

## Abstract

### Quantum Error Correction Beyond Break-Even

Volodymyr Sivak

2023

The ambition of harnessing the quantum for computation is at odds with the fundamental phenomenon of decoherence. The purpose of quantum error correction (QEC) is to counteract the natural tendency of a complex system to decohere. This cooperative process, which requires participation of multiple quantum and classical components, creates a special type of dissipation that removes the entropy caused by the errors faster than the rate at which these errors corrupt the stored quantum information. Previous experimental attempts to engineer such a process faced an excessive generation of errors that overwhelmed the error-correcting capability of the process itself. Whether it is practically possible to utilize QEC for extending quantum coherence thus remains an open question. The goal of this thesis work is to answer this question. Our main result is the first demonstration of a fully stabilized and error-corrected logical qubit whose quantum coherence is significantly longer than that of all the imperfect quantum components involved in the QEC process, beating the best of them with a coherence gain of  $G = 2.27 \pm 0.07$ . We achieved this performance by combining innovations in several domains including the fabrication of superconducting quantum circuits and model-free reinforcement learning. The experiment reported in this thesis uplifts quantum error correction from proof-of-principle studies to a practical tool for extending quantum coherence.

# Quantum Error Correction Beyond Break-Even

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

by  
Volodymyr Sivak

Dissertation Director: Michel H. Devoret

May 2023

© 2023 by Volodymyr Sivak  
All rights reserved.

To my middle school teacher of mathematics,  
Natalia Stanislavivna Bila.

# Contents

<b>List of Figures</b>	<b>iii</b>
<b>List of Tables</b>	<b>iv</b>
<b>List of Symbols</b>	<b>v</b>
<b>Acknowledgements</b>	<b>ix</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation and central question . . . . .	1
1.2 Overview of the field of QEC . . . . .	2
1.3 The structure of this thesis . . . . .	9
<b>2 Main results</b>	<b>11</b>
2.1 Introduction to the problem . . . . .	11
2.2 Engineering error correction . . . . .	12
2.3 Learning QEC circuit parameters . . . . .	15
2.4 Observing QEC beyond break-even . . . . .	18
2.5 QEC process characterization . . . . .	19
2.6 Conclusion and outlook . . . . .	22
2.7 Additional information . . . . .	23
<b>3 Experimental setup and sample</b>	<b>26</b>
3.1 Assembly . . . . .	26
3.2 Ancilla chip fabrication . . . . .	27
3.3 System parameters . . . . .	28
3.4 Control wiring . . . . .	28
<b>4 Calibration and characterization</b>	<b>31</b>
4.1 Primitive pulses . . . . .	31
4.2 Hamiltonian parameters . . . . .	32
4.3 Readout and reset . . . . .	34
4.4 Conditional displacement . . . . .	36
4.5 Oscillator error channels . . . . .	40
4.6 Active oscillator cooling . . . . .	41
<b>5 Quantum control optimization</b>	<b>45</b>
5.1 Model-based optimization of control circuits . . . . .	45
5.2 Model-free reinforcement learning for QEC . . . . .	49
<b>6 Quantum error correction</b>	<b>52</b>
6.1 Brief introduction to grid code . . . . .	52

6.2	Small-Big-Small (SBS) protocol . . . . .	54
6.3	QEC cycle: implementation details . . . . .	58
6.4	Learned and scripted parameters . . . . .	62
6.5	Syndrome measurement statistics . . . . .	65
6.6	Analysis of transmon leakage . . . . .	68
6.7	Post-selection of errors . . . . .	70
6.8	Wigner tomography of logical states . . . . .	71
6.9	Sensitivity to ancilla errors . . . . .	78
6.10	Long-time system stability . . . . .	82
6.11	Average channel fidelity . . . . .	84
<b>7</b>	<b>Outlook</b>	<b>88</b>
7.1	Complete characterization of a single logical qubit . . . . .	88
7.2	Promising variations of the QEC protocol . . . . .	89
7.3	Machine learning for quantum control . . . . .	91
	<b>Appendices</b>	<b>93</b>
<b>A</b>	<b>Phase space representations</b>	<b>93</b>
A.1	Characteristic function . . . . .	93
A.2	Wigner function . . . . .	93
	<b>References</b>	<b>95</b>

# List of Figures

1.1	Comparison of QEC experiments . . . . .	7
2.1	Experimental system . . . . .	13
2.2	QEC implementation and optimization . . . . .	16
2.3	System coherence . . . . .	18
2.4	Analysis of error syndromes . . . . .	20
3.1	Sample assembly . . . . .	26
3.2	Experimental setup . . . . .	30
4.1	Calibration of primitive pulses . . . . .	32
4.2	Calibration of Hamiltonian parameters . . . . .	33
4.3	Ancilla measurement . . . . .	35
4.4	Calibration of the ECD gate . . . . .	38
4.5	Fluctuating error channels . . . . .	41
4.6	System cooling . . . . .	42
5.1	Circuit optimization . . . . .	47
5.2	Waveform for grid state preparation . . . . .	48
5.3	Reinforcement learning . . . . .	50
6.1	Wigner functions of grid states . . . . .	53
6.2	SBS protocol I . . . . .	55
6.3	SBS protocol II . . . . .	56
6.4	Evolution of QEC circuit parameters during the training . . . . .	63
6.5	Syndrome measurement outcomes . . . . .	66
6.6	Analysis of syndrome measurement outcomes . . . . .	67
6.7	Correlation of syndrome measurement outcomes . . . . .	69
6.8	State reconstruction . . . . .	72
6.9	Wigner tomography after QEC . . . . .	75
6.10	QEC subspaces . . . . .	76
6.11	Analysis of state evolution . . . . .	77
6.12	Effect of ancilla errors . . . . .	80
6.13	Noise verification experiments . . . . .	81
6.14	Quantum system stability . . . . .	83
6.15	Average channel fidelity . . . . .	86

# List of Tables

3.1	Measured system parameters . . . . .	28
5.1	Circuit parameters . . . . .	46
6.1	Timing of the cycle components . . . . .	59
6.2	Post-selection results . . . . .	71

# List of Symbols

## Acronyms

ADC	analog to digital converter
DAC	digital to analog converter
DIO	digital input and output
ECD	echoed conditional displacement
FPGA	field programmable gate array
GKP	Gottesman, Kitaev and Preskill
GPU	graphics processing unit
JAMPA	Josephson array-mode parametric amplifier
LO	local oscillator
ML	machine learning
NISQ	noisy intermediate scale quantum
NN	neural network
PPO	proximal policy optimization
QEC	quantum error correction
QED	quantum electrodynamics
RL	reinforcement learning
ReLU	rectifier linear unit
RB	randomized benchmarking
SPAM	state preparation and measurement
SNR	signal to noise ratio
SBS	small-big-small protocol

## Qubit

$\sigma_x, \sigma_y, \sigma_z$	$2 \times 2$ Pauli matrices
$\sigma_+, \sigma_-$	spin creation and annihilation operators
$X, Y, Z$	Pauli operators
$ \pm X\rangle,  \pm Y\rangle,  \pm Z\rangle$	eigenstates of Pauli operators

## Oscillator

$a, a^\dagger$	bosonic annihilation and creation operators
----------------	---

$ n\rangle$	Fock state with $n$ photons; $a^\dagger a n\rangle = n n\rangle$
$ \alpha\rangle$	coherent state with amplitude $\alpha$ ; $a \alpha\rangle = \alpha \alpha\rangle$
$\Pi$	photon number parity operator; $\Pi = \exp(i\pi a^\dagger a)$
$\Pi_\alpha$	displaced parity operator; $\Pi_\alpha = D(\alpha)\Pi D^\dagger(\alpha)$
$W(\alpha)$	Wigner function; $W(\alpha) = (2/\pi)\langle\Pi_\alpha\rangle$

## Transmon-cavity system

$ g\rangle,  e\rangle,  f\rangle,  h\rangle$	transmon energy eigenstates
$\omega_t$	ancilla qubit frequency ( $ g\rangle -  e\rangle$ transition)
$\omega_r$	readout resonator frequency
$\omega_c$	cavity mode (oscillator) frequency
$\chi$	dispersive shift between qubit and oscillator
$\chi'$	second-order dispersive shift
$K$	oscillator Kerr nonlinearity
$H_d$	dispersive coupling Hamiltonian; $H_d = \chi\sigma_z a^\dagger a/2$

## Control primitives

$R_\varphi(\theta)$	qubit rotation; $R_\varphi(\theta) = \exp[-i(\theta/2)(\cos\varphi\sigma_x + \sin\varphi\sigma_y)]$
$D(\alpha)$	oscillator displacement; $D(\alpha) = \exp(\alpha a^\dagger - \alpha^* a)$
$R_V(\vartheta)$	virtual rotation; $R_V(\vartheta) = \exp(i\vartheta a^\dagger a)$
$\text{CR}(\theta)$	conditional rotation; $\text{CR}(\theta) = \exp(i\theta\sigma_z a^\dagger a)$
$\text{ECD}(\beta)$	echoed conditional displacement; $\text{ECD}(\beta) = \sigma_x D(\sigma_z \beta/2)$

## Logical qubit

$\mathcal{C}_0, \mathcal{C}_1, \mathcal{C}_2, \dots$	code subspace and its replicas (error subspaces)
$\Pi_i$	projector onto the subspace $\mathcal{C}_i$
$M$	symplectic matrix defining the grid code
$l_S$	size of the square grid code unit cell; $l_S = \sqrt{2\pi}$
$S_0^X, S_0^Z$	stabilizers of the ideal grid code; $S_0^X = D(l_S), S_0^Z = D(il_S)$
$X_L, Y_L, Z_L$	logical Pauli operators of the ideal grid code
$\Delta$	size of the grid code envelope
$N_\Delta$	normalizing envelope operator; $N_\Delta = \exp(-\Delta^2 a^\dagger a)$

$S_{\Delta}^x, S_{\Delta}^z$	stabilizers of finite-energy grid code; $S_{\Delta}^{x/z} = N_{\Delta} S_0^{x/z} N_{\Delta}^{-1}$
$\Pi_{\Delta}$	projector onto the finite-energy code space

## Error-correcting dissipation

$\mathcal{R}_{\Delta}^x, \mathcal{R}_{\Delta}^z$	dissipators to the +1 eigenspace of $S_{\Delta}^x$ and $S_{\Delta}^z$
$\mathcal{R}_{\Delta}$	dissipator to the finite-energy grid code; $\mathcal{R}_{\Delta} = \mathcal{R}_{\Delta}^x \circ \mathcal{R}_{\Delta}^z$
$U_{\emptyset}$	entangling unitary operation; used in the decomposition of any rank-2 dissipator, such as $\mathcal{R}_{\Delta}^x$ or $\mathcal{R}_{\Delta}^z$
$U_b$	unitary operation on the oscillator Hilbert space conditioned on the ancilla measurement outcome $b$ ; used in the decomposition of any rank-2 dissipator
$t_c$	duration of the full QEC or cooling cycle
$K_{gg}, K_{ge}, K_{eg}, K_{ee}$	Kraus operators of the SBS protocol corresponding to ancilla measurement outcomes listed in the subscript

## Error channels

$T_1^t, \bar{T}_1^t$	transmon lifetime and its week-long average
$T_1^t(\bar{n})$	transmon lifetime in the presence of $\bar{n}$ steady-state photons in the readout resonator
$T_{2E}^t, \bar{T}_{2E}^t$	transmon echo coherence time and its week-long average
$T_{\varphi}^t$	transmon pure dephasing time
$T_1^c, \bar{T}_1^c$	cavity mode lifetime and its week-long average
$T_2^c, \bar{T}_2^c$	cavity mode coherence time and its week-long average
$T_{\varphi}^c$	cavity mode pure dephasing time
$T_X, T_Y, T_Z$	lifetime of the logical $X, Y, Z$ eigenstates
$p_X, p_Y, p_Z$	probability of logical Pauli $X, Y, Z$ errors per QEC cycle
$p_{\text{err}}$	probability of leaving the code space per QEC cycle
$\gamma_{\downarrow}^t, \gamma_{\uparrow}^t$	transmon photon loss and photon gain rates in the $\{ g\rangle,  e\rangle\}$ subspace
$\gamma_1^t$	transmon energy relaxation rate; $\gamma_1^t = \gamma_{\downarrow}^t + \gamma_{\uparrow}^t$
$\gamma_2^t, \gamma_{\varphi}^t$	transmon dephasing and pure dephasing rates in the $\{ g\rangle,  e\rangle\}$ subspace; $\gamma_{\varphi}^t = \gamma_2^t - \gamma_1^t/2$
$\gamma_{\downarrow}^c, \gamma_{\uparrow}^c$	cavity mode photon loss and photon gain rates in the $\{ 0\rangle,  1\rangle\}$ subspace

$\gamma_1^c$	cavity mode energy relaxation rate; $\gamma_1^t = \gamma_\downarrow^t + \gamma_\uparrow^c$
$\gamma_2^c, \gamma_\varphi^c$	cavity mode dephasing and pure dephasing rates in the $\{ 0\rangle,  1\rangle\}$ subspace; $\gamma_\varphi^c = \gamma_2^c - \gamma_1^c/2$
$\gamma_\varphi^{c,t}$	cavity mode dephasing rate induced by transitions of a dispersively coupled transmon $\{ g\rangle,  e\rangle\}$ qubit
$\gamma_?$	unexplained contribution to cavity mode dephasing rate
$\gamma_X, \gamma_Y, \gamma_Z$	decoherence rates of the logical $X, Y, Z$ eigenstates
$\Gamma_{\{01\}}$	decay rate of the average channel fidelity of the $\{ 0\rangle,  1\rangle\}$ qubit; $\Gamma_{\{01\}} = (\gamma_1^c + 2\gamma_2^c)/3$
$\Gamma_{\{ge\}}$	decay rate of the average channel fidelity of the $\{ g\rangle,  e\rangle\}$ qubit; $\Gamma_{\{ge\}} = (\gamma_1^t + 2\gamma_{2E}^t)/3$
$\Gamma_{\text{GKP}}$	decay rate of average channel fidelity of an error-corrected logical GKP qubit; $\Gamma_{\text{GKP}} = (\gamma_X + \gamma_Y + \gamma_Z)/3$

## Reinforcement learning

$\mathcal{N}(\vec{\mu}, \vec{\sigma})$	multivariate normal action distribution with mean $\vec{\mu}$ and covariance $\text{diag}[\vec{\sigma}]^2$
$\theta$	parameters of the neural network
$\vec{p}$	parametrization of the QEC circuit (“actions”)
$B$	training batch size

## Miscellaneous

$\rho$	density matrix
$\mathcal{E}$	arbitrary error channel; $\mathcal{E} : \rho \rightarrow \mathcal{E}(\rho)$
$\overline{\mathcal{F}}$	average channel fidelity
$G$	coherence gain of an error-corrected logical qubit over the best passive qubit encoding in the system
$I_{th}, Q_{th}$	readout classification thresholds
$(s_0, s_1)$	classified measurement outcome (2 bits)
$\mathcal{F}_r^{(g)}, \mathcal{F}_r^{(e)}$	readout fidelity of the transmon $ g\rangle$ and $ e\rangle$ states
Y90	rotation by $\pi/2$ around $Y$ axis of the $\{ 0\rangle,  1\rangle\}$ qubit
$r_{ij}$	time-correlation matrix for measurement outcomes

# Acknowledgements

Scientist is not just what I am, it is who I am. An entry in my diary dated as far back as 10/16/2007 (when I was 13 years old) reads “I decided to become a Doctor of Science in mathematics and now diligently go towards this” – and it only took 15 years to get there (well, almost there)! My prior childhood aspirations ran through everything from Paleontology to Astronomy. This early passion for science was lovingly cultivated by my parents, Olena Sivak and Volodymyr Sivak. I want to express my deepest gratitude to them for choosing this path for me before I was able to consciously choose it myself.

My middle-school teacher Natalia Stanislavivna Bila played the biggest role in setting me up for the life trajectory that I love so much now. She was preparing me to compete in mathematics Olympiads, which eventually led me into physics. For her selfless commitment to being a Teacher and for countless hours that we spent together solving extracurricular math problems, I devote this thesis to her. My teachers in the Ukrainian Physics and Mathematics Lyceum (UPML) in Kyiv, 2009–2012, and in Moscow Institute of Physics and Technology (MIPT), 2012–2016, guided me through the turbulent young years and further strengthened my intention of becoming a scientist. I am indebted to Dmytro Nomirovsky, Roman Ransky, Oleksandr Shevchuk, Mikhail Feigelman, and Alexey Ioselevich.

The six years at Yale, 2016–2022, have been the realization of my scientific dreams. The flourishing, creative environment that Michel Devoret, Rob Schoelkopf, and Steve Girvin created here is truly unique. Not only because of the quality of science done here, but also because of the community of people that do it. We all live and work here as one big friendly family, supporting each other and genuinely wishing each other success. It has been a great pleasure to be a part of this family.

I cannot imagine a better Ph.D. adviser than Michel Devoret. His childlike curiosity about nature resonates with me so much, yet his contribution to science is difficult for me to even comprehend. He is an inspiration to me, a real example of a Scientist I always wanted to become. I am honored to have had an opportunity to work with Michel, and I hope that the valuable lessons about the culture of science-making that I learned from him will continue guiding me in the future.

My research trajectory at Yale took unexpected turns, which allowed me to learn a range of topics from many amazing people. Leonid Glazman was my adviser during the first year here. He helped me transition from the previous background that I acquired in Moscow to the new field of quantum information. I admire his sharp intelligence, precision, and devotion to developing his student’s skills.

When I joined Michel's lab in 2017, I started working on quantum-limited superconducting amplifiers together with Nick Frattini and under the management of Shyam Shankar. Nick taught me how to fab, how to write Python scripts, how to make measurements, how to cool down the fridge, and how to have fun in the lab! Shyam was making sure that Nick and I don't blow up everything. His thoughtful advise and guidance on technical topics wonderfully complemented Michel's more abstract and far-looking supervision. Shyam knew the lab inside out, and it was not easy to readjust after he left to become a professor at Austin.

This event marked the beginning of my more independent thinking. I deep-dived into machine learning, which Michel enthusiastically supported, and was looking for the right problem to apply it to. At the same time, the first grid-code experiment in our lab has just achieved successful preparation and error correction of these states. It fascinated me immensely, and I finally found a goal for the rest of my Ph.D. – to apply the tools from the modern machine learning for improving the error correction of the grid states, and in particular for reaching beyond the break-even performance. And this is when I started collaborating with Alec Eickbusch. Our sparse conversations grew into a close partnership, very pleasant and fruitful, which lasted until the end of my career here at Yale. Alec taught me all the basic qubit experiments: how to measure  $T_1$ ,  $T_2$ , how to calibrate the Rabi pulses, and much more. Bouncing ideas off each other with Alec was always a great joy to me, I felt that we connected really well intellectually. Without him, the experiment described in my thesis would never get off the ground.

Around the same time, I met Baptiste Royer, a theorist from YQI. In my experience, most theorists are not able speak the language of experimentalists, but Baptiste was extremely good at this. Many experiments described in my thesis are direct implementations of his creative ideas, including, of course, the error correction protocol itself. I always felt that I can rely on him, that he is genuinely involved and is doing his best to help me out with all my questions and problems. When Baptiste left to become a professor at Sherbrooke, Shraddha Singh became my new go-to theorist. After the experiment was already finished, we spent weeks at YQI working together and trying to find a new language to describe the error correction of the grid code. We hit multiple dead ends, and eventually, with remote help of Baptiste, we converged on the new understanding of the grid code QEC which is presented in this thesis. This digging into the details was so exciting for me, I only regret that we haven't started the collaboration earlier!

The results demonstrated in my thesis would not be possible without the fabri-

cation and sample assembly help from Ioannis Tsioutsios. I won't forget the crazy heart beat I had when we sat in the cleanroom together trying to clamp the single transmon chip that survived the complicated fabrication process. Besides all the lab help, Ioannis was my good friend who rescued me from all sorts of trouble that I managed to get into during the Ph.D. years.

The final data presented in this thesis was taken with a tantalum-based transmon fabricated using a new technique developed at Yale by Suhas Ganjam. Without his contribution, the maximal achieved QEC gain would only be  $G = 1.45$  (using the aluminum-based transmon from Ioannis). I am grateful to Suhas for agreeing to fabricate this chip despite his very busy schedule at the time. It is interesting that of the whole batch of samples only a single transmon survived, yet this was sufficient to take the experiment to new heights.

In the middle of my Ph.D., I met Sandro Miano. We share a delightful experience of tearing down the experimental setup of the previous grad student who used the Dreadnought dilution fridge. Afterwards, together we built a new setup for the present experiment, which was a great learning experience. Sando became my very close friend outside the lab and I am grateful to him for all the emotional (in addition to technical) support.

Another Italian mafioso that I want to thank is Luigi Frunzio. He cares about all of us students like his own children, always giving helpful advise and providing a shoulder during the difficult times. Conversations with Luigi have been so much fun, I really enjoyed learning about his diverse life experiences, and I hope my future life can be as saturated and interesting as his.

Six years is a long time, during which I collaborated and interacted with many more colleagues, some of whom also became my good friends. I want to thank Shantanu Mundhada, Gangqiang Liu, Jaya Venkatraman, Zlatko Minev, Henry Liu, Allen Mi, Andrew Lingenfelter, Kevin Chou, Ben Brock, Valla Fatemi, Max Hays, Kyle Serniak, Jacob Curtis, Ben Chapman, Chris Wang, Wei Dai, Vidul Joshi, Spencer Diamond, Jahan Claes, and Akshay Koottandavida.

I want to thank my closest friends: Pavel Kurilovich, Vlad Kurilovich, and Rodrigo Cortiñas. I met "the Twins" over 10 years ago during my time in Moscow. For such a long period, our lives were filled with common physics discussions, outdoor adventures, political protests, and personal life dramas. They are the kindest people I know and devoted friends who I can count on in any situation. I love them, period. Rodrigo entered my life only recently, but we shared so much over the past two years that I have a hard time imagining how I would live without him.

My respect, admiration, and gratitude towards Rodrigo are boundless.

I want to say a sincere thanks for unforgettable experiences to Yuting Liu who was my life partner for many years at Yale. We learned dancing together, went on so many adventurous trips, and enjoyed life to the fullest, eventually parting our ways but always remaining in each others memories.

Many more friends enriched my Ph.D. journey. I want to thank Arpit Dua, Peeyush Khare, Brian Peterson, Lesya Kosteeva, Maria Kharchenko, Yaroslav Buts, Ivan Mykhailovskiy, Clara Ma, Lyuba Hubarenko, Pasha Stroganov, Rita Davydova, and my sister Masha Sivak.

The past six years was an incredible period of my life, where creative exploration, hard work, and fun adventures intertwined thanks to all these people.

# Introduction

## 1.1 Motivation and central question

A minimal information-carrying system consists of two states, conventionally labeled 0 and 1, and it is called a bit. The quantum analog of a bit is a qubit. In addition to states  $|0\rangle$  and  $|1\rangle$ , a qubit can be in any superposition state, such as  $|0\rangle + |1\rangle$ . This abstract two-dimensional object can be realized in Nature, for example, as polarization of a photon, or spin of an electron. We will refer to the qubit state as “logical” when it is viewed from the computational perspective, to distinguish it from the “physical” state of an object that hosts the qubit. For example, a logical state  $|0\rangle + |1\rangle$  can be realized as a physical state  $|\downarrow\rangle + |\uparrow\rangle$  of a single spin. The goodness of a physical implementation of a qubit from the computational perspective is related to how reliably it can maintain the logical state and change this state when required by the computation.

Any spurious interaction of the object with its environment, which is inevitable in the real world, would interfere with the computation by changing in some way the physical state of the object. If logical state was equivalent to a physical state, e.g.  $|0\rangle = |\downarrow\rangle$  and  $|1\rangle = |\uparrow\rangle$  for a spin, it would be impossible to distinguish if the change was caused by the spurious interaction or by the computation itself. However, to achieve fault-tolerant computation, there needs to be a mechanism for making such a distinction. In other words, errors should leave some unique signatures that would never be generated by the normal computation flow. Hence, instead of encoding the logical state into a physical state of a two-dimensional object, we need an object with an extended state space that has enough room to store the information about the error – a kind of “meta-qubit”.

The computation can then be intertwined with quantum error correction (QEC) where the error signatures are checked, the errors are identified, corrected, and the logical state is recovered. On a first glance, in quantum mechanics such a goal seems unattainable, since simply checking for errors would “collapse” the quantum state. A beautiful realization in the theory of quantum error correction is that the collapse of a physical state due to error check need not affect the logical state, and hence fault-tolerant quantum computation is in principle possible.

This can be achieved by distributing the logical information in the state space of an object in such a way that whenever collapse due to error check happens,

the physical state finds itself in one of the orthogonal error subspaces which are all replicas of the same logical subspace. The index of the subspace would then serve as the error signature, while the collapsed state, although physically different, would still contain the same logical information. The state can then be brought back to the logical subspace, and the computation can continue as if the error did not happen.

This simple idea of quantum error correction has been formalized and theoretically developed over the past three decades. However, most theoretical proposals of the QEC protocols operate at the level of abstraction insufficient to capture the real-world complexity of the physical implementations. So far, all attempts to experimentally realize a QEC process faced an excessive generation of errors that overwhelmed the error-correcting ability of the process itself. **Consequently, a question of whether it is *practically* possible to harness QEC for extending the coherence of a logical quantum state remains unanswered. The goal of this thesis is to answer this question affirmatively and definitively.**

## 1.2 Overview of the field of QEC

The first example of a quantum error correcting code was given by P. Shor in 1995 [1]. This code, constructed as two concatenated layers of a classical repetition code, uses a register of nine two-level systems (henceforth called “physical qubits”) to create a single logical qubit that can be protected against bit- and phase-flip errors on any single qubit in the register. Around the same time, A. Steane provided an example of a minimal code that achieves similar protection with only seven physical qubits [2]. Early progress in the theory of quantum error correction was driven primarily by the ideas already developed in the classical information theory since the seminal work of R. Hamming in 1950 [3]. As described earlier, QEC was treated as a three-step process of detecting the occurrence of an error, decoding its signature, and applying a recovery operation. The so-called stabilizer formalism [4] emerged as a powerful general framework for constructing codes and QEC schemes that follow this pattern. The most widely studied example of a stabilizer code is the surface code [5, 6], which emerged from the ideas introduced by A. Kitaev in 2003 [7]. To generate the redundant state space, this code relies on a register of physical qubits arranged on a two-dimensional lattice. It has a remarkable phase transition in the space of error probabilities – if the rate of physical errors is below a certain threshold (and if a number of other assumptions are satisfied), then the logical er-

ror rate could be suppressed exponentially in the size of the lattice with QEC – we call this property the threshold theorem. Many other QEC codes also have such a property.

At the time of writing, the surface code is considered by the majority of the QEC community to be the most promising candidate for achieving large-scale fault-tolerant quantum computation. Without subscribing to this particular view, it is still instructive to consider the history of the development of the field of experimental QEC from a perspective of approaching closer towards the technological feasibility of a surface code. I would like to adopt such a perspective in this Chapter, since the rest of the thesis is focused on a somewhat different approach to QEC, which, according to my impression, is currently treated as merely an interesting alternative to the mainstream. Towards the end of this Chapter, we will discuss what is the central difference between these approaches.

The experimental effort in QEC has undergone tremendous progress over the years (although it lags far behind the theory to this day). Across different physical qubit modalities from traveling photons, trapped ions, and cold atoms, to superconducting circuits, the initial experimental efforts were focused on developing the control primitives required for quantum computing: single-qubit gates, entangling operations, projective measurements [8].

An extraordinary progress has happened in the field of quantum superconducting circuits, where collective degrees of freedom of billions of atoms making up the circuit can display the properties of a single “artificial atom”. The first detection of quantum tunneling of a macroscopic degree of freedom and the first observation of the quantized energy levels of a macroscopic electrical circuit were done in the experiments of M. Devoret, J. Martinis, and J. Clarke in 1985 [9, 10], giving birth to an experimental direction that would decades later become the leading candidate for realizing practical quantum computing. Starting with a demonstration of the first superconducting qubit by Y. Nakamura *et al.* in 1999 [11], the quantum coherence of these circuits has followed rapid progress over the next two decades [12, 13]. The improvement by many orders of magnitude was driven by compounded innovations in circuit design and fabrication techniques.

Once it became clear that error rates in superconducting circuits are at the level from which is conceivable to reach the threshold of the surface code, big tech companies IBM and Google launched a parallel effort to scale the size of the superconducting qubit registers. When I started on a PhD journey in 2016, registers of 5 qubits were state-of-the-art [14], and over the past six years the size of these quan-

tum systems has increased to 54 qubits in 2018 [15] and 72 qubits in 2022 [16], with individual control over all of them. This remarkable technological leap in such a short time fuels the optimism of the quantum computing community that one day the surface code can become a reality.

Powerful as it is, the surface code threshold theorem has a flip side: while all of the surface code model assumptions are not met, such qubit register systems will not display an improvement of the quality of a logical qubit compared to the constituent physical qubits – such is the nature of phase transitions. A few recent demonstrations of small surface code instances [16–18] have made explicit the massive technological challenges involved in driving progress along this direction. Therefore, in the current era of noisy intermediate-scale quantum (NISQ) devices [19], from a point of view of executing quantum computations it is advantageous to simply use the physical qubit state to represent the logical state.

The realization that useful quantum error correction might be far away in the future drives the search for practical applications of imperfect but still impressively capable quantum devices of the day. During my PhD, one of the greatest breakthroughs in this realm was the demonstration of “quantum supremacy” [15] – an ability of a small quantum processor with only 54 qubits to perform a task that the most powerful supercomputers on Earth would take orders of magnitude longer time to emulate. This task, the generation of random samples from a specific probability distribution, does not have any known applications, but it gives strong evidence that quantum entanglement could indeed be a powerful computing resource. The hope of the field is that quantum computing in the NISQ era might be used for simulations of few-but-already-too-many-body systems such as organic molecules and perhaps even bigger systems that fall under the conventional many-body physics umbrella. The first steps in this direction have already been taken, see review [20]. Of course, there is no real boundary between NISQ and beyond-NISQ applications: the more reliable logical qubits become, the wider range of problems can benefit from utilizing them.

My personal opinion is that technology required for the surface code will be developed in the foreseeable future. My optimism is based on the achievements by scientists and engineers of seemingly impossible technological feats, such as detection of gravitational waves by the Laser Interferometer Gravitational-wave Observatory (LIGO) [21] or the discovery of the Higgs boson at the Large Hadron Collider (LHC) [22], both of which happened during my scientific upbringing. The theoretical foundations for these experiments did not raise much doubt, but

it was not clear from the beginning how to overcome the involved engineering challenges, and if the immense complexity of these systems would be ultimately controllable. In light of such accomplishments, building the surface code for fault-tolerant quantum computing appears possible given enough resources, human effort, and time. However, it would be naive to believe that everything about the physics of large-scale quantum coherent devices is already understood, and all that remains is straightforward linear optimization and scaling. To some degree, classical computing has followed this linear path after the invention of a transistor. However, it remains absolutely unclear at this stage what would be the “transistor” of the surface code. Currently, in quantum processors built by IBM and Google with a forward looking goal of creating the surface code, the role of a base qubit is taken by the transmon [23, 24]. At the same time, another circuit, the fluxonium [25], is rapidly gaining popularity and showing great promise. New kinds of superconducting circuits appear every year [26–28], and the old kinds are being improved at a steady pace [29, 30]. Given this diversity, the honest assessment of the field leads to a conclusion that the base element of the future surface code has likely not even been invented yet.

Maybe this base qubit, instead of being represented with two energy levels of a static superconducting circuit, should itself be a dynamically-protected mini logical qubit? Viewed through the lens of the broader effort of the field to engineer better qubits, the contribution of this thesis consists in developing an exotic qubit of such a kind, error-corrected in real time using an efficient QEC process tailored to the properties of the system in which such a qubit is realized – the harmonic oscillator.

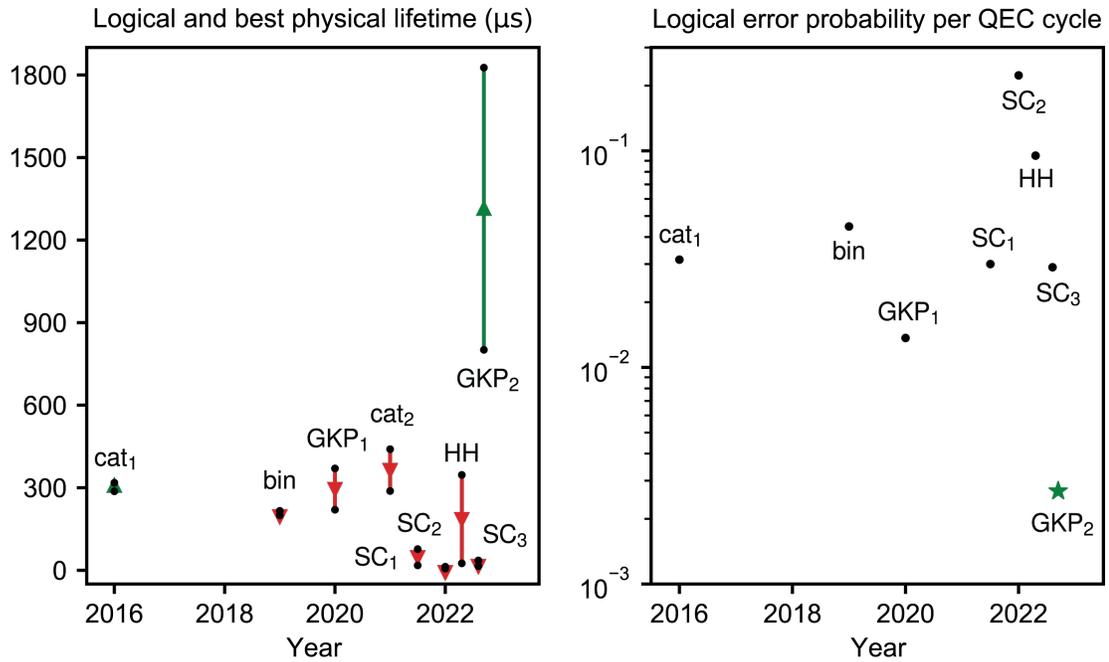
The encoding of a qubit into the large state space of a harmonic oscillator, which is implemented in this thesis, was first proposed back in 2001, when I went to the first grade of the primary school, by Gottesman, Kitaev and Preskill [31]. Other “bosonic codes” were developed since [32–34]. Using an oscillator to realize QEC proved quite successful in practice due to the inherent hardware efficiency – a large state space is created with only a single degree of freedom. This efficiency can be appreciated by comparison with a QEC approach based on registers of physical qubits, such as the surface code, where orchestrating the quantum evolution of multiple individually-controllable degrees of freedom is poised to introduce a lot of detrimental overhead. Harmonic oscillator is also used in the latter approach, although it plays a different role there. It is therefore instructive to briefly review its history in the context of quantum information processing.

When I started college in 2012, the Nobel prize in physics was awarded to

D. Wineland and S. Haroche “for ground-breaking experimental methods that enable measuring and manipulation of individual quantum systems”. In experiments by S. Haroche with cavity quantum electrodynamics (QED) systems [35], oscillator is realized as an electromagnetic mode of a three-dimensional cavity, and its state is controlled via the interaction with atoms that are sent through this cavity. Early on, it became clear that superconducting circuits are not only suitable to realize the physics of cavity QED, but that they also provide unprecedented engineering flexibility and access to new parameter regimes. This seminal insight by A. Blais *et al.* [36] and experiment by A. Wallraff *et al.* [37] started the field of circuit QED [38].

Similarly to cavity QED, the quantum state of an oscillator in circuit QED is controlled via the interaction with an atom, albeit artificial. Given the large non-linearity of an atom, it is possible to restrict its dynamics to the lowest two energy levels, hence realizing an approximate physical qubit. Different kinds of superconducting circuits realize different artificial atoms and result in qubits with different properties [12, 13, 39]. Some of them have the advantage of long coherence times, some offer fast quantum control, some meet both these criteria to a certain degree. In the multi-dimensional design problem with various tradeoffs, a circuit that strikes a good balance on many criteria while remaining fairly simple is the transmon [23]. Transmon is currently used to realize qubit registers in the surface code approach to QEC [16–18], but is also used to realize an ancilla qubit for control of an oscillator in the bosonic approach [40–43].

The experiment described in this thesis is built upon an architecture similar to [40–43]. The two crucial building blocks in this architecture are the superconducting cavity whose electromagnetic mode realizes a harmonic oscillator, and a transmon circuit whose mode of collective motion of Cooper pairs across the Josephson junction realizes an artificial atom used as a qubit. That the cooperative phenomenon of quantum error correction can be achieved with just two parties, an oscillator and a qubit, is remarkable. As of now, there is no threshold theorem that would guarantee an exponential suppression of the logical error probability by scaling the code size in such a minimal system. Moreover, there are good reasons to believe that this is not possible given realistic error channels of any oscillator, in which physical error probability increases linearly with the size of the occupied Hilbert space (instead of logarithmically, as in qubit registers). However, the threshold theorem of a surface code, beautiful as it is, from a practical standpoint is nothing more than a recipe for creating a reliable logical qubit. If a desired level of reliability can be achieved by other means, without the phase transition, then it



**Figure 1.1 | Comparison of QEC experiments.** In the first panel, lifetime is defined as the decay time constant of the average channel fidelity to identity channel. Arrows point from the best physical lifetime to the achieved logical lifetime. They are colored green if QEC has exceeded the break-even point, and red otherwise. In the second panel, logical error probability per cycle is obtained using logical lifetime and cycle duration. Experiment labels are the following:  $\text{cat}_1$ , four-legged cat code [40];  $\text{bin}$ , binomial code [41];  $\text{cat}_2$ , truncated four-legged cat code [43];  $\text{GKP}_1$ , grid code [42];  $\text{SC}_1$ , distance-three surface code [17];  $\text{SC}_2$ , distance-three surface code [18];  $\text{SC}_3$ , distance-five surface code [16];  $\text{HH}$ , heavy-hexagonal subsystem code [44]; and  $\text{GKP}_2$  is the grid code experiment from this thesis.

is just as good for the quantum computation.

What is meant by reliability of a logical qubit is a difficult multi-dimensional question. The ultimate judge of the reliability is the capacity of the system to provide correct answers to computational problems. None of the existing experimental approaches to quantum computing is developed enough to be judged in such a way. Hence, we need to look at different aspects of the logical qubit performance that we expect to be correlated with the ultimate future ability to solve computational problems. Among these aspects, the most informative is the fidelity of logical operations, and in particular the logical two-qubit gates. While there has been no demonstrations of error-corrected gates between logical qubits yet, we focus on a simpler task of using quantum error correction to improve logical qubit memory, i.e. to preserve the logical quantum state through time. In this regard, one can identify three metrics:

1. The lifetime of a logical qubit.
2. The gain  $G$  due to quantum error correction, i.e. an improvement of the lifetime of a logical qubit over the lifetime of the best physical qubit in the system.
3. The logical error probability per cycle of quantum error correction.

In Fig. 1.1, we compare the existing demonstrations of error-corrected logical qubits using these three criteria. The work described here achieved a significant advancement of the state-of-the-art on all these metrics. **The demonstration of gain  $G = 2.27 \pm 0.07$  beyond break-even is the most important contribution of this thesis. It marks a transition of QEC from proof-of-principle to a practical tool for extending quantum coherence.** This was accomplished through a combination of multiple recent innovations in domains ranging from the fabrication of superconducting quantum circuits, to low-latency digital electronics, and modern machine learning.

A lot of these innovations were done here at Yale. In particular, the circuit QED architecture for realizing long-lived oscillators in three-dimensional aluminum cavities was developed in Refs. [45, 46], and the cavity used in this work was recycled from an earlier experiment in Ref. [47]. The use of transmon circuit to realize high-fidelity quantum control of an oscillator in the dispersive regime of circuit QED was developed in Refs. [48–51]. Fabrication of a transmon with state-of-the-art quantum coherence was done with a technique developed in [52] that was adopted from an earlier work in Ref. [29]. Orchestration of the QEC process in this experiment

is realized with a field-programmable gate array (FPGA) controller developed and first applied for quantum error correction in Ref. [40]. The first error correction experiment with the same type of logical qubit as used in this thesis was done in Ref. [42]; that work in particular has made key conceptual contributions that later enabled the present experiment. The QEC protocol realized here was theoretically proposed in Ref. [53] (and independently in Ref. [54], where it was also experimentally realized with a trapped ion system). While this body of previous research by my colleagues has built a solid foundation of experimental and theoretical techniques verified in multiple experiments at Yale and other research institutions, a novel aspect brought in by my own work is the framework of model-free reinforcement learning for quantum control [55]. It is deployed here to learn the parameters of the QEC circuit *in-situ* to maximize the lifetime of a logical qubit. This framework is derived from recent progress in the field of artificial intelligence, which captivates me just as much as quantum physics. The reinforcement learning algorithm used here, the proximal policy optimization (PPO), was introduced in Ref. [56].

### 1.3 The structure of this thesis

In this thesis, I do not explain the fundamental notions of quantum physics, such as the Bloch sphere representation of a qubit or the second quantization framework for an oscillator, since a multitude of excellent introductory texts exists on these topics, for example [35]. An introduction to open quantum systems, including the quantum trajectory formalism, is presented in [57]. An introduction to quantum information science, including the basics of quantum error correction, is given in [58]. In the context of circuit QED with transmons and superconducting cavities, the theses of my predecessors at Yale [59–62] and elsewhere [63, 64] provide an exposition of the material in a way that would most closely match the background of a potential reader. An accessible overview of the grid code is provided in [65]. For those interested in machine learning, I compiled a selection of resources [66] that proved especially useful in my self-education on this topic. To mention a few, the introductions to artificial intelligence in general [67] and to reinforcement learning [68] in particular, are my favorite.

This thesis is organized as follows. For the benefit of time-constrained readers, Chapter 2 describes the key contributions distilled to about 10 pages. The rest contains details of the experimental and theoretical methods developed or adopted in this experiment. In particular, Chapter 3 describes the experimental setup and

---

the sample. Chapter 4 introduces system calibration and characterization experiments, Chapter 5 describes numerical optimization tools for quantum control and error correction. Finally, Chapter 6 focuses on the grid code: it contains a brief introduction to the theory of this code, details of the QEC protocol and its implementation, and several experiments that address various aspects of the logical qubit performance. The outlook in Chapter 7 speculates on the future development of the grid-code QEC.

# Main results

## 2.1 Introduction to the problem

Implementing a single correctable logical qubit requires a physical system with a large state space. It should accommodate the code subspace and its redundant replicas where the logical information will be displaced without distortion when physical errors occur [69]. This redundancy is inextricably associated with an additional operational cost of QEC, known as the control overhead. In the search for an efficient way to alleviate its detrimental effects, bosonic codes [31–34] based on the state space of a harmonic oscillator have been proposed as a promising alternative to the standard approach based on registers of physical qubits [1, 2, 5]. In hybrid architectures, these two approaches are complementary, with qubit-register codes built upon logical qubits dynamically protected with efficient base-layer bosonic QEC [70–72].

Although some aspects of QEC have been demonstrated with superconducting circuits [16–18, 40–44], trapped ions [54, 73, 74], and spins in solid-state systems [75–77], the control overhead prevents current-day experiments from getting to the heart of what QEC promises to achieve – extending the lifetime of quantum information stored in the system. This extension is quantified by the gain  $G$ , defined as the ratio between the coherence time of an actively error-corrected logical qubit and the best passive qubit encoding in the same system. The break-even point is reached at  $G = 1$ . A bosonic cat-code experiment [40] managed to achieve  $G = 1.1$ , but with a code that continuously shrinks to the vacuum state. Other experiments with various bosonic codes [41–43] and qubit-register codes [16–18, 44] have achieved  $G = 0.1 - 0.9$ .

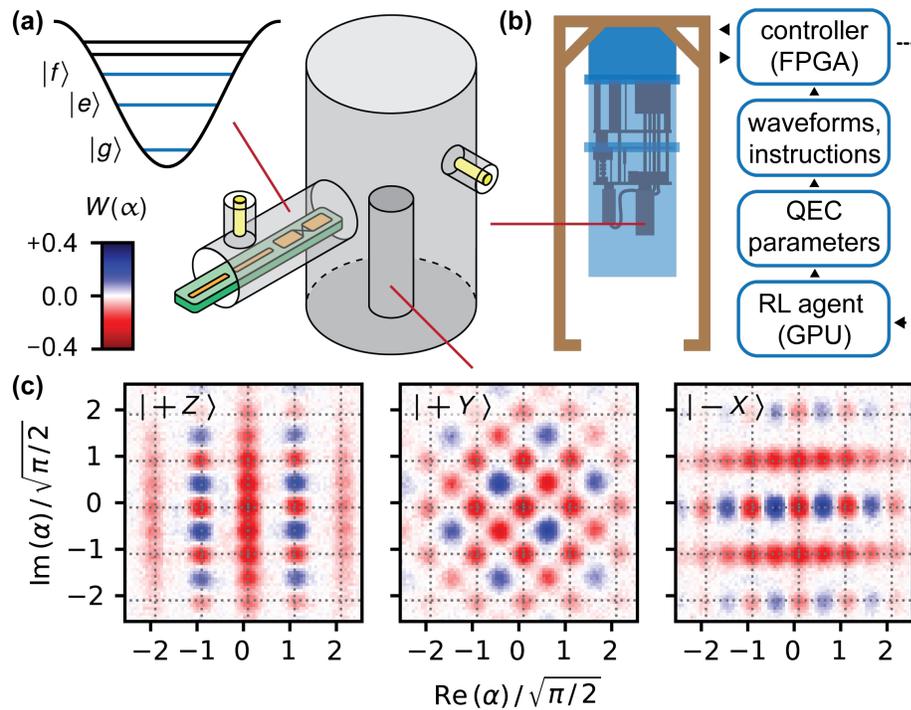
We demonstrate full code stabilization and error correction with gain  $G = 2.27 \pm 0.07$  using the Gottesman-Kitaev-Preskill (GKP) encoding [31] of a logical qubit into grid states of an oscillator. The QEC of this code was previously realized in superconducting circuits [42] and trapped ions [54]. In our work, similarly to [42], the oscillator is an electromagnetic mode of a superconducting cavity whose quantum state is manipulated using a transmon ancilla, see Fig. 2.1(a). Our system has an average relaxation and dephasing time of  $\bar{T}_1^t = 280 \mu s$  and  $\bar{T}_{2E}^t = 240 \mu s$  for the tantalum-based transmon [29], and  $\bar{T}_1^c = 610 \mu s$  and  $\bar{T}_2^c = 980 \mu s$  for the high-purity aluminum cavity [45]. We implement in this system a “trickle-down” QEC scheme

based on the proposals in Refs. [53, 54], which includes real-time classical processing and measurement-based feedback. We train the QEC circuit parameters in-situ with reinforcement learning (RL) [56, 68, 78], ensuring their adaptation to the real error channels and control imperfections of our system. At peak performance, the achieved lifetimes of logical Pauli eigenstates are  $T_X = T_Z = 2.20 \pm 0.03$  ms and  $T_Y = 1.36 \pm 0.03$  ms, and the logical Pauli error probabilities per QEC cycle are  $p_Y = (4.3 \pm 0.4) \times 10^{-4}$  and  $p_X = p_Z = (1.81 \pm 0.04) \times 10^{-3}$ . With such low logical error probabilities, we explore the QEC process on a previously inaccessible time scale of thousands of cycles, subjecting to scrutiny the standard assumptions of the theory of QEC, such as the stationarity of error rates and absence of leakage-induced correlations. Finally, we perform error-injection experiments to identify the major factors limiting logical performance and chart the path towards the next-generation logical qubit.

## 2.2 Engineering error correction

We now explain the principles of our experiment. Its core idea is to realize an artificial error-correcting dissipation that removes the entropy from the system in an efficient manner by prioritizing the correction of frequent small errors, while not neglecting rare large errors. This idea is illustrated in Fig. 2.2(a) for a cartoon system in which redundancy is achieved with only 4 orthogonal subspaces in total, where  $\mathcal{C}_0$  is the code subspace and  $\mathcal{C}_1 - \mathcal{C}_3$  are the error subspaces. In this example, the dissipation scheme #1 is maximally efficient from the perspective of entropy removal, since it corrects any error in a single step. Such an approach is taken by all qubit-register stabilizer codes, where measurement of the stabilizers, syndrome decoding, and recovery, when composed, realize a dissipation channel of high Kraus rank. Although this approach can also be applied to the oscillator grid code, see Chapter 2.7, its implementation entails large control overhead, which in practice might bring more errors than it is designed to correct. By contrast, the trickle-down dissipation scheme #2 has the capacity to correct all the same errors, but it is not able to do so in a single step. Importantly, the most probable small errors, corresponding to the error space  $\mathcal{C}_1$ , are still corrected in a single step. Owing to this simplification, such an approach reduces control overhead in the grid code, and therefore it was adopted in our work. The continuous-time version of approach #2 was also demonstrated for other bosonic codes in [43, 79].

The stabilizer generators of an *ideal* square grid code are  $S_0^x = D(l_S)$  and  $S_0^z =$



**Figure 2.1 | Experimental system.** (a) The sample consists of a superconducting aluminum cavity and a sapphire chip with a transmon circuit, readout resonator and Purcell filter. The electromagnetic mode of the cavity implements a harmonic oscillator, and  $\{|g\rangle, |e\rangle\}$  levels of the transmon are used as ancilla qubit to assist in oscillator QEC. (b) The sample is cooled in a dilution refrigerator and controlled with microwave and digital electronics. The QEC process is orchestrated by a field-programmable gate array (FPGA), and its parameters are optimized *in-situ* by a reinforcement learning agent implemented on a graphics processing unit (GPU). (c) Experimental Wigner functions of the Pauli eigenstates of a grid code with  $\Delta = 0.34$  measured after six QEC cycles.

$D(il_S)$ , where  $l_S = \sqrt{2\pi}$  is the length of a grid unit cell, and  $D(\alpha) = \exp(\alpha a^\dagger - \alpha^* a)$  is the displacement operator for an oscillator with creation and annihilation operators  $a^\dagger$  and  $a$ . Logical Pauli operators of the ideal code are defined as  $X_L = \sqrt{S_0^x}$  and  $Z_L = \sqrt{S_0^z}$ . The ideal codewords obey perfect translation symmetry in phase space and thus contain an infinite amount of energy. The finite-energy code is obtained by applying a normalizing envelope operator  $N_\Delta = \exp(-\Delta^2 a^\dagger a)$  to the ideal codewords, where  $\Delta$  parametrizes the code family that approaches the ideal code in the  $\Delta \rightarrow 0$  limit. In phase space, this parameter controls the extent of the codewords and the squeezing of their probability peaks. Our experimental Wigner functions of the codewords with  $\Delta = 0.34$  are shown in Fig. 2.1(c). The operators of the finite-energy code are obtained through the similarity transformation induced by the envelope operator [53], for example,  $S_\Delta^{x/z} = N_\Delta S_0^{x/z} N_\Delta^{-1}$ .

To realize an error-correcting dissipation channel  $\mathcal{R}_\Delta$  for the finite-energy code, there is at our disposal a single ancilla qubit and a classical controller. In principle, with such resources, it is possible to implement arbitrary quantum channels of Kraus rank  $2^M$  by recycling the ancilla  $M$  times and using feedback operations conditioned on the state of the classical  $M$ -bit memory of the controller [80, 81]. Here, we construct a rank-4 error correction channel as a composition of two rank-2 dissipators  $\mathcal{R}_\Delta = \mathcal{R}_\Delta^x \circ \mathcal{R}_\Delta^z$  that drive the system towards the +1 eigenspace of the finite-energy code stabilizers  $S_\Delta^{x/z}$ . A general rank-2 dissipation can be implemented as a unitary  $U_\emptyset$  that entangles the system with the ancilla qubit, followed by ancilla projective measurement with outcome  $b$ , and a classically-conditioned unitary  $U_b$ , see Fig. 2.2(b).

In our experiment, any unitary is compiled down to a set of primitive operations: qubit rotations around any equatorial axis  $R_\varphi(\theta) = \exp[-i(\theta/2)(\cos \varphi \sigma_x + \sin \varphi \sigma_y)]$  implemented as 32 ns Gaussian pulses with spectral corrections [82]; oscillator displacements  $D(\alpha)$  implemented as 40 ns Gaussian pulses; relatively slow conditional rotations  $\text{CR}(\theta) = \exp(i\theta \sigma_z a^\dagger a)$  implemented by simply waiting a certain amount of time under the dispersive coupling Hamiltonian  $H_d/\hbar = \chi \sigma_z a^\dagger a/2$  with  $\chi = 2\pi \times 46.5$  kHz; and virtual oscillator rotations  $R_V(\vartheta) = \exp(i\vartheta a^\dagger a)$  implemented dynamically on the field-programmable gate array (FPGA) in 448 ns. These primitives are used to construct a fast echoed conditional displacement gate  $\text{ECD}(\beta) = \sigma_x D(\sigma_z \beta/2)$  as shown in Fig. 2.2(b), whose speed  $\partial_t |\beta| = |\alpha| \chi$  is enhanced compared to the native interaction strength  $\chi$  by a large factor  $|\alpha|$  – magnitude of the intermediate displacement in phase space [42, 51].

Both rank-2 dissipators are then implemented as follows: the unitary  $U_\emptyset$  is de-

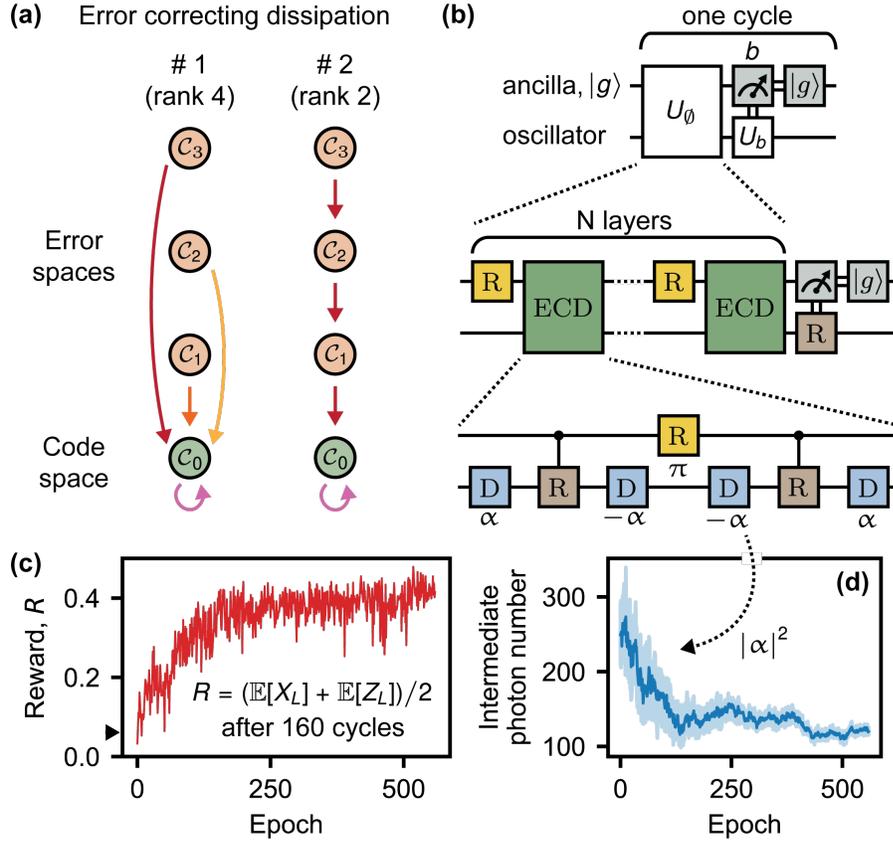
composed as a parametrized circuit consisting of layers of qubit rotations  $R_\varphi(\theta)$  and entangling ECD( $\beta$ ) gates, while the unitary  $U_b$  is realized as only a virtual rotation, see Fig. 2.2(b). The role of  $U_b$  is twofold: to implement switching between  $\mathcal{R}_\Delta^x$  and  $\mathcal{R}_\Delta^z$  by changing the quadrature of the oscillator by  $\pi/2$ , and to compensate for a spurious rotation due to the always-on dispersive coupling  $H_d$ . The role of  $U_\emptyset$  is to approximate the mapping of the finite-energy stabilizer onto the state of the ancilla together with autonomous back-action that pushes the state from the error spaces towards the code space. Several ansätze for decomposition of  $U_\emptyset$  were proposed in Ref. [53]. We adopt a modified version of the so-called small-big-small (SBS) protocol, named to reflect the relative amplitudes of the three conditional displacement gates that it contains:  $\vec{\beta} = l_S \times (i\Delta^2/2, 1, i\Delta^2/2)$ , see Chapter 6.3 for further details.

A single application of the resulting composite dissipator  $\mathcal{R}_\Delta$  realizes a QEC cycle; we refer to applications of constituent dissipators  $\mathcal{R}_\Delta^{x/z}$  as *even/odd cycles*. In our implementation, the duration of a QEC cycle is  $t_c = 2 \times 4.924 \mu\text{s}$ , which includes execution of unitary gates, ancilla measurements, and real-time processing and decision making by the controller.

## 2.3 Learning QEC circuit parameters

While the SBS ansatz and gate calibrations lead to a functioning QEC process, the highest level of performance cannot be achieved with a crude model of the system based on a few independently calibrated parameters – any such model will inevitably contain unrealistic assumptions. Some model inaccuracies and unknown control imperfections can be compensated by closed-loop optimization with direct feedback from the experimental setup. Previously, pulse-level optimization was successfully utilized to improve gate fidelities [83–86], but it was never applied to enhance the performance of QEC. Here, for the first time, we apply a real-time reinforcement learning agent to this task, as illustrated in Fig. 2.1(b). We use the proximal policy optimization (PPO) algorithm [56, 78], which was shown in simulations to outperform other approaches on high-dimensional problems with stochastic objective that arise in quantum control [55]. We parametrize the QEC circuit with  $P = 45$  parameters that include the amplitudes of various primitive pulses in the circuit decomposition, parameters of the ancilla reset, etc.

The training episodes begin with dissipative pre-cooling of the oscillator followed by feedback cooling to prepare the system ground state  $|g\rangle|0\rangle$ , see Chapter 2.7. Then, a logical Pauli eigenstate  $|+X\rangle$  or  $|+Z\rangle$  is initialized with a method

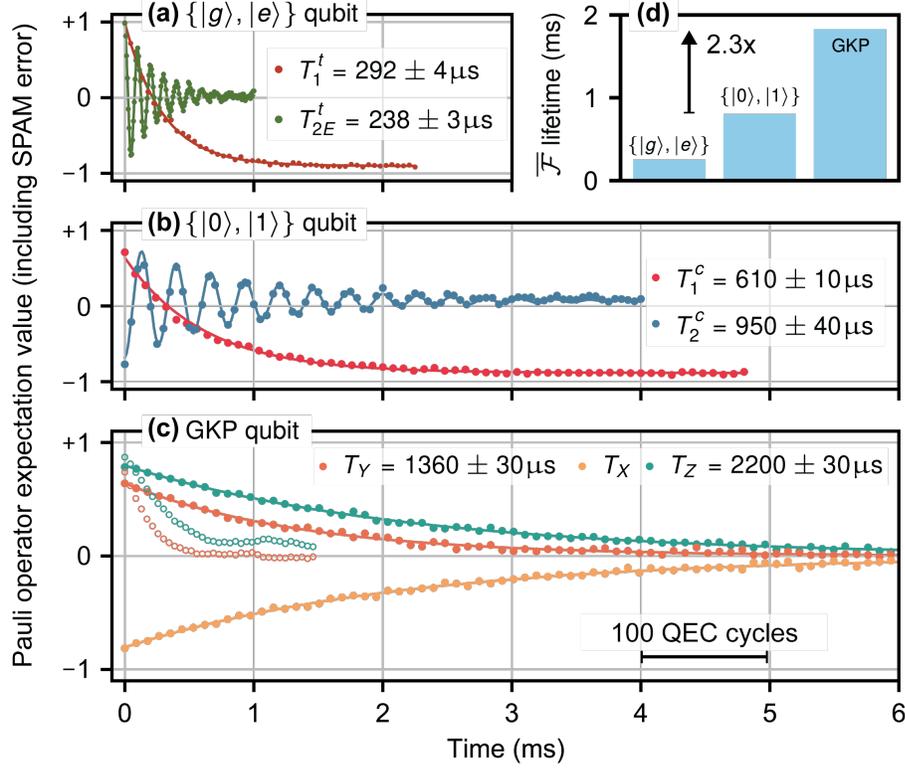


**Figure 2.2 | QEC implementation and optimization.** (a) Cartoon comparison of error-correcting dissipation channels. The standard dissipation #1 corrects any error in a single step, while the “trickle-down” dissipation #2 can be viewed as directional hopping between error spaces that eventually brings the quantum state to the code space  $\mathcal{C}_0$ . The colors of the arrows correspond to unique Kraus operators, whose number is equal to the channel rank. Higher-rank dissipation removes entropy more efficiently, but incurs larger control overhead. (b) Implementation of a general rank-2 channel on the oscillator using a single ancilla qubit. The unitary  $U_0$  is approximated as a parametrized circuit consisting of  $N$  layers of qubit rotations and oscillator conditional displacements. Each conditional displacement gate utilizes a large intermediate displacement of magnitude  $|\alpha|$  to enhance the gate speed. (c) Evolution of reward of the RL agent during the training. The black triangle indicates the start performance based on independent calibrations. Expectations of Pauli operators are taken in their respective eigenstates and include SPAM errors. (d) One realization of the learning trajectory of the intermediate photon number used to execute the big conditional displacement gate (“B” in the SBS circuit). Light blue shade shows the variance of the sampled parameter values during the training and dark blue line shows the mean.

from [51], and a candidate QEC protocol is run for  $T = 160$  cycles. We chose this duration to enhance the signal-to-noise ratio of the reward, similar to the technique used to sample randomized benchmarking cost functions [83–86]. At the end of the episode, the reward for the RL agent is obtained by measuring the logical Pauli operator  $X_L$  or  $Z_L$  (depending on the initial state), which provides a proxy for the logical lifetime. This logical measurement is done with one-bit phase estimation of the ideal-code Pauli operators [42, 87], and its fidelity is intrinsically limited to  $(1 + e^{-\pi\Delta^2/4})/2$  [72]. Although there exist methods of logical readout adapted to the finite code envelope [53, 54, 88], we use the phase estimation method to avoid biasing the RL agent towards a particular finite envelope size and to let it pick the optimal size given the error channels of our system.

By construction, the reward incentivizes the RL agent to find a QEC protocol that leads to the longest logical qubit lifetime. The typical evolution of the average reward during the training is shown in Fig. 2.2(c). The performance level indicated with a black triangle is achieved with independent calibrations of the system and control parameters, see Chapter 4. The RL agent significantly improves upon this baseline performance in two stages: in the first hundred training epochs, the agent corrects large errors in the initial parameter values, and in the subsequent few hundreds of epochs, it fine-tunes the circuit parameters to achieve the highest performance.

Several trends in the learning trajectories highlight the benefits of the model-free RL approach; we elucidate them in more detail in the Chapter 6.4. Here, we only highlight a single illustrative example. In our implementation of the ECD gate, there exists a nontrivial tradeoff between coherent and incoherent errors: the gate can be implemented faster by displacing the oscillator further in phase space, i.e., populating it with more intermediate photons, but this makes the gate more susceptible to high-order nonlinear effects [51]. Moreover, some choices of this intermediate photon number can result in a Stark shift of the ancilla into resonance with a spurious degree of freedom, e.g., a two-level defect [89, 90]. How these tradeoffs translate into logical qubit performance is difficult to model, but the RL agent can learn the optimal value of the large intermediate displacement without a model. As shown in Fig. 2.2(d), it chose to reduce the intermediate photon number, improving the performance of QEC at the cost of a much slower gate.



**Figure 2.3 | System coherence.** (a–c) For each qubit, we initialize Pauli eigenstates, let them evolve freely or under QEC for a variable amount of time, and measure the respective Pauli operators. The data for  $\{|g\rangle, |e\rangle\}$  and  $\{|0\rangle, |1\rangle\}$  qubits is fitted to amplitude damping and white-noise dephasing channel, and data for error-corrected GKP qubit is fitted to a Pauli channel. In (c), the  $|+X\rangle$  data is symmetrically reflected with respect to 0 for better visibility. Empty circles represent evolution in absence of QEC, when grid states decay towards vacuum. (d) Lifetime of average channel fidelity for these three qubits.

## 2.4 Observing QEC beyond break-even

After the training is finished, we pick the best performing QEC circuit for further characterization. Here, we focus on the ability of QEC to create a good quantum memory, i.e. to convert the effect of passage of time into an identity channel  $\mathcal{I} : \rho \rightarrow \rho$  that preserves *all* qubit states.

A metric quantifying the deviation of any quantum channel  $\mathcal{E} : \rho \rightarrow \mathcal{E}(\rho)$  from the identity is the *average channel fidelity*,  $\overline{\mathcal{F}}[\mathcal{E}] = \int d\psi \langle \psi | \mathcal{E}(|\psi\rangle\langle\psi|) | \psi \rangle$ , where the integral is over the uniform measure on the qubit state space, normalized so that  $\int d\psi = 1$ . In general, this fidelity decays over time in a nontrivial way, but to leading order it evolves as  $\overline{\mathcal{F}}(t) \approx 1 - \frac{1}{2}\Gamma t$ , where the decay rate  $\Gamma$  is equivalent to

an average decoherence rate of all pure states on the qubit Bloch sphere. Conveniently, it suffices to average across the six Pauli eigenstates alone [91], leading to an experimental procedure for extracting  $\Gamma$  that can be applied to any kind of qubit irrespective of its error channel. In Fig. 2.3, we show the results of such an experiment, conducted for three different qubit encodings in our system: the  $\{|g\rangle, |e\rangle\}$  subspace of the transmon, the  $\{|0\rangle, |1\rangle\}$  subspace of the oscillator, and grid code of the oscillator (with and without QEC).

Both the  $\{|0\rangle, |1\rangle\}$  and  $\{|g\rangle, |e\rangle\}$  qubits are subject to amplitude damping and white-noise dephasing channels, captured by their respective  $T_1$  and  $T_2$  times, with fidelity decay constant given by  $\Gamma = (1/T_1 + 2/T_2)/3$ . From the perspective of a quantum memory, the best uncorrectable physical qubit in our system is  $\{|0\rangle, |1\rangle\}$ , shown in Fig. 2.3(b), which achieves  $\Gamma_{\{01\}} = (800 \mu\text{s})^{-1}$ . The  $\{|g\rangle, |e\rangle\}$  qubit, shown for completeness in Fig. 2.3(a), only achieves  $\Gamma_{\{ge\}} = (250 \mu\text{s})^{-1}$ .

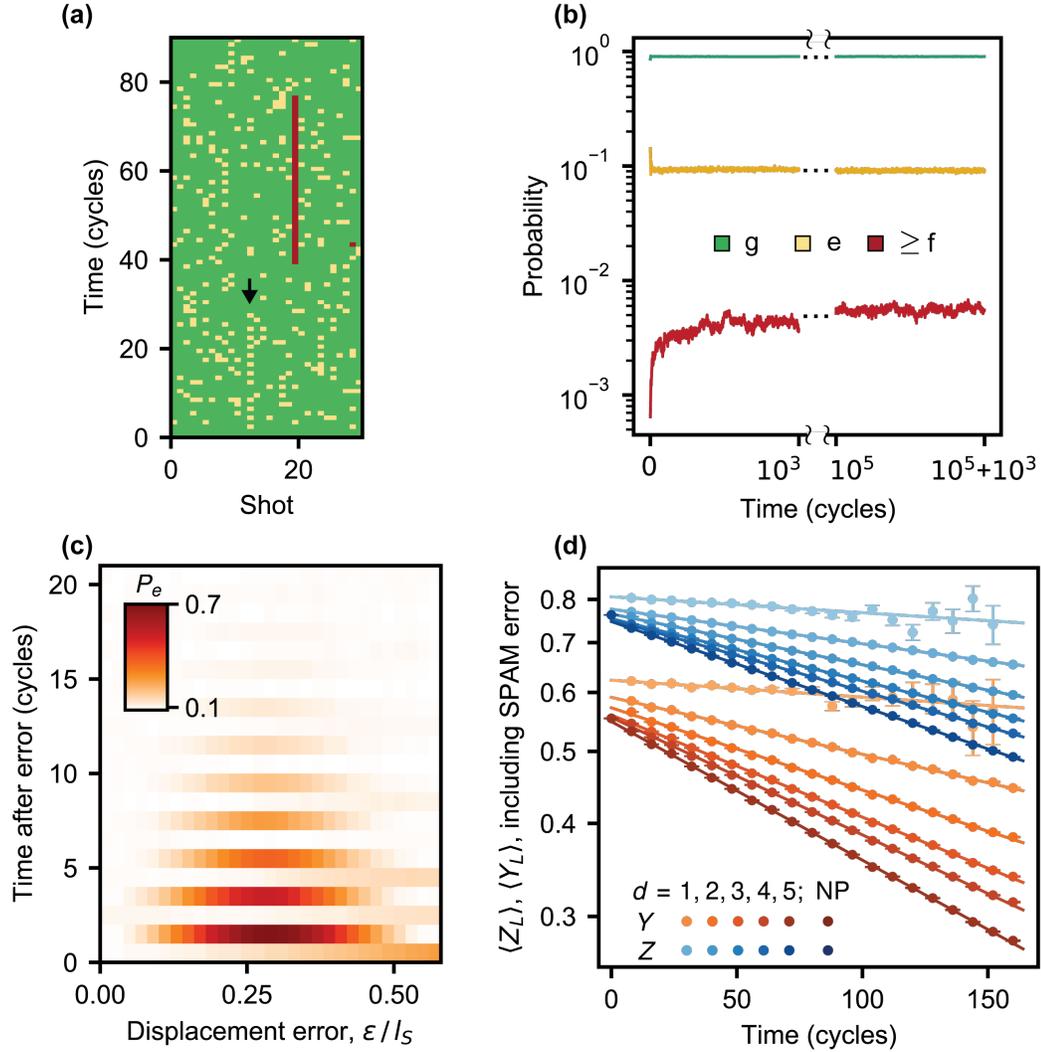
Higher excited states of the oscillator have shorter lifetime due to bosonic enhancement of spontaneous emission. Therefore, as with any QEC code, encoding a qubit using grid states incurs an immediate penalty in the fidelity decay rate. Moreover, this natural decay, shown in Fig. 2.3(c) with empty circles, takes the grid states outside the logical manifold and eventually towards the vacuum state  $|0\rangle$ .

Our error correcting dissipation stabilizes the grid code manifold and, together with naturally occurring dissipation, leads to a logical Pauli channel within this manifold, with the lifetimes of logical Pauli eigenstates of  $T_X = T_Z = 2.20 \pm 0.03$  ms and  $T_Y = 1.36 \pm 0.03$  ms. Under the Pauli channel, the fidelity decay constant is given by  $\Gamma = (1/T_X + 1/T_Y + 1/T_Z)/3$ , which in our experiment amounts to  $\Gamma_{\text{GKP}} = (1.82 \text{ ms})^{-1}$ .

The principal metric characterizing the quality of QEC from the perspective of quantum memory is the coherence gain of an actively error-corrected logical qubit over the best passive qubit encoding. In our experiment, the highest achieved gain is  $G = \Gamma_{\{01\}}/\Gamma_{\text{GKP}} = 2.27 \pm 0.07$ , confidently beyond break-even.

## 2.5 QEC process characterization

Having characterized the logical qubit as a quantum memory, we next examine the properties of the QEC process. Ancilla measurement outcomes, referred to as syndromes, inform us which stochastic path the QEC process has taken in each cycle. In Fig. 2.4(a) we show a (statistically unrepresentative) sample of these outcomes



**Figure 2.4 | Analysis of error syndromes.** (a) A sample of ancilla measurement outcomes during QEC. The  $e$  outcome (yellow) indicates correction of physical errors. The black arrow points to a syndrome string of the type  $eg/eg/\dots$  likely left by a large error in one oscillator quadrature. Red indicates transmon leakage out of the  $\{|g\rangle, |e\rangle\}$  qubit subspace. (b) Average probability of each measurement outcome as a function of time. After correcting state initialization errors, QEC settles into a steady state that persists for at least a hundred thousand cycles. (c) Probability of  $e$  outcome as a function of time after injecting position displacement errors of varying amplitude. Since a logical gate in the ideal code comprises a displacement of amplitude  $l_S/2$ , a displacement of amplitude  $l_S/4$  makes a large-distance error, which takes several cycles to correct with our low-rank QEC channel. (d) Decay of Pauli eigenstates after eliminating by post-selection the experimental shots with strings of  $\geq d$  consecutive  $e$  outcomes in the same-quadrature cycles. Data for  $X_L$  eigenstates is not shown; it is expected to behave similarly to  $Z_L$ . The improvement in lifetime indicates that  $e$  outcomes are indeed correlated with occurrence of errors. NP in the legend stands for “no post-selection”.

that comprise trajectories of different experimental shots. Such a dataset contains an immense amount of information about the QEC process, not available in previous experiments with the grid code QEC [42, 54].

To interpret this dataset, we adopt here a simplified model of trickle-down dissipation such as depicted in Fig. 2.2(a), which captures the essence of our QEC process. The caveats of this model and the exact Kraus decomposition of our QEC circuit are provided in Chapter 6.2. In this simplified model, the  $g$  outcome indicates that the state was projected onto the code space, while an  $e$  outcome indicates that the state was transferred one level down the error hierarchy, partially or completely correcting an error.

From the dataset in Fig. 2.4(a), we observe that the vast majority of outcomes are  $g$  (green), which means that errors are rare. The stochastic pattern of  $e$  outcomes (yellow) reflects randomly occurring errors. Most errors are small and, when corrected, leave single isolated  $e$  outcomes. An example syndrome string likely generated by a large error in *one* quadrature is indicated with an arrow: it has a characteristic  $eg/eg/\dots$  pattern. We also observe isolated ancilla leakage events (red). Leakage to  $|f\rangle$  is reset in the same cycle with high probability. Sometimes, leakage persists for multiple cycles (streak of red), due to the transmon escaping to a state higher than  $|f\rangle$ , which is not addressed in our reset scheme.

The average probability of each outcome as a function of time is shown Fig. 2.4(b), where the process starts from a  $|+X\rangle$  state. After about 10 cycles of initial state correction, the process settles into a dynamical equilibrium which persists for at least a hundred thousand cycles (the longest measured here) without any notable increase of the error rates over time. Detailed analysis reveals that the QEC process is nearly stationary, with residual deviations from stationarity caused by the transmon leakage to states higher than  $|f\rangle$  at a rate  $1.3 \times 10^{-4}$  per cycle, see Chapter 6.6.

In this dynamical equilibrium, physical errors excite the quantum state out of the code space with probability  $p_{\text{err}} = 0.13 \pm 0.02$  per QEC cycle, as deduced from the statistics of syndrome outcomes. The competition between physical errors and error-correcting dissipation results in a “thermal” distribution across the subspaces with probability  $\langle \Pi_0 \rangle = 0.82 \pm 0.02$  of occupying the code space, see Chapter 2.7. Having  $p_{\text{err}} \ll 1$  justifies the use of low-rank error-correcting dissipation in our system, which is sufficient to prevent physical errors from accumulating and causing logical errors. At the highest achieved QEC gain, the logical Pauli error probabilities per QEC cycle are  $p_Y = (4.3 \pm 0.4) \times 10^{-4}$  and  $p_X = p_Z = (1.81 \pm 0.04) \times 10^{-3}$ . By comparing the total logical error probability,  $p_X + p_Y + p_Z$ , to the physical error

probability  $p_{\text{err}}$ , we conclude that 97% of the errors are successfully corrected by our process.

Since rare large errors require several cycles to be corrected, the QEC process is weakly time-correlated with a correlation length of  $3.9 \pm 0.1$  cycles, see Chapter 6.6. To understand these correlations, in Fig. 2.4(c) we inject displacement errors along the position quadrature and monitor the syndromes that they produce as a function of time. Such errors leave traces of  $e$  outcomes in proportion to their distance to the closest logical operation. For example, a displacement of length 0, equivalent to a logical identity, leaves no syndrome trace; a displacement of length  $l_S/2$  is close to a logical bit-flip of the finite-energy code, and hence it leaves only a small syndrome trace; on the other hand, a midway displacement of length  $l_S/4$  makes a large-distance error which takes the longest time to correct with a low-rank dissipator, generating a lasting trace of  $e$  outcomes.

This displacement error injection experiment confirms that errors indeed generate the  $e$  syndromes, but do these syndromes herald the occurrence of errors? To verify this, we perform post-selection of trajectories with different syndrome patterns. In particular, we discard trajectories that have  $\geq d$  consecutive  $e$  outcomes in the same-quadrature cycles, with resulting post-selected decay of Pauli eigenstates shown in Fig. 2.4(d). In the case  $d = 5$ , post-selection eliminates rare large-distance errors and improves the fidelity lifetime only by a factor 1.2, but at the cost of rejection probability of  $7 \times 10^{-4}$  per cycle. On the other hand, in the case  $d = 1$ , post-selection eliminates relatively frequent small errors that are close to identity, as well as rare large uncorrectable errors that are close to a logical operation. It is because of the latter that the fidelity lifetime in this setting improves by a factor 6.3, but with a more severe rejection probability of  $6 \times 10^{-2}$  per cycle. These favorable post-selection results indicate that such a method can be used for probabilistic preparation of high-fidelity logical states, including the magic states required for universal quantum computing [92], which is left for future investigation.

## 2.6 Conclusion and outlook

In this work, we used real-time error correction to realize a fully stabilized logical qubit whose lifetime is more than doubled compared to the best passive qubit encoding in the system, marking the transition of QEC from proof-of-principle studies to a practical tool for enhancing quantum memories. Our work improves upon previous QEC experiments, which do not protect the logical identity operator  $I_L$

[40], protect only one of the logical Pauli operators  $X_L$  or  $Z_L$  [79, 93, 94], implement correction in post-processing [16–18], require post-selection [95], and do not reach break-even [16–18, 41–43]. Instrumental for this achievement, among other factors, was the adoption of a model-free learning framework, improved fabrication technique for the ancilla transmon, and a novel grid-code QEC protocol.

Performing additional experiments, we identified the core challenges that need to be addressed to ensure future progress of grid-code QEC. In particular, by studying long-time system stability, we found that occasional collapses of the logical performance are strongly correlated with appearance of spurious degrees of freedom in the system. Their resonant interaction with the Stark-shifted transmon ancilla degrades the fidelity of our operations, see Chapter 6.10. In the short term, this effect could be mitigated by adopting a tunable ancilla and periodically re-training the QEC circuit to find better spectral locations. In the long term, the behavior of these defects needs to be understood, as they pose even greater danger for scaled-up quantum devices [16–18].

In addition, we expect that considerable enhancement can be gained by tailoring the QEC process not only to error channels of the oscillator, but also those of the ancilla. Our QEC circuit is fault-tolerant with respect to ancilla phase-flip errors by design [53]. With the transmon ancilla used here, the sensitivity of the logical lifetime to ancilla phase flips is 65 times smaller than to ancilla bit flips, as found with noise injection experiments, see Chapter 6.9. Future development should incorporate robustness against ancilla bit flips, either through path-independent control [47, 96] or by adopting a biased-noise ancilla [97].

## 2.7 Additional information

### QEC of the ideal grid code

To understand error correcting properties of the ideal code, consider an error channel  $\mathcal{E}$  decomposed in the displacement basis. An ideal grid code with code projector  $\Pi_0$  satisfies Knill-Laflamme conditions [69]  $\Pi_0 D^\dagger(\varepsilon_\alpha) D(\varepsilon_\beta) \Pi_0 \propto \delta(\varepsilon_\alpha - \varepsilon_\beta) \Pi_0$  for all errors in a correctable set  $E_+ = \{D(\varepsilon) : |\operatorname{Re}(\varepsilon)|, |\operatorname{Im}(\varepsilon)| < l_S/4\}$ . A displacement error of amplitude  $\varepsilon$  creates an error state  $|\psi_\varepsilon\rangle = D(\varepsilon)|\psi\rangle$ , where  $|\psi\rangle$  is any state from the code space. Since a displaced grid state is still translationally invariant, it remains an eigenstate of the ideal code stabilizers, and the phase of its eigenvalue encodes a continuous error syndrome:  $S_0^z |\psi_\varepsilon\rangle = \exp(2il_S \operatorname{Re}[\varepsilon]) |\psi_\varepsilon\rangle$

and  $S_0^x |\psi_\varepsilon\rangle = \exp(-2il_S \text{Im}[\varepsilon]) |\psi_\varepsilon\rangle$ . Error correction of an ideal grid code can be done in a similar manner to any stabilizer code: first, measure the stabilizers to obtain the error syndrome, which here corresponds to phase estimation of  $S_0^{x/z}$  that yields the error amplitude  $\varepsilon$ . This step projects the state onto one of the orthogonal error spaces. Then, apply the recovery operation, here a simple displacement  $D(-\varepsilon)$ , to correct the error. Such procedure realizes an artificial dissipation  $\mathcal{R}$  of an infinite rank which corrects any error from  $E_+$  in a single cycle,  $(\mathcal{R} \circ \mathcal{E})(\rho) \propto \rho$ , analogously to the cartoon high-rank dissipation in Fig. 2.2(a). In contrast to this approach, our experiment realizes low-rank dissipation that asymptotically satisfies  $([\mathcal{R}]^{n \rightarrow \infty} \circ \mathcal{E})(\rho) \propto \rho$ .

## Dissipative cooling to vacuum

We utilize the dissipation engineering framework [98] to design fast cooling of the oscillator to vacuum state  $|0\rangle$  in the weak-coupling regime where previous known cooling methods [99] fail; we also expect this novel method to be applicable to cooling of trapped ions [54]. Similarly to error-correcting dissipation, we realize this cooling as a composition of two rank-2 channels that shrink the oscillator state in the orthogonal quadratures. The unitary  $U_\theta$  in this case is realized as a three-layer circuit obtained from first-order Trotter decomposition of  $U = \exp[-i\varepsilon(a\sigma_+ + a^\dagger\sigma_-)]$ , where  $\varepsilon \ll 1$  controls the cooling rate. This unitary swaps the excitations of the oscillator into the ancilla, which is reset in every cycle. The duration of one full cooling cycle (including both quadratures) is  $t_c = 2 \times 3.38 \text{ us}$ . With  $\varepsilon = 0.4$ , we achieve cooling at a rate 20 times faster than natural energy damping rate of the oscillator. In our experiment, 25 full cycles of such a dissipative cooling are then followed with a feedback cooling protocol adapted from [40] to remove any residual thermal population. See Chapter 4.6 for more details.

## Reinforcement learning implementation

The QEC circuit is parametrized with a vector  $\vec{p}$ . Instead of optimizing  $\vec{p}$  directly, the RL agent learns parameters of the probability distribution from which  $\vec{p}$  is stochastically sampled during the training to ensure adequate exploration of parameter space. To this end, we use a factorized multivariate Gaussian distribution  $\mathcal{N}(\vec{\mu}, \vec{\sigma})$  with mean  $\vec{\mu}$  and covariance matrix  $\text{diag}[\vec{\sigma}]^2$ . To capture the pattern of relations between different components of  $\vec{p}$ , the mean and covariance are represented as parametrized functions  $\vec{\mu}(\theta)$  and  $\vec{\sigma}(\theta)$  of common hidden variables  $\theta$ .

In this work,  $\vec{\mu}$  and  $\vec{\sigma}$  are produced at the output of a neural network with two fully connected layers of 50 and 20 rectifier linear unit (ReLU) neurons. Starting with initial vector of parameters  $\vec{\mu}_i$  found with independent calibrations, during the course of learning the agent gradually deforms the distribution and localizes it on the new vector  $\vec{\mu}_f$ , the final result of the optimization. Typically, as it proceeds, the agent also reduces the covariance of the distribution to have a finer control over the mean. These features of learning are observed in the example evolution of one component of  $\vec{p}$  in Fig. 2.2(d). During one training epoch, we evaluate 10 QEC circuit candidates with 300 episodes (i.e. experimental shots) per candidate. The collected information is used to update the neural network parameters  $\theta$  according to the PPO algorithm, which completes the epoch. One epoch takes approximately 16 seconds, with the majority of time spent on re-compilation of FPGA instruction sequences and its re-initialization. See Chapter 5.2 for more details.

### Steady state of the QEC process

We perform Wigner tomography of the logical states after varying duration of the QEC process, reconstruct the density matrix, and from its spectral decomposition extract the expectation value of the code projector  $\langle \Pi_0 \rangle = 0.825 \pm 0.003$ , where error bar represents the standard deviation with respect to different process durations of 100, 200, 400, and 800 cycles. In addition to the code space, only one error space is occupied in the steady state with an appreciable probability of  $0.170 \pm 0.005$ . The logical decoherence within this error space happens at the same rate as within the code space. For more details, see Chapter 6.8.

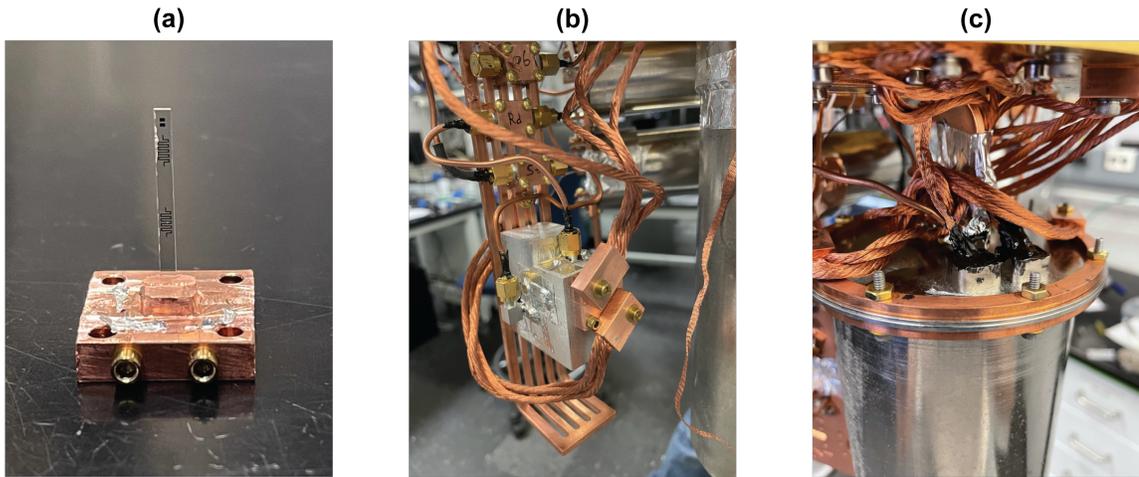
The expectation value of the code projector in the steady state can be estimated independently, using the statistics of syndrome outcomes. Under the approximations discussed in Chapter 6.5, the probability that a syndrome string of length  $2n$  consists only of  $g$  outcomes asymptotically approaches  $\langle \Pi_0 \rangle (1 - p_{\text{err}})^{n-1}$  for large  $n$ . Using this method, we extract  $\langle \Pi_0 \rangle = 0.81 \pm 0.02$  and  $p_{\text{err}} = 0.13 \pm 0.02$ . The error bar in this case represents the accuracy of the model for the string probability, which is valid to first order in  $p_{\text{err}}$ . The value of  $\langle \Pi_0 \rangle$  quoted in the main text is the average of the two methods. Constructing a detailed error budget of the aggregate error probability  $p_{\text{err}}$  based on the system-level simulation of the known error processes is an avenue left for future work.

# Experimental setup and sample

This chapter briefly describes how the sample was fabricated and what measures were taken to isolate it from the environment yet maintain the controllability. The techniques for sample design, assembly, thermalization, shielding, and wiring of the control electronics described here, are fairly standard in the field of superconducting quantum circuits. The ancilla chip fabrication technique is new (it was recently introduced in Ref. [29]), and the corresponding Chapter 3.2 was contributed by my colleague S. Ganjam who fabricated the chip.

## 3.1 Assembly

Our system design follows the hybrid planar-3D circuit QED architecture developed in [46]. The storage oscillator is realized as an electromagnetic mode hosted by a seamless superconducting coaxial stub cavity made of high-purity aluminum and treated with a chemical etch to improve surface quality. This is the same physical cavity as used in [47], although with lower coherence time due to aging during the storage time of  $\sim 2$  years. The cavity is anchored to a copper bracket inside a cryoperm shield. The ancilla chip is inserted in a tunnel waveguide that connects to the storage cavity, and is secured at one end with a copper clamp. Thermalizing copper braids run from the clamp to the base plate of the dilution refrigerator, see Fig. 3.1.



**Figure 3.1 | Sample assembly.** (a) Clamped ancilla chip. (b) Thermalization. (c) Shielding.

## 3.2 Ancilla chip fabrication

The ancilla chip contains a transmon qubit, a stripline readout resonator, and a stripline bandpass Purcell filter. The resonators and transmon are tantalum-based devices, with Josephson junction of the transmon made of a small aluminum section; they are fabricated with a process similar to [29]. Adopting a tantalum-based platform results in improved qubit coherence relative to an all-aluminum platform; however, the reasons for this are still under active investigation. Possible explanations include, but are not limited to: 1) The corrosion resistance allows for the use of rigorous acid-based cleaning techniques to be employed during the fabrication process that improves surface dielectric quality and minimizes the presence of organic residues; 2) The high melting point of tantalum allows for deposition to occur at higher temperatures, where atomic mobility is high enough to enable epitaxial film growth with a high degree of crystalline order; 3) Tantalum has a higher superconducting transition temperature than aluminum, which may lead to increased resistance to quasiparticle loss.

We use a C-plane sapphire wafer produced using the heat exchanger method (HEM), as it was shown to have smaller dielectric loss [100]. The wafer was initially cleaned with a piranha solution (2 : 1  $\text{H}_2\text{SO}_4$  :  $\text{H}_2\text{O}_2$ ) for 20 minutes and rinsed with DI water. The wafer was then annealed at  $1200^\circ\text{C}$  in an oxygen-rich environment for 1 hour. After cooling down to room temperature, the wafer was immediately transferred to a sputtering system for tantalum deposition. 150 nm of tantalum was deposited by DC magnetron sputtering with the substrate temperature being held at  $800^\circ\text{C}$ . The Purcell filter, readout resonator, and transmon pads were subtractively patterned using a positive photoresist mask and reactive ion etching. After tantalum patterning, the Josephson junction was patterned using electron-beam lithography and defined using the Dolan bridge method. The junction was deposited using electron-beam evaporation of aluminum at 2 angles with an interleaved static oxidation step to construct the tunnel barrier. Liftoff was performed in NMP heated to  $90^\circ\text{C}$ , followed by sonication in acetone, isopropanol, and DI water. Finally, the wafer was protected with a layer of photoresist before dicing into individual chips, followed by additional cleaning with NMP, acetone, and isopropanol to remove the protective photoresist.

<b>Cavity mode</b>	Frequency 1st order dispersive shift 2nd order dispersive shift Kerr nonlinearity Relaxation Dephasing	$\omega_c = 2\pi \times 4.479$ GHz $\chi = 2\pi \times 46.5$ kHz $\chi' = 2\pi \times 5.8$ Hz $K = -2\pi \times 4.8$ Hz $\bar{T}_1^c = 606 \pm 10$ us $\bar{T}_2^c = 980 \pm 30$ us
<b>Ancilla transmon</b>	Frequency Anharmonicity Relaxation Equilibrium population Dephasing (Ramsey) Dephasing (Echo)	$\omega_t = 2\pi \times 5.921$ GHz $\alpha = -2\pi \times 222$ MHz $\bar{T}_1^t = 280 \pm 30$ us $\bar{n}_{\text{th}}^t = 0.043 \pm 0.008$ $\bar{T}_{2R}^t = 62 \pm 5$ us $\bar{T}_{2E}^t = 238 \pm 8$ us
<b>Readout resonator</b>	Frequency Dispersive shift Coupling strength Internal loss	$\omega_r = 2\pi \times 9.107$ GHz $\chi_{qr} = 2\pi \times 0.60$ MHz $\kappa_{r(c)} = 2\pi \times 0.47$ MHz $\kappa_{r(i)} = 2\pi \times 0.03$ MHz

**Table 3.1 | Measured system parameters.** For transmon parameters ( $\bar{T}_1^t, \bar{T}_{2R}^t, \bar{T}_{2E}^t, \bar{n}_{\text{th}}^t$ ) and cavity parameters ( $\bar{T}_1^c, \bar{T}_2^c$ ) that appreciably fluctuate in time, we provide the mean and standard deviation over a week-long period. The definition of the Hamiltonian parameters can be found in Chapter 4.2.

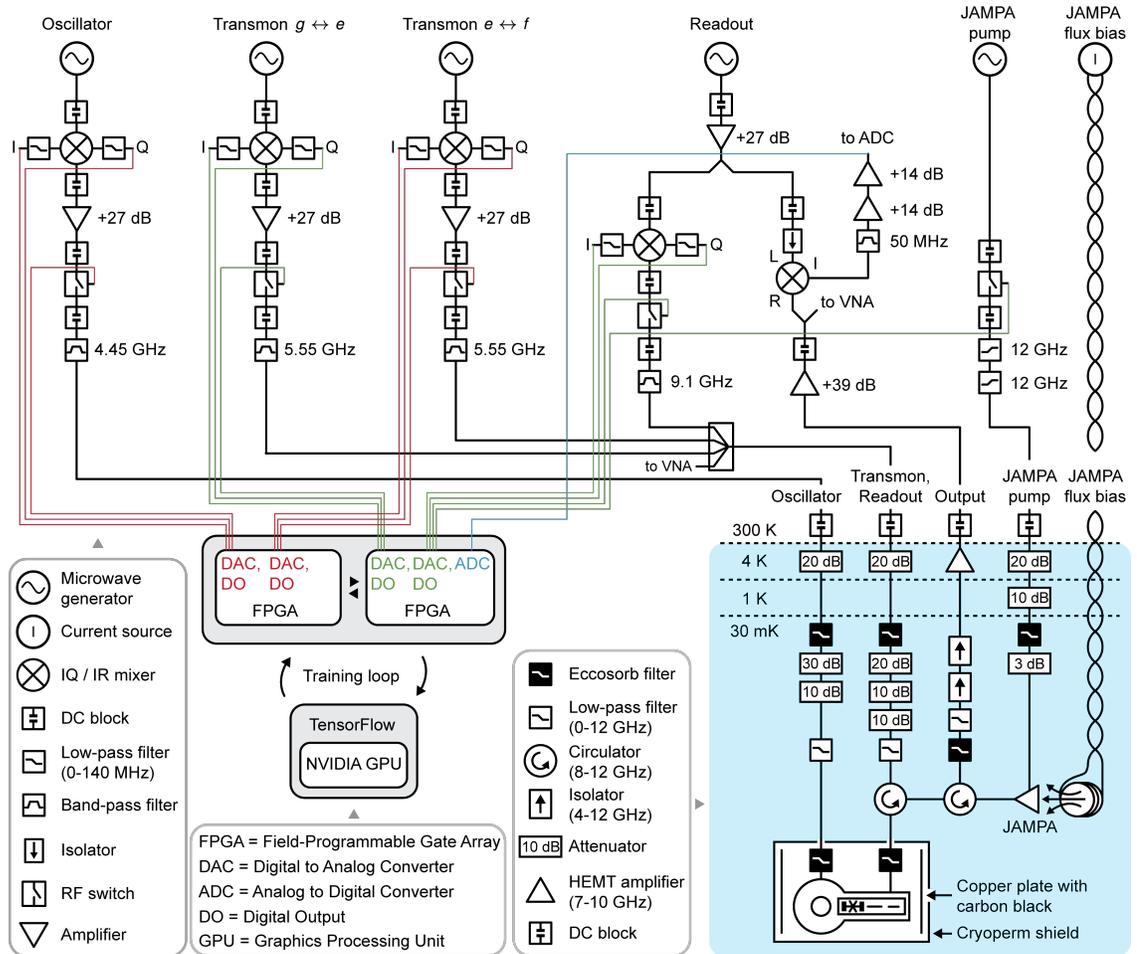
### 3.3 System parameters

The measured parameters of this system are summarized in Table 3.1.

### 3.4 Control wiring

As shown in Fig. 3.2, the quantum system is controlled by a computer (VPXI-ePC) that hosts two control cards (X6-1000M) from Innovative Integration. Each card integrates digital-to-analog converters (DAC), analog-to-digital converters (ADC), digital inputs and outputs (DIO), and a Xilinx VIRTEX-6 field-programmable gate array (FPGA). This controller was developed in [40] and used in prior bosonic QEC experiments [40, 41]. The baseband control signals are sampled from the DACs at 500 MS/s rate with 16-bit resolution and upconverted to the oscillator, qubit, and readout frequencies through single-sideband modulation of the local oscillators (Agilent N5183A). After amplification, the signals are gated with fast RF switches (9 ns rise time, 5 ns fall time) and filtered before entering the dilution refrigerator. The signals are further attenuated and filtered in the cryogenic environment.

A crucial component of the filtering is the eccosorb CR-110 infrared absorber filter [101] located inside the cryoperm shield, and the copper plate, painted with stycast epoxy mixed with black carbon powder, that wraps around the sample. On the output side, the reflected readout signal is amplified at 30 mK stage with a near-quantum-limited Josephson array-mode parametric amplifier (JAMPA) [55], followed at 4 K stage by a low-noise HEMT amplifier. Upon further amplification at 300 K stage and down-conversion to 50 MHz, the readout signal is digitized, demodulated, and integrated with a filter function to obtain  $I$  and  $Q$  quadratures. Their values are compared to the decision boundaries  $I_{th}$  and  $Q_{th}$  to obtain two bits of information  $s_0 = \Theta(I - I_{th})$  and  $s_1 = \Theta(Q - Q_{th})$ , where  $\Theta$  is the Heaviside step function. This allows to classify the measurement outcome as “ $g$ ” if  $s_0 = 0$ , “ $e$ ” if  $(s_0, s_1) = (1, 0)$ , and “ $f$ ” if  $(s_0, s_1) = (1, 1)$ . The bits  $s_0$  and  $s_1$  are redistributed to all control cards which run independent but synchronized control flows that include conditional branching on these bits. Further details of the readout subroutine are described in Chapter 4.3.



**Figure 3.2 | Experimental setup.** For simplicity, the diagram omits DC ground connections of all active components, and attenuators placed at different locations at 300 K stage to ensure power levels within specs for amplifiers and switches.

# Calibration and characterization

This chapter describes basic system calibration and characterization techniques. The calibrations of the primitive pulses and of the readout, described in Chapters 4.1, 4.3, are well-known and widely used throughout the superconducting circuits community. The calibrations of the Hamiltonian parameters, described in Chapter 4.2, rely on a mix of established techniques in the strong-coupling regime of circuit QED, and relatively new methods developed for the weak-coupling regime in Refs. [42, 51]. The calibration of the ECD gate, described in Chapter 4.4, is also based on these works, with certain modifications related to the training of the gate parameters. The novel method of dissipative oscillator cooling, described in Chapter 4.6, was theoretically proposed by my colleague B. Royer, first tested in an experiment by B. Brock, and its improved version was adopted for the present experiment and is published here for the first time.

## 4.1 Primitive pulses

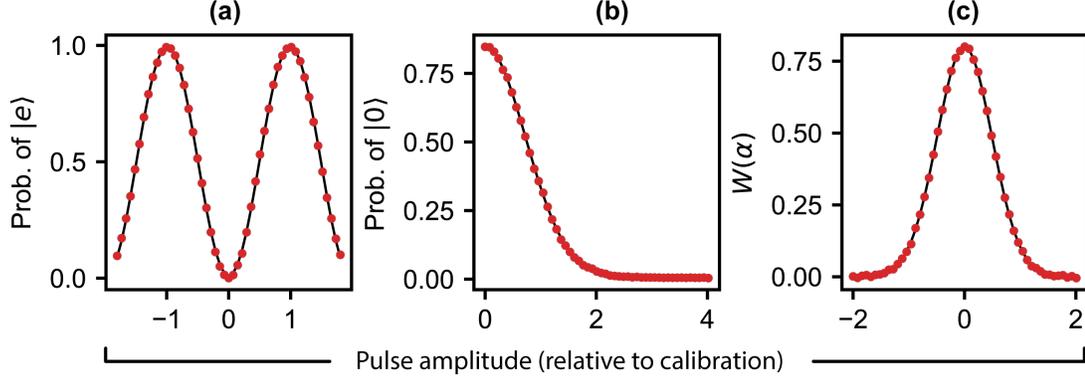
### Qubit rotation

The waveform for transmon  $g \leftrightarrow e$  and  $e \leftrightarrow f$  rotations is a Gaussian with  $\sigma = 8$  ns and symmetric chop at  $2\sigma$ . The pulse amplitude is calibrated with a standard amplitude Rabi experiment, shown in Fig. 4.1(a). We find that finite negative detuning of a few MHz is needed to maximize the Rabi contrast in both cases. In a similar manner, we calibrate a selective square pulse of duration  $2\pi/\chi \approx 22 \mu\text{s}$  that performs  $g \leftrightarrow e$  rotations conditioned on the oscillator in  $|0\rangle$ .

### Oscillator displacement

The waveform for oscillator displacements is a Gaussian with  $\sigma = 10$  ns and symmetric chop at  $2\sigma$ . It is calibrated in several steps, refining the accuracy at each step. First, before the precise value of  $\chi$  is known, we use a rough calibration by creating a coherent state of unknown amplitude  $\alpha$  and measuring the probability of  $|0\rangle$  via a selective qubit pulse, with the results shown in Fig. 4.1(b). Fitting the data to  $P(0) = e^{-|\alpha|^2}$  allows us to calibrate the DAC amplitude for displacement of  $|\alpha| = 1$ . This first-stage calibration enables us to use active oscillator cooling, see Chapter 4.6, which is important for the next calibration step that relies on

a vacuum state. Next, after determining the value of  $\chi$  (using number-resolved qubit spectroscopy, see Chapter 4.2), we measure the Wigner function of vacuum  $W(\alpha) = (2/\pi)e^{-2|\alpha|^2}$  and adjust the DAC amplitude calibration to obtain the variance of  $1/4$ , with the results shown in Fig. 4.1(c). We find that these two calibrations typically agree within 2%.



**Figure 4.1 | Calibration of primitive pulses.** (a) Amplitude Rabi experiment to calibrate qubit rotations. (b) First-step calibration of displacement: probability of  $|0\rangle$  in a coherent state,  $P(0) = e^{-|\alpha|^2}$ . (c) Second-step calibration of displacement: Wigner function of vacuum,  $W(\alpha) = (2/\pi)e^{-2|\alpha|^2}$ .

## 4.2 Hamiltonian parameters

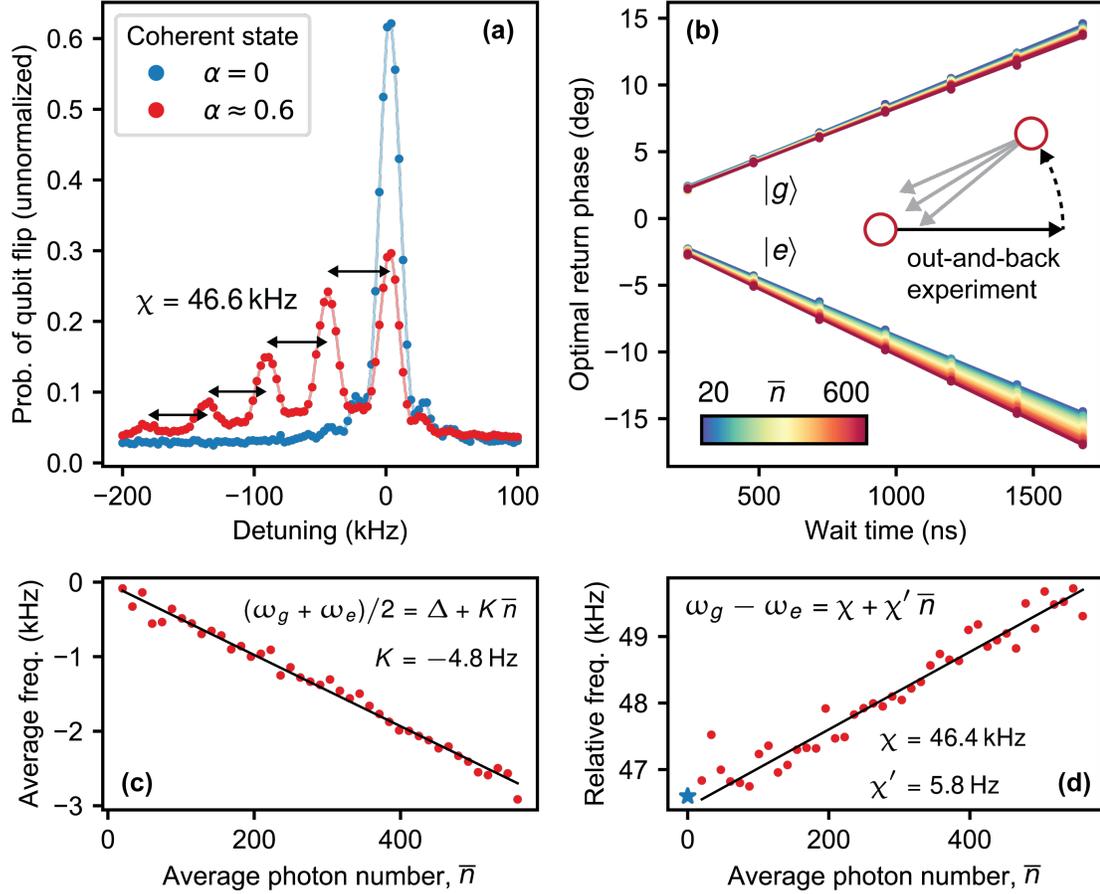
Our system is well described with the following Hamiltonian

$$H/\hbar = \Delta (a^\dagger a) + \frac{1}{2}\chi (a^\dagger a) \sigma_z + \frac{1}{2}K (a^\dagger a)^2 + \frac{1}{4}\chi' (a^\dagger a)^2 \sigma_z, \quad (4.1)$$

where  $\Delta$  is the oscillator frequency detuning in the chosen rotating frame,  $\chi$  is the dispersive shift,  $\chi'$  is the second-order dispersive shift, and  $K$  is the Kerr nonlinearity.

We calibrate  $\chi$  with number-resolved qubit spectroscopy [102] in the presence of a coherent state of small amplitude  $\alpha \approx 0.6$  in the oscillator. The spectroscopy data, shown in Fig. 4.2(a) in red, is fitted to a 5-component equal-spacing mixture of the spectroscopy lineshapes with the oscillator in vacuum, shown in blue, which results in  $\chi = 46.6$  kHz. After additionally performing the cavity mode spectroscopy (data not shown), we set the LO frequency to work in the rotating frame with  $\Delta = 0$ .

After calibrating the displacement amplitude, as described in Chapter 4.1, we perform an out-and-back experiment [51] to determine the higher order nonlinear-



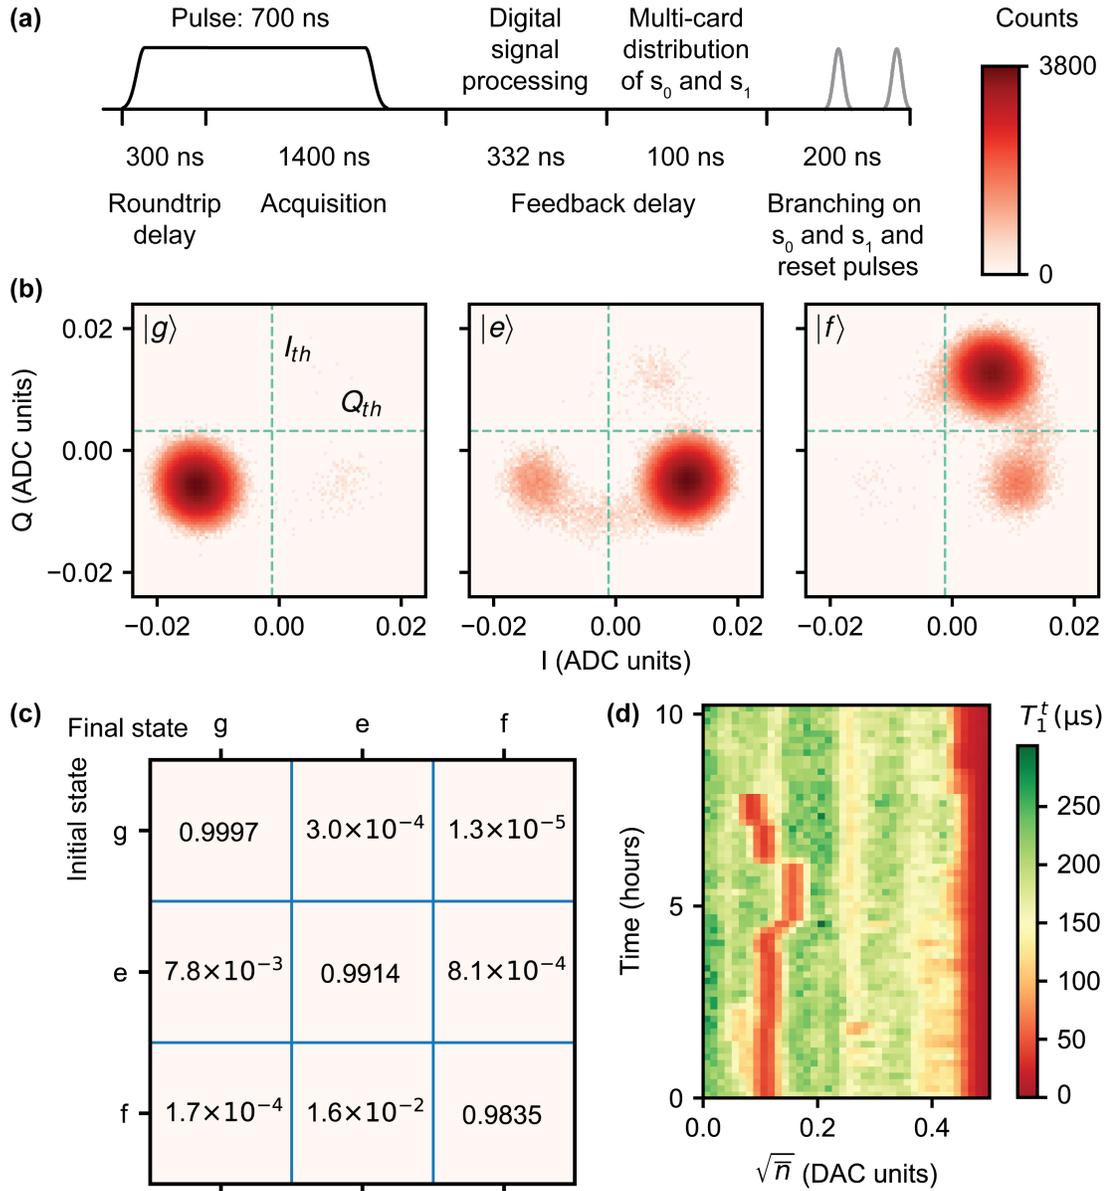
**Figure 4.2 | Calibration of Hamiltonian parameters.** (a) Number-resolved qubit spectroscopy with a selective square pulse of duration  $\sim 50 \mu\text{s}$  when the oscillator is in the vacuum state (blue) and a coherent state (red). (b) Optimal return phase in the out-and-back experiment (inset) with the qubit in  $|g\rangle$  and  $|e\rangle$ . As seen from the phase dispersion with  $\bar{n}$ , the effective oscillator nonlinearity is larger when the qubit is in  $|e\rangle$ . (c) Average oscillator rotation frequency  $(\omega_g + \omega_e)/2$ , and a linear fit to extract  $\Delta$  and  $K$ . (d) Relative oscillator rotation frequency  $\omega_g - \omega_e$ , and a linear fit to extract  $\chi$  and  $\chi'$ . The star indicates  $\chi$  obtained in (a).

ities  $K$  and  $\chi'$ . In this experiment, shown in the inset of Fig. 4.2(b), we create a coherent state  $|\alpha\rangle$  with an average number of  $\bar{n} = |\alpha|^2$  photons, wait for some time while it rotates in phase space, and attempt to return it back to the origin with a displacement of variable phase. The optimal return phase for qubit in  $|g\rangle$  and  $|e\rangle$  is shown in Fig. 4.2(b). Performing this experiment for different wait times allows to extract the effective oscillator rotation frequencies  $\omega_g(\bar{n})$  and  $\omega_e(\bar{n})$ . The linear fit of the average rotation frequency  $(\omega_g + \omega_e)/2 = \Delta + K\bar{n}$  yields the values of the detuning  $\Delta$  and Kerr nonlinearity  $K$ , as shown in Fig. 4.2(c). The linear fit of the relative rotation frequency  $\omega_g - \omega_e = \chi + \chi'\bar{n}$  yields the values of the dispersive shift  $\chi$  and the second-order dispersive shift  $\chi'$ , as shown in Fig. 4.2(d). We find that the value of  $\chi$  predicted with this method typically agrees with the value obtained via number-resolved spectroscopy to within 1%.

### 4.3 Readout and reset

The transmon measurement consists of a readout pulse of duration 700 ns with 40 ns ramp-up and ramp-down. The reflected microwave signal is acquired (after 300 ns delay to account for signal travel time) for the duration of 1400 ns. After acquiring the readout signal, FPGA performs digital signal processing, which consists of demodulation, integration of the signal with a filter function, and thresholding, all of which takes 332 ns. Next, the bits  $s_0$  and  $s_1$  that encode the measurement outcome are distributed to all control cards, which takes 100 ns. For a schematic of this measurement process, see Fig. 4.3(a). When the readout is used to realize ancilla reset, additional time is spent on branching on the  $s_0$  and  $s_1$  signals to apply appropriate feedback pulses. The branching is done as shown in Fig. 4.6(a), taking additional 200 ns to complete the reset. Due to the slow ringdown of the readout photons on a time scale of  $1/(\kappa_{r(e)} + \kappa_{r(i)}) \approx 320$  ns, the readout resonator is not fully empty when the feedback pulses are applied, partly limiting their fidelity through measurement-induced dephasing mechanism [103]. This limitation could be addressed in the future by using a strongly coupled readout resonator with photon lifetime on the order of ten nanoseconds [104], or, alternatively, by using an active resonator depletion protocol [105] as was done in the grid-code experiment [42].

To characterize the readout, we perform a two-measurement experiment in which the ancilla state is prepared with post-selection on the outcome of the first measurement [106]. The second measurement follows with a 500 ns delay after the first one. Its outcome is histogrammed, as shown in Fig. 4.3(b) for  $|g\rangle$ ,  $|e\rangle$  and  $|f\rangle$



**Figure 4.3 | Ancilla measurement.** (a) Timing of different components of readout and reset. (b) Logarithmic histogram of the integrated readout signal for different transmon initial states. (c) Markov transition matrix derived from the histogram in (b). It shows the probability of transition from any initial state to any final state during the measurement. (d) Time trace of  $|e\rangle$ -state lifetime (color coded). The resonator field amplitude in the steady state is proportional to the DAC amplitude (horizontal axis). The amplitude used for the actual readout corresponds to  $\sqrt{\bar{n}} = 0.22$ . When spurious resonance around  $\sqrt{\bar{n}} = 0.1$  reappears, the readout fidelity significantly reduces.

initial states. The SNR of our readout is large enough to not be a dominant cause of the readout infidelity. Instead, the fidelity is limited by state transitions during the finite readout time. Some transitions are expected due to the finite lifetime of the ancilla excited states, and the excess is induced by the readout pulse itself [107].

The Markov matrix in Fig. 4.3(c) describes the transition probabilities in this characterization experiment. It is obtained by integrating the parts of a histogram on various sides of two thresholds. The diagonal elements of this matrix can be interpreted as readout fidelities of different transmon states, with precision of about  $\sim 10^{-3}$  for  $|e\rangle$  and  $|f\rangle$  states due to possible decay during the 500 ns delay between the two measurements. The readout fidelity of the ground state  $\mathcal{F}_r^{(g)} = 0.9997$  is significantly better than that of the excited state  $\mathcal{F}_r^{(e)} = 0.9914$  – a crucial feature exploited in our QEC protocol, where the dominant “no error” syndrome is mapped to the  $g$  outcome

The readout fidelity of  $|g\rangle$  is also more stable over time, see Chapter 6.10. The readout fidelity of  $|e\rangle$  fluctuates over time due to fluctuations of  $T_1^t(\bar{n})$  (qubit lifetime in the presence of  $\bar{n}$  readout photons). A sample drift of  $T_1^t(\bar{n})$  is shown in Fig. 4.3(d). The exact reason for this effect is still not well understood. It is possible that the dependence  $T_1^t(\bar{n})$  comes from drive-induced hybridization of the transmon energy levels [108]. Higher levels are sensitive to offset charge, and thus fluctuations of environmental charges can affect the hybridization strength and lead to fluctuations of  $T_1^t(\bar{n})$ . Another possible explanation is that fluctuating  $T_1^t(\bar{n})$  dependence comes from a spectral overlap of the Stark-shifted qubit frequency with a spurious degree of freedom (not necessarily charged), e.g. a two-level defect, which itself fluctuates [89]. The correlation between the logical qubit performance, the readout infidelity of  $|e\rangle$  state, and the fluctuating  $T_1^t(\bar{n})$  is further discussed in Chapter 6.10.

## 4.4 Conditional displacement

We create an echoed conditional displacement gate  $\text{ECD}(\beta) = \sigma_x D(\sigma_z \beta/2)$  using the approach described in Ref. [51]. As illustrated in Fig. 4.4(a), this gate consists of the following steps: (i) the oscillator is displaced out in phase space by large amplitude  $\alpha$ ; (ii) the conditional rotation is accumulated during the time interval  $\tau$  along the arc of a large radius  $|\alpha|$  – this is equivalent to accumulation of the conditional displacement in the direction orthogonal to  $\alpha$  at an enhanced rate  $\chi|\alpha|$ ; (iii) the oscillator is returned back towards the origin of phase space with displacement

of amplitude  $-\alpha \cos(\chi\tau/2)$ ; (iv) the qubit state is flipped with an echo  $\pi$ -pulse; (v) an analogous large displacement sequence is repeated in the symmetrically opposite direction in phase space. Under the dispersive coupling model and in the limit of instantaneous rotation and displacement pulses, this protocol results in a net conditional displacement of amplitude  $\beta = -2i\alpha \sin(\chi\tau)$ . Due to deviations from this idealized scenario, such as finite pulse durations and higher order Hamiltonian terms, we need to experimentally calibrate the amplitude  $\alpha(\beta)$  of the large displacement required to achieve a desired conditional displacement.

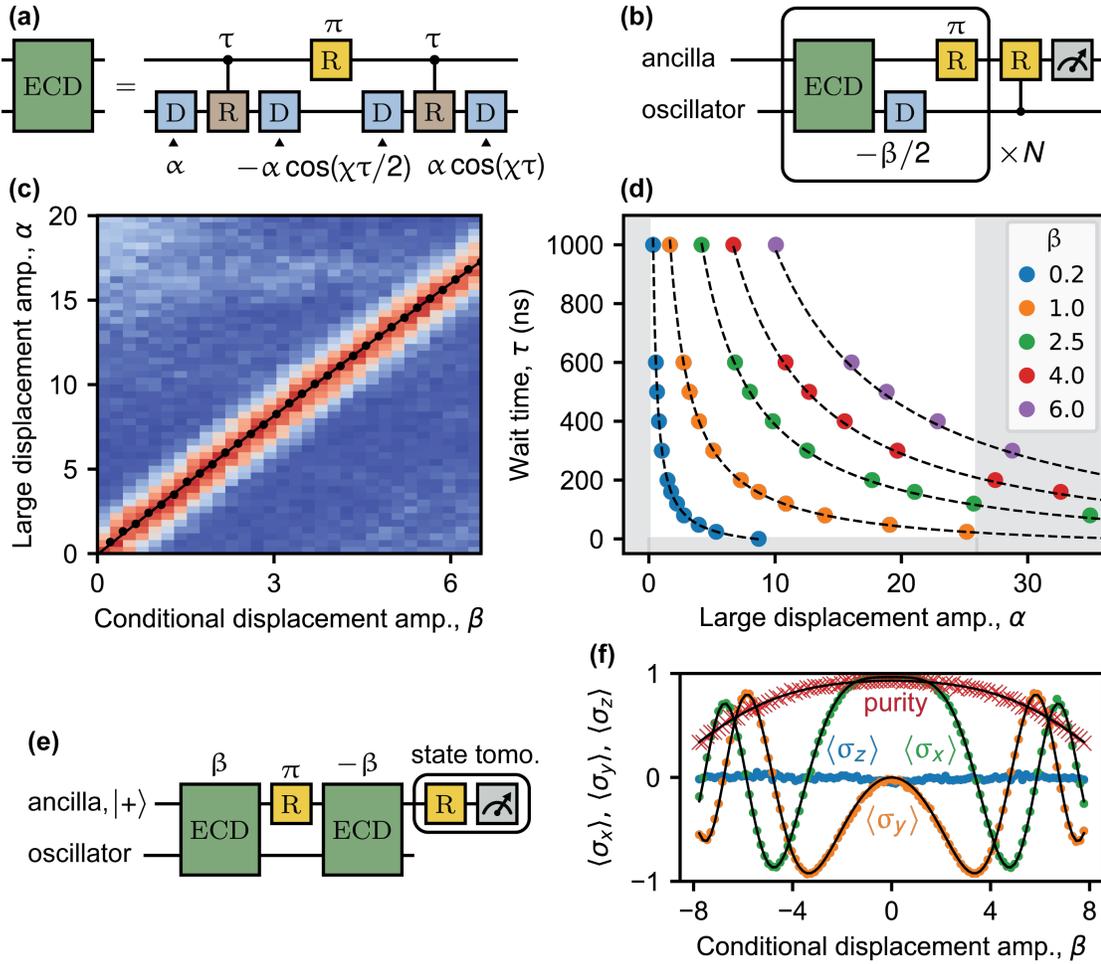
### Calibration of amplitude

Starting with qubit in  $|g\rangle$  (or  $|e\rangle$ , with similar results) and oscillator in  $|0\rangle$ , we apply the ECD gate with fixed delay  $\tau$  in the displaced state and varying amplitude  $\alpha$ , and then attempt to undo the effect of the gate and return the oscillator to vacuum with a simple displacement  $D(-\beta/2)$ . This out-and-back sequence is repeated  $N$  times to increase the resolution. At the end of the experiment, the qubit is probed with a selective pulse conditioned on oscillator in  $|0\rangle$ . The complete sequence is illustrated in Fig. 4.4(b), and the experimental data for the gate with delay  $\tau = 600$  ns is shown in Fig. 4.4(c). From this calibration measurement, we fit the dependence  $\alpha(\beta|\tau)$ , and we perform this calibration for a set of different wait times  $\tau$ .

During the optimization of the QEC performance, our RL agent is asked to pick the optimal values of the large displacement amplitude  $\alpha$  and of the conditional displacement amplitude  $\beta$ . Therefore, we need to have a calibrated inversion function  $\tau(\alpha, \beta)$  that predicts the wait time  $\tau$  to realize a gate with these parameters. We find that the empirical relation

$$\tau_e(\alpha, \beta) = \beta \left( p_0 + \frac{p_1}{2\alpha} \right) - p_2, \quad (4.2)$$

with fit parameters  $\vec{p} = \{p_0, p_1, p_2\}$ , is able to simultaneously fit all ECD calibration datasets, such as the one shown in Fig. 4.4(c), sufficiently well to be used with the training of the RL agent. Note that in the idealized model, we would have  $\vec{p} = \{0, 1/\chi, 0\}$ . The empirical fit results are shown in Fig. 4.4(d), where the shaded region indicates the prohibited parameter values, including the limited dynamic range of the DAC that allows  $\alpha \in [0, 26]$  given our choice of fixed-duration displacement pulses.



**Figure 4.4 | Calibration of the ECD gate.** (a) Realization of ECD gate using the approach from [51]. (b) Variation of the out-and-back experiment for calibration of the amplitude  $\alpha$  of the large displacement required to achieve a conditional displacement amplitude  $\beta$ . For a fixed delay  $\tau$  between out and back displacements, and a given value of  $\beta$ , we sweep  $\alpha$  to find the optimum. The out-and-back sequence is repeated  $N$  times to increase the resolution. At the end of the experiment, the qubit is probed with a selective pulse conditioned on oscillator in  $|0\rangle$ . (c) Data from the amplitude calibration experiment shown in (b), using  $\tau = 600$  ns and  $N = 4$ . (d) Simultaneous fit of the collection of fixed- $\tau$  datasets, such as the one shown in (c), to the empirical function in Eq. (4.2). The shaded region indicates the prohibited parameter values. (e) Cat-and-back experiment. Starting with a pure qubit state  $|+\rangle$  and arbitrary oscillator state, this experiment results in a phase accumulation in the equatorial plane on a qubit Bloch sphere, which is detected with qubit state tomography. (f) Results of the qubit state tomography in cat-and-back experiment with the ECD gate with wait time  $\tau = 600$  ns. In this experiment, the oscillator was initially prepared in the  $|+Z\rangle$  grid state. The data is fitted to the model in Eq. (4.4), shown with black solid lines.

## Calibration of qubit phase

As explained in Ref. [51], this experimental implementation of the  $\text{ECD}(\beta)$  gate results in additional qubit phase accumulation  $\Theta[\beta] = \xi|\beta|^2$ , i.e. we implement  $\overline{\text{ECD}}(\beta) = \exp(-i\sigma_z \Theta[\beta]/2) \text{ECD}(\beta)$ . The amplitude calibration experiment described above is not sensitive to this phase, because the qubit always remains in the eigenstate of  $\sigma_z$ . However, this phase is important when conditional displacements are concatenated, e.g. in the ECD control unitaries, as described in Chapter 5.1.

To calibrate this phase  $\Theta[\beta]$ , we perform the following “cat-and-back” experiment

$$\overline{\text{ECD}}(-\beta) R_x(\pi) \overline{\text{ECD}}(\beta), \quad (4.3)$$

also shown in Fig. 4.4(e), which is ideally equivalent to  $\sigma_x \exp(i\Theta[\beta]\sigma_z)$ . Starting with a qubit in  $|+\rangle$ , the final state will satisfy  $\langle\sigma_y\rangle = \sin(2\xi|\beta|^2)$  and  $\langle\sigma_x\rangle = \cos(2\xi|\beta|^2)$  irrespective of the initial oscillator state.

However, due to decoherence and control imperfections in the ECD implementation, we find that the ancilla qubit also experiences loss of purity. Under the assumption that the losses of purity during the two conditional displacement gates  $\overline{\text{ECD}}(\beta)$  and  $\overline{\text{ECD}}(-\beta)$  are uncorrelated and independent of the direction in phase space, we model it as a uniform contraction of the Bloch vector by  $\sqrt{1-p[\beta]}$  per ECD gate, where  $p[\beta] = \eta_0 + \eta_2|\beta|^2 + \eta_4|\beta|^4$ . Hence, we fit the cat-and-back experiment to the following model:

$$\begin{pmatrix} \langle\sigma_x\rangle \\ \langle\sigma_y\rangle \end{pmatrix} = (1-p[\beta]) \begin{pmatrix} \cos(2\Theta[\beta]) \\ \sin(2\Theta[\beta]) \end{pmatrix}, \quad (4.4)$$

with fit parameters  $\{\eta_0, \eta_2, \eta_4, \xi\}$ , of which only  $\xi$  is used in the ECD control compilation method, see Chapter 5.1.

The results of qubit state tomography together with the fit to the model in Eq. (4.4) are shown in Fig. 4.4(f) for the same ECD gate as in Fig. 4.4(c). As explained in Ref. [51], the value of  $\xi$  depends on the shape of the phase space trajectory during the ECD gate, and thus we calibrate it independently for every choice of delay time  $\tau$ .

## 4.5 Oscillator error channels

### Relaxation and excitation

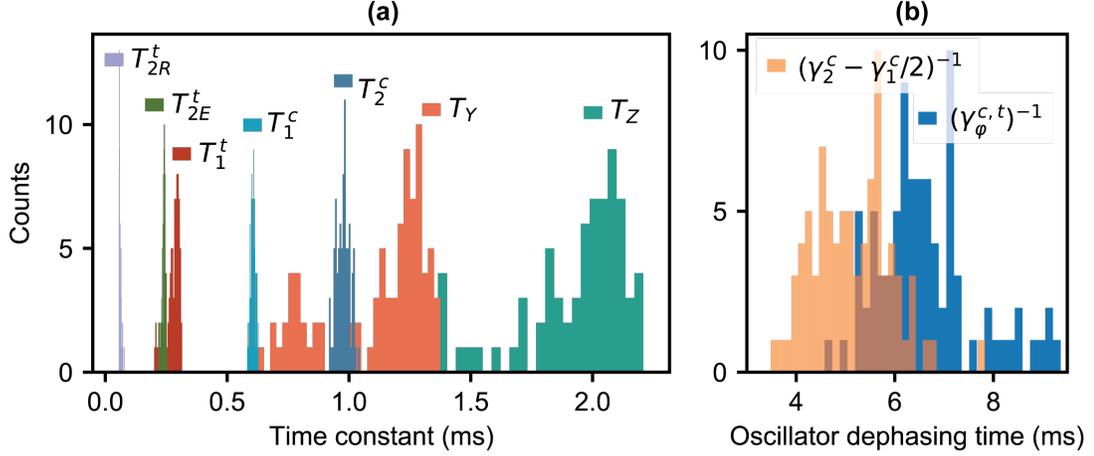
To measure the oscillator relaxation rate  $\gamma_1^c = \gamma_\downarrow^c + \gamma_\uparrow^c$ , we first prepare Fock state  $|1\rangle$  using a unitary control circuit with 5 layers, see Chapter 5.1. After a time delay of varying length, we measure the remaining occupation of  $|1\rangle$  and fit it to an exponential decay with time constant  $T_1^c = 1/\gamma_1^c$ . To measure this occupation, we apply a spectrally selective ancilla qubit pulse which flips the qubit conditioned on one photon in the oscillator. Monitoring the oscillator over a week-long period, we find the mean and standard deviation of  $\bar{T}_1^c = 606 \pm 10 \mu\text{s}$ . As seen from the histogram in Fig. 4.5(a), the relative fluctuations of  $T_1^c$  are small compared to the relative fluctuations of other error channels in the same time frame. We attribute this stability to the fact that most of the electromagnetic field of this mode resides in the vacuum of the cavity.

To bound the rate of thermal excitation  $\gamma_\uparrow^c$ , we apply the feedback cooling technique described in Chapter 4.6, to the oscillator in its steady state. Since we find no detectable difference in the qubit number-resolved spectroscopy contrast of the zeroth peak after feedback cooling, the resolution of this measurement of  $\sim 1\%$  provides a bound on the oscillator excitation rate of  $\gamma_\uparrow^c < 1/(60 \text{ ms})$ . This rate is negligible compared to all other rates in the system and is ignored in the rest of the analysis.

### Dephasing

To measure the rate of dephasing  $\gamma_2^c$  within the  $\{|0\rangle, |1\rangle\}$  manifold, we prepare a superposition  $|0\rangle + |1\rangle$  using the  $Y90$  gate realized with a unitary control circuit with 8 layers, see Table 5.1. After a time delay of varying length we apply the  $Y90$  gate again and measure the occupation of  $|0\rangle$ . In the reference frame of the LO, the oscillator state rotates with angular frequency  $\chi/2$  during the time delay, which results in Ramsey oscillations modulating the exponential decay with decay time constant  $T_2^c = 1/\gamma_2^c$ . We adjust the sampling rate to make the oscillations appear slow. We find the one-week mean and standard deviation of  $\bar{T}_2^c = 980 \pm 30 \mu\text{s}$ .

One possible source of oscillator dephasing is stochastic rotations acquired due to dispersive coupling with the transmon combined with transmon stochastic excitation and relaxation events [45]. The dephasing rate due to this effect was predicted to be  $\gamma_\phi^{c,t} \approx n_{\text{th}}^t \gamma_\downarrow^t$  in the limit  $\chi \gg \gamma_1^t$  and  $\gamma_\downarrow^t \gg \gamma_\uparrow^t$ , where  $n_{\text{th}}^t$  is the steady-

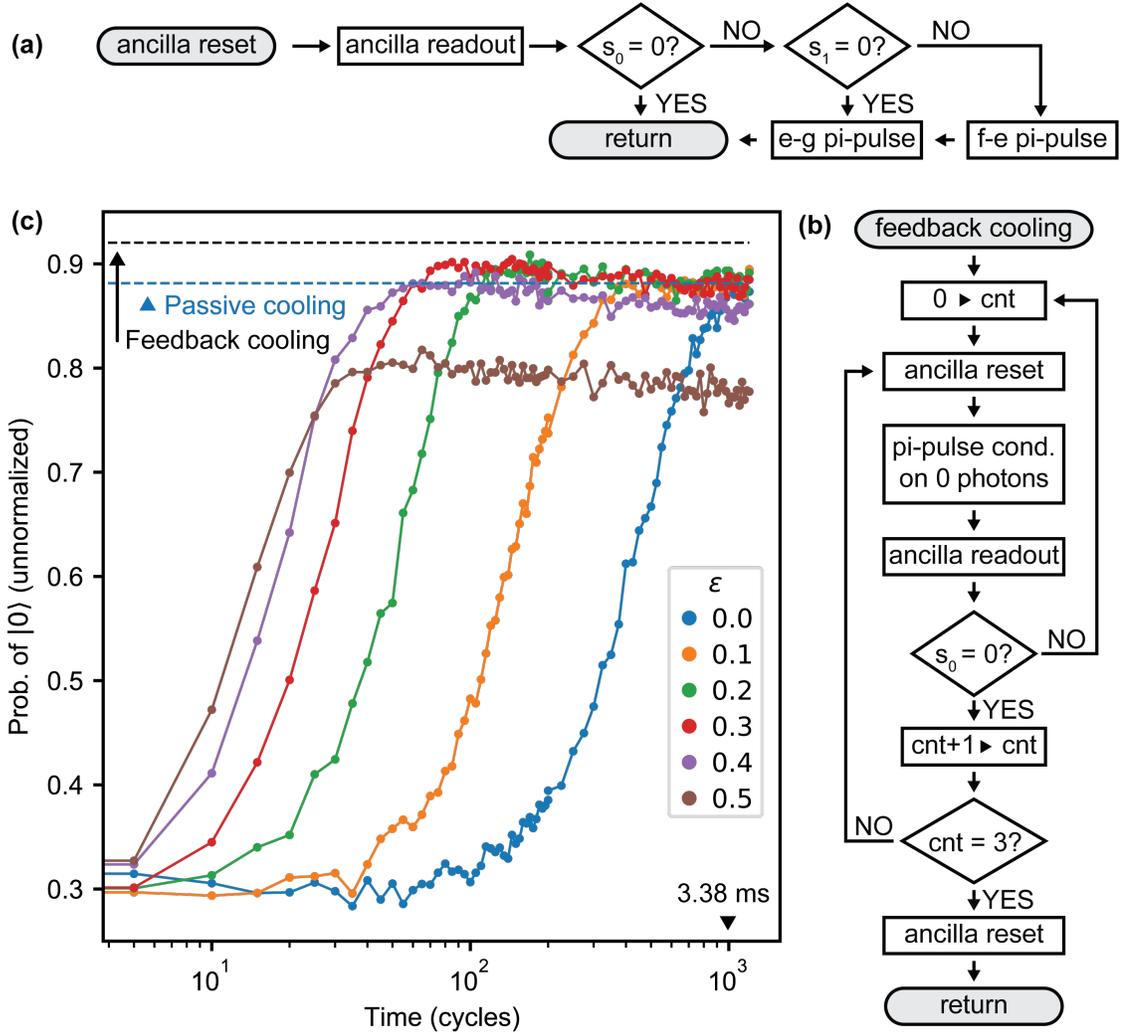


**Figure 4.5 | Fluctuating error channels.** (a) Histogram of  $T_1$  and  $T_2$  times of the transmon and the oscillator, and logical lifetimes of error-corrected grid states. The histogram is derived from a week-long scan described in Chapter 6.10. (b) Oscillator pure dephasing time extracted from the measured oscillator parameters and predicted from the dispersive coupling model.

state population of  $|e\rangle$ . In our system, the correlation between  $\gamma_\varphi^c = \gamma_2^c - \gamma_1^c/2$  and  $\gamma_\varphi^{c,t}$  is difficult to measure because these rates are small and their estimators are subject to strong relative fluctuations. By comparing the medians of their marginal distributions,  $\bar{\gamma}_\varphi^{c,t} = 1/(6.5 \text{ ms})$  and  $\bar{\gamma}_\varphi^c = 1/(5.1 \text{ ms})$ , shown in Fig. 4.5(b), we find the remaining unexplained contribution to dephasing at a rate  $\gamma_\varphi^c = 1/(24 \text{ ms})$  whose source is not yet identified. It is plausibly related to second-order excitations from  $|e\rangle$  to  $|f\rangle$  [47].

## 4.6 Active oscillator cooling

Given the long relaxation time  $\bar{T}_1^c = 606 \mu\text{s}$  of our oscillator, passive cooling that relies on the natural interaction with the cold environment is impractically long. For example, starting with a Fock state  $|1\rangle$ , it would take approximately  $4.6\bar{T}_1^c = 2.8 \text{ ms}$  to reduce the average population to 0.01 photons. In practice, since we work with the grid states, the required cooling time is even longer. Therefore, the goal of our active cooling routine is to reduce the experimental duty cycle time and also to remove any residual thermal population. We achieve these goals via a two-step procedure which consists of an engineered dissipative pre-cooling and subsequent feedback cooling.



**Figure 4.6 | System cooling.** (a) Ancilla reset subroutine with measurement-based three-state feedback. (b) Oscillator feedback cooling subroutine adapted from Ref. [40]. (c) Demonstration of dissipative cooling of the oscillator starting from GKP  $|+Z\rangle$  state with  $\Delta = 0.3$ . A single cooling cycle consists of a pulse sequence in Eq. (4.8), ancilla qubit reset as in (a), and virtual rotation gate to the orthogonal quadrature for the next cycle. The duration of a single such cycle is 3.38  $\mu$ s. The case  $\epsilon = 0$  is equivalent to passive cooling. Dashed lines represent the contrast of the zeroth photon number peak in qubit spectroscopy after passive cooling of 5 ms and after feedback cooling with  $Y = 3$ .

## Dissipative pre-cooling

We introduce a novel oscillator cooling method based on the conditional displacements, ancilla rotations, and ancilla resets. This protocol can also be realized in trapped ions, as was hinted in Ref. [54].

To derive this protocol, we apply the same dissipation engineering framework [98] as used in Ref. [53] to derive the SBS stabilization of the GKP manifold. The dissipator  $\gamma \mathcal{D}[a]$  can be approximated with a sequence of discrete entangling interactions  $U(t)$  between the ancilla and the oscillator, and ancilla resets. For  $\gamma \mathcal{D}[a]$ , the interaction should be of the form  $U(t) = \exp[-i\sqrt{\gamma t}(a\sigma_+ + a^\dagger\sigma_-)]$ , where the constraint  $\langle a^\dagger a \rangle \gamma t \ll 1$  controls the validity of this discrete approximation. To further approximate this unitary as a multi-layer circuit with gates from our gate set, we perform the first order Trotter decomposition:

$$U = \exp\left(-i\sqrt{\frac{\gamma t}{2}}(x\sigma_x + p\sigma_y)\right) \quad (4.5)$$

$$= \exp\left(-i\sqrt{\frac{\gamma t}{2}}x\sigma_x\right) \exp\left(-i\sqrt{\frac{\gamma t}{2}}p\sigma_y\right) + O(\gamma t) \quad (4.6)$$

$$\approx R_y^\dagger(\pi/2) \text{ECD}(-i\varepsilon) R_x^\dagger(\pi/2) \text{ECD}(\varepsilon) R_y(\pi/2) R_z(\pi/2), \quad (4.7)$$

where we defined the “trimming amplitude”  $\varepsilon = \sqrt{\gamma t}$ . Furthermore, since the ancilla qubit is assumed to always start in  $|g\rangle$ , we can omit the first gate  $R_z(\pi/2)$ . The resulting unitary part of the dissipative cooling circuit is:

$$R_y^\dagger(\pi/2) \text{ECD}(-i\varepsilon) R_x^\dagger(\pi/2) \text{ECD}(\varepsilon) R_y(\pi/2), \quad (4.8)$$

also summarized in Table 5.1. To achieve uniform cooling in all directions in phase space, the orientation of the ECD gates needs to cycle between position and momentum quadratures. A single cycle, including the pulse sequence in (4.8), ancilla reset, and subsequent virtual rotation gate on the FPGA, has a duration of 3.38  $\mu\text{s}$ .

To demonstrate the performance of this cooling protocol, we start with a  $|+Z\rangle$  grid state with  $\Delta = 0.3$  and apply varying number of cooling cycles, monitoring the population of  $|0\rangle$  with a selective qubit pulse. As seen in Fig. 4.6(c), dissipative cooling allows the state to shrink towards vacuum significantly faster than passive cooling. With  $\varepsilon = 0.4$ , the cooling rate is 20 times faster than energy relaxation time of the oscillator. For small  $\varepsilon \leq 0.3$  the steady-state thermal occupation after dissipative cooling is similar to passive cooling of this state of duration 5 ms.

Larger  $\varepsilon$  allows for faster cooling, but at the expense of significant residual thermal occupation.

### Feedback cooling

To remove the residual thermal photons, we further apply the feedback cooling protocol introduced in Ref. [40] and shown in Fig. 4.6(b). With the help of a selective qubit pulse conditioned on  $|0\rangle$  and qubit measurement, the protocol repetitively asks the question “Is the oscillator in vacuum?” and terminates only when it receives  $Y$  consecutive “yes” answers. It would be inefficient to run this feedback protocol starting with an arbitrary initial oscillator state, since the probability  $p_y$  of obtaining “yes” can be very small. The dissipative pre-cooling quickly boosts this probability to a level essentially limited by the fidelity of the selective qubit pulse, and thereby decreases the run time of the subsequent feedback cooling step. The run time of feedback cooling is non-deterministic, but the expected number of rounds in a model with constant  $p_y$  is

$$\overline{N}_{fc}(p_y, Y) = \frac{p_y^{-Y} - 1}{1 - p_y} \quad (4.9)$$

Our final routine, called “active cooling” throughout this work, consists of 25 cycles of dissipative pre-cooling (50 cycles, if counting each quadrature individually) with  $\varepsilon = 0.4$  followed by the feedback cooling with  $Y = 3$ . We estimate that with  $p_y = 0.87$ , achieved after the pre-cooling, the expected run time of the whole routine is approximately  $50 \times 3.38 \mu\text{s} + \overline{N}_{fc}(0.87, 3) \times 25 \mu\text{s} = 270 \mu\text{s}$ , which in our system corresponds to  $0.45 T_1^c$  (and could potentially be reduced further).

From the contrast of the zeroth photon number peak in the qubit spectroscopy [dashed lines in Fig. 4.6(c)], we see that passive cooling of duration 5 ms starting from the  $|+Z\rangle$  grid state still leaves a residual thermal population larger than what our protocol achieves in a much shorter time. However, when active cooling is applied to an oscillator in its steady state (nominally, vacuum) we find no resolvable improvement of the spectroscopy contrast, which leads us to conclude that the residual thermal population after active cooling is at the sub-percent level where it cannot be resolved with our spectroscopy. This observation is used in Chapter 4.5 to derive an upper bound on the oscillator thermal excitation rate.

# Quantum control optimization

This chapter describes model-based and model-free optimization methods used to improve the performance of our system. The model-based optimization of control pulses, described in Chapter 5.1, is largely based on the work of A. Eickbusch *et al.* in Ref. [51]. It stands out among other approaches to quantum control because of the modularity and sparse parametrization, which is in stark contrast with the dominant approach based on the gradient ascent pulse engineering (GRAPE). The framework of model-free optimization of quantum control with reinforcement learning, described in Chapter 5.2, is a novel contribution of this thesis, described in more detail in Ref. [55].

## 5.1 Model-based optimization of control circuits

### Circuit decomposition

Our control gate set consists of two parametrized gates: (i) echoed conditional displacement of the oscillator  $\text{ECD}(\beta) = \sigma_x D(\sigma_z \beta/2)$ , where  $D(\alpha) = \exp[\alpha a^\dagger - \alpha^* a]$  is the displacement operator, and (ii) rotation of the qubit around an axis in the equatorial plane  $R(\varphi, \theta) = \exp[-i(\theta/2)(\sigma_x \cos \varphi + \sigma_y \sin \varphi)]$ . Recently, it was shown that this gate set is well suited for the universal control of an oscillator with weak dispersive coupling to a qubit [51]. Most unitary operations in our experiment are decomposed as parametrized multilayer circuits of the form

$$\text{circuit}(\boldsymbol{\beta}, \boldsymbol{\varphi}, \boldsymbol{\theta}) = \underbrace{\text{ECD}(\beta_T)R(\varphi_T, \theta_T)}_{\text{layer } T} \cdots \underbrace{\text{ECD}(\beta_1)R(\varphi_1, \theta_1)}_{\text{layer } 1}, \quad (5.1)$$

where  $\boldsymbol{\beta} \in \mathbb{C}^T$  is a vector of conditional displacement amplitudes, and  $\boldsymbol{\varphi}, \boldsymbol{\theta} \in \mathbb{R}^T$  are vectors of qubit rotation phases and angles respectively. For example, we utilize this decomposition as part of the following operations:

- Dissipative cooling of the oscillator, see Chapter 4.6.
- Preparation of the GKP states, see Chapter 6.8.
- Small-Big-Small protocol, see Chapter 6.3.
- Preparation of the Fock state  $|1\rangle$ , see Chapter 4.5.

	Dissipative cooling			SBS protocol			GKP $ +Z\rangle$ preparation		
$t$	$\beta$	$\varphi$	$\theta$	$\beta$	$\varphi$	$\theta$	$\beta$	$\varphi$	$\theta$
1	+0.4	$+\pi/2$	$+\pi/2$	+0.2i	$+\pi/2$	$+\pi/2$	+0.52 + 2.54i	-1.28	+1.57
2	-0.4i	0	$-\pi/2$	$+\sqrt{2}\pi$	0	$-\pi/2$	-0.83 - 0.36i	+2.85	-2.76
3	0	$+\pi/2$	$-\pi/2$	+0.2i	0	$+\pi/2$	-0.36 + 0.85i	+0.29	+0.55
4				0	$+\pi/2$	$-\pi/2$	-0.86 + 1.61i	-0.29	+1.43
5							-2.16 + 0.12i	+0.29	+0.92
6							-0.09 + 1.73i	+2.85	-1.56
7							+2.05 + 0.73i	+0.29	+1.08
8							+0.22 - 0.66i	-0.29	-2.71
9							-0.08 - 1.56i	+0.29	+2.06
10							+0.19 + 0.04i	+2.85	+1.60
11							0	+1.86	+1.57

**Table 5.1 | Circuit parameters.** Parameters for dissipative cooling and SBS protocol are created based on the models described in Chapter 4.6 and Ref. [53] respectively. Parameters for  $|+Z\rangle$  grid state preparation are numerically optimized with Keras.

- $Y_{90}$  gate on Fock  $\{|0\rangle, |1\rangle\}$  encoding, see Chapter 4.5.

## Circuit optimization

A circuit optimization method for this gate set was developed in Ref. [51]. Here, we present a simplified modular framework based on the Keras library [109], which allows to optimize circuit parameters in a manner similar to training of the neural networks. The parametrized control circuits (5.1) are created as instances of the `tf.keras.Sequential` class which is commonly used for concatenating multiple neural network layers. Here, we instead use custom layers representing the parametrized gates  $ECD(\beta)$  and  $R(\varphi, \theta)$  as subclasses of `tf.keras.layers.Layer`. This allows us to exploit flexible and user-friendly application-programming interface of the Keras library to optimize the circuit parameters and automatically monitor various aspects of the optimization progress. To illustrate the accessibility of such an approach, in Fig. 5.1 we provide an example code for optimization of the  $Y_{90}$  gate on the  $\{|0\rangle, |1\rangle\}$  qubit. Complete code with dependencies and further examples is available in Ref. [110]. Such optimization, which is performed for a batch of  $B = 300$  circuit candidates in parallel on a graphics processing unit (GPU), takes about 10 minutes to finish. In Table 5.1, we list circuit parameters for some of the control operations in our experiment. Curiously, some of the numerically optimized parameter values are clearly interpretable, e.g. in GKP state prepara-

```

import tensorflow as tf
from quantum_control.layers import QubitRotation, ConditionalDisplacement
from utils import basis, Kronecker_product, batch_dot, normalize

N = 100 # oscillator truncation in Fock basis
T = 8   # circuit depth
B = 300 # batch size

# build the circuit as Keras model
circuit = tf.keras.Sequential(name='fockY90')
for i in range(T):
    echo_pulse = True if i<T-1 else False
    circuit.add(QubitRotation(N, [1,B]))
    circuit.add(ConditionalDisplacement(N, [1,B], echo_pulse=echo_pulse))

# Create cardinal points on the Bloch sphere
# Hilbert space convention is qubit x oscillator
states = {'+Z' : Kronecker_product([basis(0,2), basis(0,N)]),
          '-Z' : Kronecker_product([basis(0,2), basis(1,N)])}
states['+X'] = normalize(states['+Z'] + states['-Z'])[0]
states['-X'] = normalize(states['+Z'] - states['-Z'])[0]

# specify the action of Y90 gate
gate_map = {'+X' : '-Z', '-X' : '+Z'}

# define input and output states
inputs = tf.stack([states[i] for (i,f) in gate_map.items()])
targets = tf.stack([states[f] for (i,f) in gate_map.items()])

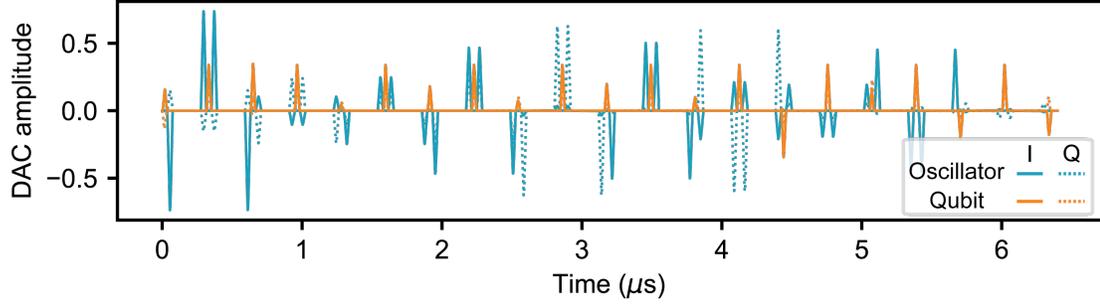
# define the loss function and the optimizer
def loss(state, target):
    return -tf.math.reduce_sum(tf.math.real(batch_dot(state, target)))

optimizer = tf.keras.optimizers.Adam(learning_rate=1e-3)

# compile this Keras model and run the optimizer
circuit.compile(optimizer=optimizer, loss=loss)
circuit.fit(x=inputs, y=targets, epochs=10000)

```

**Figure 5.1 | Circuit optimization.** Example of a Python script for optimization of the circuit parameters for Y90 gate on the  $\{|0\rangle, |1\rangle\}$  qubit. Gates are represented as custom Keras layers, and the circuit is compiled as sequential model. Optimization utilizes TensorFlow backend for automatic differentiation of the model.



**Figure 5.2 | Waveform for  $|+Z\rangle$  grid state preparation.** The parametrized control circuit is decomposed into primitive gates: qubit rotations, oscillator displacements, and conditional rotations. The waveform is compiled from this sequence of gates using experimental calibrations. Each qubit rotation and oscillator displacement is replaced with a corresponding Gaussian pulse, and the conditional rotation is replaced with a delay of certain length during which the system freely evolves under the dispersive coupling Hamiltonian.

tion circuit the rotations at steps  $t = 1, 6, 10, 11$  seem to be by an angle  $\pi$ . Detailed inspection of these circuits can lead to improved analytic constructions, which is left for future research.

## Pulse compilation

Having obtained the circuit parameters, we compile the waveforms to be played on the qubit and oscillator control lines. Such compilation requires prior calibration of the rotation  $R(\varphi, \theta)$  gate, described in Chapter 4.1, and the  $\text{ECD}(\beta)$  gate, described in Chapter 4.4.

As explained in Ref. [51] and in Chapter 4.4, our experimental implementation of the  $\text{ECD}(\beta)$  gate results in additional qubit phase accumulation  $\Theta[\beta] \propto |\beta|^2$ , i.e. we implement  $\overline{\text{ECD}}(\beta) = \exp(-i\sigma_z \Theta[\beta]/2) \text{ECD}(\beta)$ . We use the experimental calibration of this phase to adjust the numerically optimized vector  $\varphi$  according to the rule

$$\varphi_t \leftarrow \varphi_t - \sum_{\tau=1}^{t-1} (-1)^{t-\tau} \Theta[\beta_\tau], \quad t > 1. \quad (5.2)$$

In addition, in many cases of interest the ancilla qubit at the end of the circuit returns to  $|g\rangle$  and disentangles from the oscillator. In such cases, the last conditional displacement  $\text{ECD}(\beta_T)$  can be realized as a simple displacement  $D(\beta_T/2)$ . We use this simplification in state preparation circuits and in the SBS protocol.

In Fig. 5.2, we show an example waveform for unitary preparation of the  $|+Z\rangle$  grid state using a parametrized circuit with  $T = 11$  layers. Each ECD gate is decomposed via large displacements and conditional rotations. For clarity, in this example all conditional rotations are implemented with a constant wait time  $\tau = 200$  ns; hence, the whole compiled waveform has a duration of  $6.4 \mu\text{s}$ . Faster implementations are possible if the wait time is adapted to the magnitude of the conditional displacement, as described in Chapter 4.4. For example, in our system the conditional displacement of amplitude  $|\beta| < 0.5$  could, in principle, be implemented with zero wait time, see Fig. 4.4(d).

## 5.2 Model-free reinforcement learning for QEC

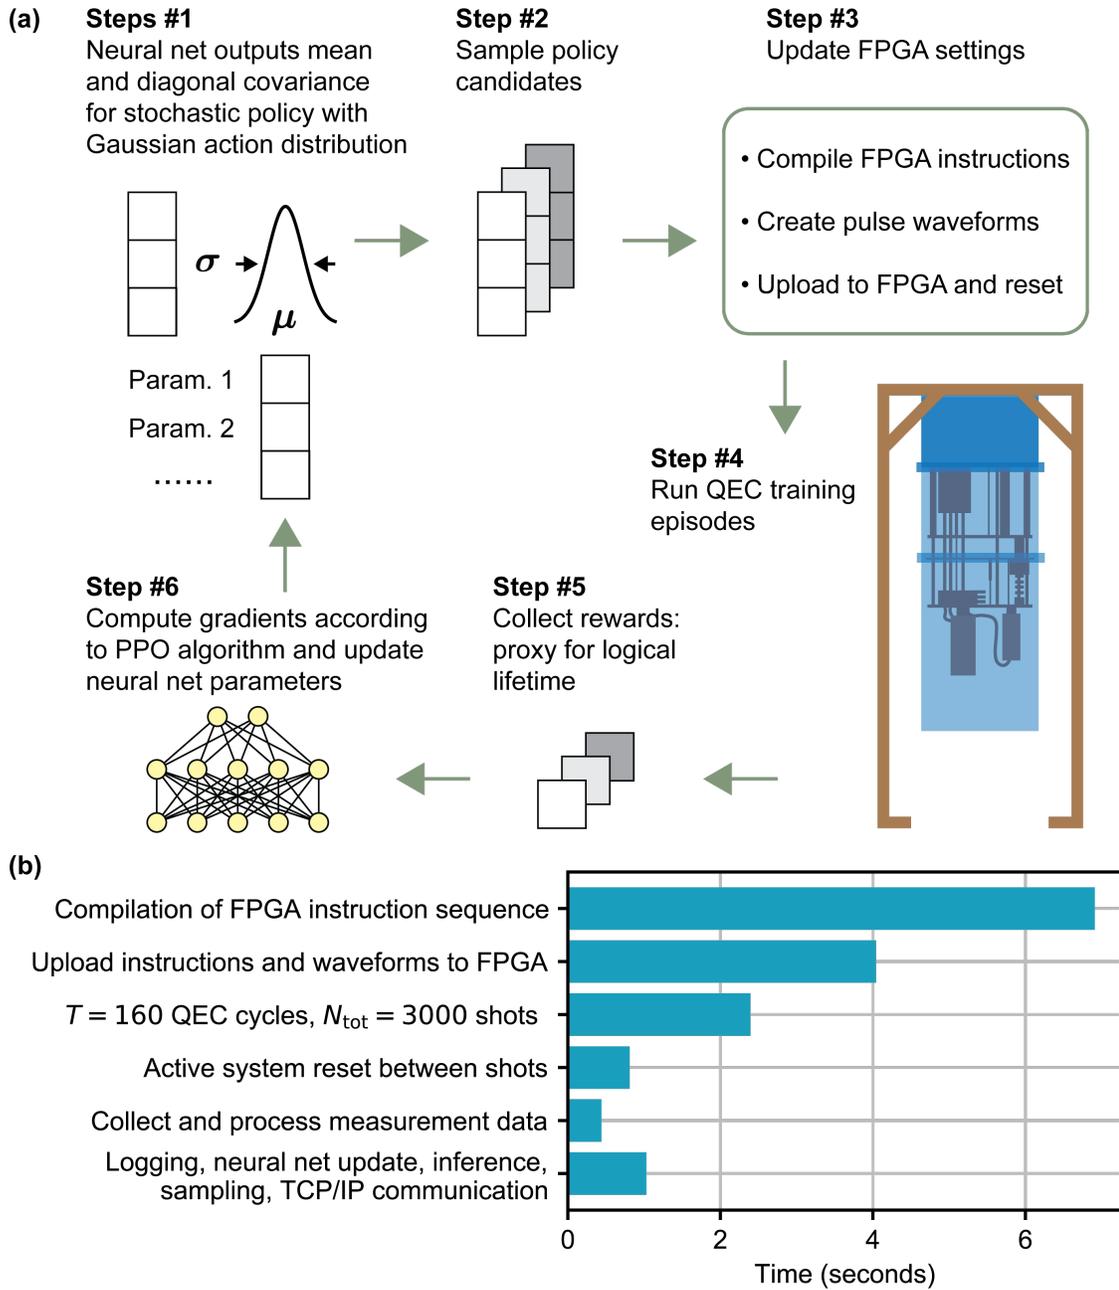
While most quantum operations in our experiment are optimized with a model-based approach described above, for quantum error correction we deploy a more powerful framework of model-free optimization. We use a reinforcement learning algorithm called proximal policy optimization (PPO) [56, 78]. For a detailed description of this algorithm in the context of quantum control we refer to Ref. [55]; here, we only provide a basic high-level picture. The complete training loop of our experiment is illustrated in Fig. 5.3; it is structured as follows:

**Step 1.** On training epoch  $t$ , neural network produces a probability distribution  $\mathcal{N}(\vec{\mu}_t, \vec{\sigma}_t)$ , where  $\vec{\mu}_t = \vec{\mu}(\theta_t)$ ,  $\vec{\sigma}_t = \vec{\sigma}(\theta_t)$ , and  $\theta_t$  summarizes the values of all weights and biases of the neural network in the current epoch.

**Step 2.** We sample a batch of  $B = 10$  parameter vectors from this distribution. They correspond to different QEC circuit candidates that should be evaluated in experiment. The neural network and sampling are implemented on NVIDIA 2080Ti graphics processing unit (GPU) in a separate computer. The sampled vectors are sent to the control computer via a local area network with negligible communication time.

**Step 3.** Based on these parameter vectors, we compile QEC circuit candidates, translated into FPGA instructions and DAC waveforms. All circuit candidates follow the same program execution flow, but the control waveforms and the content of FPGA registers is different for every candidate. The FPGA is reset and its wave memory is updated. This time-consuming step is the bottleneck of the training loop.

**Step 4.** Each candidate is evaluated in experiment. To this end, we initialize logical Pauli eigenstates  $|+Z\rangle$  and  $|+X\rangle$ , run the QEC for  $T = 160$  cycles, and then



**Figure 5.3 | Reinforcement learning.** (a) Experimental training loop. (b) Training time budget per epoch.

perform one-bit phase estimation of the corresponding logical Pauli operators. To suppress the sampling noise, we repeat this  $N_{\text{avg}} = 150$  times per Pauli and per circuit candidate. In total, one epoch of training consists of  $N_{\text{tot}} = 2BN_{\text{avg}} = 3000$  experimental shots.

**Step 5.** To produce the reward, we treat the measurement of a Pauli operator after  $T$  cycles as a proxy for logical lifetime. While averaging the measurement outcomes, we mask the experimental shots that started with incorrect state initialization, as flagged by a verification ancilla measurement after the state initialization.

**Step 6.** Once the rewards are available, PPO algorithm updates the neural network parameters  $\theta_t \rightarrow \theta_{t+1}$  for the next epoch. The gradients of these parameters are computed with automatic differentiation via back-propagation. After updating the neural network, the new training epoch begins.

The time budget of this training is shown in Fig. 5.3. All steps outlined above amount to 15.6 s per epoch. In the current implementation, the major bottleneck is Python-to-FPGA transition (step 3). Because of this, the implementation is less optimal in terms of sample efficiency than the proposal in Ref. [55]. The optimal approach would be to spend the total sample budget per epoch to evaluate more circuit candidates with minimal accuracy, instead of evaluation only a few candidates with high accuracy (achieved through averaging). In other words, based on the results of Ref. [55], we expect that a training with  $(B, N_{\text{avg}}) = (1000, 1)$  would require fewer experimental shots to reach a given performance level than a training with  $(B, N_{\text{avg}}) = (10, 100)$ . However, considering the total run time of the training, we had to compromise between bare sample efficiency (number of shots) and the overhead in step 3 of the pipeline. The overhead is independent of  $N_{\text{avg}}$  but increases with  $B$ , and due to a limited FPGA instruction sequence length we can only evaluate  $B \leq 10$  candidates per compilation. After paying the compilation overhead in step 3, a certain amount of averaging comes essentially for free and does not considerably affect the run time, hence the choices made here.

In Chapter 6.4, we describe the QEC circuit parametrization, show the evolution of parameter values during the course of training, and provide interpretation of the observed trends.

# Quantum error correction

This chapter contains additional details on the grid-code QEC that supplement the main results summarized in Chapter 2. The grid code for embedding a qubit into an oscillator, briefly described in Chapter 6.1, was known for a little over two decades [31], but its first experimental realization was achieved only recently [42]. Shortly afterwards, an improved protocol for error correction of this code was theoretically proposed in Ref. [53] and a similar protocol was (independently) proposed and realized in a trapped ion system [54]. Chapter 6.2 of this thesis describes a novel understanding of such a protocol, developed in collaboration with my colleagues B. Royer and S. Singh after the experiment was already completed.

## 6.1 Brief introduction to grid code

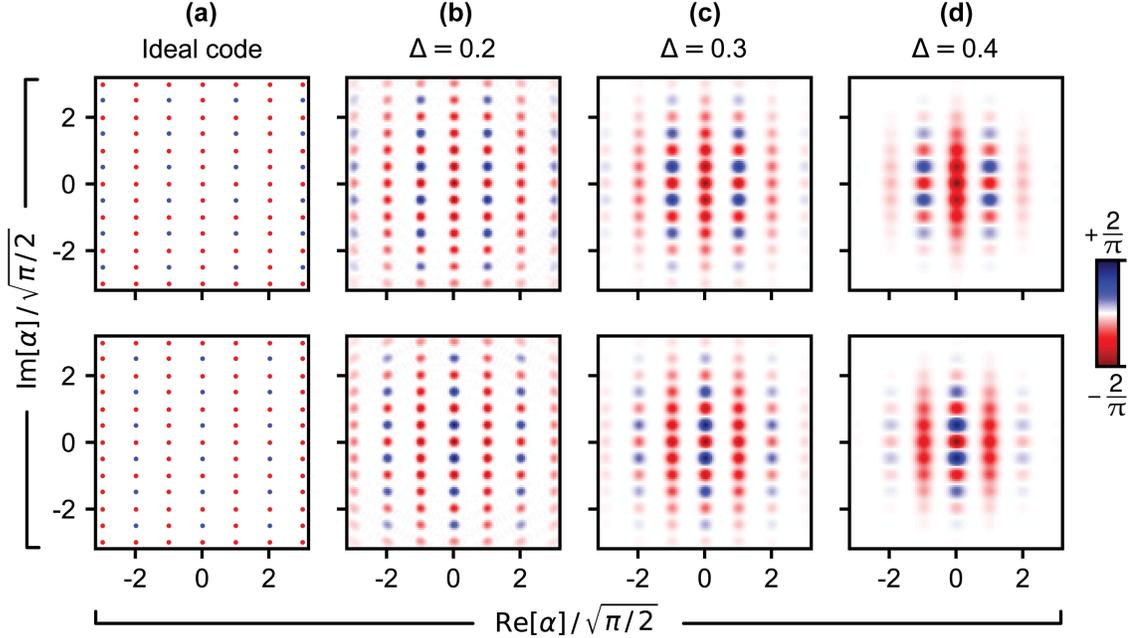
Qubit-register stabilizer codes are based on the group of Pauli operators; consider instead a stabilizer code based on the group of oscillator displacement operators. By definition, the +1 eigenstates of a displacement operator  $D(\alpha)$  are displacement-invariant in phase space along the direction of  $\alpha$  with a period  $|\alpha|$ . Having two code stabilizers  $S_0^x = D(\alpha_X)$  and  $S_0^z = D(\alpha_Z)$  imposes displacement invariance along two non-equivalent directions, which means that all codewords are grids in phase space with a unit cell defined by  $\{\alpha_X, \alpha_Z\}$ . The requirement of commutativity of  $S_0^x$  and  $S_0^z$  imposes a constraint

$$\alpha_X^* \alpha_Z - \alpha_X \alpha_Z^* = 2l_S^2 n i, \quad n \in \mathbb{Z}, \quad (6.1)$$

where  $l_S = \sqrt{2\pi}$ . Here, we consider encoding of a single logical qubit into an oscillator, which corresponds to  $n = 1$ . By parametrizing the complex-valued displacement amplitudes as  $\alpha_X = l_S[M_{22} - iM_{12}]$  and  $\alpha_Z = l_S[iM_{11} - M_{21}]$ , we obtain a grid code with the following stabilizers:

$$S_0^z = D(l_S[iM_{11} - M_{21}]), \quad (6.2)$$

$$S_0^x = D(l_S[M_{22} - iM_{12}]). \quad (6.3)$$



**Figure 6.1 | Wigner functions (numerical) of grid states.**

From the constraint (6.1) we derive a single requirement that a real matrix

$$M = \begin{bmatrix} M_{11} & M_{12} \\ M_{21} & M_{22} \end{bmatrix} \quad (6.4)$$

has a determinant  $\det M = 1$ . This matrix defines the structure of the grid in phase space. Here, we only consider the square grid code, which is obtained with  $M = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ . The hexagonal code with  $M = \sqrt{\frac{2}{\sqrt{3}}} \begin{bmatrix} 1 & 1/2 \\ 0 & \sqrt{3}/2 \end{bmatrix}$  was previously realized in Ref. [42].

The Pauli operators of the logical qubit are defined as

$$X_L = \sqrt{S_0^X} = D(l_S[M_{22} - iM_{12}]/2), \quad (6.5)$$

$$Z_L = \sqrt{S_0^Z} = D(l_S[iM_{11} - M_{21}]/2). \quad (6.6)$$

They satisfy the standard algebraic properties  $X_L^2 = I$ ,  $Z_L^2 = I$ , and  $X_L Z_L = -Z_L X_L$ , inside the code space. Using the identity  $Y_L = -iZ_L X_L$ , we find the third Pauli operator  $Y_L = -iD(l_S[iM_{11} - iM_{12} + M_{22} - M_{21}]/2)$ .

The eigenstates of Pauli  $Z_L$  of the ideal grid code are shown in Fig. 6.1(a). Finite-energy code families can be obtained by regularizing the ideal code through application of an envelope operator [53, 111], with a common choice being  $N_\Delta =$

$\exp(-\Delta^2 n)$ . We show several members of this code family in Fig. 6.1(b-d). Note that such a regularization leads to non-orthogonal states, with fidelity loss due to the finite state overlap that scales as  $\propto \exp[-(3/8)\pi/\Delta^2]/(1 - \exp(-2\Delta^2))$  (see Eq. S28 in Ref. [53]), which is negligible for our choice of  $\Delta = 0.34$ .

## 6.2 Small-Big-Small (SBS) protocol

Here, we describe the SBS protocol, first proposed in Ref. [53, 54] from a new angle. The full QEC circuit in this protocol is shown in Fig. 6.2(a) with nominal parameter values listed in Table 5.1; it implements a channel  $\mathcal{R}_\Delta(\rho) = (\mathcal{R}_\Delta^Z \circ \mathcal{R}_\Delta^X)(\rho)$ . Let  $(K_g^{x/z}, K_e^{x/z})$  denote the Kraus operators of the constituent rank-2 channels  $\mathcal{R}_\Delta^{x/z}$  (we omit the  $\Delta$  subscript from the Kraus operators for simplicity). These operators read:

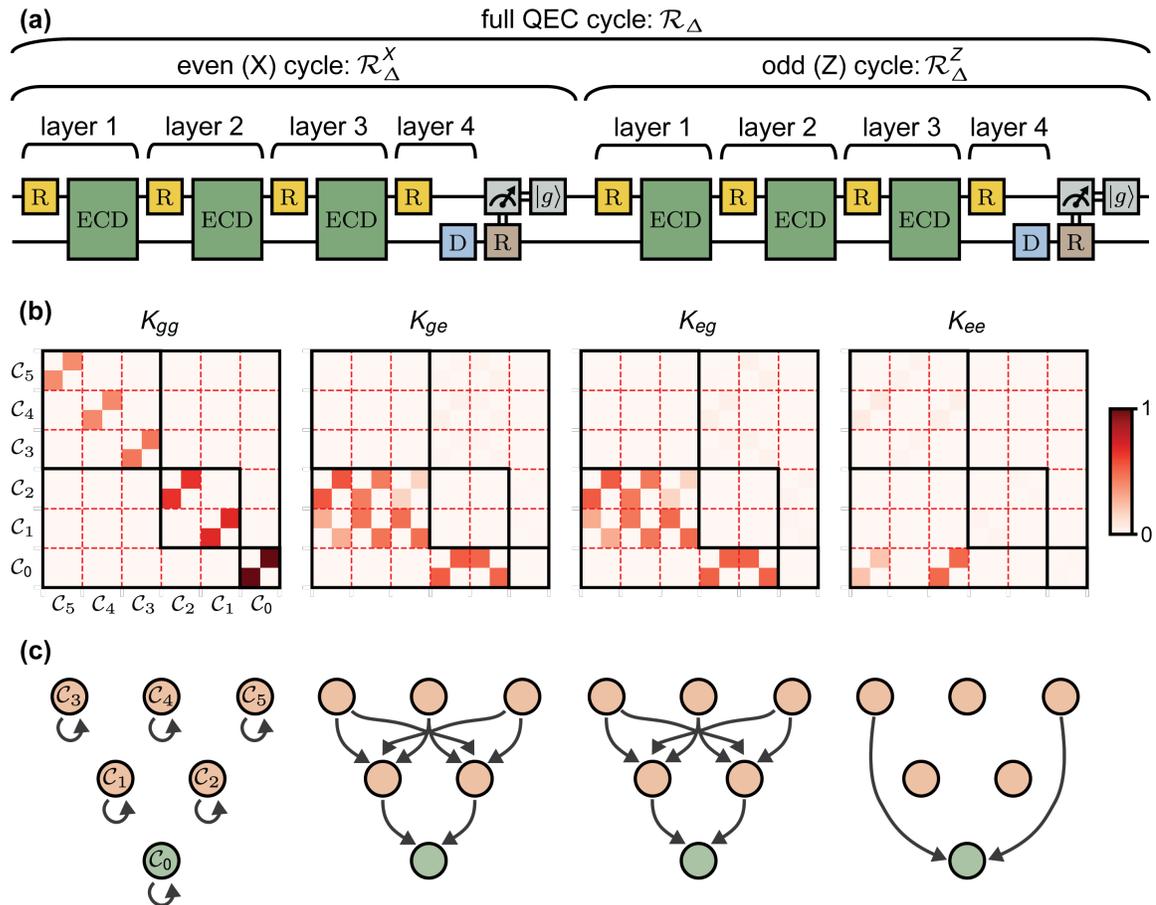
$$K_g^x = \cos(\sqrt{\pi}p) \cos(\sqrt{\pi}\Delta^2 x) + \sin(\pi\Delta^2/2) \cos(\sqrt{\pi}p), \quad (6.7)$$

$$K_e^x = -\cos(\pi\Delta^2/2) \sin(\sqrt{\pi}p) + i \cos(\sqrt{\pi}p) \sin(\sqrt{\pi}\Delta^2 x), \quad (6.8)$$

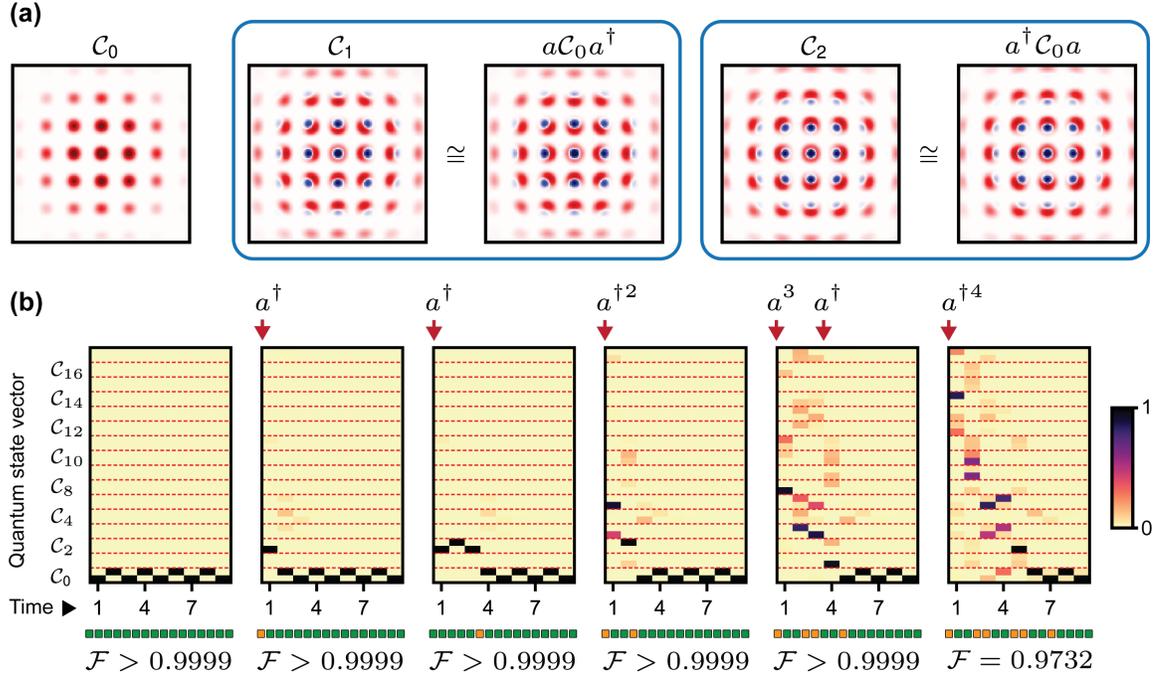
where  $x = (a + a^\dagger)/\sqrt{2}$  and  $p = i(a^\dagger - a)/\sqrt{2}$ , and  $(K_g^z, K_e^z)$  are obtained with a substitution  $(x, p) \rightarrow (-p, x)$ . Then, the Kraus operators of a composite rank-4 channel are:

$$K_{gg} = K_g^z K_g^x, \quad K_{ge} = K_g^z K_e^x, \quad K_{eg} = K_e^z K_g^x, \quad K_{ee} = K_e^z K_e^x. \quad (6.9)$$

For  $\Delta = 0.34$ , these Kraus operators are shown as matrices in the truncated eigenbasis of  $K_{gg}^\dagger K_{gg}$  in Fig. 6.2(b). This eigenbasis splits into pairs of states  $\mathcal{C}_i = \{|0_i^L\rangle, |1_i^L\rangle\}$ ,  $i \in \mathbb{N}$ , that define orthogonal replicas of the logical subspace  $\mathcal{C}_0$  generated by the errors. We show the Wigner functions of the projectors  $\Pi_0$ ,  $\Pi_1$ , and  $\Pi_2$  onto the first three subspaces in Fig. 6.3(a). Note that  $\Pi_1 \approx a \Pi_0 a^\dagger$  and  $\Pi_2 \approx a^\dagger \Pi_0 a$ , hence the errors in the first level of hierarchy resemble photon loss ( $a$ ) and gain ( $a^\dagger$ ) errors. While  $a$  and  $a^\dagger$  only approximately satisfy the Knill-Laflamme conditions [69] for the finite-energy grid code, the actual error operators that define the subspaces  $\mathcal{C}_1$  and  $\mathcal{C}_2$  satisfy these conditions exactly (since the eigenspaces of a Hermitian operator  $K_{gg}^\dagger K_{gg}$  are orthogonal). Similarly, by inspecting the Wigner functions of the projectors onto higher subspaces, we find that the second level of error hierarchy resembles  $a^2$ ,  $a^\dagger a$  and  $a^{\dagger 2}$ . The number of error subspaces in each level is given by the number of unique combinations of  $a$  and  $a^\dagger$ : two subspaces ( $a$



**Figure 6.2 | SBS protocol I.** (a) Circuit structure of one QEC cycle. (b) Kraus operators of the QEC cycle with  $\Delta = 0.34$ , written in the eigenbasis of  $K_{gg}^\dagger K_{gg}$ . This eigenbasis splits into pairs of states  $\mathcal{C}_i = \{|0_i^L\rangle, |1_i^L\rangle\}$ ,  $i \in \mathbb{N}$ , that define replicas of the logical subspace  $\mathcal{C}_0$ . Color encodes the absolute value of the matrix elements. (c) Flow diagram corresponding to each Kraus operator. Circles represent error spaces, and arrows show the most relevant matrix elements. The dynamics within the subspaces is discarded in this representation.



**Figure 6.3 | SBS protocol II.** (a) Numerical Wigner functions of the projectors onto the subspaces  $\mathcal{C}_0$ ,  $\mathcal{C}_1$ , and  $\mathcal{C}_2$ . Comparison to the subspaces generated from the code space by the errors  $a$  and  $a^\dagger$  reveals that the errors in the first level of hierarchy approximately correspond to  $a$  and  $a^\dagger$ . (b) Quantum state trajectories with errors and QEC. The state is represented in the same basis as in Fig. 6.2(b); color encodes the absolute value of the state components in this basis. Red dotted lines are guides to the eye that separate the error subspaces. The occurrence of errors is indicated with red arrows at the top. The time axis is measured in QEC cycles. The Kraus operators are applied between the time steps, and the syndrome string (bottom) encodes which Kraus operator was applied on every step. The state transfer fidelity, shown at the bottom, measures the squared overlap of the final and initial state vectors.

and  $a^\dagger$ ) in the first level, and three subspaces ( $a^2$ ,  $a^\dagger a$  and  $a^{\dagger 2}$ ) in the second level, leading to the blocks of size  $4 \times 4$  and  $6 \times 6$  in the Kraus matrices in Fig. 6.2(b). Further understanding the structure of the error hierarchy is the subject of ongoing research.

Unlike in the standard stabilizer formalism of QEC [4], Kraus operators here do not correspond to a projection of a state onto a single error subspace and its subsequent transfer to the code space. Instead, the transfer here is realized gradually, following an error hierarchy imposed by the QEC circuit. To clarify the action of the Kraus operators, their reduced representation using directional flow of a quantum state between error subspaces is shown in Fig. 6.2(c) [this representation ignores the dynamics within each subspace]. We now briefly discuss the interpretation of

the processes corresponding to each of the  $gg$ ,  $ge$ ,  $eg$ ,  $ee$  outcomes of a QEC cycle. Outcome  $gg$  heralds a process in which the state has remained in the same subspace. The probability of emitting  $gg$  from within the code space is nearly 1. This property is exploited in Chapter 6.5 to extract the expectation value of the code projector  $\langle \Pi_0 \rangle$  from the statistics of long strings of the  $gg/gg/\dots$  type. Both  $ge$  and  $eg$  outcomes herald the process in which the quantum state was transferred one level down the error hierarchy. Strings like  $eg/eg/eg/\dots$  therefore correspond to processes in which the state directionally hops level by level towards the code space. Finally, the  $ee$  outcome heralds a transfer two levels down the error hierarchy.

Besides the transfer between the error spaces, the Kraus operators apply a deterministic logical flip:  $X_L$  in the  $\mathcal{R}_\Delta^x$  cycles, and  $Z_L$  in the  $\mathcal{R}_\Delta^z$  cycles. This flip is visible in the off-diagonal structure of the sub-blocks in the Kraus matrices, see Fig. 6.2(b). For example, the lower right  $2 \times 2$  block in  $K_{gg}$  represents the code subspace, and the off-diagonal structure represents the combined effect of  $X_L Z_L = -iY_L$  on the codewords. Due to this effect, the lifetime of  $+1$  and  $-1$  logical Pauli eigenstates in our QEC protocol are exactly equal. We track the Pauli frame in software, and undo its change in the data reported in Fig. 2.3.

To demonstrate how the errors are corrected by this QEC scheme, we show several examples of quantum state trajectories in Fig. 6.3(b). In the first trajectory, the state is initialized as one of the logical basis states, and then evolved for several QEC cycles without any errors. The Pauli frame switching is apparent here from the oscillating pattern within the code space (in this picture, the phase information is not shown, but the QEC process also protects the phase of the logical qubit). In the second trajectory, an error  $a^\dagger$  was applied to the state prior to QEC, and then it was almost perfectly corrected, accompanied by the emission of  $eg/gg/\dots$  syndrome string. In the third trajectory, this error was instead corrected during the third QEC cycle, and the quantum state spent extra time in the error space  $\mathcal{C}_2$ . This example explicitly demonstrates that the Pauli frame update is applied correctly irrespective of the subspace, hence Pauli gates done in this manner are transversal. The subsequent trajectories demonstrate that even higher-order errors, such as  $a^{\dagger 2}$  or  $a^{\dagger 4}$ , can be corrected with high fidelity. Moreover, as seen in the fifth trajectory, the state can be recovered even if additional errors happen while the previous errors have not yet been fully corrected. The latter example highlights that the “slowness” of the low-rank error-correction dissipation is not a problem, as long as the error rate is sufficiently small compared to the correction rate.

A few remarks with regards to the simplified interpretation of the QEC process

in Chapter 2 are in order: (i) The correct interpretation of the action of a QEC cycle requires considering pairs of outcomes, like  $ge$ , instead of isolated outcomes, like  $g$  or  $e$ . We adopted the latter approach in Chapter 2 for simplicity of exposition. (ii) The  $gg$  outcome does not herald the projection onto the code space, as mentioned in Chapter 2, but rather a process in which “no error was corrected”. Conditioned on the state residing in the code subspace, this outcome will be emitted with probability nearly 1. However, if the state is in one of the error spaces this outcome can still occur with smaller probability starting from about 0.47 at the lowest level in the error hierarchy and reducing for higher levels. (iii) When one of the outcome  $eg$ ,  $ge$  or  $ee$  is obtained, there is a small chance that the QEC process has added an error, leading to a random walk among the error spaces that is heavily biased towards the code space.

## 6.3 QEC cycle: implementation details

In this section, we describe implementation details of a QEC cycle, whose schematic is shown in Fig. 6.2(a). The various datasets in this work were taken with several different versions of the QEC circuit. All these versions have the same overall structure, but different parameter values obtained from re-training after the system drift has appreciably affected the logical performance (this happens on a time scale of 1-2 weeks, see Chapter 6.10). Below, the quoted durations of various components of a QEC cycle refer to the circuit version that we used to collect the system lifetimes dataset and that achieved the highest reported QEC gain.

### SBS unitary

We refer to the unitary part of the circuit  $U_0$  prior to ancilla measurement as “SBS unitary” since it is based on the ansatz from Ref. [53]. The SBS unitary is compiled as a four-layer parametrized circuit with nominal parameters shown in Table 5.1, and is further translated into the pulse sequence with the method described in Chapter 5.1. The last circuit layer does not contain an ECD gate, and instead only contains a qubit rotation and oscillator displacement. Since the qubit is reset after the SBS, the function of the latter rotation is to choose the “reset axis”, which can be an arbitrary axis on the qubit Bloch sphere.

As shown in Ref. [53], without any special asymmetries between  $|g\rangle$  and  $|e\rangle$  it would not matter along which axis the ancilla reset is done – all choices result in

Component	Subcomponent	Duration (ns)
	Enter cycle	24
SBS	Enter SBS	24
	Circuit layer 1	502
	Circuit layer 2	708
	Circuit layer 3	262
	Circuit layer 4	76
	Exit SBS	24
Reset	Enter reset	24
	Roundtrip delay	300
	Acquisition	1400
	Signal processing	332
	Distribution of $s_0$ and $s_1$	100
	Branching and feedback	200
	Exit reset	24
Virtual rotation	Mixer matrix calculation	400
	Mixer update	48
Idle	Delay	452
	Exit cycle	24

**Table 6.1 | Timing of the cycle components.**

the same completely positive trace-preserving map after averaging over the measurement outcomes. However, in practice the asymmetry comes from the ancilla relaxation channel that degrades the readout fidelity of the  $|e\rangle$  state. Hence, it is advantageous to choose the reset axis that preferentially returns the  $|g\rangle$  outcome. The parameter sequence for SBS unitary in Table 5.1 takes this choice into account. The choice of reset axis also results in different unraveling of state trajectories and different Kraus operators. The choice made here enabled the interpretation of  $e$  outcomes as syndromes that signal occurrence and correction of errors, which is utilized in the post-selection experiments, described in Chapter 6.7. This is in contrast with Ref. [42], where  $g$  and  $e$  outcomes are interpreted as left or right displacement of the grid.

The duration of the SBS unitary is not fixed, because its constituent ECD gates can be implemented with different choices of the speed enhancement factor  $\alpha$  (amplitude of the intermediate displacement). Since  $\alpha$  is included in the action space of the RL agent, all circuit candidates during the training have different durations of the SBS unitary (we will soon comment on how this affects the reward comparison among them). In the final circuit that achieved the highest reported QEC gain, the duration of SBS unitary is  $t_{\text{SBS}} = 1546$  ns.

## Ancilla reset

In principle, error correction with SBS protocol could be fully autonomous (without a classical feedback loop) as was envisioned in the proposal [53] and realized in a trapped ion system [54]. The autonomous scheme has an advantage of significantly simplifying the demands on the classical co-processor (in our case, the FPGA). Moreover, there exist various dissipative reset protocols for the transmon [112–114]. However, the disadvantage of a fully autonomous implementation in our system is that it is not able to compensate for a spurious rotation of the oscillator due to the always-on dispersive coupling with the ancilla. The back-action of discarding the ancilla state during the reset is the dephasing of the oscillator – a particularly harmful error channel for the GKP code [53]. Partly because of this reason, we chose to implement ancilla reset through measurement and classical feedback, as described in Chapter 4.3, with the total duration of ancilla reset subroutine of  $t_{\text{reset}} = 2332 \text{ ns}$ .

## Virtual rotation

Due to the always-on dispersive coupling, the oscillator acquires a spurious rotation during the ancilla readout time. In experiment [42], a simple echo sequence was used to cancel this rotation. With such an approach, ancilla spends half of the time in  $|g\rangle$  and half in  $|e\rangle$  regardless of the actual syndrome measurement outcome, which is detrimental to the code due to additional error sources associated with the  $|e\rangle$  state. Here, we instead chose the reset axis which results in 0.9 probability of detecting  $|g\rangle$ . Therefore, the ability to compensate for the spurious oscillator rotation without echoing the state back to  $|e\rangle$  is crucial to maintain this advantage.

We achieve this by dynamically tracking the oscillator phase that stochastically changes due to random ancilla measurement outcomes, and compensating for it with a virtual counter-rotation. The spurious oscillator rotation angle accumulates during the reset time  $t_{\text{reset}}$ , during the time  $t_{\text{VR}}$  that it takes to execute the virtual rotation on the FPGA, and during the idle time  $t_{\text{idle}}$  when ancilla is nominally in  $|g\rangle$  (the latter will be explained shortly). Therefore, in the idealistic dispersive coupling model, the oscillator would rotate by  $\vartheta_g = \chi(t_{\text{VR}} + t_{\text{idle}} + t_{\text{reset}})/2$  if the ancilla is found in  $|g\rangle$ , and  $\vartheta_e = \chi(t_{\text{VR}} + t_{\text{idle}} - t_{\text{reset}})/2$  if it is found in  $|e\rangle$ . Although the  $|f\rangle$  state is not computational, our controller is able to reset it with an accompanying virtual rotation by angle  $\vartheta_f$ . Instead of relying on the simple dispersive coupling model, in experiment we independently calibrate the angles  $\vartheta_{g/e/f}$  with a

variation of out-and-back experiment [51] to account for additional minor timing contributions related to FPGA program entering or exiting a subroutine, etc. These calibrated angles are used to initialize the QEC circuit for training.

Another important aspect of the virtual rotation is the switching between momentum and position quadratures of the oscillator to realize  $\mathcal{R}_{\Delta}^x$  and  $\mathcal{R}_{\Delta}^z$  dissipators. Such switching can be achieved with a rotation of the SBS unitary by  $\pi/2$  in phase space. This results in additional deterministic contribution  $\vartheta_{\text{SBS}}$  in every virtual rotation gate. The value of  $\vartheta_{\text{SBS}}$  is  $\pi/2$  for the square grid code, and  $\pi/3$  for hexagonal grid code.

The virtual rotation gate utilizes a floating point register  $\vartheta$  on the FPGA. During this gate, the FPGA performs the calculation  $\vartheta \leftarrow \vartheta + \vartheta_{g/e/f} + \vartheta_{\text{SBS}}$  with subsequent reconfiguration of the dynamic mixer matrix which applies a rotation transformation to the oscillator pulses before they are being streamed at the DAC. The total duration of the virtual rotation gate that includes all these steps is independent of  $\vartheta$  and is equal to  $t_{\text{VR}} = 448$  ns.

## Idle section

During the agent training, the reward is measured after a fixed number of  $T = 160$  cycles, and not after some fixed physical duration of time. Therefore, it is necessary to keep the duration of a QEC cycle constant across different protocol candidates to ensure a fair reward comparison. At the same time, the agent is able to affect the physical duration of the SBS unitary by changing the speed enhancement factor  $\alpha$  in the ECD gates. To reconcile these two requirements, after the virtual rotation gate we add a section of idle time that is calculated based on the duration of the SBS unitary in each circuit candidate. We constrain the combined duration of SBS unitary and the idle section to be  $2 \mu\text{s}$ . In the circuit that achieved the highest reported QEC gain, the duration of the idle section was  $t_{\text{idle}} = 452$  ns. Note that in the previous QEC experiments that tried to create a long-lived quantum memory [40, 41], the idle section was inserted intentionally to avoid frequently entangling the high-quality oscillator with low-quality ancilla. Here, we find that inserting any additional idle time degrades the performance, hence we kept it nearly to a minimum while still leaving some room for change of the SBS duration by the RL agent.

In Table 6.1, we provide a detailed timing breakdown of all cycle components.

## 6.4 Learned and scripted parameters

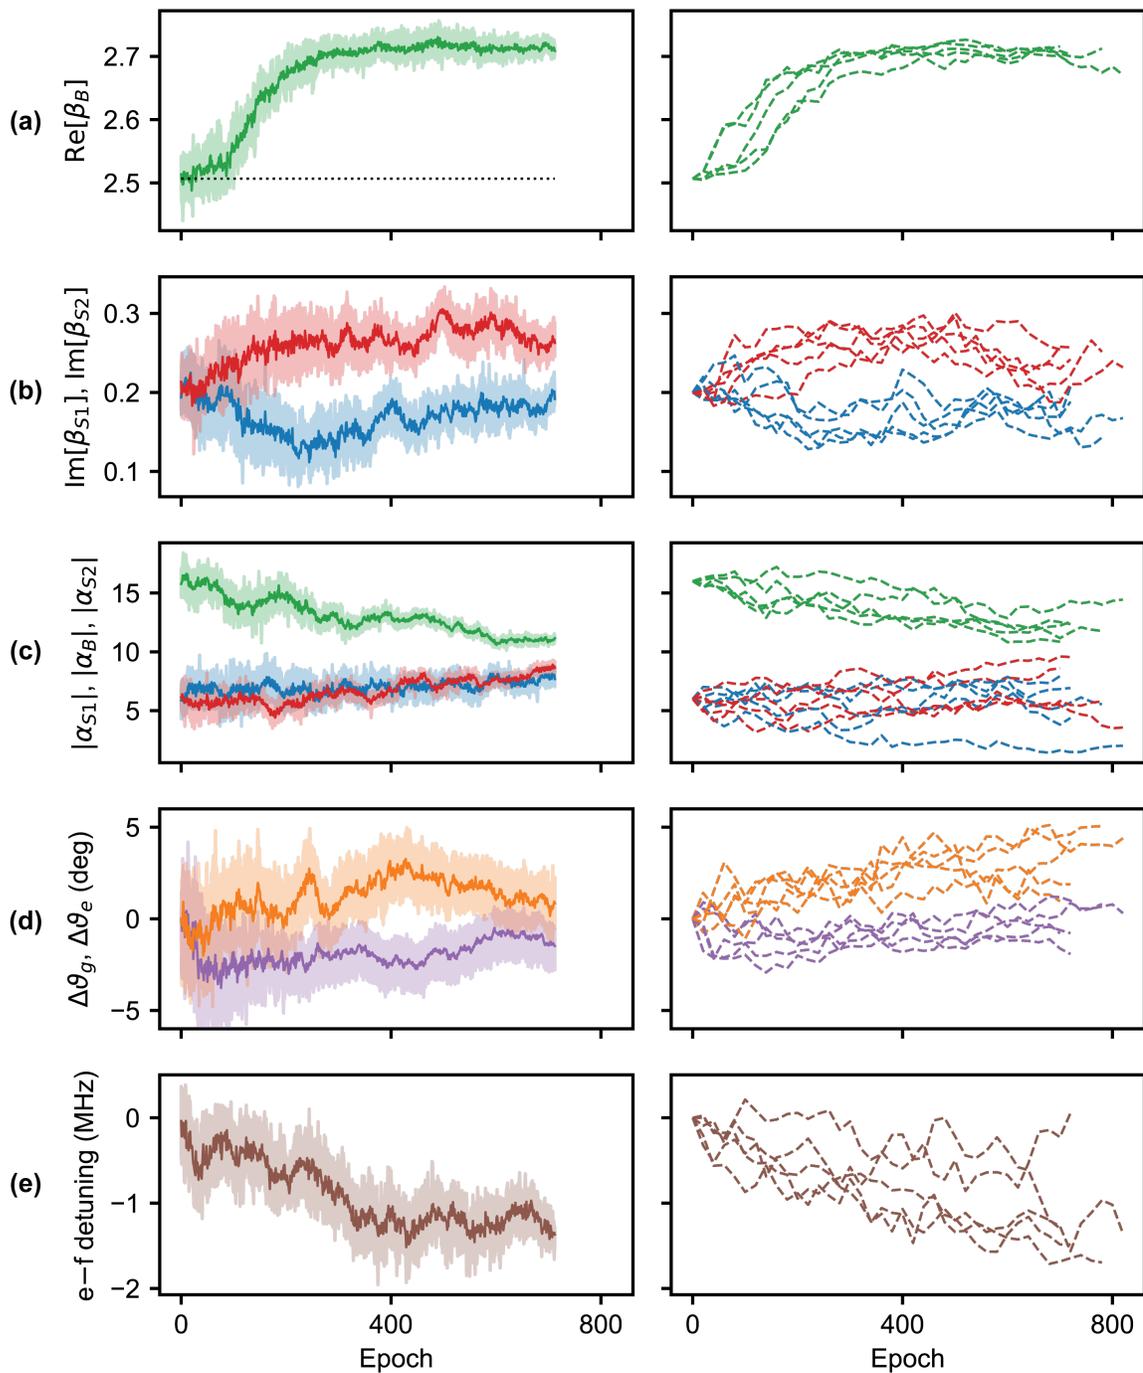
Our QEC protocol has multiple parameters which could be optimized to improve its performance. Some of these parameters are difficult to incorporate into our optimization framework in its current form, and therefore their values are chosen as an approximate compromise between various tradeoffs and then held constant. The rest  $P = 45$  of them are optimized with reinforcement learning, given a reasonable starting point obtained from independent calibrations. Here, we briefly explain the meaning of these parameters.

### Scripted parameters

- Duration, shape, and amplitude of the readout pulse.
- State classification thresholds.
- Timing of all components of the reset subroutine.
- Durations of primitive pulses (qubit rotations and oscillator displacements).
- Combined duration of the SBS unitary and the idle section.

### Learned parameters

- Virtual rotation angle  $(\vartheta_g, \vartheta_e, \vartheta_f)$  for each measurement outcome. It is initialized with a result of independent calibration using a variation of out-and-back experiment.
  - Detuning of transmon  $|g\rangle \leftrightarrow |e\rangle$  and  $|e\rangle \leftrightarrow |f\rangle$  pulses (same parameter for all pulses). It is initialized with a result of independent calibration, when oscillator is in the vacuum state (i.e. when there is no Stark shift).
  - Spectral corrections to  $|g\rangle \leftrightarrow |e\rangle$  and  $|e\rangle \leftrightarrow |f\rangle$  pulses based on derivative reduction by adiabatic gate (DRAG) scheme [82] (same parameter for all pulses). It is initialized with 0.
  - Complex-valued amplitudes of conditional displacement gates in the first three layers of the SBS unitary, and a complex-valued amplitude of the unconditional displacement in the fourth layer. Two small amplitudes are initialized with  $\beta_{S1} = \beta_{S2} = 0.2i$ , and the big amplitude is initialized with  $\beta_B = \sqrt{2\pi}$ . The unconditional displacement is initialized with 0.
  - Magnitudes of the intermediate large displacements used to execute the ECD gates in the first three layers of the SBS unitary. They are initialized with  $\alpha_{S1} = \alpha_{S2} = 6$  for small conditional displacements and  $\alpha_B = 16$  for the big conditional



**Figure 6.4 | Evolution of QEC circuit parameters during the training.** Top row: example from one particular training run. The shaded region bounds the minimal and maximal sampled parameter values. Solid lines indicate the mean. Bottom row: evolution of the mean parameter values in several independent training runs performed during a two-day period, showcasing the reproducibility of the training results.

displacement. Note that changing these parameters also influences the duration of the ECD gates.

- Angular corrections to intermediate large displacements in the first three layers of the SBS unitary. These heuristic parameters compensate for effect of second-order dispersive shift and for the fact that conditional displacement accumulates along an arc of small curvature instead of a straight line. These corrections are initialized with 0.

- Phases and angles of all ancilla rotations in the SBS circuit layers (including the echo pulses inside the ECD gates), and in the ancilla reset subroutine. These parameters are initialized with nominal values from Table 5.1.

- Detuning of the local oscillator (LO) frequency for the cavity mode. This LO is calibrated with spectroscopy and set to be half-way between the number-split oscillator frequencies when qubit is in the states  $|g\rangle$  and  $|e\rangle$ , corresponding to  $\Delta = 0$  in Eq. (4.1).

## Evolution of parameters during training

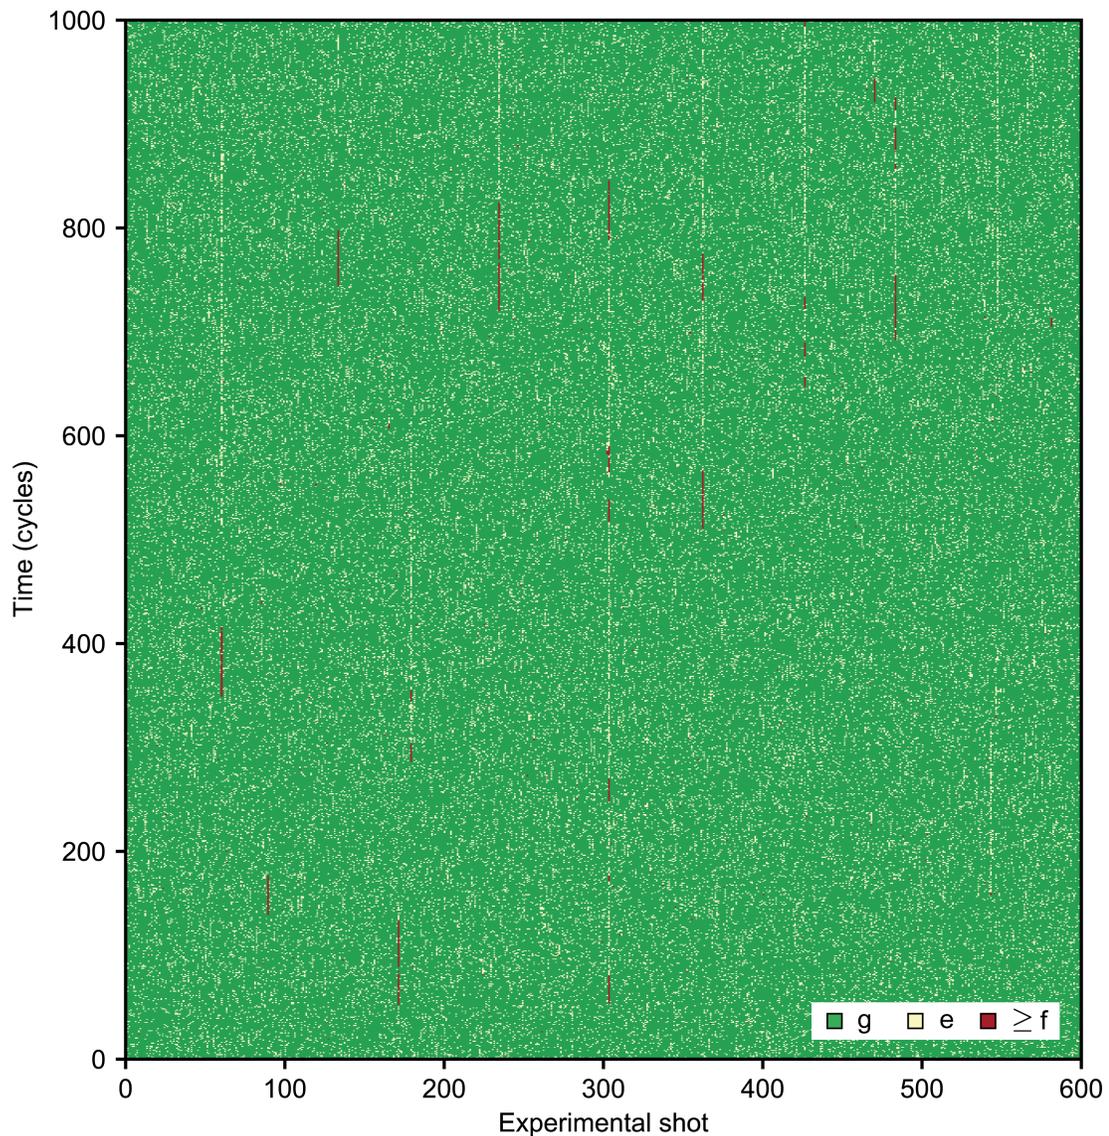
In Fig. 6.4, we show the evolution of several QEC circuit parameters during the training. Most parameters, when initialized well, merely exhibit small fluctuations around the mean. However, some parameters undergo systematic and reproducible changes, as observed in the provided examples. For instance, in Fig. 6.4(a), the big conditional displacement amplitude  $\text{Re}[\beta_B]$ , which we expect to be equal to the size of the grid unit cell, changes from a calibrated value of  $\sqrt{2\pi}$  by about 8%, likely indicating the presence of a miscalibration error (the last calibration was done several weeks prior to this training). Similarly, the trend in Fig. 6.4(e) towards the negative detuning of the  $|e\rangle \leftrightarrow |f\rangle$  pulses could be compensating for an additional Stark shift that was not present at the initial calibration stage (calibration was performed with vacuum state in the oscillator). In Fig. 6.4(b), the trend in amplitudes  $\beta_{S1}$  and  $\beta_{S2}$  of the two small conditional displacements in the SBS unitary is particularly insightful, as it helped us identify a limitation of the proposal in Ref. [53], according to which the amplitudes  $\beta_{S1}$  and  $\beta_{S2}$  should be identical and equal to  $i\Delta^2/2$ , while the RL agent systematically converges to  $|\beta_{S2}| > |\beta_{S1}|$ . Using simulations, we verified that in presence of error channels that act during the execution of the SBS unitary, this is indeed a correct inequality. The optimal ratio of these two amplitudes is found in simulations to be strongly dependent on the error channel. The agent adapts this ratio to the real error channel of our system.

In Ref. [55], it was shown that a similar RL agent is able to converge to correct solutions even starting from completely random parameter initializations. Here, we initialize the parameters close to their expected optimal values through various calibrations, but the compounded effect of small errors (at the level of a few percent) in multiple parameters results in a QEC protocol which, although fully functional, is far from optimal. In particular, we were not able to reach break-even with only the independent calibrations and educated guesses, hence model-free RL can be acknowledged as one of the most crucial factors in the success of this project.

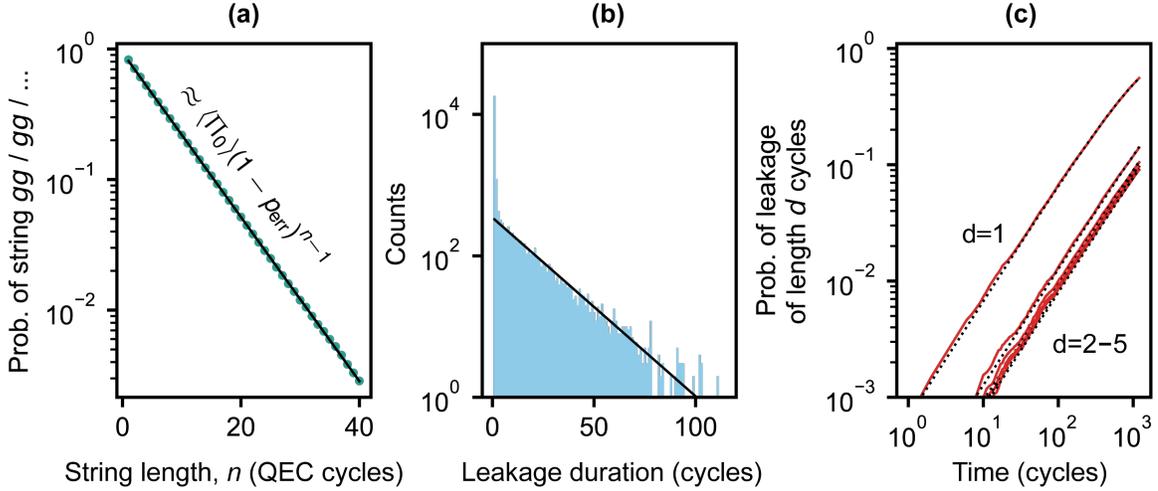
## 6.5 Syndrome measurement statistics

A sample of 600 experimental QEC shots is shown in Fig. 6.5, where QEC is ran for  $T = 1000$  cycles in each shot. Consider a string of measurement outcomes going from any chosen time step  $t_i$  to a time step  $t_{i+2n}$  (it contains  $n$  QEC cycles). The probability  $P([gg]^n)$  that this string contains only  $gg$  outcomes is shown in Fig. 6.6(a), where it is averaged over the experimental shots and over initial times  $t_i$  (the averaging over  $t_i$  is done with a sliding window method which is applicable due to process stationarity). While in the most general case the functional form of  $P([gg]^n)$  is a sum of multiple decaying exponentials, we clearly observe only a single dominant exponential contribution. Hence, we fit this probability to  $P([gg]^n) = a\lambda^n$ , obtaining  $a = 0.936 \pm 0.003$  and  $\lambda = 0.86517 \pm 0.00013$ . By adopting a model for the error process and for the QEC process, we can link the fit parameters  $\{a, \lambda\}$  to model parameters. In general, such a model would be quite complex. However, here we are interested in only two characteristic parameters of the process: the probability  $\langle \Pi_0 \rangle$  of occupying the code space in the dynamical equilibrium of the QEC process, and the probability  $p_{\text{err}}$  of having an error that transfers the state out of the code space. These parameters can be extracted with very minimal model assumptions.

Using the transfer matrix approach, it can be shown that  $p_{\text{err}} \approx 1 - \lambda$  and  $\langle \Pi_0 \rangle \approx a\lambda$  under the following assumptions: 1) The error probability is small  $p_{\text{err}} \ll 1$ , which is justified since the cycle duration is small compared to all relevant error rates in the system, and is confirmed by the fit results; 2) The conditional probability  $P_{gg}^{C_0}$  of emitting  $gg$  when the quantum state is in the code space is nearly 1. This is justified, since in the error-free model of the SBS protocol described in Chapter 6.2 this probability is 0.999; 3) The conditional probability  $P_{gg}^{\text{err}}$  of emitting  $gg$  when the quantum state is in any of the error spaces is small  $P_{gg}^{\text{err}} \ll 1$ . This assumption is



**Figure 6.5 | Syndrome measurement outcomes.** A sample of 600 experimental QEC shots of duration  $T = 1000$  cycles each. The “ $g$ ” outcome (green) is prevalent, heralding the no-error process, while occasional “ $e$ ” outcomes (yellow) indicate correction of errors, and “ $\geq f$ ” outcomes (red) indicate leakage. When transmon escapes to a state higher than  $|f\rangle$ , which is not addressed by our reset scheme, the leakage outcome persists for multiple cycles (streaks of red). In the readout IQ plane such states occur above the  $Q$  threshold, and therefore they are conveniently classified as leakage, but without further identification of the exact leakage state. After transmon stochastically drops back to  $|f\rangle$ , the controller is able to reset it and return the ancilla to the computational manifold. However, during the cycles when ancilla is effectively inactive, the code is not stabilized. Hence, leakage streaks are often followed by streaks of  $e$  outcomes where QEC re-stabilizes the code manifold.



**Figure 6.6 | Analysis of syndrome measurement outcomes.** (a) Probability of a string  $gg/gg/\dots$  as a function of its length, together with the fit to a single exponential decay. (b) Histogram of durations of leakage events. Events of duration 1 or 2 cycles are predominantly  $|f\rangle$  state; longer duration events are likely  $|h\rangle$  or higher excited states. The exponential fit gives the effective lifetime of these leakage states. (c) Fraction of shots that experienced a leakage event of duration  $d$  up to a given time. Dotted lines are fits to a constant-rate model, see Chapter 6.6.

partially justified, since in the error-free model of the SBS protocol this probability is smaller than 0.5 for the first error space, and then monotonously reduces for the higher levels of the error hierarchy; 4) The probability  $P_{gg}^{\text{corr}}$  of correcting an error and emitting  $gg$  is small  $P_{gg}^{\text{corr}} \ll 1$ . In the error-free model of the SBS protocol this probability is exactly zero, while in practice it is limited to  $\sim 10^{-2}$  due to the readout infidelity of the  $|e\rangle$  state.

This result can be intuitively understood as follows: the two most probable system trajectories that generate the string of all  $gg$ 's correspond to (i) starting in the code space and remaining there for  $n$  steps, which happens with probability  $\langle \Pi_0 \rangle (1 - p_{\text{err}})^n$ , and (ii) starting in the code space, remaining there for  $n - 1$  steps, and transitioning out on the very last step, which happens with probability  $\langle \Pi_0 \rangle (1 - p_{\text{err}})^{n-1} p_{\text{err}}$ . The sum of these two contributions equals  $\langle \Pi_0 \rangle (1 - p_{\text{err}})^{n-1} \equiv a \lambda^n$ , leading to  $p_{\text{err}} \approx 1 - \lambda$  and  $\langle \Pi_0 \rangle \approx a \lambda$ . The corrections to these formulas are of the second order in parameters  $\{p_{\text{err}}, 1 - P_{gg}^{\text{C}_0}, P_{gg}^{\text{err}}, P_{gg}^{\text{corr}}\}$ . From the fit in Fig. 6.6(b), we extract  $p_{\text{err}} = 0.13 \pm 0.02$  and  $\langle \Pi_0 \rangle = 0.81 \pm 0.02$ . Note that this result for  $\langle \Pi_0 \rangle$  agrees within the error margin with the result obtained by an independent method based on the reconstruction of the density matrix from the measured Wigner functions in Chapter 6.8. Here, the error of the fit is negligible compared to the model

approximations, hence the quoted error bars are obtained from an estimate of the second-order corrections  $\sim p_{\text{err}}^2 \approx 0.02$ .

Here, we only considered the string of a special type  $gg/gg/\dots$ ; an important avenue of future research would include learning the error channel from the full statistics of syndrome outcomes, using the dataset in Fig. 6.5.

## 6.6 Analysis of transmon leakage

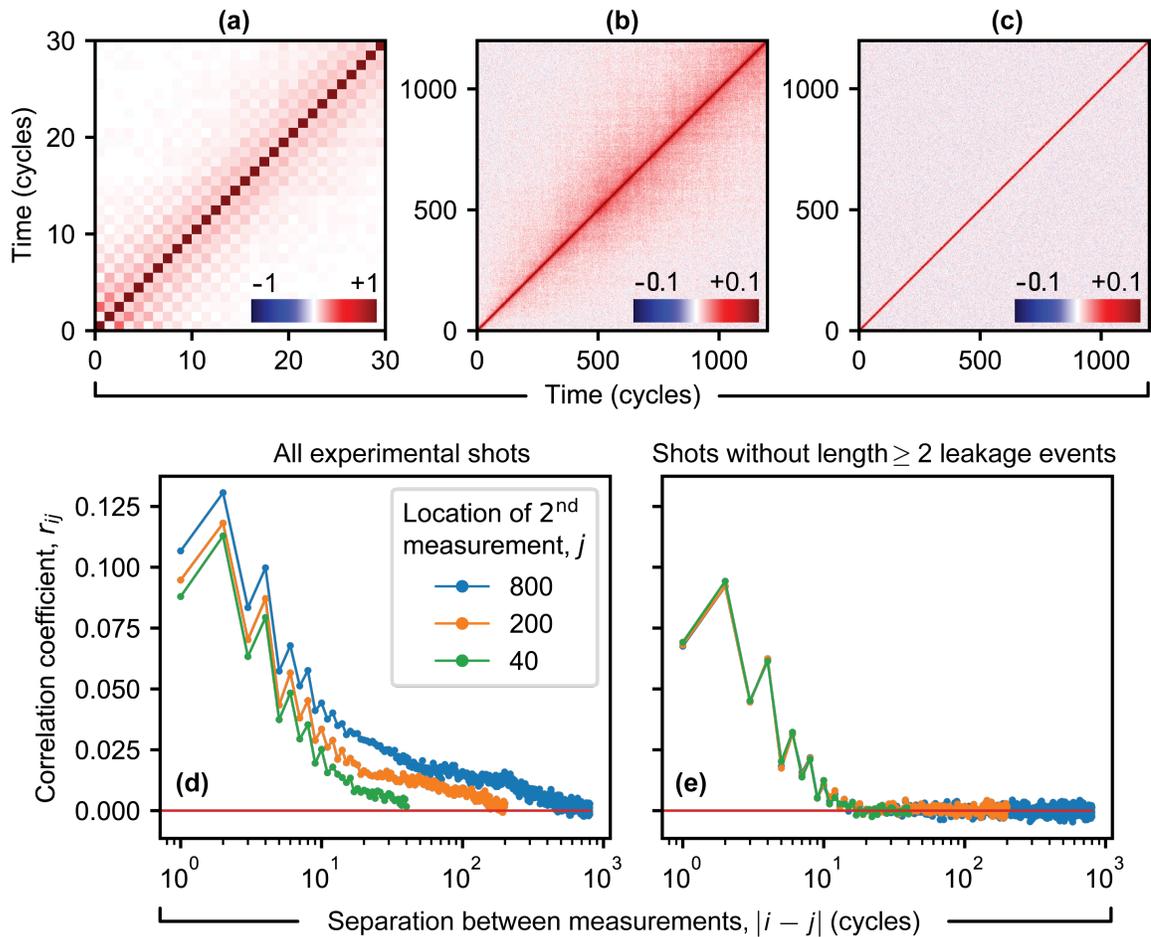
To quantify transmon leakage, we histogram its duration in Fig. 6.6(b). The most likely leakage duration is 1 cycle, since the controller resets  $|f\rangle$  to  $|g\rangle$  with high probability. However, because of the finite readout fidelity of  $|f\rangle$ , shown in Fig. 4.3(b), the controller sometimes fails to reset this state, resulting in the next most probable leakage duration of 2 cycles. After that, the histogram follows exponential distribution with decay constant of 17.2 cycles, corresponding to  $85 \mu\text{s}$ , which we attribute to the effective lifetime of higher leakage states that are not addressed by our reset scheme. This time scale is consistent with our estimate  $\sim 280 \mu\text{s}/3 = 93 \mu\text{s}$  of the  $|h\rangle$  state lifetime, derived from the bosonic statistics and the average measured lifetime of  $|e\rangle$  state.

To extract the leakage rate, in Fig. 6.6(c) we plot the fraction of shots that experienced a leakage event of a certain duration up to a given cycle. We fit the data to a constant-rate model  $L(t) = 1 - \exp(-t/\tau_l)$  where  $t$  is the cycle index and  $p_l = 1/\tau_l$  is the leakage rate (i.e. leakage probability per cycle). We find  $\tau_l = 1480 \pm 10$  cycles, corresponding to a leakage rate of  $p_l = (6.76 \pm 0.04) \times 10^{-4}$ . Similar analysis can be done for leakage events of length  $\geq 2$ ; with the fit to the same model, we find the time scale of  $\tau_{l,\geq 2} = 7820 \pm 10$  cycles, and the corresponding rate of  $p_{l,\geq 2} = (1.280 \pm 0.002) \times 10^{-4}$ . Note that this measurement was performed at the time of slightly sub-optimal performance, and therefore the leakage rate at the maximal achieved QEC gain might have been smaller.

Next, we study the correlation of syndrome measurement outcomes across time. The correlation matrix is given by

$$r_{ij} = \frac{\mathbb{E}[m_i m_j] - \mathbb{E}[m_i]\mathbb{E}[m_j]}{\sqrt{(\mathbb{E}[m_i^2] - \mathbb{E}[m_i]^2)(\mathbb{E}[m_j^2] - \mathbb{E}[m_j]^2)}}, \quad (6.10)$$

where  $m_k$  is the measurement outcome obtained at cycle  $k$ , and empirical expectation values are obtained by averaging across experimental shots. For the dataset



**Figure 6.7 | Correlation of syndrome measurement outcomes.** (a) Correlation matrix  $r_{ij}$  for the first 30 cycles, computed from the dataset in Fig. 6.6(a). (b)  $r_{ij}$  for the full QEC duration of 1200 cycles with zoomed-in color scale to resolve small numbers. (c)  $r_{ij}$  after removing leakage events of duration  $\geq 2$  cycles. (d, e) Cuts of the correlation matrices from (b) and (c) at different locations in the QEC trajectory. In a stationary process,  $r_{ij}$  would only depend on  $|i-j|$ , which is clearly not satisfied in (d). After removing length  $\geq 2$  leakage events, the deviation from stationarity is not resolvable.

of Fig. 6.5 this correlation matrix is shown in Fig. 6.7(a), where we consider only the first 30 cycles. Overall, the correlation is weak, and the correlation between  $\mathcal{R}_\Delta^x$  and  $\mathcal{R}_\Delta^z$  channels (separated by odd number of cycles) is weaker than the correlation between the same-quadrature channels (separated by even number of cycles).

By zooming in the color scale to visually resolve small numbers and considering the full duration of the trajectory of 1200 cycles, as shown in Fig. 6.7(b), it becomes evident that the process is not perfectly stationary. To emphasize this, we show in Fig. 6.7(d) the correlation coefficient  $r_{ij}$  as a function of  $|i-j|$  for several choices of  $j$ . Further along the QEC trajectory the process acquires a correlation tail. Although quite weak, this correlation stretches over hundreds of cycles.

Previously, it was demonstrated that leakage removal helps to reduce correlated errors in the arrays of transmons [115]. Our QEC protocol already contains a mechanism for leakage removal from the  $|f\rangle$  state through measurement-based feedback in every cycle. However, leakage states higher than  $|f\rangle$  are not cleared by our reset. The signature of such leakage events to higher states is two or more consecutive leakage syndrome outcomes. To check the hypothesis that this residual leakage to states higher than  $|f\rangle$  is responsible for increase of correlation, we post-select trajectories that do not have any length-two or longer leakage events. In the post-selected dataset, the correlation matrix does not display any detectable non-stationarity, as seen in Fig. 6.7(c,e), confirming the hypothesis. By fitting the remaining short-time correlations in Fig. 6.7(e) to an exponential decay, we conclude that it takes  $3.9 \pm 0.1$  cycles (approximately 2 QEC cycles) to lose the memory of a typical large error. However, the most probable small errors are corrected in a single QEC cycle.

## 6.7 Post-selection of errors

Here, we provide additional details about the the post-selection experiment that verifies the ability of our QEC scheme to faithfully identify the errors. The post-selection results are summarized in Table 6.2. Note that this experiment was performed at the time of slightly sub-optimal system performance, hence the baseline results with no post-selection are lower than in some other experiments reported here, e.g. in Chapter 6.10. The main conclusion of this post-selection experiment is that it enables significant improvement of the error probability at a cost of only a modest rejection probability.

The saturation of lifetimes in the most stringent post-selection scheme (which

Metric	NP	L	$d \geq 3$	$d \geq 2$	$d \geq 1$
Survival prob. per cycle	1.0000	0.9985	0.9972	0.9907	0.9396
Improvement of $\Gamma_{\text{GKP}}$	1.00	1.10	1.68	2.44	6.31
Lifetime of $ +Z\rangle$ (ms)	$1.874 \pm 0.004$	$2.107 \pm 0.004$	$3.13 \pm 0.02$	$4.60 \pm 0.02$	$10.0 \pm 0.6$
Lifetime of $ +Y\rangle$ (ms)	$1.147 \pm 0.004$	$1.230 \pm 0.005$	$1.93 \pm 0.01$	$2.80 \pm 0.02$	$9.4 \pm 1.0$
Lifetime of $ +Z\rangle$ (cycles)	$381 \pm 1$	$427 \pm 1$	$636 \pm 4$	$934 \pm 4$	$2000 \pm 100$
Lifetime of $ +Z\rangle$ (cycles)	$233 \pm 1$	$250 \pm 1$	$393 \pm 3$	$567 \pm 5$	$1900 \pm 200$

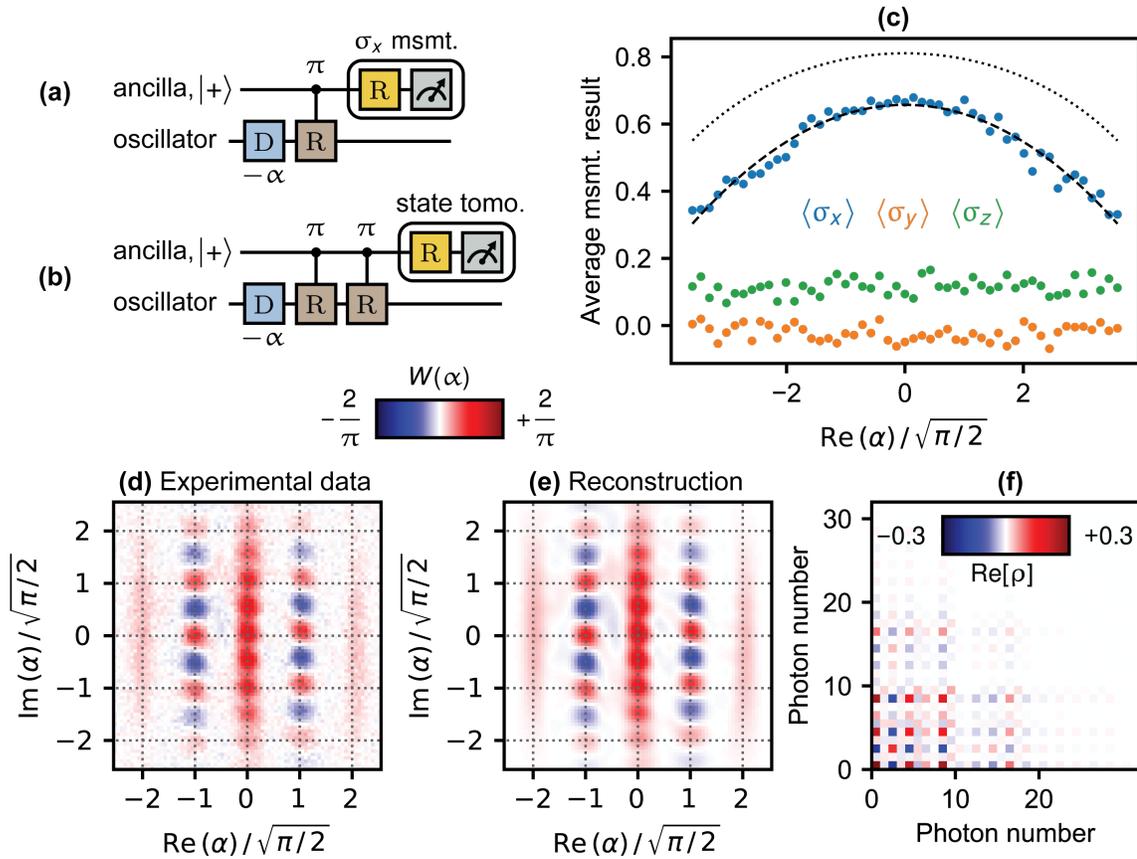
**Table 6.2 | Post-selection results.** Top row labels the post-selection schemes. NP stands for “no post-selection”; L stands for “leakage”;  $d \geq N$  means post-selection that discards trajectories containing strings of  $N$  or more consecutive  $e$  outcomes in the same-quadrature cycles.

preserves only the all- $g$  trajectories) can be related to the following mechanisms: (i) Direct logical errors, which are undetectable in any QEC scheme. (ii) Misclassification of  $e$  as  $g$  (due to ancilla decay during the measurement), which means that some of the all- $g$  trajectories that survived the post-selection actually contained errors, and some of those errors might have been close to a logical operation instead of the identity operation. (iii) The non-orthogonality of logical states, which in our case is not a limiting factor.

## 6.8 Wigner tomography of logical states

### Tomography and its calibration

Wigner tomography is derived from the expression for the Wigner function  $W(\alpha) = (2/\pi)\langle \Pi_\alpha \rangle$ , where  $\Pi_\alpha = D(\alpha)\Pi D^\dagger(\alpha)$  is the displaced parity operator, and  $\Pi = \exp(i\pi a^\dagger a)$  is the photon number parity. The displaced parity operator is unitary and can be measured with phase estimation. It is also hermitian, and hence its eigenvalues are constrained to be  $\pm 1$ . Therefore, it is particularly convenient to measure displaced parity by mapping it onto the qubit observable [116], which is achieved in our system with a circuit shown in Fig. 6.8(a). The conditional rotation gate  $\text{CR}(\pi)$  is realized with a delay of duration  $\pi/\chi$  under the dispersive coupling Hamiltonian, which amounts to approximately  $10 \mu\text{s}$ . Because of such long du-



**Figure 6.8 | State reconstruction.** (a) Wigner tomography experiment. (b) Calibration experiment to extract  $\alpha$ -dependent contrast of the Wigner tomography. This calibration relies on the assumption that the loss of contrast in tomography is primarily due to incoherent errors during the parity mapping gate  $\text{CR}(\pi)$ . (c) Results of the qubit state tomography in the calibration experiment shown in (a). The measurement of  $\langle \sigma_x \rangle$  is fit to a quadratic function of  $|\alpha|$  (black dashed line), and its square root (black dotted line) is used as a Wigner function measurement contrast in the state reconstruction. (d) Experimental Wigner function of the  $|+Z\rangle$  state immediately after initialization. (e) Wigner function of the reconstructed state. (f) Real part of the density matrix of the reconstructed state in the photon number basis.

ration, previous GKP experiments with similarly small  $\chi$  chose to perform state tomography using the characteristic function instead [42, 51]. However, the long coherence of our system allows to measure Wigner function with reasonably high fidelity.

We use several calibration techniques to improve the quality of the subsequent state reconstruction from the tomographic data. First, to symmetrize the effect of ancilla relaxation during the readout, in half of the phase estimation runs we map +1 eigenvalue of  $\Pi_\alpha$  to the  $g$  outcome, and in another half to the  $e$  outcome. This technique eliminates any finite offset in  $W(\alpha)$ , but maintains reduced contrast due to ancilla relaxation and decoherence. Next, to calibrate the contrast reduction, we perform an experiment with a similar circuit in which  $\text{CR}(\pi)$  is replaced with  $[\text{CR}(\pi)]^2 = I$ , see Fig. 6.8(b). We fit the result of this experiment, shown in Fig. 6.8(c), to  $\langle \sigma_x \rangle = 1 - p[\alpha]$ , where  $p[\alpha] = \eta_0 + \eta_2|\alpha|^2$  is the purity loss per  $\text{CR}(\pi)$  gate. Under the assumption that contrast reduction in tomography is primarily due to incoherent processes (ancilla relaxation and dephasing, and oscillator photon loss), the inferred tomography contrast is  $P(\alpha) = \sqrt{1 - p[\alpha]}$ . At  $\alpha = 0$ , this inferred contrast is equal to 0.8, which matches the measured contrast of the Wigner function of vacuum in Fig. 4.1(c), justifying the assumptions of this calibration method.

The phase space points  $\alpha_i$  for Wigner tomography are chosen on a square  $81 \times 81$  grid in a complex plane restricted to  $|\text{Re}[\alpha_i]|, |\text{Im}[\alpha_i]| \leq 3.2$ . We acquire 2400 shots per point in 6 separate acquisition time frames. Between the time frames we perform system performance checks; data acquisition is put on hold if the spurious resonance in  $T_1^t(\bar{n})$  reappears [see Fig. 4.3(d)]. A single state tomography dataset consists of 15.7 million shots, and takes a long time to acquire – from 6 hours in the case of  $T = 0$  cycles, to 26 hours in the case of  $T = 800$  cycles. Therefore, conclusions derived from the analysis of tomography data apply to long-time average system performance.

## State reconstruction

Tomographic data is used to produce a best guess for the density matrix of the state. We parametrize the density matrix as  $\rho = C^\dagger C / \text{Tr}[C^\dagger C]$ , where  $C = A + iB$ , and  $A$  and  $B$  are real-valued matrices. Such parametrization ensures that  $\rho$  is positive semi-definite with trace 1. We truncate the density matrix to dimension  $N = 32$  in photon number basis. Coefficients of matrices  $A$  and  $B$  are optimized using the least squares fit of the Wigner tomography and contrast data, with the cost function

given by

$$\text{cost} = \sum_{i=1}^{N_\alpha} \left( \frac{2}{\pi} \text{Tr}[\rho D(\alpha_i) \Pi D^\dagger(\alpha_i)] \cdot P(\alpha_i) - W(\alpha_i) \right)^2. \quad (6.11)$$

An example Wigner tomography of the  $|+Z\rangle$  grid state together with its reconstruction is shown in Fig. 6.8(d-f). As will be described shortly, using the reconstructed density matrix we extract various parameters of the state: its purity, mean photon number, and envelope size.

## Evolution of logical states

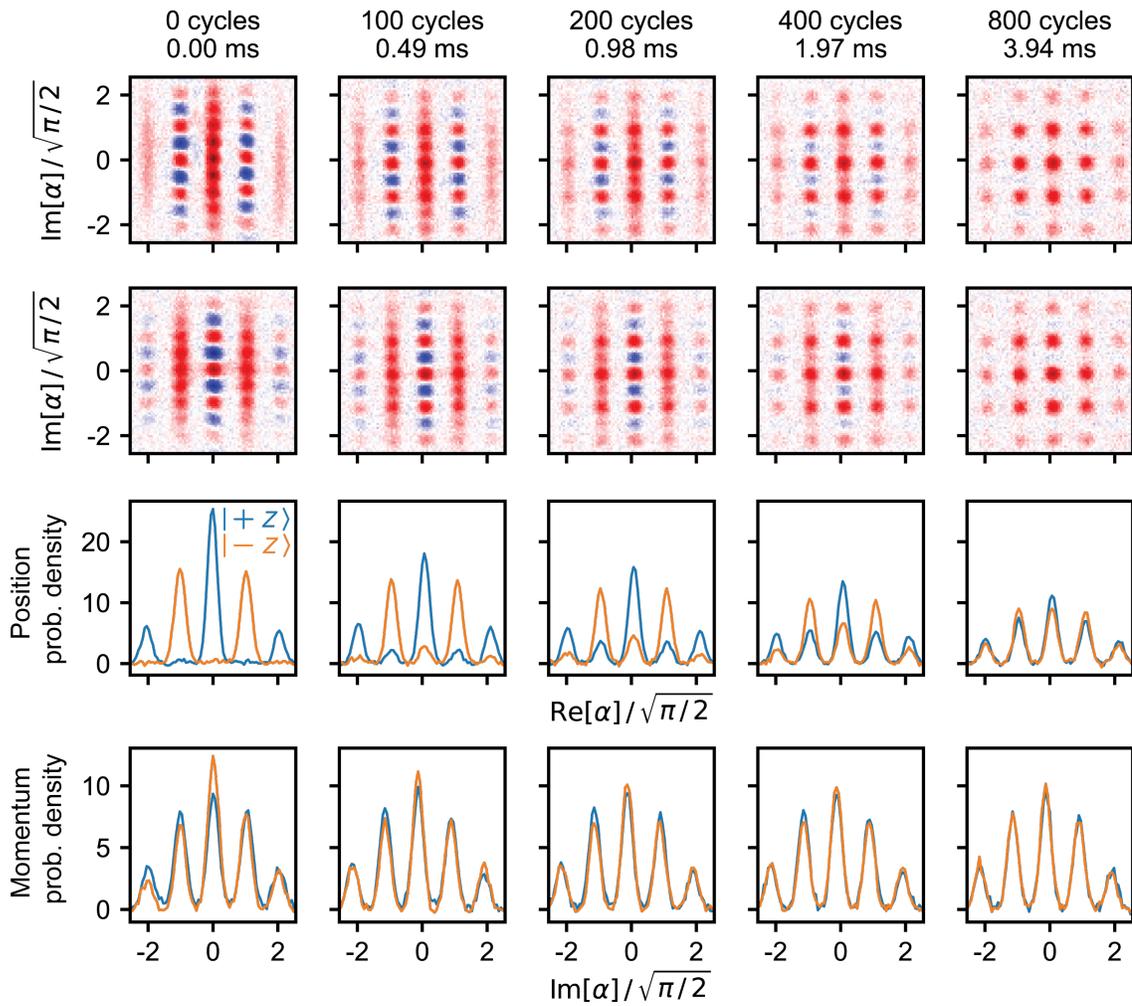
We visualize the evolution of logical  $|+Z\rangle$  and  $|-Z\rangle$  grid states during the QEC by taking Wigner tomography snapshots after 0, 100, 200, 400, and 800 cycles, with results shown in Fig. 6.9.

The marginal of the Wigner function along momentum quadrature gives the probability density of the oscillator position, shown in the third row of Fig. 6.9. The  $|+Z\rangle$  and  $|-Z\rangle$  states have non-overlapping support in position representation, clearly observed in the data at  $T = 0$  cycles. During the QEC process these basis states mix under the logical Pauli channel, which is manifested in the appearance of position peaks of the opposite state, until finally they become almost (but not completely) indistinguishable after  $T = 800$  cycles.

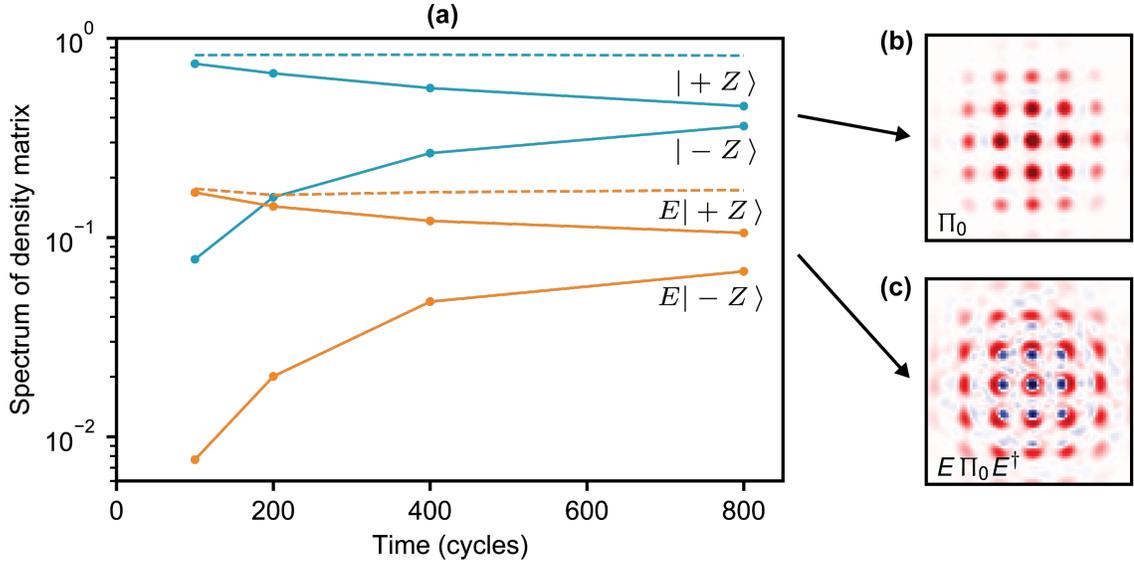
On the other hand, the marginal of the Wigner function along position quadrature gives the probability density of the oscillator momentum, shown in the last row of Fig. 6.9. In the momentum representation,  $|+Z\rangle$  and  $|-Z\rangle$  states share the same support, but have a different pattern of phases associated with the peaks of the wavefunction. The phase information is discarded in the probability density function, which looks identical for both states.

## Spectral analysis of reconstructed states

Focusing on the time evolution of the  $|+Z\rangle$  state, we perform spectral decomposition of its reconstructed density matrices at  $T = 100, 200, 400, 800$ . We find that the eigenvalues of the density matrix are arranged in pairs corresponding to the images of this state and of its complement  $|-Z\rangle$  in different subspaces of the QEC, see Fig. 6.10(a). In particular, we identify only two subspaces with a substantial presence of the state during the QEC process: the code space  $\mathcal{C}_0$ , shown in Fig. 6.10(b),



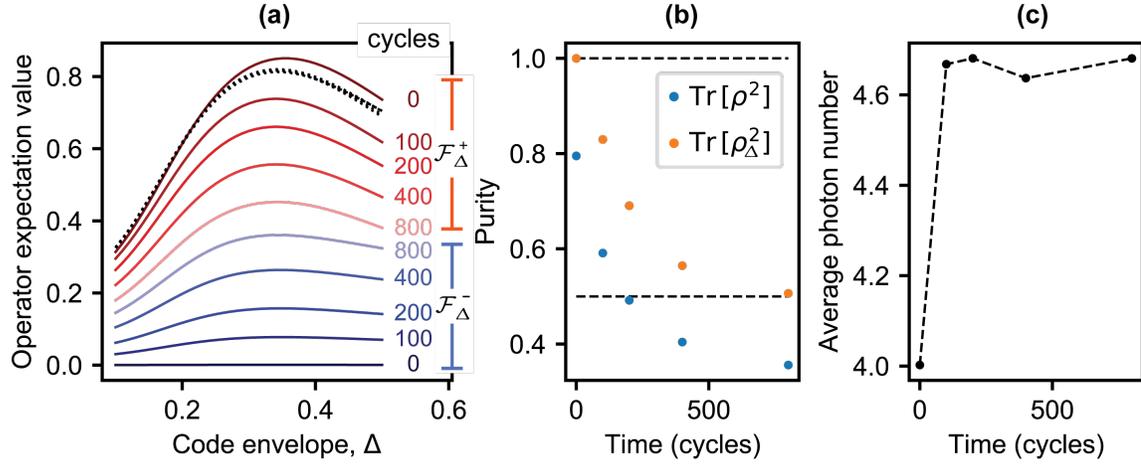
**Figure 6.9 | Wigner tomography after QEC.** Evolution of  $|+\mathcal{Z}\rangle$  state (1st row) and  $|-\mathcal{Z}\rangle$  state (2nd row) is followed for 800 cycles. Color scheme is the same as in Fig. 6.8, with the range scaled to  $[-0.63, 0.63]$ . Marginal of the Wigner function along momentum (position) quadrature, which gives probability density of the oscillator position (momentum), is shown in the 3rd (4th) row in blue for  $|+\mathcal{Z}\rangle$  state, and orange for  $|-\mathcal{Z}\rangle$  state. The probability density is not normalized.



**Figure 6.10 | QEC subspaces.** (a) Spectrum of the reconstructed density matrices of the  $|+Z\rangle$  state evolving under the QEC process. The spectrum separates into two pairs of eigenvalues: one pair corresponds to the code space  $\mathcal{C}_0$ , and another pair corresponds to an error space obtained from the code space by application of an error operator  $E$ . Dashed lines show the sum of the eigenvalues within each pair, which gives the probability of occupying the code space and the error space. (b) Wigner function of the projector  $\Pi_0$  onto the first pair of eigenvectors (taken at  $T = 800$ ), defining the code space. (c) Wigner function of the projector  $E\Pi_0E^\dagger$  onto the second pair of eigenvectors, defining the error space. This error space corresponds to  $\mathcal{C}_2$  subspace in Chapter 6.2.

and the error space  $\mathcal{C}_2$  corresponding to an error  $E$  that most closely resembles  $a^\dagger$ , see Fig. 6.10(c) and Chapter 6.2.

While the QEC circuit imposes the structure of the error subspaces, as described in Chapter 6.2, the properties of the “thermal” distribution across these subspaces in the dynamical equilibrium is defined by the strength of the various error mechanisms in our system as well as the rate at which these errors are corrected. The probability of occupying the code space, given by the sum of the first two eigenvalues, remains constant over time and equal to  $\langle \Pi_0 \rangle = 0.825 \pm 0.003$ , where error bar represents the standard deviation with respect to different durations of the QEC process. This value agrees well with an independent analysis in Chapter 6.5. Having only one relevant error subspace in the steady-state distribution also qualitatively agrees with an observation in Chapter 6.5 that errors are rare. We believe that other error subspaces are populated with probability  $< 1\%$ , which is beyond the resolution power of this method. Developing a more accurate and sample-efficient



**Figure 6.11 | Analysis of  $|+Z\rangle$  evolution.** (a) Expectation value  $\mathcal{F}_\Delta^\pm$  of the state projectors  $|\pm Z_\Delta\rangle\langle\pm Z_\Delta|$  for a range of values of  $\Delta$  and for different durations of the QEC process. The fidelity  $\mathcal{F}_\Delta^+$  decreases, while  $\mathcal{F}_\Delta^-$  increases as a function of time, as expected for a logical Pauli channel. The expectation value of the code projector  $\langle\Pi_\Delta\rangle = \mathcal{F}_\Delta^+ + \mathcal{F}_\Delta^-$  (black dotted lines) remains nearly time-independent for  $T > 0$ . (b) Purity of the reconstructed state  $\rho$  and of its projection onto the code space  $\rho_\Delta$  as a function of time. (c) Average photon number as a function of time.

reconstruction technique for characterizing the distribution across the error spaces is an important direction left for the future.

### Extracting code envelope size

For each QEC duration  $T$ , the reconstructed density matrix is used to find the fidelity of the experimental states to the family of finite-energy codewords  $\{|\pm Z_\Delta\rangle\}$  parametrized by the envelope size  $\Delta$ . For  $T > 0$ , we additionally displace the target codewords by  $(0.08 - 0.12i)\sqrt{\pi/2}$  to account for a small shift visible in the tomography. Since these target states are pure, the fidelity is given by  $\mathcal{F}_\Delta^+ = \text{Tr}[|+Z_\Delta\rangle\langle+Z_\Delta|\rho]$  and  $\mathcal{F}_\Delta^- = \text{Tr}[|-Z_\Delta\rangle\langle-Z_\Delta|\rho]$ . For experiments that start with a preparation of  $|+Z\rangle$ , these fidelities are shown in Fig. 6.11(a) for each QEC duration from the dataset in Fig. 6.9. Immediately after the initialization, the fidelity  $\mathcal{F}_\Delta^+$  is maximized for  $\Delta = 0.36$ , where it reaches 0.85. During the QEC process,  $\mathcal{F}_\Delta^+$  gradually reduces while  $\mathcal{F}_\Delta^-$  increases, consistent with the logical Pauli channel. The sum  $\langle\Pi_\Delta\rangle = \mathcal{F}_\Delta^- + \mathcal{F}_\Delta^+$ , which is equal to the expectation value of the code projector, remains nearly constant for  $T > 0$ . It is maximized at  $\Delta = 0.34$  (9.4 dB), where it is equal to  $\langle\Pi_\Delta\rangle = 0.817 \pm 0.003$ . This value is close to  $\langle\Pi_0\rangle = 0.825 \pm 0.003$  extracted from the density matrix spectrum, which indicates that the code  $\mathcal{C}_0$  stabilized in the

experiment is indeed the one defined by the envelope operator  $\exp(-\Delta^2 n)$  with  $\Delta = 0.34$ .

The purity of  $|+Z\rangle$  state is shown in Fig. 6.11(b). During the QEC process, it reduces below 0.5, since the steady state contains a mixture of the codewords with their images in the error spaces. However, the part of this mixed state that resides within the code space should approach a purity of 0.5. To confirm that this is the case, for every state  $\rho$  we define its projection onto the code space as  $\rho_\Delta = \Pi_\Delta \rho \Pi_\Delta / \text{Tr}[\Pi_\Delta \rho \Pi_\Delta]$ , where we only consider the code with the optimal envelope  $\Delta = 0.34$ . As seen in Fig. 6.11(b), the purity of  $\rho_\Delta$  after initialization is close to 1, and after hundreds of cycles it approaches 0.5, as expected for the logical Pauli channel.

Lastly, in Fig. 6.11(c) we plot the evolution of the average photon number  $\langle n \rangle = \text{Tr}[a^\dagger a \rho]$  during the QEC process. The extracted steady-state photon number is  $\langle n \rangle = 4.67 \pm 0.02$ .

## 6.9 Sensitivity to ancilla errors

### Ancilla phase flips

Our QEC circuit is fault-tolerant with respect to ancilla phase flips by design [53]. To see this, consider that if  $\sigma_z$  error happens during the ancilla readout, it would have no effect because the readout projects the ancilla onto an eigenstate of  $\sigma_z$ . Likewise, during the virtual rotation gate or the idle time the ancilla is nominally in the  $|g\rangle$  state, which is an eigenstate of  $\sigma_z$ . Finally, the effect of  $\sigma_z$  errors on the SBS unitary can be understood by propagating them through the circuit layers. For example, if such an error happens during the big conditional displacement, it is equivalent to changing the circuit parameters from  $\vec{\beta} = l_S \times (i\Delta^2/2, 1, i\Delta^2/2)$  to  $\vec{\beta} = l_S \times (i\Delta^2/2, 1, -i\Delta^2/2)$ . Since  $\Delta^2 \ll 1$ , this change is equivalent to a small error that will be corrected in the following QEC cycles.

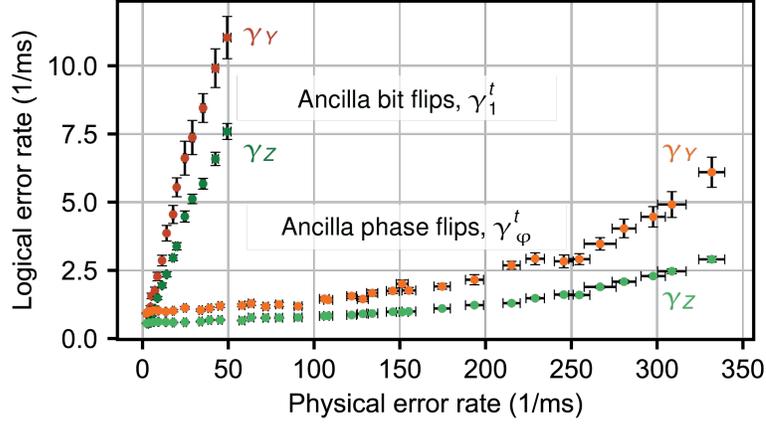
### Ancilla bit flips

In contrast, ancilla bit flips can detrimentally affect the logical qubit in several ways. If such an error happens during the middle half of the big conditional displacement of amplitude  $l_S$ , it will with high probability generate a logical error. This mechanism accounts for a significant fraction of logical errors in the experiment. For

example, its contribution to the Pauli error probability  $p_X = 1.8 \times 10^{-3}$  (per QEC cycle), is estimated to be  $\sim 0.5 \times 0.5 \times 0.5 \times (700 \text{ ns}/280 \mu\text{s}) \approx 0.3 \times 10^{-3}$ , where the factors of 0.5 account for (i) half of the superposition state being sensitive to ancilla decay, (ii) half of the QEC cycle is devoted to position quadrature, (iii) half of the big conditional displacement gate. In practice, the relaxation time of the ancilla is likely degraded during the execution of the conditional displacement due to the large number of intermediate photons in the oscillator, see [51] and evidence in Chapter 6.10. Hence, this estimate provides an optimistic lower bound. Ancilla bit flips can also create detrimental back-action on the oscillator if they happen during the readout time. Since readout outcome is used in a feedback loop to implement a virtual rotation gate, misclassification of the ancilla state generates rotational errors that the GKP code is not well suited to correct, with erroneous rotation angle distributed in the range  $0.0 - 0.6$  radians. To estimate the contribution of such errors to the logical error rate, consider that a rotation by  $\sim \Delta/(l_S/2) \approx 0.3$  radians would diminish the overlap of the blobs in the Wigner function; therefore, a significant fraction of misclassification-induced rotation errors cause large disturbance of the stabilized code space. The transmon  $|e\rangle$  state is the most prone to misclassification. Since readout fidelity of the  $|e\rangle$  state is close to 99%, and this outcome is generated 10% of the time, we estimate an additional  $\sim 0.5 \times 10^{-3}$  contribution to logical error probability per QEC cycle from this mechanism. The two contributions described here account for half of the logical error probability  $p_X$ , and the remaining half is not yet well understood.

### Transmon noise injection

To check the effect of ancilla errors on the logical performance in a controllable way, we perform noise injection experiments that selectively increase the transmon phase-flip rate  $\gamma_\varphi^t$  or bit-flip rate  $\gamma_1^t$ , with the results shown in Fig. 6.12. With noise injection, we are able to increase  $\gamma_1^t$  by a factor of 14 (spoiling  $T_1^t$  from  $290 \mu\text{s}$  to  $20 \mu\text{s}$ ), and  $\gamma_\varphi^t$  by a factor of 140 (spoiling  $T_\varphi^t$  from  $430 \mu\text{s}$  to  $3 \mu\text{s}$ ). Using linear fits in the low-error region, we extract the error sensitivities  $d\gamma_Z/d\gamma_1^t = 0.17$ ,  $d\gamma_Y/d\gamma_1^t = 0.25$ ,  $d\gamma_Z/d\gamma_\varphi^t = 0.0027$ , and  $d\gamma_Y/d\gamma_\varphi^t = 0.0050$ . The derived sensitivity of  $\Gamma_{\text{GKP}} = (\gamma_X + \gamma_Y + \gamma_Z)/3$  to ancilla phase flips is 65 times smaller than the sensitivity to ancilla bit flips, confirming the qualitative arguments provided above. Mitigating the effect of ancilla bit flips on the logical performance is one of the most important future directions in grid-code QEC.

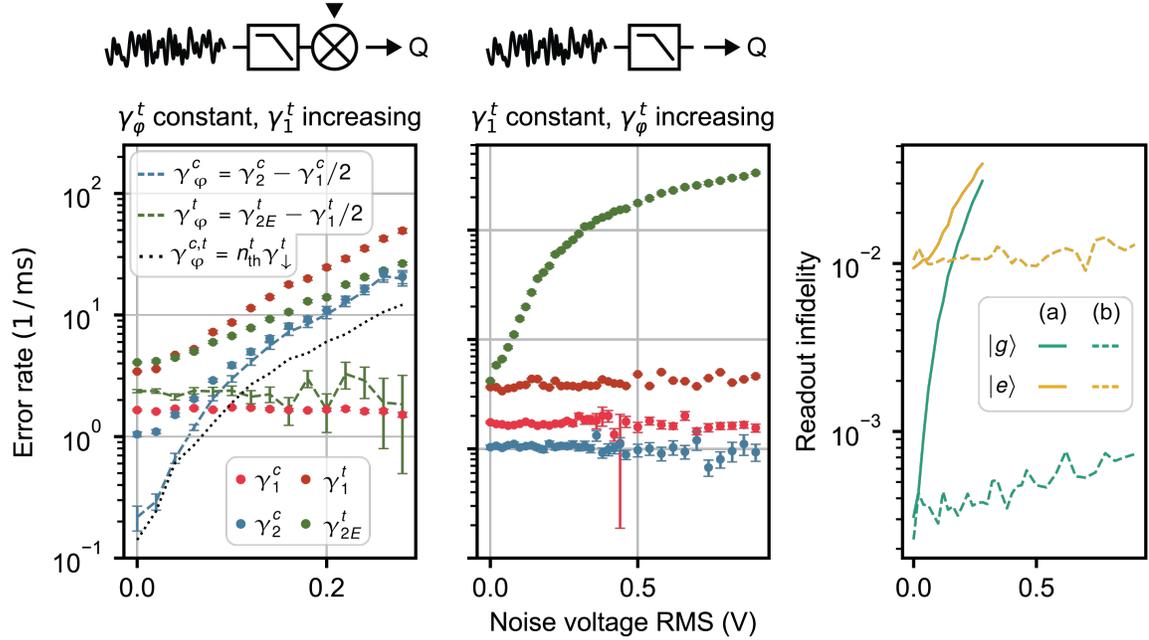


**Figure 6.12 | Effect of ancilla errors.** Logical error rates  $\gamma_Z$  and  $\gamma_Y$  as a function of physical error rates  $\gamma_\phi^t$  and  $\gamma_1^t$  of the ancilla transmon ( $\gamma_X$  is expected to behave identically to  $\gamma_Z$ ). Physical error rates are varied with noise injection. The error sensitivities  $d\gamma_Z/d\gamma_1^t = 0.17$ ,  $d\gamma_Y/d\gamma_1^t = 0.25$ ,  $d\gamma_Z/d\gamma_\phi^t = 0.0027$ , and  $d\gamma_Y/d\gamma_\phi^t = 0.0050$  are extracted by linear fits in the low-error region. The derived sensitivity of  $\Gamma_{\text{GKP}} = (\gamma_X + \gamma_Y + \gamma_Z)/3$  to phase flips is 65 times smaller than to bit flips.

### Verifying noise injection

In the following, we explain how the noise injection experiments were conducted and how we verified that the noise affects the system as intended, i.e. selectively tunes  $\gamma_\phi^t$  or  $\gamma_1^t$ . We are able to achieve high degree of selectivity, with negligibly small spurious effects. To spoil  $\gamma_1^t$ , we inject noise at the transmon frequency, and to spoil  $\gamma_\phi^t$ , we inject noise at low frequency [117]. The baseband white noise with flat spectral density up to 80 MHz is sourced from an Agilent 33250A arbitrary waveform generator. In the  $\gamma_1^t$  tuning experiment, it is upconverted to the qubit frequency using a double-balanced mixer with an LO blue-detuned by 30 MHz from the qubit frequency. After this pre-processing, the noise is filtered and combined with the qubit control line after the switch (in contrast to all control pulses, the noise is not gated).

In Fig. 6.13(a), we inject resonant noise to tune  $\gamma_1^t$ . This noise couples to the  $\sigma_x$  operator and therefore changes  $\gamma_\downarrow^t$  and  $\gamma_\uparrow^t$  symmetrically, which results in increased steady-state population of the qubit, approaching 0.5 at the largest applied noise power. Note that this noise also affects the dephasing rate  $\gamma_{2E}^t$ , but the changes in  $\gamma_{2E}^t$  are explained by changes in  $\gamma_1^t$ : the extracted pure dephasing rate  $\gamma_\phi^t = \gamma_{2E}^t - \gamma_1^t/2$  remains independent of the noise power, as intended in this experiment. The error bars on  $\gamma_\phi^t$  increase at large noise power, because this small rate is extracted as a difference of two large rates. The oscillator dephasing rate  $\gamma_2^c$  is



**Figure 6.13 | Noise verification experiments.** (a) Component error rates as a function of the root mean square (RMS) voltage of the injected noise (at the generator plane). Here, noise is up-converted to the qubit frequency to increase  $\gamma_1^t$ . (b) Same as in (a), but with the baseband noise that increases  $\gamma_\phi^t$ . (c) Readout infidelity of the transmon  $|g\rangle$  and  $|e\rangle$  states in these two noise injection settings.

also affected by the noise, which is explained by the increased rate of qubit up- and down-transitions that dephase the oscillator through the dispersive coupling [45]. The pure dephasing rate  $\gamma_\phi^c = \gamma_2^c - \gamma_1^c/2$  of the oscillator agrees reasonably well with the prediction  $\gamma_\phi^{c,t} = n_{th}^t \gamma_\downarrow^t$  derived from this mechanism (black dotted line). The disagreement at high noise power is under investigation; it likely comes from the breakdown of the simple formula for  $\gamma_\phi^{c,t}$  in the limit where  $n_{th}^t$  is not small.

In Fig. 6.13(b), we inject baseband noise to tune  $\gamma_\phi^t$ . In addition to this desired effect, within the same dynamic range of the noise we observe an undesired increase of the qubit excited state population by a factor of 2 (data not shown), likely due to the heating of the attenuators by the dissipated noise power. Since  $\gamma_\uparrow^t/\gamma_1^t \ll 1$ , the qubit lifetime is not significantly affected by this heating. The lifetime and coherence of the oscillator also remain independent of the noise power. The increase of the error bars on  $\gamma_1^c$  and  $\gamma_2^c$  with the noise power is related to strong degradation of the fidelity of the transmon selective pulse used to read out the population of the oscillator  $|0\rangle$  and  $|1\rangle$  states as described in Chapter 4.5. This pulse has a duration of  $\sim 20 \mu\text{s}$  and it is directly sensitive to the transmon coherence; at the highest injected noise power, where coherence time is spoiled down to  $3 \mu\text{s}$ , the fidelity of

this selective pulse is only a few percent.

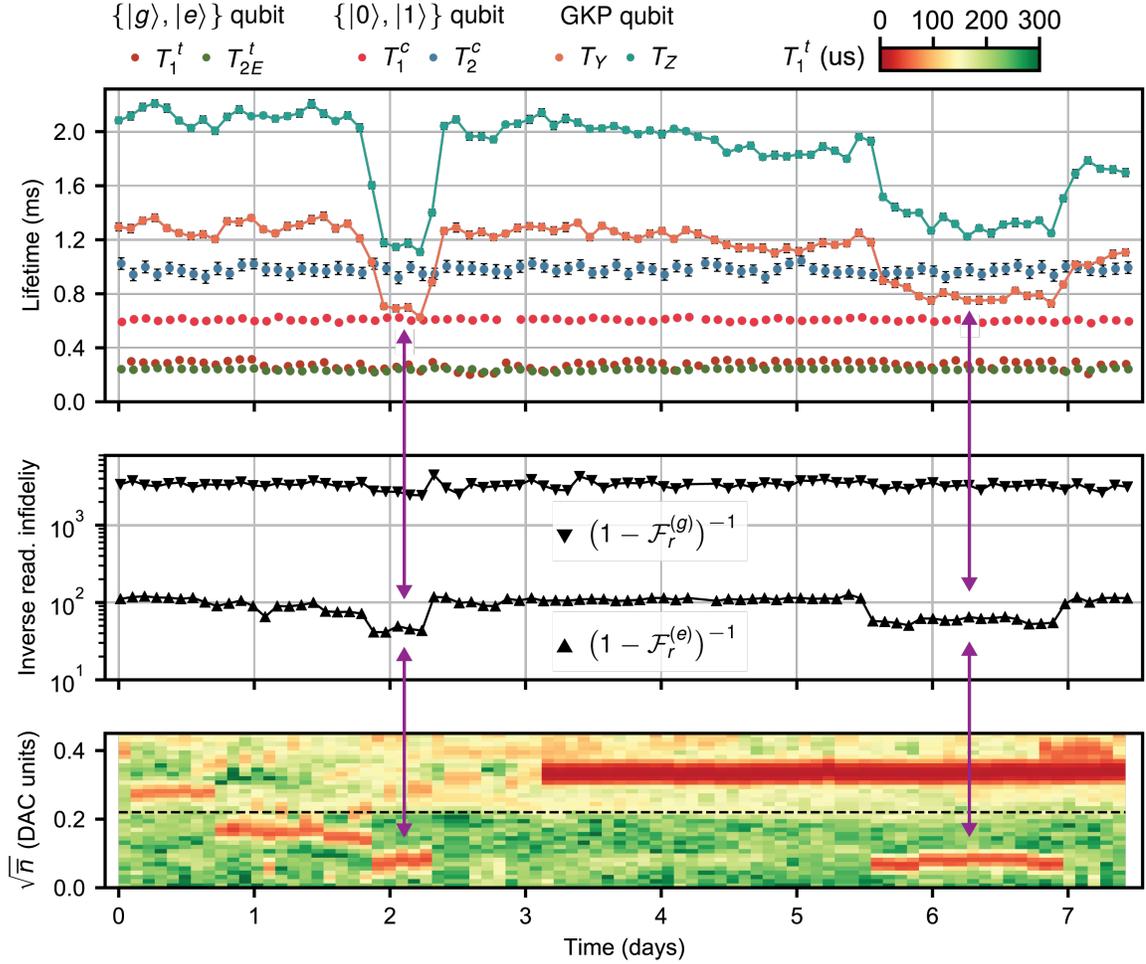
In Fig. 6.13(c), we show the effect of the noise on the readout fidelity of the transmon  $|g\rangle$  and  $|e\rangle$  states. In principle, the noise that induces phase flips ( $\sigma_z$  errors) should not affect the readout of  $\sigma_z$ . However, due to the aforementioned heating of the qubit, we observe a weak degradation of  $\mathcal{F}_r^{(g)}$ . On the other hand, noise at the qubit frequency couples to  $\sigma_x$  and results in significant degradation of both  $\mathcal{F}_r^{(g)}$  and  $\mathcal{F}_r^{(e)}$ .

## 6.10 Long-time system stability

With repetitive measurements of the lifetimes of  $\{|0\rangle, |1\rangle\}$  qubit,  $\{|g\rangle, |e\rangle\}$  qubit, and error-corrected GKP qubit, we investigate the stability of our quantum system over time, with results of a week-long scan shown in Fig. 6.14. We find that the  $\{|0\rangle, |1\rangle\}$  qubit is the most stable, which we attribute to the fact that most of the electromagnetic field is stored in the vacuum of the cavity. In contrast, the  $\{|g\rangle, |e\rangle\}$  qubit exhibits notable fluctuations of the  $|e\rangle$  state lifetime; such fluctuations are often observed in transmons [89, 118], and are typically attributed to two-level defects in the amorphous dielectric, although there are other mechanisms that could lead to such fluctuations, and their source in our system is not yet understood.

We also find significant fluctuations of the lifetime of an error-corrected GKP qubit. Periods of relative stability are regularly interrupted with sudden drops and resurgences of performance, correlated with the appearance and disappearance of a resonant feature in the  $T_1^t(\bar{n})$  dependence, see Fig. 6.14(c). The behavior of readout infidelity of the  $|e\rangle$  state is also correlated with this feature, see Fig. 6.14(b). We find that the correlation coefficient between the logical error rate and the readout infidelity is  $r = 0.81$ . However, preliminary simulations indicate that degradation of the readout fidelity alone is not sufficient to explain the collapses of the logical performance. We therefore believe that the presence of the spurious resonance affects not only the readout fidelity, but also the fidelity of the SBS unitary.

A plausible causal chain is the following: 1) for unknown reason, the spurious degrees of freedom (defects) appear and disappear; 2) when the transmon  $|g\rangle \leftrightarrow |e\rangle$  transition frequency is resonant with the defect, their interaction strength is enhanced, which reduces the lifetime of the  $|e\rangle$  state; 3) during the readout, the transmon is Stark-shifted by the readout photons into resonance with the defects, which increases the probability of readout errors; 4) when the transmon state is misclassified, the virtual rotation gate is executed with an incorrect angle, inducing a phase-



**Figure 6.14 | Quantum system stability.** (a) Lifetimes of Pauli  $Y$  and  $Z$  eigenstates of  $\{|0\rangle, |1\rangle\}$  qubit,  $\{|g\rangle, |e\rangle\}$  qubit, and an error-corrected GKP qubit. (b) Inverse readout infidelity of the transmon  $|g\rangle$  and  $|e\rangle$  states. Logical lifetime is strongly correlated with  $(1 - \mathcal{F}_r^{(e)})^{-1}$ . (c) Transmon lifetime  $T_1^t$  as a function of the number  $\bar{n}$  of the steady-state photons in the readout resonator. The dashed line denotes the DAC amplitude used for the actual readout. The correlated degradation of the system performance and appearance of a spurious resonance that degrades  $T_1^t(\bar{n})$  around  $\sqrt{\bar{n}} = 0.1$  is indicated with purple arrows.

space rotation error on the oscillator; 5) during the conditional displacement gates, the transmon is also Stark-shifted into resonance with the defects by the intermediate photons of the cavity mode; 6) when transmon decay happens during the big conditional displacement gate, it has a significant chance of inducing a logical error. This proposed connection between spurious defects and fluctuations of the logical performance could be verified with detailed system-level simulations that take into account time-dependent Stark shift of the transmon and Stark-shift-dependent degradation of  $T_1^t$ , which is left for the future analysis.

Apart from the stochastic fluctuations, we observe a systematic drift that warrants periodic retraining of the QEC circuit. This drift can be seen by comparing the initial and final data points of the scan in Fig. 6.14, where all the monitored physical error sources are similar in magnitude, but the logical lifetime is reduced in the final point as compared to the initial point. Due to this effect, the various datasets reported in our work were acquired with different version of the QEC circuit that were retrained every 1-2 weeks.

## 6.11 Average channel fidelity

The *average channel fidelity* of a quantum channel  $\mathcal{E} : \rho \rightarrow \mathcal{E}(\rho)$  to a target unitary channel  $\mathcal{U} : \rho \rightarrow U\rho U^\dagger$  is

$$\overline{\mathcal{F}}(\mathcal{E}, \mathcal{U}) = \int d\psi \langle \psi | U^\dagger \mathcal{E}(|\psi\rangle\langle\psi|) U | \psi \rangle, \quad (6.12)$$

where the integral is over the uniform measure on the state space, normalized so that  $\int d\psi = 1$ . Henceforth, we refer to this metric simply as *fidelity*.

To derive an equivalent but experimentally-compatible expression, we make use of the Pauli transfer matrix (PTM) representation of a channel

$$R_{ij}[\mathcal{E}] = \frac{1}{2} \text{Tr}(\sigma_i \mathcal{E}[\sigma_j]), \quad (6.13)$$

where  $\{\sigma_k, k = I, X, Y, Z\}$  are Pauli matrices. This representation has several useful properties, e.g. that composition of channels corresponds to a product of their PTMs [119]. In terms of the PTM, we have the following expression for fidelity:

$$\overline{\mathcal{F}} = \frac{2\mathcal{F}_e + 1}{3}, \quad \mathcal{F}_e = \frac{1}{4} \text{Tr}(R^T[\mathcal{U}] R[\mathcal{E}]), \quad (6.14)$$

where  $\mathcal{F}_e$  is often called the *entanglement fidelity*.

To benchmark a quantum error correction channel, we compare it to an identity channel  $\mathcal{I} : \rho \rightarrow \rho$  with  $R_{ij}[\mathcal{I}] = \delta_{ij}$ . In this case, Eq. (6.14) can be further simplified to

$$\bar{\mathcal{F}} = \frac{1}{12} \sum_{P=X,Y,Z} (\text{Tr}[P \mathcal{E}(|+P\rangle\langle+P|)] - \text{Tr}[P \mathcal{E}(|-P\rangle\langle-P|)]) + \frac{1}{2}, \quad (6.15)$$

where we made use of the identities  $\text{Tr}[\sigma_P \mathcal{E}(\sigma_P)] = \text{Tr}[P \mathcal{E}(|+P\rangle\langle+P|)] - \text{Tr}[P \mathcal{E}(|-P\rangle\langle-P|)]$  for  $P \in \{X, Y, Z\}$ , and  $\text{Tr}(\mathcal{E}[I]) = 2$ . The complete derivation of this formula starting from Eq. (6.12) can be found in Ref. [91]. In experiment, the right-hand side of Eq. (6.15) is measured by preparing  $\pm 1$  Pauli eigenstates, passing them through the channel  $\mathcal{E}$ , and then measuring the corresponding Pauli operator, as cartooned in Fig. 6.15(a). Such a procedure is applicable to an arbitrary duration  $t$  of the channel  $\mathcal{E}(t)$ .

We focus on the comparison of three different qubits in our system:  $\{|0\rangle, |1\rangle\}$ ,  $\{|g\rangle, |e\rangle\}$ , and an error-corrected GKP qubit. The free evolution of the two passive qubits is modeled using a composite amplitude damping and white-noise dephasing channel, while the evolution of an error-corrected GKP qubit is modeled using a logical Pauli channel. Given these well-justified assumptions on the error channels, from Eq. (6.15) we find:

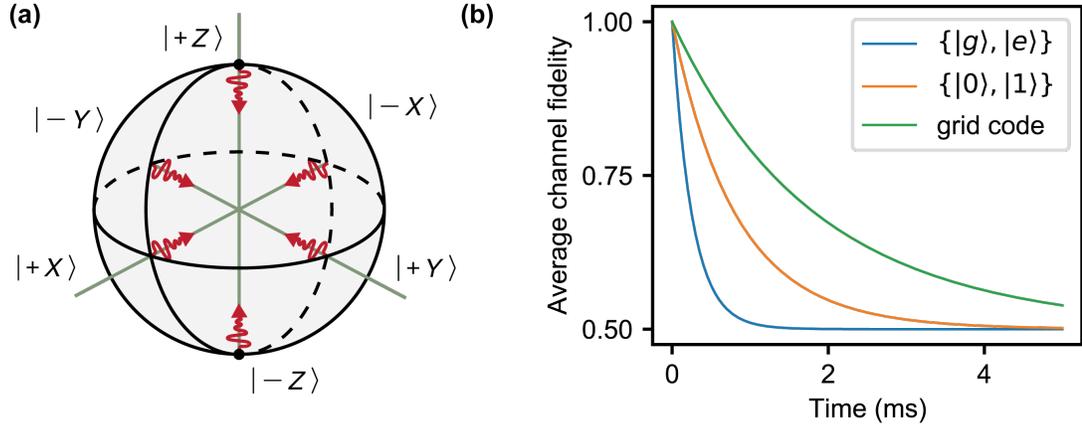
$$\bar{\mathcal{F}}_{\{|0\rangle\}}(t) = \frac{1}{6}e^{-\gamma_1^f t} + \frac{1}{3}e^{-\gamma_2^f t} + \frac{1}{2}, \quad (6.16)$$

$$\bar{\mathcal{F}}_{\{|g\rangle\}}(t) = \frac{1}{6}e^{-\gamma_1^f t} + \frac{1}{3}e^{-\gamma_2^f t} + \frac{1}{2}, \quad (6.17)$$

$$\bar{\mathcal{F}}_{\text{GKP}}(t) = \frac{1}{6}e^{-\gamma_x t} + \frac{1}{6}e^{-\gamma_y t} + \frac{1}{6}e^{-\gamma_z t} + \frac{1}{2}. \quad (6.18)$$

We show the time evolution of fidelity given by Eqs. (6.16–6.18) in Fig. 6.15(b), using experimentally extracted decay rates at the highest QEC gain measured in our experiment.

As seen above, in general the fidelity decays to its steady-state value in a way that cannot be characterized by a single time constant even in the simplest error models such as Pauli noise or amplitude damping. Therefore, fitting the fidelity decay to a single exponential is not strictly valid, although this heuristic approach was adopted in the previous works on bosonic QEC [40, 41, 43]. To avoid such an inconsistency, we consider the channel  $\mathcal{E}$  acting for only a short time  $\delta t$ . Any time dependence of the fidelity, even if it contains multiple exponentially decaying



**Figure 6.15 | Average channel fidelity.** (a) Illustration of channel action on the Pauli eigenstates. By linearity, the evolution of these six cardinal points is sufficient to predict the average effect across the whole Bloch sphere. (b) Expected time evolution of average channel fidelity for three different qubits, calculated using experimentally extracted lifetimes of Pauli eigenstates.

contributions, at short times is equivalent to a linear decay:

$$\overline{\mathcal{F}}(\delta t) = 1 - \frac{1}{2}\Gamma \delta t, \quad (6.19)$$

where  $\Gamma$  is an effective depolarization rate, and  $1/\Gamma$  is the fidelity lifetime. For a depolarizing channel  $\mathcal{E}_{\text{dep}}(\rho) = (1 - p)\rho + p\frac{I}{2}$  with a depolarization probability  $p = 1 - e^{-\gamma t}$ , we have  $\Gamma = \gamma$ , motivating the name and the coefficient  $1/2$  in Eq. (6.19). For qubits considered here, the effective depolarization rates are:

$$\Gamma_{\{01\}} = \frac{2\gamma_2^c + \gamma_1^c}{3}, \quad \Gamma_{\{ge\}} = \frac{2\gamma_{2E}^t + \gamma_1^t}{3}, \quad \Gamma_{\text{GKP}} = \frac{\gamma_X + \gamma_Y + \gamma_Z}{3}. \quad (6.20)$$

We define the coherence gain  $G$  as an improvement of the effective depolarization rate of an error-corrected logical qubit over the best physical qubit in the same system (with the break-even point corresponding to  $G = 1$ ). In a bosonic circuit QED system, the latter is typically the  $\{|0\rangle, |1\rangle\}$  qubit, hence  $G = \Gamma_{\{01\}}/\Gamma_{\text{GKP}}$ . The highest gain achieved in our experiment is  $G_{\text{max}} = 2.27 \pm 0.07$ . During a scan discussed in Chapter 6.10, gain remained above break-even 100% of this week-long time window, with a median of  $\overline{G} = 2.0$ .

Lastly, we acknowledge that the decay constant of the average channel fidelity is not the only relevant metric that we expect to be correlated with the future ability of such systems to participate in quantum computations. Other metrics, such as the

---

SPAM fidelity and the fidelity of gates, are important as well, and we leave their optimization and detailed characterization for future work.

# Outlook

Experiments described in this thesis shed light on many aspects of the grid code and of the QEC in general, but they also generate new questions and research directions. In this final Chapter, we discuss the new opportunities that emerge from our work.

## 7.1 Complete characterization of a single logical qubit

The work in this thesis was devoted to developing a quantum error correction scheme for a single logical qubit. We performed comprehensive characterization of the logical lifetime of such a qubit, but other relevant aspects of its performance were left out (due to limited time). The benchmarking of single-qubit gates, preparation of high-fidelity magic states, and optimal logical readout comprise an important extension of the current work. How can these aspects be addressed to bring to completion the characterization of a single logical GKP qubit?

1. Logical gates are essential for quantum computation. In the ideal grid code, gates from the Clifford group are “cheap” in a sense that they can be implemented by time evolution under Hamiltonians that are quadratic polynomials of  $a$  and  $a^\dagger$  [31]. However, such an approach is not suited for the *finite-energy* code, as it is agnostic to the code envelope size. Close to the completion of this work, we realized (credit goes to S. Singh) that envelope-preserving single-qubit gates can be seamlessly integrated with the QEC process: they require splitting the SBS circuit in half, which creates a Bell pair between the ancilla and the GKP qubit, then implementing the gate on the ancilla, and completing the SBS circuit. This gate teleportation scheme is generically applicable, and a simplified non envelope-preserving version of it was already demonstrated for the grid code in Ref. [42]. With transmon ancilla used in our experiment, the cost of the required ancilla gate is negligible compared to the rest of the control overhead. Hence, we expect that in our experimental system the fidelity of all single-qubit gates done with such a method would be comparable to the fidelity of a QEC cycle, i.e. it would be of the order  $10^{-3}$ . Gate fidelity could be characterized with randomized benchmarking (RB) [120, 121], and reinforcement learning could be used to optimize this fidelity *in situ* similarly to optimization of the QEC process, where the RB survival probability would serve as a reward, see [83–86].

2. Clifford gates are not sufficient for universal quantum computing, which re-

quires an additional “magic” resource [92]. This could be a  $T$  gate, or an eigenstate of this operator, the so-called “magic state”. In our system, the  $T$  gates could be implemented with a method discussed above and characterized using dedicated benchmarking techniques [122, 123]. In addition, the probabilistic preparation of high-fidelity magic states could be achieved using post-selection techniques described in Chapters 2.5 and 6.7.

3. The logical readout of a GKP qubit using phase estimation of the ideal-code Pauli operators is not adapted to the finite code envelope, resulting in limited readout fidelity, as was alluded to in Chapter 2.3. This limitation was not critical for an experiment described here, but it would become important for other information processing tasks at the logical level. High-fidelity readout of the finite-energy grid code could be achieved by explicitly adapting the readout circuit to the code size, as was proposed in Refs. [53, 54, 88]. In practice, the QEC circuit could be trained first to find the optimal code size, and then the readout circuit could be trained to improve the logical readout fidelity.

4. Given the slowness of the error-correcting dissipation, it is unavoidable that the quantum state in the dynamical equilibrium will occupy the code subspace with only a finite probability. It is therefore important that the logical operations (gates and readout) have an equivalent action on all replicas of the code subspace. This property of “transversality” needs to be understood, and there is already some evidence that the gates done in the manner described above will have such a property, see Chapter 6.2.

## 7.2 Promising variations of the QEC protocol

In the future, we expect the development of novel schemes for protecting grid states from decoherence; for example, superconducting circuits could be designed to have the grid code as a degenerate ground state manifold [124]. However, the approach developed in this thesis still leaves a lot of room for improvement. What are the limitations of the QEC scheme utilized here, and how could they be addressed to further advance the logical qubit performance?

1. Our QEC circuit is inherently fault-tolerant with respect to ancilla phase flips and non fault-tolerant with respect to ancilla bit flips, as was demonstrated in Chapter 6.9. Incorporating robustness against ancilla bit flips is crucial for reducing the logical error rate. In bosonic codes that use photon number parity as one of the stabilizers, a similar problem is currently solved with path-independent

control [47, 96]. Extending this framework to the conditional displacement gate is one way to achieve such a robustness. Another way is by using an ancilla with suppressed bit-flip rate, either based on yet-to-be-invented dissipation engineering with a transmon, or using an alternative ancilla, such as a cat qubit stabilized with artificial dissipation [32, 79, 125] or Hamiltonian confinement [93, 126, 127].

2. In the current implementation of the ancilla reset, measurement-based feedback introduces a long latency, resulting in several detrimental effects associated with ancilla excited state, see Chapters 4.3 and 6.3. While both the readout time and the feedback latency could be reduced, a realistic limitation due to the signal travel time and processing time will limit the reset latency to  $\sim 200$  ns after an enormous engineering effort. An alternative path is to utilize dissipative reset [112–114], which would lead to a fully autonomous stabilization and error correction of the grid code, as was already demonstrated in trapped ions [54]. With a dispersively coupled ancilla, discarding its state information into the uncontrolled environment leads to rotation errors on the oscillator. The reset protocol in Ref. [112] dissipates the ancilla to a  $|+\rangle$  state, which would effectively echo out the spurious rotation. This effect is also mitigated in readout techniques that rapidly swap the ancilla state into a readout mode with negligible dispersive coupling to the oscillator [128, 129]. In addition, as our work has shown, it is important to address the leakage states in the reset subroutine. A method from Ref. [114] or its generalization to include higher leakage states is useful from this perspective.

3. An important problem in the field of superconducting circuits is the presence of spurious degrees of freedom whose resonant interaction with the qubit reduces its lifetime. The detrimental effect of such defects on our experiment was demonstrated in Chapter 6.10. While a long-term solution would seek to eliminate these defects altogether, in a short term it would be interesting to try to manipulate them *in situ*. Depending on the nature of these defects, they might be responsive to quasistatic electric field [130], which could be harnessed as a control knob that pushes the defects to a different spectral location. Such a manipulation promises to improve the fidelity of the readout and of the conditional displacement gate done with the current approach [51], as both these operations entail Stark shift and hence require clean spectrum in the vicinity of the qubit frequency.

## 7.3 Machine learning for quantum control

The problem of manipulating the state of a system with the available controls to achieve a certain goal is ubiquitous in all areas of engineering. Analytical solutions to such problems are available only in the simplest cases; their value consists in providing insight, but not precision. In order to realize control of the highest quality, numerical optimization methods are indispensable, and with the recent advances in machine learning the class of problems that are solvable with numerical tools has greatly expanded. The work in this thesis provides the first example of successfully applying machine learning to the problem of quantum error correction. How else could we take advantage of this powerful framework in our attempts to achieve high-quality quantum control?

1. The techniques of dissipation engineering utilized here to implement QEC of the grid code and cooling to the vacuum state can be extended to arbitrary quantum channels. In particular, it is possible with these techniques to realize QEC of other bosonic codes, such as cat codes [32, 40] and binomial codes [33, 41], or stabilization of metrological resource states, such as squeezed states, grid sensor states [131], and states with high degree of rotational symmetry [132]. Any of these applications would involve numerical model-based optimization (for example, see Chapter 5.1 and Refs. [133, 134]), with subsequent model-free optimization (for example, using reinforcement learning, see Chapter 5.2 and Ref. [55]). The above mentioned applications could be realized in the setup developed for this experiment following the same workflow as described in Chapter 2.

2. Implementing a low-rank error correcting dissipation for the grid code allowed us to extend the logical lifetime beyond break-even. In this work, the optimization of the QEC circuit was highly constrained, with a major restriction being the channel rank: our scheme does not utilize the history of syndrome outcomes to realize a more efficient real-time correction of errors. Hence, constructing an error-correcting dissipation of high rank remains an important open problem. The technical capabilities of our FPGA controller are sufficient to implement real-time branching in decision trees of depth  $\sim 1 - 5$ , but the main challenge is to find this branching strategy in the first place. Reinforcement learning will likely be vital for this application [55, 135]. While numerical evidence that grid code QEC could benefit from utilizing the history of syndrome measurements already exists [136], it is unclear if these results extend to our experimental system. This research direction would therefore require first proving the benefits of high-rank error-correcting

dissipation in a simulation, and then porting it to a real experimental system.

3. An even more ambitious and far-reaching program would involve the development of a machine learning framework for discovering not only the QEC strategies for a given code, but also the optimal code in the first place, in a manner adapted to the control limitations and system imperfections. This problem was formulated and partially solved with reinforcement learning for simple examples in Ref. [135], but turning this proof-of-principle solution into a practically useful framework requires a lot of innovative ideas. The work described in this thesis can also be considered as a rudimentary example of such a framework, because our RL agent was able to learn the optimal size of the grid code. We extended the ideas of Ref. [135] to a model-free case, but eliminating further constraints while preserving the practicality would likely take years of interdisciplinary efforts in machine learning and quantum engineering.

4. Eliminating constraints would give the ML algorithms the freedom to explore a larger space of control solutions. However, this would inevitably entail slower learning. The state-of-the-art ML frameworks in other fields that have already undergone this transformation, such as robotics or natural language processing, require days or weeks of training on clusters of hundreds of GPUs. If quantum engineering is to undergo this transformation as well (which, in my opinion, is only a matter of time), low-latency quantum-classical communication loop is absolutely requisite. The training time of several hours in our experiment, as described in Chapter 5.2, was largely limited by the communication time, showing that even the current-day applications would greatly benefit from developing a faster interface.

# Phase space representations

## A.1 Characteristic function

– Definition:

$$C(\alpha) = \text{Tr}[\rho D(\alpha)] \quad (\text{A.1})$$

– Normalization condition:

$$\text{Tr}[\rho] = C(0) = 1 \quad (\text{A.2})$$

– Symmetry property:

$$C(-\alpha) = C^*(\alpha) \quad (\text{A.3})$$

– Overlap of two quantum states:

$$\text{Tr}[\rho_1 \rho_2] = \frac{1}{\pi} \int d^2\alpha C_1(\alpha) C_2^*(\alpha) \quad (\text{A.4})$$

– Purity of a quantum state:

$$\text{Tr}[\rho^2] = \frac{1}{\pi} \int d^2\alpha |C(\alpha)|^2 \quad (\text{A.5})$$

## A.2 Wigner function

– Definition:

$$W(\alpha) = \text{Tr}[\rho D(\alpha) \Pi D^\dagger(\alpha)], \quad \Pi = e^{i\pi a^\dagger a} \quad (\text{A.6})$$

– Normalization condition:

$$\text{Tr}[\rho] = \int d^2\alpha W(\alpha) \quad (\text{A.7})$$

– Overlap of two quantum states:

$$\text{Tr}[\rho_1 \rho_2] = \pi \int d^2\alpha W_1(\alpha) W_2(\alpha) \quad (\text{A.8})$$

– Purity of a quantum state :

$$\text{Tr}[\rho^2] = \pi \int d^2\alpha [W(\alpha)]^2 \quad (\text{A.9})$$

– Relation to characteristic function:

$$W(\beta) = \frac{1}{\pi^2} \int d^2\alpha C(\alpha) e^{\alpha\beta^* - \alpha^*\beta} \quad (\text{A.10})$$

$$C(\alpha) = \int d^2\beta W(\beta) e^{\alpha\beta^* - \alpha^*\beta} \quad (\text{A.11})$$

# Bibliography

- [1] P. W. Shor, “Scheme for reducing decoherence in quantum computer memory”, *Phys. Rev. A* **52**, R2493–R2496 (1995).
- [2] A. M. Steane, “Error correcting codes in quantum theory”, *Phys. Rev. Lett.* **77**, 793–797 (1996).
- [3] R. W. Hamming, “Error detecting and error correcting codes”, *The Bell System Technical Journal* **29**, 147–160 (1950).
- [4] D. Gottesman, “Class of quantum error-correcting codes saturating the quantum Hamming bound”, *Phys. Rev. A* **54**, 1862–1868 (1996).
- [5] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, “Surface codes: Towards practical large-scale quantum computation”, *Physical Review A* **86**, 032324 (2012).
- [6] D. Litinski, “A game of surface codes: large-scale quantum computing with lattice surgery”, *Quantum* **3**, 128 (2019).
- [7] A. Kitaev, “Fault-tolerant quantum computation by anyons”, *Annals of Physics* **303**, 2–30 (2003).
- [8] D. P. DiVincenzo, “The physical implementation of quantum computation”, *Progress of Physics* **48**, 771–783 (2000).
- [9] M. H. Devoret, J. M. Martinis, and J. Clarke, “Measurements of macroscopic quantum tunneling out of the zero-voltage state of a current-biased Josephson junction”, *Phys. Rev. Lett.* **55**, 1908–1911 (1985).
- [10] J. M. Martinis, M. H. Devoret, and J. Clarke, “Energy-level quantization in the zero-voltage state of a current-biased Josephson junction”, *Phys. Rev. Lett.* **55**, 1543–1546 (1985).
- [11] Y. Nakamura, Y. A. Pashkin, and J. Tsai, “Coherent control of macroscopic quantum states in a single-Cooper-pair box”, *Nature* **398**, 786–788 (1999).
- [12] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, “Superconducting qubits: current state of play”, *Annual Review of Condensed Matter Physics* **11**, 369–395 (2020).
- [13] I. Siddiqi, “Engineering high-coherence superconducting qubits”, *Nature Reviews Materials* **6**, 875–891 (2021).
- [14] R. Barends et al., “Superconducting quantum circuits at the surface code threshold for fault tolerance”, *Nature* **508**, 500–503 (2014).

- [15] F. Arute et al., “Quantum supremacy using a programmable superconducting processor”, *Nature* **574**, 505–510 (2019).
- [16] R. Acharya et al., “Suppressing quantum errors by scaling a surface code logical qubit”, arXiv:2207.06431 (2022).
- [17] S. Krinner, N. Lacroix, A. Remm, A. Di Paolo, E. Genois, C. Leroux, C. Hellings, S. Lazar, F. Swiadek, J. Herrmann, et al., “Realizing repeated quantum error correction in a distance-three surface code”, *Nature* **605**, 669–674 (2022).
- [18] Y. Zhao, Y. Ye, H.-L. Huang, Y. Zhang, D. Wu, H. Guan, Q. Zhu, Z. Wei, T. He, S. Cao, et al., “Realization of an error-correcting surface code with superconducting qubits”, *Physical Review Letters* **129**, 030501 (2022).
- [19] J. Preskill, “Quantum computing in the NISQ era and beyond”, *Quantum* **2**, 79 (2018).
- [20] E. Altman et al., “Quantum simulators: architectures and opportunities”, *PRX Quantum* **2**, 017003 (2021).
- [21] B. P. Abbott et al., “Observation of gravitational waves from a binary black hole merger”, *Physical Review Letters* **116**, 1–16 (2016).
- [22] G. Aad et al., “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”, *Physics Letters B* **716**, 1–29 (2012).
- [23] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, “Charge-insensitive qubit design derived from the Cooper pair box”, *Physical Review A* **76**, 042319 (2007).
- [24] J. A. Schreier, A. A. Houck, J. Koch, D. I. Schuster, B. R. Johnson, J. M. Chow, J. M. Gambetta, J. Majer, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, “Suppressing charge noise decoherence in superconducting charge qubits”, *Phys. Rev. B* **77**, 180502 (2008).
- [25] V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret, “Fluxonium: single Cooper-pair circuit free of charge offsets”, *Science* **326**, 113–116 (2009).
- [26] I. V. Pechenezhskiy, R. A. Mencia, L. B. Nguyen, Y.-H. Lin, and V. E. Manucharyan, “The superconducting quasicharge qubit”, *Nature* **585**, 368–371 (2020).

- [27] K. Kalashnikov, W. T. Hsieh, W. Zhang, W.-S. Lu, P. Kamenov, A. Di Paolo, A. Blais, M. E. Gershenson, and M. Bell, “Bifluxon: fluxon-parity-protected superconducting qubit”, *PRX Quantum* **1**, 010307 (2020).
- [28] A. Gyenis, A. Di Paolo, J. Koch, A. Blais, A. A. Houck, and D. I. Schuster, “Moving beyond the transmon: noise-protected superconducting quantum circuits”, *PRX Quantum* **2**, 030101 (2021).
- [29] A. P. Place, L. V. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman, et al., “New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds”, *Nature communications* **12**, 1–6 (2021).
- [30] A. Somoroff, Q. Ficheux, R. A. Mencia, H. Xiong, R. V. Kuzmin, and V. E. Manucharyan, “Millisecond coherence in a superconducting qubit”, arXiv:2103.08578 (2021).
- [31] D. Gottesman, A. Kitaev, and J. Preskill, “Encoding a qubit in an oscillator”, *Physical Review A* **64**, 012310 (2001).
- [32] M. Mirrahimi, Z. Leghtas, V. V. Albert, S. Touzard, R. J. Schoelkopf, L. Jiang, and M. H. Devoret, “Dynamically protected cat-qubits: a new paradigm for universal quantum computation”, *New Journal of Physics* **16**, 045014 (2014).
- [33] M. H. Michael, M. Silveri, R. T. Brierley, V. V. Albert, J. Salmilehto, L. Jiang, and S. M. Girvin, “New class of quantum error-correcting codes for a bosonic mode”, *Physical Review X* **6**, 031006 (2016).
- [34] A. L. Grimsmo, J. Combes, and B. Q. Baragiola, “Quantum computing with rotation-symmetric bosonic codes”, *Physical Review X* **10**, 011058 (2020).
- [35] S. Haroche and J.-M. Raimond, *Exploring the quantum* (Oxford University Press, New York, 2006).
- [36] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, “Cavity quantum electrodynamics for superconducting electrical circuits: an architecture for quantum computation”, *Physical Review A* **69**, 062320 (2004).
- [37] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf, “Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics”, *Nature* **431**, 162–167 (2004).
- [38] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, “Circuit quantum electrodynamics”, *Reviews of Modern Physics* **93**, 025005 (2021).

- [39] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, "A quantum engineer's guide to superconducting qubits", *Applied Physics Reviews* **6**, 021318 (2019).
- [40] N. Ofek, A. Petrenko, R. Heeres, P. Reinhold, Z. Leghtas, B. Vlastakis, Y. Liu, L. Frunzio, S. M. Girvin, L. Jiang, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, "Extending the lifetime of a quantum bit with error correction in superconducting circuits", *Nature* **536**, 441–445 (2016).
- [41] L. Hu, Y. Ma, W. Cai, X. Mu, Y. Xu, W. Wang, Y. Wu, H. Wang, Y. P. Song, C.-L. Zou, S. M. Girvin, L.-M. Duan, and L. Sun, "Quantum error correction and universal gate set operation on a binomial bosonic logical qubit", *Nature Physics* **15**, 503–508 (2019).
- [42] P. Campagne-Ibarcq, A. Eickbusch, S. Touzard, E. Zalys-Geller, N. E. Frattini, V. V. Sivak, P. Reinhold, S. Puri, S. Shankar, R. J. Schoelkopf, L. Frunzio, M. Mirrahimi, and M. H. Devoret, "Quantum error correction of a qubit encoded in grid states of an oscillator", *Nature* **584**, 368–372 (2020).
- [43] J. M. Gertler, B. Baker, J. Li, S. Shirol, J. Koch, and C. Wang, "Protecting a bosonic qubit with autonomous quantum error correction", *Nature* **590**, 243–248 (2021).
- [44] N. Sundaresan, T. J. Yoder, Y. Kim, M. Li, E. H. Chen, G. Harper, T. Thorbeck, A. W. Cross, A. D. Córcoles, and M. Takita, "Matching and maximum likelihood decoding of a multi-round subsystem quantum error correction experiment", arXiv:2203.07205 (2022).
- [45] M. Reagor, W. Pfaff, C. Axline, R. W. Heeres, N. Ofek, K. Sliwa, E. Holland, C. Wang, J. Blumoff, K. Chou, M. J. Hatridge, L. Frunzio, M. H. Devoret, L. Jiang, and R. J. Schoelkopf, "Quantum memory with millisecond coherence in circuit QED", *Physical Review B* **94**, 014506 (2016).
- [46] C. Axline, M. Reagor, R. Heeres, P. Reinhold, C. Wang, K. Shain, W. Pfaff, Y. Chu, L. Frunzio, and R. J. Schoelkopf, "An architecture for integrating planar and 3D cQED devices", *Applied Physics Letters* **109**, 042601 (2016).
- [47] S. Rosenblum, P. Reinhold, M. Mirrahimi, L. Jiang, L. Frunzio, and R. J. Schoelkopf, "Fault-tolerant detection of a quantum error", *Science* **361**, 266–270 (2018).

- [48] S. Krastanov, V. V. Albert, C. Shen, C.-L. Zou, R. W. Heeres, B. Vlastakis, R. J. Schoelkopf, and L. Jiang, “Universal control of an oscillator with dispersive coupling to a qubit”, *Physical Review A* **92**, 040303 (2015).
- [49] R. W. Heeres, B. Vlastakis, E. Holland, S. Krastanov, V. V. Albert, L. Frunzio, L. Jiang, and R. J. Schoelkopf, “Cavity state manipulation using photon-number selective phase gates”, *Physical Review Letters* **115**, 137002 (2015).
- [50] R. W. Heeres, P. Reinhold, N. Ofek, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, “Implementing a universal gate set on a logical qubit encoded in an oscillator”, *Nature Communications* **8**, 94 (2017).
- [51] A. Eickbusch, V. Sivak, A. Z. Ding, S. S. Elder, S. R. Jha, J. Venkatraman, B. Royer, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret, “Fast universal control of an oscillator with weak dispersive coupling to a qubit”, arXiv:2111.06414 (2021).
- [52] S. Ganjam, Y. Wang, Y. Lu, A. Banerjee, C. U. Lei, L. Krayzman, K. Kisslinger, C. Zhou, Y. Jia, M. Liu, L. Frunzio, and R. Schoelkopf, “A multimode approach to microwave loss characterization for on-chip superconducting quantum circuits”, in preparation (2022).
- [53] B. Royer, S. Singh, and S. M. Girvin, “Stabilization of finite-energy Gottesman-Kitaev-Preskill states”, *Physical Review Letters* **125**, 260509 (2020).
- [54] B. de Neeve, T.-L. Nguyen, T. Behrle, and J. P. Home, “Error correction of a logical grid state qubit by dissipative pumping”, *Nature Physics* **18**, 296–300 (2022).
- [55] V. V. Sivak, A. Eickbusch, H. Liu, B. Royer, I. Tsioutsios, and M. H. Devoret, “Model-free quantum control with reinforcement learning”, *Physical Review X* **12**, 011059 (2022).
- [56] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms”, arXiv:1707.06347 (2017).
- [57] C. Gardiner and P. Zoller, *Quantum noise: a handbook of Markovian and non-Markovian quantum stochastic methods with applications to quantum optics* (Springer Science & Business Media, 2004).
- [58] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information* (Cambridge University Press, Cambridge, England, 2010).

- [59] D. I. Schuster, “Circuit quantum electrodynamics”, Ph.D. thesis (Yale University, 2007).
- [60] M. J. Reagor, “Superconducting cavities for circuit quantum electrodynamics”, Ph.D. thesis (Yale University, 2016).
- [61] S. Touzard, “Stabilization of bosonic codes in superconducting circuits”, Ph.D. thesis (Yale University, 2019).
- [62] P. Reinhold, “Controlling error-correctable bosonic qubits”, Ph.D. thesis (Yale University, 2019).
- [63] D. Sank, “Fast, accurate state measurement in superconducting qubits”, Ph.D. thesis (University of California, Santa Barbara, 2019).
- [64] P. Campagne-Ibarcq, “Measurement back action and feedback in superconducting circuits”, Ph.D. thesis (Ecole Normale Supérieure, 2015).
- [65] A. L. Grimsmo and S. Puri, “Quantum error correction with the Gottesman-Kitaev-Preskill code”, *PRX Quantum* **2**, 020101 (2021).
- [66] <https://v-sivak.github.io/2020/06/30/deep-learning-resources>.
- [67] S. Russel and P. Norvig, *Artificial intelligence: a modern approach, 4th ed.* (Pearson, 2020).
- [68] R. S. Sutton and A. G. Barto, *Reinforcement learning: an introduction* (A Bradford Book, 2018).
- [69] E. Knill and R. Laflamme, “Theory of quantum error-correcting codes”, *Phys. Rev. A* **55**, 900–911 (1997).
- [70] K. Noh and C. Chamberland, “Fault-tolerant bosonic quantum error correction with the surface–Gottesman-Kitaev-Preskill code”, *Phys. Rev. A* **101**, 012316 (2020).
- [71] A. S. Darmawan, B. J. Brown, A. L. Grimsmo, D. K. Tuckett, and S. Puri, “Practical quantum error correction with the XZZX code and Kerr-cat qubits”, *PRX Quantum* **2**, 030345 (2021).
- [72] B. M. Terhal, J. Conrad, and C. Vuillot, “Towards scalable bosonic quantum error correction”, *Quantum Science and Technology* **5**, 043001 (2020).
- [73] C. Ryan-Anderson, J. Bohnet, K. Lee, D. Gresh, A. Hankin, J. Gaebler, D. Francois, A. Chernoguzov, D. Lucchetti, N. Brown, et al., “Realization of real-time fault-tolerant quantum error correction”, *Physical Review X* **11**, 041058 (2021).

- [74] L. Egan, D. M. Debroy, C. Noel, A. Risinger, D. Zhu, D. Biswas, M. Newman, M. Li, K. R. Brown, M. Cetina, et al., “Fault-tolerant control of an error-corrected qubit”, *Nature* **598**, 281–286 (2021).
- [75] G. Waldherr, Y. Wang, S. Zaiser, M. Jamali, T. Schulte-Herbrüggen, H. Abe, T. Ohshima, J. Isoya, J. Du, P. Neumann, et al., “Quantum error correction in a solid-state hybrid spin register”, *Nature* **506**, 204–207 (2014).
- [76] M. Abobeih, Y. Wang, J. Randall, S. Loenen, C. Bradley, M. Markham, D. Twitchen, B. Terhal, and T. Taminiau, “Fault-tolerant operation of a logical qubit in a diamond quantum processor”, *Nature* **606**, 884–889 (2022).
- [77] X. Xue, M. Russ, N. Samkharadze, B. Undseth, A. Sammak, G. Scappucci, and L. M. Vandersypen, “Quantum logic with spin qubits crossing the surface code threshold”, *Nature* **601**, 343–347 (2022).
- [78] S. Guadarrama et al., *TF-Agents: a library for reinforcement learning in tensorflow*, <https://github.com/tensorflow/agents>, 2018.
- [79] R. Lescanne, M. Villiers, T. Peronnin, A. Sarlette, M. Delbecq, B. Huard, T. Kontos, M. Mirrahimi, and Z. Leghtas, “Exponential suppression of bit-flips in a qubit encoded in an oscillator”, *Nature Physics* **16**, 509–513 (2020).
- [80] S. Lloyd and L. Viola, “Engineering quantum dynamics”, *Phys. Rev. A* **65**, 010101 (2001).
- [81] C. Shen, K. Noh, V. V. Albert, S. Krastanov, M. H. Devoret, R. J. Schoelkopf, S. M. Girvin, and L. Jiang, “Quantum channel construction with circuit quantum electrodynamics”, *Physical Review B* **95**, 134501 (2017).
- [82] Z. Chen et al., “Measuring and suppressing quantum state leakage in a superconducting qubit”, *Physical Review Letters* **116**, 020501 (2016).
- [83] J. Kelly et al., “Optimal quantum control using randomized benchmarking”, *Physical Review Letters* **112**, 240504 (2014).
- [84] M. A. Rol, C. C. Bultink, T. E. O’Brien, S. R. de Jong, L. S. Theis, X. Fu, F. Luthi, R. F. L. Vermeulen, J. C. de Sterke, A. Bruno, D. Deurloo, R. N. Schouten, F. K. Wilhelm, and L. DiCarlo, “Restless tuneup of high-fidelity qubit gates”, *Physical Review Applied* **7**, 041001 (2017).
- [85] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, “Leakage reduction in fast superconducting qubit gates via optimal control”, *npj Quantum Information* **7**, 14 (2021).

- [86] Y. Baum, M. Amico, S. Howell, M. Hush, M. Liuzzi, P. Mundada, T. Merkh, A. R. Carvalho, and M. J. Biercuk, “Experimental deep reinforcement learning for error-robust gate-set design on a superconducting quantum computer”, *PRX Quantum* **2**, 040324 (2021).
- [87] B. M. Terhal and D. Weigand, “Encoding a qubit into a cavity mode in circuit QED using phase estimation”, *Physical Review A* **93**, 012315 (2016).
- [88] J. Hastrup and U. L. Andersen, “Improved readout of qubit-coupled Gottesman-Kitaev-Preskill states”, *Quantum Science and Technology* **6**, 035016 (2021).
- [89] P. V. Klimov et al., “Fluctuations of energy-relaxation times in superconducting qubits”, *Physical Review Letters* **121**, 90502 (2018).
- [90] J. Lisenfeld, A. Bilmes, A. Megrant, R. Barends, J. Kelly, P. Klimov, G. Weiss, J. M. Martinis, and A. V. Ustinov, “Electric field spectroscopy of material defects in transmon qubits”, *npj Quantum Information* **5**, 1–6 (2019).
- [91] M. A. Nielsen, “A simple formula for the average gate fidelity of a quantum dynamical operation”, *Physics Letters A* **303**, 249–252 (2002).
- [92] S. Bravyi and A. Kitaev, “Universal quantum computation with ideal Clifford gates and noisy ancillas”, *Physical Review A* **71**, 022316 (2005).
- [93] A. Grimm, N. E. Frattini, S. Puri, S. O. Mundhada, S. Touzard, M. Mirrahimi, S. M. Girvin, S. Shankar, and M. H. Devoret, “Stabilization and operation of a Kerr-cat qubit”, *Nature* **584**, 205–209 (2020).
- [94] Z. Chen, K. J. Satzinger, J. Atalaya, A. N. Korotkov, A. Dunsworth, D. Sank, C. Quintana, M. McEwen, R. Barends, P. V. Klimov, et al., “Exponential suppression of bit or phase errors with cyclic error correction”, *Nature* **595**, 383–387 (2021).
- [95] C. K. Andersen, A. Remm, S. Lazar, S. Krinner, N. Lacroix, G. J. Norris, M. Gabureac, C. Eichler, and A. Wallraff, “Repeated quantum error detection in a surface code”, *Nature Physics* **16**, 875–880 (2020).
- [96] W.-L. Ma, M. Zhang, Y. Wong, K. Noh, S. Rosenblum, P. Reinhold, R. J. Schoelkopf, and L. Jiang, “Path-independent quantum gates with noisy ancilla”, *Physical Review Letters* **125**, 110503 (2020).

- [97] S. Puri, A. Grimm, P. Campagne-Ibarcq, A. Eickbusch, K. Noh, G. Roberts, L. Jiang, M. Mirrahimi, M. H. Devoret, and S. M. Girvin, “Stabilized cat in a driven nonlinear cavity: a fault-tolerant error syndrome detector”, *Phys. Rev. X* **9**, 041009 (2019).
- [98] J. A. Gross, C. M. Caves, G. J. Milburn, and J. Combes, “Qubit models of weak continuous measurements: markovian conditional and open-system dynamics”, *Quantum Science and Technology* **3**, 024005 (2018).
- [99] W. Pfaff, C. J. Axline, L. D. Burkhardt, U. Vool, P. Reinhold, L. Frunzio, L. Jiang, M. H. Devoret, and R. J. Schoelkopf, “Controlled release of multiphoton quantum states from a microwave cavity memory”, *Nature Physics* **13**, 882–887 (2017).
- [100] A. P. Read, B. J. Chapman, C. U. Lei, J. C. Curtis, S. Ganjam, L. Krayzman, L. Frunzio, and R. J. Schoelkopf, “Precision measurement of the microwave dielectric loss of sapphire in the quantum regime with parts-per-billion sensitivity”, arXiv:2206.14334 (2022).
- [101] K. Serniak, S. Diamond, M. Hays, V. Fatemi, S. Shankar, L. Frunzio, R. Schoelkopf, and M. Devoret, “Direct dispersive monitoring of charge parity in offset-charge-sensitive transmons”, *Physical Review Applied* **12**, 014052 (2019).
- [102] D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, J. Majer, B. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, “Resolving photon number states in a superconducting circuit”, *Nature* **445**, 515–518 (2007).
- [103] J. Gambetta, A. Blais, D. I. Schuster, A. Wallraff, L. Frunzio, J. Majer, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, “Qubit-photon interactions in a cavity: measurement-induced dephasing and number splitting”, *Physical Review A* **74**, 042318 (2006).
- [104] 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”, *Physical Review Applied* **7**, 054020 (2017).
- [105] D. T. McClure, H. Paik, L. S. Bishop, M. Steffen, J. M. Chow, and J. M. Gambetta, “Rapid driven reset of a qubit readout resonator”, *Physical Review Applied* **5**, 011001 (2016).

- [106] S. Touzard, A. Kou, N. E. Frattini, V. V. Sivak, S. Puri, A. Grimm, L. Frunzio, S. Shankar, and M. H. Devoret, “Gated conditional displacement readout of superconducting qubits”, *Physical Review Letters* **122**, 080502 (2019).
- [107] D. Sank et al., “Measurement-induced state transitions in a superconducting qubit: beyond the rotating wave approximation”, *Physical Review Letters* **117**, 190503 (2016).
- [108] R. Shillito, A. Petrescu, J. Cohen, J. Beall, M. Hauru, M. Ganahl, A. G. Lewis, G. Vidal, and A. Blais, “Dynamics of transmon ionization”, *Physical Review Applied* **18**, 034031 (2022).
- [109] <https://keras.io>.
- [110] [https://github.com/v-sivak/tf\\_quantum\\_simulator](https://github.com/v-sivak/tf_quantum_simulator).
- [111] I. Tzitrin, J. E. Bourassa, N. C. Menicucci, and K. K. Sabapathy, “Progress towards practical qubit computation using approximate Gottesman-Kitaev-Preskill codes”, *Physical Review A* **101**, 032315 (2020).
- [112] K. W. Murch, U. Vool, D. Zhou, S. J. Weber, S. M. Girvin, and I. Siddiqi, “Cavity-assisted quantum bath engineering”, *Physical Review Letters* **109**, 1–5 (2012).
- [113] K. Geerlings, Z. Leghtas, I. M. Pop, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret, “Demonstrating a driven reset protocol for a superconducting qubit”, *Physical Review Letters* **110**, 1–5 (2013).
- [114] P. Magnard, P. Kurpiers, B. Royer, T. Walter, J. C. Besse, S. Gasparinetti, M. Pechal, J. Heinsoo, S. Storz, A. Blais, and A. Wallraff, “Fast and unconditional all-microwave reset of a superconducting qubit”, *Physical Review Letters* **121**, 60502 (2018).
- [115] M. McEwen, D. Kafri, Z. Chen, J. Atalaya, K. Satzinger, C. Quintana, P. V. Klimov, D. Sank, C. Gidney, A. Fowler, et al., “Removing leakage-induced correlated errors in superconducting quantum error correction”, *Nature communications* **12**, 1–7 (2021).
- [116] B. Vlastakis, G. Kirchmair, Z. Leghtas, S. E. Nigg, L. Frunzio, S. M. Girvin, M. Mirrahimi, M. H. Devoret, and R. J. Schoelkopf, “Deterministically encoding quantum information using 100-photon Schrodinger cat states”, *Science* **342**, 607–610 (2013).

- [117] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, “Introduction to quantum noise, measurement, and amplification”, *Reviews of Modern Physics* **82**, 1155–1208 (2010).
- [118] M. Carroll, S. Rosenblatt, P. Jurcevic, I. Lauer, and A. Kandala, “Dynamics of superconducting qubit relaxation times”, arXiv:2105.15201 (2021).
- [119] D. Greenbaum, “Introduction to quantum gate set tomography”, arXiv:1509.02921 (2015).
- [120] E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, M. P. Da Silva, G. A. Keefe, M. B. Rothwell, T. A. Ohki, M. B. Ketchen, and M. Steffen, “Efficient measurement of quantum gate error by interleaved randomized benchmarking”, *Physical Review Letters* **109**, 1–5 (2012).
- [121] M. Kliesch and I. Roth, “Theory of quantum system certification”, *PRX Quantum* **2**, 010201 (2021).
- [122] R. Harper and S. T. Flammia, “Estimating the fidelity of T gates using standard interleaved randomized benchmarking”, *Quantum Science and Technology* **2**, 015008 (2017).
- [123] J. Helsen, X. Xue, L. M. K. Vandersypen, and S. Wehner, “A new class of efficient randomized benchmarking protocols”, *npj Quantum Information* **5**, 71 (2019).
- [124] M. Rymarz, S. Bosco, A. Ciani, and D. P. DiVincenzo, “Hardware-encoding grid states in a nonreciprocal superconducting circuit”, *Physical Review X* **11**, 011032 (2021).
- [125] Z. Leghtas, S. Touzard, I. M. Pop, A. Kou, B. Vlastakis, A. Petrenko, K. M. Sliwa, A. Narla, S. Shankar, M. J. Hatridge, et al., “Confining the state of light to a quantum manifold by engineered two-photon loss”, *Science* **347**, 853–857 (2015).
- [126] S. Puri, S. Boutin, and A. Blais, “Engineering the quantum states of light in a Kerr-nonlinear resonator by two-photon driving”, *npj Quantum Information* **3**, 1–7 (2017).
- [127] N. E. Frattini, R. G. Cortiñas, J. Venkatraman, X. Xiao, Q. Su, C. U. Lei, B. J. Chapman, V. R. Joshi, S. Girvin, R. J. Schoelkopf, et al., “The squeezed Kerr oscillator: spectral kissing and phase-flip robustness”, arXiv:2209.03934 (2022).

- [128] T. Peronnin, D. Marković, Q. Ficheux, and B. Huard, “Sequential dispersive measurement of a superconducting qubit”, *Physical Review Letters* **124**, 180502 (2020).
- [129] E. I. Rosenthal, C. M. Schneider, M. Malnou, Z. Zhao, F. Leditzky, B. J. Chapman, W. Wustmann, X. Ma, D. A. Palken, M. F. Zanner, et al., “Efficient and low-backaction quantum measurement using a chip-scale detector”, *Physical Review Letters* **126**, 090503 (2021).
- [130] J. Lisenfeld, A. Bilmes, and A. V. Ustinov, “Enhancing the coherence of superconducting quantum bits with electric fields”, *arXiv:2208.01570* (2022).
- [131] K. Duivenvoorden, B. M. Terhal, and D. Weigand, “Single-mode displacement sensor”, *Physical Review A* **95**, 012305 (2017).
- [132] W. Wang, Y. Wu, Y. Ma, W. Cai, L. Hu, X. Mu, Y. Xu, Z.-J. Chen, H. Wang, Y. P. Song, H. Yuan, C.-L. Zou, L.-M. Duan, and L. Sun, “Heisenberg-limited single-mode quantum metrology in a superconducting circuit”, *Nature Communications* **10**, 4382 (2019).
- [133] T. Fösel, S. Krastanov, F. Marquardt, and L. Jiang, “Efficient cavity control with SNAP gates”, *arXiv:2004.14256* (2022).
- [134] R. Porotti, V. Peano, and F. Marquardt, “Gradient ascent pulse engineering with feedback”, *arXiv:2203.04271* (2022).
- [135] T. Fösel, P. Tighineanu, T. Weiss, and F. Marquardt, “Reinforcement learning with neural networks for quantum feedback”, *Physical Review X* **8**, 031084 (2018).
- [136] K. H. Wan, A. Neville, and S. Kolthammer, “Memory-assisted decoder for approximate Gottesman-Kitaev-Preskill codes”, *Physical Review Research* **2**, 043280 (2020).