

UNIVERSITY of CALIFORNIA
Santa Barbara

Superconducting Qubits: Dephasing and Quantum Chemistry

A dissertation submitted in partial satisfaction of the
requirements for the degree of

Doctor of Philosophy

in

Physics

by

Peter James Joyce O'Malley

Committee in charge:

Professor John Martinis, Chair

Professor David Weld

Professor Chetan Nayak

June 2016

The dissertation of Peter James Joyce O'Malley is approved:

Professor David Weld

Professor Chetan Nayak

Professor John Martinis, Chair

June 2016

Copyright © 2016
by Peter James Joyce O'Malley

Any work that aims to further human knowledge is inherently dedicated
to future generations.

There is one particular member of the next generation to which I
dedicate this particular work.

Acknowledgements

It is a truth universally acknowledged that a dissertation is not the work of a single person.

Without John Martinis, of course, this work would not exist in any form. I will be eternally indebted to him for ideas, guidance, resources, and—perhaps most importantly— assembling a truly great group of people to surround myself with.

To these people I must extend my gratitude, insufficient though it may be; thank you for helping me as I ventured away from superconducting qubits and welcoming me back as I returned. While the nature of a university research group is to always be in flux, this group is lucky enough to have the possibility to continue to work together to build something great, and perhaps an order of magnitude luckier that we should wish to remain so. It has been an honor.

Also indispensable on this journey have been all the members of the physics department who have provided the support I needed (and PCS, I apologize for repeatedly ending up, somehow, on your naughty list). My friends have allowed me that rare treasure for a graduate student: relaxation. I would never have gotten off the ground without B, C, & C O'Malley to push me to my best. And I do not know where I would be without Anna and Iris to come home to each day.

Curriculum Vitæ

Peter James Joyce O'Malley

Education

- 2016 Ph.D., Physics, University of California, Santa Barbara
- 2008 B.S., Physics and Astronomy, Haverford College, Haverford
- 2004 Bishop O'Dowd High School, Oakland

First author publications

“Qubit Metrology of Ultralow Phase Noise Using Randomized Benchmarking“, P. J. J. O'Malley, J. Kelly, R. Barends, et al., *Physical Review Applied*, 3(4), 044009 (2015).

“Scalable Quantum Simulation of Molecular Energies“, P. J. J. O'Malley, R. Babbush, et al., *submitted* (2016).

Abstract

Superconducting Qubits: Dephasing and Quantum Chemistry

by

Peter James Joyce O'Malley

One of the most exciting potential applications of a quantum computer is the ability to efficiently simulate quantum systems, a task that is out of the reach of even the largest classical supercomputers. Such simulations require a quantum algorithm capable of efficiently representing and manipulating a quantum system, as well as a device with sufficient coherence to execute it. In this work, we describe experiments advancing both of these goals. First, we discuss dephasing—currently a leading cause of decoherence in superconducting qubits—and present measurements accurately quantifying both low- and high-frequency phase noise sources. We then discuss two quantum algorithms for the simulation of chemical Hamiltonians, and experimentally contrast their performance. These results show that with continuing improvement in quantum devices we may soon be able to apply quantum computers to practical chemistry problems.

Contents

1	Introduction	1
1.1	Quantum computation: motivation and applications	1
1.1.1	Errors and error correction	2
1.1.2	Potential applications of quantum computing	4
1.2	Superconducting qubits	5
1.2.1	History and overview	5
1.2.2	Advantages and disadvantages	6
2	Dephasing	9
2.1	A Bloch vector picture of dephasing	9
2.2	Frequency noise and the superconducting qubit	10
2.2.1	Ramsey measurements	11
2.2.2	Spin echo measurements	13
2.2.3	Rabi measurements	14
2.3	Forms of frequency noise	14
2.3.1	White noise	15
2.3.2	1/f noise	16
2.3.3	Telegraph noise	17
2.4	Sources of noise	17
2.4.1	Flux noise	18
2.4.2	Charge noise and quasiparticle tunneling	19
2.4.3	Critical current noise	20
2.4.4	Capacitance noise	20
2.4.5	Resonator induced dephasing	20
2.5	Conclusion	21
3	The Ramsey Tomography Oscilloscope	23
3.1	Introduction of the Ramsey Tomography Oscilloscope	24
3.2	Ramsey Tomography Oscilloscope results in transmon-type qubits	26
3.2.1	Flux noise in Xmons	26
3.2.2	Flux noise in gmons	27
3.2.3	Flux noise in early SiXmons	28

3.2.4	Comparison with theory	29
4	Qubit metrology of ultralow phase noise using randomized benchmarking	33
4.1	Introduction	34
4.2	RB Ramsey	37
4.3	Measuring telegraph noise	39
4.4	Measuring error from coherent qubit-qubit interactions	42
4.5	Measuring different gate implementations	44
4.6	Summary	45
5	Chemistry on a quantum computer: a brief introduction	47
5.1	Representing electrons with qubits: The Bravyi-Kitaev Transform	48
5.2	The canonical quantum chemistry algorithm	51
5.3	The Variational Quantum Eigensolver	55
6	Scalable Quantum Simulation of Molecular Energies	57
6.1	Variational quantum eigensolver	60
6.2	Phase estimation algorithm	64
6.3	Experimental Methods	67
6.4	Conclusion	68
A	Perturbative treatment of the Xmon Hamiltonian	71
B	Spectral density of phase noise	75
C	Ramsey Tomography Oscilloscope data processing	79
D	Appendices for Chapter 4	81
D.1	Theoretical relation of RB error to $\langle \phi^2 \rangle$	81
D.2	Types of phase noise	84
D.3	T_1 , Ramsey, and spin echo fits	86
D.4	Flux noise	87
D.5	RB Ramsey across the qubit spectrum	90
D.6	Charge noise	91
D.7	Calculation of Ω_{ZZ}	91
D.8	Fits to gate errors in Figure 4.4	93
D.9	Telegraph noise measured in other devices	94
E	Appendices for Chapter 5	97
E.1	The electronic structure problem	97
E.2	Experimental methods for VQE	100
E.3	Experimental methods for PEA	103
E.4	Unitary coupled cluster	106

Chapter 1

Introduction

1.1 Quantum computation: motivation and applications

The continued increase in computing power over the past decades has been an unprecedented boon to society in general, and scientific research in particular. From early analog simulations of nuclear physics in the Manhattan project to modern attempts to model the human brain [67] or the entire universe [129], nearly all areas of research make use of computing power in one way or another. Computers can solve problems that are easy to understand, yet difficult to execute—such as computing electromagnetic fields at all points in a space, or searching a massive database—as well as problems for which the solution does not have an intuitive explanation—such as neural networks capable of outperforming the best humans in the game of Go [113].

However, there are problems that will take even the most powerful supercomputer longer than the age of the universe to solve. Frustratingly, such a problem is presented by the scientific theory that underpins all of physics: quantum mechanics. The fact that particles become entangled—meaning their state cannot be described individually, but

that the quantum state must describe the system as a whole—and that the superposition of such states are equally valid themselves, means that increasing the size of a quantum system increases the resources required to simulate it exponentially. Concretely, if we can simulate a system of n electrons, then simulating a system of $n + 1$ electrons will be twice as hard. Even if computing power continues to double every few years, it will take many, many such doublings to significantly increase the size of a quantum system we can simulate. For this reason, we say that an algorithm that solves a problem (for example, simulating a physical system) is “efficient” or “scalable” if the resources it requires scale polynomially with the size of the problem.

However, in 1982 Richard Feynman proposed [38] a solution to this: use the very quantum mechanical systems that are so hard to simulate as a simulator themselves. Since then, the idea of a universal quantum computer made of quantum bits (“qubits”) has been developed, and many algorithms for such a computer have been shown to outperform their classical counterparts. While it is unknown whether an arbitrary physical system can be efficiently simulated with a universal quantum computer, the range of problems that can be efficiently solved is important enough to merit decades of work (so far!) to its realization.

1.1.1 Errors and error correction

As with so many things, though, proposing a universal quantum computer is vastly easier than building one. Noise is present in all physical systems, and it is particularly difficult for a quantum computer to tolerate. Consider a classical bit defined by the voltage in a wire: 0 V for 0, and 5 V for 1. If noise causes the wire’s voltage to fluctuate by even a volt, we can still use it as a bit by simply considering anything less than 2.5 V to be 0; all of the microscopic states of the wire—i.e. electron configurations, and so on—that result in the macroscopic property of voltage being less than 2.5 V are valid.

This leniency is lost on qubits. A qubit can be a 1 or a 0 or any superposition thereof: we represent a qubit's state as $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where α and β are complex numbers with the restriction that $|\alpha|^2 + |\beta|^2 = 1$. If the value of β changes by some amount ϵ , it is now a different computational state; for a qubit each microscopic state is unique.

It is useful to write the qubit state alternatively as $|\psi\rangle = \cos(\theta/2)|0\rangle + e^{i\phi} \sin(\theta/2)|1\rangle$, where $0 < \theta < \pi$ and $0 < \phi < 2\pi$. This is the Bloch sphere representation, where the qubit state is represented by a vector on the unit sphere, with θ as the angle to the Z-axis and ϕ the angle in the X-Y plane. The north pole is $|0\rangle$ ($\theta = 0$) and the south pole is $|1\rangle$ ($\theta = \pi$), with points on the equator representing an equal superposition of $|0\rangle$ and $|1\rangle$. When the qubit is in $|0\rangle$ or $|1\rangle$, ϕ —called the phase—is undefined. We can then consider two different types of noise: decoherence, or energy relaxation, in which the qubit transitions from $|1\rangle$ to $|0\rangle$ ¹; and dephasing, where the phase ϕ is blurred. Decoherence is governed by a timescale known as T_1 , the relaxation rate, and the process is described by a simple exponential. The dephasing timescale is sometimes called T_2 , but dephasing processes are more complicated, and discussed starting in Chapter 2.

Both decoherence and dephasing cause errors in the execution of quantum algorithms. As the number of operations required to perform a useful calculation is on the order of 10^{20} , we have a rough estimate for a necessary error rate of 10^{-20} per operation. Decades of research in experimental quantum computation has resulted in record error rates of 10^{-6} for single qubit operations [46] and 10^{-3} [16, 99, 11, 9] for two qubit operations. It therefore seems that more than just “building better qubits” will be necessary to realize a quantum computer. Fortunately, quantum error-correcting codes have been developed for just this [112, 61, 39]. The idea is to use multiple physical, error-prone qubits to encode one logical, error-free qubit. In a sense, this provides some of the leniency of the

¹By convention, the higher energy state of the qubit is denoted $|1\rangle$, so decoherence is the physical relaxation of the qubit to its ground state; qubit realizations where the two states have the same energy, such as topological qubits, are thus said to be immune to decoherence.

classical bit to the qubit, as now we have multiple “microscopic” states (states of the physical qubits) that represent a valid “macroscopic” state (the logical qubit’s state).

Error-correcting codes allow for more error-prone qubits to be used at the cost of requiring greater numbers of them, given that the qubit’s base error rate is below some threshold value which depends on the particular code. While superconducting qubits have recently reached the threshold for the surface code scheme [39] in isolated benchmarks [95, 11, 94] and implemented partial versions of it [58, 97, 30, 87], this does not mean that our work is done. First, adding ever more qubits to a device seems likely to increase error rates to some degree, due to issues of crosstalk, fabrication complications, and so on; and second, pushing the error rate further below threshold allows for using fewer physical qubits to realize a single logical qubit. As the number of (logical) qubits necessary to perform useful computations is ranges from dozens to thousands, this is a necessity. The first half of this thesis describes efforts to quantify and understand dephasing, currently the leading source of error in the Xmon superconducting qubit.

1.1.2 Potential applications of quantum computing

In addition to simulating other quantum systems, several quantum algorithms have been proposed with the potential to efficiently solve classically intractable problems. The most well known of these is Shor’s algorithm for prime factorization [111]. The best known classical algorithm for prime factorization is the number field sieve, which scales sub-exponentially in the number of bits of the integer to be factored; the difficulty of this problem has motivated the development of widely-used public-key cryptography systems based on prime factorization. Shor’s algorithm, by contrast, runs in polynomial time on a quantum computer, meaning that these cryptography systems could be efficiently broken with a large enough quantum computer. As modern cryptography systems typically use thousands of bits for their keys, a quantum computer would require thousands of (fully

error-corrected) qubits to break them, meaning that this potential security hole is still many years in the future².

Another important application for quantum computers is in the field of machine learning. Many machine learning problems are essentially very hard optimization problems, and it has been recently shown that quantum computers can potentially speed up the solution of such problems [84, 32]³.

One of the most scientifically interesting applications of a quantum computer is that for which it was originally proposed: quantum simulation. This application is of particular interest to the field of quantum chemistry, where despite the inherent intractability of the problem of simulating many electrons, significant progress has been made since the advent of quantum mechanics with the creation of various approximation methods bringing ever-larger molecules within reach of classical computation. In recent years, however, there has also been rapid progress in the development of quantum algorithms for chemistry [7], such that a quantum computer with merely dozens of qubits would allow the simulation of molecules impractical for study with classical computers; this has led to quantum chemistry being called the “killer app” for a quantum computer [123]. The second half of this thesis describes the implementation of two quantum algorithms for theoretical chemistry.

1.2 Superconducting qubits

1.2.1 History and overview

The implementation of a quantum computer requires a quantum system to serve as a qubit. In principle, any two-level system (or individually addressable two-level subspace

²Though this algorithm is still considered by many in the field to be the primary motivation for significant investments in quantum computation research by governments worldwide.

³By contrast, this application has motivated industrial investment in quantum computation research.

of a many-level system) may be used as a qubit, but as a practical matter, in 2000 DiVincenzo proposed a set of criteria necessary for a physical system to serve as a quantum computer [34]. There are many candidate systems with the potential to fulfill these criteria: photons controlled with linear optical elements, nuclear spins addressed through nuclear magnetic resonance techniques, electronic spins in quantum dots, neutral atoms confined in optical lattices, trapped ions, nitrogen-vacancy centers in diamond, topological quasiparticles braided in two dimensions, superconducting circuits based on Josephson junctions, and more. All of these systems must navigate a tradeoff between coherence and control. The more isolated a qubit is from its environment, the more resistant it is to unwanted noise, but the harder it is to precisely control. One consequence of this is that qubits that may perform extremely well on their own may suffer a drastic reduction in coherence times when coupled together.

This thesis focuses on superconducting qubits, which reside on the “better control” side of the balance. These qubits use superconducting circuits as quantum LC oscillators with flux ($\hat{\Phi}$) and charge (\hat{Q}) as the conjugate variables. The qubit’s $|0\rangle$ and $|1\rangle$ states are the ground and first excited state of the oscillator. To make this transition uniquely accessible (that is, to be able to ignore the higher levels of the oscillator) we require nonlinearity in the oscillator; using a Josephson junction for the inductor provides this. Using a pair of Josephson junctions as a superconducting quantum interference device (SQUID) allows the inductance to be tuned, and while this can be very experimentally useful, it is not required, and many superconducting qubit designs are non-tuneable.

1.2.2 Advantages and disadvantages

Superconducting qubits are attractive platforms for a quantum computer for several reasons. Their fabrication relies on well-known microfabrication techniques, making it straightforward to build more of them by simply adding additional qubits to the design.

They are controlled with microwave electronics, which again have been developed for other industrial applications. Their ease of fabrication and control also makes them natural candidates for the surface code error correction scheme.

With their strong performance on the “control” side, one might expect that superconducting qubits are lacking on the “coherence” end. Indeed, their coherence times of up to about a hundred microseconds are less than other systems, which have seen record coherence times of greater than a second. This is ameliorated by the speed with which operations are possible: the error rate from decoherence is governed by the ratio of operation time to coherence time, which is approaching 10^{-4} . However, it seems likely that further materials research will be necessary to improve coherence times. Finally, superconducting qubits also require dilution refrigerators to maintain temperatures low enough to remove thermal decoherence processes. The capacity of such cryostats is not currently a limiting factor, but advances will need to be made before we have superconducting quantum computers with thousands or millions of qubits.

The qubit used in the experiments described in this thesis is the Xmon [10]. This is a variant of the transmon qubit, proposed in 2007 [63] with the aim of reducing susceptibility to charge noise while maintaining sufficient nonlinearity for operation as a two-level system. In the last few years, transmon qubits have seen drastic improvements in operation fidelities [95, 11, 94], reductions in noise mechanisms [104], and even demonstrations of initial error correction algorithms [58, 97, 30, 87]. However, we are not yet at the level where we can simply place hundreds of qubits on a device and have a quantum computer, so we now turn to analysis of the leading cause of error in the Xmon: dephasing.

Chapter 2

Dephasing




Figure 2.1: Dephasing represented on the Bloch sphere. For an ensemble of qubit states (a) starting on the equator, over time (b) each element rotates at a slightly different frequency, causing the ensemble to spread. (In this picture, we rotate the frame at $f_{10}/2\pi$, so a state without frequency noise remains fixed.) (c) The averaged Bloch vector shrinks in from the equator.

2.1 A Bloch vector picture of dephasing

Dephasing simply means the loss of phase coherence. However, as the phase is undefined for a qubit in either of the pure basis states (i.e. $|0\rangle$ and $|1\rangle$), techniques for measuring and dealing with dephasing are more complicated than those for energy decoherence. In

the Bloch vector picture of the qubit, the qubit state rotates about the Z-axis at the qubit frequency, f_{10} ; that is, the qubit phase ϕ advances in time according to $\phi(t) = 2\pi f_{10}t$. A frequency offset of δf for a period t therefore creates a phase offset $\delta\phi = \delta f t$. (Of course, if we had a constant frequency offset we could just change or re-measure our bare frequency and eliminate the phase offset.) Thus frequency noise ($\delta f(t)$) over time produces phase noise ($\delta\phi$). Over an ensemble of experiments, frequency noise results in the Bloch vectors of each experiment spreading out from the average as some rotate with greater frequency and some with lower. Therefore, dephasing is said to shrink the (averaged) Bloch vector (see Figure 2.1). Usually, we are interested in the variance of the phase, $\langle\phi^2(t)\rangle$. In general use, “frequency noise”, “phase noise”, and “dephasing” are used interchangeably¹.

2.2 Frequency noise and the superconducting qubit

As the frequency of a superconducting qubit is set by device parameters ($f_{10} \propto 1/\sqrt{LC}$), it might be expected that frequency noise presents a particular problem if these parameters are variable. For tunable superconducting qubits, such as the Xmon, this is often the case. In this section, we begin with a discussion of basic methods for measuring dephasing, followed by an overview of different types of frequency noise, and then consider the microscopic source of these types of noise.

Ultimately, in order to measure phase noise, the qubit must be sensitive to it. Therefore, all such measurements involve preparing the qubit in a superposition state (on the equator in the Bloch sphere representation), acquiring a signal by allowing the qubit to dephase, and then measuring. This can be contrasted with an energy coherence (T_1) measurement, for example, where the qubit is prepared in the $|1\rangle$ state, which is insensitive

¹It is our opinion that “dephasing” most properly should refer to the variance $\langle\phi^2(t)\rangle$, but in the literature it is not usually used to refer to any specific quantity.

to frequency noise. Although there are several ways of classifying these measurements, one useful one is to divide them experiments where the qubit is allowed to freely evolve during the acquisition period (“free evolution”) and experiments where it is driven during that time (“driven evolution”).

2.2.1 Ramsey measurements

The most basic free evolution dephasing measurement is the Ramsey fringe experiment, first introduced by Norman Ramsey in 1950 in the context of atomic spectroscopy [93]. The qubit is prepared in a superposition state with a $\pi/2$ pulse, allowed to idle for some time t , and then rotated back with another $\pi/2$ pulse and measured. Absent any dephasing (or other source of error), the probability of measuring $|1\rangle$ varies sinusoidally in time, from 1 to 0, with the same frequency as the qubit. This is most easily understood with the Bloch sphere representation: the initial $\pi/2$ pulse rotates the Bloch vector by $\pi/2$ about the X axis, and during the evolution it precesses at the qubit frequency about the Z axis. The final $\pi/2$ rotation about the X axis and measurement then projects the Y coordinate into the measurement basis (see Figure 2.2). In the presence of dephasing, however, each iteration of the experiment will have a slightly different frequency. Considering the experiment as an ensemble of Bloch vectors, prior to the final $\pi/2$ rotation this ensemble will be spread out around the mean value of the frequency. Therefore, the averaged Bloch vector is reduced in length, reducing the amplitude of the measured sinusoidal signal. For longer t , the spread in the ensemble is increased, until finally all signal is lost and the measured probability remains at 0.5. It is thus the change in the amplitude of the signal—the envelope—that contains information about dephasing.

As it is only the envelope we are interested in, a slight modification to the experimental procedure can greatly aid in data analysis. For each iteration, we perform four different sequences, with the final $\pi/2$ pulse phase shifted by 0 , $\pi/2$, π , and $3\pi/2$ radians, changing

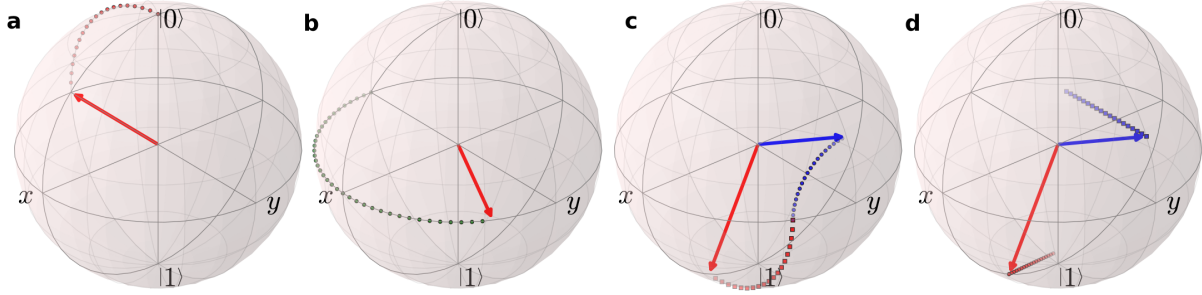


Figure 2.2: Bloch sphere diagram of a Ramsey experiment. (Unlike Figure 2.1, we here show only one experiment, and are not in the qubit’s rotating frame.) (a) The initial $\pi/2$ pulse puts the qubit on the equator. (b) The qubit evolves for some time t . (c) The recovery $\pi/2$ pulse rotates about the X axis (red) or about the Y axis (blue), so that (d) the measurement projects the Y (red) or X (coordinate) to the probability of $|1\rangle$.

the axis of rotation (that is, the final pulse is rotated between $X/2$, $Y/2$, $-X/2$, and $-Y/2$). The envelope, sometimes called the visibility, V , can then be calculated directly as

$$V = \sqrt{(P_{X/2} - P_{-X/2})^2 + (P_{Y/2} - P_{-Y/2})^2}, \quad (2.1)$$

where P_g is the measured probability for the experiment with the final gate g .

In fact, the Ramsey envelope visibility is a direct measurement of dephasing. For a free-evolution time t , the visibility is given by (see Appendix B)

$$V(t) = A \exp \left[-\frac{1}{2} \langle \phi^2(t) \rangle \right] + B, \quad (2.2)$$

where the fit parameters A (initial visibility) and B (final population) reflect errors in state preparation and measurement²; for details on $\langle \phi^2(t) \rangle$, see Section 2.3. We see that the Ramsey experiment is a fairly straightforward way to measure the dephasing $\langle \phi^2(t) \rangle$ as a function of time.

²In our devices, A is dominated by readout error and B by thermal population.

2.2.2 Spin echo measurements

The spin echo sequence (introduced by Hahn in 1950 [45]) is a modification to the Ramsey experiment designed to lengthen phase coherence. A π pulse is inserted at the midpoint of the free-evolution time to counteract the effect of offsets to the qubit frequency. Considering again the ensemble of Bloch vectors, if each vector precesses at a slightly different rate due to frequency noise, the spread of the vectors will be reversed by the π pulse, and at the end of the experiment the averaged Bloch vector will be “refocused”³.

The echoing pulse will only fully counteract frequency noise when the sequences in the ensemble have different but static precession frequencies. In the context of a single qubit experiment, this means that each run of the sequence has a constant frequency, but the frequency can change between sequences; in other words, spin echo is effective for low-frequency noise (specifically, lower than the inverse of the sequence timescale). In the context of a true ensemble of systems—magnetic moments in an NMR experiment, for example—if each system has a slight frequency offset due to local effects, spin echo techniques become a necessity, as the Ramsey envelope decays exceedingly quickly⁴.

The visibility of a spin echo envelope is analogous to Eq. (2.2):

$$V(t) = A \exp \left[-\frac{1}{2} \langle \tilde{\phi}^2(t) \rangle \right] + B, \quad (2.3)$$

where $\langle \tilde{\phi}^2(t) \rangle$ indicates that the dephasing is modified by the presence of an echo pulse. For details, see Appendix B and Section 2.3.

Furthermore, it can be desirable to add more than one echo pulse to suppress noise at a greater range of timescales. The qubit becomes sensitive to noise mainly at (or above) the frequency given by the separation between the pulses: $f_e = N/2t$, where N

³The refocusing only occurs at $t = 2t_E$, where t_E is the time of the echo pulse. That is, if you were to scan the measurement time while fixing t_E , you would find a spike in $t = 2t_E$. This spike is called the “echo”, hence the name “spin echo”.

⁴For this reason, it is common to see the coherence time measured by spin echo called T_2 and that measured by Ramsey T_2^* .

is the number of pulses and t is the sequence duration. There are different schemes for adding more pulses to the sequence; two common ones are known as Uhrig Dynamical Decoupling [124] (UDD) and Car-Purcell-Meiboom-Gill [24, 81] (CPMG) sequences. The precise sensitivity to phase noise is given by a spectral weight function, which is detailed in Appendix B.

2.2.3 Rabi measurements

Another method for measuring dephasing is with driven evolution experiments, the prototypical example of which is the Rabi sequence. In a Rabi experiment, the qubit is driven with by a pulse at the transition frequency f_{10} , causing oscillation between $|0\rangle$ and $|1\rangle$ at the Rabi frequency f_R , which is proportional to the amplitude of the driving pulse. Like spin echo sequences, the Rabi experiment is insensitive to low-frequency noise; it is primarily measures noise at f_R . The spectral weight function for Rabi measurements is given in Appendix B. A full treatment of Rabi oscillations is given in many quantum mechanics textbooks; see, for example [100].

As both Rabi and CPMG/UDD measurements are sensitive to noise at a narrow range of frequencies, they can be used for noise spectroscopy. However, the analysis is somewhat complicated by the fact that the noise spectrum is essentially convolved with the various spectral weight functions, and they are also sensitive to different noise channels. Nevertheless, noise spectroscopy over a large frequency range has been carried out with these methods [23, 138, 139, 141].

2.3 Forms of frequency noise

Now that we have a few basic ways of measuring the dephasing, $\langle\phi^2(t)\rangle$, we consider the different types of dephasing and their functional form. It is useful to consider the spectral

density $S_f(f)$ of the qubit frequency fluctuation; the mean square phase noise is related to the spectral density by (see Appendix B)

$$\langle \phi^2(t) \rangle = \int_0^\infty df S_f(f) \frac{\sin^2(\pi ft)}{(\pi f)^2}. \quad (2.4)$$

When the integrand is nonconvergent, the lower limit of integration is taken to be the inverse of the experiment duration (e.g. 1 hour⁻¹), and the upper limit is the qubit frequency, as noise power at or above that frequency drives state transitions rather than dephasing the qubit. Note that the units of S are power per bandwidth; that is, Hz² / Hz for S_f , the spectral density of frequency noise. Spectral density is often quoted in other units; for example, if the spectral density of current noise in the bias line is known, we convert from $S_{I_{\text{bias}}}$ (in A²/Hz): $S_f = (df/dI_{\text{bias}})^2 S_{I_{\text{bias}}}$.

2.3.1 White noise

The simplest form of noise is white (or uncorrelated) noise, where the noise is spectrally flat, with a constant spectral density $S_f(f) = S_0$. From Eq. (2.4), we have

$$\langle \phi^2(t) \rangle = \frac{S_0}{2} t \equiv 2 \frac{t}{T_{\phi 1}}, \quad (2.5)$$

where we here define $T_{\phi 1}$, the white noise dephasing time; we can also define a dephasing rate $\Gamma_{\phi 1} = 1/T_{\phi 1}$. For a Ramsey experiment with only white noise present, we can put Eq. (2.5) into Eq. (2.2) and get the envelope $V_{\text{white}}(t) = A \exp(t/T_{\phi 1}) + B$. (This is the reason for the factor of 2 in the definition of $T_{\phi 1}$.) The value of $T_{\phi 1}$ can be crudely thought of as “about how long you can use the qubit for before it loses phase coherence”; in a Ramsey experiment, for example, at $t = T_{\phi 1}$ the visibility will have decayed $1/e$ of its original value. Of course, depending on the coherence requirements of the experiment, the actual “use time” of the qubit may be much less.

T_1 decay as phase noise

Energy decay of a qubit (T_1) also manifests in dephasing measurements. The $|0\rangle$ state is insensitive to T_1 decay, while the $|1\rangle$ state decays at a rate $1/T_1$; the superposition state $(|0\rangle + |1\rangle)/\sqrt{2}$ therefore decays at a rate $1/2T_1$ ⁵. Even in the absence of dephasing, then, Ramsey and echo envelopes (and other dephasing measurements) will be limited by an exponential decay with time constant $2T_1$; in this situation the qubit dephasing is said to be “ T_1 limited”. Therefore, when measuring white noise dephasing, the effects of energy decay must be separated out; fortunately, this is straightforward as T_1 is easily measured in a separate experiment.

2.3.2 1/f noise

1/f noise (sometimes called pink noise or flicker noise) can be found in a number of physical and biological systems [92], but for present purposes it is well-known in electronics [15] (such as those driving the qubit), as well as Josephson junctions and SQUID loops [64, 134]. As the name suggests, the spectrum of 1/f noise is of the form $S_f(f) = S_{1/f}/f$. For a Ramsey experiment, again using Eq. (2.4), we find

$$\langle\phi^2(t)\rangle = S_{1/f}t^2 \ln \frac{0.4007}{f_c t}, \quad (2.6)$$

where f_c is the low-frequency cutoff. Typically, the inverse of the experiment’s duration is used for this. Because the logarithmic part varies slowly (not usually more than 10-20% for even large variations between experiments), it is commonly ignored, leaving

$$\langle\phi^2(t)\rangle \approx S_{1/f}^* t^2 = 2 \left(\frac{t}{T_{\phi 2}} \right)^2, \quad (2.7)$$

where here we define the correlated noise time constant $T_{\phi 2}$. We call this the “correlated noise” time constant because Eq. (2.7) is the same as the result for a noise source that

⁵This simple averaging is valid because energy decay is independent of the quantum phase.

is correlated over very long times, $S_f(f) = 2\sigma_{qb}^2\delta(f)$, where σ_{qb} is the standard deviation of the qubit frequency.

2.3.3 Telegraph noise

Telegraph noise (also called burst noise or a random telegraph signal) is seen when a qubit switches between two stable frequencies. The dephasing due to telegraph noise with switching timescale T_s and effective magnitude Δf_{10} is given by (see Appendix B)

$$\langle\phi^2(t)\rangle = (2\pi\Delta f_{10})^2 T_s [t - T_s(1 - e^{-t/T_s})], \quad (2.8)$$

where, for simplicity, we have assumed that the up and down switching rates are identical. For short timescales ($t \ll T_s$), telegraph noise looks like correlated noise; that is, $\langle\phi^2(t)\rangle \propto t^2$. Conversely, for long timescales ($t \gg T_s$) it is similar to white noise ($\langle\phi^2(t)\rangle \propto t$). This means that a system dominated by telegraph noise will display the opposite behavior of the common case where white noise dominates at short timescales and correlated (usually $1/f$) noise dominates at long timescales.

2.4 Sources of noise

We now discuss the physical sources of frequency noise in the qubit. First, to frame the discussion of noise sources, we state the equation for the frequency. The resonant frequency of an LC circuit is given by $f = 1/\sqrt{LC}$; by treating the Xmon as an anharmonic oscillator⁶, we have (see Appendix A)

$$f_{10} = \sqrt{8f_J f_C} \sqrt{\cos \frac{\pi\Phi}{\Phi_0}} - f_C \quad (2.9)$$

where now f_{10} is the frequency of only the $0 \rightarrow 1$ transition; f_J , and f_C are device parameters deriving from the inductance and capacitance, respectively; Φ_0 is the magnetic flux quantum; and Φ is the applied magnetic flux threading the SQUID loop. Noise in

f_{10} , therefore, can be considered as noise in one of these parameters.

2.4.1 Flux noise

Perhaps the most straightforward source of frequency noise comes from noise in the magnetic flux through the SQUID loop; the flux Φ of Eq. (2.9) is one of the experimental knobs, so it is natural to consider what happens when that knob shakes. Taking the derivative of Eq. (2.9), we see the flux sensitivity $\delta f_{10}/\delta\phi$ is

$$\frac{\delta f_{10}}{\delta\phi} = \frac{\pi}{2\Phi_0} \sqrt{8f_J f_C} \frac{\sin \frac{\pi\Phi}{\Phi_0}}{\sqrt{\cos \frac{\pi\Phi}{\Phi_0}}}. \quad (2.10)$$

From an experimentalist’s perspective, it is often useful to solve Eq. (2.9) for Φ and substitute this into Eq. (2.10) to get

$$\frac{\delta f_{10}}{\delta\phi} = -\frac{\pi}{2} f_p \sqrt{\left(\frac{f_p}{f_{10} + f_C}\right)^4 - 1}, \quad (2.11)$$

where we have defined $f_p \equiv \sqrt{8f_J f_C}$. We can then measure $f_{10}(\Phi)$ to fit Eq. (2.9), and immediately calculate the flux sensitivity given only the operating point f_{10} . In either case, we are now prepared to analyze noise in the flux Φ .

One potential source of flux noise is the electronics controlling the flux bias; the output noise of a signal generator, for example, will typically have a $1/f$ spectrum up to a white noise floor at high frequencies. As the electronics are easily separable from the qubit, the importance of this noise is relatively straightforward to compute. The noise spectrum can be measured with a spectrum analyzer, from which we can extract the magnitude of the $1/f$ noise (typically quoted as “noise at 1 Hz in $\text{nV}/\sqrt{\text{Hz}}$ ”) and the noise floor (in $\text{V}/\sqrt{\text{Hz}}$). Then we calculate the sensitivity $\delta f_{10}/\delta V = (\delta f_{10}/\delta\Phi)(\delta\Phi/\delta V)$, where we know $\delta\phi/\delta V$ from the wiring and inductance connecting the electronics to the

⁶The “ $-f_C$ ” is the modification resulting from treating the nonlinear nature of the Josephson inductance as a (first-order) perturbation to the simple LC oscillator; see Appendix A.

qubit. This allows us to use Eq. (2.6) or Eq. (2.7) (for the $1/f$ noise) and Eq. (2.5) (for the noise floor) together with Eq. (2.2) to predict the limit to a Ramsey experiment imposed by our electronics.

Typically, however, it is not the electronics that limit phase coherence. Flux noise in SQUIDs has been studied for three decades [134], and $1/f$ noise magnitudes of a few $\mu\Phi_0/\sqrt{\text{Hz}}$ (at 1 Hz) have been consistently found across a wide range of SQUID materials, geometries, and fabrication processes [140, 107, 17, 102, 1]. While there is no single theory to conclusively explain all the data, ongoing work points to a possible geometric understanding of SQUID flux noise (see Section 3.2.4).

Furthermore, flux noise is not necessarily $1/f$ -like. Telegraph noise can produce “bumps” or “plateaus” in Ramsey and spin echo sequences [41], and there is evidence at least in flux qubits for flux noise of this form [139]. However, there are also other sources of telegraph noise liable to confound the situation (see below).

2.4.2 Charge noise and quasiparticle tunneling

Charge noise—fluctuations in the gate (or offset) charge—and quasiparticles tunneling across the junction will both dephase Josephson junction-based qubits [73, 72, 63]; in terms of Eq. (2.9), this corresponds to noise in the f_C term (see Appendix A). The sensitivity of the qubit to these processes is governed by the charge dispersion $\delta E_{ij}/\delta n_g$, the shift in the energy spacing of between levels i and j caused by a change in the offset charge (measured in number of Cooper pairs n_g) of the device. As a transmon-based qubit, however, the Xmon is designed to have $\delta E_{ij}/\delta n_g$ exponentially suppressed in its operating regime; see [63] for details. Such frequency noise has been seen, manifesting as telegraph noise in qubits with a timescale of milliseconds [96].

2.4.3 Critical current noise

Noise in the Josephson junctions' critical current I_0 will also dephase the qubit. This noise is $1/f$ in nature, typically with amplitudes on the order of $10^{-6} I_0/\sqrt{\text{Hz}}$ at 1 Hz [126, 86]. The dephasing time caused by such variation is given by [63, 105]

$$T_{\phi 2, I_0} = \frac{2}{B f_{10}}, \quad (2.12)$$

where B is the dimensionless amplitude of the $1/f$ noise spectrum of critical current fluctuations at 1 Hz. This gives an expected $T_{\phi 2}$ in the range of tens of μs , which is not currently limiting qubit performance. However, most of the studies of critical current noise take place at temperatures well above typical qubit operating temperatures (i.e. 300 mK and above); noise is thought to decrease with decreasing temperatures. Given that we hope to soon see dephasing times in the tens of μs , the effects of critical current noise should be further studied.

2.4.4 Capacitance noise

Another potential source of noise is changes in the capacitance of the circuit (this would correspond to fluctuations in f_C in Eq. (2.9)). In fact, due to the tradeoffs made to suppress charge dispersion, the Xmon may be uniquely sensitive to such noise. However, there is as yet no experimental evidence that this poses a problem [63].

2.4.5 Resonator induced dephasing

The final potential source of dephasing we consider is dephasing caused by the coupled measurement resonator. As the act of measuring a quantum state causes its complete dephasing (it is projected to $|1\rangle$ or $|0\rangle$), any unintentional partial measurement will partially dephase the qubit. More precisely, the dispersive shift χ , the frequency shift of

the resonator in response to the qubit’s state change, can also be interpreted as the shift in the qubit frequency for each additional photon in the resonator [101]⁷. The dephasing rate for a resonator with a thermal population is given by [105]

$$\Gamma_{\phi,\text{res}} = T_{\phi,\text{res}}^{-1} = 4\bar{n}(\bar{n} + 1)\frac{\chi^2}{\kappa}, \quad (2.13)$$

where $\bar{n} = 1/(\exp(\hbar\omega_c/kT) - 1)$ is the mean thermal population of a resonator with frequency ω_c at temperature T , and κ is the resonator decay rate. Note that this expression holds in the weakly dispersive limit, $\chi/\kappa \ll 1$. In the strongly dispersive case, this dephasing has been modeled precisely in [104] and [105].

2.5 Conclusion

We have discussed basic methods for measuring dephasing as well as the forms dephasing might take and some of the most common sources. A basic Ramsey, spin echo, or Rabi measurement can be used to determine the dominant form of phase noise, which can then be used to determine its potential source. We now proceed to examine this problem in more depth. Chapter 3 describes the Ramsey Tomography Oscilloscope, useful for characterizing low-frequency noise, and discusses its application to distinct qubit types. Chapter 4 introduces the measurement of dephasing with randomized benchmarking, capable of measuring phase noise on the timescale of quantum gates, relevant to performing more complex algorithms.

⁷For this reason, such dephasing is sometimes called “photon-induced dephasing”. Additionally, the readout resonator can be a cavity (e.g. in 3-D transmons), and is then referred to as such.

Chapter 3

The Ramsey Tomography Oscilloscope

We now turn to more advanced techniques for measuring frequency noise in superconducting qubits. This chapter discusses the Ramsey Tomography Oscilloscope (RTO), useful for measuring the low-frequency end of the noise spectrum. We then present measurements on three different types of transmon-style qubits and compare the results to theory on flux noise in superconducting qubits.

Flux noise is the leading cause of dephasing in tunable superconducting qubits. It has been long studied in SQUIDs and superconducting qubits, and has been consistently found to have a $1/f^\alpha$ spectrum, with $0.8 \lesssim \alpha \lesssim 1.2$, and a magnitude of a few $\mu\Phi_0^2/\text{Hz}$ at 1 Hz, despite differences in device design, materials, and fabrication [134, 140, 54, 17, 107, 108, 23]. Several explanations for this noise have been proposed [65, 37, 26, 68] but none have satisfactorily explained all the data. However, it seems likely that the ultimate source of the noise is free magnetic fluctuators on the SQUID surface, and recent theoretical and experimental evidence strongly indicates that molecular oxygen is responsible [130, 66].

As the $1/f$ slope of flux noise in SQUIDs is found to be consistent over a wide frequency range—less 1 Hz to tens of MHz [23, 139, 141]—a measurement at low frequencies should allow us to explore the behavior of the noise across devices.

3.1 Introduction of the Ramsey Tomography Oscilloscope

Perhaps the most straightforward way to measure frequency noise in a qubit is simply to repeatedly measure the frequency and then analyze its variation. The Ramsey Tomography Oscilloscope does just that. Introduced in 2012 both at UCSB [102] and MIT Lincoln Laboratory [138], the algorithm repeats a Ramsey experiment with a fixed free evolution duration τ . This is depicted on the Bloch sphere in Figure 3.1. The angle $\theta = 2\pi f_{10}\tau$ that the qubit state rotates through is measured tomographically; that is, a final $Y/2$ rotation and measurement gives the x projection of θ , and similarly an $X/2$ rotation and measurement gives the y projection. The x and y sequences are repeated (in our case, on the order of several hundred times) and the measurement results averaged to get a single value for x or y . We then compute $f_{10} = \theta/(2\pi\tau) = \arctan(y/x)/(2\pi\tau)$.

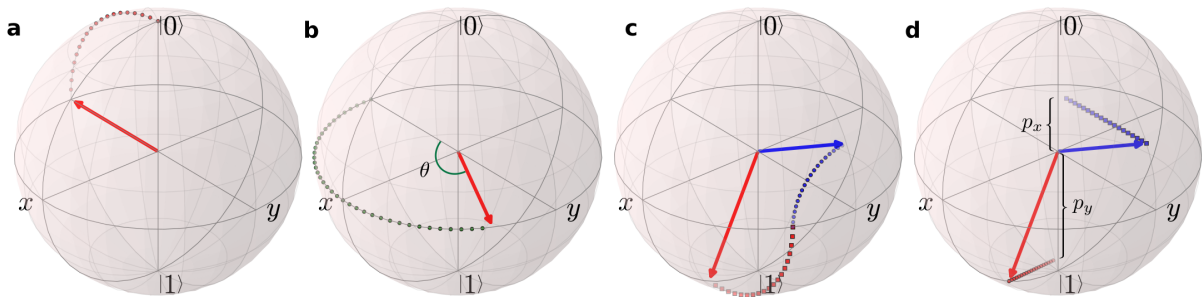


Figure 3.1: The RTO sequence, represented here on a Bloch Sphere, is identical to a Ramsey sequence (see Figure 2.2), repeated with a fixed free-evolution time τ . During the evolution (b), the qubit state acquires a phase $\theta = 2\pi f_{10}\tau$ (green). The x (blue) and y (red) components of the phase are tomographically measured (c, d) and averaged, and the phase recovered as $\theta = \arctan(p_y/p_x)$.

This experiment is repeated at regular intervals, as fast as possible, to generate a time series $f_{10}(t)$. The measurement limit imposed by the hardware is the qubit reset time (typically $150 \mu\text{s}$) combined with the number of repetitions: $150 \mu\text{s} \times 300 \text{ repetitions} \times 2 \text{ measurement angles} = 0.09 \text{ ms}$ per data point. However, ensuring that the experiment is repeated regularly over the entire data-taking period of ~ 8 hours is done in software, which has a latency of about 0.5 s , resulting in a typical Nyquist frequency f_n of 1 Hz . The data are then Fourier transformed to create the power spectral density (PSD) of the frequency noise, $S_f(f)$; the details of the data processing are given in Appendix C. For a given flux sensitivity $df_{10}/d\Phi$, this is converted to a flux noise power spectrum $S_\Phi(f) = S_f(f)/(df_{10}/d\Phi)^2$.




Figure 3.2: A processed RTO power spectral density (blue dots) fit to Eq. (3.1) (solid red line). The $1/f$ portion of the fit (dotted red line) has slope $\alpha = 0.95$ and magnitude $S_\Phi^* = 1.75 \mu\Phi_0^2/\text{Hz}$.

An example of a fully-processed noise power spectrum is given in Figure 3.2. In addition to the dominant $1/f$ component, the data show two additional effects: the

signal aliasing and a white noise floor. The fit to the data contains all three components,

$$S_{\Phi}(f) = \frac{S_{\Phi}^*}{f^{\alpha}} + S_{\Phi,0} + \frac{2S_{\Phi}^*}{(2f_n - f)^{\alpha}} \quad (3.1)$$

with two free parameters for the $1/f$ noise: the slope α and the magnitude at 1 Hz, S_{Φ}^* ; and one for the white noise: its magnitude $S_{\Phi,0}$. The third term of Eq. (3.1) represents the aliasing of the $1/f$ signal; it is relatively unimportant and only affects a few of the highest-frequency data points. The white noise is principally due to measurement error (and in some cases qubit decoherence (T_1) and white noise dephasing sources); as such, for studying flux noise, the two figures of merit extracted from the fit are the parameters to the $1/f$ fit, α and S_{Φ}^* .

Figure 3.2 shows that while the data clearly display a $1/f$ dependence, for accurate extraction of the key parameters α and S_{Φ}^* , we should reduce the noise in the PSD. The most straightforward way to do this is to take multiple RTO datasets and average them together. We find that for a given device type, the magnitude and slope of the $1/f$ noise are consistent, across different operating frequencies (provided we correctly account for the flux sensitivity $df_{10}/d\Phi$), and even across different devices on the same chip.

3.2 Ramsey Tomography Oscilloscope results in transmon-type qubits

We now present data from RTO measurements on three different qubit types.

3.2.1 Flux noise in Xmons

Four RTO datasets were taken on “standard” Xmons and are shown in Figure 3.3; these qubits were initially used to demonstrate gate fidelities suitable for quantum error correction, and are described in [11]. Two different qubits were used (q2 and q3), and




Figure 3.3: (left) Four RTO datasets (solid markers) taken on two Xmons of the device from [11]. (right) Averaged PSD (open circles) superimposed, with fit (solid line) to Eq. (3.1) and $1/f$ portion of the fit (dashed line), with $\alpha = 0.88$ and $S_{\Phi}^* = 3.0 \mu\Phi_0^2/\text{Hz}$.

two different operating frequencies and thus flux sensitivities were measured on q2. The magnitude of the averaged flux noise, $S_{\Phi}^* = 3.0 \mu\Phi_0^2/\text{Hz}$ at 1 Hz, is in good agreement with other results for aluminum superconducting qubits and SQUIDs [107, 102]. The slope, $\alpha = 0.88$, is slightly less than 1 (the case for “pure” $1/f$ noise). This is also not uncommon [17, 23, 139, 66], but is less well-reported. The Xmon RTO data will serve as a baseline for comparison with further measurements.

3.2.2 Flux noise in gmons

The gmon is a modification to the Xmon incorporating flux-tunable coupling between qubits¹. While the fabrication process of the gmon was similar to that of the Xmon, the shape of the SQUID loop is significantly modified to accommodate the coupler. Eight RTO datasets, shown in Figure 3.4, were taken on q2 of the gmon device introduced and

¹The variable for the coupling strength is g , hence the name “gmon”.




Figure 3.4: (left) Eight RTO datasets (solid markers) taken on gmon q2 of the device introduced in [98] and [25]. (right) Averaged PSD (open circles) superimposed, with fit (solid line) to Eq. (3.1) and $1/f$ portion of the fit (dashed line), with $\alpha = 0.92$ and $S_{\Phi}^* = 14.2 \mu\Phi_0^2/\text{Hz}$.

described in [98] and [25], at a variety of operating frequencies. The slope $\alpha = 0.92$ is comparable to that measured on Xmons, while the magnitude, $S_{\Phi}^* = 14.2 \mu\Phi_0^2/\text{Hz}$ at 1 Hz, is roughly five times greater.

3.2.3 Flux noise in early SiXmons

Due to a large increase in quality factor seen in superconducting aluminum resonators on silicon substrates [80], efforts are currently underway to translate this gain to Xmons by fabricating them on substrates made from silicon, rather than sapphire. RTO data were taken on an early version of these “SiXmon” qubits, described in [80]. The device design—in particular the SQUID loop—of the SiXmon is identical to the Xmon, but due to the change in materials, the fabrication process was very different. The ten datasets, taken across four different qubits, are shown in Figure 3.5. The magnitude of the $1/f$




Figure 3.5: (left) Ten RTO datasets (solid markers) taken across four SiXmon qubits. (right) Averaged PSD (open circles) superimposed, with fit (solid line) to Eq. (3.1) and $1/f$ portion of the fit (dashed line), with $\alpha = 1.13$ and $S_{\Phi}^* = 2.4 \mu\Phi_0^2/\text{Hz}$.

noise, $S_{\Phi}^* = 2.4 \mu\Phi_0^2/\text{Hz}$ at 1 Hz, is in good agreement with the Xmon, but the slope $\alpha = 1.13$ significantly different from that seen in both the Xmon and gmon devices.

3.2.4 Comparison with theory

The most likely explanation for the presence of $1/f$ noise in SQUIDs is the presence of magnetically active fluctuators on the surface of the superconductor. While there is currently no microscopic theory for the type and behavior of these fluctuators that successfully explains all the evidence, we can still use general arguments to gain insight into how the flux noise may vary across different device geometries. This method follows Ref. [17].

The effect of a magnetic fluctuator on a conducting loop may be calculated through reciprocity: the flux from a fluctuator of moment m is given by $(B \times m)/I$, where B is the magnetic field experienced by the fluctuator produced by a test current I through

the loop. The magnetic field on the surface of the loop is proportional to the surface current density; for a superconducting thin film strip of width W , the surface current density J varies across the width of the strip as [125]

$$J(x) = \frac{J(0)}{\sqrt{1 - (2x/W)^2}}, \quad (3.2)$$

where $J(0)$ is the current density at the center and x is the horizontal position. (This is for the case where $R \gg W \gg b$ and $Wb \gg \lambda^2$ for length R , film thickness b , and superconducting penetration depth λ .) Now taking $B \propto J(x)$ and $I = \int J(x)dx$, the mean-square flux induced by a fluctuator at position x is

$$\langle \Phi^2 \rangle = C \frac{J^2(x)}{(\int J(x)dx)^2}, \quad (3.3)$$

where C contains the unknown microscopic details of the fluctuator. For a surface density of fluctuators σ , the total flux is

$$\langle \Phi^2 \rangle = \int \sigma \langle \Phi^2 \rangle dA = C' \int \ln \left(\frac{2bW(l)}{\lambda^2} \right) dl, \quad (3.4)$$

where $C' = \sigma C$ now contains all microscopic details, and we have integrated $J^2(x)$ across the width of the wire and taken the first-order approximation of $\int J(x)^2 dx / (\int J(x)dx)^2$, and the remaining integral is along the length of the wire (that is, around the SQUID loop), allowing the width $W(l)$ to vary². To compare geometric effects, therefore, we can evaluate the ratio of the integrals in last part of Eq. (3.4), making the assumption that the microscopic details of the fluctuators are the same.

The data for the three different qubit types are plotted together in Figure 3.6. Comparing the SQUID geometries between the gmon and Xmon according to Eq. (3.4), we compute $\langle \Phi_{\text{gmon}}^2 \rangle / \langle \Phi_{\text{Xmon}}^2 \rangle \approx 1.7$. The relative noise magnitude, however, is $S_{\Phi, \text{gmon}}^* / S_{\Phi, \text{Xmon}}^* \approx 4.7$. While the data do not convincingly rule out the geometric argument, they also indi-

²It is common when analyzing SQUIDS to evaluate this integral for a circular washer of radius R and fixed width W and compare flux noise in terms of the ratio R/W ; for qubits with less regularly shaped SQUID loops, we attempt a more fine-grained approach here.




Figure 3.6: The averages (circles), fits (solid lines) and $1/f$ portion of the fit (dashed line) of the Xmon (blue), gmon (red), and SiXmon (green) RTO data, from Figures 3.3–3.5, shown together for comparison.

cate that such arguments may not provide a detailed guide to reducing flux noise during qubit design. As the trend in flux noise agrees with the trend predicted by geometry, however, further exploration of qubit design on flux noise magnitudes is warranted. Other flux noise results have also provided mixed support for such geometric arguments [102, 1].

For SiXmons, however, as the SQUID geometry is identical, we expect any difference to be due to materials or fabrication processes—and, indeed, the materials and fabrication are quite different. The significantly steeper slope seen in the SiXmon data is particularly compelling because the effect of flux noise on the fidelity of a sequence of quantum operations is proportional to the integral of the flux noise over many orders of magnitude, and small changes in α greatly change the value of the integral. Again, we recommend further research in this regard.

Chapter 4

Qubit metrology of ultralow phase noise using randomized benchmarking¹

We now introduce a method for measuring phase noise at the high frequency end of the spectrum—on the same timescales as quantum gates, which are tens of nanoseconds in our devices. Spectra of phase noise have been measured spectroscopically up to hundreds of megahertz [23, 139, 141] using techniques discussed in Chapter 2. These techniques acquire data over long timescales and use various methods remain sensitive to only certain frequencies of noise, making it difficult to determine precisely how much the measured noise will affect a quantum algorithm consisting of gates sequences taking place on short timescales. Furthermore, as the performance superconducting qubits improves to the level required for error correction [11, 58], the need arises for a method to accurately quantify the now minute errors induced by dephasing. The approach presented here

¹After the first paragraph, this chapter was previously published as: “Qubit Metrology of Ultralow Phase Noise Using Randomized Benchmarking”, Peter O’Malley, Julian Kelly, Rami Barends, et al. *Phys. Rev. Applied*, **3**, 044009 (2015).

makes use of the well-understood randomized benchmarking (RB) protocol [99, 74] to measure gate errors on short timescales, rapidly gathering large statistics while remaining insensitive to errors caused by state preparation and measurement. We first introduce the protocol, called “RB Ramsey”, by comparing it to standard Ramsey and spin echo sequences and describe how it is used to measure dephasing. We then use it to measure a telegraph noise mechanism in our qubit, as well as decoherence caused by two-qubit interactions. Finally, we demonstrate the separation various dephasing effects: T_1 , white noise, and correlated (e.g. $1/f$) noise.

4.1 Introduction

One of the main challenges in quantum information is maintaining precise control over the phase of a superposition state. Long-term phase stability is threatened by frequency drifts due to non-Markovian noise, which arises naturally in solid-state quantum systems [134, 44]. Fortunately, correlated noise can be suppressed using Hahn spin echo [45]. In practice, Ramsey and spin echo measurements of dephasing [31, 18, 23] characterize the dominant noise source for large error rates (0.1 to 0.5) and long times, but are fundamentally inappropriate for understanding noise dominant on the timescales and error rates needed for fault-tolerant gate operations ($< 10^{-2}$).

We introduce a metrological tool based on randomized benchmarking [62, 99, 74, 22, 42, 29] to quantify noise on timescales relevant for quantum gates. Whereas other measurement techniques based on Ramsey [31, 18, 23] and Rabi [139] measurements measure noise over long timescales and filter low frequency noise to infer gate performance at short timescales, we measure gate fidelity directly, providing immediate feedback on the impact of noise on gate performance. We apply it on a SQUID-based qubit, and show that this method determines that $1/f$ flux noise [134, 102, 138] is not currently a limiting

factor in our device. This tool also provides a powerful probe of anomalous telegraph noise sources seen in superconducting devices. We also show that undesired coherent interactions can be understood as an effective correlated noise. Finally, we demonstrate how this method allows for error budgeting and direct selection of ideal gate parameters in the presence of non-Markovian noise.

Quantum systems based on ion traps, spin qubits, and superconducting circuits are rapidly maturing, with individual operation fidelity at the levels required for fault-tolerant quantum computing [46, 22, 88, 95, 27, 103, 3, 50, 11, 58]. These systems are often limited by environmentally-induced phase noise, which can manifest as qubit frequency jitter. Noise in the phase ϕ is characterized by variance $\langle \phi^2(\tau) \rangle$, increasing linearly with time τ for white noise, and with higher power for correlated noise [76]. Ramsey and spin echo experiments measure the decay of phase coherence for large magnitudes over long timescales; at much shorter timescales, which are relevant to quantum gates but still slower than the qubit frequency, dephasing errors are small and thus hard to measure, making physical mechanisms difficult to directly identify. Here, we quantify phase noise by using RB to measure the decoherence of an identity gate versus its duration, providing an unprecedented metrological tool.

We use a superconducting quantum system based on the planar transmon qubit variant, the Xmon [10], cooled to 20 mK in a dilution refrigerator. This qubit consists of a SQUID, which serves as a tunable non-linear inductor, and a large X-shaped shunt capacitor. It is well-suited for characterizing phase noise as the qubit has long energy relaxation times, and the SQUID gives a controllable susceptibility to flux noise. These qubits have frequencies that can be tuned to 6 GHz and below and have typical nonlinearities of $\eta/2\pi = -0.22$ GHz, and capacitive coupling strengths between qubits of $2g/2\pi = 30$ MHz. Single qubit rotations are performed with microwave pulses and tuned using closed-loop optimization with RB [57]. We use a dispersive readout scheme with

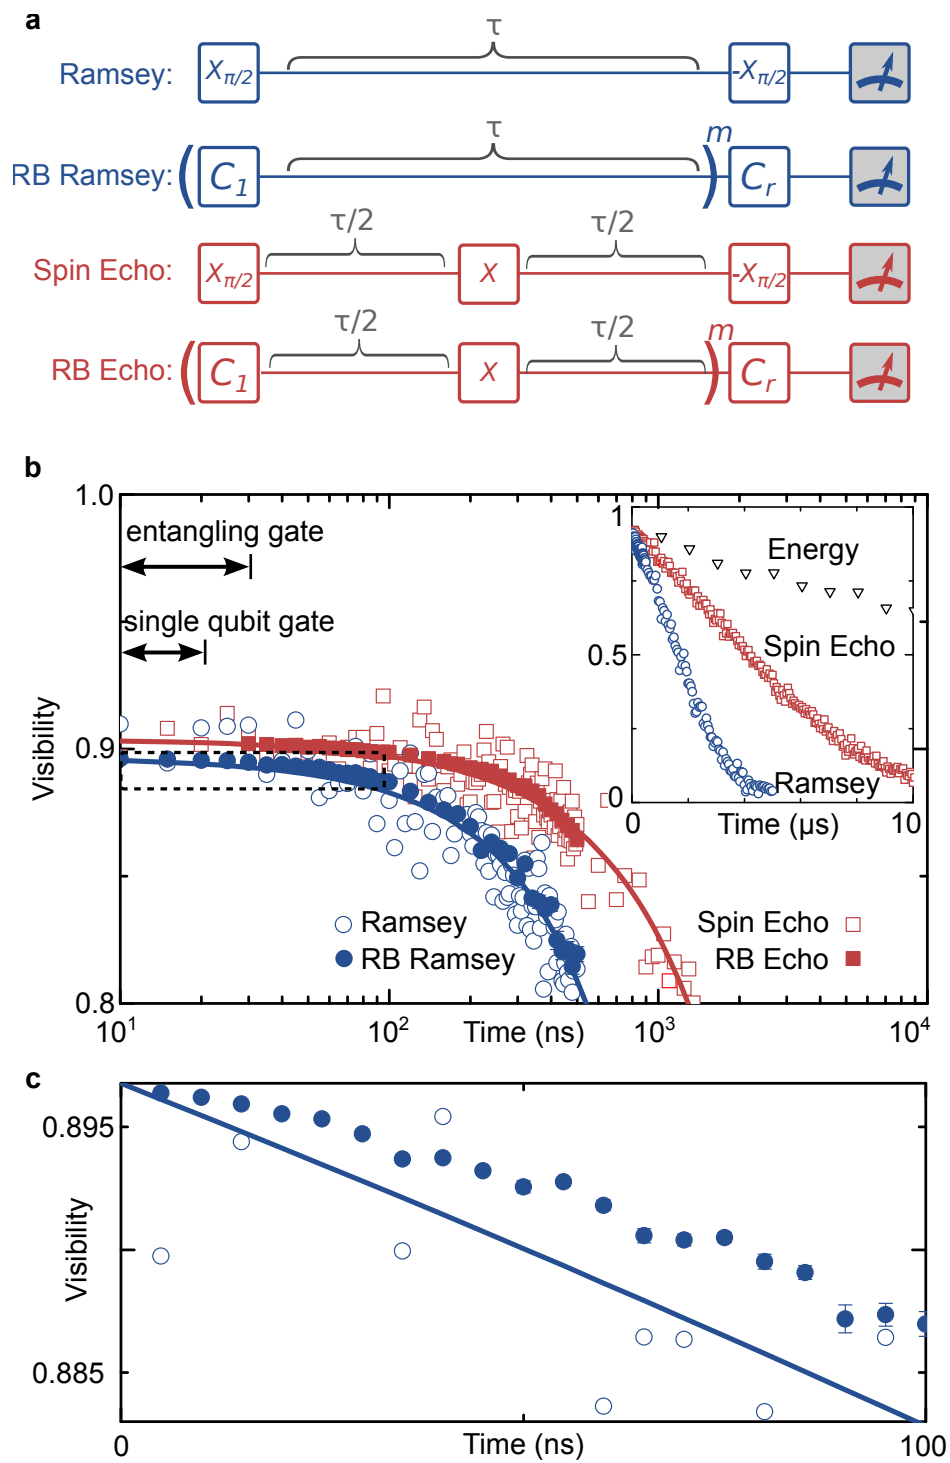


Figure 4.1: (Caption next page.)

Figure 4.1: (Previous page) (a) Gate diagram for Ramsey and Hahn spin echo sequences, and their RB equivalents. For RB Ramsey, instead of inserting an idle between $X_{\pi/2}$ pulses, we interleave the idle between m randomly selected single-qubit Clifford gates (C_1), after which the qubit is rotated back (C_r) to the pole and measured. For spin echo and RB echo, an X gate is inserted at the center of the idle. The range of m is 21 for the longest τ to 300 for the shortest. (b) (inset) T_1 (energy decay), Ramsey, and spin echo envelopes. (main) Ramsey (open circle) and spin echo (open square) envelopes at short times. RB decay envelopes are inferred from $\langle\phi^2(\tau)\rangle$ measured by RB Ramsey (solid circle) and RB echo (solid square); see text for details. Single qubit and entangling gate durations are shown for reference. Note the significantly lower noise of the RB sequences, which take approximately the same measurement time as the Ramsey and echo experiments. (c) Magnification of the dashed area in (b), showing timescales important for gates. The RB Ramsey data show a trend different from that predicted by the Ramsey fit.

capacitively coupled resonators at 6.6–6.8 GHz for state measurements [51]. For details of the experimental setup and fabrication process, see [58].

4.2 RB Ramsey

Figure 4.1a shows gate sequences for Ramsey and spin echo measurements, as well as their RB equivalents that we have called “RB Ramsey” and “RB echo”. The Ramsey experiment accumulates phase error from a single period τ , whereas the RB Ramsey experiment accumulates phase error from m applications of τ , with m typically of order 100. In RB, gate error is measured directly by interleaving gates with random Clifford group operators, which depolarize errors by evenly sampling the Hilbert space, such that repeated gate applications add error incoherently [75]. Thus, RB Ramsey has a factor m higher sensitivity than Ramsey when errors and times τ are small. The error of an idle gate, $r_{I(\tau)}$, is directly related to the variance of the phase noise by (see Appendix D.1)

$$r_{I(\tau)} = \frac{1}{6}\langle\phi^2(\tau)\rangle. \quad (4.1)$$

We infer and plot the equivalent Ramsey decay envelope visibility data V (solid circles) with $V(\tau) = A \exp(-\langle\phi^2(\tau)\rangle/2) + B$ in Figure 4.1b, with state preparation and measurement error parameters A and B extracted from the Ramsey fit as described in Appendix D.3, and $\langle\phi^2(\tau)\rangle$ measured by RB Ramsey according to Eq. (4.1). We likewise show the equivalent spin echo decay envelope from RB echo data as solid squares. The Ramsey and spin echo measurements over the same timescale are shown for comparison as open circles and open squares, respectively. We label the length of a single qubit and two-qubit entangling gate [11] to emphasize the relevant timescale. The full Ramsey and spin echo measurements are shown on the typical linear scale, together with energy relaxation, in the inset of Figure 4.1b.

As shown in Figure 4.1b, the RB Ramsey and RB echo data are consistent with the Ramsey and spin echo measurements, respectively, at short to moderate time scales, while measuring $\langle\phi^2\rangle$ with much greater precision. Any structure to short-time dephasing is obscured in the Ramsey data, whereas the RB Ramsey data reveal a time dependence that we will show is consistent with telegraph noise. The use of RB greatly improves the precision of phase noise measurements; the uncertainty of the measured Ramsey visibility for $\tau < 300$ ns is reduced by an order of magnitude. We note that the total time taken to perform the Ramsey and RB Ramsey measurements is approximately the same, and that precision would be increased for a higher-fidelity qubit by simply choosing larger m 's. Because of the imprecision of the Ramsey data at short time scales, the amount of noise present can only be inferred from the fit to the entire Ramsey dataset. However, Figure 4.1c shows that the phase noise measured by RB Ramsey can differ significantly from that expected by the Ramsey fit. The trend in their difference indicates that there is behavior to the noise at short times that Ramsey measurements miss. We examine this in Figure 4.2.

4.3 Measuring telegraph noise




Figure 4.2: RB Ramsey measurement (circles) for short timescales; note that the small error from T_1 decay, which is 9×10^{-4} at 450 ns, has been subtracted (see Appendix D.1). We fit to a telegraph noise model, Eq. (4.4); the dotted (dashed) lines give the short (long) time limit of the noise model. The inferred but negligible contribution from $1/f$ noise as measured for this qubit (see Appendix D.4) is shown as a thick line. The inset shows the experiment used to extract the 40 ns data point.

To identify the dominant noise mechanism, we examine the dependence of idle gate error on time and compare against different noise models in Figure 4.2. Whereas in Figure 4.1 we infer an equivalent Ramsey envelope, here we plot the idle gate error

directly, as measured by RB Ramsey (with small T_1 effects subtracted, see Appendix D.1). For short times, we see a non-linear increase of error with gate duration which transitions into a linear behavior for lengths above approximately 100 ns. The inset shows the sequence fidelity *vs.* number of Cliffords, with and without interleaved idles, used to extract the idle error for $\tau = 40$ ns.

While it has long been known that SQUIDs are susceptible to $1/f$ flux noise [89, 134, 110, 140, 17, 107, 102], we find this a negligible contribution to gate error. A system limited by $1/f$ and white noise would see a linear increase in error at short, and quadratic increase at long times as the $1/f$ component begins to dominate. The data exhibit the opposite trend. Moreover, the expected contribution to gate error from $1/f$ noise, as measured for this system below 1 Hz using the Ramsey Tomography Oscilloscope protocol (see [102] and Appendix D.4), is significantly less than observed here (Figure 4.2 thick solid line).

The trend observed in Figure 4.2 is consistent with telegraph noise. For a random telegraph switching of the qubit frequency, the phase noise is given by

$$\langle \phi_{\text{tel}}^2(\tau) \rangle = (2\pi\Delta f_{10})^2 T_{\text{sw}} \left(\tau - T_{\text{sw}} \left[1 - \exp\left(-\frac{\tau}{T_{\text{sw}}}\right) \right] \right), \quad (4.2)$$

where Δf_{10} is the effective switching amplitude of the qubit frequency and T_{sw} is the switching timescale. We make the simplifying assumption of symmetric telegraph noise as the measurement is unable to differentiate up and down switching rates, and note that while telegraph noise is not Gaussian, Eq. (4.2) is still approximately correct for use in Ramsey and spin echo analyses (see Appendix D.1). In a more general case, the error rate for an idle of length τ , $r_{I(\tau)}$, can be fit to a combination of error sources: white,

long-time correlated, $1/f$, and telegraph phase noise, as well as T_1 decay,

$$r_{I(\tau)} = \frac{\tau}{3T_1} + \frac{1}{6} \left(\langle \phi_{\text{white}}^2(\tau) \rangle + \langle \phi_{\text{corr}}^2(\tau) \rangle + \langle \phi_{1/f}^2(\tau) \rangle + \langle \phi_{\text{tel}}^2(\tau) \rangle \right), \quad (4.3)$$

where the derivation for $\langle \phi_{\text{white}}^2(\tau) \rangle = 2\tau/T_{\phi 1}$, $\langle \phi_{\text{corr}}^2(\tau) \rangle = 2(\tau/T_{\phi 2})^2$, and $\langle \phi_{1/f}^2(\tau) \rangle$ are given in Appendix D.2, and we assume correlated noise has a longer timescale than the experiment. The data here are fitted to a noise model featuring only T_1 decay (measured independently) and telegraph noise,

$$r_{I(\tau)} = \frac{\tau}{3T_1} + \frac{1}{6} \langle \phi_{\text{tel}}^2(\tau) \rangle, \quad (4.4)$$

indicating that $1/f$ and white noise do not dominate the error for this qubit. We extract $T_{\text{sw}} = 84 \pm 14$ ns and $\Delta f_{10} = 479 \pm 30$ kHz from the fit. The dotted (dashed) line shows this noise model in the short (long) time limit. Perhaps surprisingly, this measurement directly shows that gates of duration 20 ns can achieve fidelity > 0.999 in a system with characteristic Ramsey scale of $T_{\phi 2} = 2.0 \mu\text{s}$ (see Appendix D.3).

Telegraph noise has been studied in superconducting circuits with a variety of methods. Frequency fluctuations due to quasiparticle (QP) tunneling have been characterized by Rabi oscillations [8] and repeated direct frequency measurement [96]. For our qubit, the calculated frequency splitting due to QP tunneling ranges from 1 Hz to 14 kHz (see Appendix D.6), well below the magnitude necessary to explain the data. Photon shot noise in a coupled resonator has been shown to cause dephasing in both transmon [88, 104, 95] and flux [116] qubits. In our case the magnitude of the telegraph noise decreases as the qubit-resonator frequency difference decreases, indicating that resonator photon noise induced dephasing is not the cause. A more elusive telegraph-like noise has been measured by $T_{1\rho}$ Rabi spectroscopy in flux qubits [139], hypothesized to be due to

two sets of coupled coherent two-level states. This noise is similar in frequency to the noise measured here, with spectroscopic signatures at 1 and 20 MHz, compared to $1/84 \text{ ns} = 11 \text{ MHz}$ for this measurement. However, it is much larger in magnitude, presenting as a “dip” (or “plateau”) in spin echo measurements, which is known to happen in the presence of strong telegraph noise [41], and seen in other systems [49, 116, 96]. In our device, the telegraph noise is only dominant at short timescales, as any evidence of it in longer measurements like Ramsey and spin echo is masked by $1/f$ flux noise.

4.4 Measuring error from coherent qubit-qubit interactions

We now apply RB to coherent errors arising from unwanted qubit-qubit interactions, which can also contribute to dephasing [33]. In Figure 4.3, we explore these effects in our system. Figure 4.3a shows an energy level diagram for capacitively coupled qubits, where the fundamental entangling rate Ω_{ZZ} [40] arises from an avoided level crossing between the $|11\rangle$ state and the $|02\rangle$ and $|20\rangle$ states. This interaction manifests as a state-dependent frequency shift, falling off with detuning Δ , as measured in Figure 4.3b. We note that for a qubit coupled to a resonator, Ω_{ZZ} is equivalent to the dispersive shift [20] 2χ as defined in [63]. The inability to turn this interaction off completely results in additional errors when operating qubits simultaneously. Figure 4.3c shows average gate error *vs.* duration, when a qubit is operated in isolation or simultaneously with a coupled qubit ($\Omega_{ZZ}/2\pi = 0.4 \text{ MHz}$). Error for single qubit or simultaneous operation is inferred from the RB reference error per Clifford, divided by the average of 1.875 physical gates per Clifford [11]. The difference between isolated and simultaneous operation gives the added error from the Ω_{ZZ} interaction, which is fit to a quadratic.

This interaction is correlated, and therefore the errors are quadratic with gate du-




Figure 4.3: (a) Energy level diagram for two capacitively coupled qubits with coupling strength $2g/2\pi = 30$ MHz, detuned by frequency Δ . The avoided level crossing between the $|11\rangle$ and $|02\rangle/|20\rangle$ states repels the $|11\rangle$ frequency from the sum of $|01\rangle$ and $|10\rangle$ frequencies by the amount Ω_{ZZ} . (b) This entangling interaction causes the phase of one qubit to precess, conditional on the state of its neighbor (cartoon and inset). The Ω_{ZZ} interaction decreases with Δ , to a level of $\Omega_{ZZ}/2\pi = 0.4$ MHz at $\Delta/2\pi = 750$ MHz. (c) RB data isolating the Ω_{ZZ} interaction. Gate error is measured *vs.* gate duration for a single qubit and when qubits are operated simultaneously (inset). The difference (main figure) measures the error contribution from the Ω_{ZZ} interaction, and is fit to $1.86(\Omega_{ZZ}t_{\text{gate}}/2\pi)^2 + 1.4 \times 10^{-4}$.

ration; specifically, the error per gate due to the Ω_{ZZ} interaction between two qubits simultaneously undergoing RB is

$$E = \frac{\pi^2}{6} \left(\frac{\Omega_{ZZ}}{2\pi} t_{\text{gate}} \right)^2, \quad (4.5)$$

where $\Omega_{ZZ}/2\pi$ is the interaction magnitude and t_{gate} is the RB gate duration (see Appendix D.7). The fit to the data has a quadratic coefficient of 1.86 ± 0.1 , while $\pi^2/6 \approx 1.64$. Here, the careful application of RB both distinguishes these errors at the $1 \cdot 10^{-4}$ level, and indicates that short gates are effective in suppressing them.

4.5 Measuring different gate implementations




Figure 4.4: (color online) Operation error of σ_I and σ_Z , implemented with (closed symbols) and without (open symbols) echoing, as measured with interleaved RB. The data are fitted to a linear and quadratic form, representing uncorrelated and correlated noise. The dark gray region indicates error attributed to T_1 , the medium gray region uncorrelated noise, and the light gray region non-Markovian (e.g., telegraph) noise. Note that the I data are RB Ramsey data, the same as Figure 4.2.

We now examine the gate fidelity for a variety of gates in the presence of the non-Markovian noise we have measured. Figure 4.4 shows gate fidelity *vs.* gate length for two implementations each of two different gates: for σ_I , an idle and two microwave

pulses (X, X), and for σ_Z , a frequency detuning pulse and two microwave pulses (Y, X). The errors of these operations *vs.* duration are determined with interleaved RB. In agreement with previous measurements, we find that the error of operations without X or Y pulses (open symbols) follow a quadratic-like dependence with gate duration at these timescales. Using X or Y pulses (closed symbols), we observe a linear-like dependence at longer durations, indicating that the correlated phase noise has been suppressed. Below 40 ns, we find an increased error which we attribute to the population of higher levels due to spectral leakage [71]. The solid (dashed) lines are linear (linear and quadratic) fits to the data. For full details of the fits, see Appendix D.3.

Using the functional forms of the different error types given in Eq. (4.3), we can determine an error budget for our operations. For a typical entangling gate duration of 40 ns, T_1 contributes an error of 5×10^{-4} , and telegraph noise an error of 5×10^{-4} . With echoing pulses, the total error is 8×10^{-4} , indicating that the added echoing pulses are either not completely suppressing the phase noise or are contributing error of their own. Using a combination of RB Ramsey and RB echo, we have determined the relative contribution of different noise sources to operational error, and we can also immediately see that either short gates, or long gates with intrinsic echoing, are effective at remedying non-Markovian noise, and by how much.

4.6 Summary

RB Ramsey provides a direct measurement of phase noise in the regime most relevant to quantum gates. While previous noise spectroscopy has relied on accumulating noise over longer timescales while filtering out low-frequency noise with additional pulses, our technique directly measures small amounts of noise with repeated incoherent additions. It does not require extensive calibration, and is also robust against state preparation and

measurement error. As a gate-based measurement, it is useful in a variety of situations: measuring noise due to the environment as RB Ramsey, measuring filtered environmental noise as RB echo, and measuring dephasing induced by coherent qubit-qubit interactions. As the measurement output is gate fidelity, it is also immediately applicable as a tool to determine the highest-fidelity implementation of different quantum gates in the presence of noise. We show here that RB Ramsey is the metrological tool best suited for measuring noise in high-fidelity qubits.

We have taken RB, a protocol for determining the fidelity of gates, and applied it as a metrological tool for identifying noise processes. Applied to a superconducting qubit system, we have found a telegraph noise mechanism in a regime inaccessible to previous measurements, accurately characterized dephasing caused by coherent qubit-qubit interactions, and determined the highest-fidelity implementation of different quantum gates. Our results demonstrate that RB Ramsey is capable of measuring small noise processes at short timescales that are directly relevant to gate fidelity, and show that understanding this non-Markovian phase noise can be lead to its effective suppression through short gates and echoing.

Chapter 5

Chemistry on a quantum computer: a brief introduction

Previous chapters introduced the problem of dephasing—currently the leading source of error in Xmon qubits—and presented the results of two different techniques for characterizing opposite ends of the phase noise power spectrum. While this is clearly an immediate problem to be solved in the effort to build a quantum computer, it is not necessarily of great interest to someone who wants to use such a device for practical purposes. The remainder of this thesis turns to the opposite case: we present and demonstrate two quantum algorithms for computations of great practical interest outside of the field of quantum computing, finding molecular ground state energies. This problem has been the focus of much of theoretical chemistry since the formulation of quantum mechanics itself. However, the computational complexity of the full solution of a molecular Hamiltonian—known as the full configuration interaction, or FCI—scales exponentially with the number of electrons in the system. Quantum algorithms have the potential to make such calculations possible; indeed, it was for this purpose that Feynman proposed quantum computers [38].

This is not entirely unrelated to the work on dephasing presented earlier, as it may seem at first glance. The first algorithm discussed, the phase estimation algorithm, uses the phase of a qubit to store the result of the computation, making it particularly susceptible to any phase noise. In contrast, the variational quantum eigensolver is robust to some systematic errors, and in our implementation this robustness allows its success in the presence of dephasing.

This chapter first provides a brief overview of the chemistry necessary to encode a molecule on a quantum computer. We then describe the phase estimation algorithm, and its application to chemistry when combined with Hamiltonian Trotterization. Finally, we discuss the variational quantum eigensolver, an hybrid quantum-classical algorithm for chemistry potentially useful on quantum devices without error correction.

5.1 Representing electrons with qubits: The Bravyi-Kitaev Transform

Our goal is to compute the lowest energy eigenvalue of a given molecular Hamiltonian: the electronic structure problem. The precision to which we must measure the energy eigenvalue is known as “chemical accuracy”, which is 1.6×10^{-3} Hartree¹, 1 kcal/mol, or 0.043 eV. Chemical reaction rates are proportional to the exponential of the ratio energy difference to thermal energy kT ; a relative error equal to chemical accuracy results in a chemical rate change by an order of magnitude.

The first task in simulating a molecule, therefore, is writing down the Hamiltonian in a scalable way. This is summarized in Appendix E.1 (and references therein), and outlined in Figure E.1; the end result is a Hamiltonian written in the second quantized

¹One Hartree, the preferred energy unit of theoretical chemists, is $\hbar^2/m_e e^2 a_0^2$, where m_e , e , and a_0 are the electron mass, electron charge, and Bohr radius, respectively.

formalism:

$$H = \sum_{pq} h_{pq} a_p^\dagger a_q + \sum_{pqrs} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s \quad (5.1)$$

where a_i^\dagger and a_i are the fermionic creation and annihilation operators for an electron in molecular orbital² i and the h_{ij} coefficients are (efficiently, classically) computed from the orbital overlaps (see Eq. (E.3)). However, we cannot simply use qubits to represent fermionic operators directly, as they do not obey fermionic commutation relations; we must therefore carefully map between the two. The most commonly cited mapping is the Jordan-Wigner transform, first introduced in 1928 as a mapping between spins and fermions [53], and formulated in 2002 for use in quantum computation [115].

The Jordan-Wigner transform represents the occupation of each orbital with the state of a single qubit—this is called the “occupation basis”. Considering a state with n orbitals where f_i is the occupation of orbital i , $|f_{n-1} \dots f_0\rangle$, when applying the creation operator a_j^\dagger , the state acquires a phase of -1 for each occupied orbital with index less than j ; that is,

$$a_j^\dagger |f_{n-1} \dots f_{j+1} 0 f_{j-1} \dots f_0\rangle = (-1)^{\sum_{k=0}^{j-1} f_k} |f_{n-1} \dots f_{j+1} 1 f_{j-1} \dots f_0\rangle. \quad (5.2)$$

This means that the qubit implementation of a_j^\dagger must contain not only a gate to change the value of qubit j but also must compute the parity of all qubits with index less than j ; this is accomplished with a Z gate on each. Therefore, one fermionic operation requires $O(n)$ gates, and furthermore, these gates are nonlocal.

An alternative encoding is the “parity basis”, where qubit j now stores the parity of all orbitals with index $\leq j$. In this case, when applying a_j^\dagger , the phase acquired due to the parity of orbitals $< j$ is easily computed with a Z gate on qubit $j - 1$. However, when changing the occupation of orbital j , all qubits with index $\geq j$ must be changed;

²By “orbital” we mean a combined spin and spatial orbital, so each orbital can be occupied by at most one electron.

therefore, the computational requirements in the parity basis are equivalent to that of the occupation basis.

The Bravyi-Kitaev (BK) transform was first introduced in 2002 for the purpose of using fermions to perform quantum computation [21]; the inverse case, of representing fermions on qubits, is described in [106] and [120]. Rather than storing the occupation locally and the parity nonlocally, as in the occupation basis (and vice versa for the parity basis), the BK transformation stores both of them nonlocally. For even index j , qubit j holds the occupation of orbital j . For odd j , qubit j stores the parity of a set of orbitals with index $< j$. Following [120], the transformation is defined by an n -by- n matrix β_n , where $\beta_n \vec{f}_n = \vec{b}_n$ with the occupation number basis vector \vec{f}_n and the BK basis \vec{b}_n :

$$\beta_{2^x} = \left(\begin{array}{c|c} \beta_{2^{x-1}} & 0 \\ \hline 0 & \beta_{2^{x-1}} \\ \leftarrow 1 \rightarrow & \end{array} \right), \quad (5.3)$$

where $\leftarrow 1 \rightarrow$ represents a row of ones, and $\beta_1 = (1)$.

This is best seen with an example; here we give $\beta_8 \vec{f}_8 = \vec{b}_8$:

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 1 & 1 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \end{pmatrix} \begin{pmatrix} f_0 \\ f_1 \\ f_2 \\ f_3 \\ f_4 \\ f_5 \\ f_6 \\ f_7 \end{pmatrix} = \begin{pmatrix} f_0 \\ f_0 + f_1 \\ f_2 \\ f_0 + f_1 + f_2 + f_3 \\ f_4 \\ f_4 + f_5 \\ f_6 \\ f_0 + f_1 + f_2 + f_3 + f_4 + f_5 + f_6 + f_7 \end{pmatrix} \quad (5.4)$$

In this case, to apply a_7^\dagger , we must check qubits 3, 5, and 6 to determine the parity of orbitals < 7 ; this is called the “parity set” of orbital 7, or $P(7) = \{3, 5, 6\}$. We must then change the value of qubit 7, but as f_7 does not appear in any other qubit, we do not need to update any other qubits; thus the “update set” of orbital 7 is empty, or $U(7) = \emptyset$. Finally, to determine whether the occupation of orbital 7 is equal to the value of qubit 7, or its inverse, we consult the “flip set” of orbital 7: $F(7) = \{3, 5, 6\}$. (In general, the flip

set is a subset of the parity set; for even numbered orbitals in particular, it is empty.) For orbital 0, we note by contrast that $P(0) = F(0) = \emptyset$, while $U(0) = \{1, 3, 7\}$.

The full details of the BK transform are beyond the scope of this treatment; the transform is covered pedagogically in [106] and [120]³. We note that the size of the parity, update, and flip sets scales logarithmically with the number of orbitals, n , whereas for the Jordan-Wigner transform, the parity set scales linearly with the number of orbitals ($P_{\text{JW}}(j) = \{i | i < j\}$), while the update set is empty (and vice versa for the parity basis). This means that the BK basis requires fewer gates to implement fermionic operators than the occupation or parity bases. This is shown numerically for hydrogen in [106] and methane in [120]. Crucially, the greater locality of the BK mapping reduces the number of required CNOT gates, when compared to the Jordan-Wigner mapping; in some cases the number of single-qubit gates was greater, but the total number of gates is always fewer, indicating that there is no overhead that may make the BK transform impractical for small systems.

Furthermore, the increased locality of the BK transform is of particular importance for certain device architectures. When the qubit coupling is restricted to nearest neighbors, a logical CNOT between non-adjacent qubits requires “swapping through” intervening qubits. A SWAP gate is conventionally implemented as three CNOTs, meaning that non-local CNOTs greatly increase gate counts.

5.2 The canonical quantum chemistry algorithm

First described in 2005 [2], what has come to be called “the canonical quantum chemistry algorithm” is a combination of quantum phase estimation with Hamiltonian Trotterization. The quantum phase estimation algorithm is a standard method in the quantum

³Note that matrix row- and column-labeling conventions are reversed between the two, however.




Figure 5.1: Schematic for one iteration of phase estimation. The top qubit is the ancilla, with input state $|0\rangle$; the register qubits’ input state is the ground state, denoted $|\psi_0\rangle$. $e^{-i\hat{H}t}$ is the Trotterized version of the molecular Hamiltonian applied for time t .

information literature for measuring the eigenvalue of a unitary operator [85]; the chemistry algorithm (henceforth “PEA”) is the application of this to a molecular Hamiltonian. In particular, the problem Hamiltonian must first be encoded in the qubits, with a method such as the Jordan-Wigner or Bravyi-Kitaev transforms as discussed in Section 5.1. This encoding produces a series of gates that implement the Hamiltonian; however, the terms in this Hamiltonian will likely not all commute, so its application will require use of the Trotter-Suzuki expansion [117] (also known as “Trotter decomposition” or just “Trotterization”) to approximate the time evolution of a set of non-commuting operators, at the cost of some error. This Trotterization is then applied to a register set of qubits encoding the ground state controlled by a set of ancilla qubits. The inverse quantum Fourier transform is then used to read out the phase acquired by the ancilla qubits, which is proportional to the eigenvalue of the Trotterized Hamiltonian, i.e., the ground state energy.

For full details of the standard phase estimation algorithm, see [85]; an alternative version, iterative phase estimation, introduced by Kitaev in 1995 [60], uses fewer qubits in the ancilla at the cost of repeated measurements. This was proposed for use on superconducting qubits in 2007 [35] and described in detail in 2010 [136] in the context of molecular Hamiltonians, and as it is the method used in the experiment of Chapter 6, we describe it briefly here. In particular, for the case of one ancilla qubit, the circuit schematic is given in Figure 5.2. The action of this circuit to apply a phase to the register

qubit,

$$\begin{aligned} |0\rangle|\psi_0\rangle + \exp(-i\hat{H}t)|1\rangle|\psi_0\rangle &= |0\rangle|\psi_0\rangle + \exp(-iE_0t)|1\rangle|\psi_0\rangle \\ &= (|0\rangle + \exp(-iE_0t)|1\rangle)|\psi_0\rangle, \end{aligned} \quad (5.5)$$

where the controlled unitary evolution $\exp(-i\hat{H}t)$ is the Trotterized molecular Hamiltonian applied for time t , which applied to the ground state $|\psi_0\rangle$ produces a phase proportional to the ground state energy E_0 only if the ancilla is in state $|1\rangle$. After the final Hadamard, the probability of measuring the ancilla qubit to be 0 is $P_0 = \cos^2(\pi E_0 t)$.

One could imagine performing N measurements to determine P_0 with accuracy $1/\sqrt{N}$; however, the accuracy of the measurement of $E_0 t$ would require an exponential number of measurements in terms of the number of bits of precision required. Instead, the iterative approach uses a single measurement⁴ of the ancilla to determine whether the first bit (i.e. most significant bit) of the binary representation of $E_0 t$ is 0 or 1. The experiment is then repeated, with the Hamiltonian applied for duration $2t$. An additional Z rotation on the ancilla based on the first measurement subtracts the first bit of the energy eigenvalue; the second measurement then measures the second bit of $E_0 2t$. By doubling the Hamiltonian duration and feeding back the previous measurement, each subsequent measurement reads out the next bit of the energy eigenvalue⁵.

We now note that the PEA is an efficient algorithm; that is, the resources required scale at most polynomially with system size. For a given error threshold (e.g. chemical accuracy), the implementation of the Hamiltonian has been shown to be efficient in terms of the number of gates required [2], and the phase estimation algorithm is known to be efficient [85]. As discussed above, the encoding scheme (e.g. the Jordan-Wigner

⁴In the actual experiment, we use a majority voting scheme: the experiment is repeated many (~ 1000) times and the majority of the measurements determines 0 or 1. This allows for a certain amount of experimental error, provided the “true” value of P_0 is not too close to 50%.

⁵The iterative phase estimation algorithm has previously been presented (e.g. in [35] and [136]) as reading out the phase in reverse, from least- to most-significant bit. Both approaches are valid, but we find the most-significant bit first approach to be simpler.

or Bravyi-Kitaev transforms) is requires a polynomial number of gates, and extension to larger molecules requires (more orbitals) only a linear increase in the number of qubits.

Details of our implementation of PEA are given in Section 6.2 and Appendix E.3; we end this section by mentioning a few important qualifications to the simple algorithm described above. First, we have assumed that the register qubits are prepared in the ground state of the Hamiltonian, $|\psi_0\rangle$. In fact, it is necessary only that the state of the register has significant overlap with the true ground state. In principle, with an error-free quantum computer, one could repeat the PEA experiment until the measurement projects the register into the ground state; that is, repeat the experiment and take the lowest measured energy. However, in a device with errors this is not possible. In the case of the hydrogen molecule, the Hartree-Fock ground state overlap with the true ground state is large ($\langle\psi_0|\psi_{HF}\rangle > 0.5$), allowing us to use a majority voting scheme. For other systems—particularly ones with strong electron–electron correlations for which classical theoretical chemistry methods work poorly—the overlap of a standard approximate ground state (e.g. from Hartree-Fock) with the true ground state may be exponentially small. While other methods, such as adiabatic state preparation, have been proposed [136, 56], it is not clear that these methods will scale favorably with the number of electrons simulated. This is not surprising, however, as finding the ground state energy of an arbitrary Hamiltonian is known to be in the QMA complexity class, the quantum analog of NP [59].

The main disadvantage of the PEA is that it requires long, coherent evolution, and the effect of gate errors is not well understood. Recent proposals of more efficient phase estimation algorithms using Bayesian techniques [118, 137] reduce the resource requirements and even allow the possibility of learning about error mechanisms in the quantum device, but these have yet to be tested experimentally. Instead, we now describe a different algorithm more suitable for imperfect devices with limited coherence times, the Variational Quantum Eigensolver.

5.3 The Variational Quantum Eigensolver

The Variational Quantum Eigensolver (VQE), first proposed and demonstrated on an optical system in 2013 [90], is a hybrid algorithm that uses shallow quantum circuits at the cost of an increased number of repetitions to variationally find the ground state of a problem Hamiltonian. The basic idea is straightforward: use a parameterized ansatz to approximate the ground state of a problem Hamiltonian on the quantum computer, and then measure each of the terms of the Hamiltonian and classically compute the energy. This energy is fed to a (classical) minimization algorithm, which then proposes new parameter values. The process is repeated to minimize the energy.

For the VQE to be scalable, its constituent parts must be scalable. The classical preparation (see Appendix E.1) and fermion-to-qubit mapping (see Section 5.1) are the same as used for the PEA. The choice of ansatz used is key to the performance of the VQE; in this work, we use the unitary coupled cluster (UCC) ansatz. There is no efficient method of computing it on a classically, but efficient quantum preparations have been shown, which is evidence that the UCC may allow for quantum speedup with VQE. For details of the UCC ansatz, see Appendix E.4. Finally, the number of measurements necessary to compute the energy must scale efficiently. For chemical Hamiltonians, this has been shown to be the case [79].

One additional property of the VQE that is particularly attractive for current quantum devices is its potential robustness to certain systematic errors. This is due to the variational nature of the algorithm: the parameters that the classical minimizer determines as optimal will necessarily reflect any errors in the implementation of the ansatz, as long as they do not take the system out of the variational subspace. In our experiment in particular, the single variational parameter is the magnitude of a Z-rotation. Phase noise throughout the experiment has the experiment of an unwanted Z-rotation;

therefore, the optimization procedure automatically compensates for any correlated dephasing. This property of the VQE is what allows our experiment to measure the dissociation energy of hydrogen to within chemical accuracy. Additionally, before the advent of a fully error-corrected quantum computer there will be devices with more qubits than can be simulated on a classical supercomputer⁶; its shallow circuits and resistance to some systematic errors means that the VQE could provide industrially useful, classically incomputable results.

Finally, the variational nature of the VQE also means that it is not guaranteed to find the true ground state energy. While this is a disadvantage compared to the PEA, to be practically useful we must only find a ground state energy that is lower than what is possible to compute with classical methods. Furthermore, the VQE may be able to efficiently prepare an input state for the PEA for interesting systems where classical chemistry fails. The ultimate quantum algorithm for chemistry may therefore be a combination of the VQE and PEA.

⁶In fact, possibly quite soon.

Chapter 6

Scalable Quantum Simulation of Molecular Energies¹

Universal and efficient simulation of physical systems [70] is among the most compelling applications of quantum computing. In particular, quantum simulation of molecular energies [2], which enables numerically exact prediction of chemical reaction rates, promises significant advances in our understanding of chemistry and could enable *in silico* design of new catalysts, pharmaceuticals and materials. As scalable quantum hardware becomes increasingly viable [11, 58, 30, 97, 12], chemistry simulation has attracted significant attention [136, 55, 28, 133, 47, 91, 7, 14, 135, 127, 128, 120, 6, 90, 142, 78, 132, 79, 119, 4, 5, 122, 82, 43] since classically intractable molecules require a relatively modest number of qubits and because solutions have commercial value associated with their chemical applications [83].

The fundamental challenge in building a quantum computer is realizing high-fidelity operations in a scalable architecture [77]. Superconducting qubits have made rapid progress in recent years [11, 58, 30, 97] and can be fabricated in microchip foundries

¹This chapter has been submitted as: “Scalable Quantum Simulation of Molecular Energies”, Peter O’Malley, Ryan Babbush, et al. (2016).

and manufactured at scale [52]. Recent experiments have shown logic gate fidelities at the threshold required for quantum error correction [11] and dynamical suppression of bit-flip errors [58]. Here, we use the device reported in [58, 12, 13] to implement and compare two quantum algorithms for chemistry.

Our first experiment demonstrates the recently-proposed variational quantum eigensolver (VQE), introduced in [90]. Our VQE experiment achieves chemical accuracy and is the first scalable quantum simulation of molecular energies performed on quantum hardware, in the sense that our algorithm is efficient and does not benefit from exponentially costly precompilation [114]. When implemented using a unitary coupled cluster ansatz, VQE cannot be efficiently simulated classically and empirical evidence suggests that answers are accurate enough to predict chemical rates [90, 142, 132, 78, 79]. Because VQE only requires short state preparation and measurement sequences, it has been suggested that classically intractable computations might be possible using VQE without the overhead of error correction [79, 132]. Our experiments substantiate this notion by providing clear evidence that VQE is robust to systematic errors.

Our second experiment realizes the original algorithm for the quantum simulation of chemistry, introduced in [2]. This approach involves Trotterized simulation [121] and the quantum phase estimation algorithm (PEA) [60]. We experimentally perform this entire algorithm, including both key components, for the first time. While PEA has asymptotically better scaling in terms of precision than VQE, long and coherent gate sequences are required for its accurate implementation.

Several prior experiments have demonstrated subroutines of quantum chemistry simulations but none have shown a full algorithm that is scalable at the logical level. The phase estimation component of the canonical quantum chemistry algorithm has been demonstrated in a photonic system [69], a nuclear magnetic resonance system [36], and a nitrogen-vacancy center system [131]. While all three experiments obtained molecular

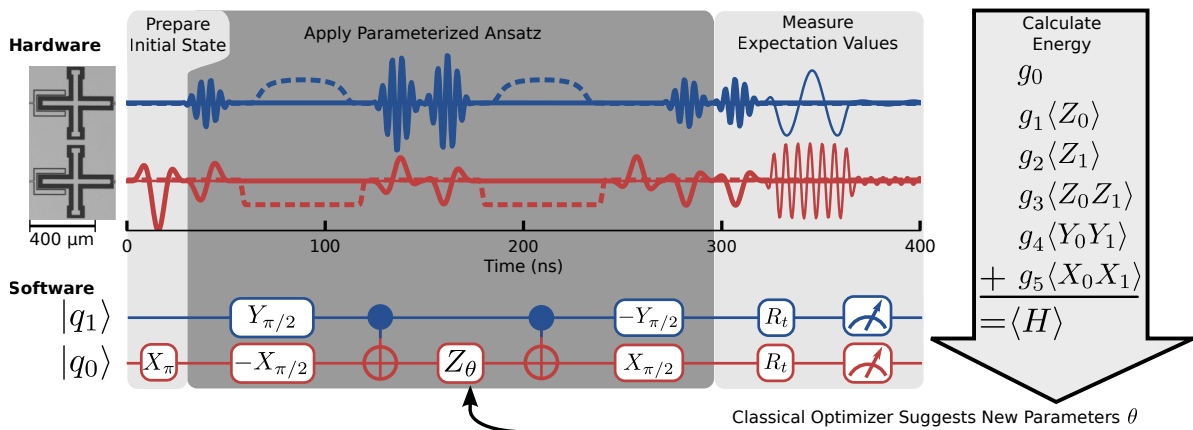


Figure 6.1: **Hardware and software schematic of the variational quantum eigensolver.** (Hardware) micrograph shows two Xmon transmon qubits and microwave pulse sequences to perform single-qubit rotations (thick lines), DC pulses for two-qubit entangling gates (dashed lines), and microwave spectroscopy tones for qubit measurements (thin lines). (Software) quantum circuit diagram shows preparation of Hartree-Fock state, followed by application of the unitary coupled cluster ansatz in Eq. (6.3) and efficient partial tomography (R_t) to measure the expectation values in Eq. (6.1). Finally, the total energy is computed according to Eq. (6.4) and provided to a classical optimizer which suggests new parameters.

energies to incredibly high precision, none of the experiments implemented the propagator in a scalable fashion (e.g. using Trotterization) as doing so requires long coherent evolutions. There have been two previous experimental demonstrations of VQE: first in a photonic system [90] and later in an ion trap [109]. Both experiments validated the variational approach and the latter implemented an ansatz based on unitary coupled cluster. All prior experiments focused on either molecular hydrogen [69, 36] or helium hydride [90, 131, 109] but none of these prior experiments employed a scalable qubit representation such as second quantization. Instead, all five prior experiments represent the Hamiltonian in a configuration basis that cannot be efficiently decomposed as a sum of local Hamiltonians, and then exponentiate this exponentially large matrix as a classical preprocessing step [69, 36, 131, 90, 109].

Until this work, important aspects of scalable chemistry simulation such as the

Jordan-Wigner transformation [115] or the Bravyi-Kitaev transformation [21, 106] had never been used to represent a molecule in an experiment. In both experiments presented here, we simulate the dissociation of molecular hydrogen in the minimal basis of Hartree-Fock orbitals, represented using the Bravyi-Kitaev transformation of the second quantized molecular Hamiltonian [120]. As shown in Appendix E.1, the molecular hydrogen Hamiltonian can be scalably written as

$$H = g_0 1 + g_1 Z_0 + g_2 Z_1 + g_3 Z_0 Z_1 + g_4 X_0 X_1 + g_5 Y_0 Y_1 \quad (6.1)$$

where $\{X_i, Z_i, Y_i\}$ denote Pauli matrices acting on the i^{th} qubit and the real scalars $\{g_\gamma\}$ are efficiently computable functions of the hydrogen-hydrogen bond length, R . The ground state energy of Eq. (6.1) as a function of R defines an energy surface. Such energy surfaces are used to compute chemical reaction rates which are exponentially sensitive to changes in energy. At room temperature, a relative error in energy of 1.6×10^{-3} Hartree (1 kcal/mol or 0.043 eV) translates to a chemical rate that differs from the true value by an order of magnitude; therefore, 1.6×10^{-3} Hartree is known as “chemical accuracy” [48]. Our goal then is to compute the lowest energy eigenvalues of Eq. (6.1) as a function of R , to within chemical accuracy.

6.1 Variational quantum eigensolver

Many popular classical approximation methods for the electronic structure problem involve optimizing a parameterized guess wavefunction (known as an “ansatz”) according to the variational principle [48]. If we parameterize an ansatz $|\varphi(\vec{\theta})\rangle$ by the vector $\vec{\theta}$ then the variational principle holds that

$$\frac{\langle \varphi(\vec{\theta}) | H | \varphi(\vec{\theta}) \rangle}{\langle \varphi(\vec{\theta}) | \varphi(\vec{\theta}) \rangle} \geq E_0, \quad (6.2)$$

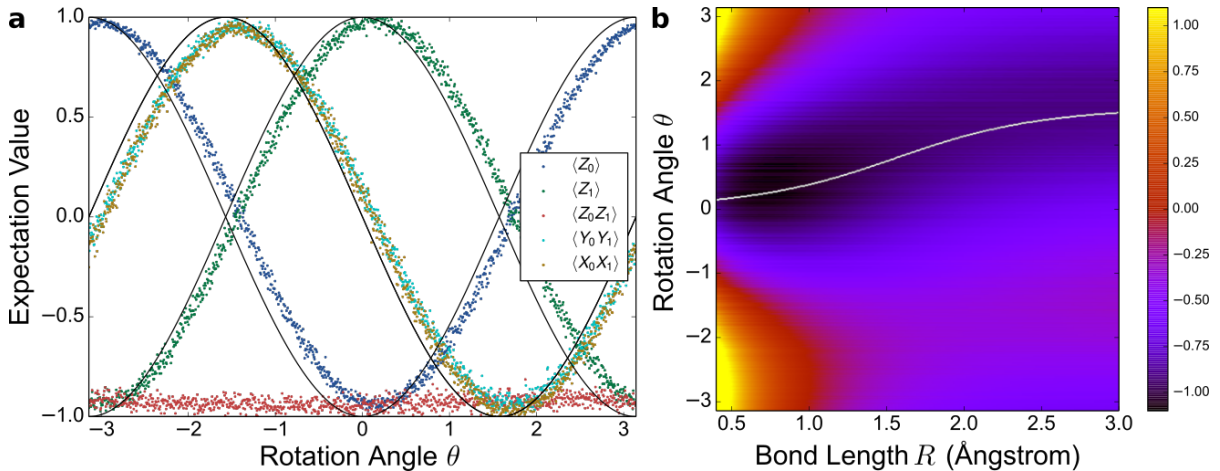


Figure 6.2: **Variational quantum eigensolver: raw data and computed energy surface.** (a) Data showing the expectation values of terms in Eq. (6.1) as a function of θ , as in Eq. (6.3). Black lines nearest to the data show the theoretical values. While such systematic phase errors would prove disastrous for PEA, our VQE experiment is robust to this effect. (b) Experimentally measured energies (in Hartree) as a function of θ and R . This surface is computed from Figure 6.2a according to Eq. (6.4). The white curve traces the theoretical minimum energy; the values of theoretical and experimental minima at each R are plotted in Figure 6.3a. Errors in this surface are plotted in Appendix E.2 in Figure E.2.

where E_0 is the smallest eigenvalue of the Hamiltonian H . Accordingly, E_0 can be estimated by selecting the parameters $\vec{\theta}$ which minimize the left-hand side of Eq. (6.2).

While the ground state wavefunction is likely to be in superposition over an exponential number of states in the basis of molecular orbitals, most classical approaches restrict the ansatz to the support of polynomially many basis elements due to memory limitations. However, quantum circuits can prepare entangled states which are not known to be efficiently representable classically. In VQE, the state $|\varphi(\vec{\theta})\rangle$ is parameterized by the action of a quantum circuit $U(\vec{\theta})$ on an initial state $|\phi\rangle$, i.e. $|\varphi(\vec{\theta})\rangle \equiv U(\vec{\theta})|\phi\rangle$. Even if $|\phi\rangle$ is a simple product state and $U(\vec{\theta})$ is a very shallow circuit, $|\varphi(\vec{\theta})\rangle$ can contain complex many-body correlations and span an exponential number of standard basis states.

We can express the mapping $U(\vec{\theta})$ as a concatenation of parameterized quantum gates, $U_1(\theta_1)U_2(\theta_2)\cdots U_n(\theta_n)$. In this work, we parameterize our circuit according to unitary

coupled cluster theory [142, 79, 132]. As described in Appendix E.4, unitary coupled cluster predicts that the ground state of Eq. (6.1) can be expressed as

$$|\varphi(\theta)\rangle = e^{-i\theta X_0 Y_1} |01\rangle, \quad (6.3)$$

where $|\phi\rangle = |01\rangle$ is the Hartree-Fock (mean-field) state of molecular hydrogen in the representation of Eq. (6.1). The gate model circuit that performs this unitary mapping is shown in the software section of Figure 6.1.

VQE solves for the parameter vector $\vec{\theta}$ with a classical optimization routine. First, one prepares an initial ansatz $|\varphi(\vec{\theta}_0)\rangle$ and then estimates the ansatz energy $E(\vec{\theta}_0)$ by measuring the expectation values of each term in Eq. (6.1) and summing these values together as

$$E(\vec{\theta}) = \sum_{\gamma} g_{\gamma} \langle \varphi(\vec{\theta}) | H_{\gamma} | \varphi(\vec{\theta}) \rangle, \quad (6.4)$$

where the g_{γ} are scalars and the H_{γ} are local Hamiltonians as in Eq. (6.1). The initial guess $\vec{\theta}_0$ and the corresponding objective value $E(\vec{\theta}_0)$ are then fed to a classical greedy minimization routine (e.g. gradient descent), which then suggests a new setting of the parameters $\vec{\theta}_1$. The energy $E(\vec{\theta}_1)$ is then measured and returned to the classical outer loop. This continues for m iterations until the energy converges to a minimum value $E(\vec{\theta}_m)$ which represents the VQE approximation to E_0 .

Because our experiment requires only a single variational parameter, as in Eq. (6.3), we elected to scan a thousand different values of $\theta \in [-\pi, \pi)$ in order to obtain expectation values which define the entire energy landscape. We did this to simplify the classical feedback routine but at the cost of needing slightly more experimental trials. These expectation values are shown in Figure 6.2a and the corresponding energy surfaces at different bond lengths are shown in Figure 6.2b. The energy surface in Figure 6.2b was locally optimized at each bond length to emulate an on-the-fly implementation.

Figure 6.3a shows the exact and experimentally determined energies of molecular

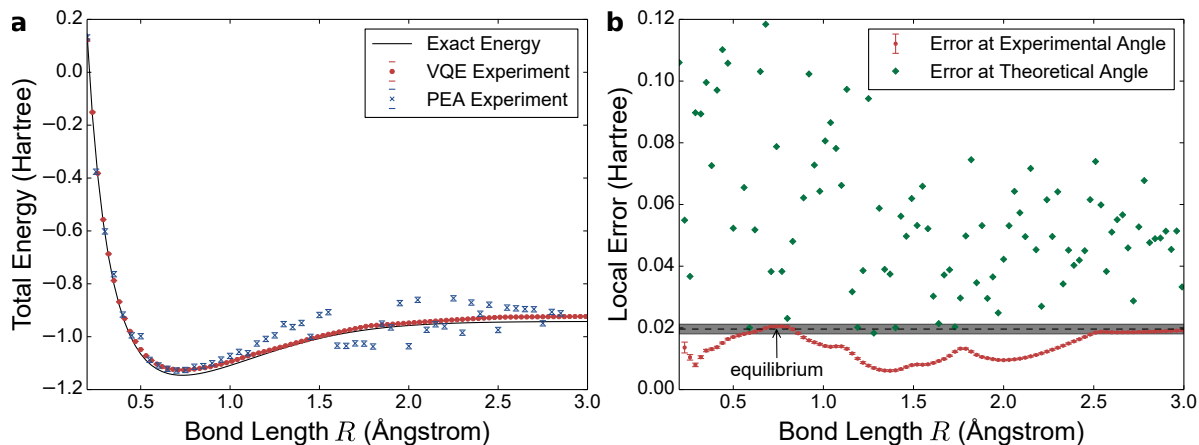


Figure 6.3: **Computed H_2 energy curve and errors.** (a) Energy surface of molecular hydrogen as determined by both VQE and PEA. VQE approach shows dissociation energy error of $(8 \pm 5) \times 10^{-4}$ Hartree (error bars on VQE data are smaller than markers). PEA approach shows dissociation energy error of $(1 \pm 1) \times 10^{-2}$ Hartree. (b) Errors in VQE energy surface. Red dots show error in the experimentally determined energies. Green diamonds show the error in the energies that would have been obtained experimentally by running the circuit at the theoretically optimal θ instead of the experimentally optimal θ . The discrepancy between blue and red dots is evidence for the robustness of VQE. The gray band encloses the chemically accurate region relative to the experimental energy of the atomized molecule. The dissociation energy is relative to the equilibrium geometry, which falls within this envelope.

hydrogen at different bond lengths. The minimum energy bond length corresponds to the equilibrium bond length, $R = 0.741 \text{ \AA}$, whereas the asymptote on the right part of the curve corresponds to dissociation into two hydrogen atoms. The energy difference between these points is the dissociation energy, and the exponential of this quantity determines the chemical dissociation rate. Our VQE experiment correctly predicts this quantity with an error of $(8 \pm 5) \times 10^{-4}$ Hartree, which is below the chemical accuracy threshold. Error bars are computed with Gaussian process regression [19] which interpolates the energy surface based on local errors from the shot-noise limited expectation value measurements in Figure 6.2a.

Errors in our simulation as a function of R are shown in Figure 6.3b. The curve in Figure 6.3b becomes nearly flat past $R = 2.5 \text{ \AA}$ because the same angle is experimentally

chosen for each R past this point. Note that the experimental energies are always greater than or equal to the exact energies due to the variational principle. Figure 6.3b shows that VQE has substantial robustness to systematic errors. While this possibility had been previously hypothesized [79], we report the first experimental signature of robustness. By performing (inefficient) classical simulations of the circuit in Figure 6.1, we identify the theoretically optimal value of θ at each R . In fact, for this system, at every value of R there exists θ such that $E(\theta) = E_0$. However, due to experimental error, the theoretically optimal value of θ differs substantially from the experimentally optimal value of θ . This can be seen in Figure 6.3b from the large discrepancy between the green diamonds (experimental energy errors at theoretically optimal θ) and the red dots (experimental energy errors at experimentally optimal θ). The experimental energy curve at theoretically optimal θ shows an error in the dissociation energy of 1.1×10^{-2} Hartree, which is more than an order of magnitude worse. One could anticipate this discrepancy by looking at the raw data in Figure 6.2a which shows that the experimentally measured expectation values deviate considerably from the predictions of theory. In a sense, the green diamonds in Figure 6.3b show the performance of a non-variational algorithm, which in theory gives the exact answer, but in practice fails due to systematic errors.

6.2 Phase estimation algorithm

We also report an experimental demonstration of the original quantum algorithm for chemistry [2]. Similar to VQE, the first step of this algorithm is to prepare the system register in a state having good overlap with the ground state of the Hamiltonian H . In our case, we begin with the Hartree-Fock state, $|\phi\rangle$. We then evolve this state under H using a Trotterized approximation to the time-evolution operator. The execution of this unitary is controlled on an ancilla initialized in the superposition state $(|0\rangle + |1\rangle)/\sqrt{2}$.

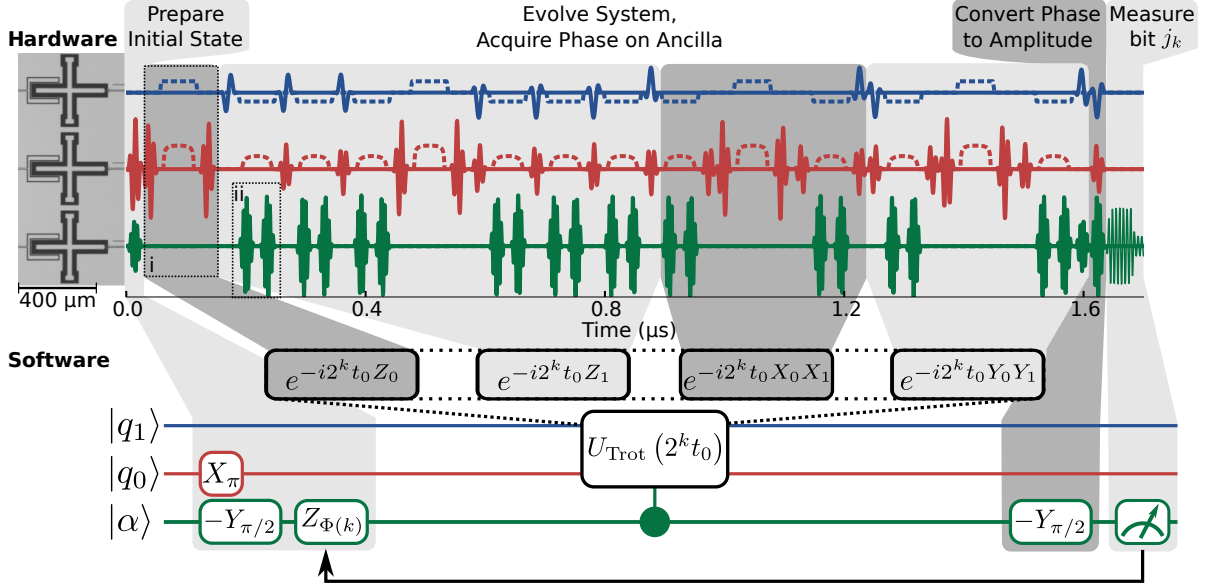


Figure 6.4: **Hardware and software schematic of the Trotterized phase estimation algorithm.** (Hardware) a micrograph shows three Xmon transmon qubits and microwave pulse sequences, including (i) the variable amplitude CZ_ϕ (not used in Figure 6.1) and (ii) dynamical decoupling pulses not shown in logical circuit. (Software) state preparation includes putting the ancilla in a superposition state and compensating for previously measured bits of the phase using the gate Z_{Φ_k} (see text). The bulk of the circuit is the evolution of the system under a Trotterized Hamiltonian controlled by the ancilla. Bit j_k is determined by a majority vote of the ancilla state over one thousand repetitions.

The time-evolution operator can be approximated using Trotterization [121] as

$$e^{-iHt} = e^{-it \sum_\gamma g_\gamma H_\gamma} \approx U_{\text{Trot}}(t) \equiv \left(\prod_\gamma e^{-ig_\gamma H_\gamma t/\rho} \right)^\rho \quad (6.5)$$

where the H_γ are local Hamiltonians as in Eq. (6.1) and the error in this approximation depends linearly on the time step ρ^{-1} [121]. Application of the propagator induces a phase on the system register so that

$$e^{-iHt}|\phi\rangle = \left(\sum_n e^{-iE_n t} |n\rangle\langle n| \right) |\phi\rangle = \sum_n a_n e^{-iE_n t} |n\rangle \quad (6.6)$$

where $|n\rangle$ are eigenstates of the Hamiltonian such that $H|n\rangle = E_n|n\rangle$ and $a_n = \langle n|\phi\rangle$. By controlling this evolution on the ancilla superposition state, one entangles the system

register with the ancilla. Accordingly, by measuring the phase between the $|0\rangle$ state and $|1\rangle$ state of the ancilla, one measures the phase $E_n t$ and collapses the system register to the state $|n\rangle$ with probability $|a_n|^2$.

Our PEA implementation is based on a modification of Kitaev’s iterative phase estimation algorithm [136, 60]. The circuit used is shown in Figure 6.4 and detailed descriptions of the subroutines used to control $U_{\text{Trrot}}(2^k t_0)$ on an ancilla are shown in Appendix E.3. The rotation $Z_{\Phi(k)}$ in Figure 6.4 feeds back classical information from the prior $k - 1$ measurements using phase kickback as

$$\Phi(k) = \pi \sum_{\ell=0}^{k-1} \frac{j_\ell}{2^{\ell-k+1}}. \quad (6.7)$$

With iterative phase estimation, one measures the phase accumulated on the system one bit at a time. Even when a_0 is very small, one can use iterative phase estimation to measure eigenvalues if the system register remains coherent throughout the entire phase determination. Since the Hartree-Fock state has strong overlap with the ground state of molecular hydrogen (i.e. $|\langle 0|\phi\rangle|^2 > 0.5$) we were able to measure each bit independently with a majority-voting scheme, reducing coherence requirements. For b bits, the ground state energy is digitally computed as a binary expansion of the measurement outcomes,

$$E_0^b = -\frac{\pi}{t_0} \sum_{k=0}^{b-1} \frac{j_k}{2^{k+1}}. \quad (6.8)$$

Experimentally computed energies are plotted alongside VQE results in Figure 6.3a. Because energies are measured digitally in iterative phase estimation, the experimentally determined PEA energies in Figure 6.3a agree exactly with theoretical simulations of Figure 6.4, which differ from the exact energies due to the approximation of Eq. (6.5). The primary difficulty of the PEA experiment is that the controlled application of $U_{\text{Trrot}}(2^k t_0)$ requires complex quantum circuitry and long coherent evolutions. Accordingly, we ap-

proximated the propagator in Eq. (6.5) using a single Trotter step ($\rho = 1$), which is not sufficient for chemical accuracy. Our PEA experiment shows an error in the dissociation energy of $(1 \pm 1) \times 10^{-2}$ Hartree.

In addition to taking only one Trotter step, we performed classical simulations of the error in Eq. (6.5) under different orderings of the H_γ in order to find the optimal Trotter sequences at each value of R . The Trotter sequences used in our experiment as well as parameters such as t_0 are reported in Appendix E.3. Since this optimization is intractable for larger molecules, our PEA protocol benefited from inefficient classical preprocessing (unlike our VQE implementation). Nevertheless, this is the first time the canonical quantum algorithm for chemistry has been executed in its entirety and as such, represents a significant step towards scalable implementations.

6.3 Experimental Methods

Both algorithms are implemented with a superconducting quantum system based on the Xmon [10], a variant of the planar transmon qubit [63], in a dilution refrigerator with a base temperature of 20 mK. Each qubit consists of a SQUID (superconducting quantum interference device), which provides a tunable nonlinear inductance, and a large X-shaped capacitor; qubit frequencies are tunable up to 6 GHz and have a nonlinearity of $(\omega_{21} - \omega_{10}) = -0.22$ GHz. The qubits are capacitively coupled to their nearest neighbors in a linear chain pattern, with coupling strengths of 30 MHz. Single-qubit quantum gates are implemented with microwave pulses and tuned using closed-loop optimization with randomized benchmarking [57]. Qubit state measurement is performed in a dispersive readout scheme with capacitively coupled resonators at 6.6-6.8 GHz [58]. For details of the device fabrication and conventions for reporting qubit parameters, see [58].

Our entangling operation is a controlled-phase (CZ_ϕ) gate, accomplished by holding

one of the qubits at a fixed frequency while adiabatically tuning the other close to an avoided level crossing of the $|11\rangle$ and $|02\rangle$ states [11]. To produce the correct phase change ϕ , the acquired phase is measured with quantum state tomography versus the amplitude of the trajectory, and the amplitude for any given ϕ is then determined via interpolation [12]. To minimize leakage out of the computation subspace during this operation, we increase the gate duration from the previously used 40 ns to 50 ns, and then shape the pulse trajectory. The CZ_ϕ gate as implemented here has a range of approximately 0.25 to 5.0 rad; for smaller values of ϕ , parasitic interactions with other qubits become nontrivial, and for larger ϕ , leakage is significant. For ϕ outside this range, the total rotation is accomplished with two physical gates. For CZ_ϕ gates with $\phi = \pi$, the amplitude and shape of the trajectory are further optimized with ORBIT [57]. $CZ_{\phi \neq \pi}$ is only necessary in the PEA experiment (see Appendix E.3).

The gates used to implement both VQE and PEA are shown in Appendix E.2 and Appendix E.3, respectively. A single VQE sequence consists of 11 single-qubit gates and two CZ_π gates. A PEA sequence has at least 51 single-qubit gates, four $CZ_{\phi \neq \pi}$ gates, and ten CZ_π gates; more were required when not all ϕ values are within the range that could be performed with a single physical gate.

6.4 Conclusion

We report the use of quantum hardware to experimentally compute the energy landscape of molecular hydrogen using both PEA and VQE. We perform the first experimental implementation of the Trotterized molecular time-evolution operator and then measure energies using PEA. Due to the costly nature of Trotterization, we are able to implement only a single Trotter step, which is not enough to achieve chemical accuracy. By contrast, our VQE experiment achieves chemical accuracy and shows significant robustness to

certain types of error.

The robustness of VQE is partially a consequence of the adaptive nature of the algorithm; the classical outer loop of VQE helps to avoid systematic errors by acting similarly to the calibration loops used to tune individual quantum gates [57]. This minimization procedure treats the energy functional as a black box in that no assumptions are made about the actual circuit ansatz being implemented. Thus, VQE seeks to find the optimal parameters in a fashion that is blind to control errors, such as pulse imperfection, crosstalk and stray coupling in the device. We observe a remarkable increase in precision by using the experimentally optimal parameters rather than the theoretically optimal parameters. This finding inspires hope that VQE may provide solutions to classically intractable problems even without error correction. Additionally, these results motivate future experiments which take “sublogical” hardware calibration parameters, e.g. microwave pulse shapes, as variational parameters.

Appendix A

Perturbative treatment of the Xmon Hamiltonian

Here we give a brief derivation of Eq. (2.9) by treating the Xmon as a first-order perturbation of an LC oscillator.

The Hamiltonian of the Xmon simply consists of the energies of the capacitor and Josephson junction,

$$H = \frac{\hat{Q}^2}{2C} - E_J \cos(2\pi\hat{\Phi}/\Phi_0), \quad (\text{A.1})$$

where C is the total capacitance, the junction energy $E_J = I_0\Phi_0/2\pi$, I_0 is the junction critical current, and Φ_0 is the magnetic flux quantum. The conjugate variables \hat{Q} (charge) and $\hat{\Phi}$ (flux) have the commutation relation $[\hat{\Phi}, \hat{Q}] = i\hbar$. It is common to write the Hamiltonian in terms of the number of Cooper pairs that have tunneled through the junction, $\hat{n} = \hat{Q}/2e$ (e being the electron charge),

$$H = 4E_C\hat{n}^2 - E_J \cos(\hat{\delta}), \quad (\text{A.2})$$

where we have also defined the capacitor energy $E_C = e^2/2C$ and $\hat{\delta} = 2\pi\hat{\Phi}/\Phi_0$. In the

standard Xmon regime, $E_J/E_C \gg 1$, and we expand cosine to fourth power,

$$H = 4E_c n^2 - E_J(1 - \delta^2/2 + \delta^4/24 + \dots). \quad (\text{A.3})$$

We can now treat $H = 4E_c \hat{n}^2 + E_J \hat{\delta}^2/2$ as a standard harmonic oscillator Hamiltonian with a perturbing term $H_1 = -\frac{E_J}{24} \hat{\delta}^4$. The solution to the harmonic oscillator is well known; see [100] for a standard treatment, or Appendix B of [101] for a particularly clear statement of the results, which we now use. The (unperturbed) frequency of the oscillator f_p is

$$f_p = \sqrt{8E_J E_C}/\hbar = \sqrt{8f_J f_C} = \frac{1}{2\pi} \sqrt{\frac{1}{L_J C}}, \quad (\text{A.4})$$

where we introduce f_J and f_C as the junction and capacitor energies in frequency units, and for the last equality we introduce the Josephson inductance $L_J = \Phi_0/2\pi I_0$. We can write $\hat{\delta}$ and H_1 in terms of the harmonic oscillator's raising and lowering operators,

$$\hat{\delta} = \left(\frac{8E_C}{E_J}\right)^{1/4} \frac{a + a^\dagger}{\sqrt{2}} \quad (\text{A.5})$$

$$H_1 = -\frac{E_C}{12}(a + a^\dagger)^4, \quad (\text{A.6})$$

which allows us to solve for the energy shifts with first-order perturbation theory, $E_n^{(1)} = \langle n^{(0)} | H_1 | n^{(0)} \rangle$ (remembering, of course, that $[a, a^\dagger] \neq 0$ when expanding the to the fourth power). This results in the level shifts $\delta E_0^{(1)} = -E_C/4$ and $\delta E_1^{(1)} = -5E_C/4$, leaving us with

$$f_{10} = f_p - f_C. \quad (\text{A.7})$$

We now consider the fact that the Josephson junction is actually a SQUID, with a tunable inductance $L_J(\Phi) = L_{J0}/\cos(\pi\Phi/\Phi_0)$, where Φ is the flux through the SQUID

loop. Switching $L_J \rightarrow L_J(\Phi)$ in Eq. (A.4) and inserting that into Eq. (A.7), we have

$$f_{10}(\Phi) = \frac{1}{2\pi\sqrt{L_{J0}C}} \sqrt{\cos \frac{\pi\Phi}{\Phi_0}} - f_C = \sqrt{8f_J f_C} \sqrt{\cos \frac{\pi\Phi}{\Phi_0}} - f_C, \quad (\text{A.8})$$

which is Eq. (2.9).

Appendix B

Spectral density of phase noise

In this appendix, we state several results useful for calculating the impact of dephasing sources on typical dephasing measurements. These have all been previously derived; for details, see chapter 3 of [31], [76], and [23].

We first note that the spectral density we use $S(f)$ is one-sided and the frequency is given in Hertz (not radians). In the most general case, the mean-squared phase noise $\langle\phi^2(t)\rangle$ is related to the spectral density $S_\lambda(f)$ of a noise source λ by¹

$$\langle\phi^2(t)\rangle = \left(2\pi\frac{df_{10}}{d\lambda}\right)^2 \int_0^\infty df S_\lambda(f)W(f,t), \quad (\text{B.1})$$

where $df_{10}/d\lambda$ is the sensitivity of the qubit frequency to the noise source λ , and $W(f,t)$ is a spectral weight function, which accounts for the sensitivity of the particular experiment to different ranges of frequency noise. For an experiment sensitive to the full noise spectrum (for example, a Ramsey or RB Ramsey experiment), the weight function is

$$W_0(f,t) = \frac{\sin^2(\pi ft)}{(\pi f)^2}. \quad (\text{B.2})$$

¹This is essentially an example of the Wiener-Khinchin theorem.

For a spin echo sequence (a single additional π pulse at time $t/2$), the spectral weight is given by

$$W_{SE}(f, t) = \tan^2\left(\frac{\pi ft}{2}\right) \frac{\sin^2(\pi ft)}{(\pi f)^2}. \quad (\text{B.3})$$

As expected, noise at DC is eliminated and low-frequency noise suppressed by the tangent term. With additional echoing pulses, the low-frequency noise is further suppressed. The exact form of W depends on the exact implementation of the pulse sequence; the supplement of [23] presents one methodology to calculate it. An example given in [76], for a $(2N + 1)$ -pulse spin echo sequence, is

$$W_{SEN}(f, t) = \tan^2\left(\frac{\pi ft}{2N}\right) \frac{\sin^2(\pi ft)}{(\pi f)^2}. \quad (\text{B.4})$$

The noise sensitivity peaks at $f = N/t$; [23] uses this fact to perform noise spectroscopy by varying the number of pulses.

Finally, for a Rabi experiment with Rabi frequency f_r , the spectral weight function is given by

$$W_R(f, t) = \left(\frac{f_r f}{f_r^2 - f^2}\right)^2 \frac{\sin^2(\pi ft)}{(\pi f)^2}. \quad (\text{B.5})$$

Again the noise sensitivity peaks at a particular frequency, f_r ; this is used for spectroscopy in [141].

We now consider different forms of the noise spectrum, $S_\lambda(f)$. For white noise, $S_\lambda(f) = S_{\lambda,0}$. Using this and Eq. (B.2) to integrate Eq. (B.1), we compute the sensitivity of a Ramsey experiment to white noise:

$$\langle \phi^2(t) \rangle = \left(2\pi \frac{df_{10}}{d_\lambda}\right)^2 \frac{S_{\lambda,0}}{2}, \quad (\text{B.6})$$

which is Eq. (2.5). We also note that the same expression holds for echo sequences with a spectral weight given by Eq. (B.3).

For noise that is correlated over a long timescale, the spectral density is $S_f(f) =$

$2\sigma_{qb}^2\delta(f)$, where σ_{qb} is the standard deviation of the qubit frequency. This is the case for very slowly varying qubit frequency; it can be thought of as varying only between repetitions of the experiment. Again using Eq. (B.2) and Eq. (B.1), we find

$$\langle\phi^2(t)\rangle = \sigma_{qb}^2 t^2. \quad (\text{B.7})$$

Similarly, for $1/f$ noise, the frequency spectrum $S_\lambda(f) = S_{1/f,\lambda}/f$ gives

$$\langle\phi^2(t)\rangle = \left(2\pi\frac{df_{10}}{d\lambda}\right)^2 S_{1/f,\lambda} t^2 \ln \frac{0.4007}{f_c t}, \quad (\text{B.8})$$

where we now introduce a low-frequency cutoff f_c because the integral diverges as $f_c \rightarrow 0$. The divergence, though, is logarithmic, meaning that the exact value of f_c is not important; typically the inverse of the total experimental time is used. For the same reason, the logarithmic part of Eq. (B.8) is frequently ignored altogether, and Eq. (B.7) is used instead. For a spin echo sequence, the dephasing from $1/f$ noise is given by

$$\langle\phi^2(t)\rangle = \left(2\pi\frac{df_{10}}{d\lambda}\right)^2 S_{1/f,\lambda} t^2 \ln 2. \quad (\text{B.9})$$

The value of the logarithmic factor can vary between the Ramsey and the spin echo case by more than an order of magnitude.

For telegraph noise, the spectral density is given by

$$S_f(f) = \frac{4(2\pi\Delta f)^2\Gamma_\uparrow\Gamma_\downarrow}{\Gamma_\Sigma((2\pi f)^2 + \Gamma_\Sigma^2)}, \quad (\text{B.10})$$

where Δf is the magnitude of the switching, Γ_\uparrow and Γ_\downarrow are the up and down switching rates, and $\Gamma_\Sigma = \Gamma_\uparrow + \Gamma_\downarrow$. Now using Eq. (2.4), we find

$$\langle\phi^2(t)\rangle = 2\frac{(2\pi\Delta f)^2}{\Gamma_\Sigma}\frac{\Gamma_\uparrow\Gamma_\downarrow}{\Gamma_\Sigma^2}\left(t - \frac{1 - \exp(-\Gamma_\Sigma t)}{\Gamma_\Sigma}\right). \quad (\text{B.11})$$

At short times ($t \ll \Gamma_\Sigma^{-1}$), Eq. (B.11) reduces to that of correlated noise with $T_{\phi 2} = \Gamma_\Sigma/(\sqrt{2\Gamma_\uparrow\Gamma_\downarrow}\pi\Delta f)$; when the experimental duration is less than the switching timescale, the noise becomes just a constant offset from the average frequency. Similarly, for times

much greater than the switching timescale, telegraph noise looks like white noise; and as expected, for $t \ll \Gamma_\Sigma^{-1}$ we see $T_{\phi 1} = \Gamma_\Sigma^3 / [\Gamma_\uparrow \Gamma_\downarrow (2\pi \Delta f)^2]$. For simplicity, if we assume that $\Gamma_\uparrow = \Gamma_\downarrow$, and let the effective switching amplitude $\Delta f_{10} = 2\Delta f \sqrt{\Gamma_\uparrow \Gamma_\downarrow} / \Gamma_\Sigma$, we get

$$\langle \phi^2(t) \rangle = (2\pi \Delta f_{10})^2 T_s [t - T_s(1 - e^{-t/T_s})], \quad (\text{B.12})$$

where $T_s = 1/\Gamma_\Sigma$ is the switching timescale, which is Eq. (2.8). Finally, for a spin echo sequence, the dephasing due to telegraph noise is

$$\langle \phi^2(t) \rangle = 2 \frac{(2\pi \Delta f)^2 \Gamma_\uparrow \Gamma_\downarrow}{\Gamma_\Sigma \Gamma_\Sigma^2} \left(t - \frac{3 + \exp(-\Gamma_\Sigma t) - 4 \exp(-\Gamma_\Sigma t/2)}{\Gamma_\Sigma} \right). \quad (\text{B.13})$$

Finally, we will now give a brief derivation of the Ramsey visibility in terms of $\langle \phi^2(t) \rangle$, Eq. (2.2). For a single Ramsey experiment, the expected value is $V_0 = \exp(i\phi)$, where here ϕ is the deviation of the phase from the mean value. Assuming Gaussian noise, and using $\langle \phi \rangle = 0$, we have

$$\begin{aligned} \langle e^{i\phi} \rangle &= \langle 1 + i\phi - \phi^2/2 + \dots \rangle \\ &= 1 - (1/2)\langle \phi^2 \rangle + \dots \\ &= e^{-\langle \phi^2 \rangle/2}. \end{aligned} \quad (\text{B.14})$$

In the case that $\langle \phi^2(t) \rangle$ is described by white or correlated noise, we therefore have $V(t) = e^{-t/T_{\phi 1}}$ or $V(t) = e^{-(t/T_{\phi 2})^2}$, analogous to the T_1 of energy relaxation.

Appendix C

Ramsey Tomography Oscilloscope data processing

In this appendix we describe the data processing used to extract the frequency noise power spectrum from the time series $f_{10}(t)$.

Prior to computing the power spectral density (PSD), we pre-process the time series $f_{10}(t)$. First, the mean is subtracted. The time series data are taken at regular intervals with time step t_s . Occasionally during the data taking irregular software lag causes a the system to be unable to maintain this spacing; in this case a single data point is skipped. The second pre-processing step is filling in the skipped data points with 0. This may slightly reduce the measured noise, but the fraction of missed data points is typically about 0.1%, so the inaccuracy introduced is less than what would result from taking the discrete Fourier transform of unevenly spaced data. Due to the long durations of these scans—8 hours or greater—sometimes it is necessary to truncate the beginning or end of the time series if there was an interfering event that caused the sample temperature to fluctuate¹.

¹Typically someone will fill cryogenics in the refrigerator without knowing an experiment is underway.

The PSD is then computed by taking the real, discrete Fourier transform of the time series², squaring it, and normalizing by multiplying by $2t_s n$, where n is the number of points in the time series. The PSD is then averaged by binning in log space; the data shown in Chapter 3 use 100 bins logarithmically spaced from 10^{-4} Hz to 1 Hz. We use the same binning across all datasets so they may be easily averaged together and compared. This is the PSD of the frequency noise, $S_f(f)$, used in Section 3.1.

For completeness, we also note that the fits to Eq. (3.1) were done using the binned data with the standard least-squares method, with each bin weighted logarithmically so that each frequency decade contributes equally to the fit. Finally, when averaging PSDs from multiple RTO datasets, the highest and lowest frequency bins are dropped if not all constituent PSDs contributed to them (if, for example, the datasets were of different total length or t_s).

²Specifically, we use the `fft.rfft` function of the `numpy` Python library, which, depending on the specific installation, typically uses a C or FORTRAN library such as `FFTPACK`.

Appendix D

Appendices for Chapter 4

D.1 Theoretical relation of RB error to $\langle \phi^2 \rangle$

In order to determine the effect of various dephasing mechanisms on an RB sequence, we first consider the following simplified model: a single qubit begins in $|\psi_0\rangle = |0\rangle$, then a randomly chosen perfect Clifford rotation C_1 is applied, and then a phase $\phi_{g,n}$ is accumulated by application of a Z rotation to simulate phase noise. The random Clifford and noise gate pair are repeated N times, after which the single Clifford C_r that is the inverse of all the previous Cliffords is applied to rotate back to (nearly) $|0\rangle$ and we measure the probability of error, $P_{err} = |\langle 1|\psi_N\rangle|^2$.

The value of $\phi_{g,n}$ depends on the dephasing model employed. For example, for static dephasing (e.g., a frequency offset), it is constant: $\phi_{g,n} = \phi_{g,st}$. For white noise, $\phi_{g,n}$ is randomly sampled from a symmetric Gaussian distribution. In general, $\phi_{g,n}$ is arbitrary, but we assume $|\phi_{g,n}| \ll 1$. The average square of $\phi_{g,n}$ is denoted $\langle \phi_g^2 \rangle$.

We now consider the “error angle”, $\Delta\phi$, the angular separation of $|\psi_N\rangle$ from $|0\rangle$ in the Bloch sphere picture of a single qubit, noting that $P_{err} = \langle (\Delta\phi/2)^2 \rangle$, assuming $|\Delta\phi| \ll 1$. Because $|\phi_{g,n}| \ll 1$ and N is not too large, after each rotation $|\psi\rangle$ is close to one of the

six axes ($\pm X, \pm Y, \pm Z$), and the angular distance from the axis is $\Delta\phi$. There is a 1/3 chance that the qubit is near the pole (i.e., Z axis) and then the rotation $\phi_{g,n}$ does not change $\Delta\phi$, while with 2/3 probability the qubit is near the equator and $\Delta\phi$ is changed.

For any dephasing model, it is straightforward to see that the evolution of $\Delta\phi$ is essentially a random walk in two dimensions, and that

$$\langle(\Delta\phi)^2\rangle = \frac{2}{3}N\langle\phi_g^2\rangle, \quad (\text{D.1})$$

assuming $N\langle\phi_g^2\rangle \ll 1$. The RB error is then

$$P_{err} = \langle(\Delta\phi/2)^2\rangle = \frac{1}{6}N\langle\phi_g^2\rangle. \quad (\text{D.2})$$

It might be expected that in the static dephasing case—when there are correlated phase contributions—there can be some sort of echoing effect; for example, if a Clifford takes $|\psi\rangle$ to the +Y axis and it is rotated by $\phi_{g,st}$, then if the next Clifford is an X rotation, putting $|\psi\rangle$ near the -Y axis, the following rotation also by $\phi_{g,st}$ will cancel the previous noise rotation. However, when the full set of Clifford rotations is used, there are four rotations that take $|\psi\rangle$ near the -Y axis, and each orients the previous $\Delta\phi$ in a different direction relative to the axis, resulting in equal probability of canceling the previous rotation, doubling it, or moving in one of the two perpendicular directions. The noise accumulated between rotations is therefore uncorrelated with previous or future noise; the Clifford set is error depolarizing. Therefore, Eq. (D.1) and Eq. (D.2) hold regardless of the noise model.

This simplified model has been confirmed with simulation, for both a static and an uncorrelated noise model with $\phi_{g,n} = \pm\phi_g$.

This implies that RB is an effective way to measure dephasing, if the sequence error occurring between the gates is attributable to dephasing. This can be done easily by comparing the sequence fidelity of an RB sequence with interleaved idling time to that

of a reference RB sequence, effectively subtracting out errors due to the Clifford gates themselves—in other words, measuring the fidelity of an idle using interleaved RB, as in [11]. We can therefore measure the dephasing that takes place during an idle, and by varying the length τ of an idle, measure dephasing as a function of time, $\langle\phi^2(\tau)\rangle$ (for brevity we remove the subscript g). With $r_{I(\tau)}$ being the error rate (i.e., error per gate) of an idle, we thus arrive at Eq. (4.1):

$$P_{err}/N = r_{I(\tau)} = \frac{1}{6}\langle\phi^2(\tau)\rangle. \quad (\text{D.3})$$

For completeness, we also mention here the effect of energy relaxation (T_1 decay) on the fidelity of RB sequences. After each Clifford, the qubit state $|\psi\rangle$ is near the equator of the Bloch sphere with probability $2/3$. In this case the probability of the energy relaxation event is $\tau/2T_1$ (we assume $\tau \ll T_1$); such an event moves $|\psi\rangle$ by approximately the angle $\pi/2$ on the Bloch sphere, thus leading to the error probability $1/2$ at the end of the RB sequence. The corresponding contribution to the RB error per gate is $(2/3) \times (\tau/2T_1) \times (1/2) = \tau/6T_1$. With probability $1/6$ the qubit state after a Clifford is close to the North pole (state $|0\rangle$); then there is no energy relaxation. Finally, with probability $1/6$ the qubit state is close to the South pole $|1\rangle$; then the probability of the energy relaxation event is τ/T_1 , which moves the state by approximately the angle π , thus almost certainly leading to the RB error. The corresponding contribution to the RB error per gate is $(1/6) \times (\tau/T_1) \times 1 = \tau/6T_1$. Adding together the two contributions, we arrive at

$$P_{err}/N = \frac{\tau}{3T_1}. \quad (\text{D.4})$$

Since T_1 can be measured independently, the effects of T_1 decay can be calculated and subtracted from the results obtained with RB, much as it can be subtracted from Ramsey visibility decays as well. In our experiment T_1 is relatively large, and therefore

this correction is small.

D.2 Types of phase noise

We now discuss the form of $\langle \phi^2(\tau) \rangle$ for different sources of noise. For completeness, we also show the similar characteristic, $\langle \tilde{\phi}^2(\tau) \rangle$, for the echo sequence of duration τ (with π pulse at $\tau/2$). Most of results discussed here were presented earlier, e.g., in Refs. [31, 76, 140].

The average values $\langle \phi^2(\tau) \rangle$ and $\langle \tilde{\phi}^2(\tau) \rangle$ for the idle and echo sequence, respectively, can be calculated via the spectral density $S(\omega)$ of the qubit frequency fluctuation,

$$\langle \phi^2(\tau) \rangle = \tau^2 \int_0^\infty S(\omega) \left(\frac{\sin(\omega\tau/2)}{\omega\tau/2} \right)^2 \frac{d\omega}{2\pi}, \quad (\text{D.5})$$

$$\langle \tilde{\phi}^2(\tau) \rangle = \tau^2 \int_0^\infty S(\omega) \frac{\sin^4(\omega\tau/4)}{(\omega\tau/4)^2} \frac{d\omega}{2\pi}, \quad (\text{D.6})$$

where $S(\omega)$ is single-sided and the average frequency fluctuation is assumed to be zero.

For the white noise with a flat spectral density, $S(\omega) = S_0$, we find

$$\langle \phi_{\text{white}}^2(\tau) \rangle = \langle \tilde{\phi}_{\text{white}}^2(\tau) \rangle = \frac{S_0}{2} \tau = 2 \frac{\tau}{T_{\phi 1}}, \quad (\text{D.7})$$

where $T_{\phi 1} = 4/S_0$ is the dephasing time due to white noise. Note that the factor of 2 in the last expression cancels when the corresponding visibility of a Ramsey or echo sequence, $\exp(-\tau/T_{\phi 1})$, is calculated.

For noise that is correlated over very long times (very slowly fluctuating qubit frequency), $S(\omega) = 4\pi\sigma_{qb}^2\delta(\omega)$, where σ_{qb} is the standard deviation of the qubit frequency $2\pi f_{10}$. In this case

$$\langle \phi_{\text{corr}}^2(\tau) \rangle = \sigma_{qb}^2 \tau^2 = 2 \left(\frac{\tau}{T_{\phi 2}} \right)^2, \quad \langle \tilde{\phi}_{\text{corr}}^2(\tau) \rangle = 0, \quad (\text{D.8})$$

where $T_{\phi 2} = \sqrt{2}/\sigma_{qb}$ is the Ramsey dephasing timescale due to such correlated noise.

Obviously, in this case there is no dephasing in the echo sequence visibility.

For $1/f$ noise let us use $S(\omega) = \frac{S_{1/f}}{\omega/2\pi}$, then [76, 140]

$$\langle \phi_{1/f}^2(\tau) \rangle = S_{1/f} \tau^2 \ln \frac{0.4007}{f_c \tau}, \quad (\text{D.9})$$

$$\langle \tilde{\phi}_{1/f}^2(\tau) \rangle = S_{1/f} \tau^2 \ln 2, \quad (\text{D.10})$$

where $f_c = \omega_c/2\pi$ is the low-frequency cutoff of the $1/f$ noise (e.g., the inverse of the total duration of the experiment), which is introduced as the lower limit of integration in Eq. (D.5). Note that in Eq. (D.9) we assumed $f_c \tau \lesssim 0.2$. As the log part in Eq. (D.9) varies slowly, typically it is ignored and $1/f$ noise for $\langle \phi^2(\tau) \rangle$ is treated with Eq. (D.8). Note that the factors in Eq. (D.9) and (D.10) are different, resulting in different effective dephasing times $T_{\phi 2}$ for the Ramsey and echo sequences.

Finally, let us consider a telegraph noise, for which the qubit frequency $2\pi f_{10}$ switches between two values separated by $\Delta\omega_{qb}$, with up (down) switching rate of Γ_{\uparrow} (Γ_{\downarrow}). In this case

$$S(\omega) = \frac{4(\Delta\omega_{qb})^2 \Gamma_{\uparrow} \Gamma_{\downarrow}}{\Gamma_{\Sigma}(\omega^2 + \Gamma_{\Sigma}^2)}, \quad \Gamma_{\Sigma} = \Gamma_{\uparrow} + \Gamma_{\downarrow}, \quad (\text{D.11})$$

so using Eqs. (D.5) and (D.6) we obtain

$$\langle \phi_{\text{tel}}^2(\tau) \rangle = 2 \frac{(\Delta\omega_{qb})^2 \Gamma_{\uparrow} \Gamma_{\downarrow}}{\Gamma_{\Sigma}} \frac{1}{\Gamma_{\Sigma}^2} \left(\tau - \frac{1 - e^{-\Gamma_{\Sigma} \tau}}{\Gamma_{\Sigma}} \right), \quad (\text{D.12})$$

$$\langle \tilde{\phi}_{\text{tel}}^2(\tau) \rangle = 2 \frac{(\Delta\omega_{qb})^2 \Gamma_{\uparrow} \Gamma_{\downarrow}}{\Gamma_{\Sigma}} \frac{1}{\Gamma_{\Sigma}^2} \left(\tau - \frac{3 + e^{-\Gamma_{\Sigma} \tau} - 4e^{-\Gamma_{\Sigma} \tau/2}}{\Gamma_{\Sigma}} \right). \quad (\text{D.13})$$

Note that at short time, $\tau \ll \Gamma_{\Sigma}^{-1}$, the effect of the telegraph noise is similar to the effect of the correlated noise with $T_{\phi 2} = \sqrt{2} \Gamma_{\Sigma} / (\sqrt{\Gamma_{\uparrow} \Gamma_{\downarrow}} \Delta\omega_{qb})$, while at long time, $\tau \gg \Gamma_{\Sigma}^{-1}$ it is similar to the effect of white noise with $T_{\phi 1} = \Gamma_{\Sigma}^3 / [\Gamma_{\uparrow} \Gamma_{\downarrow} (\Delta\omega_{qb})^2]$.

Defining the effective switching amplitude as $2\pi \Delta f_{10} = 2\Delta\omega_{qb} \sqrt{\Gamma_{\uparrow} \Gamma_{\downarrow}} / \Gamma_{\Sigma}$ and intro-

ducing notation $T_{\text{sw}} = 1/\Gamma_{\Sigma}$, we can rewrite Eq. (D.12) as

$$\langle \phi_{\text{tel}}^2(\tau) \rangle = (2\pi\Delta f_{10})^2 T_{\text{sw}} [\tau - T_{\text{sw}}(1 - e^{-\tau/T_{\text{sw}}})], \quad (\text{D.14})$$

which is Eq. (4.2). In the case where $\Gamma_{\uparrow} = \Gamma_{\downarrow}$, as we have assumed, $2\pi\Delta f_{10}$ provides a lower bound on $\Delta\omega_{qb}$. Note that the telegraph noise is not Gaussian. Therefore, while the obtained equations can be used to find the RB error per gate, they cannot, strictly speaking, be used to find the visibility of the standard Ramsey and echo sequences. Nevertheless, they can be used approximately if $|\Delta\omega_{qb}|/\min(\Gamma_{\uparrow}, \Gamma_{\downarrow}) \ll 1$, because at short time the accumulated phase shift is small and the assumption of Gaussianity is not needed, while at longer time, when the phase becomes comparable to 1, the probability distribution for the phase becomes Gaussian due to a large number of switching events.

D.3 T_1 , Ramsey, and spin echo fits

The T_1 data are fit to a simple exponential, $P_1(t) = A \exp(t/T_1) + B$, and we find $T_1 = 26.7 \mu\text{s}$. The Ramsey and spin echo envelopes are each fit to a noise model that includes white and correlated components,

$$V(t) = A \exp \left[-t/T_{\phi 1} - (t/T_{\phi 2})^2 \right] + B, \quad (\text{D.15})$$

where $V(t)$ is the Ramsey/echo visibility, t is the length of the idle as shown in Figure 4.1, $T_{\phi 1}$ is the white noise dephasing timescale, $T_{\phi 2}$ is the correlated noise dephasing timescale, and A and B are the result of state preparation and measurement errors. The fit parameters are given below. Note that each of the fits includes the full range of data, from $0 < t < 5.0 \mu\text{s}$ for Ramsey and $0 < t < 12.0 \mu\text{s}$ for Echo.

Table D.1: Ramsey and spin-echo fit parameters.

Sequence	$T_{\phi 1}$ (μs)	$T_{\phi 2}$ (μs)	A	B
Ramsey	6.8	2.8	0.88	0.015
Spin echo	15.1	7.5	0.88	0.021

D.4 Flux noise

Flux noise on this device, plotted in Figure D.1, has been measured over the frequency range $10^{-4} < f < 1$ Hz, using the Ramsey Tomography Oscilloscope (RTO) protocol of repeated frequency measurements as described in [102]. Four measurements were made on this device (open markers), at three different operating points, and then each measurement was binned in log-space, and the binned measurements averaged together (closed squares). This average is fit (solid line) to an aliased $1/f$ and white noise model, given by

$$S_{\phi}(f) = S_{\phi}^*/f^{\alpha} + S_{\phi}^*/(2f_n - f)^{\alpha} + S_{\text{white}}, \quad (\text{D.16})$$

where $S_{\phi}(f)$ is the flux noise power, expressed in $(\mu\Phi_0)^2/\text{Hz}$, f is the noise frequency, α is the slope of the noise (1 for pure $1/f$ noise), S_{ϕ}^* is the flux noise power at 1 Hz, $f_n = 1$ Hz is the Nyquist frequency of the measurement, and S_{white} is the white noise floor. From the fit we extract $S_{\phi}^* = 2.4(\mu\Phi_0)^2$, $\alpha = 0.99$, and $S_{\text{white}} = 9.7(\mu\Phi_0)^2/\text{Hz}$. We attribute the white noise to state preparation and measurement error. The dashed line shows the $1/f$ fit extended to 1 Hz, where the value of the y-intercept is S_{ϕ}^* .

To plot the inferred flux noise contribution in Figure 4.2 and Figure D.2 below, we use Eq. (D.9), with $S_{1/f} = \partial f/\partial\phi \cdot S_{\phi}^*$ taken from the measurements above, and $f_c = 10$ min, the length of the experiment. The value of the log factor of Eq. (D.9) varies from 13 to 7 for $1 < \tau < 450$ ns.

This analysis assumes that the low frequency flux noise measured here can be ex-

trapolated to high frequencies. In Figure 4.2, however, we see that this calculation underestimates the amount of high frequency noise, and furthermore, that the noise is telegraph in nature, not $1/f$.




Figure D.1: Flux noise as measured with RTO [102]

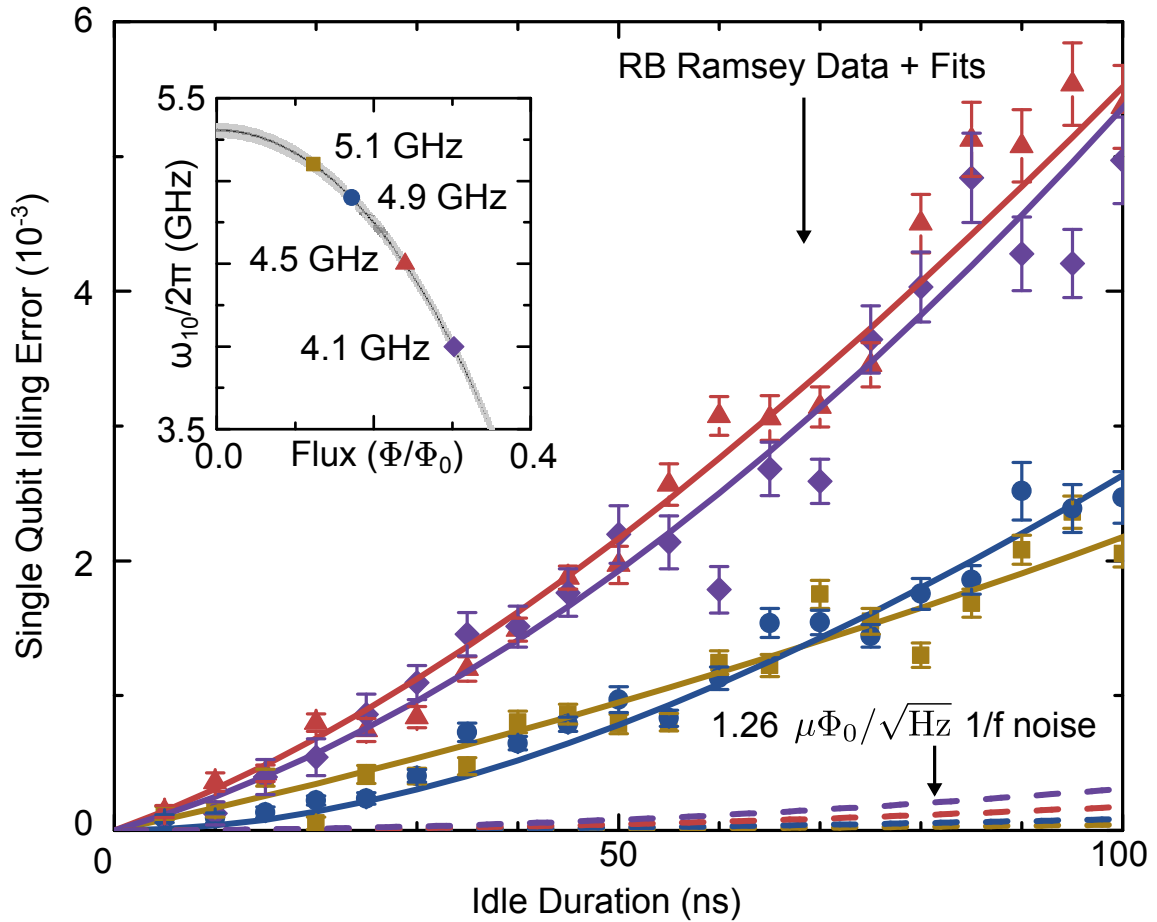


Figure D.2: (color online) RB Ramsey idling error *vs.* duration, for various frequencies; T_1 effects have been subtracted according to Eq. (D.4). The dashed lines denote the inferred contribution from $1/f$ flux noise at the four different operating points. The inset shows frequency spectroscopy *vs.* applied flux, following the expected dependence [63]; the four operating points are shown.

Table D.2: Telegraph fit parameters.

f_{10} (GHz)	$df/d\Phi$ (GHz/ Φ_0)	T_1 (μ s)	T_{ϕ_1} (μ s)	T_{sw} (ns)	Δf_{10} (kHz)
5.1	3.39	30.6	20.6	182000	184
4.9	4.81	26.7	-	84	479
4.5	6.95	31.3	12.4	98	484
4.0	9.23	36.2	15.5	263	469

D.5 RB Ramsey across the qubit spectrum

Figure D.2 shows RB Ramsey measurements at three additional qubit frequencies; the data for the 4.9 GHz operating point are the same as in Figure 4.2. The inset shows the frequency-flux relation for this qubit, with the four operating points denoted by symbols; $df/d\phi$ changes by a factor of 2.7 between the operating points, to explore different susceptibilities to flux noise. The remaining three datasets are fit to a noise model incorporating telegraph and white noise; that is,

$$r_{I(\tau)} = \frac{1}{6} (\langle \phi_{\text{tel}}^2(\tau) \rangle + \langle \phi_{\text{white}}^2(\tau) \rangle) \quad (\text{D.17})$$

[see Eqs. (D.7) and (D.14)]. We show the fit parameters here. We note that at the highest qubit frequency, the large T_{sw} indicates that the telegraph noise model is not needed here and can be replaced by the correlated noise model with $T_{\phi_2} = \sqrt{2}/[2\pi\Delta f_{10}] = 1.2 \mu$ s. The Ramsey data for this frequency, fit to Eq. (D.15), give $T_{\phi_1} = 10.7 \mu$ s and $T_{\phi_2} = 3.6 \mu$ s, which indicates that even though the telegraph dephasing source is not present at this operating point, the dephasing magnitude measured by Ramsey still does not match that found with RB.

Despite tuning the flux Φ/Φ_0 over most of its range, we find that $1/f$ noise does not contribute appreciably to gate errors. For typical gates of length 20ns, idle fidelities greater than 0.999 are seen over the frequency range, demonstrating that tunable qubits

can achieve high fidelity even when biased significantly away from the flux-insensitive operating point.

D.6 Charge noise

To calculate the expected frequency fluctuation due to charge noise, we use Eq. (2.5) from [63]

$$\epsilon_m \simeq (-1)^m E_C \frac{2^{4m+5}}{m!} \sqrt{\frac{2}{\pi}} \left(\frac{E_J}{2E_C} \right)^{\frac{m}{2} + \frac{3}{4}} e^{-\sqrt{8E_J/E_C}}, \quad (\text{D.18})$$

where ϵ_m is the charge dispersion for energy level m , and E_J and E_C are the Josephson energy and charging energy, respectively, of the qubit. Note that we can also write $E_J/E_C \approx (\omega_{01}/\eta - 1)^2/8$ (following from Eq. (2.11)), where $\omega_{01}/2\pi$ is the qubit frequency and the qubit anharmonicity $f_{12} - f_{01} = \eta/2\pi = -215$ MHz. We then calculate ϵ_1 for the two ends of the qubit spectrum; we find $\epsilon_1(\omega_{01}/2\pi = 6 \text{ GHz}) = 3.6 \text{ Hz}$ and $\epsilon_1(\omega_{01}/2\pi = 4 \text{ GHz}) = 14.4 \text{ kHz}$, both of which are far below the measured charge noise fluctuation frequency of $\approx 500 \text{ kHz}$. We also note the qubits used in Ref. [96] have charge noise fluctuations of the same order as the telegraph noise measured here, but charge noise of that magnitude is expected, as explained by the different parameter range of those qubits: $\omega_{01}/2\pi = 4.387 \text{ GHz}$ and $\eta/2\pi = -334 \text{ MHz}$, giving $\epsilon_1 \approx 2 \text{ MHz}$.

D.7 Calculation of Ω_{ZZ}

Two capacitively coupled qubits have an XX-type coupling of the form $g(|01\rangle\langle 10| + |10\rangle\langle 01|)$, where the coupling constant g is half the swap rate between the qubits. The interaction between the higher levels, $\sqrt{2}g(|11\rangle\langle 20| + |02\rangle\langle 11|) + \sqrt{2}g(|11\rangle\langle 20| + |02\rangle\langle 11|)$, results in a repulsion of the $|11\rangle$ level from the $|02\rangle$ and $|20\rangle$ levels; this energy shift in the $|11\rangle$ level produces a ZZ-type interaction between the qubits. In the far-detuned limit,

neglecting the XX-coupling, the two-qubit Hamiltonian becomes

$$H = \omega_1|10\rangle\langle 10| + \omega_2|01\rangle\langle 01| + (\omega_1 + \omega_2 + \Omega_{ZZ})|11\rangle\langle 11|, \quad (\text{D.19})$$

$$\Omega_{ZZ} = \frac{2g^2}{\Delta - \eta_2} + \frac{2g^2}{-\Delta - \eta_1}, \quad (\text{D.20})$$

where ω_n and η_n are the qubit frequencies and nonlinearities, respectively, and $\Delta = \omega_1 - \omega_2$. In our system, $\eta_1 = \eta_2 \equiv \eta$, giving

$$\Omega_{ZZ} = \frac{4g^2\eta}{\Delta^2 - \eta^2}. \quad (\text{D.21})$$

When both qubits are simultaneously performing an RB sequence, phase error ϕ per idle gate in qubit A is

$$\phi = \pm \frac{\Omega_{ZZ}}{2} t_{\text{gate}} \quad (\text{D.22})$$

where t_{gate} is the idle gate duration, and the frequency shift $\pm\Omega_{ZZ}/2$ assumes centering the qubit frequency. This gives $\langle\phi^2\rangle = (\Omega_{ZZ}t_{\text{gate}})^2/4$, and since for RB the error per gate is $E = \langle\phi^2\rangle/6$ [see Eq. (D.3)], we arrive at Eq. (4.5) for the error per gate due to the Ω_{ZZ} interaction,

$$E = \frac{\pi^2}{6} \left(\frac{\Omega_{ZZ}}{2\pi} t_{\text{gate}} \right)^2. \quad (\text{D.23})$$

Table D.3: Gate error fit parameters

Gate	Linear Term (10^{-6} error/ns)	Quadratic Term (10^{-6} error/ns ²)
I	17	0.22
XX	20	-
Z	24	0.18
YX	22	-

D.8 Fits to gate errors in Figure 4.4

For the data in the Figure 4.4, the fits are made either to a simple linear model in the case of Markovian noise (the XX and YX cases) or to a quadratic and linear model in the case of non-Markovian noise (the I and Z cases). There is no offset in any fit. Note that the contribution from $T_1 = 26.7 \mu\text{s}$ to the linear portion of the error, given by Eq. (D.4), is 9.3×10^{-6} error/ns, or roughly half of the error measured. The remainder is equivalent to a white noise dephasing with time constant $T_{\text{white}} \approx 30 \mu\text{s}$, according to Eqs. D.3 and D.7. The quadratic terms correspond with $T_{\phi_2} \approx 1 \mu\text{s}$.

D.9 Telegraph noise measured in other devices

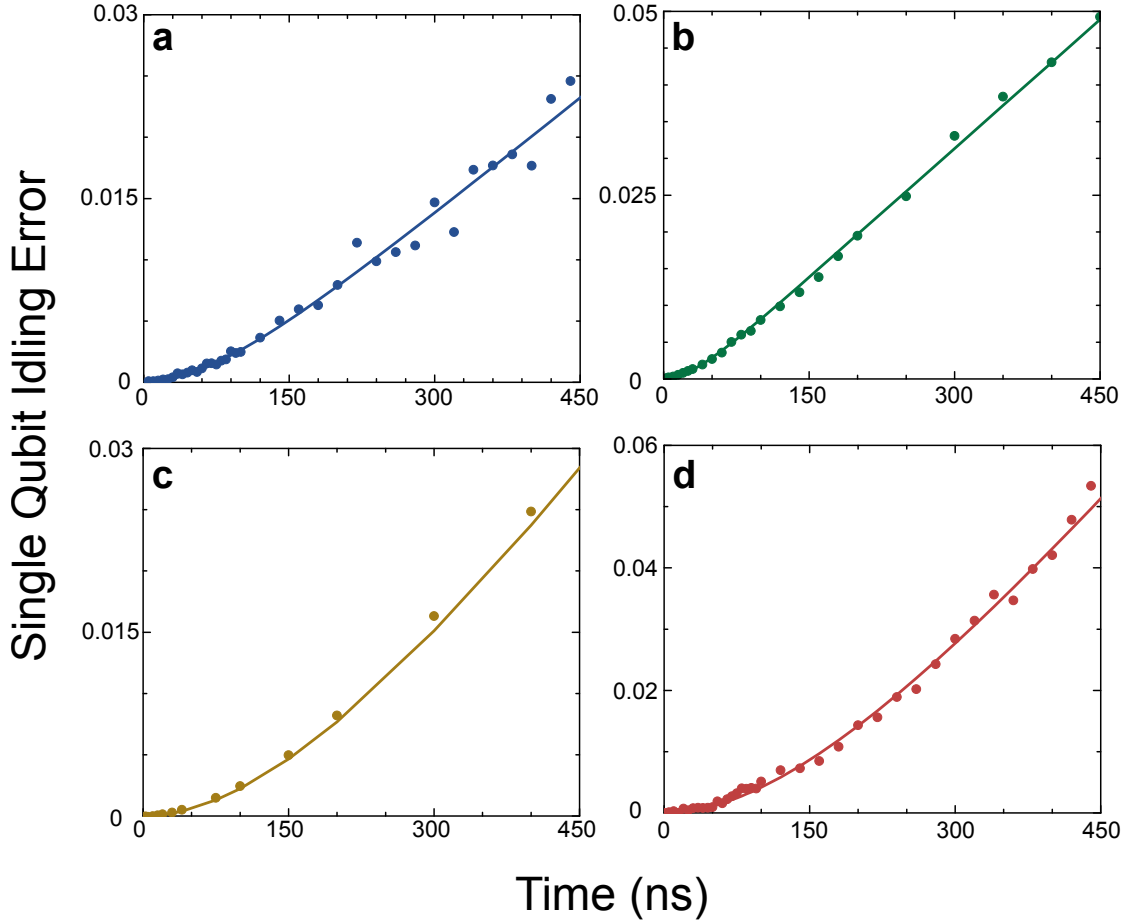


Figure D.3: Telegraph noise measured with RB Ramsey in other devices. All fits included T_1 and telegraph noise only (Eq. (4.4)). (a) A reproduction of Figure 4.2 data for reference. (b) Measurement of another Xmon on the same chip as the device. (c) Measurement of an Xmon qubit from another sample; see [11] for device details. (d) Measurement of a gmon qubit; see [98] for device details.

Telegraph noise has been observed in many other devices. In Figure D.3, we present RB Ramsey measurements of three other devices that show telegraph noise, with the data from Figure 4.2 reproduced for reference (a); one is another device on the same chip (b), one another Xmon with different parameters [11] (c), and the last a gmon qubit [98] (d). All fits were to T_1 and telegraph noise only, Eq. (4.4), with fit parameters given in

Table D.4: Fits for telegraph noise measured in other devices (Figure D.3); see text and references for sample details.

Sample (see text)	f_{10} (GHz)	$df/d\Phi$ (GHz/ Φ_0)	T_1 (μ s)	T_{sw} (ns)	Δf_{10} (kHz)	Device Details
a	4.9	4.81	26.7	84	479	q2 of [58]
b	4.8	5.36	15.7	183	274	q0 of [58]
c	5.5	3.96	22.2	201	199	q2 of [11]
d	4.9	6.62	15.7	32	528	q1 of [98]

Table D.4.

Appendix E

Appendices for Chapter 5

E.1 The electronic structure problem

The central problem of quantum chemistry is to compute the lowest energy eigenvalue of the molecular electronic structure Hamiltonian. The eigenstates of this Hamiltonian determine almost all of the properties of interest in a molecule or material, and as the gap between the ground and first excited state is often much smaller than the thermal energy at room temperature, the ground state is of particular interest. To arrive at the standard form of this Hamiltonian used in quantum computation, one begins from a collection of nuclear charges Z_i and a number of electrons in the system N for which the corresponding Hamiltonian is written

$$\begin{aligned} H = & - \sum_i \frac{\nabla_{R_i}^2}{2M_i} - \sum_i \frac{\nabla_{r_i}^2}{2} - \sum_{i,j} \frac{Z_i}{|R_i - r_j|} \\ & + \sum_{i,j>i} \frac{Z_i Z_j}{|R_i - R_j|} + \sum_{i,j>i} \frac{1}{|r_i - r_j|} \end{aligned} \tag{E.1}$$

where the positions, masses, and charges of the nuclei are R_i , M_i , Z_i , and the positions of the electrons are r_i . Here, the Hamiltonian is in atomic units of energy known as Hartree.

One Hartree is $\hbar^2/m_e e^2 a_0^2$ (630 kcal/mol or 27.2 eV) where m_e , e and a_0 denote the mass of an electron, charge of an electron and Bohr radius, respectively.

This form of the Hamiltonian and its real-space discretization are often referred to as the first quantized formulation of quantum chemistry. Several approaches have been developed for treating this form of the problem on a quantum computer [55]; however, the focus of this work is the second quantized formulation. To reach the second quantized formulation, one typically first approximates the nuclei as fixed classical point charges under the Born-Oppenheimer approximation, chooses a basis ϕ_i in which to represent the wavefunction, and enforces anti-symmetry with the fermion creation and annihilation operators a_i^\dagger and a_j to give

$$H = \sum_{pq} h_{pq} a_p^\dagger a_q + \sum_{pqrs} h_{pqrs} a_p^\dagger a_q^\dagger a_r a_s \quad (\text{E.2})$$

with

$$h_{pq} = \int d\sigma \phi_p^*(\sigma) \left(\frac{\nabla_r^2}{2} - \sum_i \frac{Z_i}{|R_i - r|} \right) \phi_q(\sigma) \quad (\text{E.3})$$

$$h_{pqrs} = \int d\sigma_1 d\sigma_2 \frac{\phi_p^*(\sigma_1) \phi_q^*(\sigma_2) \phi_s(\sigma_1) \phi_r(\sigma_2)}{|r_1 - r_2|} \quad (\text{E.4})$$

where σ_i is now a spatial and spin coordinate with $\sigma_i = (r_i, s_i)$, and the standard anti-commutation relations that determine the action of a_i^\dagger and a_j are $\{a_i^\dagger, a_j\} = \delta_{ij}$ and $\{a_i^\dagger, a_j^\dagger\} = \{a_i, a_j\} = 0$. Finally, the second quantized Hamiltonian must be mapped into qubits for implementation on a quantum device. The most common mappings used for this purpose are the Jordan-Wigner transformation [115] and the Bravyi-Kitaev transformation [21, 106, 120].

Using the Bravyi-Kitaev transformation, the spin Hamiltonian for molecular hydrogen

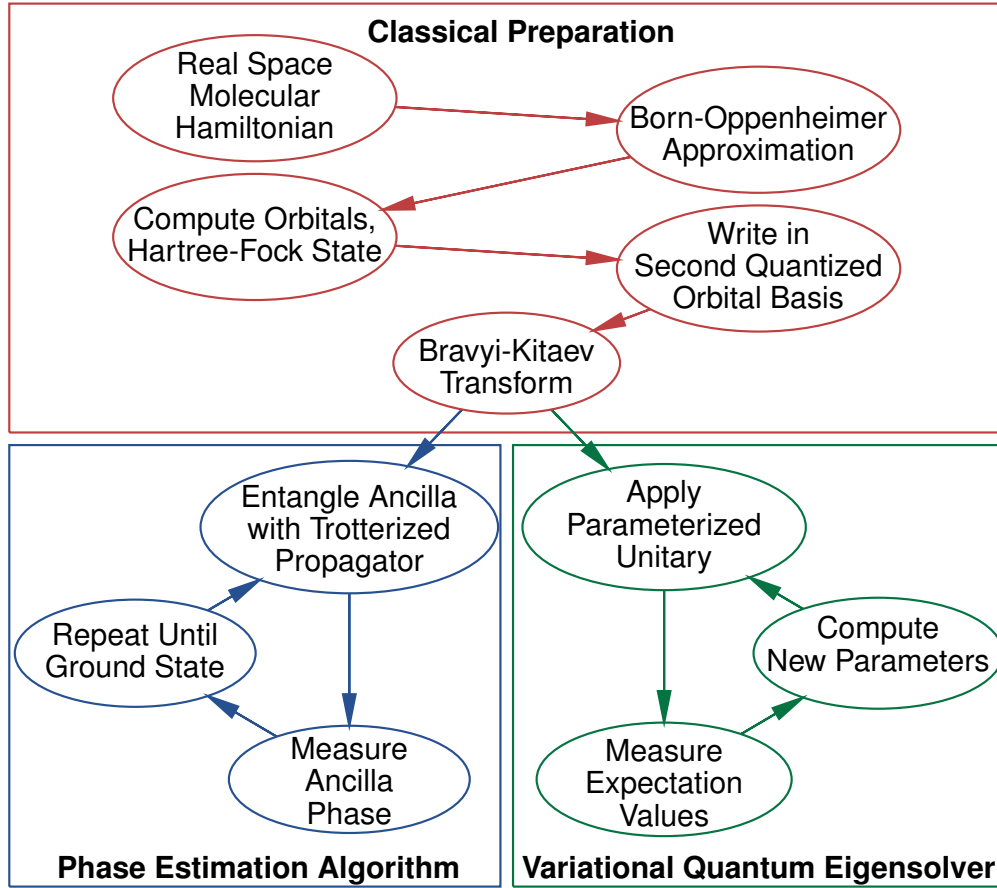


Figure E.1: A flowchart describing steps required to quantum compute molecular energies using both PEA and VQE.

in the minimal (STO-6G) basis, as reported in [106], is given by

$$\begin{aligned}
 H = & f_0 1 + f_1 Z_0 + f_2 Z_1 + f_3 Z_2 + f_1 Z_0 Z_1 \\
 & + f_4 Z_0 Z_2 + f_5 Z_1 Z_3 + f_6 X_0 Z_1 X_2 + f_6 Y_0 Z_1 Y_2 \\
 & + f_7 Z_0 Z_1 Z_2 + f_4 Z_0 Z_2 Z_3 + f_3 Z_1 Z_2 Z_3 \\
 & + f_6 X_0 Z_1 X_2 Z_3 + f_6 Y_0 Z_1 Y_2 Z_3 + f_7 Z_0 Z_1 Z_2 Z_3
 \end{aligned} \tag{E.5}$$

where the values $\{f_i\}$ depend on the fixed bond length of the molecule. We notice that this Hamiltonian acts off-diagonally on only two qubits (the ones having tensor factors

of 0 and 2), those colored in red in Eq. (E.5). Because we start our simulations in the Hartree-Fock state, a classical basis state, we see that the Hamiltonian stabilizes qubits 1 and 3 so that they are never flipped throughout the simulation. We can use this symmetry to scalably reduce the Hamiltonian of interest to the following effective Hamiltonian which acts only on two qubits,

$$\tilde{H} = g_0 1 + g_1 Z_0 + g_2 Z_1 + g_3 Z_0 Z_1 + g_4 X_0 X_1 + g_5 Y_0 Y_1 \quad (\text{E.6})$$

where the values $\{g_i\}$ depend on the fixed bond lengths of the molecule. We further note that the term $Z_0 Z_1$ commutes with all other terms in the Hamiltonian. Since the ground state of the total Hamiltonian certainly has support on the Hartree-Fock state, we know the contribution to the total energy of $Z_0 Z_1$ (it is given by the expectation of those terms with the Hartree-Fock state). Steps to prepare this Hamiltonian are summarized in the upper-half of Figure E.1.

E.2 Experimental methods for VQE

For the VQE experiment, the qubits q_0 and q_1 are used, at 4.49 and 5.53 GHz, respectively, while all the other qubits are detuned to 3 GHz and below. X_π , Y_π , $\pm X_{\pi/2}$, and $\pm Y_{\pi/2}$ gates are 25 ns long, and pulse amplitudes and detunings from f_{10} are optimized with ORBIT; for these parameters, additional pulse shaping (e.g. DRAG) proved unnecessary (see [13] for details of pulse detuning and shaping). The amplitude, trajectory, and compensating single-qubit phases of the CZ_π gate are optimized with ORBIT as well. The duration of the CZ_π is 55 ns, during which the frequency of q_0 is fixed and q_1 is moved. The rotation Z_θ (the adjustable parameter in Eq. (6.3)) is implemented as a phase shift on all subsequent gates. As operated here, q_0 and q_1 have energy relaxation times $T_1 = 62.8$ and $21.4 \mu\text{s}$, and Ramsey decay times $T_2^* = 1.1$ and $1.9 \mu\text{s}$, respectively.

The expectation values used to calculate the energy of the prepared state are mea-

sured with partial tomography; for example, X_1X_0 is measured by applying $Y_{\pi/2}$ gates to each qubit prior to measurement. We emphasize that for chemistry problems, the number of measurements scales polynomially [79]. Readout duration is set to 1000 ns for higher fidelity (compared to [58], where the “measure”/odd-numbered qubits utilized much shorter readout). In addition to discriminating between $|0\rangle$ and $|1\rangle$, higher level qubit states were also measured (called $|2\rangle$ for simplicity). Readout fidelities are typically $>99\%$ for $|0\rangle$, and $\sim 95\%$ for $|1\rangle$ and $|2\rangle$, and measurement probabilities are corrected for readout error. After readout correction, experiments where one of the qubits is measured in $|2\rangle$ are dropped; any probability to be in $|2\rangle$ is set to zero and remaining probabilities are renormalized.

The circuit pulse sequence used to implement the UCC sequence in Eq. (6.3) is shown in Figure 6.1. The experiment is performed in different gauges of the Bravyi-Kitaev transform; these correspond to the $|0\rangle$ ($|1\rangle$) state of q_0 representing the first orbital being unoccupied (occupied) or occupied (unoccupied), and similarly for q_1 representing the parity of the first two orbitals being even (odd) or odd (even). In practice, a gauge change means a flip of the value of one or both qubits in the Hartree-Fock (HF) input state, and a sign change on the relevant terms of the Hamiltonian. In the standard gauge, the HF state is $|01\rangle$ and is prepared with an X_π gate on q_0 . Statistics from the experiment in these gauges are then averaged together. We also drop the first $-Y_{\pi/2}$ on q_0 ; for an input state of either $|0\rangle$ or $|1\rangle$, it has no effect given that $X_{\pi/2}$ is the only gate preceding it.

The energy for a given nuclear separation R is computed by calculating the value of the Hamiltonian with the expectation values measured for each θ and choosing the smallest energy. This is done for all values of R to construct the energy surface. Figure 6.2a shows the raw expectation values (after readout correction); Figure 6.2b shows the measured energy versus θ for each value of R and Figure E.2 shows the errors in that surface. Error bars were computed from a Gaussian process regression [19] applied to the energy

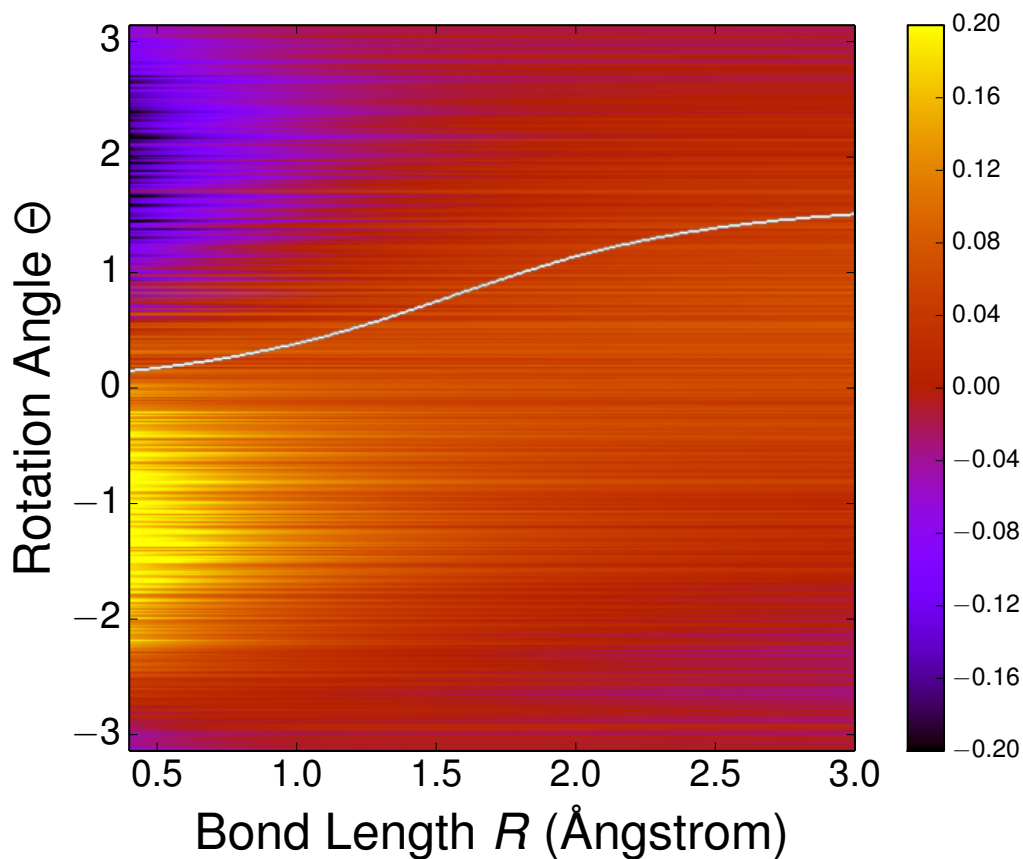


Figure E.2: Errors in the VQE energy surface (in Hartree) as a function of bond length and rotation angle. This plot looks somewhat like the derivative of Figure 6.2b with respect to R and θ because errors are greatest where the energy is most sensitive to changes in system parameters. As in Figure 6.2b, the white curve traces the theoretical minimum energy which is seen to be in good agreement with the data. Note that while errors in the energy surface are sometimes negative, all energies are bounded from below by the variational minimum.

landscape obtained from Figure 6.2b using error estimates propagated from the shot-noise limited measurements shown in Figure 6.2a.

E.3 Experimental methods for PEA

The PEA experiment uses three qubits: q_0 for the ancilla, and q_1 and q_2 for the register. Operating frequencies are 4.56, 5.65, and 4.80 GHz for q_0 , q_1 , and q_2 , respectively. Pulse tune-up is the same as for the VQE experiment. For the entangling gates (CZ_ϕ between q_0 and q_1 , and CZ_π between q_1 and q_2), however, the adjacent non-interacting qubit must be decoupled from the interaction. For the CZ_π , q_0 is decoupled with paired X_π and $-X_\pi$ pulses; this has the effect of “echoing out” any acquired state-dependent phase on q_0 from q_1 and vice versa, while minimizing stray single-qubit phases on q_0 by keeping its frequency stationary. For the CZ_ϕ , however, q_2 is detuned to frequencies significantly below the q_0 - q_1 interaction; while this makes single-qubit phases on q_2 harder to compensate, it is more effective at minimizing the impact of q_2 on the CZ_ϕ gate. This combination of decoupling methods was found to be optimal to minimize error on the phase of q_0 , which is the critical parameter in the PEA experiment.

As the CZ_ϕ gate varies the amplitude of q_1 's frequency trajectory over a wide range (approximately 200 MHz to 950 MHz) particular values of ϕ can be more sensitive lossy parts of the q_1 's frequency spectrum that are rapidly swept past and easily compensated for in the standard case of only tuning up $\phi = \pi$. Therefore, for some values of ϕ it is necessary to individually tune in compensating phases on q_0 . This is implemented by executing the individual term of the Hamiltonian, varying the compensating phase on q_0 , and fitting for the value that minimizes the error of that term. After performing this careful compensation when necessary, the experiment produces the bit values (0 or 1) for each different Hamiltonian (i.e. each separation R) at each evolution time t that match those predicted by numerical simulation.

As operated in this experiment, q_0 , q_1 , and q_2 have T_1 values of 48.1, 23.7, and 43.0 μs , and T_2^* times of 1.3, 1.6, and 0.8 μs , respectively. Figure 6.4 shows the pulses for

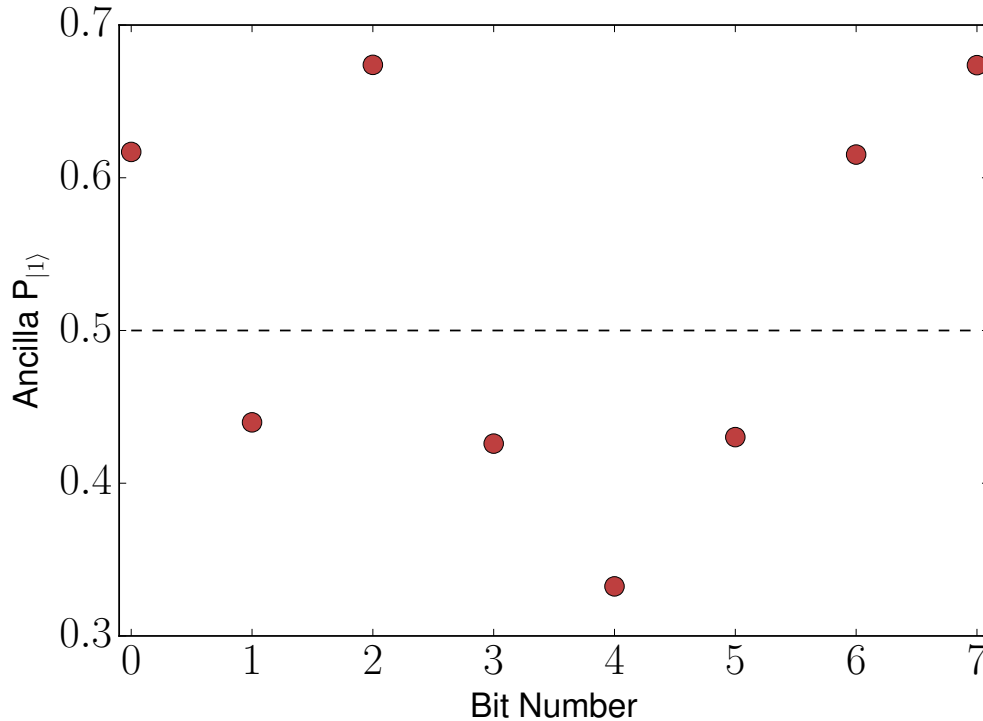



Figure E.3: Example data for a single PEA experiment, run at $R = 1.55 \text{ \AA}$. The results are shown without phase kickback from the measurements of the previous bit. The line at $P_{|1\rangle} = 0.5$ discriminates a measurement of 1 from 0.

one iteration of the PEA experiment; Figure E.3 shows an example of the measurement results for one value of R . The parameters at each R are given in Table E.1. For reference, included in this section are the implementations of all the terms of the Bravyi-Kitaev Hamiltonian for molecular hydrogen. In the following diagrams, α is the ancilla qubit (q_0 in the experiment), and 0 and 1 are the register qubits (q_1 and q_2 in the experiment). We must always be aware that representing our terms in terms of these gates, and then in terms of the actual basis, is not necessarily the most efficient approach.


CNOT

CNOT is implemented as a CZ_π and two rotations.




SWAP

SWAP is implemented as three consecutive CZ_π gates with intermediate rotations.




Controlled evolution under Z_0

Z_0 is implemented as CZ_ϕ and a z rotation on the control qubit.




Controlled evolution under Z_1

Z_1 is the same as Z_0 , but surrounded by SWAP gates so that the ancilla interacts with the other qubit.




Controlled evolution under X_0X_1

For X_0X_1 , we first change bases with $Y_{\pi/2}$ gates, then compute the parity of the register qubits with a CNOT, then apply the controlled phase, and finally undo the parity computation and basis change. Note that the $Y_{\pi/2}$ gates will cancel on the middle qubit with our CNOT implementation.



Controlled evolution under Y_0Y_1

Y_0Y_1 is the same as X_0X_1 with a different basis change.



E.4 Unitary coupled cluster

The application of VQE requires the choice of an ansatz, and in this work we have focused on the unitary coupled cluster (UCC) ansatz. This ansatz is a unitary variant of the method sometimes referred to as the “gold standard of quantum chemistry”, namely coupled cluster with single and double excitations with perturbative triples excitations

[48]. The unitary variant has the advantage of satisfying a variational principle with respect to all possible parameterizations. While the unitary variant has no efficient preparation scheme on a classical computer, scalable methods of preparation for a fixed set of parameters on a quantum device have now been documented several times [90, 142, 132, 79].

The UCC ansatz $|\varphi(\vec{\theta})\rangle$ is defined with respect to a reference state, which in this work we take to be the Hartree-Fock state $|\phi\rangle$,

$$|\varphi(\vec{\theta})\rangle = U(\vec{\theta})|\varphi\rangle = e^{T(\vec{\theta})-T(\vec{\theta})^\dagger}|\phi\rangle \quad (\text{E.7})$$

where $T(\vec{\theta})$ is the anti-Hermitian cluster operator:

$$T = \sum_k {}^{(k)}T(\vec{\theta}) \quad (\text{E.8})$$

$${}^{(1)}T(\vec{\theta}) = \sum_{\substack{i_1 \in \text{occ} \\ a_1 \in \text{virt}}} \theta_{i_1}^{a_1} a_{a_1}^\dagger a_{i_1} \quad (\text{E.9})$$

$${}^{(2)}T(\vec{\theta}) = \sum_{\substack{i_1, i_2 \in \text{occ} \\ a_1, a_2 \in \text{virt}}} \theta_{i_1, i_2}^{a_1, a_2} a_{a_2}^\dagger a_{i_2} a_{a_1}^\dagger a_{i_1} \quad (\text{E.10})$$

where the occ and virt spaces are defined as the occupied and unoccupied sites in the Hartree-Fock state and the definition of higher-order cluster operators ${}^{(k)}T$ follows naturally. When only including up to the first two terms in the cluster expansion, we term the ansatz unitary coupled cluster with single and doubles excitations (UCCSD) [48].

The task within VQE is to determine the optimal values of the cluster amplitudes $\theta_{i_1}^{a_1}$, which are determined by the variational minimum of a nonlinear function. As with all nonlinear minimizations, the choice of starting parameters is key to algorithmic performance. As in classical coupled cluster, we can determine the starting amplitudes perturbatively through Möller-Plesset perturbation theory (MP2) [48]. For molecular

hydrogen, there is exactly one term in the UCCSD ansatz.

The MP2 guess amplitudes are given by the equations

$$\theta_i^a = 0, \quad \theta_{ij}^{ab} = \frac{h_{ijba} - h_{ijab}}{\epsilon_i + \epsilon_j - \epsilon_a - \epsilon_b} \quad (\text{E.11})$$

where ϵ_a refer to the 1-electron occupied and virtual orbital energies from the Hartree-Fock calculation and the h_{ijab} are computed as in Eq. (E.3). In the MP2 guess, the vanishing of the singles amplitudes is a result of the fact that single excitations away from the Hartree-Fock reference do not couple through the Hamiltonian as a consequence of Brillouin’s theorem [48]. As the solution of the classical coupled cluster equations is also efficient, it is possible to use amplitudes from a method like CCSD as starting values as well. We note in both cases however, that the single-reference, perturbative nature of these constructions may lead to poor initial guesses for systems with strong multireference character or entanglement. In these cases the amplitudes may represent poor guesses, requiring more iterations for convergence. As such, a better initial guess in such problems may be a related optimization, such as a different molecular geometry of the same system. In cases where the perturbative estimates are accurate, one can discard operations related to very small amplitudes in the state preparation circuit, leading to computational savings.

R	1	Z_0	Z_1	Z_0Z_1	X_0X_1	Y_0Y_1	t_0	Ordering	Trotter Error
0.20	2.8489	0.5678	-1.4508	0.6799	0.0791	0.0791	1.500	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0124
0.25	2.1868	0.5449	-1.2870	0.6719	0.0798	0.0798	1.590	$Z_0 \cdot Y_0Y_1 \cdot Z_1 \cdot X_0X_1$	0.0521
0.30	1.7252	0.5215	-1.1458	0.6631	0.0806	0.0806	1.770	$X_0X_1 \cdot Z_0 \cdot Y_0Y_1 \cdot Z_1$	0.0111
0.35	1.3827	0.4982	-1.0226	0.6537	0.0815	0.0815	2.080	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0368
0.40	-0.2047	0.4754	-0.9145	0.6438	0.0825	0.0825	2.100	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0088
0.45	-0.2677	0.4534	-0.8194	0.6336	0.0835	0.0835	2.310	$X_0X_1 \cdot Z_0 \cdot Y_0Y_1 \cdot Z_1$	0.0141
0.50	-0.3202	0.4325	-0.7355	0.6233	0.0846	0.0846	2.580	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0672
0.55	-0.3642	0.4125	-0.6612	0.6129	0.0858	0.0858	2.700	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0147
0.60	-0.4012	0.3937	-0.5950	0.6025	0.0870	0.0870	2.250	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0167
0.65	-0.4323	0.3760	-0.5358	0.5921	0.0883	0.0883	3.340	$Z_1 \cdot X_0X_1 \cdot Z_0 \cdot Y_0Y_1$	0.0175
0.70	-0.4584	0.3593	-0.4826	0.5818	0.0896	0.0896	0.640	$Z_0 \cdot Y_0Y_1 \cdot Z_1 \cdot X_0X_1$	0.0171
0.75	-0.4804	0.3435	-0.4347	0.5716	0.0910	0.0910	0.740	$Z_0 \cdot Y_0Y_1 \cdot Z_1 \cdot X_0X_1$	0.0199
0.80	-0.4989	0.3288	-0.3915	0.5616	0.0925	0.0925	0.790	$Z_0 \cdot Y_0Y_1 \cdot Z_1 \cdot X_0X_1$	0.0291
0.85	-0.5143	0.3149	-0.3523	0.5518	0.0939	0.0939	3.510	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0254
0.90	-0.5271	0.3018	-0.3168	0.5421	0.0954	0.0954	3.330	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0283
0.95	-0.5377	0.2895	-0.2845	0.5327	0.0970	0.0970	4.090	$X_0X_1 \cdot Z_0 \cdot Y_0Y_1 \cdot Z_1$	0.0328
1.00	-0.5463	0.2779	-0.2550	0.5235	0.0986	0.0986	4.360	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0362
1.05	-0.5533	0.2669	-0.2282	0.5146	0.1002	0.1002	4.650	$Z_1 \cdot X_0X_1 \cdot Z_0 \cdot Y_0Y_1$	0.0405
1.10	-0.5588	0.2565	-0.2036	0.5059	0.1018	0.1018	4.280	$Z_1 \cdot X_0X_1 \cdot Z_0 \cdot Y_0Y_1$	0.0243
1.15	-0.5631	0.2467	-0.1810	0.4974	0.1034	0.1034	5.510	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0497
1.20	-0.5663	0.2374	-0.1603	0.4892	0.1050	0.1050	5.950	$Z_0 \cdot Y_0Y_1 \cdot Z_1 \cdot X_0X_1$	0.0559
1.25	-0.5685	0.2286	-0.1413	0.4812	0.1067	0.1067	6.360	$X_0X_1 \cdot Z_1 \cdot Y_0Y_1 \cdot Z_0$	0.0585
1.30	-0.5699	0.2203	-0.1238	0.4735	0.1083	0.1083	0.660	$Z_1 \cdot X_0X_1 \cdot Z_0 \cdot Y_0Y_1$	0.0905
1.35	-0.5706	0.2123	-0.1077	0.4660	0.1100	0.1100	9.810	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0694
1.40	-0.5707	0.2048	-0.0929	0.4588	0.1116	0.1116	9.930	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0755
1.45	-0.5702	0.1976	-0.0792	0.4518	0.1133	0.1133	5.680	$Y_0Y_1 \cdot Z_0 \cdot X_0X_1 \cdot Z_1$	0.0142
1.50	-0.5693	0.1908	-0.0666	0.4451	0.1149	0.1149	10.200	$Z_1 \cdot X_0X_1 \cdot Z_0 \cdot Y_0Y_1$	0.0885
1.55	-0.5679	0.1843	-0.0549	0.4386	0.1165	0.1165	9.830	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0917
1.60	-0.5663	0.1782	-0.0442	0.4323	0.1181	0.1181	8.150	$Z_0 \cdot Y_0Y_1 \cdot Z_1 \cdot X_0X_1$	0.0416
1.65	-0.5643	0.1723	-0.0342	0.4262	0.1196	0.1196	8.240	$X_0X_1 \cdot Z_0 \cdot Y_0Y_1 \cdot Z_1$	0.0488
1.70	-0.5621	0.1667	-0.0251	0.4204	0.1211	0.1211	0.520	$Z_1 \cdot X_0X_1 \cdot Z_0 \cdot Y_0Y_1$	0.0450
1.75	-0.5597	0.1615	-0.0166	0.4148	0.1226	0.1226	0.520	$Z_0 \cdot Y_0Y_1 \cdot Z_1 \cdot X_0X_1$	0.0509
1.80	-0.5571	0.1565	-0.0088	0.4094	0.1241	0.1241	1.010	$Z_0 \cdot X_0X_1 \cdot Z_1 \cdot Y_0Y_1$	0.0663
1.85	-0.5544	0.1517	-0.0015	0.4042	0.1256	0.1256	0.530	$Z_1 \cdot X_0X_1 \cdot Z_0 \cdot Y_0Y_1$	0.0163
1.90	-0.5516	0.1472	0.0052	0.3992	0.1270	0.1270	1.090	$X_0X_1 \cdot Z_0 \cdot Z_1 \cdot Y_0Y_1$	0.0017
1.95	-0.5488	0.1430	0.0114	0.3944	0.1284	0.1284	0.610	$X_0X_1 \cdot Z_1 \cdot Z_0 \cdot Y_0Y_1$	0.0873
2.00	-0.5458	0.1390	0.0171	0.3898	0.1297	0.1297	1.950	$Z_1 \cdot Z_0 \cdot X_0X_1 \cdot Y_0Y_1$	0.0784
2.05	-0.5429	0.1352	0.0223	0.3853	0.1310	0.1310	4.830	$X_0X_1 \cdot Y_0Y_1 \cdot Z_0 \cdot Z_1$	0.0947
2.10	-0.5399	0.1316	0.0272	0.3811	0.1323	0.1323	1.690	$Y_0Y_1 \cdot X_0X_1 \cdot Z_0 \cdot Z_1$	0.0206
2.15	-0.5369	0.1282	0.0317	0.3769	0.1335	0.1335	0.430	$X_0X_1 \cdot Y_0Y_1 \cdot Z_0 \cdot Z_1$	0.0014
2.20	-0.5339	0.1251	0.0359	0.3730	0.1347	0.1347	1.750	$Z_0 \cdot Z_1 \cdot X_0X_1 \cdot Y_0Y_1$	0.0107
2.25	-0.5310	0.1221	0.0397	0.3692	0.1359	0.1359	11.500	$X_0X_1 \cdot Z_1 \cdot Z_0 \cdot Y_0Y_1$	0.0946
2.30	-0.5280	0.1193	0.0432	0.3655	0.1370	0.1370	0.420	$Z_0 \cdot Z_1 \cdot X_0X_1 \cdot Y_0Y_1$	0.0370
2.35	-0.5251	0.1167	0.0465	0.3620	0.1381	0.1381	0.470	$Z_1 \cdot Z_0 \cdot Y_0Y_1 \cdot X_0X_1$	0.0762
2.40	-0.5223	0.1142	0.0495	0.3586	0.1392	0.1392	10.100	$X_0X_1 \cdot Z_1 \cdot Z_0 \cdot Y_0Y_1$	0.0334
2.45	-0.5195	0.1119	0.0523	0.3553	0.1402	0.1402	11.200	$Z_0 \cdot Z_1 \cdot X_0X_1 \cdot Y_0Y_1$	0.0663
2.50	-0.5168	0.1098	0.0549	0.3521	0.1412	0.1412	0.580	$Z_0 \cdot Y_0Y_1 \cdot X_0X_1 \cdot Z_1$	0.0296
2.55	-0.5141	0.1078	0.0572	0.3491	0.1422	0.1422	11.000	$Z_0 \cdot Z_1 \cdot X_0X_1 \cdot Y_0Y_1$	0.0550
2.60	-0.5114	0.1059	0.0594	0.3461	0.1432	0.1432	11.000	$Z_0 \cdot X_0X_1 \cdot Y_0Y_1 \cdot Z_1$	0.0507
2.65	-0.5089	0.1042	0.0614	0.3433	0.1441	0.1441	11.040	$Z_1 \cdot X_0X_1 \cdot Y_0Y_1 \cdot Z_0$	0.0490
2.70	-0.5064	0.1026	0.0632	0.3406	0.1450	0.1450	0.400	$Z_0 \cdot Z_1 \cdot Y_0Y_1 \cdot X_0X_1$	0.0471
2.75	-0.5039	0.1011	0.0649	0.3379	0.1458	0.1458	0.450	$Y_0Y_1 \cdot Z_0 \cdot Z_1 \cdot X_0X_1$	0.0061
2.80	-0.5015	0.0997	0.0665	0.3354	0.1467	0.1467	0.950	$Z_0 \cdot Y_0Y_1 \cdot X_0X_1 \cdot Z_1$	0.0368
2.85	-0.4992	0.0984	0.0679	0.3329	0.1475	0.1475	10.600	$Z_0 \cdot X_0X_1 \cdot Y_0Y_1 \cdot Z_1$	0.0324

Table E.1: The Hamiltonian coefficients for Eq. (6.1) and parameters (see text) for the PEA experiment for each value of R .

Bibliography

- [1] S. M. Anton, J. S. Birenbaum, S. R. O’Kelley, V. Bolkhovsky, D. A. Braje, G. Fitch, M. Neeley, G. C. Hilton, H. M. Cho, K. D. Irwin, F. C. Wellstood, W. D. Oliver, A. Shnirman, and John Clarke. Magnetic flux noise in dc SQUIDS: Temperature and geometry dependence. *Physical Review Letters*, 110(14):147002, apr 2013.
- [2] Alán Aspuru-Guzik, Anthony D Dutoi, Peter J Love, and Martin Head-Gordon. Simulated quantum computation of molecular energies. *Science (New York, N. Y.)*, 309(5741):1704–1707, 2005.
- [3] DD Awschalom, LC Bassett, and AS Dzurak. Quantum spintronics: engineering and manipulating atom-like spins in semiconductors. *Science*, 339(March):1174–1179, 2013.
- [4] Ryan Babbush, Dominic W Berry, Ian D Kivlichan, Annie Y Wei, Peter J Love, and Alán Aspuru-Guzik. Exponentially More Precise Quantum Simulation of Fermions I: Quantum Chemistry in Second Quantization. *e-print arXiv: 1506.01020*, 2015.
- [5] Ryan Babbush, Dominic W Berry, Ian D Kivlichan, Annie Y Wei, Peter J Love, and Alán Aspuru-Guzik. Exponentially More Precise Quantum Simulation of Fermions II: Quantum Chemistry in the CI Matrix Representation. *e-print arXiv: 1506.01029*, jun 2015.
- [6] Ryan Babbush, Peter J Love, and Alán Aspuru-Guzik. Adiabatic Quantum Simulation of Quantum Chemistry. *Scientific Reports*, 4(6603):1–11, 2014.
- [7] Ryan Babbush, Jarrod McClean, Dave Wecker, Alán Aspuru-Guzik, and Nathan Wiebe. Chemical Basis of Trotter-Suzuki Errors in Chemistry Simulation. *Physical Review A*, 91(2):22311, 2015.
- [8] M Bal, M H Ansari, J. L. Orgiazzi, R M Lutchyn, and a Lupascu. Dynamics of parametric fluctuations induced by quasiparticle tunneling in superconducting flux qubits. page 12, 2014.

- [9] CJ Ballance, TP Harty, NM Linke, and DM Lucas. High-fidelity two-qubit quantum logic gates using trapped calcium-43 ions. *arXiv preprint arXiv: ...*, pages 4–5, 2014.
- [10] R. Barends, J. Kelly, A. Megrant, D. Sank, E. Jeffrey, Y. Chen, Y. Yin, B. Chiaro, J. Mutus, C. Neill, P. O’Malley, P. Roushan, J. Wenner, T. C. White, A. N. Cleland, and John M. Martinis. Coherent Josephson Qubit Suitable for Scalable Quantum Integrated Circuits. *Physical Review Letters*, 111(8):080502, aug 2013.
- [11] R Barends, J Kelly, A Megrant, A Veitia, D Sank, E Jeffrey, T C White, J Mutus, A G Fowler, B Campbell, Y Chen, Z Chen, B Chiaro, A Dunsworth, C Neill, P O’Malley, P Roushan, A Vainsencher, J Wenner, A N Korotkov, A N Cleland, and John M Martinis. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–3, apr 2014.
- [12] R. Barends, L. Lamata, J. Kelly, L. García-Álvarez, A. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J Y Mutus, B Campbell, Yu Chen, Z Chen, B Chiaro, A Dunsworth, I. C. Hoi, C Neill, P. J. J. O’Malley, C. Quintana, P Roushan, A Vainsencher, J Wenner, E. Solano, and John M Martinis. Digital quantum simulation of fermionic models with a superconducting circuit. *arXiv*, page 5, 2015.
- [13] R Barends, A Shabani, L Lamata, J Kelly, A Mezzacapo, U Las Heras, R Babush, A Fowler, B Campbell, Y Chen, Z Chen, B Chiaro, A Dunsworth, E Jeffrey, E Lucero, A Megrant, J Mutus, M Neeley, C Neill, P O’Malley, C Quintana, P Roushan, D Sank, A Vainsencher, J Wenner, T White, E Solano, H Neven, and J Martinis. Digitized Adiabatic Quantum Computing with a Superconducting Circuit. *e-print arXiv: 1511.03316*, 2015.
- [14] Bela Bauer, Dave Wecker, Andrew J Millis, Matthew B Hastings, and M Troyer. Hybrid quantum-classical approach to correlated materials. *e-print arXiv:1510.03859*, oct 2015.
- [15] D A Bell. A survey of $1/f$ noise in electrical conductors. *Journal of Physics C: Solid State Physics*, 13(24):4425–4437, aug 1980.
- [16] Jan Benhelm, Gerhard Kirchmair, Christian F. Roos, and Rainer Blatt. Towards fault-tolerant quantum computing with trapped ions. *Nature Physics*, 4(6):463–466, jun 2008.
- [17] Radoslaw Bialczak, R. McDermott, M. Ansmann, M. Hofheinz, N. Katz, Erik Lucero, Matthew Neeley, a. O’Connell, H. Wang, a. Cleland, and John Martinis. $1/f$ Flux Noise in Josephson Phase Qubits. *Physical Review Letters*, 99(18):187006, nov 2007.

- [18] Michael Biercuk, Hermann Uys, Aaron VanDevender, Nobuyasu Shiga, Wayne Itano, and John Bollinger. Experimental Uhrig dynamical decoupling using trapped ions. *Physical Review A*, 79(6):062324, jun 2009.
- [19] Christopher M Bishop. *Pattern Recognition and Machine Learning*. Springer, New York, 2006.
- [20] Alexandre Blais, Ren-Shou Huang, Andreas Wallraff, S. Girvin, and R. Schoelkopf. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Physical Review A*, 69(6):062320, jun 2004.
- [21] Sergey B Bravyi and Alexei Yu Kitaev. Fermionic Quantum Computation. *Annals of Physics*, 298:210–226, 2002.
- [22] K. R. Brown, a. C. Wilson, Y. Colombe, C. Ospelkaus, a. M. Meier, E. Knill, D. Leibfried, and D. J. Wineland. Single-qubit-gate error below 10^{-4} in a trapped ion. *Physical Review A*, 84(3):030303, sep 2011.
- [23] Jonas Bylander, Simon Gustavsson, Fei Yan, Fumiki Yoshihara, Khalil Harrabi, George Fitch, David G. Cory, Yasunobu Nakamura, Jaw-Shen Tsai, and William D. Oliver. Noise spectroscopy through dynamical decoupling with a superconducting flux qubit. *Nature Physics*, 7(7):565–570, may 2011.
- [24] H. Y. Carr and E. M. Purcell. Effects of diffusion on free precession in nuclear magnetic resonance experiments. *Physical Review*, 94(3):630–638, may 1954.
- [25] Yu Chen, C Neill, P Roushan, N Leung, M Fang, R Barends, J Kelly, B Campbell, Z Chen, B Chiaro, A Dunsworth, E Jeffrey, A Megrant, J Y Mutus, P. J J O’Malley, C M Quintana, D Sank, A Vainsencher, J Wenner, T C White, Michael R Geller, A N Cleland, and John M Martinis. Qubit architecture with high coherence and fast tunable coupling. *Physical Review Letters*, 113(22), 2014.
- [26] Sangkook Choi, Dung Hai Lee, Steven G Louie, and John Clarke. Localization of metal-induced gap states at the metal-insulator interface: Origin of flux noise in SQUIDS and superconducting qubits. *Physical Review Letters*, 103(19), 2009.
- [27] T. Choi, S. Debnath, T.a. Manning, C. Figgatt, Z.-X. Gong, L.-M. Duan, and C. Monroe. Optimal Quantum Control of Multimode Couplings between Trapped Ion Qubits for Scalable Entanglement. *Physical Review Letters*, 112(19):190502, may 2014.
- [28] N Cody Jones, James D Whitfield, Peter L McMahon, Man-Hong Yung, Rodney Van Meter, Alán Aspuru-Guzik, and Yoshihisa Yamamoto. Faster quantum chemistry simulation on fault-tolerant quantum computers. *New Journal of Physics*, 14(11):115023, 2012.

- [29] A. D. Córcoles, Jay Gambetta, Jerry Chow, John Smolin, Matthew Ware, Joel Strand, B. Plourde, and M. Steffen. Process verification of two-qubit quantum gates by randomized benchmarking. *Physical Review A*, 87(3):030301, mar 2013.
- [30] A. D. Córcoles, Easwar Magesan, Srikanth J. Srinivasan, Andrew W. Cross, M. Steffen, Jay M. Gambetta, and Jerry M. Chow. Demonstration of a quantum error detection code using a square lattice of four superconducting qubits. *Nature Communications*, 6:6979, 2015.
- [31] Audrey Cottet. *Implementation of a quantum bit in a superconducting circuit*. Phd thesis, Université Paris VI, 2002.
- [32] Vasil S Denchev, Sergio Boixo, Sergei V Isakov, Nan Ding, Ryan Babbush, Vadim Smelyanskiy, John Martinis, and Hartmut Neven. What is the Computational Value of Finite Range Tunneling? *arXiv*, (3):1–17, 2015.
- [33] L DiCarlo, J M Chow, J M Gambetta, Lev S Bishop, B R Johnson, D I Schuster, J Majer, a Blais, L Frunzio, S M Girvin, and R J Schoelkopf. Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature*, 460(7252):240–4, jul 2009.
- [34] David P. DiVincenzo. The Physical Implementation of Quantum Computation. *Fortschritte der Physik*, 48(9-11):771–783, 2000.
- [35] Miroslav Dobšíček, Göran Johansson, Vitaly Shumeiko, and Göran Wendin. Arbitrary accuracy iterative quantum phase estimation algorithm using a single ancillary qubit: A two-qubit benchmark. *Physical Review A - Atomic, Molecular, and Optical Physics*, 76(3):030306, sep 2007.
- [36] Jiangfeng Du, Nanyang Xu, Xinhua Peng, Pengfei Wang, Sanfeng Wu, and Dawei Lu. NMR implementation of a molecular hydrogen quantum simulation with adiabatic state preparation. *Physical Review Letters*, 104(3):30502, 2010.
- [37] Lara Faoro and Lev B Ioffe. Microscopic origin of low-frequency flux noise in josephson circuits. *Physical Review Letters*, 100(22), 2008.
- [38] Richard P. Feynman. Simulating Physics with Computers. *International Journal of Theoretical Physics*, 21(6), 1982.
- [39] Austin G. Fowler, Matteo Mariantoni, John M. Martinis, and Andrew N. Cleland. Surface codes: Towards practical large-scale quantum computation. *Physical Review A*, 86(3):032324, sep 2012.
- [40] Andrei Galiutdinov, Alexander N. Korotkov, and John M. Martinis. Resonatorzero-qubit architecture for superconducting qubits. *Physical Review A*, 85(4):042321, apr 2012.

- [41] Y. Galperin, B. Altshuler, J. Bergli, and D. Shantsev. Non-Gaussian Low-Frequency Noise as a Source of Qubit Decoherence. *Physical Review Letters*, 96(9):097009, mar 2006.
- [42] Jay M. Gambetta, a. D. Córcoles, S. T. Merkel, B. R. Johnson, John a. Smolin, Jerry M. Chow, Colm a. Ryan, Chad Rigetti, S. Poletto, Thomas a. Ohki, Mark B. Ketchen, and M. Steffen. Characterization of Addressability by Simultaneous Randomized Benchmarking. *Physical Review Letters*, 109(24):240504, dec 2012.
- [43] L Garc\'{i}a-Álvarez, U Las Heras, A Mezzacapo, M Sanz, E Solano, and L Lamata. Quantum chemistry and charge transport in biomolecules with superconducting circuits. *e-print arXiv: 1511.09355*, 2015.
- [44] Crispin W Gardiner and Peter Zoller. *Quantum Noise: A Handbook of Markovian and Non-Markovian Quantum Stochastic Methods with Applications to Quantum Optics*. Springer, New York, third edition, 2004.
- [45] E. L. Hahn. Spin echoes. *Physical Review*, 80(4):580–594, nov 1950.
- [46] TP Harty, DTC Allcock, and CJ Ballance. High-fidelity preparation, gates, memory and readout of a trapped-ion quantum bit. *arXiv preprint arXiv: ...*, pages 1–5, 2014.
- [47] Matthew B Hastings, Dave Wecker, Bela Bauer, and Matthias Troyer. Improving Quantum Algorithms for Quantum Chemistry. *Quantum Information & Computation*, 15(1-2):1–21, 2015.
- [48] T Helgaker, P Jorgensen, and J Olsen. *Molecular Electronic Structure Theory*. Wiley, 2002.
- [49] G. Ithier, E. Collin, P. Joyez, P. Meeson, D. Vion, D. Esteve, F. Chiarello, A. Shnirman, Y. Makhlin, J. Schrieffer, and G. Schön. Decoherence in a superconducting quantum bit circuit. *Physical Review B*, 72(13):134519, oct 2005.
- [50] K M Itoh, D N Jamieson, J C McCallum, A S Dzurak, and A Morello. Quantifying the quantum gate fidelity of single-atom spin qubits in silicon by randomized benchmarking. pages 1–8, 2014.
- [51] Evan Jeffrey, Daniel Sank, J.Y. Mutus, T.C. White, J. Kelly, R. Barends, Y. Chen, Z. Chen, B. Chiaro, a. Dunsworth, a. Megrant, P.J.J. O’Malley, C. Neill, P. Roushan, a. Vainsencher, J. Wenner, a.N. Cleland, and John M. Martinis. Fast Accurate State Measurement with Superconducting Qubits. *Physical Review Letters*, 112(19):190504, may 2014.
- [52] M W Johnson, M H S Amin, S Gildert, T Lanting, F Hamze, N Dickson, R Harris, A J Berkley, J Johansson, P Bunyk, E M Chapple, C Enderud, J P Hilton,

- K Karimi, E Ladizinsky, N Ladizinsky, T Oh, I Perminov, C Rich, M C Thom, E Tolkacheva, C J S Truncik, S Uchaikin, J Wang, B Wilson, and G Rose. Quantum annealing with manufactured spins. *Nature*, 473(7346):194–198, 2011.
- [53] P Jordan and E Wigner. Über das Paulische Äquivalenzverbot. *Zeitschrift für Physik*, 47(9):631–651, 1928.
- [54] K Kakuyanagi, T Meno, S Saito, H Nakano, K Semba, H Takayanagi, F Deppe, and A Shnirman. Dephasing of a superconducting flux qubit. *Physical Review Letters*, 98(4), 2007.
- [55] Ivan Kassal, James Whitfield, Alejandro Perdomo-Ortiz, Man-Hong Yung, and Alán Aspuru-Guzik. Simulating Chemistry Using Quantum Computers. *Ann. Rev. Phys. Chem.*, 62(1):185–207, 2010.
- [56] Ivan Kassal, James D. Whitfield, Alejandro Perdomo-Ortiz, Man-Hong Yung, and Alán Aspuru-Guzik. Simulating chemistry using quantum computers. *Annual review of physical chemistry*, 62(1):185–207, may 2011.
- [57] J. Kelly, R. Barends, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, a. Dunsworth, a.G. Fowler, I.-C. Hoi, E. Jeffrey, a. Megrant, J. Mutus, C. Neill, P.J.J. O’Malley, C. Quintana, P. Roushan, D. Sank, a. Vainsencher, J. Wenner, T.C. White, a.N. Cleland, and John M. Martinis. Optimal Quantum Control Using Randomized Benchmarking. *Physical Review Letters*, 112(24):240504, jun 2014.
- [58] J. Kelly, R. Barends, a. G. Fowler, A. Megrant, E. Jeffrey, T. C. White, D. Sank, J. Y. Mutus, B. Campbell, Yu Chen, Z. Chen, B. Chiaro, A. Dunsworth, I.-C. Hoi, C. Neill, P. J. J. O’Malley, C. Quintana, P. Roushan, A. Vainsencher, J. Wenner, a. N. Cleland, and John M. Martinis. State preservation by repetitive error detection in a superconducting quantum circuit. *Nature*, 519(7541):66–69, 2015.
- [59] Julia Kempe, Alexei Kitaev, and Oded Regev. The Complexity of the k-Local Hamiltonian Problem. *SIAM Journal on Computing*, 35(5):1070–1097, 2006.
- [60] A. Yu. Kitaev. Quantum measurements and the Abelian stabilizer problem. *arXiv preprint quant-ph/9511026*, pages 1–22, 1995.
- [61] A.Yu. Kitaev. Fault-tolerant quantum computation by anyons. *Annals of Physics*, 303(1):2–30, 2003.
- [62] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. Blakestad, J. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. Wineland. Randomized benchmarking of quantum gates. *Physical Review A*, 77(1):012307, jan 2008.
- [63] Jens Koch, Terri Yu, Jay Gambetta, A. Houck, D. Schuster, J. Majer, Alexandre Blais, M. Devoret, S. Girvin, and R. Schoelkopf. Charge-insensitive qubit design derived from the Cooper pair box. *Physical Review A*, 76(4):042319, oct 2007.

- [64] Roger H Koch, John Clarke, W M Goubau, J M Martinis, C M Pegrum, and D. J. van Harlingen. Flicker (1/f) noise in tunnel junction dc SQUIDS. *Journal of Low Temperature Physics*, 51(1-2):207–224, 1983.
- [65] Roger H Koch, David P Divincenzo, and John Clarke. Model for 1/f flux noise in SQUIDS and qubits. *Physical Review Letters*, 98(26), 2007.
- [66] P Kumar, S Sendelbach, M A Beck, J W Freeland, Zhe Wang, Hui Wang, C C Yu, R Q Wu, D P Pappas, and R. McDermott. Origin and Suppression of 1/f Magnetic Flux Noise. 2016.
- [67] Susanne Kunkel, Maximilian Schmidt, Jochen Martin Eppler, Hans Ekkehard Plesser, Gen Masumoto, Jun Igarashi, Shin Ishii, Tomoki Fukai, Abigail Morrison, Markus Diesmann, and Moritz Helias. Spiking network simulation code for petascale computers. *Frontiers in Neuroinformatics*, 8(78), 2014.
- [68] S. LaForest and Rogério de Sousa. Flux-vector model of spin noise in superconducting circuits: Electron versus nuclear spins and role of phase transition. *Physical Review B*, 92(5):054502, 2015.
- [69] B P Lanyon, J D Whitfield, G G Gillett, M E Goggin, M P Almeida, I Kassal, J D Biamonte, M Mohseni, B J Powell, M Barbieri, A Aspuru-Guzik, and a G White. Towards quantum chemistry on a quantum computer. *Nature chemistry*, 2(2):106–111, feb 2010.
- [70] Seth Lloyd. Universal Quantum Simulators. *Science*, 273(5):1073–1078, aug 1996.
- [71] Erik Lucero, Julian Kelly, Radoslaw C. Bialczak, Mike Lenander, Matteo Mariantoni, Matthew Neeley, A. D. O’Connell, Daniel Sank, H. Wang, Martin Weides, James Wenner, Tsuyoshi Yamamoto, A. N. Cleland, and John M. Martinis. Reduced phase error through optimized control of a superconducting qubit. *Physical Review A*, 82(4):042339, oct 2010.
- [72] R. M. Lutchyn, L. I. Glazman, and A. I. Larkin. Kinetics of the superconducting charge qubit in the presence of a quasiparticle. *Physical Review B - Condensed Matter and Materials Physics*, 74(6):064515, aug 2006.
- [73] Roman Lutchyn, Leonid Glazman, and Anatoly Larkin. Quasiparticle decay rate of Josephson charge qubit oscillations. *Physical Review B - Condensed Matter and Materials Physics*, 72(1):014517, jul 2005.
- [74] Easwar Magesan, J. M. Gambetta, and Joseph Emerson. Scalable and Robust Randomized Benchmarking of Quantum Processes. *Physical Review Letters*, 106(18):180504, may 2011.

- [75] Easwar Magesan, Jay M. Gambetta, B. R. Johnson, Colm a. Ryan, Jerry M. Chow, Seth T. Merkel, Marcus P. da Silva, George a. Keefe, Mary B. Rothwell, Thomas a. Ohki, Mark B. Ketchen, and M. Steffen. Efficient Measurement of Quantum Gate Error by Interleaved Randomized Benchmarking. *Physical Review Letters*, 109(8):080505, aug 2012.
- [76] John Martinis, S. Nam, J. Aumentado, K. Lang, and C. Urbina. Decoherence of a superconducting qubit due to bias noise. *Physical Review B*, 67(9):094510, mar 2003.
- [77] John M. Martinis. Qubit metrology for building a fault-tolerant quantum computer. (2):1–4, 2015.
- [78] Jarrod R McClean, Ryan Babbush, Peter J Love, and Alán Aspuru-Guzik. Exploiting locality in quantum computation for quantum chemistry. *The Journal of Physical Chemistry Letters*, 5(24):4368–4380, 2014.
- [79] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics*, 18(2):1–20, feb 2016.
- [80] A. Megrant. *No Title*. PhD thesis, University of California, Santa Barbara, 2016.
- [81] S. Meiboom and D. Gill. Modified Spin-Echo Method for Measuring Nuclear Relaxation Times, 1958.
- [82] Nikolaj Moll, Andreas Fuhrer, Peter Staar, and Ivano Tavernelli. Optimizing Qubit Resources for Quantum Chemistry Simulations in Second Quantization on a Quantum Computer. *e-print arXiv:1510.04048*, 2015.
- [83] Leonie Mueck. Quantum reform. *Nature Chemistry*, 7(5):361–363, 2015.
- [84] Hartmut Neven, Vasil S Denchev, Geordie Rose, and William G Macready. Training a Large Scale Classifier with the Quantum Adiabatic Algorithm. *e-print arXiv:09120779*, page 14, 2009.
- [85] Michael A Nielsen and Isaac L Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, New York, NY, USA, 10th edition, 2011.
- [86] C. D. Nugroho, V. Orlyanchik, and D. J. Van Harlingen. Low frequency resistance and critical current fluctuations in Al-based Josephson junctions. *Applied Physics Letters*, 102(14):142602, 2013.
- [87] Nissim Ofek, Andrei Petrenko, Reinier Heeres, Philip Reinhold, Zaki Leghtas, Brian Vlastakis, Yehan Liu, Luigi Frunzio, S M Girvin, Liang Jiang, Mazyar Mirrahimi,

- M H Devoret, and R J Schoelkopf. Demonstrating Quantum Error Correction that Extends the Lifetime of Quantum Information. 2016.
- [88] Hanhee Paik, D. I. Schuster, Lev S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. Observation of High Coherence in Josephson Junction Qubits Measured in a Three-Dimensional Circuit QED Architecture. *Physical Review Letters*, 107(24):240501, dec 2011.
- [89] E. Paladino, L. Faoro, G. Falci, and Rosario Fazio. Decoherence and $1/f$ Noise in Josephson Qubits. *Physical Review Letters*, 88(22):228304, may 2002.
- [90] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J Love, Alán Aspuru-Guzik, and Jeremy L O’Brien. A variational eigenvalue solver on a quantum processor. *arXiv.org*, quant-ph(May):1–10, 2013.
- [91] David Poulin, M B Hastings, Dave Wecker, Nathan Wiebe, Andrew C Doherty, and Matthias Troyer. The Trotter Step Size Required for Accurate Quantum Simulation of Quantum Chemistry. *Quantum Information & Computation*, 15(5-6):361–384, 2015.
- [92] WH Press. Flicker noises in astronomy and elsewhere. *Comments on Astrophysics*, 7(4):103–119, 1978.
- [93] Norman F. Ramsey. A molecular beam resonance method with separated oscillating fields. *Physical Review*, 78(6):695–699, jun 1950.
- [94] Matthew Reagor, Wolfgang Pfaff, Christopher Axline, Reinier W Heeres, Nissim Ofek, Katrina Sliwa, Eric Holland, Chen Wang, Jacob Blumoff, Kevin Chou, J Michael, Luigi Frunzio, Michel H Devoret, Liang Jiang, and Robert J Schoelkopf. A quantum memory with near-millisecond coherence in circuit QED. *arXiv*, page 1508.05882v2, 2015.
- [95] Chad Rigetti, Jay M. Gambetta, Stefano Poletto, B. L. T. Plourde, Jerry M. Chow, a. D. Córcoles, John a. Smolin, Seth T. Merkel, J. R. Rozen, George a. Keefe, Mary B. Rothwell, Mark B. Ketchen, and M. Steffen. Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms. *Physical Review B*, 86(10):100506, sep 2012.
- [96] D Ristè, C C Bultink, M J Tiggelman, R N Schouten, K W Lehnert, and L DiCarlo. Millisecond charge-parity fluctuations and induced decoherence in a superconducting transmon qubit. *Nature communications*, 4:1913, jan 2013.
- [97] D Ristè, S Poletto, M-Z Huang, A Bruno, V Vesterinen, O-P Saira, and L DiCarlo. Detecting bit-flip errors in a logical qubit using stabilizer measurements. *Nature communications*, 6:6983, 2015.

- [98] P Roushan, C Neill, and Y Chen. Observation of topological transitions in interacting quantum circuits. *Nature*, 515(7526):241–244, 2014.
- [99] C a Ryan, M Laforest, and R Laflamme. Randomized benchmarking of single- and multi-qubit control in liquid-state NMR quantum information processing. *New Journal of Physics*, 11(1):013034, jan 2009.
- [100] J. J. Sakurai. *Modern_Quantum_Mechanics _J_J_Sakurai.pdf*. Addison-Wesley, revised edition, 1994.
- [101] Daniel Sank. *Fast, Accurate State Measurement in Superconducting Qubits*. PhD thesis, University of California, Santa Barbara, 2014.
- [102] Daniel Sank, R. Barends, Radoslaw C. Bialczak, Yu Chen, J. Kelly, M. Lenander, E. Lucero, Matteo Mariantoni, A. Megrant, M. Neeley, P. J. J. O’Malley, A. Vainsencher, H. Wang, J. Wenner, T. C. White, T. Yamamoto, Yi Yin, A. N. Cleland, and John M. Martinis. Flux Noise Probed with Real Time Qubit Tomography in a Josephson Phase Qubit. *Physical Review Letters*, 109(6):067001, aug 2012.
- [103] Philipp Schindler, Daniel Nigg, Thomas Monz, Julio T. Barreiro, Esteban Martinez, Shannon X. Wang, Stephan Quint, Matthias F. Brandl, Volckmar Nebendahl, Christian F. Roos, Michael Chwalla, Markus Hennrich, and Rainer Blatt. A quantum information processor with trapped ions. *New Journal of Physics*, 15(12):123012, dec 2013.
- [104] a. P. Sears, A. Petrenko, G. Catelani, L. Sun, Hanhee Paik, G. Kirchmair, L. Frunzio, L. I. Glazman, S. M. Girvin, and R. J. Schoelkopf. Photon shot noise dephasing in the strong-dispersive limit of circuit QED. *Physical Review B*, 86(18):180504, nov 2012.
- [105] Adam Patrick Sears. *Extending Coherence in Superconducting Qubits : from microseconds to milliseconds*. PhD thesis, Yale University, 2013.
- [106] Jacob T. Seeley, Martin J. Richard, and Peter J. Love. The Bravyi-Kitaev transformation for quantum computation of electronic structure. *Journal of Chemical Physics*, 137(22), 2012.
- [107] S. Sendelbach, D. Hover, a. Kittel, M. Mück, John Martinis, and R. McDermott. Magnetism in SQUIDs at Millikelvin Temperatures. *Physical Review Letters*, 100(22):227006, jun 2008.
- [108] S Sendelbach, D Hover, M Mück, and R. McDermott. Complex inductance, excess noise, and surface magnetism in dc SQUIDs. *Physical Review Letters*, 103(11), 2009.

- [109] Yangchao Shen, Xiang Zhang, Shuaining Zhang, Jing-Ning Zhang, Man-Hong Yung, and Kihwan Kim. Quantum implementation of unitary coupled cluster for simulating molecular electronic structure. *arXiv: 1506.00443*, pages 1–24, 2015.
- [110] Alexander Shnirman, Gerd Schön, Ivar Martin, and Yuriy Makhlin. Low- and High-Frequency Noise from Coherent Two-Level Systems. *Physical Review Letters*, 94(12):127002, apr 2005.
- [111] Peter W Shor. Algorithms for Quantum Computation : Discrete Log and Factoring Extended Abstract 1 Introduction. *Proceedings of the 35th Annual Symposium on the Foundations of Computer Science*, page 124, 1994.
- [112] Peter W. Shor. Scheme for reducing decoherence in quantum computer memory. *Physical Review A*, 52(4):R2493–R2496, oct 1995.
- [113] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529:484, 2016.
- [114] John A Smolin, Graeme Smith, and Alexander Vargo. Oversimplifying quantum factoring. *Nature*, 499(7457):163–165, 2013.
- [115] R. Somma, G. Ortiz, J. Gubernatis, E. Knill, and R. Laflamme. Simulating physical phenomena by quantum networks. *Physical Review A*, 65(4):042323, 2002.
- [116] M Stern, G Catelani, Y Kubo, C Grezes, A Bienfait, D Vion, D Esteve, and P Bertet. Flux Qubits with Long Coherence Times for Hybrid Quantum Circuits. pages 1–8, 2014.
- [117] Masuo Suzuki. Generalized Trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Communications in Mathematical Physics*, 51(2):183–190, 1976.
- [118] Krysta M Svore, Matthew B Hastings, and Michael Freedman. Faster phase estimation. *Quantum Information & Computation*, 14(3):306–328, 2014.
- [119] Borzu Toloui and Peter J Love. Quantum Algorithms for Quantum Chemistry based on the sparsity of the CI-matrix. *e-print arXiv: 1312.2579*, 2013.
- [120] Andrew Tranter, Sarah Sofia, Jake Seeley, Michael Kaicher, Jarrod McClean, Ryan Babbush, Peter V. Coveney, Florian Mintert, Frank Wilhelm, and Peter J. Love. The Bravyi-Kitaev transformation: Properties and applications. *International Journal of Quantum Chemistry*, 115(19):1431–1441, 2015.

- [121] Hale F Trotter. On the product of semi-groups of operators. *Proc. Am. Math. Soc.*, 10(4):545–551, 1959.
- [122] Colin J Trout and Kenneth R Brown. Magic state distillation and gate compilation in quantum algorithms for quantum chemistry. *International Journal of Quantum Chemistry*, 115(19):1296–1304, 2015.
- [123] Matthias Troyer. High Performance Quantum Computing, 2014.
- [124] Götz S Uhrig. Exact results on dynamical decoupling by π pulses in quantum information processes. *New Journal of Physics*, 10(8):083024, aug 2008.
- [125] T Van Duzer and C W Turner. *Principles of superconductive devices and circuits*. Elsevier, 1981.
- [126] D. J. Van Harlingen, T. L. Robertson, B. L T Plourde, P. A. Reichardt, T. A. Crane, and John Clarke. Decoherence in Josephson-junction qubits due to critical-current fluctuations. *Physical Review B - Condensed Matter and Materials Physics*, 70(6), 2004.
- [127] Libor Veis and Jiri Pittner. Adiabatic state preparation study of methylene. *The Journal of Chemical Physics*, 140(214111):1–21, 2014.
- [128] Libor Veis, Jakub Višňák, Hiromi Nakai, and Jiri Pittner. Quantum chemistry beyond Born-Oppenheimer approximation on a quantum computer: a simulated phase estimation study. *e-print arXiv: 1507.03271*, 2015.
- [129] M. Vogelsberger, S. Genel, V. Springel, P. Torrey, D. Sijacki, D. Xu, G. Snyder, S. Bird, D. Nelson, and L. Hernquist. Properties of galaxies reproduced by a hydrodynamic simulation. *Nature*, 509(7499):177–82, may 2014.
- [130] Hui Wang, Chuntai Shi, Jun Hu, Sungho Han, Clare C. Yu, and R. Q. Wu. Candidate Source of Flux Noise in SQUIDs: Adsorbed Oxygen Molecules. *Physical Review Letters*, 115(7):1–5, 2015.
- [131] Ya Wang, Florian Dolde, Jacob Biamonte, Ryan Babbush, Ville Bergholm, Sen Yang, Ingmar Jakobi, Philipp Neumann, Alán Aspuru-Guzik, James D Whitfield, and Jörg Wrachtrup. Quantum Simulation of Helium Hydride Cation in a Solid-State Spin Register. *ACS Nano*, 9(8):7769–7774, 2015.
- [132] Dave Wecker, Matthew B Hastings, and Matthias Troyer. Progress towards practical quantum variational algorithms. *Physical Review A*, 92(4):42303, 2015.
- [133] David Wecker, Bela Bauer, Bryan K Clark, Matthew B Hastings, and Matthias Troyer. Gate-count estimates for performing quantum chemistry on small quantum computers. *Physical Review A*, 90(022305):1–13, 2014.

- [134] Frederick C. Wellstood, Cristian Urbina, and John Clarke. Low-frequency noise in dc superconducting quantum interference devices below 1 K. *Applied Physics Letters*, 50(12):772, mar 1987.
- [135] James D Whitfield. Spin-free quantum computational simulations and symmetry adapted states. *Journal of Chemical Physics*, 139(2):21105, 2013.
- [136] James D. Whitfield, Jacob Biamonte, and Alán Aspuru-Guzik. Simulation of Electronic Structure Hamiltonians Using Quantum Computers. (March 2015):22, 2010.
- [137] Nathan Wiebe and Christopher E. Granade. Efficient Bayesian phase estimation. *arXiv: 1508.00869*, 2(2):1–12, 2015.
- [138] Fei Yan, Jonas Bylander, Simon Gustavsson, Fumiki Yoshihara, Khalil Harrabi, David G. Cory, Terry P. Orlando, Yasunobu Nakamura, Jaw-Shen Tsai, and William D. Oliver. Spectroscopy of low-frequency noise and its temperature dependence in a superconducting qubit. *Physical Review B*, 85(17):174521, may 2012.
- [139] Fei Yan, Simon Gustavsson, Jonas Bylander, Xiaoyue Jin, Fumiki Yoshihara, David G Cory, Yasunobu Nakamura, Terry P Orlando, and William D Oliver. Rotating-frame relaxation as a noise spectrum analyser of a superconducting qubit undergoing driven evolution. *Nature communications*, 4:2337, jan 2013.
- [140] F. Yoshihara, K. Harrabi, a. Niskanen, Y. Nakamura, and J. Tsai. Decoherence of Flux Qubits due to $1/f$ Flux Noise. *Physical Review Letters*, 97(16):167001, oct 2006.
- [141] Fumiki Yoshihara, Yasunobu Nakamura, Fei Yan, Simon Gustavsson, Jonas Bylander, William D. Oliver, and Jaw-Shen Tsai. Flux qubit noise spectroscopy using Rabi oscillations under strong driving conditions. *Physical Review B*, 89(2):020503, jan 2014.
- [142] Man-Hong Yung, J Casanova, A Mezzacapo, J McClean, L Lamata, A Aspuru-Guzik, and E Solano. From transistor to trapped-ion computers for quantum chemistry. *Scientific Reports*, 4(3589):9, 2014.