

Enhancing the Lifetime of Quantum Information with Cat States in Superconducting Cavities

A. A. Petrenko



Abstract

Enhancing the Lifetime of Quantum Information with Cat States in Superconducting Cavities

Andrei Alexei Petrenko
2017

The field of quantum computation faces a central challenge that has thus far impeded the full-scale realization of quantum computing machines: decoherence. Decoherence is a general process by which quantum bits, or qubits, interact in unknown ways with their environment and thereby corrupt the information encoded within them. Remarkably, however, protocols for Quantum Error Correction (QEC) exist, and their discovery was a critical advance in the pursuit to build practical quantum computers. To implement QEC, one redundantly encodes a qubit in a higher dimensional space using quantum states with carefully tailored symmetry properties. Projective measurements of these parity-type observables provide error syndrome information with which errors can be corrected via simple operations. Reaching the “break-even” point, at which a logical qubit’s lifetime exceeds the lifetime of the system’s qubit constituents, has thus far remained an outstanding goal. In this work, we implement QEC within a superconducting cavity Quantum Electrodynamics (cQED) architecture that exploits the advantages of encoding quantum information in superpositions of coherent states, or cat states, in highly coherent superconducting cavities. This hardware-efficient approach, termed the cat code, simplifies the encoding scheme and requires the extraction of just one error syndrome via single-shot photon number parity measurements. By implementing the cat code within a full QEC system, we demonstrate for the first time quantum computing that reaches the break-even point. Beyond applications to error correction, logical qubit encodings based on the cat code paradigm can be used to probe more fundamental questions of quantum entanglement between physical qubits and coherent states. Specifically, we demonstrate the violation of a Bell inequality in such a setup, which underscores our ability to efficiently extract information from continuous variables encodings. These results highlight the power of novel, hardware-efficient qubit encodings over traditional QEC schemes. Furthermore, they advance the field of experimental error correction from confirming the basic concepts to exploring the metrics that drive system performance and the challenges in implementing a fault-tolerant system.

Enhancing the Lifetime of Quantum Information with
Cat States in Superconducting Cavities

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

by
Andrei Alexei Petrenko

Dissertation Director: Dr. Robert J. Schoelkopf

May 2017

© 2017 by Andrei Alexei Petrenko

All rights reserved.

To Мама, Тато, Наталка, та Sierra

Acknowledgements

The list of people to whom I owe a great debt is long indeed. My colleagues in lab, both past and present, have guided me through problems ranging from fridge cooldown procedures to the fundamentals of quantum computing with endless reserves of patience and generosity. Hanhee Paik, Gerhard Kirchmair, and Adam Sears introduced me to the lab and its inner workings. It is through them that I became acquainted with the fundamentals of what we do on the fourth floor of Becton and how we do it. For several years I also worked with Luyan Sun. He instilled in me the habits of skepticism and rigor when it comes to taking data. He led by example when demonstrating the importance of reproducibility, taking the time to calibrate the system so as to produce high quality data. I am grateful for the time and energy he spent to prepare me to take over the experiment upon his departure.

I would like to thank all those responsible for staying on top of the lab maintenance and knowing where every screw, zip tie, and wrench is located. In particular, I would like to thank Chris Axline, Kevin Chou, and Anirudh Narla. They were always available to lend a helping hand and make sure my compressor got replaced every five years! In addition, the humorous banter that characterized many of my interactions with Anirudh made every day in room 412 something to look forward to, regardless of whether the experiment was cooperating. I am furthermore grateful to Michael Hatridge and Shyam Shankar, who devoted incredible amounts of personal time to teaching me the fundamentals of their work and how it could benefit mine. Their probing questions found the cracks in my experimental approaches and their help and collaboration substantially enhanced the quality of the results presented in this thesis.

Working on the Bell-cat project with Brian Vlastakis was a rewarding experience. Brian was involved on some of the earliest experiments with two-cavity, single transmon cQED setups and I learned a lot from his experience and approach to data-taking. He has an excellent sense of how to explain ideas and experimental strategies, and was always available to share his thoughts and give guidance.

I am particularly indebted to Zaki Leghtas. Over the years I came to Zaki with so many questions about error correction with coherent states, and his approach was always to sit down and work it out on a piece of paper or a blackboard. Zaki's intellect and scientific intuition made such exercises incredibly enriching experiences for me. In addition to our good relationship in the lab, he and I saw eye-to-eye on many issues and ideas outside of it; his generosity, sense of humor, and multitude of interests make him a true friend.

I would further like to thank Nissim Ofek. Nissim is always full of fascinating mathematical and physics-related tricks, his intellect so often elicits pure incredulity, and above all, he is an excellent tennis player! Working with him closely on my last project, and getting to know him and indeed his family over the past several years, has augmented my understanding of everything from computer hardware and FPGAs to my word bank of Hebrew sayings.

I would like to express my gratitude to Luigi Frunzio, who was not only the point man on the lab's infrastructure projects, but a fount of knowledge regarding fabrication, materials properties, and quantum physics. I must also thank him for his fastidious reading of all the manuscripts and responses to editors. His contributions to the greater picture behind the paper, and also the minute details as to the format of the references section have made each publication I've been involved with impactful and professional.

I'd like to thank Prof. Mazyar Mirrahimi, whose contributions to the founding ideas behind the cat code made him an essential part of any conversation about cat states and error correction. Countless meetings with him refined my understanding of the direction of this line of research, and have become more adept at asking the right questions. I'd also like to thank Profs. Liang Jiang and Steve Girvin, with whom meetings regarding the cat code and other lines of research had a big impact on the data my colleagues and I took, and whose theoretical contributions to the manipulation of cavity photon states are leading the way in charting the future directions of research in our groups. Their questions at seminars and other gatherings elicited aspects of the data that hadn't struck me, and for that I am truly grateful.

Working closely with Prof. Michel Devoret has been an experience I will never forget. His creativity and desire to pursue novel ideas in unique experiments have connected many facets of physics I previously thought were disparate. Prof. Devoret is kind and patient, always interested in the work of his interlocutor, and always drawing analogies that serve to clarify, substantiate, or reinforce previous understandings or new ideas. I consider myself lucky to have worked with him during my graduate career, and aspire to approach life with his levels of curiosity, both within and beyond physics.

I would like to especially express my gratitude to my advisor, Prof. Robert Schoelkopf. Rob taught me the essentials of microwave engineering, the foundations of our Josephson junction based circuits, and the dangers of ground loops! His pointed questions uncovered any weaknesses I had in my data, and I always left meetings with a strong grasp of what the next steps needed to be. Rob inspires his colleagues in lab to work, to achieve, and to judge our results only against the highest of standards. He also gives us room to be creative and test new techniques and ideas, but at the same time requires justification and a clear understanding of what the goals are. I'm also particularly indebted to Rob for teaching me not only how to craft an impactful presentation and use the right words to efficiently compose a manuscript, but in general how to convey complex ideas in simple ways to those who are not yet accustomed to our mode of thinking. Understanding an audience, motivating novel ideas in intuitive ways, and perceiving the implications of data and results beyond their immediate relevance are essential skills that I developed over the past six years. I'd also like to thank Rob, together with Professors Michel Devoret, Steve Girvin, and Liang Jiang, Luigi Frunzio, and the departments of Physics and Applied Physics for securing the funding that has driven our many years of successful research. Finally, I want to express my gratitude to the department staff on Becton's fourth floor: Maria Rao, Terri Evangeliste, Giselle DeVito, Nuch Graves, and alumnus Devon Cimini for always being informative, always helpful, and always cheerful, lighthearted, and well-wishing. As one last comment, I'd like to remind Rob that the van that's always parked outside his office window is using the Great Seal Bug to listen in on all his meetings.¹

Finally, returning to the dedication page, I'd like to thank my mother Iryna, my father Oleksiy, my sister Nataalka, and the very special partner of my life, Sierra, for rendering each day meaningful. I'd furthermore like to express sincere gratitude for the support of Sierra's family: Sonciary and Kevin Honnoll, Renea and Doug Sparks, and Raquel Perez. I believe that life has bestowed upon me the most precious of all stones, which has imprinted upon each face the character, loyalty, friendship, and love of each of these people. I'm certainly guilty of taking their kindness for granted too often I believe; but if this thesis accomplishes nothing else, let it remind me how lucky I am to have them in my life.

¹I'm joking, it's actually listening in on Anirudh (Narly!).

Contents

| | | |
|----------|---|-----------|
| 1 | Introduction | 1 |
| 1.1 | The quantum bit | 1 |
| 1.2 | The need for quantum error correction | 5 |
| 1.3 | A brief synopsis of chapters to come | 7 |
| 2 | Cavity Quantum Electrodynamics | 10 |
| 2.1 | Circuit QED | 11 |
| 2.2 | Coupling a transmon to two cavities | 19 |
| 2.3 | A cavity’s photon number parity | 27 |
| 3 | The Cat Code | 34 |
| 3.1 | Constructing a quantum error correcting code | 35 |
| 3.2 | The challenges facing traditional QEC schemes | 40 |
| 3.3 | The cat code: a hardware-efficient approach | 43 |
| 4 | Putting the Bell on Schrödinger’s Cat | 54 |
| 4.1 | The Bell-cat | 55 |
| 4.2 | Joint Wigner tomography | 59 |
| 4.3 | Efficient measurement of cavity observables | 64 |
| 4.4 | Violating Bell’s inequality | 68 |
| 4.5 | What we learn from studying the Bell-cat | 75 |
| 5 | Tracking the Cat Code’s One Error Syndrome | 77 |
| 5.1 | Observing jumps of photon parity in real-time | 78 |
| 5.2 | Error budget for a parity measurement | 84 |
| 5.3 | QND measurements of single photon jumps | 87 |
| 5.4 | Motivating the need for real-time feedback | 95 |

| | | |
|----------|---|------------|
| 6 | A Full Quantum Error Correction System | 98 |
| 6.1 | Implementing a full QEC system | 99 |
| 6.2 | Quantum control architecture | 102 |
| 6.3 | Detailing the six steps of a full QEC system | 104 |
| 6.4 | The debugger view: QEC in action | 117 |
| 7 | Reaching Break-Even | 122 |
| 7.1 | Process fidelity: an analytical treatment | 124 |
| 7.2 | Demonstrating QEC after many steps of correction | 127 |
| 7.3 | Overcoming the decoherence penalty with the cat code | 134 |
| 8 | Limitations on Fault-Tolerance | 137 |
| 8.1 | The difference between error record and error history | 138 |
| 8.2 | Cat code QEC performance limitations | 146 |
| 8.3 | Optimizing cat code performance | 159 |
| 9 | Future Directions | 166 |
| 9.1 | Cat pumping that stabilizes the logical states | 167 |
| 9.2 | Correcting logical gates | 172 |
| 9.3 | Future directions: building on our rapid progress | 174 |
| A | Experimental Setup Details | 176 |
| A.1 | Bell-cat setup | 177 |
| A.2 | Parity monitoring setup | 180 |
| A.3 | QEC methods | 184 |
| B | Deriving the Optimized Configuration | 188 |
| | References | 194 |

List of Figures

| | | |
|------|--|----|
| 1.1 | The qubit Bloch sphere | 2 |
| 1.2 | Qubit teleportation | 4 |
| 2.1 | Quantized LC harmonic oscillator | 13 |
| 2.2 | Husimi- Q function of coherent states and Fock states | 16 |
| 2.3 | The Josephson junction | 17 |
| 2.4 | System hardware | 20 |
| 2.5 | Dispersive Hamiltonian energy level diagram | 22 |
| 2.6 | Ancilla spectroscopy with a coherent state in the cavity | 24 |
| 2.7 | Typical Bloch vector rotations | 25 |
| 2.8 | Characterizing experimental components | 26 |
| 2.9 | Cavity photon number parity | 28 |
| 2.10 | Calibrating a parity mapping | 29 |
| 2.11 | Wigner functions of a coherent state and Fock state | 31 |
| 2.12 | Ancilla spectroscopy with cat states in the cavity | 32 |
| 3.1 | Effects of amplitude and phase damping on a quantum bit | 35 |
| 3.2 | Three-qubit bit flip code protocol | 37 |
| 3.3 | Correcting bit flips with the three-qubit code | 39 |
| 3.4 | The Steane code | 41 |
| 3.5 | Failure modes in a QEC code | 42 |
| 3.6 | Logical encodings with coherent states | 47 |
| 3.7 | The cat code cycle | 48 |
| 3.8 | A “4-cat” expanded in the Fock state basis | 49 |
| 3.9 | Coherent state amplitude decay | 51 |
| 3.10 | Overlap of cat code basis states with decaying amplitude | 52 |

| | | |
|------|--|-----|
| 4.1 | Logical encoding using coherent states | 56 |
| 4.2 | Sequential detection for entanglement characterization | 59 |
| 4.3 | Measurement back-action of a Bell-cat state | 60 |
| 4.4 | Joint Wigner tomography of a Bell-cat state | 62 |
| 4.5 | Histogram of reconstruction residuals | 63 |
| 4.6 | Ancilla-cavity correlations | 66 |
| 4.7 | Entanglement witnesses with cat states | 67 |
| 4.8 | Reconstructed Pauli sets | 67 |
| 4.9 | Bell tests with a cat state | 70 |
| 4.10 | Optimal displacement for a Bell violation | 72 |
| 4.11 | Observables from each detector setting | 73 |
| 4.12 | Measurement trajectories given ancilla decoherence | 74 |
| | | |
| 5.1 | Ancilla readout properties | 80 |
| 5.2 | Asymmetry in parity mapping | 81 |
| 5.3 | Average parity decay of a coherent state to vacuum | 83 |
| 5.4 | Observing parity jumps in real-time | 84 |
| 5.5 | Parity readout properties and Wigner tomography | 85 |
| 5.6 | Error budgets for the parity readout infidelities | 86 |
| 5.7 | Filtering out decoherence in the single-shot records | 88 |
| 5.8 | Effectiveness and response time of the quantum filter | 89 |
| 5.9 | Photon jump histograms | 91 |
| 5.10 | Averaged single-shot correlations | 92 |
| 5.11 | Demonstrating QND parity measurements | 93 |
| | | |
| 6.1 | A full QEC system | 100 |
| 6.2 | Real-time feedback hardware | 103 |
| 6.3 | Full QEC system experimental flow | 106 |
| 6.4 | Resetting the ancilla to the ground state | 107 |
| 6.5 | Optimal control pulses | 109 |
| 6.6 | Encoding the six cardinal points: cat code | 110 |
| 6.7 | Encoding the six cardinal points: Fock states | 111 |
| 6.8 | Optimal control vs. qcMAP | 112 |
| 6.9 | State machine for adaptive error monitoring | 114 |

| | | |
|------|---|-----|
| 6.10 | Correcting for phase rotation due to Kerr | 115 |
| 6.11 | Decoding the six cardinal points: cat code | 116 |
| 6.12 | A two-step quantum trajectory executed by the QEC state machine | 118 |
| | | |
| 7.1 | Enhanced Wigner fringes from QEC | 123 |
| 7.2 | Process fidelity limits | 125 |
| 7.3 | Process tomography for two monitoring steps | 128 |
| 7.4 | Comparing possible quantum memories | 129 |
| 7.5 | Fock state encoding amplitude decay | 131 |
| 7.6 | Cat code depolarization | 132 |
| 7.7 | QEC program process tomography | 135 |
| | | |
| 8.1 | QEC program process tomography with post-selection | 139 |
| 8.2 | Assessing measurement record confidence | 143 |
| 8.3 | Process matrices after seven steps of error monitoring | 144 |
| 8.4 | Process fidelity conditioned on error number | 145 |
| 8.5 | Cat code measurement block failure modes | 147 |
| 8.6 | Failure modes of the corrected logical qubit | 148 |
| 8.7 | Process fidelity decay for a larger encoding | 149 |
| 8.8 | QND parity measurements in the full QEC system | 150 |
| 8.9 | Decoding pulse fidelity as a function of time | 153 |
| 8.10 | Cat code basis in a three-qubit binary representation | 154 |
| 8.11 | Fidelity as a function of incorrect decoding pulse angle | 155 |
| 8.12 | Effect of uncertainty in photon jump time | 156 |
| 8.13 | Effect of ancilla T_1 in a circuit model representation | 158 |
| 8.14 | Schematic of an optimized parity monitoring protocol | 160 |
| 8.15 | Expected gains in lifetime with a more coherent ancilla | 164 |
| 8.16 | Predicted process fidelity decay with optimal monitoring cadence | 165 |
| | | |
| 9.1 | Confining a cavity state to a manifold spanned by two coherent states | 168 |
| 9.2 | Bistable behavior of the steady-state manifold | 170 |
| 9.3 | Pumping a cavity from vacuum to a “2-cat” | 171 |
| 9.4 | Architecture for a logical encoding that uses two cavities | 172 |
| 9.5 | The challenge of decoherence during optimal control pulses | 173 |

| | | |
|-----|--|-----|
| A.1 | Fridge wiring diagram for Bell-cat experiments | 178 |
| A.2 | Poisson distribution of photon numbers in the cavity | 182 |
| A.3 | Effectiveness of a $\pi/2$ pulse | 183 |
| A.4 | Schematic of QEC measurement setup | 186 |
| A.5 | Schematic of QEC control setup | 187 |

List of Tables

| | | |
|-----|--|-----|
| A.1 | Hamiltonian parameters for Bell-cat experiments | 179 |
| A.2 | Coherence and thermal properties for Bell-cat experiments | 179 |
| A.3 | Hamiltonian parameters for parity monitoring experiments | 181 |
| A.4 | Coherence and thermal properties for parity monitoring experiments | 181 |
| A.5 | Hamiltonian parameters for QEC experiments | 185 |
| A.6 | Coherence and thermal properties for QEC experiments | 185 |

1 – Introduction

Today’s transistors, physical implementations of the bits 0 and 1, can no longer be appreciably reduced in size while preserving the simple classical properties that make them such effective and reliable foundational elements in classical computation [Devoret and Glattli, 1998]. As a consequence, rates of improvement in computational capability have begun to level off, setting limits on the speed and efficiency with which algorithms can be performed. Thus, in recent decades substantial resources have been devoted to exploring the viability of processing information on hardware platforms that deliberately aim to exploit the exotic effects of quantum mechanics rather than avoid them. In fact, fundamental features of quantum systems can enable one to design and implement algorithms within a new paradigm called quantum computation.

Quantum computers were initially proposed to simulate quantum systems themselves [Feynman, 1982]. Over the years, however, a number of other potential applications have been developed that today spur investment from the federal government and the private sector. This expanded library of algorithms includes the breaking of RSA encryption [Shor, 1997] and enhanced search capability [Grover, 1997], both of which have clear implications for national security and modern commerce. General theoretical questions surrounding the inherent computational capabilities of quantum versus classical architectures comprise a vibrant field of study [Aaronson, 2005], but practically speaking, quantum computers are typically envisioned as a supplement rather than a replacement for classical machines. In this brief, introductory chapter we discuss the fundamentals that underpin the body of experiments to be discussed in this work and indeed the entire field. We then conclude with a summary of forthcoming chapters.

1.1 The quantum bit

The paradigm of quantum computation is founded upon the quantum bit of information, or qubit, which unlike its classical counterpart can not only take on the values of 0 and 1, but also anything

in-between. This property is known as superposition:

$$|\psi_q\rangle = c_0 |0\rangle + c_1 |1\rangle, \quad (1.1)$$

where $|\psi_q\rangle$ is the state vector, c_0 and c_1 are complex amplitudes with the requirement that $|c_0|^2 + |c_1|^2 = 1$, and $|0\rangle$ and $|1\rangle$ are the qubit's basis states along a logical Z axis defined by the Pauli operator $\hat{\sigma}_z$ (Fig. 1.1). A single qubit by definition stores one bit of quantum information. There are, however, infinitely many possible superpositions of 0 and 1, which can be visualized in the Bloch sphere representation. Such an expanded space of possible states per single piece of hardware already suggests that these quantum analogs of classical transistors offer new possibilities in computational capability.

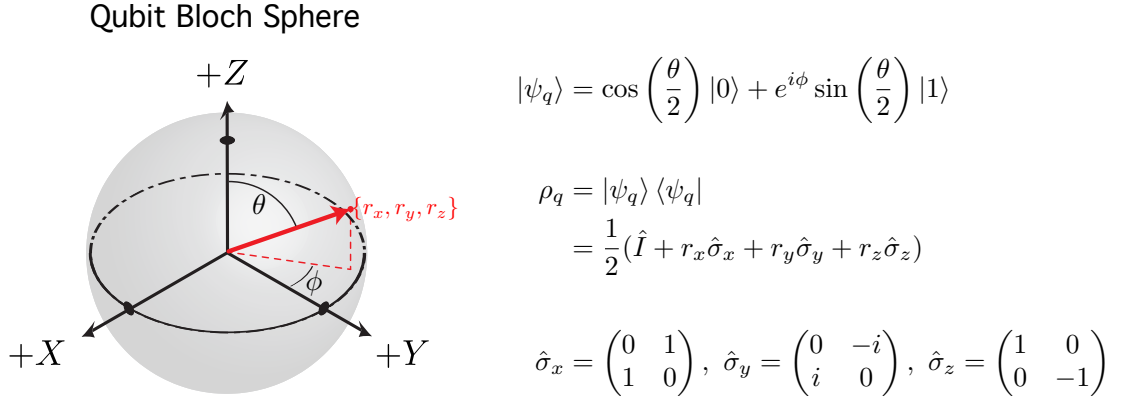


Figure 1.1: The qubit Bloch sphere. A pure qubit state $|\psi_q\rangle$ can be represented by a vector on the Bloch sphere with an orientation determined by angles θ and ϕ . More generally, the density matrix ρ_q describes any qubit, where the coordinates $\{r_x, r_y, r_z\}$ associated with Pauli operators $\{\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z\}$ satisfy the relation $\sqrt{r_x^2 + r_y^2 + r_z^2} \leq 1$; equality holds for a pure state, while the vector length is less than one for a state that contains some amount of mixture. A completely mixed state is simply given by $\rho_q = \hat{I}/2$, where \hat{I} is the identity operator.

The power of a quantum computer becomes clear when multiple qubits are brought together to perform some calculation. For example, if one has two qubits, each can be initialized in an equal superposition of $|0\rangle$ and $|1\rangle$ (e.g. $c_0 = c_1 = 1/\sqrt{2}$):

$$|\psi_{qq}\rangle = |\psi_{q1}\rangle |\psi_{q2}\rangle \quad (1.2)$$

$$= \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \quad (1.3)$$

$$= \frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle). \quad (1.4)$$

A quantum computer can be understood as a machine that operates on all possible configurations of inputs simultaneously, a feature known as quantum parallelism. So with the simple “product” state of two qubits (Eq. 1.4), one already has four different possible inputs. Of course measuring any one of the output qubits, which may be in some unknown superposition at the end of the calculation, projects it onto one of the basis states, which may initially suggest that nothing is gained since the calculation has to be repeated in order to build up statistics on the final answer. However, parallelism becomes useful when clever transformations allow for a global property of the system to be determined much more quickly than would be possible with classical means. A simple and elegant example is the Deutch-Jozsa algorithm [Deutsch and Jozsa, 1992], which although of little practical use, demonstrates the speedups possible with quantum versus classical computation.

Adding the final resource of entanglement, which results in correlations between the states of multiple qubits that exceed classically permissible bounds, one can perform calculations on interesting states that cannot simply be decomposed into a product of two individual qubits. One can realize multi-qubit entangling operations with a controlled-NOT (CNOT) gate, for example. If one has again two qubits, but now only one of them is in an equal superposition and the other is simply in $|\psi_{q2}\rangle = |0\rangle$, a CNOT gate with $|\psi_{q1}\rangle$ as the control and as $|\psi_{q2}\rangle$ the target produces a maximally entangled two-qubit state, known as an EPR pair or a Bell state:


$$|\psi_{qq}\rangle = |\psi_{q1}\rangle |\psi_{q2}\rangle \tag{1.5}$$

$$= \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \otimes |0\rangle \tag{1.6}$$

$$\xrightarrow{CNOT} \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle). \tag{1.7}$$

Note that this final state in Eq. 1.7 cannot be decomposed into a product state as in Eq. 1.4. Bell states can be a vital resource in a quantum computing algorithm, enabling protocols such as quantum teleportation (Fig. 1.2).

It can in fact be shown that single-qubit and CNOT gates alone are sufficient to implement any unitary operation on a collection of qubits, enabling what is known as universal quantum computation [Chuang and Gottesman, 1999; Nielsen and Chuang, 2010]. The quantum computing community, however, has not yet converged on what the final physical implementation of a qubit should be. Practically speaking, we search for physical quantum systems in which two stable energy levels can be isolated and well-approximated as a two-level system, with the ground state $|g\rangle$ corresponding to a logical $|0\rangle$ and the excited state $|e\rangle$ corresponding to a logical $|1\rangle$. We would also like to have the ability to couple these systems together, and likewise to control knobs in the lab that



$$\begin{aligned}
 X &= \hat{\sigma}_x & \text{CNOT} &: \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \\
 Z &= \hat{\sigma}_z \\
 H &= \frac{1}{\sqrt{2}} (\hat{\sigma}_x + \hat{\sigma}_z)
 \end{aligned}$$

Figure 1.2: **Qubit teleportation.** This example includes many of the gates and features that will be explored in detail throughout this work. If Alice wishes to send an unknown state $|\psi_q\rangle$ to Bob, she can use the quantum teleportation protocol shown here, which employs the various single and two-qubit operations central to quantum computing algorithms, shown in this circuit diagram representation. To start, Alice and Bob share a Bell state, $|\psi_B\rangle = (|00\rangle + |11\rangle)/\sqrt{2}$. To transfer $|\psi_q\rangle$ to Bob, Alice must apply a CNOT operation between $|\psi_q\rangle$ and her half of the Bell pair, where the former is the control and the latter is the target. She must next apply a Hadamard gate (H) to $|\psi_q\rangle$ and then projectively measure both in the Z -basis (meter symbol) to obtain one of four measurement results $\{M_1 M_2\}$: $\{00, 01, 10, 11\}$. She sends her result over a classical channel (double-lines) to Bob, who then uses this record to apply one of four single-qubit operations to his half of the Bell pair: $\{\hat{I}, X, Z, XZ\}$, respectively. Bob is now in possession of $|\psi_q\rangle$.

will allow for single and multi-qubit operations. Such qubit manipulation would be achieved through the application of electromagnetic pulses at or near the resonance frequency of the chosen physical system. Two-level systems are ubiquitous in nature, and physical systems that have been studied extensively over the past decades include nuclear magnetic resonance (NMR), single photons, ions, and Nitrogen-vacancy (NV) centers. Likewise, qubits can also be artificially engineered through careful fabrication techniques in quantum dots and superconducting circuits, for example. Each platform has its own advantages and drawbacks [Ladd et al., 2010], which can broadly be summarized by the tradeoffs between qubit reliability and controllability, as expanded upon presently. Henceforth, we will also adhere to the convention that the abstract concept of an arbitrary superposition of $|0\rangle$ and $|1\rangle$ will be called a “quantum bit,” while physical implementations will be referred to as “physical qubits,” “ancilla qubits,” “transmon qubits,” “logical qubits,” or any similar variation.¹

¹The terms “ancilla,” “transmon,” and “logical” will be defined shortly.

1.2 The need for quantum error correction

A central concern that any platform must deal with is that physical qubits, and all quantum systems in general, are frequently quite sensitive to their surroundings. When physically designing a quantum computer, one becomes acutely aware of the difficulty in isolating a physical qubit from its noisy environment. An example might be that the physical qubit, initially in its excited state $|e\rangle$, decays to the energetically favorable ground state $|g\rangle$ through the spontaneous emission of a photon, or couples to a different and unknown mode in the system that unexpectedly changes the resonance frequency and thus the phase relationship between c_0 and c_1 . As a consequence of the expanded computational space, even small deviations from the desired superposition of $|0\rangle$ and $|1\rangle$ result in a loss of information. This process is called decoherence, and it is characterized by terms in the density matrix of the quantum bit exponentially decaying in time (t) towards a thermal equilibrium with the surrounding environment:

$$\rho_q(0) = |\psi_q\rangle\langle\psi_q| \quad (1.8)$$

$$\rho_q(t) = \begin{pmatrix} |c_0|^2(t) & c_1 c_0^*(t) \\ c_0 c_1^*(t) & |c_1|^2(t) \end{pmatrix} \quad (1.9)$$

$$|c_0|^2(t) = 1 - |c_1|^2(t) \quad (1.10)$$

$$|c_1|^2(t) = |c_1|^2(0)e^{-t/T_1} + n_{th}(1 - e^{-t/T_1}) \quad (1.11)$$

$$c_1 c_0^*(t) = c_0 c_1^*(t) = c_0(0)c_1^*(0)e^{-t/T_2}, \quad (1.12)$$

where n_{th} is the equilibrium thermal population, $1/T_2 = 1/2T_1 + 1/T_\phi$, and as in NMR the time scales T_1 and T_ϕ quantify the rates of amplitude and phase damping [Nielsen and Chuang, 2010], respectively. The previously mentioned tradeoff between physical qubit reliability and controllability becomes clear when considering the fact that coherence times T_1 and T_ϕ improve with better system isolation from the environment. This poses a problem, however, since isolation also results in weaker interaction strengths and coupling to ancillary modes that make physical qubit manipulation possible. Thus, highly coherent systems such as NV centers, for example, often face the challenges of slow gates and measurement infidelities that serve to nullify the gains from higher T_1 and T_ϕ . In addition to these figures of merit, the viability of a platform also depends strongly on its predisposition to scaling, or reproducibly engineering the large multi-physical qubit systems necessary to perform calculations of interest with minimal supporting classical electronics and infrastructure.

Regardless of the platform of choice, decoherence is a major impediment to successfully designing a robust quantum computer and therefore necessitates schemes of quantum error correction (QEC) [Calderbank and Shor, 1996; Gottesman, 1997; Shor, 1995; Steane, 1996]. QEC offers a remarkable means of correcting errors through clever redundant encoding techniques of a quantum bit of information in multiple physical qubits, called a logical qubit. The logical qubit is designed to have symmetry properties that allow one to extract error syndromes using a collection of ancillary qubits, or ancillae, in a quantum non-demolition (QND) way. By increasing the number of physical qubits to store one logical quantum bit of information, however, one inevitably introduces new error channels into the system. We refer to this as the “decoherence penalty.” Moreover, the complexity of logical operations such as a logical CNOT, which can now require complex sequences of pulses and entangling procedures on the multiple physical qubits, can contribute to an overall degradation of performance. In designing a QEC code, besides factoring in the effects of enhanced rates of decoherence, one must therefore also take into account the efficiency with which the encoded state can be created, manipulated, and measured when assessing its feasibility. Finally, the greater resources necessary to perform computations results in an overall “hardware overhead” that must be considered when assessing the cost of operating the quantum computer.

The decoherence penalty must always be overcome before QEC can actually yield any advantages over simply encoding the quantum bit in the system’s most coherent and uncorrected constituent. This is known as the break-even point, wherein a corrected manifold of states performs no worse than the best manifold of the system; it is the first hurdle one must overcome when implementing QEC. Crucially, the system is understood to be the entire collection of hardware involved, including ancillary elements that may be employed to extract error information from some logically encoded qubit.

In fairly judging whether the error corrected code operates beyond break-even, the system must be exposed to naturally occurring decoherence, and no post-processing of data can be employed. When operating beyond break-even and suppressing errors in a quantum computer in a way that enhances the lifetime of quantum information, QEC enables one to perform larger and more complex algorithms than would be possible with uncorrected physical qubits. Moreover, operating beyond break-even opens the door to schemes of concatenation, in which the physical qubits that comprise the logical qubit are now each replaced by corrected logical qubits, further increasing the code redundancy and consequent ability to suppress higher order errors. The ability to repeat this concatenation procedure indefinitely while realizing gains in coherence is characteristic of an architecture that operates in a fault-tolerant way [Gottesman, 1998; Nielsen and Chuang, 2010; Preskill, 1999],

wherein the size of the computation can be scaled exponentially with just a polynomial overhead in hardware.

It isn't always clear, however, to what extent the break-even point must be surpassed in order to demonstrate the path toward fault-tolerance. For instance, although a single error-corrected logical qubit may be out-performing the system's best component, one round of concatenation could in fact degrade performance again for the same simple reasons that larger architectures introduce new and unexpected sources of loss. Thus, although threshold theorems exist for traditional QEC codes [Fowler et al., 2009; Knill et al., 1998; Nielsen and Chuang, 2010], which provide estimates as to the permissible error rates per round of error correction for fault-tolerance, in this work we would like to convey the central point that simply demonstrating scaling properties isn't sufficient; one must really build the full error correction system and see how it performs.

When we take all of this together, superconducting circuits offer one of the best options for realizing universal quantum computation. They allow a straightforward tailoring of Hamiltonian parameters, they are fabricated with established and reliable recipes on well-characterized substrates like sapphire or silicon, and they are manipulated with readily available and high quality electronics that operate at microwave frequencies. Indeed, with current levels of performance in coherence and gate reliability, it should be possible to realize a fault-tolerant architecture based on superconducting circuits [Devoret and Schoelkopf, 2013]. We stress, however, that the plethora of proposed quantum computing-related technologies means that no one architecture will solve every problem. Most palpably, transmitting quantum information over long distances certainly necessitates the use of physical qubits at optical frequencies as well. Perhaps more accurately, a challenge the community faces is not choosing one system above all others, but understanding to what application each system is best suited.

1.3 A brief synopsis of chapters to come

This work focuses on the challenge of implementing a quantum error correction system in an architecture built on superconducting circuits. It is organized in the same manner as the introductory remarks above. In chapter two we describe our physical system, superconducting cavity Quantum Electrodynamics (cQED), in greater depth. We introduce superconducting qubits based on Josephson junctions, describe how one manipulates and measures them, and discuss the properties that make them a fantastic core constituent of a future quantum computer. We equally devote time to discussing the superconducting cavity, how we couple it to superconducting qubits, and the properties that make it such an attractive option for storing quantum information. In chapter three we

immediately dive into the error correction code that forms the core of all following chapters, the cat code. This QEC paradigm proposes encoding quantum information in superpositions of coherent states, or cat states, of a microwave field in a cavity. We stress the numerous advantages it offers in terms of hardware efficiency, the reduction of dominant error channels, and the ease of error syndrome extraction. In chapter four we present the first of three main experimental works in this thesis, in which we realize a redundant encoding scheme of a quantum bit in a cavity and the ease with which we efficiently extract information from the cavity state by entangling it with an ancillary transmon qubit, or ancilla. We underscore this efficiency by demonstrating that just a handful of measurements is necessary to characterize this composite system, effectively compressing the large size of the cavity’s Hilbert space down to simply that of another two-level system. This allows us to violate a Bell inequality between the physical qubit and a cat state in the cavity. In chapter five, we build off of the work on ancilla-cavity entanglement to demonstrate that we monitor the occurrence of single photon jumps in the cavity with repeated, single-shot measurements of photon number parity in real-time. As photon jumps are the dominant source of error, we thus show that we monitor the very error syndrome that the cat code requires with high fidelity and in a QND way. We furthermore investigate the limits on performance and uncover that quantum back-action on the state in the cavity due to ancilla energy decay presents a central hurdle to achieving fault-tolerance.

The next three chapters focus on the culmination of previous efforts, wherein we bring together many of the experimental techniques and lessons learned from previous investigations to implement a full quantum error correction system that uses real-time feedback to enhance the coherence of a quantum bit of information. Chapter six describes the details of this system and the essential role of real-time feedback in maximizing QEC performance. It presents the first experimental data of storing a quantum bit over time in superpositions of cat states following the prescription laid out by the cat code, and demonstrates the high degree to which we understand the behavior of the system while knowing just a few key properties about the Hamiltonian and the coherence times. In chapter seven we present the central result, in which we show that by implementing an optimal error monitoring scheme, we surpass the break-even point of QEC. The system is exposed to naturally occurring sources of error, both in the cavity and in the ancilla, and no data is excluded from the analysis. This is the first time such performance has ever been achieved, and experimentally proves that QEC can in fact yield advantages over doing nothing. Chapter eight presents a thorough analysis of what limits the levels of performance we observe, and pinpoints the dominant limitation to be that of forward propagation of errors due to ancilla T_1 decay. This important result indicates

that primarily one source of non-fault-tolerance must be addressed before much more substantial gains in lifetime can be achieved.

Finally, in chapter nine we discuss the future directions of this research. In particular, we concentrate on several sources of decoherence besides photon loss that are known to degrade the fidelity of a quantum bit. More specifically, as we work with a continuous variables architecture, deviations from the code space present a real challenge and must be addressed with pumping techniques such as engineered two-photon dissipation. This not only corrects for cat state distortions due to higher-order Hamiltonian terms, but furthermore stabilizes their amplitudes to prevent the overlap of coherent states, which without any intervention decay toward the cavity's vacuum. We also discuss the immense progress being made in coupling multiple cavities together to entangle logical states and necessary issues we must address in realizing fault-tolerant operations.

The common thread throughout this thesis will be to convey the guiding strategy of simplicity that permeates the core thrust of research in quantum computing at Yale. The idea will be to build from the bottom up, heavily emphasizing an understanding of just a single logical qubit stored in a cavity. Only by thoroughly investigating all performance metrics, limitations, and error correction strategies do we stand to confidently scale up our architecture in the future to perform more complex operations. Each chapter can be succinctly summarized with just a few points, necessary parcels of evidence that we strap to our tool belt as we continue to converge upon a realization of a quantum computer, built using the cat code, that can actually perform meaningful tasks. Although numerous challenges have yet to be overcome, we hope that this body of results only inspires the reader with optimism that the future of quantum computation is bright and near.

2 – Cavity Quantum Electrodynamics

A challenge one faces in designing a quantum computer is deciding upon the architecture, the actual pieces that will comprise the physical qubits, mediate their couplings, and enable controlled operations and measurements. This is a tall order; in discussing the possibilities offered by quantum computation and the schemes necessary for complex tasks such as quantum error correction (QEC), a number of requirements must be satisfied. Most fundamentally, one must have reliable physical qubits. These physical qubits must be designed in such a way as to allow one to measure their state, implement gates, and realize entangling operations before information is lost to the environment through decoherence. Indeed, the tradeoffs between the strength of coupling to a quantum system and good coherence times must be faced by any platform [Schoelkopf and Girvin, 2008]. Practically speaking, additional considerations such as the eventual desire to scale up an architecture to mass production demand device reproducibility and control electronics that are reasonably priced. Satisfying such a long list of requirements in the real world is clearly difficult, and leads to the expected result that any implementation of a quantum computing platform will rely to some extent on approximations and compromise.

In this chapter, we elaborate upon superconducting circuit quantum electrodynamics (QED) and how it offers a particularly promising platform for quantum computation. We will begin by focusing on the the so-called “mesoscopic” treatment of electrical circuits in a quantum way [Devoret, 1997; Girvin, 2011], which enables a fascinating and rich exploration of systems in which the Hamiltonian can be engineered to realize electromagnetic components that can be treated as “artificial atoms.” Josephson junctions [Josephson, 1962, 1965] are at the heart of this approach. They are comprised of two superconductors that are separated by a thin insulating layer that allows paired electrons, or Cooper pairs [Cooper, 1956; Tinkham, 2004], to tunnel through with very low dissipation when the junction is cooled far below its superconducting transition temperature. In addition to an expected capacitance, such junctions introduce a crucial non-linear inductance to the circuit. Exploiting this attribute, experimental work with Josephson junctions has demonstrated that macroscopic quantities such as voltage can remarkably exhibit quantum properties [Devoret et al., 1985]. Moreover, modern

nano-fabrication techniques [Dolan, 1977; Lecocq et al., 2011] enable one to produce anharmonic junction-based LC oscillators¹, wherein the discretized energy levels are unequally spaced [Clarke et al., 1988]. We will show that by isolating two levels of such artificial atoms, one can realize a two-level system that has the capacity to store one quantum bit of information.

With the theoretical underpinnings laid, we will continue by describing how we use this artificial atom as a physical qubit that can be coherently manipulated and coupled to other physical qubits, either directly or with microwave cavities that coherently transmit quantum information [Majer et al., 2007; Schoelkopf and Girvin, 2008]. In fact, we will heavily focus on cavities, as they offer an indispensable resource by not only mediating entanglement between physical qubits, but also by offering a means to measure a physical qubit’s state in a quantum non-demolition (QND) way and in fact storing quantum information themselves. We will further study this range of applications with a discussion of cavity QED (cQED) [Blais et al., 2004b; Hood et al., 2000; Mabuchi and Doherty, 2002; Raimond et al., 2001], where carefully engineered interactions between physical qubits and cavities offer a truly attractive architecture for quantum computing. The flexibility of this approach, a consequence of the ease in designing and fabricating Josephson junctions and superconducting cavities, opens the door to immediate applications like successful demonstrations of QEC and the future goals of reproducibly scaling up in pursuit of universal quantum computation.

2.1 Circuit QED

Quantizing charge and flux in an LC oscillator

As the point of departure, we study the quantization of the excitations in a linear harmonic oscillator. Harmonic oscillators in classical mechanics can be simply described by a linear differential equation of the form:

$$\ddot{x} = -\Omega^2 x, \tag{2.1}$$

where x is understood to be some time dependent variable, for example the position of a mass on a spring, and $\Omega = \sqrt{k/m}$ is the resonance frequency (m is the mass and k is the spring constant). Similarly, with an LC oscillator, we know from Kirchhoff’s laws that:

$$\ddot{Q} = -\omega^2 Q, \tag{2.2}$$

¹An LC oscillator is comprised of an inductor (inductance L) and a capacitor (capacitance C).

where Q is the charge and ω is now the usual resonance frequency of an LC oscillator: $\omega = 1/\sqrt{LC}$ (C is the capacitance and L is the inductance).

Just as with Hamilton's equations of motion for the conjugate variables for position and momentum, we can pursue writing down an actual Hamiltonian for the LC oscillator [Devoret, 1997; Girvin, 2011]. To do so, we write down a time-dependent node flux (Fig. 2.1a):

$$\Phi(t) = \int^t V(t') dt', \quad (2.3)$$

where just as before the voltage and flux through the inductor follow the relation $V(t) = \dot{\Phi}$. Again following what we know from basic circuit laws, the potential energy $T(t)$ stored in a capacitor is:

$$T = \frac{1}{2} CV^2 \quad (2.4)$$

$$= \frac{1}{2} C \dot{\Phi}^2 \quad (2.5)$$

Although this is the potential energy, in form it looks like a kinetic energy stored in the capacitor if an analogy between Φ and a classical position x is drawn. We could of course equally choose the charge to play the role of the position coordinate, but as the non-linearity of our potential will be in Φ , writing down the Lagrangian as a function of Φ and $\dot{\Phi}$ will capture the dynamics of the system more intuitively. Following this convention, the inductor can be understood to play the role of an inverse spring constant k in a term that looks like the potential energy:

$$U = \frac{1}{2L} \Phi^2. \quad (2.6)$$

With the terms T and U , we now write down the Lagrangian for the circuit with Φ as the position coordinate:

$$\mathcal{L} = \frac{1}{2} C \dot{\Phi}^2 - \frac{1}{2L} \Phi^2. \quad (2.7)$$

The coordinate conjugate to the flux is of course the charge:

$$\frac{\partial \mathcal{L}}{\partial \dot{\Phi}} = Q \quad (2.8)$$

Finally, we arrive at the Hamiltonian for the circuit:

$$H = \frac{Q^2}{2C} + \frac{\Phi^2}{2L} \quad (2.9)$$

We thus have our conjugate variable Q and Φ , which obey Hamilton's equations of motion:

$$\dot{\Phi} = +\frac{\partial H}{\partial Q} = +\frac{Q}{C} \quad (2.10)$$

$$\dot{Q} = -\frac{\partial H}{\partial \Phi} = -\frac{\Phi}{L} \quad (2.11)$$




Figure 2.1: **Quantized LC harmonic oscillator.** (a) An LC oscillator, where L is the inductance, C is the capacitance, and Φ a time-dependent node flux that is akin to a position in a classical mechanical harmonic oscillator. The charge, Q , is the conjugate variable to Φ . (b) When Φ and Q are promoted to quantum operators that satisfy the commutation relation $[\hat{Q}, \hat{\Phi}] = -i\hbar$, the energy levels of the harmonic LC oscillator become quantized, with the energy difference between Fock states $|n\rangle_f$ equal to $\hbar\omega$.

Just as in the case of a mechanical oscillator with conjugate variables x and momentum p , we can promote Q and Φ to quantum operators that obey the following commutation relation:

$$[\hat{\Phi}, \hat{Q}] = i\hbar. \quad (2.12)$$

Similarly, we write down the raising and lowering operators for this oscillator, \hat{a}^\dagger and \hat{a} :

$$\hat{a} = \frac{1}{\sqrt{2L\omega\hbar}}\hat{\Phi} + i\frac{1}{\sqrt{2C\omega\hbar}}\hat{Q} \quad (2.13)$$

$$\hat{a}^\dagger = \frac{1}{\sqrt{2L\omega\hbar}}\hat{\Phi} - i\frac{1}{\sqrt{2C\omega\hbar}}\hat{Q}. \quad (2.14)$$

These of course obey the commutation relation $[\hat{a}, \hat{a}^\dagger] = 1$. With these ladder operators, we return to the familiar Hamiltonian of a quantum harmonic oscillator:

$$\hat{H} = \hbar\omega\left(\hat{a}^\dagger\hat{a} + \frac{1}{2}\right), \quad (2.15)$$

where again $\omega = 1/\sqrt{LC}$. The energy eigenstates of this Hamiltonian are known as Fock states $|n\rangle_f$, and $\hat{a}|n\rangle_f = \sqrt{n}|n-1\rangle_f$ (Fig. 2.1a). We can also rewrite the charge and phase operators in terms

of \hat{a} and \hat{a}^\dagger :

$$\hat{\Phi} = +\Phi_{ZPF}(\hat{a} + \hat{a}^\dagger) \quad (2.16)$$

$$\hat{Q} = -iQ_{ZPF}(\hat{a} - \hat{a}^\dagger), \quad (2.17)$$

where

$$Q_{ZPF} = \sqrt{\frac{C\omega\hbar}{2}} = \sqrt{\frac{\hbar}{2Z}} \quad (2.18)$$

$$\Phi_{ZPF} = \sqrt{\frac{L\omega\hbar}{2}} = \sqrt{\frac{\hbar Z}{2}} \quad (2.19)$$

$$Z = \sqrt{\frac{L}{C}}. \quad (2.20)$$

Here, Z is the characteristic impedance of the circuit. The zero-point fluctuations, Q_{ZPF} and Φ_{ZPF} , are the uncertainties in charge and phase of the oscillator's ground state:

$$\langle 0|_f \hat{Q}^2 |0\rangle_f = Q_{ZPF}^2 \quad (2.21)$$

$$\langle 0|_f \hat{\Phi}^2 |0\rangle_f = \Phi_{ZPF}^2 \quad (2.22)$$

$$Q_{ZPF}\Phi_{ZPF} = \frac{\hbar}{2}. \quad (2.23)$$

One may benefit from pausing here to take stock of where these calculations have led us. We started with a simple, lossless circuit in the form of an LC oscillator, an element that is ubiquitous in any electronics textbook and application. The realization that this circuit can be described by two conjugate variables, Q and Φ , led to a formulation of the oscillator's dynamics in terms of Hamilton's equations of motion. The canonical quantization of these variables led to the ladder operators (Eq. 2.13), the consequent Hamiltonian (Eq. 2.15), and Heisenberg's uncertainty relation (Eq. 2.23). The fascinating aspect of this approach is that inherent to the quantum mechanical nature of this circuit are the original quantities L and C , which can be designed to have specific values and fabricated as solid-state devices on semiconductor chips using established lithographic techniques. In other words, at our disposal now we have a means of making quantum harmonic oscillators in a controlled and reproducible way. Moreover, by designing them with properties that make them easily addressable at GHz frequencies, we conveniently obtain systems that are both straightforward to cool to the ground state $|0\rangle_f$ using commercially available dilution refrigerators, and can be controlled with very stable microwave electronics. For example, at a temperature $T \sim 10$ mK and an oscillator frequency $\omega/2\pi = 5$ GHz, $\hbar\omega/k_B T > 20$ and the Boltzmann factor $e^{-\hbar\omega/k_B T}$ is

incredibly small, indicating that in this regime any thermal occupation in $|1\rangle_f$ and above should be completely negligible.

The challenge of addressing specific energy levels

Experimentally, when driving a harmonic oscillator with a classical source of radiation, rather than addressing single Fock states one instead creates a coherent state, an excitation of the oscillator that has a well-defined amplitude and phase (to within the limits imposed by Heisenberg's uncertainty relation). A coherent state is described by a complex amplitude α , the real and imaginary components of which are given by the conjugate variables of the system. For example, returning to position (x) and momentum (p), we can write: $\alpha = x + ip$. Quantum mechanically, one can define a coherent state as the application of a unitary displacement operator $\hat{D}(\alpha)$ on the oscillator's vacuum, or the Fock state $|0\rangle_f$ [Glauber, 1963; Haroche and Raimond, 2006]:

$$\hat{D}(\alpha) = e^{(\alpha\hat{a}^\dagger - \alpha^*\hat{a})} \quad (2.24)$$

$$|\alpha\rangle = \hat{D}(\alpha) |0\rangle_f \quad (2.25)$$

$$= e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle_f. \quad (2.26)$$

A coherent state is thus a superposition of the oscillator's Fock states weighted by a Poisson distribution, where the average photon number $\bar{n} = |\alpha|^2$ equals its variance. In addition, a coherent state is in fact an eigenstate of the lowering operator: $\hat{a}|\alpha\rangle = \alpha|\alpha\rangle$, as can be verified with a simple calculation.

We can represent a state of the oscillator in phase-space with a Husimi- Q function $Q^{[\rho]}(\beta)$ [Haroche and Raimond, 2006]:

$$Q^{[\rho]}(\beta) = \frac{1}{\pi} \langle \beta | \rho | \beta \rangle \quad (2.27)$$

$$= \frac{1}{\pi} \langle 0 |_f \hat{D}(-\beta) \rho \hat{D}(\beta) | 0 \rangle_f, \quad (2.28)$$

where ρ is the density matrix of the the oscillator. The Q function returns the average projection of the oscillator state coherently displaced by β onto $|0\rangle_f$ for all points in phase space. The Q function for $|\alpha\rangle$ (Fig. 2.2a) is therefore a $2D$ Gaussian centered at α with a finite width (a result of the uncertainty relations). The oscillator's vacuum, $|0\rangle_f$, is a Gaussian with a variance in the two quadratures given by Φ_{ZPF} and Q_{ZPF} . A single Fock state has a Q function that underscores the absence of any phase information: a ring peaked at α that is weighted by a Gaussian envelope and

symmetric about the origin (Fig. 2.2b). Although linear oscillators can exhibit quantum mechanical properties, they nonetheless do not provide us with the two-level system we seek. We must introduce a non-linearity into the system in order to lift the degeneracy of the energy spectrum.

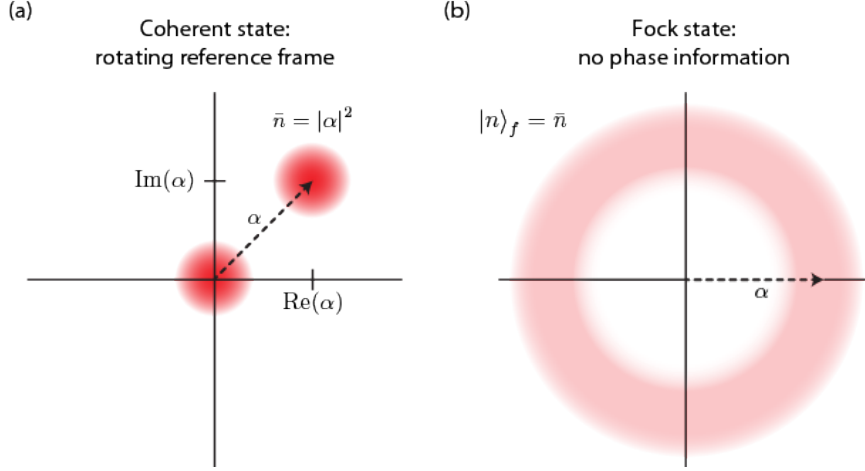


Figure 2.2: Husimi- Q function of coherent states and Fock states. (a) This Q function shows that displacing the oscillator vacuum creates a coherent state of amplitude α ($\hat{D}(\alpha)|0\rangle_f = |\alpha\rangle$), where the average photon number $\bar{n} = |\alpha|^2$. In this example, we are in the rotating reference frame, and so the coherent state stands still in phase space. Experimentally, this requires tuning the the classical source of radiation that performs the displacements to measure the Q function to the oscillator's resonance frequency ω . (b) The Q function of a single Fock state, a ring centered about the origin, contains no phase information, as expected. In this example, the Fock state $|n\rangle_f = \bar{n}$ is depicted.

Introducing anharmonicity with the Josephson junction

We introduce an anharmonicity into the oscillator by replacing the linear inductor with a Josephson junction, a structure that is comprised of two superconducting islands separated by a thin insulating layer (Fig. 2.3a). One can show that the Josephson junction behaves as a non-linear inductor by invoking the two remarkable Josephson relations [Josephson, 1962, 1965; Tinkham, 2004]:

$$I(\phi) = I_c \sin \phi \quad (2.29)$$

$$\hbar \frac{\partial \phi}{\partial t} = 2eV, \quad (2.30)$$

where $I(\phi)$ is the current through the junction; ϕ is the phase difference between the wave functions on either side of the barrier; I_c is the superconductor's critical current; e is the electron charge; and V is an applied voltage across the junction from some external source.

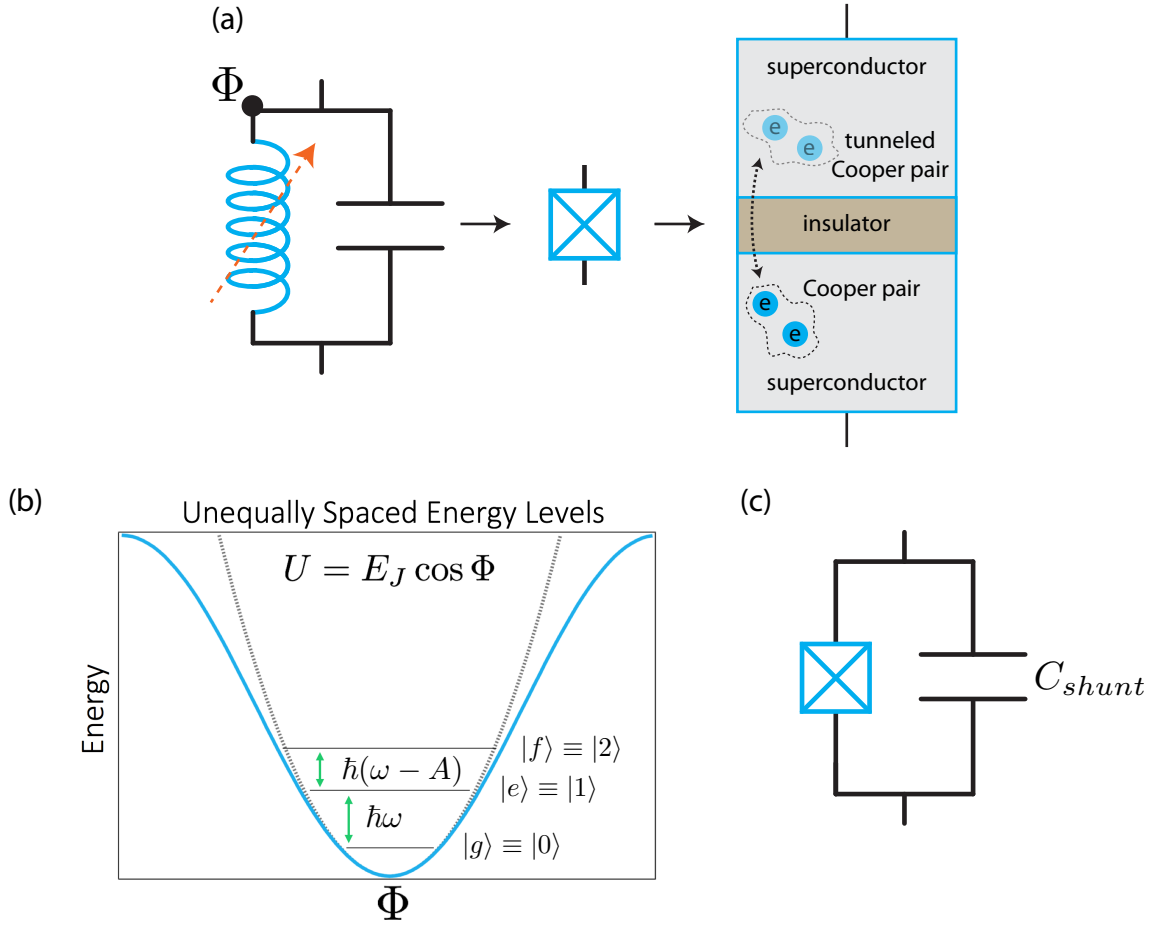


Figure 2.3: The Josephson junction. (a) The Josephson junction is comprised of two superconductors that are separated by a thin insulating layer that allows the tunneling of Cooper pairs. When cooled far below the temperature at which the superconducting phase transition occurs, this element (depicted as a square with a cross inside) behaves as a nearly lossless non-linear inductor. (b) The non-linear inductance changes the quadratic potential (dotted gray line) to a cosine potential (blue) that scales with the Josephson energy E_J . This has the effect of lifting the degeneracy of the energy spectrum, introducing an anharmonicity (A) that allows us to isolate the first two energy levels (ground state $|g\rangle$; first excited state $|e\rangle$) as our physical qubit. In a real experimental implementation, unwanted excitations to higher excited states (e.g. $|f\rangle$) break the two-level approximation, resulting in code-space leakage. (c) The transmon circuit consists of a Josephson junction shunted by a large capacitance C_{shunt} , which lowers the charging energy and thus decreases the sensitivity of the transmon frequency to charge noise.

We first recall the relation shown in Eq. 2.3, and see that in conjunction with Eq. 2.30 the phase ϕ and the node flux Φ are simply related:

$$\phi = 2\pi \frac{\Phi}{\Phi_0}, \quad (2.31)$$

where $\Phi_0 = h/2e$ is the superconducting flux quantum. The junction inductance can then simply be shown to be a non-linear function of Φ :

$$V = L \frac{d\Phi}{dt} \tag{2.32}$$

$$= L \frac{dI(\Phi)}{d\Phi} \frac{d\Phi}{dt} \tag{2.33}$$

$$L(\Phi) = \frac{\Phi_0}{2\pi I_c \cos(2\pi \frac{\Phi}{\Phi_0})} \tag{2.34}$$

Likewise, the resulting energy stored in the junction can be calculated from the two Josephson relations as well [Devoret and Martinis, 2004]:

$$E(t) = \int_{-\infty}^t I(t')V(t')dt' \tag{2.35}$$

$$= E_J \cos(2\pi \frac{\Phi}{\Phi_0}) \tag{2.36}$$

where $E_J = \Phi_0 I_c / 2\pi$ is the Josephson energy. We see that now the oscillator potential is no longer quadratic, but rather a cosine function of Φ (Fig. 2.3b), lifting the degeneracy of the energy spectrum that makes linear LC oscillators difficult to manipulate in a non-classical way.

With the Josephson junction characterized, we can finally move on to the specific realization of the physical qubit that we will heavily focus on for the remainder of this work: the transmon. By introducing a large shunting capacitance in parallel with the junction (Fig. 2.3c), we lower the charging energy and transform the device into an anharmonic oscillator that is highly insensitive to charge noise [Houck et al., 2009; Koch et al., 2007; Schreier et al., 2008]. There are a number of other superconducting qubit flavors, such as the Cooper pair box [Bouchiat et al., 1998; Nakamura et al., 1999], the flux qubit [Mooij et al., 1999], the phase qubit [Martinis et al., 2002], and the fluxonium qubit [Manucharyan et al., 2009]. The transmon, however, offers an excellent resource for strong coupling to cavities for the purposes of state readout and entangling operations, all while maintaining good coherence properties and large anharmonicities that enable the use of fast gates. Their hamiltonians are well understood and can be custom-tailored by changes in the fabrication recipes. Moreover, they are reproducible, easy to package, and have the straightforward path toward scalability. In the next section, we introduce a cavity QED (cQED) architecture that couples a single transmon to two cavities at the same time, and begin to explore the enormous capacity such a flexible setup offers for studying numerous quantum optics phenomena in the microwave frequency domain.

2.2 Coupling a transmon to two cavities

Cavities are structures that support standing electromagnetic waves at certain indexed frequencies, or modes, which are set by the cavity geometry. A cavity can be understood as an infinite series of parallel LCR oscillators, where the inductances and capacitances are chosen to match the frequencies and characteristic impedances of the resonant modes, and the resistances reflect the finite quality factors of each mode [Pozar, 1998]. They can take on numerous different forms, including Fabry-Perot cavities, transmission lines, and in the context of transmons and the research goals outlined previously, superconducting 3D cavities. By coupling a transmon to a superconducting cavity, we first gain the ability to inhibit spontaneous emission by designing the transmon to have a resonance frequency that is far detuned from that of any cavity. Assuming no other sources of decay afflict the system, for example, the coherence times of the transmon and cavity are related through the Purcell effect [Purcell, 1946]. Indeed, it has been shown in 2D geometries that coherence times can be enhanced [Houck et al., 2008] by tuning a transmon out of resonance with a cavity mode to which it is capacitively coupled, thereby indicating the suppression of Purcell-induced spontaneous emission into multiple lossy cavity modes. Furthermore, we can use the cavity as an effective device to read out the state of a superconducting qubit in a QND way, as was shown first in a 2D geometry with the Cooper pair box [Blais et al., 2004a; Wallraff et al., 2004, 2005]. Finally, by off-resonantly coupling the transmon to the cavity, we can apply the dispersive and rotating wave approximations to the Jaynes-Cummings Hamiltonian [Haroche and Raimond, 2006; Jaynes and Cummings, 1963] to place our system into the strong dispersive regime of cQED [Schuster et al., 2007], wherein due to the a.c. Stark shift the transmon frequency depends (to first order) linearly on the number of photons in the cavity, and vice-versa. This regime, for example, allows us to entangle the transmon and cavity for the purpose of using the former to probe the photon number in the latter in a QND way [Johnson et al., 2010].

Despite the advantages of precise fabrication and a small form factor, 2D geometries in cQED systems with transmons and superconducting cavities suffer from low coherence times, typically on the order of several microseconds². Although no one source of loss seems to be responsible for this performance [Martinis and Megrant, 2014], there is strong evidence that dielectric losses arising from interfaces between the superconductors and their substrates contribute substantially to the degradation of T_1 times [Martinis et al., 2005; Wang et al., 2015]³. The change to a 3D

²Although T_1 times can be several tens of microseconds [Barends et al., 2013], T_ϕ has not been shown to be as high.

³Although the latter study was carried out in a 3D geometry, its conclusions are nonetheless valid for 2D systems as well.

geometry, wherein the majority of the electric field of every mode in the system resides in vacuum, thus presented a major leap forward in terms of both transmon and cavity lifetimes [Paik et al., 2011; Reagor et al., 2013]. Further advances in understanding how to diagonalize Hamiltonians with arbitrary numbers of physical qubits and cavities with a so-called “black box quantization” (BBQ) approach opened the door to building complex multi-physical qubit and multi-cavity systems in which the coupling strengths and anharmonicities could be predicted with great accuracy [Nigg et al., 2012]. In particular, by using BBQ to understand how to design an architecture in which two 3D cavities are bridged by a single transmon (Fig. 2.4), numerous experiments have demonstrated unprecedented control over manipulating coherent states and understanding the stochastic evolution of cavity electromagnetic fields in time [Blumoff et al., 2016; Heeres et al., 2015; Holland et al., 2015; Kirchmair et al., 2013; Leghtas et al., 2015a; Reagor et al., 2016; Vlastakis et al., 2013; Wang et al., 2016]. In this work, we will expand on three results in particular [Ofek et al., 2016; Sun et al., 2013; Vlastakis et al., 2015]. Before doing so, however, we first explore in detail the full Hamiltonian of this architecture and simple examples of how the transmon and cavities are characterized.

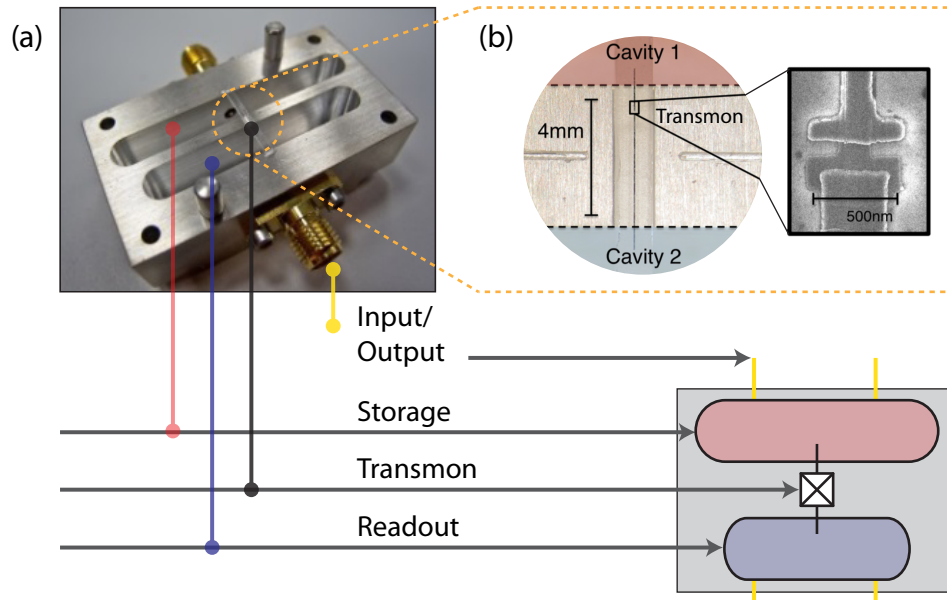


Figure 2.4: **System hardware.** (a) The two-cavity, single transmon architecture that sits in a 10 mK environment inside a dilution refrigerator. The storage cavity has a high quality factor, the transmon physical qubit is used as an ancillary system (ancilla) to probe and manipulate the storage cavity, and the readout cavity is used to perform projective measurements of the ancilla state. Coupling pins are used to apply pulses to each component from microwave electronics at room temperature. (b) Zoom-in images of the trench region between the two cavities, which show the two coupling antennas of the transmon and the Josephson junction that separates them.

Building the system Hamiltonian

Earlier we arrived at the result that the energy of a Josephson junction is the cosine of the node flux, scaled by the Josephson energy (Eq. 2.36). As shown in sec. 2.1, this flux can be promoted to a quantum operator, which allows us to write down the Hamiltonian for the circuit:

$$H = -E_J \cos[\hat{\Phi}] \quad (2.37)$$

$$= -E_J \cos[\Phi_{ZPF}(\hat{a} + \hat{a}^\dagger)], \quad (2.38)$$

where we have used the result in Eq. 2.16. When coupling other oscillators to this circuit, or cavities in this case, the modes of the system become hybridized, and therefore an excitation in any of the modes necessarily results in some current flowing through the junction and thus some contribution to the flux. For a two-cavity, single transmon architecture, we can thus write down the full Hamiltonian as:

$$H = -E_J \cos[\Phi_{ZPF}^a(\hat{b} + \hat{b}^\dagger) + \Phi_{ZPF}^s(\hat{a}_s + \hat{a}_s^\dagger) + \Phi_{ZPF}^r(\hat{a}_r + \hat{a}_r^\dagger)] \quad (2.39)$$

With this Hamiltonian we introduce some new notation. With the architecture shown in Fig. 2.4, both cavities are designed to have high quality factors, with single-photon lifetimes ranging from $\sim 100 \mu\text{s}$ to $\sim 1 \text{ms}$. One of these cavities, the “storage” (subscripts “s”), will be kept high-Q, while the second, referred to as the “readout” (subscripts “r”) will be intentionally over-coupled to a 50Ω output line and used to quickly probe the state of the transmon. Finally, the transmon will be used as an ancillary system, or ancilla (subscripts “a”), to aid in the manipulation of the storage cavity in ideally a QND way. To understand how these functions can be performed, we Taylor expand the Hamiltonian (Eq. 2.39), and invoking the rotating wave approximation (RWA) [Cohen-Tannoudji et al., 1992; Haroche and Raimond, 2006] we keep only non-rotating terms:

$$\begin{aligned} \hat{H}/\hbar &\approx \tilde{\omega}_a \hat{b}^\dagger \hat{b} + \tilde{\omega}_s \hat{a}_s^\dagger \hat{a}_s + \tilde{\omega}_r \hat{a}_r^\dagger \hat{a}_r \\ &- \chi_{sa} \hat{b}^\dagger \hat{b} \hat{a}_s^\dagger \hat{a}_s - \chi_{ra} \hat{b}^\dagger \hat{b} \hat{a}_r^\dagger \hat{a}_r - \chi_{sr} \hat{a}_s^\dagger \hat{a}_s \hat{a}_r^\dagger \hat{a}_r \\ &- \frac{K_a}{2} \hat{b}^{\dagger 2} \hat{b}^2 - \frac{K_s}{2} \hat{a}_s^{\dagger 2} \hat{a}_s^2 - \frac{K_r}{2} \hat{a}_r^{\dagger 2} \hat{a}_r^2. \end{aligned} \quad (2.40)$$

In this expression, the $\tilde{\omega}$ are the modified resonance frequencies when the storage, readout, and ancilla are coupled together. This coupling comes about due to the transmon antennae that protrude into the interior of each cavity (Fig. 2.4b), resulting in a dipole coupling strength g that is typically on the

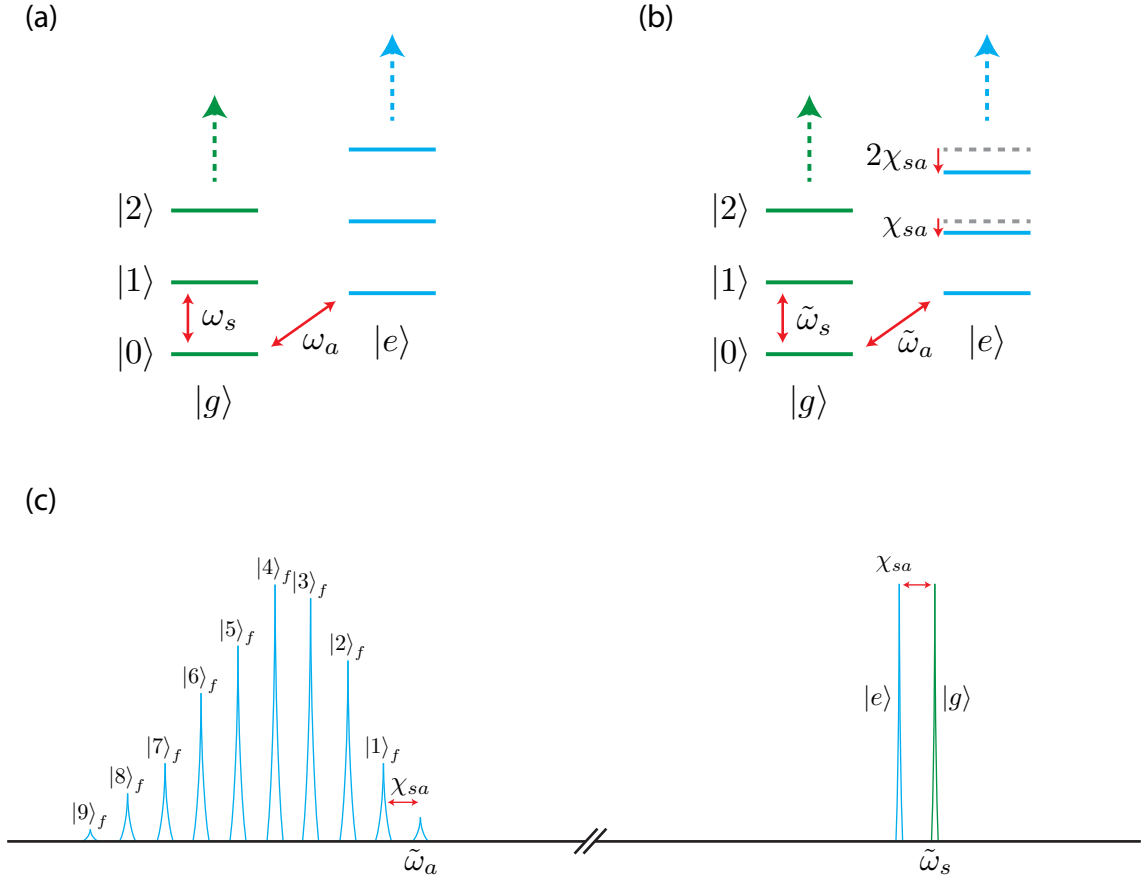


Figure 2.5: Dispersive Hamiltonian energy level diagram. (a) This diagram shows the total energy in the system. An uncoupled cavity and transmon have independent energy level ladders. The resonance frequency of the cavity ω_s is independent of the number of excitations in the transmon (resonance frequency ω_a), and vice versa. (b) When the transmon and cavity are coupled and the detuning between the two is large, not only do the bare resonance frequencies shift ($\omega_s \rightarrow \tilde{\omega}_s$, $\omega_a \rightarrow \tilde{\omega}_a$) due to the interaction, but the resonance frequency of each depends linearly (to first order) on the number of excitations in the other, with a difference in frequency of $-\chi_{sa}$ per excitation. (c) This diagram can also be represented in a spectroscopy-like depiction, which shows the shift in the resonance frequency of the cavity when the transmon is in $|e\rangle$ versus $|g\rangle$, and likewise an example of transmon qubit spectroscopy with a coherent state of light in the cavity, which shows in this case about ten different peaks with Poisson-weighted heights, each shifted down in frequency by the product of χ_{sa} and the Fock state number.

order of 100–200 MHz. If the cavity and ancilla are on resonance with one another, the parameter g sets the frequency of vacuum Rabi oscillations [Schoelkopf and Girvin, 2008]. The dispersive shift in the resonance frequencies for a large detuning $\Delta \gg g$ between the cavity and transmon is second order in $(g/\Delta)^2$ [Schuster et al., 2007]. As a result, the bare mode frequencies are modified ($\omega_s \rightarrow \tilde{\omega}_s$, $\omega_a \rightarrow \tilde{\omega}_a$) and the frequency of each mode depends linearly (to first order) on the number of excitations in the other (Fig. 2.5), resulting in Hamiltonian terms such as $-\chi_{sa} \hat{b}^\dagger \hat{b} \hat{a}_s^\dagger \hat{a}_s$ above.

Crucially, we can use these terms in the Hamiltonian to perform conditional operations in our system. Namely, we use the frequency shift of the readout cavity χ_{ra} to detect the ancilla state [Blais et al., 2004a; Wallraff et al., 2004, 2005], and likewise use the shift χ_{sa} to perform conditional gates between the ancilla and the storage [Johnson et al., 2010]. For example, we can flip the state of the ancilla if and only if there are 0 photons in the cavity by driving the ancilla at $\tilde{\omega}_a$ with a long pulse with a narrow spectral content.

The remaining terms in the third row of Eq. 2.40 are the anharmonicities, or Kerrs, of each component. As discussed, when the cavities and transmon are coupled together, all modes become hybridized. As a result, they have non-zero participation in the energy stored in the junction, and so an excitation in any mode always causes some current to flow in the junction. The degree of this participation is set by the magnitude of the zero-point fluctuations and the Josephson energy, and so the transmon has the highest anharmonicity of the three by upwards five orders of magnitude; typically $K_a \approx 200$ MHz while $K_s, K_r \approx 1\text{--}10$ kHz. The anharmonicity of the transmon is essential, as it allows us to isolate two levels of the physical qubit and perform the operations necessary for quantum computation. For example, if $K_a \approx 200$ MHz, microwave pulses (e.g. with a Gaussian envelope) as short as 10 ns can be used to pulse the transmon from the ground to excited state ($|g\rangle \rightarrow |e\rangle$) with minimal spectral content at the next transition from the excited to second excited state⁴ ($|e\rangle \rightarrow |f\rangle$). The anharmonicities of the modes that ideally should remain harmonic, in particular those of the storage cavity, are one of the prices we pay for incorporating the Josephson junction. The greater the non-linearity, the larger the dispersive shifts and faster the operations, but also the larger the cavity Kerrs. As discussed later in this work, K_s in particular leads to some degradation of stored information. Indeed, the expansion of the cosine has infinitely many terms, and so as coherence times increase and durations of experiments become long enough, small corrections that include terms on the order of hundreds of hertz or less become important as well, and significantly, can be understood and corrected [Nigg et al., 2012; Reagor et al., 2015].

For many of the experiments we perform, we make the approximation that the transmon really is a two-level system, replacing the operators \hat{b}, \hat{b}^\dagger with simply the projector $|e\rangle\langle e|$ onto the excited state. In addition, the readout cavity is also typically in vacuum⁵ during most of the pulse sequence, except during the measurement of the transmon state, which can in many instances be treated as simply an ideal projective measurement onto $|g\rangle\langle g|$; and so we exclude any Hamiltonian term containing $\hat{a}_r, \hat{a}_r^\dagger$. Thus, the final simplified Hamiltonian that well-approximates the dynamics of

⁴In practice, corrections such as DRAG [Chow et al., 2010a; Motzoi et al., 2009] are necessary for fast pulses to prevent leakage to $|f\rangle$.

⁵The small finite thermal population of the readout cavity always contributes to some degree of transmon dephasing. This is discussed in detail in [Sears et al., 2012].

the two-cavity, one ancilla transmon system is given by the following expression:

$$\hat{H}/\hbar = \tilde{\omega}_s \hat{a}_s^\dagger \hat{a}_s + (\tilde{\omega}_a - \chi_{sa} \hat{a}_s^\dagger \hat{a}_s) |e\rangle \langle e| - \frac{K_s}{2} \hat{a}_s^{\dagger 2} \hat{a}_s^2, \quad (2.41)$$

where we keep the Kerr of the storage cavity due to its importance in the QEC experiments to be described in later chapters.

Characterizing a transmon and storage cavity

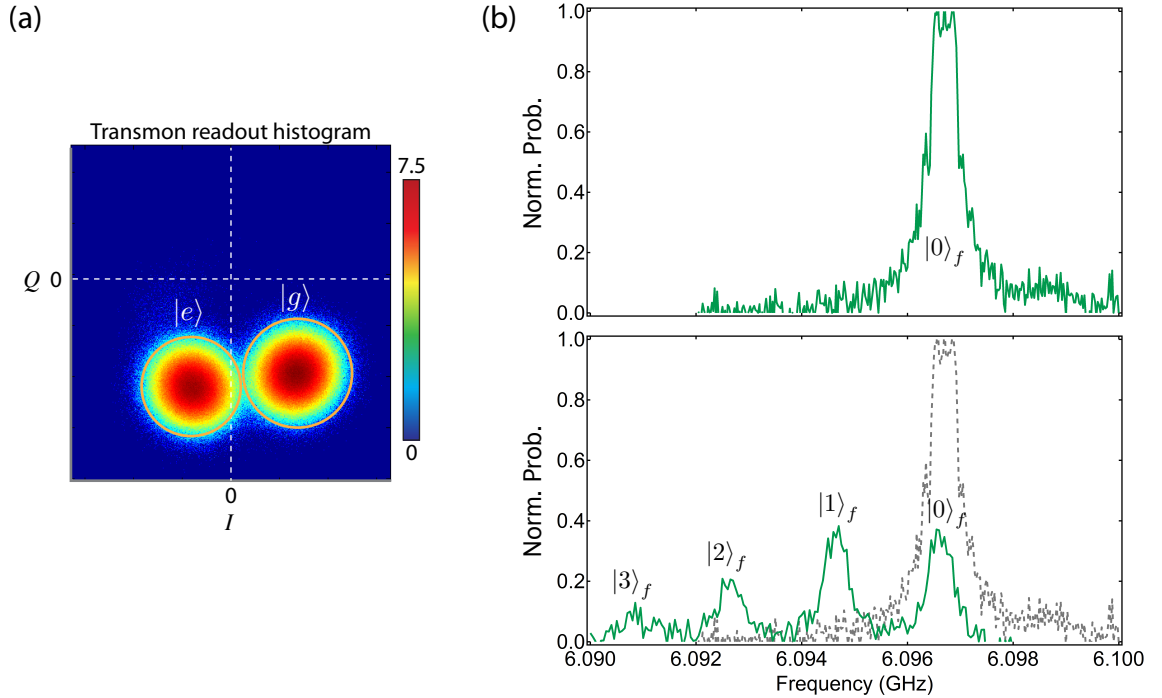


Figure 2.6: Ancilla spectroscopy with a coherent state in the cavity. (a) This 2D histogram shows an example of demodulated and integrated in-phase (I) and quadrature (Q) values of readout pulses. Note the logarithmic scale (base 2). This demonstrates our ability to learn the state of the ancilla transmon in a singleshot way using a parametric amplifier (here a JPC), achieving discrimination fidelities of over 99%. (b) (Top) When the storage cavity is in the vacuum ($|0\rangle_f$), ensemble-averaged ancilla spectroscopy returns a single peak at the frequency $\tilde{\omega}_a \approx 6.097$ GHz. (Bottom) When the storage is displaced with a coherent state of amplitude $\alpha \approx 1$, resulting in an average photon number of $\bar{n} = |\alpha|^2 \approx 1$, the resulting spectroscopy shows that the ancilla is dispersively shifted down in frequency in increments of $\chi_{sa} \approx 2$ MHz per Fock state. This “number splitting” spectroscopy is a vivid example of the uncertainty in photon number associated with coherent states. The noisiness of the traces is a result of a finite integration length of the readout signal and the number of averaged traces. There is a trade-off: longer integration lengths result in better signal-to-noise, but increase the risk of transmon decay during the readout, while averaging more risks broadening the spectroscopy peaks if the transmon frequency drifts. The width of the peak is set by the transmon’s T_2 , corresponding essentially to the quality factor of the transmon mode.

Any experiment begins with device characterization. Although the BBQ procedure accurately predicts the frequencies and couplings between the transmon and the cavities, it is our job to have

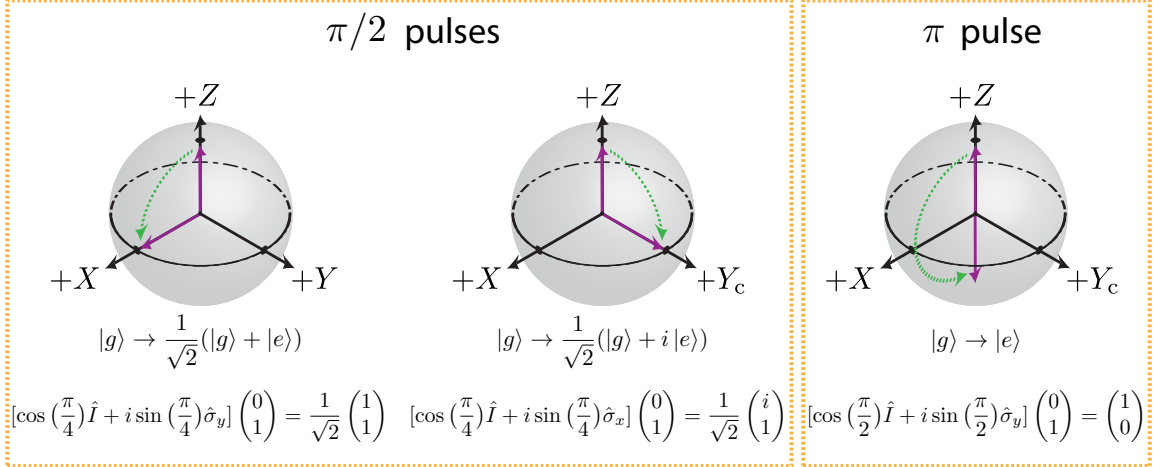


Figure 2.7: Typical Bloch vector rotations. In most experiments we perform predominantly a small set of pulses that take the Bloch vector to and from the equator (“ $\pi/2$ ” pulses: left, middle), or flip the physical qubit (“ π ” pulse: right). The images show examples of these pulses in action on a physical qubit initialized in $|g\rangle$. Note that we can imprint a phase between $|g\rangle$ and $|e\rangle$ by changing the phase of the microwave drive at specific points in a given experimental sequence.

a good estimate of E_J , which sets the Hamiltonian energy scale. To do so without cooling down the device, we use the Ambegaokar-Baratoff formula [Ambegaokar and Baratoff, 1963], which to a good approximation relates the junction conductance at room temperature and energy gap of the superconductor to E_J :

$$E_J = \frac{1}{2} \frac{\hbar}{(2e)^2} G_N \Delta, \quad (2.42)$$

where G_N is the normal state conductance and Δ is the gap⁶. For aluminum, the material of choice for our junctions, G_N is in fact about 15% lower at base as opposed to room temperature, and $2\Delta/\hbar \approx 100$ GHz. Having cooled down the setup with an accurate estimate of parameters in hand, we first test the quality of the readout signal. Crucial to the success of all the QEC results to follow is the ability to learn the state of the ancilla in a single shot [Hatridge et al., 2013; Vijay et al., 2011], which requires excellent signal amplification at the quantum limit. We accomplish this by using a Josephson parametric amplifier, either a phase-sensitive (JBA) [Kamal et al., 2009; Murch et al., 2013; Vijay et al., 2009] or a phase-preserving (JPC) [Bergeal et al., 2010], as the first stage of amplification. The resulting histogram shows a clear bimodal distribution between (Fig. 2.6a) ancilla ground $|g\rangle$ and excited $|e\rangle$ states.

⁶We stress the difference between the choice of symbol (Δ) for the superconducting energy gap and the detuning used above. These are, perhaps unfortunately, the standard notations, but should clearly not be confused here.

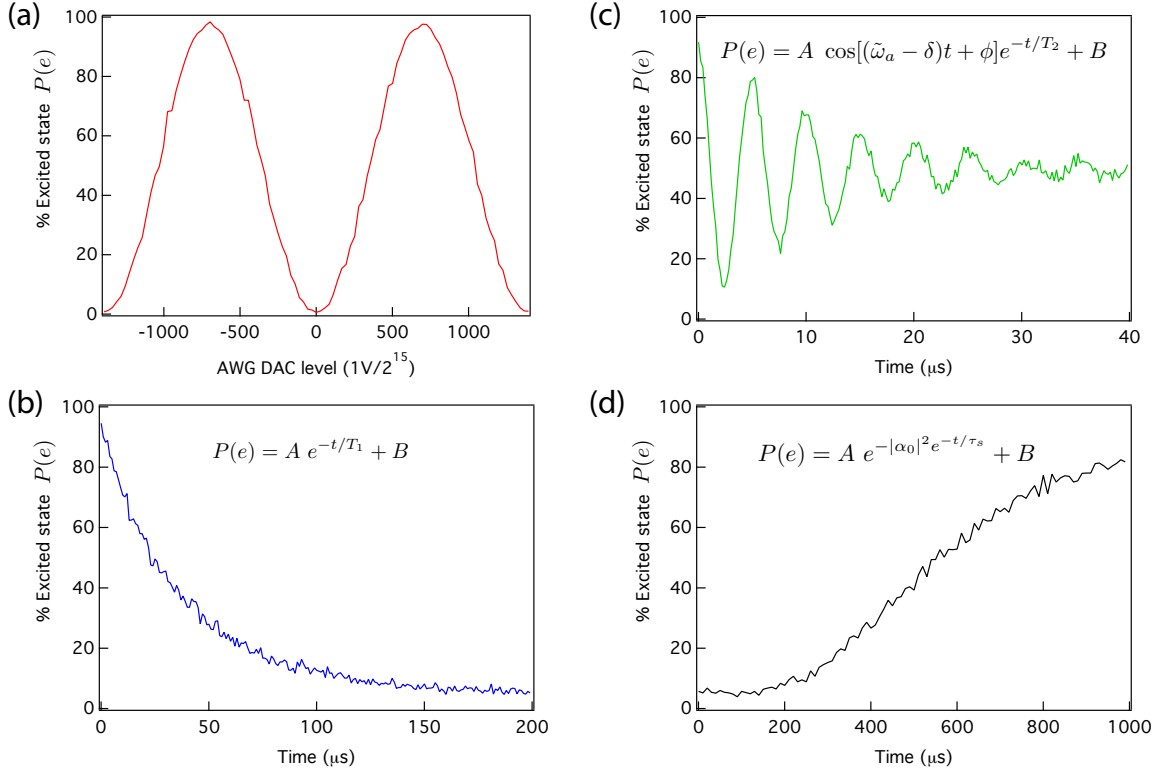


Figure 2.8: Characterizing experimental components. (a) We perform power Rabi experiments to determine the DAC values on our AWGs that correspond to π and $\pi/2$ pulses. (b) Using this information, we perform T_1 experiments by π pulsing the ancilla from $|g\rangle \rightarrow |e\rangle$, varying a delay time, and then reading out the state. In this case, $T_1 \approx 30 \mu s$. The exponential curve has an amplitude A that is ideally +1 and is set by the initial conditions of the transmon state before the π pulse. The curve asymptotes to a value B , which gives the equilibrium thermal population of the transmon. (c) We also perform T_2 experiments, which are comprised of a Ramsey sequence of two $\pi/2$ pulses separated by a varying delay time. In addition to providing us with a coherence time, this experiment also indicates how far detuned we are from the transmon mode. In this example, we purposely introduced a $\delta \sim 250$ kHz detuning, and see that the T_2 is roughly $10 \mu s$. There may also additionally be some phase offset ϕ . (d) We can also find the single-photon lifetime τ_s of the storage cavity by displacing the cavity with a coherent state of amplitude α_0 , waiting varying lengths of time, applying a number-selective π pulse on the ancilla such that it is flipped from $|g\rangle$ to $|e\rangle$ only if there are 0 photons remaining, and then performing a projective readout of the ancilla. As the coherent state decays toward vacuum, the probability of ancilla excitation increases. The amplitude A is set by how effective the long, number-selective π pulse is. With a finite T_2 , the transmon dephasing during this pulse is non-negligible, limiting the maximum signal. The offset B is in fact determined by both the thermal population of the cavity and the transmon. In this example, $\tau_s \approx 250 \mu s$.

Using this readout capability, we can then perform a series of spectroscopy measurements to find the various resonant modes of the system. Figure 2.6b, for example, shows transmon spectroscopy with and without a coherent state in the storage cavity, which demonstrates that we can find the $|g\rangle \rightarrow |e\rangle$ transition in the presence of many photons in the storage, thus extracting χ_{sa} . With the resonant frequencies known, we can proceed to manipulate the transmon in various ways (Figs. 2.7, 2.8), including performing Rabi-flopping experiments to determine the power needed to

flip the transmon from $|g\rangle$ to $|e\rangle$; T_1 experiments to extract the average rate of amplitude decay; and T_2 experiments to extract the frequency detuning and also the average rate of dephasing $\Gamma_\phi = 1/T_\phi$. We again note the relationship: $T_2^{-1} = (2T_1)^{-1} + T_\phi^{-1}$. Additionally we perform an experiment to extract the average single-photon lifetime of the storage cavity, τ_s , which involves performing conditional pulses on the ancilla that reveal whether there are any photons remaining in the cavity after a coherent displacement and a varying delay time. This is of course identical to measuring the Q function at the origin of phase space as a function of time. The exact sources of decoherence that limit T_1 , T_2 , and τ_s cannot be pinpointed with full certainty, but likely candidates include photon shot-noise dephasing through finite residual thermal population in the low-Q readout cavity [Sears et al., 2012], dielectric surface and interface losses [Wang et al., 2015], and quasi-particles [Pop et al., 2014], among others.

2.3 A cavity's photon number parity

Eigenstates of photon number parity can be understood as states of the quantum harmonic oscillator in which only either even-numbered or odd-numbered Fock states are occupied (Fig. 2.9):

$$|\psi\rangle_e = \sum_n c_n^e |2n\rangle_f \quad (2.43)$$

$$|\psi\rangle_o = \sum_n c_n^o |2n+1\rangle_f. \quad (2.44)$$

Although perhaps trivial to see, the loss of single excitations from these states changes the even states to odd states and vice-versa:

$$\hat{a} |\psi\rangle_e = \sum_n c_n^{o'} |2n+1\rangle_f \quad (2.45)$$

$$= |\psi\rangle_o' \quad (2.46)$$

$$\hat{a} |\psi\rangle_o = \sum_n c_n^{e'} |2n\rangle_f \quad (2.47)$$

$$= |\psi\rangle_e'. \quad (2.48)$$

Furthermore, we can define a parity operator $\hat{P} = e^{i\pi\hat{a}^\dagger\hat{a}}$, of which any $|\psi\rangle_e$ or $|\psi\rangle_o$ are eigenstates with eigenvalues $+1$ or -1 , respectively:

$$\hat{P}|\psi\rangle_e = +|\psi\rangle_e \quad (2.49)$$

$$\hat{P}|\psi\rangle_o = -|\psi\rangle_o. \quad (2.50)$$

We note, however, that in general mixed states can also satisfy the property $\text{Tr}(\hat{P}\rho_e) = +1$ and $\text{Tr}(\hat{P}\rho_o) = -1$, where ρ_e (ρ_o) is an incoherent mixture of even (odd) Fock states. These simple yet important properties of the quantum harmonic oscillator will in fact play an important role in the quantum error correction code we will be investigating in later chapters.

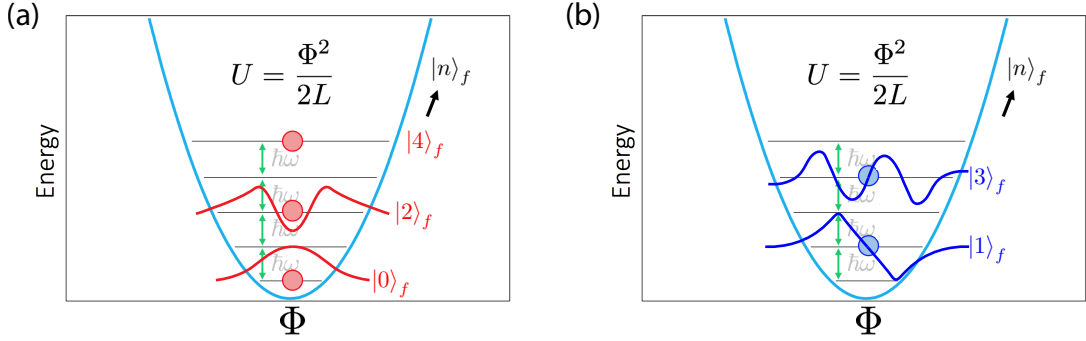


Figure 2.9: Cavity photon number parity. (a) An arbitrary cavity state that is comprised of only even Fock states is said to be of even parity (red disks). The wave functions are symmetric about 0 in the position coordinate Φ . (d) Conversely, a cavity state with only odd Fock states is of odd parity (blue disks). The wave functions are anti-symmetric about 0 in the position coordinate Φ . Note that applying the lowering operator \hat{a} on an odd parity state transforms it to an even parity state and vice-versa.

A parity measurement can be realized in the experiment with a simple Ramsey pulse sequence. It consists of two $\pi/2$ pulses on the ancilla separated by a waiting time of $t = \pi/\chi_{sa}$. In order to calibrate this waiting time, a Ramsey experiment is performed, wherein the two pulses are separated by varying delays (Fig. 2.10). Assuming the rotating reference frame of the ancilla is set at its resonance frequency when the storage cavity is in vacuum, with no displacement the Ramsey contrast decays with a time constant given by T_2 of the transmon. Increasing displacements at first result in a more rapid decay of the signal as the state in the ancilla and the storage cavity become entangled at rate χ_{sa} , where due to the dispersive interaction ancilla Bloch vectors associated with different Fock states $|n\rangle_f$ precess at rates proportional to $n\chi_{sa}$. After a time of $2\pi/\chi_{sa}$ one witnesses a revival in the Ramsey signal as the ancilla and cavity become disentangled; each Bloch vector acquires a phase proportional to 2π . This demonstrates the coherent nature of the interaction. At a time $t = \pi/\chi_{sa}$, opposite parities (even: red, odd: blue) refocus individually and point in opposite directions as all

Bloch vectors associated with odd photon numbers acquire a phase $(2n+1)\pi$. Likewise, in the cavity picture, by waiting for $t = \pi/\chi_{sa}$, one realizes a controlled-phase gate $C_\pi = \hat{I} \otimes |g\rangle\langle g| + e^{i\pi \hat{a}_s^\dagger \hat{a}_s} \otimes |e\rangle\langle e|$, adding a π phase shift per photon on the cavity state conditioned on the ancilla state [Bertet et al., 2002; Vlastakis et al., 2013]. A $\pi/2$ pulse at this point in time maps opposite parities to different ancilla states. In the example shown in Fig. 2.10, even maps to excited state $|e\rangle$ and odd maps to ground state $|g\rangle$.




Figure 2.10: Calibrating a parity mapping. The five different plots show the resulting Ramsey contrast for five different initial displacements of the storage cavity: $\hat{D}(\alpha = 0, 0.75, 1.5, 1.75, 2)$. Cartoon Bloch spheres at different points in time show the distribution of even (red) and odd (blue) Bloch vectors precessing at multiples of χ_{sa} as dictated by the dispersive interaction in the Hamiltonian, where the rotating frame has been set at the ancilla frequency when the cavity is in vacuum. For large enough displacements, note that at $t = \pi/\chi_{sa}$ the average signal becomes flat, indicating that the average parity is zero. At $t = 2\pi/\chi_{sa}$, we witness the characteristic revival of the Ramsey signal as the ancilla and cavity momentarily return to a product state.

The Husimi- Q function offers us a good way to understand the distribution of photons in phase space, but poorly characterizes coherence between Fock state components, since interference fringes in the Q function tomograms are exponentially suppressed with increasing state size [Haroche and Raimond, 2006]. For the remainder of this work, we will instead concentrate on the Wigner function,

which quantifies the average photon number parity of the cavity at every point in phase space [Cahill and Glauber, 1969; Haroche and Raimond, 2006; Lutterbach and Davidovich, 1997]. In addition to providing information as to the distribution of Fock states, the Wigner function has the capability to depict interference fringes, the hallmark of a quantum state, with high contrast. This will be an indispensable tool for characterizing not only the levels of entanglement we can see between the storage cavity and ancilla, but extracting an error syndrome from a logical qubit encoded in a cavity state.

Quantitatively, the Wigner function is the expectation value of the parity operator after the cavity state ρ_s is coherently displaced with an amplitude β :

$$W(\beta) = \frac{2}{\pi} \text{Tr}[\hat{D}(\beta)^\dagger \rho_s \hat{D}(\beta) \hat{P}] \quad (2.51)$$

Knowing the Wigner function is equivalent to knowing the density matrix of the cavity state ρ_s [Cahill and Glauber, 1969; Haroche and Raimond, 2006] (up to a maximum photon number set by the largest value of the displacement used to characterize the state in the IQ plane) and reconstructions may be performed to determine ρ_s from a Wigner tomogram [Kirchmair et al., 2013]. As in the Q function, coherent states look like disks in phase space with a (more narrow) Gaussian envelope (Fig. 2.11). Unlike the Q function, however, the Wigner function can be negative, a feature that indicates the non-classical nature of a particular state. In this sense, it is a quasi-probability distribution, and thus just as the trace of ρ_s must be +1, the integral of the Wigner function over both position, dx (equivalently, along the I axis), and momentum, dp (equivalently, along the Q axis), must equal +1. In fact, one can also use the Wigner function to determine the marginal probability distribution along x or p by integrating $W(x, p)$ along the conjugate variable. For example, if the cavity is in a pure state ($\rho_s = |\psi_s\rangle\langle\psi_s|$):

$$|\psi_s(x)|^2 = \int dp W(x, p) \quad (2.52)$$

$$|\psi_s(p)|^2 = \int dx W(x, p) \quad (2.53)$$

The negativity of the Wigner function is clear if one looks at the tomogram of a Fock state; as with the Q function, the circular symmetry arises due to the lack of phase information, while the oscillations into regions of negative average parity demonstrate that Fock states are quantum in nature. Note that the origin of the Wigner function simply reveals the actual photon number parity of the state itself, with no displacements applied. Thus, the Fock state shown in Fig. 2.11 shows the greatest negativity precisely at the origin, since the parity of $|3\rangle_f = -1$. A number of other works have

demonstrated excellent control of photon states and Wigner tomography in cQED systems [Bertet et al., 2002; Brune et al., 1992; Deléglise et al., 2008; Hofheinz et al., 2009; Kirchmair et al., 2013; Vlastakis et al., 2013].

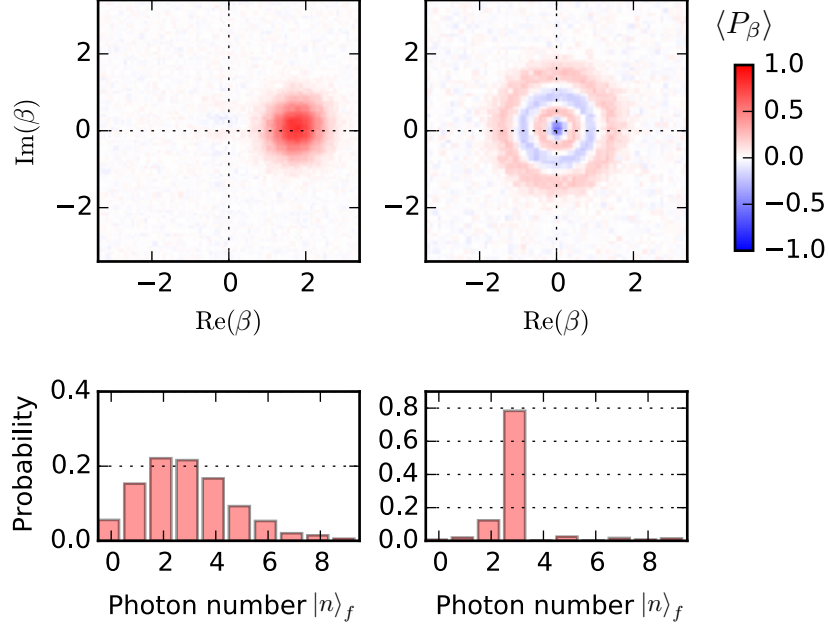


Figure 2.11: **Wigner functions of a coherent state and Fock state.** (a) A measured Wigner function of a coherent state $|\alpha\rangle$ where $\alpha = \sqrt{3}$. The photon number distribution shows the expected Poisson envelope. The color scale shows the average parity: $\langle \hat{P}_\beta \rangle = \langle \hat{D}(\beta) \hat{P} \hat{D}(\beta)^\dagger \rangle$. (b) The measured Wigner function of a 3-photon Fock state. The photon number distribution is dominated by $|n\rangle_f = 3$, with experimental non-idealities such as photon loss contributing to other spurious components.

Although Fock states may on the surface be easier to understand in terms of applications to quantum computation, given their discrete nature, as described earlier they aren't easy to create in a real experimental setting with the cQED setup introduced above. Coherent states, on the other hand, require a simple and short pulse of a microwave generator. The work of [Vlastakis et al., 2013] demonstrated how a quantum bit of information can be mapped onto superpositions of coherent states, or cat states:

$$|\psi\rangle_{2cat}^\pm \approx \frac{1}{\sqrt{2}}(|\alpha\rangle \pm |-\alpha\rangle) \quad (2.54)$$

These coherent superpositions are in fact eigenstates of photon number parity (Fig. 2.12). Furthermore, since coherent states are eigenstates of \hat{a}_s , these cat states are in fact eigenstates of \hat{a}_s^2 , a suggestive feature useful to QEC and to which we will return in the next chapter. Figure 2.12

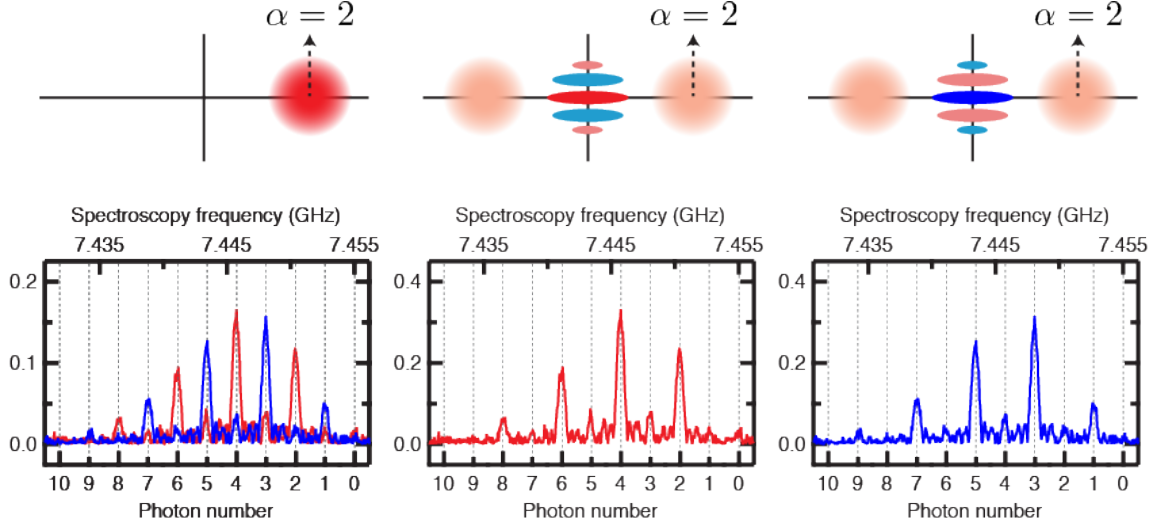


Figure 2.12: **Ancilla spectroscopy with cat states in the cavity.** The Wigner functions in the top panel are cartoons; the spectroscopy in the bottom panel is actual data from [Vlastakis et al., 2013]. With a coherent state in the cavity (left), the ancilla transmon spectroscopy shows the number splitting that is characteristic of the dispersive interaction. The height of each peak matches the Poisson coefficient for the corresponding Fock state component of the coherent state expansion in the Fock basis. When even and odd (middle, right) cat states are in the cavity, however, the spectroscopy shows that only even (red) or only odd (blue) Fock state components are detected. Weak signals of opposite parity are due to single photon loss during the measurement. The cartoons of the cat states illustrate the expected fringe pattern for the cat states. Note that the fringes are π out of phase between even and odd cats, and furthermore that the origin of the Wigner function indicates the parity of the state (even: bright red, $\langle \hat{P}_0 \rangle = +1$; odd: bright blue, $\langle \hat{P}_0 \rangle = -1$).

shows the spectroscopy of the ancilla transmon when cat states of opposite parity are created in the storage cavity. As we will show, the highly coherent nature of cavities makes them a natural choice for quantum computation applications beyond reading out the state of a physical qubit. Their structural simplicity cuts down on the complexity, and thus the number of potential decoherence channels in the setup. Such attributes are often under-appreciated in the field, leading to unrealistic expectations and predictions as to future scalability and performance of large multi-physical qubit systems. In this work, we instead strive to convince the reader of the virtues of hardware efficiency.

In cavity QED with Josephson junction-based physical qubits we find an excellent platform to pursue our goal of realizing quantum computation. This architecture offers a clear physical element that can be well-approximated as a two level system, the transmon; through single-shot measurements we can learn the transmon state with high fidelity and then proceed to apply arbitrary pulses to prepare this physical qubit in any state we desire; gates can be tuned to excellent accuracy; and the coherence properties range from tens of microseconds to upwards of milliseconds in the case of cavities, offering ample time for gate operations that can be performed on time scales of tens of nanoseconds. These five DiVincenzo criteria must be satisfied by a system before it can be

scaled up to perform meaningful calculations as a quantum computer [Nielsen and Chuang, 2010]. Current physical qubit performance in cQED is indeed approaching levels where scaling to larger multi-physical qubit architectures, which can perform quantum error correction in a fault-tolerant manner, can be explored [Barends et al., 2014]. These results, although promising and exciting, however, must be interpreted with caution. The architectures in which they were demonstrated were modest in size and complexity. As will be shown and emphasized throughout this work, the true metric of performance is only obtained when the entire system is up and running. For example, while illustrating how bit-flip error correction ostensibly suppresses transmon decay from $|e\rangle$ to $|g\rangle$ when the number of physical qubits increases from five to nine [Kelly et al., 2015], one cannot simply assume that one will continue to gain with greater physical qubit numbers. One must prove it experimentally. The following chapters lay the groundwork for what is the culmination of the work presented here: using the simple two-cavity, single transmon setup introduced previously to implement a full QEC system that protects an arbitrary bit of quantum information while exposed to every possible source of decoherence. It is an exciting time for the field of cQED. The results we describe here present a real step forward toward the ultimate goal of building the first practical quantum computer.

3 – The Cat Code

The basic principles of quantum mechanics prevent the application of classical schemes of error correction to a quantum system in order to restore the integrity of the information. Indeed, by simply measuring a quantum state to check for the occurrence of errors, one necessarily projects it out of a possible superposition in which it may be, thereby losing any information encoded in relative quantum amplitudes and phases. Furthermore, the no-cloning theorem [Wootters and Zurek, 1982] prevents one from making copies of an unknown quantum state. Overcoming these hurdles requires more sophisticated techniques in quantum error correction (QEC), ones which borrow from the basic classical strategies of measurement and code redundancy while at the same time adhering to the core constraint that in the process nothing must be learned about the actual quantum bit of information.

Redundantly encoding quantum information need not require copying; by using the capability of entangling gates as a resource, one can encode a quantum bit in a higher dimensional space while still maintaining the quantum amplitudes of the initial state. The challenge is to devise an encoding scheme that transfers information from a single physical qubit onto a collection of physical qubits, or logical qubit, which is endowed with cleverly chosen symmetry properties that allow one to extract an error syndrome without disturbing the original information. Implementing QEC in the laboratory is challenging, requiring a complex system with many physical qubits. Even for a perfectly realized QEC system of finite size, there will always be unrecoverable errors or failure modes, resulting in an exponential decay of the information over time. In fact, error correction first introduces a decoherence penalty, since an uncorrected logical qubit consisting of n physical qubits (for typical first order codes $n \sim 5 - 10$ [Steane, 1996]) will experience decoherence that is of order n times faster. A central goal of QEC is to suppress the naturally occurring errors and surpass the break-even point, where the lifetime gain due to error correction is larger than this penalty. These considerations motivate exploring a hardware-efficient approach to QEC, with which it may be more tractable to not only overcome the decoherence penalty, but furthermore to pinpoint the leading limitations to fault-tolerance.

In this chapter, we provide an overview of the traditional error correction codes that require many physical qubits to realize. In so doing, we will motivate the need for a hardware-efficient approach to QEC, one that exploits many of the properties of superconducting cavities that make them attractive candidates for a quantum memory. We will then conclude with a particular QEC scheme called the cat code, which proposes a logical encoding of a quantum bit in superpositions of coherent states in a cavity and calls for measuring just one error syndrome in the system: changes in photon number parity.

3.1 Constructing a quantum error correcting code

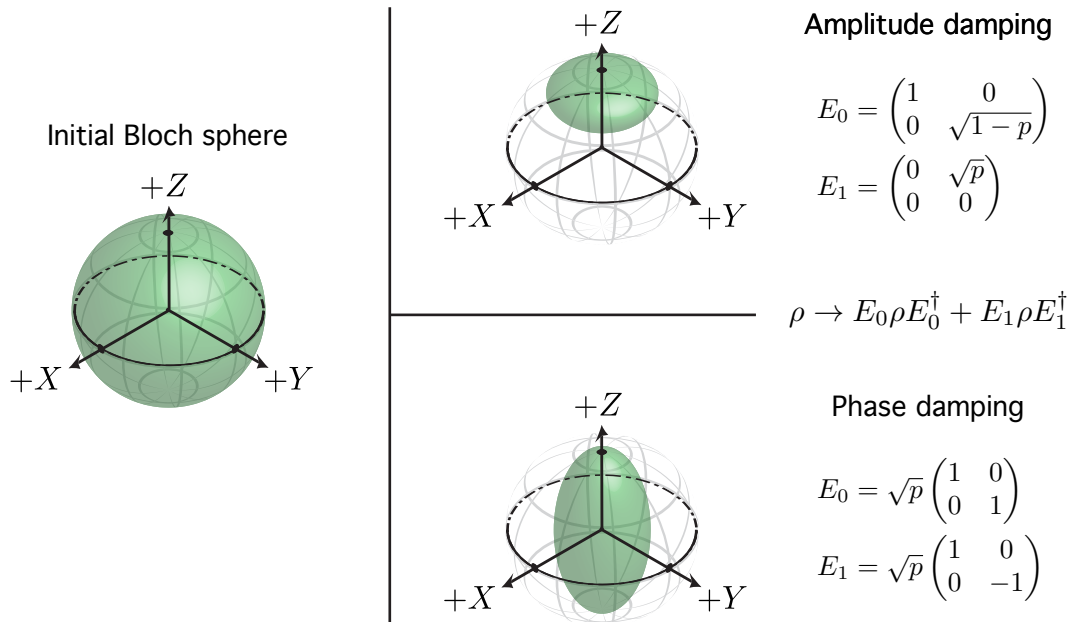


Figure 3.1: **Effects of amplitude and phase damping on a quantum bit.** Cartoon Bloch spheres show the effects of the two dominant error channels on a quantum bit. The initial Bloch sphere shaded in green represents any pure state, where the length of the Bloch vector is 1. Under the action of amplitude damping, with some probability $p = e^{-t/T_1}$ the environment measures the loss of a photon due to spontaneous emission, modeled by the error channel \hat{E}_1 . Not detecting photon loss constitutes a measurement as well, and so the probability of the quantum bit occupying the excited state decreases in time, as modeled by \hat{E}_0 . The time t is chosen to allow for appreciable decay. In the case of phase damping, where $p = e^{-t/T_\phi}$, every vector shrinks towards the Z -axis; only the ground and excited states $|g\rangle$ and $|e\rangle$ remain pure.

The remarkable discovery of QEC protocols [Shor, 1995] was a critical advance that gives hope for eventually realizing useful quantum computers. These protocols are predicated on the understanding that physical two-level systems (e.g. the transmon, to a good approximation), are generally susceptible to two sources of decoherence: amplitude damping and pure dephasing [Nielsen and Chuang,

2010]. The former can be thought of as resulting from the spontaneous emission of a photon¹, or the application of the lowering operator $\hat{\sigma}_- = \hat{\sigma}_x + i\hat{\sigma}_y$ on the Bloch vector with a probability characterized by T_1 . The latter, although describing a continuous process, can nonetheless be modeled by the probabilistic occurrence of phase-flips, where the discrete error channel is instead the stochastic application of $\hat{\sigma}_z$ with a probability characterized by T_ϕ . The evolution of a state exposed to these error channels can be visualized as a shrinking of the Bloch sphere toward the energetically preferential ground state for amplitude damping, or toward the Z -axis for phase damping (Fig. 3.1). The crucial point is that by discretizing the continuous errors to the probabilistic occurrence of either bit-flips ($\hat{\sigma}_x$), phase-flips ($\hat{\sigma}_z$), or both ($\hat{\sigma}_y = i\hat{\sigma}_x\hat{\sigma}_z$), the challenge of implementing QEC to protect a quantum bit from any error in principle becomes tractable.

In designing an error correction code, the logical encoding must be chosen carefully. One must first ensure that different errors in the code are mapped to orthogonal subspaces to allow for error detection, and second that no measurement back-action occurs from the occurrence of an error, where the measurement is performed by the environment. These requirements are formalized by the Knill-Laflamme condition [Knill and Laflamme, 1997], which states:

$$\langle L_k | \hat{E}_n^\dagger \hat{E}_{n'} | L_{k'} \rangle = \lambda_n \delta_{nn'} \delta_{kk'}, \quad (3.1)$$

where $|L_k\rangle$ is one of K logical states, \hat{E}_n is one of N possible errors, and λ_n is proportional to the probability of that error occurring. In the case of a logical qubit, $K = 2$ ($|0_L\rangle \equiv |L_0\rangle$ and $|1_L\rangle \equiv |L_1\rangle$) and for the physical qubits considered in this work $N = 4$ for the four different possible error operators: \hat{I} , $\hat{\sigma}_x$, $\hat{\sigma}_y$, $\hat{\sigma}_z$. Assuming for the moment that $k = k'$ in Eq. 3.1 so that $\delta_{kk'} = 1$, the term $\delta_{nn'}$ requires that each error operator takes the same state $|L_k\rangle$ to orthogonal subspaces, satisfying the first requirement. Now assuming that $\delta_{nn'} = 1$, not only are different logical states taken to orthogonal subspaces when acted upon with the same error operator, but also that the probability of error, proportional to λ_n , is independent of the logical state chosen. For example, if $\hat{E}_n = \hat{\sigma}_x^L$ (some logical bit flip $|0_L\rangle \leftrightarrow |1_L\rangle$), and $\langle 0_L | \hat{E}_n^\dagger \hat{E}_n | 0_L \rangle \neq \langle 1_L | \hat{E}_n^\dagger \hat{E}_n | 1_L \rangle$, then $\hat{\sigma}_x^L(c_0 |0_L\rangle + c_1 |1_L\rangle) = c'_0 |0_L\rangle + c'_1 |1_L\rangle$; the changes $c_0 \rightarrow c'_0$ and $c_1 \rightarrow c'_1$ are the consequences of measurement back-action and result in code failure since information about the logical qubit has been gained.

¹In a realistic experimental setting, the finite temperature of the physical qubit additionally results in a probability of excitation, or the application of the raising operator $\hat{\sigma}_+ = \hat{\sigma}_x - i\hat{\sigma}_y$. Additionally, when the two-level approximation breaks down, code-space leakage occurs. This effect is not included in the discussion here.

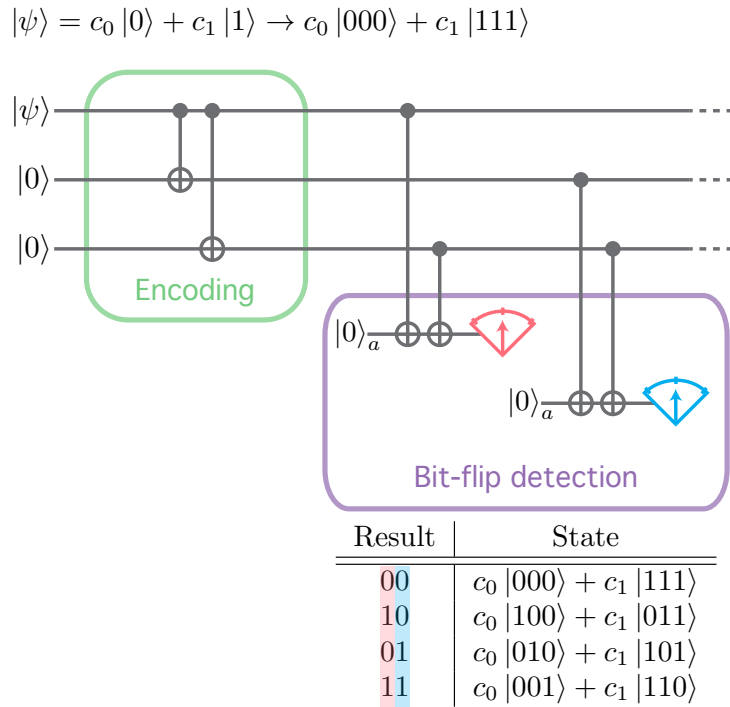


Figure 3.2: Three-qubit bit flip code protocol. Encoding consists of two CNOT operations between the first physical qubit and two additional physical qubits, which maps the state $|\psi_q\rangle$ into a larger Hilbert space where the logical basis states become $|000\rangle$ and $|111\rangle$. Error syndromes are measured by mapping the parities of the encoding onto two ancilla qubits ($|0\rangle_a$). Measurements of the ancillae project the code either back into the initial code space or into one of three error spaces; each space is associated with a unique measurement pattern. This code can be modified to detect phase flips as well, but needs greater redundancy in order to correct for bit and phase flips together.

The example of a three-qubit bit flip code illustrates the basic idea [Nielsen and Chuang, 2010]. Let’s assume a state $|\psi\rangle_{init} = c_0 |0\rangle + c_1 |1\rangle$ is encoded in a single physical qubit, and that this physical qubit can only experience continuous errors that rotate the Bloch vector about an unknown axis in the xy plane of the Bloch sphere. Cloning is not allowed, and measurement destroys the superposition, so a redundant encoding is necessary to allow one to discretize the continuous space of errors into errors correctable by unitary operations. As shown in Fig. 3.2, the redundant encoding can be realized by entangling the first physical qubit with two extra physical qubits using two CNOT gates:

$$|\psi\rangle_{init} = c_0 |0\rangle + c_1 |1\rangle \rightarrow c_0 |000\rangle + c_1 |111\rangle, \quad (3.2)$$

where the logical qubits are: $|0_L\rangle \equiv |000\rangle$ and $|1_L\rangle \equiv |111\rangle$. This redundancy and ability to correct comes with two related “costs:” an increased complexity, or “resource overhead,” and a “decoherence

penalty.” The resource overhead is related to the number of physical qubits required per logical qubit; it characterizes the difficulty of the experimenter’s task to realize a certain calculation. The greater the resource overhead, the more hardware is required to operate the quantum computer and thus the greater the cost. The decoherence penalty is the enhanced rate of decoherence the experimenter faces within the logical space due to the increased size of the redundant encoding. It is a ubiquitous facet of any QEC code. For example, assuming the three physical qubits in the bit flip code above are identical, the error rate experienced by $|\psi\rangle_{init}$ is three times greater than with a single physical qubit. The advantage of course is that error correction is possible in principle, given that $|\psi\rangle_{init}$ now has symmetry properties, or parities, that allow error syndrome extraction without learning anything about c_0 and c_1 .

Returning to the bit flip code, we complete the example by working through how the error syndromes can be extracted. Assuming that each physical qubit has the same probability p of deviation from its original state, the following eight projectors show all possible error permutations with measurement probabilities given in Fig. 3.3:

$$P_0 \equiv |000\rangle\langle 000| + |111\rangle\langle 111| \quad \text{no error} \quad (3.3)$$

$$P_1 \equiv |100\rangle\langle 100| + |011\rangle\langle 011| \quad \text{bit flip on first physical qubit} \quad (3.4)$$

$$P_2 \equiv |010\rangle\langle 010| + |101\rangle\langle 101| \quad \text{bit flip on second physical qubit} \quad (3.5)$$

$$P_3 \equiv |001\rangle\langle 001| + |110\rangle\langle 110| \quad \text{bit flip on third physical qubit} \quad (3.6)$$

$$P_4 \equiv |110\rangle\langle 110| + |100\rangle\langle 100| \quad \text{bit flips on first and second physical qubits} \quad (3.7)$$

$$P_5 \equiv |101\rangle\langle 101| + |010\rangle\langle 010| \quad \text{bit flips on first and third physical qubits} \quad (3.8)$$

$$P_6 \equiv |011\rangle\langle 011| + |100\rangle\langle 100| \quad \text{bit flips on second and third physical qubits} \quad (3.9)$$

$$P_7 \equiv |111\rangle\langle 111| + |000\rangle\langle 000| \quad \text{bit flips on all physical qubits.} \quad (3.10)$$

Concentrating on P_0 through P_3 for the moment, note that with two bits of information one can learn uniquely which of the three physical qubits had a bit flip. For example, these bits can be the outcomes Z_1Z_3 and Z_2Z_3 of measuring the operators $\hat{\sigma}_{z,1}\hat{\sigma}_{z,3}$ and $\hat{\sigma}_{z,2}\hat{\sigma}_{z,3}$ respectively, where the indices refer to the three physical qubits in the encoding. If one could measure these parities, with probability $\approx 1 - 3p$ (assuming p is small) the state would be projected back to the original $|\psi\rangle_{init} = c_0|000\rangle + c_1|111\rangle$, while with a probability $\approx p$ the state would be projected into one of three error spaces given by a bit flip on one of the three physical qubits in the logical encoding.

In fact, these parities Z_1Z_3 and Z_2Z_3 can be mapped onto two ancillary physical qubits, or ancillae, again using CNOT gates, as shown in Fig. 3.2. By projectively measuring the states of the




Figure 3.3: **Correcting bit flips with the three-qubit code.** This plot shows an example logical encoding in which the probability of a bit flip on each of the three physical qubits the comprise the logical is 20%. With about 90% probability, the state either remains in the initial code space (black bar), or in one of three error spaces in which one of the three physical qubits flipped (red bars). With the remaining 10%, however, multiple errors occur and result in code failure.

ancillae one can use the resulting measurement pattern to learn which of the code physical qubits, if any, had an error; one thus forces the system to choose whether an error has occurred. These error syndrome measurements are QND on the encoded information as they reveal nothing about the amplitudes c_0 and c_1 and serve only to project the state either back into the code space or into one of three unique error spaces. One can check easily that the Knill-Laflamme condition in Eq. 3.1 is satisfied for this encoding. In principle, the error need not be corrected either, as the basis of the code space can be redefined to take into account that a bit flip has occurred. This further simplifies the task of error correction to monitoring changes in the parity with repeated error syndrome measurements in time.

The approximation that p is small of course means that with some small probability that scales initially like p^2 this code will fail; it can be quickly seen that bit flips on two of the physical qubits cannot be distinguished from a single bit flip on the third with the parity measurement scheme above. This code is therefore known as a first-order error correcting code as it can handle only one bit flip per round of correction. Further redundancy, and thus a larger hardware overhead, is necessary to promote this code to higher orders. The effect of this quadratic behavior is evident

in a number of QEC experiments, including those on cQED platforms [Reed et al., 2012]. From these results we can already glean the following key points: for QEC to be successful, a single round of error correction must be fast compared to the average error rate of the logical to keep p low; in addition, the gates and measurements must be high-fidelity so that error syndromes faithfully reflect the true error trajectory of the quantum bit. However, one must be prepared for the fact that there will still always be a finite probability that an error will be missed. As a result, the fidelity of the state is destined to decay exponentially in time if this probability is constant per round of correction.

3.2 The challenges facing traditional QEC schemes

With the three-qubit code above, which just corrects for bit flips,² we already see that fast error syndrome measurements may not be so easy to implement experimentally. Although the four CNOT operations may be implemented quickly and even with the requisite fidelity [Barends et al., 2014; Chow et al., 2012, 2013], the entire error correction step may still be challenging to realize quickly compared to T_1 times in typical cQED systems. In particular, dispersive ancilla readout durations, ancilla resets, and any necessary calculations that may need to be done in real-time quickly add up. Furthermore, in order to suppress both bit and phase flips, traditional quantum error correction schemes [Fowler et al., 2012; Steane, 1996] in fact require an even greater overhead than in the example presented above. In the case of the Steane code shown in Fig. 3.4, the encoding consists of seven physical qubits and two quite complicated logical states:

$$|0_L\rangle = \frac{1}{\sqrt{8}}(|0000000\rangle + |1010101\rangle + |0110011\rangle + |1100110\rangle \tag{3.11}$$

$$+ |0001111\rangle + |1011010\rangle + |0111100\rangle + |1101001\rangle)$$

$$|1_L\rangle = \frac{1}{\sqrt{8}}(|1111111\rangle + |0101010\rangle + |1001100\rangle + |0011001\rangle \tag{3.12}$$

$$+ |1110000\rangle + |0100101\rangle + |1000011\rangle + |0010110\rangle).$$

Furthermore, the Steane code calls for six ancillae to measure six different parities all within a single round of correction. This first order code can handle errors only on one physical qubit in the code space per round of correction. This of course renders the task of realizing gains from using QEC challenging, as the extra decoherence penalty introduced by the larger encoding increases the logical error rate within the code by at least a factor of seven. This means that to begin

²This code can also be easily modified to correct phase flips instead of bit flips, but it cannot correct both bit and phase flips at the same time.

with the coherence times of the constituent physical qubits must be high enough such that all six syndrome measurements can be performed before multiple errors start to degrade the integrity of the information. Higher coherence times, however, are a result of weaker couplings of a physical qubit to its environment, which unfortunately means that the strength with which one can intentionally couple this physical qubit to other quantum systems (e.g. cavities) or experimental control lines may also be reduced. As a result, the ratio of coherence times to the speed with which single and multi-physical qubit gates can be performed is more difficult to increase than each property independently, rendering six fast parity measurements in one round of correction quite challenging in typical quantum information architectures.




Figure 3.4: The Steane code. This first order code can correct both bit and phase flips in one of the seven code physical qubits within one round of QEC. Using a series of controlled-NOT (CNOT) gates, multi-physical qubit parities are mapped onto a number of ancillae, projective readouts of which are used to determine which, if any, of the code physical qubits ($|q_1\rangle$ through $|q_7\rangle$) needs to be corrected and with which operation. Here, $|+\rangle = H|0\rangle$.

Compounding this challenge is that multiple errors within a single round of correction are by no means the only source of infidelity in any real experimental implementation. When looking at how such a code can fail (Fig. 3.5), other general mechanisms include uncorrectable errors, such as code space leakage to states that nullify the assumption that the physical implementation of a physical qubit is a true two-level system; readout and ancilla preparation errors that provide the wrong syndrome information; undesired couplings between physical qubits in the codespace and indeed any other component in the architecture that can induce unknown phases; and perhaps most destructively of all is the forward propagation of errors, where by the very action of measuring the error syndrome new errors are introduced into the code that one cannot correct. This is a

central vehicle of non-fault-tolerance in a code such as this and requires yet further code redundancy to suppress. Note that the first five sources can in principle be suppressed by faster syndrome measurement rates, whereas the sixth favors less frequent interrogation due to the probability of propagating errors from the ancillae into the code. This will be a central consideration for an implementation the error correction system discussed in later chapters.

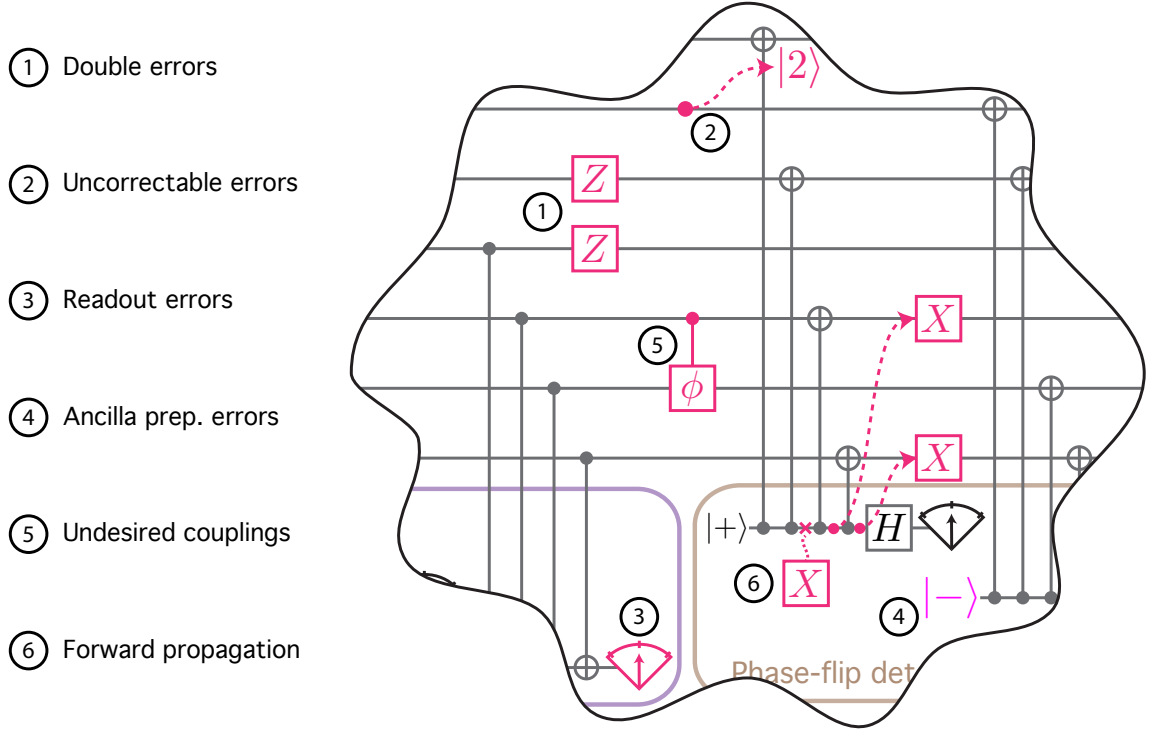


Figure 3.5: **Failure modes in a QEC code.** This zoom-in of a portion of the Steane code from Fig. 3.4 depicts six different possible failure modes that can result in physical qubit decoherence. These general mechanisms are applicable to any measurement-based error correction scheme and must be carefully considered when implementing an actual QEC system experimentally.

In principle, a system that implements QEC can actually pass a “break-even” point, overcoming the decoherence penalty and thereby preserving quantum information for longer than the lifetime of its constituent parts. Returning to the example of the Steane code, however, certain estimates predict that the error rate per gate must be on the order of 1×10^{-4} [Nielsen and Chuang, 2010]. Given that current QEC implementations in superconducting cQED experiments can optimistically measure the six parities in 1–10 μs , the coherence times also need to be upwards of milliseconds to prevent the accumulation of multiple errors, a feat that has been achieved only in fluxonium qubits with T_1 (although not yet T_ϕ) [Pop et al., 2014], but not yet in transmons, which have typically been the physical qubit of choice in recent QEC experiments [Córcoles et al., 2015; Kelly et al., 2015; Reed et al., 2012; Ristè et al., 2015]. Reaching the break-even point and demonstrating the

extension of a physical qubit’s lifetime with such multi-physical qubit paradigms has thus remained an outstanding and challenging goal [Devoret and Schoelkopf, 2013].

Several previous works have demonstrated elements of QEC on platforms such as NMR [Cory et al., 1998; Knill et al., 2001a; Leung et al., 1999; Moussa et al., 2011], ions [Chiaverini et al., 2004; Nigg et al., 2014; Schindler et al., 2011], nitrogen vacancy (NV) centers [Cramer et al., 2016; Taminiau et al., 2014; Waldherr et al., 2014], photons [Aoki et al., 2009; Pittman et al., 2005], and superconducting transmons [Córcoles et al., 2015; Kelly et al., 2015; Reed et al., 2012; Risté et al., 2015]. These works, however, primarily illustrate the signatures or scaling properties of QEC codes rather than test the capacity of the system to extend the lifetime of quantum information over time. Far from suggesting any shortcomings or a lack of sophistication in the experimental techniques employed in those experiments, their inability to preserve a quantum bit of information over time from natural errors underscores the pivotal role played by all failure modes shown in Fig. 3.5 in contributing to decoherence. Indeed, practically speaking, in each implementation it was in fact more advantageous to simply store the quantum bit in the system’s most coherent element, realizing just a passive and uncorrectable quantum memory. What we learn from such endeavors is that system complexity that leads to large decoherence penalties necessitates physical qubit performance levels that simply cannot be currently achieved. Furthermore, it remains unclear whether the reported high fidelities of the universal gate set can be sustained with ever larger and more complex microwave connectivity schemes. We thus strive to find a hardware-efficient solution, one that circumvents paradigms akin to the Steane code by cutting down on the overhead involved, and reduces the impact of each error channel listed in Fig. 3.5 on QEC performance.

3.3 The cat code: a hardware-efficient approach

Although typically used for the purposes of measurement [Blais et al., 2004a; Wallraff et al., 2005] or to mediate interactions between physical qubits [Majer et al., 2007], cavities in superconducting cQED architectures can be a valuable resource for storing a logical qubit in a continuous-variable system [Braunstein, 2005]. It has been shown, in fact, that universal quantum computation with continuous-variable encoding schemes is possible [Gottesman et al., 2001; Lloyd and Braunstein, 1999; Menicucci et al., 2006] and can in fact offer advantages over those employing collections of discrete physical qubits. Although a continuous-variable quantum computer formally has the same power as a discrete-variable system, there are regimes in which it could be more efficient. For example, a single oscillator can in principle accommodate an unlimited amount of information, owing to the infinite size of its Hilbert space. The hardware requirements can be more favorable as well, call-

ing for linear elements and photon detectors in the optical platforms [Knill et al., 2001b], or simple microwave cavities with long coherence times that are easy to assemble in superconducting cQED systems [Reagor et al., 2013, 2016]. Furthermore, the natural relation between continuous variables and communication can in principle simplify the transmission of quantum information for the purposes of teleportation [Braunstein and Kimble, 1998; Lloyd and Slotine, 1998], cryptography [Ralph, 1999], and dense coding [Braunstein and Kimble, 2000], to name a few.

There are trade-offs as well, however, which include challenges resulting from possible non-orthogonality of basis states in experimental realizations, the possibility of continuous excursions from a logical sub-space, and manipulating encoded states with high fidelity. Nonetheless, several promising continuous variable QEC protocols exist [Braunstein, 1998; Leghtas et al., 2013; Lund et al., 2008; Michael et al., 2016; Mirrahimi et al., 2014; Ralph, 2011], and substantial progress has been made in demonstrating that continuous-variable encodings can be a powerful resource for the storage, control, and measurement of quantum information [Aoki et al., 2009; Deléglise et al., 2008; Heeres et al., 2015; Hofheinz et al., 2009; Jensen et al., 2011; Leghtas et al., 2015a; Pittman et al., 2005; Sun et al., 2013; Vlastakis et al., 2013, 2015]. Moreover, with recent progress in demonstrating cQED architectures that offer a natural path toward scalability [Brecht et al., 2015, 2016; Mineev et al., 2016], we see continuous variables systems as a promising platform for realizing a practical quantum computer.

A cavity's one dominant error channel

The time evolution of the density matrix ρ_s of a field in a cavity, which has some equilibrium thermal photon population n_{th}^s , a rate κ_s associated with single photon creation and annihilation operators \hat{a}_s^\dagger and \hat{a}_s , is well-modeled by the following Lindblad operators in the master equation formalism [Haroche and Raimond, 2006]:

$$L_- = \sqrt{\kappa_s(1 + n_{th}^s)}\hat{a}_s \quad (3.13)$$

$$L_+ = \sqrt{\kappa_s n_{th}^s}\hat{a}_s^\dagger, \quad (3.14)$$

where the master equation reads:

$$\begin{aligned} \frac{d\rho_s}{dt} = & -i\tilde{\omega}_s[\hat{a}_s^\dagger\hat{a}_s, \rho_s] - \frac{\kappa_s(1 + n_{th}^s)}{2}(\hat{a}_s^\dagger\hat{a}_s\rho_s + \rho_s\hat{a}_s^\dagger\hat{a}_s - 2\hat{a}_s\rho_s\hat{a}_s^\dagger) \\ & - \frac{\kappa_s n_{th}^s}{2}(\hat{a}_s\hat{a}_s^\dagger\rho_s + \rho_s\hat{a}_s\hat{a}_s^\dagger - 2\hat{a}_s^\dagger\rho_s\hat{a}_s) \end{aligned} \quad (3.15)$$

One can show that such a formulation returns the expected prediction that on average the occupation of the cavity mode $\bar{n} = \text{Tr}[\rho_s \hat{a}_s^\dagger \hat{a}_s]$ simply decays exponentially in time to thermal equilibrium with a characteristic time constant $\tau_s = 1/\kappa_s$:

$$\bar{n}(t) = n_0 e^{-t/\tau_s} + n_{th}^s (1 - e^{-t/\tau_s}) \quad (3.16)$$

Treating n_{th}^s as negligible for the remainder of this discussion, one may conclude that the evolution of the density matrix ρ_s can be simply described by the stochastic application of the lowering operator \hat{a}_s on the cavity field. This assertion, that there are essentially just two dominant processes within the cavity, the application of $e^{-\frac{\kappa_s}{2} \hat{a}_s^\dagger \hat{a}_s \Delta t}$ (for time steps Δt) and the stochastic application of \hat{a}_s , is a powerful incentive to consider storing quantum information in a superconducting cavity rather than a two level system, such as a transmon, which is susceptible to both amplitude (σ_-) and phase damping (σ_z).

One may raise the concern that in a real experiment, if the cavity is coupled to a transmon for example, fluctuations in the the critical current of the Josephson junction may result in some frequency jitter of the cavity mode. The current that runs across the junction when this mode is excited, however, is much smaller than in the case of excitations in the transmon. In other words, the cavity inherits a much smaller non-linearity from the junction than does the transmon, where the anharmonicity of the former is typically four orders of magnitude smaller than that of the latter. Indeed, it has been experimentally demonstrated [Reagor et al., 2016] that 3D superconducting cavities have no currently measurable source of inherent dephasing arising from a Lindblad operator of the form $L_\phi = \sqrt{\kappa_\phi} \hat{a}_s^\dagger \hat{a}_s$ or higher order photon loss mechanisms such as $L_{2ph} = \sqrt{\kappa_{2ph}} \hat{a}_s^2$ [Sun et al., 2013]. In practice, some dephasing is induced by its dispersive coupling to occupation fluctuations of other modes in the system, particularly to that of the transmon; this is a central concern that will be discussed in later chapters.

The governing goal is to thus construct a code that can track the occurrence of single photon jumps, as this would correct for the dominant error channel in the system. In implementing a QEC system to realize this goal, we look to translate the discretized energy dissipation of the cavity field into a unitary operation on an encoded state, the occurrence of which can be deduced from an appropriate error syndrome measurement. It turns out that the simple symmetry property of cavities introduced in chapter 2, photon number parity, is naturally suited to play the role of an error syndrome in such an architecture. As will be described presently, the remarkable feature that cavities have just one dominant error channel in the form a single photon loss allows for a clever encoding scheme that uses coherent states, which are eigenstates of \hat{a}_s . The task is to construct a

logical encoding using these states that satisfies the requirements for quantum error correction as outlined above while minimizing the necessary decoherence penalty. In the remaining chapters we concentrate on one particular QEC proposal called the cat code [Leghtas et al., 2013; Mirrahimi et al., 2014]. We use the experimental system shown in Fig. 2.4 and described in greater detail in A.3. The result is a substantially simplified proposal for a quantum memory that suppresses the one dominant source of natural decoherence in the system. Although facing the set of challenges that comes with moving from a discrete to a continuous-variables system, it nonetheless opens the door to realizing error correction at the break-even point.

A logical basis of cat states

Coherent states $|\alpha\rangle$ are an attractive option for a logical encoding scheme (Fig. 3.6) as they are eigenstates of \hat{a}_s , where $|\alpha\rangle = e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^n}{\sqrt{n!}} |n\rangle_f$ for a complex amplitude α , and $\hat{a}_s |\alpha\rangle = \alpha |\alpha\rangle$. This feature suggests that one could try encoding a quantum bit in a superposition of coherent states: $|\psi_q\rangle = c_0 |0\rangle + c_1 |1\rangle \rightarrow c_0 |\alpha\rangle + c_1 |-\alpha\rangle$. As the overlap between two coherent states falls off exponentially in the difference of their amplitudes [Haroche and Raimond, 2006], choosing an $|\alpha|^2 = \bar{n} \gtrsim 1.5$ would be sufficient for basis states $|\alpha\rangle$ and $|-\alpha\rangle$ to be almost completely orthogonal, with an overlap of $\sim 0.2\%$. The penalty we pay is that the rate of photon jumps γ scales with the mean photon number $\bar{n} = |\alpha|^2$ [Haroche and Raimond, 2006]. This is the QEC decoherence penalty in such a scheme. When $c_0 = c_1 = \pm 1/\sqrt{2}$, $|\psi_q\rangle \rightarrow 1/\sqrt{2}(|\alpha\rangle \pm |-\alpha\rangle)$, the logical encoding is an equal superposition of coherent states that we refer to in this work as “2-cat” states, which are eigenstates of the even (+) or odd (−) photon number parity operator $\hat{P} = e^{i\pi\hat{a}_s^\dagger\hat{a}_s}$:

$$|C_\alpha^+\rangle = \frac{1}{\sqrt{2}}(|\alpha\rangle + |-\alpha\rangle) = \sqrt{2}e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^{2n}}{\sqrt{(2n)!}} |2n\rangle_f \quad (3.17)$$

$$|C_\alpha^-\rangle = \frac{1}{\sqrt{2}}(|\alpha\rangle - |-\alpha\rangle) = \sqrt{2}e^{-|\alpha|^2/2} \sum_{n=0}^{\infty} \frac{\alpha^{2n+1}}{\sqrt{(2n+1)!}} |2n+1\rangle_f, \quad (3.18)$$

where $\langle C_\alpha^\pm | \hat{P} | C_\alpha^\pm \rangle = \pm 1$. This parity is a quantity that can be measured with high fidelity in our system with a simple Ramsey-style pulse sequence [Bertet et al., 2002; Haroche et al., 2007], as was first shown in chapter 2 and will be used extensively in chapters 4 and 5. The problem with this encoding, however, is that aside from the special case of a “2-cat,” for arbitrary c_0 and c_1 there is no parity symmetry and no other measurable symmetry property that would indicate the loss of a photon. For example, as seen in Fig. 3.6a, the cardinal point along $+Y_c$ is a cat state of zero parity.

This requires us to move on to the cat code [Leghtas et al., 2013; Mirrahimi et al., 2014], wherein we access a larger part of the cavity’s Hilbert space in order to accommodate an encoding scheme

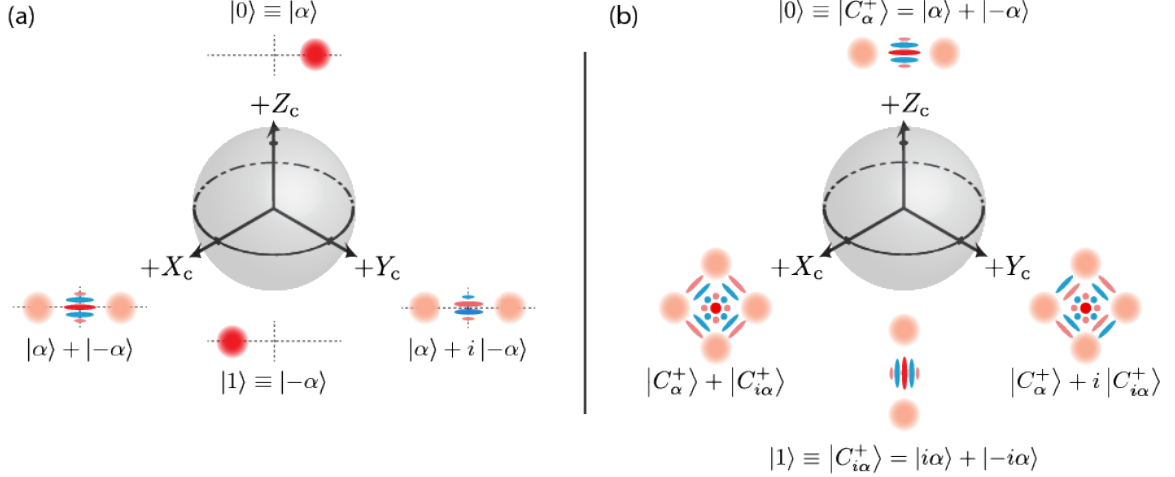


Figure 3.6: Logical encodings with coherent states. (a) The logical $|0\rangle$ and logical $|1\rangle$ are defined by coherent states $|\alpha\rangle$ and $|- \alpha\rangle$, respectively. Note that while the superposition of the two along $+X_c$ is an even parity “2-cat,” any other state in this encoding lacks such a symmetry property. (b) In order to have any arbitrary state on the Bloch sphere an eigenstate of photon number parity, the cat code uses this “2-cat” itself as the logical $|0\rangle$ and the same “2-cat” $\pi/2$ out of phase as the logical $|1\rangle$. Superpositions of these states are called “4-cats.”

where the individual basis states are themselves “2-cats” along the real and imaginary axes in phase space (Fig. 3.6b): $|C_\alpha^\pm\rangle \equiv \mathcal{N}_\alpha^\pm(|\alpha\rangle \pm |-\alpha\rangle)$ and $|C_{i\alpha}^\pm\rangle \equiv \mathcal{N}_\alpha^\pm(|i\alpha\rangle \pm |-i\alpha\rangle)$, where

$$\mathcal{N}_\alpha^\pm(t) = \frac{1}{\sqrt{2(1 \pm e^{-2|\alpha(t)|^2})}}, \quad (3.19)$$

and $\mathcal{N}_\alpha^\pm \rightarrow 1/\sqrt{2}$ for large α . To prevent appreciable basis overlap, one must now have $|\alpha|^2 = \bar{n} \gtrsim 2$ and thus $\gamma \gtrsim 2\kappa_s$, a slightly larger penalty than in the previous scheme, but still substantially smaller than the factor of seven in the Steane code. Such a modification allows us to encode a quantum state with arbitrary c_0 and c_1 in an eigenstate of photon number parity:

$$|\psi\rangle_{init} = c_0 |0\rangle + c_1 |1\rangle \rightarrow c_0 |C_\alpha^+\rangle + c_1 |C_{i\alpha}^+\rangle. \quad (3.20)$$

This in turn allows changes in parity to serve as the error syndrome for the loss of a photon in a logical qubit (Fig. 3.7). After the first photon jump we have:

$$\begin{aligned} \hat{a}_s(c_0 |C_\alpha^+\rangle + c_1 |C_{i\alpha}^+\rangle) &= \mathcal{N}_\alpha^- [c_0(|\alpha\rangle - |-\alpha\rangle) + ic_1(|i\alpha\rangle - |-i\alpha\rangle)] \\ &= c_0 |C_\alpha^-\rangle + ic_1 |C_{i\alpha}^-\rangle, \end{aligned} \quad (3.21)$$

and similarly after the second photon jump:

$$\begin{aligned} \hat{a}_s(c_0 |C_\alpha^- \rangle + ic_1 |C_{i\alpha}^- \rangle) &= \mathcal{N}_\alpha^+ [c_0(|\alpha\rangle + |-\alpha\rangle) - c_1(i|\alpha\rangle + |-i\alpha\rangle)] \\ &= c_0 |C_\alpha^+ \rangle - c_1 |C_{i\alpha}^+ \rangle \end{aligned} \quad (3.22)$$

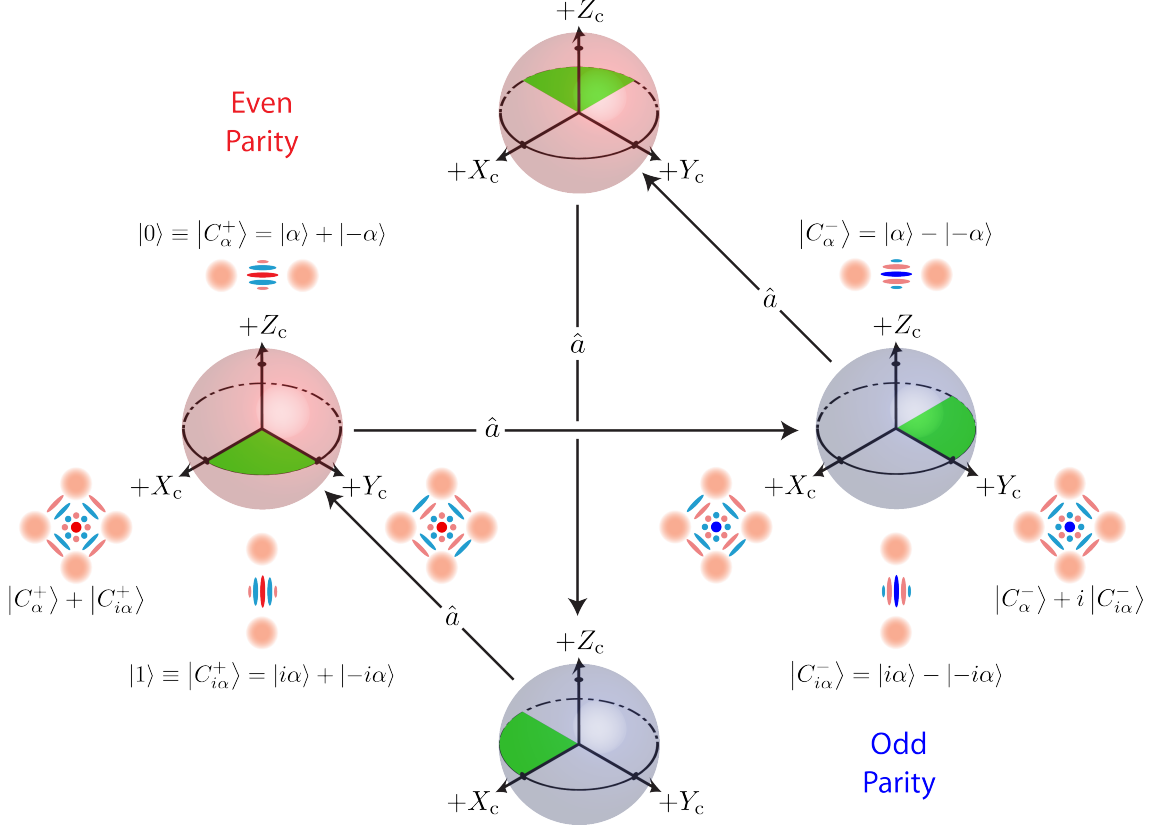
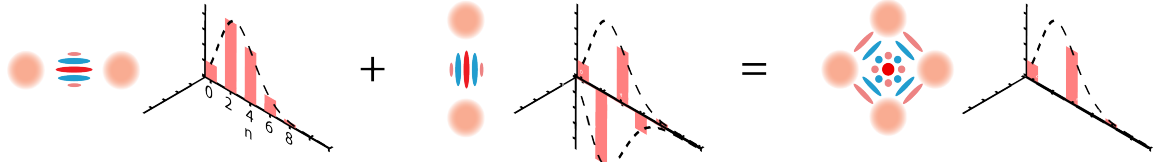


Figure 3.7: The cat code cycle. In the logical encoding of $|0_L^\pm\rangle \equiv |C_\alpha^\pm\rangle = |\alpha\rangle \pm |-\alpha\rangle$ and $|1_L^\pm\rangle \equiv |C_{i\alpha}^\pm\rangle = |i\alpha\rangle \pm |-i\alpha\rangle$ (normalizations omitted), the two “2-cats” $|C_\alpha^\pm\rangle$ and $|C_{i\alpha}^\pm\rangle$ are eigenstates of either even (+) or odd (−) photon number parity (an “ n -cat” is a superposition of n coherent states). For large enough $|\alpha|$ they are effectively orthogonal to one another. In this basis, the states along $+X_c$ and $+Y_c$ are both “4-cats” of even or odd parity as well. The different patterns in the fringes of their cartoon Wigner functions signify the different phase relationship between the basis states. These features allow one to store a quantum bit in a superposition of “2-cats,” $|\psi\rangle_{init} = c_0 |C_\alpha^\pm\rangle + c_1 |C_{i\alpha}^\pm\rangle$, and at the same time monitor the parity as the error syndrome without learning anything about c_0 or c_1 . In this example, we choose to encode in the even parity basis ($|0_L^+\rangle$ and $|1_L^+\rangle$), although the odd basis can equally be chosen. The loss of a single photon changes not just the parity of the basis states (red shading: even; blue shading: odd), but the phase relationship between them by a factor of i as well: $\hat{a}_s(c_0 |C_\alpha^+\rangle + c_1 |C_{i\alpha}^+\rangle) = c_0 |C_\alpha^-\rangle + ic_1 |C_{i\alpha}^-\rangle$. Thus, after one photon jump, one finds the initial quantum bit rotated by $\pi/2$ about the Z_c axis. With each subsequent application of \hat{a}_s , the encoded state cycles between the even and odd parity subspaces, while due to each consequent multiplication of the coefficient c_1 by i , the encoded information rotates about the Z_c axis by $\pi/2$, as indicated by the rotation of the green shaded slice. Between the stochastic applications of \hat{a}_s , the cat states deterministically decay toward vacuum: $\alpha \rightarrow \alpha e^{-\kappa_s t/2}$ (not depicted here), indicating that the logical basis changes in time.

Equations 3.20-3.22 show that the cat code maps a photon loss error in the cavity field onto a rotation by $\pi/2$ about the logical Z axis, as seen from the factor of i that comes out in front of c_1 . With each successive error the parity of the basis states cycles between even and odd, while the encoded information continues to rotate about Z in increments of $\pi/2$, returning to the initial state after four errors. This modulo-four behavior can also be understood from the perspective of a Fock state expansion, shown in Fig. 3.8. For example, due to destructive interference, an equal superposition of the basis states of the form $|C_\alpha^+\rangle + |C_{i\alpha}^+\rangle$ has non-zero Fock state components $|0\rangle_f, |4\rangle_f, |8\rangle_f, \dots$ with the appropriate Poisson coefficients. One can see that after four applications of \hat{a}_s , $|0\rangle_f$ disappears while $|4\rangle_f \rightarrow |0\rangle_f, |8\rangle_f \rightarrow |4\rangle_f$, etc. The relation $\hat{a}_s |n\rangle_f = \sqrt{n} |n-1\rangle_f$ ensures that the resulting state still has the proper Poisson coefficients to be the original superposition of $|C_\alpha^+\rangle$ and $|C_{i\alpha}^+\rangle$. One thus returns to the original state $|C_\alpha^+\rangle + |C_{i\alpha}^+\rangle$.

Left uncorrected, an encoded state devolves into a mixture of cat states at the enhanced rate γ . By performing single-shot parity measurements, however, we will repeatedly update our knowledge as to the parity of the state and infer the occurrence of an error when the parity changes [Sun et al., 2013]. We will thereby follow the stochastic evolution of the cavity state through the cat code cycle, maintaining the coherence of the quantum bit despite errors in the encoding. Note that the cat code also satisfies the Knill-Laflamme condition for QEC. The two dominant operators that act within the logical code space are either \hat{I} or \hat{a}_s , and so one can check that for these two errors and the logical states $|0_L^\pm\rangle \equiv |C_\alpha^\pm\rangle$ and $|1_L^\pm\rangle \equiv |C_{i\alpha}^\pm\rangle$ Eq. 3.1 is satisfied.



$$\begin{aligned}
|0_L^+\rangle &= \frac{1}{\sqrt{2}}(|\alpha\rangle + |-\alpha\rangle) \\
&= \sqrt{2}e^{-\frac{|\alpha|^2}{2}} \sum_{n=0}^{\infty} \frac{\alpha^{2n}}{\sqrt{(2n)!}} |2n\rangle \\
|1_L^+\rangle &= \frac{1}{\sqrt{2}}(i|\alpha\rangle + |-i\alpha\rangle) \\
&= \sqrt{2}e^{-\frac{|\alpha|^2}{2}} \sum_{n=0}^{\infty} \frac{(i\alpha)^{2n}}{\sqrt{(2n)!}} |2n\rangle \\
|X_L^+\rangle &= \frac{1}{\sqrt{2}}(|C_\alpha^+\rangle + |C_{i\alpha}^+\rangle) \\
&= 2e^{-\frac{|\alpha|^2}{2}} \sum_{n=0}^{\infty} \frac{\alpha^{4n}}{\sqrt{(4n)!}} |4n\rangle
\end{aligned}$$

Figure 3.8: A “4-cat” expanded in the Fock state basis. The modulo-four behavior of the cat code can be seen by writing out the “4-cat” in this example in the Fock state basis. This superposition of two “2-cats” $\pi/2$ out of phase results in the destructive interference of all Fock state components except for those that are multiples of four: $|0\rangle_f, |4\rangle_f, |8\rangle_f, \dots$. Applying \hat{a}_s on this state takes $|X_L^+\rangle = \sum_{n=0}^{\infty} c_{4n} |4n\rangle_f \rightarrow \sum_{n=0}^{\infty} c_{4n+3} |4n+3\rangle_f$, where the c_n are the Poisson coefficients for a coherent state of amplitude $|\alpha|$. After three more applications of \hat{a}_s , the state returns to $|X_L^+\rangle$.

Basis orthogonality

In addition to the stochastic loss of single photons, as shown in Eq. 3.16, the energy of the cavity field decays deterministically to vacuum at a rate κ_s , $\alpha(t) = \alpha e^{-\kappa_s t/2}$, as shown in Fig. 3.9. Therefore, in the experimental implementation of the cat code, we must always take into account the decay of the cat state amplitude after a finite time of monitoring t in order to keep track of the evolving basis states:

$$|C_\alpha^\pm\rangle \rightarrow |C_{\alpha e^{-\kappa_s t/2}}^\pm\rangle = \mathcal{N}_\alpha^\pm(t) (|\alpha(t)\rangle \pm |-\alpha(t)\rangle) \quad (3.23)$$

Of course, without any intervention the state stored in the cavity eventually decays to vacuum, thereby erasing any stored information. This effect is not irreversible, however, as energy can be periodically re-pumped into the cavity using dissipative Hamiltonian engineering schemes, for example, wherein the application of off-resonant pumps at carefully chosen frequencies stabilizes the superpositions of coherent states in phase-space [Leghtas et al., 2015a]. Perhaps surprisingly, the loss of a photon has no effect on the amplitude of a coherent state, as can be seen from a simple argument in [Haroche and Raimond, 2006], section 4.4.4. The authors explain that losing single photons simply updates one’s knowledge that there must have been more photons in the cavity immediately prior to the jump. By virtue of this curious property of coherent states, the amplitude of our logical states is independent of the number of photon jumps detected.

The non-orthogonality of the basis states in the cat code is an important consideration in deciding the initial amplitude of the encoded state. Larger cat states mean that the cat code can be employed for longer periods of time without applying any unitary gates or dissipative pumps [Leghtas et al., 2015a] to restore the amplitude. As this is largely a technical point that has been demonstrated not to be a fundamental limitation, the more salient question is at what point does increasing the cat state amplitude begin to adversely affect the performance of the code due to the increased rate of errors. The trade-off between non-orthogonality and average error rate is in fact very generous, however, since the overlap between two coherent states falls off exponentially with the difference between them in a cavity’s phase space [Haroche and Raimond, 2006] while the error rate increases linearly in \bar{n} :

$$\langle \alpha | \beta \rangle = e^{-|\alpha|^2/2} e^{-|\beta|^2/2} e^{-\alpha^* \beta} \quad (3.24)$$

$$|\langle \alpha | \beta \rangle|^2 = e^{-|\alpha - \beta|^2} \quad (3.25)$$

$$|\psi\rangle = \frac{1}{2}(|\alpha\rangle + |-\alpha\rangle + |i\alpha\rangle + |-i\alpha\rangle)$$




Figure 3.9: Coherent state amplitude decay. Regardless of the number of photon jumps, the amplitude of the cat states always decays toward vacuum at the rate $\kappa_s/2$, eventually becoming so small that the overlap between constituent states of the encoding becomes non-negligible and results in a loss of fidelity. Shown here is one particular “4-cat” with different amplitudes, starting from $\alpha = \sqrt{3}$ when the overlap between the nearest coherent states is much less than 1% all the way to vacuum, where all the information is lost.

Using Eq. 3.24, one can perform a similar calculation for the cat code basis states to obtain the following overlaps:

$$|\langle C_\alpha^+ | C_{i\alpha}^+ \rangle|^2 = \left(\frac{2e^{-\alpha^2} \cos(\alpha^2)}{1 + e^{-2\alpha^2}} \right)^2 \quad (3.26)$$

$$|\langle C_\alpha^- | C_{i\alpha}^- \rangle|^2 = \left(\frac{2e^{-\alpha^2} \sin(\alpha^2)}{1 - e^{-2\alpha^2}} \right)^2, \quad (3.27)$$




Figure 3.10: **Overlap of cat code basis states with decaying amplitude.** This plot shows the overlap between the even basis states $|C_\alpha^+\rangle$, $|C_{i\alpha}^+\rangle$; the odd basis states $|C_\alpha^-\rangle$, $|C_{i\alpha}^-\rangle$; and as a reference two coherent states at $|\alpha\rangle$ and $|-\alpha\rangle$. The overlap between $|C_\alpha^-\rangle$ and $|C_{i\alpha}^-\rangle$ increases more rapidly with decreasing α since $|1\rangle_f$ is the lowest odd parity Fock state available, rather than $|0\rangle_f$ as in the case of even parity states. These curves are plotted for a $\kappa_s = 1/250 \mu\text{s}$, and green dotted lines indicate a window in time wherein the basis states decay starting from $\bar{n}_0 = 2$.

where α is understood to be a real number here. The trigonometric terms are a result of the interference changing with α , or equivalently, with time. Figure 3.10 shows how the overlap of the basis states in equations 3.26 and 3.27 changes as α decays. One can see that at an $\bar{n}_0 = 2$ ($\alpha \sim 1.4$), the overlap of the even parity basis states in particular is quite low, thus making it advantageous in this case to encode the quantum bit in a superposition of $|C_\alpha^+\rangle$ and $|C_{i\alpha}^+\rangle$. The overlap continues to be only on the order of a couple percent as time elapses, in fact decreasing at first for even parity states due to the aforementioned interference phenomenon.

The cat code versus traditional QEC

Comparing traditional QEC schemes to the cat code, the former typically protect a physical qubit from decoherence by projecting components of the redundant encoding into spaces defined by four unitary operators: identity \hat{I} , and the Pauli operators $\hat{\sigma}_x$, $\hat{\sigma}_y$ and $\hat{\sigma}_z$. In the latter, however, within the logical encoding of the cat code there are only two such operators: \hat{I} (no photon loss and amplitude decay $|\alpha\rangle \rightarrow |\alpha e^{-\kappa_s \delta t/2}\rangle$) and $(\hat{I} + i\hat{\sigma}_z)/\sqrt{2}$ (application of \hat{a}). Furthermore, given that

the overlap between coherent states $|\langle \alpha | i\alpha \rangle|$ falls off exponentially with increasing $|\alpha|$ [Haroche and Raimond, 2006], we can use basis states of average photon number $\bar{n} \approx 2$, which increases the error rate within the encoding by only a factor of ~ 2 [Haroche and Raimond, 2006] rather than by orders of magnitude as in the traditional schemes presented earlier [Fowler et al., 2012; Steane, 1996].

Having described the details of how the cat code works in theory, how do we realize it in experiment? We will simply use the two-cavity, single transmon setup first introduced in chapter 2. We stand to benefit enormously from the sheer simplicity of such an architecture; 3D cavities are simple to build, provide excellent isolation for physical qubits from their environments, have the best coherence properties to date in cQED systems, and require only one ancilla to measure their dominant error syndrome. Furthermore, for a dispersive shift on the order of several MHz, the parity mapping time takes just a few hundred nanoseconds, upwards of four orders of magnitude faster than the single photon lifetime of recent 3D cavities. As will be seen in later chapters, the speed of this single effective CNOT gate between the ancilla and the logical qubit allows for very high error syndrome measurement fidelities, performance that current implementations of error correction in multi-physical qubit systems aiming to implement the Steane or surface codes cannot currently match.

Beyond offering a correctable quantum memory that suppresses the decoherence of a quantum bit, the cat code can be used to implement universal quantum computing as well [Mirrahimi et al., 2014]. Crucial to the viability of this paradigm must be a successful demonstration that superpositions of cat states can be manipulated in order to realize single and multi-logical qubit gates. This goal is currently a substantial focus of current research and indeed impressive progress has already been made [Leghtas et al., 2015a; Wang et al., 2016] in demonstrating key aspects of the proposal in [Mirrahimi et al., 2014]. For the remainder of this work we will concentrate primarily on demonstrating that information stored in the quantum state in the cavity can be extracted efficiently and furthermore used to enhance the lifetime of quantum information. Interspersed throughout the experimental demonstrations and discussions thereof, however, we will remark upon a number of capabilities we already have at our disposal to create and manipulate states in a cavity, thereby primarily introducing some key concepts that future works will tackle in a rigorous way. In so doing, we hope to present a cohesive story on the current capabilities we have in realizing QEC systems and the necessary future steps for moving towards fault-tolerant architectures.

4 – Putting the Bell on Schrödinger’s Cat

The cat code offers a scheme to store quantum information in superpositions of coherent states and to protect the system from decoherence through the measurement of a single error syndrome with a single ancilla qubit. The success of this proposal hinges on the expectation that redundantly encoding a quantum bit in a cavity can be done with high fidelity and that extracting information (e.g. photon number parity) from the state can be accomplished efficiently. As this cavity is intentionally decoupled from every component of the system aside from the ancilla, at our disposal we essentially have just the dispersive interaction, gates on the ancilla, and the application of coherent displacements as the only tools with which to demonstrate these necessary capabilities. Using such means, the work on deterministically mapping a quantum bit onto cat states through a protocol called qcMAP [Leghtas et al., 2013; Vlastakis et al., 2013] quantified the fidelity of a possible encoding protocol.

In this chapter, however, we would like to go further and characterize the entanglement between the ancilla and a logical qubit comprised of coherent states with various entanglement metrics, or witnesses, that include a CHSH version [Clauser et al., 1969] of the Bell inequality [Bell, 1964]. We will assess the extent to which the logical state continuously deviates from the code-space and how efficiently we can measure the density matrix of the joint ancilla-cavity system. Furthermore, by using coherent states in this composite system, we will show the effects of decoherence by continuously varying the size of prepared entangled states [Brune et al., 1996], something unachievable with discrete systems. The degree to which the resulting data conforms with the assumption that single photon loss is the dominant error channel will dictate future strategies for implementing a quantum error correction system with this architecture. The techniques we propose here provide an important set of analytical tools for quantum systems comprised of entangled physical qubits and cavities [Deléglise et al., 2008; Hofheinz et al., 2009; Jeong and Kim, 2002; Leibfried et al., 1996; Sun et al., 2013], and serve to benchmark the suitability of storing quantum information redundantly in cavities as opposed to large collections of coupled two-level systems.

4.1 The Bell-cat

In this chapter we focus on the simple logical encoding where the basis states along the logical Z axis are defined by coherent states of opposite phase, $|\alpha\rangle$ and $|- \alpha\rangle$, introduced in sec. 3.3 and shown again in Fig. 4.1a. This logical qubit has the following Pauli operators:

$$\begin{aligned}
 X_s &= |-\alpha\rangle \langle \alpha| + |\alpha\rangle \langle -\alpha| \\
 Y_s &= j |-\alpha\rangle \langle \alpha| - j |\alpha\rangle \langle -\alpha| \\
 Z_s &= |\alpha\rangle \langle \alpha| - |-\alpha\rangle \langle -\alpha| \\
 I_s &= |\alpha\rangle \langle \alpha| + |-\alpha\rangle \langle -\alpha|.
 \end{aligned} \tag{4.1}$$

Before moving on to the encoding and measurement schemes, we must first confront the issue of non-orthogonality within the code when the overlap between $|\alpha\rangle$ and $|- \alpha\rangle$ is non-negligible. In this regime the system deviates from the two-level system approximation, resulting in a loss of fidelity of any encoded information. Seeking to identify a figure of merit for sufficient basis orthogonality, we can calculate the maximum Von-Neumann entropy of the encoded space to determine its capacity to store information:

$$\begin{aligned}
 S &= -\text{Tr} [\rho_{\max} \log_2 \rho_{\max}] \\
 &= -\sum_i \eta_i \log_2 \eta_i
 \end{aligned} \tag{4.2}$$

where $\rho_{\max} = \frac{1}{2}(|\alpha\rangle \langle \alpha| + |-\alpha\rangle \langle -\alpha|)$ is the density matrix for a complete mixture of the logical subspace and η is its set of eigenvalues. Rewriting ρ_{\max} in the even/odd cat state basis:

$$\rho_{\max} = \frac{1}{2}(1 + e^{-2|\alpha|^2}) |E\rangle \langle E| + \frac{1}{2}(1 - e^{-2|\alpha|^2}) |O\rangle \langle O| \tag{4.3}$$

where $|E\rangle, |O\rangle = \frac{1}{\sqrt{2(1 \pm e^{-2|\alpha|^2})}}(|\alpha\rangle \pm |-\alpha\rangle)$. Recall that $\langle E|O\rangle = 0$ for all coherent state amplitudes α . This gives the following entropy relation:

$$\begin{aligned}
 S &= -\frac{1}{2}(1 + e^{-2|\alpha|^2}) \log_2 \left(\frac{1}{2}(1 + e^{-2|\alpha|^2}) \right) \\
 &\quad - \frac{1}{2}(1 - e^{-2|\alpha|^2}) \log_2 \left(\frac{1}{2}(1 - e^{-2|\alpha|^2}) \right)
 \end{aligned} \tag{4.4}$$

The entropy varies from zero bits to a value asymptotically approaching a single bit with increasing coherent state amplitudes α , as seen in Fig. 4.1b. The orthogonality between logical states

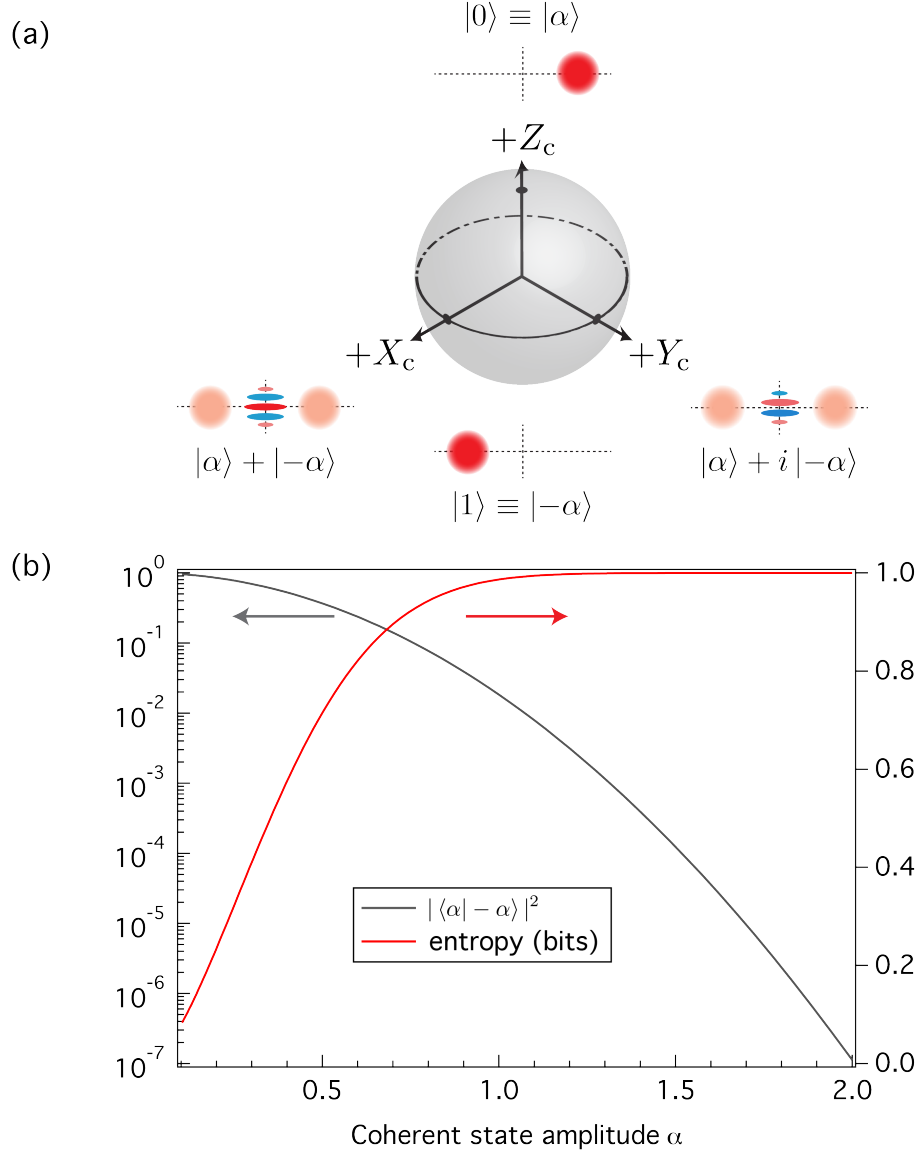


Figure 4.1: Logical encoding using coherent states. (a) The space spanned by the superposition of quasi-orthogonal coherent states $|\alpha\rangle$, $|- \alpha\rangle$ constitutes an encoded quantum bit in the cavity. While the cavity state can be represented by its Wigner function, this logical state is also described by a vector within its encoded Bloch sphere. Shown is the logical qubit Bloch sphere denoting the $+X_c$, $+Y_c$, and $+Z_c$ encoded states; a diagram of the cavity Wigner function accompanies each of these three states. (b) The capacity to store information in a cat state is determined by the orthogonality of its logical states $|\alpha\rangle$, $|- \alpha\rangle$. Shown is a comparison between the coherent state overlap (gray, left axis) and the maximum Von Neumann entropy Eq. 4.4 (red, right axis) for this logical encoding. Notice that the entropy rapidly approaches one bit for $\alpha > 1$, ensuring that information can be reliably encoded into the coherent states with manageable separations.

$|\langle \alpha | -\alpha \rangle|^2$ is directly related to this information capacity and serves as a proxy for validating the two-level system approximation of the cavity state. We see that for amplitudes $\alpha \gtrsim 1$, the entropy is very nearly equal to one, indicating that one quantum bit of information can be stored in encodings

of at least this size. This conclusion is of course consistent with the numerical result presented in Fig. 3.10.

Having decided upon the minimal size of a reliable logical encoding, the next step is to actually create the entangled ancilla-cavity state we aim to characterize, called the Bell-cat. The experiment utilizes the same setup introduced in chapter 2, namely a circuit QED architecture [Paik et al., 2011; Wallraff et al., 2004] consisting of two waveguide cavities coupled to a single transmon ancilla [Sun et al., 2013; Vlastakis et al., 2013], where the hardware is identical to that shown in Fig. 2.4. One long-lived cavity, with a relaxation time $\tau_s = 55 \mu\text{s}$, is used for quantum information storage, while the second cavity, strongly over-coupled to a 50Ω transmission line to allow fast field decay ($\tau_r = 30 \text{ ns}$) into the measurement chain is used to realize repeated projective measurements of the ancilla. The transmon ancilla, with relaxation and decoherence times $T_1, T_2 \approx 10 \mu\text{s}$ is coupled to both cavity modes and mediates entanglement and measurement of the state in the storage. All modes have transition frequencies between 5–8 GHz and are off-resonantly coupled (see A.1 for more details). We use a Josephson bifurcation amplifier (JBA) [Vijay et al., 2009] in a double-pumped configuration in combination [Kamal et al., 2009; Murch et al., 2013] with a dispersive readout to perform repeated ancilla measurements with a detection fidelity of 98.0% at a minimum of 800 ns between measurements.

The storage cavity and ancilla modes are well-described by the dispersive Hamiltonian:

$$\hat{H}/\hbar = \tilde{\omega}_s \hat{a}_s^\dagger \hat{a}_s + (\tilde{\omega}_a - \chi_{sa} \hat{a}_s^\dagger \hat{a}_s) |e\rangle \langle e|, \quad (4.5)$$

where \hat{a}_s is the storage cavity ladder operator, $|e\rangle \langle e|$ is the excited state ancilla projector, $\tilde{\omega}_s, \tilde{\omega}_a$ are the storage cavity and ancilla transition frequencies, and χ_{sa} is the dispersive interaction strength between the two modes (1.4 MHz). This interaction creates a shift in the transition frequency of one mode dependent on the other's excitation number, enabling ancilla-cavity entanglement [Brune et al., 1992]. As described in Fig. 4.2a, the system is first prepared in a product state $|\psi\rangle = \frac{1}{\sqrt{2}}(|g\rangle + |e\rangle) \otimes |\alpha\rangle$, where $|g\rangle, |e\rangle$ are the ground and excited states of the ancilla and $|\alpha\rangle$ is a coherent state of the cavity mode. Under the dispersive interaction we allow the system to evolve for a time $t = \pi/\chi_{sa}$, creating a state that mirrors the form of an EPR pair [Einstein et al., 1935], or Bell state (e.g. $|\psi_B\rangle = \frac{1}{\sqrt{2}}(|gg\rangle + |ee\rangle)$):

$$|\psi_{BC}\rangle = \frac{1}{\sqrt{2}}(|g, \alpha\rangle + |e, -\alpha\rangle). \quad (4.6)$$

We call $|\psi_{BC}\rangle$ the Bell-cat [Brune et al., 1992; Vlastakis et al., 2013].

If one recalls a typical introductory quantum mechanics textbook, the ‘Schrödinger’s cat’ thought experiment highlights the counterintuitive concept of entanglement in macroscopically distinguishable systems: an atom in an excited state and a cat are in a box; if the atom decays, it uncorks a vial of poison that kills the cat, while if it does not decay, the cat lives. Does this mean that the state of the atom and the cat are actually entangled to the extent that they can violate a Bell inequality? Intuitively we of course know that this cannot be the case, as enormous objects such as cats are not well-represented quantum mechanically. Yet we also know that a loophole-free Bell test has been performed in NV centers to prove that entanglement beyond classical bounds can only be explained with a quantum physical theory [Hensen et al., 2015], so where do we draw the line? At what point does a quantum mechanical system devolve into a classical one?

The Bell-cat offers us the perfect platform to explore the answers to this question. The transmon of course plays the role of the atom, and the coherent state plays the role of the cat. We can make this coherent state as small or as large as we wish. For very small amplitudes α , when the coherent states are largely dominated by a few Fock states, with sufficient measurement efficiency we should expect (or hope) to witness entanglement. For very large α , we expect no entanglement. Additionally, we have the capability to explore the fascinating regime in which α is of some intermediate size – macroscopically distinguishable, yet still exhibiting signatures of quantum behavior. By continuously varying α , a capability unique to our system, we will show how the Bell-cat straddles the quantum-to-classical regime, and conclude that decoherence resolves the ‘Schrödinger’s cat’ paradox.

In the typical Bell state $|\psi_B\rangle$, quantum measurement back-action is a clear manifestation of the entanglement between the quantum bits, demonstrating correlations between measurement outcomes that can be reformulated into entanglement witness metrics. Likewise, after preparing $|\psi_{BC}\rangle$ we can measure along one of the three ancilla axes $M_a \in \{X_a, Y_a, Z_a\}$ and obtain one of two possible outcomes, either $|g\rangle$ or $|e\rangle$; due to measurement back-action, we expect the state in the cavity $|\psi_s\rangle$ to depend on the outcome of this measurement according to the following prescription (normalizations omitted):

| M_a | Pre-meas. state | $ \psi_s\rangle$ if result $ g\rangle$ | $ \psi_s\rangle$ if result $ e\rangle$ |
|-------|---|--|--|
| X_a | $ g\rangle(\alpha\rangle + \alpha\rangle) + e\rangle(\alpha\rangle - \alpha\rangle)$ | $+X_s$ | $-X_s$ |
| Y_a | $ g\rangle(\alpha\rangle - j \alpha\rangle) + e\rangle(\alpha\rangle + j \alpha\rangle)$ | $-Y_s$ | $+Y_s$ |
| Z_a | $ g\rangle \alpha\rangle + e\rangle -\alpha\rangle$ | $+Z_s$ | $-Z_s$ |

(4.7)

In the first step in the sequential detection scheme, the value of the ancilla tomography measurement (Fig. 4.2b) is recorded and the ancilla is reset to $|g\rangle$ using real-time feedback (see A.1). The displaced

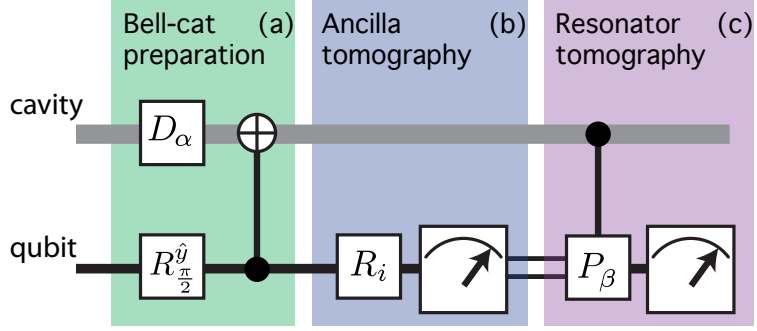


Figure 4.2: **Sequential detection for entanglement characterization.** (a) A quantum circuit outlines the method to prepare and measure entanglement between the ancilla and cavity state using sequential detection. State preparation is performed by first creating a product state $|\psi\rangle = \frac{1}{\sqrt{2}}(|g\rangle + |e\rangle) \otimes |\alpha\rangle$ with a cavity displacement $\hat{D}(\alpha)$ of amplitude α and an ancilla gate $R_{\hat{y}, \frac{\pi}{2}}$ corresponding to a $\frac{\pi}{2}$ rotation around the \hat{y} -axis. A conditional gate using the dispersive interaction produces the entangled state $|\psi_{BC}\rangle = \frac{1}{\sqrt{2}}(|g, \alpha\rangle + |e, -\alpha\rangle)$. (b), (c) Tomography is performed by measuring an observable of both the ancilla and cavity with sequential QND measurements. A pre-rotation R_i allows ancilla detection along one of three basis vectors X , Y , and Z . The ancilla is reset and a cavity observable \hat{P}_β is mapped to the ancilla for a subsequent measurement, where $\hat{P}_\beta = \hat{D}(\beta)\hat{P}\hat{D}(\beta)^\dagger$ is the displaced photon number parity operator. Sequential detections are binary results compared shot-by-shot to determine ancilla-cavity correlations.

photon-number parity observable \hat{P}_β of the cavity is subsequently mapped onto the ancilla using Ramsey interferometry [Bertet et al., 2002], followed by a second ancilla state detection (Fig. 4.2c); here, $\hat{P}_\beta = \hat{D}(\beta)\hat{P}\hat{D}(\beta)^\dagger$, where $\hat{D}(\beta)$ is the displacement operator ($\hat{D}(\beta)|0\rangle_f = |\beta\rangle$) and \hat{P} the photon number parity operator, the eigenvalue of which is detected with 95.5% fidelity (see A.1). The Wigner function $W(\beta) = \frac{2}{\pi}\langle\hat{P}_\beta\rangle$ is constructed from an ensemble of such measurements with different displacement amplitudes β .

Wigner tomography (Fig. 4.2c) demonstrates the quantum back-action. Indeed, we see the expected superpositions of coherent states in the cavity conditioned on the result of the ancilla measurement (Fig. 4.3)¹. Moreover, qualitatively these cavity states appear to largely remain in the code-space defined above (Fig. 4.1a) given the locations of the coherent states and the patterns of the interference fringes. With these results we are now in the position to benchmark the entanglement in the system by quantifying the correlations that we observe.

4.2 Joint Wigner tomography

As the density matrix of a cavity state can be completely described by direct measurements in the continuous-variable basis with the Wigner function [Lutterbach and Davidovich, 1997], we can

¹The method of using strong projective measurements for the create of cat states has been demonstrated in previous works [Deléglise et al., 2008] as well.

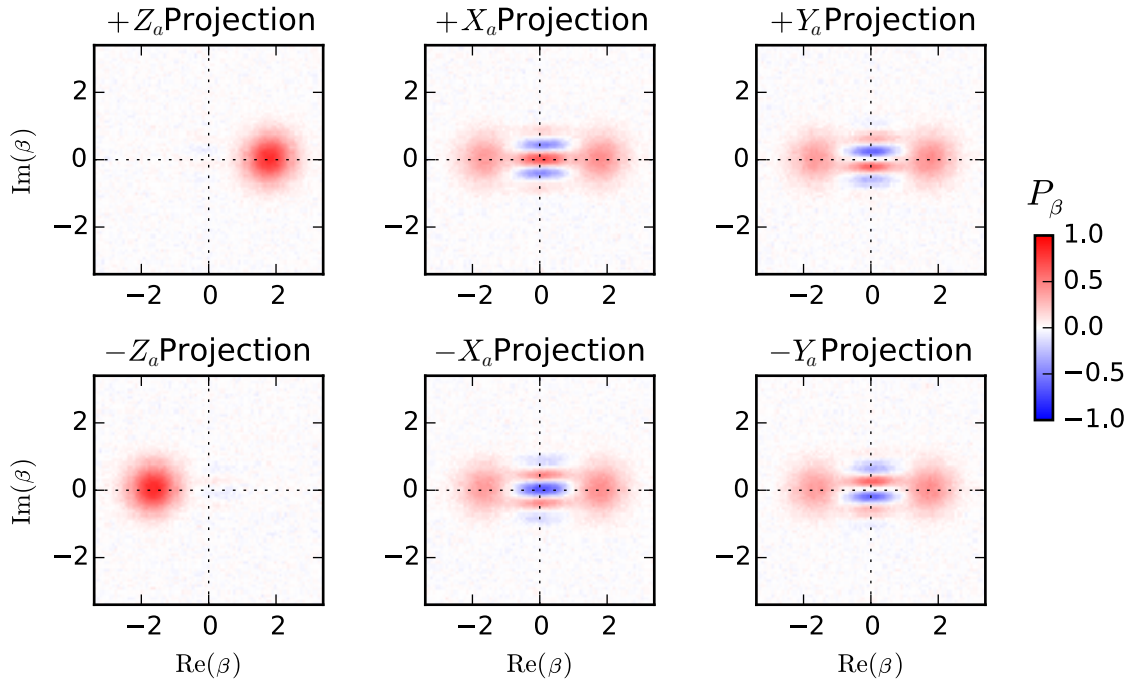


Figure 4.3: **Measurement back-action of a Bell-cat state.** Shown are the resulting projections of the cavity state when preparing $|\psi_{BC}\rangle$ and measuring the ancilla along one of its three axes $M_a \in \{X_a, Y_a, Z_a\}$. The initial coherent state amplitude is $\alpha = \sqrt{3}$. While a measurement along Z_a results in a projected coherent state with opposite phases $|\pm\alpha\rangle$, measuring along the X_a and Y_a axes results in a projected cat state each with different inference fringe phases. Combining these measurements with the probability to obtain each result allows us to construct the state of the entire system and is used to create the joint Wigner function representation in Fig. 4.4.

use the outcomes of the sequential detection scheme to express an entangled ancilla-cavity state in what we call the joint Wigner representation. We construct this representation by correlating ancilla observables $\hat{\sigma}_i \in \{\hat{J}_a, \hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z\}$ with the displaced photon number parity at every point in the cavity's phase space:

$$W_i(\beta) = \frac{2}{\pi} \langle \hat{\sigma}_i \hat{P}_\beta \rangle \quad (4.8)$$

These four distributions are a complete representation of the combined ancilla-cavity quantum state (Fig. 4.4a). While other representations exist for similar systems [Eichler et al., 2012; Jeong et al., 2014; LinPeng et al., 2013; Morin et al., 2014], $W_i(\beta)$ is directly measured with this detection scheme and does not require a density matrix reconstruction.

The efficiency of our detection scheme can be quantified with the visibility \mathcal{V} of the unnormalized joint Wigner measurements. We determine \mathcal{V} by tracing over both the ancilla and cavity states and

comparing this to its ideal value $\int W_I^{\text{ideal}}(\alpha) d^2\beta = 1$:

$$\begin{aligned} \mathcal{V} &= \int W_I(\beta) d^2\beta \\ &= \frac{2}{\pi} \int \langle \hat{I} \hat{P}_\beta \rangle d^2\beta \\ &= 0.85 \pm 0.01. \end{aligned} \tag{4.9}$$

The visibility is primarily limited by measurement fidelity and ancilla decoherence between detection events. We postpone the discussion of ancilla decoherence to later chapters in which its effects on quantum error correction are described in detail.

We are also interested in finding the fidelity of the ancilla-cavity density matrix to the ideal Bell-cat. To do so, we first note that an ancilla operator A can be written in the Pauli basis as $A = \sum_i A_i \hat{\sigma}_i$ where $A_i = \frac{1}{2} \text{Tr}[A \hat{\sigma}_i]$. Likewise, a bounded cavity observable can be represented in the continuous variables basis as $B = \frac{1}{\pi} \int B(\beta) \hat{P}_\beta d^2\beta$, where $B(\beta) = \text{Tr}[B \hat{P}_\beta]$ [Cahill and Glauber, 1969]. Finally, the composite ancilla-cavity density matrix can be written as:

$$\rho = \pi \sum_i \int W_i(\beta) \hat{\sigma}_i \hat{P}_\beta d^2\beta \tag{4.10}$$

Note that for separable states $\rho = \rho_a \otimes \rho_s$, this relation can be split up into discrete and continuous components:

$$\rho = \frac{1}{2} \sum_i \text{Tr}[\rho_a \hat{\sigma}_i] \hat{\sigma}_i \otimes 2\pi \int \frac{2}{\pi} \text{Tr}[\rho_s \hat{P}_\beta] \hat{P}_\beta d^2\beta \tag{4.11}$$

For any state ρ , we can write the mean value of an observable for the combined system with the following relation:

$$\begin{aligned} \langle AB \rangle &= \text{Tr}[AB\rho] \\ &= \frac{1}{2} \text{Tr} \left[\sum_{i,j} \int A_i B(\beta) W_j(\beta') \hat{\sigma}_i \hat{\sigma}_j \hat{P}_\beta \hat{P}_{\beta'} d^2\beta d^2\beta' \right] \end{aligned} \tag{4.12}$$

Using the following operator rules $\text{Tr}[\hat{\sigma}_i \hat{\sigma}_j] = 2\delta_{ij}$ and $\text{Tr}[\hat{P}_\beta \hat{P}_{\beta'}] = \delta^2(\beta - \beta')$ we can simplify Eq. 4.12:

$$\langle AB \rangle = \sum_i \int A_i B(\beta) W_i(\beta) d^2\beta \tag{4.13}$$

The overlap integral used in this calculation is similar to descriptions of the standard Wigner function [Cahill and Glauber, 1969; Haroche and Raimond, 2006]. We use the result of the derivation above to determine the fidelity to a target state $\mathcal{F} = \text{Tr}[\rho_{BC}\rho] = \frac{\pi}{2} \sum_i \int W_i^{BC}(\beta)W_i(\beta)d^2\beta$, where $W_i^{BC}(\beta)$ are the joint Wigner functions of the ideal state $\rho_{BC} = |\psi_{BC}\rangle\langle\psi_{BC}|$ and $W_i(\beta)$ are the measured joint Wigner functions (normalized by \mathcal{V}). We obtain a state fidelity $\mathcal{F} = 87 \pm 2\%$ for a displacement amplitude $\alpha = \sqrt{3}$. This amplitude was chosen not just to ensure orthogonality between logical states ($|\langle\alpha|-\alpha\rangle|^2 = 6 \times 10^{-5} \ll 1$), but also to highlight the interference fringes in the Wigner tomograms, which increase in frequency with larger amplitudes [Haroche and Raimond, 2006].

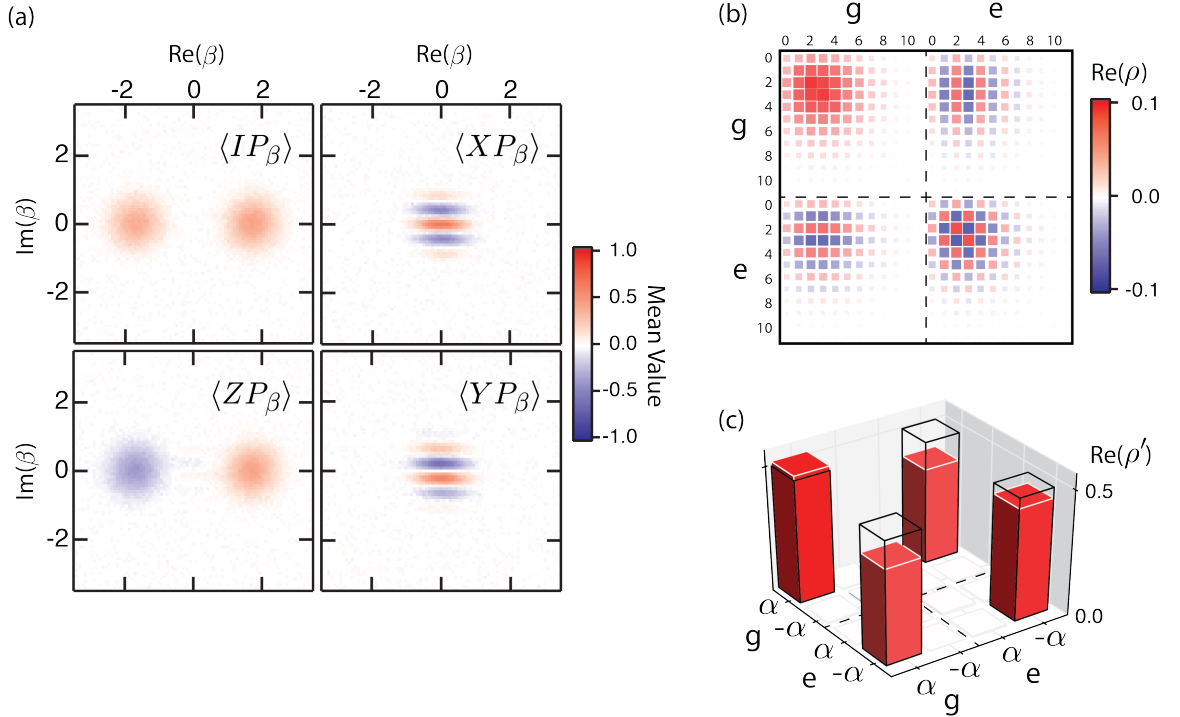


Figure 4.4: Joint Wigner tomography of a Bell-cat state. (a) The set of joint Wigner functions $W_i(\beta) = \frac{2}{\pi} \langle \hat{\sigma}_i \hat{P}_\beta \rangle$ represents the state of an ancilla-cavity system with correlations between the ancilla observables $\hat{\sigma}_i = \{\hat{I}_a, \hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z\}$ and cavity observable \hat{P}_β reported for a state $|\psi_{BC}\rangle$ and displacement amplitude $\alpha = \sqrt{3}$. Shown are measurements comprised of four panels $\langle I_a \hat{P}_\beta \rangle$, $\langle X_a \hat{P}_\beta \rangle$, $\langle Y_a \hat{P}_\beta \rangle$, $\langle Z_a \hat{P}_\beta \rangle$ of 6500 correlations each between the ancilla and cavity states. Interference fringes in $\langle X_a \hat{P}_\beta \rangle$ and $\langle Y_a \hat{P}_\beta \rangle$ reveal quantum coherence in the entangled state. (b) From the set of joint Wigner functions we perform a density matrix reconstruction to show the combined ancilla-cavity state ρ in the Fock state basis. (c) Projecting ρ onto the logical basis $\Phi = |\alpha\rangle\langle\alpha| + |-\alpha\rangle\langle-\alpha|$ produces the reduced, unnormalized density matrix $\rho' = \Phi\rho\Phi^\dagger$ in the form of a traditional Bell state. The reduction in contrast of the off-diagonal components in ρ' is due to decoherence in the physical system during preparation and measurement.

In Fig. 4.4b, we show the reconstructed density matrix of a target Bell-cat state. We perform this reconstruction with a-priori assumptions that the cavity state is truncated to twelve occupied photon number states $N_{\max} = 12$, the resulting noise of each averaged measurement is Gaussian-distributed (Fig. 4.5), and the reconstructed density matrix is positive semidefinite with trace equal

to one. Under these constraints, we perform a least squares regression using a Maximum likelihood estimation [Smolin et al., 2012]. We can go one step further and project ρ onto our logical basis $\Phi = |\alpha\rangle\langle\alpha| + |-\alpha\rangle\langle-\alpha|$. We obtain an unnormalized matrix $\rho' = \Phi\rho\Phi^\dagger$, which clearly resembles the density matrix of a Bell state $\rho_B = |\psi_B\rangle\langle\phi_B|$ (Fig. 4.4c).

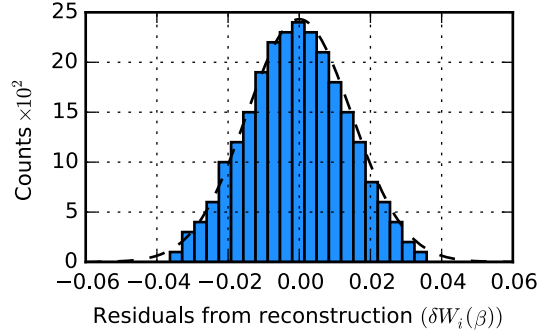


Figure 4.5: Histogram of reconstruction residuals. Plotted are the residuals corresponding to the density matrix reconstruction of the Bell-cat state shown in Fig. 4.4. This Histogram shows the distribution of the 25,000 residuals from the joint Wigner function which gives a Gaussian distribution (mean value $\mu = 7.0 \times 10^{-4}$, standard deviation $\sigma = 0.015$), which agree with our expectation for statistical error $\sigma_{\text{est}} = \frac{1}{\sqrt{N}} \approx 0.015$.

These results demonstrate the possibility of understanding the joint entangled state of the ancilla and cavity in a simple four dimensional, rather than infinite dimensional, space. Such a dramatic simplification is exactly what we require in looking to encode a quantum bit in a cavity. Although we work with a continuous variables system, we have nonetheless shown that one can simply understand the logical encoding in coherent states of opposite phase as just another two-level system, much like the ancilla with which it is entangled. A key difference, however, is that the trace of ρ' does not equal 1. Rather than implying that something in our analysis has led to an unphysical state, this discrepancy instead reveals that the composite system must be afflicted by code-space leakage, a result of decoherence during the conditional gate in particular². We will return to this important point in later chapters. Having quantified what we surmised earlier and see in Fig. 4.3, that the cavity state primarily remains in the code space after ancilla measurement back-action, we learn the following key piece of information: the majority of the information provided by the Wigner tomogram is in fact concentrated only in a few locations in the cavity's phase space: around the origin and at $|\alpha\rangle$ and $|-\alpha\rangle$; these are the points we must focus on measuring with high fidelity in order to extract the necessary information about the entangled Bell-cat state.

²Code-space leakage is of course a problem in any experimental implementation of a two-level system; in the case of the transmon, this takes the form of excitations to higher states (e.g. $|f\rangle$).

4.3 Efficient measurement of cavity observables

The number of measurement settings required to perform cavity state tomography can be resource intensive. State tomography of the ancilla requires just four pre-rotations followed by a projective measurement. With the cavity, on the other hand, we have thus far required a large grid of photon number parity measurements after varying displacements to construct the full Wigner tomogram. Each tomogram, however, clearly contains useful information in only a small portion of phase space, as seen in Fig. 4.3 where the value of $\langle \hat{P}_\beta \rangle$ is approximately zero outside of the regions near $\beta = \{0, \alpha, -\alpha\}$. This suggests that complete cavity state tomography need not be required, and in fact many fewer measurements could be used to characterize a state when operating in a smaller, encoded subspace.

Restricting ourselves to the logical qubit subspace shown in Fig. 4.1, in fact only four values of the cavity Wigner function $W(\beta)$ are required to reconstruct the state, an example of direct fidelity estimation (DFE) [da Silva et al., 2011; Flammia and Liu, 2011]. For large cat states $|\langle \alpha | -\alpha \rangle|^2 \ll 1$, the encoded state observables map to cavity observables as:

$$\begin{aligned} X_s &= \hat{P}_0 & I_s &= \hat{P}_\alpha + \hat{P}_{-\alpha} \\ Y_s &= \hat{P}_{\frac{j\pi}{8\alpha}} & Z_s &= \hat{P}_\alpha - \hat{P}_{-\alpha} \end{aligned} \quad (4.14)$$

where $\{I_s, X_s, Y_s, Z_s\}$ form the Pauli set for the logical qubit in the cavity. We justify such a choice by writing down the following photon-number parity \hat{P} relations:

$$\begin{aligned} \langle \alpha | \hat{P}_0 | \alpha \rangle &= \langle \alpha | -\alpha \rangle \ll 1 \\ \langle \alpha | \hat{P}_0 | -\alpha \rangle &= \langle \alpha | \alpha \rangle = 1 \\ \langle \alpha | \hat{P}_\alpha | \alpha \rangle &= \langle \alpha - \alpha | \alpha - \alpha \rangle \ll 1 \\ \langle \alpha | \hat{P}_\alpha | -\alpha \rangle &= e^{2(\alpha\alpha^* - \alpha^*\alpha)} \langle \alpha | -\alpha \rangle, \end{aligned} \quad (4.15)$$

where again $\hat{P}_\alpha = \hat{D}(\alpha)\hat{P}\hat{D}(\alpha)^\dagger$ for some displacement amplitude α and $|\langle \alpha | -\alpha \rangle| \ll 1$. Now taking the projector $M = |\alpha\rangle\langle\alpha| + |-\alpha\rangle\langle-\alpha|$, we derive the encoded state's Pauli operators from the cavity

state observables reported in Eq. 4.15:

$$\begin{aligned}
M\hat{P}_0M^\dagger &\approx |-\alpha\rangle\langle\alpha| + |\alpha\rangle\langle-\alpha| \\
M\hat{P}_\alpha M^\dagger &\approx |\alpha\rangle\langle\alpha| \\
M\hat{P}_{-\alpha}M^\dagger &\approx |-\alpha\rangle\langle-\alpha| \\
M\hat{P}_{\frac{j\pi}{8\alpha}}M^\dagger &\approx j|-\alpha\rangle\langle\alpha| - j|\alpha\rangle\langle-\alpha|
\end{aligned} \tag{4.16}$$

Putting these relationships together builds the encoded state observables $\{I_s, X_s, Y_s, Z_s\}$ and reveals that these observables can be efficiently measured using Wigner tomography. Observables I_s and Z_s require a comparison between two different observables. For true single-shot readout of these logical observable Z_s , measuring a single value in the cavity state Husimi-Q distribution $Q(\alpha) = \frac{1}{\pi}\langle\alpha|\rho|\alpha\rangle$ [Haroche and Raimond, 2006] can be employed where $Z_s = 2\pi Q(\alpha) - 1$. This is being explored in future experiments.

Cuts in the joint Wigner function (Fig. 4.6) show these observables and their correlations to the ancilla as a function of cat state size. As the superposition state is made larger, interference fringe oscillations increase while fringe amplitude decreases due to photon loss. For a state $|\psi_{BC}\rangle$ with $|\alpha| = \sqrt{3}$, we estimate the direct fidelity \mathcal{F}_{DFE} :

$$\begin{aligned}
\mathcal{F}_{\text{DFE}} &= \frac{1}{4}(\langle I_a I_s \rangle + \langle X_a X_s \rangle - \langle Y_a Y_s \rangle + \langle Z_a Z_s \rangle) \\
&= 0.72 \pm 0.02,
\end{aligned} \tag{4.17}$$

putting a fidelity bound on the target state with no corrections for visibility (Fig. 4.7). This estimate is related to the benchmarks reported above $\mathcal{F}_{\text{DFE}} \approx \mathcal{V} \times \mathcal{F}$ and far surpasses the 50% threshold for a classically correlated state. This indicates both high fidelity state-preparation and measurement, and demonstrates that strong correlations are directly detectable using joint Wigner tomography.

We see that by choosing an encoding scheme where states of a quantum bit are mapped onto a superposition of coherent states $|\alpha\rangle$ and $|-\alpha\rangle$, we can condense the joint Wigner representation down to just sixteen correlations, equivalent to a two-quantum bit measurement set that constitutes the permutation of $\{I, X, Y, Z\}$ for each quantum bit. We can determine this Pauli set from the reconstructed ancilla-cavity state by projecting onto the encoded basis of $\{I_s, X_s, Y_s, Z_s\}$. Fig. 4.8 shows the resulting two-quantum bit Pauli set for the ancilla and an logical qubit in the cavity mode, a variant of the reduced density matrix representation shown in Fig. 4.4. These results highlight our

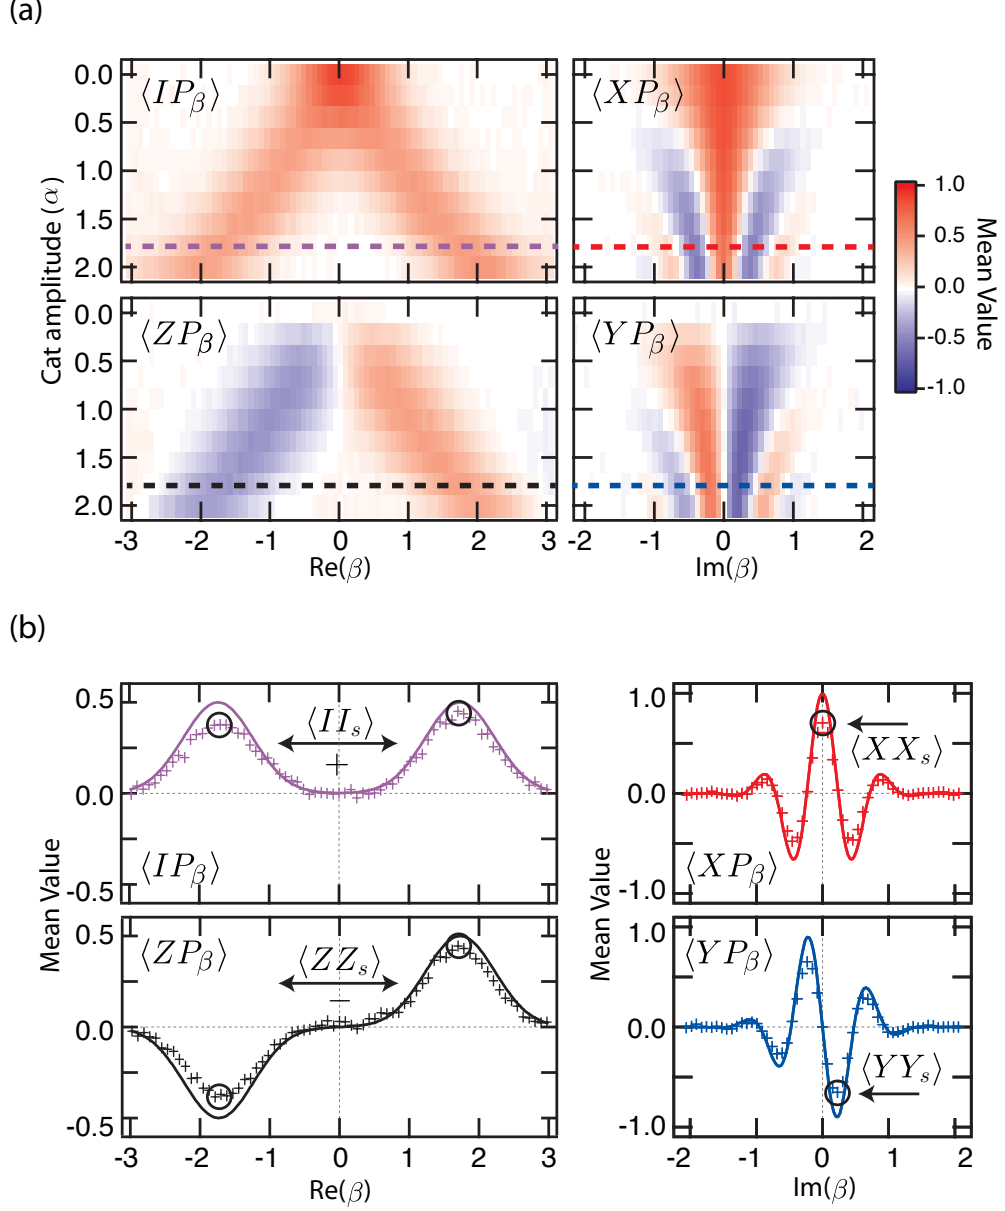


Figure 4.6: **Ancilla-cavity correlations.** (a) Correlations are measured for entangled states $|\psi_{BC}\rangle$ with cat state amplitudes ranging from $\alpha = 0$ to 2. Cuts in joint Wigner functions $\langle IP_\beta \rangle$ and $\langle ZP_\beta \rangle$ at $\text{Im}(\beta) = 0$ show the increasing separation of the coherent state superpositions. Cuts in the joint Wigner functions $\langle XP_\beta \rangle$ and $\langle YP_\beta \rangle$ at $\text{Re}(\beta) = 0$ reveal the interference fringe oscillations dependence on cat state size, which increase in frequency with increasing cat state amplitude. (b) By viewing just single cuts at $\alpha = \sqrt{3}$, we see single-shot correlations (crosses) as compared to what is expected from an ideal system with perfect preparation and measurement (solid line). From the cuts we see the individual measurement settings used to determine joint encoded observables $\{II_s, XX_s, YY_s, ZZ_s\}$. While $\langle XX_s \rangle$ and $\langle YY_s \rangle$ can be determined from a single measurement setting, $\langle II_s \rangle$ and $\langle ZZ_s \rangle$ are determined from the sum and difference of two different settings. From these four correlations we immediately find a fidelity to an entangled state $\mathcal{F}_{\text{DFE}} = \frac{1}{4}(\langle II_s \rangle + \langle XX_s \rangle - \langle YY_s \rangle + \langle ZZ_s \rangle) = (72 \pm 2)\%$.

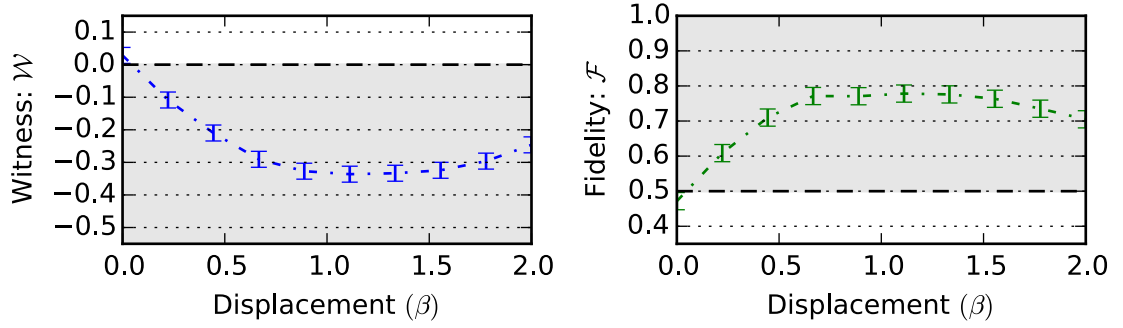


Figure 4.7: **Entanglement witnesses with cat states.** An entanglement witness and direct fidelity estimation (DFE) are determined by measuring four correlations between the ancilla and logical qubit. (a) The entanglement witness $\mathcal{W} = II - ZZ - XX + YY$ shows entanglement for all negative values (gray shading). (b) DFE to a target Bell state $\mathcal{F} = II + XX - YY + ZZ$ is also shown where entanglement can be confirmed for values above $\mathcal{F} > 0.5$. Error bars show the standard deviation due to random error limited by the total number of samples N taken at each displacement ($N > 4000$).




Figure 4.8: **Reconstructed Pauli sets.** The set of sixteen joint Pauli operators spans the Hilbert space of the ancilla and the logical qubit. Shown is the Pauli set for the entangled target state $|\psi_B\rangle$ obtained in two ways. (Red) is the reconstructed Pauli set using a density matrix reconstruction of the full quantum state with no normalization constraint, then projecting onto the encoded subspace. (Blue) shows the values discerned from an overlap integral of the measured joint-Wigner functions (Eq. 4.12). These measurements agree with each other within statistical errors.

ability to recast the state encoded in the cavity as one that has a small, simple set of observables that directly mirrors that of a physical qubit.

In characterizing the entanglement between the ancilla and the cavity, the compression of the cavity's phase space down to just four points demonstrates the efficiency with which information can be encoded within and extracted from such a large Hilbert space. These two points are crucial to the success of the cat code, and once more highlight the advantages of a redundant encoding in

superpositions of coherent states. Indeed, recalling the example of the bit-flip code (Fig. 3.2) or the Steane code (Fig. 3.4), creating an analog of the Bell-cat and benchmarking the entanglement between code and ancillae would require many more operations and tomography measurements.

With the ability to extract correlations in a single-shot way, we can in fact go one step further and test to what extent, if any, the Bell-cat violates the Bell inequality, which sets the most stringent bound for benchmarking entanglement. In so doing, we will be able to directly evaluate the effect of code size on the levels of entanglement we can detect by continuously varying the size of the cat states. This will allow us to identify a regime in which we both maintain the orthogonality between the coherent states in the logical basis and observe a Bell inequality violation, an important piece of information we will use when designing an error correction system that uses the cat code.

4.4 Violating Bell’s inequality

The hallmark of entanglement is the detection of strong correlations between systems, most starkly demonstrated by the violation of a Bell inequality. Although originally proposed to investigate local hidden variable theory [Bell, 1987], a Bell inequality can be used to benchmark the ability to entangle and extract information from an entangled two-quantum bit system [van Enk et al., 2007]. Using the Clauser-Horne-Shimony-Holt (CHSH) variant [Clauser et al., 1969] of the Bell test, this violation has been demonstrated with photons [Aspect et al., 1981; Freedman and Clauser, 1972], atoms [Hofmann et al., 2012; Rowe et al., 2001], solid-state spins [Hensen et al., 2015; Pfaff et al., 2012], and artificial atoms in superconducting circuits [Ansmann et al., 2009; Chow et al., 2010b]. Here, we use the CHSH formulation of a Bell test to characterize entanglement between the ancilla and a cat state. Although Bell tests using homodyne measurements have been proposed [Gilchrist et al., 1998; Leonhardt and Vaccaro, 1995], here we instead choose the CHSH Bell test that states that the sum of four classical correlations will be bounded such that:

$$-2 \leq \mathcal{O} = \langle A_a A_s \rangle + \langle A_a B_s \rangle - \langle B_a A_s \rangle + \langle B_a B_s \rangle \leq 2 \quad (4.18)$$

where, in this experiment, A_a, B_a are two ancilla observables and A_s, B_s are two cavity observables. We perform two Bell tests (Fig. 4.9) with correlations taken shot-by-shot with no post-selection or compensation for detector inefficiencies.

Test 1

In the first test we choose cavity observables Z_s, X_s and ancilla observables $Z_a(\theta), X_a(\theta)$ where:

$$Z_a(\theta) = Z_a \cos \frac{\theta}{2} - X_a \sin \frac{\theta}{2} \quad X_a(\theta) = X_a \cos \frac{\theta}{2} + Z_a \sin \frac{\theta}{2} \quad (4.19)$$

The angle θ corresponds to a rotation of the ancilla state before detection. In Fig. 4.9a, we plot \mathcal{O} for each of the four permutations of the joint observables and find a maximum Bell violation for an angle $\theta = -\pi/4$ as expected [Clauser et al., 1969] with the observables:

$$\begin{aligned} A_a &= \frac{X_a + Z_a}{\sqrt{2}}; & B_a &= \frac{X_a - Z_a}{\sqrt{2}} \\ A_s &= Z_s; & B_s &= X_s \end{aligned} \quad (4.20)$$

We investigate this system's susceptibility to decoherence by continuously increasing the cat state amplitude α and observe a Bell signal with a maximal value $\mathcal{O}_1 = 2.30 \pm 0.04$ for $\alpha = 1$ (Fig. 4.9b). We see a Bell signal surpassing bounded values up to cat states of size $|\alpha - (-\alpha)|^2 = 16$ photons [Deléglise et al., 2008; Vlastakis et al., 2013].

We can model the effects of photon loss and measurement inefficiency on the maximum violation, as shown in Fig. 4.9c. For the ideal case without decoherence, an overlap of the coherent state superposition decreases the contrast in $\langle A_a Z_s \rangle$ and $\langle B_a Z_s \rangle$, limiting the maximum Bell signal following the relation:

$$\mathcal{O}_{\text{ideal}} = \sqrt{2}(2 - e^{-8|\alpha|^2})$$

Now first introducing measurement inefficiency, the contrast of this maximum Bell signal is expected to scale with the visibility \mathcal{V} as:

$$\mathcal{O}_{\text{vis}} = \sqrt{2}\mathcal{V}(2 - e^{-8|\alpha|^2})$$

Next, photon loss will also have an effect on the maximum Bell signal by reducing the measured contrast of all correlations for $\langle A_a X_s \rangle$ and $\langle B_a X_s \rangle$. This produces an amplitude-dependent maximum Bell Signal:

$$\mathcal{O}_{\text{loss}} = \sqrt{2}(1 - e^{-8|\alpha|^2} - e^{-2|\alpha|^2 t_{eff}/\tau_s})$$

where τ_s is the photon decay time constant and t_{eff} is the effective time to create and measure the Bell-cat state. Finally taking into account both visibility and photon loss produces the expected maximum Bell signal for this first test:

$$\mathcal{O}_{\text{pred}} = \sqrt{2}\mathcal{V}(1 - e^{-8|\alpha|^2} - e^{-2t_{eff}/\tau_s|\alpha|^2})$$

This predicted Bell signal is shown in Fig. 4.9b using the measured joint-Wigner contrast $\mathcal{V} = 0.85$ and time between cavity state creation and detection $t_{eff} = 1.24\mu s$. As is evident, the agreement between the data and this simple analytical treatment is excellent.

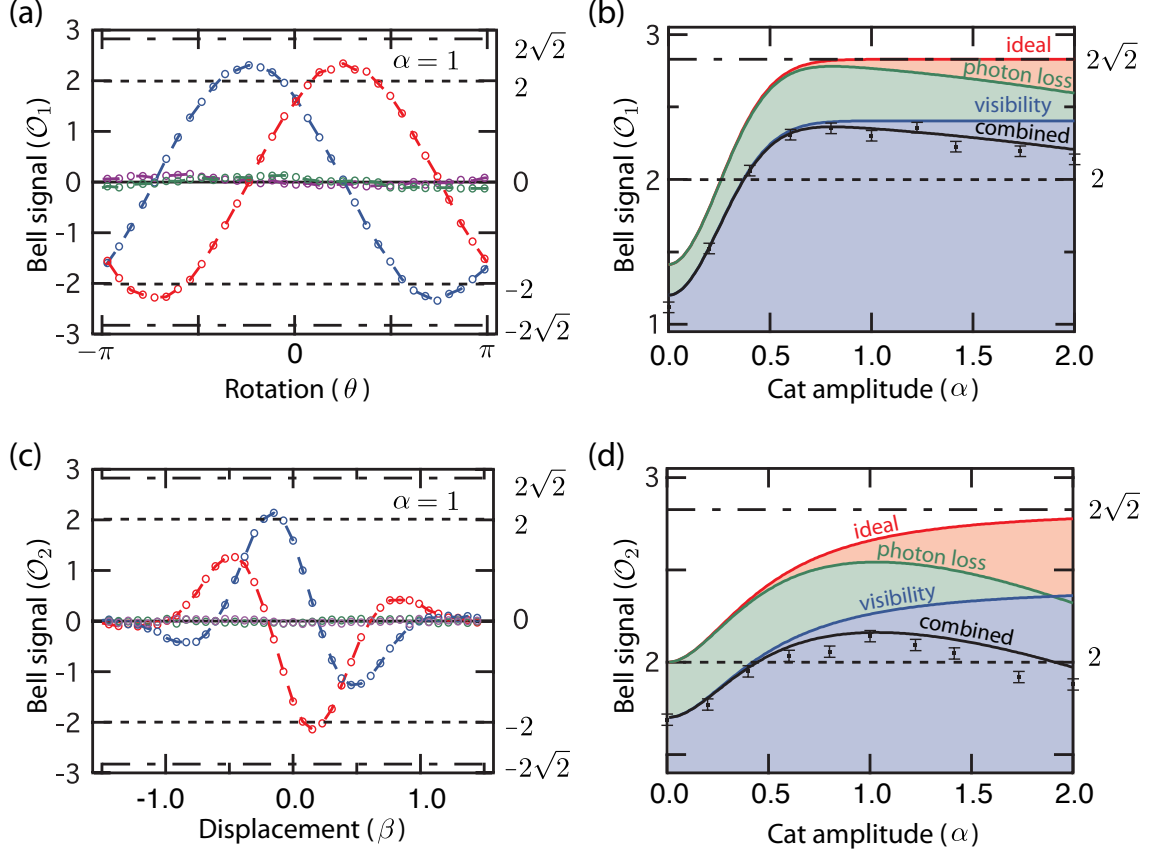


Figure 4.9: Bell tests with a cat state. A CHSH Bell test between an ancilla and cavity is the sum of four correlations $\mathcal{O} = \langle AA_s \rangle + \langle AB_s \rangle - \langle BA_s \rangle + \langle BB_s \rangle$, where A, B are observables of the ancilla and A_s, B_s are observables of the cavity. **(a)** We use correlations between ancilla state observables $X(\theta) = X \cos(\theta/2) + Z \sin(\theta/2)$ and $Z(\theta) = Z \cos(\theta/2) - X \sin(\theta/2)$ and encoded state observables X_s, Z_s to perform a CHSH Bell test as a function of ancilla detector angle θ . Shown in (a) are four traces that are the result of every possible combination of X, Z, X_s, Z_s . A maximum Bell signal is found at $\theta = -\frac{\pi}{4}$. **(b)** We report this maximum Bell signal for different cat state amplitudes α . Plotted points (black) are the average Bell signal for a given amplitude and show the dependence of the entangled state with photon loss and detector visibility. Error bars denote the standard deviation of the average signal due to random error as a consequence of a limited sample size ($N = 4,000$). Solid lines describe the predicted trends given the measured cavity decay rate and detection visibility. While the ideal behavior (red) for an entangled state approaches $\mathcal{O} = 2\sqrt{2}$, photon loss (green), detector visibility (blue), and their combined effects (black) ultimately limit the maximum Bell signal achieved. **(c), (d)** Furthermore, we realize a second Bell test using ancilla observables X, Y and cavity state observables $X_s(\alpha) = \hat{D}(j\alpha)X_s\hat{D}(j\alpha)^\dagger$, $Y_s(\alpha) = \hat{D}(j\alpha)Y_s\hat{D}(j\alpha)^\dagger$, where α corresponds to a tomography displacement amplitude serving as a rotation of the effective cavity detector angle. There is a mismatch in the maxima obtained in the two different Bell tests due to increased susceptibility to photon loss in the second test. Both, however, show a violation at least four standard deviations beyond the classical limit defined by the CHSH Bell inequality.

Test 2

Measurements along Z_s require assumptions of the symmetry of the prepared state. The Bell violations demonstrated with the first test are not single-shot, strictly speaking, as a difference between the signals at $\beta = \alpha$ and $\beta = -\alpha$ is used to calculate Z_s . The CHSH inequality does not constrain one to pick specific observables, however, and so we can instead employ an alternative Bell test (using a scheme similar to Ref. [Park et al., 2012]), in which we choose ancilla observables X_a, Y_a and cavity observables $X_s(\beta), Y_s(\beta)$:

$$\begin{aligned} X_s(\beta) &= \hat{D}(j\beta)\hat{P}_0\hat{D}(j\beta)^\dagger \approx X_s \cos \frac{\beta}{4\beta} + Y_s \sin \frac{\beta}{4\beta} \\ Y_s(\beta) &= \hat{D}(j\beta)\hat{P}_{\frac{j\pi}{8\beta}}\hat{D}(j\beta)^\dagger \approx Y_s \cos \frac{\beta}{4\beta} - X_s \sin \frac{\beta}{4\beta}. \end{aligned} \quad (4.21)$$

The displacement amplitude β corresponds to an approximate rotation of the encoded cavity state before detection. In Fig. 4.9c, we plot \mathcal{O}_2 for each of the four permutations of the joint observables and find a maximum Bell violation $\mathcal{O}_2 = 2.14 \pm 0.03$ for a displacement $\beta = 0.15$ and $\alpha = 1$ (Fig. 4.9d), which produces the approximate observables:

$$\begin{aligned} A_a &= X_a; & B_a &= Y_a \\ A_s &= \frac{X_s + Y_s}{\sqrt{2}} & B_s &= \frac{X_s - Y_s}{\sqrt{2}} \end{aligned} \quad (4.22)$$

Eq. 4.15 describes the ideal observables that efficiently determine a logical qubit observable, using a superposition state with $|\alpha| \gg 1$. For smaller coherent displacements, a modification must be introduced in order to observe the maximum Bell signal in this second CHSH experiment. The optimal observable $\hat{P}_{\pm j\beta_0} \sim \frac{1}{\sqrt{2}}(\hat{X}_c \pm \hat{Y}_c)$ in fact follows the relation:

$$\frac{\alpha - \beta_0}{\alpha + \beta_0} = \tan 4\beta_0\alpha, \quad (4.23)$$

as shown in [Park et al., 2012], sec. III; here, β_0 is the amplitude for a coherent displacement $\hat{D}(j\beta_0)$ to perform the measurement $\hat{P}_{j\beta_0}$ given α . In the large α limit, the observable corresponds to the logical qubit state observable $\frac{1}{\sqrt{2}}(\hat{X}_c + \hat{Y}_c)$ and follows the relationship $\hat{P}_{\alpha = \frac{j\pi}{16\beta}}$, the analog of a $\pi/4$ rotation in the cavity's phase space. The predicted and chosen optimal values for a maximum CHSH Bell signal in the second test are shown in Fig. 4.10.

As with the first test, we can similarly model the effects of photon loss and measurement inefficiency for the second test, as shown in Fig. 4.9d. The ideal case is the result of four summed joint

Wigner values represented as:

$$\mathcal{O}_{\text{ideal}} = 2(\cos 4\alpha_0\alpha + \sin 4\alpha_0\alpha)e^{-2|\alpha_0|^2}, \quad (4.24)$$

where β_0 is an optimal displacement for maximum violation which can be calculated from Eq. 4.23 and in detail in Ref. [Park et al., 2012]. Taking into account photon loss and measurement inefficiency produces the following relationship:

$$\mathcal{O}_{\text{pred}} = 2\mathcal{V}e^{-2t_{\text{eff}}/\tau_s|\alpha|^2}(\cos 4\alpha_0\alpha + \sin 4\alpha_0\alpha)e^{-2|\alpha_0|^2} \quad (4.25)$$

This predicted Bell signal is shown in Fig. 4.9d using the measured joint-Wigner contrast $\mathcal{V} = 0.85$ and an effective time $t_{\text{eff}} = 1.24 \mu\text{s}$, and again demonstrates excellent agreement with the measured data.




Figure 4.10: Optimal displacement for a Bell violation. For performing the second Bell test, the optimal observables to measure maximum correlations depend on the size of the Bell-cat state Eq. 4.23. The dashed black line shows numerically calculated optimal displacement points as a function of coherent state amplitude α . Shown in circles are the experimentally determined optimal displacement values used to measure a maximum Bell violation. Differences between chosen and ideal values are a result of the discretization of our measurement settings. The dashed green line is the approximate trend $|\alpha_{\text{opt}}| = |\frac{j\pi}{16\alpha}|$ for large cat states, which diverge at small α .

A lower Bell signal is observed in the second test due to its greater sensitivity to photon loss, yet in both tests two regimes are evident. For small cat state amplitudes, the initial Bell signal is limited by the non-orthogonality of the coherent state superpositions, while for large displacements the system's sensitivity to photon loss results in a reduction of the Bell signal. Larger, more distinguishable states quickly devolve into a classical mixture due to the onset of decoherence, corresponding to the resolution of Schrödinger's thought experiment. However, for intermediate cat state sizes, we observe Bell signals surpassing classical predictions larger than statistical uncertainties in both tests. The results provide direct evidence that no dominant error processes in the composite system besides single photon loss in the cavity and ancilla decoherence due to T_1 and T_2 limit the degree of

entanglement we can observe. Although perhaps not surprising, this result is certainly reassuring, and provides the first indication of the strong grasp we have in understanding the dynamics of our hardware-efficient system.

Scrutinizing detector bias

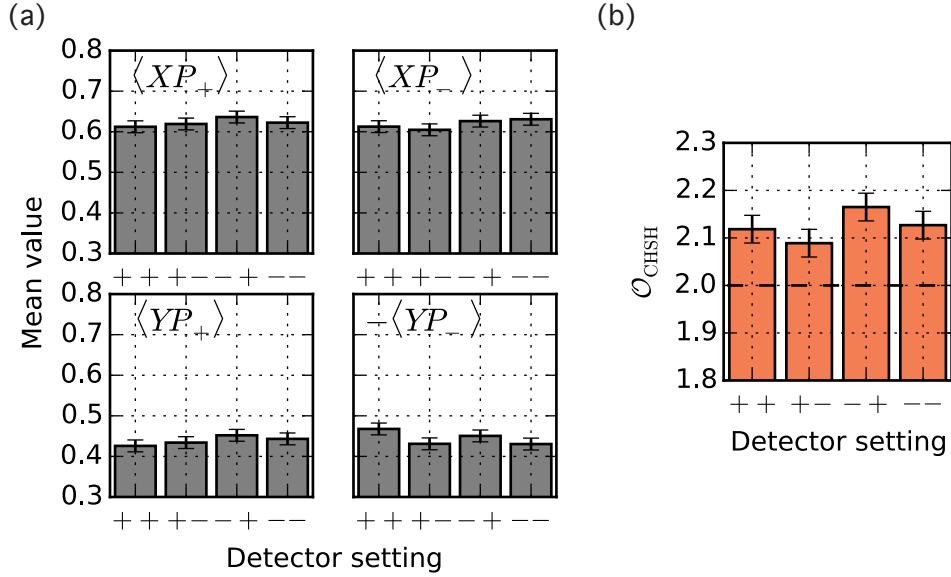


Figure 4.11: **Observables from each detector setting.** (a) To ensure that a particular detector setting is not producing systematic errors we have not taken into account, we report a Bell test for each detector setting used to observe our maximum violation in the second test. The expectation value of each observable used in that Bell test is shown for the four detector settings used. Significant deviations due to unexpected systematic errors are not observed. (b) A Bell test is analyzed for each detector setting to determine the effects of possible systematic errors. Each of these subtests violate Bell’s inequality by more than three standard deviations of their statistical error.

The sequential detection protocol in this experiment uses the same detector to perform first an ancilla measurement followed by a cavity measurement. To minimize unwanted systematic errors due to detector cross-talk between measurements, we perform each experiment under four detector setting permutations (Fig. 4.11). Two settings are used for the ancilla measurement: a pre-rotation which maps an ancilla eigenstate $|\pm\rangle$ to detector values $\pm M_1^a$ and another which maps $|\pm\rangle$ to $\mp M_1^a$. Two settings are used for the cavity measurement: a Ramsey experiment that maps a cavity eigenstate $|\pm\rangle$ to detector values $\pm M_2^s$ and another which maps $|\pm\rangle$ to $\mp M_2^s$. Each detector setting is performed an equal number of times³ and results are combined to remove unwanted correlations between detector readings and measured quantum observables. Systematic errors are shown to be

³By not randomly selecting detector orientations, we close a freedom-of-choice loophole.

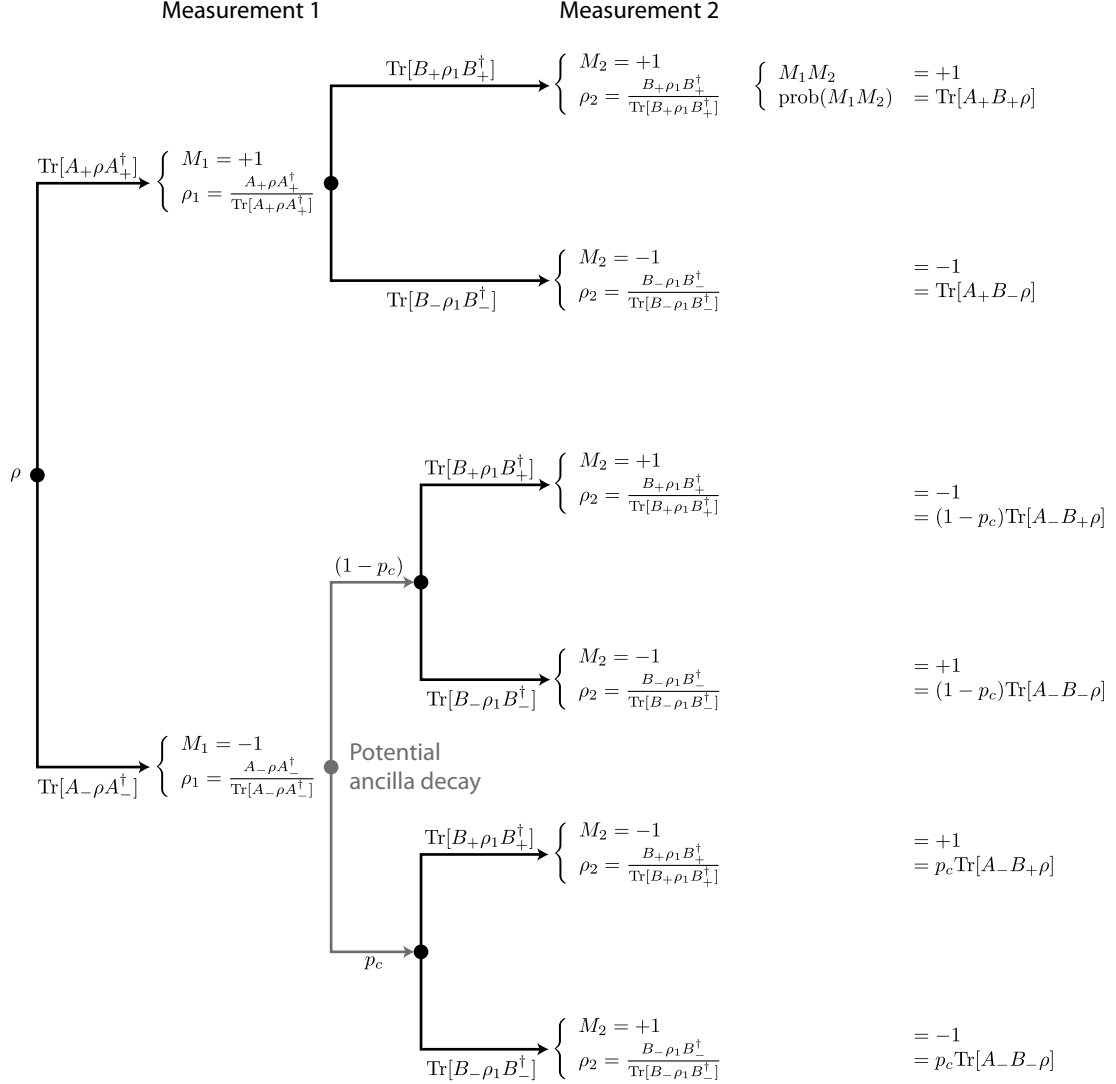


Figure 4.12: **Measurement trajectories given ancilla decoherence.** We can model the behavior of ancilla decoherence in a single measurement trajectory. Ancilla decay (which occurs with a probability p) can lead to an improper initialization of the second detection and in turn produces an incorrect measurement result. This form of detector cross-talk can lead to a reduction in visibility and potential systematic offset of the measured ancilla-cavity observable: $\langle AB \rangle \rightarrow (1 - p_s) \langle AB \rangle - p_s \langle B \rangle$.

within statistical bounds of the experiment and each detector setting violates Bell's inequality by at least three standard deviations.

The dominant form of cross-talk for this experiment is due to ancilla state decoherence between measurements. To realize the cavity state measurement, the ancilla must be initialized in $|g\rangle$, which we perform using active feedback (see A.1). Ancilla decay can occur during this reset process causing an incorrect initialization for cavity state detection. We can model this error by observing the possible trajectories of each measurement outcome (Fig. 4.12). This modifies the average measurement

of the observable $A_a B_s$ where A_a, B_s are ancilla and cavity operators respectively, which can be decomposed into ancilla projectors $A_a B_s = (A_{a+} - A_{a-}) B_s$, where $A_{a+} + A_{a-} = I$. Due to ancilla decay, the measured value $\langle A_{a+} B_s \rangle$ will be modified to $(1 - 2p_s) \langle A_{a+} B_s \rangle$ where p_s is the probability of ancilla decay in the time between the first measurement and the feedback rotation. This relation changes the measurement into:

$$\begin{aligned}
\langle AB \rangle &\rightarrow (1 - 2p_s) \langle A_+ B \rangle - \langle A_- B \rangle \\
&= (1 - p_s) \langle A_+ B - A_- B \rangle - p_s \langle A_+ B + A_- B \rangle \\
&= (1 - p_s) \langle AB \rangle - p_s \langle B \rangle
\end{aligned} \tag{4.26}$$

For measuring $B = X_s, Y_s, Z_s$ of the Bell-cat state $|\psi_{BC}\rangle$, we expect $\langle B \rangle = 0$, which gives merely a reduction in the visibility of the observable $\langle AB \rangle$ by a factor $(1 - p_s)$ without systematic offsets. We estimate in this experiment that $p_s = 1 - e^{-\frac{\tau_{\text{wait}}}{T_1}} \approx 0.06$, where $\tau_{\text{wait}} \approx 740$ ns, the total length of the readout, starting from when photons first enter the readout cavity to when the feedback resets the ancilla to $|g\rangle$. With this justification we can predict the additional loss in visibility \mathcal{V} mentioned in the previous section which gives a visibility $\mathcal{V}_{\text{pred}} = (1 - p_s) \mathcal{V} = 82\%$. The experimentally obtained visibility \mathcal{V} is 85%; we believe the discrepancy between predicted and measured values is due to an overestimate in the time the ancilla is susceptible to energy decay during measurement.

4.5 What we learn from studying the Bell-cat

Besides characterizing the high degree of entanglement in our Bell-cat, the tests detailed above demonstrate that simple encoding techniques allow for the efficient extraction of information from states stored in a cavity, illustrating the viability of measuring redundantly encoded states in multi-level systems [Gottesman et al., 2001]. Furthermore, this implementation provides a resource for quantum state tomography and quantum process tomography of continuous-variable systems and creates a platform for measurement-based quantum computation and quantum error correction using superconducting cavities [Leghtas et al., 2013]. Finally, these features can be extended to multi-cavity systems [Milman et al., 2004], which will require entanglement detection between continuous variable degrees of freedom and entanglement distribution of complex oscillator states.

From a physical perspective, the series of experiments described in this chapter offers a unique perspective on the ‘Schrödinger’s cat’ thought experiment. The artificial atom is entangled with a macroscopically distinguishable state in the cavity, the coherent state of finite amplitude. With the understanding that the degree to which the system violates the Bell inequality is dictated primarily

by the atom’s decoherence and photon loss, we can imagine a situation in which the quality factor of both components becomes arbitrarily large, wherein entanglement would be confirmed for coherent states large enough to blur the boundary between what is quantum and what is classical.

From a more practical point of view, we have already demonstrated clear advantages of hardware efficiency without even exploring error correction in detail. In contrast to large multi-physical qubit systems, efficient encodings and measurements of information redundantly stored in a cavity can be performed with two simple operations: conditional gates comprised of just a standard Ramsey-type pulse sequence and displacements of the cavity field. Furthermore, the good agreement between the data and models of entanglement witness decay with increasing encoding size confirms the assertion that single photon loss is the primary source of decoherence in the cavity. Thus demonstrating key requirements laid out in the cat code proposal with simple and straightforward experimental techniques engenders the confidence that motivates us to continue exploring its feasibility.

Yet the logical encoding used in this study of ancilla-cavity entanglement (Fig. 4.1) does not lend itself naturally to quantum error correction. As described in 3.3, a basis in which the Z_s axis is defined by coherent states of opposite phase lacks an inherent symmetry property that can be exploited to measure an error syndrome to correct for single photon loss. Indeed, the only two states that are left unchanged by photon number parity measurements are those that lie along $\pm X_s$. However, these are of course the very same states that are used to build even and odd parity bases for a logical qubit encoded in “4-cats.” So another noteworthy conclusion we draw from the study of Bell-cat entanglement is that measurements of \hat{X}_s , or the very error syndrome the cat code requires, can be extracted from our system with high fidelity and in a single-shot way without projecting the state out of the code space. This motivates the following question: what will we learn by now repeating this measurement in time?

5 – Tracking the Cat Code’s One Error Syndrome

The cat code lays forth a protocol to measure the occurrence of photon jumps through a mapping of the logical qubit’s symmetry property, photon number parity, onto an ancilla whose state can be projectively measured. The system that can be used to realize such a scheme is hardware-efficient, and requires just two cavities bridged by a single transmon ancilla in a superconducting cQED architecture. The Hamiltonian is well understood and the parity mapping scheme requires just a simple Ramsey-style pulse sequence. So in principle, all the ingredients are there to measure the parity in real-time with repeated single-shot measurements in order to protect a quantum bit of information from the cavity’s dominant source of decoherence, single photon loss.

The question, however, is not whether the parity can be measured. By showing a violation of a Bell inequality in the previous chapter we clearly demonstrate that the parity information of the cavity state can be extracted efficiently. Yet that demonstration employs just one such measurement and thus leaves several questions outstanding, answers to which are crucial in determining the feasibility of the cat code for future experiments. Namely, we would first like to know how quickly and with what fidelity the parity can be extracted from the system. In any practical implementation, one will require an accumulation of measurement statistics to acquire a best estimate of the cavity state, and an approach to learn this information in lieu of experimental imperfections must be developed. Secondly, we must know what effect the syndrome measurement has on the logical qubit to see if by the very act of checking for photon jumps do we induce them to occur with greater frequency. In this chapter, we will quantify to what degree the parity monitoring scheme is quantum non-demolition (QND), or to what extent a parity measurement leaves the parity unchanged.

Along the same lines, maintaining the coherence of the superposition of cat states requires not just knowing the parity with high confidence, but also maintaining the correct phase of the coherent states in phase space. Thus, in a more general sense, we will introduce the importance of quantifying to what extent ancilla errors during syndrome measurements propagate to errors in the cavity state. This last consideration in particular will be a recurring theme throughout the remaining chapters due to its integral role in determining a host of strategies meant to maximize the performance of an

error correcting system. This chapter primarily focuses on the first two questions: understanding the speed and fidelity with which we extract the error syndrome, and how QND is our measurement scheme. We will then conclude with a brief discussion of the third: what factors must be taken into account to minimize the back-action on the field in the cavity from errors in the ancilla. Our conclusions lay the groundwork for the subsequent three chapters, in which we demonstrate an application of repeated single-shot parity measurements in a full quantum error correction system that is capable of enhancing the lifetime of a quantum bit of information.

5.1 Observing jumps of photon parity in real-time

Besides their necessity in quantum error correction and quantum information, QND measurements play a central role in quantum mechanics. The application of an ideal projective QND measurement yields a result corresponding to an eigenvalue of the measured operator, and projects the system onto the eigenstate associated with that eigenvalue. Moreover, the measurement must leave the system in that state, so that subsequent measurements always return the same result. The hallmark of a continuously repeated high fidelity QND measurement is that it demonstrates a canonical *Gedankenexperiment*: individual quantum jumps between eigenstates are resolved in time on a single quantum system. This ideal measurement capability has only been experimentally realized in the last few decades. The jumps of a two-level system between its energy eigenstates were first observed for single trapped ions [Bergquist et al., 1986; Nagourney et al., 1986; Sauter et al., 1986], and later in single NV centers in diamond [Jelezko et al., 2002; Neumann et al., 2010; Robledo et al., 2011]. The jumps of an oscillator between eigenstates with different numbers of excitations (Fock states), were first observed for the motion of an electron in a Penning trap [Peil and Gabrielse, 1999]. More recently, the observation of quantum jumps of light in cavity QED [Gleyzes et al., 2007; Guerlin et al., 2007], where the number of microwave photons in a cavity is probed with Rydberg atoms, has enabled a range of new experiments in quantum feedback and control [Deléglise et al., 2008; Sayrin et al., 2011]. With the advent of quantum-limited parametric amplifiers [Bergeal et al., 2010; Castellanos-Beltran et al., 2008], measurement techniques for superconducting devices have rapidly advanced as well. For instance, the frequency shift of a cavity has been used to observe the quantum jumps of an ancilla between energy eigenstates [Hatridge et al., 2013; Vijay et al., 2011].

In this work, we demonstrate the first observation of quantum jumps of an error syndrome rather than an energy eigenstate. In our system, the error syndrome is a change in the photon number parity of the cavity field, which indicates the occurrence of single photon loss. With each change of parity, we only learn that the photon number changes between even and odd, but not how many

photons there are. This is a crucial distinction from projecting a cavity, or any other quantum system, onto a definite energy eigenstate. The measurements maintain the field in a superposition of coherent states, and the jumps of the error syndrome simply indicate a known change in the phase relationship between them. As described in chapter 3, this feature is exactly what quantum error correction, and in particular the cat code, requires. What we will find is that even in the presence of rapidly repeated measurements the field in the cavity is on average largely unperturbed, exhibiting the same time constant of decay toward vacuum as when it is left to evolve freely with no outside intervention. At the same time, when examining individual single-shot time records of the measurement, we will see clear signatures of parity jumps that with high fidelity reveal the occurrence of single photon loss in real-time.

Setting the measurement strategy

In this experiment, we employ the same three-dimensional circuit QED architecture [Paik et al., 2011; Wallraff et al., 2004] as introduced previously (Fig. 2.4), with a single superconducting transmon ancilla coupled to two waveguide cavities [Kirchmair et al., 2013; Vlastakis et al., 2013]. The transmon ancilla has a transition frequency of $\tilde{\omega}_a/2\pi = 5.938$ GHz, an energy relaxation time $T_1 = 8 \mu\text{s}$, a Ramsey time $T_2 = 5 \mu\text{s}$, and is found to be in the ground state $|g\rangle$ 86% of the time when in thermal equilibrium with the environment. The high frequency cavity, with $\tilde{\omega}_r/2\pi = 8.174$ GHz and a lifetime of 30 ns, serves only as a fast readout of the ancilla state. In order to perform a high-fidelity single-shot dispersive readout of the ancilla, we use a Josephson bifurcation amplifier (JBA) operating in a double-pumped mode [Kamal et al., 2009; Murch et al., 2013] as the first stage of amplification. The low frequency cavity, with $\tilde{\omega}_s/2\pi = 7.216$ GHz and a lifetime of $\tau_s = 1/\kappa_s = 55 \mu\text{s}$, stores the photon states that are measured and manipulated; its thermal population is $\sim 2\%$ (see sec. A.2 for more details). For simplicity, we will refer to the storage cavity as “the cavity” henceforth.

The ancilla and cavity are in the strong dispersive coupling regime, which again can be described by the Hamiltonian:

$$\hat{H}/\hbar = \tilde{\omega}_a |e\rangle \langle e| + (\tilde{\omega}_s - \chi_{sa} |e\rangle \langle e|) \hat{a}_s^\dagger \hat{a}_s \quad (5.1)$$

where \hat{a}_s and \hat{a}_s^\dagger are the annihilation and creation operators respectively, $|e\rangle$ is the excited state of the ancilla, and $\chi_{sa}/2\pi = 1.789$ MHz is the ancilla state-dependent frequency shift of the cavity. The readout cavity Hamiltonian terms are omitted because this component remains in the ground state while the parity mapping is performed. We have adjusted the phase between the JBA readout signal and the pump such that $|g\rangle$, $|e\rangle$, and $|f\rangle$ states can be distinguished with optimal contrast, where $|f\rangle$ corresponds to all ancilla states above $|e\rangle$. Figure 5.1 shows the trimodal histogram of the

ancilla readout. Thresholds between $|g\rangle$ and $|e\rangle$, and between $|e\rangle$ and $|f\rangle$ states have been chosen to digitize the readout signal to $+1, -1$, and 0 for $|g\rangle, |e\rangle$, and $|f\rangle$ states respectively. We assign a zero to the $|f\rangle$ states to indicate a “failed” measurement with no useful information about the parity. The basic ancilla readout properties are shown with the cavity left in its thermal state. The $|g\rangle$ state is prepared through a post-selection of an initial ancilla measurement, while $|e\rangle$ and $|f\rangle$ are prepared by properly pulsing the selected $|g\rangle$ state. The loss of fidelity predominantly comes from the T_1 process during both the waiting time of the initialization measurement (500 ns) and the ancilla readout time (300 ns).

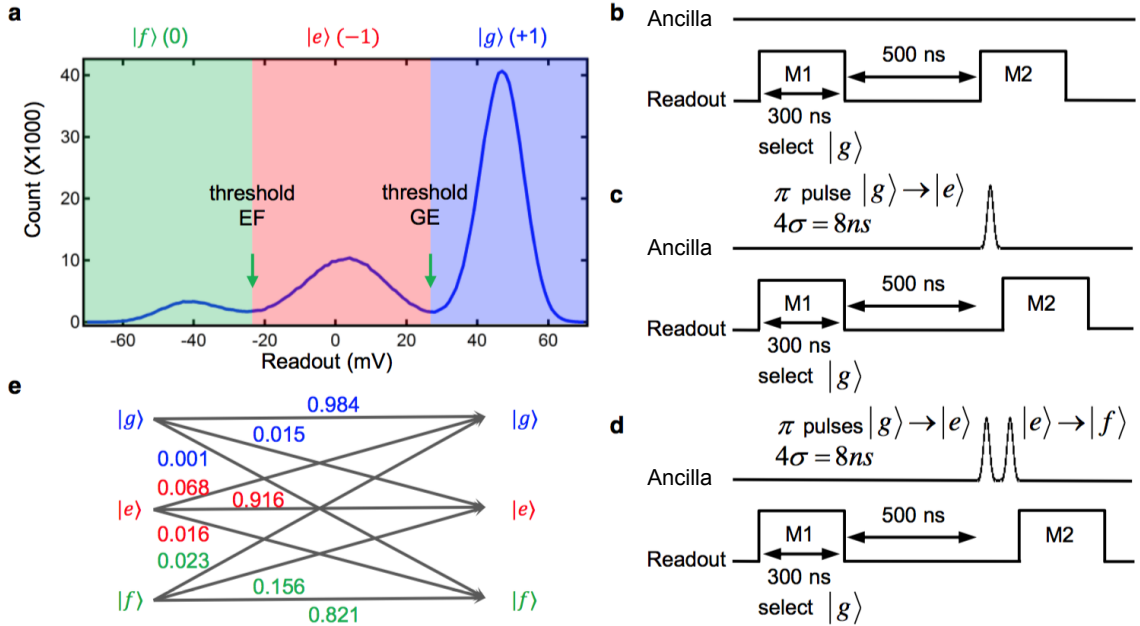


Figure 5.1: Ancilla readout properties. (a) Histogram of single-shot ancilla readout. The phase between the JBA readout and the pump tone has been adjusted such that $|g\rangle, |e\rangle$, and $|f\rangle$ states can be distinguished with optimal spacings. The signal has been digitized such that a $+1$ is assigned to $|g\rangle$, -1 to $|e\rangle$ and 0 to $|f\rangle$. (b-d) Illustration of pulse sequence (not to scale) producing an ancilla readout matrix with the storage cavity left in vacuum. The $|g\rangle$ state (b) is prepared through a post-selection of an initial ancilla measurement M1, while $|e\rangle$ (c) and $|f\rangle$ (d) are prepared by properly pulsing the selected $|g\rangle$ state. A histogram of the second measurement M2 gives the ancilla readout properties. (e) The resulting ancilla readout matrix.

As the first step in the experimental sequence here, we choose to initialize the cavity in an eigenstate of parity. We do not use the qcMAP protocol [Leghtas et al., 2013; Vlastakis et al., 2013], but instead create a “2-cat” stochastically through the method of measurement back-action described in sec. 4.1. We first create the Bell-cat $|\psi_{BC}\rangle = \frac{1}{\sqrt{2}}(|g, \alpha\rangle + |e, -\alpha\rangle)$, as shown in Fig. 4.2, and then measure the ancilla along $-X_a$. If the ancilla is observed to be in $|g\rangle$, we have an even cat in the cavity, while if the result is $|e\rangle$ we have an odd cat. We then continue with repeated

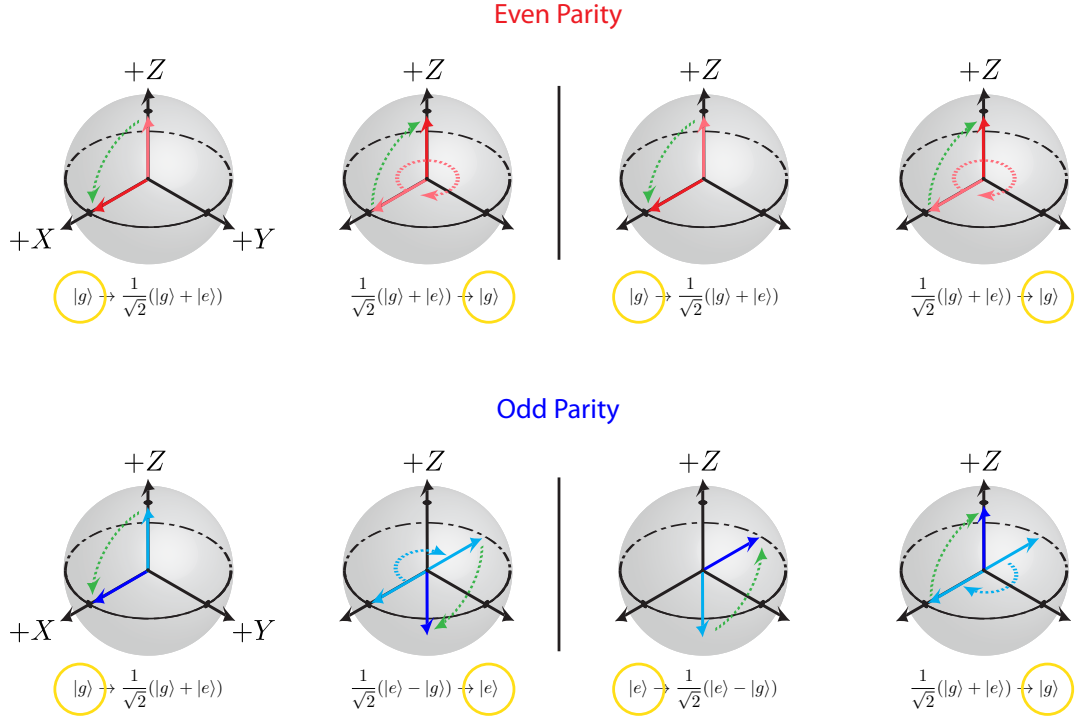


Figure 5.2: Asymmetry in parity mapping. Two consecutive parity mappings, separated by vertical black lines, are shown. Green arrows indicate the action of the pulse in rotating the Bloch vectors. A parity mapping consists of two $\pi/2$ pulses separated by a waiting time $t = \pi/\chi_{sa}$. Light and dark shades of arrows indicate the initial (prior to first $\pi/2$ pulse) and final (after to second $\pi/2$ pulse) Bloch vectors of the ancilla for each parity. Red (blue) arrows correspond to the time evolution of an ancilla state with two (one) photons in the cavity. In this example, the two pulses are opposite in sign, which means that if the parity is even and the ancilla is initialized in $|g\rangle$ ($|e\rangle$), the resulting measurement pattern will be “flat,” with $|g\rangle$ ($|e\rangle$) the only final outcome. On the other hand, should the cavity parity be odd, using the same protocol maps the ancilla to the opposite pole of the Bloch sphere as compared with its initial state. In this example, in the first mapping the ancilla starts in $|g\rangle$ and ends in $|e\rangle$; without a reset pulse, the next parity mapping begins with the ancilla in $|e\rangle$, after which it returns to $|g\rangle$, thus exhibiting an “oscillating” pattern. We therefore expect the signature of a photon jump in the single-shot traces to be a change between the flat and oscillating patterns.

measurements of parity at a rate that far exceeds the average lifetime of a single photon in the cavity. With the parameters above, we achieve a parity mapping and subsequent projective ancilla measurement in $1 \mu\text{s} \ll 55 \mu\text{s}$. Such high cadence is an important requirement for successful QEC, as it both reduces the risk of two errors that occur in quick succession and allows one to rapidly build up measurement statistics.

Regarding the creation of the initial cavity state and the pulse pattern of the parity mapping, we make two somewhat arbitrary choices. First, we employ a measurement-based strategy to create the “2-cats” rather than qcMAP. Although the resulting cavity state is obtained probabilistically, this method is easier to implement and faster than qcMAP at making parity eigenstates. Second, upon creating $|\psi_{BC}\rangle$ we measure the ancilla along $-X_a$ rather than $+X_a$. When repeated in time,

a sequence of parity measurements returns two contrasting measurement outcome patterns, either flat or oscillating (Fig. 5.2). A change in the measurement record from one pattern to the other indicates a change in parity and the loss of a photon. The correspondence of each pattern to a particular parity, however, depends on the choice of sign in the second $\pi/2$ pulse of the Ramsey sequence, as depicted in Fig. 5.2. Since the parity of the vacuum state $|0\rangle_f$ is even, on average the parity $\langle \hat{P} \rangle$ will always be biased towards even as the field decays in time, regardless of the initial coherent displacement α in the cavity (Fig. 5.3):

$$\langle \hat{P} \rangle = e^{-2|\alpha(t)|^2} \quad (5.2)$$

$$= e^{-2\bar{n}_0 e^{-\kappa_s t}}. \quad (5.3)$$

As measuring the ancilla in $|g\rangle$ rather than $|e\rangle$ is notionally preferable due to the effects of T_1 decay, measuring along $-X_a$ increases the probability of a flat pattern with the ancilla in $|g\rangle$ for even parity. As will be shown later, this consequently increases the fidelity of the parity measurement for the majority of the single-shot traces.

Outside of this experiment, however, this strategy will not be the future approach we will take. For large cat state amplitudes ($|\alpha| \gtrsim \sqrt{2}-\sqrt{4}$) stabilized for long periods of time [Leghtas et al., 2015a], effectively no bias towards any one parity exists and so neither parity mapping protocol is more capable in keeping the ancilla in $|g\rangle$ over the other. As will be shown in later chapters, this consideration provides one of the first motivations for the application of real-time feedback to the experimental implementation of cat code in our system, one which puts both protocols on an equal footing, maximizes the parity measurement fidelity for both parity eigenstates, and substantially reduces the probability of ancilla excitation to states $|f\rangle$.

Single-shot trajectories of cavity parity

We now proceed to demonstrate the real-time changes in cavity parity through repeated single-shot measurements of X_s . Figure 5.4 shows a simple schematic of the pulse sequence and typical resulting 400 μs single-shot traces. The initial displacement is $|\alpha| = 1.0$ and the repetition interval of the parity measurements is 1 μs , much lower than the single photon lifetime $\tau_s = 55 \mu\text{s}$ (Fig. 5.3). We observe a variety of measurement records, from quiet traces that last for hundreds of microseconds with no apparent changes in parity, to those that have as many as five jumps. The clear dichotomy between the flat and oscillating patterns in the raw data of our traces suggests a strong sensitivity of single-shot parity measurements to discrete changes in the state of the cavity. At this point, however,




Figure 5.3: **Average parity decay of a coherent state to vacuum.** The lifetime of the cavity is characterized by measuring the average parity of a coherent state $|\alpha\rangle$ left to decay freely to vacuum as a function of time. The measurement protocol is shown in the inset; the cavity is displaced with coherent states of various amplitudes (distinguished by the different colors), is left to evolve freely for varying delay times, and then has its average parity measured (rectangle P). All measured curves saturate at the same value, which deviates from exactly +1 due to a finite thermal population in the cavity. The solid lines are global fits using Eq. 5.3 that give a time constant $\tau_s = 55 \mu\text{s}$. The saturation level has been forced to 0.96 (due to $n_{th}^s = 0.02$), represented by the dashed horizontal line.

we cannot yet claim with any confidence that these traces indicate that single photon jumps are being monitored in a QND way.

In pursuing this goal, we are immediately faced with the challenge that qualitatively one can see that over the course of 400 μs every trace has a number of intermittent and usually brief changes in measurement correlation. At certain times, for example, flat patterns of consecutive $|g\rangle$ results jump to a flat pattern of repeated measurements of $|e\rangle$. At other times, the ancilla seems to settle in the $|f\rangle$ state for some length of time and then returns back to either $|e\rangle$ or $|g\rangle$. One may explain some of these anomalies away by claiming that there always exists a probability that two jumps may happen in quick succession, perhaps explaining the occurrence of single spikes between $|g\rangle$ and $|e\rangle$. The probability to have so many of these, however, is exceedingly low according to the Poisson distribution, given that the average time between photon jumps in these traces should be on the order of $\sim 1 \mu\text{s}/\tau_s \approx 0.02$ for an $\alpha \approx 1$. Instead, the litany of all other possible sources of error, which include ancilla T_1 and T_ϕ , photon jumps during parity mappings, and indeed perhaps the non-QND nature of parity measurements themselves, are the likely culprits here. But before we draw any

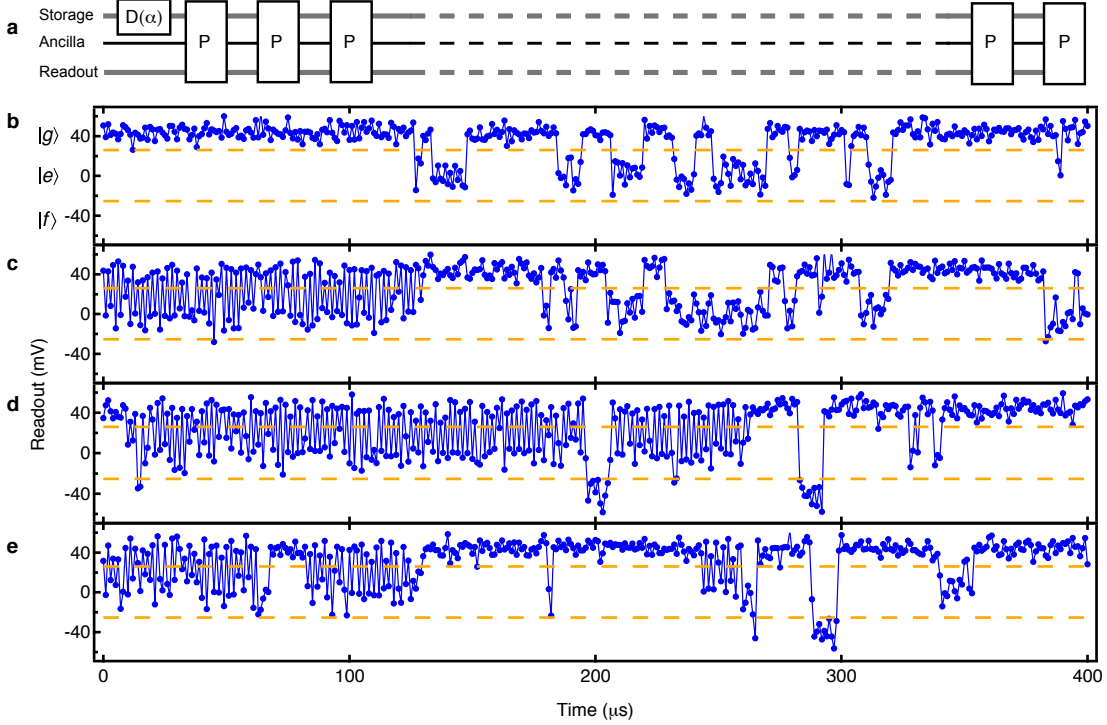


Figure 5.4: Observing parity jumps in real-time. These traces reveal the occurrence of parity jumps in real-time. Horizontal dashed orange lines delineate the thresholds to distinguish $|g\rangle$, $|e\rangle$, and higher excited states of the ancilla, denoted as $|f\rangle$. Blue dots are the actual measured voltage signal. **(a)** In this experiment, we first displace the storage cavity and then proceed to measure the parity (P rectangle) repeatedly in time in a single-shot way. The sign of the second $\pi/2$ pulse is chosen to be negative to measure a flat pattern when the cavity returns to vacuum (even parity). The repetition time of the parity measurement is $1 \mu\text{s}$, and traces in (b) through (e) all have an initial displacement of $|\alpha| = 1$. **(b)** No obvious changes between flat and sustained oscillating measurement patterns are observed. For the most part the cavity seems to be in an even parity eigenstate for the entire duration of $400 \mu\text{s}$, although multiple occurrences of multiple jumps in quick succession cannot be ruled out. **(c)** One parity jump is clearly discerned, as seen from the change in the measurement pattern (oscillating vs. flat) at about $130 \mu\text{s}$. **(d)** At least two parity jumps are recorded at about $10 \mu\text{s}$ and then again at $260 \mu\text{s}$. The ancilla leaves the computational space ($|f\rangle$) at about $200 \mu\text{s}$, a feature that disables the parity measurement. **(e)** A trace with at least five discernible jumps.

conclusions as to the QND nature of the parity measurements from these qualitative observations alone, we first must understand the fidelity of a single parity measurement given the preparation of a known parity state in the cavity.

5.2 Error budget for a parity measurement

Rigorously understanding what limits the performance of a single parity measurement plays an important role in assessing the future viability of the parity mapping protocol. More specifically in the context of this experiment, we would like to be sure that our knowledge of the Hamiltonian parameters and coherence properties of the system's components is sufficient to predict the results demonstrated thus far. This approach sets an important precedent in benchmarking the performance

of an entire error correcting system, as will be seen in chapter 6. The levels of detail we delve into here also underscore the virtue of the cat code’s simplicity, which offers the luxury of having to unravel the dynamics of just a single error syndrome measurement.

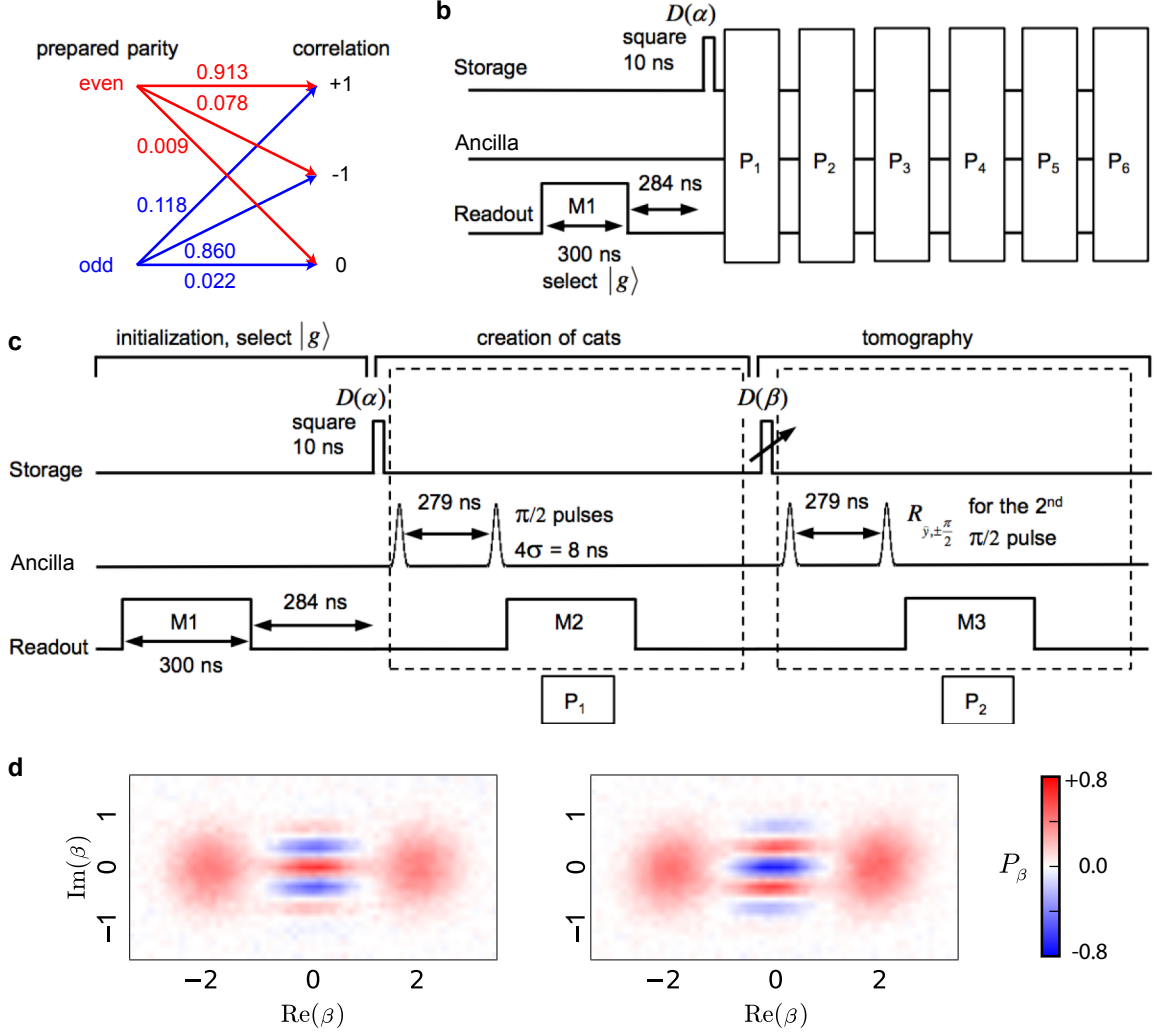


Figure 5.5: Parity readout properties and Wigner tomography. (a) Conditional probabilities for the correlations between neighboring parity measurements given initial even and odd cat states. These are found using the Ramsey mapping in which the second $\pi/2$ pulse is negative and a coherent displacement of $\alpha = 1$; $\mathbb{P}(+1|even) = 0.913$, $\mathbb{P}(-1|odd) = 0.860$, etc. (b) To start the experiment we post-select $|g\rangle$ after an initial ancilla readout (M_1); then perform five parity measurements (P_1, \dots, P_5) and post-select on a flat (oscillating) ancilla measurement pattern, thus purifying even (odd) parity states. Conditional probabilities $\mathbb{P}(\pm 1, 0|even/odd)$ are determined from a histogram of the results of the sixth parity measurement. (c) Illustration of the pulse sequence (not to scale) for producing the cat states and, additionally, Wigner tomography. A 280 ns waiting time after each measurement has been chosen to ensure that the readout cavity is at the vacuum state and likewise to coincide with π/χ_{sa} so that even and odd cats are aligned in phase space post ancilla readout. Tomograms are symmetrized to remove background offsets due to the cavity’s anharmonicity (see A.2). The ancilla pulses have a Gaussian envelope truncated to $4\sigma = 8$ ns and the displacement pulses on the storage cavity are 10 ns square pulses. The dashed line enclosure represents the pulse sequence for a parity measurement. (d) Wigner tomography of even (left) and odd (right) parity “2-cats” demonstrates the same back-action as seen in Fig. 4.3.

We would like to quantitatively know what stages of the parity measurements are most sensitive to the dominant sources of ancilla and cavity decoherence. Figure 5.5 shows the parity readout properties of our system. The overall performance of the parity measurement is characterized by looking at the distribution of ancilla states after a parity mapping conditioned on either an initial even or odd cat state in the cavity. Taking the example of starting in an even cat, the system is first initialized in the state $|\psi\rangle_{init} = (|\alpha\rangle + |-\alpha\rangle)|g\rangle/\sqrt{2}$ to the best of our ability (Fig. 5.5b). This state is created by first measuring the ancilla and through post-selection choosing only when the result is $|g\rangle$; then displacing the cavity ($|0\rangle_f \rightarrow |\alpha = 1\rangle$); and then purifying the cavity state into one of even parity by performing five consecutive parity measurements (P_1, \dots, P_5) and again post-selecting only when all results are $|g\rangle$. We then find the probability of measuring each ancilla state in the sixth parity measurement P_6 . The correlated outcomes can be $+1$ ($|g\rangle, |g\rangle$), -1 ($|g\rangle, |e\rangle$), and 0 ($|g\rangle, |f\rangle$). Ideally, conditional probabilities $\mathbb{P}(+1|even)$ and $\mathbb{P}(-1|odd)$ both equal $+1$, while the remaining $\mathbb{P}(+1|odd)$, $\mathbb{P}(-1|even)$, $\mathbb{P}(0|even)$, and $\mathbb{P}(0|odd)$, are 0. In practice, $\mathbb{P}(+1|even)$ and $\mathbb{P}(-1|odd)$ are on the order of 90%.

| Source of errors | Duration (μs) | Coherence time (μs) | Parity meas. errors |
|--|----------------------|----------------------------|---------------------|
| ancilla T_1 between two $\pi/2$ pulses | 0.28 | 8.0 | 1.8% |
| ancilla T_ϕ between two $\pi/2$ pulses | 0.28 | 7.3 | 3.8% |
| ancilla excitation to $ f\rangle$ between two $\pi/2$ pulses | 0.28 | 56 | 0.5% |
| photon jump between two $\pi/2$ pulses ($\bar{n} = 1$) | 0.28 | 50 | 0.3% |
| ancilla measurement inaccuracy in P_5 | NA | NA | 0.8% |
| ancilla excitation to $ e\rangle$ after measuring $ g\rangle$ in P_5 | 0.28 | 56 | 0.5% |

Figure 5.6: **Error budgets for the parity readout infidelities.** The error budget for the parity readout infidelity accurately predicts $\mathbb{P}(+1|even)$ from Fig. 5.5. .

The deviation of the conditional probabilities from their ideal values primarily arises from ancilla decoherence. The error budgets for the parity measurement fidelity can be estimated by breaking down a parity measurement into all of its constituent pulses and delays (e.g. Fig. 5.5c). For the following we again assume that we start in $|\psi\rangle_{init} = (|\alpha\rangle + |-\alpha\rangle)|g\rangle/\sqrt{2}$ prior to the sixth measurement P_6 . Should an ancilla T_1 decay event occur between the two $\pi/2$ pulses of the Ramsey mapping of P_6 , the second $\pi/2$ pulse simply takes the ancilla back to the equator of the Bloch sphere, and the subsequent projective readout has a 50% chance of reporting an ancilla result that correctly corresponds to the cavity parity. An ancilla T_ϕ event, or phase flip, between the two $\pi/2$

pulses flips the ancilla state on the equator and leads to the final state $|\psi\rangle_{fin} = (|\alpha\rangle + |-\alpha\rangle)|e\rangle/\sqrt{2}$ at the end of the sequence, resulting in an error 100% of the time. Additional decoherence arising from excitations to the $|f\rangle$ state between the two $\pi/2$ pulses in the sixth parity measurement also contributes to infidelity. The photon jump process between the two $\pi/2$ pulses leads to $|\psi\rangle_{fin} = (|\alpha\rangle - |-\alpha\rangle)[(1 - e^{i\chi_{sa}t_j})|g\rangle - (1 + e^{i\chi_{sa}t_j})|e\rangle]/2\sqrt{2}$, where t_j is the time of the photon jump. On average this gives a 50% chance of an incorrect answer. Finally, the ancilla measurement inaccuracy in the fifth parity measurement P_5 (due to measurement infidelity as shown by the overlapping histograms in Fig. 5.1) plus again ancilla excitation both serve to decrease $\mathbb{P}(+1|even)$ as they break the assumption that $|\psi\rangle_{init} = (|\alpha\rangle + |-\alpha\rangle)|g\rangle/\sqrt{2}$. All of the above sources of errors can be estimated given our knowledge of the sequence duration and the coherence times of the ancilla and cavity. Tabulated together, they add up to 7.7% (Fig. 5.6), in a good agreement with the $\mathbb{P}(+1|even) = 91.3\%$.

5.3 QND measurements of single photon jumps

The quantum filter

The parity measurement infidelities we found in the previous section indicate that with high probability many of the abrupt features in the traces from Fig. 5.4, which briefly break a well-defined measurement pattern, may in fact be due to ancilla and cavity decoherence. In an effort to obtain a cleaner signal that is perhaps more faithful to the cavity's true error trajectory, we design a quantum filter that astutely estimates the actual parity of the cavity at a time t . It takes as inputs the entire raw measurement record up to t , the initial cavity displacement, the ancilla measurement fidelities (Fig. 5.1), and the conditional probabilities to measure certain correlations for given prepared parity eigenstates (Fig. 5.5). The details of how exactly this filter is designed are described in detail in A.2; its restorative effects, however, are plainly clear in Fig. 5.7 and in greater detail in Fig. 5.8a, where it is shown in red. In an ideal system with no ancilla decoherence and perfect measurement fidelity, the cavity parity can be derived exactly from a single-shot trajectory by correlating neighboring points of raw data. Using the Ramsey sequence chosen here, the correlation for an oscillating pattern is always -1 (odd parity), and $+1$ for a flat pattern (even parity). In our system many of these correlations are spurious due to errant measurements. The quantum filter, however, is clearly much less sensitive to such abrupt changes in the measurement record and $|f\rangle$ state occupation, and in particular favors changes to even parity after long time intervals since the likelihood of the cavity being in vacuum becomes increasingly high. Due to the smoothing effect of the quantum filter we

actually can achieve nearly unity detection sensitivity to changes in photon parity. Intermittent spikes, small and large, are symptoms either of the effects of decoherence on the filter’s estimate of parity, or simply photon jumps in quick succession that the filter lacks the bandwidth to catch.




Figure 5.7: **Filtering out decoherence in the single-shot records.** Applying a quantum filter, in red, to the single-shot traces in Fig. 5.4. (a) The filter confirms what we earlier suspected, that changes in the measurement pattern are most likely due to measurement error rather than actual jumps of parity. (b) The filter indicates a single parity jump, given by the complete roll-over from -1 to $+1$ (right axis). Spikes in the filter indicate likely ancilla decoherence, although multiple parity jumps in quick succession cannot be ruled out. These spikes highlight the filter’s finite bandwidth. (c) The measurement record is filtered to show two parity jumps. (d) A trace with at least five parity jumps. Despite intermittent spikes, the filter achieves nearly complete confidence in the parity for most of the trace.

To illustrate the importance of filter bandwidth on jump sensitivity, we show in Fig. 5.8b a zoom-in of the behavior of the quantum filter applied to typical photon jump events. Green and cyan curves are fits of the filter at the transition based on a tanh function, giving a transition time constant of less than $1 \mu\text{s}$. We can also use a simpler definition for τ_f as just the total time it takes the filter to roll over from one threshold to the other (± 0.9). This approach yields a longer time constant, $\tau_f \approx 2 \mu\text{s}$, but one that more accurately captures the dynamics of the measurement results. The probability of missing a photon jump is equal to the probability of a photon jump occurring within τ_f . This is given by:

$$\mathbb{P}_{jump} = \frac{\bar{n}}{\tau_{tot}} \int_0^{\tau_f} e^{-t\bar{n}/\tau_{tot}} dt \quad (5.4)$$

$$= 1 - e^{-\tau_f \bar{n}/\tau_{tot}}. \quad (5.5)$$




Figure 5.8: **Effectiveness and response time of the quantum filter.** (a) Effectiveness of the quantum filter. Blue (raw) and red (filtered) curves are the same as that shown in Fig. 5.7. The green curve is the direct correlation of ancilla states between neighboring parity measurements. The red curve is clearly much smoother and can reject the brief changes in the green curve. (b) Response time of the quantum filter applied to typical photon jump events. The blue curve is the raw data from a repeated parity measurement. The red curve is the corresponding filter based on the quantum filter. Green (cyan) curves are fits to tanh functions of the filter at the transition down (up), giving a transition time constant of less than $1 \mu\text{s}$. However, the response time of the filter to make a transition between ± 0.9 is $\tau_f \approx 2 \mu\text{s}$.

For $\bar{n} = 1$ and $\tau_{tot} \approx 50 \mu\text{s}$, the probability of missing both jumps is thus $\mathbb{P}_{jump} \approx 4\%$, and $\approx 15\%$ for $\bar{n} = 4$. This performance may strike one as too poor for the cat code error correction protocol to ever succeed in a real experiment. Using cavities and ancillae with coherence times of an order of magnitude greater [Reagor et al., 2016], however, substantially reduces these percentages. Indeed, as will be seen in later chapters, hardware improvements render such double errors as one of the lowest contributions to cat code failure.

We have thus answered the first key question introduced at the outset of the chapter. By measuring the photon parity with a frequency that far exceeds the average error rate and intelligently interpreting the results with the use of a quantum filter, we can be very confident about the parity of the state in the cavity at any point in time. Not surprisingly, however, we also find that the sensitivity to jumps relies on the coherence properties of the system's components. In a system with ancillae and cavities that have superior coherence properties to those used here, this performance can certainly be boosted, as shown in chapters 6 to 8. For now, however, we must demonstrate that changes in parity indeed correspond to the error process for which the cat code is designed to correct.

Counting changes in parity

In order to test how faithfully ostensible jumps in parity correspond to QND measurements of single photon loss, we first simply count the number of jumps extracted from the quantum filter during $500 \mu\text{s}$ of repeated parity measurements. We have applied a Schmitt trigger [Schmitt, 1938] to the filtered data in order to digitize the filter and reject the unavoidable noise (spikes in the filter) coming from ancilla decoherence and erroneous parity readout. The two thresholds for the Schmitt trigger are chosen to be ± 0.9 to allow for a large discrimination. The number of parity jumps is inferred from the number of transitions in the digital data after the Schmitt trigger is applied. Figure 5.9 shows the histograms of the extracted number of jumps for an initial even or odd cat state by post-selection. As expected, the even and odd distributions in Fig. 5.9 indeed show a 98% vs. 2% mixing, in accord with the knowledge that $n_{th}^s = 2\%$. Due to the finite bandwidth of the filter, however, if two photon jumps occur within a filter response time τ_f , the Schmitt trigger does not catch both jumps.

Although we have no way of knowing the true number of photon jumps for each parity measurement trajectory, we can test how well the distribution of measured jumps agrees with the Poisson distribution we expect. Due to the complication of background thermal excitation and finite response time of the filter, finding an analytical solution becomes difficult. Instead, we compare a numerical Monte Carlo simulation with the experiment. In the simulation, we use a coherent state as the initial state without distinguishing the parity. Each simulation trajectory is $500 \mu\text{s}$ long and includes a transition probability $n \rightarrow n + 1$ from the background thermal excitation. We also neglect trajectories that have neighboring jumps within the response time τ_f of the quantum filter. Then, for each trajectory we count the number of jumps, and finally we make a histogram of those numbers based on 100,000 trajectories (black solid lines in Fig. 5.9). Such good agreement between simulation and data provides the first concrete evidence that repeated parity measurements do in fact faithfully track the loss of single photons in real-time.

Quantifying how QND parity measurements are

The extent to which these results agree with the expected Poisson distribution of photon jump number already strongly suggests that the parity measurements are highly QND as well. We can confirm this quantitatively by examining the decay rate of the average parity of a coherent state subjected to repeated single-shot parity measurements. Although the quantum filter appears to predict the jump statistics with good accuracy, we nonetheless would first like to confirm that we understand the system dynamics without heavily processing the raw data. We briefly return to

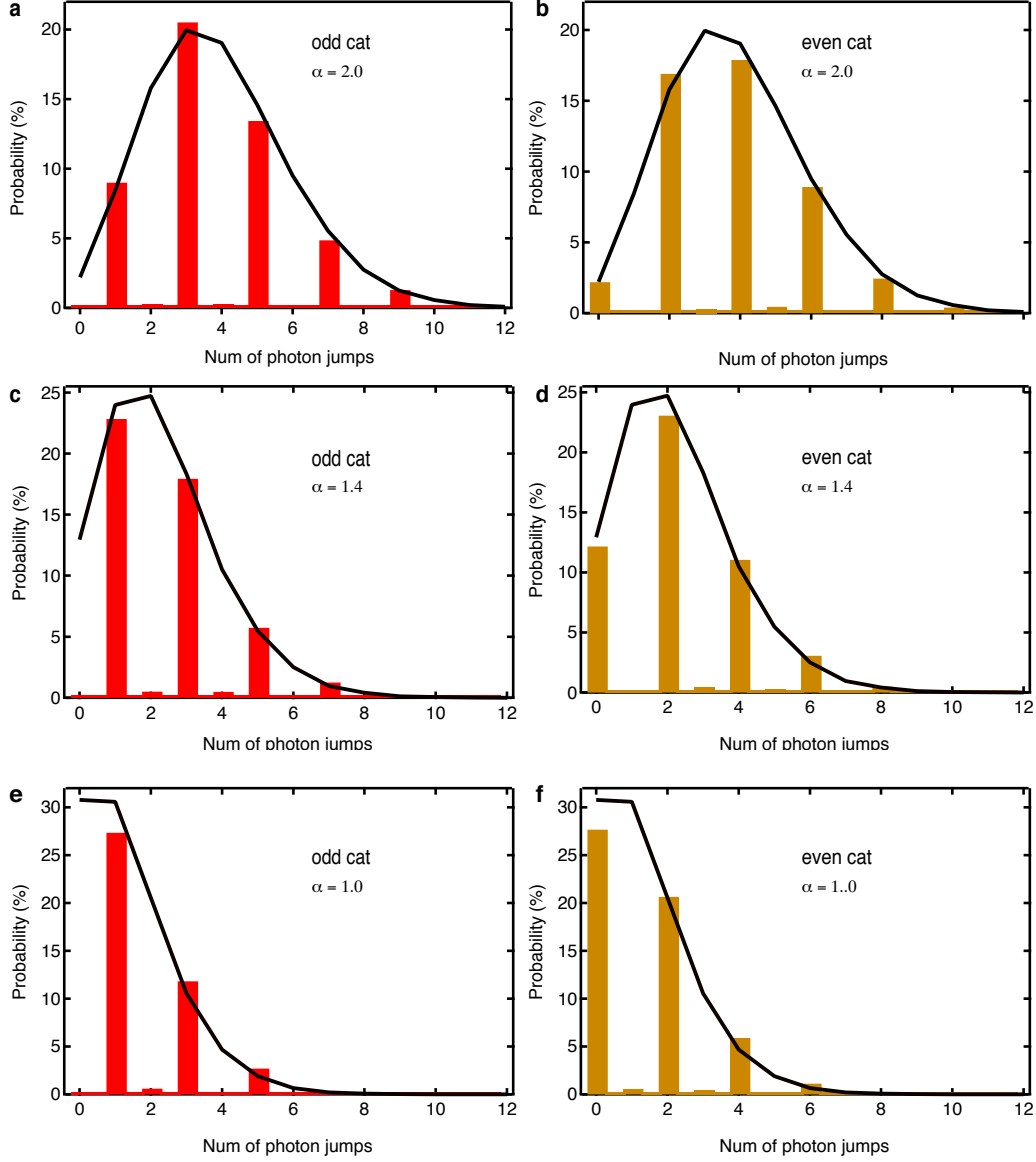


Figure 5.9: **Photon jump histograms.** Histograms of the number of jumps extracted from the quantum filter during $500 \mu\text{s}$ of repeated parity measurements for an initial even or odd cat state, as chosen through post-selection on the first parity measurement. (a), (b) $|\alpha| = 2.0$. (c), (d) $|\alpha| = 1.4$. (e), (f) $|\alpha| = 1.0$. Solid lines are numerical simulations that include n_{th}^s and the finite response time of the filter. In the simulation, we use a coherent state as the initial state without distinguishing the parity.

analyzing the data from the perspective of correlations only, and assume that the repeated parity measurement process is Markovian. The ensemble averaged parity dynamics obtained from the average correlation $\langle C_{cor}(t) \rangle$ under repeated parity monitoring are then simply:

$$\langle C_{cor}(t) \rangle = \mathbb{P}(+1, t) - \mathbb{P}(-1, t) + \mathbb{P}(0, t) \frac{\mathbb{P}(+1, t) - \mathbb{P}(-1, t)}{\mathbb{P}(+1, t) + \mathbb{P}(-1, t)} \quad (5.6)$$

where $\mathbb{P}(+1, t)$, $\mathbb{P}(-1, t)$, and $\mathbb{P}(0, t)$ are the probabilities of measuring positive, negative, and zero correlations at time t . The third term comes from the fact that the cases with zero correlation are assigned to previously measured non-zero correlations $+1$ and -1 , whose probabilities are $\mathbb{P}(+1, t - \Delta t)$ and $\mathbb{P}(-1, t - \Delta t)$, respectively. For small Δt , $\mathbb{P}(\pm 1, t - \Delta t) \approx \mathbb{P}(\pm 1, t)$.




Figure 5.10: Averaged single-shot correlations. Ensemble averaged raw data of single-shot records obtained directly from the correlation of ancilla states between neighboring parity measurements, using 20,000 traces similar to those shown in Fig. 5.4. Seven different cavity displacements are shown. Solid lines are predictions based on Eq. 5.6, in excellent agreement with the measured data. The offset of the averaged parity at $t = 0$ comes from the asymmetric parity readout fidelities between the even and odd cat states. The fact that the saturated parity value in the long time limit is evidence that ancilla decoherence and other parity measurement errors degrade the quality of the raw signal.

The probabilities $\mathbb{P}(+1, t)$, $\mathbb{P}(-1, t)$, and $\mathbb{P}(0, t)$ depend on both the measured parity readout property $\mathbb{P}(\pm 1, 0|even/odd)$ and the even and odd parity evolution $P_e(t)$ and $P_o(t)$ of the photon state:

$$\begin{aligned}
 \mathbb{P}(+1, t) &= \mathbb{P}(+1|even)P_e(t) + \mathbb{P}(+1|odd)P_o(t) \\
 \mathbb{P}(-1, t) &= \mathbb{P}(-1|even)P_e(t) + \mathbb{P}(-1|odd)P_o(t) \\
 \mathbb{P}(0, t) &= \mathbb{P}(0|even)P_e(t) + \mathbb{P}(0|odd)P_o(t)
 \end{aligned} \tag{5.7}$$

where $P_e(t) = (e^{-2|\alpha|^2 e^{-\kappa_s^{\text{tot}} t}} + 1)/2$, $P_o(t) = (1 - e^{-2|\alpha|^2 e^{-\kappa_s^{\text{tot}} t}})/2$, and $\kappa_s^{\text{tot}} \geq \kappa_s$ is a modified decay rate that could be greater than κ_s if the parity measurements are not QND. With all the parameters in Eq. 5.6 known, $\langle C_{cor}(t) \rangle$ can then be predicted.

Comparing the averages of correlated single-shot records with the model derived above (Fig. 5.10), the agreement is excellent provided $1/\kappa_s^{\text{tot}} = 1/49 \mu\text{s}$. We obtain an identical time constant if we

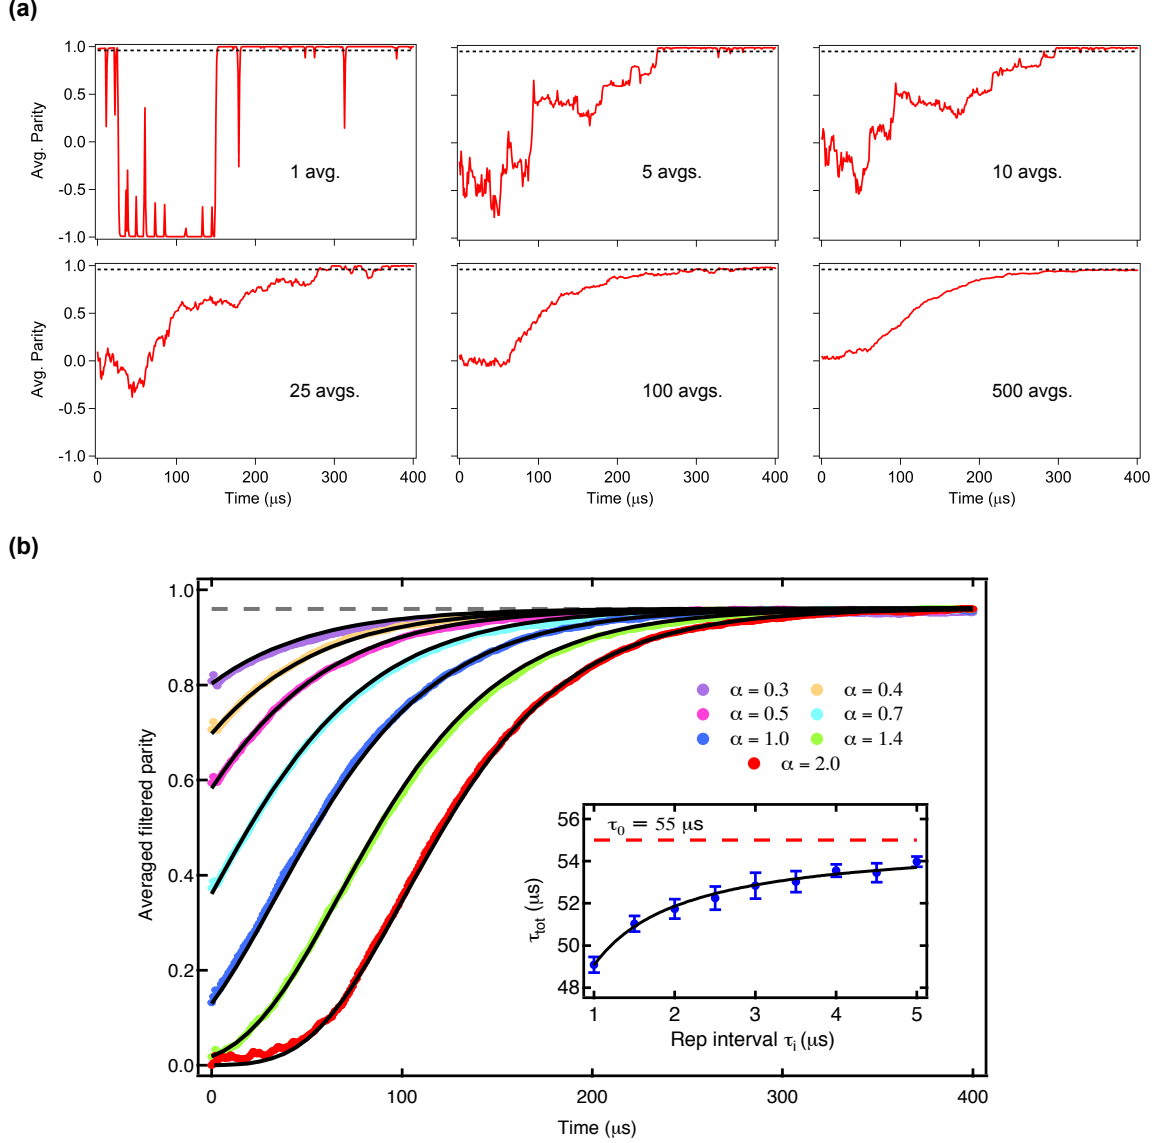


Figure 5.11: Demonstrating QND parity measurements. (a) An example of averaging filtered single-shot traces similar to those shown in Fig. 5.7. In this example, $\alpha = 2.0$. After 500 averages the curve already follows the trend predicted by Eq. 5.8. The average parity is expected to saturate at $\frac{1}{1+2n_{th}^2} = 0.96$ and indeed does so, as shown by the dashed horizontal line. (b) Ensemble averages of 20,000 filtered single-shot parity traces for different cavity displacements α . Solid lines are fitted theoretical curves for coherent states using Eq. 5.8, where n_{th}^s and the corresponding $|\alpha|$ are fixed, leaving $1/\kappa_s^{tot} = \tau_{tot}$ as the only free parameter. Here $\tau_{tot} = 49 \mu s$ from a global fit, and closely matches the lifetime $\tau_s = 55 \mu s$ as obtained from a free evolution measurement. This indicates that repeated parity measurements do not significantly perturb the cavity state, inducing few extra jumps and no measurable increase in cavity thermal population. Inset: extracted time constants as a function of different parity measurement repetition intervals. The error bars indicate one standard deviation from repeated measurements. The decay time τ_{tot} is modeled as $1/\tau_{tot} = 1/\tau_s + P_D/\tau_{rep}$, where P_D is the probability of inducing an extra parity change, or photon jump. A fit (solid line) gives $P_D = 2 \times 10^{-3}$, indicating that a single parity measurement is 99.8% QND.

average the filtered data as well (Fig. 5.11). In form, the curves in Figs. 5.10 and 5.11 mirror those shown in the case of the free decay of a coherent state (Fig. 5.3). The filtered data in fact demon-

strates the ability of the quantum filter to correct for measurement infidelity by removing the initial offset at $t = 0$ and saturating at the predicted thermal population of the cavity. The decay curves for seven different displacements are all well-modeled by a simple equation (similar to Eq. 5.3) that accounts for the finite thermal population of the cavity and a decay of average parity where \hat{a}_s is the only loss operator:

$$\langle \hat{P} \rangle = \frac{1}{1 + 2n_{th}^s} e^{-2|\alpha|^2 e^{-\kappa_s^{tot} t} / (1 + 2n_{th}^s)}. \quad (5.8)$$

We thus arrive at the conclusion that aside from a modified time constant, the dynamics of the system under repeated error syndrome measurements are modeled well by the simple stochastic loss of single photons in time.

The discrepancy in the time constants, $1/\kappa_s^{tot} = 49 \mu s$ versus $1/\kappa_s = 55 \mu s$, however, indicates that the parity measurements are not 100% QND. In order to quantify how destructive each parity measurement is, we extract the total decay rate of the parity ($1/\tau_{tot} = \kappa_s^{tot}$) from the ensemble averaged parity dynamics obtained with the quantum filter for nine different repetition intervals (τ_{rep}). The total decay rate is well modeled by the parallel combination of the free decay time ($\tau_s = 55 \mu s$) plus a constant demolition probability P_D :

$$\frac{1}{\tau_{tot}} = \frac{1}{\tau_s} + \frac{P_D}{\tau_{rep}}. \quad (5.9)$$

We find that $P_D = 0.002$ per measurement interval, as shown by the fit in the inset of Fig. 5.11. In other words, a single parity measurement is 99.8% QND, leaving the parity of the cavity state largely unperturbed. Subsequent parity measurements, however, can catch this event with high fidelity. This result has important consequences on the future viability of the cat code. If the very act of measuring parity were to induce photon jumps without our knowledge, the parity tracking protocol itself would be flawed. Performing the parity mapping and the subsequent strong projective readout of the ancilla, however, barely disturbs the logical qubit, a key ingredient for a working QEC system that allows us to repeat the procedure many times per photon jump.

Throughout the preceding analysis, we glossed over one crucial point. As shown in Fig. 5.5a, the conditional probability $\mathbb{P}(-1|odd)$ is a full 5% lower than $\mathbb{P}(+1|even)$. Yet, the only difference between the two protocols used to obtain these numbers lies in the parity of the initial state purified by the first five parity measurements; $\mathbb{P}(+1|even)$ corresponds to taking even parity states to $|g\rangle$ and $\mathbb{P}(-1|odd)$ corresponds to taking odd states to $|e\rangle$. We thus come full circle to the issue we encountered early on in sec. 5.1 when choosing the sign of the second $\pi/2$ pulse in the Ramsey

mapping to be negative. The problem arises from the substantial probability of ancilla T_1 decay during a projective readout of the ancilla in $|e\rangle$, resulting in parity information of inherently lower confidence for odd states with this choice of pulse sequence. Indeed, the quantum filter picks up on this problem as well; it is no coincidence that the majority of the spikes in the filtered data in Fig. 5.7 occur during the oscillating patterns that correspond to odd parity. At first one may brush this aside and claim that with higher coherence times this effect won't be as pronounced, and that even now the filter is so effective that this asymmetry does not play that great of a role in reducing the fidelity of the logical qubit. But this of course assumes that only missed photon jumps dephase the cat states. In the following section, confronting the third and final consideration introduced at the beginning of this chapter, we see why this in fact is not the case and begin to consider what is required for a parity measurement to be QND on the entire state in the cavity.

5.4 Motivating the need for real-time feedback

Finding that the parity measurements we employ to detect errors in the cavity are highly QND encourages us to continue the study of the cat code as a viable route towards protecting a quantum bit. This positive result, however, must not eclipse an equally pressing concern: the degree to which we may still be perturbing the information stored in the cat states. Referencing the Wigner tomograms (Fig. 5.5d), although we can confidently claim that we are QND as far as the point at the origin is concerned, a priori the same cannot be said regarding rest of the information present at other points in phase-space. In other words, although the parity monitoring doesn't change photon number probabilities it could still change the relative phases between the Fock states superpositions that constitute the cat states. This more general issue of the forward propagation of errors from the syndrome measurements to the logical qubit stored in the cavity must now be understood in detail and addressed.

We now look at the effects the dominant sources of decoherence for the ancilla, T_1 and T_ϕ , have on the state in the cavity. The very term in the Hamiltonian that enables the parity measurement, namely the dispersive shift in frequency χ_{sa} per photon, leaves the state in the cavity vulnerable to dephasing due to its coupling to a lossy ancilla. The detrimental effects of T_1 decay are thus apparent when recalling the entanglement between the cavity and the ancilla during the parity mapping protocol, where the cat state begins to acquire a phase at a rate $q\chi_{sa}$ that depends on the ancilla state $\{|g\rangle (q = 0), |e\rangle (q = 1), |f\rangle (q \approx 2)\}$. In the following example, we fix the rotating frame to be the one rotating at the cavity frequency $\tilde{\omega}_s$ when the ancilla is in $|g\rangle$. Specifically, if ancilla T_1 relaxation occurs at t during the parity protocol waiting time π/χ_{sa} , an ideal cat acquires

a phase proportional to t :

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|\alpha\rangle \pm |-\alpha\rangle) \rightarrow \frac{1}{\sqrt{2}}(|\alpha e^{i\chi_{sa}t}\rangle \pm |-\alpha e^{i\chi_{sa}t}\rangle) \quad (5.10)$$

We thus see that if the ancilla state changes at a random time without our knowledge, the cat state acquires an unknown phase proportional to χ_{sa} and the amount of time the state is out of the chosen rotating frame, resulting in a complete loss of information. In the Wigner tomography, this would manifest itself as a washing out of the cat state's features, and unlike at the origin, successive measurements can only further reduce the fidelity. Thus, in a given single-shot record, ancilla T_1 decay and excitations to higher ancilla states impart an arbitrary phase on the cat states from which it would be impossible to recover from without some auxiliary correction protocol.

The contribution of ancilla dephasing T_ϕ enters in a subtle way. Without loss of generality, let us assume the system is initially at state $|\psi\rangle = (|\alpha\rangle + |-\alpha\rangle)|g\rangle/\sqrt{2}$. After the first $\pi/2$ pulse, the dispersive interaction kicks in and the state becomes at some time $t < \pi/\chi_{sa}$:

$$|\psi\rangle = \frac{1}{2}[(|\alpha\rangle + |-\alpha\rangle)|g\rangle + (|\alpha e^{i\chi_{sa}t}\rangle + |-\alpha e^{i\chi_{sa}t}\rangle)|e\rangle]. \quad (5.11)$$

Again let's assume that at time t an error in the ancilla occurs, but this time a phase flip. The state becomes:

$$|\psi\rangle = \frac{1}{2}[(|\alpha\rangle + |-\alpha\rangle)|g\rangle - (|\alpha e^{i\chi_{sa}t}\rangle + |-\alpha e^{i\chi_{sa}t}\rangle)|e\rangle], \quad (5.12)$$

and then continues to evolve in the expected way, regardless of the change in sign of the term associated with $|e\rangle$. At the end of the waiting time π/χ_{sa} , the system becomes a product state again $|\psi\rangle = (|\alpha\rangle + |-\alpha\rangle)(|g\rangle - |e\rangle)/2$. The second $\pi/2$ pulse just takes the ancilla to the other pole on the Bloch sphere opposite to the case without the phase flip, erroneously indicating a change in parity, but otherwise leaving the cavity phase unchanged. Another way to understand the difference between an ancilla T_1 vs. T_ϕ error would be to recall the dispersive interaction between ancilla and cavity as formulated in the Hamiltonian: $(\tilde{\omega}_s - \chi_{sa}|e\rangle\langle e|)\hat{a}_s^\dagger\hat{a}_s$. Any change of the ancilla energy results in a change of the cavity frequency, thereby dephasing the cavity state. A phase flip by contrast does not change the ancilla energy, and therefore the cavity phase remains unaffected. Any extra phase acquired by the cavity state contingent on the final ancilla state post-measurement would be completely deterministic (to within our ability to discriminate between $|e\rangle$ vs. $|g\rangle$). By contributing to the parity measurement infidelity, however, the effect of T_ϕ necessitates more measurements to

achieve the high degree of confidence in the estimation of the cavity state's parity. As increasing the number of measurements increases the likelihood of T_1 decay, ancilla dephasing still facilitates an overall degradation of cat state fidelity.

What remains to be studied is how the possible failure modes, mirroring those introduced in chapter 3, contribute to the most important figure of merit: the degree of fault-tolerance the current experimental implementation of the cat code currently offers. The subtle relationship between errors in the code space and errors in the ancilla will in fact prompt us to completely rethink the optimal strategy in implementing QEC, and will challenge the typical wisdom that syndrome measurements performed at the highest frequency possible maximize the benefits of error correction. Aside from the natural desire to improve ancilla and cavity lifetimes, further characterizing these types of error processes are important next steps, and the benefits of doing so will become apparent in later chapters.

The study of parity measurement infidelity in sec. 5.2, together with the aforementioned consequences of potential ancilla back-action on the state in the cavity, motivate us to enhance the technological capabilities of our experiment to improve the performance of the parity monitoring protocol. For example, unmistakable signs of excitation to the higher excited states of the ancilla, as seen in Fig. 5.7, demonstrate the high probability of losing parity information over long tracking times if the ancilla occupies $|e\rangle$ too frequently. This inconvenient detail of the experiment can clearly be overcome with a straightforward application of real-time feedback. If one considers, for example, the possibility of modifying the pulse sequence in real-time to maximize the probability to measure the ancilla in its ground state after every measurement, already substantial inroads in enhancing parity mapping fidelities can be envisioned. We thus postpone the discussion on the levels of performance we can achieve until we describe in the next chapter a more sophisticated system in which real-time feedback plays a central role.

6 – A Full Quantum Error Correction System

A key challenge in demonstrating QEC that actually yields advantages when implemented within an experimental setup is including every step of the process to ensure that the dominant forms of decoherence that actually occur will be handled. One must show that *all quantum states* can be protected from *naturally occurring errors without the use of post-selection*. Previous works have demonstrated parts of a full QEC solution, but have never realized a system that can satisfy each of these three requirements at once. Typically, the approach has been either to correct artificially induced errors; focus on one out of several dominant error processes; assess the performance of known protected or particularly vulnerable states; employ post-selection as a means to study specific decoherence mechanisms; or implement some combination thereof [Aoki et al., 2009; Chiaverini et al., 2004; Córcoles et al., 2015; Cory et al., 1998; Cramer et al., 2016; Kelly et al., 2015; Knill et al., 2001a; Leung et al., 1999; Moussa et al., 2011; Nigg et al., 2014; Pittman et al., 2005; Reed et al., 2012; Ristè et al., 2015; Schindler et al., 2011; Taminiau et al., 2014; Waldherr et al., 2014]. By isolating subsets of a general error process to study the viability of a QEC scheme, however, these works indicate a correction of specific errors under restricted circumstances. Indeed, they do not quantify the exponential decay in time one would see in the process fidelity of a quantum bit subjected to repeated rounds of error correction.

In large part the central challenge for these efforts has been the resource overhead required by traditional QEC proposals: the large numbers of imperfect physical qubits and couplings involved that very quickly degrade QEC performance, in particular due to the decoherence penalty, as introduced in chapter 3. Due to this penalty, the coherence time of the logical qubit in fact drops at first due to the greater number of physical qubits involved in the encoding, typically by a factor of five to ten for traditional codes [Steane, 1996]. The overhead has prevented demonstrations of QEC beyond the break-even point, wherein one can fully make up for the decoherence penalty. Specifically, previous works have been unable to realize a system in which applying error correction actually yields a logical qubit with a lifetime longer than the system’s most coherent constituent. In this work, the hardware requirements of the cat code are minimal, and there is just one dominant source of

decoherence that we must combat: single photon loss. The prospects for achieving break-even are thus more realistic.

We have shown in the previous chapters that with fast, repeated photon number parity measurements, the natural occurrence of single photon loss can be tracked in real-time. We have understood the effects of an imperfect ancilla on the levels of fidelity achievable and that we must be wary of the possibility of the back-action of ancilla energy decay on the cavity state. In this chapter, we demonstrate the implementation of a full QEC system that protects a logical qubit exposed to its natural environment. We detail every step, from encoding a qubit, to monitoring the error syndrome, to decoding back, and we analyze the performance with full process tomography. We show, and indeed stress, the crucial role played by real-time feedback, which not only performs the more routine functions of ancilla reset, but moreover responds to the stochastic evolution of the cavity field and determines the necessary decoding procedure that depends on the error history. Finally, we set the stage for a discussion in subsequent chapters of what it takes to overcome the decoherence penalty incurred by a redundant encoding in a larger logical space.

6.1 Implementing a full QEC system

In a practical setting, in order to demonstrate QEC successfully, one must implement all elements of a full QEC system within a single shot of the experiment and benchmark the performance without the use of post-selection. Our strategy is based on the assumption that single photon loss is the dominant source of decoherence in the system, as justified in the introduction of the cat code (sec. 3.3). Consequently, as a first-order approach to enhancing the qubit's lifetime, we track just a single error syndrome: changes in photon parity that indicate the occurrence of photon loss. We are of course cognizant of sources of error besides photon loss, which serve to degrade QEC performance. By not performing post-selection or restricting the quantum states to be protected to specific and known points on the Bloch sphere, however, we do not artificially exclude such non-idealities and inflate our results. Instead, we employ a full QEC system in order to ensure that we in fact *include* all sources of decoherence. We thereby highlight the real challenges in demonstrating a system that can actually offer gains in preserving the integrity of quantum information over time.

Broadly speaking, a full QEC system is composed of the steps illustrated in Fig. 6.1, where the first and sixth steps, necessary for process tomography, serve to benchmark the performance of four in-between. The flow begins with an initialization of the quantum bit. In our experiment this involves applying pulses on the ancilla transmon to encode one of the six cardinal points, shown as forming an octahedron on the Bloch sphere. In a general application, initialization may simply

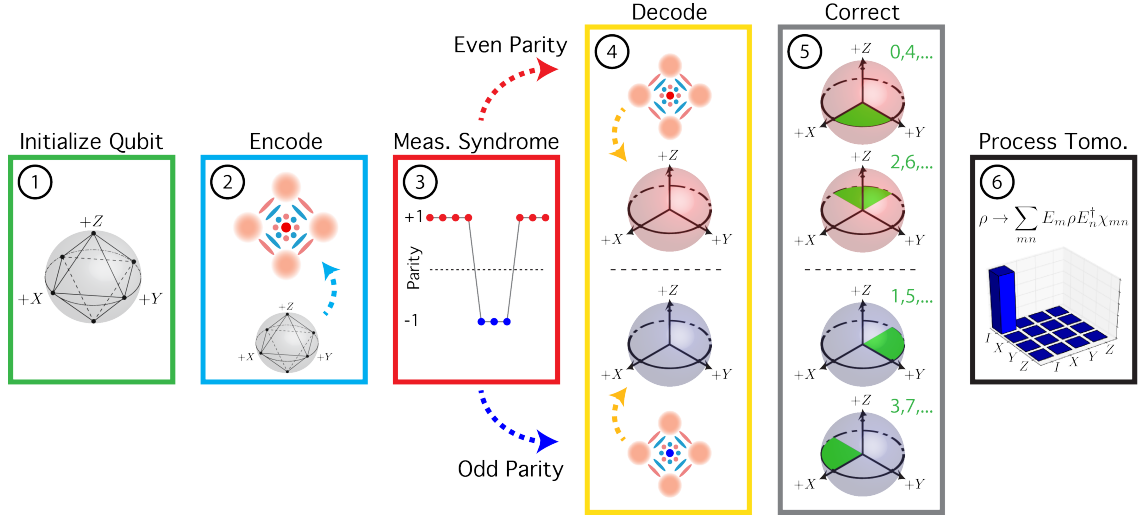


Figure 6.1: A full QEC system. Steps two through five take an input quantum bit and return a corrected quantum bit at the output. The first and sixth step benchmark the performance of the four in-between. This formulation is not specific to the cat code and generally describes the steps involved in any active, measurement-based error correction scheme. The process begins with an initialization of the quantum bit in one of the six cardinal points on the Bloch sphere, enough to characterize the action of the system on an arbitrary input state. Next, the state is encoded in a logical qubit, here comprised of superpositions of coherent states. The error syndromes are then monitored, and the resulting error history is used to decode and finally correct the state. Ideally, the corrected state will be the same as the one initialized. In reality, due to not only decoherence in the logical space but also any imperfections in the operations required in each of these steps, the final quantum bit will necessarily contain some degree of incoherent mixture, evidenced by deviations of the process fidelity X matrix from the ideal form depicted here.

involve the receipt of a quantum bit from some preceding step in an algorithm. The point is that this step must reflect the fact that one has no prior knowledge of the state that will be logically encoded; hence, studying the action of error correction on the six cardinal points amounts to benchmarking the process on any arbitrary input [Nielsen and Chuang, 2010]. The second step, encoding, involves applying some sequence of pulses on the system to transfer the information to a larger Hilbert space, thereby redundantly mapping the information onto a logical qubit whose symmetry properties allow one to monitor error syndromes. In this realization, we map the information from the ancilla onto a superposition of coherent states in the cavity following the prescription of the cat code. Next, one repeatedly monitors the errors due to the natural sources of decoherence in the system with single-shot QND syndrome measurements. As described several times, we infer the occurrence of errors in the cavity due to single photon loss from changes in the photon number parity of the cat states. Decoding and then correcting returns a density matrix that is ideally the same as the one that was initialized. In our system, depending on the best estimate of the number of jumps at the end of the monitoring step, different decoding pulses are applied to transfer the information back from the cavity back onto the ancilla, where simple corrections either with actual gates or in software can be made to account for the number of photon jumps that were detected. Taking

each one of the six cardinal points through these steps, from encoding to correction, allows us to perform process tomography (see chapter 7), which quantifies how well the QEC system implements the identity \hat{I} operation. Stated more simply, when the dust settles we obtain a single number that ranges from +1, which reflects a perfect quantum memory, to +0.25, which indicates the system returns a completely incoherent mixture.

The experimental components and setup of our system are again similar to those used in chapters 4 and 5, and are presented in greater detail in A.3. We employ a superconducting cQED architecture [Wallraff et al., 2004] using 3D waveguide cavities [Paik et al., 2011]. Specifically, a single transmon qubit is coupled to two such cavities [Kirchmair et al., 2013; Vlastakis et al., 2013] and is used as an ancilla to both interrogate the error syndrome and to encode and decode the logical states. One of the cavities stores the logical states while the other is used for ancilla readout. The dominant interaction terms are described by the following Hamiltonian:

$$\hat{H}/\hbar = \tilde{\omega}_s \hat{a}_s^\dagger \hat{a}_s + (\tilde{\omega}_a - \chi_{sa} \hat{a}_s^\dagger \hat{a}_s) |e\rangle \langle e| - \frac{K}{2} \hat{a}_s^{\dagger 2} \hat{a}_s^2, \quad (6.1)$$

with $|e\rangle \langle e|$ being the ancilla excited state projector; $\tilde{\omega}_s$, $\tilde{\omega}_a$ the storage cavity (henceforth just the cavity) and ancilla transition frequencies; $\chi_{sa}/2\pi \sim 1.95$ MHz the dispersive interaction strength between the two; and $K/2\pi \sim 4.5$ kHz the cavity anharmonicity, or Kerr. The readout cavity is excluded as it is only used to measure the ancilla at the end of each parity-check. The coherence times are much improved in this realization as compared with the experiments reported in previous chapters; $T_1 \approx 35 \mu\text{s}$ and $T_2 \approx 13 \mu\text{s}$, while the cavity has a single-photon Fock state relaxation time of $\tau_s \approx 250\mu\text{s}$, and $T_2^s \approx 330 \mu\text{s}$. To perform high-fidelity single-shot measurements of the ancilla [Hatridge et al., 2013; Vijay et al., 2011], we set the readout cavity to have a 1 MHz bandwidth and use a nearly quantum-limited phase-preserving amplifier, the Josephson Parametric Converter (JPC) [Bergeal et al., 2010], as the first stage of amplification, which allows for a readout fidelity of 99.3% and an error syndrome measurement fidelity of 98.5% (see sec. A.3 for more details).

Real-time feedback plays a crucial role in maximizing the performance of the full QEC system in our experiment and indeed enables us to demonstrate all the six steps outlined above without the use of post-selection. Thus, before delving into the detailed experimental implementation, we provide a high-level overview of the quantum control architecture in order to outline the basic hardware components that enable the fast and sophisticated reactions to the stochastic evolution of the quantum system. The discussion then turns to the specific applications of feedback and how encoding and decoding pulses are implemented. Finally, an operational example of the full QEC system is presented within a debugger view, which demonstrates not only the successful execution

of feedback responsibilities, but moreover our capability to predict the evolution of our system in time knowing only the parity measurement fidelity and the Hamiltonian parameters.

6.2 Quantum control architecture

Future experiments in quantum computation will necessitate modifying the pulse generation in real-time as a response to returning signals that indicate the current state of a system (real-time feedback), in particular to minimize qubit decoherence. Crucial to the success of the quantum experiment is the efficiency of collecting, interpreting, and reacting to the returned signals with sub-microsecond pulse lengths and feedback latencies. Previous results demonstrated that real-time feedback can be a powerful tool for realizing control and performance enhancements of a quantum system. For example, weak measurements with phase-sensitive amplification and a phase-locked loop setup were used to stabilize Rabi oscillations in a cQED setting [Vijay et al., 2012]. Using simple control electronics, the results of weak measurements with Rydberg atoms of a microwave field in a Fabry-Perot cavity were fed back into an actuator to stabilize a cavity state [Sayrin et al., 2011]. Additionally, FPGA-based control electronics were used in a cQED setting to apply a π pulse conditioned on a measurement result [Ristè et al., 2013], and enhance the performance of Bell-state stabilization using two FPGA cards that communicate their results to one another [Liu et al., 2016]. As in ref. [Liu et al., 2016], the architecture demonstrated here is comprised of four major components all on one piece of hardware (Fig. 6.2): Digital-to-Analog converters (DACs) that output pulses; Analog-to-Digital converters (ADCs) that sample input signals; digital inputs/outputs (DIG-IOs) that enable inter-card communication as well as the triggering of certain digital RF components; and finally a Field Programmable Gate Array (FPGA) that dictates the flow of the experiment in time, orchestrating the three previous components to steer the quantum system to some desired state in real-time. This is our quantum controller. It is a new, multi-purpose computer designed to execute programs for quantum control. In this experiment, we translate the cat code into a program the controller understands, load it onto the FPGA, and press “play.” A carefully choreographed stream of pulses on the ancilla, cavity, and readout cavity ensues to take a quantum bit through a series of error syndrome measurements and return a corrected quantum bit, thus realizing a full QEC system.

Each hardware unit, or “card,” is an independent agent. It combines the functionality of an instrument like a commercially available Arbitrary Waveform Generator (AWG), a data sampling card, and certain data analysis functions crucial for efficient feedback all on one piece of equipment. Such a design dramatically enhances the possible levels of control, sophistication, and complexity of

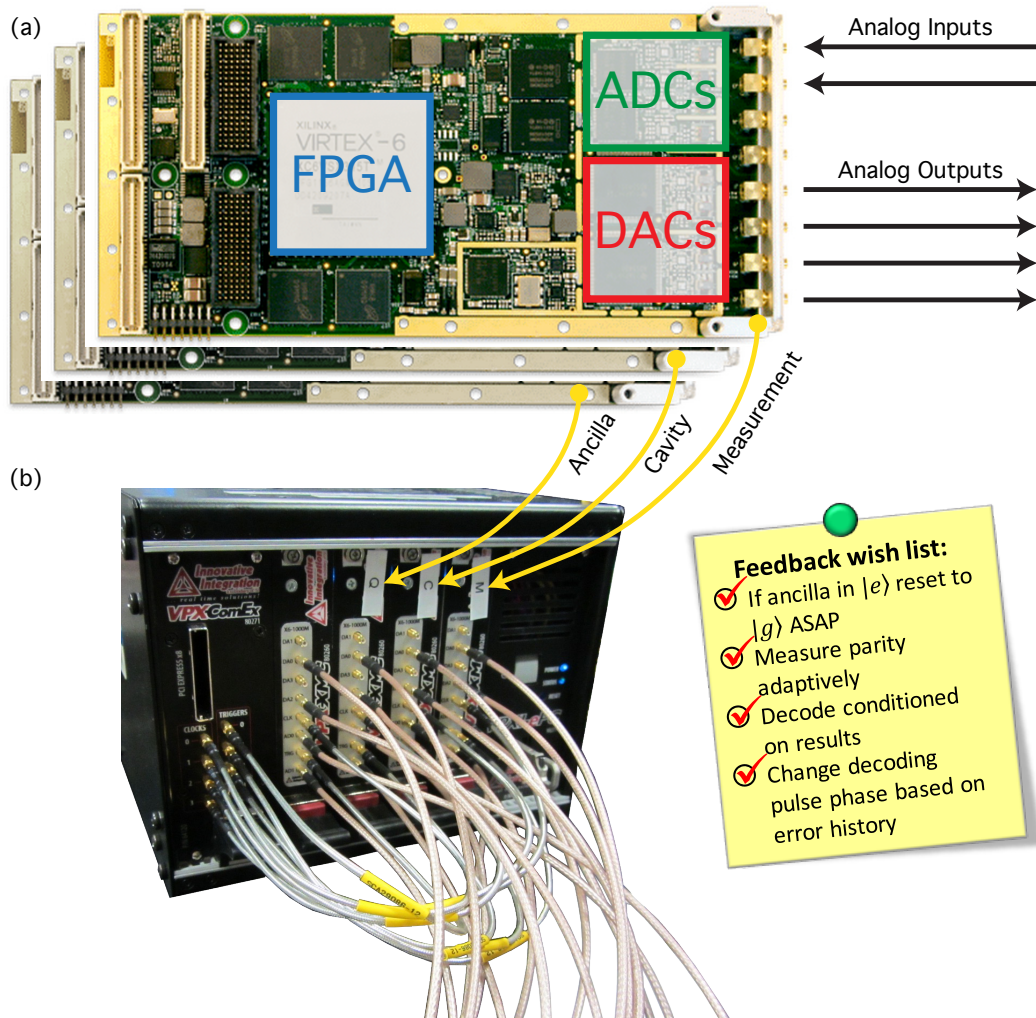


Figure 6.2: **Real-time feedback hardware.** (a) A picture of an Innovative Integration X6-1000M card and the layout of the main components, with arrows indicating where signals from the experiment come in, and pulses to the experiment go out. (b) A picture of the VPXI-ePC chassis with a VPX-COMEX module and four X6-1000M cards, three of which are used in the experiment (M: measurement, C: storage cavity, Q: ancilla). The feedback wish list briefly summarizes the intended implementations of real-time feedback.

a quantum experiment. Furthermore, all cards run in parallel with no inherent dependency on each other. They may produce pulses that are sent to manipulate a particular component of the quantum system. Incoming signals from the quantum system may also be routed as inputs to the cards in some pre-defined way. Each card thus produces and analyzes different signals, and it can then distribute its findings among the other cards in real-time through a dedicated digital communication layer.

The common denominator in this scheme is the set of instructions loaded onto each card prior to the experiment, which coordinates how the cards work together. Once the experiment starts there is no one master card that must dictate the flow; each card can decide what to do independently. The result is a decentralized network of classical computation that provides a fast, efficient, and flexible

platform to interface with the quantum system. By properly coordinating the signals sent and received by this network of cards, the user ultimately coordinates the interactions between distinct entities of the quantum system, all accomplished on time scales of just a few hundred nanoseconds.

We use three Innovative Integration X6-1000M boards housed in a VPXI-ePC chassis with a VPX-COMEX module, shown in Fig. 6.2, which produces a 1 GHz clock and triggers synchronized to within 50 ps. Each board contains two 1 GS/s ADCs, two 1 GS/s DAC channels, and digital inputs/outputs that are controlled by a Xilinx VIRTEX-6 FPGA loaded with in-house logic. The three boards are synchronized to control the storage cavity, readout cavity, and ancilla transmon. The readout signals are routed to the ADCs on the readout cavity board, whereafter the FPGA demodulates and thresholds the signal to determine the state of the ancilla ($|g\rangle$, $|e\rangle$, and higher). The feedback latency between the last sample arriving at the input to the ADC to the first sample out of the DAC is approximately 200 ns, providing us with a powerful tool to mitigate the deleterious effects of the forward propagation of errors due to ancilla decoherence, as described in the following section.

6.3 Detailing the six steps of a full QEC system

The role of real-time feedback can be divided into two general parts: maximizing the probability that the ancilla remains in its ground state throughout the error monitoring, and modifying the decoding pulses in real-time to take into account the error history. These two points summarize the entries in the feedback wish list of Fig. 6.2. Using real-time feedback to not only reset the ancilla whenever it is measured to be in the excited state, but also changing the measurement sequence to preferentially map the parity to the ancilla ground state minimizes the probability of this dephasing, owing to the asymmetry in the rates of ancilla decay versus ancilla excitation. Secondly, no one single set of pulses can map back the information from cavity states of different parity onto the ancilla. Moreover, the exact pattern of photon jumps dictates the final phases of the cat states due to an interesting effect of the cavity anharmonicity. These considerations are elaborated upon in detail below. To start, we summarize our full protocol (Fig. 6.3), wherein every run of the experiment cycles through the following steps:

1. System and ancilla reset – using feedback, we make sure that the ancilla is in ground state $|g\rangle$ and the cavity is in vacuum $|0\rangle_f$.
2. Quantum bit initialization – in each realization of the experiment we apply a gate on the ancilla to encode one of the six cardinal points on the Bloch sphere. This over-complete set

of states allows us to perform process tomography of the QEC system (see sec. 7.1), and is equivalent to characterizing the action of a system on the quantum bit $|\psi\rangle_{init} = c_0 |0\rangle + c_1 |1\rangle$.

3. Encoding – we transfer the quantum bit from the ancilla onto a superposition of cat states in the cavity. At the end of this step the state of the cavity is $|\psi\rangle_{init} \rightarrow c_0 |C_\alpha^+\rangle + c_1 |C_{i\alpha}^+\rangle$, while the ancilla, to the best of our ability given experimental realities (see sec. 6.3), ends in $|g\rangle$, ideally completely disentangled from the cavity state.
4. Parity monitoring – we identify photon jumps, or errors in our logical qubit, by monitoring the parity of the logical state in the cavity. This is done using an adaptive parity monitoring scheme. The total duration of each error syndrome measurement is just $1 \mu\text{s}$, or $\sim 0.8\%$ of the average time between photon jumps for cat states of $\bar{n} = 2$.
5. Decoding and correction – the quantum bit of information is brought back onto the ancilla using the knowledge we gather while monitoring the error syndrome (see sec. 6.3). A different decoding pulse is used for each point in time due to the changing amplitude and Kerr evolution of the cat states. At the end of the decoding pulse, the cavity should ideally be completely in vacuum and with the measured error record the quantum bit is corrected following the cat code prescription.
6. Tomography – we perform state tomography on the ancilla to compare the final quantum $|\psi\rangle_{fin}$ with the initial state $|\psi\rangle_{init}$. Using the results we fully characterize the QEC system process (see sec. 7.1).

Steps 1-2 prepare the system in its ground state with high accuracy. This part does not assume any knowledge about the quantum state it is designed to protect or the decoherence mechanisms. Steps 3-5 are the error correction part of the experiment. In the final step we measure the ancilla that is ideally back in the initial state. Any deviation leads to a decay of the process fidelity in time. While the entire experiment is implemented as one big state machine, only the exact durations of the system and ancilla reset step are not predetermined.

Ancilla reset

Ancilla reset is more than 99.8% effective and effectively no residual thermal population is measured after cavity reset (less than 0.1%). As detailed in Tab. A.5, the thermal populations of the ancilla and the cavity are $\sim 4\%$ and $< 2\%$, respectively. This is enough to adversely affect not only our encoding pulses, but also subsequent error syndrome detection. The protocol starts with the quantum controller measuring the state of the ancilla. If the result is the excited state $|e\rangle$, the

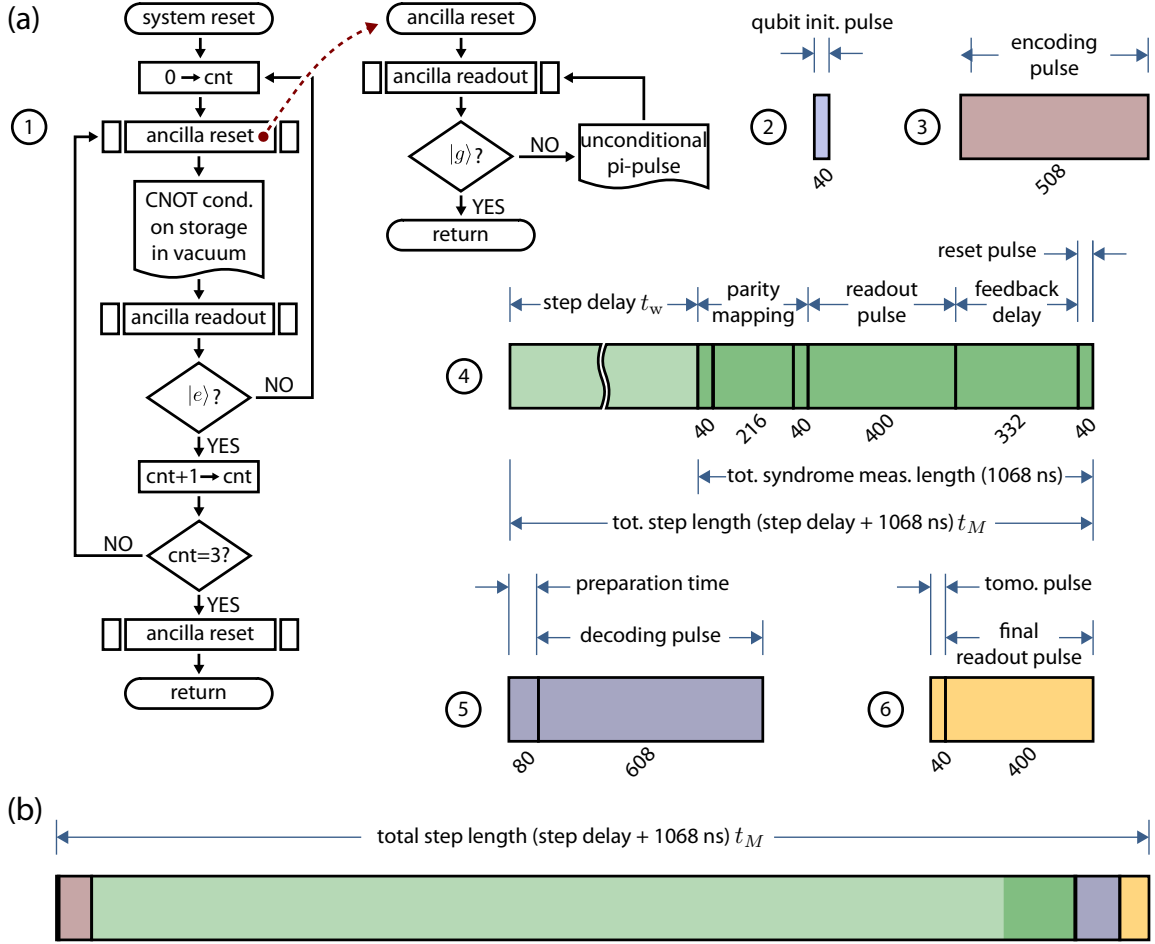


Figure 6.3: Full QEC system experimental flow. (a) The six steps of the QEC protocol using a standard flow-chart convention. 1. System reset: we use a long selective π pulse on the ancilla ($\sigma \sim 600$ ns) that addresses only the Fock state $|0\rangle_f$ of the cavity [Schuster et al., 2007] to verify that it is indeed in vacuum. In order to boost our confidence, we require three consecutive verifications to trust the results (counter “cnt” must be incremented from 0 to 3 for the process to continue). We then perform the ancilla reset protocol by measuring it and applying a short pulse ($\sigma = 2$ ns) to return to $|g\rangle$ if it is found to be in $|e\rangle$. 2. Ancilla initialization: we apply a short pulse ($\sigma = 2$ ns) to encode the ancilla into one of 6 cardinal points on the Bloch sphere. 3. Encoding: an optimized control pulse of length 508ns transfers the quantum information from the ancilla to the cavity, leaving the ancilla in $|g\rangle$. 4. Parity monitoring: we repeat the adaptive parity monitoring protocol. Each monitoring step begins with a delay of some duration, followed by the right Ramsey-like sequence that maps the ancilla back to $|g\rangle$ if there was no photon jump during the delay. We then measure the ancilla; if we find it in $|e\rangle$, we reset it as soon as possible. This happens 332 ns from the moment the readout pulse ends (200 ns of FPGA calculation latency, plus experimental delays such as finite microwave cable lengths). 5. Decoding: after a short delay to finalize the estimation of current state in the cavity, the decoding pulse is chosen in real-time and is played with a best estimate of a corrected cavity phase to account for Kerr-induced phase rotation for non-zero error cases (see sec. 6.3). 6. State tomography: measuring the ancilla after a pre-rotation to find the final density matrix of the quantum bit. (b) The whole protocol set to scale, shown to emphasize that we interrogate the system for only a fraction of the entire sequence duration.

controller applies a fast π pulse (Gaussian envelope with $\sigma = 2$ ns) to return the ancilla to $|g\rangle$ and measures again; if the pulse is not successful the loop is repeated, while if the pulse is successful the experiment continues. With feedback latencies of just ~ 200 ns (last sample in, first sample out), a readout pulse duration of $\tau_{\text{meas}} \approx 400$ ns, and latencies due to cables into and out of the experimental setup totaling ~ 100 ns, we are able to reset the ancilla to $> 99.8\%$ in $|g\rangle$. This protocol was also demonstrated in [Ristè et al., 2012].

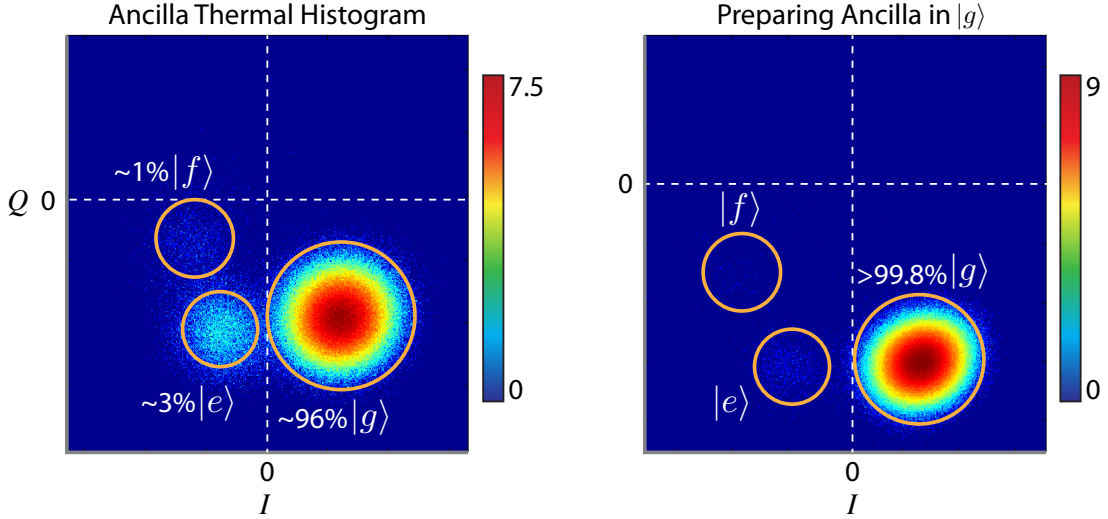


Figure 6.4: **Resetting the ancilla to the ground state.** Shown are 2D histograms of single-shot measurements of the ancilla transmon, where the logarithm (base 2) of the number of counts per bin is given by the color. The numerical values along the x and y axes are the result of the quantum controller’s integration of the demodulated in-phase (I) and quadrature (Q) readout signals, and are thus in scaled voltage units. The value of this scaling factor is determined by details of the integration implementation in the FPGA and are thus presented as arbitrary here. In thermal equilibrium (5×10^5 shots), the ancilla is measured to be in $|g\rangle \sim 96\%$ of the time, $|e\rangle \sim 3\%$, and states $|f\rangle$ and higher $\sim 1\%$; orange circles outline where the majority of counts for each state lie. After running the reset protocol (1×10^6 shots), the ancilla is measured in $|g\rangle$ more than $\sim 99.8\%$ of the time.

Second, we use the now initialized ancilla to project the cavity state to the vacuum by applying long π pulses on the ancilla (Gaussian envelope with $\sigma = 600$ ns $> 1/\chi_{sa}$) that address only the $|0\rangle_f$ Fock state [Schuster et al., 2007]. If the result of a subsequent ancilla measurement is $|e\rangle$, with high probability the cavity is in vacuum. These pulses, however, have a lower fidelity ($\sim 90 - 95\%$) owing to the ancilla T_2 , and so we repeat this experiment until we measure $|e\rangle$ three times consecutively. Once this occurs, we once again employ the protocol above to reset the ancilla to $|g\rangle$ and continue to the encoding step. With a thermal population of $\sim 2\%$ and a lifetime of $250 \mu\text{s}$, the average rate of excitation of the cavity from $|0\rangle_f \rightarrow |1\rangle_f$ is on the order of 10 ms, and so we are unable to measure any residual population in $|1\rangle_f$ after this purification procedure.

Encoding a logical qubit

We employ optimal control pulses to encode and decode logical states in the cavity based on the Gradient Ascent Pulse Engineering (GRAPE) algorithm originally developed for pulse sequences in NMR spectroscopy [de Fouquieres et al., 2011; Khaneja et al., 2005]. This algorithm is designed to numerically find a set of pulses that most accurately realizes a unitary operation or state transfer, taking an initial state $|\psi(t=0)\rangle$ to a final state $|\psi(T)\rangle$. We define the fidelity of the simulated state transfer F_{oc} to be:

$$F_{oc} = \frac{1}{K^2} \left| \sum_k^K \langle \psi_k(T) | \psi_k^{\text{tar}} \rangle \right|^2, \quad (6.2)$$

for a target state $|\psi^{\text{tar}}\rangle$, where K is the total number of state transfers we wish to realize. In order to model the physical limits in output amplitude imposed by our electronics hardware, we add an amplitude constraint of the form $\lambda \sum_{t=1}^T e^{(a_t/h)^2}$, where λ is an overall scaling; a_t is the amplitude at each point in time of the pulse (discretized into 1 ns steps); and h is an amplitude threshold, which we choose to be slightly below the maximum output amplitude our waveform generators can produce. This penalty term turns on sharply when a_t reaches h . The scaling factor λ is a proportionality constant that makes the total penalty much smaller than 1 for pulses that have all amplitudes below h . We also include a derivative penalty to give preference to smoother pulses, similarly defined as $\lambda_d \sum_{t=1}^T e^{(a_t - a_{t-1})^2 / h_d^2}$. With such a term included, the simulation favors pulses with changes smaller than h_d between neighboring control points. The criterion we enforce is that F_{oc} must exceed a value typically set to be 98%, although this constraint is relaxed when the overlap of basis states becomes non-negligible (see sec. 8.2).

In our implementation, we expand the Hamiltonian introduced earlier (Eq. 6.1) to include higher order terms, which may become more influential in dictating the time evolution of the state during the pulse:

$$\hat{H}/\hbar = \tilde{\omega}_a \hat{b}^\dagger \hat{b} + \tilde{\omega}_s \hat{a}_s^\dagger \hat{a}_s - \chi_{sa} \hat{b}^\dagger \hat{b} \hat{a}_s^\dagger \hat{a}_s - \frac{K_a}{2} \hat{b}^{\dagger 2} \hat{b}^2 - \frac{K_s}{2} \hat{a}_s^{\dagger 2} \hat{a}_s^2 + \chi'_{sa} \hat{b}^{\dagger 2} \hat{b}^2 \hat{a}_s^\dagger \hat{a}_s. \quad (6.3)$$

We furthermore include driving terms on the ancilla and cavity of the form $\varepsilon_a(t) \hat{b}^\dagger + \varepsilon_a^*(t) \hat{b}$ and $\varepsilon_s(t) \hat{a}_s^\dagger + \varepsilon_s^*(t) \hat{a}_s$. Temporal envelopes $\varepsilon_a(t)$ and $\varepsilon_s(t)$, which specify $\hat{U}(t)$, are discretized into 1 ns pieces. It is the shape and amplitude of these envelopes that we wish to numerically optimize in order to realize the desired state transfer. More explicitly, for the encoding pulses we wish to find a

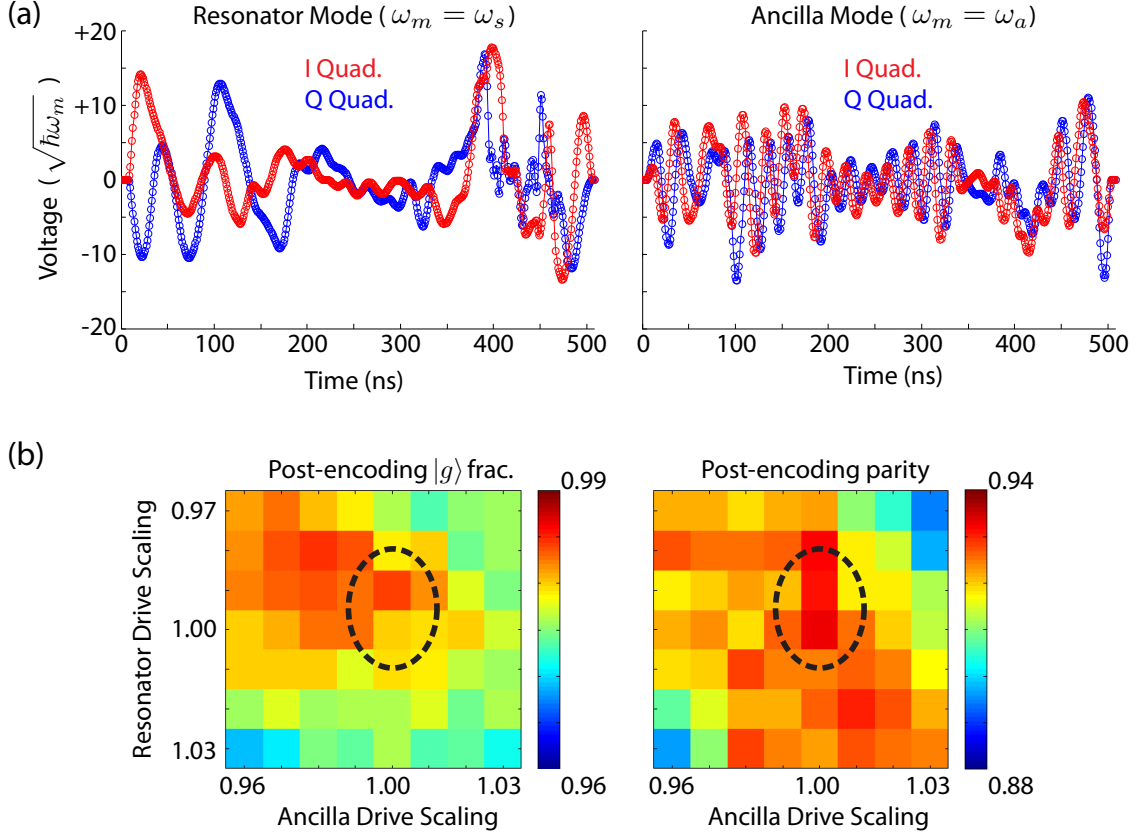


Figure 6.5: **Optimal control pulses.** (a) Example encoding pulses for cat states of initial size $\bar{n}_0 = 3$. The pulses are played by the controller at the same time, minus a 2 ns offset to account different lengths of line going down to the cavity versus the ancilla (determined experimentally by stepping relative pulse delays in 1 ns increments). The y-axis is given in units of the voltage that needs to be applied on a particular mode (frequency $\tilde{\omega}_m$) for the same amount of time to insert one quantum of energy. For the cavity, this value is determined by finding the voltage needed to create a coherent state of amplitude $|\alpha| = 1$ for a square pulse of duration 508 ns; likewise, with the ancilla it's the voltage needed to perform a π pulse in 508 ns. Mixer quadratures I and Q are shown in red and blue, respectively. (b) In executing the numerical optimization to find the optimal control pulses we do not include as simulation inputs experimental imperfections, such as frequency-dependent reflections in our microwave lines and components. When implementing the pulses experimentally, we calibrate a scaling factor on the amplitudes for both the cavity and ancilla drives by performing a 2D voltage sweep to see at which scalings we find the maximum parity of the cavity state (ideally +1), which should also coincide with the maximum occupation of the ancilla in its ground state $|g\rangle$ (ideally 100%). Shown here are images where we have already found the optimal range of drive scalings, outlined in the black dotted ellipses.

single unitary \hat{U}_{tar} such that for all c_0 and c_1 we have:

$$\hat{U}_{\text{tar}}(c_0 |g\rangle + c_1 |e\rangle) |0\rangle \rightarrow |g\rangle (c_0 |C_\alpha^+\rangle + c_1 |C_{i\alpha}^+\rangle) \quad (6.4)$$

This unitary takes a quantum bit initially stored in the ancilla (with cavity in the vacuum) to a superposition of cat states in the cavity with the same amplitudes c_0 and c_1 (returning the ancilla to $|g\rangle$). Figure 6.5a shows a set of such encoding pulses on the ancilla and cavity.




Figure 6.6: **Encoding the six cardinal points: cat code.** Using joint Wigner tomography [Vlastakis et al., 2015], we plot the experimentally obtained conditional Wigner functions ($W_z(\beta) = \langle \hat{\sigma}_z \hat{P}(\beta) \rangle$) to demonstrate the capability of a single pair of pulses to encode an arbitrary vector on the Bloch sphere. The only difference in the pulse sequence is the initial qubit preparation pulse. The encode-decode process fidelity is 93%.

Crucial to the success of finding an optimal control pulse with high fidelity is an accurate knowledge of the dominant Hamiltonian parameters. Furthermore, careful microwave hygiene at all points in the experimental chain is necessary to prevent undesired reflections and dispersions that can distort the pulse as it goes from room temperature to the setup inside the dilution refrigerator. Figure 6.5b demonstrates a calibration sequence we use to tune the amplitudes on individual ancilla and cavity drives. Ideally, the encoding pulse returns the ancilla to the ground state and creates a cat state with parity of +1. In practice, both the parity and the final ground state occupation are slightly lower than their ideal values, and are sensitive to errors in pulse power. By sweeping the relative powers for both the ancilla and cavity drives, we find the maximum ground state occupation and parity value to occur at roughly equal scalings. The full Wigner tomography shown in Fig. 6.6 illustrates visually how we do indeed have the capability to encode any arbitrary state in a superposition of cat states ($\bar{n} = 3$) with the same pulses, where the only difference lies in the preparation of the initial quantum bit. Shown in Fig. 6.7 is an example of encoding all six cardinal points in superpositions of Fock states $|0\rangle_f$, $|1\rangle_f$, respectively.




Figure 6.7: **Encoding the six cardinal points: Fock state superpositions.** These Wigner tomograms instead show all six cardinal points encoded in the Fock states $|0\rangle_f, |1\rangle_f$. Again, the only difference in the pulse sequence is the initial preparation pulse.

The use of such optimal control pulses is a departure from the theoretically proposed [Leghtas et al., 2013] and experimentally demonstrated [Vlastakis et al., 2013] qcMAP protocol, which just as the encoding pulses described above, deterministically maps a quantum bit from the ancilla onto a superposition of coherent states in the cavity. Although beautiful in its simplicity and intuitive design, qcMAP suffers from generally longer total pulse lengths than an optimal control equivalent, and additionally lacks the capacity to correct for Hamiltonian terms of higher order than the dispersive shift χ_{sa} , such as the cavity Kerr. Indeed, by comparing the two schemes side-by-side, as shown in figure 6.8, one sees that encoding a state using qcMAP requires nearly a three-fold increase in pulse duration, decreasing the final fidelity of the logical qubit due to ancilla decoherence. Likewise, one can see the clear evidence of distortion to a subset of the four coherent states, which primarily arises due to the asymmetry in the location of the coherent states during different stages of the sequence, resulting in consequent asymmetric smearing due to the Kerr. Importantly, such distortion does not imply that the information is lost; rather, the difficulty this introduces is that decoding the state back onto the ancilla using the reverse of qcMAP becomes more challenging and less successful. Using the optimal control pulses addresses this challenge by symmetrizing the cavity

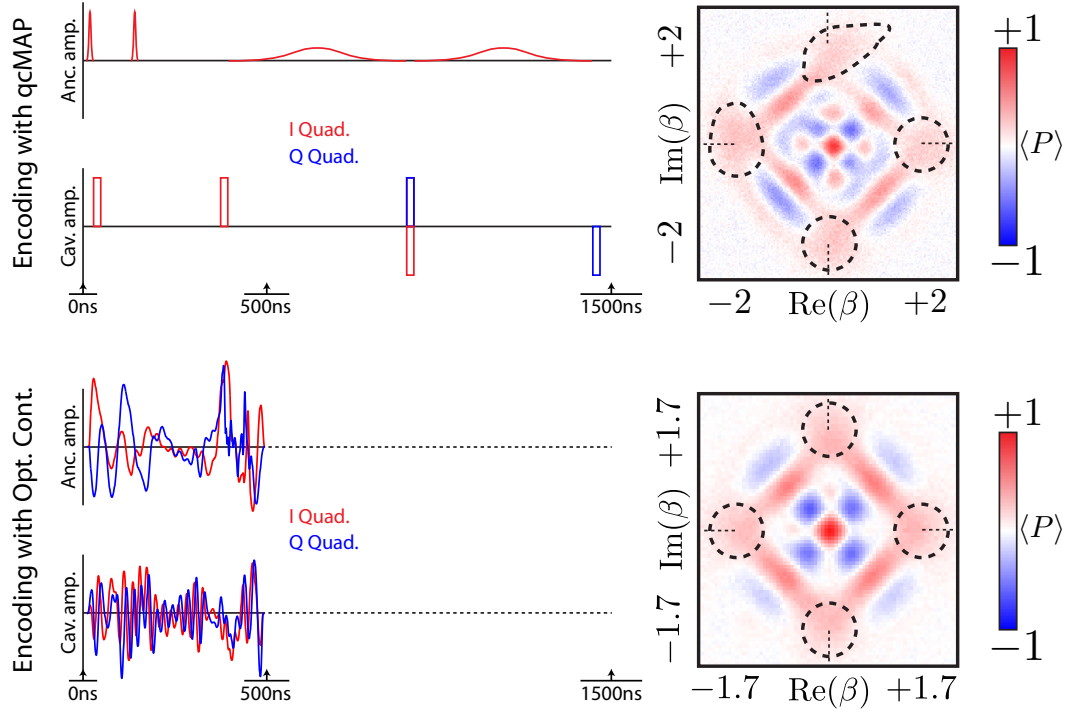


Figure 6.8: **Optimal control vs. qcMAP.** Top: a typical qcMAP sequence that maps a quantum bit onto a superposition of coherent states in the cat code paradigm. The total sequence is comprised of well-known and easily characterizable pulses, such as unselective displacements (rectangles), unselective ancilla pulses (sharp Gaussians), and ancilla pulses selective on the cavity $|0\rangle_f$ Fock state (long, broad Gaussians). Such simplicity makes qcMAP both easy to calibrate and straightforward to diagnose from the perspective of ancilla and cavity decoherence. Bottom: in contrast, the optimal control pulses are much faster and result in cat states with less distortion, as emphasized by the black dotted circles. Their transient action on the joint system is, however, difficult to intuitively grasp.

state, as seen in the Wigner function in the lower panel of Fig. 6.8. We do, however, lose the intuitive understanding of the time evolution of the joint ancilla-cavity state.

Adaptive parity monitoring

This quantum state machine implements an adaptive parity monitoring scheme in which the parity mapping protocol is updated in real-time to maximize the probability to measure $|g\rangle$ (Fig. 6.9). This substantially reduces the time the ancilla spends in $|e\rangle$, enhancing parity measurement fidelities and making it an essential component in the experimental workflow. Recalling the discussion in sec. 5.1, and more specifically Fig. 5.2, in measuring photon number parity as the error syndrome, two protocols may be used that both employ a Ramsey-style pulse sequence to map opposite parities to opposite poles of the ancilla Bloch sphere; they differ only in the sign of the second $\pi/2$ pulse. For example, when the cavity starts in an even parity cat state and the sign of the second $\pi/2$ pulse is

positive, the ancilla ends up in $|e\rangle$ at the end of the protocol, while if the parity is odd, the ancilla ends up in $|g\rangle$. Likewise, if the parity is even but the sign of the second pulse is negative, the ancilla again ends up in $|g\rangle$. Thus, in implementing our QEC system, simply repeating just one of the two protocols during the error syndrome monitoring does not suffice, since with either one the ancilla spends much more time in the excited state for one of the two parities. This asymmetry provides a strong motivation for using real-time feedback.

The prescription is simple: assuming the parity of cavity state is known, one must employ the pulse sequence that maps the ancilla to $|g\rangle$ for that known parity; if instead $|e\rangle$ is measured, one should assume that a photon has jumped and in the following parity measurement flip the sign of the second pulse to invert the interpretation of the subsequent mapping (Fig. 6.9). With such a scheme, one no longer expects the flat and oscillating patterns seen in Fig. 5.4. Rather, a change in parity is inferred only from one measurement of the ancilla in $|e\rangle$.

The benefits of employing this adaptive protocol cannot be overstated. Feedback latencies of just ~ 200 ns mean that the ancilla spends just ~ 700 ns in $|e\rangle$ per error. Without feedback, this time can be far greater, perhaps as much as ~ 50 μ s per error in a 100 μ s-long experiment, effectively guaranteeing cat state dephasing with our ancilla coherence times (see sec. 8.2). As shown in Fig. 6.3, with the adaptive parity monitoring scheme and the ancilla reset described above, the full timeline of our measurement sequence is designed to have the ancilla in the ground state as much as possible. We therefore regard the role of the quantum controller to be crucial to our goal of realizing a QEC scheme without the use of any post-selection or corrections for measurement inefficiencies.

Error history-dependent state rotation

The adaptive monitoring protocol is supplemented with the addition of a second application of real-time feedback: an instruction to record not just the occurrence, but the time at which an error occurs. The necessity of this feature stems from the non-commutativity of the Kerr Hamiltonian $\frac{K_s}{2}\hat{a}_s^\dagger{}^2\hat{a}_s^2$ and the dissipation operator \hat{a}_s , which leads to an extra undesired effect of random photon jumps: a phase shift of the cavity state in phase space that is proportional to K and the time at which the jump occurs, t_j : $|\psi(t)\rangle \xrightarrow{\hat{a}_s} e^{i\theta\hat{a}_s^\dagger\hat{a}_s}(\hat{a}_s|\psi(t))$ [Leghtas et al., 2013]. When the difference in time between syndrome measurements t_w is non-zero, the uncertainty in jump time grows with increasing t_w , leading to the aforementioned dephasing. For n total syndrome measurements spaced by t_w there is, however, still a known average angle of rotation for a jump that is measured at step j : $\bar{\theta}_j = K_s(j - 1/2)t_w$. In other words, $K_s(j - 1/2)t_w$ is our best estimate of t_j given the measurement

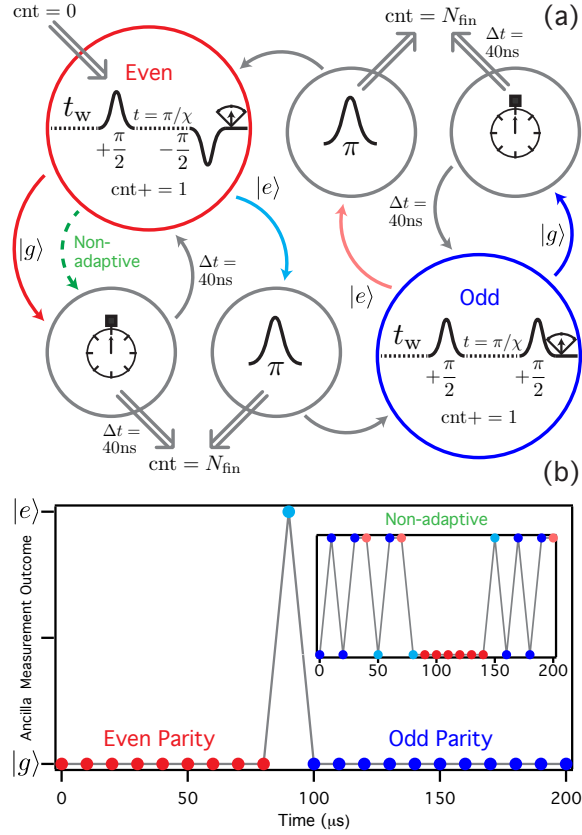


Figure 6.9: State machine for adaptive error monitoring. (a) This state machine implements an adaptive parity monitoring scheme. A parity mapping is comprised of two $\pi/2$ pulses separated by a waiting time of $t = \pi/\chi_{sa}$ [Bertet et al., 2002]. If the sign of the second pulse is positive (negative), even parity is mapped to ancilla $|e\rangle$ ($|g\rangle$) and odd parity is mapped to $|g\rangle$ ($|e\rangle$). The state machine updates the sign of the second pulse in real-time to maximize the probability to measure $|g\rangle$, reducing the probability of changes in ancilla energy during the projective measurement. Entering the state machine (double-arrow pointing inward), a counter is initialized ($\text{cnt} = 0$); the system idles for t_w ; even parity is mapped to $|g\rangle$ (red circle); the ancilla is then measured and the counter is incremented. If the result is $|g\rangle$, the system idles for 40 ns (stopwatch) and then returns to the previous state. The result $|e\rangle$ indicates a photon has likely jumped; a π pulse is then applied (Gaussian envelope; $\sigma = 2$ ns; duration 40 ns); the system thenceforth uses a pulse sequence that maps odd parity to $|g\rangle$ (blue circle). The controller returns to the initial state after another photon jump. When the counter reaches $\text{cnt} = N_{\text{fin}}$ (pre-loaded), the system exits (double-arrows pointing out). Measurement infidelities are emphasized by lighter shades of red and blue, corresponding to lower confidence when the meter measures $|e\rangle$. The non-adaptive protocol simply cycles between using a fixed parity mapping sequence and the short idling time (green dotted arrow). (b) Results. An example single-shot record of parity measurement results with $t_w = 9 \mu\text{s}$ demonstrates the difference between the adaptive and non-adaptive protocols. In the former, the ancilla is found to be in $|e\rangle$ once out of 20 measurements. With the latter, $|e\rangle$ is measured 8 times; given $t_w/T_1 \approx 0.3$, the odds of ancilla decay in this trace are so high that it is unclear how many errors occurred.

cadence. We can and must take this angle into account to prevent substantially greater excursions out of the logical subspace.

In order to do so, our controller must record the step in the monitoring at which the photon jump occurs, or equivalently the time. Then, in real-time it must apply a rotation to the coordinate

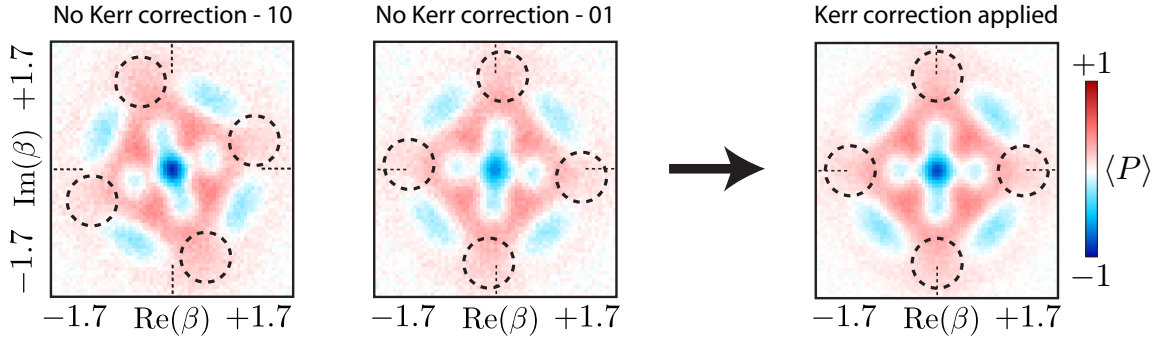


Figure 6.10: Correcting for phase rotation due to Kerr. Starting from the left, the first two Wigner tomograms show the cavity state post-selecting on measuring trajectories 10 and 01, respectively. The right-most tomogram shows just a single “1-error” state, where all 10 and 01 trajectories have been aligned in real-time and averaged together using the knowledge of the Kerr and the time at which the jump occurred. Dotted circles indicate the approximate location of the coherent states. Differences in parity and fringe contrast for the different trajectories are explained in sec. 8.1

system of the cavity’s phase space by an angle $\bar{\theta}(k) = \sum_{j=1}^k \bar{\theta}_j$ for k jumps so that the decoding pulse at the end of the sequence is applied correctly. For n monitoring steps there are $l = \frac{n!}{k!(n-k)!}$ different combinations of jump times for $k \leq n$ errors, and thus the controller must individually align all l error trajectories that correspond to k photon jumps on top of one another. For example, when $n = 3$ one can have $2^n = 8$ different monitoring outcomes ($0 \equiv$ “no error” and $1 \equiv$ “error”): 000, 100, 010, 001, 110, 101, 011, 111; in this case, the feedback rotates 100 by $\bar{\theta}_1 = K_s t_w / 2$, 010 by $\bar{\theta}_2 = 3K_s t_w / 2$, and 001 by $\bar{\theta}_3 = 5K_s t_w / 2$ so all three can be decoded with a single pulse. Prior to decoding, the feedback takes just 100 ns to align all frames of reference to the orientation defined at the outset, consolidating trajectories of equal error number yet different error timestamp into a single effective cavity state. Figure 6.10 shows an example of how such a correction succeeds in aligning all trajectories to the same phase for two tracking steps. The Wigner function on the left (center) shows the cavity state after an error history of 10 (01) and the right-most tomogram shows how these two can be aligned and averaged together not in post-processing but rather in real-time. This aspect of the feedback highlights the complexity of the calculations that the controller does in real-time. Furthermore, it can in principle handle an unlimited number of steps; as the number of combinations of jump times grows exponentially, it is a testament to the capability of the logic to efficiently perform and store the results of such calculations. In the future, when measurement rates become much faster, this will be an indispensable feature.

Adaptive quantum bit decoding

The decoding pulses simply reverse the encoding described above, up to a modification that accounts for the effects of amplitude damping and Kerr evolution. A single decoding pulse, however, cannot take two cavity states of different parity back to the same state of the ancilla since a unitary operation cannot bring two orthogonal states to a single state. For a given monitoring time we thus prepare two sets of decoding pulses, one for even states and one for odd states. Based on the parity of the final state, the controller decides in real-time to apply one of two decoding pulses, depending on even or odd parity, to map the cavity state back onto the ancilla. Although this feature is simple to implement, it is in some sense the most crucial; applying the wrong pulse does not disentangle the cavity and ancilla at the end of the decoding, leading to a completely incoherent quantum bit mixture when tracing over the cavity state.

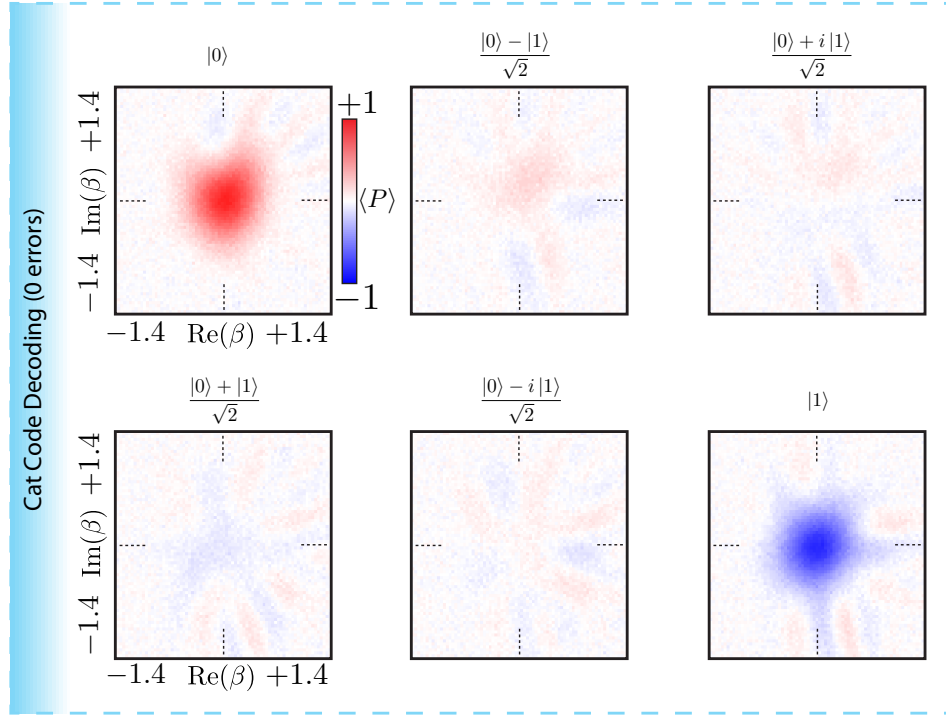


Figure 6.11: **Decoding the six cardinal points: cat code.** These Wigner tomograms show the action of cat state decoding pulses that immediately follow the encoding shown in the first panel. The first and sixth tomogram demonstrate that we map the “2-cat” along the real axis back to vacuum with the ancilla in $|g\rangle$ ($W(\beta) = +\frac{2}{\pi} \langle \hat{D}(\beta) \hat{P} \hat{D}(\beta)^\dagger \rangle$) and a “2-cat” along the imaginary axis back to vacuum with the ancilla in $|e\rangle$ ($W(\beta) = -\frac{2}{\pi} \langle \hat{D}(\beta) \hat{P} \hat{D}(\beta)^\dagger \rangle$). The difference in sign (and consequently color, red vs. blue) between the two is a result of different ancilla initialization ($|g\rangle$ vs. $|e\rangle$) prior to Wigner tomography. The remaining four joint tomograms should ideally have no visible features for perfect encoding and decoding since $\langle \hat{\sigma}_z \rangle = 0$. Experimental imperfections and primarily ancilla decoherence, however, result in residual cavity-ancilla entanglement at the end of the sequence and thus slightly visible interference features in the Wigner functions.

After monitoring errors for arbitrary lengths of time the decoding pulse must also take into account the substantial deviations of the coefficients in the Fock state expansion of the basis states from their original Poisson values. We thus use these pulses to remove any distortions in the cavity state due to the deterministic action of the Kerr Hamiltonian and deterministic amplitude damping due to energy decay. For example, after a monitoring time T the decoding pulse for even parity realizes the following state transfer:

$$\begin{aligned}
& |g\rangle \{e^{-i\frac{K_s}{2}\hat{a}_s^{\dagger 2}\hat{a}_s^2 T} \mathcal{N}_\alpha^+ [c_0(|\alpha(T)\rangle + |-\alpha(T)\rangle) + c_1(|i\alpha(T)\rangle + |-i\alpha(T)\rangle)]\} \\
& \rightarrow (c_0 |g\rangle + c_1 |e\rangle) |0\rangle_f
\end{aligned} \tag{6.5}$$

For the data in which we perform many monitoring steps to study the decay of the process fidelity in time, presented in sec. 7.3, we require a different pair of decoding pulses for each of the nine data points in the plot. The feedback controller stores these in memory and applies them at the appropriate time. These results demonstrate that beyond offering the convenience of fast encoding and decoding that take into account distortions due to higher order Hamiltonian parameters, optimal control pulses provide a striking example of the levels of control possible with continuous-variable systems in a cQED framework.

6.4 The debugger view: QEC in action

With each step of the full QEC system detailed, the stage is now set to see how all of the pieces fit together with a debugger view of two error syndrome steps over $\sim 28 \mu\text{s}$ of monitoring, as seen in Fig. 6.12. With this approach we check at every step of the program whether the pulses we apply and measurement results we obtain correspond to the state predicted by the cat code and our knowledge of the Hamiltonian parameters. We seek to not only understand the central figure of merit in the work, the final process fidelity, but moreover whether the parity, changes in cavity phase, and jump statistics all match what we expect. By finding excellent agreement with simulation across many metrics, we can be confident that the behavior coincides with our understanding of the full QEC system.

The density matrix, ρ_{init} , is initialized in one of the six cardinal points and then encoded onto a superposition of coherent states in the cavity. Wigner tomography is obtained by direct measurements of the cavity Wigner function in the continuous variables basis [Lutterbach and Davidovich, 1997] to give a snapshot of the cavity state at any point in time. In this example, only one of the six cardinal points, $\frac{1}{\sqrt{2}}(|C_\alpha^+\rangle - |C_{i\alpha}^+\rangle)$, is depicted. As intended, the measured parity of the initial state

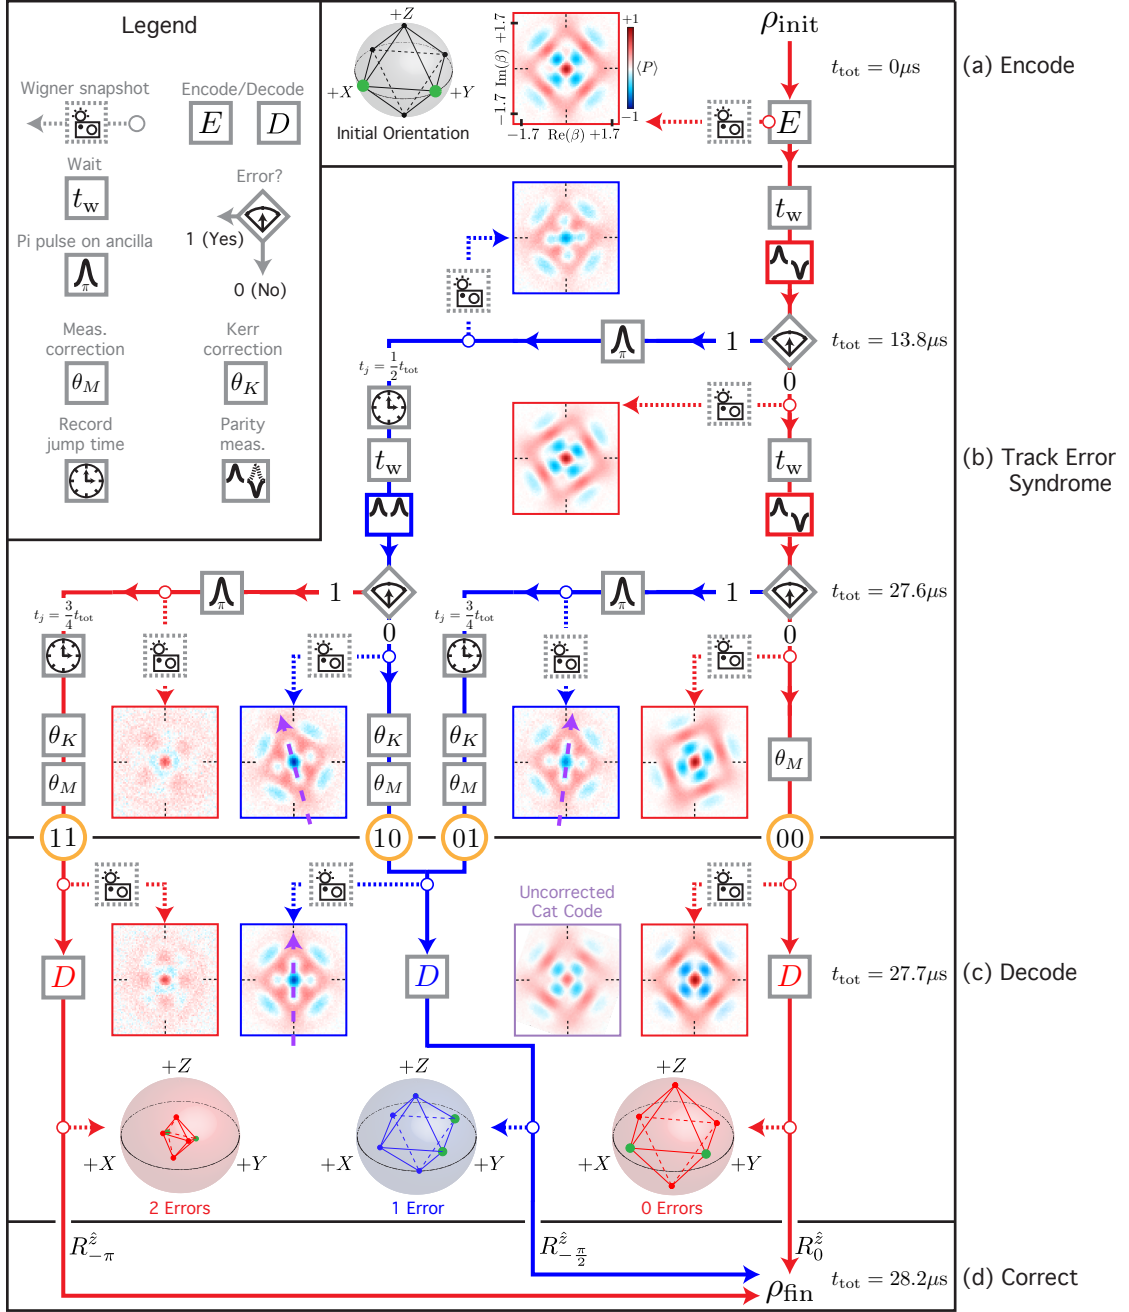


Figure 6.12: **Example of a two-step quantum trajectory executed by the QEC state machine.** (a) Green markers indicate the initial coordinate system orientation. A “Wigner [tomography] snapshot” is shown for an initial state $\frac{1}{\sqrt{2}}(|C_{\alpha}^{+}\rangle - |C_{i\alpha}^{+}\rangle)$; $\bar{n}_0 = |\alpha|^2 = 3$. (b). Parity measurement rectangles: pulse patterns that map even (odd: dotted line) parity onto ancilla $|g\rangle$ ($|e\rangle$); diamonds: branching on the ancilla measurement result ($0 \rightarrow$ no error, $|g\rangle$; $1 \rightarrow$ error, $|e\rangle$); π pulse rectangle: ancilla reset; clock icon: record jump time t_j ; purple arrows demonstrate real-time phase rotation capabilities. Deterministic rotations θ_M are due to cross-Kerr interactions between the readout and storage cavities during projective measurements. (c) A real-time decision is made to apply decoding pulses for even (red D), or odd (blue D) parity pulses. (d) Correction step returns ρ_{fin} , with rotations by 0 (0 errors), $-\pi/2$ (1 error), and $-\pi$ (2 errors) made in software.

starts off even, and the collective fringe pattern seen in the tomogram closely matches the expected form as seen in simulation for an initial average photon number $\bar{n} = 3$ in the cavity. Remember furthermore that the Wigner function is equivalent to the density matrix of the encoded state [Cahill and Glauber, 1969], so at each subsequent point where the tomogram snapshots are presented, we can verify the parity, phase, and occurrence probability of the cavity state.

The controller interrogates the error syndrome by employing the adaptive parity monitoring protocol, with delays of $t_w \sim 13 \mu\text{s}$ between each measurement. Starting in even parity, and given the low probability ($\sim 20\%$) to have an error between two consecutive syndrome measurements for this measurement cadence, the controller plays the pulse sequence that maps even parity to $|g\rangle$ and odd parity to $|e\rangle$. When an error occurs and the parity changes, the controller pulses the ancilla from $|e\rangle$ back to $|g\rangle$ and then continues monitoring errors by employing the opposite protocol, which instead maps odd parity to $|g\rangle$ and even parity to $|e\rangle$. Therefore, again, throughout a single measurement trajectory, counting the number of errors amounts to just counting the number of times $|e\rangle$ is measured.

One may immediately ask where this seemingly arbitrary choice of delay t_w comes from, as typically the natural premise in error correction is that errors are best suppressed with a higher measurement cadence. Such an assumption overlooks, however, the potentially destructive effect of the syndrome measurement itself, most notably through the potential of forward propagating errors from the ancilla into the code space. Indeed, as outlined above, this consideration is central to motivating several implementations of the real-time feedback meant to maximize the occupation of the ancilla in the ground state. As discussed in 5.4, entangling the ancilla with the cavity too frequently instead leads to a net loss in the performance of the system. Despite the greater statistics and consequently the improved estimate of the photon number parity in time that come with an increased measurement rate, the phase information of the cavity state increasingly runs the risk of being lost if the ancilla undergoes T_1 decay. The choice of delay time can therefore be motivated by the need to strike a balance between dephasing due to missing photon jumps and the probability of forward propagation of errors, as elaborated upon in greater detail in upcoming chapters (see sec. 8.3).

The program stores in memory a single measurement record of 0s (no error) and 1s (error) that specifies the monitoring history of the encoded state; Fig. 6.12 shows the four possibilities for two steps: $\{00, 01, 10, 11\}$. Conditioned on obtaining one of these four records, each tomogram provides a striking visual demonstration of the cat code in action. Interference fringes, signatures of quantum coherence [Haroche and Raimond, 2006], continue to be sharp and extremal as the

program proceeds in time, as compared with the case of performing no parity monitoring at all. Indeed, at each point in the program, the tomograms agree well in parity contrast, phase, and amplitude as seen in simulations. More specifically, with each case the quantities of interest can be predicted with high accuracy: the parity (origin of the Wigner function) matches the controller’s best estimate at any time (border color); the difference in phase between each trajectory matches the expected extra rotation $\theta_K \approx 20^\circ \times t_j$ for every jump occurring at time t_j , emphasized by the purple arrows for cases 01 and 10; $\theta_M \approx 20^\circ$ of cavity phase rotation due to measurement cross-Kerr arising from the introduction of photons into the readout cavity; and finally the jump probabilities $\{70.4\%, 13.7\%, 11.8\%, 4.1\%\}$ of each trajectory $\{00, 01, 10, 11\}$, which can be accurately predicted with a Bayesian analysis by knowing the parity measurement infidelity and probability of photon loss (see sec. 8.1). Furthermore, ancilla state tomography after decoding (Fig. 6.12c) conditioned on the number of errors returns octahedrons similar to the one in the initialization step. The rotation of the six cardinal points by $\pi/2$ for one error and π for two errors indicates that the cat code successfully maps photon loss errors in the cavity onto a unitary operation on the logical qubit. Upon completion of the program’s execution, the knowledge of how many errors occurred is equivalent to having corrected the state. Although aligning the Bloch spheres of all error trajectories to the same orientation requires a simple phase adjustment on the ancilla drive in the decoding sequence, here we instead choose to explicitly emphasize how the cat code maps errors to logical rotations. The program thus returns the corrected quantum bit, now stored again in the ancilla, completing the full QEC cycle.

Besides illustrating self-consistency, these results also highlight the successful implementation of each intended application of real-time feedback. Looking at the 1-error case, for example, one sees that in order to obtain the measured jump statistics, both the ancilla reset and adaptive monitoring must have been executed properly; moreover, the consolidation of 01 and 10 trajectories into a single 1-error case demonstrates that the necessary changes in decoding pulse phase were introduced; and finally, as the final octahedron exhibits the necessary $\pi/2$ rotation about Z , the correct decoding pulses based on the final best-estimate of parity must have been applied. Coordinating so many steps in the QEC system is a testament to the levels of sophistication the quantum controller can achieve. The high levels of predictability throughout the entire debugger view again highlight the advantages of the hardware-efficient scheme, as in the previous chapter: knowing the Hamiltonian parameters together with a measurement fidelity of a single error syndrome is sufficient to encapsulate the evolution of an error-corrected logical qubit.

We would also like to note here that each Wigner snapshot is a 71×71 grid of parity measurements. As shown in chapter 4, this granularity is sufficient for accurate cavity state reconstruction. After the first step of error correction, joint ancilla-cavity state tomography therefore requires about 20,000 measurements per repetition. On the other hand, the Steane code with seven physical data qubits and six ancillae (introduced in chapter 3) would instead require millions. In this sense, the hardware efficiency of the cat code translates not to just a simplified error correction scheme, but furthermore enables full state characterization after the syndrome measurements.

With a convincing demonstration of a full QEC system in hand, we now turn to performing more than just two steps of error correction and look beyond the multitude of statistics presented above to converge on a single metric for the performance of the system as a whole: the process fidelity. The next chapter is devoted to this sixth step in Fig. 6.1, which assigns just a single number to the system and serves to benchmark its performance against the available passive schemes for the storage of quantum information.

7 – Reaching Break-Even

The ability of QEC to suppress errors empowers a quantum computer to perform larger and more complex algorithms. If one implements an error correction system within the framework of a fault-tolerant architecture, in which the performance of all the components also exceeds the architecture's threshold, the error in a fixed-length computation can be reduced exponentially with only a polynomial overhead in hardware. Challenges abound, given the complexity of a typical QEC circuit, the stipulation that errors do not propagate, and the demanding error thresholds [Fowler et al., 2012; Steane, 1996]. In particular, recalling the discussion in sec. 3.2, an uncorrected logical qubit consisting of n physical qubits (for typical first order codes $n \sim 5 - 10$ [Steane, 1996]) incurs a decoherence penalty, wherein the decoherence is of order n times faster. Theoretical models estimate that gate infidelities on the physical qubits must be on the order of 10^{-4} for traditional QEC schemes to be successful. Works that demonstrate performance approaching these levels [Barends et al., 2014; Chow et al., 2012] do so in simplified settings, with few physical qubits in small and well-controlled microwave environments. As implementations such as these will eventually require substantially greater resources, it isn't at all clear that currently reported gate fidelities can be sustained as systems become more complex. Indeed, symptoms of such challenges are already apparent in a recent experiment on correcting a classical bit with repeated error correction that uses quantum operations [Kelly et al., 2015]. Although the group reports CNOT fidelities ostensibly near the thresholds for fault-tolerance (on the order of 99%), employing two such gates to entangle a single physical qubit with two others results in a GHZ state with a fidelity that is completely incommensurate with their previous claims. In all likelihood, their losses come from not just the increased decoherence penalty due to the greater number of physical qubits, but also from the resource overhead that comes with the supporting architecture required to control them, which elevates the significance of perhaps previously overlooked error mechanisms. As discussed in chapter 3, the cat code offers enormous benefits in this regard.

The debugger view shown in Fig. 6.12 already demonstrates that the full QEC system implements the cat code as intended. Indeed, errors within the logical space due to single photon loss are mapped

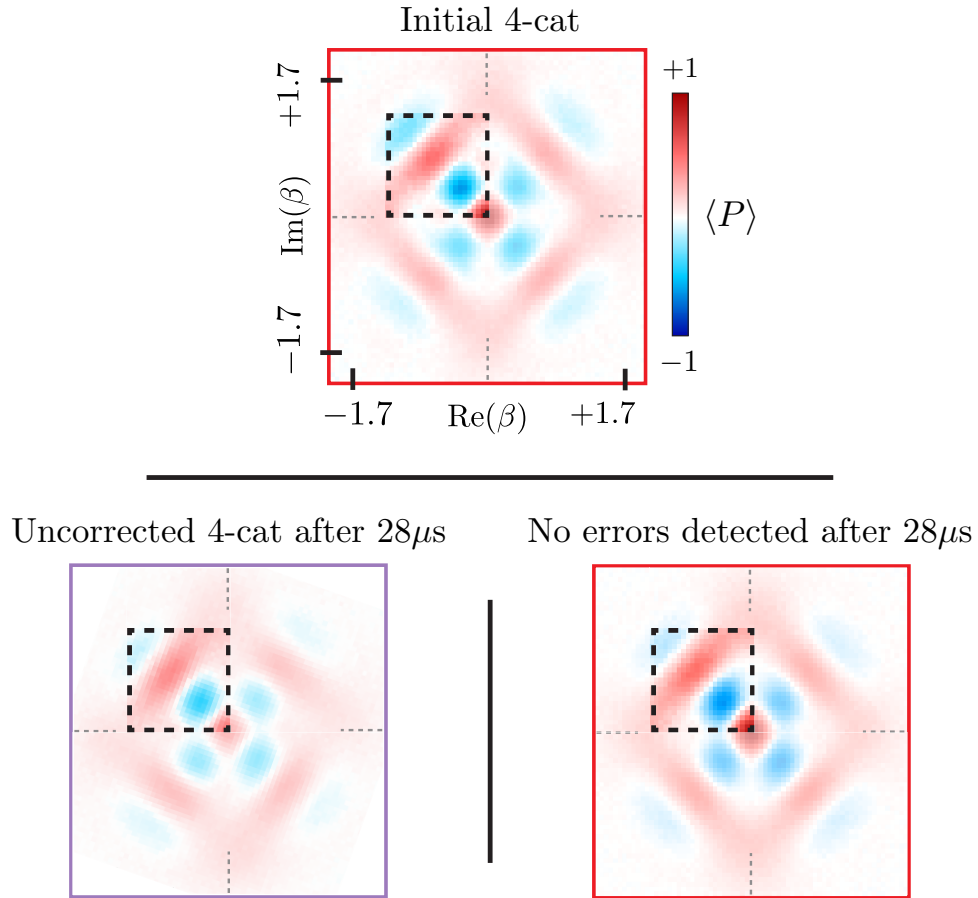


Figure 7.1: **Enhanced Wigner fringes from QEC.** These Wigner functions are enlarged versions of those in Fig. 6.12. The top tomogram shows the initial “4-cat,” the bottom-left shows the same state left to evolve naturally with no QEC for $28 \mu\text{s}$, and the bottom-right shows the same state after two error syndrome measurements that return the result record 00. The black dotted-line boxes emphasize the difference in the quality of the fringe patterns between corrected and uncorrected states. Note that the initial state and the corrected one both have diagonal fringes between the coherent states that are higher in contrast than in the uncorrected case. This example shows how coherence between the superpositions of coherent states is encoded in not just the parity, but in other key places in the cavity’s phase space.

onto discrete rotations in increments of $\pi/2$ about the logical qubit’s Z -axis per photon jump. It is also qualitatively clear that for at least some portion of the data, for example the 00 trajectory, one can already gain by employing the QEC system rather than just letting the logical qubit evolve and lose coherence on its own. This is evident from the higher coherence of the corrected cavity state, as seen in Fig. 7.1. Such improved purity is exactly what we seek in having our measurements be QND on the entire state, as discussed in sec. 5.4. We can perhaps surmise the same about the 10 trajectory, but a priori it isn’t so clear that the Wigner functions of 01 and 11 demonstrate any benefit from QEC given their less pronounced features.

The desire to understand exactly how well our QEC system performs motivates the next step and the focus of this chapter: calculating the process fidelity of the full QEC system from start to finish. Rather than deducing state fidelities from the Wigner tomograms, we instead decode the logical qubit back onto the ancilla, perform the necessary corrections to take into account the number of detected changes in parity, and simply compare the resulting quantum bit to the initial quantum bit for all possible input states. We thus quantify the performance of the QEC system in suppressing natural qubit decoherence, and moreover do so without the use of post-selection.

Observing how the process fidelity decays over time with more applications of error syndrome measurements allows us to finally answer the question posed at the end of the previous chapter. By comparing the integrity of an actively protected quantum bit to one that is redundantly encoded yet evolving freely with no interruption, we show in this chapter the first demonstration of a QEC system that actually overcomes the decoherence penalty inherent to QEC. The significance of reaching this break-even point lies in the remarkable property that a logical qubit can remain coherent for longer periods of time if aided by a supporting architecture constructed with human hands. With this result, we prove that QEC is possible in a real-world setting, and set the stage for understanding what further steps must be taken to improve performance. We conclude the chapter by launching a serious discussion on what fault-tolerance requires.

7.1 Process fidelity: an analytical treatment

We seek to benchmark the entire process of the QEC system. The density matrix of the final state, ρ_{fin} , is the output of the entire QEC system process (see sec. 6.4) $\mathcal{E}(\rho_{init})$: $\rho_{fin} = \mathcal{E}(\rho_{init})$, where $\rho_{init} = |\psi\rangle_{init}\langle\psi|_{init}$. Ideally, $\mathcal{E}(\rho_{init}) = \rho_{init}$, where the process is simply given by the identity operator \hat{I} , corresponding to perfect error correction. In reality, however, due to decoherence in conjunction with experimental imperfections, $\mathcal{E}(\rho_{init})$ is a combination of non-unitary and unitary operations on the encoded state.

In order to characterize the full process $\mathcal{E}(\rho_{init})$, we find ρ_{fin} by performing state tomography of the ancilla following the correction step to measure the components of the final Bloch vector $\vec{r} = \{r_x, r_y, r_z\}$: $\rho_{fin} = (\hat{I} + r_x\hat{\sigma}_x + r_y\hat{\sigma}_y + r_z\hat{\sigma}_z)/2$, where $\hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z$ are the Pauli operators. The results allow us to represent $\mathcal{E}(\rho_{init})$ in the chi (X) matrix representation using the operator-sum notation [Nielsen and Chuang, 2010]: $\mathcal{E}(\rho_{init}) = \sum_{jk} \tilde{E}_j \rho_{init} \tilde{E}_k^\dagger X_{jk}$, where for a single quantum bit $\tilde{E}_0 = \hat{I}, \tilde{E}_1 = \hat{\sigma}_x, \tilde{E}_2 = -i\hat{\sigma}_y, \tilde{E}_3 = \hat{\sigma}_z$ and the coefficients X_{jk} comprise the process matrix X . This is a complex 4×4 matrix of trace $\text{Tr}(X) = 1$ that completely describes the action of our QEC system on an arbitrary input state. We define the fidelity F to be the overlap of the measured chi matrix,

X^M , with X_0 , the ideal identity process: $F = \text{Tr}(X^M X_0)$. In principle, only four cardinal points are needed to determine X^M , the two at the poles of the Bloch sphere ($+\vec{z}$, $-\vec{z}$) and those along $\hat{\sigma}_x$ ($+\vec{x}$) and $\hat{\sigma}_y$ ($+\vec{y}$). Following the derivation presented in [Nielsen and Chuang, 2010], we can also find a simple formula for the $(0,0)$ entry of X^M , X_{00}^M , which is equivalent to the expression above for the fidelity to the identity process. It requires the results of state tomography, $\vec{r}^{\hat{n}} = \{r_x^{\hat{n}}, r_y^{\hat{n}}, r_z^{\hat{n}}\}$, for the four cardinal points ($\hat{n} = +\vec{x}, +\vec{y}, +\vec{z}, -\vec{z}$):

$$X_{00} = \frac{1}{4} \left(1 + \left(r_x^{+\vec{x}} - \frac{r_x^{+\vec{z}} + r_x^{-\vec{z}}}{2} \right) + \left(r_y^{+\vec{y}} - \frac{r_y^{+\vec{z}} + r_y^{-\vec{z}}}{2} \right) + \frac{r_z^{+\vec{z}} - r_z^{-\vec{z}}}{2} \right) \quad (7.1)$$

We perform these calculations with both $+\vec{x}, +\vec{y}$ and $-\vec{x}, -\vec{y}$, however, to verify that there are no unexpected asymmetries in the cat code. Figure. 7.2 shows a simple example of a X matrix for an ideal identity process and one for complete depolarization [Nielsen and Chuang, 2010], in which $r_x^{+\vec{x}}, r_y^{+\vec{y}}, r_z^{+\vec{z}}$, and $r_z^{-\vec{z}}$ all decay to 0.

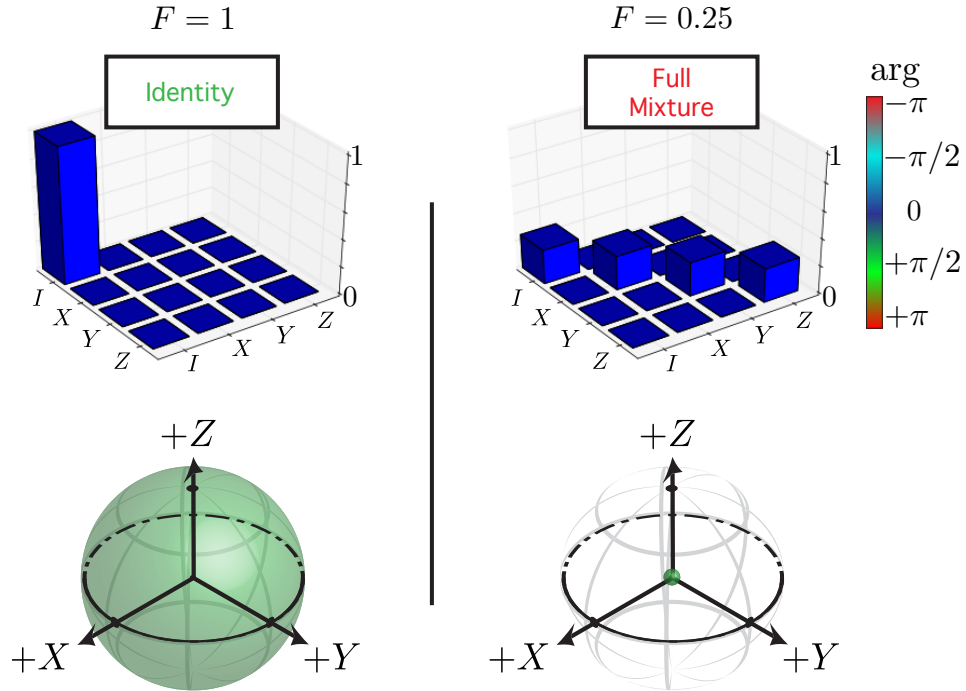


Figure 7.2: Process fidelity limits. The X matrix on the left corresponds to the identity process \hat{I} . The matrix on the right shows an example of a depolarization process that leads to a completely mixed state. Note the minimum process fidelity is equal to 0.25 rather than 0, indicating that all information about the initial Bloch vector has been lost. This simple matrix reveals that the process one has implemented ended up applying one of four operations with 25% probability: either \hat{I} , or rotations by π around one of the three Bloch sphere axes. The left Bloch sphere shows a pure state after a known unitary operation, such as \hat{I} , while the right Bloch sphere (small green sphere at the origin) shows that under depolarization all Bloch vector components go to 0.

In order to develop some intuition about this somewhat abstract formulation, we can first consider a simple case where a cavity state $\rho_s(t) = |\alpha(t)\rangle\langle\alpha(t)|$ is left to evolve without intrusion for a short time δt . If we exclude ancilla non-idealities, undesired couplings such as Kerr, and only consider the possibility of single errors within our code space, we can write the cavity density matrix $\rho_s(\delta t)$ as:

$$\rho_s(\delta t) = p_0\rho_0 + p_1\rho_1, \quad (7.2)$$

where p_0 and p_1 are the probabilities associated with obtaining density matrices ρ_0 and ρ_1 after δt . In the specific example here, ρ_0 is the state of the cavity after some amount of energy decay yet no detected jump in photon number, while ρ_1 is the state of the cavity if a photon jump does occur. The probability p_1 is just $\bar{n}_0\kappa_s\delta t$, and for a small δt , $p_0 \approx 1 - p_1$. The natural question of why a field's energy decays only when no photon jump is recorded is elegantly explained in [Haroche and Raimond, 2006], section 4.4.4. Recalling that the energy decay of a coherent state after δt takes the form $|\alpha\rangle \rightarrow |\alpha e^{-\kappa_s\delta t/2}\rangle$ and that a photon jump is equivalent to applying the lowering operator \hat{a}_s on the cavity field, we can rewrite ρ_0 and ρ_1 as:

$$\rho_0 = \frac{e^{-(\kappa_s\delta t/2)\hat{a}_s^\dagger\hat{a}_s}\rho_s(0)e^{-(\kappa_s\delta t/2)\hat{a}_s^\dagger\hat{a}_s}}{\text{Tr}(e^{-(\kappa_s\delta t/2)\hat{a}_s^\dagger\hat{a}_s}\rho_s(0)e^{-(\kappa_s\delta t/2)\hat{a}_s^\dagger\hat{a}_s})} \quad (7.3)$$

$$\rho_1 = \frac{\sqrt{\kappa_s\delta t}\hat{a}_s\rho_s(0)\sqrt{\kappa_s\delta t}\hat{a}_s^\dagger}{\text{Tr}(\sqrt{\kappa_s\delta t}\hat{a}_s\rho_s(0)\sqrt{\kappa_s\delta t}\hat{a}_s^\dagger)} \quad (7.4)$$

We see, however, that:

$$\text{Tr}(\sqrt{\kappa_s\delta t}\hat{a}_s\rho_s(0)\sqrt{\kappa_s\delta t}\hat{a}_s^\dagger) = \bar{n}_0\kappa_s\delta t = p_1 \quad (7.5)$$

$$\text{Tr}(e^{-(\kappa_s\delta t/2)\hat{a}_s^\dagger\hat{a}_s}\rho_s(0)e^{-(\kappa_s\delta t/2)\hat{a}_s^\dagger\hat{a}_s}) = p_0 \quad (7.6)$$

The state $\rho_s(\delta t)$ can thus be written as:

$$\begin{aligned} \rho_s(\delta t) &= e^{-(\kappa_s\delta t/2)\hat{a}_s^\dagger\hat{a}_s}\rho_s(0)e^{-(\kappa_s\delta t/2)\hat{a}_s^\dagger\hat{a}_s} + \sqrt{\kappa_s\delta t}\hat{a}_s\rho_s(0)\sqrt{\kappa_s\delta t}\hat{a}_s^\dagger \\ &= \sum_{k=0}^1 \hat{E}_k\rho_s(0)\hat{E}_k^\dagger, \end{aligned} \quad (7.7)$$

where $\hat{E}_0 = e^{-(\kappa_s \delta t / 2) \hat{a}_s^\dagger \hat{a}_s}$ and $\hat{E}_1 = \sqrt{\kappa_s \delta t} \hat{a}_s$. This convention, known as the operator-sum notation [Nielsen and Chuang, 2010], allows us to finally define the process for the time step δt :

$$\mathcal{E}(\rho_s(0)) = \sum_{k=0}^1 \hat{E}_k \rho_s(0) \hat{E}_k^\dagger \quad (7.8)$$

Note for small δt , $\sum_{k=0}^1 \hat{E}_k^\dagger \hat{E}_k \approx \hat{I}$, as required. Translating these error processes on the actual cavity field to those of Pauli operations on the logical qubit returns the X matrix formulation introduced above.

7.2 Demonstrating QEC after many steps of correction

With the analytical definitions of process fidelity in place, we begin quantifying the performance of the cat code by first again returning to the debugger view presented in 6.12, and more specifically to the decoding step that shows the transfer of the quantum bit from the superposition of cat states to the ancilla transmon. Shown in Fig. 7.3 are the process matrices for the QEC program demonstrated in Fig. 6.12. Just by looking at these images and the associated fidelities we can already glean some important information about the QEC process we implement. Firstly, the form of X_j^M (X^M for $j = 0, 1$, and 2 errors) matches the process matrix for ideal rotations about the Z axis by $j \cdot \pi/2$, $X_{j\pi/2}$. Secondly, the fidelities are lower for higher detected error numbers, an expected trend given the non-zero syndrome measurement infidelity. Signatures of developing incoherent mixture are evident from the non-zero values in all diagonal elements, which approach 0.25 for a fully mixed state. In fact, the equal height of the diagonal elements of X^M that should ideally be zero suggests that depolarization is the dominant error channel in the system; the following section explores this in greater detail. Trajectories with two errors, however, have a substantially lower fidelity and have a number of non-negligible off-diagonal entries in X^M , which suggests that the decoding pulses were frequently applied at incorrect angles due to misinformation from errant syndrome measurements. Although at first perhaps surprising, this feature can in fact be explained by the low confidence that a record with consecutive error detections (...11...) faithfully reflects the true error history of the encoded state (discussed in greater detail in the following chapter, sec. 8.1). Such trajectories occur just 4.1% of the time after 28 μs , however, and thus have little bearing on the final output. Upon correction, the resemblance of X^M to X_0 reflects the simplicity of the cat code; by accounting for single photon jumps we witness no dominant processes besides \hat{I} and emerging depolarization within the logical space.

We note that after the decoding step we do perform a small correction to the data in software prior to calculating X_M . Given that the optimal control encoding and decoding pulses do not realize the intended unitary perfectly at each time step, there could be some overall unintended rotation of the final quantum bit. Following the approach in [Schindler et al., 2011], we allow for one and the same rotation to be applied in software to all six cardinal points simultaneously that maximizes the overlap of X^M with X_0 . This is a simple change of reference frame that in no way compensates for measurement infidelity or an artificial enhancement of performance. It is equivalent to applying a fixed pre-determined unitary operation on the decoded quantum bit to adjust its orientation that most closely matches that of the initial state. In practice, rotations do not exceed several degrees around each axis and have only a small effect on the reported results.

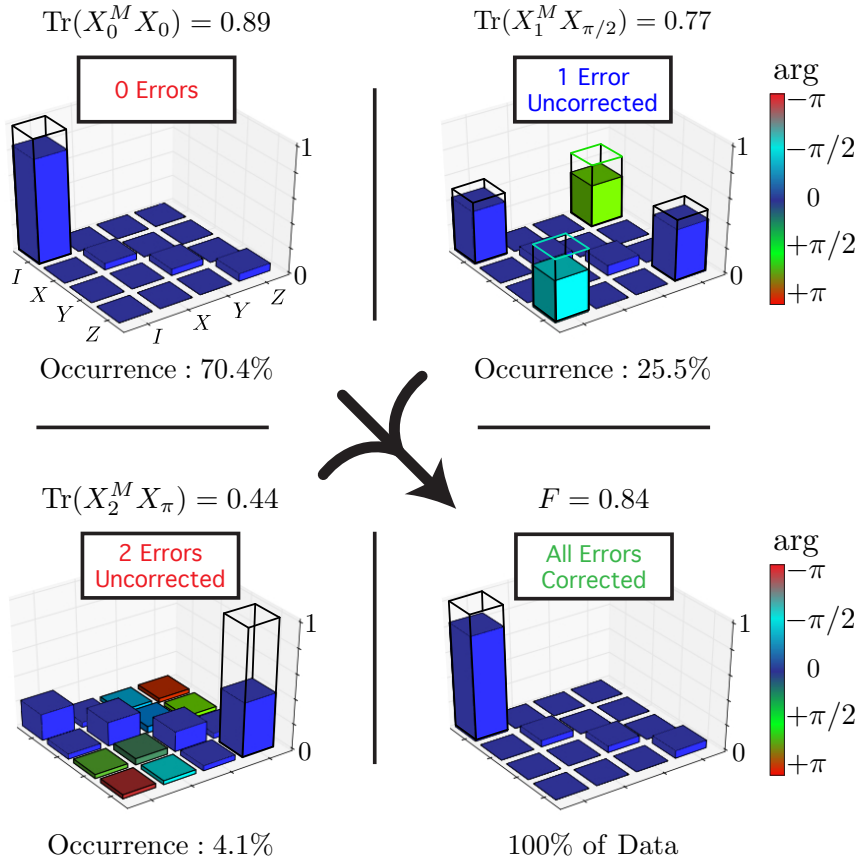


Figure 7.3: **Process tomography for two monitoring steps.** Results for $j = 0, 1$, and 2 errors prior to correction. Experimental data for X_j^M is shown in solid bars, and is obtained using the results of state tomography conditioned on error number as shown in the debugger view of the full QEC system (Fig. 6.12); the values are complex numbers with amplitude on the vertical axis and an argument specified by the bar color. Amplitudes less than 0.01 are not depicted. Ideal $X_{j\pi/2}$ process matrices are shown in wire-outlined bars. Process tomography after correction is shown following the arrow, where the final fidelity $F = 0.84$. Non-zero diagonal elements in the X matrix indicate developing incoherent mixture in the final state. Note that there is no post-selection of data.

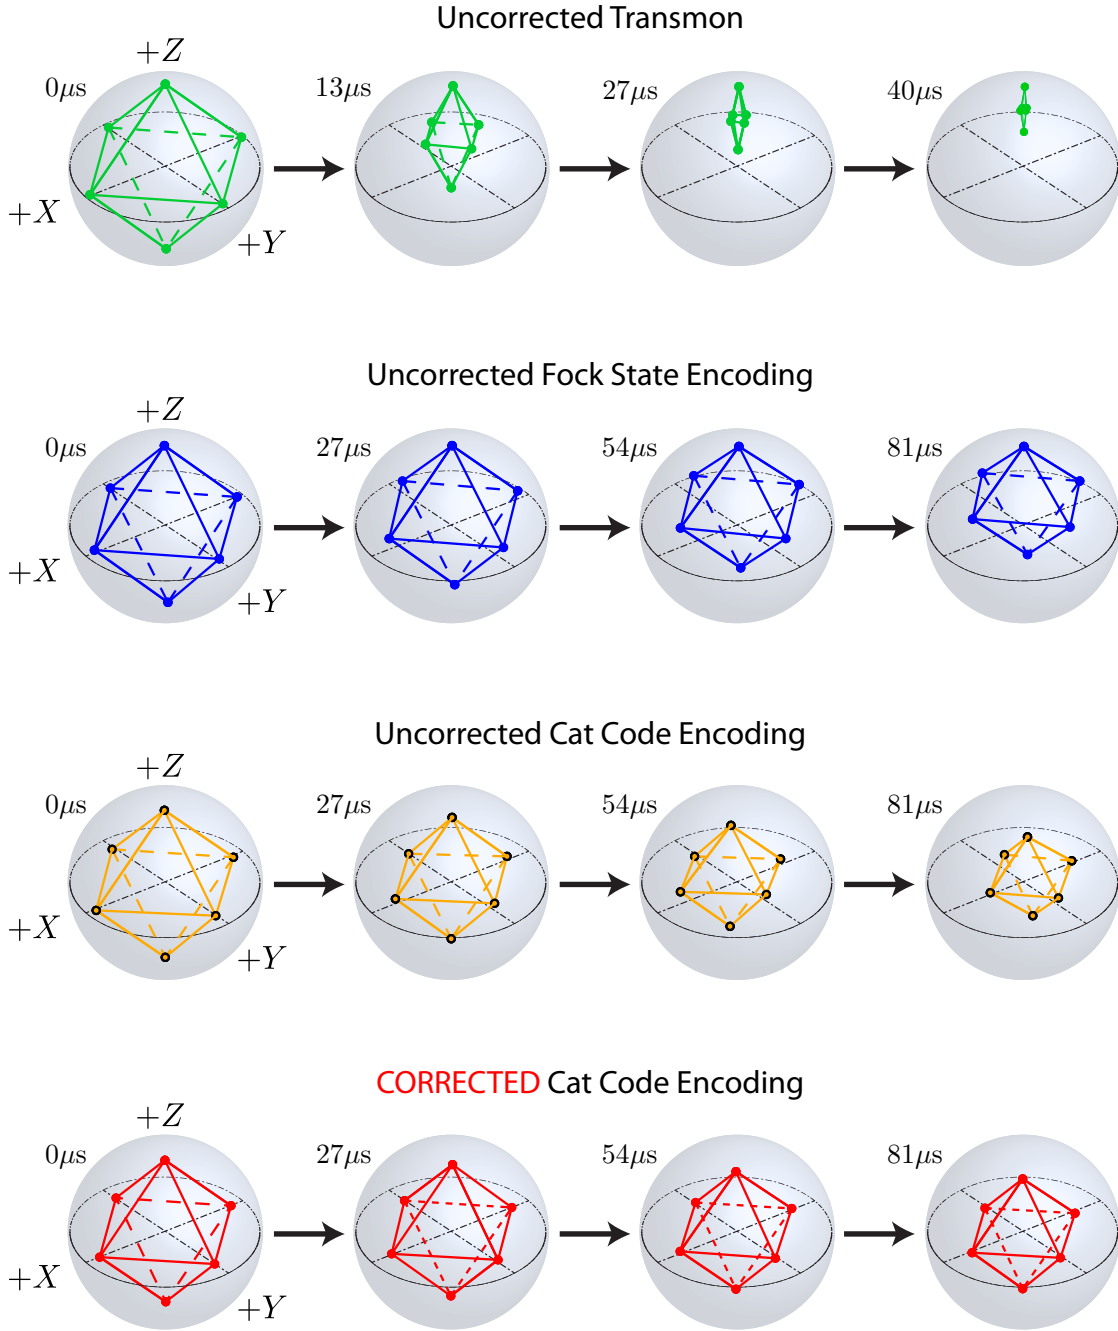


Figure 7.4: **Comparing possible quantum memories.** State tomography, after the decoding pulses that transfer the information from the cat states onto the ancilla transmon, is shown for four examples of quantum memory implementations in our system. The uncorrected transmon and Fock state encoding show the expected trends characteristic of amplitude decay; the transmon in particular, shown on a different time-scale, decays very quickly. In the cases of the uncorrected and corrected cat codes ($\bar{n}_0 = 2$), the octahedrons decay toward fully mixed states at the origin. The six cardinal points in the latter, however, clearly exhibit greater contrast, indicative of the benefits of active error correction. The uniform shrinking of each point in time demonstrates that the loss of fidelity is well-represented by a depolarization channel.

The final fidelity $F = 0.84$ after two steps of monitoring for $28 \mu\text{s}$ tells us nothing about how the process fidelity decays over time. Our feedback controller, however, can in fact perform an unlimited number of consecutive syndrome measurements while still maintaining all aforementioned feedback capabilities. Moving to an initial encoding size of $\bar{n}_0 = 2$ in order to reduce the probability of errors in the code space, we implement the cat code with multiple repetitions of the QEC program's monitoring step over $\sim 110 \mu\text{s}$, starting with zero time delay where decoding immediately follows encoding. Each syndrome measurement is separated by an optimized waiting time ranging from $t_w \approx 15 \mu\text{s}$ to $t_w \approx 25 \mu\text{s}$, which as briefly described in sec. 6.4 balances contributions from competing decoherence mechanisms (see the following chapter, sec. 8.3 for a derivation). In setting the stage for a discussion of what it means to overcome the decoherence penalty introduced by the redundant encoding, we compare the decay of the quantum bit protected by the cat code to a quantum bit stored using three different passive quantum memories: an uncorrected transmon, an uncorrected Fock state encoding (the system's most coherent piece of hardware), and the cat code with no QEC. As shown in Fig. 7.4, the six cardinal points in each of the four storage schemes decay in time. Even without quantifying these results yet, we can already visually see that the contrast of the octahedron corresponding to the corrected cat code clearly exceeds that of the uncorrected encoding at each point in time. This indicates that the photon number parity measurements do indeed offer benefits over simply allowing the cat states to undergo uninterrupted decay. We include the octahedron corresponding to the transmon to underscore how rapidly it decays toward a single point as compared to any storage scheme in the cavity. This finding is of course certainly expected given the disparity in coherence times between the two. The significance of this result, however, lies in the realization that even by using an ancilla with a T_2 that is a factor of ~ 25 lower than that of the cavity, applying QEC with the cat code nonetheless increases the lifetime of the uncorrected logical qubit.

Discerning a qualitative difference between the performance of the Fock states and the cat code, however, is difficult. Beyond the fact that the final octahedra for the two cases look to be similar in contrast, the encoding penalty for the cat code is slightly inferior due to the greater complexity of the initial states (and thus longer length of the optimal control pulses), and the decoherence processes are clearly very different as well. In particular, both the Fock state and the transmon encodings are susceptible to generalized amplitude damping, which is given by the following process [Nielsen and Chuang, 2010]:

$$\mathcal{E}(\rho) = E_0\rho E_0^\dagger + E_1\rho E_1^\dagger + E_2\rho E_2^\dagger + E_3\rho E_3^\dagger \quad (7.9)$$

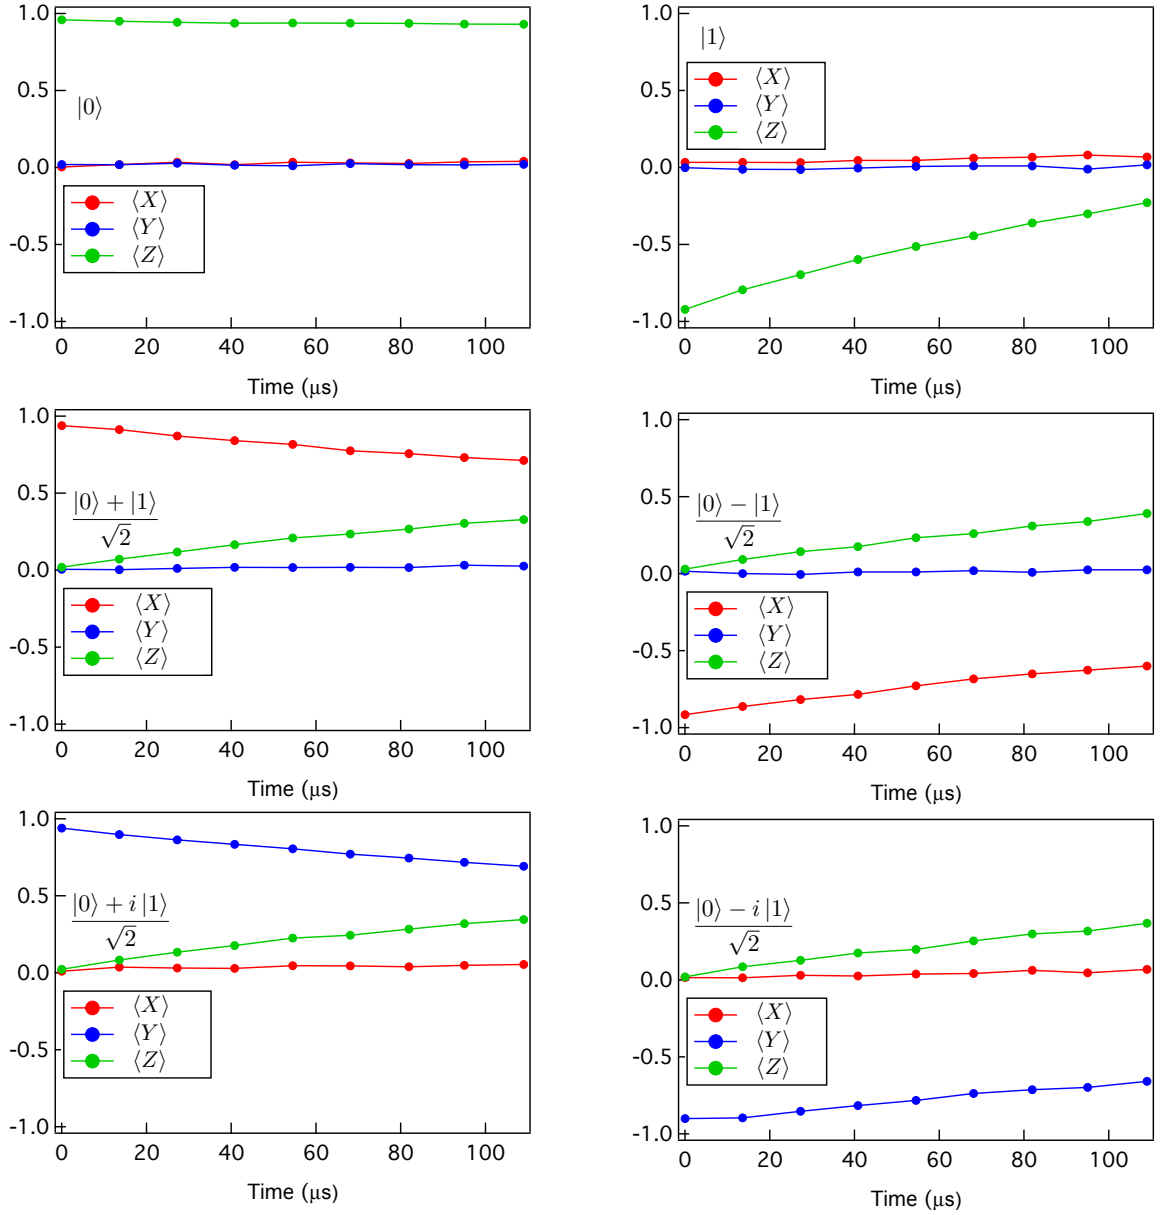


Figure 7.5: **Fock state encoding amplitude decay.** The Fock state encoding shows decay curves typical of amplitude damping, or T_1 -type decoherence, in which all coordinates preferentially decay toward the energetically favorable cavity ground state $|0\rangle_f$. One can see in each plot that the value of the $\langle Z \rangle$ coordinate monotonically increases towards (or stays at) +1 regardless of the initial state, while every other coordinate decays to 0.

$$E_0 = \sqrt{1 - n_{th}} \begin{pmatrix} 1 & 0 \\ 0 & \sqrt{1 - f(t)} \end{pmatrix}$$

$$E_2 = \sqrt{n_{th}} \begin{pmatrix} \sqrt{1 - f(t)} & 0 \\ 0 & 1 \end{pmatrix}$$

$$E_1 = \sqrt{1 - n_{th}} \begin{pmatrix} 0 & \sqrt{f(t)} \\ 0 & 0 \end{pmatrix}$$

$$E_3 = \sqrt{n_{th}} \begin{pmatrix} 0 & 0 \\ \sqrt{f(t)} & 0 \end{pmatrix}$$

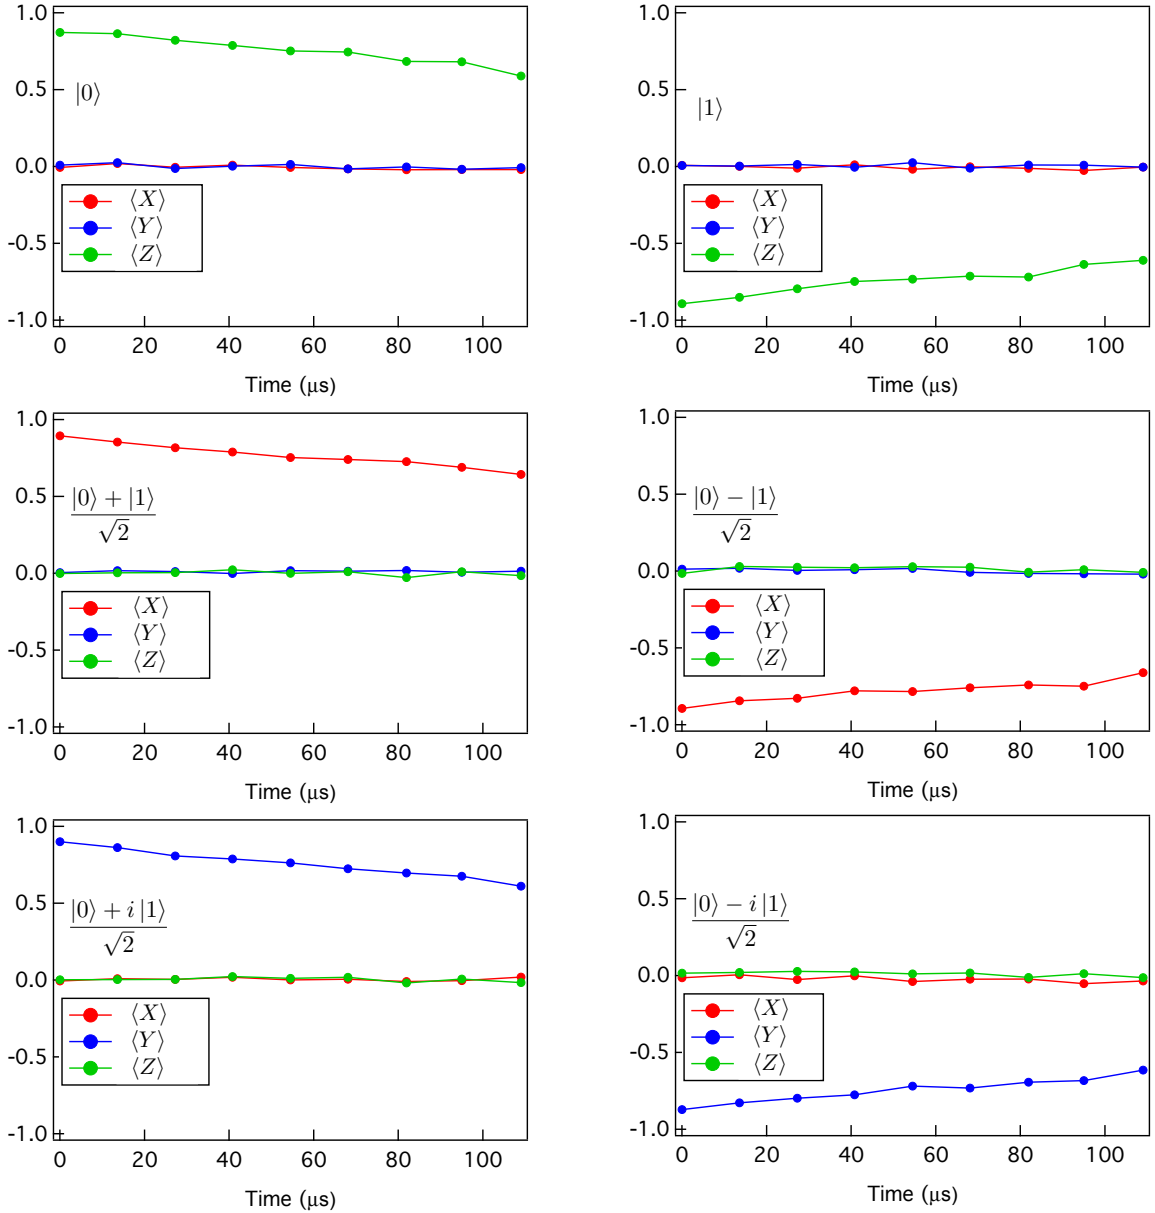


Figure 7.6: Cat code depolarization This data demonstrates that depolarization, in which every Bloch vector shrinks uniformly toward a fully mixed state at the origin, is the dominant error channel in the full QEC system. The six plots show the decay in time of the average X , Y , and Z components of the Bloch vector ($\langle X \rangle$, $\langle Y \rangle$, $\langle Z \rangle$ respectively) for each cardinal point after using the cat code QEC system. These six points are initialized by applying the identity (I), π pulse about the Y axis (R_π^y), or $\pm\pi/2$ rotations about the Y or X axes ($R_{\pi/2}^y, R_{-\pi/2}^y, R_{\pi/2}^x, R_{-\pi/2}^x$) on the ancilla prior to the error monitoring. This data is used to calculate the process matrix X^M of the corrected quantum bit. In each of the six cases, only the non-zero coordinate of the Bloch vector at zero time decays while the other two remain at 0 throughout the entire tracking duration. We find the decay rate of cat states along $|\pm X_L^\pm\rangle$ to be slightly more robust as these states are symmetric about both axes in the cavity's phase space, while $|C_\alpha^+\rangle$, $|C_{i\alpha}^+\rangle$ and $|\pm Y_L^\pm\rangle$ are symmetric about only one. Thus, rotations in phase space are somewhat less detrimental for $|\pm X_L^\pm\rangle$.

where $f(t)$ is a function of time of the form $f(t) = 1 - e^{-t/t_0}$; for the transmon $t_0 = T_1$ and $n_{th} = n_{th}^a$; for the Fock state $t_0 = \tau_s$ and $n_{th} = n_{th}^s$ (see sec. A.3). As seen in Fig. 7.5b, all Bloch sphere vectors preferentially decay toward the energetically favorable ground state of the encoding. The off-diagonal elements in the density matrices of these encodings decay with a time constant that in addition to the amplitude damping also includes the pure dephasing in the system; these combined rates are T_2 for the transmon and T_2^s for the Fock state components. Thus, we find that $r_x^{+\bar{x}}$ and $r_y^{+\bar{y}}$ in eq. 7.1 decay at T_2^s (T_2) for the Fock state (transmon) encodings, while $(r_z^{+\bar{z}} - r_z^{-\bar{z}})/2$ decays at τ_s (T_1). The process fidelity as a function of time $X_{00}^M(t)$ therefore decays at two different rates, resulting in a double-exponential behavior. In the decay of the cardinal points in the Fock state encoding (Fig. 7.5) one can only discern a single exponential behavior, however, as the maximum duration of $\sim 110 \mu\text{s}$ is still short compared to the cavity's coherence times.

In contrast, with the cat code we find that all of the cardinal points on the Bloch sphere shrink uniformly toward the fully mixed state $\rho_{fin} = \hat{I}/2$ (Fig. 7.6). Such decoherence can be described by the process of depolarization, in which one essentially loses knowledge as to the orientation of the Bloch vector:

$$\begin{aligned}
\mathcal{E}(\rho) &= \frac{p\hat{I}}{2} + (1-p)\rho \\
&= (1 - \frac{3p}{4})\rho + \frac{p}{4}(\hat{\sigma}_x\rho\hat{\sigma}_x + \hat{\sigma}_y\rho\hat{\sigma}_y + \hat{\sigma}_z\rho\hat{\sigma}_z) \\
&= \sum_j \tilde{E}_j\rho\tilde{E}_j^\dagger X_{jj}
\end{aligned}
\quad X = \begin{pmatrix} 1 - \frac{3p}{4} & 0 & 0 & 0 \\ 0 & \frac{p}{4} & 0 & 0 \\ 0 & 0 & \frac{p}{4} & 0 \\ 0 & 0 & 0 & \frac{p}{4} \end{pmatrix} \quad (7.10)$$

This simple formula shows that the signature of depolarization errors is a diagonal process matrix X in which the value of X_{00} decreases with p , while the remaining diagonal components increase with p , for p a probability between 0 and 1. The dominant source of depolarization in our work stems from the incorrect application of the decoding pulse at the end of the QEC sequence, a consequence of both syndrome measurement results that incorrectly indicate the occurrence of an error, and of dephasing due to the back-action of ancilla T_1 decay. If the decoding pulse is applied at either the wrong angle (quantified in sec. 8.2) or attempts to decode the wrong parity (see sec. 6.3), at the end of the decoding sequence we are left with a completely mixed quantum bit after tracing over the cavity state. Although at first glance the basis states along the logical Z axis, “2-cats,” should be immune to double-errors within the waiting time between syndrome measurements, the resulting Kerr phase accumulation after two photon jumps is sufficient to appreciably rotate the cat states out of the logical space.

7.3 Overcoming the decoherence penalty with the cat code

Using the results of the state tomography shown in the octahedra of Fig. 7.4 to find how the process fidelity of each quantum bit storage scheme decays over time, we arrive at a central result of our work. Specifically, we look at how well the full QEC system implements the identity operation, where the process fidelity is defined as $F = X_{00}^M$ (or equivalently, $F = \text{Tr}(X^M X_0)$). The benefits of employing QEC are once again apparent, as seen in Fig. 7.7. Without any post-selection or renormalization to remove encoding-associated penalties, the corrected cat code overcomes the decoherence penalty by a factor of 2.2.

In many ways this is the culmination of the goals first introduced in chapter 3. We began by stressing the challenges associated with traditional error correction protocols by framing the discussion around the decoherence penalty as the dominant limitation; constructing a logical space from so many physical qubits simply makes things worse with current levels of performance across all quantum computing platforms. In our experiments, however, we have shown that a very simple architecture introduces a much smaller resource overhead to begin with for a single logical qubit: an enhanced error rate that scales linearly with the average photon number of the cat states. A logical qubit in which the non-orthogonality between the basis states is very nearly 0 requires an $\bar{n} \approx 4$ and a single ancilla for the error syndrome, versus the Steane code’s seven physical qubits and six ancillae. Even with an ancilla that lacks the cavity’s coherence properties, the encoding and decoding losses are just $\sim 4\%$ per operation, as seen in the initial offsets in Fig. 7.7; this state-of-the-art performance demonstrates a manipulation of a full logical quantum bit that to-date no other system has been able to achieve. Finally, sparsely applied parity measurements that entangle the encoded state with the ancilla only briefly allow us to extract an error syndrome with high enough fidelity to not just make up for the decoherence penalty of $\bar{n}_0 = 2$, but in fact realize gains in performance by slowing down the rate of decay by a factor of 2.2. We thus prove that one can actually realize gains with QEC while satisfying the three criteria set out in chapter 6: *all quantum states* are protected from *naturally occurring errors without the use of post-selection*.

Additionally, the performance of the cat code matches that of a simple encoding in the $|0\rangle_f, |1\rangle_f$ Fock states after $\sim 80 \mu\text{s}$, and by decaying with a time constant that exceeds that of the Fock state by a factor of 1.1, it protects the information from decoherence more effectively than using the Fock state encoding as a passive quantum memory. At this point one may ask if this decrease in the rate of process fidelity decay when using the cat code with active error correction sufficiently supports the claim that the system surpasses the break-even point of QEC. For instance, one may hold the




Figure 7.7: QEC program process tomography. To implement QEC, we redundantly encode the quantum bit in cat states ($\bar{n}_0 = 2$) and pay the required decoherence penalty, ubiquitous to QEC. This leads initially to worse performance; the process fidelity of the uncorrected cat code (orange circles), where cat states are left to decay freely between encoding and decoding, exhibits faster decay as compared to the Fock states $|0\rangle_f, |1\rangle_f$ (gray circles). Substantial improvements in performance are realized with the full QEC system; the corrected cat code (red triangles) surpasses the uncorrected transmon (green squares) by a factor of ~ 20 , makes up for the QEC decoherence penalty by a factor of ~ 2.2 , and offers an improvement over the Fock state encoding by a factor of ~ 1.1 . With only high-confidence trajectories (blue diamonds), the decay time increases to $\tau > 0.5$ ms. Top axis: number of syndrome measurements used for each point in the corrected cat code. Cat code data: 100,000 averages per point; transmon, Fock states: 50,000 averages per point; error bars are smaller than marker sizes. Although no data exhibits strictly single-exponential decay, all curves are well modeled by $F(t) = 0.25 + Ae^{-t/\tau}$; dotted lines are fits; $F = 0.25$ (dashed line) implies a complete loss of information; uncertainties are given by errors in the fit. Fluctuations in the uncorrected cat code are explained by the Kerr effect and are reproduced in simulation.

pedantic point of view that as the curve of the corrected cat code never actually crosses the Fock state encoding, one cannot realize any actual advantage to using the former in a practical sense: given this exact system, the cat code at best matches the fidelity of the Fock state encoding only at one point in the monitoring sequence, at $\sim 80 \mu\text{s}$. The claim of surpassing break-even, however, should rather be understood as characterizing the capacity of the QEC system to yield advantages in a future realization of a quantum computer. For example, immediate issues concerning encoding and decoding losses, responsible for the initial drop of $\sim 8\%$ in the process fidelity at time $t = 0$,

must of course be addressed before any practical application becomes a reality. Indeed, thus far no architecture on any platform has managed to demonstrate a redundant encoding with higher fidelity than that of the single, physical qubit. This problem will be solved once operations on and between cat states approach fidelities necessary for fault-tolerance. If the experiment reported here were then performed in the exact same way, each curve in Fig. 7.7 would start at some value very close to +1, at which point the cat code would always outperform every element in the system. Thus, although it is true that as presented here the cat code doesn't currently provide any practical advantages, the point is that these results demonstrate that it will once the necessary gate fidelities are achieved. The full QEC system implemented here completely overcomes the decoherence penalty associated with using the full QEC system, a feature we argue is more germane to the discussion of what break-even requires.

To conclude, our results show that QEC can actually protect an unknown bit of quantum information, and extend its lifetime by active means. By employing the cat code as the foundation of a novel QEC system, we demonstrate the advantages of the hardware efficiency: the capacity of a single cavity to store a logical qubit; its intrinsically high coherence times; and the need to monitor just one error syndrome with just one ancilla. Furthermore, we demonstrate the crucial role of real-time feedback, an addition to the experimental setup that vastly improves error correction performance and allows us to realize the cat code at the break-even point of QEC. Future goals include combining the cat code with mechanisms to re-inflate cat state amplitudes and to equip the parity monitoring protocol to handle changes in ancilla energy, thereby addressing issues of non-fault-tolerance head-on. We believe our results, however, already motivate the adaptation of QEC schemes to exploit the efficiencies of hardware platforms beyond the purview of traditional architectures, and the promise of cat states as components integral to future quantum computing applications.

8 – Limitations on Fault-Tolerance

Now that we've demonstrated a full QEC system that overcomes its decoherence penalty while sweeping nothing under the rug, we would like to understand what limits its final performance. The goal in this work is not to demonstrate anything that is fault-tolerant. Indeed, by design, the current implementation of the cat code suppresses errors to first order, so at the outset we are prepared to be limited at least by the occurrence of multiple errors within the encoding that the error syndrome measurements are not designed to capture. Moreover, working within a continuous variable space, we are destined to lose fidelity simply because sources of dephasing arising from smaller terms in the Hamiltonian become non-negligible over long periods of monitoring and obviously result in some form of decoherence simply due to leakage out of the code space. In other words, photon number parity measurements teach us only about the occurrence of photon loss rather than any change in the phase of the state in the cavity. These sources must still of course be quantified and their effect on future implementations of the cat code assessed.

To varying degrees, every step in the full QEC system described in sec. 6.4 lowers the fidelity of the final state. Starting with the quantum bit initialization, which is limited by pulse fidelities, and moving through the encoding, monitoring, and decoding stages, which all suffer from decoherence in the code space and in the ancilla, we inevitably end up losing some amount of information when the experiment concludes. The questions are: how much, and what are the dominant limitations?

The most pressing consideration is the contribution to dephasing due to forward propagation of errors. Every time we perform an error syndrome measurement, with some probability ancilla T_1 decay occurs, the back-action of which makes the state in the cavity irrecoverable (discussed in sec. 5.4). This motivates a substantial pivot in the strategy of implementing the full error correction system. We can largely suppress the contributions to infidelity arising from higher order error processes, pulse and readout imperfections, and deviations of the cat states from their defined basis with faster measurement rates, which allow for a build-up of confidence in the true error trajectories. Forward propagation, on the other hand, pulls us in the opposite direction. Entangling the imperfect ancilla with the logical qubit increases the probability of information loss. The more frequently error

syndrome measurements are performed, the greater the probability per unit time the final quantum bit will be unrecoverable. If the ancilla’s coherence properties are sufficiently poor, the optimal strategy calls for a measurement cadence with which accruing meaningful statistics to perform some type of majority voting no longer becomes possible. In this case, we must take every measurement result at face value.

This chapter is devoted to a discussion of the dominant sources of error that afflict the system and how they contribute to the non-fault-tolerance of the current implementation of the cat code. We present each source of decoherence in the general context of the dominant loss mechanisms ubiquitous in any measurement-based error correction scheme, as first introduced in chapter 3. By understanding how these different fault mechanisms enter into the calculus of optimizing the performance of the cat code, we predict the impact better ancilla coherence properties will have on our ability to protect a quantum bit from decoherence. Thus, in step with the general guiding philosophy in this work, we scrutinize each aspect of the QEC system in great detail before thinking about how to scale up to more complex architectures. We ultimately arrive at a second, equally important result of this work: the forward propagation of errors from ancilla T_1 decay is by far the dominant limitation that currently prevents our QEC system from offering much greater gains in lifetime. In many ways this conclusion reflects yet another crucial advantage of the hardware efficiency of the code. By pinpointing the dominant source of non-fault-tolerance, we know where to focus our efforts to improve the performance of the error correcting system.

8.1 The difference between error record and error history

The simplicity of the curves shown in Fig. 7.7 belies the trove of information we can garner from looking more closely at the statistics of these measurements. The exact pattern of the error syndrome measurement results also presents us with a valuable measure of confidence as to our knowledge of the actual trajectory of the encoded state. For example, using the post-selection strategy that accepts only high-confidence trajectories yields a decay constant of the QEC system’s process fidelity of over half a millisecond while still keeping a majority of the data (Fig. 8.1). This suggests of course that with faster syndrome measurement rates, one can approach a regime in which real-time statistical analysis of the error syndromes will be possible. Simple strategies such as majority voting on the measurement result, for example, will offer a means to rule out results with errant parity mappings. As a result, this will improve the confidence in the knowledge of how many photon jumps actually occurred. We thus seek to understand the statistics of the ostensible jumps in photon number we detect, and use the resulting analysis to deduce the optimal measurement strategy.

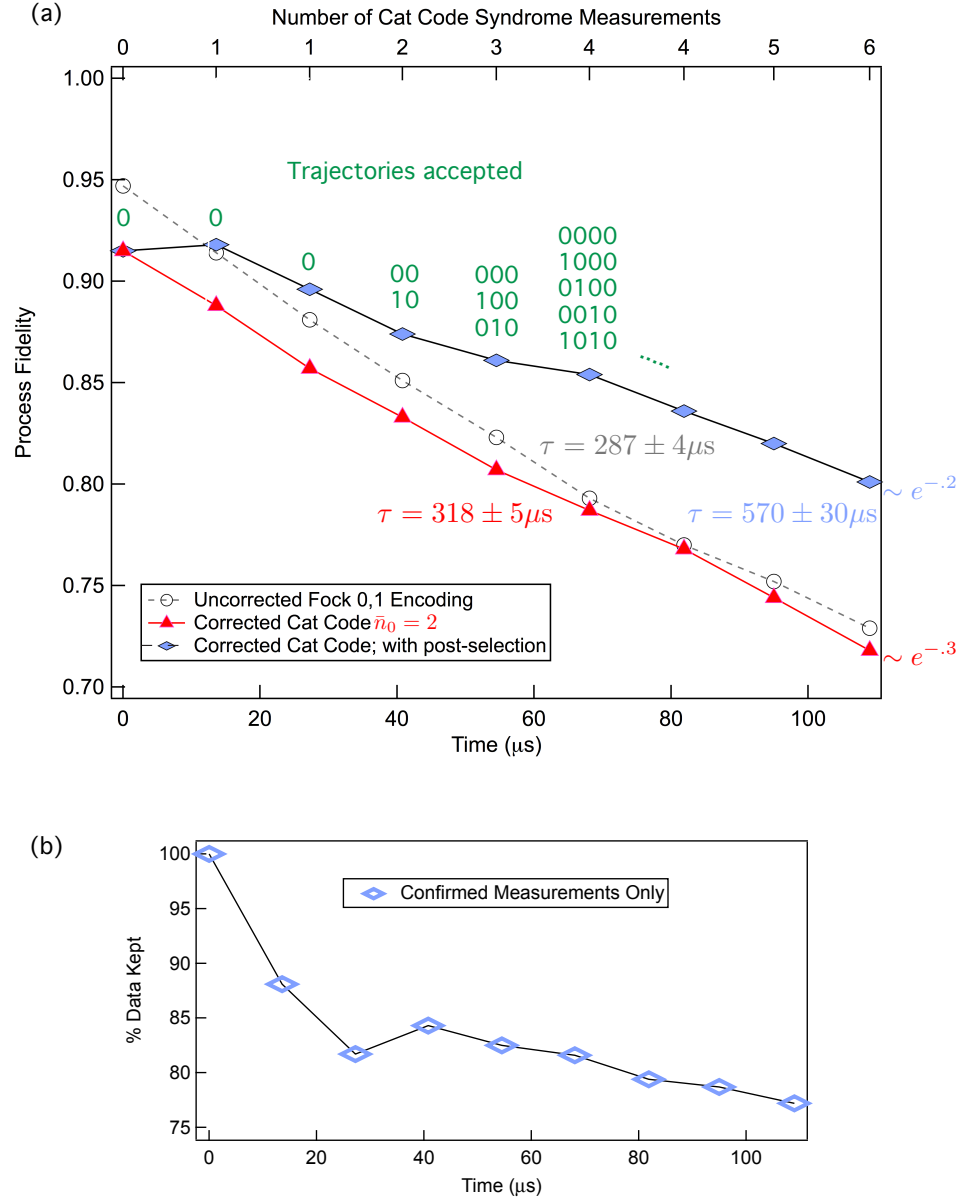


Figure 8.1: **QEC program process tomography with post-selection.** (a) This plot demonstrates the further enhancements in performance possible with post-selection as compared with the data shown in Fig. 7.7. The process fidelity decay of the cat code QEC system with $\bar{n}_0 = 2$ improves substantially if low confidence trajectories are omitted from the data (purple diamonds). Binary representations of high confidence trajectories; the first time a two jump record is included in the averaging occurs at 68 μs (1010). (b) This plot shows the measured occurrence of high confidence trajectories. Note that after $\sim 100 \mu\text{s}$ we still keep nearly 80% of the data.

With the calculations below, we expand upon the disparity in fidelity between high and low-confidence records. In particular, we discuss why the process fidelity of the 2-error case after two syndrome measurements, shown in Fig. 7.3, differs so substantially from that of the 0 and 1 error

cases. We proceed to show that with more syndrome measurements, the fidelity of the two-error case in fact increases toward a maximal value attained at a later time. This seemingly paradoxical observation should not come as a surprise if one recalls the discussion of the quantum filter in sec. 5.3. The improving fidelity of multi-error cases in time simply underscores the point that the confidence associated with a given measured error record relies on a build-up of statistics. Despite the paucity of such statistics here as compared to the hundreds of repeated measurements presented in chapter 5, we can nonetheless quantify why records indicating consecutive errors are so untrustworthy.

In order to understand this effect, we invoke Bayes' rule. The measured fidelity of successfully mapping parity in the presence of $\bar{n} = 3$ photons in the cavity is 97.7%. Although very high, this number still leads to the somewhat surprising differences in the confidence of certain measurement results over others. We start by assuming that we have encoded a quantum bit in cat states of size $\bar{n} = 3$ and that the ancilla is in $|g\rangle$ prior to the parity mapping. This is the initial system state (Fig. 6.12a). After a round of error monitoring in which we use the parity protocol that maps even parity to $|g\rangle$ and odd parity to $|e\rangle$ (indicated by the superscript “-”), the probabilities to measure ancilla $|g\rangle, |e\rangle$ are given by:

$$p(g) = p^-(g|0\varepsilon)p(0\varepsilon) + p^-(g|1\varepsilon)p(1\varepsilon) \quad (8.1)$$

$$p(e) = p^-(e|0\varepsilon)p(0\varepsilon) + p^-(e|1\varepsilon)p(1\varepsilon), \quad (8.2)$$

where $p(0\varepsilon) = e^{-(\bar{n}e^{-t/\tau_s})t_w/\tau_s}$ is the probability that the cavity state had 0 parity jumps, $p(1\varepsilon) = 1 - p(0\varepsilon)$, and $p^-(g|0\varepsilon)$ and $p^-(e|1\varepsilon)$ are respectively the probabilities to measure $|g\rangle$ when the cavity state had 0 parity jumps and $|e\rangle$ when the cavity had 1 parity jump. Likewise, when we use the parity protocol that maps odd parity to $|g\rangle$ and even parity to $|e\rangle$ (indicated by the superscript “+”), the probabilities to measure ancilla $|g\rangle, |e\rangle$ are given by:

$$p(g) = p^+(g|0\varepsilon)p(0\varepsilon) + p^+(g|1\varepsilon)p(1\varepsilon) \quad (8.3)$$

$$p(e) = p^+(e|0\varepsilon)p(0\varepsilon) + p^+(e|1\varepsilon)p(1\varepsilon), \quad (8.4)$$

where $p^+(g|0\varepsilon)$ and $p^+(e|1\varepsilon)$ are respectively the probabilities to measure $|g\rangle$ when the cavity had 0 parity jumps and $|e\rangle$ when the cavity had 1 parity jump. In our system, for average photon number $\bar{n} = 3$ in the cavity, $p^-(g|0\varepsilon) = p^+(g|0\varepsilon) = 0.983$ and $p^+(e|1\varepsilon) = p^-(e|1\varepsilon) = 0.971$.

We seek to predict the statistics for monitoring errors for two steps and an initial cat size of average photon number $\bar{n} = 3$, as presented in Fig. 7.3. Following the flow of Fig. 3a-d, assuming at time $t = 0$ we start with an even parity state in the cavity and perform the first round of error

monitoring that maps even parity to $|g\rangle$, after $13.8 \mu\text{s}$ $p(g) = 0.841$ and $p(e) = 0.159$. Using Bayes' rule, we can now calculate the conditional probabilities for the cavity to be in a certain parity state given the measurement outcome:

$$p^-(0\varepsilon|g) = \frac{p^-(g|0\varepsilon)p(0\varepsilon)}{p(g)} = 0.995 \quad (8.5)$$

$$p^-(1\varepsilon|e) = \frac{p^-(e|1\varepsilon)p(1\varepsilon)}{p(e)} = 0.910 \quad (8.6)$$

The key point here is the difference in the confidence as to the true occurrence of an error when the ancilla ends up in $|e\rangle$. The small measurement infidelities together with the relatively low probability to have an error in the first place leads to a considerable difference of $\sim 8\%$ between the two conditional probabilities $p^-(0\varepsilon|g)$ and $p^-(1\varepsilon|e)$. This difference leads to a higher likelihood for the parity meter to suggest the occurrence of another error in the encoded state in the subsequent measurement, and thus leads to a substantially lower confidence in any trajectory that indicates consecutive errors, as we show next.

If we measure $|g\rangle$, we continue using the same protocol, but now to obtain $p^\pm(g)$ and $p^\pm(e)$ (probabilities to measure $|g\rangle$ and $|e\rangle$ for the two different parity mapping protocols) we no longer have the luxury of knowing that we start in an even state and thus must use the conditional probabilities obtained above for the following:

$$\begin{aligned} p^-(g) &= [p^-(g|0\varepsilon)p(0\varepsilon) + p^-(g|1\varepsilon)p(1\varepsilon)]p^-(0\varepsilon|g) \\ &\quad + [p^-(g|1\varepsilon)p(0\varepsilon) + p^-(g|0\varepsilon)p(1\varepsilon)][1 - p^-(0\varepsilon|g)] = 0.834, \end{aligned} \quad (8.7)$$

and $p^-(e) = 1 - p^-(g) = 0.166$. Similarly for the case where we instead measure $|e\rangle$ and the protocol is flipped in the second round:

$$\begin{aligned} p^+(g) &= [p^+(g|0\varepsilon)p(0\varepsilon) + p^+(g|1\varepsilon)p(1\varepsilon)]p^-(1\varepsilon|e) \\ &\quad + [p^+(g|1\varepsilon)p(0\varepsilon) + p^+(g|0\varepsilon)p(1\varepsilon)][1 - p^-(1\varepsilon|e)] = 0.777, \end{aligned} \quad (8.8)$$

and $p^+(e) = 1 - p^+(g) = 0.223$.

We now have the probabilities to obtain the following measurement records, which closely match those presented in Fig. 7.3:

$$p_{0\varepsilon} = p(g)p^-(g) = 0.841 \times 0.834 \quad (8.9)$$

$$= 0.701$$

$$p_{1\varepsilon} = p(g)p^-(e) + p(e)p^+(g) = 0.841 \times 0.166 + 0.159 \times 0.777 \quad (8.10)$$

$$= 0.263$$

$$p_{2\varepsilon} = p(e)p^+(e) = 0.159 \times 0.223 \quad (8.11)$$

$$= 0.036$$

Beyond telling us that we understand the statistics of our system, this calculation also provides crucial information as to the confidence of certain trajectories over others. First, one may immediately note the slight asymmetry between measuring $|g\rangle$ and then $|e\rangle$ ($0.841 \times 0.166 = 0.140$) vs. the reverse order ($0.159 \times 0.777 = 0.124$). Indeed, with the following conditional probabilities for all possible error histories (gg, eg, ge, ee) we see the huge benefit of a “confirmation” g measurement on the probability that the measured trajectory faithfully reflects the error trajectory of the encoded state:

$$p^-(0\varepsilon|gg) = \frac{p^-(g|0\varepsilon)p(0\varepsilon)p^-(0\varepsilon|g)}{p^-(g)} = 0.993 \quad (8.12)$$

$$p^+(1\varepsilon|eg) = \frac{p^+(g|0\varepsilon)p(0\varepsilon)p^-(1\varepsilon|e)}{p^+(g)} = 0.978 \quad (8.13)$$

$$p^-(1\varepsilon|ge) = \frac{p^-(e|1\varepsilon)p(1\varepsilon)p^-(0\varepsilon|g)}{p^-(e)} = 0.869 \quad (8.14)$$

$$p^+(2\varepsilon|ee) = \frac{p^+(e|1\varepsilon)p(1\varepsilon)p^-(1\varepsilon|e)}{p^+(e)} = 0.592 \quad (8.15)$$

It becomes clear that measurements of “error” confirmed by subsequent measurements of “no error” have a $\sim 10\%$ higher fidelity than those without such a confirmation. Moreover, if two consecutive “error” measurements are recorded, the probability drops substantially by $\sim 20 - 30\%$. With these findings, many of the features in the data fall into place. One can see the effects of these conditional probabilities by looking at the 1-error cases in the Wigner tomography in Figs. 6.10 and 6.12, where the parity and fringe contrast of the 01 case appear to be less negative and sharp than that of 10 ($0 \equiv$ “no error” and $1 \equiv$ “error”). Along the same lines, the case 11 has by far the lowest fidelity, as confirmed by the process tomography in Fig. 7.3. The single-shot records

that come with each repetition of the monitoring sequence thus provide us with crucial information beyond simply how to bin each result.




Figure 8.2: **Assessing measurement record confidence.** Predicted statistics and confidence for the corrected cat code after four syndrome measurements over $68\ \mu\text{s}$ of monitoring; $\bar{n}_0 = 2$. After four rounds of error correction there are sixteen possible result records: 0000, 1000, \dots , 1111. The left plot shows the predicted probability to measure each of these records individually (green bars, top axis), and the cumulative probability (bottom axis). In the right plot, we also show the predicted process fidelity conditioned on measuring each record (red bars are even parity, blue bars are odd parity). This conditioned fidelity corresponds to our confidence in the output. In the column separating the plots, the numbers in green correspond to the actual values of each individual green bar (left plot), and the numbers in red and blue correspond to the values of the red and blue bars (right plot). The axis of the right plot, cumulative process fidelity, is the cumulative sum of the predicted conditional probabilities weighted by the trajectory occurrence probability. It is interesting to compare the records 1010 and 0001. The first suggests two photon jumps (during the first and third steps) and the second suggests a single photon jump during the last step. The conditional process fidelity for 1010 is actually higher. This is because measuring ancilla $|g\rangle$, which indicates no change in parity, has a higher probability of being correct than measuring $|e\rangle$, which does indicate a change of parity. Thus, every error (1) in 1010 is “verified” by the subsequent measurement of no error (0), while an error in the last step of 0001 has a higher likelihood to be a faulty measurement. Outlined in purple is the set of data we accept as “high-confidence” trajectories (Fig. 8.1), wherein every 1 is confirmed by a subsequent 0.

We can extend these calculations to a general simulation that handles an arbitrary number of correction steps to show that this simple approach captures many of the features we see in our data. A detailed look at tracking for $\sim 68\ \mu\text{s}$ in Fig. 8.2, for example, shows the individual process fidelities we expect to measure for every possible measurement record. A particularly noteworthy conclusion is that trajectories such as 1010 have a higher expected fidelity than 0001. Although the former suggests more errors in the encoded state, each measurement result 1 is confirmed by a subsequent 0, whereas this is not the case in the latter.

In summary, with each error syndrome measurement, all measurement infidelities are pushed onto records that report higher and higher error numbers, and with time, these records become more and more common. As seen in Fig. 8.2, the occurrence of high-confidence trajectories falls in time, albeit slowly. In this sense, although the post-selection substantially improves the quality of the final quantum bit, it nonetheless results in an exponential decay of acceptable trajectories, an expected trade-off.

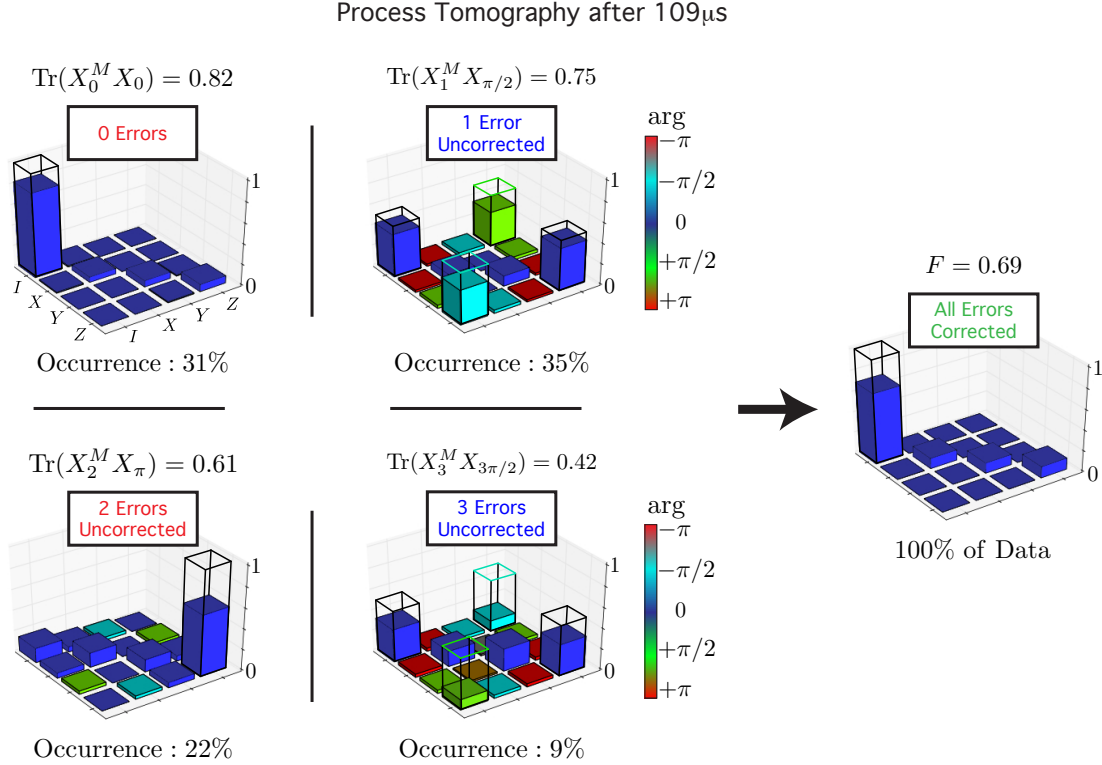


Figure 8.3: Process matrices after seven steps of error monitoring. Measured process matrices X_j^M for $j = 0, 1, 2,$ and 3 errors after $109 \mu\text{s}$ for an initial encoding size of $\bar{n}_0 = 3$. Ideal processes are given by X_j and are wire-outlined. As compared with the data in Fig. 7.3, note that the 0 error case has the greatest drop in fidelity, the 1 error case goes down slightly, and the 2 error case increases substantially. Note the 3 error case also exhibits clear signatures of the correct X_3 form. The substantial drop in fidelity from 0.84 of Fig. 7.3 to 0.69 here is primarily due to the drop in the occurrence of 0 error cases with time.

Thus, the notion that an error record and an actual error history are not simply related by some overall scaling corresponding to measurement infidelity explains why we should in fact expect the fidelities of multi-error cases to not simply decay in time. As we've seen, building up statistics over time provides valuable information in estimating the properties of the final state in the cavity. Continuing the monitoring shown in the debugger, where $\bar{n}_0 = 3$, (Fig. 6.12) for more than two steps provides a palpable example. After seven steps over $\sim 110 \mu\text{s}$, the X^M matrices conditioned on zero to three photon jumps (Fig. 8.3) show a similar form as to those in Fig. 7.3. Interestingly, and as




Figure 8.4: **Process fidelity conditioned on error number.** This plot shows the process fidelities conditioned on individual error trajectories for $j = 0, 1,$ and 2 errors (red circles, blue squares, and green triangles respectively). Unsurprisingly, the 0 error case has the highest fidelity, followed by the 1 and 2 error cases. The initially surprising feature here is that the process fidelity of the 2 -error case increases with time. This is precisely a consequence of error syndrome measurement fidelity, wherein with more 2 error trajectories we have more result records that have the aforementioned confirmation measurements. In other words, with more statistics (inset) we have greater knowledge that measured 2 error cases in fact correspond to actual 2 error cases in the encoded state. Non-monotonic variations in the data points throughout the entire curve are attributed to variations in the efficacy of the decoding pulses at different points in time.

expected, the zero jumps cases experiences the greatest drop in fidelity, from $F = 0.89$ to $F = 0.82$; despite the high fidelity associated with so many consecutive “ 0 ” measurements, a probability of missing photon jumps or misreading the parity meter still exists (discussed next in sec. 8.2). This means that the quality of the estimate of the cavity parity and phase can only deteriorate. The fidelity of the one-jump case in fact stays more-or-less the same, while that of the two-jump case substantially improves as more high-confidence trajectories become available. The lower final process fidelity is due to the exponentially diminishing number of 0 -error cases.

Looking at the data for all points in time and plotting the fidelity decay conditioned on the numbers of errors detected, as shown in Fig. 8.4, demonstrates the contrasting trends for the $0, 1,$ and 2 -error case. The 0 -error case decays with a time constant of $\sim 630 \mu\text{s}$; this is the longest time constant in our error correction system, and demonstrates the high fidelity we can achieve if we use the cat code only as an error indicator. The fidelity of the 2 -error case instead slowly increases

with time as the occurrence of two-error trajectories that contain higher confidence “confirmation” measurements within them increases as well, leveling off after about $80 \mu\text{s}$ of monitoring. Crucially, these results do not suggest that the cat code is ill-equipped to handle multiple errors throughout a monitoring sequence. Rather, as explained in the following sections, they highlight the trade-offs we make between mitigating the effects of ancilla back-action and the consequent dearth of statistics that results from the requirement of a substantially reduced syndrome measurement rate.

8.2 Cat code QEC performance limitations

In discussing what limits our performance, we return to the original motivation for the cat code itself, namely the variety of errors one must inevitably face when implementing any QEC system and the rewards one reaps by moving to a hardware-efficient approach. The cat code still faces the same dominant avenues of code failure as those introduced in the context of the Steane code in Fig. 3.5. The need for only one ancilla and only one error syndrome measurement, however, drastically simplifies the analysis of where things can go wrong (Fig. 8.5) and enables us to estimate the severity of each source of loss. Table 1 details an infidelity budget for the cat code, which lists the main contributions to depolarization. Although the contribution of each source is small, the sources are many and add up quickly, bluntly encapsulating the challenges one faces in realizing fault-tolerant QEC. Many of the effects can be mitigated by measuring more quickly and perhaps employing a quantum filter to retrieve a best estimate of the parity at any given time (see sec. 5.3). For example, infidelities due to missing the occurrence of an error; ancilla preparation and readout; and excursions out of the logical space due to Kerr, all approach zero as the measurement rate increases. These thus impose no intrinsic limitation on cat code performance.

Crucially, however, errors due to the ancilla T_1 still persist regardless of measurement rate. Due to its dispersive coupling to the cavity, a change in the energy of the ancilla at an unknown time imparts an unknown rotation to the cavity state in phase space (introduced in sec. 5.4); this is the forward propagation of an error. Any such rotation commutes with the parity measurement, indicating that we learn nothing about these errors from our one syndrome measurement and that the cat code in its current implementation is not fault-tolerant. Measuring the syndrome more frequently only increases the likelihood of ancilla-induced dephasing, necessitating an optimized measurement cadence to balance the risk of missing errors during t_w with the probability of error propagation during each syndrome interrogation. By slowing down the measurement rate, however, we unfortunately elevate the contributions from the remaining sources of decoherence.

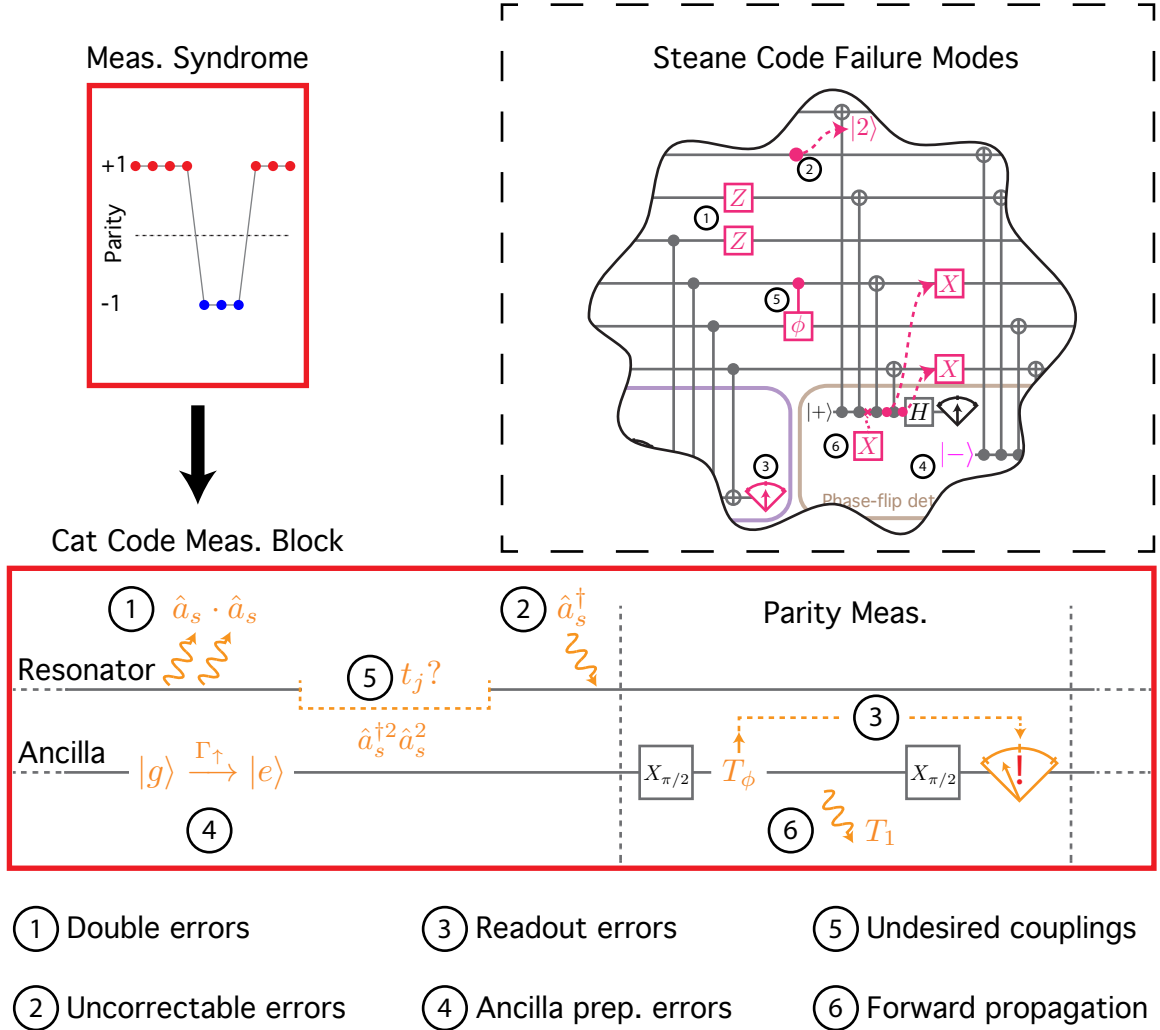


Figure 8.5: **Cat code measurement block failure modes.** The syndrome measurement step, shown in Fig. 6.1, can be unfolded to analyze the sources of loss that contribute to decoherence in the full QEC system. Just as in the case of the Steane code, there are six dominant failure modes in the cat code: (1) the loss of two photons between consecutive syndrome measurements, $\hat{a}_s \cdot \hat{a}_s$ (double outgoing lines); (2) possible errors that the cat code does not address, such as additions of a single photon, \hat{a}_s^\dagger (single incoming line); (3) a failed parity mapping resulting from ancilla dephasing (ancilla T_ϕ); (4) incorrect ancilla initialization prior to syndrome measurement resulting from unknown excitations, ancilla $\Gamma_\uparrow (|g\rangle \rightarrow |e\rangle)$; (5) undesired couplings that result in dephasing due to Kerr, $\hat{a}_s^{\dagger 2} \hat{a}_s^2$ (uncertainty in the time of a photon jump t_j); and (6) ancilla decoherence that directly propagates to unrecoverable errors in the logical qubit, a result of ancilla decay or excitation, T_1 (single outgoing line). A repeat of the image from Fig. 3.5 highlights the correspondence between the Steane and cat codes.

In the remainder of this section we expand upon the calculations that produce the predicted gains in lifetime of the cat code over the Fock state encoding listed in Table 1. Each dominant source of decoherence, whether arising from errors in the encoded state or in the syndrome interrogation, contributes to the probability that the cat code will fail in a single round of error correction. When isolated in a hypothetical situation as the only source of loss in the system, it can be quantified with simple estimates based on coherence times, thermal populations, and measurement fidelities.

| Failure Mode | Dominant Source | Max. Rate $t_w \approx 0\mu s$ | Opt. Rate $t_w \approx 20\mu s$ |
|--------------------------------------|--|-----------------------------------|------------------------------------|
| | | Predicted τ | Predicted τ |
| Double Errors | cavity $\hat{a} \cdot \hat{a}$ | 40 ms | 1.7 ms |
| Uncorrectable Errors | cavity \hat{a}^\dagger | 6 ms | 6 ms |
| Readout Error | transmon T_ϕ | 7 ms | 2 ms |
| Ancilla Preparation | transmon Γ_\uparrow | 300 ms | 900 μs |
| Undesired Couplings | cavity $\hat{a}^{\dagger 2} \hat{a}^2$ | 600 ms | 3 ms |
| Forward Propagation | transmon T_1 | 200 μs | 600 μs |
| Net Lifetime | predicted | 200 μs | 320 μs |
| | measured | - | 318 μs |
| Gain Over Uncorrected Logical Qubit: | | 1.4 | 2.2 |
| Gain Over Best Physical Qubit: | | 0.7 | 1.1 |

Table 1. Failure modes of the corrected logical qubit. This table shows the predicted decay time constant (τ) of quantum information stored in a corrected logical qubit using the cat code paradigm under a scenario in which each individual failure mode is the only source of loss. Dominant modes of failure in the cat code are double errors (\hat{a}_s followed by \hat{a}_s) between consecutive syndrome measurements separated by a time t_w ; possible errors that the cat code does not address, such as additions of a single photon (\hat{a}_s^\dagger); a failed parity mapping resulting from ancilla dephasing (T_ϕ); incorrect ancilla initialization prior to syndrome measurement resulting from unknown excitations (Γ_\uparrow) of the ancilla during t_w ; undesired couplings that result in dephasing due to Kerr ($\hat{a}_s^{\dagger 2} \hat{a}_s^2$); and finally ancilla decoherence that directly propagates to unrecoverable errors in the cavity state, a result of ancilla decay or excitation (T_1). Two different measurement strategies are shown for an initial $\bar{n}_0 = 2$: as quickly as possible ($t_w \approx 0\mu s$) and the optimal monitoring time ($t_w \approx 20\mu s$). The lowest two rows show the multiplicative gains of cat code performance over the decay constants of the uncorrected logical qubit (147 μs) and the system’s best physical qubit (287 μs , Fock states $|0\rangle_f, |1\rangle_f$). These gains reflect the combined effects of all loss channels acting together. We predict all numbers using an analytical model derived in the following section (sec. 8.3) and show that for the net gains the failure modes do not contribute independently. Using the optimal measurement strategy, we find that the predicted gains in lifetime over the uncorrected logical qubit and over the Fock state encoding match the measured performance of the corrected cat code (318 μs). Lifetimes of at least 6 ms would be possible if the forward propagation of errors from the syndrome measurements were abated.

We stress, however, that the sources of loss detailed in the sections below do not act independently when considering the experimental reality. Indeed, simply adding all of the rates in parallel leads to an underestimate of the cat code performance. In sec. 8.3 we analyze the system as a whole and show that we can analytically predict the data for the cat code decay shown in Fig. 7.7.




Figure 8.7: **Process fidelity decay for a larger encoding.** An identical plot to that shown in Fig. 7.7, except the initial cat size here is $\bar{n}_0 = 3$. Every time constant involving cat code is lower than for $\bar{n}_0 = 2$ due to the increased error rate in the cat states for higher photon numbers, $\gamma = \bar{n}\kappa_s \approx 3\kappa_s$ versus $\gamma \approx 2\kappa_s$. On the top axis we plot the number of syndrome measurements used for each point in the corrected cat code; note that for this larger encoding we typically use more measurements for a given time step than for cat states of $\bar{n}_0 = 2$.

Double-errors

The cat code is a first-order code, which means that the error syndrome we employ cannot detect the occurrence of multiple errors between two consecutive measurements. The probability of such events, $p_{2\varepsilon}$, can be calculated from the Poisson distribution:

$$p_{2\varepsilon}(t_M) = \frac{(\bar{n}\kappa_s t_M)^2}{2} e^{-\bar{n}\kappa_s t_M}, \quad (8.16)$$

where we take the approximation that \bar{n} is constant throughout the small time interval t_M . In this expression, $t_M \approx t_w + 1\mu\text{s}$ is the total measurement time; t_w is the time delay between the end of one syndrome measurement and the beginning of the next; and the parity mapping together with ancilla readout totals $\sim 1\mu\text{s}$ (see Fig. 6.3 for exact timings).

The average time between photon jumps is given by $1/\bar{n}\kappa_s$. A simple calculation using Eq. 8.16 for measurement intervals $t_M \approx 1\mu\text{s}$ and $t_M \approx 21\mu\text{s}$ returns the predicted gains in the process




Figure 8.8: **QND parity measurements in the full QEC system.** We repeat the same experiment demonstrated in 5.3, which quantifies the effect of measuring parity on the effective cavity decay rate $1/\tau_{tot}$. The main plot shows three decay curves of average parity $\langle \hat{P} \rangle$ versus time for three different measurement repetition intervals: $2 \mu s$, $5 \mu s$, and $20 \mu s$; the initial displacement is $|\alpha| = 2$. The inset shows a plot of the extracted time constants for these three and other points. The data fits well to a model in which the natural decay rate acts in parallel with an induced decay rate p_d/τ_{rep} , for τ_{rep} the repetition interval. In this experiment we find p_d to be 0.1% per measurement, in agreement with the results in [Sun et al., 2013].

fidelity lifetime over a Fock state encoding, τ_{f01} , one would expect to see if missing such events were the only source of error. Defining a gain $G_{2\varepsilon}(t_M) = t_M/(p_{2\varepsilon}\tau_{f01})$, we find:

$$G_{2\varepsilon}(1 \mu s) \approx 125 \quad (8.17)$$

$$G_{2\varepsilon}(21 \mu s) \approx 6, \quad (8.18)$$

reproducing the results presented in Tab. 8.6. Figure 8.7 shows an example of the effect of greater redundancy on the performance of the cat code. With a cat states of $\bar{n}_0 = 3$, the greater rate of errors in the logical space reduces the fidelity of each syndrome measurement and likewise necessitates more frequent syndrome interrogation. As a result, the corrected cat code underperforms the Fock state encoding. Once again, however, it overcomes the decoherence penalty, this time by a factor of ~ 3.3 .

We also quantify how Quantum Non-Demolition (QND) our parity measurements are, just as in chapter 5. Ideally, for perfectly QND measurements $1/\tau_{tot}$ should match the natural decay rate $1/\tau_s$, regardless of the measurement rate. In reality, there is a small probability P_D that by measuring parity we induce more photon jumps. Using the methods studied extensively in 5.3, we find that $P_D = 0.001$ in this system (Fig. 8.8), comparable to our previous result [Sun et al., 2013]. The

probability of dephasing, however, is P_D^2 , since this is the probability of inducing two jumps in a row; this effect is negligible.

Note that regardless of measurement rate each curve saturates at a measured $\langle \hat{P} \rangle \approx 0.95$. This is consistent with a thermal population of the storage cavity $n_{th}^s < 2\%$, which reduces the average parity from ideally +1 of the vacuum to $\langle \hat{P} \rangle \approx 1 - 2n_{th}^s$, and a parity measurement fidelity of 98.5% for no photons in the cavity. Such performance would be impossible without the crucial application of the adaptive parity monitoring protocol (see sec. 6.3) and highlights the advantages of always maximizing the probability that the ancilla returns to the ground state.

Uncorrectable errors

The cat code cannot distinguish between photon loss (\hat{a}_s) and photon gain (\hat{a}_s^\dagger). The probability of excitation due to \hat{a}_s^\dagger is given by $p_{\uparrow s}(t_M) = t_M n_{th}^s \bar{n} / \tau_s$. Given the low thermal population and high coherence of the cavity, we expect an \hat{a}_s^\dagger event on average every ~ 6 ms for $\bar{n} = 2$, a rate of thermal excitation that is negligible compared to all other sources of loss. If this were the only source of code failure, the gain $G_{\uparrow s} = t_M / (p_{\uparrow s} \tau_{f01})$ would be independent of t_M and equal to approximately 20, as given in Tab. 8.6.

When these currently uncorrectable sources of error become dominant, the redundancy of the cat code can be augmented by increasing the size of the logical basis states to superpositions of three coherent states (and higher) [Leghtas et al., 2013]. Although coherent states are not eigenstates of \hat{a}_s^\dagger , for large enough amplitudes the addition of a single photon results in a distortion in the Poisson coefficients that can still be corrected by the pumping scheme described in [Leghtas et al., 2015a].

Readout error

During the parity mapping sequence, ancilla dephasing due to T_2 is the dominant contribution to the overall drop in parity measurement fidelity. The parity of the state of course does not change upon an errant measurement, but our reaction to the result within the experimental flow does (see sec. 6.4). Were it not for the detrimental effects of ancilla back-action (see sec. 8.2), the optimal approach would be to measure as quickly as possible to build up measurement statistics (see sec. 8.1). Indeed this was the strategy implemented in chapter 5, where the goal was to understand the dynamics of photon jumps and so the forward propagation of ancilla decoherence was not considered. A quantum filter used Bayesian statistics to best estimate the parity of the state at any given time, and it would take about three consecutive agreeing measurements for the filter to completely switch from one parity to another. In this work, with improved fidelities and lifetimes we understand that with a

measurement cadence of $t_M \approx 1 \mu\text{s}$ it would take roughly two measurements ($2 \mu\text{s}$ total) for an equivalent filter to converge on a parity with high probability. If a photon jump occurs within this effective bandwidth, the filter will not detect it, resulting in a readout error. With average photon jump times on the order of $140 \mu\text{s}$ for an $\bar{n} = 2$ in the cavity, the probability of missing a jump is therefore $p_{mj} \approx 2/140 \approx 1.5\%$. The gain is therefore approximately equal to $140 \mu\text{s}/(p_{mj}\tau_{f01}) \approx 25$, as in Tab. 8.6.

For the optimal measurement cadence $t_M \approx 21 \mu\text{s}$ for $\bar{n} = 2$, given that the probability to have photon jumps within t_M is approximately 20%, the optimal strategy is to trust each result implicitly (see sec. 8.3). If ancilla dephasing were the only source of error, after the syndrome mapping time π/χ_{sa} the purity of the ancilla state would decrease to approximately $\pi/(\chi_{sa}T_2) \approx 0.98$. The remaining 2%, which is incoherent mixture, would be measured to have the correct syndrome mapping result half of the time. The probability to dephase the logical qubit due to an errant syndrome measurement result becomes $p_\phi \approx 1\%$. The gain $G_\phi(t_M) = t_M/(p_\phi\tau_{f01}) \approx 7$ returns the result presented in Tab. 8.6.

Ancilla preparation

After every syndrome measurement, we reinitialize the ancilla to $|g\rangle$ regardless of the result (see sec. 6.3). Given its finite rate of excitation Γ_\uparrow , after t_M the ancilla may no longer be in the ground state with a probability $p_{\uparrow a} = \Gamma_\uparrow t_M$, which leads to an errant subsequent syndrome measurement. With a maximal syndrome measurement rate, ancilla preparation errors are negligible as some type of majority voting can be performed on groups of measurements to filter out this effect. Errors only occur on the order of $p_{\uparrow a}^2$ for majority voting in groups of three. The gain is therefore $G_{\uparrow a}(t_M) = t_M/(p_{\uparrow a}^2\tau_{f01}) \approx 2000$, reproducing the result in Tab. 8.6. For the optimal rate, however, this gain is limited by Γ_\uparrow only, and thus $G_{\uparrow a}(t_M) = 1/(\Gamma_\uparrow\tau_{f01}) \approx 3$. This mechanism is of course also responsible for cavity dephasing, as is described in section 8.2, and could be mitigated by stabilizing the ancilla ground state during t_w .

Orthogonality of basis states

Using a Python quantum simulation software package called QuTiP [Johansson et al., 2012, 2013], we simulate the effect of the increasing non-orthogonality on the efficacy of the optimized decoding pulses to faithfully transfer an encoded state in the cavity back onto the ancilla (Fig. 8.9). We find that with an initial encoding size of $\bar{n}_0 = 2$ and after a time of $\sim 100 \mu\text{s}$, the error in the decoding pulses for even parity states approximately equals 2-3%, while for the odd parity states it

approximately equals 6-7%. Using the Poisson distribution to calculate the percentage of even and odd parity states after $\sim 100 \mu\text{s}$, we find that the resulting infidelity due to overlapping basis states at the end of the tracking sequence amounts to roughly 4-5%. For earlier times, this error rapidly decreases toward 0, indicating that even for small cat sizes, the approximation that the basis states are orthogonal is still quite accurate.




Figure 8.9: **Decoding pulse fidelity as a function of time.** These numbers are obtained from the fidelities predicted by the Python QuTip [Johansson et al., 2012, 2013] quantum simulation package. Consistent with the analytical expressions for cat code basis overlap plotted in Fig. 3.10, the optimal control map back pulses have a lower fidelity for odd parity states. As expected, the fidelity is also lower when the size of the initial encoding is $\bar{n}_0 = 2$ rather than $\bar{n}_0 = 3$.

Code-space leakage

To a good approximation, for $\bar{n} \lesssim 2$ the basis states of the cat code can be interpreted as superpositions of only the Fock states $|0\rangle_f \rightarrow |7\rangle_f$. In turn, this restricted space can be described in a binary representation that requires just three physical qubits (Fig. 8.10), with coefficients given by the Poisson distribution of a coherent state of $\alpha \lesssim \sqrt{2}$ and with d_0 the least significant bit in $|d_2 d_1 d_0\rangle$ (e.g. $|110\rangle \equiv |6\rangle$). In this formulation, d_0 is the “Parity Bit:” when $d_0 = 0$ the parity of the state is even and when $d_0 = 1$ the parity is odd. Note that the even and odd logical basis states are still all mutually orthogonal.

Although in principle such an encoding scheme can be fashioned in a cQED system with three transmons, the error processes would be completely different, dominated by single transmon energy decay and dephasing rather than the correlated shift of bits arising from the action of some effective lowering operator. The utility of this representation, however, is that it emphasizes the possibility of excursions out of the code space; for increasing deviations ϵ_n of the coefficients c_n from their ideal

3-Qubit Binary Representation (valid for $\bar{n} = |\alpha|^2 \lesssim 2$)

$$\begin{array}{l}
 |C_\alpha^+\rangle \approx c_0 |000\rangle + c_2 |010\rangle + c_4 |100\rangle + c_6 |110\rangle \\
 |C_{i\alpha}^+\rangle \approx c_0 |000\rangle - c_2 |010\rangle + c_4 |100\rangle - c_6 |110\rangle \\
 |C_\alpha^-\rangle \approx c_1 |001\rangle + c_3 |011\rangle + c_5 |101\rangle + c_7 |111\rangle \\
 |C_{i\alpha}^-\rangle \approx c_1 |001\rangle - c_3 |011\rangle + c_5 |101\rangle - c_7 |111\rangle
 \end{array}
 \quad
 \begin{array}{l}
 \text{Fock states in binary} \\
 |n\rangle \rightarrow |d_2 d_1 d_0\rangle \\
 \text{Poisson Coefficients} \\
 c_n = e^{-\frac{|\alpha|^2}{2}} \frac{\alpha^n}{n!}
 \end{array}$$

Figure 8.10: Cat code basis in a three-qubit binary representation. The cat code logical basis states $|C_\alpha^\pm\rangle$ and $|C_{i\alpha}^\pm\rangle$ can be expanded in the Fock basis, where each component is rewritten in binary. For small cat sizes of $|\alpha|^2 \lesssim 2$, only three physical qubits are necessary to realize such a representation to high accuracy, as the Poisson coefficients c_n for Fock states greater than $|111\rangle \equiv |7\rangle$ become vanishingly small. Should these coefficients deviate from their specified values without our knowledge, the integrity of the quantum information may start to suffer.

values as specified by the Poisson distribution, the overlap $(\langle C_\alpha^+ | + \sum_{n=0}^7 \epsilon_n \langle n |) |C_\alpha^+\rangle \rightarrow 0$. This effect is of course continuous.

One may note that the Kerr of the cavity immediately changes these coefficients at a rate K_s . This effect, however, is deterministic, does not change the parity of the state, and in fact periodically brings the coefficients back to their original values (minus the effect of amplitude decay) [Kirchmair et al., 2013]. It therefore does not constitute a source of dephasing since it can be taken as just a continuous change of basis in time. As long as we take this into account when decoding the encoded state back onto the ancilla at the end of our protocol, no information is lost (sec. 6.3). There are, however, several non-deterministic effects that do constitute dephasing, all arising from undesired interactions of the cavity with the ancilla, with the readout cavity, and again with itself (a second effect of Kerr to be described shortly). Some of these sources of loss are possible to partially recover from even in the current implementation of the experiment, while others are a central vehicle of non-fault-tolerance in this system.

Undesired couplings – self-Kerr & cross-Kerr

The non-commutativity of the cavity’s Kerr Hamiltonian and the annihilation operator $[\frac{K_s}{2} \hat{a}_s^{2\dagger} \hat{a}_s^2, \hat{a}_s] \neq 0$ results in the following relation [Leghtas et al., 2013] (excluding an irrelevant global phase):

$$\hat{a}_s e^{-i\frac{K_s}{2} t \hat{a}_s^{2\dagger} \hat{a}_s^2} = e^{-i\frac{K_s}{2} t \hat{a}_s^{2\dagger} \hat{a}_s^2} e^{-iK t \hat{a}_s^\dagger \hat{a}_s} \hat{a}_s. \quad (8.19)$$

This means that every time a photon jumps at a time t_j , the cavity state is rotated in phase space by an angle $\theta = K_s t_j$. Without an infinite cadence of measurement, however, there is always some finite uncertainty in t_j , δt_j , and consequently in θ , $\delta\theta$. A non-zero $\delta\theta$ results in an angular spread of cat

code basis states in phase space. On a shot-by-shot case, it means that we lose track of the phase of the cavity state within the angular window defined by $K_s t_M$. In other words, the state leaks out of the code space. We study the effects of such leakage by encoding a quantum bit into our code-space and then immediately thereafter intentionally decoding back at the wrong phase (Fig. 8.11). The resulting Gaussian curve allows us to quantify the sensitivity of the cat code to uncertainties in the jump time.




Figure 8.11: Fidelity as a function of incorrect decoding pulse angle. Measured loss of fidelity as a function of an intentional phase offset θ of a decoding pulse that immediately follows quantum bit encoding. With a standard deviation of $\sim 24^\circ$, the broad Gaussian fit shows that for small rotations the fidelity suffers only quadratically. Hence, the Kerr-induced rotation per photon jump is not a major source of dephasing for low jump numbers even with a $t_w \approx 20 \mu\text{s}$. However, as ancilla decay during mapping can rotate the state by any angle, it causes a substantial degradation in process fidelity.

If we were to measure the error syndrome very quickly, we would know t_j to high accuracy and the resulting error due to uncertainty in jump time would be negligible given the second-order dependence of process fidelity on decoding angle. With a $t_M \approx 1 \mu\text{s}$, the average angle of rotation $\bar{\theta} \approx 1^\circ$. For this cadence, we can assume a uniform probability distribution that gives an uncertainty in angle of $\delta\theta \approx 0.5^\circ$. Using the result of the fit in Fig. 8.11, we average over the Gaussian distribution within a window of $\pm 0.8^\circ$ to find a probability of dephasing $p_K \approx 0.02\%$, an expectedly minor contribution. Weighting p_K by the probability to have a jump within t_M , which is on the order of 1%, the predicted gain in such a scenario where this is the only source of error is consequently very high: $G_K = t_M / (p_K \cdot 0.01 \cdot \tau_{f01}) \approx 2000$, as in Tab. 8.6.

Given the necessity of spacing out parity measurements in time by t_w in order to maximize lifetime gain, however, the absolute time of the jump and thus the value of θ inherit some non-

negligible uncertainty, resulting in code space leakage. In other words, the coefficients c_n now deviate from the Poisson distribution by ϵ_n that are unknown. For a typical $t_w \approx 20 \mu\text{s}$ and the value of the cavity's Kerr, the uncertainty in jump angle is $\sim 10^\circ$, resulting in a $\sim 3\%$ loss of fidelity. This loss of course increases with a greater error count. Assuming the probability of detecting a photon jump is again $\sim 20\%$ per step, the loss in process fidelity is $p_K \approx 0.2 \times 0.03$. The gain $G_K = t_M/(p_K\tau_{f01}) \approx 10$, as in Tab. 8.6. The rate K_s , however, is on the order of several kHz, and so in principle we can completely recover from this minor dephasing by interleaving the parity measurements with the dissipative pumping scheme demonstrated in [Leghtas et al., 2015a], which pumps and refocuses slightly dephased cat states back to the original logical basis (restoring their amplitude as well).

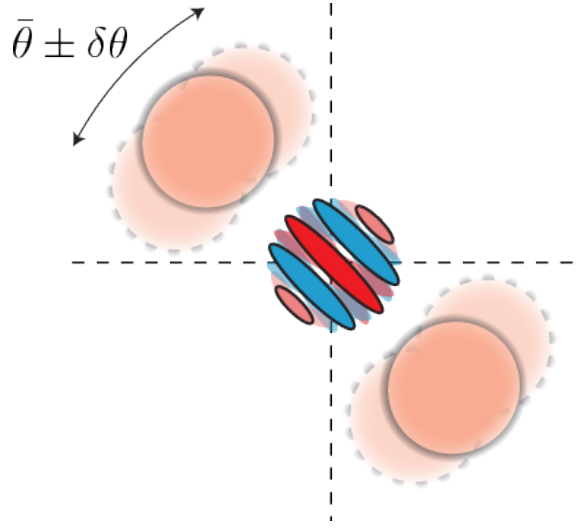


Figure 8.12: Effect of uncertainty in photon jump time. The state in the cavity acquires an extra phase θ every time a photon jump occurs that is proportional to the Kerr and the absolute time of the jump t_j : $\theta = K_s t_j$. As the parity measurements are performed approximately every $20 \mu\text{s}$, there is an inherent uncertainty associated with the absolute time of the jump, which causes some degree of dephasing. This cartoon illustrates this feature; although we know the average phase rotation $\bar{\theta}$, the unknown exact jump time results in an uncertainty in this angle $\delta\theta$. If this uncertainty is small, as seen in Fig. 8.11 the efficacy of the decoding pulse is not substantially degraded. This effect encourages a greater syndrome measurement frequency.

Likewise, code space leakage occurs whenever a coherent state is injected into the readout cavity to measure the state of the ancilla. As coherent states have an uncertainty in their average photon number of $\sqrt{\bar{n}}$, the cross-Kerr interaction leads to a dephasing of the encoded state at a rate proportional to χ_{sr} . Quantitatively, per measurement we see an average rotation of $\bar{\phi} = \bar{n}\chi_{sr}\tau_{\text{meas}} \approx 20^\circ$ of the cavity state in phase space for a readout pulse duration τ_{meas} . Given the value of χ_{sr} , we estimate that each readout pulse contains $\bar{n} \approx 70$ photons. The uncertainty in the rotation scales as the square root of \bar{n} : $\delta\bar{\phi} = \sqrt{\bar{n}}\chi_{sr}\tau_{\text{meas}} \approx 2^\circ$. With n measurements, the total uncertainty in the

angle is $\sqrt{n}\delta\bar{\phi}$. For example, after ten measurements this uncertainty is still much smaller than the standard deviation of the Gaussian in Fig. 8.11. Given the minimal effect on the process fidelity, this source of dephasing is excluded from the discussion. We therefore treat the cavity anharmonicity and coupling to the readout cavity not as sources of non-fault-tolerance, but rather as necessary technical trade-offs that can in principle be effectively suppressed.

Forward propagation

Infidelities due to ancilla dephasing outlined in sec. 8.2 and the forward propagation of errors to be discussed in this section have a common denominator: ancilla decoherence. In the former, both phase flips and amplitude decay of the ancilla contribute to a decrease in parity measurement fidelity. In the latter, one can see by looking at a simplified system Hamiltonian that the frequency of the cavity depends on the state of the transmon:

$$\hat{H}/\hbar = \tilde{\omega}_a |e\rangle \langle e| + (\tilde{\omega}_s - \chi_{sa} |e\rangle \langle e|) a_s^\dagger a_s \quad (8.20)$$

Figure 8.13 depicts how one can model a parity measurement in a digitized version (Fig. 8.10) of the cat code. Employing a single ancilla, the parity measurement is nothing more than a CNOT gate between this ancilla and the parity bit d_0 , which specifies the state’s symmetry with respect to a 180° rotation. The CNOT is written here equivalently as a controlled phase gate between two Hadamard gates (H) [Nielsen and Chuang, 2010]. The higher parity bits, d_1 and d_2 , provide further information about the state’s symmetry properties with respect to 90° and 45° degrees. The first panel shows that with no ancilla energy decay, the parity mapping is perfect since it does nothing to coefficients c_n at the end of the protocol.

The length of this mapping, $\pi/\chi_{sa} \approx 250$ ns, however, is a small but non-negligible fraction of the ancilla T_1 . One can model this finite gate time by splitting the phase gate into two “controlled- $\pi/2$ ” gates and adding two phase gates to the next parity bit d_1 . With a perfect parity mapping one obtains the exact same results as in the first row. If the ancilla decays exactly half-way through the sequence, however, the cavity state inherits a phase of $\pi/2$ in phase space; this is a logical bit flip in our basis. One can continue splitting the gate into smaller and smaller pieces (e.g. third panel), where now the ancilla T_1 decay rotates the cavity state by an arbitrary angle that is known only if the time of ancilla decay is known. Of course, as we depend on the entangling interaction between the ancilla and cavity throughout the parity mapping time, in this implementation we have no way of detecting when such decay occurs. Equivalently, the photon number parity operator $\hat{P} = e^{i\pi\hat{a}_s^\dagger\hat{a}_s}$

Digitizing the Parity Measurement

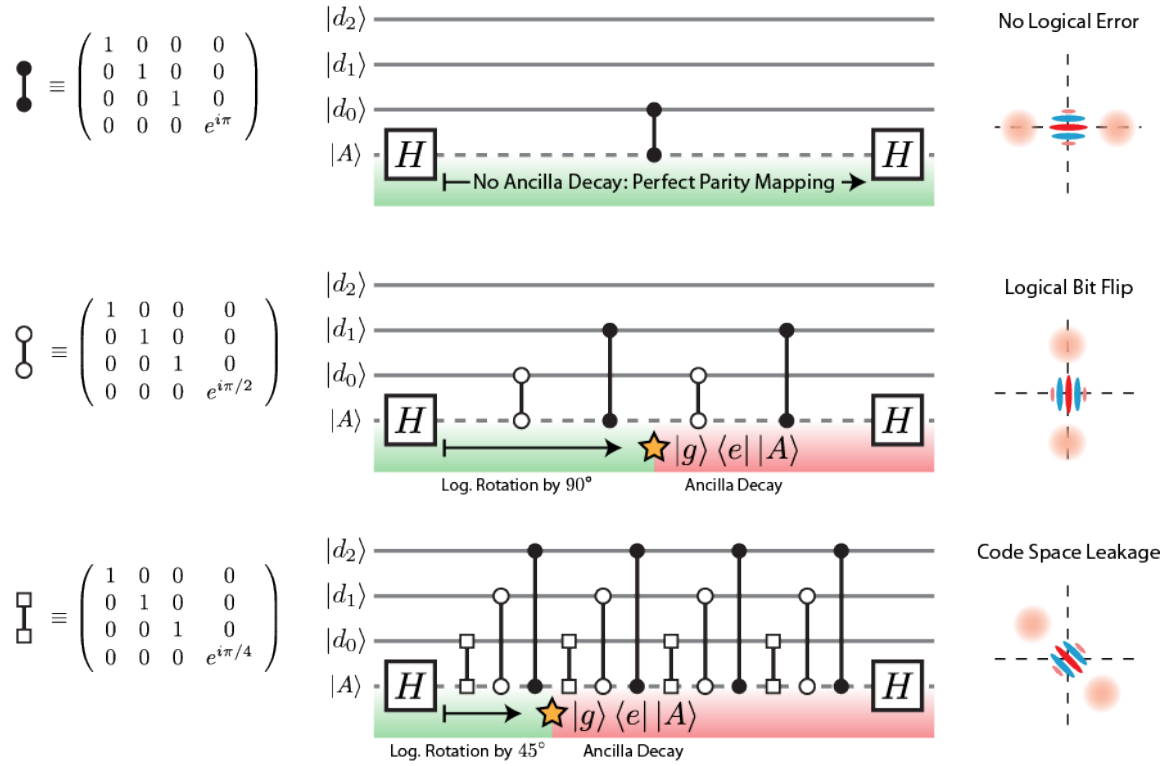


Figure 8.13: Effect of ancilla T_1 in a circuit model representation. Code space leakage can be particularly acute if the ancilla $|A\rangle$ undergoes energy decay (or excitation) during the parity mapping. Shown in the first panel is an ideal parity mapping using the binary representation. The least significant bit, d_0 , is the parity bit; the parity is even for $d_0 = 0$ and odd for $d_0 = 1$, regardless of d_1 and d_2 . The parity mapping is thus a simple CNOT gate, depicted here as a controlled phase (solid black circles, π phase shift) between two Hadamard gates (H). Such a circuit representation belies the fact that the mapping is finite in time, lasting $\pi/\chi_{sa} \approx 250$ ns. To obtain a better approximation of the true dynamics, we split the CPHASE into two pieces, where the two empty circles are now controlled phase gates with a $\pi/2$ phase shift. A simple calculation in the Fock basis demonstrates that if the ancilla suddenly decays to $|g\rangle$ exactly halfway through the mapping, a logical bit flip occurs. For arbitrary decay times, we witness code space leakage, where the cat state is aligned with neither the real nor the imaginary axis. The third row shows an example of this for one more layer of granulation.

commutes with any rotation in phase space: $[\hat{P}, e^{i\theta\hat{a}_s^\dagger\hat{a}_s}] = 0$. The environment gains information that we do not. Beyond the risk of T_1 decay during the parity mapping, the equally detrimental effect of T_1 decay of the ancilla during the readout pulse (before ancilla reset, see sec. 6.3) reduces the fidelity of all trajectories in which one or more photon jumps occur. In addition, unknown ancilla excitations from $|g\rangle$ to $|e\rangle$ (or higher states), which occur at a rate proportional to Γ_\uparrow , dephase the cavity state similarly.

We can comb through the probabilities of a T_1 event in each of the three main steps throughout the entire duration of the sequence and calculate a gain in a manner similar to calculations in previous sections. During the parity mapping, the probability of ancilla decay is $p_{\downarrow,a,1} \approx \pi/(\chi_{sa} \cdot 2T_1) = 0.004$; the probability of measuring ancilla $|e\rangle$ equals the probability of measuring a photon jump, and so

the contribution to dephasing is $p_{\downarrow a,2} \approx (\bar{n}\kappa_s t_M)\tau_{\text{meas}}/T_1 = 0.0001$, the probability of a photon jump times the probability of T_1 decay in a duration τ_{meas} ; finally, during t_M there is always the risk of ancilla excitation, with a probability $p_{\uparrow a,3} \approx t_M\Gamma_{\uparrow} = 0.001$. The total probability of forward propagation is $p_{fp}(t_M) \approx p_{\downarrow a,1} + p_{\downarrow a,2} + p_{\uparrow a,3}$. For $\bar{n} = 2$, we find that $p_{fp}(1 \mu\text{s}) \approx 0.5\%$. Defining the gain $G_{fp}(t_M) = t_M/(p_{fp}\tau_{f01})$, we find $G_{fp}(1 \mu\text{s}) \approx 0.7$. Performing the same calculation for $t_M \approx 21 \mu\text{s}$, we find $G_{fp}(21 \mu\text{s}) \approx 2$. We thus obtain the numbers in the final row of Tab. 8.6, and arrive at the key constraint of our system: measuring more frequently enhances the likelihood of forward propagation of errors. As seen in the preceding sections, by mitigating this decoherence channel we stand to gain substantially in all other aspects with faster syndrome measurements.

8.3 Optimizing cat code performance

The presence of forward propagating errors in our system due to ancilla decoherence substantially alters the strategy one normally seeks to employ in a QEC system. Typically, the goal is to suppress the occurrence of errors within the code-space to second order with measurements performed at the maximum rate permitted by the parameters of the system. With such a strategy, however, one necessarily entangles the logical states with the ancillary systems needed to extract the error syndromes for a substantial fraction of the QEC protocol's duration. As the rate of photon jumps in our system is much lower than $1/\tau_{\text{meas}}$, the probability that two errors occur within τ_{meas} is considerably lower than the probability of ancilla-induced decoherence. Such a strategy thus results in a net-loss.

We thus explore a different approach, one that instead slows down the syndrome measurement cadence to find an optimal balance between errors in the code and ancilla induced dephasing. We take the point of view that experimentally our task is to preserve a quantum bit of information for a total time T . The analytical model we present below then calculates the optimal measurement cadence and the predicted lifetime. It takes into consideration the basic measured parameters in our system: cavity and ancilla coherence properties ($\tau_s = 250\mu\text{s}$, $T_1 = 35\mu\text{s}$, $T_2 = 12\mu\text{s}$), thermal populations, and measurement fidelities. The predictions we arrive at closely match the data we present in sec. 7.3. Such results indicate that not only can we successfully optimize and employ a measurement strategy that preserves a quantum bit beyond the break-even point, but given only a basic set of assumptions about the sources of loss in the system, namely those outlined in sec. 8.2, we capture the dominant mechanisms that set the performance metrics of our QEC system. Using this model, we can then predict the potential gains we expect to witness when certain key system parameters are enhanced. In particular, by improving ancilla T_1 to levels of $\sim 100 \mu\text{s}$ (well within the

range of current transmon technology), we enter the regime in which the optimal strategy encourages measurements that are fast enough to warrant majority voting in real-time, thereby substantially mitigating decoherence due to ancilla T_2 . With this single improvement, the cat code promises to provide gains of over an order of magnitude.

A simplified model

As a brief primer, we can define the effective cavity decay rate, κ_{eff} , which predicts the gain one would expect to see in a parity tracking protocol over the natural cavity photon jump rate $\bar{n}\kappa_s$ (or the Fock state encoding, if scaled appropriately). This rate be described by the sum of all dominant error rates including the three sources of infidelity described above plus a fourth to include readout errors. This κ_{eff} should also be a function of the number of consecutive parity measurements, N , and a waiting time τ_W between each one of these N -measurement “packets” (Fig. 8.14). Quantifying the parity tracking performance as the foundation for a QEC protocol comes down to finding the optimal N and τ_W that give the lowest κ_{eff} , given the realistic experimental parameters at hand. Explicitly, we have:

$$\kappa_{eff} = \left[\frac{(\bar{n}\kappa_s)^2(N\tau_M + \tau_W)^2}{2} + N\epsilon_{T_1} + \mathcal{O}(\epsilon_{T_\phi}^m) + \mathcal{O}(\epsilon_M^m) \right] \frac{1}{N\tau_M + \tau_W}, \quad (8.21)$$

where τ_M is the time it takes to perform one parity measurement; $\epsilon_{T_1} \sim \tau_M/T_1$; $\epsilon_{T_\phi} = \pi/\chi_{sa}T_\phi$; ϵ_M are the loss of cat fidelity due to ancilla T_1 , T_ϕ , and measurement inaccuracy respectively; and $m = \frac{N+1}{2}$.




Figure 8.14: **Schematic of an optimized parity monitoring protocol.** An optimized parity tracking scheme would involve performing packets of N measurements, each lasting a time τ_M , followed by a waiting time of τ_W .

Central to this paradigm is understanding that the incorrect knowledge of the number of parity jumps is the first major contribution to cat state decoherence. Single photon jump events can be tracked as demonstrated in the previous chapters, and therefore the corresponding phase errors that arise from the stochastic application of \hat{a}_s on the logical qubit can be corrected. The term that goes

like $(\bar{n}\kappa_s)^2$ determines the probability instead to have two or more jumps within the time window $(N\tau_M + \tau_W)$. Measurement infidelity in the form of ancilla dephasing ϵ_{T_ϕ} and ϵ_M contributes to misleading the observer that a jump has occurred, which is equivalent to losing track of the number of jumps. These terms contribute to the κ_{eff} expression exponentially, and in a “majority vote” fashion (assuming N is odd), which explains the choice of exponent in ϵ_{T_ϕ} and ϵ_M . Therefore, in principle the errors due to ϵ_{T_ϕ} and ϵ_M can drop out by either simple majority voting or the more elaborate quantum filtering. Collectively, these three terms correspond to cat state dephasing in the event of the incorrect tabulation of the number of jumps the cat state has undergone during its evolution. The second major contribution to decoherence, ϵ_{T_1} , comes from cat state dephasing due to T_1 . It only takes one ancilla decay event within a string of N measurements to completely dephase a cat state, which is why ϵ_{T_1} grows linearly with N . Therefore, an optimal balance needs to be struck between gaining information about photon jumps and dephasing the cavity state by inquiring too frequently.

The full optimal strategy

In our experimental implementation, the optimal strategy is clearly to measure so infrequently that each measurement result must be taken at face value (Tab. 8.6). At the moment immediately prior to the application of the decoding pulse, all we have in our possession is a record of measurements that we can use to deduce the error trajectory of the state that was encoded in the cavity. As a function of this classical information we then act on the system and decode the state to the best of our knowledge. There are two questions we can ask:

- a. For a given trajectory of photon jumps, what is the probability that the conclusion we obtained, based on the measurement record, is correct?
- b. What is the probability distribution of possible trajectories that may produce a given measurement record?

The first question relates to the optimization process. We wish to maximize the probability of correctly identifying the actual cavity error trajectory. Trajectories with either many errors or consecutive errors result in a lower probability of success, as shown in sec. 8.1.

The second question relates to our confidence in the output. There are different error trajectories that can produce the same measurement record. The best strategy is to simply choose the most probable one. Our confidence in the output is then the probability that this scenario occurred conditioned on the measured record. The output is thus not only the final state, but also a measure

of confidence. We can either ignore this extra classical information and treat the whole process as a black box (red curve, Fig. 7.7) or we can also use this information to post-select the data relative to some required confidence constraints (blue curve, Fig. 8.1).

Below we summarize our findings from an analytical model fully derived in appendix B. For a fixed, desired time T for which we would like to correct a quantum bit, we define a configuration to be the combination of the following parameters:

1. The initial Cat size, $\bar{n}_0 = |\alpha|^2$.
2. The number of parity tracking steps S .
3. The step durations $\{t_1, t_2, \dots, t_S\}$, where $\sum_{k=1}^S t_k = T$.

The process fidelity, $F_{process}$ decays exponentially from 1 to 1/4. Here we derive and optimize the scaled version which decays from 1 to 0 at exactly the same rate, which we denote as the *fidelity*:

$$F \equiv \frac{F_{process} - 1/4}{3/4}. \quad (8.22)$$

This is the probability we successfully corrected the state. We can write the total fidelity as a product of four terms:

$$F(T, \bar{n}_0, S, \{t_k\}_{t=1}^S) = F_{\Gamma_\uparrow}(T) \cdot F_{ED}(T, \bar{n}_0) \cdot F_T(T, \bar{n}_0, S, \{t_k\}_{t=1}^S) \cdot F_{KD}(T, \bar{n}_0, S, \{t_k\}_{t=1}^S). \quad (8.23)$$

F_{Γ_\uparrow} Whenever the ancilla is excited to $|e\rangle$ we lose the encoded information. This affects our protocol and also any other storage scheme that uses the cavity equally, since the ancilla's rate of excitation should be independent of the state in the cavity. The term F_{Γ_\uparrow} depends on T alone and equals $e^{-T \cdot \Gamma_\uparrow}$.

F_{ED} The fidelities of the encoding and decoding pulses depend on the initial and final cat sizes, which are \bar{n}_0 and $\bar{n}_0 e^{-\kappa_s T}$. The non-orthogonality is simulated and taken into account numerically (Fig. 8.9).

F_T The loss of fidelity due to the monitoring itself depends on the ancilla's figures of merit throughout the time-scale of the parity mapping and projective measurement. It also depends on the cat size and κ_s through the probability to miss photon jumps during a single step.

F_{KD} The uncertainty in the angle due to the rotation from the non-commutativity of the Kerr Hamiltonian and \hat{a}_s decreases the fidelity of the decoding pulse. We calculate this Kerr rotation

distribution from the number of expected photon jumps and the step lengths, together with the measured fidelity of the decoding pulse as a function of the angle (Fig. 8.11).

Ignoring F_{KD} for the moment, we can show that the optimal fidelity can be written in the following form:

$$F^{OPT}(T, \bar{n}_0) = e^{-T\Gamma_\uparrow} \cdot F_{ED}(T, \bar{n}_0) \cdot e^{-\bar{n}_0[1-e^{-\kappa_s T}]/G}, \quad (8.24)$$

where G , the system gain, is a function of the other system parameters (χ_{sa}, T_1, \dots) and is a constant of the system. When $\kappa_s T \ll 1$ we can approximate the optimized fidelity as:

$$F^{OPT}(T, \bar{n}_0) = e^{-T\Gamma_\uparrow} \cdot F_{ED}(0, \bar{n}_0) \cdot e^{-\kappa_s T \cdot \frac{\bar{n}_0}{G}}, \quad (8.25)$$

which shows that the decay rate of the quantum error corrected information is G/\bar{n}_0 slower compared to storage cavity decay rate κ_s . The process fidelity decay rate of the $|0\rangle_f, |1\rangle_f$ Fock state encoding decays 3/2 slower than κ_s . The break-even condition is therefore:

$$\frac{2G}{3\bar{n}_0} > 1. \quad (8.26)$$

As we increase \bar{n}_0 , the gain in lifetime decreases. On the other hand, in order to have sufficient orthogonality between the logical basis states, \bar{n}_0 has to be high enough. Hence, an optimal \bar{n}_0 exists. Since eq. 8.24 ignores F_{KD} , it expresses an upper bound for the fidelity, and thus G has to be even larger in order to get an actual gain in lifetime. Better ancilla coherence times will increase the optimal measurement cadence and make F_{KD} approach unity (see sec. 8.2).

Figure 8.15 shows how G depends on T_1 and T_ϕ of the ancilla. With our ancilla's coherence times, G is about 5. With $\bar{n}_0 = 2$, the ratio $2G/3\bar{n}_0$ equals 1.65. The actual gain is lower, however, since we need to take into account the effect of the Kerr rotation and the degradation of the decoding pulse due to loss of orthogonality. We can take these effects, F_{KD} and $F_{ED}(T)$, into account and optimize the configuration when fixing the number of steps. Figure 8.16 displays the measured fidelity and our model's prediction for the same configurations. It also displays the optimal expectations when forcing different numbers of steps. Our model predicts the measured data accurately and lets us find the optimal configuration as a function of the total duration T . We end up with a total predicted gain of 10%, in line with the results in Fig. 7.7.

Just as in the work on the Bell-cat (chapter 4) and the repeated monitoring of photon number parity (chapter 5), we again arrive at a point in the story where we can feel comfortable that little, if anything, has been excluded from our understanding of this full QEC system. Taking stock of




Figure 8.15: **Expected gains in lifetime with a more coherent ancilla.** The gain G as a function of T_1 and T_ϕ of the ancilla. With the parameters of our ancilla we get $G = 4.96$, and the predicted gain over the Fock state encoding is 1.65. This value is higher than the measured 10% improvement since this plot does not include the effects of dephasing due to Kerr or the degradation of information due to overlapping logical basis states. A key point is that with ancilla coherence on the order of $100 \mu s$, we already expect to see gains of an order of magnitude.

what we've just learned through the analysis of the dominant sources of decoherence in the system and their effects on the decay of process fidelity with time, we are encouraged to press forward and confront the issue of fault-tolerance head-on. Far from being discouraged that once again it seems like the T_1 of a physical qubit stands in the way of realizing a QEC system that can enhance the lifetime of a quantum bit for periods of time to be practically applicable, we instead find our




Figure 8.16: **Predicted process fidelity decay with optimal monitoring cadence.** The optimal predicted process fidelity for our system given ancilla and cavity coherence times, and cavity Kerr (squares). Comparing this simulation with the data (circles) shown in Fig. 7.7, we find that our model faithfully predicts the measurement results at each point. We also display the expected process fidelity we would have obtained had we fixed the number of steps, in colored curves. Commensurate with the top axis Fig. 7.7, we chose optimal configuration for each of the total time durations.

task to be simpler and more focused than before. There is no inherent limitation to the cat code. Even with the basic hardware that we’ve been using for several years now we can actually overcome the decoherence penalty of a bonafide QEC code. Instead of having to understand the dizzying array of (perhaps unknown) couplings and Hamiltonian parameters in large, multi-physical qubit systems, we take our rigorous understanding with us towards the natural and necessary future steps of increasing the redundancy of the cat code to handle ancilla back-action and the remaining entries in Tab. 8.6. The fact that our predictions are in such good agreement with the data gives us confidence in designing future hardware upgrades as a means to increasing the complexity of our quantum devices to the point that they can perform meaningful calculations. In conclusion, perhaps the most important contribution of our results thus far is that they have contributed to crystallizing our understanding of what fault-tolerance with the cat code will require and where to go next.

9 – Future Directions

In implementing the first quantum memory that outperforms the system’s most coherent element, we’ve taken a big step towards realizing fault-tolerant quantum computation [Devoret and Schoelkopf, 2013]. Under the umbrella of this single figure of merit, however, many important lessons serve to direct future endeavors in demonstrating the cat code as a fully fledged logical qubit on which logical operations can be performed and algorithms can be executed. We have shown the capability to efficiently encode and extract information from a continuous variable system through the violation of a Bell-inequality, where the choice of coherent states as the basis allows us to treat the state in the cavity to a good approximation as just another two-level system. The experimental verification of how photon loss leads to the degradation of the CHSH entanglement witness with increasing cat size confirms the assumption that single-photon loss is by far the dominant error mechanism in such cavities. A time constant of parity decay of a state in the cavity subjected to repeated error syndrome interrogation that matches the case of uninterrupted free evolution indicates that single-shot photon number parity measurement provide a fast, high fidelity, and QND means of monitoring the occurrence of errors. Increasing the complexity and sophistication of the architecture to a full quantum error correction system, in which real-time feedback plays a crucial role, has allowed us to benchmark for the first time the severity of each dominant source of decoherence that can afflict any QEC system. This result provides crucial information as to the single dominant source of non-fault-tolerance in this implementation. Given the ability to not only make up for the decoherence penalty associated with the redundant encoding of a quantum bit using the cat code, but moreover realize gains in lifetime over the system’s most coherent element, we have underscored the advantages a hardware-efficient architecture yields. Indeed, we’ve argued that ancillae with far inferior coherence properties can be utilized to great effect in enhancing the lifetime of quantum information.

In this chapter we explore some of the next pressing challenges that need to be faced by future endeavors in order to improve the performance of the QEC system further. We discuss the recent progress that has been made in combating sources of dephasing that parity measurements do not

address, namely the decay of the coherent states to vacuum, and an uncertainty in the location of the coherent states in phase space that comes about to the effects of higher order Hamiltonian terms, such as the cavity’s anharmonicity and its interaction with the readout cavity during the projective ancilla measurement. Furthermore, we discuss the necessity to move beyond implementing a quantum memory with a single logical qubit in one cavity to coupling multiple cavities to perform logical operations and the crucial requirement of correcting errors during logical gates. We then conclude with an outlook of how a modular architecture designed to incorporate all of these ideas presents an encouraging, viable, and exciting path toward realizing a quantum computer in the near-future.

9.1 Cat pumping that stabilizes the logical states

In order to move beyond the corrected quantum memory demonstrated here towards a scalable architecture with which we can perform actual computations, we need several advancements. First, we must suppress many other sources of dephasing in the system that fall under the general heading of code-space leakage. The photon number parity measurements provide essential information as to the occurrence of discrete errors in the state of the cavity. They do not, however, reveal anything about the code-space leakage due to ancilla state measurement, uncertainty in the photon jump time, and the overlap of basis states that constantly increases in time as the coherent states decay to vacuum. The goal is to thus confine any state in the cavity to a quantum manifold that is spanned by the coherent states that comprise the components of “2-cats” or “4-cats,” for example. Such stabilization, crucial to future cat code applications, has already been proposed [Mirrahimi et al., 2014] and successfully demonstrated in the impressive experimental work shown in [Leghtas et al., 2015a], in which the basis states $|\alpha\rangle$ and $|\!-\alpha\rangle$ were chosen. The idea is to engineer a driven dissipative process in which pairs of photons are added to and subtracted from the storage cavity. This can be achieved by using a Josephson junction as a 4-wave mixer, which when stimulated by an off-resonant pump tone converts photons from the readout cavity into two photons in the storage, and vice-versa (Fig. 9.1). The frequency of the pump tone must be chosen to obey a specific frequency matching condition to enable this process:

$$\omega_p = 2\tilde{\omega}_s - \tilde{\omega}_r, \tag{9.1}$$

where ω_p , $\tilde{\omega}_s$, and $\tilde{\omega}_r$ the pump, storage, and readout frequencies respectively. Applying the pump together with a weak drive on the readout cavity, one realizes a system that is well-described by the

following Hamiltonian in the RWA approximation:

$$\hat{H}_{sr}/\hbar = g_2^* \hat{a}_s^2 \hat{a}_r^\dagger + g_2 \hat{a}_s^\dagger \hat{a}_r + \epsilon_d \hat{a}_r^\dagger + \epsilon_d^* \hat{a}_r - \chi_{sr} \hat{a}_r^\dagger \hat{a}_r \hat{a}_s^\dagger \hat{a}_s - \frac{K_s}{2} \hat{a}_s^{\dagger 2} \hat{a}_s^2 - \frac{K_r}{2} \hat{a}_r^{\dagger 2} \hat{a}_r^2 \quad (9.2)$$

$$g_2 = \frac{\chi_{sr} \xi_p^*}{2} \quad (9.3)$$

$$\xi_p \approx \frac{-i\epsilon_p}{\frac{\kappa_r}{2} + i(\tilde{\omega}_r - \omega_p)}, \quad (9.4)$$

where κ_r is the decay rate of the readout cavity, ϵ_r and ϵ_p are the respective drive strengths for the readout and pump, and g_2 reflects the conversion of pairs of photons in the storage to single photons in the readout.

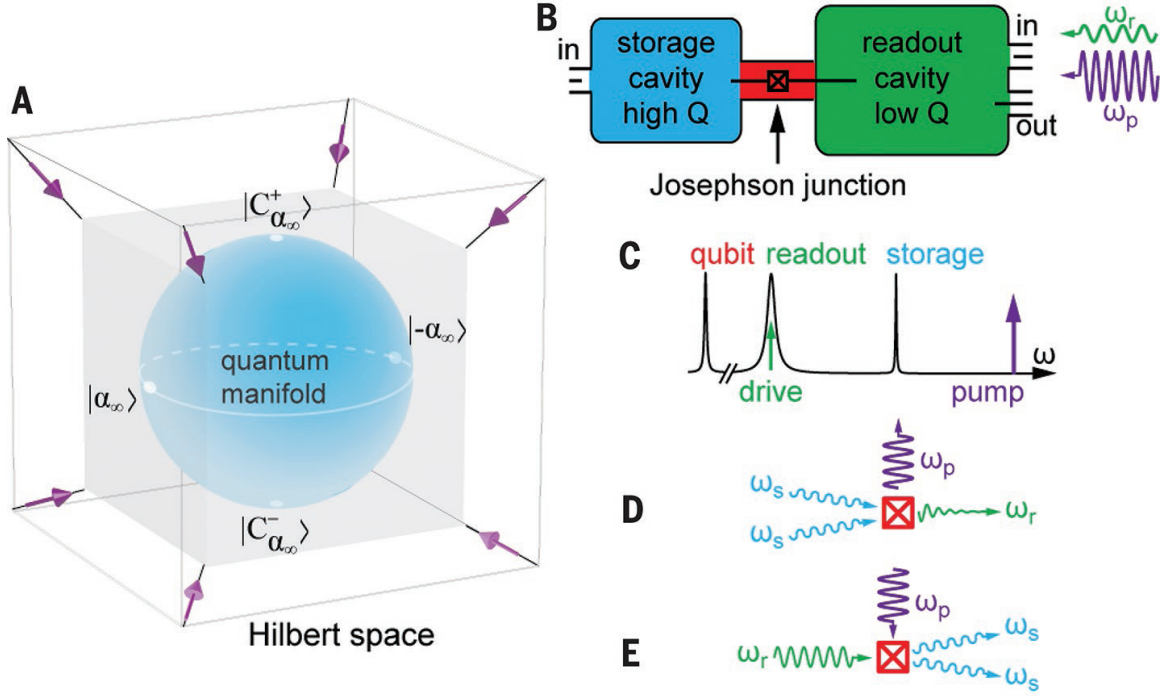


Figure 9.1: **Confining a cavity state to a manifold spanned by two coherent states.** (a). Confinement of a quantum state belonging to a large Hilbert space into a 2D quantum manifold. The outer and inner cubes form a hypercube representing a multidimensional Hilbert space. The inner blue sphere represents the manifold of states spanned by the two coherent states. Stabilizing forces direct all states toward the inner sphere without inducing any rotation in this sub-space, as indicated by the purple arrows. (b) Similar to the setups described in the previous experiments in this work, the long-lived storage cavity (high Q) and fast readout cavity (low Q) are bridged by an ancilla transmon. The pump, readout, and ancilla manipulation drives are applied through the readout input port. The input port of the storage is used to map out the Wigner function. (c) Cartoon spectrum of the relevant frequencies in the setup. (d), (e) Depictions of the 4-wave mixing process mediated by the Josephson junctions, where two storage photons are converted into a readout photon that is lost to the 50Ω environment of the output transmission line and two storage photons are created as a result of a drive on the readout.

One can then adiabatically eliminate terms including the readout in 9.2 to obtain the following, simplified form that describes the dynamics that govern the storage cavity [Leghtas et al., 2015b]:

$$\hat{H}_s/\hbar = \epsilon_2^* \hat{a}_s^2 + \epsilon_2 \hat{a}_s^{\dagger 2} - \frac{K_s}{2} \hat{a}_s^{\dagger 2} \hat{a}_s^2 \quad (9.5)$$

$$\epsilon_2 = -i \frac{\chi_{sr}}{\kappa_r} \xi_p^* \epsilon_d \quad (9.6)$$

$$\kappa_2 = \frac{\chi_{sr}^2}{\kappa_r} |\xi_p|^2, \quad (9.7)$$

where ϵ_2 is the nonlinear drive that inserts photon pairs into the storage and κ_2 is the decay rate of photon pairs. With this Hamiltonian, the storage cavity converges into the two-dimensional manifold spanned by the coherent states $|\pm\alpha_\infty\rangle$, where

$$\alpha_\infty|_{K_s, \kappa_s=0} = i \sqrt{\frac{2\epsilon_d}{\xi_p \chi_{sr}}}. \quad (9.8)$$

In the actual experiment in reference [Leghtas et al., 2015a], $\chi_{sr}/2\pi = 206$ kHz, $K_r = 2.14$ MHz, $K_s = 4$ kHz, $1/\kappa_s = 20$ μ s, and $1/\kappa_r = 25$ ns. Additionally, the applied pump strength resulted in $|\xi_p|^2 = 1.2$, and thus $g_2/2\pi = 111$ kHz and $\kappa_2/\kappa_s = 1.0$. With such parameters, the effects of the engineered bistability of the system are evident. Starting from coherent states at different angles in phase space, applying the pumping sequence stabilizes the states at $|\pm\alpha_\infty\rangle$ (Fig. 9.2). Moreover, by looking at snapshots of the cavity state during its evolution from vacuum to an equal superposition of $|\alpha_\infty\rangle$ and $|\alpha_\infty\rangle$, an even parity “2-cat,” one sees clear signatures of a quantum state from the negative values of the interference fringes in the Wigner function (Fig. 9.3).

As described in chapter 3, the cat code of course requires at least a “4-cat” for a single logical qubit in one cavity, which means that the pumping architecture must be modified to allow for the stabilization of a quantum manifold spanned by the four coherent states $|\alpha_\infty\rangle$, $|\alpha_\infty\rangle$, $|i\alpha_\infty\rangle$, and $|-i\alpha_\infty\rangle$. Such modifications, however, call for a rather involved redesign of the Josephson junction geometry. Perhaps a more advantageous scheme could instead rely on the “2-cats” that form the logical basis residing in two different cavities. Indeed, such a device has already been demonstrated [Wang et al., 2016] to show the entanglement between cat states in two cavities and the ability to measure a joint parity error syndrome with a new transmon qubit design (Fig. 9.4).

These results offer ample evidence that using such a pumping scheme to stabilize superpositions of coherent states can offer an effective means of suppressing forms of dephasing within the cat code paradigm aside from single photon loss. The parameters of the experiment described above also indicate, however, that one necessarily pays a price by employing such an architecture. The

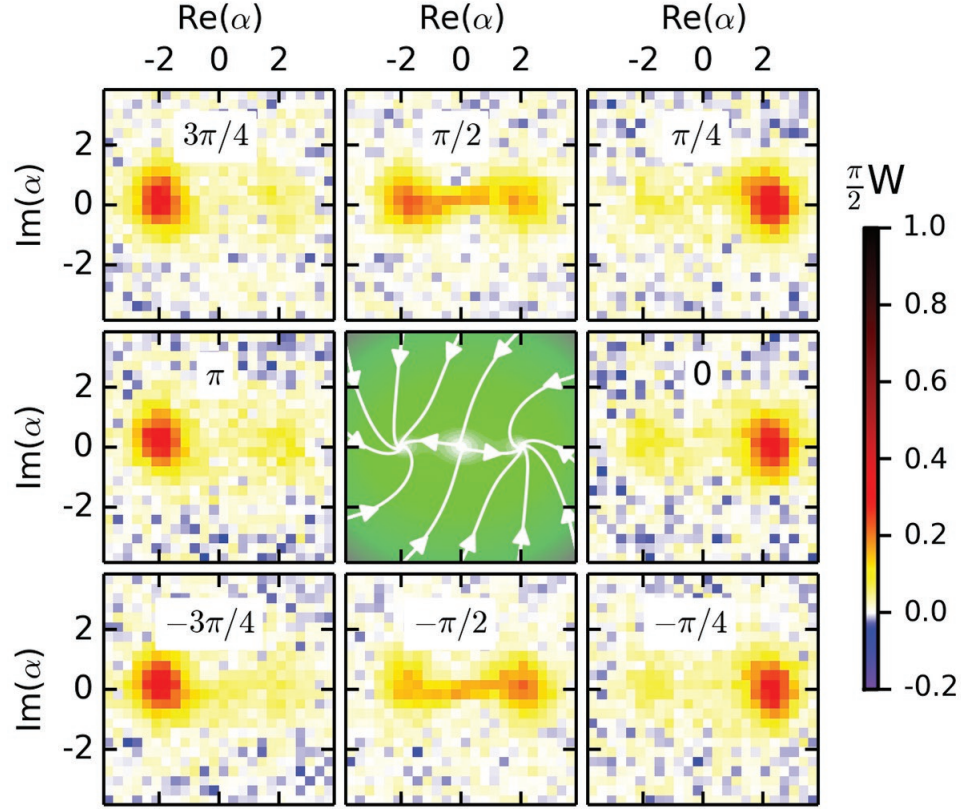


Figure 9.2: **Bistable behavior of the steady-state manifold.** The central panel shows the theoretical classical equivalent of a potential of the storage nonlinear dynamics. Shown are the locations of two stable steady states at $|\pm\alpha_\infty\rangle$ and a saddle point when the cavity is in vacuum ($|0\rangle_f$). The trajectories given by the white lines and arrows indicate the how an initialized state at the border converges to one of the two steady states; the curvature is a result of the cavity anharmonicity. The remaining eight outside panels show the measured Wigner functions after initial cavity displacements of $|\alpha| = 2.6$ at different angles followed by $10\ \mu\text{s}$ of pumping. For angles $0, \pm\pi/4$, the state mainly evolves to $|\alpha_\infty\rangle$, while for angles $\pi, \pm\pi/4$ the state evolves to mainly $|\alpha_\infty\rangle$. At angles $\pm\pi/2$, as the initial state is almost symmetrically positioned with respect to $|\pm\alpha_\infty\rangle$ the state evolves to a mixture of $|\pm\alpha_\infty\rangle$.

stronger coupling between the storage and readout cavities, necessary for the pumping to convert photons on useful time scales, means that measuring the state of the ancilla during repeated error syndrome interrogation must be as efficient as possible to mitigate the resulting dephasing (see sec. 8.2). Additionally, this larger cross-Kerr requires a much larger coupling strength between the transmon qubit and the readout cavity ($\sim 35\ \text{MHz}$) than used, for example, in the experiments described in previous chapters. In turn this requires the use of a Purcell filter on the output of the readout cavity to suppress Purcell-enhanced decay of the transmon mode, an extra addition of hardware that can potentially adversely affect system performance. Likewise, using a stronger pump to enhance the conversion process can inadvertently turn on undesired hamiltonian terms that have been excluded in the approximations above. Such new sources of loss are unavoidable though, and are another example of how extra resource overhead must be introduced in order to perform error

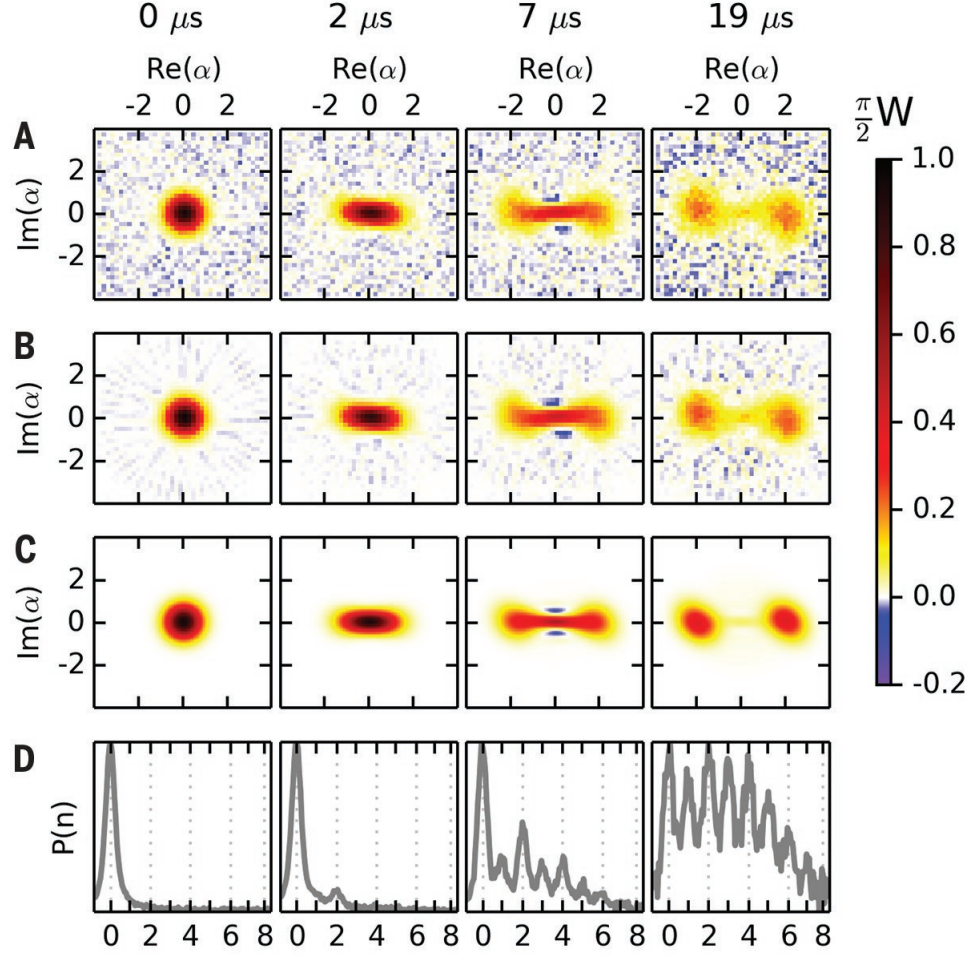


Figure 9.3: **Pumping a cavity from vacuum to a “2-cat.”** (a) Measured data. (b) Reconstructed density matrices. (c) Numerical simulations. The state is initialized in vacuum at time $t = 0 \mu\text{s}$. After $2 \mu\text{s}$ of pumping, the state becomes squeezed in the Q quadrature. Negative fringes are visible after $7 \mu\text{s}$, indicating a coherent superposition of the two steady states that approximate a “2-cat.” At $t = 19 \mu\text{s}$, the state in the storage cavity is stabilized at $|\pm\alpha_\infty\rangle$, but the fringes have disappeared due to unavoidable single photon loss. (d) Storage photon number distribution $P(n)$ as measured by spectroscopy of the transmon qubit after pumping. The nonclassical nature of the states in the storage cavity at $t = 2 \mu\text{s}$ and $t = 7 \mu\text{s}$ is confirmed by the non-Poissonian distribution of the photon numbers.

correction, much as in the case of increasing the redundancy of a quantum bit encoding in order to identify a measurable error syndrome. As with the experiments leading up to the result presented in chapter 7, the work described above is already being improved upon in redesigned setups, and an imminent demonstration of concurrent pumping and single-shot parity measurements promises to offer even greater enhancements to quantum bit coherence.




Figure 9.4: **Architecture for a logical encoding that uses two cavities.** (a) A cartoon image of the device, which shows two coaxial $\lambda/4$ cavities [Reagar et al., 2016], Alice and Bob, both coupled to a single Y-shaped transmon ancilla that itself is also coupled to a 2D stripline cavity used for projective readout. (b) A top view of the same device, which shows how each component of the system is positioned. (c) A photograph of the actual device with a quarter coin for size-scale reference.

9.2 Correcting logical gates

Beyond stabilizing a state in the cavity to a quantum manifold, which as presented can still be thought of as enhancing the performance of a quantum memory, we must also be able to implement single and multi-physical qubit gates that are universal. The remarkable levels of control we already have over the states in the cavity as evidenced by the action of the optimal control pulses (see sec. 6.3) in encoding and decoding a quantum bit from and back onto a transmon ancilla suggests that this could be a promising path forward. As with the parity measurements, however, the challenge is to combat the effects of ancilla decoherence that reduces the fidelity of the mappings on time scales roughly dictated by the product $\chi_{sa}T_2$, where χ_{sa} sets the speed limit of the pulse and ancilla T_2 sets the maximum fidelity that can be achieved after a certain interval of time. In

our current implementation, we are still working to understand how to make optimal control pulses fault-tolerant. Indeed, as seen Wigner snapshots of the cavity state in Fig. 9.5, which shows an example encoding pulse for $|-X_L\rangle$ (sec. 6.3), while the pulse is being played the evolution of the state is extremely non-trivial. If at any point the ancilla undergoes either a bit flip or a phase flip, the encoding, or any other operation for that matter, will fail.

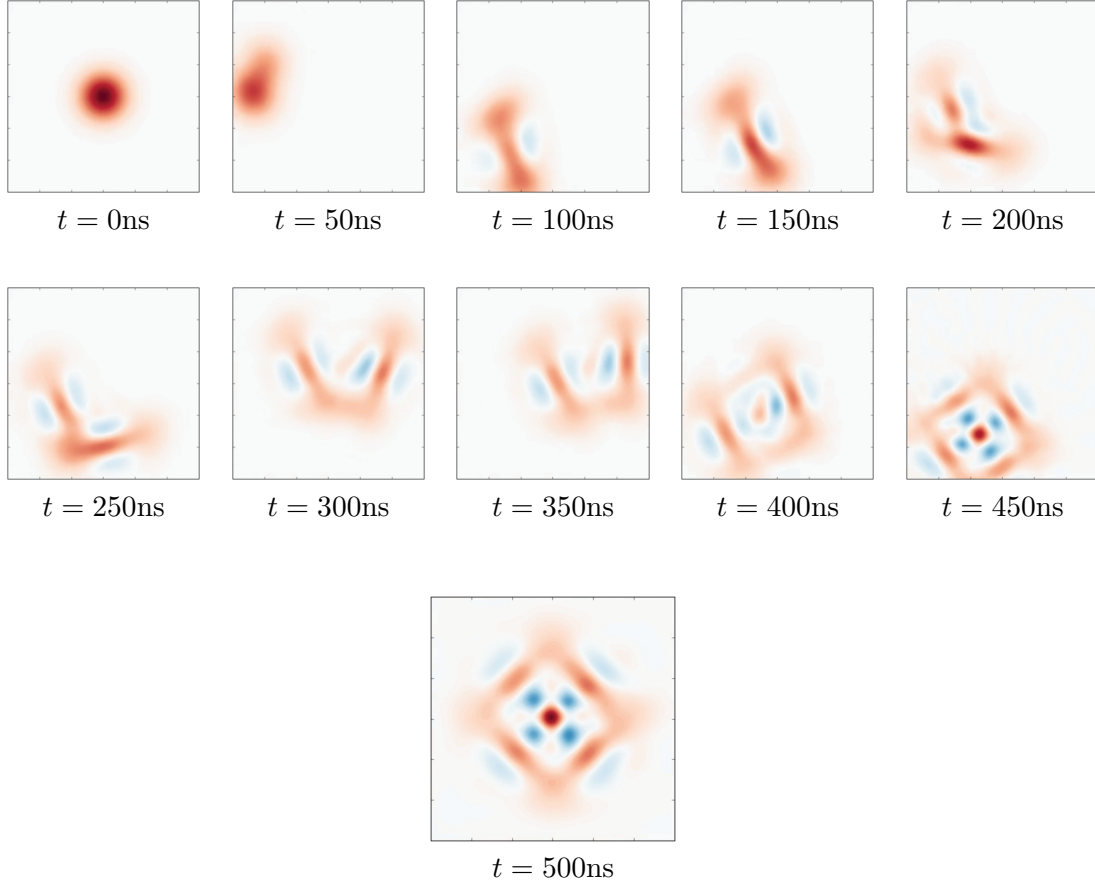


Figure 9.5: The challenge of decoherence during optimal control pulses. Snapshots in steps of 50 ns of cavity state Wigner functions during an optimal control pulse that maps the quantum bit $|\psi_q\rangle = 1/\sqrt{2}(|0\rangle - |1\rangle)$ from the ancilla onto $|-X_L\rangle$. Between the initial vacuum and final “4-cat” the cavity and ancilla are entangled and evolving in completely unintuitive ways. If the ancilla undergoes any form of decoherence at any point during this pulse, with near certainty the quantum bit will be completely unrecoverable; by the time the pulses finishes, the state in the cavity will have left the code space.

The proposal in [Mirrahimi et al., 2014] in fact also outlines in detail how supplementing the cat pumping with operations such as small displacements, turning on intentional Kerr evolution, and realizing beam-splitter Hamiltonians between neighboring cavities can implement the necessary set of gates for universal quantum computation (perhaps in a fault-tolerant fashion). Thus, although optimal control pulses as described seem to have run into the limits imposed by ancilla coherence,

our experience in understanding their strengths and drawbacks motivates exploring their potential applications in regimes where the transmon is used merely to supply the nonlinearity to the system without its mode participating in the actual manipulation of the logical state. Together with the pumping schemes, this can offer a promising path forward in performing gates that in principle are not intrinsically limited by the system's least coherent constituent.

9.3 Future directions: building on our rapid progress

Further necessary advances include achieving a high degree of device reproducibility, motivating cavities that can be integrated into the fabrication process of the physical qubits. Substantial progress towards this end has been made recently [Brecht et al., 2015, 2016; Mineev et al., 2016] and remains a central focus of future work, where key challenges include reaching the same internal quality factors in micro-fabricated cavities as in their 3D-machined predecessors. Finally, we also need to integrate a host of components that can efficiently route signals in a modular architecture and enable processes such as remote entanglement and distillation. Recent experiments have demonstrated remarkable progress towards these ends as well [Narla et al., 2016; Sliwa et al., 2015].

The sum total of this work over the past four years illustrates the at times painstaking effort required to understand, design, and optimize a quantum system to perform in ways conducive towards future applications. The resulting breadth of knowledge acquired, however, now puts down a flag in this journey, a checkpoint at which results make sense and to which we can always return to regroup. In fact, the system as it stands is quite simple: three quantum controller cards hooked up to a two-cavity setup with standard microwave drive and readout lines; the sequences that realize the full error correction system are written; they can now only be improved and built upon to run new and interesting experiments. These experiments can now exploit the results of previous efforts to grasp, and to face, the host of challenges yet to be overcome.

The first experiment that used a two-cavity, single transmon qubit 3D architecture was published in March of 2013 [Kirchmair et al., 2013]. By 2016, just three years later, the groups at Yale are using real-time feedback to implement quantum memories that surpass break-even, employing dissipative pumping schemes to stabilize a quantum manifold, entangling logical qubits across multiple cavities, performing arbitrary gates on arbitrary states in a cavity, remotely entangling physical qubits, rerouting signals on command, autonomously stabilizing Bell states with real-time feedback, as well as constantly refining their theoretical understanding of encodings and error correction schemes with logical states stored in a cavity. Such rates of innovation are truly inspiring, and herald an era of progress in quantum computation that invigorates the entire field, sparks new discussions and

collaborations, and challenges the mainstream perspective that multi-physical qubit architectures are the only conceivable way forward towards a practical, fault-tolerant device. In truth, from sophisticated fabrication techniques at MIT Lincoln Labs and impressive advances in coherence times in planar structures at IBM, to alternative paradigms based on quantum annealing with D-Wave and Google that raise important questions of quantum versus classical computation capabilities, with any approach comes a wide array of insights, successes, and failures that we can always learn from. By doing so we stand to maintain our position at the forefront of the field. Exciting challenges abound. And they are our future directions.

A – Experimental Setup Details

Experiments are performed in a cryogen-free dilution refrigerator at a base temperature of ~ 10 mK. Our output signal amplification chain consists of two stages. The first stage is comprised of a Josephson parametric amplifier. For the experiments in which we create a Bell-cat and quantify how QND single-shot parity measurements are (chapters 4 and 5, respectively), we used a Josephson bifurcation amplifier (JBA) [Vijay et al., 2009] operating in a double-pumping configuration [Kamal et al., 2009; Murch et al., 2013], which provides phase-sensitive detection. For the experiments with the full QEC system (chapters 6 through 9), we used a phase-preserving JPC. In all setups, the parametric amplifiers are followed by a high electron mobility transistor (HEMT) amplifier bolted to the 4K plate of the fridge.

The transmon ancilla is fabricated on a c -plane sapphire (Al_2O_3) substrate with a double-angle evaporation of aluminum after a single electron-beam lithography step. The two-cavity, single transmon system is well described by the approximate dispersive Hamiltonian (Eq. A.1):

$$\begin{aligned} \hat{H}/\hbar \approx & \tilde{\omega}_a \hat{b}^\dagger \hat{b} + \tilde{\omega}_s \hat{a}_s^\dagger \hat{a}_s + \tilde{\omega}_r \hat{a}_r^\dagger \hat{a}_r \\ & - \chi_{sa} \hat{b}^\dagger \hat{b} \hat{a}_s^\dagger \hat{a}_s - \chi_{ra} \hat{b}^\dagger \hat{b} \hat{a}_r^\dagger \hat{a}_r - \chi_{sr} \hat{a}_s^\dagger \hat{a}_s \hat{a}_r^\dagger \hat{a}_r \\ & - \frac{K_a}{2} \hat{b}^{\dagger 2} \hat{b}^2 - \frac{K_s}{2} \hat{a}_s^{\dagger 2} \hat{a}_s^2 - \frac{K_r}{2} \hat{a}_r^{\dagger 2} \hat{a}_r^2, \end{aligned} \quad (\text{A.1})$$

where $\tilde{\omega}_{s,r,a}$ are the storage, readout, and ancilla transmon transition frequencies, $\hat{a}_s, \hat{a}_r, \hat{b}$ are the associated ladder operators; and $K_{s,r,a}, \chi_{sa,ra,sr}$ are the modal anharmonicities and dispersive shifts respectively. The resonant frequency of the readout cavity is determined by transmission spectroscopy. The transmon and storage cavity frequencies are found using two-tone spectroscopy. The transmon anharmonicity K_a is measured using two-tone spectroscopy to observe the $0 - 2$ two-photon transition [Paik et al., 2011]. The dispersive shift between the transmon and the readout cavity χ_{ra} is found by taking the difference in frequency between the readout resonance when the transmon is in the ground and the excited states. The dispersive shift between the transmon and the storage cavity χ_{sa} is found using two methods: photon number dependent ancilla spectroscopy

[Schuster et al., 2007], and observing transmon state revival using Ramsey interferometry [Vlastakis et al., 2013]. The storage-readout cross-Kerr χ_{sr} is predicted using its approximate relationship between K_s and K_r [Nigg et al., 2012]. Finally, we predict the readout cavity anharmonicity K_r using its approximate dependence on the measured values of K_a and the transmon-readout dispersive shift χ_{ra} [Nigg et al., 2012]. When performing Wigner tomography, in order to remove the cross-Kerr effect between the readout cavity and the storage cavity, which skews the readout signal for large storage cavity displacements, we have followed the identical procedure as in the Supplementary Material in Ref. [Vlastakis et al., 2013]. The idea is to perform tomography using two pulse sequences, which both measure parity but differ in the sign of the second $\pi/2$ pulse (Fig. 5.2). We then take the difference between the two measurements to remove any bias in the measurement.

The lifetime of the readout cavity is found from its line-width. The thermal population of the transmon is determined from a histogram of one million single-shot measurements of the transmon thermal state, where the signal-to-noise ratio provided by the JBA (or JPC) allows discrimination between $|g\rangle$ and all states not $|g\rangle$. The thermal population of the storage cavity is found by taking the difference between parity measurements of the thermal and vacuum states of the cavity. A vacuum state is prepared by first performing two parity measurements on the thermal state and then post-selecting such that all results give even parity, projecting the thermal state onto $|0\rangle_f$. Finally, the known thermal population of the readout cavity is bounded by the dephasing rate Γ_ϕ of the ancilla: $\Gamma_\phi = \bar{n}_{th} \kappa_r$, where \bar{n}_{th} is the readout cavity's thermal occupation and κ_r is the readout single-photon decay rate [Sears et al., 2012].

A.1 Bell-cat setup

The refrigerator wiring (Fig. A.1), including the filters and attenuators used, is similar to that of [Sun et al., 2013], but with the addition of a feedback system used for ancilla reset. The Hamiltonian parameters are summarized in Tab. A.1. The storage cavity anharmonicity K_s is determined by displacing the cavity with a coherent state and observing its time evolution with Wigner tomography. The resulting dynamics are characterized by state reconstruction and K_s is observed by the state's quadratic dependence of phase on photon number. The coherence times and thermal occupations of the system's three modes are given in Tab. A.2.

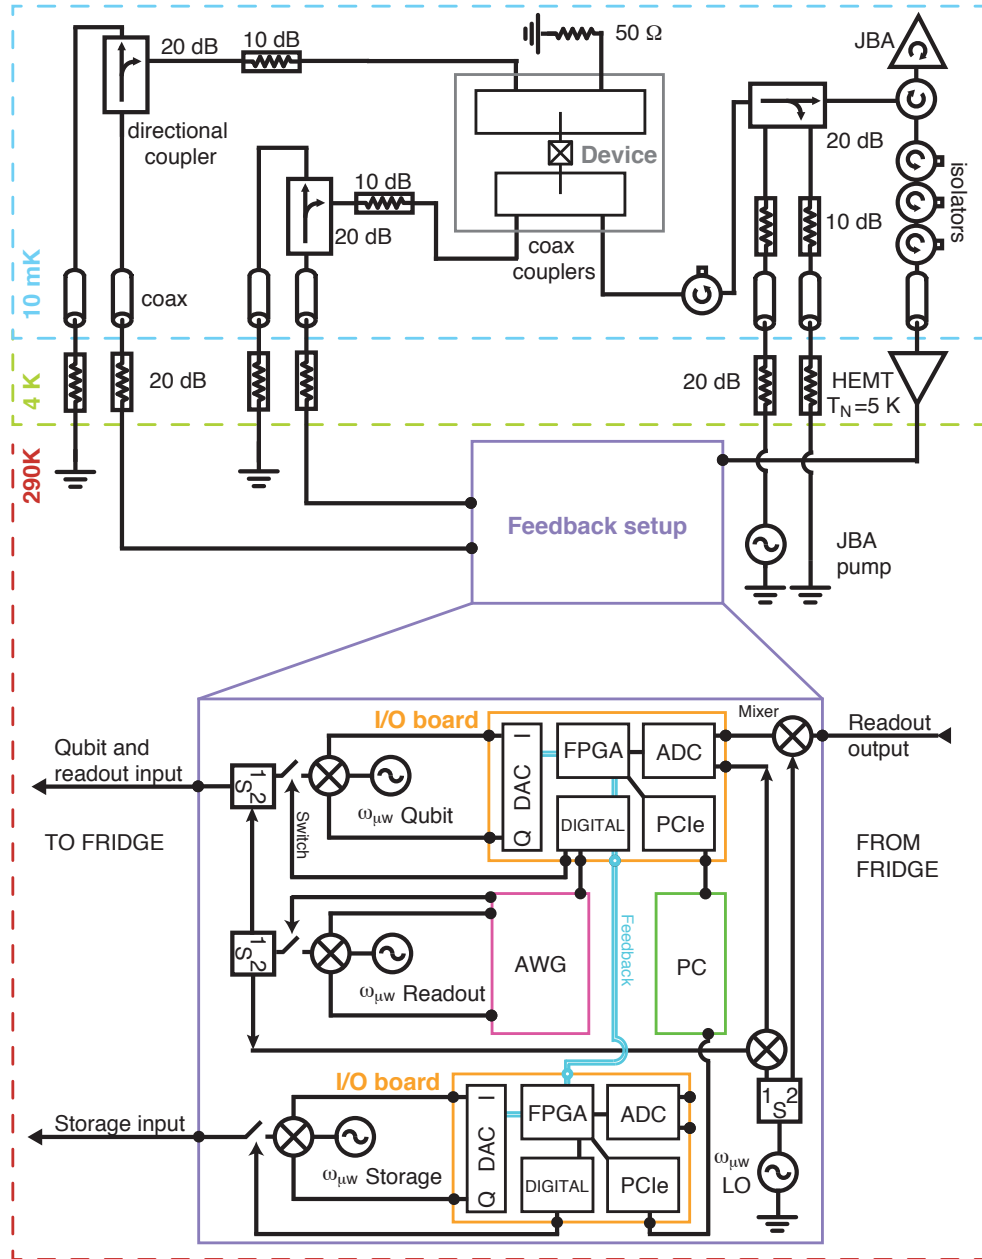


Figure A.1: **Fridge wiring diagram for Bell-cat experiments.** The schematic of the experiment shows a two-cavity one-transmon device identical to that used in [Sun et al., 2013] with the addition of a feedback setup for active ancilla reset. The feedback setup uses two input-output (I/O) boards for ancilla and storage cavity control and one arbitrary waveform generator (AWG) for readout cavity control. All have a dedicated microwave generator and mixer for amplitude and phase modulation. Each I/O board has five main components: 1) a digital-to-analog converter (DAC) for pulse generation; 2) digital outputs serving as marker channels; 3) an analog-to-digital converter (ADC) that samples input signals; 4) an FPGA that demodulates the signals from the ADC and based on predefined thresholds determines the measured ancilla state, $|g\rangle$ or $|e\rangle$; and 5) a PCIe connection that transfers FPGA data to a computer (PC) for analysis. In this setup, the top I/O board serves as the master, which accepts the readout signal, returns ancilla state information, and using digital output signals, triggers the AWG and the second I/O card given a particular ancilla measurement result.

| Term | Measured (Prediction) |
|-------------------------|--------------------------|
| $\tilde{\omega}_a/2\pi$ | 5.7651 GHz |
| $\tilde{\omega}_s/2\pi$ | 7.2164 GHz |
| $\tilde{\omega}_r/2\pi$ | 8.1740 GHz |
| $K_a/2\pi$ | 240 MHz |
| $K_s/2\pi$ | 1.5 kHz |
| $K_r/2\pi$ | (2 kHz) |
| $\chi_{sa}/2\pi$ | 1.43 MHz |
| $\chi_{ra}/2\pi$ | 1 MHz |
| $\chi_{sr}/2\pi$ | (1.7 kHz) |

Table A.1: **Hamiltonian parameters for Bell-cat experiments.**

| | Ancilla | Storage | Readout |
|------------------|------------|------------|---------|
| T_1 | 10 μ s | - | - |
| T_2 | 10 μ s | - | - |
| τ_s | - | 55 μ s | 30 ns |
| ground state (%) | 90% | > 98% | > 99.8% |

Table A.2: **Coherence and thermal properties for Bell-cat experiments.**

Measurement fidelities

We define single-shot measurement fidelity as $F_a = \frac{P(g|g)+P(e|e)}{2}$, where $P(g|g)$ and $P(e|e)$ are the probabilities to get $|g\rangle$ ($|e\rangle$) knowing that we start with $|g\rangle$ ($|e\rangle$). The state $|g\rangle$ is prepared through purification of the ancilla thermal state with realtime feedback (see the following section). Given a preparation of $|g\rangle$, we have a 98.5% chance of measuring $|g\rangle$ again ($P(g|g) = 0.985$). Likewise, we find $P(e|e) = 0.975$ by preparing $|g\rangle$ and rotating the state to $|e\rangle$. This gives a single-shot measurement fidelity of $F_q = 98\%$. We find our cavity parity measurement fidelity by purifying the storage cavity thermal state into $|0\rangle_f$ then performing one of two kinds of parity measurement. We report a parity measurement fidelity for $n = 0$ photons as $F_c = \frac{P(g|E_1)+P(e|E_2)}{2} = 95.5\%$, where $P(g|E_1)$ ($P(e|E_2)$) is the probability to measure $|g\rangle$ ($|e\rangle$) given that the parity is even for each of the two measurement settings. We expect F_c to decrease with increasing numbers of photons in the cavity due to single photon loss during the measurement sequence.

I/O control parameters

We employ a field-programmable gate array (FPGA) in order to implement an active feedback scheme. We use an X6-1000M board from Innovative Integration that contains two 1 GS/s ADCs, two 1 GS/s DAC channels, and digital inputs/outputs all controlled by a Xilinx VIRTEX-6 FPGA

loaded with custom logic. We synchronize two such boards in a master/slave configuration to have IQ control of both the ancilla/storage cavity. IQ control over the readout cavity is performed with a Tektronix AWG, which is triggered by the master board. The readout and reference signals are routed to the ADCs on the master board, whereafter the FPGA demodulates the signal and decides whether the ancilla is in $|g\rangle$ or $|e\rangle$. The feedback latency of the FPGA logic (last in, first out LIFO) is 320 ns. Additional delays for active feedback include cable delay (~ 100 ns) and readout pulse length with cavity decay time (320 ns). Thus, in total the ancilla waits $\tau_{\text{wait}} \sim 740$ ns between the time photons first enter the readout cavity and the time at which the feedback pulse resets the ancilla.

Implementations of feedback

Feedback is used three times during a single iteration of the experiment. Prior to the Bell-cat preparation, we purify the ancilla state to $|g\rangle$ by measuring the ancilla and applying a π pulse if the result is $|e\rangle$. We succeed in preparing $|g\rangle$ with a probability of 99%. Secondly, when performing ancilla tomography we reset the ancilla to $|g\rangle$ if it is measured to be in $|e\rangle$. Since we must wait τ_{wait} before feedback can be applied, the cavity state will acquire an additional phase $\chi_{sa}\tau_{\text{wait}}$ if the ancilla is in $|e\rangle$. In this case, in addition to resetting the ancilla, the FPGA applies an equivalent phase shift on the subsequent Wigner tomography pulse. This feedback implementation does not close the ‘locality’ loophole for a CHSH Bell test and therefore cannot be used to test local realism.

A.2 Parity monitoring setup

The refrigerator wiring is similar to that shown in Fig. A.1, only without the feedback setup. The Hamiltonian parameters are summarized in Tab. A.3. The coherence times and thermal occupations of the system’s three modes are given in Tab. A.4.

Figure A.2 shows the probability of the first eight Fock states $n = 0, 1, 2, \dots, 7$ as a function of displacement amplitude $|\alpha|$. The results are in excellent agreement with a Poisson distribution, indicating a good control of the coherent state in the cavity. We scale the x-axis from the voltage amplitude of the displacement pulse applied from an arbitrary waveform generator and use this scaling as a calibration. There is a small residual amplitude for the $n = 1$ peak even with no displacement (point near origin), allowing us to infer that there is a background photon population $n_{th}^s = 0.02$ in the cavity. The inset of Fig. A.2 shows the so-called number splitting peaks of the ancilla due to different photon numbers in the cavity, which is displaced with a 10 ns square pulse

| Term | Measured (Prediction) |
|-------------------------|--------------------------|
| $\tilde{\omega}_a/2\pi$ | 5.938 GHz |
| $\tilde{\omega}_s/2\pi$ | 7.2164 GHz |
| $\tilde{\omega}_r/2\pi$ | 8.1740 GHz |
| $K_a/2\pi$ | 240 MHz |
| $K_s/2\pi$ | 1.5 kHz |
| $K_r/2\pi$ | (2 kHz) |
| $\chi_{sa}/2\pi$ | 1.789 MHz |
| $\chi_{ra}/2\pi$ | 0.93 MHz |
| $\chi_{sr}/2\pi$ | (1.7 kHz) |

Table A.3: **Hamiltonian parameters for parity monitoring experiments.**

| | Ancilla | Storage | Readout |
|------------------|----------|-----------|---------|
| T_1 | $8\mu s$ | - | - |
| T_2 | $5\mu s$ | - | - |
| τ_s | - | $55\mu s$ | 30 ns |
| ground state (%) | 86% | 98% | > 99.5% |

Table A.4: **Coherence and thermal properties for parity monitoring experiments.**

right before the spectroscopy measurement. A second order polynomial fit $\chi(n) = -\chi_{sa}n + \chi'_{sa}n^2$, where n is the peak number, gives a non-linear correction to the dispersive shift [Vlastakis et al., 2013] $\chi'_{sa}/2\pi = 1.9 \pm 0.1$ kHz, which is small enough to be neglected in the cavity dynamics.

To perform a good parity measurement, the $\pi/2$ pulses in the Ramsey parity mapping should equally cover as many number splitting peaks as possible without significantly exciting the $|f\rangle$ state. We choose a Gaussian envelope pulse truncated to $4\sigma = 8$ ns ($\sigma_f = 80$ MHz) for a good compromise. Figure A.3 shows the effectiveness of the $\pi/2$ pulses as a function of \bar{n} in the cavity. The curvature for $\bar{n} > 4$ is due to the finite bandwidth of those pulses in the frequency domain.

The quantum filter

In estimating the parity in time, in order to mitigate the effects due to ancilla decoherence, $|f\rangle$ state of the ancilla (an undesirable state that obscures the parity measurement), and other imperfections in the ancilla readout in extracting the parity, we have applied a quantum filter [Belavkin, 1992; Bouten et al., 2007]. The quantum filter can loosely be thought of as a more rigorous application of majority voting on the measurement results, wherein it integrates the quantum stochastic master equation and returns a result that depends on the entire measurement record. At each point in time, the quantum filter is updated in two steps: first, a new density matrix $\tilde{\rho}_s(C_{t+dt})$ is calculated from the




Figure A.2: **Poisson distribution of photon numbers in the cavity.** Dotted color lines are data for the first eight Fock states $n = 0, 1, 2, \dots, 7$ as a function of displacement amplitude $|\alpha|$. The measurements are performed with a selective π pulse on each number splitting peak and the resulting signal amplitude should be proportional to the corresponding number population. These oscillation amplitudes have been normalized to probabilities such that the sum of the amplitudes corresponding to $n = 0$ and $n = 1$ equals unity. Dashed lines are theoretical curves with a Poisson distribution $\mathbb{P}(|\alpha|, n) = |\alpha|^{2n} e^{-|\alpha|^2} / n!$ where the x-axis has had a single scale factor adjusted to fit all these probabilities. The excellent agreement indicates good control over the coherent state in the cavity and also gives a good calibration of the cavity displacement amplitude. Based on the probability of $n = 1$ at $|\alpha| = 0$, we find a background photon population $n_{th}^s = 0.02$ in the cavity. Inset: spectroscopy (left axis) of the number splitting peaks of the ancilla when populating different photon numbers in the cavity. Top panel shows the difference between peak positions and a linear fit. The curvature necessitates a second order polynomial fit resulting a linear dispersive shift $\chi_{sa}/2\pi = 1.789 \pm 0.002$ MHz and a non-linear dispersive shift $\chi'_{sa}/2\pi = 1.9 \pm 0.1$ kHz.

best estimate $\rho_s(C_t)$ at the previous point based only on the decoherence of the cavity; second, the density matrix $\tilde{\rho}_s(C_{t+dt})$ gets updated as a best estimate $\rho_s(C_{t+dt})$ according to Bayes' rule based on the newly acquired knowledge from the current parity measurement. This density matrix $\rho_s(C_{t+dt})$ is then used as the input for the next iteration. We have truncated the dimension of the density matrix to $N = 5\bar{n} \approx 20$, which is large enough to cover all relevant number states. To initialize the density matrix after a displacement $D(\alpha)$, we have set $\rho_s(t = 0) = (1 - n_{th}^s)D(\alpha) |0\rangle_f \langle 0|_f D^\dagger(\alpha) + n_{th}^s D(\alpha) |1\rangle_f \langle 1|_f D^\dagger(\alpha)$, taking into account the background photon population in the limit $n_{th}^s \ll 1$.

At time t , the density matrix of the photon state is $\rho_s(C_t)$, which depends on all previous correlations. At $t + dt$, only considering the decoherence of the cavity, the expected density matrix from free evolution becomes $\tilde{\rho}_s(C_{t+dt}) = M_{down}\rho_s(C_t)M_{down}^\dagger + M_{up}\rho_s(C_t)M_{up}^\dagger + M_{no}\rho_s(C_t)M_{no}^\dagger$, where $M_{down} = \sqrt{\kappa_{down}dt}a$, $M_{up} = \sqrt{\kappa_{up}dt}a^\dagger$, and $M_{no} = I - (M_{down}^\dagger M_{down} + M_{up}^\dagger M_{up})/2$ are the Kraus operators for photon loss, absorption of thermal photons, and no jump events respectively.




Figure A.3: **Effectiveness of a $\pi/2$ pulse.** Blue and red data (bottom axis) is ensemble averaged ancilla readout after consecutively (with no wait time) applying $(\pi/2, \pi/2)$ and $(\pi/2, -\pi/2)$ pulses, respectively as a function of different \bar{n} introduced into the cavity. The curvature for $\bar{n} > 4$ comes from the finite bandwidth of the pulses in the frequency domain. The green curve (top axis) is a time Rabi trace for an amplitude comparison with no initial cavity displacement.

We have $\kappa_{down} = (n_{th}^s + 1)\kappa$ and $\kappa_{up} = n_{th}^s\kappa$, and $\kappa = 1/\tau_{tot}$ is the energy decay rate in the cavity under repeated parity measurements. The additional information C_{t+dt} acquired from the parity measurement at $t + dt$ changes the quantum state according to:

$$\rho_s(C_{t+dt}) = \begin{cases} \mathbb{P}(even|C_{t+dt}) \frac{\hat{P}_{even}\tilde{\rho}_s(C_{t+dt})\hat{P}_{even}}{\text{Tr}(\hat{P}_{even}\tilde{\rho}_s(C_{t+dt})\hat{P}_{even})} \\ \quad + \mathbb{P}(odd|C_{t+dt}) \frac{\hat{P}_{odd}\tilde{\rho}_s(C_{t+dt})\hat{P}_{odd}}{\text{Tr}(\hat{P}_{odd}\tilde{\rho}_s(C_{t+dt})\hat{P}_{odd})}, & \text{if } C_{t+dt} \neq 0, \\ \tilde{\rho}_s(C_{t+dt}) & \text{if } C_{t+dt} = 0, \end{cases} \quad (\text{A.2})$$

where \hat{P}_{even} and \hat{P}_{odd} are the projectors onto the even and odd manifolds, $\hat{P} = \hat{P}_{even} - \hat{P}_{odd} = e^{i\pi\hat{a}_s^\dagger\hat{a}_s}$ is the parity operator, $\mathbb{P}(even|C_{t+dt})$ and $\mathbb{P}(odd|C_{t+dt})$ are the probabilities of being in the even and odd parity respectively for a measured C_{t+dt} . To simplify the quantum filter, we assume that the event of the ancilla jumping to the $|f\rangle$ states is independent of the cavity parity being even or odd. Hence, if the measured correlation is zero, the density matrix of the photon state is assigned to the

expected one from free evolution. Based on Bayes' rule, Eq. A.2 becomes:

$$\rho_s(C_{t+dt}) = \begin{cases} \frac{\mathbb{P}(C_{t+dt}|even)\hat{P}_{even}\tilde{\rho}_s(C_{t+dt})\hat{P}_{even} + \mathbb{P}(C_{t+dt}|odd)\hat{P}_{odd}\tilde{\rho}_s(C_{t+dt})\hat{P}_{odd}}{\mathbb{P}(C_{t+dt})}, & \text{if } C_{t+dt} \neq 0, \\ \tilde{\rho}_s(C_{t+dt}) & \text{if } C_{t+dt} = 0. \end{cases} \quad (\text{A.3})$$

where $\mathbb{P}(C_{t+dt}) = \mathbb{P}(C_{t+dt}|even)\text{Tr}[\hat{P}_{even}\tilde{\rho}_s(C_{t+dt})\hat{P}_{even}] + \mathbb{P}(C_{t+dt}|odd)\text{Tr}[\hat{P}_{odd}\tilde{\rho}_s(C_{t+dt})\hat{P}_{odd}]$. The best parity estimation of the resonator state is then:

$$P(t+dt) = \text{Tr}[\rho_s(C_{t+dt})\hat{P}] \quad (\text{A.4})$$

A.3 QEC methods

Our setup is identical to that described in [Vlastakis et al., 2015] and [Sun et al., 2013], aside from a 2 mm rather than 4 mm wall separating the storage and readout cavities, a Josephson Parametric Converter (JPC) [Bergeral et al., 2010] replacing a Josephson Bifurcation Amplifier (JBA) [Vijay et al., 2009] as the first stage of amplification, and a quantum control architecture replacing the Tektronix AWG configuration (Figs. A.4, A.5). The Hamiltonian parameters are summarized in Tab. A.5. The coherence times and thermal occupations of the system's three modes are given in Tab. A.6.

The Kerr interaction of the storage cavity K_s is measured by monitoring the errors of a cat code logical state (see sec. 6.4) and finding the difference in phase between trajectories where errors are measured to occur at different times: $\Delta\theta = K_s\Delta t_j$. In Fig. 6.10, we show two Wigner functions for the case of a single parity jump: 01 and 10, where $0 \equiv$ "no error" and $1 \equiv$ "error." We Fourier transform circular cuts at a fixed radius of these Wigner functions that show pronounced interference fringes to compare the phase of the oscillations for 01 vs. 10 and in so doing find $\Delta\theta$. On average, photon jumps for 01 versus 10 are separated in time by t_M , where t_M is the total time between syndrome measurements; we thus find the average difference between jump time $\overline{\Delta t_j} = t_M$ to find K_s .

A lower bound on the dephasing rate of the storage cavity, Γ_ϕ^s , is given by $\Gamma_\phi^s = \Gamma_\uparrow$ [Reagor et al., 2016], akin to the dephasing one expects of an ancilla (e.g. transmon) coupled to a low-Q readout cavity with a finite thermal population [Sears et al., 2012]. The storage cavity coherence time T_2^s is thus given by $(T_2^s)^{-1} = (2\tau_s)^{-1} + \Gamma_\phi^s$, where τ_s is the average lifetime of the single photon Fock state $|1\rangle_f$. The coherence time T_2^s is consistent with the observed time constant in the decay of the process fidelity of a ancilla stored in Fock states $|0\rangle_f, |1\rangle_f$.

The single-shot ancilla measurement fidelity is 99.3%. The parity measurement fidelity is 98.5% for no photons in the storage cavity, 98.1% for an average photon number $\bar{n} = 2$, and 97.7% for $\bar{n} = 3$; these fidelities are also obtained using the methods in sec. A.1.

| Term | Measured (Prediction) |
|-------------------------|--------------------------|
| $\tilde{\omega}_a/2\pi$ | 6.2815 GHz |
| $\tilde{\omega}_s/2\pi$ | 8.3056 GHz |
| $\tilde{\omega}_r/2\pi$ | 9.3149 GHz |
| $K_a/2\pi$ | 297 MHz |
| $K_s/2\pi$ | 4.5 kHz |
| $K_r/2\pi$ | (0.5 kHz) |
| $\chi_{sa}/2\pi$ | 1.97 MHz |
| $\chi_{ra}/2\pi$ | 1 MHz |
| $\chi_{sr}/2\pi$ | (2 kHz) |
| $\chi'_{sa}/2\pi$ | 1 kHz |

Table A.5: **Hamiltonian parameters for QEC experiments.**

| | Ancilla | Storage | Readout |
|------------------|-----------|------------|---------|
| T_1 | $35\mu s$ | - | - |
| T_2 | $12\mu s$ | - | - |
| τ_s | - | $250\mu s$ | 100 ns |
| T_2^s | - | $330\mu s$ | - |
| ground state (%) | 96% | > 98% | > 99.3% |

Table A.6: **Coherence and thermal properties for QEC experiments.**

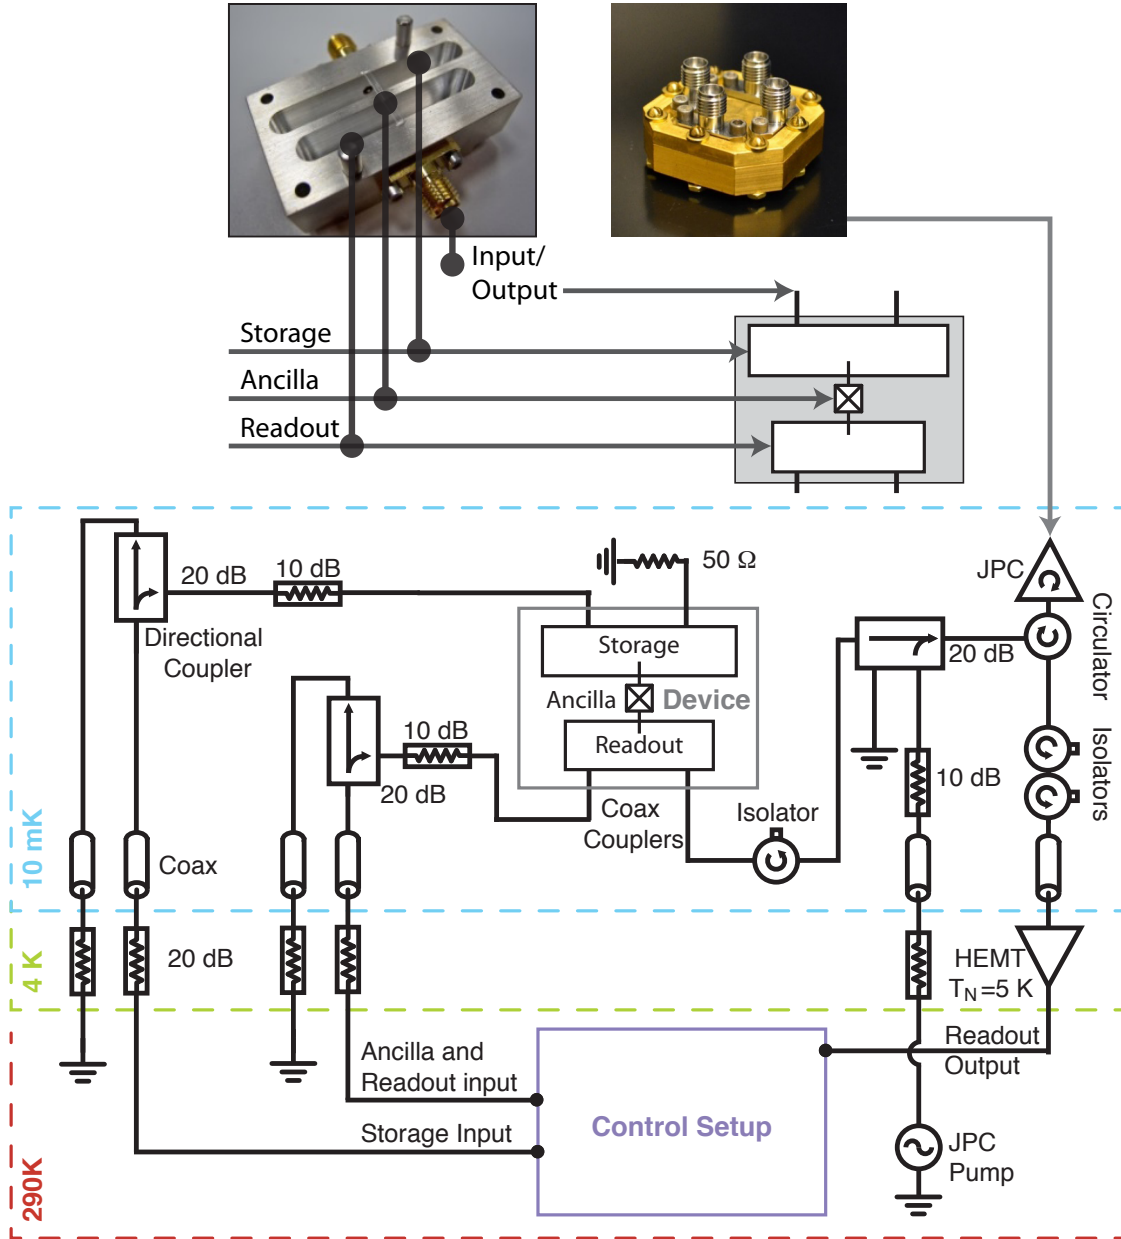


Figure A.4: **Schematic of QEC measurement setup.** The schematic of the experiment shows a two-cavity one-ancilla device similar to that used in [Sun et al., 2013]. We use a JPC amplifier and a control setup that allows us to implement a full QEC system.

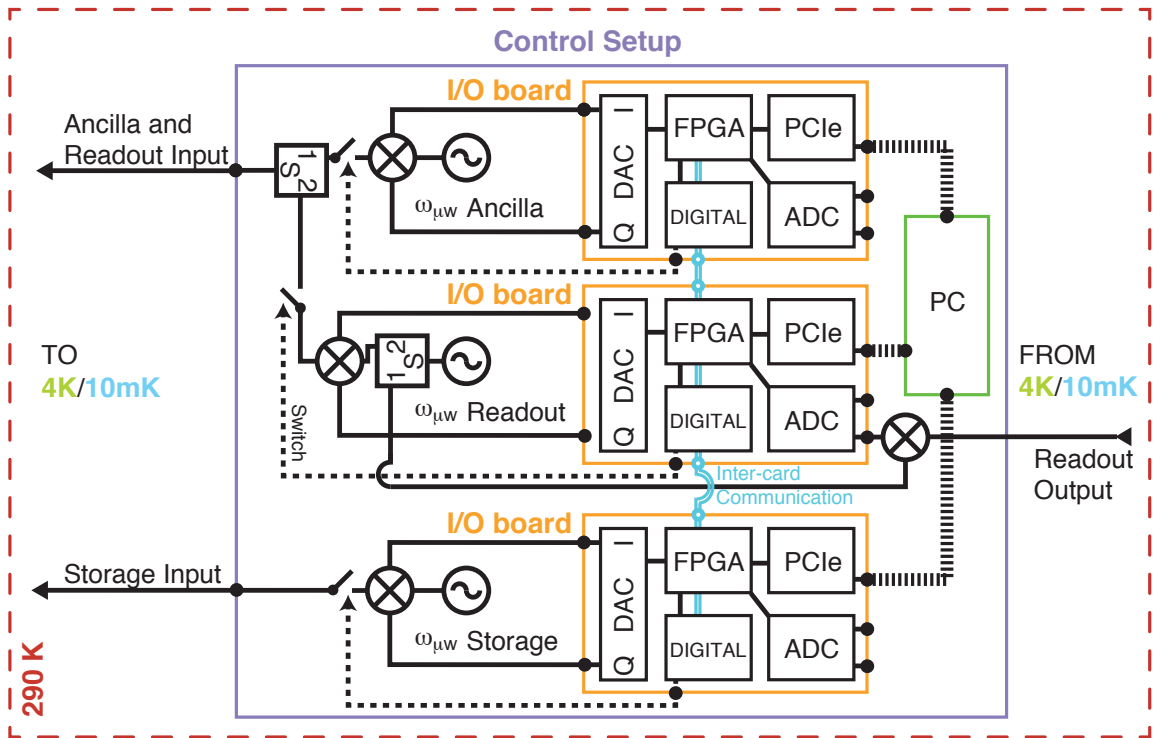


Figure A.5: **Schematic of QEC control setup.** A schematic of the control setup used in the full QEC system. Three independent I/O cards control the ancilla, readout cavity, and storage cavity with single-sideband modulated and shaped waveforms; the outputs and feedback steps are controlled by the on-board FPGAs. Inter-card communication channels synchronize the execution of all pulses.

B – Deriving the Optimized Configuration

Here we derive the optimal fidelity F assuming $F_{KD} = 1$. We follow a pessimistic approach. Any error that may happen will be regarded as a total loss of information. An error in the parity measurement could, in principle, be corrected by repeating it several times and taking a majority vote. In practice, although our measurement fidelities are very high, the errors that these extra measurements introduce are larger than the those we wish to correct for, so it is better to blindly trust any result. This approach simplifies the analysis greatly, and we can calculate the probability to successfully keep the information for each step independently and take the product to get the final total fidelity.

A successful step starts with the ancilla in $|g\rangle$, followed by either no jumps or a single photon jump during the delay, an accurate parity measurement, and the ancilla back in $|g\rangle$ at the end. There are several failure mechanisms:

1. While waiting, the ancilla may have been excited to $|e\rangle$
2. Two or more photon jumps occurred during the step delay
3. The parity measurement returned the wrong answer
4. A successful parity measurement brought the ancilla to $|e\rangle$, but the following reset pulse failed to return it back to $|g\rangle$

The probabilities to have zero, one, or more jumps are a function of the cat size at the beginning of the step and the step length. We express the step success probability as:

$$F_k = e^{-t_k \Gamma_{\uparrow}} [P_k(0) \cdot f_0 + P_k(1) \cdot f_1], \quad (\text{B.1})$$

where $P_k(n)$ is the probability to have n photon jumps in the k^{th} step, f_0 (f_1) is the conditional success probability when no (a single) photon jump occurred. The final success probability is then:

$$F/F_{ED} = \prod_{k=1}^S e^{-t_k \Gamma_{\uparrow}} [P_k(0) \cdot f_0 + P_k(1) \cdot f_1] = \underbrace{e^{-T \Gamma_{\uparrow}}}_{F_{\Gamma_{\uparrow}}} \cdot \underbrace{\prod_{k=1}^S [P_k(0) \cdot f_0 + P_k(1) \cdot f_1]}_{F_T}. \quad (\text{B.2})$$

From this point onward we focus only on F_T .

The success of the parity measurement depends primarily on the ancilla's T_2 . On top of that, there is the readout fidelity (which is different for the ground and the excited states). When a single photon jump occurs, the ancilla ends up in $|e\rangle$. It may decay back to $|g\rangle$ before the reset pulse, which means that the reset pulse inadvertently returns it back to $|e\rangle$. This is a critical period of time when we are vulnerable to T_1 decay of the ancilla:

$$f_0 \approx e^{-\frac{\pi}{\chi_{sa} T_2}} \cdot M_{gg} \quad (\text{B.3})$$

$$f_1 \approx e^{-\frac{\pi}{\chi_{sa} T_2}} \cdot M_{ee} \cdot e^{-\frac{\tau_{\text{meas}} + T_{FB}}{T_1}}, \quad (\text{B.4})$$

where M_{gg} (M_{ee}) is the probability to measure correctly $|g\rangle$ ($|e\rangle$), $\tau_{\text{meas}} = 400$ ns is the readout pulse length and $T_{FB} = 332$ ns is the feedback latency that includes delays due to the experimental setup (cables, etc.). Ancilla T_1 decay causes code failure no matter when it happens, which is why we take into account the whole duration until the ancilla is back in $|g\rangle$.

At the beginning of the k^{th} step, the averaged photon number in the cavity is given by (\bar{n}_0 is the cat size at the beginning of the $k = 1$ step):

$$\bar{n}_{k-1} = \bar{n}_0 \cdot \exp\left(-\kappa_s \sum_{i=1}^{k-1} t_i\right). \quad (\text{B.5})$$

During the step delay, the number of photon jumps follows a Poisson distribution with the following mean value (the decay during the step itself is taken into account):

$$\lambda_k = \bar{n}_{k-1} \cdot [1 - e^{-\kappa_s t_i}]. \quad (\text{B.6})$$

We expect no photon jumps with probability $e^{-\lambda_k}$ and a single jump with probability $\lambda_k e^{-\lambda_k}$. Hence, the step fidelity is simply given by:

$$F_T^{(k)} = f_0 e^{-\lambda_k} + f_1 \lambda_k e^{-\lambda_k} = (f_0 + f_1 \lambda_k) e^{-\lambda_k} \quad (\text{B.7})$$

↓

$$F_T(T, \bar{n}_0, S, \{t_k\}_{t=1}^S) = \prod_{k=1}^S [(f_0 + f_1 \lambda_k) e^{-\lambda_k}] \quad (\text{B.8})$$

We can now optimize the step length for a fixed number of steps. For $S = 1$, a single step, the solution is forced to be $t_1 = T$. For two steps we will need to optimize the following expression:

$$\max_{t_1} F_T(T, \bar{n}_0, S, \{t_1, t_2 = T - t_1\}) = [(f_0 + f_1 \lambda_1) e^{-\lambda_1}] \cdot [(f_0 + f_1 \lambda_2) e^{-\lambda_2}] \quad (\text{B.9})$$

$$= (f_0 + f_1 \lambda_1) \cdot (f_0 + f_1 \lambda_2) e^{-(\lambda_1 + \lambda_2)} \quad (\text{B.10})$$

$$= (f_0 + f_1 \lambda_1) \cdot (f_0 + f_1 \lambda_2) e^{-(1 - e^{-\kappa_s T})} \quad (\text{B.11})$$

The expected number of photon jumps during the two steps sums up to the expected number of jumps during the whole duration. This is independent of how we partition the whole duration into two steps. We can simply maximize the multiplication of $(f_0 + f_1 \lambda_1)$ and $(f_0 + f_1 \lambda_2)$. Since the sum of these terms is constant, the maximum is achieved when they are equal, meaning $\lambda_1 = \lambda_2$. In other words, we need to maintain a constant rate of photon jumps between the steps. This will be true for any number of steps; for S steps the mean number of photon jumps per step is given by:

$$\lambda_k = \frac{\bar{n}_0 [1 - e^{-\kappa_s T}]}{S}. \quad (\text{B.12})$$

Substituting this expression into eq. B.8 and using the optimal step lengths, we obtain:

$$F_T(T, \bar{n}_0, S) = \left[\left(f_0 + f_1 \frac{\bar{n}_0 [1 - e^{-\kappa_s T}]}{S} \right) \cdot e^{-\bar{n}_0 [1 - e^{-\kappa_s T}] / S} \right]^S \quad (\text{B.13})$$

$$= \underbrace{\left(f_0 + f_1 \frac{\bar{n}_0 [1 - e^{-\kappa_s T}]}{S} \right)^S}_{\equiv F'_T} \cdot e^{-\bar{n}_0 [1 - e^{-\kappa_s T}]} \quad (\text{B.14})$$

As an exercise, we can see how this expression behaves for S much greater than the number of expected photon jumps during the whole duration, $S \gg \bar{n}_0[1 - e^{-\kappa_s T}]$:

$$\begin{aligned}
F_T(\bar{n}_0, T, S) &= \lim_{S \rightarrow \infty} f_0^S \left(1 + \frac{f_1 \bar{n}_0 [1 - e^{-\kappa_s T}]}{f_0 S} \right)^S e^{-\bar{n}_0 [1 - e^{-\kappa_s T}]} \\
&\approx f_0^S e^{\bar{n}_0 [1 - e^{-\kappa_s T}] \frac{f_1}{f_0}} e^{-\bar{n}_0 [1 - e^{-\kappa_s T}]} \\
&= f_0^S e^{-(1 - f_1/f_0) \bar{n}_0 [1 - e^{-\kappa_s T}]} \\
&\approx f_0^S \left(\frac{f_1}{f_0} \right)^{\bar{n}_0 [1 - e^{-\kappa_s T}]} \\
&= f_0^{S - \bar{n}_0 [1 - e^{-\kappa_s T}]} f_1^{\bar{n}_0 [1 - e^{-\kappa_s T}]} \tag{B.15}
\end{aligned}$$

What this limit means is that when we measure frequently enough, the fidelity will fall off by a factor f_0 for any steps when no jumps happened and by f_1 when a single jump occurs. As the step size is so short, errors due to double photon jumps are negligible and therefore excluded.

We continue with F'_T :

$$n_j \equiv \bar{n}_0 [1 - e^{-\kappa_s T}] \cdot \frac{f_1}{f_0} \tag{B.16}$$

↓

$$F'_T(T, \bar{n}_0, S) = f_0^S \left(1 + \frac{n_j}{S} \right)^S, \tag{B.17}$$

where n_j is the number of expected photon jumps during the whole duration up to a scale factor of order 1. We will treat the number of steps as a continuous variable and find its optimum. In practice, we will use the closest integer:

$$\begin{aligned}
\frac{d}{dS} F'_T &= F'_T \cdot \frac{d}{dS} \left[S \cdot \log(f_0) + S \cdot \log \left(1 + \frac{n_j}{S} \right) \right] \\
&= F'_T \cdot \left[\log(f_0) + \log \left(1 + \frac{n_j}{S} \right) + \frac{S}{\left(1 + \frac{n_j}{S} \right)} \cdot \frac{-n_j}{S^2} \right] \\
&= F'_T \cdot \left[\log(f_0) + \log \left(1 + \frac{n_j}{S} \right) - \frac{n_j}{S + n_j} \right] \stackrel{\text{want}}{=} 0 \tag{B.18}
\end{aligned}$$

We need to solve:

$$\log(f_0) + \log \left(1 + \frac{n_j}{S} \right) - \frac{n_j}{S + n_j} = 0 \tag{B.19}$$

↓

$$\log(f_0) + \log \left(1 + \frac{1}{r} \right) - \frac{1}{1 + r} = 0 \tag{B.20}$$

where $r \equiv S/n_j$ (the number of steps per photon jump up to small correction). Although we cannot solve for r explicitly, it is a function of f_0 , the success probability of the parity measurement conditioned on no photon jumps ($r(f_0)$). The optimal number of steps is then:

$$S^{OPT} = r(f_0) \cdot n_j = \frac{r(f_0) \cdot f_1}{f_0} \bar{n}_0 [1 - e^{-\kappa_s T}], \quad (\text{B.21})$$

leading to the optimal average number of photon jumps per step:

$$\lambda_k = \frac{\bar{n}_0 [1 - e^{-\kappa_s T}]}{S^{OPT}} = \frac{f_0}{r(f_0) \cdot f_1}. \quad (\text{B.22})$$

In the limit $f_0 \rightarrow 1$, $r(f_0)$ approaches infinity. We can expand eq. B.20 in $1/r$:

$$\begin{aligned} \log(f_0) &= \frac{1}{1+r} - \log\left(1 + \frac{1}{r}\right) \\ \log(1 - [1 - f_0]) &= \frac{1}{r} \cdot \frac{1}{1 + \frac{1}{r}} - \left[\frac{1}{r} - \frac{1}{2r^2} + \dots\right] \\ -[1 - f_0] - \frac{1}{2}[1 - f_0]^2 - \dots &= \frac{1}{r} - \frac{1}{r^2} + \dots - \left[\frac{1}{r} - \frac{1}{2r^2} + \dots\right] \\ 1 - f_0 &\approx \frac{1}{2r^2} \end{aligned} \quad (\text{B.23})$$

For $f_0, f_1 \sim 1$, the optimized average number of jumps per step is simply $1/r$. Hence, $1/2r^2$ is the Poissonian probability to have two jumps. The last approximation states that in this limit we need to match the parity measurement infidelity with the double-jump probability. As long as f_0 is above 90%, this approximation will be correct within 50% of the right value.

We can now substitute the optimal steps number and get the maximal success probability as a function of the total duration and initial cat size:

$$\begin{aligned} F_T'^{OPT}(T, \bar{n}_0) \equiv F_T'(T, \bar{n}_0, S^{OPT}) &= \left[f_0 \left(1 + \frac{n_j}{S^{OPT}} \right) \right]^{S^{OPT}} \\ \log(F_T'^{OPT}) &= S^{OPT} \underbrace{\left[\log(f_0) + \log\left(1 + \frac{1}{r}\right) \right]}_{= \frac{1}{1+r}} \\ &= \frac{r(f_0)}{1+r(f_0)} \frac{f_1}{f_0} \cdot \bar{n}_0 [1 - e^{-\kappa_s T}] \end{aligned} \quad (\text{B.24})$$

$$\begin{aligned} \downarrow \\ F_T^{OPT}(T, \bar{n}_0) = F_T'^{OPT}(T, \bar{n}_0) \cdot e^{-\bar{n}_0[1 - e^{-\kappa_s T}]} &= e^{+\bar{n}_0[1 - e^{-\kappa_s T}]} \frac{r(f_0)}{1+r(f_0)} \frac{f_1}{f_0} \cdot e^{-\bar{n}_0[1 - e^{-\kappa_s T}]} \end{aligned} \quad (\text{B.25})$$

The second term is the uncorrected cat state decay, which is \bar{n}_0 faster than κ_s . The first term counteracts this decay, and so we identify it as the action of the QEC.

We define the unit-less parameter G as follows:

$$G \equiv \frac{1}{1 - \frac{f_1}{f_0} \cdot \frac{r(f_0)}{1+r(f_0)}} \quad (\text{B.26})$$

$$\begin{aligned} &\downarrow \\ F_T^{OPT} &= e^{-\frac{\bar{n}_0[1-e^{-\kappa_s T}]}{G}} \underset{\kappa_s T \ll 1}{\approx} e^{-\kappa_s T \cdot \frac{\bar{n}_0}{G}}. \end{aligned} \quad (\text{B.27})$$

In other words, we slow down the decay of the cat state by a factor of G . This number is a constant of the system, a function of the various infidelities. It approaches infinity as f_0, f_1 get closer to 1.

References

- Aaronson, S. Guest column: Np-complete problems and physical reality. *SIGACT News*, 36(1):30–52, 2005.
- Ambegaokar, V. and Baratoff, A. Tunneling between superconductors. *Phys. Rev. Lett.*, 10:486–489, 1963.
- Ansmann, M., Wang, H., Bialczak, R. C., Hofheinz, M., Lucero, E., Neeley, M., O’Connell, A. D., Sank, D., Weides, M., Wenner, J., Cleland, A. N., and Martinis, J. M. Violation of Bell’s inequality in Josephson phase qubits. *Nature*, 461(7263):504–506, 2009.
- Aoki, T., Takahashi, G., Kajiya, T., Yoshikawa, J., Braunstein, S. L., van Loock, P., and Furusawa, A. Quantum error correction beyond qubits. *Nat. Phys.*, 5(8):541–546, 2009.
- Aspect, A., Grangier, P., and Roger, G. Experimental tests of realistic local theories via Bell’s theorem. *Phys. Rev. Lett.*, 47(7):460–463, 1981.
- Barends, R., Kelly, J., Megrant, A., Sank, D., Jeffrey, E., Chen, Y., Yin, Y., Chiaro, B., Mutus, J., Neill, C., O’Malley, P., Roushan, P., Wenner, J., White, T. C., Cleland, A. N., and Martinis, J. M. Coherent josephson qubit suitable for scalable quantum integrated circuits. *Phys. Rev. Lett.*, 111:080502, 2013.
- Barends, R., Kelly, J., Megrant, A., Veitia, A., Sank, D., Jeffrey, E., White, T. C., Mutus, J., Fowler, A. G., Campbell, B., Chen, Y., Chen, Z., Chiaro, B., Dunsworth, A., Neill, C., O’Malley, P., Roushan, P., Vainsencher, A., Wenner, J., Korotkov, A. N., Cleland, A. N., and Martinis, J. M. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, 2014.
- Belavkin, V. P. Quantum stochastic calculus and quantum nonlinear filtering. *J. Multi. Anal.*, 42:171, 1992.
- Bell, J. S. On the einstein-podolsky-rosen paradox. *Physics*, 1:195–200, 1964.
- Bell, J. S. *Speakable and Unsayable in Quantum Mechanics*. Camb. Univ. Press, 1987.
- Bergeal, N., Schackert, F., Metcalfe, M., Vijay, R., Manucharyan, V. E., Frunzio, L., Prober, D. E., Schoelkopf, R. J., Girvin, S. M., and Devoret, M. H. Phase-preserving amplification near the quantum limit with a Josephson ring modulator. *Nature*, 465(7294):64–68, 2010.

- Bergquist, J. C., Hulet, R. G., Itano, W. M., and Wineland, D. J. Observation of quantum jumps in a single atom. *Phys. Rev. Lett.*, 57:1699–1702, 1986.
- Bertet, P., Auffeves, P., A .and Maioli, Osnaghi, S., Meunier, T., Brune, M., Raimond, J. M., and Haroche, S. Direct Measurement of the Wigner Function of a One-Photon Fock State in a Cavity. *Phys. Rev. Lett.*, 89(20):200402, 2002.
- Blais, A., Huang, R., Wallraff, A., Girvin, S. M., and Schoelkopf, R. J. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Phys. Rev. A*, 69:062320, 2004a.
- Blais, A., Huang, R., Wallraff, A., Girvin, S. M., and Schoelkopf, R. J. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Phys. Rev. A*, 69:062320, 2004b.
- Blumoff, J. Z., Chou, K., Shen, C., Reagor, M. J., Axline, C., Brierley, R. T., Silveri, M. P., Wang, C., Vlastakis, B., Nigg, S. E., Frunzio, L., Devoret, M. H., Jiang, L., Girvin, S. M., and Schoelkopf, R. J. Implementing and Characterizing Precise Multiqubit Measurements. *Phys. Rev. X*, 6:031041, 2016.
- Bouchiat, V., Vion, D., Joyez, P., Esteve, D., and Devoret, M. H. Quantum coherence with a single cooper pair. *Phys. Scripta*, 1998(T76):165, 1998.
- Bouten, L., van Handel, R., and James, M. An introduction to quantum filtering. *SIAM J. Cont. Optim.*, 46:2199, 2007.
- Braunstein, S. L. Quantum error correction for communication with linear optics. *Nature*, 394(6688):47–49, 1998.
- Braunstein, S. L. Quantum information with continuous variables. *Rev. Mod. Phys.*, 77(2):513–577, 2005.
- Braunstein, S. L. and Kimble, H. J. Teleportation of continuous quantum variables. *Phys. Rev. Lett.*, 80: 869–872, 1998.
- Braunstein, S. L. and Kimble, H. J. Dense coding for continuous variables. *Phys. Rev. A*, 61(4):042302, 2000.
- Brecht, T., Reagor, M. J., Chu, Y., Pfaff, W., Wang, C., Frunzio, L., Devoret, M. H., and Schoelkopf, R. J. Demonstration of superconducting micromachined cavities. *Applied Physics Letters*, 107(19), 2015.
- Brecht, T., Pfaff, W., Wang, C., Chu, Y., Frunzio, L., Devoret, M. H., and Schoelkopf, R. J. Multilayer microwave integrated quantum circuits for scalable quantum computing. *Npj Quant. Inf.*, 2:16002 EP, 2016.

- Brune, M., Haroche, S., Raimond, J. M., Davidovich, L., and Zagury, N. Manipulation of photons in a cavity by dispersive atom-field coupling: Quantum-nondemolition measurements and generation of “Schrödinger cat” states. *Phys. Rev. A*, 45(7):5193–5214, 1992.
- Brune, M., Hagley, E., Dreyer, J., Maître, X., Maali, A., Wunderlich, C., Raimond, J. M., and Haroche, S. Observing the Progressive Decoherence of the “Meter” in a Quantum Measurement. *Phys. Rev. Lett.*, 77(24):4887–4890, 1996.
- Cahill, K. E. and Glauber, R. J. Density operators and quasiprobability distributions. *Phys. Rev.*, 177:1882–1902, 1969.
- Calderbank, A. R. and Shor, P. W. Good quantum error-correcting codes exist. *Phys. Rev. A*, 54:1098–1105, 1996.
- Castellanos-Beltran, M. A., Irwin, K. D., Hilton, G. C., Vale, L. R., and Lehnert, K. W. Amplification and squeezing of quantum noise with a tunable Josephson metamaterial. *Nat. Phys.*, 4(12):929–931, 2008.
- Chiaverini, J., Leibfried, D., Schaetz, T., Barrett, M. D., Blakestad, R. B., Britton, J., Itano, W. M., Jost, J. D., Knill, E., Langer, C., Ozeri, R., and Wineland, D. J. Realization of quantum error correction. *Nature*, 432(7017):602–605, 2004.
- Chow, J. M., DiCarlo, L., Gambetta, J. M., Motzoi, F., Frunzio, L., Girvin, S. M., and Schoelkopf, R. J. Optimized driving of superconducting artificial atoms for improved single-qubit gates. *Phys. Rev. A*, 82:040305, 2010a.
- Chow, J. M., DiCarlo, L., Gambetta, J. M., Nunnenkamp, A., Bishop, L. S., Frunzio, L., Devoret, M. H., Girvin, S. M., and Schoelkopf, R. J. Detecting highly entangled states with a joint qubit readout. *Phys. Rev. A*, 81(6):062325, 2010b.
- Chow, J. M., Gambetta, J. M., Córcoles, A. D., Merkel, S. T., Smolin, J. A., Rigetti, C., Poletto, S., Keefe, G. A., Rothwell, M. B., Rozen, J. R., Ketchen, M. B., and Steffen, M. Universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits. *Phys. Rev. Lett.*, 109:060501, 2012.
- Chow, J. M., Gambetta, J. M., Magesan, E., Srinivasan, S. J., Cross, A. W., Abraham, D. W., Masluk, N. A., Johnson, B. R., Ryan, C. A., and Steffen, M. Implementing a strand of a scalable fault-tolerant quantum computing fabric. *Nat. Comm.*, 5:4015, 2013.
- Chuang, I. and Gottesman, D. Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations. *Nature*, 402(6760):390–393, 1999.
- Clarke, J., Cleland, A. N., Devoret, M. H., Esteve, D., and Martinis, J. M. Quantum mechanics of a macroscopic variable: The phase difference of a Josephson junction. *Science*, 239(4843):992–997, 1988.

- Clauser, J. F., Horne, M. A., Shimony, A., and Holt, R. A. Proposed experiment to test local hidden-variable theories. *Phys. Rev. Lett.*, 23(15):880–884, 1969.
- Cohen-Tannoudji, C., Dupont-Roc, J., and Grynberg, G. *Photons and Atoms*. Wiley, New York, 1992.
- Cooper, L. N. Bound electron pairs in a degenerate fermi gas. *Phys. Rev.*, 104:1189–1190, 1956.
- Córcoles, A. D., Magesan, E., Srinivasan, S. J., Cross, A. W., Steffen, M., Gambetta, J. M., and Chow, J. M. Demonstration of a quantum error detection code using a square lattice of four superconducting qubits. *Nat. Comm.*, 6:6979, 2015.
- Cory, D., Price, M., Maas, W., Knill, E., Laflamme, R., Zurek, W. H., Havel, T., and Somaroo, S. Experimental Quantum Error Correction. *Phys. Rev. Lett.*, 81(10):2152–2155, 1998.
- Cramer, J., Kalb, N., Rol, M. A., Hensen, B., Blok, M. S., Markham, M., Twitchen, D. J., Hanson, R., and Taminiau, T. H. Repeated quantum error correction on a continuously encoded qubit by real-time feedback. *Nat. Comm.*, 7:11526, 2016.
- da Silva, M., Landon-Cardinal, O., and Poulin, D. Practical Characterization of Quantum Devices without Tomography. *Phys. Rev. Lett.*, 107(21):210404, 2011.
- de Fouquieres, P., Schirmer, S. G., Glaser, S. J., and Kuprov, I. Second order gradient ascent pulse engineering. *J. Mag. Res.*, 212(2):412 – 417, 2011.
- Deléglise, S., Dotsenko, I., Sayrin, C., Bernu, J., Brune, M., Raimond, J. M., and Haroche, S. Reconstruction of non-classical cavity field states with snapshots of their decoherence. *Nature*, 455(7212):510–514, 2008.
- Deutsch, D. and Jozsa, R. Rapid solution of problems by quantum computation. *Proc. R. S. Lon. A: Math., Phys. Eng. Sci.*, 439(1907):553–558, 1992.
- Devoret, M. H. Quantum fluctuations in electrical circuits. In Giacobino, E. and David, F., editors, *Quantum Fluctuations*, volume 63, chapter 10. Elsevier, 1997.
- Devoret, M. H. and Glattli, C. Single-electron transistors. *Phys. World*, 11(9):29, 1998.
- Devoret, M. H. and Martinis, J. M. Implementing qubits with superconducting integrated circuits. *Quant. Inf. Proc.*, 3(1):163–203, 2004.
- Devoret, M. H. and Schoelkopf, R. J. Superconducting Circuits for Quantum Information: An Outlook. *Science*, 339(6124):1169–1174, 2013.
- Devoret, M. H., Martinis, J. M., and Clarke, J. Measurements of macroscopic quantum tunneling out of the zero-voltage state of a current-biased josephson junction. *Phys. Rev. Lett.*, 55:1908–1911, 1985.

- Dolan, G. J. Offset masks for lift-off photoprocessing. *Appl. Phys. Lett.*, 31(5):337–339, 1977.
- Eichler, C., Lang, C., Fink, J. M., Govenius, J., Philipp, S., and Wallraff, A. Observation of Entanglement Between Itinerant Microwave Photons and a Superconducting Qubit. *Phys. Rev. Lett.*, 109(24):240501, 2012.
- Einstein, A., Podolsky, B., and Rosen, N. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev.*, 47:777–780, 1935.
- Feynman, R. P. Simulating physics with computers. *Inter. J. Theor. Phys.*, 21(6):467–488, 1982.
- Flammia, S. T. and Liu, Y. Direct Fidelity Estimation from Few Pauli Measurements. *Phys. Rev. Lett.*, 106(23):230501, 2011.
- Fowler, A. G., Stephens, A., and Groszkowski, P. High-threshold universal quantum computation on the surface code. *Phys. Rev. A*, 80(5):052312, 2009.
- Fowler, A. G., Mariantoni, M., Martinis, J. M., and Cleland, A. N. Surface codes: Towards practical large-scale quantum computation. *Phys. Rev. A*, 86(3):032324, 2012.
- Freedman, S. J. and Clauser, J. F. Experimental test of local hidden-variable theories. *Phys. Rev. Lett.*, 28(14):938–941, 1972.
- Gilchrist, A., Deuar, P., and Reid, M. D. Contradiction of quantum mechanics with local hidden variables for quadrature phase amplitude measurements. *Phys. Rev. Lett.*, 80:3196, 1998.
- Girvin, S. M. Circuit QED: Superconducting qubits coupled to microwave photons. In Devoret, M. H., Huard, B., Schoelkopf, R. J., and Cugliandolo, L. F., editors, *Quantum Machines: Measurement Control of Engineered Quantum Systems*, volume 96, chapter 3. Oxf. Univ. Press, 2011.
- Glauber, R. J. Coherent and incoherent states of the radiation field. *Phys. Rev.*, 131:2766–2788, 1963.
- Gleyzes, S., Kuhr, S., Guerlin, C., Bernu, J., Deleglise, S., Hoff, U. B., Brune, M., Raimond, J. M., and Haroche, S. Quantum jumps of light recording the birth and death of a photon in a cavity. *Nature*, 446(7133):297, 2007.
- Gottesman, D. Stabilizer codes and quantum error correction. *Doc. Diss.*, 1997.
- Gottesman, D. Theory of fault-tolerant quantum computation. *Phys. Rev. A*, 57(1):127–137, 1998.
- Gottesman, D., Kitaev, A., and Preskill, J. Encoding a qubit in an oscillator. *Phys. Rev. A*, 64(1):012310, 2001.
- Grover, L. K. Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.*, 79(2):325–328, 1997.

- Guerlin, C., Bernu, J., Deleglise, S., Sayrin, C., Gleyzes, S., S., K., Brune, M., Raimond, J. M., and Haroche, S. Progressive field-state collapse and quantum non-demolition photon counting. *Nature*, 448(7156):889, 2007.
- Haroche, S. and Raimond, J. M. *Exploring the Quantum: Atoms, Cavities, and Photons*. Oxf. Univ. Press, 2006.
- Haroche, S., Brune, M., and Raimond, J. M. Measuring the photon number parity in a cavity: from light quantum jumps to the tomography of non-classical field states. *J. Mod. Opt.*, 54(13-15):2101–2114, 2007.
- Hatridge, M. J., Shankar, S., Mirrahimi, M., Schackert, F., Geerlings, K., Brecht, T., Sliwa, K. M., Abdo, B., Frunzio, L., Girvin, S. M., Schoelkopf, R. J., and Devoret, M. H. Quantum Back-Action of an Individual Variable-Strength Measurement. *Science*, 339(6116):178–181, 2013.
- Heeres, R. W., Vlastakis, B., Holland, E. T., Krastanov, S., Albert, V. V., Frunzio, L., Jiang, L., and Schoelkopf, R. J. Cavity state manipulation using photon-number selective phase gates. *Phys. Rev. Lett.*, 115:137002, 2015.
- Hensen, B., Bernien, H., Dreau, A. E., Reiserer, A., Kalb, N., Blok, M. S., Ruitenbergh, J., Vermeulen, R. F. L., Schouten, R. N., Abellan, C., Amaya, W., Pruneri, V., Mitchell, M. W., Markham, M., Twitche, D. J., Elkouss, D., Wehner, S., Taminiau, T. H., and Hanson, R. Loophole-free bell inequality violation using electron spins separated by 1.3 kilometres. *Nature*, 526(7575):682–686, 2015.
- Hofheinz, M., Wang, H., Ansmann, M., Bialczak, R. C., Lucero, R., Neeley, M., O’Connell, A. D., Sank, D., Wenner, J., and Martinis, J. M. Synthesizing arbitrary quantum states in a superconducting resonator. *Nature*, 459(7246):546–549, 2009.
- Hofmann, J., Krug, M., Ortengel, N., Gérard, L., Weber, M., Rosenfeld, W., and Weinfurter, H. Heralded entanglement between widely separated atoms. *Science*, 337(6090):72–75, 2012.
- Holland, E. T., Vlastakis, B., Heeres, R. W., Reagor, M. J., Vool, U., Leghtas, Z., Frunzio, L., Kirchmair, G., Devoret, M. H., Mirrahimi, M., and Schoelkopf, R. J. Single-photon-resolved cross-kerr interaction for autonomous stabilization of photon-number states. *Phys. Rev. Lett.*, 115:180501, 2015.
- Hood, C. J., Lynn, T. W., Doherty, A. C., Parkins, A. S., and Kimble, H. J. The atom-cavity microscope: Single atoms bound in orbit by single photons. *Science*, 287(5457):1447–1453, 2000.
- Houck, A. A., Schreier, J. A., Johnson, B. R., Chow, J. M., Koch, J., Gambetta, J. M., Schuster, D. I., Frunzio, L., Devoret, M. H., Girvin, S. M., and Schoelkopf, R. J. Controlling the Spontaneous Emission of a Superconducting Transmon Qubit. *Phys. Rev. Lett.*, 101(8):080502, 2008.
- Houck, A. A., Koch, J., Devoret, M. H., Girvin, S. M., and Schoelkopf, R. J. Life after charge noise: recent results with transmon qubits. *Quant. Inf. Proc.*, 8(2):105–115, 2009.

- Jaynes, E. T. and Cummings, F. W. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proc. IEEE*, 51(1):89–109, 1963.
- Jelezko, F., Popa, I., Gruber, A., Tietz, C., Wrachtrup, J., Nizovtsev, A., and Kilin, S. Single spin states in a defect center resolved by optical spectroscopy. *Appl. Phys. Lett.*, 81(12):2160–2162, 2002.
- Jensen, K., Wasilewski, W., Krauter, H., Fernholz, T., Nielsen, B. M., Owari, M., Plenio, M. B., Serafini, A., Wolf, M. M., and Polzik, E. S. Quantum memory for entangled continuous-variable states. *Nat. Phys.*, 7(1):13–16, 2011.
- Jeong, H. and Kim, M. Efficient quantum computation using coherent states. *Phys. Rev. A*, 65(4):042305, 2002.
- Jeong, H., Zavatta, A., Kang, M., Lee, S., Costanzo, L. S., Grandi, S., Ralph, T. C., and Bellini, M. Generation of hybrid entanglement of light. *Nat. Phot.*, 8(7):564–569, 2014.
- Johansson, J. R., Nation, P. D., and Nori, F. Qutip: An open-source python framework for the dynamics of open quantum systems. *Comp. Phys. Comm.*, 184(8):1760 – 1772, 2012.
- Johansson, J. R., Nation, P. D., and Nori, F. QuTiP 2: A Python framework for the dynamics of open quantum systems. *Comp. Phys. Comm.*, 184(4):1234–1240, 2013.
- Johnson, B. R., Reed, M. D., Houck, A. A., Schuster, D. I., Bishop, L. S., Ginossar, E., Gambetta, J. M., DiCarlo, L., Frunzio, L., Girvin, S. M., and Schoelkopf, R. J. Quantum non-demolition detection of single microwave photons in a circuit. *Nat. Phys.*, 6(9):663–667, 2010.
- Josephson, B. D. Possible new effects in superconductive tunnelling. *Phys. Lett.*, 1(7):251 – 253, 1962.
- Josephson, B. D. Supercurrents through barriers. *Adv. in Phys.*, 14(56):419–451, 1965.
- Kamal, A., Marblestone, A., and Devoret, M. H. Signal-to-pump back action and self-oscillation in double-pump Josephson parametric amplifier. *Phys. Rev. B*, 79:184301, 2009.
- Kelly, J., Barends, R., Fowler, A. G., Megrant, A., Jeffrey, E., White, T. C., Sank, D., Mutus, J. Y., Campbell, B., Chen, Y., Chen, Z., Chiaro, B., Dunsworth, A., Hoi, I. C., Neill, C., O’Malley, P., Quintana, C., Roushan, P., Vainsencher, A., Wenner, J., Cleland, A. N., and Martinis, J. M. State preservation by repetitive error detection in a superconducting quantum circuit. *Nature*, 519:66–69, 2015.
- Khaneja, N., Reiss, T., Kehlet, C., Schulte-Herbrüggen, T., and Glaser, S. J. Optimal control of coupled spin dynamics: design of nmr pulse sequences by gradient ascent algorithms. *J. Mag. Res.*, 172(2):296 – 305, 2005.

- Kirchmair, G., Vlastakis, B., Leghtas, Z., Nigg, S. E., Paik, H., Ginossar, E., Mirrahimi, M., Frunzio, L., Girvin, S. M., and Schoelkopf, R. J. Observation of quantum state collapse and revival due to the single-photon Kerr effect. *Nature*, 495(7440):205–209, 2013.
- Knill, E. and Laflamme, R. Theory of quantum error-correcting codes. *Phys. Rev. A*, 55(7541):900–911, 1997.
- Knill, E., Laflamme, R., and Zurek, W. H. Resilient quantum computation: error models and thresholds. *Proc. R. S. Lon. A: Math., Phys. Eng. Sci.*, 454(1969):365–384, 1998.
- Knill, E., Laflamme, R., Martinez, R., and Negrevergne, C. Benchmarking quantum computers: The five-qubit error correcting code. *Phys. Rev. Lett.*, 86:5811–5814, 2001a.
- Knill, E., Laflamme, R., and Milburn, G. J. A scheme for efficient quantum computation with linear optics. *Nature*, 409(6816):46–52, 2001b.
- Koch, J., Yu, T., Gambetta, J. M., Houck, A. A., Schuster, D. I., Majer, J., Blais, A., Devoret, M. H., Girvin, S. M., and Schoelkopf, R. J. Charge-insensitive qubit design derived from the Cooper pair box. *Phys. Rev. A*, 76(4):042319, 2007.
- Kofman, A. G. and Korotkov, A. N. Analysis of Bell inequality violation in superconducting phase qubits. *Phys. Rev. B*, 77(10):104502, 2008.
- Ladd, T. D., Jelezko, F., Laflamme, R., Nakamura, Y., Monroe, C., and O’Brien, J. L. Quantum computers. *Nature*, 464(7285):45–53, 2010.
- Lecocq, F., Pop, I. M., Peng, Z., Matei, I., Crozes, T., Fournier, T., Naud, C., Guichard, W., and Buisson, O. Junction fabrication by shadow evaporation without a suspended bridge. *Nanotechnology*, 22(31):315302, 2011.
- Leghtas, Z., Kirchmair, G., Vlastakis, B., Schoelkopf, R. J., Devoret, M. H., and Mirrahimi, M. Hardware-Efficient Autonomous Quantum Memory Protection. *Phys. Rev. Lett.*, 111(12):120501, 2013.
- Leghtas, Z., Touzard, S., Pop, I. M., Kou, A., Vlastakis, B., Petrenko, A., Sliwa, K. M., Narla, A., Shankar, S., Hatridge, M. J., Reagor, M. J., Frunzio, L., Schoelkopf, R. J., Mirrahimi, M., and Devoret, M. H. Confining the state of light to a quantum manifold by engineered two-photon loss. *Science*, 347(6224):853–857, 2015a.
- Leghtas, Z., Touzard, S., Pop, I. M., Kou, A., Vlastakis, B., Petrenko, A., Sliwa, K. M., Narla, A., Shankar, S., Hatridge, M. J., Reagor, M. J., Frunzio, L., Schoelkopf, R. J., Mirrahimi, M., and Devoret, M. H. Confining the state of light to a quantum manifold by engineered two-photon loss: Supplement. *Science*, 347(6224):853–857, 2015b.

- Leibfried, D., Meekhof, D., King, B., Monroe, C., Itano, W. M., and Wineland, D. J. Experimental Determination of the Motional Quantum State of a Trapped Atom. *Phys. Rev. Lett.*, 77(21):4281–4285, 1996.
- Leonhardt, U. and Vaccaro, J. A. Bell correlations in phase space: application to quantum optics. *J. Mod. Opt.*, 42(5):939–943, 1995.
- Leung, D., Vandersypen, L., Zhou, X., Sherwood, M., Yannoni, C., Kubinec, M., and Chuang, I. Experimental realization of a two-bit phase damping quantum code. *Phys. Rev. A*, 60:1924–1943, 1999.
- LinPeng, X. Y., Zhang, H. Z., Xu, K., Li, C. Y., Zhong, Y. P., Wang, Z. L., Wang, H., and Xie, Q. W. Joint quantum state tomography of an entangled qubit–resonator hybrid. *New J. Phys.*, 15(12):125027, 2013.
- Liu, Y., Shankar, S., Ofek, N., Hatridge, M. J., Narla, A., Sliwa, K. M., Frunzio, L., Schoelkopf, R. J., and Devoret, M. H. Comparing and combining measurement-based and driven-dissipative entanglement stabilization. *Phys. Rev. X*, 6:011022, 2016.
- Lloyd, S. and Braunstein, S. L. Quantum computation over continuous variables. *Phys. Rev. Lett.*, 82:1784–1787, 1999.
- Lloyd, S. and Slotine, J. J. E. Analog quantum error correction. *Phys. Rev. Lett.*, 80:4088–4091, 1998.
- Lund, A. P., Ralph, T. C., and Haselgrove, H. L. Fault-tolerant linear optical quantum computing with small-amplitude coherent states. *Phys. Rev. Lett.*, 100:030503, 2008.
- Lutterbach, L. G. and Davidovich, L. Method for direct measurement of the Wigner function in cavity QED and ion traps. *Phys. Rev. Lett.*, 78(13):2547–2550, 1997.
- Mabuchi, H. and Doherty, A. C. Cavity quantum electrodynamics: Coherence in context. *Science*, 298(5597):1372–1377, 2002.
- Majer, J., Chow, J. M., Gambetta, J. M., Koch, J., Johnson, B. R., Schreier, J. A., Frunzio, L., Schuster, D. I., Houck, A. A., Wallraff, A., Blais, A., Devoret, M. H., Girvin, S. M., and Schoelkopf, R. J. Coupling superconducting qubits via a cavity bus. *Nature*, 449(7161):443–447, 2007.
- Manucharyan, V. E., Koch, J., Glazman, L. I., and Devoret, M. H. Fluxonium: Single cooper-pair circuit free of charge offsets. *Science*, 326(5949):113–116, 2009.
- Martinis, J. M. and Megrant, A. UCSB final report for the CSQ program: Review of decoherence and materials physics for superconducting qubits. *arXiv:quant-ph/1410.5793*, 2014.
- Martinis, J. M., Nam, S., Aumentado, J., and Urbina, C. Rabi oscillations in a large josephson-junction qubit. *Phys. Rev. Lett.*, 89:117901, 2002.

- Martinis, J. M., Cooper, K. B., McDermott, R., Steffen, M., Ansmann, M., Osborn, K. D., Cicak, K., Oh, S., Pappas, D. P., Simmonds, R. W., and Yu, C. C. Decoherence in josephson qubits from dielectric loss. *Phys. Rev. Lett.*, 95:210503, 2005.
- Menicucci, N., van Loock, P., Gu, M., Weedbrook, C., Ralph, T. C., and Nielsen, M. Universal Quantum Computation with Continuous-Variable Cluster States. *Phys. Rev. Lett.*, 97(11):110501, 2006.
- Michael, M., Silveri, M. P., Brierley, R. T., Albert, V. V., Salmilehto, J., Jiang, L., and Girvin, S. M. New Class of Quantum Error-Correcting Codes for a Bosonic Mode. *Phys. Rev. X*, 6:031006, 2016.
- Milman, P., Auffeves, A., Yamaguchi, F., Brune, M., Raimond, J. M., and Haroche, S. A proposal to test Bell's inequalities with mesoscopic non-local states in cavity QED. *The European Physical Journal D*, 32(2):233–239, 2004.
- Mineev, Z. K., Serniak, K., Pop, I. M., Leghtas, Z., Sliwa, K. M., Hatridge, M. J., Frunzio, L., Schoelkopf, R. J., and Devoret, M. H. Planar multilayer circuit quantum electrodynamics. *Phys. Rev. Appl.*, 5:044021, 2016.
- Mirrahimi, M., Leghtas, Z., Albert, V. V., Touzard, S., Schoelkopf, R. J., Jiang, L., and Devoret, M. H. Dynamically protected cat-qubits: a new paradigm for universal quantum computation. *New J. Phys.*, 16(4):045014, 2014.
- Mooij, J. E., Orlando, T. P., Levitov, L., Tian, L., van der Wal, C. H., and Lloyd, S. Josephson persistent-current qubit. *Science*, 285(5430):1036–1039, 1999.
- Morin, O., Huang, K., Liu, J., Le Jeannic, H., Fabre, C., and Laurat, J. Remote creation of hybrid entanglement between particle-like and wave-like optical qubits. *Nat. Phot.*, 8(7):570–574, 2014.
- Motzoi, F., Gambetta, J. M., Rebentrost, P., and Wilhelm, F. K. Simple pulses for elimination of leakage in weakly nonlinear qubits. *Phys. Rev. Lett.*, 103:110501, 2009.
- Moussa, O., Baugh, J., Ryan, C. A., and Laflamme, R. Demonstration of Sufficient Control for Two Rounds of Quantum Error Correction in a Solid State Ensemble Quantum Information Processor. *Phys. Rev. Lett.*, 107(16):160501, 2011.
- Murch, K. W., Weber, S. J., Macklin, C., and Siddiqi, I. Observing single quantum trajectories of a superconducting quantum bit. *Nature*, 502(7470):211–214, 2013.
- Nagourney, W., Sandberg, J., and Dehmelt, H. Shelved optical electron amplifier: Observation of quantum jumps. *Phys. Rev. Lett.*, 56:2797–2799, 1986.
- Nakamura, Y., Pashkin, Y. A., and Tsai, J. S. Coherent control of macroscopic quantum states in a single-cooper-pair box. *Nature*, 398(6730):786–788, 1999.

- Narla, A., Shankar, S., Hatridge, M. J., Leghtas, Z., Sliwa, K. M., Zaly-Geller, E., Mundhada, S. O., Pfaff, W., Frunzio, L., Schoelkopf, R. J., and Devoret, M. H. Robust concurrent remote entanglement between two superconducting qubits. *Phys. Rev. X*, 6:031036, 2016.
- Neumann, P., Beck, J., Steiner, M., Rempp, F., Fedder, H., Hemmer, P. R., Wrachtrup, J., and Jelezko, F. Single-shot readout of a single nuclear spin. *Science*, 329:542, 2010.
- Nielsen, M. and Chuang, I. *Quantum computation and quantum information*. Camb. Univ. Press, 2010.
- Nigg, D., Müller, M., Martinez, E. A., Schindler, P., Hennrich, M., Monz, T., Martin-Delgado, M. A., and Blatt, R. Quantum computations on a topologically encoded qubit. *Science*, 345(6194):302–305, 2014.
- Nigg, S. E., Paik, H., Vlastakis, B., Kirchmair, G., Shankar, S., Frunzio, L., Devoret, M. H., Schoelkopf, R. J., and Girvin, S. M. Black-box superconducting circuit quantization. *Phys. Rev. Lett.*, 108:240502, 2012.
- Ofek, N., Petrenko, A., Heeres, R., Reinhold, P., Leghtas, Z., Vlastakis, B., Liu, Y., Frunzio, L., Devoret, M. H., Jiang, L., and Schoelkopf, R. J. Extending the Lifetime of a Quantum Bit with Error Correction in Superconducting Circuits. *Nature*, 536(7616):441–445, 2016.
- Paik, H., Schuster, D. I., Bishop, L. S., Kirchmair, G., Catelani, G., Sears, A. P., Johnson, B. R., Reagor, M. J., Frunzio, L., Glazman, L. I., Girvin, S. M., Devoret, M. H., and Schoelkopf, R. J. Observation of High Coherence in Josephson Junction Qubits Measured in a Three-Dimensional Circuit QED Architecture. *Phys. Rev. Lett.*, 107(24):240501, 2011.
- Park, J., Saunders, M., Shin, Y., An, K., and Jeong, H. Bell-inequality tests with entanglement between an atom and a coherent state in a cavity. *Phys. Rev. A*, 85(2):022120, 2012.
- Peil, S. and Gabrielse, G. Observing the quantum limit of an electron cyclotron: Qnd measurements of quantum jumps between fock states. *Phys. Rev. Lett.*, 83:1287–1290, 1999.
- Pfaff, W., Taminiau, T. H., Robledo, L., Bernien, H., Markham, M., Twitchen, D. J., and Hanson, R. Demonstration of entanglement-by-measurement of solid-state qubits. *Nat. Phys.*, 9(1):29–32, 2012.
- Pittman, T. B., Jacobs, B. C., and Franson, J. D. Demonstration of quantum error correction using linear optics. *Phys. Rev. A*, 71(5):052332, 2005.
- Pop, I. M., Geerlings, K., Catelani, G., Schoelkopf, R. J., Glazman, L. I., and Devoret, M. H. Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles. *Nature*, 508(7496):369–372, 2014.
- Pozar, D. M. *Microwave Engineering*. John Wiley & Sons, Inc., second edition, 1998.

- Preskill, J. Battling decoherence: the fault-tolerant quantum computer. *Phys. Today*, 1999.
- Purcell, E. M. Spontaneous emission probabilities at radio frequencies. *Phys. Rev.*, 69:681, 1946.
- Raimond, J. M., Brune, M., and Haroche, S. Manipulating quantum entanglement with atoms and photons in a cavity. *Rev. Mod. Phys.*, 73:565–582, 2001.
- Ralph, T. C. Continuous variable quantum cryptography. *Phys. Rev. A*, 61(1):010303, 1999.
- Ralph, T. C. Quantum error correction of continuous-variable states against Gaussian noise. *Phys. Rev. A*, 84(2):022339, 2011.
- Reagor, M. J., Paik, H., Catelani, G., Sun, L., Axline, C., Holland, E. T., Pop, I. M., Masluk, N. A., Brecht, T., Frunzio, L., Devoret, M. H., Glazman, L. I., and Schoelkopf, R. J. Ten milliseconds for aluminum cavities in the quantum regime. *Appl. Phys. Lett.*, 102(19), 2013.
- Reagor, M. J., Pfaff, W., Heeres, R. W., Ofek, N., Chou, K., Blumoff, J. Z., Leghtas, Z., Touzard, S., Sliwa, K. M., Holland, E. T., Albert, V. V., Frunzio, L., Devoret, M. H., Jiang, L., and Schoelkopf, R. J. Engineering non-linear resonator mode interactions in circuit qed by continuous driving: Manipulation of a photonic quantum memory. Abstract BAPS.2015.MAR.Y39.11 presented at APS March Meeting 2015, 2015.
- Reagor, M. J., Pfaff, W., Axline, C., Heeres, R. W., Ofek, N., Sliwa, K. M., Holland, E. T., Wang, C., Blumoff, J. Z., Chou, K., Hatridge, M. J., Frunzio, L., Devoret, M. H., Jiang, L., and Schoelkopf, R. J. Quantum memory with millisecond coherence in circuit QED. *Phys. Rev. B*, 94:014506, 2016.
- Reed, M. D., DiCarlo, L., Nigg, S. E., Sun, L., Frunzio, L., Girvin, S. M., and Schoelkopf, R. J. Realization of three-qubit quantum error correction with superconducting circuits. *Nature*, 482(7385):382–385, 2012.
- Ristè, D., Bultink, C. C., Lehnert, K. W., and DiCarlo, L. Feedback control of a solid-state qubit using high-fidelity projective measurement. *Phys. Rev. Lett.*, 109:240502, 2012.
- Ristè, D., Dukalski, M., Watson, C. A., de Lange, G., Tiggelman, M. J., Blanter, Y. M., Lehnert, K. W., Schouten, R. N., and DiCarlo, L. Deterministic entanglement of superconducting qubits by parity measurement and feedback. *Nature*, 502(7471):350–354, 2013.
- Ristè, D., Poletto, S., Huang, M. Z., Bruno, A., Vesterinen, V., Saira, O. P., and DiCarlo, L. Detecting bit-flip errors in a logical qubit using stabilizer measurements. *Nat. Comm.*, 6:6983, 2015.
- Robledo, L., Childress, L., Bernien, H., Hensen, B., Alkemade, P. F. A., and Hanson, R. High-fidelity projective read-out of a solid-state spin quantum register. *Nature*, 477(7366):574–578, 2011.
- Rowe, M. A., Kielpinski, D., Meyer, V., Sackett, C. A., Itano, W. M., Monroe, C., and Wineland, D. J. Experimental violation of a Bell’s inequality with efficient detection. *Nature*, 409(6822):791–794, 2001.

- Sauter, T., Neuhauser, W., Blatt, R., and Toschek, P. E. Observation of quantum jumps. *Phys. Rev. Lett.*, 57:1696–1698, 1986.
- Sayrin, C., Dotsenko, I., Zhou, X., Peaudecerf, B., Rybarczyk, T., Gleyzes, S., Rouchon, P., Mirrahimi, M., Amini, H., Brune, M., Raimond, J. M., and Haroche, S. Real-time quantum feedback prepares and stabilizes photon number states. *Nature*, 477(7362):73–77, 2011.
- Schindler, P., Barreiro, J. T., Monz, T., Nebendahl, V., Nigg, D., Chwalla, M., Hennrich, M., and Blatt, R. Experimental Repetitive Quantum Error Correction. *Science*, 332(6033):1059–1061, 2011.
- Schmitt, O. H. A thermionic trigger. *Jour. of Sci. Inst.*, 15(1), 1938.
- Schoelkopf, R. J. and Girvin, S. M. Wiring up quantum systems. *Nature*, 451(7179):664–669, 2008.
- Schreier, J. A., Houck, A. A., Koch, J., Schuster, D. I., Johnson, B. R., Chow, J. M., Gambetta, J. M., Majer, J., Frunzio, L., Devoret, M. H., Girvin, S. M., and Schoelkopf, R. J. Suppressing charge noise decoherence in superconducting charge qubits. *Phys. Rev. B*, 77(18):180502, 2008.
- Schuster, D. I., Houck, A. A., Schreier, J. A., Wallraff, A., Gambetta, J. M., Blais, A., Frunzio, L., Majer, J., Johnson, B. R., Devoret, M. H., Girvin, S. M., and Schoelkopf, R. J. Resolving photon number states in a superconducting circuit. *Nature*, 445(7127):515–518, 2007.
- Sears, A. P., Petrenko, A., Catelani, G., Sun, L., Paik, H., Kirchmair, G., Frunzio, L., Glazman, L. I., Girvin, S. M., and Schoelkopf, R. J. Photon shot noise dephasing in the strong-dispersive limit of circuit QED. *Phys. Rev. B*, 86(18):180504, 2012.
- Shor, P. W. Scheme for reducing decoherence in quantum computer memory. *Phys. Rev. A*, 52:R2493–R2496, 1995.
- Shor, P. W. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM Jour. Comp.*, 26(5):1484–1509, 1997.
- Sliwa, K. M., Hatridge, M. J., Narla, A., Shankar, S., Frunzio, L., Schoelkopf, R. J., and Devoret, M. H. Reconfigurable josephson circulator/directional amplifier. *Phys. Rev. X*, 5:041020, 2015.
- Smolin, J. A., Gambetta, J. M., and Smith, G. Efficient method for computing the maximum-likelihood quantum state from measurements with additive Gaussian noise. *Phys. Rev. Lett.*, 108(7):070502, 2012.
- Steane, A. M. Error correcting codes in quantum theory. *Phys. Rev. Lett.*, 77:793–797, 1996.
- Sun, L., Petrenko, A., Leghtas, Z., Vlastakis, B., Kirchmair, G., Sliwa, K. M., Narla, A., Hatridge, M. J., Shankar, S., Blumoff, J. Z., Frunzio, L., Mirrahimi, M., Devoret, M. H., and Schoelkopf, R. J. Tracking Photon Jumps with Repeated Quantum Non-Demolition Parity Measurements. *Nature*, 511(7510):444–448, 2013.

- Taminiau, T. H., Cramer, J., van der Sar, T., Dobrovitski, V. V., and Hanson, R. Universal control and error correction in multi-qubit spin registers in diamond. *Nat. Nano.*, 9:171–176, 2014.
- Tinkham, M. *Introduction to Superconductivity: Second Edition (Dover Books on Physics) (Vol i)*. Dover Publications, second edition, 2004.
- van Enk, S. J., Lütkenhaus, N., and Kimble, H. J. Experimental procedures for entanglement verification. *Phys. Rev. A*, 75:052318, 2007.
- Vijay, R., Devoret, M. H., and Siddiqi, I. The Josephson bifurcation amplifier. *Rev. Sci. Inst.*, 80(11):111101, 2009.
- Vijay, R., Slichter, D. H., and Siddiqi, I. Observation of quantum jumps in a superconducting artificial atom. *Phys. Rev. Lett.*, 106:110502, 2011.
- Vijay, R., Macklin, C., Slichter, D. H., Weber, S. J., Murch, K. W., Naik, R., Korotkov, A. N., and Siddiqi, I. Stabilizing Rabi oscillations in a superconducting qubit using quantum feedback. *Nature*, 490(7418):77–80, 2012.
- Vlastakis, B., Kirchmair, G., Leghtas, Z., Nigg, S. E., Frunzio, L., Girvin, S. M., Mirrahimi, M., Devoret, M. H., and Schoelkopf, R. J. Deterministically Encoding Quantum Information Using 100-Photon Schrödinger Cat States. *Science*, 342(6158):607–610, 2013.
- Vlastakis, B., Petrenko, A., Ofek, N., Sun, L., Leghtas, Z., Sliwa, K. M., Liu, Y., Hatridge, M. J., Blumoff, J. Z., Frunzio, L., Mirrahimi, M., Jiang, L., Devoret, M. H., and Schoelkopf, R. J. Characterizing entanglement of an artificial atom and a cavity cat state with Bell’s inequality. *Nat. Comm.*, 6:8970, 2015.
- Waldherr, G., Wang, Y., Zaiser, S., Jamali, M., Schulte-Herbrüggen, T., Abe, H., Ohshima, T., Isoya, J., Du, J. F., Neumann, P., and Wrachtrup, J. Quantum error correction in a solid-state hybrid spin register. *Nature*, 506(7487):204–207, 2014.
- Wallraff, A., Schuster, D. I., Blais, A., Frunzio, L., Huang, R., Majer, J., Kumar, S., Girvin, S. M., and Schoelkopf, R. J. Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature*, 431(7005):162–167, 2004.
- Wallraff, A., Schuster, D. I., Blais, A., Frunzio, L., Majer, J., Devoret, M. H., Girvin, S. M., and Schoelkopf, R. J. Approaching Unit Visibility for Control of a Superconducting Qubit with Dispersive Readout. *Phys. Rev. Lett.*, 95(6):060501, 2005.
- Wang, C., Axline, C., Gao, Y. Y., Brecht, T., Chu, Y., Frunzio, L., Devoret, M. H., and Schoelkopf, R. J. Surface participation and dielectric loss in superconducting qubits. *Appl. Phys. Lett.*, 107(16), 2015.

Wang, C., Gao, Y. Y., Reinhold, P., Heeres, R. W., Ofek, N., Chou, K., Axline, C., Reagor, M. J., Blumoff, J. Z., Sliwa, K. M., Frunzio, L., Girvin, S. M., Jiang, L., Mirrahimi, M., Devoret, M. H., and Schoelkopf, R. J. A schrödinger cat living in two boxes. *Science*, 352(6289):1087–1091, 2016.

Wootters, W. K. and Zurek, W. H. A single quantum cannot be cloned. *Nature*, 299(5886):802–803, 1982.