often introduced to mediate the exchange of photon pairs between the oscillator and its environment, as depicted in Figure 1.8. This results in a two-mode Lindblad master equation,

$$\frac{d\rho}{dt} = -i[\mathbf{H}_{AB}, \rho] + \kappa_b \mathcal{D}[\mathbf{b}]\rho \quad (1.38)$$

where $\mathbf{H}_{AB} = g_2(\mathbf{a}^2 - \alpha^2)\mathbf{b}^\dagger + \text{h.c.}$ is a driven two-to-one photon exchange Hamiltonian with \mathbf{b} the annihilation operator of the ancillary buffer mode. In this context, the cat qubit mode a is sometimes also referred to as the




Figure 1.8: Schematic of a cat qubit mode a stabilized by a two-to-one photon exchange with buffer mode b , mediated by a nonlinear circuit element (middle) and a pump at the appropriate differential frequency. The buffer mode is driven and dissipates photons to the environment.

memory mode. Note that, similarly as before with the two-photon pump, this interaction Hamiltonian can be rewritten as

$$\mathbf{H}_{AB} = g_2 \mathbf{a}^2 \mathbf{b}^\dagger + g_2^* \mathbf{a}^{\dagger 2} \mathbf{b} + \varepsilon_b \mathbf{b}^\dagger + \varepsilon_b^* \mathbf{b} \quad (1.39)$$

where $\varepsilon_b = -\alpha^2 g_2$. The cat state amplitude α can thus be tuned dynamically through the drive on the buffer mode.

In the limit of $\kappa_b \gg g_2$, this additional mode can be adiabatically eliminated to retrieve the single-mode model of (1.34) with a typical two-photon loss rate $\kappa_2 \equiv 4g_2^2/\kappa_b$ [Azouit *et al.* 2017, Forni *et al.* 2018]. This follows the general recipe of adiabatic elimination for bi-partite systems, with an interaction Hamiltonian of the form $g\mathbf{L}\mathbf{b}^\dagger + g^*\mathbf{L}^\dagger\mathbf{b}$ and a strong single-mode dissipation on the buffer. Such a setup results in an effective dissipation in $\mathcal{D}[\mathbf{L}]$ on the initial subsystem.

While the buffer mode was initially introduced to engineer two-photon dissipation and often ignored in subsequent studies, this thesis will reconsider its role in the overall cat qubit implementation. Indeed, purity of the memory mode can be lost through buffer mode damping, such that it is interesting to understand how the buffer is brought out of equilibrium when controlling the memory. In addition, the buffer provides inertia in the system that can be leveraged to improve physical gate designs on cat qubits, as further discussed in Chapter 3. Recent experiments such as [Réglade *et al.* 2023, Marquet *et al.* 2023] step away from the adiabatic elimination regime (characterized by $\kappa_b \gg g_2$) to increase the effective rates of cat confinement. In this case, one cannot assume the single-mode model of (1.34) and non-trivial dynamics between both modes can be observed. In [Réglade *et al.* 2023], we also demonstrate how to leverage the buffer mode to perform a full tomography readout of the memory mode. Such a design gets rid of the readout resonator and transmon of previous experiments that was

solely used for state readout at the final stage. This new setup is highly advantageous to avoid unwanted transmon ionization [Shillito *et al.* 2022] that eventually limits the bit-flip lifetime of cat qubits.

1.4.2.3 Engineering with an ATS

The two-to-one photon exchange of (1.38) was first demonstrated in [Leghtas *et al.* 2015] with the use of a Josephson junction coupler between memory and buffer modes. Because this setup featured large spurious self-Kerr terms, Ref. [Lescanne *et al.* 2020b] introduced a new Kerr-free superconducting dipole to mediate this interaction.

The Asymmetrically Threaded SQUID (ATS) is composed of two Josephson junctions of equal energy E_J (in the ideal case) shunted by a large inductance of energy E_L . It is represented in Figure 1.7(b). Similarly to a Superconducting Quantum Interference Device (SQUID), this setup allows for two closed magnetic loops that each can be threaded through an external flux. The resulting inductive potential of this dipole reads

$$U_{\text{ATS}}(\varphi) = \frac{1}{2}E_L\varphi^2 - E_J \cos(\varphi + \varphi_{\text{ext},1}) - E_J \cos(\varphi - \varphi_{\text{ext},2}) \quad (1.40)$$

Rearranging the cosine, we get

$$U_{\text{ATS}}(\varphi) = \frac{1}{2}E_L\varphi^2 - 2E_J \cos(\varphi_\Sigma) \cos(\varphi + \varphi_\Delta) \quad (1.41)$$

where the sum and differential external fluxes were defined as $\varphi_\Sigma = (\varphi_{\text{ext},1} + \varphi_{\text{ext},2})/2$ and $\varphi_\Delta = (\varphi_{\text{ext},1} - \varphi_{\text{ext},2})/2$. When biasing the ATS at the asymmetric flux point

$$\begin{aligned} \varphi_\Sigma &= \pi/2 + \varepsilon(t) \\ \varphi_\Delta &= \pi/2 \end{aligned} \quad (1.42)$$

where $\varepsilon(t)$ is an additional radiofrequency pump, and expanding the potential to first order in $|\varepsilon(t)|$, we get

$$U_{\text{ATS}}(\varphi) = \frac{1}{2}E_L\varphi^2 - 2E_J\varepsilon(t) \sin(\varphi). \quad (1.43)$$

From this sine potential, one can get odd powers of annihilation operators which are required for the two-to-one photon exchange of (1.38). Indeed, for a circuit with both a memory mode \mathbf{a} and buffer \mathbf{b} participating in the ATS, and Taylor expanding the sine in powers of the flux, we get

$$U_{\text{ATS}}(\varphi) = \frac{1}{2}E_L\varphi^2 - 2E_J\varepsilon(t)\left(\varphi - \frac{1}{6}\varphi^3 + \mathcal{O}(\varphi^5)\right) \quad (1.44)$$

where $\varphi = \varphi_{\text{zpf},a}(\mathbf{a} + \mathbf{a}^\dagger) + \varphi_{\text{zpf},b}(\mathbf{b} + \mathbf{b}^\dagger)$. In particular, by pumping at the differential frequency $\omega_p = 2\omega_a - \omega_b$ where $\omega_{a/b}$ are the memory and buffer mode frequencies respectively, the cubic flux term yields the two-to-one photon exchange Hamiltonian after a RWA.

In recent experiments, one of the main challenge is to achieve a large cat qubit confinement rate, which requires a large two-to-one photon coupling g_2 in the first place. To do so, the flux of the memory mode should have a large participation ratio into the total flux, $\varphi_{\text{zpf},a}$, crossing the ATS. However, this also implies that higher-order nonlinear terms of the ATS — often detrimental to the overall protection — could become significant when probed with additional drives and pumps. In [Marquet *et al.* 2023], an autoparametric resonance was engineered in which $2\omega_a \approx \omega_b$ by design such that no additional pump was required to make the coupling resonant. For the same two-to-one photon coupling, this allows a smaller participation ratio and thus a smaller Purcell decay rate. With this setup, the largest two-to-one photon coupling to date was achieved, with $g_2/2\pi \approx 6$ MHz.

1.4.2.4 Semiclassical analysis

Similarly to our analysis of Kerr cat qubits, this section tackles the semiclassical analysis of dissipative cat qubits taking into consideration the buffer mode [Réglade *et al.* 2023]. Starting from (1.38), we write down the Heisenberg picture master equation on the classical fields $a = \langle \mathbf{a} \rangle$ and $b = \langle \mathbf{b} \rangle$. This yields

$$\dot{a} = -2ig_2 a^* b \quad (1.45a)$$

$$\dot{b} = -ig_2(a^2 - \alpha^2) - \frac{\kappa_b}{2}b \quad (1.45b)$$

where we have taken $g_2, \alpha \in \mathbb{R}$ for simplicity. The first thing we can note from these equations is that we can retrieve the adiabatic elimination result in the limit where the b mode dynamics is much faster than that of the a mode. Indeed, setting $db/dt = 0$ and reinserting the steady state solution $b = -2ig_2(a^2 - \alpha^2)/\kappa_b$ in (1.45a) yields the single-mode equation,

$$\dot{a} = -\frac{4g_2^2}{\kappa_b} a^* (a^2 - \alpha^2). \quad (1.46)$$

This is indeed the master equation corresponding to two-photon dissipation with rate $\kappa_2 = 4g_2^2/\kappa_b$. In addition, the dissipative semi-classical potential corresponding to this equation of motion, obtained as before from $dx/dt = -\partial_x \mathcal{V}(x, y)$ and $dy/dt = -\partial_y \mathcal{V}(x, y)$, reads

$$\mathcal{V}(x, y) = \frac{\kappa_2}{4}(x^2 + y^2)^2 - \frac{\kappa_2 \alpha^2}{2}(x^2 - y^2) \quad (1.47)$$

We recognize here the same semiclassical potential as in (1.27a) for Kerr cat qubits, but this time dissipative instead of Hamiltonian. As such, any initial condition (x_0, y_0) that is not a fixed point of the dynamics would verify that $\mathcal{V}(x(t), y(t)) < \mathcal{V}(x_0, y_0)$ for $t > t_0$, hence ensuring the asymptotic convergence towards stable minimums of the potential.

Coming back to (1.45), three fixed points of the dynamics can be identified on the two-mode system. The first is unstable and reads $(a, b) = (0, 2ig_2\alpha^2/\kappa_b)$. Indeed, plugging the corresponding classical states of this fixed point in the full master equation of (1.38) shows that it is not a steady state. This is a direct consequence of the uncertainty principle, since the state cannot be localized on a single point of phase space. The other two fixed points are, as expected, $(a, b) = (\pm\alpha, 0)$. Counter-intuitively, note that the b mode is exactly in vacuum even though it is actively driven to maintain the steady state population in the memory. Linearizing the equations of motion (1.45) around these two fixed points yields

$$\dot{\tilde{a}} = \mp 2ig_2\alpha b \quad (1.48a)$$

$$\dot{b} = \mp 2ig_2\alpha\tilde{a} - \frac{\kappa_b}{2}b \quad (1.48b)$$

where $\tilde{a} = a \mp \alpha$. Defining the two-mode variable $X = [\tilde{a}, b]^T$, we can rewrite these equations as

$$i\dot{X}(t) = AX(t) \quad (1.49)$$

where

$$A = \begin{pmatrix} 0 & \pm 2g_2\alpha \\ \pm 2g_2\alpha & -i\kappa_b/2 \end{pmatrix} \quad (1.50)$$

of which the solution is $X(t) = e^{-itA}X(0)$. The local convergence rate towards fixed points of the dynamics, often referred to as confinement rate, is thus given by the eigenvalue of largest imaginary part of A . To determine this confinement rate, there are three possible regimes of interest depending on the ratio between buffer mode dissipation and rate of two-to-one photon exchange.

- *Overdamped regime*, $\kappa_b/2 > 4\alpha g_2$. This first regime overlaps with the adiabatic elimination regime of $\kappa_b \gg g_2$. The eigenvalues of A are then given by $\lambda_{\pm} = -i\kappa_b(1 \pm \sqrt{1 - (8\alpha g_2/\kappa_b)^2})/4$, and the corresponding confinement rate by $\kappa_{\text{conf}} = 2|\text{Im}[\lambda_-]|$. In particular, in the limit where $\kappa_b \gg g_2$, we retrieve the familiar confinement rate

$$\kappa_{\text{conf}} = 16\alpha^2 g_2^2 / \kappa_b = 4\alpha^2 \kappa_2 \quad (1.51)$$

where $\kappa_2 = 4g_2^2/\kappa_b$, as we showed with our previous SFB analysis.




Figure 1.9: Confinement rate κ_{conf} of a dissipative cat qubit as a function of the buffer mode dissipation rate κ_b . The maximal confinement rate $\kappa_{\text{conf}}^* = 4\alpha g_2$ is obtained at the critical regime for $\kappa_b = 8\alpha g_2$. The red line is obtained from semiclassical analysis. Blue lines are obtained from exact diagonalization of the Liouvillian.

- *Critical regime*, $\kappa_b/2 = 4\alpha g_2$. At the critical regime, the eigenvalues of A are degenerate and given by $\lambda_{\pm} = -i\kappa_b/2$. This yields the maximal possible confinement rate $\kappa_{\text{conf}}^* = 4\alpha g_2$ for a fixed value of the two-to-one photon exchange rate g_2 .
- *Underdamped regime*, $\kappa_b/2 < 4\alpha g_2$. Finally, the underdamped regime is characterized by damped oscillations when the solution reconverges to the system steady states. In this case, the eigenvalues of A read $\lambda_{\pm} = -i\kappa_b(1 \pm i\sqrt{1 - (8\alpha g_2/\kappa_b)^2})/4$. The confinement rate is then given by their shared imaginary part, and reads $\kappa_{\text{conf}} = 2|\text{Im}[\lambda_{\pm}]| = \kappa_b/2$.

Figure 1.9 shows the confinement rate for these three regimes as a function of the buffer mode dissipation κ_b . The underdamped regime (left) thus shows a linear increase of the confinement rate with κ_b , and the overdamped regime (right) a $1/\kappa_b$ dependence. Since the buffer mode dissipation can be tuned through capacitive coupling to the transmission line, the ideal experimental regime is likely that of the critical regime which maximizes confinement rate of the memory. As a final remark, note that this semiclassical analysis breaks down for $\alpha = 0$, for which one of the steady states of two-photon dissipation is purely quantum.

In the following section, we turn to the methods of gate engineering for both Kerr and dissipative cat qubits.

1.4.3 Engineering gates on protected qubits

For quantum computation with cat qubits, a set of physical gates has to be devised and tailored to the available hardware. This set should be universal at

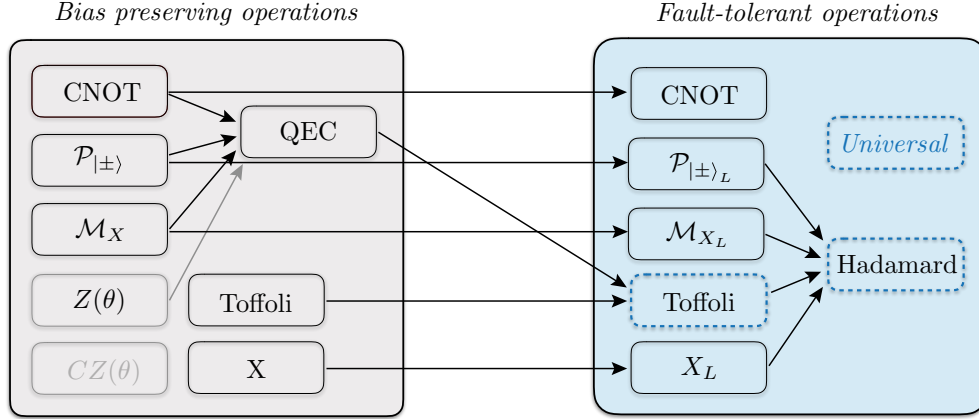


Figure 1.10: Overall scheme for achieving fault-tolerant universal quantum computation using cat qubits in a repetition code. Physical and bias-preserving gates are depicted on the left-hand side. Fault-tolerant logical operations acting on the repetition cat qubits (right-hand side) are built out of these operations, as depicted by the arrows. Figure borrowed from [Guillaud & Mirrahimi 2019].

the logical level, such that when these physical gates are forward-concatenated with an error-correcting code, the available logical operations can carry out any quantum algorithm. Since cat qubits are highly biased towards phase-flip errors, all physical gates should further be *bias-preserving*, which means that no phase-flip error should be converted into a bit-flip error during the operation of the device. This would otherwise break fault-tolerance of the error-correcting code that features no bit value redundancy at the logical level.

In [Guillaud & Mirrahimi 2019], such a set of bias-preserving physical operations is proposed for universal quantum computation with cat qubits. This set is comprised of only four gates — Pauli X , Z rotation, CNOT and Toffoli —, in addition to the preparation and measurement of cat states in the $|\pm\rangle_L$ computational basis. This scheme is represented in Figure 1.10 together with the fault-tolerant logical operations it can achieve. In this section, we go over the current proposals for the design of the four physical gates with Kerr and dissipative cat qubits, following [Mirrahimi *et al.* 2014], [Puri *et al.* 2020], and [Guillaud & Mirrahimi 2019]. Reference [Chamberland *et al.* 2022] also provides an extensive review of gate designs in the dissipative case with the introduction of analytic formulas for gate errors.

1.4.3.1 Pauli X

The Pauli X gate corresponds to a swap of $|0\rangle \approx |\alpha\rangle$ and $|1\rangle \approx |-\alpha\rangle$. This can be achieved with a π -phase delay on the oscillator, or equivalently, with a cat qubit codespace rotation in phase space for some time $T = \pi/\Delta$, where Δ is a frequency detuning with respect to the cat-qubit frame. This exchanges the two computational basis states, as illustrated in Figure 1.11. From a complementary viewpoint, the engineered dynamics should confine the cat qubit codespace with (1.19) (Kerr) or (1.34) (dissipative), but in a frame that rotates at the detuned frequency. Working this back to cat-qubit frame, we see that we must engineer the dynamics of

$$\begin{aligned} \text{Kerr: } \frac{d\rho}{dt} &= -i[\Delta \mathbf{a}^\dagger \mathbf{a}, \rho] + i[K(\mathbf{a}^{\dagger 2} - \alpha^{*2} e^{2i\Delta t})(\mathbf{a}^2 - \alpha^2 e^{-2i\Delta t}), \rho] \\ \text{Dissipative: } \frac{d\rho}{dt} &= -i[\Delta \mathbf{a}^\dagger \mathbf{a}, \rho] + \kappa_2 \mathcal{D}[\mathbf{a}^2 - \alpha^2 e^{-2i\Delta t}] \rho \end{aligned} \quad (1.52)$$

Since these master equations preserve photon parity, no phase errors are induced by this gate design.

While the implementation of both terms of these two equations is well-established, one can in principle choose to engineer only one of the two terms. If only the cat confinement term is kept, the cat state is pulled by the time-varying setpoint of the Kerr Hamiltonian or two-photon dissipation, incurring negligible leakage and bit-flips for a sufficiently slow codespace rotation. Note in particular that for an adiabatic codespace deformation, the Kerr Hamiltonian implementation would benefit from the adiabatic theorem and achieve exponentially low leakage during the process. This is not true of two-photon dissipation.

With the detuning Hamiltonian alone, the rotation is performed exactly, but the confining effect of the Kerr Hamiltonian or two-photon dissipation is turned off during the gate. For fast-enough gates however, the design is still bias-preserving under local errors, and codespace leakage can be suppressed after the gate in the dissipative case.

1.4.3.2 $Z(\theta)$ rotations

Cat-qubit $Z(\theta)$ rotations require the accumulation of a different phase on the $|\alpha\rangle$ and $|-\alpha\rangle$ components of the codespace. Since this exact evolution is not directly accessible with simple experimental means, the standard proposal is an approximate one based on the adiabatic theorem for Kerr cat qubits, or the Zeno effect for dissipative cat qubits.

The combination of a small drive displacing the state in phase-space and of the confinement scheme pulling the state back to the codespace induces




Figure 1.11: Codespace deformation. (a) Idling cat qubit, with fixed points of the confinement in $\pm\alpha$ (red diamonds). (b) Codespace deformation during an X gate. The confinement fixed points are rotated along the red circle. The cat qubit is "pulled" towards these fixed points for sufficiently slow rotations. Figure adapted from [Guillaud & Mirrahimi 2019].

an effective phase-shift inside a slightly deformed codespace, and hence the desired gate. The master equation to be engineered takes the standard form,

$$\text{Kerr: } \frac{d\rho}{dt} = -i[\mathbf{H}_Z(t), \rho] + i[K(\mathbf{a}^{\dagger 2} - \alpha^{*2})(\mathbf{a}^2 - \alpha^2), \rho] \quad (1.53a)$$

$$\text{Dissipative: } \frac{d\rho}{dt} = -i[\mathbf{H}_Z(t), \rho] + \kappa_2 \mathcal{D}[\mathbf{a}^2 - \alpha^2]\rho \quad (1.53b)$$

where

$$\mathbf{H}_Z(t) = \varepsilon_Z(t)\mathbf{a}^\dagger + \varepsilon_Z^*(t)\mathbf{a} \quad (1.54)$$

is a single-photon drive and where the gate angle reads $\theta = 4 \int \text{Re}[\alpha \varepsilon_Z] dt$. Contrary to the case of the X gate, it is here essential to keep the confinement on during the operation. Otherwise, a standard displacement would be induced without any particular gate induced. Once again, it is instructive to plug in the SFB change of basis formula. This yields (for $\alpha, \varepsilon_Z \in \mathbb{R}$),

$$\mathbf{H}_Z(t) = \varepsilon_Z(t)\sigma_z \otimes (2\alpha + \tilde{\mathbf{a}}^\dagger + \tilde{\mathbf{a}}) \quad (1.55)$$

from which we can identify two terms. The first term in $2\alpha\varepsilon_Z(t)\sigma_z$ is proportional to the Pauli operator and thus engineers the required gate. The second in $\sigma_z \otimes (\tilde{\mathbf{a}} + \tilde{\mathbf{a}}^\dagger)$ is a spurious leakage term that induces codespace excitations correlated with Pauli Z errors on the qubit. Chapter 3 explores solutions to mitigate the negative effects of this spurious term.

1.4.3.3 CNOT and Toffoli

By definition, a CNOT gate is a Pauli X gate on a target qubit conditioned on the state of a control qubit along its Z axis, or equivalently, a Pauli Z gate on a control qubit conditioned on the state of a target qubit along its X axis. Therefore, the standard design of CNOT gates for cat qubits involves a combination of the X and $Z(\pi)$ gate implementations, and of their respective issues. From the viewpoint of the X gate, the dynamics of (1.52) should be applied conditionally on the computational state of a control qubit. This can be achieved with the two-mode master equation

$$\text{Kerr: } \frac{d\rho}{dt} = -i[\mathbf{H}_{CX}, \rho] + i[\mathbf{H}_{\text{Kerr},C}, \rho] + i[\mathbf{H}_{\text{Kerr},T}(t), \rho] \quad (1.56a)$$

$$\text{Dissipative: } \frac{d\rho}{dt} = -i[\mathbf{H}_{CX}, \rho] + \kappa_2 \mathcal{D}[\mathbf{L}_C]\rho + \kappa_2 \mathcal{D}[\mathbf{L}_T(t)]\rho \quad (1.56b)$$

where $\mathbf{H}_{\text{Kerr},C} = -K(\mathbf{a}_C^{\dagger 2} - \alpha^{*2})(\mathbf{a}_C^2 - \alpha^2)$, and $\mathbf{L}_C \equiv \mathbf{a}_C^2 - \alpha^2$ are the standard Kerr Hamiltonian and two-photon dissipation on the control qubit. On the target qubit, $\mathbf{H}_{\text{Kerr},T} \equiv -K\mathbf{L}_T^\dagger(t)\mathbf{L}_T(t)$ is a time-dependent Kerr-based Hamiltonian, and $\mathbf{L}_T(t)$ a time-dependent two-photon dissipation that reads

$$\mathbf{L}_T(t) \equiv \mathbf{a}_T^2 - \frac{\alpha}{2}(\mathbf{a}_C + \alpha) + \frac{\alpha}{2}(\mathbf{a}_C - \alpha)e^{-2i\Delta t}. \quad (1.57)$$

Finally, the coupling Hamiltonian reads

$$\mathbf{H}_{CX} \equiv \varepsilon_{CX}(\mathbf{a}_C^\dagger + \mathbf{a}_C - 2\alpha)(\mathbf{a}_T^\dagger \mathbf{a}_T - n_p). \quad (1.58)$$

Here $\mathbf{a}_{C/T}$ denote the control and target qubit modes respectively, n_p is any even integer close to $|\alpha|^2$, and $\varepsilon_{CX}(t) = \Delta(t)/4\alpha$. Similarly to the X gate, Hamiltonian (1.58) may be dropped to ease experimental requirements as long as the time-dependent fixed points of the cat confinement in (1.56) rotates slowly enough to adiabatically pull the state. Alternatively, the target qubit confinement of (1.56) can also be dropped, in which case the Hamiltonian (1.58) alone induces the intended target qubit rotation, but at the cost of turning off target qubit confinement temporarily during the gate.

The viewpoint of the Z gate clarifies the impact of conditioning on the control qubit. The link is most direct when the target confinement of (1.56) is dropped, as proposed in Chapter 2. Indeed, Hamiltonian (1.58) together with confinement on the control qubit amounts to the same dynamics as for the $Z(\theta)$ rotation; the only difference being that the drive amplitude ε_Z is now conditioned on the photon number in the target qubit state. In some way, one can identify $\varepsilon_Z \sim \varepsilon_{CX}(\mathbf{a}_T^\dagger \mathbf{a}_T - n_p)$. This achieves the required gate since the control qubit undergoes an even number of $Z(\pi)$ gates for all even Fock states of the target qubit, and respectively an odd number of $Z(\pi)$ gates for

all odd Fock states, hence refocusing on the even and odd cat codespaces at the end of the CNOT gate.

For Toffoli gates, the discussion is very similar to that of the CNOT gate, but a second control mode should be added to the equations to engineer, symmetrically to the first control mode. We refer to [Guillaud & Mirrahimi 2019] for further details.

1.4.3.4 Gate errors

During each of these gates, there are two main effects that induce phase errors. The first is spontaneous emission of photons from the oscillator, with dissipation operator $\kappa_a \mathcal{D}[\mathbf{a}]$. This causes phase errors linearly in time and in $|\alpha|^2$ independently of the ongoing gate, and also results in correlated errors for multi-qubit gates [Chamberland *et al.* 2022]. Indeed, plugging in the SFB into this dissipator and neglecting second-order terms yields

$$\kappa_a \mathcal{D}[\mathbf{a}] \approx \kappa_a |\alpha|^2 \mathcal{D}[\boldsymbol{\sigma}_z]. \quad (1.59)$$

This directly highlights the loss of phase information induced by single-photon losses on the cat qubit, with a rate $\Gamma_Z \approx \kappa_a |\alpha|^2$. This source of decoherence can only be mitigated by increasing oscillator lifetime or by changing the overall qubit encoding. Single-photon losses thus motivate the need for faster gate operations (increasing overall fidelities) and limit the benefits granted by increasing $|\alpha|^2$ (providing an optimal working point).

Phase errors are also induced directly by gate processes, but only on modes for which parity-switching dynamics is engineered. In particular, target qubits of multi-qubit gates do not suffer from such errors, and hence most of the dynamics of interest occurs on control qubits. Consequently, $Z(\theta)$ rotation gates provide much of the important physics, and strategies to mitigate cat qubit gate errors can all be understood within the scope of this single-qubit gate. Therefore, this thesis often begins with the design of $Z(\theta)$ gates before generalizing.

1.5 Outline and Contributions

In this introductory chapter, we have outlined some of the key elements in the path towards quantum computation with protected cat qubits. Starting from the early development of quantum physics (1900s), quantum computation was imagined by Richard Feynman only much later in the century (1982). We then reviewed progress in superconducting circuits (1990s) and their typical models and Hamiltonians. Continuous-variable qubits were then theorized as a promising path towards fault-tolerance (2001), from which the idea

of passively protecting a cat qubit with dissipation arose several years later (2014). This approach being actively pursued experimentally, another variety of cat qubit protection based on the Kerr effect was invented (2017). Finally, we discussed the current proposals for bias-preserving operations with both kinds of protected cat qubits, and their path towards fault-tolerant quantum computation (2019).

Regrettably, both dissipative and Kerr cat qubits suffer from respective limitations, namely the low-order scaling of gate fidelities and a noise bias limited by thermal-induced leakage. The goal of this thesis will thus be to design new ways of controlling protected cat qubits in order to get rid of these limitations. This manuscript is organized as follows.

In Chapter 2, we provide a detailed analysis of the reason why Kerr-based confinement shows a bounded scaling in noise bias in the presence of thermal noise. We then introduce a novel confinement mechanism based on a two-photon exchange from the memory to an ancillary mode, and show how it can be combined in a highly efficient manner with dissipative cat qubits. Such a combined confinement mechanism displays improved single- and multi-qubit gate fidelities while retaining the advantageous scaling of bit-flip errors. It can further be implemented in a superconducting circuit setup with minimal overhead.

In Chapter 3, we dive into the intricate interactions between memory and buffer mode, and describe the role of the buffer mode in the making of gate errors during Zeno gates. From this insight, we devise four separate designs of high-fidelity gates for dissipative cat qubits. Two of these designs are based on feeding information leaked out of the buffer mode back into the memory. The two other designs rely on high-order Hamiltonian or dissipative engineering such that no information is lost to the environment in the first place. For all of these designs, we go over specific implementations in superconducting circuits, and how eventual spurious terms could limit viability of the proposals.

In Chapter 4, we introduce a novel method to differentiate through a Lindblad master equation with a constant memory cost, based on the reverse time integration of an adjoint state master equation. Such a tool is useful in many contexts such as parameter fitting, state tomography or quantum optimal control. In the context of this thesis, it could be used for the optimization of gate or readout processes with cat qubits. To be more general, we demonstrate it in the context of transmon readout. By performing a gradient descent on the readout pulse within a transmon-resonator-filter model, we show how pulse shaping can improve readout fidelities compared to standard readout protocols. It also provides interpretable pulses, and learns how to fuse operations together within a protocol where dissipation is key to the process being optimized.

Chapter 5 then concludes this thesis. Throughout the manuscript, we briefly mention some other contributions. Namely, we have provided some theoretical support on the work of [Ruiz *et al.* 2023] in which a detuned Kerr cat qubit at specific working points is studied for the multiplicity of its spectral degeneracies, we helped devise the transmon-free readout protocol of [Réglade *et al.* 2023] that was then experimentally demonstrated, and we co-developed the open-source library *dynamiqs* that provides end-to-end differentiable solvers of quantum differential equations, available at github.com/dynamiqs/dynamiqs.

Combined confinement of cat qubits

This chapter covers the work that was published in [Gautier *et al.* 2022] regarding hybrid confinement schemes for the protection of cat qubits.

Contents

2.1	Introduction	39
2.2	Cat confinement and bit-flip suppression	41
2.2.1	Confinement and decoherence	41
2.2.2	Bit-flip suppression	43
2.2.3	Estimating bit-flip errors for an idling qubit	46
2.2.4	Combined Kerr and two-photon dissipation confinement	50
2.3	Combined confinement with two-photon exchange . .	52
2.3.1	The two-photon exchange Hamiltonian	52
2.3.2	Combined TPE and two-photon dissipation	54
2.4	Engineering gates under combined confinement . . .	58
2.4.1	Zeno Z gate	58
2.4.2	CNOT gate	66
2.5	Towards experimental realization	72
2.5.1	Superconducting circuit implementation	72
2.5.2	Spurious terms	75
2.5.3	Circuit Hamiltonian derivation	77
2.6	Conclusion	81

2.1 Introduction

In the previous chapter, we reviewed the two main approaches that have been proposed in order to confine a cat qubit to its two-dimensional codespace. The

first approach is based on engineered two-photon dissipation which attracts any initial state towards the cat qubit manifold [Mirrahimi *et al.* 2014], and the second on a Hamiltonian confinement with two-photon driving and a Kerr nonlinearity [Puri *et al.* 2017]. In this case, confinement is ensured by the degenerate ground eigenspace of the Hamiltonian and by its spectral gap to the rest of the spectrum.

The dissipative confinement approach benefits from the advantage that the states in the cat qubit manifold correspond to the only steady states of the dissipative mechanism. Any potential leakage outside the code space is therefore suppressed actively and any undesired perturbation to the ideal dynamics lead only to bit-flip and/or phase-flip errors in the code space. Furthermore, as discussed in [Cohen 2017] and experimentally demonstrated in [Lescanne *et al.* 2020b], the associated bit-flip errors are exponentially suppressed with the oscillator mean number of photons. The main inconvenience of this approach is currently the limited performance of the physical gates, which leads to challenging requirements for the ratio between engineered two-photon dissipation and spurious relaxation rates in order to reach fault-tolerance thresholds.

With Kerr-based confinement, gate performances can be improved through the application of superadiabatic pulse designs that take advantage of the purely Hamiltonian evolution to suppress gate-induced leakage out of the code space [Xu *et al.* 2022a]. However, in the absence of dissipative stabilization of the cat qubit manifold, leakage induced by perturbations different from quasi-static Hamiltonians — e.g. from thermal excitation or photon dephasing — is not countered. As shown in [Putterman *et al.* 2021] and further argued in this chapter, this leakage can in turn lead to significant bit-flip errors that are not exponentially suppressed with the number of photons, or at least not within the usual range of parameters.

The authors of [Putterman *et al.* 2021] proposed a promising approach to remedy this problem and to ensure the suppression of bit-flips. This approach consists in the addition of a colored relaxation to compensate for the leakage out of the cat manifold. Through this addition, the bit-flip suppression would be re-established for Kerr cat qubits, but to reach the same level of performance as in the case of the two-photon dissipation, one would need a careful hardware engineering of the bath beyond the Purcell filters that are routinely used in superconducting devices.

It is tempting to think that by combining the two mechanisms of two-photon dissipation and Kerr Hamiltonian confinement, we might benefit from the best of both worlds. This idea is valid but it comes with some important limitations. First, the realization of the bias-preserving operations would require the combination of the engineered Hamiltonians and dissipators as laid out in [Guillaud & Mirrahimi 2019] and [Puri *et al.* 2020]. This is a daunting

experimental task which could lead to unanticipated roadblocks. Second, as it will become clear in the following analysis, in order to re-establish a similar level of bit-flip suppression to the purely dissipative mechanism, the Kerr strength cannot surpass the two-photon decay rate. This significantly reduces the interest of the Kerr confinement in performing fast and high-fidelity gates.

To overcome this difficulty, we introduce an alternative cat qubit confinement Hamiltonian that does not suffer from the limitations of Kerr confinement, and which can be combined with a dissipative protection mechanism. While our focus here is on bosonic cat qubits, the idea is rather general. Indeed, the protection of quantum information generally relies on either a Hamiltonian gap, where the information is encoded in a degenerate eigenspace of an engineered Hamiltonian, or a dissipative mechanism, usually based on reservoir engineering or feedback control. We show that the combination of these two techniques is not straightforward. More precisely, for such combinations one cannot simply rely on the energy gap but it is also needed to carefully engineer the spectrum of excited energy levels. We therefore hope that similar methods could be replicated on other quantum information devices.

The purpose of the present chapter is thus to propose a highly efficient scheme for combined Hamiltonian and dissipative cat qubit confinement, and to analyze its performance compared to other schemes. Section 2.2 first investigates the exponential suppression of bit-flip errors for two existing confinement schemes: two-photon dissipation and Kerr confinement. Section 2.3 introduces the Two-Photon Exchange (TPE) Hamiltonian, a nonlinear coupling between a resonator and a two-level buffer system that features uniformly spectral gaps, and shows that it is an ideal candidate for combination with a two-photon dissipative stabilization. In Section 2.4, the performance of single-qubit Z gates and a two-qubit CNOT gates under such a combined confinement is studied, and gate fidelity improvements by factors of up to 400 compared to standard designs are shown without compromising bit-flip protection. Finally, Section 2.5 presents a proposal for the experimental realization of this combined TPE and two-photon dissipation confinement scheme, and investigates the influence of various experimental imperfections on the scheme's performance. We conclude in Section 2.6.

2.2 Cat confinement and bit-flip suppression

2.2.1 Confinement and decoherence

Let us consider a quantum harmonic oscillator under some cat qubit confinement scheme and typical decoherence effects. The master equation governing

the evolution of the oscillator is given by

$$\frac{d\rho}{dt} = \mathcal{L}_{\text{conf}}\rho + \kappa_- \mathcal{D}[\mathbf{a}]\rho + \kappa_+ \mathcal{D}[\mathbf{a}^\dagger]\rho + \kappa_\phi \mathcal{D}[\mathbf{a}^\dagger \mathbf{a}]\rho. \quad (2.1)$$

Here $\mathcal{D}[\mathbf{L}]\rho = \mathbf{L}\rho\mathbf{L}^\dagger - \{\mathbf{L}^\dagger \mathbf{L}, \rho\}/2$ is the dissipation super-operator associated to channel \mathbf{L} , applied with pure dephasing rate κ_ϕ , single-photon relaxation and excitation rates $\kappa_- = \kappa_1(1 + n_{\text{th}})$ and $\kappa_+ = \kappa_1 n_{\text{th}}$, with n_{th} the average number of thermal photons. The superoperator $\mathcal{L}_{\text{conf}}$ denotes a confinement scheme, i.e. a super-operator with a two-dimensional steady state manifold equal to the cat qubit codespace.

For dissipative cat qubits, the confinement scheme takes the form of a two-photon driven dissipation $\mathcal{L}_{\text{conf}} = \kappa_2 \mathcal{D}[\mathbf{a}^2 - \alpha^2]$ where the bare confinement rate is the two-photon dissipation rate κ_2 [Mirrahimi *et al.* 2014]. This confinement scheme is an asymptotic stabilization scheme in the sense of dynamical systems theory, meaning that any initial state will asymptotically converge to the cat qubit manifold. The stabilization is furthermore exponential since any initial state in a close neighborhood of the cat qubit manifold will converge exponentially fast to a steady state in this manifold. In the limit $\kappa_2 \gg \kappa_1$, the typical time scale of this convergence is given by a re-scaled confinement rate $\kappa_{\text{conf}} = 4|\alpha|^2 \kappa_2$ [Azouit *et al.* 2016].

For Kerr cat qubits, the confinement scheme is instead an energy-conserving superoperator $\mathcal{L}_{\text{conf}} = -i[\mathbf{H}_{\text{Kerr}}, \cdot]$ where $\mathbf{H}_{\text{Kerr}} = -K(\mathbf{a}^{\dagger 2} - \alpha^{*2})(\mathbf{a}^2 - \alpha^2)$ with Kerr non-linearity K [Puri *et al.* 2017]. This Hamiltonian features an exactly degenerate ground eigenspace given by the cat qubit manifold, which is gapped from other eigenstates with an energy approximately given by $4|\alpha|^2 K$ in the limit of large α , as shown in Figure 2.1(a). This energy gap protects the qubit from spurious weak and slowly varying Hamiltonian perturbations, as the perturbed eigenspace will remain very close to the cat qubit manifold. However, when the state leaks out, no process ensures its asymptotic convergence back to the cat qubit manifold.

The excited eigenstates of the Kerr Hamiltonian can be separated into two branches of even and odd photon number parities, which asymptotically converge towards $|\mathcal{C}_\alpha^+\rangle$ and $|\mathcal{C}_\alpha^-\rangle$ respectively under two-photon dissipation. Since modifying the phase between even and odd cat states corresponds to bit-flips, we would ideally want these two excited branches to be degenerate like the cat states. However, the Kerr Hamiltonian excited eigenstates are all non-degenerate, with energy level spacing between consecutive energy levels increasing with the excitation number, as shown on Figure 2.1(a). Although each spacing can be suppressed exponentially with the mean number of photons $|\alpha|^2$, this suppression only kicks in at $|\alpha|^2 \gtrsim 4n$ for the n -th pair of excited states [Puri *et al.* 2017, Putterman *et al.* 2021]. As discussed in [Putterman *et al.* 2021] and further argued here, this fact has major con-

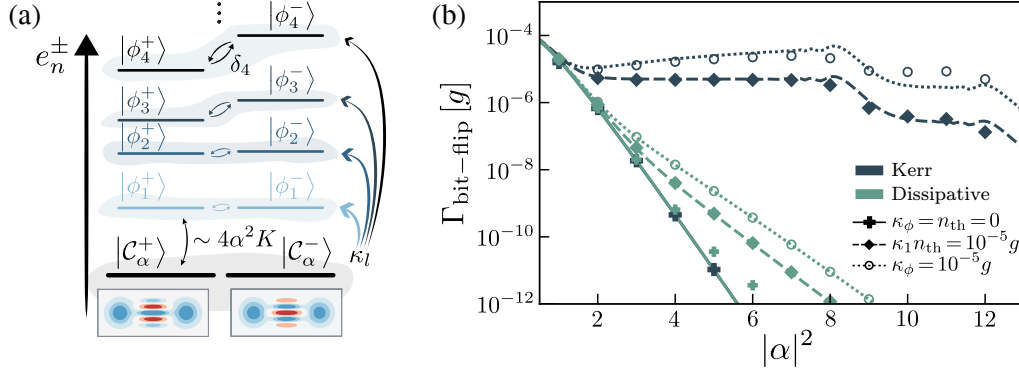


Figure 2.1: (a) Energy spectrum of the Kerr Hamiltonian separated in two branches of even and odd parity eigenstates. The ground eigenspace is the cat qubit codespace, gapped from excited eigenstates by an energy $|e_1^\pm - e_0| \approx 4|\alpha|^2 K$. Leakage out of the codespace at rate $\kappa_l = \kappa_1 n_{th} + |\alpha|^2 \kappa_\phi$ readily brings the system towards a superposition of all excited states $|\phi_n^\pm\rangle$, where even and odd excited eigenstates dephase at rate $\delta_n = |e_n^- - e_n^+|$. (b) Bit-flip error rate of an idle cat qubit confined with a Kerr Hamiltonian (blue, $g \equiv K$) or with two-photon dissipation (green, $g \equiv \kappa_2$), in units of g . The cat qubit is subject to single-photon loss with rate $\kappa_1 = 10^{-3}g$ (crosses, solid lines), with additional thermal noise $\kappa_1 n_{th} = 10^{-5}g$ (diamonds, dashed lines) or pure dephasing $\kappa_\phi = 10^{-5}g$ (circles, dotted lines). Markers indicate numerical data, lines indicate analytical fits.

sequences on the effective bit-flip suppression for such a confinement.

2.2.2 Bit-flip suppression

In Figure 2.1(b), the bit-flip error rates of idling dissipative cat qubits and idling Kerr cat qubits are presented in the presence of various noise processes. While markers show simulation data at large idling times, lines show analytical predictions as detailed in Section 2.2.3. For $\kappa_\phi = n_{th} = 0$, both schemes feature an exponential bit-flip error suppression, proportional to $\kappa_1 |\alpha|^2 \exp(-4|\alpha|^2)$. In the presence of thermal or dephasing noise, dissipative cat qubits retain an exponential error bias with exponential suppression proportional to $\kappa_l \exp(-2|\alpha|^2)$ where $\kappa_l = \kappa_1 n_{th} + |\alpha|^2 \kappa_\phi$ is the leakage rate out of the cat qubit manifold.

Kerr cat qubits on the other hand feature an approximately constant bit-flip error rate given by the leakage rate κ_l . The predictions fit simulation data very closely, and we now detail the physical process of Kerr-induced bit-flip errors behind the models used here. Thermal and dephasing noise induce leakage out of the codespace at a rate denoted by κ_l , which is *a priori* small

compared to the Kerr nonlinearity, $\kappa_l \ll K$. This leakage is slow but, in the absence of any mechanism attracting it back to the cat qubit manifold, over a time t it deforms a state initially in the cat qubit manifold into a new state that is about $\kappa_l t$ -removed from the steady state manifold. This new state has a finite overlap with all excited eigenstates of the Kerr Hamiltonian (denoted as $|\phi_n^\pm\rangle$) which form an orthogonal basis of the Hilbert space. Let us consider the n -th pair of Kerr eigenstates. These two eigenstates dephase at a rate given by their energy level spacing $\delta_n = e_n^- - e_n^+$. Since they are by definition eigenstates of even and odd parities, some dephasing of phase angle π will result in an effective bit-flip error.

With respect to the above discussion, single-photon loss is in fact a confinement mechanism attracting the leaked state back to the cat qubit manifold at the rate κ_1 . However, as $\kappa_1 \ll K$, the system will potentially undergo bit-flip errors before this reconvergence to the code space. In Section 2.2.3, this model is detailed and, over a time t short compared to rate of leakage, $\kappa_l t \ll 1$, and short compared to the rate of reconvergence, $\kappa_1 t \ll 1$, it yields the following bit-flip error probability,

$$p_X(t) = \kappa_l t \sum_{n>0} \lambda_n \left[1 - \frac{\sin(\delta_n t)}{\delta_n t} \right]. \quad (2.2)$$

Here $\lambda_n = \sum_{\pm} |\langle \phi_n^\pm | \alpha, 1 \rangle|^2 / 2$ where $|\alpha, 1\rangle$ is the first displaced Fock state $\hat{D}(\alpha) |\hat{n} = 1\rangle$, i.e. the state obtained at first order under thermal excitation or photon dephasing when starting at $|\alpha\rangle$ in the cat qubit codespace. In this sense, $\kappa_l \lambda_n$ represents the rate at which each pair of Kerr excited eigenstates is populated by codespace leakage. In addition, the right-hand side bracket of Equation (2.2) represents how the dephasing of the even and odd parity eigenstates, accumulated over time t before getting back to the codespace, translates to bit-flip errors. This bracket vanishes for small time or energy level spacing $\delta_n t \ll 1$ and equals 1 for large time or energy level spacing $\delta_n t \gg 1$, as expected.

With this analytical formula, we can put some numbers on bit-flip error probabilities under Kerr confinement by looking at Figure 2.2 in which panels (a) and (b) show respectively δ_n and λ_n with respect to $|\alpha|^2$ for the first pairs of Kerr eigenstates, while panel (c) shows each term of Equation (2.2) with respect to time for $|\alpha|^2 = 4, 8$ and 12. Together, these figures show how the various eigenstates contribute to bit-flip errors via Equation (2.2) for each mean number of photons and each time. Note thus that according to Figure 2.2(b), a distribution over all Kerr eigenstates is induced by a single thermal photon or dephasing perturbation. In turn, the associated δ_n increases with n and features two regimes as a function of α (constant, then exponentially suppressed) as shown on Figure 2.2(a); a log-scale version of the

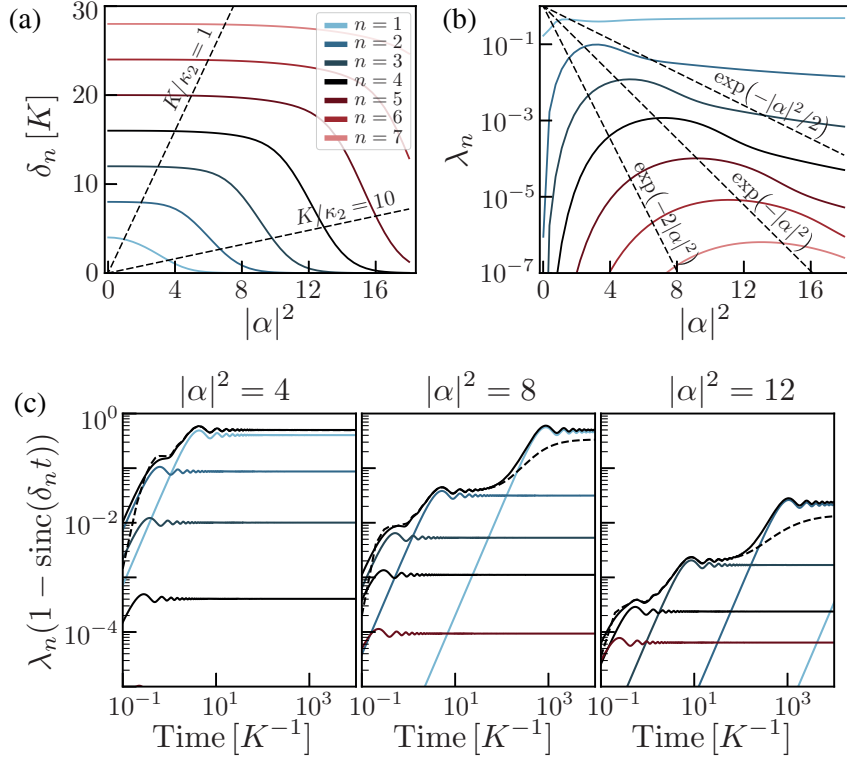


Figure 2.2: (a) Energy level spacing between even and odd parity eigenstates of the Kerr Hamiltonian $\delta_n = |e_n^- - e_n^+|$. Dashed lines represent the threshold $\delta_n = \kappa_{\text{conf}}$, below which eigenstates benefit from the additional two-photon dissipation scheme. (b) Overlap between the n -th pair of excited Kerr eigenstates and a cat qubit state deformed by thermal or dephasing noise, $\lambda_n = \sum_{\pm} |\langle \phi_n^{\pm} | \alpha, 1 \rangle|^2 / 2$. Dashed lines represent the threshold $\kappa_l \lambda_n = \kappa_l \exp(-c|\alpha|^2)$ under which exponential suppression of bit-flip errors is retrieved with suppression exponent c . (c) Individual terms of the sum of (2.2) with respect to time. The solid black line represents the sum of all terms. The dashed black line represents $p_X(t)/\kappa_l t$ extracted from a numerical simulation of Kerr confinement with single-photon loss $\kappa_1 = 10^{-3}K$ and thermal noise $n_{\text{th}} = 10^{-2}$.

latter is also shown in Figure 2.3 for better readability. The overall bit-flip rate results from the combination of these two effects.

For instance, at $|\alpha|^2 = 12$ and considering a typical time $t \sim 10/K$, Figure 2.2(c) shows that the $n = 3$ pair of excited Kerr eigenstates contributes the most to the sum of Equation (2.2). Indeed, we find that $\delta_3 t \approx 5.2$, thus contributing to a bit-flip error dominated by $\kappa_l t \lambda_3$ with $\lambda_3 \approx 1.7 \cdot 10^{-3}$. In comparison, the dissipative cat qubit error probability is also proportional to $\kappa_l t$ but with an exponential prefactor $\exp(-2|\alpha|^2) \approx 4 \cdot 10^{-11}$. Terms with lower n indices contribute less to Kerr-induced bit-flip errors because the δ_n

are well inside the exponentially suppressed regime, e.g. for the second pair of excited states, $1 - \text{sinc}(\delta_2 t) \approx 2.5 \cdot 10^{-4}$. On the other hand, terms with larger n indices contribute with $1 - \text{sinc}(\delta_n t) \approx 1$ even at shorter time scales, but with smaller effect since λ_n decreases exponentially with n [Putterman *et al.* 2021]. The study of [Putterman *et al.* 2021] shows that at very large values of $|\alpha|^2$ (well above the average photon numbers considered here) one can hope to retrieve an approximate exponential suppression given by $p_X \propto \exp(-0.94|\alpha|^2)$ which is still less significant than the prefactor $\exp(-2|\alpha|^2)$ achieved for dissipative cat qubits and with all values of $|\alpha|^2$. Note furthermore that, working at very large values of $|\alpha|^2$ induces significant phase-flip errors and can lead to several other unanticipated difficulties [Chamberland *et al.* 2022].

2.2.3 Estimating bit-flip errors for an idling qubit

In this section, we consider an idling cat qubit subject to one of the confinement schemes discussed previously, and derive a simplified model for the estimation of bit-flip errors that leads to Equation (2.2).

2.2.3.1 General Framework

Let us consider a quantum harmonic oscillator initialized in the pure state

$$\boldsymbol{\rho}(t=0) = |0_L\rangle\langle 0_L| = |\alpha\rangle\langle\alpha| + \mathcal{O}(e^{-2|\alpha|^2}). \quad (2.3)$$

Following the definition of Chapter 1, the bit-flip error probability after an idling time t is given by

$$p_X(t) = (1 - \text{Tr}[\mathbf{J}_z \boldsymbol{\rho}(t)]) / 2 \quad (2.4)$$

To compute this bit-flip error probability, let us estimate $\boldsymbol{\rho}(t)$ at any given time, and then reinsert it in the above formula. In order to simplify the analytical derivation of $\boldsymbol{\rho}(t)$, we will neglect the exponential corrections of Equation (2.3) and assume that the initial state is coherent. Any resulting bit-flip estimation will thus be correct up to this exponential correction. By symmetry, we can then treat the case in which the system is initialized in the $|1_L\rangle$ logical state.

As previously, we assume that the system is subject to a cat qubit confinement scheme and to usual decoherence effects such as single-photon relaxation, thermal excitation and pure dephasing,

$$\frac{d\boldsymbol{\rho}}{dt} = \mathcal{L}_{\text{conf}}\boldsymbol{\rho} + \kappa_- \mathcal{D}[\mathbf{a}]\boldsymbol{\rho} + \kappa_+ \mathcal{D}[\mathbf{a}^\dagger]\boldsymbol{\rho} + \kappa_\phi \mathcal{D}[\mathbf{a}^\dagger \mathbf{a}]\boldsymbol{\rho} \quad (2.5)$$

where $\mathcal{L}_{\text{conf}}$ is a confinement superoperator with confinement rate g , while $\kappa_- = \kappa_1(1 + n_{\text{th}})$ and $\kappa_+ = \kappa_1 n_{\text{th}}$ are the single-photon relaxation rates, n_{th}

the average number of thermal photons at the resonator frequency and κ_ϕ is the pure dephasing noise rate. At all times, let us separate the density matrix in two according to

$$\boldsymbol{\rho}(t) = (1 - \varepsilon_l(t)) |\alpha\rangle\langle\alpha| + \boldsymbol{\rho}_l(t). \quad (2.6)$$

Thus, $\boldsymbol{\rho}_l$ is the part of the density matrix that represents population leaked outside the coherent subspace $|\alpha\rangle\langle\alpha|$. It is such that $\text{Tr}[\boldsymbol{\rho}_l(t)] = \varepsilon_l(t)$. Suppose decoherence rates are small compared to the confinement amplitude, $\kappa_+, \kappa_-, \kappa_\phi \ll g$, and the time of evolution is short compared to the time scale of population leakage out of the coherent state, $\kappa_l t \ll 1$ where $\kappa_l = \kappa_+ + |\alpha|^2 \kappa_\phi$ is the leakage rate. In this limit, leakage is small and $|\varepsilon_l(t)| \ll 1$. We first differentiate Equation (2.6) with respect to time, and second reinsert Equation (2.6) in Equation (2.5), to obtain the two following equations

$$\frac{d\boldsymbol{\rho}}{dt} = -\dot{\varepsilon}_l |\alpha\rangle\langle\alpha| + \frac{d\boldsymbol{\rho}_l}{dt} \quad (2.7)$$

and

$$\begin{aligned} \frac{d\boldsymbol{\rho}}{dt} &= \mathcal{L}_{\text{conf}} \boldsymbol{\rho}_l + \kappa_- \mathcal{D}[\mathbf{a}] |\alpha\rangle\langle\alpha| + \kappa_+ \mathcal{D}[\mathbf{a}^\dagger] |\alpha\rangle\langle\alpha| + \kappa_\phi \mathcal{D}[\mathbf{a}^\dagger \mathbf{a}] |\alpha\rangle\langle\alpha| + \mathcal{O}(\varepsilon_l^2) \\ &= \mathcal{L}_{\text{conf}} \boldsymbol{\rho}_l + \kappa_l (|\alpha, 1\rangle\langle\alpha, 1| - |\alpha\rangle\langle\alpha|) - \frac{\alpha}{2} (\kappa_1 + \kappa_\phi) (|\alpha, 1\rangle\langle\alpha| + |\alpha\rangle\langle\alpha, 1|) \\ &\quad - |\alpha|^2 \kappa_\phi (|\alpha, 2\rangle\langle\alpha| + |\alpha\rangle\langle\alpha, 2|) + \mathcal{O}(\varepsilon_l^2) \end{aligned} \quad (2.8)$$

where $|\alpha, n\rangle = D(\alpha)(\mathbf{a}^\dagger)^n |0\rangle$ is the n -th displaced Fock state, and where we have used that $\mathcal{L}_{\text{conf}} |\alpha\rangle\langle\alpha| = 0$ by definition. At this point, we project Equations (2.7) and (2.8) inside the $|\alpha\rangle\langle\alpha|$ subspace, which yields

$$\dot{\varepsilon}_l = \kappa_l - \langle\alpha| \mathcal{L}_{\text{conf}} \boldsymbol{\rho}_l |\alpha\rangle. \quad (2.9)$$

We have thus found the true rate of leakage outside of the codespace, $\dot{\varepsilon}_l$, which is bounded by κ_l . In particular, for a purely Hamiltonian confinement mechanism, we have $\mathbf{H}_{\text{conf}} |\pm\alpha\rangle = 0$ by definition such that

$$\langle\alpha| \mathcal{L}_{\text{conf}} \boldsymbol{\rho}_l |\alpha\rangle = -i \langle\alpha| (\mathbf{H}_{\text{conf}} \boldsymbol{\rho}_l - \boldsymbol{\rho}_l \mathbf{H}_{\text{conf}}) |\alpha\rangle = 0. \quad (2.10)$$

Therefore, Equation (2.9) simplifies to $\dot{\varepsilon}_l = \kappa_l$. For a purely Hamiltonian confinement, all leakage out of the codespace due to dissipative noise builds up and never reconverges back to the codespace. Projecting Equations (2.7) and (2.8) outside of the $|\alpha\rangle\langle\alpha|$ subspace further yields the master equation on the leaked density matrix. For a purely Hamiltonian confinement, it reduces to

$$\begin{aligned} \frac{d\boldsymbol{\rho}_l}{dt} &= -i [\mathbf{H}_{\text{conf}}, \boldsymbol{\rho}_l] + \kappa_l |\alpha, 1\rangle\langle\alpha, 1| \\ &\quad - \frac{\alpha}{2} (\kappa_1 + \kappa_\phi) (|\alpha, 1\rangle\langle\alpha| + |\alpha\rangle\langle\alpha, 1|) - |\alpha|^2 \kappa_\phi (|\alpha, 2\rangle\langle\alpha| + |\alpha\rangle\langle\alpha, 2|). \end{aligned} \quad (2.11)$$

2.2.3.2 Kerr Confinement

We now consider a Kerr Hamiltonian confinement such that $\mathbf{H}_{\text{conf}} = \mathbf{H}_{\text{Kerr}}$, and assume the dynamics of Equation (2.11). To investigate the associated evolution of ρ_l , we expand quantum operators into the Kerr Hamiltonian eigenbasis with eigenstates $\{|\phi_n^\pm\rangle\}_{n=0}^\infty$ and eigenenergies $\{e_n^\pm\}_{n=0}^\infty$. This basis has a degenerate ground eigenspace of energy $e_0^\pm = 0$ such that $|\phi_0^\pm\rangle = |\mathcal{C}_\alpha^\pm\rangle$ is a cat qubit basis state. All other eigenstates are separated in two branches of even and odd parities with $0 < e_n^- - e_n^+ \ll e_{n+1}^\pm - e_n^\pm$. However, $e_n^- - e_n^+$ grows unbounded in n . These energy level spacings $\delta_n = |e_n^- - e_n^+|$ are represented in vertical axis logarithmic scale on Figure 2.3. The expansion thus yields

$$|\alpha\rangle = (\mathcal{N}_+ |\phi_0^+\rangle + \mathcal{N}_- |\phi_0^-\rangle)/2 \quad (2.12a)$$

$$|\alpha, 1\rangle = \sum_{n>0} (\lambda_n^+ |\phi_n^+\rangle + \lambda_n^- |\phi_n^-\rangle) \quad (2.12b)$$

$$\rho_l(t) = \sum_{n,m} \sum_{s,r=\pm} \tau_{nm}^{sr}(t) |\phi_n^s\rangle\langle\phi_m^r| \quad (2.12c)$$

where $\mathcal{N}_\pm = \sqrt{2(1 \pm e^{-2|\alpha|^2})}$ and τ_{nm}^{sr} is the quantity to be investigated. The λ_n^\pm coefficients are defined by (2.12b) and are real-valued for α real. Reinjecting in Equation (2.11) and projecting onto the $|\phi_n^s\rangle\langle\phi_m^r|$ density matrix element yields (for $n + m > 0$)

$$\frac{d}{dt} \tau_{nm}^{sr} = -i\tau_{nm}^{sr}(e_n^s - e_m^r) + \kappa_{nm}^{sr} \quad (2.13)$$

where κ_{nm}^{sr} is the rate at which the corresponding density matrix coefficient is populated, and given by $\kappa_{nn}^{sr} = \kappa_l \lambda_n^s \lambda_n^r$ for $n = m$. We can solve each equation individually for null initial conditions, yielding

$$\tau_{nn}^{ss}(t) = (\lambda_n^s)^2 \kappa_l t \quad (2.14a)$$

$$\tau_{nm}^{sr}(t) = i\kappa_{nm}^{sr} \frac{e^{-i(e_n^s - e_m^r)t} - 1}{e_n^s - e_m^r}. \quad (2.14b)$$

In this Kerr eigenbasis expansion of ρ_l , the dominating terms are the diagonal terms τ_{nn}^{ss} and the terms $\tau_{nn}^{s\bar{s}}$ where $\bar{s} = -s$. All other terms are much smaller due to their $1/(e_n^s - e_m^r)$ dependence and we therefore neglect them. The density matrix ρ_l is thus almost block diagonal in the Kerr eigenbasis, with blocks of the form

$$\begin{bmatrix} \tau_{nn}^{++}(t) & \tau_{nn}^{+-}(t) \\ \tau_{nn}^{-+}(t) & \tau_{nn}^{--}(t) \end{bmatrix} = \kappa_l t \begin{bmatrix} (\lambda_n^+)^2 & \lambda_n^+ \lambda_n^- \frac{e^{i\delta_n t} - 1}{i\delta_n t} \\ \lambda_n^- \lambda_n^+ \frac{e^{-i\delta_n t} - 1}{-i\delta_n t} & (\lambda_n^-)^2 \end{bmatrix} \quad (2.15)$$

where $\delta_n = e_n^- - e_n^+ > 0$. After idling time $T \sim 1/\delta_n$, coherences of this diagonal block have flipped sign thus inducing bit flip. The bit-flip error is




Figure 2.3: Energy level spacing between even and odd parity eigenstates of the Kerr Hamiltonian $\delta_n = |e_n^- - e_n^+|$. Similar plot as in Figure 2.2(a) but with a logarithmic scale on the vertical axis.

then given by

$$\begin{aligned}
p_X(t) &= \frac{1}{2} \left(1 - \text{Tr}[\hat{J}_z \boldsymbol{\rho}] \right) \\
&= \frac{1}{2} \left(1 - (1 - \kappa_l t) \text{Tr}[\hat{J}_z |\alpha\rangle\langle\alpha|] - \text{Tr}[\hat{J}_z \boldsymbol{\rho}_l] \right) \\
&= \frac{1}{2} \left(\kappa_l t - \text{Tr}[\hat{J}_z \boldsymbol{\rho}_l] + \mathcal{O}(e^{-2\alpha^2}) \right) \\
&\approx \frac{1}{2} \kappa_l t - \frac{1}{2} \sum_{n>0} \left(\tau_{nn}^{--}(t) \langle \phi_n^+ | \hat{J}_{+-} | \phi_n^- \rangle + \text{h.c.} \right) \\
&\approx \frac{1}{2} \kappa_l t - \frac{1}{2} \sum_{n>0} \left(\tau_{nn}^{--}(t) + \text{h.c.} \right) \\
&\approx \frac{1}{2} \kappa_l t \sum_{n>0} \left((\lambda_n^+)^2 + (\lambda_n^-)^2 - 2\lambda_n^+ \lambda_n^- \frac{\sin(\delta_n t)}{\delta_n t} \right) \\
&\approx \kappa_l t \sum_{n>0} \lambda_n \left(1 - \frac{\sin(\delta_n t)}{\delta_n t} \right)
\end{aligned} \tag{2.16}$$

where $\lambda_n = [(\lambda_n^+)^2 + (\lambda_n^-)^2]/2$ is the value used and shown in the previous section. In this derivation, three different assumptions were made. The first is that $\text{Tr}[\hat{J}_z |\alpha\rangle\langle\alpha|] = 1$, which is valid up to corrections of order $\exp(-2|\alpha|^2)$. The other two are that $\langle \phi_n^- | \hat{J}_{+-} | \phi_n^+ \rangle = 1$ and that $(\lambda_n^+)^2 + (\lambda_n^-)^2 = 2\lambda_n^+ \lambda_n^-$. These two assumptions are valid in the limit of large α , and would otherwise only result in an additional prefactor in front of the sine term which does not change the overall estimation.

We have thus obtained the previously used analytical formula from Equation (2.2). In the following section, we move on to the discussion of combined Kerr and two-photon dissipation confinement, and investigate whether it is possible to retrieve the exponential suppression with Kerr cat qubits.

2.2.4 Combined Kerr and two-photon dissipation confinement

In order to better protect a Kerr cat qubit against thermal and dephasing noise while still benefiting from its excellent performance for gate engineering [Xu *et al.* 2022a], it is interesting to consider combining it with a dissipative confinement scheme. In this thesis, we only consider using a two-photon dissipation scheme, but our conclusions hold also for other schemes such as the colored dissipation introduced in [Putterman *et al.* 2021].

In order to recover the exponential suppression of bit-flip errors, the dissipative stabilization rate should be larger than any significant Kerr-induced dephasing between odd and even photon number subspaces. If this is verified, then leakage out of the cat qubit manifold will reconverge faster towards its initial state thanks to dissipation than it will dephase towards the other side of phase space due to Kerr effects. To make this argument clearer, let us assume a target exponential suppression of bit-flips given by $\Gamma_{\text{bit-flip}} \propto \exp(-c|\alpha|^2)$ with $c \sim 1$. Then all eigenstates with λ_n above the corresponding dashed line on Figure 2.2(b) could contribute to more bit-flip errors than the targeted exponential. To avoid these contributions, those eigenstates should dephase at a rate δ_n smaller than the dissipative confinement rate used to reconverge the state. For instance, at $|\alpha|^2 = 12$, the $n \leq 5$ eigenstates are above the line $\exp(-|\alpha|^2)$, and their largest degeneracy is $\delta_5 \approx 20K$. Therefore, the dissipative confinement rate should be $\kappa_{\text{conf}} = 4|\alpha|^2\kappa_2 \gg 20K$, or equivalently, $\kappa_2 \gg 0.42K$. This reasoning holds for any value of α and any target exponential rate. This explains how to set a lower bound on the ratio κ_2/K .

Conversely, we may start with an upper bound on κ_2/K , e.g. from the motivation of sufficiently maintaining the advantages of Kerr Hamiltonian confinement towards implementing gates. Consider for instance that we limit ourselves to $\kappa_2 \leq 0.1K$ (e.g. as proposed in [Puri *et al.* 2020]). Then Kerr-induced dephasing will be avoided only for the eigenstates below the dashed line $K/\kappa_2 = 10$ on Figure 2.2(a). At $|\alpha|^2 = 12$, these states read $n \leq 3$. Eigenstates with $n > 3$ will each contribute to bit-flip errors with a rate $\kappa_l\lambda_n$. Figure 2.2(b) indicates that $\lambda_4 \approx 2 \cdot 10^{-4}$, so we have to expect a bit-flip rate $\Gamma_{\text{bit-flip}} \gtrsim 2 \cdot 10^{-4}\kappa_l$.

Figure 2.4 shows the actual bit-flip rate when combining Kerr and dissipative confinement schemes. In these simulations, the cat qubit is idling and the amplitude of the Kerr Hamiltonian is varied. The resonator is subject to single-photon loss at rate $\kappa_1 = 10^{-3}\kappa_2$ with either thermal excitations at rate $\kappa_1 n_{\text{th}} = 10^{-5}\kappa_2$ (diamonds) or pure dephasing at rate $\kappa_\phi = 10^{-5}\kappa_2$ (circles). The noise rates considered here and in the rest of this work are within one order of magnitude of already achieved noise rates, both for Kerr [Grimm *et al.* 2020] and dissipative cat qubits [Touzard *et al.* 2018], and

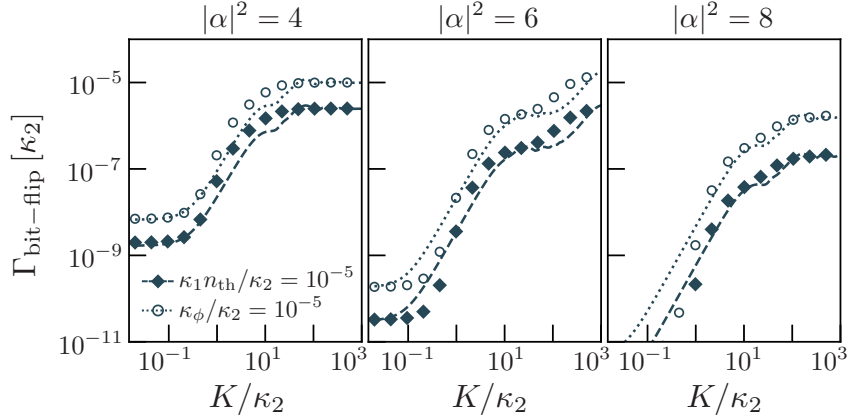


Figure 2.4: Bit-flip error rate of an idle cat qubit confined by a combined Kerr and two-photon dissipation for increasing cat sizes. The cat qubit is subject to single-photon loss of amplitude $\kappa_1 = 10^{-3}\kappa_2$, with additional thermal excitation $\kappa_1 n_{\text{th}} = 10^{-5}\kappa_2$ (diamonds, dashed lines) or pure dephasing $\kappa_\phi = 10^{-5}\kappa_2$ (circles, dotted lines). At low enough K , the exponential bit-flip suppression is restored. Markers indicate numerical simulation data, lines represent Equation (2.17).

are expected to be in a feasible range for the next generation of cat qubit experiments. In this Figure, each plot shows a different number of photons $|\alpha|^2 = 4, 6, 8$. In the $K/\kappa_2 \ll 1$ regime, the dissipative cat qubit regime is retrieved with a clear exponential suppression at rate $\exp(-2|\alpha|^2)$. In the opposite regime of $K/\kappa_2 \gg 1$, the Kerr cat qubit regime is retrieved. Between these two regimes, a smooth transition occurs. The analytical predictions represented with dashed and dotted lines in this figure follow directly from Equation (2.2) evaluated at the Kerr excited level characteristic lifetime $t = 1/\kappa_{\text{conf}}$, giving

$$\begin{aligned} \Gamma_{\text{bit-flip}} = & \kappa_1 |\alpha|^2 \exp(-4|\alpha|^2) \\ & + \kappa_l \exp(-2|\alpha|^2) \\ & + \kappa_l \sum_{n>0} \lambda_n \left[1 - \frac{\sin(\delta_n/\kappa_{\text{conf}})}{\delta_n/\kappa_{\text{conf}}} \right]. \end{aligned} \quad (2.17)$$

The first two terms are the contribution of two-photon dissipation, whose induced bit-flip errors are exponentially suppressed in $|\alpha|^2$, while the third term is inherited from the trade-off between Kerr dephasing and dissipative confinement.

The purely dissipative stabilization scheme is ideal for an idling qubit under Markovian noise, as shown here, but Hamiltonian confinement will show its benefits as we add gates or small Hamiltonian perturbations. From this perspective, an ideal working point for the combined confinement scheme would

be the largest value of K for which the dissipative exponential suppression remains intact. For the $|\alpha|^2$ values shown here and as further argued in the following section, this corresponds approximately to $K/\kappa_2 \sim 0.3$. Unfortunately, at this working point the Kerr non-linearity turns out to be too small to provide any practical advantage over a purely dissipative scheme. As we have tried to explain, the deterioration of bit-flip protection for larger K values is partly inherited from the increasing energy level spacing δ_n between even and odd photon number eigenstates of the Kerr Hamiltonian. Indeed, to avoid bit-flips, the dissipative stabilization scheme has to counter the increasing dephasing rates associated to δ_n , for any eigenstate n having significant overlap with a slightly deformed/displaced cat state.

In the following section, we introduce a new cat qubit confinement Hamiltonian that is Kerr-like, but that features bounded energy level spacing between even and odd photon number eigenstates. We observe that this new confinement scheme is indeed much more compatible with a dissipative stabilization. It is furthermore very similar in experimental implementation to two-photon dissipation, thus opening the door to a natural combination of Hamiltonian and dissipative cat qubit confinement.

2.3 Combined confinement with two-photon exchange

2.3.1 The two-photon exchange Hamiltonian

Two-photon dissipation can be engineered by coupling the cat qubit resonator with a two-photon exchange Hamiltonian $g_2(\mathbf{a}^2\mathbf{b}^\dagger + \mathbf{a}^{\dagger 2}\mathbf{b})$ to a low-Q buffer mode, that is a resonator \mathbf{b} undergoing strong single-photon loss [Leghtas *et al.* 2015]. When single-photon relaxation κ_b of the buffer mode is large compared to the two-photon exchange rate, the buffer mode can be adiabatically eliminated thus resulting in two-photon dissipation $\kappa_2\mathcal{D}[\mathbf{a}^2]$ on the cat qubit mode, with rate $\kappa_2 = 4g_2^2/\kappa_b$ [Azouit *et al.* 2017]. To engineer the additional $-\alpha^2$ term of the two-photon dissipator, one can drive the buffer mode at its resonance with Hamiltonian $-\alpha^2 g_2(\mathbf{b} + \mathbf{b}^\dagger)$, leading effectively to a two-photon drive on the cat qubit mode.

Instead of coupling the cat qubit resonator to a low-Q buffer mode, it is interesting to consider what would happen using a high-Q and highly anharmonic buffer mode. Considering only the two lowest energy levels of the buffer mode, the Two-Photon Exchange (TPE) Hamiltonian resulting from such a setup reads

$$\mathbf{H}_{\text{TPE}} = g_2(\mathbf{a}^2 - \alpha^2)\sigma_+ + g_2^*(\mathbf{a}^{\dagger 2} - \alpha^{*2})\sigma_- \quad (2.18)$$




Figure 2.5: (a) Energy spectrum of the TPE Hamiltonian separated in two branches of even (left) and odd (right) parity eigenstates. The zero-energy eigenspace is the cat qubit manifold, gapped from the rest of the spectrum by an energy $|E_1^\pm - E_0| \approx 2|\alpha|g_2$. The spectrum is symmetric with respect to the zero-energy point. (b) Schematic of a combined TPE and two-photon dissipation confinement scheme. The cat qubit resonator (blue) is coupled to both a low-Q buffer mode (green) and an anharmonic high-Q buffer mode (red) through an element that exchanges pairs of photons of the cat qubit mode with single photons from either buffer mode.

where σ_\pm are the lowering and raising operators of the two-level buffer system. The TPE Hamiltonian features an exactly degenerate zero-energy eigenspace given by the cat qubit codespace with the buffer in its ground state, $\text{span}\{|\Phi_0^+\rangle = |\mathcal{C}_\alpha^+\rangle |g\rangle, |\Phi_0^-\rangle = |\mathcal{C}_\alpha^-\rangle |g\rangle\}$. It also preserves the photon number parity in the cat qubit resonator, such that its eigen-spectrum can be separated in two branches of even and odd parities. Solving the eigenvalue equation $\mathbf{H}_{\text{TPE}} |\Phi_n^\pm\rangle = E_n^\pm |\Phi_n^\pm\rangle$, the other eigenstates of the TPE Hamiltonian are found to be

$$\begin{aligned} |\Phi_n^\pm\rangle &= \frac{1}{\sqrt{2}} (|\phi_n^\pm\rangle |g\rangle + |\psi_n^\pm\rangle |e\rangle), \quad E_n^\pm/g_2 = \sqrt{|e_n^\pm|/K} \\ |\Phi_{-n}^\pm\rangle &= \frac{1}{\sqrt{2}} (|\phi_n^\pm\rangle |g\rangle - |\psi_n^\pm\rangle |e\rangle), \quad E_{-n}^\pm/g_2 = -\sqrt{|e_n^\pm|/K} \end{aligned} \quad (2.19)$$

where $\{|\phi_n^\pm\rangle\}_{n=0}^\infty$ are excited eigenstates of the Kerr Hamiltonian $\mathbf{H}_{\text{Kerr}} = -K(\mathbf{a}^{\dagger 2} - \alpha^{*2})(\mathbf{a}^2 - \alpha^2)$ associated to eigenvalues e_n^\pm , thus $\mathbf{H}_{\text{Kerr}} |\phi_n^\pm\rangle = e_n^\pm |\phi_n^\pm\rangle$. Furthermore, $\{|\psi_n^\pm\rangle\}_{n=1}^\infty$ are eigenstates of the reversed Kerr Hamiltonian $\mathbf{H}'_{\text{Kerr}} = -K(\mathbf{a}^2 - \alpha^2)(\mathbf{a}^{\dagger 2} - \alpha^{*2})$ such that $\mathbf{H}'_{\text{Kerr}} |\psi_n^\pm\rangle = e_n^\pm |\psi_n^\pm\rangle$, with the same eigenvalues. Eigenstates of these two Hamiltonians are linked by the relation $|\psi_n^\pm\rangle = (\mathbf{a}^2 - \alpha^2) |\phi_n^\pm\rangle / \sqrt{e_n^\pm}$. The spectrum of the TPE Hamiltonian, depicted on Figure 2.5(a) for $|\alpha|^2 = 8$, is thus closely related to the Kerr

spectrum through (2.19).

First and foremost, the TPE Hamiltonian thus inherits an energy gap of approximately $2|\alpha|g_2$ between the cat space and the rest of the spectrum. As such, it can be used for Hamiltonian cat qubit confinement. In the notations of Equation (2.1), the TPE confinement super-operator reads $\mathcal{L}_{\text{conf}} = -i[\mathbf{H}_{\text{TPE}}, \cdot]$ with confinement rate $g \equiv g_2$. While in our analysis and simulations we consider the buffer mode to be a two-level system, in practice, it can be realized using a sufficiently anharmonic mode. As shown in Section 2.5, this anharmonicity should be large enough compared to the TPE Hamiltonian gap to avoid the appearance of spurious eigenstates inside the gap.

Second, like for the Kerr Hamiltonian, excited eigenstates of even and odd photon number parity have different energies, with each spacing closing exponentially at sufficiently large $|\alpha|$. This contributes to bit-flip errors via the same mechanism as the Kerr Hamiltonian, i.e. with an expression similar to (2.2). However, unlike the Kerr Hamiltonian, the associated energy gaps $\Delta_n = \sqrt{e_n^-} - \sqrt{e_n^+}$ for the TPE Hamiltonian remain uniformly bounded, as shown on Figure 2.6(a). This can be understood as a direct consequence of the square root: denoting $e_n^\pm = \bar{e}_n \pm \delta_n/2$ and considering $\bar{e}_n \gg \delta_n$ to expand the square root, we get $\Delta_n/g_2 \approx \delta_n/2\sqrt{K|\bar{e}_n|}$; thus δ_n/K increasing more slowly than $\sqrt{|\bar{e}_n|/K}$ is sufficient to keep Δ_n bounded. Figure 2.6(b) then shows the associated relative overlap between a cat qubit state deformed by leakage and the n -th pair of TPE eigenstates. It features a similar profile as with in Kerr Hamiltonian case.

2.3.2 Combined TPE and two-photon dissipation

The boundedness of Δ_n has major consequences when considering the combination of Hamiltonian confinement with two-photon dissipative confinement, allowing us to fully benefit from the advantages of both schemes. Indeed, consider the bit-flip rate estimation

$$\begin{aligned} \Gamma_{\text{bit-flip}} &= \kappa_1 |\alpha|^2 \exp(-4|\alpha|^2) \\ &\quad + \kappa_l \exp(-2|\alpha|^2) \\ &\quad + \kappa_l \sum_{n \neq 0} \Lambda_n \left[1 - \frac{\sin(\Delta_n/\kappa_{\text{conf}})}{\Delta_n/\kappa_{\text{conf}}} \right] \end{aligned} \tag{2.20}$$

which is the equivalent of (2.17) for TPE confinement. Similarly, $\kappa_l \Lambda_n$ is the rate at which pairs of TPE eigenstates are populated by codespace leakage, where $\Lambda_n = \sum_{\pm} |\langle \Phi_n^\pm | \alpha, 1 \rangle|^2 / 2 = \lambda_n / 2$. Indeed, these overlaps are exactly half of their Kerr equivalents but they are twice as many due to the symmetry of the TPE energy spectrum (see Section 2.2.3).




Figure 2.6: (a) Energy level spacing between even and odd parity eigenstates of the TPE Hamiltonian, $\Delta_n = E_n^- - E_n^+$. Dashed lines represent the threshold $\Delta_n = \kappa_{\text{conf}}$, below which eigenstates benefit from the additional two-photon dissipation scheme. (b) Overlap between the n -th pair of excited TPE eigenstates and a cat qubit state deformed by thermal or dephasing noise, $\Lambda_n = \sum_{\pm} |\langle \Phi_n^{\pm} | \alpha, 1 \rangle|^2 / 2$. Dashed lines represent the threshold $\kappa_l \Lambda_n = \kappa_l \exp(-c|\alpha|^2)$ under which exponential suppression of bit-flip errors is retrieved with suppression exponent c . (c) Individual terms of the sum of (2.20) with respect to time. The solid black line represents the sum of all terms. The dashed black line represents $p_X(t)/\kappa_l t$ extracted from a numerical simulation of TPE confinement with single-photon loss $\kappa_1 = 10^{-3}g_2$ and thermal noise $n_{\text{th}} = 10^{-2}$.

Contrarily to the Kerr case, the last bracket can now be made small uniformly for all n by taking $\kappa_{\text{conf}} \gg g_2$. This criterion recasts into $g_2/\kappa_2 \ll 4|\alpha|^2$. Concretely, whereas we have argued that an exponential suppression of bit-flip errors at a rate $\exp(-|\alpha|^2)$ and for $|\alpha|^2 = 12$ requires $\kappa_2 \gg 0.42K$ for a Kerr confinement (see Section 2.2.4), the same kind of computation leads to a requirement of $\kappa_2 \gg .014g_2$ in the case of TPE confinement.

These arguments are validated by numerical simulations. Figure 2.7 shows the bit-flip rate associated to the combined TPE and two-photon dissipation




Figure 2.7: Bit-flip error rate of an idle cat qubit confined by a combined TPE Hamiltonian and two-photon dissipation scheme for increasing cat sizes. The cat qubit is subject to single-photon loss of amplitude $\kappa_1 = 10^{-3}\kappa_2$, with additional thermal excitations $\kappa_1 n_{\text{th}} = 10^{-5}\kappa_2$ (diamonds, dashed lines) or pure dephasing $\kappa_\phi = 10^{-5}\kappa_2$ (circles, dotted lines). The exponential bit-flip suppression is restored at low enough g_2 , an order of magnitude earlier than for the Kerr Hamiltonian. Markers indicate numerical simulation data, lines represent the formula (2.20).

confinement schemes. Similarly to the combined Kerr scheme shown in Figure 2.4, the cat qubit is idling and the amplitude of the TPE confinement is varied. The same noise parameters and number of photons are used. The markers indicate numerical simulation data, the analytical predictions represented with dashed and dotted lines correspond to (2.20), which appear to match quite well. Qualitatively, the plots resemble Figure 2.4 of the Kerr Hamiltonian. For $g_2/\kappa_2 \ll 1$, the purely dissipative cat qubit regime is retrieved with a clear exponential suppression at rate $\exp(-2|\alpha|^2)$. In the opposite regime of $g_2/\kappa_2 \gg 1$, the TPE cat qubit regime is retrieved. Between these two regimes, a smooth transition occurs. Quantitatively however, the transitions occur at much larger g_2/κ_2 values, as the horizontal scale on Figure 2.7 is shifted to the left by more than one order of magnitude.

Another way to look at the data is presented in Figure 2.8. The vertical axis represents the exponential suppression factor γ such that $\Gamma_{\text{bit-flip}} \propto \exp(-\gamma|\alpha|^2)$, for an exponential fit over the range $2 \leq |\alpha|^2 \leq 12$. This exponential suppression factor is evaluated as a function of g_2/κ_2 and K/κ_2 respectively for the TPE and Kerr Hamiltonian confinements. Both cases are thus simulated in combination with dissipative confinement, and with the same error channels as on the previous figures. For both combined Kerr and combined TPE confinement, a smooth transition is observed from the $\gamma \geq 2$ regime (purely dissipative confinement) to the $\gamma \rightarrow 0$ regime (bit-flip




Figure 2.8: Exponential suppression factor of bit-flip errors γ , such that $\Gamma_{\text{bit-flip}} \propto \exp(-\gamma|\alpha|^2)$, for an idle cat qubit confined with a combined dissipative and Hamiltonian scheme, for either a TPE Hamiltonian (red) or a Kerr Hamiltonian (dark blue). The γ value is computed from an exponential fit over the range $2 \leq |\alpha|^2 \leq 12$. The cat qubit is subject to single-photon loss of amplitude $\kappa_1 = 10^{-3}\kappa_2$, in addition with thermal excitations $\kappa_1 n_{\text{th}} = 10^{-5}\kappa_2$ (diamonds) or pure dephasing $\kappa_\phi = 10^{-5}\kappa_2$ (circles). Dashed lines show the approximate boundary of the $\gamma \geq 2$ exponential suppression regime, at $g_2/\kappa_2 = 10$ and $K/\kappa_2 = 0.3$ respectively.

of order κ_l induced by strong Hamiltonian contribution in Equations (2.17) and (2.20)). Working points which maximize the Hamiltonian confinement without compromising the bit-flip protection can be identified just before this transition starts, respectively at $g_2/\kappa_2 \approx 10$ and $K/\kappa_2 \approx 0.3$. As already discussed, the Kerr Hamiltonian working point is significantly lower than the TPE one due to the worse scaling of its energy level spacing.

Note that the Hamiltonian gaps of the Kerr and TPE Hamiltonian scale as $4|\alpha|^2$ and $2|\alpha|$ respectively. At the working points computed here, this implies that the TPE Hamiltonian gap is still 5 to 10 times larger than the Kerr one for a mean number of photon $2 \leq |\alpha|^2 \leq 12$. This should enable us to truly benefit from the combination of both confinement strategies, namely retaining the exponential bit-flip suppression induced by the dissipative stabilization, with a TPE Hamiltonian strong enough to drastically improve the speed of quantum gates. Simulations in the following sections further indicate that gate performances put K and g_2 on an equal footing instead of the Hamiltonian gaps, thus further favoring the TPE confinement scheme.

Before ending this section, we must say a few words about the experimental implementation of the TPE confinement. The TPE Hamiltonian is also interesting in this respect. Its similarity to the two-photon dissipation indicates that they can be engineered simultaneously and through simple modifications of the dissipative setup [Lescanne *et al.* 2020b]. Figure 2.5(b) shows

a schematic of a cat qubit with both types of confinement schemes. The resonator that encodes the cat qubit (blue) is coupled to a low-Q buffer (green) and to a nonlinear high-Q buffer (red) through a two photon exchanger. A single element can be used to engineer the two TPE couplings, with the low-Q and high-Q buffers, thus inducing a combined dissipative-conservative confinement which has the benefit of protecting against leakage from the code space due to thermal and dephasing noise, but also against spurious Hamiltonians with energy scales small compared to the TPE energy gap.

In the following sections, bias-preserving single-qubit Z gates and two-qubit CNOT gates with combined confinement schemes are investigated at the working points identified on Figure 2.8. The Hamiltonian confinements will then show their benefits in rejecting the non-ideal terms present in the gate Hamiltonians, while the dissipative confinement maintains protection against perturbation channels like κ_1 and κ_ϕ . Note that the addition of a buffer mode into the system leads to additional noise sources for TPE Hamiltonian confinement. The effect of these sources is thoroughly investigated in Section 2.5, confirming that our conclusions should still hold under realistic conditions.

2.4 Engineering gates under combined confinement

In quantum computing architectures with biased-noise qubits, it is essential that the error bias is conserved throughout the operation of the device. A quantum gate with cat qubits should therefore preserve the exponential suppression of bit-flip errors. In this section, we focus on the bias-preserving realization of single-qubit Z gates and two-qubit CNOT gates [Guillaud & Mirrahimi 2019, Puri *et al.* 2020]. Similarly to [Guillaud & Mirrahimi 2019], the proposed CNOT gate can be extended to a three qubit CCNOT (Toffoli) gate which together with appropriate state preparation and measurements provide the building blocks of a hardware-efficient fault-tolerant universal quantum computer based on repetition cat qubits [Guillaud & Mirrahimi 2019, Chamberland *et al.* 2022]. Single-qubit X gates of angle π are also required for universality but are trivial to implement either in software, by commuting the Pauli gate with the circuit, or in hardware, with a half-period delay on the oscillator.

2.4.1 Zeno Z gate

In this section, we focus on the case of single-qubit Z gate with a combined dissipative and Hamiltonian confinement. Section 2.4.2 introduces a new design

for the CNOT gate under such combined confinement. The addition of a Kerr or TPE confinement compared to fully dissipative cat qubits offers the perspective of improved gate speed, while the two-photon dissipation stabilizes the code manifold and protects it from leakage-induced bit-flip errors.

2.4.1.1 Master equation

Similarly to [Mirrahimi *et al.* 2014, Puri *et al.* 2017] and the experimental realizations of [Touzard *et al.* 2018, Grimm *et al.* 2020], a rotation around the Z axis of the cat qubit can be performed via the application of a resonant drive on the cat qubit resonator. Indeed, taken alone this would displace the cat in phase space, but adding it to a confinement process essentially induces a phase accumulation between the different regions of phase space where the state remains confined, thus essentially between $|\alpha\rangle \approx |0_L\rangle$ and $|\alpha\rangle \approx |1_L\rangle$. In an appropriate rotating frame, the system is therefore subject to the following master equation during the gate,

$$\frac{d\rho}{dt} = -i [\varepsilon_Z(t)\mathbf{a}^\dagger + \varepsilon_Z^*(t)\mathbf{a}, \rho] - ig [\mathbf{H}_{\text{conf}}, \rho] + \kappa_2 \mathcal{D}[\mathbf{a}^2 - \alpha^2]\rho, \quad (2.21)$$

where $g\mathbf{H}_{\text{conf}} \equiv -K\mathbf{H}_{\text{Kerr}}$ or $g\mathbf{H}_{\text{conf}} \equiv g_2\mathbf{H}_{\text{TPE}}$ for combined Kerr and combined TPE confinements respectively. Furthermore $\varepsilon_Z(t)$ denotes the complex amplitude of the resonant drive that can be slowly varying in time, turning on and off the logical Z rotation.

A fully dissipative gate is retrieved at $g = 0$. This situation can be analyzed with Zeno dynamics at various orders of the small parameter $|\varepsilon_Z/\kappa_2|$. The state always remains $|\varepsilon_Z/\kappa_2|$ -close to the cat qubit subspace and, as argued in [Guillaud & Mirrahimi 2019], the bit-flip errors remain exponentially suppressed during the gate. At first order, the ε_Z drive induces an effective Zeno dynamics that rotates the state of the qubit around its Z axis of the Bloch sphere at a speed given by $4\text{Re}(\alpha^*\varepsilon_Z(t))$ [Mirrahimi *et al.* 2014], thus performing the gate. At the second order, it leads to effective phase decoherence $\mathcal{D}[\sigma_z]$, in other words phase-flip errors. These so-called non-adiabatic errors thus scale as $T\varepsilon_Z^2/\kappa_2 \propto 1/(T\kappa_2)$ since $\varepsilon_Z \propto 1/T$ where T is the gate time. However, typical decoherence channels on the cat qubit resonator, like the dissipation channels with κ_1 and n_{th} described in the simulations of Sections 2.2 and 2.3, induce direct phase errors, which scale linearly with the gate time. The trade-off between these two effects yields an optimal gate time at which gate fidelity is maximal [Chamberland *et al.* 2022].

A fully Hamiltonian gate is instead retrieved at $\kappa_2 = 0$. This situation can be analyzed from the viewpoint of an adiabatically varying Hamiltonian. Taking advantage of the Hamiltonian gap, the adiabatic theorem ensures that transitions outside the cat qubit subspace can be suppressed exponentially

in the gate time [Nenciu 1993], if the Hamiltonian varies smoothly enough. Since the evolution is purely Hamiltonian, this implies that gate-induced errors can be suppressed exponentially in T . The prefactors of these errors can be further improved drastically by using superadiabatic pulse designs as proposed in [Xu *et al.* 2022a]. It thus seems that gates can be performed orders of magnitudes faster with Hamiltonian confinement than with a fully dissipative confinement scheme, for which non-adiabatic errors scale linearly in $1/T$. However, as we have shown in the previous sections, a fully Hamiltonian confinement does not preserve the error bias of cat qubits under typical dissipation channels, and is therefore irrelevant on its own.

2.4.1.2 Gate errors

With a combined dissipative and Hamiltonian confinement, the phase-flip performance of Z gates can be improved without compromising bit-flip errors. From a general viewpoint, the phase flip error probability of a Z gate can be modeled as

$$p_Z = \kappa_1 |\alpha|^2 T + p_Z^{\text{NA}} \quad (2.22)$$

where the first term is identical to the idling qubit and represents errors due to the dominant dissipation channel, namely single-photon loss in the resonator mode, while p_Z^{NA} corresponds to “non-adiabatic” phase errors induced by the gate operation. For a Z gate with combined Kerr and two-photon dissipation, the gate-induced phase-flip error probability is derived in Section 2.4.1.3, assuming a constant Hamiltonian drive throughout the gate and neglecting fast transients when switching it on and off. The gate-induced phase-flip error is then given by

$$p_Z^{\text{NA}} = \frac{1}{1 + \frac{4K^2}{\kappa_2}} \frac{\theta^2}{16|\alpha|^4 \kappa_2 T} \quad (2.23)$$

where θ is the angle of the Z gate. This expression contains the dominant error, which is closer to the fully dissipative case than to the purely Hamiltonian case, with a linear scaling in $1/T$. Compared to the model of [Chamberland *et al.* 2022] for fully dissipative gates, this equation features an additional prefactor that is tunable through the ratio between Kerr non-linearity and dissipative confinement rate. Consequently, non-adiabatic phase errors are suppressed quadratically with the addition of a Kerr confinement Hamiltonian compared to the fully dissipative design. One should bear in mind though that this additional confinement, in presence of noise mechanisms such as thermal excitation and photon dephasing, induces additional bit-flip errors as discussed in previous sections, which are kept in check up to the working point provided in Figure 2.8. For larger Hamiltonian confinement rates than this working point, phase errors can be further reduced at

the cost of an increase in bit-flip errors. In this case, an optimization could be performed at the level of the logical quantum error correcting code to find an optimal working point.

For the combined TPE and two-photon dissipation confinement, the gate-induced phase-flip error probabilities are numerically fitted with

$$p_Z^{\text{NA}} = \frac{1}{1 + \frac{4g_2^2}{\kappa_2^2}} \frac{\theta^2}{16|\alpha|^4 \kappa_2 T} + \frac{\theta^2}{32|\alpha|^4 g_2^2 T^2} . \quad (2.24)$$

The first term is similar to the combined Kerr gate, with a quadratic suppression of non-adiabatic phase errors in the ratio between Hamiltonian and dissipative confinement rates. The second term appears to be specific to the TPE confinement and decreases quadratically in $g_2 T$ independently of κ_2 , as long as $g_2 \gtrsim \kappa_2$.

2.4.1.3 Analytical derivation of gate errors

In this section, we derive the phase-flip error probability caused by the operation of a Z gate under combined Hamiltonian and dissipative confinement, as discussed in the previous section. More precisely, we calculate first-order phase errors for the combined Kerr scheme. The calculation for the combined TPE scheme appear to be more complex and require higher order derivations. In this case, we only provide the formula that has been used to analytically fit the numerical results of the main text.

The master equation that describes a Z gate with a combined Kerr and dissipative confinement scheme reads

$$\frac{d\rho}{dt} = -i\varepsilon_Z [\mathbf{a} + \mathbf{a}^\dagger, \rho] + iK [(\mathbf{a}^{\dagger 2} - \alpha^{*2})(\mathbf{a}^2 - \alpha^2), \rho] + \kappa_2 \mathcal{D}[\mathbf{a}^2 - \alpha^2] \rho \quad (2.25)$$

For the purpose of the following analysis, we move into the SFB using the change of basis definition, $\mathbf{a} \rightarrow \sigma_z \otimes (\tilde{\mathbf{a}} + \alpha)$, where σ_z is the Pauli Z operator of the logical qubit and $\tilde{\mathbf{a}}$ the photon annihilation operator of the gauge mode. The master equation therefore reads

$$\begin{aligned} \frac{d\rho}{dt} = & -i\varepsilon_Z [\sigma_z(\tilde{\mathbf{a}} + \tilde{\mathbf{a}}^\dagger + 2\text{Re}[\alpha], \rho] + \kappa_2 \mathcal{D}[\tilde{\mathbf{a}}^2 + 2\alpha\tilde{\mathbf{a}}] \rho \\ & + iK [\tilde{\mathbf{a}}^{\dagger 2} \tilde{\mathbf{a}}^2 + 2\alpha^* \tilde{\mathbf{a}}^\dagger \tilde{\mathbf{a}}^2 + 2\alpha \tilde{\mathbf{a}}^{\dagger 2} \tilde{\mathbf{a}} + 4|\alpha|^2 \tilde{\mathbf{a}}^\dagger \tilde{\mathbf{a}}, \rho] \end{aligned} \quad (2.26)$$

Like in [Chamberland *et al.* 2022], in the limit of large $|\alpha|$, we can keep only the terms of order $|\alpha|^2$ and neglect all terms of order α or 1. This leading-order approximation is found to be valid for the estimation of phase-flip errors, which are the dominating errors for cat qubit Z gates. Furthermore, we move into

the rotating frame with the ideal Z gate Hamiltonian $\mathbf{H}_{Z,\text{ideal}} = 2 \text{Re}[\alpha] \varepsilon_Z \boldsymbol{\sigma}_z$. The master equation, thus describing the remaining error, is then given by:

$$\frac{d\rho}{dt} = iK_{\text{conf}}[\tilde{\mathbf{a}}^\dagger \tilde{\mathbf{a}}, \rho] + \kappa_{\text{conf}} \mathcal{D}[\tilde{\mathbf{a}}]\rho - i\varepsilon_Z(t) [\boldsymbol{\sigma}_z(\tilde{\mathbf{a}} + \tilde{\mathbf{a}}^\dagger), \rho] \quad (2.27)$$

where $K_{\text{conf}} = 4|\alpha|^2 K$ and $\kappa_{\text{conf}} = 4|\alpha|^2 \kappa_2$. It is now desired to adiabatically eliminate the gauge mode and keep only the effective dynamics on the qubit mode. To do so, we could use the generic and mathematically grounded approach of [Azouit *et al.* 2017]. Here, for simplicity, we write the Langevin equations in the Heisenberg picture of (2.27). For the logical qubit part, it reads $d\boldsymbol{\sigma}_z/dt = 0$, while for the gauge mode it reads

$$\frac{d\tilde{\mathbf{a}}}{dt} = -i\varepsilon_Z(t)\boldsymbol{\sigma}_z + iK_{\text{conf}}\tilde{\mathbf{a}} - \frac{1}{2}\kappa_{\text{conf}}\tilde{\mathbf{a}}. \quad (2.28)$$

The solution to this equation is given by

$$\tilde{\mathbf{a}}(t) = \tilde{\mathbf{a}}(0)e^{-(\frac{1}{2}\kappa_{\text{conf}} - iK_{\text{conf}})t} - i\boldsymbol{\sigma}_z \int_0^t e^{-(\frac{1}{2}\kappa_{\text{conf}} - iK_{\text{conf}})(t-t')} \varepsilon_Z(t') dt' \quad (2.29)$$

The first term vanishes at a rate κ_{conf} before the drive is applied. If $\varepsilon_Z(t)$ varies sufficiently slowly or jumps between constant values on which it stays sufficiently long compared to κ_{conf} , then the gauge mode follows the dynamics of the drive up to negligible transients. In this case, (2.29) simplifies to

$$\tilde{\mathbf{a}}(t) = \frac{i\varepsilon_Z(t)}{iK_{\text{conf}} - \kappa_{\text{conf}}/2} \boldsymbol{\sigma}_z \quad (2.30)$$

and the gauge mode dissipation of (2.27) results in equivalent $\boldsymbol{\sigma}_z$ dissipation, i.e.

$$\kappa_{\text{conf}} \mathcal{D}[\tilde{\mathbf{a}}(t)] = \kappa_{\text{conf}} \frac{\varepsilon_Z(t)^2}{K_{\text{conf}}^2 + \kappa_{\text{conf}}^2/4} \mathcal{D}[\boldsymbol{\sigma}_z] \quad (2.31)$$

From a more physical viewpoint, any displacement out of the code subspace will result in equivalent dephasing of the qubit due to two-photon dissipation bringing the state back in the cat qubit manifold with a phase error. This effect is mitigated with the additional Kerr Hamiltonian as it limits the displacement amplitude of the gauge mode.

We have thus derived an instantaneous dephasing term on the qubit mode. The total phase-flip error probability on the qubit after a $Z(\pi)$ gate is thus simply given by the integral of the dephasing rate over the gate time, and reads

$$p_Z^{\text{NA}} = \frac{1}{1 + 4K^2/\kappa_2^2} \frac{\overline{\varepsilon_Z^2} T}{|\alpha|^2 \kappa_2} \quad (2.32)$$

where $\overline{\varepsilon_Z^2} = 1/T \int_0^T \varepsilon_Z(t)^2 dt$ is average squared amplitude of the drive over gate time T . This is the same result as in [Chamberland *et al.* 2022], with the

additional prefactor involving K/κ_2 which can be used to reduce non-adiabatic phase errors thanks to the Kerr confinement.

Taking into account the single-photon loss of the cat qubit resonator with rate κ_1 , and assuming a constant drive $\varepsilon_Z(t) = \theta/4 \operatorname{Re}(\alpha)T$ during the gate, thus discarding the short transients in (2.30) when switching the gate Hamiltonian on and off, the total phase error for the Z gate with angle θ is given by

$$p_Z = \kappa_1 |\alpha|^2 T + \frac{1}{1 + 4K^2/\kappa_2^2} \frac{\theta^2}{16|\alpha|^4 T \kappa_2} \quad (2.33)$$

In the case of the combined TPE confinement, the phase-flip error probability is well fitted with the formula

$$p_Z = \kappa_1 |\alpha|^2 T + \frac{1}{1 + 4g_2^2/\kappa_2^2} \frac{\theta^2}{16|\alpha|^4 T \kappa_2} + \frac{\theta^2}{32|\alpha|^4 T^2 g_2^2}. \quad (2.34)$$

In this formula, the first contribution due to photon loss is the same gate-independent contribution as for all other confinement schemes. The second contribution is similar to the term found for the combined Kerr confinement. The third contribution is however new and specific to the TPE Hamiltonian, where the Z gate Hamiltonian induces hybridization of the codespace with a high-Q buffer state. We leave the derivation of this fitting formula for future work.

2.4.1.4 Numerical analysis

In Figure 2.9(a), these gate-induced phase errors are represented for three confinement schemes: a two-photon dissipation scheme (in green), a combined Kerr and two-photon dissipation scheme at the working point $K/\kappa_2 = 0.3$ (in dark blue) and a combined TPE and two-photon dissipation scheme at the working point $g_2/\kappa_2 = 10$ (in red). We remind that, as shown in Figure 2.8, these working points correspond to maximal Hamiltonian confinement strengths at which the bit-flip errors remain suppressed with a rate of order $\exp(-2|\alpha|^2)$ for $2 \leq |\alpha|^2 \leq 12$. The mean number of photons is fixed at $|\alpha|^2 = 8$ and the gate time varies along the horizontal axis. Markers show numerical data, while lines show analytical fits as given by Equations (2.23) and (2.24).

For the gate with a fully dissipative confinement, we find similar results as in the literature [Guillaud & Mirrahimi 2019, Chamberland *et al.* 2022, Xu *et al.* 2022a] with non-adiabatic errors suppressed linearly in the gate time. For the combined Kerr and dissipative confinement, almost no gain in performance is found compared to the dissipative one because of the low Kerr working point at $K/\kappa_2 = 0.3$. The model of (2.23) predicts a reduction by a

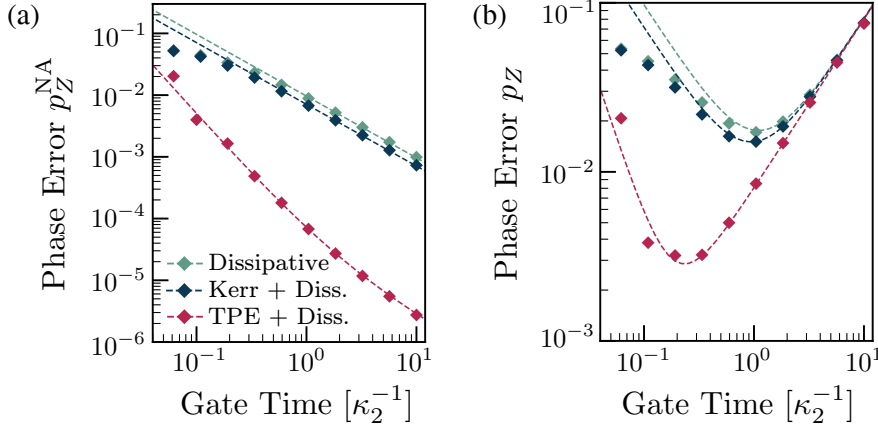


Figure 2.9: (a) Gate-induced phase errors for a Z gate at $|\alpha|^2 = 8$ with different confinement schemes: two-photon dissipation (green), combined Kerr and two-photon dissipation at $K/\kappa_2 = 0.3$ (dark blue), and combined TPE and two-photon dissipation at $g_2/\kappa_2 = 10$ (red). (b) Total phase errors for the same Z gates, assuming single-photon loss at rate $\kappa_1 = 10^{-3}\kappa_2$, thermal noise $n_{\text{th}} = 10^{-2}$ and pure dephasing $\kappa_\phi = 10^{-5}\kappa_2$. The combination of TPE and dissipative confinement allows us to win almost one order of magnitude on the phase error, and operates optimally almost one order of magnitude faster. For both plots, a constant Hamiltonian drive ε_Z is used and gate times are in units of $1/\kappa_2$. Markers indicate numerical data, dashed lines indicate analytical fits.

factor of about 1.36 on non-adiabatic errors. On the other hand, the combined TPE and dissipative confinement shows a significant gain in performance with a factor of about 400 compared to the fully dissipative confinement, in the limit of large gate times.

Figure 2.9(b) shows the complete phase-flip error probability in the presence of the other noise sources on the cat qubit resonator, and in particular with single-photon losses at rate $\kappa_1 = 10^{-3}\kappa_2$. In this case, the optimal phase-flip error probability for Z gate drops from $p_Z^* \approx 2\%$ in the fully dissipative scheme to $p_Z^* \approx 0.3\%$ in the combined TPE and dissipative scheme. The optimal gate time T^* that minimizes phase errors is also much smaller with the combined TPE scheme than with the fully dissipative one, hence leading to faster gate designs. For lower values of κ_1/κ_2 , the performance gain compared to fully dissipative gates is expected to increase, up to the factor of 400 characterizing the non-adiabaticity error.

In Figure 2.10, we further explore different working points for the combined TPE and dissipative confinement, at increasing g_2/κ_2 values. Plot (a) shows non-adiabatic phase flip errors while plot (b) shows the corresponding bit-flip errors, with same noise sources as in the Figure 2.9(b). As expected, non-

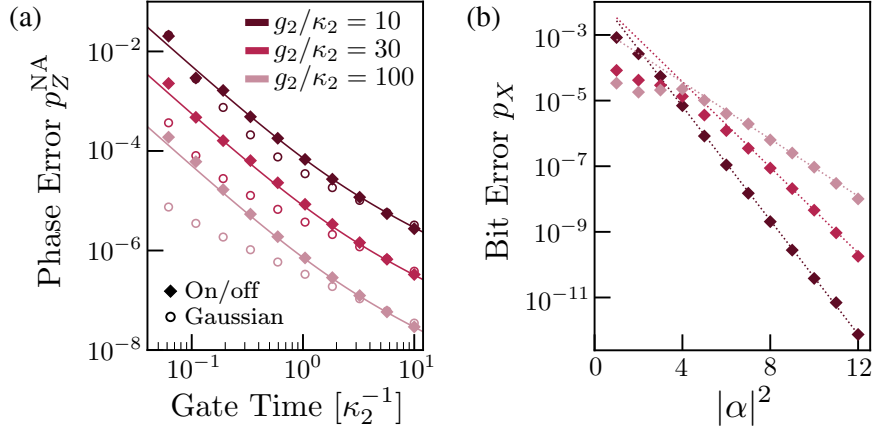


Figure 2.10: (a) Gate-induced phase errors for a Z gate at $|\alpha|^2 = 8$ for a combined TPE and two-photon dissipation confinement, at varying confinement ratios of g_2/κ_2 . Diamonds indicate a constant gate Hamiltonian drive, while circles indicate a gaussian-type gate Hamiltonian drive with first order superadiabatic correction, as in [Xu *et al.* 2022a]. Lines indicate the analytical fit (2.24) in the constant drive regime. Gate times are in units of $1/\kappa_2$. (b) Total bit-flip errors for a Z gate under the same confinement schemes, with single-photon loss at rate $\kappa_1 = 10^{-3}\kappa_2$, thermal noise $n_{\text{th}} = 10^{-2}$ and pure dephasing $\kappa_\phi = 10^{-5}\kappa_2$. Markers indicate numerical data, and dotted lines indicate exponential fits for $2 \leq |\alpha|^2 \leq 12$; the corresponding exponential suppression factors are the values γ represented on Figure 2.8.

adiabatic phase flip errors are suppressed quadratically with g_2/κ_2 , while the exponential suppression of bit-flip errors is continuously degraded from $\gamma \geq 2$ as g_2/κ_2 increases above the working point. The dotted lines on plot (b) follow the values of γ shown in Figure 2.8. To find the optimal error rate on the logical level, an optimization on the level of the quantum error correcting code should be performed.

Figure 2.10(a) also shows Z gate simulations with a superadiabatic pulse design (carved out circles). The applied drive amplitude is an adaptation of [Xu *et al.* 2022a] to the TPE Hamiltonian confinement and reads $\varepsilon_Z(t) \propto \Omega_G(t) + \ddot{\Omega}_G(t)/E_1^2$ where Ω_G is a second-order gaussian pulse and E_1 is the TPE Hamiltonian gap. The simulations show that such drive designs can further improve the performance of Z gates, but their benefits especially show up beyond the working point $g_2/\kappa_2 = 10$. In other words, the presence of the necessary dissipative confinement in order to counter bit flips induced by κ_1 , n_{th} and κ_ϕ processes, appears to limit the benefits of superadiabatic drives in this setting. Whether the compatibility of combined confinement with superadiabatic drives can be improved, remains a question for future research.

2.4.2 CNOT gate

The two-qubit CNOT gate is one of the most elementary entangling gates and holds a major role in the operation of a quantum computer. In error correcting codes, it is used to map the error syndromes on ancilla qubits that are subsequently measured. The success of an error correction process is therefore mainly limited by the infidelity of CNOT gates. Indeed, the effective logical error probability is more sensitive to physical CNOT fidelity than to any other component of the error correcting code, such as ancilla preparation or measurement [Fowler *et al.* 2012]. It is therefore crucial to engineer CNOT gates that can reach fidelities well above the error correction threshold. In this section, we first review the proposals on the design of CNOT gates for confined cat qubits. We then provide a new proposal that can achieve fast gate speeds, with drastically improved phase fidelity compared to dissipative schemes, and without compromising the bit-flip protection like fully Hamiltonian schemes.

2.4.2.1 Gate engineering through code deformation

For dissipative cat qubits, CNOT gates can be engineered through a slow variation of the dissipative confinement parameters, such that the target cat qubit manifold experiences a slow rotation in the harmonic oscillator phase space conditionally on the state of the control cat qubit [Guillaud & Mirrahimi 2019]. Once the conditional rotation reaches an angle π , the rotated and original cat qubit manifolds are again superimposed, but with a logical X-rotation between the states that they encode. More concretely, this process is effectively described by the time-dependent Lindblad operators

$$\mathbf{L}_2^{(C)} = \mathbf{a}_C^2 - \alpha^2, \quad (2.35a)$$

$$\mathbf{L}_2^{(T)}(t) = \mathbf{a}_T^2 - \alpha^2 + \frac{\alpha}{2}(e^{2i\varphi(t)} - 1)(\mathbf{a}_C - \alpha) \quad (2.35b)$$

where $\mathbf{a}_{C/T}$ are the annihilation operators of the control and target cat qubit resonators, $\varphi(t) = \pi t/T$ is the rotation angle of the target cat qubit codespace in the harmonic oscillator phase space, and T is the gate time. Indeed, at fixed t , the steady state of these dissipators is a four-dimensional subspace spanned by the coherent states $|\alpha\rangle_C \otimes |\pm\alpha\rangle_T$ and $|\alpha e^{i\varphi(t)}\rangle_C \otimes |\pm\alpha e^{i\varphi(t)}\rangle_T$. Assuming that the state of the system follows exactly the steady state subspace throughout the gate, a state $|\alpha\rangle_C \otimes |\pm\alpha\rangle_T \approx |0_L\rangle_C \otimes |0/1_L\rangle_T$ would thus not move, while a state $|\alpha e^{i\varphi(t)}\rangle_C \otimes |\pm\alpha\rangle_T \approx |1_L\rangle_C \otimes |0/1_L\rangle_T$ would have moved after time T to $|\alpha e^{i\varphi(t)}\rangle_C \otimes |\mp\alpha\rangle_T \approx |1_L\rangle_C \otimes |1/0_L\rangle_T$, effectively achieving a CNOT gate. However, such an assumption is only fulfilled in the limit of infinitely slow gates. In realistic conditions, the state of the harmonic oscillator lags behind the stabilized manifold, inducing errors.

To effectively reduce these errors, a so-called feedforward Hamiltonian can be added throughout the gate. This approach, which can be seen as a shortcut to adiabaticity for open quantum systems [Alipour *et al.* 2020], drives the oscillator state in such a way as to maintain it within the time-dependent dissipators steady state manifold, hence suppressing non-adiabatic errors. In its ideal form, a feedforward Hamiltonian for the cat qubit CNOT gate would read

$$\mathbf{H}_{\text{CX,ideal}} = \dot{\varphi} |1_L\rangle\langle 1_L|_C \otimes \mathbf{a}_T^\dagger \mathbf{a}_T + \mathbf{H}_s, \quad (2.36)$$

where the first term makes the target qubit rotate conditionally on the control being in the $|1_L\rangle_C$ logical state, and \mathbf{H}_s with $\mathbf{H}_s |1_L\rangle_C = \mathbf{H}_s |0_L\rangle_C = 0$ allows an arbitrary Hamiltonian when the control qubit is outside the code space [Guillaud & Mirrahimi 2019, Chamberland *et al.* 2022]. To approach such ideal feedforward Hamiltonian with realistic Hamiltonians, an approximate version is proposed [Guillaud & Mirrahimi 2019, Puri *et al.* 2020]

$$\mathbf{H}_{\text{CX}} = -\varepsilon_{\text{CX}}(\mathbf{a}_C + \mathbf{a}_C^\dagger - 2\text{Re}(\alpha))(\mathbf{a}_T^\dagger \mathbf{a}_T - |\alpha|^2). \quad (2.37)$$

where $\varepsilon_{\text{CX}} = \dot{\varphi}/4\text{Re}(\alpha)$. Indeed, at least for large α , the left-hand side bracket approximately evaluates to 0 if the control qubit is in the $|\alpha\rangle_C$ state, and to $-4\text{Re}(\alpha)$ if the control qubit is in the $|\alpha\rangle_C$ state (to see this, consider e.g. the mean value of the quadrature operator $\mathbf{a}_C + \mathbf{a}_C^\dagger$). This Hamiltonian can further be engineered through a four-wave mixing element coupled to cat qubit resonators as demonstrated in [Touzard *et al.* 2019]. The similarity of the Hamiltonian acting on the control qubit with the Z gate Hamiltonian is no coincidence. Indeed, in a dual viewpoint the CNOT gate corresponds to a Z rotation of the control qubit, conditioned on the logical phase of the target qubit. Since logical phase here corresponds to photon-number parity of the cat qubit, this dual viewpoint is clearly visible in \mathbf{H}_{CX} .

For Kerr cat qubits [Puri *et al.* 2020], the CNOT gate can be realized through the slowly varying Hamiltonian

$$\mathbf{H} = -K\mathbf{L}_2^{(\text{C})\dagger}\mathbf{L}_2^{(\text{C})} - K\mathbf{L}_2^{(\text{T})\dagger}(t)\mathbf{L}_2^{(\text{T})}(t) + \mathbf{H}_{\text{CX}}. \quad (2.38)$$

The first two terms feature a time-dependent four-dimensional subspace of ground states following the same evolution as for the dissipatively confined gate, and the third term is the feedforward Hamiltonian designed to drive the oscillator state along with this time-dependent steady state manifold. In this case, the adiabaticity condition ensuring accurate gate operation reads $\dot{\varphi} \ll 4|\alpha|^2 K$ where the right-hand side is the Kerr Hamiltonian gap [Puri *et al.* 2020]. For a fully Hamiltonian-based gate with sufficiently smooth variation of ε_{CX} , this should ensure exponential suppression of gate-induced errors as a function of the gate time, like for the Z gate. Its scaling can be further improved by using superadiabatic pulse designs as proposed in [Xu *et al.* 2022a], adding a feedforward Hamiltonian term

$H_{\text{CX}}^{(\text{sa})} \propto \dot{\varepsilon}_{\text{CX}} i(\mathbf{a}_C - \mathbf{a}_C^\dagger)(\mathbf{a}_T^\dagger \mathbf{a}_T - |\alpha|^2)$. The exponential suppression of non-adiabatic errors in the gate time is a drastic improvement from the linear scalings observed both for purely dissipative gates, and for Kerr-based CNOT gates with discontinuous on/off operation of ε_{CX} . Fully Hamiltonian gates are however limited by the same effects as discussed in previous sections, i.e. the important rate of bit-flip errors induced by thermal and dephasing noise.

2.4.2.2 Target stabilization-free gate protocol

Like for the Z gate, engineering CNOT gates with a combined Hamiltonian and two-photon dissipation confinement thus arises as a promising approach to combine the best of both worlds. However, at first glance, such a scheme appears experimentally very challenging due to the large number of Hamiltonians and dissipators that need to be engineered. For a combined Kerr and two-photon dissipation CNOT gate for instance, control and target qubit should be coupled through various mixing terms that engineer the target qubit Kerr confinement as a function of the control qubit, and also to a common buffer mode to engineer the target qubit two-photon dissipation term as depending on the control qubit. With the additional feedforward Hamiltonian, Kerr nonlinearities on each mode, and various single and two-photon drives that appear in Hamiltonian (2.38), this leads to a daunting experimental task.

To avoid the above complexities in engineering various Hamiltonians and dissipators, we consider a much simpler scheme where the confinement is turned off on the target cat qubit during the operation of the CNOT gate [Puri *et al.* 2019]. This is motivated by the fact that the feedforward Hamiltonian approximately preserves the coherent states composing the target cat qubit, by making them rotate in phase space by an angle that depends on the rather well-defined quadrature value of the control qubit, and it also rigorously preserves the phase of the target qubit (as it preserves photon number). During a fast gate, the main gate-induced imperfections thus take place on the control qubit, where the feedforward drive takes a similar form as for the Z gate. Leaving the target qubit unprotected only for the short gate time is therefore not too detrimental. Indeed, local leakage will be corrected as soon as the gate ends and confinement is turned back on, well before significant bit-flips can be induced. Regarding performance, it is not even clear how much could be gained by leaving on a coupled combined confinement scheme. In absence of an efficient design for TPE-type Hamiltonian confinement, with the typical parameters of the present work, gate performances are in fact improved when switching off κ_2 on the target qubit.

For a combined TPE and dissipative confinement scheme, the master equation describing the evolution of the system during the CNOT gate with our

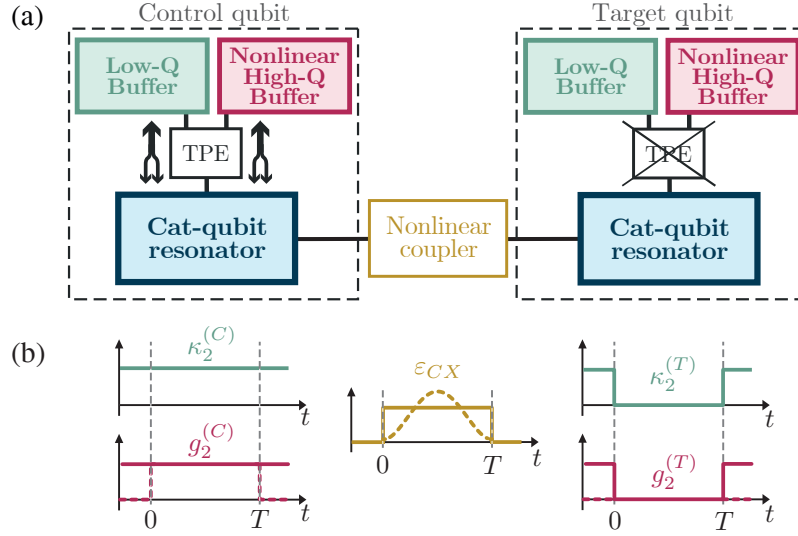


Figure 2.11: (a) Circuit schematic of a CNOT gate using a combined TPE Hamiltonian and two-photon dissipation confinement on both qubits. The control qubit resonator is connected to the target qubit resonator through a nonlinear coupler that engineers the feedforward Hamiltonian of (2.37). During the CNOT gate, both Hamiltonian and dissipative confinements are turned off on the target qubit by turning off the pump of the TPE element. (b) Amplitudes of the different dissipation and Hamiltonian terms during the CNOT gate. While the TPE confinement of the control qubit is essential during the gate, it can be turned off during the idling time. Keeping it on can however be useful for compensating spurious Hamiltonian perturbations. For the feedforward Hamiltonian, it is possible to benefit from superadiabatic pulse designs to further accelerate gates [Xu *et al.* 2022a].

proposal is thus given by

$$\frac{d\rho}{dt} = -i[\mathbf{H}_{CX}, \rho] - i[\mathbf{H}_{\text{TPE}}^{(C)}, \rho] + \kappa_2 \mathcal{D}[\mathbf{L}_2^{(C)}]\rho \quad (2.39)$$

where $\mathbf{H}_{\text{TPE}}^{(C)} = g_2(\mathbf{a}_C^2 - \alpha^2)\sigma_+ + \text{h.c.}$. A schematic of such a CNOT implementation is shown in Figure 2.11. This proposal should greatly simplify the experimental design of CNOT gates as the only coupling term that remains between control and target qubits is the feedforward Hamiltonian, experimentally demonstrated in [Touzard *et al.* 2019]. We next quantify the associated bit-flip and phase-flip errors.

2.4.2.3 Numerical analysis

In Figure 2.12, we plot the performance of CNOT gates with three different confinement schemes. The first one is the fully dissipative scheme with

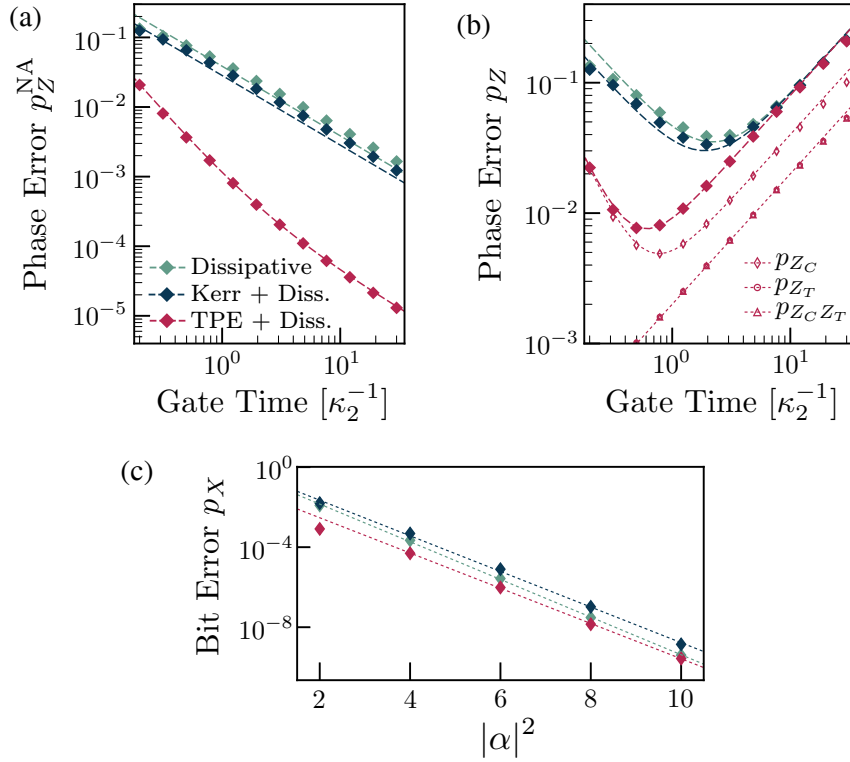


Figure 2.12: (a) Non-adiabatic phase errors of a CNOT gate at $|\alpha|^2 = 4$, with different confinement schemes: two-photon dissipation (green), combined Kerr and two-photon dissipation at $K/\kappa_2 = 0.3$ (dark blue), and combined TPE and two-photon dissipation at $g_2/\kappa_2 = 10$ (red). For this last scheme, the confinement is turned off on the target cat qubit during the gate. (b) Phase errors of a CNOT gate for the same confinement schemes, with single-photon loss at rate $\kappa_1 = 10^{-3}\kappa_2$, thermal noise $n_{\text{th}} = 10^{-2}$ and pure dephasing $\kappa_\phi = 10^{-5}\kappa_2$ on both control and target qubits. Full markers show total phase error, empty markers show the breakdown of phase errors for the combined TPE scheme. (c) Total bit-flip errors of a CNOT gate for the same confinement schemes, with noise rates as in plot (b), and at gate time $T = 1/\kappa_2$. For all plots, a constant gate Hamiltonian drive is used, and markers indicate numerical data. Dashed lines indicate analytical fits, and solid lines indicate numerical fits of exponential suppression.

both a control dissipator and a time-dependent target dissipator as proposed in Ref. [Guillaud & Mirrahimi 2019] (green), the second one is the combined Kerr Hamiltonian and two-photon dissipation acting on both control and target qubits as proposed in Ref. [Puri *et al.* 2020] but at the working point $K/\kappa_2 = 0.3$ (dark blue), and the third one is the combined TPE and two-photon dissipation scheme as described by Equation (2.39) at the working point $g_2/\kappa_2 = 10$ (red).

Similarly to the Z gate, Figure 2.12(a) first compares the ‘non-adiabatic’ phase errors, induced by the gate operation itself in absence of any perturbations on the two qubits. As explained above, the gate only affects the phase of the control qubit. We recall that the operating points have been selected in Section 2.3 to anticipate reasonable protection against typical error sources. Like for the Z gate, the combined TPE scheme clearly outperforms the other two, up to a factor about 100. For the three confinement schemes, non-adiabatic errors are fitted by the following formulas:

(Two-Photon Dissipation Confinement)

$$p_Z^{\text{NA}} = \frac{\pi^2}{64|\alpha|^2\kappa_2 T} \quad (2.40)$$

(Combined Kerr and Two-Photon Dissipation Confinement)

$$p_Z^{\text{NA}} = \frac{1}{1 + \frac{4K^2}{\kappa_2^2}} \frac{\pi^2}{64|\alpha|^2\kappa_2 T} \quad (2.41)$$

(Combined TPE and Two-Photon Dissipation Confinement)

$$p_Z^{\text{NA}} = \frac{1}{1 + \frac{4g_2^2}{\kappa_2^2}} \frac{\pi^2}{16|\alpha|^2\kappa_2 T} + \frac{\pi^2}{32|\alpha|^2 g_2^2 T^2}. \quad (2.42)$$

The first two formulas can be derived following the SFB approach introduced in [Chamberland *et al.* 2022], similarly to the case of Z gates (Section 2.4.1.3), while the third one is numerically fitted and valid as long as $g_2 \gtrsim \kappa_2$. Its analytical derivation is left for future work.

All three formulas scale as $|\alpha|^{-2}$ while their Z gate equivalents scale as $|\alpha|^{-4}$. This comes from the fact that the feedforward Hamiltonian ($\hat{a}_C + \hat{a}_C^\dagger$) acting on the control qubit like for the Z gate, is here multiplied by an amplitude that scales with the variance of the photon number in the target qubit state. It can be explicitly computed in the SFB as explained in [Chamberland *et al.* 2022]. Furthermore, the constant prefactor in the first term of Equation (2.42) involves $1/16$ instead of $1/64$ as in Equation (2.40) or Equation (2.41). This is an indirect effect of turning off the target two-photon dissipation, which lets \mathbf{H}_{CX} induce more entanglement of the control qubit phase information with target qubit leakage. At first glance, keeping two-photon dissipation on the target qubit may thus seem beneficial to gain this prefactor. However, keeping this dissipation would also introduce additional channels for control qubit phase dissipation, reducing the benefits of its Hamiltonian confinement towards performing fast gates. As such, compared to the purely dissipative scheme, the quadratic prefactor in g_2/κ_2 largely compensates the constant prefactor with the settings and working point of our proposal.

In Figure 2.12(b), the total phase-flip error probability of CNOT gates with the same three confinement schemes are shown in the presence of typical noise sources. Similarly to [Guillaud & Mirrahimi 2019], the latter induces phase-flip errors not only on the control qubit, but also on the target qubit, and correlated between control and target qubit. The figure hence shows the breakdown of those errors. Solely the control qubit undergoes non-adiabatic errors and the formulas are the same as in [Guillaud & Mirrahimi 2019] for purely dissipative confinement:

$$\begin{aligned} p_{Z_c} &= \kappa_1 |\alpha|^2 T + p_Z^{\text{NA}} \\ p_{Z_t} &= p_{Z_c Z_t} = \frac{1}{2} \kappa_1 |\alpha|^2 T, \end{aligned} \quad (2.43)$$

up to the expression of p_Z^{NA} . The optimal phase-flip error probability p_{Z_c} remains independent of α when using (2.42) instead of (2.40). For a realistic single-photon loss rate of $\kappa_1 = 10^{-3} \kappa_2$, this optimum is about $p_{Z_c}^* \approx 0.8\%$ for our scheme, to be compared with $p_{Z_c}^* \approx 5\%$ in the dissipative scheme. Like in the Z gate simulations, this optimal gate fidelity is furthermore obtained at a typical gate time about 5 times faster for the combined TPE scheme. The working point selection for the combined Kerr and dissipative confinement, appears to make it barely better than purely dissipative confinement in terms of reducing non-adiabatic and hence total phase flip error. Once again, for lower values of κ_1/κ_2 , further performance gains can be expected compared to fully dissipative gates, up to the factor of about 100 gained on the non-adiabaticity error at large gate times.

In Figure 2.12(c), total bit flip error probabilities for the same three confinement schemes are shown in the presence of the same noise sources as in plot (b). As expected from the working point selection, all three CNOT designs are bias-preserving with an exponential error suppression factor greater than 2. In particular, as the gate is fast enough, turning off the confinement of the target cat qubit during the operation does not have a significant impact on the bit-flip suppression. Dotted lines show exponential numerical fits of these bit-flip error probabilities. The exponential suppression factors are given by $\gamma \approx 2.17$ (fully dissipative scheme), $\gamma \approx 2.05$ (combined Kerr), and $\gamma \approx 2.03$ (combined TPE).

2.5 Towards experimental realization

2.5.1 Superconducting circuit implementation

The exponential suppression of bit-flips in dissipative cat qubits was experimentally demonstrated in [Lescanne *et al.* 2020b] using superconducting circuits as a physical platform. In this experiment the cat qubit resonator was




Figure 2.13: Superconducting circuit schematic of a combined TPE and two-photon dissipation confinement. The cat qubit resonator (blue) is capacitively coupled to an ATS (black) that mediates two-photon exchanges both with its self high-Q mode (red) and with a low-Q buffer mode (green). The ATS is DC-biased at the normalized flux $(0, \pi)$ working point and pumped at frequencies $2\omega_a - \omega_h$ and $2\omega_a - \omega_l$, where ω_a , ω_h and ω_l are the resonant frequencies of the cat qubit resonator, high-Q buffer and low-Q buffer respectively. To ensure the anharmonicity of the ATS mode (red), the ATS junctions are designed to admit slightly different Josephson energies. In addition, both buffers are resonantly driven to control the cat qubit mean number of photons.

coupled to a low-Q buffer mode with a two-photon exchange Hamiltonian, resulting in an effective two-photon dissipation on the cat qubit resonator. The two-photon exchange Hamiltonian was engineered with an ATS, a non-linear coupling element made up of two (ideally) identical junctions shunted by a large inductance. Threading both loops of the ATS at normalized DC flux biases of 0 and π magnetic flux quanta respectively allows to keep only the odd parity mixing terms of the participating modes. Then flux pumping the ATS at frequency $2\omega_a - \omega_b$, where $\omega_{a/b}$ are respectively the cat qubit and buffer resonator frequencies, creates the required two-photon exchange coupling. Finally, in [Lescanne *et al.* 2020b], not only did the ATS implement the two-photon exchange coupling but it also hosted the low-Q buffer mode.

The combined TPE and dissipative confinement scheme proposed in this thesis can be implemented with minor modifications of this experimental setup. A superconducting circuit schematic of a setup implementing the confinement is represented in Figure 2.13. Similarly to [Lescanne *et al.* 2020b], the cat qubit resonator (in blue) is capacitively coupled to the ATS (in black and red) which is DC-biased at the same working point of 0 and π . The ATS

is further coupled to a linear resonator (in green) which will serve as the low-Q buffer for two-photon dissipation engineering. This leaves the self-mode of the ATS open for other purposes and we use it as the high-Q buffer mode of the device. Thus, instead of having a qubit as the high-Q buffer, this implementation proposes to use a sufficiently anharmonic oscillator mode. In simulations shown later in this Section, we observe that performance degradation remains moderate even with a fully harmonic mode as high-Q buffer, but this option would require further analysis if it were to be considered. Note also that other circuit QED implementations of the TPE Hamiltonian at $\alpha = 0$ were proposed in [Felicetti *et al.* 2018a, Felicetti *et al.* 2018b].

To make the high-Q buffer mode anharmonic in this implementation scheme, it is sufficient that the ATS junction energies are designed to be not identical. Note that even in [Lescanne *et al.* 2020b] these junctions were not perfectly identical and the amount of asymmetry in that setup might be sufficient to ensure the buffer mode to be anharmonic enough for the design of the present thesis. The circuit Hamiltonian engineered with such a setup reads

$$\mathbf{H} = \omega_a \mathbf{a}^\dagger \mathbf{a} + \omega_h \mathbf{b}_h^\dagger \mathbf{b}_h + \omega_l \mathbf{b}_l^\dagger \mathbf{b}_l - 2E_J [\varepsilon(t) \sin(\varphi) + \eta \widetilde{\cos}(\varphi)] \quad (2.44)$$

where \mathbf{a} , \mathbf{b}_h and \mathbf{b}_l are the hybridized modes corresponding to the cat qubit resonator, high-Q and low-Q buffers respectively with ω_a , ω_h and ω_l their corresponding resonance frequencies, and where $\widetilde{\cos}(\varphi) = \cos(\varphi) + \varphi^2/2$. The total phase across the ATS dipole element reads $\varphi = \varphi_a(\mathbf{a} + \mathbf{a}^\dagger) + \varphi_h(\mathbf{b}_h + \mathbf{b}_h^\dagger) + \varphi_l(\mathbf{b}_l + \mathbf{b}_l^\dagger)$. The energy E_J is the average of the two junction energies, $E_J = (E_{J,1} + E_{J,2})/2$, and η is the asymmetry of the junctions, $\eta = (E_{J,1} - E_{J,2})/(E_{J,1} + E_{J,2})$. Finally, $\varepsilon(t)$ is a Radiofrequency (RF) flux pumped at two frequencies $2\omega_a - \omega_h$ and $2\omega_a - \omega_l$ with respective amplitudes ε_h and ε_l .

In Section 2.5.3, we show that this Hamiltonian, after an RWA and adding resonant drives on both buffer modes, reduces to

$$\begin{aligned} \mathbf{H} = & g_{2,l}(\mathbf{a}^2 - \alpha^2)\mathbf{b}_l^\dagger + g_{2,h}(\mathbf{a}^2 - \alpha^2)\mathbf{b}_h^\dagger + \text{h.c.} \\ & - \chi_{hh}\mathbf{b}_h^{\dagger 2}\mathbf{b}_h^2 - \chi_{ll}\mathbf{b}_l^{\dagger 2}\mathbf{b}_l^2 - \chi_{aa}\hat{a}^{\dagger 2}\mathbf{a}^2 \\ & - \chi_{ah}\mathbf{a}^\dagger\mathbf{a}\mathbf{b}_h^\dagger\mathbf{b}_h - \chi_{al}\mathbf{a}^\dagger\mathbf{a}\mathbf{b}_l^\dagger\mathbf{b}_l - \chi_{lh}\mathbf{b}_l^\dagger\mathbf{b}_l\mathbf{b}_h^\dagger\mathbf{b}_h \end{aligned} \quad (2.45)$$

where TPE amplitudes are given by $g_{2,x} = E_J\varphi_a^2\varphi_x\varepsilon_x/2$, and where self Kerr and cross Kerr terms are given by $\chi_{xx} = \eta E_J\varphi_x^4/2$ and $\chi_{xy} = \eta E_J\varphi_x^2\varphi_y^2$ with $x, y = a, h$ or l . The parameters can be chosen in such a way that χ_{hh} strongly dominates all the other coupling terms and that $|g_{2,l}|$ and $|g_{2,h}|$ strongly dominate all the remaining terms. In this manner, the dominant Hamiltonian to be considered is

$$\mathbf{H}_0 = g_{2,l}(\mathbf{a}^2 - \alpha^2)\mathbf{b}_l^\dagger + g_{2,h}(\mathbf{a}^2 - \alpha^2)\mathbf{b}_h^\dagger + \text{h.c.} - \chi_{hh}\mathbf{b}_h^{\dagger 2}\mathbf{b}_h^2 \quad (2.46)$$

and all other self-Kerr and cross-Kerr terms can be considered as perturbations to this dynamics. All these perturbation terms are parity preserving but they could in principle induce bit-flip errors on the cat qubit. However, thanks to the combined Hamiltonian and dissipative confinement, it is sufficient to have the amplitude of these terms small compared to the confinement rate to exponentially suppress their effects. This suppression is facilitated for all the coupling terms, because both buffer modes would nominally be in their ground state for an idling qubit with no further noise.

The fast decay of the low-Q buffer mode at a rate $\kappa_{b,l}$, together with the first term of Equation (2.46) yields effectively the two-photon dissipation $\kappa_2 \mathcal{D}[\mathbf{a}^2 - \alpha^2]$ on the cat qubit resonator. The effective rate κ_2 is given by $4g_{2,l}^2/\kappa_{b,l}$ in the limit of $g_{2,l}/\kappa_{b,l} \ll 1$. The second term is the TPE Hamiltonian term. Note that the ratio of TPE to two-photon dissipation confinement rates can actively be tuned through the RF flux amplitudes ε_h and ε_l . Finally, the third term of Equation (2.46) is the Kerr nonlinearity on the buffer mode that is essential to make it anharmonic. This Kerr nonlinearity should be large compared to the TPE Hamiltonian gap, to ensure that terms associated to higher excitations of the high-Q buffer would rotate at a much faster rate than the TPE Hamiltonian contribution, and thus be rejected through RWA.

In summary, the typical orders of magnitude necessary for the design are given by

$$\chi_{ah}, \chi_{aa}, \chi_{al} \ll g_{2,h}, g_{2,l} \ll \chi_{hh} \quad (2.47)$$

which, by noting that $\varphi_a, \varphi_l < \varphi_h$, implies

$$\eta\varphi_h^2\varphi_a^2 \ll \varepsilon_h\varphi_h\varphi_a^2, \varepsilon_l\varphi_l\varphi_a^2 \ll \eta\varphi_h^4. \quad (2.48)$$

By design one can choose φ_h sufficiently large with respect to φ_l and φ_a , and next tune the junction asymmetry η to fulfill the above requirements. An example of a set of experimental parameters that can fulfill these requirements is given by $\varphi_h = 0.2$, $\varphi_a = 0.04$, $\varphi_l = 0.08$, $\eta = 1\%$, $\varepsilon_h = 0.02$, $\varepsilon_l = 0.05$. With these parameters, the condition of Equation (2.48) reads $6.4 \cdot 10^{-7} \ll 64 \cdot 10^{-7} \ll 160 \cdot 10^{-7}$. Further assuming an average Josephson energy of $E_J/2\pi = 80$ GHz would yield two-photon exchange rates at $g_{2,l}/2\pi = g_{2,h}/2\pi = 250$ kHz.

2.5.2 Spurious terms

In this section, we attempt to numerically characterize the impact of spurious Hamiltonian and Lindblad terms related to the high-Q buffer in the TPE confinement. In particular, in Figure 2.14, we investigate a Z gate where we estimate the probability of phase-flip errors at a fixed mean number of photons $|\alpha|^2 = 8$ (top) and bit-flip errors at a fixed gate time $T = 1/\kappa_2$ (bottom).

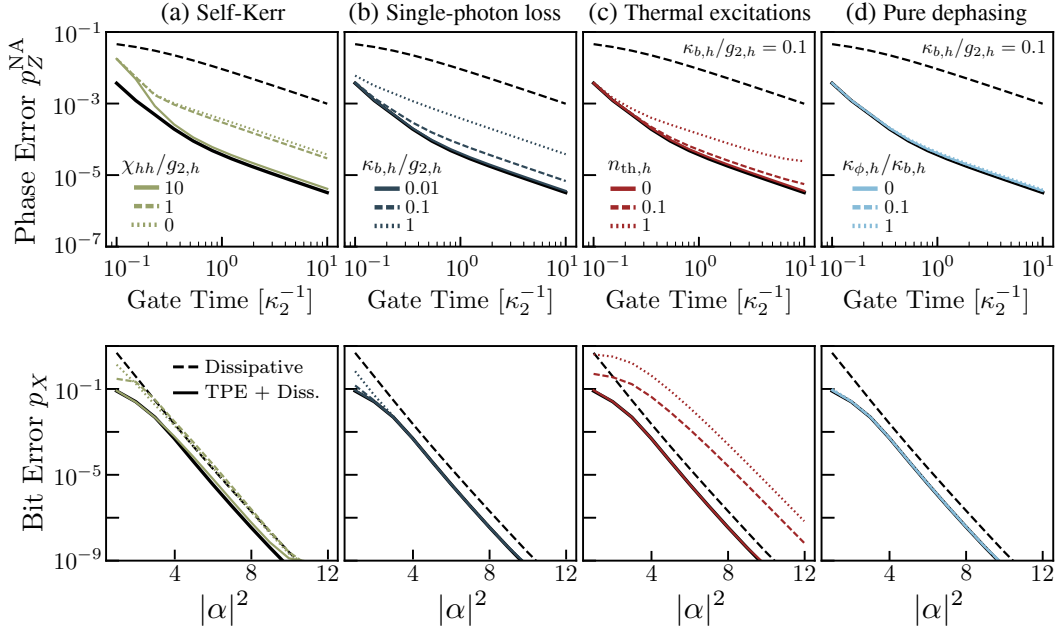


Figure 2.14: Characterization of the spurious contributions induced by the high-Q buffer mode for a Z gate. Non-adiabatic phase-flip error probability for $|\alpha|^2 = 8$ (top) and bit-flip error probability at $T = 1/\kappa_2$ (bottom). In all plots, reference Z gate simulations for two confinement schemes are shown: the fully dissipative confinement (dashed black lines) and the combined TPE and dissipative confinement at the working point $g_{2,h}/\kappa_2 = 10$. Other simulations are the same as for the combined confinement reference, with additional terms as typically resulting from an experimental realization. (a) The buffer mode is modeled as an anharmonic oscillator, with Hamiltonian $\omega_{b,h}\mathbf{b}_h^\dagger\mathbf{b}_h - \chi_{hh}\mathbf{b}_h^{\dagger 2}\mathbf{b}_h^2$. (b) Assuming a perfect two-level buffer, relaxation is considered as a Lindblad term $\kappa_{b,h}\mathcal{D}[\sigma_-]$. (c) Thermal excitation is added to the situation of plot (b), thus relaxation is now modeled by the Lindbladian $\kappa_{b,h}(1 + n_{th,h})\mathcal{D}[\sigma_-] + \kappa_{b,h}n_{th,h}\mathcal{D}[\sigma_+]$, for a fixed relaxation rate $\kappa_{b,h}/g_{2,h} = 10^{-2}$. (d) Pure dephasing relaxation is added to the situation of plot (b), modeled by a Lindbladian $\kappa_{b,h}\mathcal{D}[\sigma_-] + \kappa_{\phi,h}\mathcal{D}[\sigma_z]$, with a fixed single-photon relaxation rate $\kappa_{b,h}/g_{2,h} = 10^{-2}$.

The first effect investigated in Figure 2.14(a) is the impact of having an oscillator mode of finite anharmonicity as a buffer, instead of a strict two-level system. Considering a Kerr-type non-linearity for this buffer mode, thus with Hamiltonian $-\chi_{hh}\mathbf{b}_h^{\dagger 2}\mathbf{b}_h^2$, we vary the Kerr strength χ_{hh} to make the mode more or less anharmonic. This Kerr strength should be large compared to all other frequencies in the system, in order to prevent transitions to higher excited states of the buffer mode. This is demonstrated in panel (a) since only large enough Kerr strengths compared to g_2 yield phase errors as good as the

reference simulation in solid black lines (where the buffer mode is assumed to be a qubit). Note however that reasonable phase-flip performances are obtained even with perfectly harmonic buffer modes and that the bit-flip error probability is unaffected, thanks to two-photon dissipation. The option of a harmonic buffer mode would require more careful analysis though, since the corresponding confinement Hamiltonian appears to have a high-dimensional 0-energy eigenspace.

The second effect, investigated in Figure 2.14(b), is high-Q buffer relaxation i.e. finite lifetime of the qubit. This is taken into account in simulations with an additional Lindblad term $\kappa_{b,h}\mathcal{D}[\sigma_-]$. For large single-photon loss rates $\kappa_{b,h} \gg g_{2,h}$, the high-Q mode in fact becomes low-Q and induces a new two-photon dissipation channel just like \mathbf{b}_l , leading to an additional two photon dissipation rate $\tilde{\kappa}_2 = 4g_{2,h}^2/\kappa_{b,h}$. For $\kappa_{b,h} \simeq g_{2,h}$, we can expect the buffer to combine high-Q and low-Q effects. This only shows up as an increase of phase-flip error probabilities during the gate, yet still much better than the performance reached for the fully dissipative gate. In the associated figure, we find that $\kappa_{b,h}/g_{2,h} \leq 0.1$ is enough to obtain performances similar to the reference one.

The third effect investigated, in Figure 2.14(c) is relaxation of the buffer mode in a non-zero temperature environment, corresponding to a Lindbladian $\kappa_{b,h}(1 + n_{\text{th},h})\mathcal{D}[\sigma_-] + \kappa_{b,h}n_{\text{th},h}\mathcal{D}[\sigma_+]$. Thermal excitation induces jumps from $|g\rangle$ to $|e\rangle$ which translate to the cat qubit resonator as excitation of the form $(\mathbf{a}^2 - \alpha^2)^\dagger$. Such excitation has a similar effect as direct thermal excitation of the cat qubit mode already considered in the main text. As far as $g_{2,h}$ is fixed below the working point shown in Figure 2.8, the induced leakage is compensated by the two-photon dissipation. Therefore the exponential bit-flip suppression is maintained even though the error probability is increased by a constant factor. This is shown in the bottom plot of Figure 2.14(c). There only remains an indirect effect on phase-flips due to larger leakage during the gate operation.

Finally, the fourth effect, investigated in Figure 2.14(d), is pure dephasing of the buffer mode, corresponding to a Lindbladian $\kappa_{b,h}\mathcal{D}[\sigma_-] + \kappa_{\phi,h}\mathcal{D}[\sigma_z]$. For the values investigated in this figure, i.e. up to $\kappa_{\phi,h}/g_{2,h} = 0.01$, pure dephasing on the buffer mode has no impact on either phase-flip or bit-flip errors during Z gates.

2.5.3 Circuit Hamiltonian derivation

In this section, we derive the Hamiltonian for the superconducting circuit implementation shown in (2.45). Its equivalent circuit is shown in Figure 2.15 with an ATS acting as a high-Q buffer, capacitively coupled to the cat qubit




Figure 2.15: Equivalent circuit diagram of the experimental proposal of Figure 2.13. The cat qubit LC resonator (blue) is capacitively coupled to the ATS circuit element (red), itself capacitively coupled to the low-Q buffer LC resonator (green). The high-Q buffer is the self-mode of the ATS, made of two Josephson junctions (nonlinear inductance represented by a cross & associated capacitor) and an inductance in parallel, and threaded with two flux biases $\varphi_{\text{ext},1}$ and $\varphi_{\text{ext},2}$. Not shown here: the low-Q buffer resonator is capacitively coupled to a dissipative bath.

resonator and to the buffer resonator. The full Hamiltonian of this circuit reads

$$\mathbf{H} = \omega_{a,0} \mathbf{a}^\dagger \mathbf{a} + \omega_{h,0} \mathbf{b}_h^\dagger \mathbf{b}_h + \omega_{l,0} \mathbf{b}_l^\dagger \mathbf{b}_l - E_{J,1} \cos(\varphi + \varphi_{\text{ext},1}) - E_{J,2} \cos(\varphi - \varphi_{\text{ext},2}) \quad (2.49)$$

where \mathbf{a} , \mathbf{b}_h and \mathbf{b}_l are slightly hybridized modes corresponding to the cat qubit, high-Q buffer and low-Q buffer modes respectively, with corresponding mode frequencies $\omega_{a,0}$, $\omega_{h,0}$ and $\omega_{l,0}$. $E_{J,1}$ and $E_{J,2}$ denote the junction energies of the ATS, and $\varphi_{\text{ext},1}$ and $\varphi_{\text{ext},2}$ are the tunable flux biases in each ATS loop as shown in Figure 2.15. The total phase φ across the ATS inductance is given by $\varphi = \varphi_a(\mathbf{a}^\dagger + \mathbf{a}) + \varphi_h(\mathbf{b}_h^\dagger + \mathbf{b}_h) + \varphi_l(\mathbf{b}_l^\dagger + \mathbf{b}_l)$ where φ_a , φ_h and φ_l are the participation ratios of the three modes, deduced from Kirchhoff's laws and the modes hybridization. The corresponding Hamiltonian can be rewritten as

$$\begin{aligned} \mathbf{H} = & \omega_{a,0} \mathbf{a}^\dagger \mathbf{a} + \omega_{h,0} \mathbf{b}_h^\dagger \mathbf{b}_h + \omega_{l,0} \mathbf{b}_l^\dagger \mathbf{b}_l \\ & - 2E_J [\cos(\varphi_\Sigma) \cos(\varphi + \varphi_\Delta) + \eta \sin(\varphi_\Sigma) \sin(\varphi + \varphi_\Delta)] \end{aligned} \quad (2.50)$$

where $\varphi_\Sigma = (\varphi_{\text{ext},1} + \varphi_{\text{ext},2})/2$ and $\varphi_\Delta = (\varphi_{\text{ext},1} - \varphi_{\text{ext},2})/2$ are the average and difference of the two flux biases, $E_J = (E_{J,1} + E_{J,2})/2$ is the average junction energy, and $\eta = (E_{J,1} - E_{J,2})/(E_{J,1} + E_{J,2})$ is the junction asymmetry. Our working point considers the ATS to be DC biased at the flux bias point defined by

$$\varphi_\Sigma = \frac{\pi}{2} + \varepsilon(t), \quad \varphi_\Delta = \frac{\pi}{2} \quad (2.51)$$

with two additional RF fluxes applied on φ_Σ such that $\varepsilon(t) = \varepsilon_1 \cos(\omega_{p,1}t) + \varepsilon_2 \cos(\omega_{p,2}t)$. This bias point is chosen such that only odd powers of φ are turned on for $\eta = 0$. The Hamiltonian now reads

$$\mathbf{H} = \omega_{a,0} \mathbf{a}^\dagger \mathbf{a} + \omega_{h,0} \mathbf{b}_h^\dagger \mathbf{b}_h + \omega_{l,0} \mathbf{b}_l^\dagger \mathbf{b}_l - 2E_J [\varepsilon(t) \sin(\varphi) + \eta \cos(\varphi)] \quad (2.52)$$

where we have assumed $|\varepsilon(t)| \ll 1$. By expanding the cosine and sine terms up to fourth order and absorbing the quadratic terms of the cosine in the mode frequencies, the Hamiltonian becomes

$$\mathbf{H} = \omega_a \mathbf{a}^\dagger \mathbf{a} + \omega_h \mathbf{b}_h^\dagger \mathbf{b}_h + \omega_l \mathbf{b}_l^\dagger \mathbf{b}_l - 2E_J [\varepsilon(t) \varphi - \varepsilon(t) \varphi^3/6 + \eta \varphi^4/24] \quad (2.53)$$

We can move to a displaced frame for the three modes to absorb the linear term $-2E_J \varepsilon(t) \varphi$ in the above Hamiltonian. Assuming single-photon decay rates κ_a , $\kappa_{b,h}$ and $\kappa_{b,l}$, the displacements are given by

$$\xi_x = \sum_{k=1,2} \frac{-iE_J \varphi_x \varepsilon_k}{i(\omega_x - \omega_{p,k}) + \kappa_x/2} e^{-i\omega_{p,k}t}, \quad x = a, h, l. \quad (2.54)$$

The Hamiltonian in the displaced frame is given by

$$\tilde{\mathbf{H}} = \omega_a \mathbf{a}^\dagger \mathbf{a} + \omega_h \mathbf{b}_h^\dagger \mathbf{b}_h + \omega_l \mathbf{b}_l^\dagger \mathbf{b}_l + E_J [\varepsilon(t) \tilde{\varphi}^3/3 - \eta \tilde{\varphi}^4/12]. \quad (2.55)$$

Here,

$$\begin{aligned} \tilde{\varphi} &= \varphi_a (\mathbf{a}^\dagger + \mathbf{a}) + \varphi_h (\mathbf{b}_h^\dagger + \mathbf{b}_h) + \varphi_l (\mathbf{b}_l^\dagger + \mathbf{b}_l) \\ &+ s_1 e^{-i\omega_{p,1}t} + s_1^* e^{i\omega_{p,1}t} + s_2 e^{-i\omega_{p,2}t} + s_2^* e^{i\omega_{p,2}t} \end{aligned} \quad (2.56)$$

with

$$s_k = \sum_{x=a,h,l} \frac{iE_J \varepsilon_k \varphi_x^2}{i(\omega_x - \omega_{p,k}) + \kappa_x/2}, \quad k = 1, 2. \quad (2.57)$$

Finally, we consider the addition of resonant microwave drives on the buffer modes with amplitudes $\zeta_h(t) = \zeta_h e^{-i\omega_h t}$ and $\zeta_l(t) = \zeta_l e^{-i\omega_l t}$. The Hamiltonian becomes

$$\begin{aligned} \tilde{\mathbf{H}} &= \omega_a \mathbf{a}^\dagger \mathbf{a} + \omega_h \mathbf{b}_h^\dagger \mathbf{b}_h + \omega_l \mathbf{b}_l^\dagger \mathbf{b}_l + [\zeta_h(t) \mathbf{b}_h^\dagger + \text{h.c.}] + [\zeta_l(t) \mathbf{b}_l^\dagger + \text{h.c.}] \\ &+ E_J [\varepsilon(t) \tilde{\varphi}^3/3 - \eta \tilde{\varphi}^4/12] \end{aligned} \quad (2.58)$$

We set the pumping frequencies at $\omega_{p,1} = 2\tilde{\omega}_a - \tilde{\omega}_h$ and $\omega_{p,2} = 2\tilde{\omega}_a - \tilde{\omega}_l$ where $\tilde{\omega}_x$ are AC stark shifted frequencies, i.e. including the effects of powers of $\tilde{\varphi}$ on the actual mode frequencies. By going to the rotating frame of each mode, and performing a RWA, we obtain the effective Hamiltonian

$$\begin{aligned} \tilde{\mathbf{H}} &= g_{2,h} (\mathbf{a}^2 - \alpha^2) \mathbf{b}_h^\dagger + g_{2,l} (\mathbf{a}^2 - \alpha^2) \mathbf{b}_l^\dagger + \text{h.c.} \\ &- \chi_{hh} (\mathbf{b}_h^\dagger)^2 \mathbf{b}_h^2 - \chi_{ll} (\mathbf{b}_l^\dagger)^2 \mathbf{b}_l^2 - \chi_{aa} (\mathbf{a}^\dagger)^2 \mathbf{a}^2 \\ &- \chi_{ah} \mathbf{a}^\dagger \mathbf{a} \mathbf{b}_h^\dagger \mathbf{b}_h - \chi_{al} \mathbf{a}^\dagger \mathbf{a} \mathbf{b}_l^\dagger \mathbf{b}_l - \chi_{lh} \mathbf{b}_l^\dagger \mathbf{b}_l \mathbf{b}_h^\dagger \mathbf{b}_h \end{aligned} \quad (2.59)$$

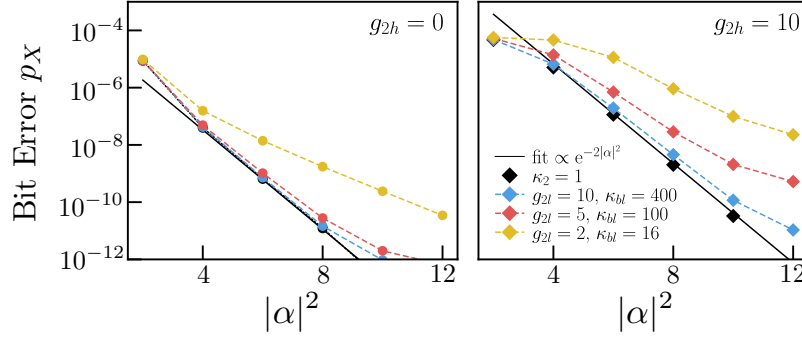


Figure 2.16: Bit-flip error probability of an idling qubit for a combined TPE and dissipative confinement with single-photon loss $\kappa_1 = 10^{-3}$ and thermal noise $n_{\text{th}} = 10^{-2}$. The left panel shows numerical simulations without the TPE Hamiltonian and with a single low-Q buffer mode, while the right panel shows numerical simulations with the TPE Hamiltonian and two buffer modes. In both panels, colored lines show different values of $g_{2,l}$ and $\kappa_{b,l}$ but at fixed effective two-photon dissipation, $\kappa_{2,\text{eff}} = 4g_{2,l}^2/\kappa_{b,l}$, computed using the adiabatic elimination formula valid at $g_{2,l} \ll \kappa_{b,l}$. Black diamonds show reference simulations with the low-Q buffer mode completely eliminated according to this adiabatic elimination formula, thus actually corresponding to two-photon dissipation $\kappa_2 \mathcal{D}[a^2 - \alpha^2]$. The black line is a numerical fit of the black diamonds with exponential ratio $\exp(-2|\alpha|^2)$.

where the resonant drive amplitudes were set as $\zeta_h = -\alpha^2 g_{2,h}$ and $\zeta_l = -\alpha^2 g_{2,l}$. The various coupling strengths in this Hamiltonian are given by

$$\begin{aligned}
 g_{2,h} &= E_J \varphi_a^2 \varphi_h (\varepsilon_1/2 - \eta s_1) \\
 g_{2,l} &= E_J \varphi_a^2 \varphi_l (\varepsilon_2/2 - \eta s_2) \\
 \chi_{xx} &= \eta E_J \varphi_x^4 / 2 \\
 \chi_{xy} &= \eta E_J \varphi_x^2 \varphi_y^2
 \end{aligned} \tag{2.60}$$

with $x, y = a, h, l$ and $x \neq y$.

Assuming the Hamiltonian of Equation (2.59) with a strong single-photon dissipation on the low-Q buffer mode, $\kappa_{b,l} \mathcal{D}[\mathbf{b}_l]$, it is possible to adiabatically eliminate the buffer mode to obtain an effective two-photon dissipation of amplitude $\kappa_{2,\text{eff}} = 4g_{2,l}^2/\kappa_{b,l}$ in the limit of $\kappa_{b,l} \gg g_{2,l}$. Figure 2.16 showcases a numerical simulation to validate this adiabatic elimination. On the left panel, the usual dissipative cat qubit situation [Lescanne *et al.* 2020b, Touzard *et al.* 2018] with a single buffer mode is shown. On the right panel, a dissipative cat qubit with the additional TPE confinement at $g_{2,h}/\kappa_2 = 10$, and thus with two buffer modes, is shown. In both panels, a reference simulation is shown in black where the Hamiltonian coupling to the low-Q buffer mode is replaced by the dissipation with $\kappa_{2,\text{eff}}$ resulting from the adiabatic

elimination formula. Colored lines then show the results of the full system, including the low-Q buffer mode, for increasing values of $\kappa_{b,l}/g_{2,l}$ at a fixed effective two-photon dissipation rate $\kappa_{2,\text{eff}}$. As can be seen, in both cases the bit-flip error is efficiently suppressed, with a rate tending towards the adiabatic elimination formula once $\kappa_{b,l}/g_{2,l}$ reaches values about 20 to 40.

2.6 Conclusion

Bosonic encoding and biased-noise qubits have been gaining a lot of attention for their promise of hardware-efficient fault-tolerance. Along these lines, cat qubits have been investigated with a particular interest and following two confinement approaches: first, a two-photon driven dissipative stabilization scheme and second a two-photon driven Kerr Hamiltonian confinement. In this chapter, we have proposed a new confinement design that benefits from the advantages of both dissipative and Hamiltonian approaches. This design, combining the previously known driven dissipative scheme with a Hamiltonian confinement using a TPE Hamiltonian, provides a robust exponential suppression of bit-flip errors with the cat size together with enhanced gate performances, both in speed and fidelity.

Furthermore, this proposal is implementable with only minor modifications of known experimental designs of dissipative cat qubits. We expect that this approach can further reduce the hardware overhead for fault-tolerance with cat qubits. Additional mathematical analysis of these combined confinement schemes and numerical studies of three-qubit Toffoli gates, which are essential for universal quantum computation with repetition cat qubits, will be subject of forthcoming research work.

Designing high-fidelity Zeno gates

This chapter covers the work that was published in [Gautier *et al.* 2023b], proposing shortcuts for Zeno gate designs of dissipative cat qubits.

Contents

3.1	Introduction	84
3.2	Zeno gate errors	85
3.3	Summary of gate designs	88
3.3.1	Buffer photodetection with classical feedback	89
3.3.2	Cat-buffer autonomous feedback	89
3.3.3	Locally flat Hamiltonian	90
3.3.4	Discrete jump	90
3.3.5	Combined dissipation and two-photon exchange Hamiltonian	91
3.3.6	Combining designs	91
3.4	Gate designs	91
3.4.1	Buffer photodetection with classical feedback	91
3.4.2	Cat-buffer autonomous feedback	100
3.4.3	Locally flat Hamiltonian	105
3.4.4	Discrete jump	112
3.5	Robustness to noise	119
3.5.1	On the impact of noise and spurious Hamiltonians	119
3.5.2	Thermal noise in the buffer mode	121
3.6	Conclusion	123

3.1 Introduction

The promise of quantum computing relies on the unpleasant predicament that a quantum system should be freely controllable but also very long-lived, two often conflicting requirements. Dissipative cat qubits are no exception to this rule. Thanks to their autonomous bit-flip protection and dynamical phase-flip protection once concatenated with a repetition code [Guillaud & Mirrahimi 2019], they are long-lived qubits. However, their controllability will likely be one of the main limitations towards achieving error correction below threshold [Shor 1996]. Indeed, current gate proposals with cat qubits feature a poor scaling with the relevant experimental parameters that may limit their practical use, and the few experimental implementations of such gates have yet to demonstrate high-fidelity gates [Touzard *et al.* 2018, Réglade *et al.* 2023]. In this aim, several proposals have been put forward during the last two years to improve the fidelity of cat-qubit gates — e.g. [Xu *et al.* 2022a, Xu *et al.* 2022b] or Chapter 2, with respective benefits and feasibility strongly depending on the particular experimental setup envisioned.

The present work pursues with these research efforts, proposing alternatives whose performance is competitive at least in some contexts, and which may sometimes be naturally combined with these other strategies. This chapter thus introduces four new designs for $Z(\theta)$, CNOT and Toffoli gates on dissipatively stabilized cat qubits to help mitigate the incoherent phase errors induced by those gates.

To answer this limitation, our main approach is to reduce the logical information which the two-photon dissipation carries away to the environment under Zeno driving and which induces phase-flip backaction. Our first two designs maintain the same Zeno drives and interaction with the buffer mode mediating two-photon dissipation, but they feed information back from the buffer state to the cat-qubit system before it can leak out to the environment. With this principle, we were able to improve gate fidelities by up to two orders of magnitude with realistic experimental parameters. Our other two designs instead modify the Zeno gate drive, to avoid pushing any information from the cat qubit to the dissipative buffer in the first place. One solution, based on locally flat Hamiltonians, promises a polynomial improvement in the scaling of gate errors with the cat size. Our final design demonstrates exponentially small $Z(\theta)$ gate errors with a tailored dissipation to an ancillary qubit. While its generalization to CNOT gates is beyond the current state of the art, it provides a new way to leverage dissipation for gate engineering.

This Chapter is organized as follows. We begin in Section 3.2 by providing a new perspective on the origin of gate errors using the buffer mode. Section 3.3 gives a short summary of the different gate error mitigation de-

signs introduced in this chapter, and each design is then detailed separately in Section 3.4. Finally, Section 3.5 explores the experimental realization of these designs and their main limitations. We conclude in Section 3.6.

3.2 Zeno gate errors

In Chapter 1 and then in Chapter 2, we reviewed the standard way of engineering gates with dissipative cat qubits, based on the Zeno effect. By weakly driving the cat qubit mode while monitoring it through a two-photon dissipative process, one can navigate the computational space and engineer gates. However, such gate designs induce gate errors that scale only linearly with the relevant parameters, here the two-photon dissipation rate and the gate time. In this section, we provide a new perspective on the origin of these gate errors by considering the buffer mode into the derivation. This will provide intuition for the following sections to find gate designs that can mitigate these large gate errors.

Let us consider the full Hamiltonian engineered for $Z(\theta)$ gates in the presence of the buffer mode,

$$\mathbf{H} = \mathbf{H}_{AB} + \mathbf{H}_Z. \quad (3.1)$$

where $\mathbf{H}_{AB} = g_2(\mathbf{a}^2 - \alpha^2)\mathbf{b}^\dagger + \text{h.c.}$ is a two-to-one photon exchange Hamiltonian between memory and buffer, and $\mathbf{H}_Z = \varepsilon_Z \mathbf{a}^\dagger + \varepsilon_Z^* \mathbf{a}$ a gate drive. Together with the high damping rate of the buffer, the first term mediates the two-photon dissipation while the second term drives the required gate. Let us move into the SFB as introduced in [Chamberland *et al.* 2022] (see also Section 1.3.2 for a short review). This change of basis reads

$$\mathbf{a} \rightarrow \boldsymbol{\sigma}_z \otimes (\tilde{\mathbf{a}} + \alpha) \quad (3.2)$$

and effectively represents a displaced oscillator where the displacement is conditional on $\boldsymbol{\sigma}_z$, the Pauli operator corresponding to the cat qubit logical state. This change of basis is non-orthonormal and approaches degeneracy at high Fock states, but it is sufficiently close to a regular change of coordinates for states close to $|\pm\alpha\rangle$, and thus adequate for our investigation of local errors in the short time limit. Then, $\tilde{\mathbf{a}}$ is a gauge mode that models a local oscillator around the $|\pm\alpha\rangle$ coherent states; in particular, $\tilde{\mathbf{a}}$ in vacuum is equivalent to being perfectly inside the cat-qubit encoding space. With this definition, the Hamiltonian (3.1) reads ($\alpha \in \mathbb{R}$)

$$\tilde{\mathbf{H}} = g_2(\tilde{\mathbf{a}}^2 + 2\alpha\tilde{\mathbf{a}})\mathbf{b}^\dagger + \varepsilon_Z \boldsymbol{\sigma}_z (\tilde{\mathbf{a}}^\dagger + \alpha) + \text{h.c.} \quad (3.3)$$

The ideal Hamiltonian that implements the $Z(\theta)$ rotation of the qubit (and nothing else) now appears as the term $(\alpha\varepsilon_Z \boldsymbol{\sigma}_z + \text{h.c.})$. In the following, we

move into the rotating frame of this Pauli Hamiltonian to simplify the analysis. In addition, in the limit of a small Zeno drive, the effective displacements on the gauge mode $\tilde{\mathbf{a}}$ and on the buffer mode \mathbf{b} are small, such that we can neglect the second-order term $\tilde{\mathbf{a}}^2\mathbf{b}^\dagger + \text{h.c.}$. The corresponding Hamiltonian thus reads

$$\tilde{H}' \approx 2\alpha g_2 \tilde{\mathbf{a}}\mathbf{b}^\dagger + \varepsilon_Z \sigma_z \tilde{\mathbf{a}}^\dagger + \text{h.c.} \quad (3.4)$$

Writing the master equation with the Hamiltonian of (3.4) and in the Heisenberg picture [Breuer *et al.* 2007] for both $\tilde{\mathbf{a}}$ and \mathbf{b} yields a set of coupled equations,

$$\dot{\tilde{\mathbf{a}}} = -2i\alpha g_2 \mathbf{b} - i\varepsilon_Z \sigma_z \quad (3.5a)$$

$$\dot{\mathbf{b}} = -2i\alpha g_2 \tilde{\mathbf{a}} - \kappa_b \mathbf{b}/2 \quad (3.5b)$$

Decoupling these equations gives the second-order differential equations

$$\ddot{\tilde{\mathbf{a}}} + \frac{1}{2}\kappa_b \dot{\tilde{\mathbf{a}}} + \nu^2 \tilde{\mathbf{a}} = -\frac{i}{2}\kappa_b \varepsilon_Z \sigma_z - i\dot{\varepsilon}_Z \sigma_z \quad (3.6a)$$

$$\ddot{\mathbf{b}} + \frac{1}{2}\kappa_b \dot{\mathbf{b}} + \nu^2 \mathbf{b} = -\nu \varepsilon_Z \sigma_z \quad (3.6b)$$

where $\nu \equiv 2\alpha g_2$, and we have used that $d\sigma_z/dt = 0$ since σ_z commutes with (3.4). Finally, making the inverse change of basis on the cat qubit mode, i.e. $\sigma_z \otimes (\tilde{\mathbf{a}} + \alpha) \rightarrow \mathbf{a}$ yields

$$\ddot{\mathbf{a}} + \frac{1}{2}\kappa_b \dot{\mathbf{a}} + \nu^2 \mathbf{a} = \nu^2 \alpha \sigma_z - \frac{i}{2}\kappa_b \varepsilon_Z - i\dot{\varepsilon}_Z \quad (3.7a)$$

$$\ddot{\mathbf{b}} + \frac{1}{2}\kappa_b \dot{\mathbf{b}} + \nu^2 \mathbf{b} = -\nu \varepsilon_Z \sigma_z \quad (3.7b)$$

Both the cat qubit and buffer mode are thus described by a damped harmonic oscillator equation with natural frequency ν and damping rate $\kappa_b/2$. These oscillators are further driven out of equilibrium by the ε_Z drive.

On the cat qubit mode, we naturally find that equilibrium is given by $\mathbf{a}_{eq} = \alpha \sigma_z$ when $\varepsilon_Z = 0$, which corresponds to the computational states of cat qubits up to exponentially small corrections that were neglected by introducing the SFB. We also find that, for $\varepsilon_Z > 0$, the mode is displaced along the $\langle \text{Im}(\mathbf{a}) \rangle$ quadrature independently of σ_z , as expected.

On the buffer mode, it is quite interesting to note that the right-hand side term is proportional to σ_z , such that the buffer mode is displaced in opposite directions depending on the computational state of the cat qubit mode. In other words, the equilibrium position for the buffer mode is $\mathbf{b}_{eq} = -\varepsilon_Z/\nu \sigma_z$ and more importantly, $\mathbf{b} \propto \sigma_z$ at all times. Because the buffer mode is largely damped, the environment obtains information about the state of the cat qubit mode through this effect. As a consequence, measurement of the bit value by the environment dissolves the information contained in the superposition of these bit states, and thus induces phase errors. In fact, it is simple to re-derive




Figure 3.1: (a) Cat qubit mode a stabilized by a two-to-one photon exchange with buffer mode b , mediated by a nonlinear circuit element (middle). The gate drive, by displacing the cat mode, induces a small splitting of the buffer mode conditioned on the cat state. Information is then lost to the environment through the high damping rate of the buffer, hence inducing incoherent gate errors. (b) Average displacement of the buffer (color, left scale) and cat mode (black, right scale) before, during and after a Zeno $Z(\pi)$ gate of duration T , with a cat mode initialized in $|0_L/1_L\rangle \approx |\pm\alpha\rangle$ and buffer initialized in vacuum. Markers show numerical data obtained from integration of the full master equation. Lines show (3.7).

the well-known result of the $Z(\theta)$ gate non-adiabatic error rate by replacing \mathbf{b} with \mathbf{b}_{eq} in the dissipator of (1.38) and integrating over the gate duration. This yields

$$p_Z = p_Z^{(0)} \equiv \frac{\theta^2}{16|\alpha|^4 T} \frac{\kappa_b}{4g_2^2} \quad (3.8)$$

which is the same result as in [Chamberland *et al.* 2022] with $\kappa_2 \equiv 4g_2^2/\kappa_b$.

Figure 3.1(b) shows the average displacement of the buffer mode along the $\langle \text{Re}(\mathbf{b}) \rangle$ quadrature before, during, and after a $Z(\pi)$ gate of duration T . Depending on the initial state of the cat qubit mode, either $|0_L\rangle$ (blue) or $|1_L\rangle$ (yellow), the buffer mode is displaced in one direction or the other. In fact, we observe the dynamics of a damped oscillator with equilibrium position $\mathbf{b}_{eq} = \pm\epsilon_Z/\nu$ for $0 < t < T$ and $\mathbf{b}_{eq} = 0$ for $t > T$, in excellent agreement with

Table 3.1: Comparison of different designs of $Z(\theta)$ gates for dissipative cat qubits. All designs but the last one can be generalized to two- and three-qubit CNOT and Toffoli gates. For the second and last designs, σ_+ denotes a Pauli creation operator on some ancillary qubit. \mathbf{H}_{AB} and \mathbf{H}_Z are defined in Equations (1.38) and (1.54) respectively, and $\mathbf{a}_\theta = \sin(\theta/2)\mathbf{a} + i\cos(\theta/2)\alpha$. In the last column, p_Z denotes the probability of gate errors over a single $Z(\pi)$ gate of duration T .

	Hamiltonian \mathbf{H}	Dissipator \mathcal{D}	Gate Errors
Mirrahimi 2014, Guillaud 2019	$\mathbf{H}_{AB} + \mathbf{H}_Z$	$\kappa_b \mathcal{D}[\mathbf{b}]$	$p_Z^{(0)} \equiv \frac{\pi^2}{16 \alpha ^4 T} \frac{\kappa_b}{4g_2^2}$
Chapter 2	$\mathbf{H}_{AB} + \mathbf{H}_Z + \mathbf{H}_{TPE}$ with $\mathbf{H}_{TPE} \equiv g_2'(\mathbf{a}^2 - \alpha^2)\sigma_+ + \text{h.c.}$	$\kappa_b \mathcal{D}[\mathbf{b}]$	$p_Z = \frac{1}{1 + (2g_2'/\kappa_2)^2} p_Z^{(0)}$
Section 3.4.1	$\mathbf{H}_{AB} + \mathbf{H}_Z$	$\kappa_b \mathcal{D}[\mathbf{b}]$ (photodetected)	$p_Z \gtrsim (1 - \eta)p_Z^{(0)}$ (detection efficiency η)
Section 3.4.2	$\mathbf{H}_{AB} + \mathbf{H}_Z$	$\kappa_{ab} \mathcal{D}[\mathbf{ab}]$	$p_Z = \mu p_Z^{(0)}$ with $\mu \gtrsim 0.02$
Section 3.4.3	$\mathbf{H}_{AB} + \mathbf{H}_{Z,N}$ with $\mathbf{H}_{Z,N} \propto \sum_{n=0}^N c_n (\mathbf{a} + \mathbf{a}^\dagger)^{2n+1}$	$\kappa_b \mathcal{D}[\mathbf{b}]$	$p_Z = \nu \alpha ^{-2N} p_Z^{(0)}$ with $\nu \sim 1$
Section 3.4.4	\mathbf{H}_{AB}	$\kappa_b \mathcal{D}[\mathbf{b}]$ $\kappa_Z \mathcal{D}[\mathbf{a}_\theta \sigma_+]$	$p_Z = e^{-\kappa_Z \alpha ^2 T}$

our analysis based on the approximate model (3.7). The displacement of the cat qubit mode is also shown in black and is independent of the initial state, as expected.

This derivation is meant to provide intuition to the reader about the origin of gate errors. With this intuition in mind, the following sections will introduce multiple gate designs to reduce the errors induced by cat qubit Zeno dynamics, beginning with a summary.

3.3 Summary of gate designs

With the analysis of the previous section in mind, we understand that the loss of phase information during gates is due to the conditional displacement of the buffer mode which is then measured by the environment. This is indeed the only non-unitary channel, hence how quantum information can be lost. Compared to a system with no ancillary buffer in which the information would be directly lost to the environment, the delay provided by the buffer mode can be exploited to reduce gate-induced phase errors without tampering with the ongoing gate.

In this chapter, we introduce two methods that rely on kicking back qubit information that has been transferred into the buffer mode, and two methods that rely on the reduction of information transfer to the buffer mode in the first place. They are all summarized in Table 3.1 along with the regular Zeno-based gate of [Mirrahimi *et al.* 2014] and the combined confinement method of Chapter 2. Although the table only tackles these methods in the scope of the $Z(\theta)$ gate, they can be generalized to multi-qubit gates and aim to represent a wide range of ideas for the mitigation of dissipative cat qubit gate errors. Furthermore, some of these ideas can in principle be combined to attain even higher fidelities.

3.3.1 Buffer photodetection with classical feedback

In the first method, detailed in Section 3.4.1, the principle is to retrieve the information leaking out by directly measuring the field coming out of the buffer mode, instead of letting it get lost to the environment. An appropriate feedback action can then restore this information back into the cat qubit system. The measurement is a photon counter, free of bit value information contained in the field quadratures. The feedback action corresponds to additional Pauli Z gates performed either in software, or through a modification of the gate drive amplitude and/or duration. Alternatively, one can pursue a heralded gate. An immediate limitation of this technique is detector efficiency which will directly limit the proportion of information loss which we can counter with respect to (3.8). Despite this limitation, this method could become viable with the rapid improvement in circuit-integrated photodetectors, and is in any case instructive for the following design.

3.3.2 Cat-buffer autonomous feedback

A second method based on the buffer information is presented in Section 3.4.2. The idea is again that feedback on the cat qubit photon-number parity is applied after detection of a buffer mode photon. However, instead of actually applying a measurement and action, this loop is now applied autonomously thanks to a tailored dissipation operator. This tailored dissipation takes the form $\mathcal{D}[ab]$ such that any time the buffer loses a photon to the environment — and by doing so swaps the cat qubit parity, as we showed in Section 3.2 —, a second parity-switch is applied on the cat qubit through the loss of a single cavity photon. With this autonomous feedback, “detection efficiency” is in principle perfect and a parameter-independent improvement in fidelity of about two orders of magnitude is numerically demonstrated. This design can further be generalized to any multi-qubit C^nX gate with no additional multi-qubit interactions. The residual phase errors with this design are due

to second-order effects that cause imperfections in the applied feedback, and thermal noise in the buffer.

3.3.3 Locally flat Hamiltonian

This third method is the first of a second strategy which consists in minimizing the amount of computational information transferred from the memory to the buffer mode by the gate process. Doing so, the environment cannot in turn receive information by measuring buffer output photons, so qubit coherence is preserved. This gate design is presented in Section 3.4.3. It introduces drive Hamiltonians of the form $\mathbf{H} = f(\mathbf{x})$ where $\mathbf{x} = \mathbf{a} + \mathbf{a}^\dagger$ is the real field quadrature operator on the cat qubit mode, and such that $f(x)$ describes a quasi-potential that is locally flat around $x = \pm\alpha$ but with different mean values $f(\alpha) \neq f(-\alpha)$. The first condition ensures that the drive Hamiltonian is approximately constant over the \mathbf{x} -eigenstates spanned by each computational state $|\pm\alpha\rangle$, such that it induces almost no dynamics on them and they stay inside the codespace throughout the gate. The second condition ensures that each computational state picks up a different phase, hence a rotation about the Z axis. Such Hamiltonians can be engineered with various orders of odd polynomials in \mathbf{x} , and, in the limit of a polynomial of infinite order, $f(x)$ would essentially become the sign function; then exponentially low gate errors in $|\alpha|^2$ are demonstrated, only limited by the finite overlap of coherent states $|\alpha\rangle$ and $|\alpha\rangle$.

3.3.4 Discrete jump

Section 3.4.4 introduces a cat qubit gate design based on a tailored ‘discrete’ dissipation. Concretely, through interaction with an ancillary qubit mode, exactly a single photon is subtracted from the system and the cat state is mapped onto the same state with a gate applied, and so in a discrete manner. Furthermore, for the specific angle $\theta = \pi$, the system stays exactly within its codespace during gates, so no information is transmitted to the environment, opening the door to exponentially low gate errors.

The main limitation of this gate design lies in the introduced ancillary qubit. Indeed, ancillary qubit relaxation would result in additional $Z(\theta)$ rotations for $Z(\theta)$ gates, or additional $Z(\pi)$ rotations for CZ gates. In the case of CNOT or Toffoli gates, it would induce X gate errors, thus killing the error bias. While the dissipator to be engineered for the $Z(\theta)$ and CZ gates is feasible with current state of the art experiments, a viable way to engineer the required CNOT and Toffoli dissipators thus remains to be found.

3.3.5 Combined dissipation and two-photon exchange Hamiltonian

Finally, the combined two-photon exchange Hamiltonian and two-photon dissipation method introduced in Chapter 2 can also yield high-fidelity gates. Indeed, thanks to the additional two-photon exchange Hamiltonian confinement, the effective displacement of the cat qubit mode during a gate is greatly reduced such that less information is transmitted to the buffer mode and thus to the environment. This provides a very simple design for gate error mitigation, especially considering the similarity between this Hamiltonian and the one required for the usual dissipative confinement. This design is not further treated in this chapter, and we refer to Chapter 2 for more details.

3.3.6 Combining designs

As a final remark to this summary, let us note that the gate designs introduced here can be combined together to further improve gate fidelities. As an example, it could be highly favorable to engineer a dissipative cat qubit with the correlated dissipator in $\mathcal{D}[\mathbf{ab}]$, together with a locally flat Hamiltonian to drive gates, plus possibly an optimized activation profile like in [Xu *et al.* 2022a]. We however treat each design separately for clarity.

3.4 Gate designs

In this section, we go over each gate design introduced in this chapter individually, and provide details on their actual implementations, performance and possible limitations.

3.4.1 Buffer photodetection with classical feedback

3.4.1.1 Design principle

Consider the model of (1.38) with a photodetector measuring the output field of the buffer mode. It is governed by a stochastic master equation (SME) that reads [Steck 2007]

$$d\rho = -i[\mathbf{H}_{AB} + \mathbf{H}_Z, \rho] dt + \kappa_b \mathcal{D}_\eta[\mathbf{b}]\rho dt + \mathcal{J}[\mathbf{b}]\rho dN_\eta \quad (3.9)$$

where \mathbf{H}_{AB} denotes two-photon exchange between cat and buffer modes, \mathbf{H}_Z the drive Hamiltonian,

$$\mathcal{D}_\eta[\mathbf{b}]\rho = \mathcal{D}[\mathbf{b}]\rho - \eta(\mathbf{b}\rho\mathbf{b}^\dagger - \langle\mathbf{b}^\dagger\mathbf{b}\rangle\rho) \quad (3.10)$$

is a corrected dissipation that accounts for the backaction of no-detection events, and

$$\mathcal{J}[\mathbf{b}]\rho = \frac{\mathbf{b}\rho\mathbf{b}^\dagger}{\langle\mathbf{b}^\dagger\mathbf{b}\rangle} - \rho \quad (3.11)$$

is a stochastic jump process that accounts for detection events. Here, $\eta \in [0, 1]$ is the detector efficiency and dN_η denotes a stochastic counting process such that it is unity with probability $\langle dN_\eta \rangle = \eta\kappa_b\langle\mathbf{b}^\dagger\mathbf{b}\rangle dt$ and zero otherwise. Like the non-stochastic master equation, the SME of (3.9) features the cat qubit codespace with the buffer in vacuum as its only subspace of steady states. If the system steers away from these steady states — a process otherwise known as ‘codespace leakage’ which is in particular induced during gates — then the buffer mode will get populated through the two-photon exchange interaction and it will output photons to the detector through its large damping rate. During this process, the detector may click, depending on the average buffer mode population and on detection efficiency.

An alternate explanation to this process is found from the viewpoint of four-wave mixing, as represented in Figure 3.2(a). Indeed, during the dynamics, it is possible for 1 photon of the gate drive and 1 photon of the cavity (both at frequency ω_a) to be converted into 1 photon of the buffer (at frequency ω_b) and 1 photon of the microwave pump (at frequency $2\omega_a - \omega_b$). Since this buffer photon is then emitted to the environment, the cat qubit cavity is effectively subject to single-photon dissipation events, thus inducing exact parity swaps on the memory.

During a $Z(\theta)$ gate, we thus have the following situation. First, if no buffer photons are detected during the process, the system follows the dynamics of the no-jump deterministic equation, i.e. (3.9) with $dN_\eta = 0$. During the gate, modes \mathbf{a} and \mathbf{b} get entangled, but after the gate, \mathbf{b} asymptotically relaxes back to vacuum under the no-detection backaction. The cat qubit then gets back to a pure state in which its phase information has been perfectly preserved (in addition to the bit value, which is exponentially protected).

Alternatively, one or several buffer photons can be detected during the process. This time, the system also follows the no-detection dynamics, up until the first detection event, at which point it is projected according to $\rho \rightarrow \mathbf{b}\rho\mathbf{b}^\dagger$. This corresponds to a phase jump of approximately π , with the exact angle depending on jump time and gate parameters. This can be roughly understood by recalling the analysis of Section 3.2, where for an approximate model and in absence of photodetector, we observed that $\mathbf{b}(t)$ in the Heisenberg picture is proportional to σ_z . After a detection-induced jump, we can thus perform classical feedback on the qubit to correct for this π phase shift, hence improving gate fidelities.

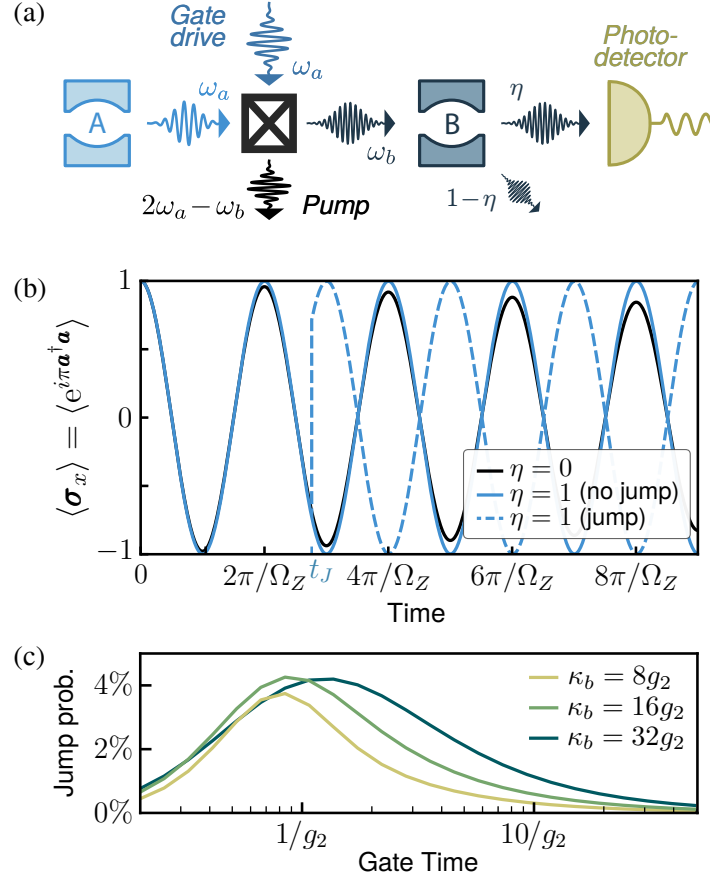


Figure 3.2: (a) Four-wave mixing process showing single cavity photons converted upwards into buffer photons when the gate drive is on. They are then measured by a photodetector on the buffer output. (b) Rabi oscillations with a photodetector on the buffer mode output. The system is initialized in $|+\rangle_L \otimes |0\rangle$ and a constant single-photon drive (1.54) is turned on at $t = 0$. Parity of the cat qubit mode is monitored against time for $\eta = 0$ (no detector) and $\eta = 1$ (perfect detector) both for a no-jump trajectory and a single-jump trajectory. In this numerical simulation, $\Omega_Z \equiv 4\alpha\varepsilon_Z = \pi/g_2$, $\kappa_b = 8g_2$ and $|\alpha|^2 = 8$. (c) Total probability of at least one jump to occur during a $Z(\pi)$ gate, for $|\alpha|^2 = 8$.

3.4.1.2 Measurement strategy

The choice of a photodetector on **b** — instead of e.g. homodyne or heterodyne detection — is motivated as follows. If $\eta = 1$, then no information is lost to the environment and the combined qubit-buffer state remains pure at all times, such that the final cat-qubit state is pure. However, this does not automatically imply that we would be able to perfectly restore the qubit state *before* measurement; indeed, like in a standard perfect measurement, if the

detections contain information about the qubit state, then the complementary qubit information is scrambled by back-action. According to the analysis of Section 3.2, the buffer real quadrature contains information about the qubit being in $|0_L\rangle$ or $|1_L\rangle$, and thus measuring this quadrature would necessarily induce qubit phase decoherence. Therefore, we choose to measure the energy (photon number) of the buffer, which erases this qubit logical information from detection results and hence should imply preservation of the qubit phase (and bit value) for $\eta = 1$. In other words, by measuring the photon number, we prevent that the environment would induce detrimental backaction due to measuring the real quadrature of the buffer.

This picture is in fact exact. Both the Hamiltonian $\mathbf{H}_{AB} + \mathbf{H}_Z$ and the dissipation in $\mathcal{D}[\mathbf{b}]$ commute with a joint x -axis conjugation of both phase spaces. On the output channel \mathbf{b} this conjugation involves a minus sign, which a quadrature measurement could in principle detect, but when measuring $\mathbf{b}^\dagger \mathbf{b}$ this sign strictly disappears from the equations. Then the output signal contains zero information about the logical bit value of the cat qubit; hence, for $\eta = 1$, the phase information of the cat qubit must be perfectly preserved. More details on this invariance of the master equation under joint phase conjugation can be found in Appendix A.

3.4.1.3 Jump and no-jump trajectories

Before discussing the full performance of the design with classical feedback, we examine jump and no-jump trajectories of the scheme separately.

Figure 3.2(a) shows a numerical simulation of Z -axis Rabi oscillations both for a zero-photon-detected trajectory (solid blue) and for a single-photon-detected trajectory (dashed blue), as well as for the standard Zeno gate (black). To evaluate the qubit σ_x expectation value when the system is not exactly in codespace, we take the photon-number parity of mode \mathbf{a} . For the standard Zeno design, gate errors accumulate over time as shown by the decreasing amplitude of oscillations. For the photodetection design however, the qubit phase converges to an indeterminacy of order 10^{-3} (not visible) — reflecting the steady state entanglement of \mathbf{a} and \mathbf{b} modes during gate operation (see Figure 3.1(b)) —, and then keeps oscillating without loss. The single-jump trajectory clearly shows a full dephasing of angle $\simeq \pi$ after the detection event at a random time $t = t_J$, and further keeps oscillating without phase loss. Figure 3.2(b) shows the probability that a detection event occurs during a $Z(\pi)$ gate, as a function of gate time and assuming an ideal photodetector. This jump probability is evaluated as $p_J = 1 - \exp(-\int \langle dN_\eta \rangle(t))$. The low jump probability ensures that trajectories with more than 2 or 3 detection events will be extremely rare.

The perfect preservation of cat qubit phase by gate operation holds after

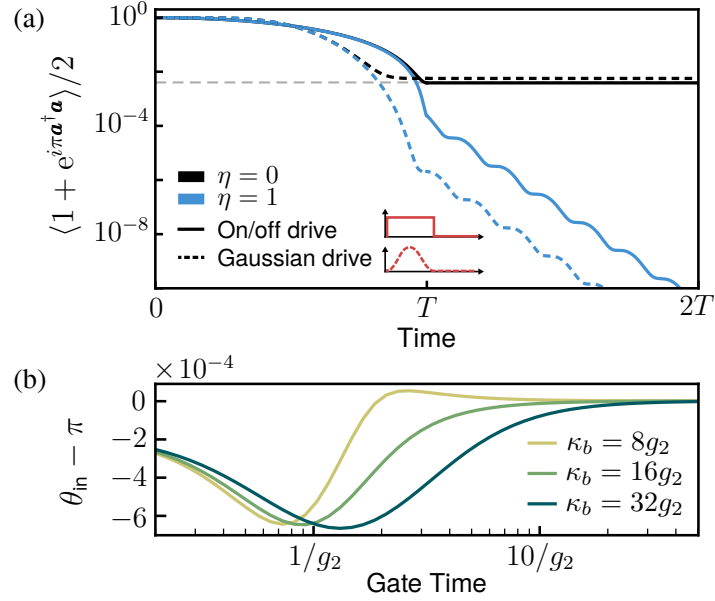


Figure 3.3: (a) Time-evolution of the parity during a $Z(\pi)$ gate with an ideal photodetector on the buffer mode output ($\eta = 1$) and without ($\eta = 0$). Only the no-detection trajectory is shown. The system is initialized in $|+\rangle_L \otimes |0\rangle$ and a single-photon drive (1.54) is on for $0 < t < T$. For $t > T$, the buffer mode reconverges to vacuum as it disentangles from the cat qubit mode. The dashed gray horizontal line shows (3.8). In this numerical simulation, $T = 4/g_2$, $\kappa_b = 8g_2$ and $|\alpha|^2 = 8$. (b) Input angle $\theta_{\text{in}} = 4\alpha \int \varepsilon_Z dt$, such that an exact $Z(\pi)$ gate is produced after infinite-time reconvergence to the buffer mode vacuum.

full disentanglement of the memory and buffer modes, which includes a reconvergence phase to the cat-qubit codespace after the gate drive has been turned off while still monitoring \mathbf{b} with the photodetector. This is illustrated on Figure 3.3(a), for two time-dependent shapes of single-photon drives. For $t \in [0, T]$ i.e. while the gate drive is on, the build-up of a π -angle Rabi oscillation is observed, with better phase precision in presence of a photodetector ($\eta = 1$, blue) than without detector ($\eta = 0$, black). Also, as already observed in [Xu *et al.* 2022a], the gaussian-like time-dependent drive offers significantly better performance because the system mode is closer to its drive-less steady state at the end of the gate. For $t > T$, the single-photon drive is turned off and the reconvergence begins. In absence of a photodetector, the qubit phase remains perfectly constant during reconvergence since photon-number parity is conserved by the two-photon dissipation dynamics. This is shown on Figure 3.3(a) for the black curves. In contrast, when keeping the photodetector on the \mathbf{b} mode for $t > T$, the SME does not preserve photon-number parity due to the non-linear backaction terms featured in (3.11), and there-

fore it is indeed possible to further improve the phase precision as the modes progressively disentangle. The oscillations observed during this reconvergence are due to the damped harmonic oscillator behavior of the mode, as shown in (3.7b).

For $\eta = 1$, in principle there is no limit to phase precision after full disentanglement. To check this, we have adjusted the drive amplitude numerically to perform an ideal $Z(\pi)$ gate at $t \rightarrow \infty$, as shown in Figure 3.3(a). Then, we have checked that the same drive amplitude indeed performs rotations of the same angle whichever the initial cat qubit state. Figure 3.3(b) shows the typical small correction to be applied on the drive amplitude for trajectories without measurement detections.

3.4.1.4 Design performance

We now describe the performance of the full design with buffer mode photodetection and classical feedback. To quantify this, we simulate the following master equation,

$$d\rho = -i[\mathbf{H}_{AB} + \mathbf{H}_Z, \rho] dt + \kappa_b \mathcal{D}_\eta[\mathbf{b}]\rho dt + \mathcal{J}[\mathbf{Z}(\pi)\mathbf{b}]\rho dN_\eta \quad (3.12)$$

which is the same as (3.9) but with a jump operator in $\mathbf{Z}(\pi)\mathbf{b}$ that indicates a Pauli correction on the memory mode for every buffer photon detected. Note that this stochastic master equation does not correspond to any physical model. We only introduce it to explain and quantify the idea behind the design. The actual feedback should be applied separately from the photodetection, for instance in software before any non-Clifford gate or with a subsequent $Z(k\pi)$ gate where k is the number of detected photons. Alternatively, erasure errors can be included in the model, in which case qubits with at least one buffer photon detected during gates are discarded. Hence, with this plethora of possible feedback strategies, we limit our study to (3.12).

The $Z(\pi)$ gate performance achieved after a finite time with this design is shown on Figure 3.4(a). The curves result from an average over several realizations for which the photodetector may have clicked at different times, and the drive is optimized for the zero-detection trajectory to achieve an ideal gate after full reconvergence to the codespace. For this feedback scheme, the gate fidelity is limited by finite detection efficiency as shown by the linear scaling of dashed blue lines that correspond to $\eta = 0.5$ and $\eta = 0.9$. For the ideal photodetector $\eta = 1$, gate fidelities scale with a high-order polynomial in the gate time. In principle, towards ultimate precision, one could perform an ideal feedback of angle $\phi(t_J) \approx \pi$ that depends on the exact jump time, and obtain an error-less gate. This is discussed further in Section 3.4.1.6.

Generalization of the design to multi-qubit CNOT and Toffoli gates is quite straightforward. Considering only the simpler scheme of [Gautier *et al.* 2022]




Figure 3.4: Phase errors of $Z(\pi)$ and CNOT gates for the standard Zeno design ($\eta = 0$, black) and with a photodetector on the buffer mode output of the cat qubits (blue). The figure shows an average over every possible stochastic trajectory for the photodetection design of (3.12). Gate drives are Hamiltonians (1.54) and (1.58) with gaussian time-dependence. Dashed gray lines show analytical gate errors [Chamberland *et al.* 2022]. In these numerical simulations, $\kappa_b = 8g_2$ and $|\alpha|^2 = 8$.

for which target mode stabilization is turned off during the gate process, photodetection should be performed on the buffer mode output of the control qubit(s) only. The numerical performance of this design for the CNOT gate is shown on Figure 3.4(b). Here, feedback is also assumed to be perfectly applied following every buffer mode photodetection, according to (3.12). Similarly to the single-qubit gate, we find several orders of magnitude fidelity improvement in the ideal photodetector case, and otherwise a fidelity improvement limited by detection efficiency. In the next section, we discuss non-ideal photodetectors in more details.

3.4.1.5 Non-ideal photodetector

A realistic photodetector is never ideal and features a finite detection efficiency η . In this case, only part of the information lost to the environment is retrieved, and the resulting gate features dynamics in between the two regimes $\eta = 0$ and $\eta = 1$. A lower bound on the error is then given by

$$p_Z \gtrsim (1 - \eta) p_Z^{(0)}, \quad (3.13)$$

where $p_Z^{(0)}$ is the phase error of the regular Zeno gate, for instance as given by (3.8) for the $Z(\theta)$ gate. Indeed, a fraction $(1 - \eta)$ of the state would behave as in the absence of a photodetector. We denote this as approximate because, when the photodetector does click, we can herald a successful detection trajectory if such heralding is compatible with the rest of the architecture.




Figure 3.5: Scaling of phase errors of a $Z(\pi)$ gate against detection inefficiency $1 - \eta$, for increasing values of $|\alpha|^2$. Gate errors are evaluated at $t = T = 4/g_2$ without reconvergence ($T_c = 0$). Markers show numerical integration of the SME with classical feedback. Solid lines show (3.13), with a linear scaling with $1 - \eta$. In this numerical simulation, $\kappa_b = 8g_2$.

Figure 3.5 shows the numerical scaling of phase errors of a $Z(\pi)$ gate against the detection inefficiency $1 - \eta$ in log-log scale. We indeed find a linear scaling (+1 slope in log scale) with the amount of information lost to the environment according to (3.13), that eventually saturates for large enough detection efficiencies. This saturation results from the imperfect reconvergence to the codespace, as previously discussed.

In practice, a non-ideal photodetector does not only feature a non-unity detection efficiency, it may also have a finite dark count rate (the detector clicks without any photon measured) and a finite uncertainty on the detection time. The former would result in the erroneous application of a π -angle feedback at the dark count rate similar to thermal photons in the buffer mode (see Section 3.5.2). The latter has no impact on the design fidelity as a $Z(\pi)$ correction gate is applied independently of the detection time. In any case, both of these effects are widely negligible compared to detection inefficiency, which even with state of the art photodetectors is typically in the $\eta = 0.1$ – 0.5 range. With the rapid improvement in photodetector efficiencies [Albertinale *et al.* 2021, Dassonneville *et al.* 2020, Lescanne *et al.* 2020a], we may however expect this gate design to become viable in the coming years (see also Figure 3.14(a)). The design can also inspire other feedback methods in which feedback is hardware-efficient or made autonomously, such as the one introduced in the following section.




Figure 3.6: (Black) Angle of the Z rotation feedback to perform after a buffer mode photodetection at any given time, for a $Z(\pi)$ gate of duration T . (Blue) Buffer mode population for the no-jump trajectory, proportional to the photodetection probability $\langle dN_\eta \rangle$. In these numerical simulations, $\kappa_b = 8g_2$ and $|\alpha|^2 = 8$.

3.4.1.6 Ideal feedback action

In this section, we discuss the ideal feedback action to be applied assuming a perfect photodetector with detection efficiency $\eta = 1$ and instantaneous feedback. Such an ideal feedback can in principle yield error-less gates. Indeed, if $\eta = 1$ and the measurement does not probe the computational state, then no information is lost to the environment and the combined qubit-buffer state remains pure at all times, such that the final cat-qubit state is pure. An appropriate feedback action can thus retrieve the desired angle of rotation.

Figure 3.6 shows the profiles of optimal feedback angles $\pi + \delta(t_J)$ upon photon detection, as a function of detection time t_J , for a targeted $Z(\pi)$ gate. The mean photon population of the buffer mode is also shown as a function of time, which is proportional to the instantaneous photodetection probability $\langle dN_\eta \rangle$. Note that close to $t = T$ in the middle plot, the angle δ diverges to negative values, but $\langle b^\dagger b \rangle \ll 1$ at this time such that it is excessively unlikely to ever have to perform this feedback action.

To perform this classical feedback upon detection of buffer photons, multiple solutions exist. The first and most straightforward is to actually add this $Z(\pi + \delta)$ gate to the actuation, which would *occasionally* increase the gate time depending on the input gate angle θ and detection time t_J . For instance, if a photon is detected at the beginning of a $Z(0.9\pi)$ gate, then a shorter $Z(-0.1\pi + \delta)$ gate should be performed instead, while if it is detected

towards the end then a longer $Z(1.9\pi + \delta)$ gate would be realized. While this happens only for photon-detected trajectories, and thus with low probability, it causes unpredictable gate times.

As a second option, the feedback could adjust the final $Z(\theta)$ gate up to an integer multiple of π (similarly, adjust the CNOT gate up to an integer number of $Z(\pi)$ on the control qubit). The remaining gate $Z(n\pi)$ is a Pauli gate, which either a specific hardware operation can implement or the software can keep track of by adapting remaining computer operations until the next non-Clifford gate is reached.

For CNOT gates, the output of the target buffer mode should in principle also be monitored due to the reconvergence phase in which the target mode undergoes significant dynamics. Importantly, the effect of this monitoring — whatever the detection results — is only on the phase of the control qubit. Indeed, the bit degrees of freedom remain exponentially protected on both qubits; and since no operator ever couples to the bit-value information of the target qubit (unlike the control qubit, for which $\mathbf{a} + \mathbf{a}^\dagger$ is sensitive to the bit value), it means that its bit-value information can never leak out in this perfectly monitored scheme, and hence no phase-blurring backaction can happen on the target qubit. We have indeed verified numerically that, for $\eta = 1$ and no buffer photons detected, an ideal CNOT gate is retrieved after an infinite time reconvergence by adjusting a $Z(\delta)$ gate on the control qubit with $\delta \ll 1$. A detection on the control qubit buffer output at time t_J requires an additional $Z(\pi + \delta(t_J))$ gate, as for the previously discussed $Z(\theta)$ gate scheme. In contrast, a detection event on the output of the target buffer mode still requires corrections of order δ only. This is consistent with a first-order Heisenberg picture analysis as in Section 3.2, for which \mathbf{b}_C is proportional to $\boldsymbol{\sigma}_{z,C}$ while \mathbf{b}_T does not carry any qubit information.

3.4.2 Cat-buffer autonomous feedback

3.4.2.1 Design principle

This second gate design, inspired by the previous one, introduces a correlated dissipator to remove entropy from the system instead of the standard single-photon dissipation on the buffer mode. The corresponding master equation to be engineered on the cavity-buffer system then reads

$$\frac{d\rho}{dt} = -i[\mathbf{H}_{AB} + \mathbf{H}_Z, \rho] + \kappa_{ab}\mathcal{D}[\mathbf{ab}]\rho \quad (3.14)$$

with the main addition of a two-mode dissipation operator that was recently realized in [Gertler *et al.* 2023] in the context of pair-cat codes stabilization [Albert *et al.* 2019]. The main idea behind this peculiar dissipation arises

from the photodetector scheme of the previous section. When the cat-qubit leaks out of its codespace under the gate drive, the buffer mode is populated through the \mathbf{H}_{AB} interaction, then inducing an eventual phase-flip on the cavity mode as explained in Section 3.2. The correlated dissipation thus ensures that whenever the buffer mode loses a photon (hence inducing a $Z(\pi)$ error on the logical cat qubit), a direct photon loss on the cavity mode is also produced thus switching the cat qubit parity a second time and correcting for the error autonomously.

Mathematically, this can be understood using the SFB [Chamberland *et al.* 2022] transformation as defined in (3.2), which yields

$$\begin{aligned}\kappa_{ab}\mathcal{D}[\mathbf{ab}] &\rightarrow \kappa_{ab}\mathcal{D}[\boldsymbol{\sigma}_z(\tilde{\mathbf{a}} + \alpha)\mathbf{b}] \\ &= |\alpha|^2\kappa_{ab}\mathcal{D}[\boldsymbol{\sigma}_z\mathbf{b}] + \mathcal{O}(|\tilde{\mathbf{a}}^\dagger\tilde{\mathbf{a}}|^{1/2})\end{aligned}\tag{3.15}$$

where the second line approximation holds since $|\tilde{\mathbf{a}}^\dagger\tilde{\mathbf{a}}| \ll 1$ in the limit of a small amount of leakage. Keeping only the leading order term in (3.15), and taking $\kappa_b \equiv |\alpha|^2\kappa_{ab}$, it is possible to perform the same gate error derivation as in Section 3.2. This yields the exact same resulting set of equations on \mathbf{a} and \mathbf{b} , but with a different dissipation operator in $\boldsymbol{\sigma}_z\mathbf{b} \propto \boldsymbol{\sigma}_z^2 = \mathbf{I}$ where \mathbf{I} is the identity on the two-level cat qubit mode; i.e. the environment does not receive any information about the qubit state, as required. We emphasize once again that this derivation is only first-order and that non-linear effects have been neglected, for instance with terms in $\tilde{\mathbf{a}}^2\mathbf{b}^\dagger + \text{h.c.}$ that may eventually limit gate performances.

To engineer the correlated dissipation (3.14), an ancillary low-Q reservoir mode \mathbf{r} can be introduced into the setup [Gertler *et al.* 2023]. By engineering four-wave mixing between those three modes and a classical pump, it is then possible to enable Hamiltonian interaction of the form $g_{ab}\mathbf{abr}^\dagger + g_{ab}^*\mathbf{a}^\dagger\mathbf{b}^\dagger\mathbf{r}$. Together with a large reservoir damping of the form $\kappa_r\mathcal{D}[\mathbf{r}]$ and in the limit of $\kappa_r \gg g_{ab}$, it is possible to eliminate the fast dynamics of the reservoir mode. On the reduced system of the cat and buffer modes, the required correlated dissipator is then obtained, with typical amplitude $\kappa_{ab} = 4g_{ab}^2/\kappa_r$ [Azouit *et al.* 2017].

In the limit of $\kappa_{ab}|\alpha|^2 \gg g_2$, it is further possible to eliminate the fast dynamics of the buffer mode to obtain an effective single-mode master equation on the cat qubit. Using the effective operator formalism of [Reiter & Sørensen 2012] yields the following dynamics on the cat qubit mode,

$$\frac{d\rho}{dt} = \frac{4g_2^2}{\kappa_{ab}}\mathcal{D}[\mathbf{a}(\mathbf{a}^\dagger\mathbf{a})^{-1}(\mathbf{a}^2 - \alpha^2)]\rho\tag{3.16}$$

where $(\mathbf{a}^\dagger\mathbf{a})^{-1}$ is the pseudo-inverse of the photon number operator and describes the effective difference in dynamics undergone during transient buffer

excitation. This master equation indeed describes a parity-switching stabilization of the cat mode, but it differs from the master equation targeted in [Xu *et al.* 2022b] which instead reads $\mathcal{D}[(c_1\mathbf{a} + c_2\mathbf{a}^\dagger)(\mathbf{a}^2 - \alpha^2)]$ for non-squeezed cat states, with appropriate coefficients $c_{1/2}$. However, both describe a similar dynamics and autonomous correction principle. More details on this model reduction can be found in Section 3.4.2.5.

3.4.2.2 Parity-switching dynamics

This dissipation with $\mathcal{D}[\mathbf{a}\mathbf{b}]$ is activated more generally when the cat-qubit leaks out of its codespace, triggering a reaction of its buffer. The effect will be beneficial whenever the leakage source is associated to parity-switching. For standard cat qubits, this is mainly the case for the user-induced Zeno dynamics as just discussed (for $Z(\theta)$ gates, on the control qubit of CNOT gates, or other similar gates). It can also occur by thermal excitation of the cat mode, of the form $\mathcal{D}[\mathbf{a}^\dagger]$. In contrast, leakage that preserves the photon-number parity would then induce phase errors by following the same process. This is for instance true of pure dephasing, of the form $\mathcal{D}[\mathbf{a}^\dagger\mathbf{a}]$. However, the amplitude of such effects is often negligible compared to other sources of phase errors such as gates and finite resonator lifetime.

The recent preprint of [Xu *et al.* 2022b] also explores how single-photon losses, of the form $\mathcal{D}[\mathbf{a}]$, can induce parity-switching leakage on a squeezed cat qubit. Indeed, while the annihilation operator leaves the coherent states constituting regular cats in place, it does induce leakage on squeezed coherent states [Schlegel *et al.* 2022]. As such, similarly to [Xu *et al.* 2022b], the correlated dissipator introduced in our work performs an autonomous correction of single-photon annihilation on squeezed cats, to first order; a better correction would be obtained with dissipation in $\mathcal{D}[S(\mathbf{a})\mathbf{b}]$ where $S(\mathbf{a})$ is the squeezed annihilation operator, featuring the squeezed coherent state as an eigenstate.

One of the main limitations of such approaches stems from thermal noise in the buffer mode, since any buffer excitation would decay by triggering a phase flip on the cat qubit. This effect is further discussed in Section 3.5.2.

3.4.2.3 Design performance

Let us now focus on the performance of this gate design, starting with the single-qubit $Z(\theta)$ gate. Figure 3.7 shows the errors induced by a $Z(\theta)$ gate for the regular Zeno gate (black) and for the correlated dissipator design of (3.14) (blue) for both phase-flip and bit-flip errors. Following our previous analysis, the comparison between both designs is made at fixed $\kappa_b = |\alpha|^2\kappa_{ab} = 8g_2$ in order to keep the same damping rate. Left-side plots show the scaling with




Figure 3.7: Gate errors of a $Z(\pi)$ gate with the autonomous feedback design of (3.14) at $\kappa_{ab}/|\alpha|^2 = 8g_2$ (blue) and with the regular Zeno design at $\kappa_b = 8g_2$ (black). Left: fixed cat size, $|\alpha|^2 = 8$. Right: fixed gate time, $T = 10/g_2$. Gate errors are evaluated at $t = T$ or after full reconvergence to the steady state ($T_c = \infty$). While the cat qubit is entangled with its buffer and leaks out of codespace, the logical phase value is evaluated with photon-number parity on \mathbf{a} and the logical bit value is evaluated with two-photon dissipation invariant, see Chapter 1.

the gate time, while the right-side is plotted against cat size. For phase-flip errors, an improvement by a constant factor of about $\mu \approx 0.02$ is found, independent of both gate time and cat size, and limited by second-order effects as discussed previously. In the regime of short gate times, the gate drive is large in amplitude, and so the buffer is largely entangled with the cat mode at the end of the gate. For this reason, a reconvergence time is required to reach the constant fidelity gain of μ . In other words, the buffer has not had enough time to lose its excitations to the environment, and so the correlated dissipator has not yet corrected for the coherent errors that occurred during the gate.

The correlated dissipator of (3.14) can raise concerns about bit-flip errors since it would, by itself, stabilize the vacuum state in both resonators. Hence, the bottom plots of Figure 3.7 investigate this bit-flip error and shows that the exponential scaling in $|\alpha|^2$ is preserved, although slightly degraded from that of the regular Zeno gate. This is likely due to the additional leakage induced by the correlated dissipator. Indeed, whenever the buffer mode is populated, the cat mode is pushed towards its own vacuum state, at a rate proportional to




Figure 3.8: Gate errors of a CNOT gate with the autonomous feedback design of (3.14) at $\kappa_{ab}/|\alpha|^2 = 8g_2$ (blue) and with the regular Zeno design at $\kappa_b = 8g_2$ (black). Left: fixed cat size, $|\alpha|^2 = 8$. Right: fixed gate time, $T = 10/g_2$. Gate errors are evaluated at $t = T$ or after full reconvergence to the steady state ($T_c = \infty$). Dashed gray lines show analytical CNOT gate errors [Chamberland *et al.* 2022].

the buffer population. When stopping operation before reconvergence (solid blue curve), the bit-flip rate looks artificially reduced for very short gate times, as the buffer had no time to dissipate yet.

3.4.2.4 Multi-qubit gates

For CNOT and Toffoli gates, the gate drive is parity-preserving on target qubit. As explained for the buffer photodetection gate design, in principle then detecting all buffer relaxations would induce a combination of phase corrections on the control qubit; to first order however, this comes down to only requiring the correlated dissipation on control qubits. Figure 3.8 investigates the performance of such a setup for two-qubit CNOT gates. The correlated dissipator of (3.14) is activated on the control qubit, and the drive Hamiltonian of (1.58) is switched on. Here, the performance is very similar to that of the single-qubit $Z(\theta)$ gate, with a phase fidelity improvement of $1/\mu \approx 50$ in the best case scenario. This is obtained at large gate times when the drive is small compared to the damping rate, $\varepsilon_{CX} \ll \kappa_{ab}|\alpha|^2$.

3.4.2.5 Adiabatic elimination of the buffer mode

In this section, we derive the result presented in Section 3.4.2.1 on the adiabatic elimination of the buffer mode in the presence of a correlated dissipator. The two-mode master equation is initially given by

$$\frac{d\rho}{dt} = -i[\mathbf{H}_{AB}, \rho] + \kappa_{ab}\mathcal{D}[ab]\rho. \quad (3.17)$$

This equation describes an exchange term between cat and buffer modes, as well as a strong correlated dissipation. Here, the fast dynamics corresponds to the deexcitation of the buffer mode thanks to the dissipation term, while the slow dynamics is that of the exchange Hamiltonian. The goal is thus to adiabatically eliminate the fast dynamics when the buffer is excited, and to derive an effective single-mode equation for the A mode. Taking the notations of [Reiter & Sørensen 2012], we have

$$\begin{aligned} \mathbf{V}_+ &= g_2(\mathbf{a}^2 - \alpha^2)\mathbf{b}^\dagger \\ \mathbf{V}_- &= g_2^*(\mathbf{a}^{\dagger 2} - \alpha^{*2})\mathbf{b} \end{aligned} \quad (3.18)$$

which are perturbative (de-)excitations of the system, $\mathbf{L} = \sqrt{\kappa_{ab}}\mathbf{a}\mathbf{b}$ is the jump operator from excited to ground subspaces, and $\mathbf{H}_g = \mathbf{H}_e = 0$ are the block diagonal Hamiltonians in the ground and excited subspaces. In addition, the non-hermitian Hamiltonian in the excited subspace reads

$$\mathbf{H}_{\text{NH}} = -\frac{i}{2}\kappa_{ab}\mathbf{a}^\dagger\mathbf{a}\mathbf{b}^\dagger\mathbf{b}. \quad (3.19)$$

Therefore, the effective single-mode dynamics reads

$$\frac{d\rho}{dt} = -i[\mathbf{H}_{\text{eff}}, \rho] + \mathcal{D}[\mathbf{L}_{\text{eff}}]\rho \quad (3.20)$$

where $\mathbf{H}_{\text{eff}} \propto \mathbf{H}_{\text{NH}}^{-1} + \mathbf{H}_{\text{NH}}^{\dagger -1} = 0$, and

$$\mathbf{L}_{\text{eff}} = \mathbf{L}(\mathbf{H}_{\text{NH}})^{-1}\mathbf{V}_+ = \frac{2ig_2}{\sqrt{\kappa_{ab}}}\mathbf{a}(\mathbf{a}^\dagger\mathbf{a})^{-1}(\mathbf{a}^2 - \alpha^2) \quad (3.21)$$

This describes a parity-switching jump operator with cat qubit steady states. In the semi-classical limit, $\mathbf{a}(\mathbf{a}^\dagger\mathbf{a})^{-1} \sim \alpha^{-1}$, such that the effective two-photon dissipation rate is indeed given by $\kappa_2 \equiv 4g_2^2/\alpha^2\kappa_{ab}$.

3.4.3 Locally flat Hamiltonian

3.4.3.1 Intuition

Quantum states feature a dispersion in position and momentum in their phase-space representation, with an equal variance in the case of coherent states. A single-photon drive, such as the gate drive of (1.54), thus acts differently along this dispersion and according to the specific position and momentum values of the state. This induces a phase-space displacement that can be frozen by a continuous measurement, such as the one of two-photon dissipation. This is the so-called Zeno effect.

From the wavefunction perspective, a coherent state is gaussian in position representation¹ and, for a real coherent amplitude, reads $\psi(x) \propto \exp[-(x - \langle x \rangle)^2/2]$. Since the Zeno drive of (1.54) reads $H_Z(x) = \varepsilon_Z x$ in this same representation, the time-evolution of the wavefunction under H_Z can be trivially obtained as after some time t , the wavefunction reads $\psi(x, t) \propto \exp[-ix\varepsilon_Z t]\psi(x)$. This corresponds to a position-dependent phase shift, and therefore a displacement along the momentum axis p . However, two-photon dissipation prevents this displacement, and only the average phase shift $\exp[-i\langle x \rangle \varepsilon_Z t]$ remains, thus driving a cat qubit $Z(\theta)$ gate together with a logical phase blurring reflecting the variance of $\exp[-ix\varepsilon_Z t]$ over each coherent state.

From this viewpoint, one way to improve the precision of the $Z(\theta)$ gate is to use a drive Hamiltonian with little dispersion over each coherent state. Moreover, this property should be robust to all effects which the cat-qubit is meant to cover, i.e. most prominently local displacements in phase space. Conversely, such a Hamiltonian would induce almost no displacement of the coherent states, nor deformation of any kind. Hence the phase gate could be implemented without relying on two-photon dissipation, and the evolution can be purely unitary which is another way to see that the gate would induce no phase losses.

3.4.3.2 Design principle

In this sense, an ideal² drive Hamiltonian for single-qubit $Z(\theta)$ gates reads

$$\mathbf{H}_{Z,\infty} \equiv \varepsilon_Z \text{sign}(\mathbf{x}) \quad (3.22)$$

where sign denotes the sign function and $\mathbf{x} = \mathbf{a} + \mathbf{a}^\dagger$. This Hamiltonian yields a global phase difference for each half-plane of phase space, and thus engineers the required gate without any phase loss while being robust to any local error — e.g. small displacements or distortions of cat states. It is represented in thin black lines on Figure 3.9(a), together with the position representation of a superposition cat state in the background. Figure 3.9(b) and Figure 3.9(c) show the phase errors induced by such a single-qubit $Z(\pi)$ gate with the Hamiltonian of (3.22) in black lines, both for varying gate time (left) and varying cat size (right). We indeed find that this Hamiltonian provides phase errors that scale exponentially in the cat size, and that are drastically smaller than with the regular Zeno Hamiltonian of (1.54), represented in yellow lines.

¹Here we use the mathematical equivalence with a mechanical harmonic oscillator, to call position x and momentum p the real and imaginary quadratures of the electromagnetic mode.

²This drive Hamiltonian is ideal up to exponentially small corrections, due to the orthonormalization of the computational basis — two coherent states are never exactly or-




Figure 3.9: (a) Position distribution of locally flat drive Hamiltonians of (3.22) and (3.23) that minimize the variance over the quantum fluctuations of cat states. (b, c) Gate phase errors for a $Z(\pi)$ gate with a locally flat drive Hamiltonian at (b) fixed cat size, $|\alpha|^2 = 8$, and (c) fixed gate time, $T = 10/g_2$. In these simulations, $\kappa_b = 8g_2$. Lines show numerical fits of the form $p_Z \propto T^{-1}|\alpha|^{-2(2+N)}$ (color) and $p_Z \propto T^{-1}\exp(-2|\alpha|^2)$ (black). Markers show numerical data.

The Hamiltonian of (3.22) is highly non-linear in \mathbf{x} and is therefore not accessible for state-of-the-art superconducting circuits, nor for other quantum computation platforms. It can however be approximated. Let us define the set of Hamiltonians

$$\mathbf{H}_{Z,N} \equiv \varepsilon_Z \sum_{n=0}^N c_n \mathbf{x}^{2n+1} \quad (3.23)$$

for $N \geq 0$, where c_n are constants to be determined. These Hamiltonians have an odd distribution in position space. With the proper definition of c_n , they can approximate the sign Hamiltonian of (3.22) locally around each cat qubit coherent state, even if the approximation cannot be global. The optimal c_n constants should therefore minimize the gate errors induced by the combination of two-photon dissipation and drive during a given gate. Loosely speaking, these errors scale with the amount of non-flatness of the drive Hamiltonian of (3.23) — a perfectly flat drive Hamiltonian would not

thogonal — or in other words, to the gaussian tails of coherent states that cross the $x = 0$ line in their phase space distributions.

displace the state and would thus result in an error-free gate. Therefore, we minimize the variance of the Hamiltonian across a single coherent state, given by

$$V_N(\{c_n\}) \equiv \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} H_{Z,N}(x)^2 e^{-\frac{1}{2}(x-2\alpha)^2} dx, \quad (3.24)$$

under the constraint of a fixed angle of gate rotation,

$$\frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} H_{Z,N}(x) e^{-\frac{1}{2}(x-2\alpha)^2} dx = \frac{\theta}{2T} \quad (3.25)$$

where $H_{Z,N}(x)$ is the position distribution of Hamiltonian (3.23). This optimization problem is solved numerically using a Lagrange multiplier which is differentiated analytically and then minimized through matrix inversion. More details on this minimization can be found in Section 3.4.3.6.

Figure 3.9(a) shows the first five of these Hamiltonians in colored lines, as determined by the previously-described minimization process. While the $N = 0$ corresponds to the regular Zeno Hamiltonian of (1.54) with a linear distribution, the $N > 0$ Hamiltonians show locally flat distributions around both coherent states, with increasing flatness as N grows.

3.4.3.3 Results

The performance of these Hamiltonians is then evaluated for the single-qubit $Z(\pi)$ gate in Figure 3.9(b) and Figure 3.9(c), with the gate phase error as a function of gate time and cat size respectively. The first plot shows a constant improvement in phase fidelity of the gate as N increases, while the scaling with the gate time stays linear. The second plot however shows an improved scaling of phase errors with the cat size as N grows. Although we lack an analytical derivation of phase errors for this set of drive Hamiltonians, a numerical fit of the form $p_Z \propto |\alpha|^{-2(2+N)}$ is performed as highlighted by the colored lines, and matches the numerical simulations particularly well in the large cat size limit. We attribute this scaling to the fact that, as both coherent states come further apart in phase space, the variance minimization process can achieve flatter distributions with each additional degree of freedom provided by the increasing N .

3.4.3.4 Multi-qubit gates

Generalization of this single-qubit gate design to multi-qubit gates is quite straightforward. Similarly as for the regular Zeno gate, the single-qubit drive of (3.23) should be multiplied by a phase-space rotation on the target qubit,




Figure 3.10: Gate-induced control phase errors for a CNOT gate with a locally flat drive Hamiltonian as in (3.26). Left: fixed cat size, $|\alpha|^2 = 8$. Right: fixed gate time $T = 10/g_2$. Dissipative stabilization is acting on the control qubit only, with $\kappa_b = 8g_2$. Lines show numerical fits of the form $p_{Z_C} \propto T^{-1}|\alpha|^{-2N}$ (color) and $p_{Z_C} \propto T^{-1} \exp(-2|\alpha|^2)$ (black). Markers show numerical data.

which yields

$$\mathbf{H}_{CX,N} \equiv \varepsilon_{CX} \left(\sum_{n=0}^N c_n \mathbf{x}_C^{2n+1} \right) \otimes \left(\mathbf{a}_T^\dagger \mathbf{a}_T - n_p \right) \quad (3.26)$$

where $\mathbf{x}_C = \mathbf{a}_C + \mathbf{a}_C^\dagger$, the $\mathbf{a}_{C/T}$ are annihilation operators on control and target qubits respectively, n_p is any even integer close to $|\alpha|^2$. Together with a static two-photon dissipation on the control qubit, and optionally with the correlated two-photon dissipation on the target qubit [Guillaud & Mirrahimi 2019], this process achieves a CNOT gate.

Figure 3.10 investigates the phase errors induced by this two-qubit gate design on the control qubit, both against gate time (left) and cat size (right). Similar conclusions as for the single-qubit $Z(\theta)$ gate are reached. Gate errors scale linearly with time, and numerically we fit these errors according to $p_{Z_C} \propto |\alpha|^{-2(1+N)}$. The $|\alpha|^2$ difference in scaling compared to the single-qubit gate is due to the target qubit term in the drive Hamiltonian (3.26). Indeed, the operator $\mathbf{a}_T^\dagger \mathbf{a}_T - n_p$ induces a different (integer) number of $Z(\pi)$ gates on the control qubit for each Fock state of the target qubit, and the dispersion on this number increases with α (see [Chamberland *et al.* 2022] for details).

3.4.3.5 Engineering with an ATS

Overall, the high-order Hamiltonians introduced in this section can greatly improve the performance of dissipative cat qubit gates, at the cost of additional difficulties to engineer the required Hamiltonians. Let us however note that high-order non-linearities are always present in superconducting circuits,

even if they are often neglected due to their low amplitudes. Since the c_n coefficients scale as $c_n \propto |\alpha|^{-2n}$, even moderate superconducting non-linearities could be enough to engineer these Hamiltonians, and so especially in the large cat size limit.

To illustrate this possibility, let us consider a realistic set of parameters for the current experimental proposal of dissipative cat qubits based on the ATS [Lescanne *et al.* 2020b, Berdou *et al.* 2022]. The ATS is a nonlinear circuit element made of a SQUID shunted by an inductance, which creates two flux loops that are then threaded at 0 and π flux bias respectively. The ideal Hamiltonian resulting from this setup reads

$$\mathbf{H} = \omega_a \mathbf{a}^\dagger \mathbf{a} + \omega_b \mathbf{b}^\dagger \mathbf{b} - 2E_J \varepsilon(t) \sin(\varphi_a(\mathbf{a} + \mathbf{a}^\dagger) + \varphi_b(\mathbf{b} + \mathbf{b}^\dagger)) \quad (3.27)$$

for an ATS coupled capacitively to both the cat qubit and buffer modes. Here, the $\varepsilon(t)$ term corresponds to a differential flux drive that can be frequency tuned to make specific terms in the sine Hamiltonian resonant. Also, $\varphi_{a/b}$ denote the energy participation of each mode into the ATS. Typically, it is desired to engineer φ_a as large as possible to create strong two-to-one photon exchange rates, bearing in mind that a large φ_a also induces increased high-order non-linear effects and increased single-photon losses by coupling to the buffer transmission line.

For concreteness, let us assume that we want to implement a $Z(\pi)$ gate as in (3.23) with $N = 2$ (i.e. up to 5-th order Hamiltonian terms), with a gate time $T = 500$ ns and for a cat of size $|\alpha|^2 = 8$. In this case, the global drive Hamiltonian amplitude to be engineered reads $\varepsilon_Z/2\pi = 1/8\alpha T \approx 88$ KHz, and the polynomial coefficients of order 1, 3 and 5 read $c_0 \approx 0.66$, $c_1 \approx -0.055$ and $c_2 \approx 0.0021$ respectively. To make each of these coefficients match the ATS Hamiltonian terms, the following identities should be met,

$$\varepsilon_Z c_0 = 2E_J \varepsilon_0 \varphi_a + \mathcal{O}(\varphi_a^3) \quad (3.28a)$$

$$\varepsilon_Z c_1 = 2E_J \varepsilon_1 \varphi_a^3/3! + \mathcal{O}(\varphi_a^5) \quad (3.28b)$$

$$\varepsilon_Z c_2 = 2E_J \varepsilon_2 \varphi_a^5/5! + \mathcal{O}(\varphi_a^7) \quad (3.28c)$$

where $\varepsilon_k \ll 1$ are the flux drives amplitudes such that $\varepsilon(t) = \sum_k \varepsilon_k \cos((2k+1)\omega_a t)$. The right-hand side terms simply result from a Taylor expansion of the sine up to 5-th order in $\varphi_{a/b} \ll 1$. Further assuming realistic experimental parameters of $\varepsilon_k = 0.01$, $\varphi_a = 0.1$ and $E_J/2\pi = 90$ GHz [Lescanne *et al.* 2020b], we find that the right-hand side terms of (3.28) are at least 10 times larger than required by the left-hand side terms. In other words, the achievable experimental parameters are 10 times larger than the actual drive amplitudes to be engineered. This gives some leeway to either implement even higher-order gate designs, or to release experimental constraints. In addition, for the same parameters, a factor of

$|\alpha|^4 = 64$ improvement in gate fidelities can be expected, making the design particularly attractive.

As a final remark, one may also seek to engineer such nonlinear terms e.g. with time-dependent schemes and rotating wave approximation, or by involving auxiliary systems. Although, such attempts should keep in mind the utmost importance of preserving the noise bias throughout gate operation.

3.4.3.6 Optimal drive Hamiltonians

In Section 3.4.3.2, we define drive Hamiltonians with odd-power polynomials of the position operator \mathbf{x} . For cat qubit gate engineering, it is required that these high-order drives are locally flat around both cat qubit coherent components. Mathematically, this corresponds to a minimization of the variance over one coherent state, defined by

$$V_N(\{c_n\}) \equiv \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} H_{Z,N}(x)^2 e^{-\frac{1}{2}(x-2\alpha)^2} dx, \quad (3.29)$$

under the constraint of a fixed mean value of the drive over this coherent state, defined by

$$E_N(\{c_n\}) \equiv \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} H_{Z,N}(x) e^{-\frac{1}{2}(x-2\alpha)^2} dx \quad (3.30)$$

where $H_{Z,N}(x) = \varepsilon_Z \sum_{n=0}^N c_n x^{2n+1}$ is the potential to be optimized, with $N+1$ constants to be determined.

To perform this constrained minimization problem, we use Lagrange multipliers, and define the Lagrangian function as

$$\mathcal{L}_N(\{c_n\}, \lambda) = V_N(\{c_n\}) - \lambda(E_N(\{c_n\}) - \varepsilon_0) \quad (3.31)$$

where λ is a Lagrange multiplier, and ε_0 is the fixed mean value of E_N . Thanks to the simple form of (3.29) and (3.30), it is quite simple to find the global minimum of this Lagrangian function exactly. Differentiating (3.31) with respect to all $N+2$ variables yields

$$\begin{cases} \frac{\partial \mathcal{L}_N}{\partial c_k} = 2 \sum_{n=0}^N I_{2(n+k+1)} c_n - \lambda I_{2k+1} \\ \frac{\partial \mathcal{L}_N}{\partial \lambda} = 1 - \sum_{n=0}^N I_{2n+1} c_n \end{cases} \quad (3.32)$$

where we have defined

$$I_k \equiv \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} x^k e^{-\frac{1}{2}(x-2\alpha)^2} dx. \quad (3.33)$$

Through an integration by parts, a recurrence relation can be obtained for I_k . It reads

$$I_{k+1} = 2\alpha I_k + k I_{k-1}, \quad (3.34)$$

with $I_0(\alpha) = 1$ and $I_1(\alpha) = 2\alpha$. As such, I_k is a k -th order polynomial in α . Then, the global minimum of (3.31) is such that $\partial \mathcal{L}_N / \partial c_k = 0$ for all k , and $\partial \mathcal{L}_N / \partial \lambda = 0$. This corresponds to a linear set of equations in c_k and λ , and can therefore be rewritten as a problem of the form $Ax = y$ where A is a matrix of the I_k integrals, $x = (c_0, \dots, c_N, \lambda)$ and $y = (0, 0, \dots, 0, -1)$. Such a system is easily solved numerically through matrix inversion, thus yielding the solution of the initial problem.

3.4.4 Discrete jump

For this final design, we shift away from the Zeno effect that has for now been the common basis for all physical cat gate implementations with a Z component. This section introduces discrete gates that rely on a dissipative coupling to an ancillary mode. We begin with single-qubit $Z(\theta)$ gates, and generalize to CZ and CNOT gates afterwards.

3.4.4.1 Design principle

Consider the following master equation with the buffer mode adiabatically eliminated,

$$\frac{d\rho}{dt} = \kappa_2 \mathcal{D}[\mathbf{a}^2 - \alpha^2] \rho + \kappa_Z \mathcal{D}[\mathbf{a}_\theta \sigma_+] \rho \quad (3.35)$$

where σ_+ denotes the creation operator of an ancillary qubit mode, and

$$\mathbf{a}_\theta = \cos(\theta/2) \mathbf{a} + i \sin(\theta/2) \mathbf{a} \quad (3.36)$$

is a modified annihilation operator such that $\mathbf{a}_\pi = i\mathbf{a}$, and $\mathbf{a}_\theta |\pm\alpha\rangle = \exp(\pm i\theta/2) |\pm\alpha\rangle$. From this identity, it is immediate to verify that projecting any cat qubit state with \mathbf{a}_θ achieves a $Z(\theta)$ rotation of this state, up to exponentially small corrections in $|\alpha|^2$ to account for state normalization. Furthermore, assuming that this ancillary mode is initialized in its ground state $|g\rangle$ and does not suffer from errors, the correlated dissipator in $\mathbf{a}_\theta \sigma_+$ engineers the loss of exactly one \mathbf{a}_θ photon from the cat mode, after which nothing more can happen with this dissipation. This can thus be seen as a kind of photon blockade technique, which can further be made error-resilient by, for instance, sending the additional photon to the meta-stable level of a three-level Λ -system or into an infinite transmission line.

For the particular value $\theta = \pi$, (3.35) induces the loss of exactly one \mathbf{a} photon. Since the logical $|\pm_L\rangle$ cat states feature even and odd photon-number

parities respectively and the logical bit value is exponentially protected under such losses thanks to the two-photon dissipation, this peculiar dissipator results in a Pauli $Z(\pi)$ gate on the cat qubit mode after infinite time evolution.

To estimate gate fidelities at finite times analytically, let us assume that the complete system is initialized in $\rho(0) = \rho_{A,0} \otimes |g\rangle\langle g|$ with $\rho_{A,0}$ the initial density matrix on the cat mode. The full system density matrix at time t can then be separated along the diagonal matrix elements of the ancillary mode according to $\rho = \rho_g \otimes |g\rangle\langle g| + \rho_e \otimes |e\rangle\langle e|$ (the correlated dissipator of (3.35) does not induce diagonal to off-diagonal transitions). Inserting this expression in (3.35) thus yields

$$\frac{d\rho_g}{dt} = \kappa_2 \mathcal{D}[\mathbf{a}^2 - \alpha^2] \rho_g - \frac{1}{2} \kappa_Z \left\{ \mathbf{a}_\theta^\dagger \mathbf{a}_\theta, \rho_g \right\} \quad (3.37a)$$

$$\frac{d\rho_e}{dt} = \kappa_2 \mathcal{D}[\mathbf{a}^2 - \alpha^2] \rho_e + \kappa_Z \mathbf{a}_\theta \rho_g \mathbf{a}_\theta^\dagger \quad (3.37b)$$

where $\rho_g(0) = \rho_{A,0}$ and $\rho_e(0) = 0$. Let us first discuss these coupled equations without the $\{\mathbf{a}_\theta^\dagger \mathbf{a}_\theta, \cdot\}$ term of (3.37a), which will be shown to be the main limitation of the design. As can be seen from the right-hand side term in (3.37b), a transfer of population is made from $|g\rangle$ to $|e\rangle$ at rate $\kappa_Z |\alpha|^2$. Upon this population transfer, a single \mathbf{a}_θ projection of the cat mode is performed thanks to the $\mathbf{a}_\theta \rho_g \mathbf{a}_\theta^\dagger$ term, thus achieving the required gate. The minimal gate-induced phase errors that can be achieved at finite time thus come down to the fraction of states which have not undergone the jump from $|g\rangle$ to $|e\rangle$, and read

$$p_Z = \exp(-|\alpha|^2 \kappa_Z t) \quad (3.38)$$

where t is the time of evolution under the correlated dissipator. The phase error would thus decrease exponentially with $|\alpha|^2$ and with the effective gate time $\kappa_Z t$.

Let us now consider the effect of the right-hand side term of (3.37a).

For the particular case $\theta = \pi$, since $\mathbf{a}_\pi = i\mathbf{a}$ and thus $\mathbf{a}_\pi^\dagger \mathbf{a}_\pi = \mathbf{a}^\dagger \mathbf{a}$, this term remains a parity-preserving operator. Therefore, the evolution with (3.37a) and (3.37b) perfectly preserves the phase of the logical qubit, up to inducing the desired phase-flip gate from ρ_g to ρ_e . Besides this, the last term of (3.37a) acts much like a dispersive dissipation $\mathcal{D}[\mathbf{a}^\dagger \mathbf{a}]$, inducing noisy rotation of the cat-qubit out of its codespace. As long as $\alpha \kappa_Z \ll 4|\alpha|^2 \kappa_2$, this effect is countered by the two-photon dissipation. Since the mean energies of even and odd cats are exponentially close, the induced logical bit-flip is exponentially small in $|\alpha|^2$. Up to an upper bound on κ_Z/α , the exponentially scaling phase gate thus indeed holds.

For gate angles $\theta \neq \pi$, to consider the effect of the last term in (3.37a) we write out

$$\mathbf{a}_\theta^\dagger \mathbf{a}_\theta = \cos^2(\theta/2) |\alpha|^2 + \sin^2(\theta/2) \mathbf{a}^\dagger \mathbf{a} - i \sin(\theta) (\mathbf{a}^\dagger - \mathbf{a})/2. \quad (3.39)$$

The second line of this operator would induce parity-switching leakage from the cat codespace, and therefore parity errors at a rate proportional to $\sin^2(\theta)$ once brought back to the codespace by two-photon dissipation. The gate fidelities obtained from this scheme are still competitive with other gate designs, and particularly in the limit of $\kappa_Z \ll \kappa_2$; see simulation results below. This can be understood as the last term in (3.37a) has an effect roughly similar to a Hamiltonian in $i \sin(\theta)(\mathbf{a}^\dagger - \mathbf{a})$. Hence, much like the analysis of Section 3.2, it induces phase errors at a rate proportional to κ_Z^2/κ_2 , while the phase gate happens at rate κ_Z .

3.4.4.2 Qutrit design

But in fact, we can do better and retrieve for any θ the same performance as the $\theta = \pi$ case with a slightly more involved gate design. Consider the following master equation,

$$\begin{aligned} \frac{d\boldsymbol{\rho}}{dt} = & \kappa_2 \mathcal{D}[\mathbf{a}^2 - \alpha^2] \boldsymbol{\rho} + \kappa_Z \mathcal{D}[\mathbf{a}_\theta |e\rangle\langle g|] \\ & + \kappa_Z \mathcal{D}[\mathbf{a}_{\theta+\pi} |f\rangle\langle g|] + \kappa'_Z \mathcal{D}[\mathbf{a}_\pi |e\rangle\langle f|] \end{aligned} \quad (3.40)$$

where $|g\rangle$, $|e\rangle$ and $|f\rangle$ denote the three lowest energy levels of a qutrit, for instance that of a transmon. This master equation now involves three population transfers. The first is the same as in (3.35) and performs a $Z(\theta)$ gate with a $|g\rangle$ to $|e\rangle$ transfer. The other two terms also describe a $Z(\theta)$ gate, but made in two steps by first transferring to $|f\rangle$ and then to $|e\rangle$. The main goal of adding these two terms is to cancel out parity-switching dynamics in the $|g\rangle\langle g|$ subspace. Indeed, we now have that

$$\mathbf{a}_\theta^\dagger \mathbf{a}_\theta + \mathbf{a}_{\theta+\pi}^\dagger \mathbf{a}_{\theta+\pi} \propto \cos^2(\theta/2) |\alpha|^2 + \sin^2(\theta/2) \mathbf{a}^\dagger \mathbf{a}. \quad (3.41)$$

In the $|f\rangle\langle f|$ subspace, the dynamics is also parity-preserving since a $Z(\pi)$ gate is performed. In other words, two paths have been constructed which both feed into the same final state with the required Z rotation angle of θ , and whose interference cancels the parity-switching term in (3.39). This scheme would mainly be limited by qutrit characteristics, to be specified from experimental implementation, and by the requirement that both paths should feature the same rate of dissipation. In the following, we focus on the qubit-enabled $Z(\pi)$ gate, and also compare qubit- and qutrit-enabled $Z(\theta)$ gate designs.

3.4.4.3 Results

First, we investigate $Z(\pi)$ gates. Figure 3.11 shows the phase errors induced by this discrete $Z(\pi)$ gate without any additional error processes on the cat

or ancillary modes. On plot (a), phase errors are shown in semi-log scale as a function of time for varying κ_Z/κ_2 , and fixed $|\alpha|^2 = 4$ or 8 in the master equation (3.35). For each value of κ_Z/κ_2 , both numerical simulations (solid) and the minimal gate error formula of (3.38) (dashed) are shown. We retrieve that, in the limit of $\kappa_Z \ll \kappa_2$, the expected formula fits numerical simulations perfectly and exponentially small gate errors are achieved. As κ_Z is increased, we find a deviation from the optimal gate errors due to the competition with two-photon dissipation, as the cat significantly leaves the codespace of intended size $|\alpha|^2$ during the gate, getting drawn closer to the vacuum. An exponential scaling with time is maintained, although with a smaller exponential rate, as discussed in the previous subsection.

In Figure 3.11(c), we extract this exponential scaling rate from a linear fit of $\gamma(t) = -\ln(p_Z(t))$, which yields γ_Z such that $p_Z(t) \sim \exp(-\gamma_Z t)$ for sufficiently large time values. This exponential rate is then plotted as a function of κ_Z/κ_2 for varying values of $|\alpha|^2$, and also compared to the ideal rate of (3.38) given by $\gamma_Z = |\alpha|^2 \kappa_Z$. Again, a transition from the optimal gate error regime to a sub-optimal regime is found as κ_Z/κ_2 is increased. In addition, the point of transition scales as α , which would confirm the limit of validity of the optimal regime given by $\alpha \kappa_Z \ll 4|\alpha|^2 \kappa_2$, or equivalently, $\kappa_Z/\kappa_2 \ll 4\alpha$.

The effect of κ_Z/κ_2 on *bit-flip* errors is investigated on Figure 3.12(b). In this particular simulation, a relatively large gate time is fixed at $T = 10/\kappa_2$ to ensure the relevance of the study. We find that bit errors are indeed suppressed exponentially according to $p_X \propto \exp(-2|\alpha|^2)$ for all values of κ_Z investigated, thus preserving the error bias of cat qubits. However, as the effective cat size is reduced by the κ_Z term of (3.37a) during the gate time, the prefactor of this exponential increases linearly with κ_Z/κ_2 .

The main side-process limiting the fidelity of discrete gates would be ancillary mode lifetime. Since the gate is based on a transition from the ground to the excited state of the ancilla that will serve as a photon blocker, any unwanted transition between these two ancillary states will perturb the gate process. Figure 3.12 investigates finite qubit lifetimes for the two-level ancillary mode. Here, a discrete $Z(\pi)$ gate is simulated numerically in the presence of single-photon losses on the ancillary qubit of the form $\kappa_q \mathcal{D}[\sigma_-]$, with varying values of κ_q/κ_2 . An analytical fit is further shown for each numerical simulation, which reads $p_Z(t) = \kappa_q t + \exp(-|\alpha|^2 \kappa_Z t)$ and matches each line very well. Indeed, ancillary mode losses induce unwanted $|e\rangle \rightarrow |g\rangle$ transitions at a constant rate which are followed exponentially quickly by a $|g\rangle \rightarrow |e\rangle$ transition of the discrete gate correlated dissipator, inducing a second phase-flip and hence cancelling the $Z(\pi)$ gate. Thermal excitations of the ancillary mode would have a similar effect by activating the $|g\rangle \rightarrow |e\rangle$ transition without a cat mode parity switch, but in a cold environment such excitations have a much smaller rate than qubit decay.

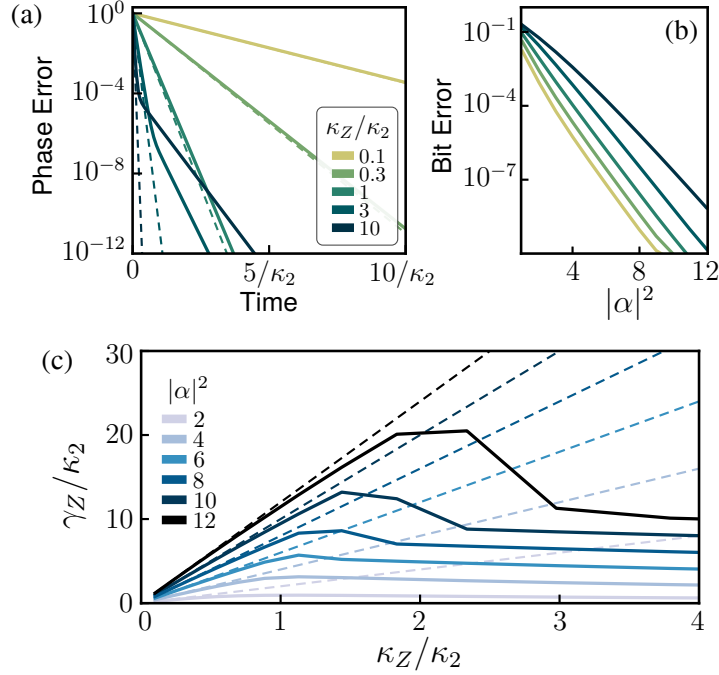


Figure 3.11: (a,b) Gate-induced phase and bit errors of a discrete $Z(\pi)$ gate as in (3.35) for multiple values of κ_Z/κ_2 . Left: fixed cat size, $|\alpha|^2 = 8$. Right: fixed gate time, $T = 10/\kappa_2$. In the left plot, dashed lines show $p_Z = \exp(-|\alpha|^2 \kappa_Z t)$. (c) Exponential rate γ_Z obtained from a linear fit of $\gamma(t) = -\ln(p_Z(t)) \propto \gamma_Z t$. Dashed lines show $\gamma_Z = |\alpha|^2 \kappa_Z$.

While this linear increase of phase errors may appear limiting for the usefulness of the gate design, we remind that finite lifetime of the cat oscillator also induce phase errors that scale linearly in time, according to $p_Z = |\alpha|^2 \kappa_1 t$ where κ_1 is the rate of single-photon losses. Those losses are inevitable without changing the overall cat qubit encoding. In contrast, for the ancilla, while a transmon-like qubit would yield a simple gate implementation, the particular gate design is compatible with more specific quantum systems. There is no requirement for the protection of the $|g\rangle + |e\rangle \leftrightarrow |g\rangle - |e\rangle$ transition, since only the diagonal elements are used by the gate design. As such, any ancilla system that can robustly implement a single transition from $|g\rangle$ to $|e\rangle$ would suffice for better gate protection. One can think of a three-level Λ system with two meta-stable ground states [Kumar *et al.* 2016, Vepsäläinen *et al.* 2019], or a system where the ancilla state would escape away (but never back to $|g\rangle$) after reaching $|e\rangle$.

Figure 3.13 investigates other gate angles, $\theta \neq \pi$, both with the ancillary two-level system design of (3.35) (left), and with the three-level system design of (3.40) that cancels out parity-switching in the $|g\rangle \times |g\rangle$ subspace through interference (right). In the first scheme, the phase errors first scale exponen-




Figure 3.12: Gate-induced phase errors of a discrete $Z(\pi)$ gate as in (3.35) with additional single-photon losses of the form $\kappa_q \mathcal{D}[\sigma_-]$ on the ancillary qubit. Dashed lines show $p_Z = \kappa_q t + \exp(-|\alpha|^2 \kappa_Z t)$. The cat size is fixed at $|\alpha|^2 = 8$ and $\kappa_Z/\kappa_2 = 1$.

tially until they hit a plateau corresponding to the parity-switching leakage of (3.39). This plateau scales according to $\sin^2(\theta)$. With the qutrit scheme, the exponential scaling of phase errors is also observed but for longer time scales, thus confirming that parity-switching dynamics is canceled out. A plateau is still hit after some time t but at lower phase errors than for the previous two-level system scheme. This new limitation is explained by the fact that, if the state is initially not coherent or with a different coherent state amplitude than α , \mathbf{a}_θ does not map exactly onto a $Z(\theta)$ gate. Since the $|g\rangle\langle g|$ subspace dynamics induces such incoherent fluctuations, gate fidelities are limited by this effect. For sufficiently small values of κ_Z/κ_2 , the state however stays inside the codespace at all times and the exponential scaling of gate errors is maintained for longer time scales.

3.4.4.4 Multi-qubit gates

Regarding multi-qubit gates, the single-qubit discrete method of (3.35) can be generalized but not necessarily in a trivial manner. For CZ gates, consider the two-qubit operator,

$$\mathbf{L}_{CZ} = -\mathbf{a}_1(\mathbf{a}_2 - \alpha) + \alpha(\mathbf{a}_2 + \alpha) \quad (3.42)$$

where $\mathbf{a}_{1/2}$ are the annihilation operators on each mode involved. This operator is such that $\mathbf{L}_{CZ} |\pm\alpha\rangle_1 |\alpha\rangle_2 = 4\alpha |\pm\alpha\rangle_1 |\alpha\rangle_2$ and $\mathbf{L}_{CZ} |\pm\alpha\rangle_1 |-\alpha\rangle_2 = \pm 4\alpha |\pm\alpha\rangle_1 |\alpha\rangle_2$ which thus meets the CZ gate requirements. Note in particular the inter-exchangeability of \mathbf{a}_1 and \mathbf{a}_2 in this operator, which shows the CZ gate symmetry. By engineering a dissipation of the form $\mathcal{D}[\mathbf{L}_{CZ}\sigma_+]$, an exponentially scaling CZ(π) gate can be achieved. It would however suffer




Figure 3.13: Gate-induced phase errors of a discrete $Z(\pi/3)$, with the designs of (3.35) using an ancillary two-level system (left) and of (3.40) using an ancillary three-level system (right). Dashed lines show $p_Z = \exp(-|\alpha|^2 \kappa_Z t)$. Insets show the gate transitions between ancillary states for either design.

from the same issues as qubit-enabled $Z(\theta)$ gates due to the $|g\rangle\langle g|$ parity-switching dynamics, since $\mathbf{L}_{CZ}^\dagger \mathbf{L}_{CZ}$ contains parity-switching terms both in \mathbf{a}_1 or \mathbf{a}_2 . Similarly as for single-qubit $Z(\theta)$ rotations, it is possible to cancel out these terms through interference by instead using four operators similar to \mathbf{L}_{CZ} and a higher-dimensional ancillary system.

For CNOT gates, consider the two-qubit operator,

$$\mathbf{L}_{CX} = \alpha \mathbf{I} \otimes \mathbf{P}_+ + \mathbf{a} \otimes \mathbf{P}_- \quad (3.43)$$

where the first and second modes in this outer product are control and target qubits respectively, where \mathbf{I} is the oscillator identity, and $\mathbf{P}_\pm = (e^{i\pi \mathbf{a}^\dagger \mathbf{a}} \pm 1)/2$ are projectors on the even and odd parity subspaces of the target mode respectively. Similarly as for the single-mode method, engineering a correlated dissipator with an ancillary qubit mode of the form $\mathcal{D}[\mathbf{L}_{CX} \sigma_+]$ would achieve a discrete CNOT gate, up to exponentially small corrections. However, the projection operators \mathbf{P}_\pm are highly non-local, and therefore this scheme is of questionable use. First, such operators are not currently implementable with superconducting circuits, although there is a path towards them as high-impedance operators [Cohen 2017]. Second, if one assumes access to such “next-generation” operators, then for fairness one should view as relatively easy too the options of e.g. Section 3.4.3. Third, engineering non-local operators in a cat-qubit context is a Pandora’s box, as it introduces artificial operators against which the exponential bit-flip protection is *not* designed to work.

Other designs may take inspiration from the Zeno implementation of the CNOT gate, i.e. inducing, in superposition, $n - n_p$ discrete $Z(\theta)$ gates on the control qubit conditional on Fock state n in the target qubit. In the Zeno Hamiltonian implementation, this indeed remains compatible with ex-

ponential bit-flip protection. Generalizing this to a jump operator appears to require a complicated ancilla system. For instance, if we would rely on $n - n_p$ ancilla jumps to perform the right number of $Z(\theta)$ gates, then Fock-number information would leak out to the environment unless specific erasure actions are taken to impeach detecting the number of jumps. In absence of concrete insight on realistic experimental building blocks associated to such ancilla, we leave this for future research.

Barring the experimental difficulty of realizing an operator like \mathbf{L}_{CX} , the effect of unwanted $|g\rangle \leftrightarrow |e\rangle$ transitions of the ancillary mode that monitors the gate warrants a caveat. Indeed, an unwanted transition would mean applying the CNOT gate possibly an even number of times, which would by definition involve both phase-flip and bit-flip errors. In the traditional model, bit-flip errors remain exponentially suppressed even for CNOT gates, thanks to a smart use of the available continuous phase space which we lose with this design. On the upside, the particular use of the ancillary qubit in this discrete gate opens the door to specific designs reducing ancilla-induced errors, as discussed above for the $Z(\theta)$ gate.

3.5 Robustness to noise

In the previous section, we studied gate designs with simplified master equation models in order to capture the key elements of each proposal. However, actual experimental setups feature various sources of errors that could in principle hinder gate performances. In this section, we study a more complete model of errors to demonstrate that our proposals hold even under realistic noise processes. We begin with a set of standard error sources, and then move on to the specific case of thermal noise in the buffer mode for the first two gate designs.

3.5.1 On the impact of noise and spurious Hamiltonians

Let us consider the master equation

$$\begin{aligned} \frac{d\rho}{dt} = & \mathcal{L}_0\rho - i[\mathbf{H}_s, \rho] + \kappa_a(1 + n_{th,a})\mathcal{D}[\mathbf{a}]\rho \\ & + \kappa_a n_{th,a}\mathcal{D}[\mathbf{a}^\dagger]\rho + \kappa_{\phi,a}\mathcal{D}[\mathbf{a}^\dagger\mathbf{a}]\rho \end{aligned} \quad (3.44)$$

where \mathcal{L}_0 denotes the Liouvillian to be engineered on the system — varied from one gate design to the next — which typically contains two-to-one photon exchange, buffer mode dissipation and a gate drive on the memory mode, and

$$\mathbf{H}_s = -K_a \mathbf{a}^{\dagger 2} \mathbf{a}^2 + \chi_{ab} \mathbf{a}^\dagger \mathbf{a} \mathbf{b}^\dagger \mathbf{b} - K_b \mathbf{b}^{\dagger 2} \mathbf{b}^2 \quad (3.45)$$




Figure 3.14: Phase errors during a $Z(\pi)$ gate under a complete model of noise for each gate design introduced in the main text. Each design is compared to the standard Zeno gate (black). In all panels, $|\alpha|^2 = 4$ and energies are set to $g_2/2\pi = 1$ MHz, $K_a/2\pi = 1$ kHz, $K_b/2\pi = 810$ kHz, $\chi_{ab}/2\pi = 65$ kHz, $\kappa_a/2\pi = 53$ Hz, $n_{th,a} = 10\%$ and $\kappa_{\phi,a}/2\pi = 10$ Hz. For panels (a), (c) and (d), $\kappa_b/2\pi = 8$ MHz. For panel (b), $\kappa_{ab}/2\pi = 2$ MHz.

is a spurious Hamiltonian with Kerr and cross-Kerr terms on the memory and buffer modes. These terms typically result from the residual cosine potential in the full circuit Hamiltonian, but in practice, they can be small compared to two-photon dissipation when the system is engineered with an Asymmetrically Threaded SQUID at the appropriate flux bias point [Lescanne *et al.* 2020b]. Our master equation model also includes single-photon losses, thermal photons and pure dephasing on the memory mode with respective rates κ_a , $n_{th,a}$ and $\kappa_{\phi,a}$. Note that we do not include thermal photons on the buffer mode in this model as they are discussed separately in Section 3.5.2.

To study these noise processes, we consider specific values for each term mainly extracted from Refs. [Lescanne *et al.* 2020b, Réglade *et al.* 2023]. The two-to-one photon exchange coupling is $g_2/2\pi = 1$ MHz and combined with a buffer mode dissipation at $\kappa_b/2\pi = 8$ MHz which yields an effective two-photon dissipation rate $\kappa_2/2\pi = 500$ kHz. Kerr and cross-Kerr energies are $K_a/2\pi = 1$ kHz, $K_b/2\pi = 810$ kHz and $\chi_{ab}/2\pi = 65$ kHz. Single-photon loss is $\kappa_a/2\pi = 53$ Hz corresponding to $T_1 = 3$ ms in the memory mode, which is typical in current experiments with high-Q superconducting resonators. Note that lower cavity lifetimes could also be investigated, but would limit the potential gain in gate fidelity of our proposals compared to standard Zeno gates as phase errors would be dominated by single-photon losses in this regime. Finally, thermal noise and pure dephasing are $n_{th,a} = 10\%$ and $\kappa_{\phi,a}/2\pi = 10$ Hz.

In Figure 3.14, gate errors under this noise model are shown for each gate design from the main text. Each time, the performance is compared to the standard Zeno gate for a $Z(\pi)$ gate, shown in black lines. In all panels, phase errors eventually converge to a linear dependence in the large gate time regime in which single-photon losses dominate. In addition, there is systematically

an optimal gate time which minimizes gate errors and that represents the optimal trade-off between gate-induced and single-photon loss errors.

For the photodetection design in panel (a), we vary the photodetection efficiency from a perfect photodetector $\eta = 1$ to a non-ideal one at $\eta = 0.5$, demonstrating again the linear dependence with detection inefficiency as discussed in the main text. Panel (b) shows the performance of the autonomous feedback design with $\kappa_{ab}/2\pi = 2$ MHz and $\kappa_b = 0$. In panel (c), we show the locally flat Hamiltonian design for increasing number of odd drive terms from $N = 1$ to $N = 4$ each time gaining in optimal gate fidelity and gate time. Finally, panel (d) shows performances of the discrete jump design where the optimal gate fidelity depends on the ratio of κ_Z to κ_2 and on the actual error model studied.

3.5.2 Thermal noise in the buffer mode

In this second subsection, we discuss the role of thermal noise in the buffer mode for the first two designs based on the feedback of information introduced in the main text. Since these two schemes use the buffer mode population to detect potential phase-flips induced on the memory mode, it is quite natural that one of their main limitations should come from spurious population in the buffer.

3.5.2.1 Photodetection

For the photodetection scheme of Section 3.4.1, the feedback action to apply is a $Z(\pi)$ gate after every buffer photon detection. Therefore, the rate of thermal photons is directly linked to the rate of phase information loss on the cat state, assuming that the feedback action is perfect. From the point of view of the SFB, one can adiabatically eliminate the gauge mode in (3.12) and average out the stochastic terms. This leads to a simplified model in the absence of a gate drive,

$$\frac{d\rho}{dt} = \kappa_b\eta(1 + n_{th,b})\mathcal{D}[\mathbf{b}\sigma_z]\rho + \kappa_b(1 - \eta)(1 + n_{th,b})\mathcal{D}[\mathbf{b}]\rho + \kappa_b n_{th,b}\mathcal{D}[\mathbf{b}^\dagger]\rho \quad (3.46)$$

Here, the beamsplitter interaction between gauge and buffer modes is eliminated compared to the fast dynamics on the buffer. In particular, notice the correlated dissipation in $\mathcal{D}[\mathbf{b}\sigma_z]$ which arises from the idealized classical feedback action on the memory mode. From this master equation, we easily compute the average expectation value of the σ_x operator by going to the Heisenberg picture, which reads

$$\langle\sigma_x\rangle(t) = \langle\sigma_x\rangle(t=0) \times \exp(-\gamma t) \quad (3.47)$$




Figure 3.15: (Left) Photon-number parity of an idling cat state under the photodetection and classical feedback design of Section 3.4.1 with thermal noise in the buffer mode $n_{th,b} = 2\%$. Parity is averaged over 1000 stochastic trajectories, and fitted with an exponential decay at rate $\gamma = 2\eta\kappa_b n_{th,b}(1 + n_{th,b})$. In this simulation, $\kappa_b/g_2 = 8$, $|\alpha|^2 = 8$ and $\eta = 1$. (Right) Photon-number parity of an idling cat state under the autonomous feedback of Section 3.4.2 with thermal noise in the reservoir mode $n_{th,r} = 2\%$. Parity is fitted with an exponential decay at rate $\gamma = 2\sqrt{g_2\kappa_{ab}}n_{th,r}$. In this simulation, $\kappa_{ab}|\alpha|^2/g_2 = 8$ and $|\alpha|^2 = 8$.

where $\gamma = 2\eta\kappa_b n_{th,b}(1 + n_{th,b})$ rate of parity information loss in the presence of thermal noise in the buffer. Note that the classical feedback action is always optional, and typically it should only be turned on during gates. To know whether the gate design is beneficial in practice, this rate should be compared to the rate of single-photon losses κ_a which is the usual dominating source of loss of parity information.

On the left panel of Figure 3.15, we show the parity of an idling cat state initialized in $|\pm_L\rangle$ as obtained from numerical integration of the full stochastic master equation (i.e. including two-photon coupling, buffer mode dissipation and thermal noise in the buffer). The solution is averaged over 1000 stochastic trajectories, and we find a very good fit with the simplified model as described above.

3.5.2.2 Autonomous feedback

For the feedback design of Section 3.4.2, the feedback action is this time performed autonomously through a correlated dissipation operator in $\mathcal{D}[ab]$. Buffer excitations would then also result in spurious parity swaps on the cat mode. However, with this design, the buffer mode is *a priori* a high-Q mode since the correlated dissipation is engineered through a third reservoir mode \mathbf{r} as discussed in the main text, with a three mode coupling of the form $g(ab\mathbf{r}^\dagger + \text{h.c.})$. Therefore, the main limitation of this design does not come from buffer

mode excitations, but rather from thermal noise in the reservoir mode in the form $\kappa_r n_{th,r} \mathcal{D}[\mathbf{r}^\dagger]$. From adiabatic elimination of this reservoir with the formalism of [Azouit *et al.* 2017, Forni *et al.* 2018], we derive an effective rate of correlated excitations. The effective master equation on the memory and buffer modes for an idling cat qubit thus reads

$$\frac{d\rho}{dt} = -i[\mathbf{H}_{AB}, \rho] + \kappa_{ab}(1 + n_{th,r})\mathcal{D}[\mathbf{ab}]\rho + \kappa_{ab}n_{th,r}\mathcal{D}[\mathbf{a}^\dagger\mathbf{b}^\dagger]\rho \quad (3.48)$$

where a new term in $\mathcal{D}[\mathbf{a}^\dagger\mathbf{b}^\dagger]$ appears compared to the ideal master equation. This dissipation creates correlated excitations on the memory and buffer modes. If these correlated excitations are dissipated through the $\mathcal{D}[\mathbf{ab}]$ term, then the parity is swapped back to its original value and no phase information is lost. However, the correlated excitation can also undergo partial two-photon Rabi oscillations through the \mathbf{H}_{AB} term and then eventually dissipate away. In this case, the parity is not strictly preserved by the dynamics and phase information is lost to the environment.

On the right panel of Figure 3.15, we show the parity of an idling cat state initialized in $|\pm_L\rangle$ as obtained from the numerical integration of (3.48). We indeed find that the parity information decays exponentially with a rate γ . For all values of g_2 , κ_{ab} and $n_{th,r}$ we investigated, this rate seems to fit the formula $\gamma = 2\sqrt{g_2\kappa_{ab}n_{th,r}}$. In particular, turning off either two-to-one photon coupling or thermal noise on the reservoir mode suppresses this effect.

3.6 Conclusion

We have introduced four new designs of dissipative cat qubit gates that can help reduce gate-induced phase errors, and therefore help reach error-correction thresholds. Upon the observation that incoherent gate errors result from the entanglement between the cat qubit and its buffer mode, we have devised two designs meant to take advantage of the buffer memory in the system. The first one is based on the photodetection of the buffer mode output, and thus on the retrieval of information that can then be classically fed back. The precision of this scheme is only limited by the photodetector efficiency. The second one relies on an autonomous error correcting scheme for which first-order gate-induced parity errors are automatically corrected for as buffer photons exit the resonator. In particular, this second design is readily implementable with superconducting circuits and achieves a reduction of up to two orders of magnitude in gate errors. The same setup can also be used with squeezed cat states for the autonomous correction of single-photon losses [Xu *et al.* 2022b], currently one of the main limiting factors of bosonic qubits.

We have also described two drastically different gate designs. By engineering higher-order drive Hamiltonians that feature locally flat energy potentials in position representation, spurious effects on the cat qubit mode can be avoided and an improvement in gate error fidelities is achieved. We believe that such Hamiltonians could also be used for other purposes, such as cat state preparation. Finally, we have explored how to engineer cat qubit gates without the Zeno effect, and introduced dissipation-based gate engineering. By coupling the cat mode to an ancillary nonlinear mode that monitors the ongoing gate, a discrete π -phase gate is realized. This method can circumvent the usual linear gate time scaling of Zeno-based gates and achieve exponentially scaling gate fidelities, and associated CNOT gate designs are therefore our subject of ongoing investigation.

We hope that this chapter has been able to provide its readers intuition about the design of cat qubit gates. While we have explored many paths towards high-fidelity gates, we believe that further improvements can still be achieved, in particular for multi-mode encodings of cat qubits. Finally, while these gate designs have been particularly focused on dissipative cat qubits, they should also inspire the design of superadiabatic operations on any bosonic or dissipatively-stabilized system. In particular, most tailored dissipation operators often rely on a highly damped buffer mode together with adiabatic elimination. We have thus described for the first time how the dynamics of this buffer mode can be taken advantage of.

Optimal control of transmon readout

This chapter covers work in preparation for publication that was mainly pursued while visiting the group of Alexandre Blais in Sherbrooke, Canada [Gautier *et al.* 2023a].

Contents

4.1	Introduction	125
4.2	Adjoint state quantum optimal control	127
4.2.1	Automatic differentiation	128
4.2.2	Adjoint state method	130
4.2.3	Derivation of adjoint state method equations	133
4.3	Optimal control of transmon readout	135
4.3.1	Transmon-resonator-filter model	135
4.3.2	Dispersive readout metrics	137
4.3.3	Optimizing readout	138
4.3.4	Results	139
4.4	Conclusion	143

4.1 Introduction

In Chapter 3, we introduced new designs for the high-fidelity control of a dissipative cat qubit that leveraged Hamiltonians or dissipative mechanisms tailored to the quantum system. This chapter will also be dedicated to high-fidelity control of a quantum system, but this time through pulse engineering and in a much more general setup. Superconducting circuits are controlled with microwave pulses that are sent to the chip through transmission lines. These pulses are produced by mixing a control field arising from an Arbitrary

Waveform Generator (AWG) and a carrier field of fixed frequency, both being readily controllable from the laboratory. A question that then naturally arises is: what are the optimal pulses that can achieve fast and high-fidelity quantum operations?

Quantum optimal control [Peirce *et al.* 1988, Werschnik & Gross 2007] seeks to answer this question, through analytical formulations when the problem allows it [Marte *et al.* 1991, Motzoi *et al.* 2009], and otherwise through numerical investigation [Krotov 1995]. For the latter, a recent but popular method is Gradient Ascent Pulse Engineering (GRAPE) [Khaneja *et al.* 2005] which consists in iterating over a set of discretized pulses through a gradient-based classical optimization algorithm. The method thus relies on the ability to compute gradients of a given cost function — e.g. the fidelity of a quantum gate, or the overlap to some objective state — with respect to the discrete pulse. For closed quantum systems, this can be achieved by back-propagating the Schrödinger equation with Trotter expansions of short-time propagators. For open quantum systems, a similar back-propagation can be performed on the vectorized Lindblad master equation [Schulte-Herbrüggen *et al.* 2011], or under certain conditions, directly on the time-ordered Magnus expansion [Boutin *et al.* 2017]. However, these methods suffer from several pitfalls: they can only be implemented with a small subset of cost functions, their gradients are estimated only to low-order in the Trotter expansion which limits the convergence of the gradient descent algorithm, and they are not always suited to large open quantum systems due to memory or speed requirements.

To overcome some of these issues, one solution is to delegate gradient computations to a problem-agnostic routine such as that of automatic differentiation [Jirari 2009, Leung *et al.* 2017, Abdelhafez *et al.* 2019]. Through successive applications of the chain rule, automatic differentiation can differentiate through any numerical problem that can be expressed with standard linear algebra operations, and has thus seen tremendous use in machine learning [Baydin *et al.* 2018]. While very promising for quantum optimal control, it still suffers from one major bottleneck: memory usage. Indeed, back-propagating through the complete time integration of a Lindblad master equation requires storing a density matrix at every numerical time step, and thus a memory cost in $\mathcal{O}(nN^2)$ where n is the number of time steps, and N the Hilbert space size. For large open quantum systems ($N \gtrsim 100$) or problems that require many integration steps ($n \gtrsim 1000$), this memory cost can quickly limit the utility of these methods.

In this chapter, we introduce a new method to differentiate through a Lindblad master equation with a constant memory cost in $\mathcal{O}(N^2)$, which is based on the reverse time integration of an adjoint state master equation [Pontryagin *et al.* 1962]. This method is also numerically fast, with a time cost only about twice that of automatic differentiation-based methods,

and it is completely agnostic to the model or to the cost functions considered. This makes this method extremely promising for many practical problems in quantum optimal control, but also for parameter fitting, state tomography, or any application requiring gradients.

We then demonstrate how this method applies to the optimization of transmon readout [Koch *et al.* 2007, Mallet *et al.* 2009], currently one the bottlenecks of transmon-based quantum computers [Swiadek *et al.* 2023]. By considering a complete experimental model including the transmon, its readout resonator and a driven Purcell filter, we show that pulse shaping improves readout fidelity and readout time by up to 25% compared to standard readout pulses on the experimental setup we selected. By including shelving pulses on the transmon to probe higher-level state transitions, we show even greater improvements with readout times up to twice as fast as standard methods. In addition, the optimal control algorithm discovers interpretable pulses, e.g. pulses that fuse several operations efficiently, or take into account the Stark shift. Because transmon readout is a complex problem involving multiple modes, a large Hilbert space, and a strong dissipative process that is key to the operation to optimize, it is a prime example of a problem that was previously inaccessible with standard optimal control tools but that we can thoroughly study with our novel adjoint state method.

The chapter is organized in two main sections. After a short review of automatic differentiation-based optimal control, Section 4.2 covers the adjoint state method for open quantum systems. In particular, we derive the main equations used for the numerical computation of gradients in constant memory cost. Section 4.3 then discusses the optimization of transmon readout, first with an overview of the physical model, of the optimal control process, and then showcasing the numerical results. We conclude in Section 4.4.

4.2 Adjoint state quantum optimal control

Consider a generic quantum optimal control problem for which we want to find a set of parameters that minimize some cost function C . The cost function is in general function of the problem parameters $\theta = (\theta_1, \dots, \theta_m)$ and of the density matrix of the system at a discrete set of times $\rho(t_i)$, such that $C \equiv C(\theta, \rho(t_0), \dots, \rho(t_n))$. For gradient-based optimization, it is required to compute the derivative of the cost function with respect to each parameter, $dC/d\theta$. A standard optimization algorithm such as Adam [Kingma & Ba 2014] or Stochastic Gradient Descent [Robbins & Monro 1951] can then be employed to find a global minimum of the parameter landscape, or at least a good local minimum. In the following, we quickly review how automatic differentiation can be used to compute these gradients, before moving on to adjoint state

quantum optimal control.

4.2.1 Automatic differentiation

Automatic differentiation is a family of techniques designed to compute partial derivatives of a numeric function defined in a computer program. It is different from numerical differentiation (evaluation of derivatives through finite difference approximations) or symbolic differentiation (propagation of analytical derivative formulas) in that only the numerical value of the derivative is obtained, but as accurately and efficiently as the original function evaluation.

Automatic differentiation features two main modes of operation: forward mode and reverse mode. While both are based on a repeated application of the chain rule, they differ in the direction in which it is applied (starting from the inputs or from the outputs). Importantly, this implies a different scaling with the number of inputs and outputs of the relevant function $f : \mathbb{R}^m \rightarrow \mathbb{R}^p$. Indeed, forward mode automatic differentiation is equivalent in memory and time to $\mathcal{O}(m)$ function evaluations, whereas reverse mode automatic differentiation is equivalent to $\mathcal{O}(p)$ function evaluations. Because quantum optimal control usually features many parameters ($m \gg 1$) and only a single cost function ($p = 1$), the latter mode is well suited to the problematic at hand and will be our focus from now on. For a detailed review of both modes, see [Baydin *et al.* 2018].

To illustrate how reverse mode automatic differentiation works, let us consider the example $f(\theta_1, \theta_2) = \sin(\theta_1) + \theta_1 \sqrt{\theta_2}$, borrowed from [Leung *et al.* 2017]. Its computational graph is represented in Figure 4.1. In this graph, each node represents an elementary operation, and each arrow an intermediate value w_i with $i \in [0, 6]$ that is used by a computer to numerically evaluate the function for a given set of inputs. This list of intermediate values is shown in the table below Figure 4.1 and is referred to as an *evaluation trace*. A forward pass then refers to a single evaluation of the function f . In the example of Figure 4.1, starting from the numerical value of both parameter (θ_1, θ_2) , the forward pass will subsequently evaluate all intermediate values w_0, w_1, w_2, \dots , until the final output is reached, $y \equiv f(\theta_1, \theta_2) = w_6$.

The backward pass is then carried out to evaluate gradients. In this pass, an adjoint is associated to each intermediate variable w_i , and represents the sensitivity of the final output y to a change in w_i . It is defined as

$$\bar{w}_i = \frac{\partial y}{\partial w_i}. \quad (4.1)$$

Importantly, the chain rule yields simple relations between these adjoint intermediate values which are directly linked to the elementary operations from




Figure 4.1: Computational graph of $f(\theta_1, \theta_2) = \sin(\theta_1) + \theta_1 \sqrt{\theta_2}$ and schematic of reverse mode automatic differentiation. (a) In the forward pass, operations are performed successively starting from the parameters (θ_1, θ_2) . (b) In the backward pass, the graph is back-propagated starting from the final node. A single pass is needed to compute the gradient with respect to all parameters. The table shows the intermediate values w_i and \bar{w}_i that are used in the computation. Figure and table inspired from [Baydin *et al.* 2018].

the computational graph. As such, the backward pass will iterate successively over these adjoint states, up until the gradient with respect to all parameters is computed. In the example of Figure 4.1, we thus get,

$$\bar{w}_6 = \frac{\partial y}{\partial w_6} = \frac{\partial w_6}{\partial w_6} = 1, \quad (4.2)$$

which is the seed of the reverse-mode differentiation, followed by

$$\bar{w}_5 = \frac{\partial y}{\partial w_5} = \frac{\partial y}{\partial w_6} \frac{\partial w_6}{\partial w_5} = \bar{w}_6 \cdot 1 = 1, \quad (4.3)$$

$$\bar{w}_3 = \frac{\partial y}{\partial w_3} = \frac{\partial y}{\partial w_5} \frac{\partial w_5}{\partial w_3} = \bar{w}_5 \cdot w_1 = \theta_1, \quad (4.4)$$

and finally,

$$\bar{w}_2 = \frac{\partial y}{\partial w_2} = \frac{\partial y}{\partial w_3} \frac{\partial w_3}{\partial w_2} = \bar{w}_3 \cdot \frac{1}{2\sqrt{w_2}} = \frac{\theta_1}{2\sqrt{\theta_2}}, \quad (4.5)$$

which is also the gradient with respect to θ_2 since $w_2 = \theta_2$. Hence, following a similar approach for the upper part of the graph, the gradient with respect to all input parameters can be computed efficiently, with the same computational cost as the initial function evaluation and in a single pass. Importantly, forward pass intermediate values should be kept in memory for the backward pass, as they may be used in gradient evaluations. For instance, this was the case in (4.4) in which the w_1 value was required. In general, any intermediate value involved in a multiplicative operation will need to be stored in memory. We can already see how this could be problematic for problems involving large amounts of matrix multiplications.

Coming back to quantum optimal control, automatic differentiation can be a powerful tool for computing gradients with respect to arbitrary cost functions. For many problems of interest, the cost function is straightforwardly defined from the density matrix at some given time, and/or from the parameters. Since the density matrix is itself a function of the parameters, through

$$\rho(t) = \mathcal{T} \exp\left(\int_{t_0}^t \mathcal{L}_\theta(s) ds\right) \rho(t_0) \quad (4.6)$$

where \mathcal{T} is the time-ordering operator, and \mathcal{L}_θ is the parametrized Liouvillian, it is possible to obtain the gradients $dC/d\theta$ using reverse mode automatic differentiation [Leung *et al.* 2017]. While this has seen many applications, the main limitation of the method comes from memory requirements since reverse-mode automatic differentiation requires storage of all intermediate values during the forward pass. Here, intermediate values are density matrices so this memory requirement can quickly blow up for large quantum systems, as pointed out in [Lu *et al.* 2023]. Concretely, for a Hilbert space of size $N = 1000$, and $n = 2000$ numerical Liouvillian evaluations, the sole storage of all density matrices in double floating point precision already represents 29.5 Gb of memory. This is more than most modern GPUs or CPU RAM capacities, and does not even take into account other objects that should be also stored (e.g. Hamiltonians, jump operators and so on). To overcome this limitation, the following section introduces the adjoint state method for open quantum systems.

4.2.2 Adjoint state method

The *adjoint state method* is a numerical method used to compute gradients that was first introduced in [Pontryagin *et al.* 1962] and has seen applications

in several fields of science and engineering that rely on partial differential equations. It is based on the back-propagation of an adjoint state, a dual representation of the original state. In this section, we derive the adjoint state method for open quantum systems.

Let us consider the adjoint state of the density matrix,

$$\phi(t) \equiv \frac{dC}{d\boldsymbol{\rho}(t)} \quad (4.7)$$

where C is the previously introduced cost function, and $\boldsymbol{\rho}$ the density matrix. This adjoint state is an instantaneous representation of the sensitivity of the cost function with respect to the density matrix. For open quantum systems, the density matrix is solution to a master equation,

$$\frac{d\boldsymbol{\rho}}{dt} = \mathcal{L}\boldsymbol{\rho} \equiv -i[\mathbf{H}, \boldsymbol{\rho}] + \sum_k \mathcal{D}[\mathbf{L}_k]\boldsymbol{\rho}, \quad (4.8)$$

where \mathbf{H} is the system Hamiltonian, \mathbf{L}_k are jump operators, and $\mathcal{D}[\mathbf{L}]\boldsymbol{\rho} = \mathbf{L}\boldsymbol{\rho}\mathbf{L}^\dagger - \{\mathbf{L}^\dagger\mathbf{L}, \boldsymbol{\rho}\}/2$. It is then possible to show (see Section 4.2.3) that the adjoint state is subject to an analogous master equation,

$$\frac{d\phi}{dt} = -\mathcal{L}^\dagger\phi \equiv -i[\mathbf{H}, \phi] - \sum_k \mathcal{D}^\dagger[\mathbf{L}_k]\phi, \quad (4.9)$$

where $\mathcal{D}^\dagger[\mathbf{L}]\phi = \mathbf{L}^\dagger\phi\mathbf{L} - \{\mathbf{L}^\dagger\mathbf{L}, \phi\}/2$. The adjoint master equation (4.9) can thus be integrated numerically over the time interval of interest $[t_0, t_n]$ from an initial condition at the final time, $\phi(t_n) \equiv dC/d\boldsymbol{\rho}(t_n) = \partial C/\partial\boldsymbol{\rho}(t_n)$, where the second equality stands if the cost function is an explicit function of the final density matrix. This initial condition can either be computed analytically if possible, and otherwise through automatic differentiation. Notice in particular the overall minus sign in (4.9) which will ensure contracting dynamics in reverse time, and thus numerical stability of the solution. Gradients with respect to the problem parameters are then given by

$$\frac{dC}{d\theta} = \frac{\partial C}{\partial\theta} - \int_{t_n}^{t_0} \partial_\theta \text{Tr}[\phi^\dagger(t)\mathcal{L}(t, \theta)\boldsymbol{\rho}(t)] dt, \quad (4.10)$$

as we also derive in Section 4.2.3. Therefore, the gradients required for optimal control can be computed immediately from the knowledge of both the adjoint state and density matrix at all times.

A typical optimization loop is represented in Figure 4.2(a). From an initial set of parameters (here, a sequence of pulses), a forward pass that consists of solving the master equation once over the full integration time is performed. From the density matrices computed, the cost function corresponding to the

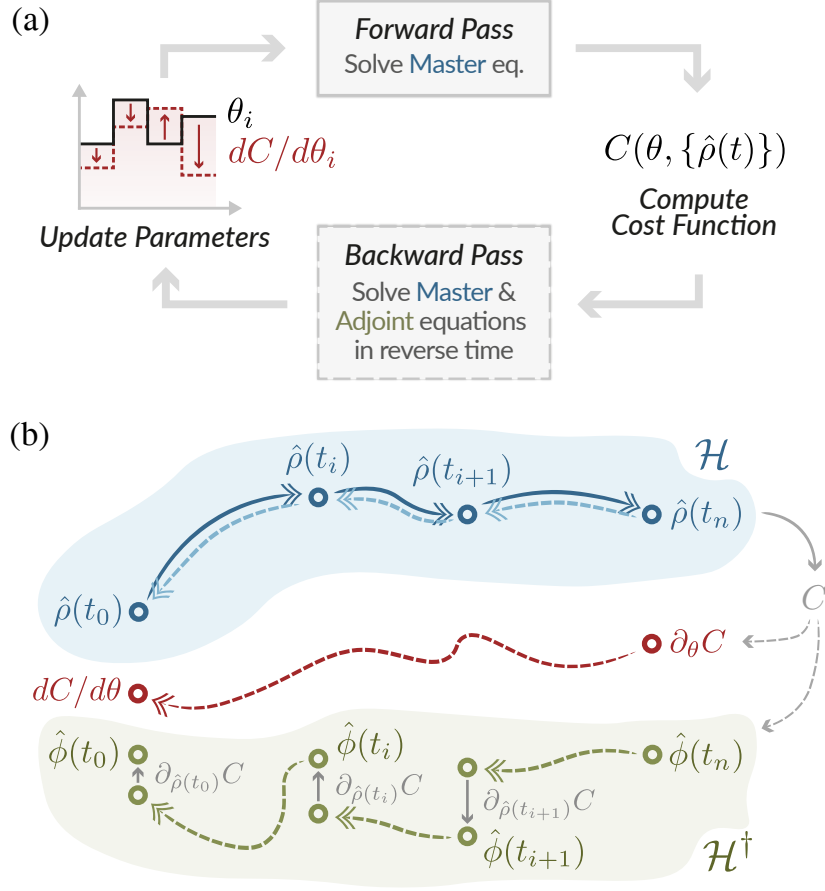


Figure 4.2: Adjoint state quantum optimal control. (a) Optimization loop. Gradients are computed during the backward pass, and fed to a gradient-descent algorithm that computes the parameter updates. (b) Schematic of the forward and backward passes. The master equation is integrated and checkpointed for several time points (dark blue). In the backward pass (dashed lines), the density matrix is recomputed in reverse time (light blue) together with the adjoint state (green) and with the gradients (red). When a checkpoint is reached, the density matrix is restored to its forward time trajectory, and the adjoint state updated with the corresponding cost function gradient.

initial set of parameters is obtained. Then, the optimization goes through a backward pass during which both the master and adjoint equations are integrated in reverse time and simultaneously [Chen *et al.* 2018]. Equation (4.10) then yields the full gradients which are fed to the optimization algorithm to iterate over the parameters.

More precisely, Figure 4.2(b) shows a schematic of both the forward and backward passes. During the forward time integration of the master equation (dark blue lines), several solutions are saved to compute the final cost function

and as checkpoints for the backward pass¹. During the backward pass, both the master (light blue) and adjoint state (green) equations are integrated in reverse time. The adjoint state is regularly updated by the partial derivative of the cost function with respect to the checkpointed density matrix to account for specific dependencies of the cost function. Finally, the integrand of (4.10) is evaluated at every time step through one of two ways. Either the partial gradient $\partial_\theta \mathcal{L}(t, \theta)$ is computed analytically and substituted inside the equation, or more generally, it can be computed from automatic differentiation of the adjoint equation integration by noting that $\partial_\theta \text{Tr}[\phi^\dagger \mathcal{L} \rho] = -\text{Tr}[\partial_\theta (d\phi/dt)^\dagger \rho]$. We utilize the latter method.

The full optimization loop features a time cost equivalent to the integration of four master equations, and a $\mathcal{O}(n_c \times N^2)$ memory footprint where N is the Hilbert space dimension and n_c the number of time points involved in the cost function definition. For a large class of problems, either only the final density matrix is needed and $n_c = 1$, or the cost function can be partially computed on the fly when time $t = t_i$ is reached in which case the cost reduces to $\mathcal{O}(N^2)$. Our method thus features a reduction of multiple orders of magnitude in memory compared to standard automatic differentiation methods for which the density matrix has to be saved at every numerical time step [Leung *et al.* 2017], a major bottleneck of the approach [Lu *et al.* 2023].

In practice, this whole optimization scheme is implemented in PyTorch for its automatic differentiation capabilities and GPU support. We are able to run optimization problems for open quantum system with hundreds of parameters and arbitrary cost functions, for Hilbert space dimensions of up to $N \lesssim 5,000$ while keeping memory requirements within GPU capacity. In the following subsection, we derive the adjoint state method equations (4.9) and (4.10) for open quantum systems. We then demonstrate this method for the optimization of transmon readout, a notoriously difficult problem to simulate numerically.

4.2.3 Derivation of adjoint state method equations

4.2.3.1 Adjoint state master equation

We will now prove that the adjoint state follows the differential equation (4.9), following a similar derivation as in [Chen *et al.* 2018]. Starting from

¹The reverse time master equation features expanding dynamics, which may induce divergence during numerical integration. Checkpointing on the time scale of the largest dissipation operator stabilizes the numerical integration.

the master equation (4.8), we can write that for an $\varepsilon \ll 1$ change in time,

$$\begin{aligned}\boldsymbol{\rho}(t + \varepsilon) &= \boldsymbol{\rho}(t) + \int_t^{t+\varepsilon} \mathcal{L}(\tau, \theta) \boldsymbol{\rho}(\tau) d\tau \\ &= \boldsymbol{\rho}(t) + \varepsilon \mathcal{L}(t, \theta) \boldsymbol{\rho}(t) + \mathcal{O}(\varepsilon^2)\end{aligned}\quad (4.11)$$

Using the chain rule and the definition of the adjoint state, we then find

$$\phi(t) = \frac{dC}{d\boldsymbol{\rho}(t)} = \frac{dC}{d\boldsymbol{\rho}(t + \varepsilon)} \frac{d\boldsymbol{\rho}(t + \varepsilon)}{d\boldsymbol{\rho}(t)} = \phi(t + \varepsilon) (1 + \varepsilon \mathcal{L}(t, \theta) + \mathcal{O}(\varepsilon^2)) \quad (4.12)$$

The proof of (4.9) follows from the definition of the derivative,

$$\begin{aligned}\frac{d\phi(t)}{dt} &= \lim_{\varepsilon \rightarrow 0} \frac{\phi(t + \varepsilon) - \phi(t)}{\varepsilon} \\ &= -\lim_{\varepsilon \rightarrow 0} (\phi(t + \varepsilon) \mathcal{L}(t, \theta) + \mathcal{O}(\varepsilon^2)) \\ &= -\phi(t) \mathcal{L}(t, \theta)\end{aligned}\quad (4.13)$$

or equivalently, using that the adjoint state is hermitian, $\phi^\dagger = dC/d\boldsymbol{\rho}^\dagger = dC/d\boldsymbol{\rho} = \phi$, we get the required result,

$$\frac{d\phi(t)}{dt} = -\mathcal{L}^\dagger(t, \theta) \phi(t). \quad (4.14)$$

4.2.3.2 Gradients

To derive an explicit expression for $dC/d\theta$, we also follow [Chen *et al.* 2018] and introduce an augmented density matrix $\boldsymbol{\rho}_{aug}(t) = [\boldsymbol{\rho}(t), \theta(t)]^T$ which includes a second block row with time-dependent parameters. In practice, parameters are constant throughout the integration, but this fictitious time-dependence will allow us to isolate the sensitivity of the cost function to parameter changes at every point in time. The augmented density matrix follows the differential equation,

$$\frac{d\boldsymbol{\rho}_{aug}(t)}{dt} = \mathcal{L}_{aug}(t, \theta) \boldsymbol{\rho}_{aug}(t) = \begin{pmatrix} \mathcal{L}(t, \theta) \boldsymbol{\rho}(t) \\ 0 \end{pmatrix} \quad (4.15)$$

We can also introduce the augmented adjoint state, $\phi_{aug}(t) = [\phi(t), dC/d\theta]$, where the parameter sensitivity $dC/d\theta$ also has a fictitious time-dependence. We then follow a similar derivation as in the previous section. The main difference comes in the application of the chain rule of (4.12) which now yields

$$\begin{aligned}\phi_{aug}(t) &= \frac{dC}{d\boldsymbol{\rho}_{aug}(t)} = \frac{dC}{d\boldsymbol{\rho}_{aug}(t + \varepsilon)} \frac{d\boldsymbol{\rho}_{aug}(t + \varepsilon)}{d\boldsymbol{\rho}_{aug}(t)} \\ &= \phi_{aug}(t + \varepsilon) + \varepsilon \phi_{aug}(t + \varepsilon) \begin{pmatrix} \mathcal{L}(t, \theta) & \partial_\theta \mathcal{L}(t, \theta) \boldsymbol{\rho}(t) \\ 0 & 0 \end{pmatrix}\end{aligned}\quad (4.16)$$

where ∂_θ denotes the partial derivative with respect to θ . Similarly as before, we can rearrange this expression, reinsert it in the definition of the derivative and take the limit of $\varepsilon \rightarrow 0$. Then, selecting only the second element of the resulting differential equation yields,

$$\frac{d}{dt} \left(\frac{dC}{d\theta} \right) = -\partial_\theta \text{Tr}[\phi^\dagger(t) \mathcal{L}(t, \theta) \rho(t)]. \quad (4.17)$$

We have thus obtained a straightforward differential equation on the parameter sensitivity. Integrating this equation from $t = 0$ to T , with the initial condition that $dC/d\theta(T) = \partial C/\partial\theta$ yields,

$$\frac{dC}{d\theta} = \frac{\partial C}{\partial\theta} - \int_T^0 \partial_\theta \text{Tr}[(\mathcal{L}^\dagger(t, \theta) \phi^\dagger(t)) \rho(t)] dt \quad (4.18)$$

which is (4.10). We can also get the derivative of the loss function with respect to the integration time T from a simple application of the chain rule,

$$\frac{dC}{dT} = \frac{dC}{d\rho(T)} \frac{d\rho(T)}{dT} = \text{Tr}[\phi^\dagger(T) \mathcal{L}(T, \theta) \rho(T)]. \quad (4.19)$$

Finally, between (4.9), (4.18) and (4.19), we get the gradients of the loss function with respect to all relevant objects involved in the computation.

4.3 Optimal control of transmon readout

Now that we have introduced the adjoint state method for open quantum systems, let us look into how it can be employed for the optimization of transmon readout pulses.

4.3.1 Transmon-resonator-filter model

Consider the following model of a transmon coupled to a resonator and Purcell filter ($\hbar = 1$), as depicted in Figure 4.3,

$$\begin{aligned} \mathbf{H}(t) = & 4E_C \mathbf{n}_t - E_J \cos(\varphi_t) + \omega_r \mathbf{a}^\dagger \mathbf{a} + \omega_f \mathbf{f}^\dagger \mathbf{f} \\ & - ig \mathbf{n}_t (\mathbf{a} - \mathbf{a}^\dagger) - J (\mathbf{a} - \mathbf{a}^\dagger) (\mathbf{f} - \mathbf{f}^\dagger) \\ & + \Omega_t \mathbf{n}_t \sin(\omega_{d,t} t) - i\Omega_f (\mathbf{f} - \mathbf{f}^\dagger) \sin(\omega_{d,f} t). \end{aligned} \quad (4.20)$$

The first two terms denote the free transmon Hamiltonian with charging energy E_C , Josephson energy E_J , and with \mathbf{n}_t and φ_t the charge and phase operators. Resonator and filter modes are denoted by \mathbf{a} and \mathbf{f} , with respective frequencies ω_r and ω_f . These three modes are capacitively coupled in




Figure 4.3: Lumped-element circuit of a transmon-resonator-filter model. The transmon (green) and readout resonator (blue) are capacitively coupled with coupling g . The resonator is itself capacitively coupled to a Purcell filter (dark blue) with coupling $J \ll g$. This filter is damped and driven through a transmission line.

series with coupling strengths $g \gg J$, and the full system can be driven either through the transmon at frequency $\omega_{d,t}$ and amplitude Ω_t , or through the Purcell filter at frequency $\omega_{d,f}$ and amplitude Ω_f . Our model also includes transmon relaxation with rate κ_t and the engineered filter loss with rate κ_f through which readout is performed. This is a standard model for dispersive readout of a transmon [Sete *et al.* 2015], where the weak dispersive coupling to the readout resonator — obtained e.g. after a Schrieffer-Wolff transformation of (4.20) — can probe the transmon through a state-dependent frequency shift, and where the Purcell filter is introduced to shield the transmon from Purcell decay [Purcell 1946]. Readout is then performed through heterodyne measurement of the filter.

For the numerical simulation of this model, we first diagonalize the free transmon Hamiltonian and identify the lowest energy eigenstates. We also diagonalize the resonator-filter subsystem yielding two normal modes, each directly coupled to the transmon. Finally, we assume the Rotating Wave Approximation (RWA) on all couplings and drives. This allows for larger numerical time steps by eliminating fast oscillating dynamics, which eventually simplifies master equation integration. However, this also implies that some non-standard dynamics are not fully captured, such as transmon chaos [Cohen *et al.* 2022] or ionization [Shillito *et al.* 2022]. In an effort to avoid those regimes, we will intentionally limit resonator populations below the critical photon number of the system [Blais *et al.* 2021]. Note that this RWA is not required for the adjoint state method, and is here only to simplify the numerical endeavor.

In the rest of this chapter, we fix model parameters to demonstrate the effectiveness of the optimal control method on a given system. Unless stated otherwise, we use $E_J/E_C = 51$, $E_C/2\pi = 315$ MHz, correspond-

ing to a bare transmon frequency $\omega_t/2\pi = 6$ GHz and to an anharmonicity $\alpha/2\pi = -349$ MHz. Resonator and filter frequencies are $\omega_r/2\pi = 7.2$ GHz and $\omega_f/2\pi = 7.21$ GHz and couplings $g/2\pi = 150$ MHz and $J/2\pi = 30$ MHz. This yields a transmon-resonator detuning of $\Delta/2\pi = 1.2$ GHz and a critical photon number $\bar{n}_{\text{crit}} = (\Delta/2g)^2 = 16$. Finally, the filter loss rate is $\kappa_f/2\pi = 30$ MHz and the transmon relaxation rate $\kappa_t/2\pi = 8$ kHz corresponding to $T_1 = 1/\kappa_t = 20$ μ s.

4.3.2 Dispersive readout metrics

Dispersive readout is a widespread superconducting circuit readout technique based on the dispersive coupling of a qubit to a resonator. By neglecting the filter in the model of (4.20) and reducing the transmon to a two-level system, this coupling takes the form $\chi\sigma_z\mathbf{a}^\dagger\mathbf{a}$ and highlights the transmon state-dependent frequency shift on the resonator that can be leveraged for readout. Therefore, a weak state-dependent rotation in phase space is induced when the resonator is not in vacuum. Then, standard heterodyne detection of the resonator can be used to distinguish between either states.

To maximize the fidelity of a dispersive readout, there are mainly two criteria to optimize. First, in order to minimize the impact of transmon decay and eventual thermalization, the readout sequence should be fast. In the limit of a fast readout compared to the transmon lifetime, the fidelity error induced by this effect grows linearly in time according to $\varepsilon_{a,0} \propto \tau_m/T_1$ where τ_m is the readout integration time and T_1 the transmon lifetime. Second, the readout contrast should be large. Defining the average field value in the filter as $\beta_{e/g} = \text{Tr}[\mathbf{f}\boldsymbol{\rho}_{g/e}]$ where $\boldsymbol{\rho}_{g/e}$ is the density matrix obtained after initialization in the transmon $|g/e\rangle$ state, this means that the $|\beta_e - \beta_g|$ distance should stay as large as possible for as long as possible. To quantify this second criteria, the standard scale is Signal-to-Noise Ratio (SNR), defined as

$$\text{SNR}^2(\tau_m) = \frac{|\langle \mathbf{M}(\tau_m) \rangle_e - \langle \mathbf{M}(\tau_m) \rangle_g|^2}{\langle \mathbf{M}_N^2(\tau_m) \rangle_e + \langle \mathbf{M}_N^2(\tau_m) \rangle_g} \quad (4.21)$$

where $\mathbf{M}(\tau_m) = \int_0^{\tau_m} (w_x \mathbf{V}_x(t) + w_p \mathbf{V}_p(t)) dt$ is the measurement signal, with tunable weighting functions $w_{x/p}$ and quadrature signals $\mathbf{V}_{x/p}$ that include the average quadrature field value and a white-noise contributions from the uncertainty principle. Also, $\mathbf{M}_N^2 = (\mathbf{M} - \langle \mathbf{M} \rangle)^2$ is the measurement noise operator that takes into account added noise and intrinsic vacuum noise. One can further show that with optimal weighting functions, the following simplified formula for the SNR holds [Bultink *et al.* 2018],

$$\text{SNR}(\tau_m) = \sqrt{2\eta\kappa_f \int_0^{\tau_m} |\beta_e(t) - \beta_g(t)|^2 dt} \quad (4.22)$$

where η is the measurement efficiency, κ_f the Purcell filter decay rate. Indeed, the measurement signal scales linearly with the integral of $|\beta_e - \beta_g|^2$, while the measurement noise scales only as the square root of the same integral, hence the overall square root scaling. Thanks to its robustness and simplicity to numerical evaluation, the rest of this work will use this formula for SNR computations.

Coming back to readout fidelity, it can be expressed in a straightforward manner from the two quantities we have introduced, i.e. transmon lifetime and SNR. Again assuming a simplified gaussian model, we have [Gambetta *et al.* 2007, Swiadek *et al.* 2023],

$$\begin{aligned}\varepsilon_a(\tau_m) &= 1 - \mathcal{F}_{g,e} = (P(g|e) + P(e|g)) / 2 \\ &\gtrsim \frac{1}{2} \left(1 - \operatorname{erf} \left(\frac{\operatorname{SNR}(\tau_m)}{2} \right) \right) + \frac{\tau_m}{2T_1}\end{aligned}\quad (4.23)$$

where ε_a is the assignment error, $\mathcal{F}_{g,e}$ the readout fidelity, and $P(i|j)$ the probability of measurement state $|i\rangle$ knowing that $|j\rangle$ was prepared. The assignment error is thus the key metric to minimize when optimizing readout. In the rest of this work, we will assume the formula of the second line when evaluating assignment errors numerically.

4.3.3 Optimizing readout

For the numerical optimization of readout pulses, the process was the following. We first fix the readout integration time τ_m , and discretize the pulse in time steps of $\tau_0 = 1$ ns. The pulse amplitude at each time step thus yields a set of $N = \lceil \tau_m / \tau_0 \rceil$ complex-valued parameters Ω_j . This discrete pulse is then fed to a gaussian filter \mathcal{F}_g of bandwidth $\omega_{\text{bw}} / 2\pi = 425.4$ MHz to replicate a realistic AWG [Motzoi *et al.* 2011]. The resulting pulse then reads

$$\Omega(t) = \mathcal{F}_g \left[\sum_{j=0}^{N-1} \Omega_j \Pi_j(t, \tau_0) \right] \quad (4.24)$$

with the rectangle function

$$\Pi_j(t, \tau_0) \equiv \Theta(t - j\tau_0) - \Theta(t - (j+1)\tau_0), \quad (4.25)$$

where Θ is the Heaviside unit step function. The pulse carrier frequency ω_d is also parameterized, but fixed throughout the $[0, \tau_m]$ time interval. Using this continuous pulse, the three-mode open quantum system model with Hamiltonian (4.20) is integrated over the readout time. We then evaluate a cost function tailored to the optimization problem. The main contribution

to this cost function is the SNR of (4.22), which is straightforward to evaluate from the integrated density matrix at evenly spaced time points sampled throughout the evolution. Other smaller contributions to the cost function include penalties for large photon populations in both resonator and filter modes, penalties for pulse amplitudes above a certain ceiling (here, 200 MHz), or population in excited transmon states. Finally, the gradient of this cost function with respect to each parameter is evaluated using the adjoint state method introduced in Section 4.2.

This whole process represents a single iteration of the optimization. In practice, several hundred iterations are performed and parameters are iterated over at each epoch using an Adam gradient descent algorithm [Kingma & Ba 2014]. The seed parameters at the first epoch are either randomly sampled, or taken from a known working pulses (e.g. a flat readout pulse). Once the optimizer has converged, we perform a final run of the optimized pulses with a larger Hilbert space and a smaller numerical time step in order to validate results. This validation run is meant to speed up the overall optimization by allowing faster iterations, yet still provide an accurate final result. In the following section, we show these validated results for several optimization runs, and demonstrate improved readout compared to standard pulses.

4.3.4 Results

Figure 4.4 shows both the SNR and the assignment error against the integration time of readout for various reference and optimized pulse sequences. Let us first discuss the reference pulses, plotted in light and dark blue color, before moving on to the optimized pulses. The first reference corresponds to a flat pulse (including a 2 ns pulse ramp-up and ramp-down) with amplitude fixed at $|\Omega_{f,0}|/2\pi = 70$ MHz, and the second to a two-step pulse with a strong initial drive at $|\Omega_{f,1}|/2\pi = 160$ MHz for the first 4 ns (designed to rapidly populate the resonator [Walter *et al.* 2017]), followed by the same flat pulse as before. These two pulse envelopes are shown in Figure 4.5(a). The $|\Omega_{f,0}|$ amplitude is chosen such that the resonator reaches a population $\bar{n} = \bar{n}_{\text{crit}}$ in its steady state. In addition, the carrier frequency of these pulses is fixed at the dressed frequency of the lowest normal mode in the resonator-Purcell filter subsystem.

As the integration time increases, we find that both references feature an increasing SNR and a decreasing assignment error up until the transmon T_1 limit is reached (materialized by the solid black line). The minimum assignment error is obtained at $\tau_m \simeq 80$ ns and $\tau_m \simeq 65$ ns respectively for the flat and two-step pulses. This is similar performance to state-of-the-art readout experiments [Walter *et al.* 2017, Sunada *et al.* 2022, Swiadek *et al.* 2023],




Figure 4.4: Optimizing transmon readout with the adjoint state method. Signal-to-noise ratio (top) and assignment error (bottom) against readout integration time for various reference and optimized pulses. Flat and two-step results are fitted with (4.26). Optimized pulses are shown in red/green. Hollow markers show optimized pulses with the addition of a transmon drive of optimized frequency at $\omega_{d,t} \approx \omega_r$ (red) and $\omega_{d,t} \approx \omega_{ef}$ (green). For all pulses, the assignment error is limited by transmon lifetime errors (black line).

as expected from our choice of experimental parameters. We further provide a fit of the SNR, shown in solid lines, according to the analytical formula [Didier *et al.* 2015],

$$\text{SNR}(\tau_m) = \alpha \sqrt{2\eta\kappa} (\sqrt{\tau_m} - \sqrt{\tau_{m,0}}) \quad (4.26)$$

where $\alpha = 2|\Omega_f \sin(2\phi)|/\kappa_f$ is the effective resonator displacement in the steady state, with $\phi = \arctan(2\chi/\kappa_f)$ and χ the dispersive shift obtained from exact diagonalization of (4.20). This formula shows a square root dependence of the SNR with the integration time in the $\tau_m \rightarrow \infty$ limit, as expected from (4.22). Furthermore, the term $\sqrt{\tau_{m,0}}$ corresponds to an initial delay for the resonator to populate. It is numerically fitted, and we find $\tau_{m,0} \simeq 19$ ns for the flat pulse, and $\tau_{m,0} \simeq 13$ ns for the two-step pulse. We find an excellent agreement between (4.26) and the numerical simulations.

Let us now compare these two reference pulses to optimized pulses. Red markers in Figure 4.4 show the SNR and assignment error for an optimized pulse. We find a modest improvement compared to the previously discussed two-step pulse, suggesting that this strategy is already almost optimal if only a drive on the filter is available. Indeed, the dispersive coupling being bounded, the speed at which the $\rho_{f,g/e}$ states in the resonator separate is bound, and it is likely that one can only improve readout by decreasing the delay for the resonator to populate, $\tau_{m,0}$. Figure 4.5(b) shows the optimized pulse for the




Figure 4.5: Example of various reference and optimized pulses for a 40 ns readout. Solid and dashed lines show the real and imaginary parts of (4.24) respectively. (a) Flat (dark blue) and two-step (light blue) pulses. (b) Optimized readout pulse. (c) Optimized readout pulses with a transmon drive at $\omega_{d,t} \approx \omega_r$. (d) Optimized readout pulses with a transmon drive at $\omega_{d,t} \approx \omega_{ef}$.

40 ns run. A two-step-like shape can be recognized, with a strong drive for $t \lesssim 5$ ns and a smaller drive for the rest of the sequence. We attribute small oscillations throughout the pulse to a gauge degree of freedom of the optimizer (the state is free to rotate in phase space).

We then study the optimization of two other readout strategies, both involving an additional drive on the transmon. The first strategy, shown in hollow red markers, involves a transmon drive at the frequency of the readout resonator, $\omega_{d,t} \approx \omega_r$ and is meant to reproduce the cloaking mechanism of [Lledó *et al.* 2022, Muñoz-Arias *et al.* 2023] for which the resonator population can be controlled independently of the transmon. While the cloaking mechanism is not retrieved here, simulations still show improvements compared to reference pulses, with the optimal assignment error obtained at $\tau_m \simeq 60$ ns. The pulses for the 40 ns run are shown in Figure 4.5(c) for completeness, although we currently are lacking an interpretation of what the optimizer is converging towards.

The second strategy, shown in hollow green markers, involves a transmon drive at the frequency of the $|e\rangle$ - $|f\rangle$ transmon transition. This time, the optimizer finds a pulse sequence that populates the $|f\rangle$ state of the transmon to probe a larger dispersive coupling χ_{fg} than the standard shift χ_{eg} [Chen *et al.* 2023]. Here, a clear improvement both in SNR and assignment error is demonstrated, with the optimal assignment error achieved at $\tau_m \simeq 35$ ns, almost twice as fast the best reference pulse and with an error




Figure 4.6: Phase space distance in the Purcell filter mode between $|g\rangle$ and $|e\rangle$ transmon states against time, for a 40 ns readout and for various reference and optimized pulses. The SNR is the integral of this quantity, up to some multiplicative factor.

about twice as low. In addition, Figure 4.5(d) shows the corresponding pulses, again for the 40 ns pulse sequence. On the filter side, the pulse is very similar to that of the optimized sequence of panel (b) with a two-step-like pulse. On the transmon side however, the optimizer converges towards a short π -pulse of the $|e\rangle$ - $|f\rangle$ transition at the beginning of the sequence, followed by no drive for the rest of readout. There are a few remarkable things to note. First, the π -pulse uses a DRAG-like sequence [Motzoi *et al.* 2009], with a drive on both quadratures to eliminate leakage to other transmon states. We find a fidelity in excess of 99% on this π -pulse achieved in less than 10 ns and while the filter is also being driven. Second, this π -pulse is performed at the same time as the resonator is populated, such that the optimizer is fusing operations together instead of doing them sequentially, a crucial time gain. And finally, the frequency of this π -pulse being also optimized, we find that the optimizer is taking into account the stark shift induced by the large drive on the Purcell filter. Overall, this optimized pulse is extremely promising and suggests that further experimental improvements on readout fidelities can still be achieved. As a final remark, note that this readout strategy is not a Quantum Non Demolition (QND) measurement, but that QND-ness could be retrieved with a subsequent $|f\rangle$ - $|e\rangle$ transition that could be requested within the optimization scheme.

To complete this results section, we show in Figure 4.6 the effective distance $|\beta_e - \beta_g|^2$ in the Purcell filter against time for the previously discussed readout schemes. The quantity being optimized is thus the integral of these curves, up to a constant prefactor. In all cases, the distance is approximately zero in the beginning of readout, and then increases in an oscillating pattern. There are a few things to note on this figure. First, optimized distances begin ramping up sooner than non-optimized distances. Second, the oscillations of

optimized distances are flattened out by the optimizer, such as to keep the largest distance at all times. This confirms that, while the dispersive coupling is bound by the experimental setup, carrying out pulse engineering on readout pulses can still provide improvements on readout contrasts.

4.4 Conclusion

In this chapter, we introduced the adjoint state method for open quantum systems, which allows for the computation of gradients of an arbitrary function with respect to any number of parameters in a constant $\mathcal{O}(N^2)$ memory cost, with N the Hilbert space size. This method is completely general, adaptable to any open quantum system, and enables applications in optimal control, experimental calibration or state tomography. The method is based on the introduction of the adjoint state of the density matrix, which represents the sensitivity of the function of interest to changes in the density matrix. By solving the adjoint state master equation in the backward pass, and simultaneously recomputing the solution to the master equation in reverse time, the memory cost of the method can be largely reduced compared to standard automatic differentiation.

To make this method widely available without the overhead of re-implementing it, we have been developing a library aimed at providing end-to-end differentiable solvers of quantum differential equations. The *dynamiqs* library is open-source and available at github.com/dynamiqs/dynamiqs. It provides solvers for Schrödinger, master and stochastic master equations with various numerical integration methods. It also supports batching over many initial states or Hamiltonians and is GPU-enabled. While still under active development, the library can already serve as a basis for similar applications as demonstrated in this work.

Here, we demonstrated how the adjoint state method applies to the optimization of transmon readout, a complex problem involving multiples modes and inherent dissipation. Through pulse engineering, we have been able to demonstrate improvements in readout fidelity on a complete readout setup that includes a many-level transmon, a readout resonator and a Purcell filter. We have also explored different readout strategies, with additional drives on the transmon meant to achieve cloaking or shelving. In particular, the shelving strategy shows a two-fold improvement in readout time and in readout fidelity on the specific experimental parameters we consider. The gradient descent optimization further found interpretable pulses that would suggest these pulses are robust to parameter drifts, and can be applied to other experimental setups.

Conclusions and perspectives

Since 2012, several research groups have been working towards engineering cat qubits in superconducting circuits thanks to their promise of hardware-efficient, fault-tolerant and universal quantum computation. This began with a proposal for the autonomous correction of cat qubit errors in a superconducting cavity [Leghtas *et al.* 2013], shortly followed by the seminal paper on their stabilization with engineered two-photon dissipation [Mirrahimi *et al.* 2014]. The paper also featured proposals for single- and multi-qubit gate implementations and opened the way to universal quantum computation with cat states. In 2014, the first experiment of dissipative stabilization was put together in a transmon-based device [Leghtas *et al.* 2015]. Single-qubit Rabi oscillations were then shown in a similar device [Touzard *et al.* 2018], although evidence of the promised exponential noise bias had yet to be presented. With the invention of the Asymmetrically Threaded SQUID (ATS) in 2019 [Lescanne *et al.* 2020b], the Leghtas group was eventually able to engineer the large non-linearity required for dissipative stabilization of cat qubits but at a Kerr-free point. This allowed for the first demonstration of exponential noise bias, with a lifetime saturating at about 10 ms. Recently, with the invention of a transmon-free readout protocol [Réglade *et al.* 2023], we were able to enhance this lifetime by up to 10 s with no apparent saturation, and to demonstrate the first bias-preserving gate on cat qubits. This is indeed a crucial step towards quantum computation with cat qubits as laid out in [Guillaud & Mirrahimi 2019], with not only the integrated protection of bit-flip errors, but also with an outer error-correction code for phase-flip errors.

Concurrently to this work on dissipative cat qubits, an alternative confinement mechanism was invented [Puri *et al.* 2017] and pursued. Using the inherent Kerr non-linearity of superconducting circuits and an additional two-photon drive, Kerr cat qubits manifested similar properties as their dissipative counterparts but arguably with lighter circuit designs. This path was also carried out, with theoretical proposals for bias-preserving gates and correction of phase-flip errors [Puri *et al.* 2020], as well as several experimental demonstrations [Grimm *et al.* 2020, Frattini *et al.* 2022]. However, for the reasons we thoroughly detailed in Chapter 2, Kerr cat qubits suffer from thermal-induced leakage that limits the scaling of their noise bias. This likely explains that experimental works have not yet been able to measure an exponential

noise bias with Kerr cat qubits, and suggests that more advanced stabilization protocols should be devised.

One such protocol was presented in this thesis, with a Hamiltonian confinement scheme based on a two-to-one photon exchange between a cavity and a buffer mode. While this scheme is qualitatively similar to Kerr Hamiltonian confinement, it is better suited to being combined with two-photon dissipation to retrieve the exponential noise scaling thanks to its bounded energy spectrum degeneracies, and implementation similarities to the dissipative scheme. We were thus able to demonstrate single- and two-qubit gates with improved fidelities while retaining the essential exponential suppression of bit-flip errors. We also showed a viable path towards the experimental realization of such protocols, with a thorough study of the various spurious effects that could limit the viability of the proposal.

In the coming years, it is expected that research on cat qubits will move towards multi-qubit studies, with the eventual realization of an error-correction code hosting a fully protected qubit with minimal hardware. But this objective would require some key advances to be achieved. First, two-qubit gates on cat qubits have yet to be demonstrated. While the theoretical proposal of [Guillaud & Mirrahimi 2019] yields a fully protected and bias-preserving CNOT gate, this thesis introduced an alternative approach to simplify the experimental CNOT apparatus by turning off target qubit stabilization during the gate. For a fast enough gate performed in a low-leakage system, this proposal also preserves the noise bias at the cost of a subsequent idling time to refocus leakage. Second, all single- and multi-qubit gates should feature fidelity errors below the threshold of the error-correction code. In this regard, standard gate designs based on the Zeno effect scale according to the square root of the main source of dissipation. This is likely to limit experimental advances towards threshold-compatible gates, and hence why Chapter 3 was fully dedicated to the design of high-fidelity Zeno gates for dissipative cat qubits.

Indeed, this thesis introduced four new designs of dissipative cat qubit gates that can be engineered with standard superconducting circuit tools. By noting that errors during Zeno gates are a result of entanglement between the cat qubit memory and its highly damped buffer mode, we have devised two gate designs based on retrieving the information leaked out to the environment. The first relied on an integrated photodetection apparatus, and the second on reservoir engineering such that information feedback was made autonomously. Then, we also presented two gate designs whose objectives were to minimize the information leaked out the environment altogether. This time, we showed how precise Hamiltonian engineering can yield high-fidelity gates while limiting the impact of second-order Zeno effects, but also presented how to circumvent the linear Zeno scaling altogether through dissipation-based

gate engineering.

Perhaps the most significant contribution of this work on Zeno gates was not the specific gate designs, but how the buffer mode was reconsidered into the big picture. Indeed, this ancillary mode was first introduced in [Leghtas *et al.* 2015] to mediate the two-to-one photon exchange coupling and eventually engineer the sought after two-photon dissipation through reservoir engineering. In the limit of a large damping of the buffer, its adiabatic elimination could be carried out to leave only the mode of interest: the memory. However, we now realize that non-trivial dynamics can take place between both cavities, which can be leveraged for gate engineering, but also for readout as we showed in [Réglade *et al.* 2023]. More than actual operations, it is likely that lowering the buffer mode damping rate could also be beneficial towards the realization of a cat qubit error-correction code by increasing effective confinement rates. In this case, it is ever more important to include the buffer mode in experiment models, since conditions for adiabatic elimination are not anymore met.

Towards the realization of an error-correction code for cat qubits, another important aspect lies in the optimization of experimental processes. Because error correction features a tipping point below which logical errors scale exponentially in the code distance, it is of utmost importance to exploit the full capabilities of every experimental apparatus. In Chapter 4 of this thesis, we introduced the adjoint state method for open quantum systems for the numerical optimization of any controllable operation on any quantum device. Through the adjoint state of the density matrix, the method makes it possible to compute the gradient of an arbitrary cost function in constant memory cost. This opens the way to a plethora of optimal control problems that were previously inaccessible, such as transmon readout as was demonstrated in this thesis. For this precise application in which dissipation is key to the problem statement, we have shown interpretable pulses that improved upon standard readout protocols by up to a factor two in speed. We however expect that the method will see applications in many quantum devices.

The adjoint state method is particularly relevant for the study of continuous-variable systems due to the inherent large dimension of their Hilbert spaces. For instance, there is little analytic intuition into the control of GKP grid states, especially with regards to multi-qubit operations, such that the ability to perform optimal control on such systems would be highly beneficial. In the context of cat qubits, one specific application to the adjoint state method that arose recently was the optimization of the readout pulse of [Réglade *et al.* 2023]. Indeed, the transmon-free readout protocol involved different pulses on memory and buffer, chained together with varying durations and amplitudes. While the short length of the protocol was key to minimize dissipation and maximize readout contrast, the precise controls

were chosen through trial and error due to a lack of analytical results. We can thus expect sizeable improvements from using the adjoint state method in this context.

In the broader sense, this thesis has been trying to improve upon the physical implementation of cat qubits, through more involved stabilization protocols, new gate designs, or even through the optimal control of operating pulses. While it is certain that progress has been made in this regard, we have also realized how improving these aspects often involves more complex experimental apparatuses and non-standard engineering. However, with the rapid progress of superconducting circuits, and in particular with how well experimental physicists can now handle the specificities of circuit Hamiltonians, we can expect that the proposals presented in this thesis may fall within standard control within the coming years. If so, we have analyzed and quantified the possible gains that could be realized with such designs. In the quest towards fault-tolerant quantum computation with cat states, it is likely that at least some of these complexities will be overcome.

More than cat qubit analysis, several conclusions of our study can be generalized in the context of reservoir engineering. First, these systems should be studied as a whole by considering the buffer mode used for dissipation engineering from start to finish. We have shown how non-trivial effects can arise from interaction with this buffer, and how it can be leveraged for gate or readout improvements. Stabilization mechanisms can then be adapted to one requirement or the other. Second, a strong robustness to leakage can be expected from reservoir engineering, as shown from our study of cat qubit confinement mechanisms. This result should generally hold for a large range of systems, since entropy is extracted from the mode of interest and transferred to the reservoir and into the environment.

Appendices

Invariance under joint phase conjugation

In this appendix, we show how the Lindblad master equation of a $Z(\theta)$ gate is invariant under a joint x-axis phase conjugation of the cat and buffer modes.

A.1 Phase conjugation superoperator

An x-axis phase conjugation consists in flipping the sign of quadrature \mathbf{x} while keeping quadrature \mathbf{p} unchanged [Cerf & Iblisdir 2001]. While this is an unphysical transformation — it does not preserve commutation relations —, it can be understood as a time-reversal operator. It is also a standard example of a superoperator which is positive but not completely positive [Wolf & Cirac 2008]. In our case, the symmetry of the master equation by joint phase conjugation implies that the logical bit information of the cat qubit (encoded in the x-axis) is also encoded in the x-axis of the buffer mode, and so in an exact manner.

Let us first define the x-axis phase conjugation superoperator on a single mode, which we denote as \mathcal{C} . By definition, this superoperator is such that $\mathcal{C} |x\rangle\langle x| = |-x\rangle\langle -x|$ and $\mathcal{C} |p\rangle\langle p| = |p\rangle\langle p|$ where $|x\rangle$ and $|p\rangle$ are quadrature eigenstates. By linearity, the identities $\mathcal{C}\mathbf{x} = -\mathbf{x}$ and $\mathcal{C}\mathbf{p} = \mathbf{p}$ directly follow since $\mathbf{x} = \int dx \cdot x |x\rangle\langle x|$ and $\mathbf{p} = \int dp \cdot p |p\rangle\langle p|$. More generally, we have that $\mathcal{C}f(\mathbf{x}) = f(-\mathbf{x})$ for any function f . Reinserting $|x\rangle = \int dp e^{-ipx} |p\rangle$ in this identity, we have that

$$\begin{aligned} & \int dx f(x) \iint dp dp' e^{-i(p-p')x} \cdot \mathcal{C} |p\rangle\langle p'| \\ &= \int dx f(x) \iint dp dp' e^{i(p-p')x} \cdot |p\rangle\langle p'| \end{aligned} \tag{A.1}$$

Rearranging the integrals and denoting $\tilde{f}(p)$ the Fourier transform of $f(x)$, (A.1) reads

$$\iint dp dp' \tilde{f}(p - p') \cdot (\mathcal{C} |p\rangle\langle p'| - |p'\rangle\langle p|) = 0 \tag{A.2}$$

Since this equation holds for any function \tilde{f} , then we have that $\mathcal{C} |p\rangle\langle p'| = |p'\rangle\langle p|$, and similarly for the other quadrature, $\mathcal{C} |x\rangle\langle x'| = |-x'\rangle\langle -x|$. In particular, notice the transposition in both of these relations, which accounts for the time-reversal property of the conjugation.

Before moving on, there is a last identity that will be useful for our upcoming derivation. For any density matrix ρ , we have

$$\begin{aligned}\mathcal{C}(\mathbf{x}\rho) &= \mathcal{C} \left(\iint dx dx' x \rho(x, x') |x\rangle\langle x'| \right) \\ &= \iint dx dx' x \rho(x, x') |-x'\rangle\langle -x| \\ &= \iint dx' dx (-x') \rho(-x', -x) |x\rangle\langle x'| \\ &= -\mathcal{C}(\rho)\mathbf{x}.\end{aligned}\tag{A.3}$$

Using that \mathcal{C}^2 yields the identity superoperator, we directly get the converse identity, $\mathcal{C}(\rho\mathbf{x}) = -\mathbf{x}\mathcal{C}(\rho)$. Similarly for the other quadrature, we have that $\mathcal{C}(\mathbf{p}\rho) = \mathcal{C}(\rho)\mathbf{p}$ and $\mathcal{C}(\rho\mathbf{p}) = \mathbf{p}\mathcal{C}(\rho)$.

A.2 Master equation invariance

Our goal is now to show that the joint phase conjugation of the cat and buffer modes is a symmetry of the $Z(\theta)$ gate dynamics. This involves showing that the joint phase conjugation superoperator $\mathcal{C}_{AB} = \mathcal{C}_A \otimes \mathcal{C}_B$ commutes with the Lindblad superoperator \mathcal{L} , where

$$\mathcal{L} = \mathcal{L}_Z + \mathcal{L}_{AB} + \mathcal{L}_{1,A} + \mathcal{L}_{1,B}\tag{A.4}$$

where $\mathcal{L}_Z = -i[\mathbf{H}_Z, \cdot]$, $\mathcal{L}_{AB} = -i[\mathbf{H}_{AB}, \cdot]$, $\mathcal{L}_{1,A} = \kappa_a \mathcal{D}[\mathbf{a}]$ and $\mathcal{L}_{1,B} = \kappa_b \mathcal{D}[\mathbf{b}]$. By linearity, it is sufficient to show the commutation relation for each term in the Lindblad superoperator. For the single-mode cavity drive with Hamiltonian $\mathbf{H}_Z = \varepsilon_Z \mathbf{x}_a$, we get

$$\begin{aligned}\mathcal{C}(\mathcal{L}_Z(\rho)) &= -i\varepsilon_Z \mathcal{C}(\mathbf{x}_a \rho - \rho \mathbf{x}_a) \\ &= -i\varepsilon_Z (-\mathcal{C}(\rho)\mathbf{x}_a + \mathbf{x}_a \mathcal{C}(\rho)) \\ &= \mathcal{L}_Z(\mathcal{C}(\rho)).\end{aligned}\tag{A.5}$$

For the two-photon exchange term, the Hamiltonian reads $\mathbf{H}_{AB} = g_2(\mathbf{a}^2 - \alpha^2)\mathbf{b}^\dagger + \text{h.c.}$. Rewriting the Hamiltonian in terms of quadrature operators yields

$$\mathbf{H}_{AB}/g_2 = (\mathbf{x}_a^2 - \mathbf{p}_a^2)\frac{\mathbf{x}_b}{4} + (\mathbf{p}_a \mathbf{x}_a + \mathbf{x}_a \mathbf{p}_a)\frac{\mathbf{p}_b}{4} - \alpha^2 \mathbf{x}_b.\tag{A.6}$$

Since this Hamiltonian features only terms with an odd number of quadrature \mathbf{x} and even number of quadrature \mathbf{p} operators, using the identities (A.3) and the related ones yields an overall minus sign and a transposition, from which we infer

$$\mathcal{C}(\mathcal{L}_{AB}(\boldsymbol{\rho})) = \mathcal{L}_{AB}(\mathcal{C}(\boldsymbol{\rho})). \quad (\text{A.7})$$

Finally, for the single-photon loss operators on either mode $\mathcal{L}_1 = \mathcal{D}[\mathbf{x} + i\mathbf{p}]$, and again using the (A.3) identities,

$$\begin{aligned} \mathcal{C}(\mathcal{L}_1(\boldsymbol{\rho})) &= \mathcal{C}((\mathbf{x} + i\mathbf{p}) \cdot \boldsymbol{\rho} \cdot (\mathbf{x} - i\mathbf{p})) \\ &\quad - \mathcal{C}((\mathbf{x}^2 + \mathbf{p}^2) \cdot \boldsymbol{\rho})/2 - \mathcal{C}(\boldsymbol{\rho} \cdot (\mathbf{x}^2 + \mathbf{p}^2))/2 \\ &= (\mathbf{x} + i\mathbf{p}) \cdot \mathcal{C}(\boldsymbol{\rho}) \cdot (\mathbf{x} - i\mathbf{p}) \\ &\quad - \mathcal{C}(\boldsymbol{\rho}) \cdot (\mathbf{x}^2 + \mathbf{p}^2)/2 - (\mathbf{x}^2 + \mathbf{p}^2) \cdot \mathcal{C}(\boldsymbol{\rho})/2 \\ &= \mathcal{L}_1(\mathcal{C}(\boldsymbol{\rho})) \end{aligned} \quad (\text{A.8})$$

By linearity, the required result is shown, $\mathcal{C}\mathcal{L} - \mathcal{L}\mathcal{C} = 0$. It shows the invariance of the $Z(\theta)$ gate master equation under a joint x-axis phase conjugation of both cat and buffer modes.

Bibliography

- [Abdelhafez *et al.* 2019] Mohamed Abdelhafez, David I Schuster and Jens Koch. [Gradient-based optimal control of open quantum systems using quantum trajectories and automatic differentiation](#). Physical Review A, 2019. (Cited on page 126.)
- [AI 2023] Google Quantum AI. [Suppressing quantum errors by scaling a surface code logical qubit](#). Nature, vol. 614, no. 7949, pages 676–681, 2023. (Cited on page 8.)
- [Albert *et al.* 2019] Victor V Albert, Shantanu O Mundhada, Alexander Grimm, Steven Touzard, Michel H Devoret and Liang Jiang. [Pair-cat codes: autonomous error-correction with low-order nonlinearity](#). Quantum Science and Technology, vol. 4, no. 3, page 035007, 2019. (Cited on pages 13 and 100.)
- [Albertinale *et al.* 2021] Emanuele Albertinale, Léo Balembois, Eric Billaud, Vishal Ranjan, Daniel Flanigan, Thomas Schenkel, Daniel Estève, Denis Vion, Patrice Bertet and Emmanuel Flurin. [Detecting spins by their fluorescence with a microwave photon counter](#). Nature, vol. 600, no. 7889, pages 434–438, 2021. (Cited on page 98.)
- [Alipour *et al.* 2020] Sahar Alipour, Aurelia Chenu, Ali T. Rezakhani and Adolfo del Campo. [Shortcuts to Adiabaticity in Driven Open Quantum Systems: Balanced Gain and Loss and Non-Markovian Evolution](#). Quantum, vol. 4, page 336, 2020. (Cited on page 67.)
- [Arute *et al.* 2019] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends *et al.* [Quantum supremacy using a programmable superconducting processor](#). Nature, vol. 574, no. 7779, pages 505–510, 2019. (Cited on pages 2, 8 and 10.)
- [Ataides *et al.* 2021] J. Pablo Bonilla Ataides, David K. Tuckett, Stephen D. Bartlett, Steven T. Flammia and Benjamin J. Brown. [The XXXX surface code](#). Nature Communications, vol. 12, no. 1, page 2172, 2021. (Cited on page 18.)
- [Azouit *et al.* 2016] Rémi Azouit, Alain Sarlette and Pierre Rouchon. [Well-posedness and convergence of the lindblad master equation for a quantum harmonic oscillator with multi-photon drive and damping](#). ESAIM: COCV, vol. 22, no. 4, pages 1353–1369, 2016. (Cited on page 42.)

- [Azouit *et al.* 2017] R Azouit, F Chittaro, A Sarlette and P Rouchon. [Towards generic adiabatic elimination for bipartite open quantum systems](#). Quantum Science and Technology, vol. 2, no. 4, page 044011, 2017. (Cited on pages 26, 52, 62, 101 and 123.)
- [Baydin *et al.* 2018] Atilim Gunes Baydin, Barak A Pearlmutter, Alexey Andreyevich Radul and Jeffrey Mark Siskind. [Automatic differentiation in machine learning: a survey](#). Journal of Machine Learning Research, vol. 18, pages 1–43, 2018. (Cited on pages 126, 128 and 129.)
- [Bellman 1953] Richard Bellman. Stability theory of differential equations. Dover Publications, 1953. (Cited on page 25.)
- [Bennett *et al.* 1993] Charles H Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres and William K Wootters. [Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels](#). Physical review letters, vol. 70, no. 13, page 1895, 1993. (Cited on page 2.)
- [Berdou *et al.* 2022] C Berdou, A Murani, U Reglade, WC Smith, M Villiers, J Palomo, M Rosticher, A Denis, P Morfin, M Delbecq *et al.* [One hundred second bit-flip time in a two-photon dissipative oscillator](#). arXiv, 2022. (Cited on page 110.)
- [Blais *et al.* 2007] Alexandre Blais, Jay Gambetta, Andreas Wallraff, David I Schuster, Steven M Girvin, Michel H Devoret and Robert J Schoelkopf. [Quantum-information processing with circuit quantum electrodynamics](#). Physical Review A, vol. 75, no. 3, page 032329, 2007. (Cited on page 7.)
- [Blais *et al.* 2021] Alexandre Blais, Arne L. Grimsmo, S. M. Girvin and Andreas Wallraff. [Circuit quantum electrodynamics](#). Rev. Mod. Phys., vol. 93, page 025005, 05 2021. (Cited on pages 4, 5, 7, 9 and 136.)
- [Boutin *et al.* 2017] Samuel Boutin, Christian Kraglund Andersen, Jayameenakshi Venkatraman, Andrew J Ferris and Alexandre Blais. [Resonator reset in circuit QED by optimal control for large open quantum systems](#). Physical Review A, vol. 96, no. 4, page 042315, 2017. (Cited on page 126.)
- [Breuer *et al.* 2007] Heinz-Peter Breuer, Francesco Petruccione *et al.* The theory of open quantum systems. Oxford University Press, 2007. (Cited on page 86.)
- [Bultink *et al.* 2018] Cornelis Christiaan Bultink, Brian Tarasinski, Niels Haandbæk, Stefano Poletto, Nadia Haider, DJ Michalak, Alessandro

- Bruno and Leonardo DiCarlo. [General method for extracting the quantum efficiency of dispersive qubit readout in circuit QED](#). *Applied Physics Letters*, vol. 112, no. 9, 2018. (Cited on page 137.)
- [Campagne-Ibarcq *et al.* 2020] Philippe Campagne-Ibarcq, Alec Eickbusch, Steven Touzard, Evan Zalys-Geller, Nicholas E Frattini, Volodymyr V Sivak, Philip Reinhold, Shruti Puri, Shyam Shankar, Robert J Schoelkopf, Luigi Frunzio, Mazyar Mirrahimi and Michel H. Devoret. [Quantum error correction of a qubit encoded in grid states of an oscillator](#). *Nature*, vol. 584, no. 7821, pages 368–372, 2020. (Cited on page 11.)
- [Cerf & Iblisdir 2001] NJ Cerf and Sofyan Iblisdir. [Phase conjugation of continuous quantum variables](#). *Physical Review A*, vol. 64, no. 3, page 032307, 2001. (Cited on page 151.)
- [Chamberland *et al.* 2022] Christopher Chamberland, Kyungjoo Noh, Patricio Arrangoiz-Arriola, Earl T Campbell, Connor T Hann, Joseph Iverson, Harald Putterman, Thomas C Bohdanowicz, Steven T Flammia and Andrew Keller. [Building a fault-tolerant quantum computer using concatenated cat codes](#). *PRX Quantum*, vol. 3, no. 1, page 010329, 2022. (Cited on pages 12, 16, 31, 35, 46, 58, 59, 60, 61, 62, 63, 67, 71, 85, 87, 97, 101, 104 and 109.)
- [Chen *et al.* 2018] Ricky T. Q. Chen, Yulia Rubanova, Jesse Bettencourt and David K Duvenaud. [Neural Ordinary Differential Equations](#). In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31. Curran Associates, Inc., 2018. (Cited on pages 132, 133 and 134.)
- [Chen *et al.* 2023] Liangyu Chen, Hang-Xi Li, Yong Lu, Christopher W Warren, Christian J Križan, Sandoko Kosen, Marcus Rommel, Shahnawaz Ahmed, Amr Osman, Janka Biznárová *et al.* [Transmon qubit readout fidelity at the threshold for quantum error correction without a quantum-limited amplifier](#). *npj Quantum Information*, vol. 9, no. 1, page 26, 2023. (Cited on page 141.)
- [Clarke & Wilhelm 2008] John Clarke and Frank K Wilhelm. [Superconducting quantum bits](#). *Nature*, vol. 453, no. 7198, pages 1031–1042, 2008. (Cited on page 7.)
- [Cochrane *et al.* 1999] P. T. Cochrane, G. J. Milburn and W. J. Munro. [Macroscopically distinct quantum-superposition states as a bosonic code for amplitude damping](#). *Physical Review A*, vol. 59, no. 4, pages 2631–2634, 1999. (Cited on pages 12 and 13.)

- [Cohen *et al.* 2022] Joachim Cohen, Alexandru Petrescu, Ross Shillito and Alexandre Blais. [Reminiscence of classical chaos in driven transmons](#). arXiv preprint arXiv:2207.09361, 2022. (Cited on page 136.)
- [Cohen 2017] Joachim Cohen. [Autonomous quantum error correction with superconducting qubits](#). PhD thesis, Université Paris Sciences et Lettres, 2017. (Cited on pages 40 and 118.)
- [Dassonneville *et al.* 2020] Rémy Dassonneville, Réouven Assouly, Théau Peronnin, Pierre Rouchon and Benjamin Huard. [Number-resolved photocounter for propagating microwave mode](#). Physical Review Applied, vol. 14, no. 4, page 044022, 2020. (Cited on page 98.)
- [De Neeve *et al.* 2022] Brennan De Neeve, Thanh-Long Nguyen, Tanja Behrle and Jonathan P Home. [Error correction of a logical grid state qubit by dissipative pumping](#). Nature Physics, vol. 18, no. 3, pages 296–300, 2022. (Cited on page 11.)
- [Deutsch 1985] David Deutsch. [Quantum theory, the Church–Turing principle and the universal quantum computer](#). Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences, vol. 400, no. 1818, pages 97–117, 1985. (Cited on page 2.)
- [Didier *et al.* 2015] Nicolas Didier, Jérôme Bourassa and Alexandre Blais. [Fast Quantum Nondemolition Readout by Parametric Modulation of Longitudinal Qubit-Oscillator Interaction](#). Phys. Rev. Lett., vol. 115, page 203601, Nov 2015. (Cited on page 140.)
- [Douçot & Vidal 2002] Benoit Douçot and Julien Vidal. [Pairing of Cooper pairs in a fully frustrated Josephson-junction chain](#). Physical review letters, vol. 88, no. 22, page 227005, 2002. (Cited on page 8.)
- [Felicetti *et al.* 2018a] S. Felicetti, D. Z. Rossatto, E. Rico, E. Solano and P. Forn-Díaz. [Two-photon quantum Rabi model with superconducting circuits](#). Physical Review A, vol. 97, page 013851, 2018. (Cited on page 74.)
- [Felicetti *et al.* 2018b] Simone Felicetti, Myung-Joong Hwang and Alexandre Le Boité. [Ultrastrong-coupling regime of nondipolar light-matter interactions](#). Physical Review A, vol. 98, page 053859, 2018. (Cited on page 74.)
- [Feynman 1982] Richard P Feynman. [Simulating physics with computers](#). In Feynman and computation, pages 133–153. CRC Press, 1982. (Cited on page 2.)

- [Forni *et al.* 2018] Paolo Forni, Alain Sarlette, Thibault Capelle, Emmanuel Flurin, Samuel Deléglise and Pierre Rouchon. [Adiabatic elimination for multi-partite open quantum systems with non-trivial zero-order dynamics](#). 2018 IEEE Conference on Decision and Control (CDC), pages 6614–6619, 2018. (Cited on pages 26 and 123.)
- [Fowler *et al.* 2012] Austin G. Fowler, Matteo Mariantoni, John M. Martinis and Andrew N. Cleland. [Surface codes: Towards practical large-scale quantum computation](#). Phys. Rev. A, vol. 86, page 032324, 09 2012. (Cited on pages 18 and 66.)
- [Frattini *et al.* 2017] NE Frattini, U Vool, S Shankar, A Narla, KM Sliwa and MH Devoret. [3-wave mixing Josephson dipole element](#). Applied Physics Letters, vol. 110, no. 22, 2017. (Cited on pages 21 and 22.)
- [Frattini *et al.* 2022] Nicholas E Frattini, Rodrigo G Cortiñas, Jayameenakshi Venkatraman, Xu Xiao, Qile Su, Chan U Lei, Benjamin J Chapman, Vidul R Joshi, SM Girvin, Robert J Schoelkopf *et al.* [The squeezed Kerr oscillator: Spectral kissing and phase-flip robustness](#). arXiv, 2022. (Cited on pages 15, 21 and 145.)
- [Gambetta *et al.* 2007] Jay Gambetta, W A Braff, A Wallraff, S M Girvin and R J Schoelkopf. [Protocols for optimal readout of qubits using a continuous quantum nondemolition measurement](#). Physical Review A, 2007. (Cited on page 138.)
- [Gautier *et al.* 2022] Ronan Gautier, Alain Sarlette and Mazyar Mirrahimi. [Combined Dissipative and Hamiltonian Confinement of Cat Qubits](#). PRX Quantum, vol. 3, page 020339, 05 2022. (Cited on pages 39 and 96.)
- [Gautier *et al.* 2023a] Ronan Gautier, Élie Gérois and Alexandre Blais. Optimal control of transmon readout and reset. In preparation, 2023. (Cited on page 125.)
- [Gautier *et al.* 2023b] Ronan Gautier, Mazyar Mirrahimi and Alain Sarlette. [Designing high-fidelity Zeno gates for dissipative cat qubits](#). arXiv, 2023. (Cited on page 83.)
- [Gertler *et al.* 2023] Jeffrey M Gertler, Sean van Geldern, Shruti Shirol, Liang Jiang and Chen Wang. [Experimental Realization and Characterization of Stabilized Pair Coherent States](#). PRX Quantum, vol. 4, page 020319, 2023. (Cited on pages 100 and 101.)
- [Gottesman *et al.* 2001] Daniel Gottesman, Alexei Kitaev and John Preskill. [Encoding a qubit in an oscillator](#). Physical Review A, vol. 64, no. 1, page 012310, 2001. (Cited on page 11.)

- [Gouzien *et al.* 2023] Élie Gouzien, Diego Ruiz, Francois-Marie Le Régent, Jérémie Guillaud and Nicolas Sangouard. [Performance Analysis of a Repetition Cat Code Architecture: Computing 256-bit Elliptic Curve Logarithm in 9 Hours with 126 133 Cat Qubits](#). Phys. Rev. Lett., vol. 131, page 040602, Jul 2023. (Cited on page 18.)
- [Grimm *et al.* 2020] A. Grimm, N. E. Frattini, S. Puri, S. O. Mundhada, S. Touzard, M. Mirrahimi, S. M. Girvin, S. Shankar and M. H. Devoret. [Stabilization and operation of a Kerr-cat qubit](#). Nature, vol. 584, no. 7820, pages 205–209, 2020. (Cited on pages 21, 50, 59 and 145.)
- [Guillaud & Mirrahimi 2019] Jérémie Guillaud and Mazyar Mirrahimi. [Repetition Cat Qubits for Fault-Tolerant Quantum Computation](#). Physical Review X, vol. 9, no. 4, page 041053, 2019. (Cited on pages 12, 18, 19, 31, 33, 35, 40, 58, 59, 63, 66, 67, 70, 72, 84, 88, 109, 145 and 146.)
- [Hamilton 1833] William Rowan Hamilton. On a general method of expressing the paths of light, & of the planets, by the coefficients of a characteristic function, volume 1. Dublin University Review and Quarterly Magazine, 1833. (Cited on page 22.)
- [Haroche & Raimond 2006] Serge Haroche and Jean-Michel Raimond. Exploring the quantum: Atoms, cavities, and photons. Oxford University Press, 08 2006. (Cited on page 9.)
- [Hu *et al.* 2019] Ling Hu, Yuwei Ma, Weizhou Cai, Xianghao Mu, Yuan Xu, Weiting Wang, Yukai Wu, Haiyan Wang, YP Song, C-L Zou *et al.* [Quantum error correction and universal gate set operation on a binomial bosonic logical qubit](#). Nature Physics, vol. 15, no. 5, pages 503–508, 2019. (Cited on page 11.)
- [Ioffe & Feigel'man 2002] LB Ioffe and MV Feigel'man. [Possible realization of an ideal quantum computer in Josephson junction array](#). Physical Review B, vol. 66, no. 22, page 224503, 2002. (Cited on page 8.)
- [Jaynes & Cummings 1963] E.T. Jaynes and F.W. Cummings. [Comparison of quantum and semiclassical radiation theories with application to the beam maser](#). Proceedings of the IEEE, vol. 51, no. 1, pages 89–109, 1963. (Cited on page 10.)
- [Jirari 2009] Hamza Jirari. [Optimal control approach to dynamical suppression of decoherence of a qubit](#). Europhysics Letters, vol. 87, no. 4, page 40003, 2009. (Cited on page 126.)
- [Josephson 1962] B. D. Josephson. [Possible new effects in superconductive tunnelling](#). Physics Letters, vol. 1, no. 7, pages 251–253, 1962. (Cited on page 6.)

- [Khaneja *et al.* 2005] Navin Khaneja, Timo Reiss, Cindie Kehlet, Thomas Schulte-Herbrüggen and Steffen J. Glaser. [Optimal control of coupled spin dynamics: design of NMR pulse sequences by gradient ascent algorithms](#). Journal of Magnetic Resonance, vol. 172, no. 2, pages 296–305, 2005. (Cited on page 126.)
- [Kim *et al.* 2023] Youngseok Kim, Andrew Eddins, Sajant Anand, Ken Xuan Wei, Ewout Van Den Berg, Sami Rosenblatt, Hasan Nayfeh, Yantao Wu, Michael Zaletel, Kristan Temme *et al.* [Evidence for the utility of quantum computing before fault tolerance](#). Nature, vol. 618, no. 7965, pages 500–505, 2023. (Cited on page 8.)
- [Kingma & Ba 2014] Diederik P Kingma and Jimmy Ba. [Adam: A method for stochastic optimization](#). arXiv, 2014. (Cited on pages 127 and 139.)
- [Kitaev 2006] Alexei Kitaev. [Protected qubit based on a superconducting current mirror](#). arXiv, 2006. (Cited on page 8.)
- [Knill *et al.* 2000] Emanuel Knill, Raymond Laflamme and Lorenza Viola. [Theory of quantum error correction for general noise](#). Physical Review Letters, vol. 84, no. 11, page 2525, 2000. (Cited on page 18.)
- [Knill 2005] Emanuel Knill. [Quantum computing with realistically noisy devices](#). Nature, vol. 434, no. 7029, pages 39–44, 2005. (Cited on page 3.)
- [Koch *et al.* 2007] Jens Koch, Terri M. Yu, Jay Gambetta, A. A. Houck, D. I. Schuster, J. Majer, Alexandre Blais, M. H. Devoret, S. M. Girvin and R. J. Schoelkopf. [Charge-insensitive qubit design derived from the Cooper pair box](#). Phys. Rev. A, vol. 76, page 042319, 10 2007. (Cited on pages 3, 7 and 127.)
- [Krotov 1995] Vadim Krotov. Global methods in optimal control theory, volume 195. CRC Press, 1995. (Cited on page 126.)
- [Kumar *et al.* 2016] KS Kumar, Antti Vepsäläinen, S Danilin and GS Paraoanu. [Stimulated Raman adiabatic passage in a three-level superconducting circuit](#). Nature communications, vol. 7, no. 1, pages 1–6, 2016. (Cited on page 116.)
- [Leghtas *et al.* 2013] Zaki Leghtas, Gerhard Kirchmair, Brian Vlastakis, Robert J. Schoelkopf, Michel H. Devoret and Mazyar Mirrahimi. [Hardware-Efficient Autonomous Quantum Memory Protection](#). Phys. Rev. Lett., vol. 111, page 120501, Sep 2013. (Cited on page 145.)
- [Leghtas *et al.* 2015] Z. Leghtas, S. Touzard, I. M. Pop, A. Kou, B. Vlastakis, A. Petrenko, K. M. Sliwa, A. Narla, S. Shankar, M. J. Hatridge,

- M. Reagor, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi and M. H. Devoret. [Confining the state of light to a quantum manifold by engineered two-photon loss](#). *Science*, vol. 347, no. 6224, pages 853–857, 2015. (Cited on pages 27, 52, 145 and 147.)
- [Lescanne *et al.* 2020a] Raphaël Lescanne, Samuel Deléglise, Emanuele Albertinale, Ulysse Réglade, Thibault Capelle, Edouard Ivanov, Thibaut Jacqmin, Zaki Leghtas and Emmanuel Flurin. [Irreversible qubit-photon coupling for the detection of itinerant microwave photons](#). *Physical Review X*, vol. 10, no. 2, page 021038, 2020. (Cited on page 98.)
- [Lescanne *et al.* 2020b] Raphaël Lescanne, Marius Villiers, Théau Peronnin, Alain Sarlette, Matthieu Delbecq, Benjamin Huard, Takis Kontos, Mazyar Mirrahimi and Zaki Leghtas. [Exponential suppression of bit-flips in a qubit encoded in an oscillator](#). *Nature Physics*, vol. 16, no. 5, pages 509–513, 2020. (Cited on pages 22, 27, 40, 57, 72, 73, 74, 80, 110, 120 and 145.)
- [Leung *et al.* 2017] Nelson Leung, Mohamed Abdelhafez, Jens Koch and David Schuster. [Speedup for quantum optimal control from automatic differentiation based on graphics processing units](#). *Physical Review A*, vol. 95, no. 4, page 042318, 2017. (Cited on pages 126, 128, 130 and 133.)
- [Lledó *et al.* 2022] Cristóbal Lledó, Rémy Dassonneville, Adrien Moulinas, Joachim Cohen, Ross Shillito, Audrey Bienfait, Benjamin Huard and Alexandre Blais. [Cloaking a qubit in a cavity](#). *arXiv*, 2022. (Cited on page 141.)
- [Lu *et al.* 2023] Yunwei Lu, Sandeep Joshi, Vinh San Dinh and Jens Koch. [Optimal control of large quantum systems: assessing memory and runtime performance of GRAPE](#). *arXiv*, 2023. (Cited on pages 130 and 133.)
- [Mallet *et al.* 2009] François Mallet, Florian R Ong, Agustin Palacios-Laloy, François Nguyen, Patrice Bertet, Denis Vion and Daniel Esteve. [Single-shot qubit readout in circuit quantum electrodynamics](#). *Nature Physics*, vol. 5, no. 11, pages 791–795, 2009. (Cited on page 127.)
- [Manucharyan *et al.* 2009] Vladimir E Manucharyan, Jens Koch, Leonid I Glazman and Michel H Devoret. [Fluxonium: Single cooper-pair circuit free of charge offsets](#). *Science*, vol. 326, no. 5949, pages 113–116, 2009. (Cited on page 8.)

- [Marquet *et al.* 2023] Antoine Marquet, Antoine Essig, Joachim Cohen, Nathanaël Cottet, Anil Murani, Emanuele Abertinale, Simon Dupouy, Audrey Bienfait, Théau Peronnin, Sébastien Jezouin *et al.* [Autoparametric resonance extending the bit-flip time of a cat qubit up to 0.3 s.](#) arXiv, 2023. (Cited on pages 26 and 28.)
- [Marte *et al.* 1991] P. Marte, P. Zoller and J. L. Hall. [Coherent atomic mirrors and beam splitters by adiabatic passage in multilevel systems.](#) Phys. Rev. A, vol. 44, pages R4118–R4121, Oct 1991. (Cited on page 126.)
- [Michael *et al.* 2016] Marios H. Michael, Matti Silveri, R. T. Brierley, Victor V. Albert, Juha Salmilehto, Liang Jiang and S. M. Girvin. [New Class of Quantum Error-Correcting Codes for a Bosonic Mode.](#) Physical Review X, vol. 6, no. 3, page 031006, 2016. (Cited on page 11.)
- [Mirrahimi *et al.* 2014] Mazyar Mirrahimi, Zaki Leghtas, Victor V Albert, Steven Touzard, Robert J Schoelkopf, Liang Jiang and Michel H Devoret. [Dynamically protected cat-qubits: a new paradigm for universal quantum computation.](#) New Journal of Physics, vol. 16, no. 4, page 045014, 2014. (Cited on pages 12, 13, 15, 16, 19, 24, 25, 31, 40, 42, 59, 88, 89 and 145.)
- [Motzoi *et al.* 2009] F. Motzoi, J. M. Gambetta, P. Rebentrost and F. K. Wilhelm. [Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits.](#) Phys. Rev. Lett., vol. 103, page 110501, Sep 2009. (Cited on pages 126 and 142.)
- [Motzoi *et al.* 2011] F Motzoi, Jay M Gambetta, ST Merkel and FK Wilhelm. [Optimal control methods for rapidly time-varying Hamiltonians.](#) Physical Review A, vol. 84, no. 2, page 022307, 2011. (Cited on page 138.)
- [Muñoz-Arias *et al.* 2023] Manuel H Muñoz-Arias, Cristóbal Lledó and Alexandre Blais. [Qubit readouts enabled by qubit cloaking.](#) arXiv, 2023. (Cited on page 141.)
- [Nakamura *et al.* 1999] Yasunobu Nakamura, Yu A Pashkin and JS Tsai. [Coherent control of macroscopic quantum states in a single-Cooper-pair box.](#) Nature, vol. 398, no. 6730, pages 786–788, 1999. (Cited on page 3.)
- [Nenciu 1993] G. Nenciu. [Linear adiabatic theory. Exponential estimates.](#) Communications in Mathematical Physics, vol. 152, pages 479–496, 1993. (Cited on pages 20 and 60.)
- [Paik *et al.* 2011] Hanhee Paik, David I Schuster, Lev S Bishop, Gerhard Kirchmair, Gianluigi Catelani, Adam P Sears, BR Johnson, MJ Reagor, Luigi Frunzio, Leonid I Glazman *et al.* [Observation of](#)

- high coherence in Josephson junction qubits measured in a three-dimensional circuit QED architecture. *Physical Review Letters*, vol. 107, no. 24, page 240501, 2011. (Cited on page 8.)
- [Peirce *et al.* 1988] Anthony P. Peirce, Mohammed A. Dahleh and Herschel Rabitz. [Optimal control of quantum-mechanical systems: Existence, numerical approximation, and applications](#). *Phys. Rev. A*, vol. 37, pages 4950–4964, Jun 1988. (Cited on page 126.)
- [Pellizzari *et al.* 1995] T. Pellizzari, S. A. Gardiner, J. I. Cirac and P. Zoller. [Decoherence, Continuous Observation, and Quantum Computing: A Cavity QED Model](#). *Phys. Rev. Lett.*, vol. 75, pages 3788–3791, 11 1995. (Cited on page 3.)
- [Pontryagin *et al.* 1962] LS Pontryagin, VG Boltyanskii, RV Gamkrelidze and EF Mishechenko. The mathematical theory of optimal processes. Routledge, 01 1962. (Cited on pages 126 and 130.)
- [Preskill 1998] John Preskill. [Reliable quantum computers](#). *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, vol. 454, no. 1969, pages 385–410, 1998. (Cited on page 3.)
- [Purcell 1946] Edward Mills Purcell. [Spontaneous emission probabilities at radio frequencies](#). *Physical Review*, vol. 69, pages 674–674, Jun 1946. (Cited on page 136.)
- [Puri *et al.* 2017] S. Puri, S. Boutin and A. Blais. [Engineering the quantum states of light in a Kerr-nonlinear resonator by two-photon driving](#). *npj Quantum Information*, vol. 3, no. 1, page 18, 2017. (Cited on pages 19, 20, 40, 42, 59 and 145.)
- [Puri *et al.* 2019] Shruti Puri, Alexander Grimm, Philippe Campagne-Ibarcq, Alec Eickbusch, Kyungjoo Noh, Gabrielle Roberts, Liang Jiang, Mazhar Mirrahimi, Michel H. Devoret and S. M. Girvin. [Stabilized Cat in a Driven Nonlinear Cavity: A Fault-Tolerant Error Syndrome Detector](#). *Physical Review X*, vol. 9, page 041009, 2019. (Cited on page 68.)
- [Puri *et al.* 2020] S. Puri, L. St-Jean, J. A. Gross, A. Grimm, N. E. Frattini, P. S. Iyer, A. Krishna, S. Touzard, L. Jiang, A. Blais, S. T. Flammia and S. M. Girvin. [Bias-preserving gates with stabilized cat qubits](#). *Science Advances*, vol. 6, no. 34, page eaay5901, 2020. (Cited on pages 19, 31, 40, 50, 58, 67, 70 and 145.)
- [Putterman *et al.* 2021] Harald Putterman, Joseph Iverson, Qian Xu, Liang Jiang, Oskar Painter, Fernando G. S. L. Brandão and Kyungjoo Noh.

- [Stabilizing a Bosonic Qubit Using Colored Dissipation](#). Physical Review Letters, vol. 128, no. 11, page 110502, 2021. (Cited on pages 40, 42, 46 and 50.)
- [Ralph *et al.* 2003] T. C. Ralph, A. Gilchrist, G. J. Milburn, W. J. Munro and S. Glancy. [Quantum computation with optical coherent states](#). Physical Review A, vol. 68, page 042319, 2003. (Cited on pages 12 and 13.)
- [Reiter & Sørensen 2012] Florentin Reiter and Anders S Sørensen. [Effective operator formalism for open quantum systems](#). Physical Review A, vol. 85, no. 3, page 032111, 2012. (Cited on pages 101 and 105.)
- [Robbins & Monro 1951] Herbert Robbins and Sutton Monro. [A Stochastic Approximation Method](#). The Annals of Mathematical Statistics, vol. 22, no. 3, pages 400 – 407, 1951. (Cited on page 127.)
- [Ruiz *et al.* 2023] Diego Ruiz, Ronan Gautier, Jérémie Guillaud and Mazyar Mirrahimi. [Two-photon driven Kerr quantum oscillator with multiple spectral degeneracies](#). Phys. Rev. A, vol. 107, page 042407, 04 2023. (Cited on page 37.)
- [Réglade *et al.* 2023] Ulysse Réglade, Adrien Bocquet, Ronan Gautier, Antoine Marquet, Emanuele Albertinale, Natalia Pankratova, Mattis Hallen, Felix Rautschke, Lev-Arcady Sellem, Pierre Rouchon, Alain Sarlette, Mazyar Mirrahimi, Philippe Campagne-Ibarcq, Raphael Lescanne, Sebastien Jezouin and Zaki Leghtas. [Quantum control of dynamical states with switching times exceeding ten seconds](#). arXiv, 2023. (Cited on pages 15, 26, 28, 37, 84, 120, 145 and 147.)
- [Schlegel *et al.* 2022] David S. Schlegel, Fabrizio Minganti and Vincenzo Savona. [Quantum error correction using squeezed Schrödinger cat states](#). Phys. Rev. A, vol. 106, page 022431, 08 2022. (Cited on page 102.)
- [Schreier *et al.* 2008] Joseph A Schreier, Andrew A Houck, Jens Koch, David I Schuster, Bradley R Johnson, Jerry M Chow, Jay M Gambetta, J Majer, Luigi Frunzio, Michel H Devoret *et al.* [Suppressing charge noise decoherence in superconducting charge qubits](#). Physical Review B, vol. 77, no. 18, page 180502, 2008. (Cited on page 8.)
- [Schulte-Herbrüggen *et al.* 2011] T Schulte-Herbrüggen, A Spörl, N Khaneja and SJ Glaser. [Optimal control for generating quantum gates in open dissipative systems](#). Journal of Physics B: Atomic, Molecular and Optical Physics, vol. 44, no. 15, page 154013, 2011. (Cited on page 126.)

- [Sete *et al.* 2015] Eyob A. Sete, John M. Martinis and Alexander N. Korotkov. [Quantum theory of a bandpass Purcell filter for qubit readout](#). Physical Review A, vol. 92, no. 1, page 012325, 2015. (Cited on page 136.)
- [Shillito *et al.* 2022] Ross Shillito, Alexandru Petrescu, Joachim Cohen, Jackson Beall, Markus Hauru, Martin Ganahl, Adam GM Lewis, Guifre Vidal and Alexandre Blais. [Dynamics of transmon ionization](#). Physical Review Applied, vol. 18, no. 3, page 034031, 2022. (Cited on pages 27 and 136.)
- [Shor 1994] Peter W Shor. [Algorithms for quantum computation: discrete logarithms and factoring](#). In Proceedings 35th annual symposium on foundations of computer science, pages 124–134. Ieee, 1994. (Cited on page 2.)
- [Shor 1996] P. W. Shor. [Fault-tolerant quantum computation](#). In Proceedings of 37th Conference on Foundations of Computer Science. IEEE Comput. Soc. Press, 1996. (Cited on page 84.)
- [Sivak *et al.* 2023] VV Sivak, A Eickbusch, B Royer, S Singh, I Tsioutsios, S Ganjam, A Miano, BL Brock, AZ Ding, L Frunzio *et al.* [Real-time quantum error correction beyond break-even](#). Nature, vol. 616, no. 7955, page 50–55, 2023. (Cited on page 11.)
- [Steck 2007] Daniel A Steck. Quantum and atom optics. available online at <http://steck.us/teaching>, 2007. (Cited on page 91.)
- [Sunada *et al.* 2022] Y. Sunada, S. Kono, J. Ilves, S. Tamate, T. Sugiyama, Y. Tabuchi and Y. Nakamura. [Fast Readout and Reset of a Superconducting Qubit Coupled to a Resonator with an Intrinsic Purcell Filter](#). Phys. Rev. Appl., vol. 17, page 044016, Apr 2022. (Cited on page 139.)
- [Swiadek *et al.* 2023] François Swiadek, Ross Shillito, Paul Magnard, Ants Remm, Christoph Hellings, Nathan Lacroix, Quentin Ficheux, Dante Colao Zanuz, Graham J Norris, Alexandre Blais *et al.* [Enhancing Dispersive Readout of Superconducting Qubits Through Dynamic Control of the Dispersive Shift: Experiment and Theory](#). arXiv, 2023. (Cited on pages 127, 138 and 139.)
- [Touzard *et al.* 2018] S. Touzard, A. Grimm, Z. Leghtas, S. O. Mundhada, P. Reinhold, C. Axline, M. Reagor, K. Chou, J. Blumoff, K. M. Sliwa, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi and M. H. Devoret. [Coherent oscillations inside a quantum manifold stabilized by dissipation](#). Physical Review X, vol. 8, no. 2, page 021005, 2018. (Cited on pages 50, 59, 80, 84 and 145.)

- [Touzard *et al.* 2019] S. Touzard, A. Kou, N. Frattini, V. Sivak, S. Puri, A. Grimm, L. Frunzio, S. Shankar and M. H. Devoret. [Gated conditional displacement readout of superconducting qubits](#). Physical Review Letters, vol. 122, no. 8, page 080502, 2019. (Cited on pages 67 and 69.)
- [Vepsäläinen *et al.* 2019] Antti Vepsäläinen, Sergey Danilin and Gheorghe Sorin Paraoanu. [Superadiabatic population transfer in a three-level superconducting circuit](#). Science advances, vol. 5, no. 2, page eaau5999, 2019. (Cited on page 116.)
- [Vion *et al.* 2002] D. Vion, A. Aassime, A. Cottet, P. Joyez, H. Pothier, C. Urbina, D. Esteve and M. H. Devoret. [Manipulating the Quantum State of an Electrical Circuit](#). Science, vol. 296, no. 5569, pages 886–889, 2002. (Cited on pages 3 and 6.)
- [Vool & Devoret 2017] Uri Vool and Michel Devoret. [Introduction to quantum electromagnetic circuits](#). International Journal of Circuit Theory and Applications, vol. 45, no. 7, pages 897–934, 2017. (Cited on page 4.)
- [Wallraff *et al.* 2004] Andreas Wallraff, David I Schuster, Alexandre Blais, Luigi Frunzio, R-S Huang, Johannes Majer, Sameer Kumar, Steven M Girvin and Robert J Schoelkopf. [Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics](#). Nature, vol. 431, no. 7005, pages 162–167, 2004. (Cited on page 6.)
- [Walter *et al.* 2017] T. Walter, P. Kurpiers, S. Gasparinetti, P. Magnard, A. Potočnik, Y. Salathé, M. Pechal, M. Mondal, M. Oppliger, C. Eichler and A. Wallraff. [Rapid High-Fidelity Single-Shot Dispersive Readout of Superconducting Qubits](#). Phys. Rev. Appl., vol. 7, page 054020, May 2017. (Cited on page 139.)
- [Werschnik & Gross 2007] J Werschnik and EKH Gross. [Quantum optimal control theory](#). Journal of Physics B: Atomic, Molecular and Optical Physics, vol. 40, no. 18, page R175, 2007. (Cited on page 126.)
- [Wolf & Cirac 2008] Michael M Wolf and J Ignacio Cirac. [Dividing quantum channels](#). Communications in Mathematical Physics, vol. 279, no. 1, pages 147–168, 2008. (Cited on page 151.)
- [Xu *et al.* 2022a] Qian Xu, Joseph K. Iverson, Fernando G. S. L. Brandão and Liang Jiang. [Engineering fast bias-preserving gates on stabilized cat qubits](#). Physical Review Research, vol. 4, no. 1, page 013082, 2022. (Cited on pages 40, 50, 60, 63, 65, 67, 69, 84, 91 and 95.)

- [Xu *et al.* 2022b] Qian Xu, Guo Zheng, Yu-Xin Wang, Peter Zoller, Aashish A. Clerk and Liang Jiang. [Autonomous quantum error correction and fault-tolerant quantum computation with squeezed cat qubits](#). arXiv, 2022. (Cited on pages 84, 102 and 123.)