

Syracuse University

SURFACE at Syracuse University

Dissertations - ALL

SURFACE at Syracuse University

11-29-2022

High Fidelity Universal Gates Performed on a Continuously-Decoupled Coherence Enhanced Transmon Qubit

Michael Senatore

Follow this and additional works at: <https://surface.syr.edu/etd>



Part of the [Quantum Physics Commons](#)

Recommended Citation

Senatore, Michael, "High Fidelity Universal Gates Performed on a Continuously-Decoupled Coherence Enhanced Transmon Qubit" (2022). *Dissertations - ALL*. 1602.
<https://surface.syr.edu/etd/1602>

This Dissertation is brought to you for free and open access by the SURFACE at Syracuse University at SURFACE at Syracuse University. It has been accepted for inclusion in Dissertations - ALL by an authorized administrator of SURFACE at Syracuse University. For more information, please contact surface@syr.edu.

ABSTRACT

Decoherence is the primary limiting factor for the utility of modern qubits and qubit networks; most chiefly, pure dephasing which limits the operational time any gate-sequence can produce a high-fidelity result. In this dissertation, I present the results of my experiment, performing fast, high fidelity, universal single-qubit gates, on a qubit which has been decoupled from pure dephasing resulting from environmental noise. This technique can expand operational ranges of qubits—such as allowing the high-coherence operation of a flux-tunable qubit far away from its flux-insensitive sweet-spot; broadening our selection of viable qubits by making otherwise low-coherence qubits operable with high coherence, or improving the coherence of higher order quantum networks which have limited coherence time due to qubit to qubit interactions producing prohibitive amounts of pure dephasing. This technique could be performed on any deterministic qubit of any modality which can receive drives of a physically similar kind as my particular test-platform, the superconducting transmon.

**HIGH FIDELITY UNIVERSAL GATES
PERFORMED ON A
CONTINUOUSLY-DECOUPLED COHERENCE
ENHANCED TRANSMON QUBIT**

By

Michael Senatore

BA, Colgate University, 2014

MS, Syracuse University, 2017

DISSERTATION

SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY IN PHYSICS

Syracuse University

December 2022

Copyright ©Michael Senatore, 2022
All rights reserved

I would like to thank several scientific mentors that helped lead me through the path I've been traversing for the last 7+ years.

I would like to thank Professor Britton Plourde for my start in the field of experimental superconducting qubit research. Were it not for Britton, I would not be studying what I am today, nor would I have had the skills or knowledge required to perform this research. Britton, for your mentorship and for the skills you led me to learn over my 4 years in your lab, and for your understanding when I needed to move on, I thank you.

I would like to thank Matthew Lahaye, the PI of the superconducting lab at the AFRL with which I performed the research detailed in this dissertation, and a good friend. Matt, you have been a skillful scientific advisor and a leader who treated my needs as a person with as much importance as the research we were performing. The amount of attention and talent that it takes to create such a supportive environment cannot be overstated, and I thank you for every ounce of time and attention you took to crafting such an encouraging environment.

I would like to thank Daniel Campbell, a research scientist in the superconducting lab at the AFRL, and an excellent colleague. Between his initial demonstration of the technique I use in my experiment, and his tutelage over the last 2 years bringing me fully up to speed on performing measurement and data analysis, I owe much of my current capability at the measurement-computer to him. Dan, thank you for the long conversations, and the painstaking effort you took to bring information across the occasionally nigh-infinite energy barrier that is my skull.

To everyone in the Physics department at Syracuse University, and everyone who works at the AFRL helping to make the lab a fun, interesting, and supportive place to be and work, I thank you.

Throughout my career as a graduate student, I have tried my best to hold, in as much importance as the research, my life outside of it. My emotional connections with people in my life have been at least as important in keeping me afloat through the sea of uncertainty as the scientific mentorship I've received.

I would like to thank my parents. Mom and Dad, were it not for you, I would not exist. Obviously. But more importantly I would never have become the kind of curious, discerning, and determined person that I am today, and I wouldn't have had the stubbornness required to make it to the end of this very long, often very difficult

path. For your decades of parenting, for your answers of "why" and "how", and for your belief and insistence that I should try and would succeed if I did, no matter at what, I thank you. I love you both very much.

I would like to thank my friends, who helped me continue to feel like a human connected to others in the world rather than just a brain in a vat who happened to have fingertips on the other end of a measurement computer. Rodney, Fey, Bethany, Holly, Andrew, Bella, Ari, Kat and so many other close friends who have each become like family to me, thank you for every moment we shared over the many years. From ballroom dancing to video gaming, cackling over memes or having hard conversations over heartache, thank you all for every second of life that you shared with me. It is thanks to you all that my good spirits are in no short-supply, and I am forever grateful.

Though often the path to a PhD is all-consuming, I found a way to push my life forward through the thick of it. Michelle Fecio, you and I met in the midst of the worst pandemic either of us have yet lived through. Our first date was a 3 hour hike at Green Lakes State Park, followed by pizza and antipasto we brought back to your basement apartment because we couldn't eat at restaurants. Since then, we bought a house together, and not too long before the writing of these very words, we got engaged. For your patience with me through the difficulties of the final steps along this path, and in the hopes of a very long and very happy life together; thank you for everything. I love you.

Contents

1	Introduction	1
2	Simulation, Design, and Fabrication	8
2.1	Readout Hardware	8
2.2	Transmon Qubits	11
2.3	Coupling	17
3	Measurement, Analysis, and Tuneup	20
3.1	Measurement Hardware	20
3.2	Spectroscopy for Qubit State Read-Out	23
3.3	Two-Tone Qubit Spectroscopy	26
3.4	Qubit State Manipulation	27
3.5	Qubit Characteristic Lifetime and Decoherence time	30
3.6	Second Excited State and Anharmonicity	41
4	Continuous Decoupling and Fast, High-Fidelity Universal gates	44
4.1	Longitudinal Noise and Continuous Decoupling	45
4.2	Mitigating the failure of the Rabi Approximation	50
4.3	Fast Universal Gates	51
4.4	Designing gates despite frequency constraints	52
4.5	State Preparation and Tuneup	55
4.5.1	Frequency and Time Domain measurements	59
4.5.2	Designing Gates	61
4.5.3	Tuning up time-evolution Z-gates	65
4.6	Data and Results	66

4.6.1	Paths for improvement	73
4.6.2	Conclusion	74
5	Hamiltonian Design	75
6	Departing Remarks and Future Work	77

List of Figures

1	A qubit is just a quantized two-state system which can be operated as binary information	2
2	The dark gray is all superconducting metal, and the white is bare-substrate. The inset images are two examples of different Josephson-elements; the top inset image is a Josephson-element of a SQUID made from 2 Josephson junctions with its accompanying flux-bias line to tune frequency, and the bottom inset image is the variant with a single Josephson junction.	9
3	A diagram of the energy level manifold of a transmon qubit, and a comparison between the energy potential of a transmon and the standard harmonic oscillator potential, illustrating the key feature that allows constraint within a qubit basis; the anharmonicity (α).	13
4	This is an approximation of the flux-tuning of the frequency of a nominally symmetric flux-tunable transmon qubit.	15
5	The room-temperature hardware send signals down lines of the dilution refrigerator, which interact with our sample at the coldest stage before being returned to the room temperature data-acquisition hardware.	21
6	This image provides a reminder of the device geometry; the signal I described above travels through the transmission line shown here, thus interacting with the rest of the on-chip circuit.	22

7	Data plot of single-tone spectroscopy performed on a readout resonator through the transmission-line – the center of the Lorentzian feature is the fundamental frequency of the readout resonator (f_R)	23
8	This is a power scan of a resonator in spectroscopy; at low powers the resonator’s fundamental frequency does not vary as a function of power, but at higher powers we see an abrupt change in the resonance.	24
9	Two scans of the readout resonator feature, the red curve is taken when the qubit coupled to it is in the ground state, and the blue curve when the qubit is in the first excited state. The black line is the marker which represents the frequency (f_R) where we would then choose to position our readout-resonator-probe-tone for future two-tone-spectroscopy. The difference in the frequencies between the center of these two features is Δ_{QR}	25
10	Two-tone spectroscopy performed by reading out the amplitude of the tone at the resonator’s un-pulled resonance frequency (f_R), while sourcing a second tone at (f_{spec}) trying to excite the qubit.	26
11	This is a data plot of coherent oscillations of the qubit as a function of drive power. A pulse of constant duration is sourced at the qubit transition frequency (f_{01}) and its amplitude is scanned, producing this oscillation between the two eigenstates.	28
12	This is a data plot of coherent oscillations of the qubit as a function of drive duration. A pulse of constant amplitude is sourced at the qubit transition frequency (f_{01}) and its duration is scanned, producing this oscillation between the two eigenstates.	29
13	This is a colormap of coherent oscillations over time made at different frequency detuning (Δ_f). The oscillations are slowest when the tone frequency matches the qubit frequency. Here, the colorscale represents the measured probe amplitude, which is the established proxy for qubit-state.	31
14	A scan of the decay time T_1 , taken 36 times and averaged. Each datapoint on this plot represents 36,000 measurements of the qubit state at each time after preparation in the first excited state all averaged together.	32

15	The conventional representation of superposition states of the qubit; the Bloch sphere.	34
16	A measurement of decoherence on an intentionally low-coherence device performed with a Ramsey experiment. The noise responsible for the dominant proportion of the decoherence is pure-dephasing from flux-frequency noise on a flux-sensitive transmon.	37
17	A demonstration of a way to visualize decoherence on the Bloch sphere. The superposition starts well-defined, and progresses from blue to orange to red, smearing out as a function of time.	38
18	A measurement of decoherence on an intentionally low-coherence device performed with a Hahn-Echo experiment. The noise responsible for the dominant proportion of the decoherence is pure-dephasing from flux-frequency noise on a flux-sensitive transmon. Some of it is mitigated by this measurement technique.	40
19	A measurement of the transition between the first excited state and the second-excited state made by exciting the qubit into the first excited state before scanning a pulse through frequency-space searching for the second-excited state transition.	42
20	The Bloch sphere of the bare transmon overlaid with a drive along the x-axis. The following demonstration of continuous decoupling is prepared with the continuous drive and Bloch-vector (red arrow) colinear along this axis.	47
21	These are the numerically-simulated energy levels of the rotating-frame splitting when driven at a gap (Ω) of about 23 MHz. The frequency-detuning-insensitivity-sweet-spot is the point of minimum gap-energy (Ω) and is detuned from the frequency of the rest-frame-qubit. This detuning Δ_ρ is marked with the vertical dashed line.	49
22	X or Y axis $\pi/2$ pulses are sufficient when combined with a single continuously variable gate-type (such as continuous phase evolution when the relative phase between the qubit and its drive evolves as a function of time), to provide universal access to the entire space. . . .	52

23	An example of a $\pi/2$ pulse and its Fourier transform (FT); crucially, the FT of this pulse has sinc-like features which present an opportunity to create regions of intentionally low excitation: wherever the FT goes to zero, interaction with modes at the same location in frequency-space are minimized. This gives us the freedom to minimize leakage into the second-excited state. The time domain pulse-width w_p in nanoseconds determines the profile of the FT.	53
24	An example of a $\pi/2$ pulse in the frequency-domain overlaid (on the right-axis) atop the level manifold of the continuously decoupled transmon. The maximum of the FT is placed at the frequency of the rotating frame computational basis ($f_{\rho 01}$), which is the continuously decoupled qubit. The frequencies at which 2nd-excited-state transitions are present are to the best of our ability positioned at minima of this FT to optimize computational basis confinement.	55
25	An example of the decoupling-drive tone (shown in blue) while a random selection of fast gates (amplitude envelope in red, tone example in orange) is performed on the rotating frame qubit. The decoupling-drive tone is a continuous tone of constant amplitude with a phase that makes it a rotation about the X-axis of the rest-frame basis, while the fast gates have a chosen phase which would make them a rotation about the orthogonal Y-axis in the rest-frame.	56
26	Continuous Decoupling preparation and mapping to and from the rest-frame and the rotating frame. Under typical transmon operation, the qubit is just the ground and first excited states of the transmon in the rest frame. Under continuously decoupled operation, the qubit is now the previously degenerate +/- X states in the rotating frame, split by the Rabi frequency of the drive along the X axis. Readout is performed in the rest-frame.	57
27	A series of scans typically used to determine $T_{1\rho}$ vs. Continuous-decoupling drive amplitude, searching for regions of drive amplitude which have suitably long $T_{1\rho}$ for coherent operation.	58
28	A series of sequences each with a different number of $\pi/2$ X-gates, allowing us to tune-up a high-fidelity $\pi/2$ X-gate.	64

29	Measurements of coherence performed with Hahn-Echo experiments of each the rest-frame transmon operated traditionally (i.e. bare transmon), and the continuously-decoupled state in the rotating-frame gated with our novel technique, along with exponential fit-curves shown in the solid blue and red lines, respectively.	67
30	Measurement of Fidelity of Clifford gates measured via randomized benchmarking, along with exponential fit-curve shown.	68
31	Typically, symmetric flux tunable transmons are operated at the blue points, where the frequency is maximum, called the Upper Sweet Spot (USS). In the case of my experiments in continuous decoupling, the device was prepared at the spot denoted by the orange dot, at a point of high-slope, and thus, low coherence.	70
32	Here, the rotating frame Hamiltonian has been reparameterized as a function of magnetic flux. For proper bookkeeping I am working in the number-state basis N for the drive field and the rotating frame energy eigenbasis for the qubit states. The top left plot shows the full flux-tuning from the USS in Frequency vs Flux to zero frequency. Due to the decoupling-drive, the eigenstates would cross at the chosen frequency of the decoupling-drive tone, and it is at this point where we see our avoided-level crossing representative of the rotating-frame splitting which is our new qubit in the top right plot. The bottom plot is merely the top right plot renormalized to show clearly the point of closest approach, the frequency-insensitive sweet-spot at which we operate the system, such that our new computational basis is made up of the two states shown, $ +\rangle$ and $ -\rangle$	72

List of Symbols

Decay Lifetime = T_1

Coherence Time = T_2

Relaxation Rate = $\Gamma_1 = 1/T_1$

Decoherence Rate = $\Gamma_2 = 1/T_2$

Fundamental Frequency = f_0

Speed of Light = c

Effective Electric Permittivity = ϵ_{eff}

Phase velocity of signals = v_{ph}

Inductance = L

Capacitance = C

Inductance per unit length = L_l

Capacitance per unit length = C_l

Anharmonicity = α

Computational Basis transition frequency in cycles per second = f_{01}

Josephson Energy = E_J

Charging Energy = E_C

Planck's Constant = h

Reduced Planck's Constant = \hbar

Euler's number = e

Linewidth = κ

Mutual Inductance = M

Frequency in radians per second = ω

Unloaded impedance of a circuit or circuit element = Z_0

Coupling capacitance between qubit and resonator = C_{QR}

Coupling constant between qubit and resonator = g_{QR}

Fundamental frequency of a readout resonator = f_R

Sum of capacitances = C_Σ

Difference in frequency between qubit and resonator = Δ
 Difference in frequency between qubit and drive tone = Δ_f
 Pulled frequency of the readout resonator = f_{RP}
 Frequency of spectroscopy tone = f_{spec}
 Frequency of driven qubit oscillations = Ω
 Computational basis frequency of qubit in radians per second = ω_{01}
 Time-dependent voltage applied to a qubit = $V_d(t)$
 Pauli evolutions associated with Pauli gates about X, Y, and Z axes = $\sigma_x, \sigma_y, \sigma_z$
 Wavefunction representing some arbitrary quantum superposition = $|\Psi\rangle$
 A polar phase = θ
 An azimuthal phase = φ
 A ground state = $|0\rangle$
 A first excited state = $|1\rangle$
 A second excited state = $|2\rangle$
 Component of decoherence contributed by pure dephasing = $\Gamma_{pure-dephasing}$
 Decay lifetime of the rotating-frame computational basis = $T_{1\rho}$
 Coherence time of rotating-frame qubit measured via Ramsey scan = $T_{2\rho Ramsey}$
 Coherence time of rotating-frame qubit measured via Hahn-Echo scan = $T_{2\rho Echo}$
 Coherence time of rotating-frame computational basis generally = T_2
 A rotation about the Z axis = Z_ϕ
 Anharmonicity in the rotating-frame = α_ρ
 Gate-pulse duration or width (in the time domain) = w_p
 Frequency-insensitivity sweet-spot detuning of decoupling drive-tone = Δ_ρ
 Absolute frequency of the rotating-frame sweet-spot = f_ρ
 Transition frequency from ground to second-excited state in radians per second = ω_{02}
 Frequency of the decoupling tone in radians per second = ω_d
 Derivative = δ
 Total time derivative = δ_t
 Unitary evolution operator = U
 Inverse transpose of unitary evolution operator = U^\dagger
 Hamiltonian = H
 Transformed Hamiltonian = H'
 Transformed wavefunction = Ψ'

Chapter 1

Introduction

Quantum hardware realizes a new paradigm of computing, networking, sensing, and timing, using fundamental interactions at the smallest energy scale in physics. These systems rely on superposition within and across discrete, quantized computing elements: any quantum element that can be constrained to 2 states and operated in a controlled fashion can be called a quantum bit, or qubit, an example of which is shown in Fig.1. While qubit modalities vary across an expansive spectrum, quantum algorithms—or the framework by which these systems can be used to process information—ideally are conceptually agnostic to the underlying method of generating these 2-state quantum systems. That is, the execution of algorithms should not be directly impacted by whether the qubit itself is comprised of telecom-wavelength photons or excitations of the internal states of trapped ions, or excitations of modes in superconducting circuits, etc. There is *in fact* a nonuniform distribution of algorithm implementations across the various qubit modalities, but this is a mere quirk of circumstance.

Practical considerations, including the susceptibility to error, connectivity, and speed of operations for qubits, do however distinguish the prospective quantum hardware from one another. As a result, each community studying each modality has been tirelessly pursuing ways to improve their qubits. These improvements could come from better fabrication methods, better classical hardware, novel control schemes, unique Hamiltonian design, or other avenues of improvement [1, 2, 3].

There are also growing efforts to mesh qubits from different modalities within the




Figure 1: A qubit is just a quantized two-state system which can be operated as binary information

same hybrid quantum network to combine their strengths (or mitigate their weaknesses) [4, 5]. These efforts seek to leverage modularity as a design philosophy so that specialized quantum processing capabilities can be independently developed and then networked to form a greater hybrid quantum computer, analogous to modular components in modern commercial classical computers of all kinds.

One strategy to mitigate certain practical weaknesses which is broadly in development in the quantum-information community is developing a "virtual" or "encoded" qubit which is separated from the non-idealities of the physical device being used to produce the information. These methods, some of which include dynamical decoupling, continuous decoupling, laser defining, and others, either decouple the computational basis from dominant sources of certain types of noise, or they generate a new computational basis which is a subset of the usual Hilbert space, which is itself less sensitive to certain dominant sources of noise [2, 6, 7, 8].

It is important to note that these "virtual" qubits are distinct from "logical qubits" which are qubits collaboratively simulated by a collection of N physical qubits, where those physical devices are operated identically to agree on a single qubit state; a sort of quantum-error-correction made necessary beyond the usual classical error correction in standard classical computers by the no-cloning theorem [9]. Instead, "virtual qubits" are still just single physical qubits operated in such a way that the physical peculiarities of the chosen device-modality become less important, and qubit coherence

(i.e. the preservation of the a quantum state over parameter-space and time) becomes limited instead by the control tones used to generate this new virtual computational basis [6, 8, 10, 11, 12, 13].

My work has primarily been in superconducting qubits, specifically transmons which themselves were born out of a rich history of innovating to demonstrate and preserve quantum coherence. The very first demonstration of coherent dynamics performed with a superconducting qubit was published in 1999 by Nakamura, Pashkin and Tsai out of the NEC group in Japan [14]. This device was a crucial piece of evidence that coherent quantum behavior could exist in fabricated macroscopic circuits which effectively gave rise to the field of superconducting qubits. The field was so nascent that time domain measurements of coherence of such a system had not yet been performed. Nakamura et. al were the first, measuring 2 ns of coherent oscillation. They predicted coherence times that “could exceed 1 μ s”, modest figures by modern standards.

Since then, vast improvements have been made in fabrication technology and design, giving these qubits weaker avenues for decay, less sensitivity to sources of noise, and controllable readout with designable coupling parameters. One key improvement was the transition from Cooper pair box(CPB) to the transmon detailed in the 2007 paper by Jens Koch et al. from Yale [15]. Like the CPB, the transmon’s dynamics are determined by the interplay of electrostatic charging and Josephson energies. Unlike the CPB whose eigenstates change drastically as a function of applied DC voltages or the movement of nearby charges, the transmon is relatively insensitive to these changes; effectively fixed as a function of charge and thereby has much longer-lived coherence due to the reduction in frequency-dither. Between advances like these, and incredible leaps in the community’s capability to communicate with and manipulate their devices – led on by works such as David Schuster’s thesis entitled Circuit Quantum Electrodynamics out of the Schoelkopf group [3] – the field has advanced several orders of magnitude within only a few decades. Today, lifetimes and coherence times can be in excess of hundreds of microseconds, with some specialized fluxonium devices even exceeding 1ms in coherence time, such as the device discussed in the published paper entitled “Millisecond coherence in a superconducting qubit” by Aaron Somoroff et al [16].

Nearly every subfield which has coalesced around each qubit modality has had

similar stories of advancement and improvement over the years. The key sought-after attribute is fidelity. That is, the ability to prepare any arbitrary sequence such that the measured result is as close to the intended result as possible. (See section 4.4 for a more quantitative treatment.) Each subfield has steadily been pushing toward more viable devices, operating with higher fidelities and building into higher order networks. The ultimate goal, of course, is a large network of many qubits with each qubit coupled to many of its neighbors, such that each qubit can be individually operated with high fidelity, and that many-qubit-gates can also be performed with similar per-qubit fidelity. Unfortunately, each subfield is still short of that goal thanks to its own set of limitations [1].

In the subfield of superconducting qubits and several others, the most prohibitive problems we encounter are additional sources of decay and decoherence by increasing the size of the quantum network or the overall coupling rates to various parameters. Coupling with the environment often takes the form of fleeting interactions with many environmental modes [17] which statistically results in an irreversible loss of information to the environment. These irreversible losses can broadly take two forms: the loss or gain of a photon of energy, which instantiates itself in our quantum systems as a transition between the two states of our quantum bit (i.e. a bit-flip error, similar to that of classical computing), or the loss of certainty in our qubit parameters (i.e. a phase-flip error, which in the domain of computing is unique to quantum bits). The source of bit-flip errors is "fast" noise near the computational basis frequency, and the source of phase-flip errors is "slow" noise well below that energy-scale. The rate at which these bit-flip errors occur is called the relaxation rate $\Gamma_1 = 1/T_1$, and the rate at which phase-flip errors occur is called the decoherence rate $\Gamma_2 = 1/T_2$, where T_1 and T_2 are respectively called the characteristic decay and decoherence lifetimes.

As a hypothetical example, suppose I have a qubit which I've coupled to a classical circuit element. This circuit element limits the lifetime of the qubit to 1 ms, such that if the qubit decays into the classical element, the quantum information has collapsed and is no longer retrievable for the purposes of quantum processing. This is a bit-flip error for which the qubit state transitions from "1" to "0", encouraged by coupling to this classical circuit element. Add nine more of these classical elements with similar effective coupling – or indeed scale up the coupling strength by the commensurate factor – and that 1 ms limit may become less than 100 μ s.

Now, instead of a qubit coupled to 10 classical elements, rather, let us imagine this system is an attempt at an all-to-all 10 qubit network. Our operating frequency range might be from about 1 GHz to about 12 GHz due to hardware limitations, so these qubits would have to be spaced 1.2 GHz apart for independent addressing, and these qubits of course need to be coupled to one-another since we want to operate multi-qubit gates. Now, the presence of these other qubits in the network—the very thing we need in order to build a quantum computer with any real efficacy—is the very feature that limits its own lifetime due to spontaneous decay from any one of these qubits into any of its other neighbors unless error correction can mitigate such mistakes.

Of course, this is a simplified picture, but what it's meant to convey is that this is a non-trivial problem. We need more qubits coupled to one another, and we need to be able to control each of these individually. We also need to perform multi-qubit gates that require a many-qubit-system with moderate coupling, which limits individual qubit lifetimes and pushes gate-fidelities lower, making the prospect steadily worse. On the face of it, it would seem that the system size-scaling we need is antithetical to its own high-fidelity operation. This issue is made yet worse because not only is the decay-lifetime influenced, but so too is the coherence-lifetime deteriorated by this scaling; that is, the characteristic time any one superposition-state can remain coherent. So, not only does building large scale networks cause more bit-flip errors, but its also replete with sources of phase-flip errors. One of the colloquial names for these problems combined with the difficulty of unique addressability is the *frequency crowding problem* [18].

Let's move one step further. So far we've really only discussed the scaling of a many-qubit-system with a single qubit modality. Imagine now we expand to a hybrid-quantum-system among several qubit modalities. Each modality has its own mode of control and coupling, different lifetimes, different coupling strengths and thereby gate-speeds, and specialized hardware. Superconducting qubits all operate around 0.5 – 12 GHz or so in cryogenic systems under ultra-high-vacuum conditions (being that the entire fridge is essentially a cryo-pump when in operation), shielded from black-body-radiation, and even errant single visible photons. Typical trapped-ion qubits operate in dedicated ultra-high-vacuum systems under the influence of large classical DC and RF fields and several high power lasers carrying photons in

the visible spectrum. Photonics qubits require high power pump-beams and several degrees of optical control. Each of these requisite environments is perhaps antithetical to each other in a way that is likely to either inject noise into each respective qubit-environment or force us to use qubits that are suboptimal for coherent operation so that we might better interface across modalities; in either case this means shorter lifetimes, shorter coherence, and worse gate-fidelity, limiting the degree to which we can usefully operate our qubits.

All of these limitations seem to demand a solution that mitigates the degree to which the environment can impact the performance of our devices or indeed, the degree to which the devices parasitically impact each other, namely something akin to dynamical or continuous decoupling. While I know of no evidence to suggest this would fix the problem of decay/lifetime, it would absolutely aid in repairing problems of dephasing/coherence. While continuous decoupling techniques have become very well studied in the experimental quantum information field over the last decade or so, the key capability of operating such decoupled systems robustly and universally has so far been out of reach.

In this dissertation, I introduce a novel method of performing fast, high fidelity, universal single-qubit gates on a system continuously decoupled from coherence-limiting noise. This new technique allows for the operation of qubits that are continuously decoupled from noise in their environment which limits its coherence. Potentially, this may grant us the freedom to operate otherwise low-coherence qubits with high-coherence, or put qubits in environments that would otherwise limit their coherence to unsuitable degrees, and operate them with nominally good coherence despite a deleterious environment. Moreover, this technique is the first demonstration of fast, universal single qubit gates on such a coherence-enhanced system, and may be relevant across the entire field of experimental qubit research. That is, this technique should apply to any qubit irrespective of the physical platform; though I performed this experiment on superconducting qubits, it should apply to any deterministic qubit platform that can be operated with drives in the way I'll describe in later chapters.

Before discussing this, however, we need to review a few topics.

First in Chapter 2, I'll walk through the details of how I design and simulate our devices for fabrication. These devices have an array of crucial parameters that have to be chosen very carefully in order for them to behave the way we want, and figuring

out exactly what geometries to use is non-trivial.

Next in Chapter 3, I'll discuss how we measure these devices and indeed what information we need in order to tuneup the continuous-decoupling control scheme. The course of measurement is important; from establishing readout to extracting the particular parameters needed to inform decisions we make about the control scheme, each step of measurement is often prerequisite for the next.

Finally in Chapter 4, I'll discuss the central result of my thesis – how we universally gate the continuously-decoupled device – and explain our data and results before briefly detailing the Hamiltonian of our system in Chapter 5, and reviewing some possible future research directions related to this novel capability in Chapter 6.

Chapter 2

Simulation, Design, and Fabrication

While this work pertains to just about any deterministic qubit system that can be subjected to arbitrary driven X , Y , and Z rotations on the Bloch sphere, mine was performed on superconducting transmon qubits. One example of transmons I've designed which are at the heart of the experiments I performed for this thesis, is shown in Fig 2. This single superconducting circuit comprised (from bottom to top) of the transmission-line through which travels all input and output signals that we measure, the resonator used for readout of the qubit, and the qubit itself comprised of its capacitor and accompanying Josephson-element, together with a drive-line and a flux-bias line.

In the remainder of this chapter, I outline how each of the key components of this device is simulated, designed, and fabricated.

2.1 Readout Hardware

In the case of the transmission line, thankfully there is very little work to do; coplanar waveguide (CPW) transmission line impedances are well understood [19]. In the case of the typical designer, a quick "CPW impedance calculator" will do the job. 50Ω impedance is the usual standard for most systems in circuit quantum electrodynamics (cQED), which in CPWs on silicon corresponds to a 10-micron wide strip of superconducting metal with 6-micron gaps to ground on either side. This may be modified to compensate for the added capacitance of grounded cross-overs or other structures which may change the capacitance per unit length of the transmission line.

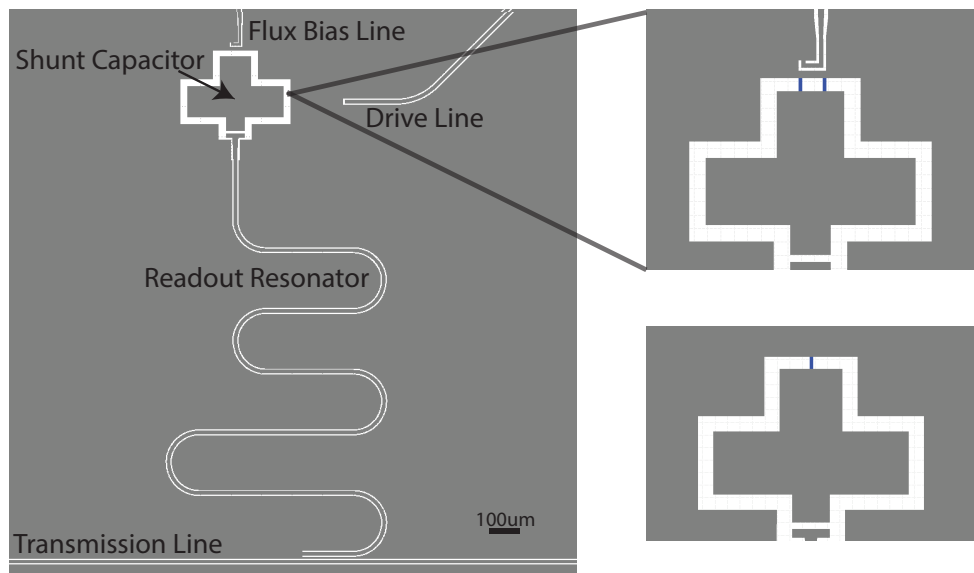


Figure 2: The dark gray is all superconducting metal, and the white is bare-substrate. The inset images are two examples of different Josephson-elements; the top inset image is a Josephson-element of a SQUID made from 2 Josephson junctions with its accompanying flux-bias line to tune frequency, and the bottom inset image is the variant with a single Josephson junction.

Similarly, if a material with high kinetic-inductance is used as the superconductor, the series inductance will also modify the expected impedance. In either case, these may require some simulation and/or calculation to figure out exactly how to compensate and maintain 50Ω impedance to match the rest of the measurement hardware.

In regard to the readout resonator, the impedance is much the same story; the rest of its parameters – most importantly the fundamental frequency which is most critical when designing a resonator for qubit readout – require a bit more work. Typically I would start using rough calculations of quarter-wave CPW resonators as a function of length to choose the total linear path-length of the waveguide.

Combining the capacitance and inductance per unit length calculated or simulated for a given CPW resonator, along with the equation for the frequency of a quarter-wave-resonator, we may find the following [19]:

$$f_0 = \frac{c}{\sqrt{\epsilon_{eff}}} \frac{1}{4l}, \quad (2.1)$$

$$v_{ph} = \frac{c}{\sqrt{\epsilon_{eff}}}, \quad (2.2)$$

$$v_{ph} = \frac{1}{\sqrt{L_l C_l}}, \quad (2.3)$$

$$f_0 = \frac{1}{\sqrt{LC}}, \quad (2.4)$$

as helpful equations for the calculation for the frequency (f_0) of the resonator as a function of its length (l), and the phase velocity (v_{ph}) of signals propagating through it – defined by the speed of light (c) and the effective dielectric constant (ϵ_{eff}), or indeed merely its inductance (L) and capacitance (C), or their per-unit-length values (L_l, C_l) which determine these quantities. This will provide a close enough result to put the drawn design through an EM simulator like COMSOL and have a reasonable frequency window within which to search for the fundamental resonance of the quarter-wave resonator.

In our case, the device I designed had four readout resonators each used to measure one of the four qubits and coupled to the transmission line. Crucially, the fundamental frequencies of these resonators should be far enough apart in frequency that they don't cross one another or overlap when being pulled by their qubits, but close enough together in frequency to make possible any desired multiplexing or to ensure they all land in a good operating regime of parameter space. Perhaps it goes without saying, but we should also be able to identify which resonator on-chip corresponds to which resonance feature in spectroscopy so that we know which qubit we're reading before we attempt qubit-spectroscopy. Being able to easily distinguish the resonance features and match them to the known resonator structure on chip makes measurement much simpler. That is, if we design 4 resonators such that resonator 1, 2, 3, and 4, are in ascending frequency order, perhaps 7 GHz, 7.05 GHz, 7.1 GHz, 7.15 GHz, then we had better be able to trust our simulations and designs within a precision tighter than 50 MHz, else resonator 1 and resonator 2 may come out on top of each other, or swapped in order, and now we can easily make a mistake about which qubit we're actually reading out in a way that can only be detected if we have local lines for individual qubits, or the qubits are distinguishable in measurement due to some other contrasting features.

2.2 Transmon Qubits

The most complicated part of the circuit, the transmon itself has several key parameters that must be chosen carefully.

First, its capacitance to every nearby isolated piece of metal needs to be taken into account. The sum of these values in the case of a grounded qubit like mine makes up the total capacitance of the transmon, which, when paired with the chosen Josephson junction inductance (L_J), defines the gap-energy between the ground-state and the first-excited state – the operating frequency of the qubit. In the case of floating designs, the calculation from capacitance matrices is much more complex than a simple sum, but in either case all must be known to moderate precision [15]. The capacitance alone also defines the anharmonicity (α) in the asymptotic limit of the transmon, which is the energy difference of the transitions between each progressively higher energy eigenstate of the transmon; the exact treatment of the anharmonicity

is detailed in figure 5 of the same 2007 paper by Jens Koch where its sensitivity to the $\frac{E_J}{E_C}$ is described [15]. In this paragraph, I've introduced several concepts at once, each of which calls for an equation (2.5, 2.6, 2.7).

The transition frequency of the computational basis—that is, the frequency of a microwave photon which carries the amount of energy required to excite the transmon from its ground-state to its first-excited state—can be defined at least two useful ways:

$$f_{01} = (\sqrt{8E_J E_c} - E_c)/h, \quad (2.5)$$

$$f_{01} = \frac{1}{2\pi\sqrt{L_J C}} - \frac{e^2}{2hC}, \quad (2.6)$$

$$\alpha = \frac{E_c}{h} = \frac{e^2}{2hC}. \quad (2.7)$$

Equations 2.5 and 2.6 are identical expressions for the computational-basis transition frequency of a fixed-frequency transmon, or indeed the transition frequency at the upper-flux-sweet-spot of a flux-tunable transmon in cycles per second. E_J in Eq. 2.5 is the Josephson energy, and E_c the capacitive energy of the transmon – known as "charging energy". Expressing these values in cycles per second is a common colloquial short-hand, and is merely taking the energy expressed in Joules or eVs and dividing out Planck's constant (h), which is my preferred convention for simplicity. That way, instead of saying that $E_c = 0.8011 \mu\text{eV}$, I would say $E_c = 193.7 \text{ MHz}$, because this value is more meaningful when discussing transmon parameters, specifically because of how we measure them. That is, if the anharmonicity is 193.7 MHz, and the actual computational basis transition frequency is 4.00 GHz, then I know that I need to avoid signals at $4.00 \text{ GHz} - 0.1937 \text{ GHz} = 3.8063 \text{ GHz}$, meanwhile the value "0.8011 μeV " doesn't convey anything to me.

While the charging energy is only a function of the capacitance, which is merely a function of metal areas and geometries, the Josephson energy is only a function of the inductance of the Josephson element as it is defined in fabrication. The Josephson junction is a trilayer of a thin insulating material sandwiched between two superconducting electrodes. When below the superconducting critical temperature, this gives rise to a large, non-linear inductance mediated by quantum tunneling of the charge carrier across the insulating barrier [20, 21] – this is the key element of these kinds of anharmonic oscillators, and it is this inductance that gives rise to the Josephson energy, as shown when comparing Eq. 2.5 and 2.6.

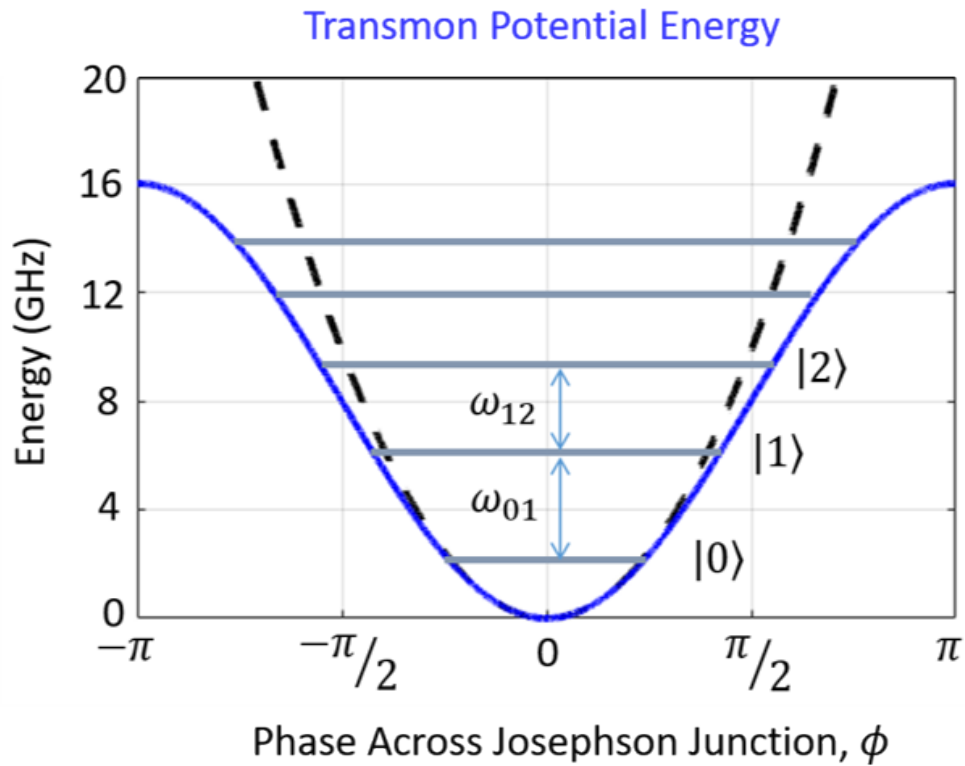


Figure 3: A diagram of the energy level manifold of a transmon qubit, and a comparison between the energy potential of a transmon and the standard harmonic oscillator potential, illustrating the key feature that allows constraint within a qubit basis; the anharmonicity (α).

Here it becomes important to mention that while the Josephson element's attributes are indeed determined in fabrication, the effective inductance can be changed in-situ as a function of local magnetic fields if the Josephson element is designed with more than one Josephson junction to form a Superconducting Quantum Interference Device (SQUID). This behavior is illustrated in Fig. 4 and outlined in detail in Ref. [15], where for an arbitrary superconducting transmon with two Josephson junctions making a SQUID, we get the following equation:

$$E_J = E_{J\Sigma} \cos\left(\frac{\pi\phi}{\phi_0}\right) \sqrt{1 + d^2 \tan^2\left(\frac{\pi\phi}{\phi_0}\right)}, \quad (2.8)$$

where $E_{J\Sigma} = E_{J1} + E_{J2}$ which is the sum of the Josephson energy of each junction in the SQUID, ϕ is the magnetic flux through the SQUID loop, ϕ_0 is a flux-quantum, and $d = (\frac{E_{J1}}{E_{J2}} - 1)/(1 - \frac{E_{J1}}{E_{J2}})$ which is basically a measure of how symmetric the SQUID is, or how similar the two Josephson inductances are. In the case of symmetric SQUIDs where both Josephson junctions are nominally the same (within fabrication limits of precision) we can approximate $d = 0$ and the equation becomes:

$$E_J = E_{J\Sigma} \cos\left(\frac{\pi\phi}{\phi_0}\right). \quad (2.9)$$

This gives rise to a flux-frequency relationship that looks similar to Figure 4. In either equation 2.5 or 2.6, the first term is merely the equation for the frequency of a classical harmonic LC oscillator (in Eq. 2.6 it is written in the standard way, and in 2.5 it is written explicitly in the context of transmon energies), and the second term is the anharmonicity of the transmon. You can see the difference this anharmonicity makes in Fig. 3 which compares the energy potentials of the harmonic oscillator and the transmon. Together they determine the actual transition frequency. The anharmonicity alone is also important at the very least because it sets a “speed limit” of qubit operations. The higher the anharmonicity, the faster one can operate the qubit while keeping it well-behaved thanks to the nature of the Fourier transform of microwave pulses.

Here I find it pertinent to interject a few key concepts and definitions I've been glossing over until this point. What is the difference between a “transmon” and a “qubit”? A qubit is merely a quantized system (i.e. with discrete energy eigenstates) with two states that can be treated as a piece of binary information which has,

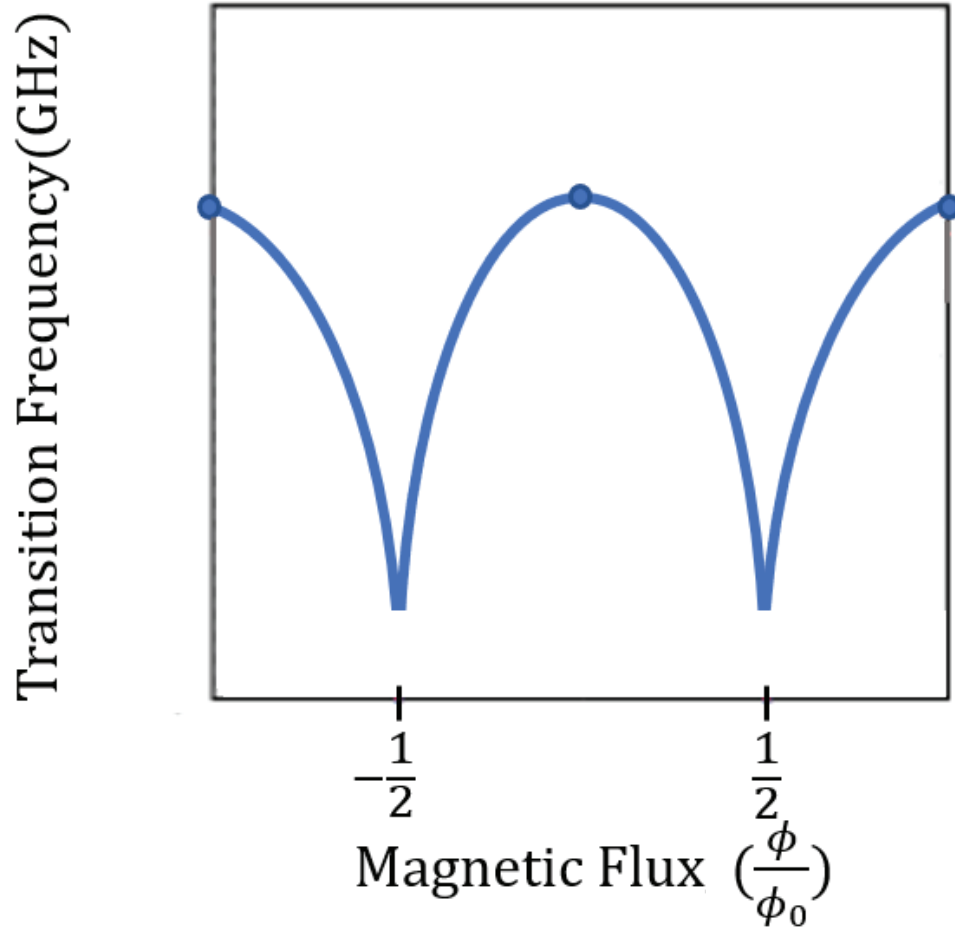


Figure 4: This is an approximation of the flux-tuning of the frequency of a nominally symmetric flux-tunable transmon qubit.

among others, the property of superposition. A transmon is a superconducting LC-oscillator with a small anharmonicity and very low charge-dispersion – i.e. very weak dependence of the energy of the eigenstates on offset charge, which results in very low charge-noise sensitivity [15]. The transmon has many energy levels, but due to its non-zero anharmonicity, it can be operated in such a way as it is forced to remain almost always in either of its two lowest energy eigenstates. As a result, we often call a transmon a “qubit” because they are intentionally designed, fabricated, and operated as a qubit. The key distinction is that technically the “qubit” is just the chosen computational basis of the transmon when it is being constrained and operated within only two particular energy eigenstates (typically the lowest two, because these are the easiest to understand and the longest-lived).

Beyond all of the above key parameters, there are a few other fine details that are important when making decisions about how to design device-geometries. The global question I ask myself when designing a circuit is “how will my decisions impact the qubit’s participation with potential sources of noise?” The answers I generate to this question guide my hand in achieving required parameters while trying to minimize unwanted ones. The following are some examples of choices I make to improve the transmon device itself. Refer to Fig. 2 for a visual example.

1. Wide gaps between transmon and ground to reduce TLS participation
2. Wide transmon electrodes to reduce junction critical current noise
3. Symmetry between the flux-bias line and parasitic M’ loops to reduce flux noise [15, 22]
4. Etched flux vortex traps especially near the transmon to reduce magnetic-field noise

These attributes should not be taken to their extremes, however. For example, if the gaps are too wide, it will be impossible to achieve suitably high shunt-capacitance to the ground-plane and other electrodes. In each of these cases, there is probably some optimum to seek between decreasing the relevant noise contribution and optimizing some other relevant target parameter.

2.3 Coupling

Of course, each of these circuit components needs to interact with each other in the right way to be controlled, manipulated, or measured, and so we have to discuss coupling.

The coupling between the transmission line and the resonator defines the linewidth (κ) of the resonator feature (in the frequency domain) in spectroscopy. This value defines the lifetime and interaction strength of the readout resonator as it relates to coupling to the transmission line. The longer the section of the resonator that runs near and parallel to the transmission line, the stronger that coupling, the larger that κ , and shorter lived and broader the resonator feature is. For any given system being designed, this feature must be chosen carefully. Specifically, in order to optimize readout SNR the value of κ should be about equal to the resonator pull, which is related to the strength of the qubit-to-resonator coupling g_{QR} . I'll define next – a typical value for κ is about 250 kHz, which just so happens to be reasonable to optimize readout SNR without setting too low a Purcell limit (discussed in depth below). In the following, M is the mutual inductance between the resonator and the transmission line, ω is the fundamental frequency of the resonator in radians per second, L is the self inductance of the resonator, and Z_0 is the impedance of the circuit, in this case designed the same between the transmission line and resonator both:

$$\kappa = \frac{M^2\omega^2}{((L + M)Z_0)}. \quad (2.10)$$

The coupling capacitance (C_{QR}) between the resonator and the transmon sets the coupling (g_{QR}) between the transmon and the resonator, and it sets the Purcell decay between the qubit and the resonator—the rate at which the qubit decays from its first-excited state to its ground-state by emitting a photon into a mode of the resonator [23]. It also sets the "resonator pull" or the frequency shift imparted by the qubit onto its resonator. This physical phenomenon is often called the "dispersive shift" between the qubit and its readout resonator, and it is defined by the Jaynes-Cummings Hamiltonian, a generalized form of which most appropriate to the transmon we find discussed in depth by Jens Koch [15]. This shift is crucial to readout, and is the basis of all qubit measurement I performed. We want the coupling

to be strong enough to be able to see the interaction between the qubit and the resonator (since that's how we measure the qubit) but weak enough that the Purcell loss doesn't meaningfully limit the lifetime of the qubit. The details of these are laid out in Chapter 3 and determined by the following expressions:

$$g_{QR} = \frac{C_{QR}\sqrt{f_{01}f_R}}{\sqrt{C_\Sigma C_R}}, \quad (2.11)$$

$$\Delta_{QR} = \frac{g_{QR}^2 \alpha}{\Delta(\Delta - \alpha)}, \quad (2.12)$$

$$T_{1 \text{ Purcell}} = \frac{\Delta^2}{\kappa g_{QR}^2}, \quad (2.13)$$

$$\Delta = f_R - f_{01}. \quad (2.14)$$

In the above equations, Δ_{QR} is the resonator pull caused by the dispersive interaction between the qubit and the resonator, f_{01} is the transition frequency of the computational basis of the qubit, f_R is the fundamental resonance frequency of the readout resonator, C_Σ is the total shunt capacitance of the transmon, C_R is the total capacitance of the readout resonator [15].

The couplings between the transmon and its control lines are also crucial. The coupling between the qubit and the drive-line (or charge-line) sets how much power is needed to be sourced at room temperature to drive the qubit at a particular Rabi frequency, but it also sets the decay-rate from the qubit into the continuum of the drive-line. The stronger this coupling is, the easier it is to drive the qubit strongly with low generator outputs, and the smaller potential crosstalk is to other qubits on the same chip, but the faster the qubit decays into the drive-line potentially limiting the lifetime. The mutual inductance between the flux-bias line and the transmon's superconducting quantum interference device (SQUID) loop sets the level of DC current required through the bias-line for a full flux-quantum of flux-tuning, but also injects additional flux-noise. The higher that mutual inductance, the less current required to achieve a particular range of flux-tuning, but the worse the flux-noise environment [1, 2, 3].

Beyond these factors, I also consider the following:

1. Ground-ground crossovers or other connections meant to make the microwave-ground more robust (and reduce slot-line-modes) and simplify return-current

paths

2. Smooth and large-radius curves in CPWs to reduce impedance mismatches
3. Try to design structures such that changing one parameter doesn't much alter its others (modularity) – an example of this would be designing a transmon such that the coupling capacitance between the qubit and the readout resonator can be altered without modifying the shunt capacitance or the capacitance of the qubit to its drive line.

When designing systems like these, it's incredibly helpful to utilize some sort of a spreadsheet that represents a series of targets and calculated or simulated parameters for a potential qubit chip being designed.

In my case, this spreadsheet combines all of the above equations and a few others so that I can update a particular target value and track the changes that update makes across all parameters rather than having to process these changes individually.

When I'm looking to establish for the first time in a brand-new design, or substantially tweak any of these parameters beyond a few percent, I always use simulation software to acquire parameter values. Equations and approximations are usually good enough to make very small tweaks (within 5 percent) on previous designs, I almost always re-simulate in my favorite EM-solver for the given application for any changes larger than that for any key parameters.

In this case, the devices used were fabricated by Rigetti Quantum Foundry using my designs submitted to their fabrication team.

Chapter 3

Measurement, Analysis, and Tuneup

3.1 Measurement Hardware

First, I should briefly discuss the measurement hardware of which there is an extensive diagram shown in Fig. 5. At room temperature there is a collection of microwave electronics assembled to source signals of desired frequencies and shape pulses of those signals with specific amplitude profiles on timescales of nanoseconds to control and measure the superconducting qubit. The signals from these pieces of room-temperature hardware are sent down the fridge through microwave-lines. These lines are silver plated copper-nickel inner-core, CuNi outer-jacket lines from room-temperature to 3 K and NbTi superconducting lines from 3 K down to the coldest stage. The former are high electrical conductivity low thermal conductivity lines, and the latter are low thermal conductivity superconducting lines at these lower temperatures. On these lines are attenuators at each temperature stage. These are meant to attenuate mostly the thermal noise from 300 K black-body radiation, and effectively cool the signal down to the milikelvin level. After these, there is also a low-pass filter to filter out any noise above 12 GHz, the assumption being that we won't be operating any of our devices with signals at higher frequency than that.

After those attenuators and the 12 GHz low-pass filter, the signal reaches the device itself shown in Fig. 6, after which that signal is amplified thrice through a traveling-parametric-wave-amplifier (TWPA) [24] fabricated by MIT Lincoln Lab, a cryo High Electron Mobility Transistor (HEMT) amplifier, and a room temperature HEMT, and fed back into the room-temperature Keysight PXI chassis and digitizers

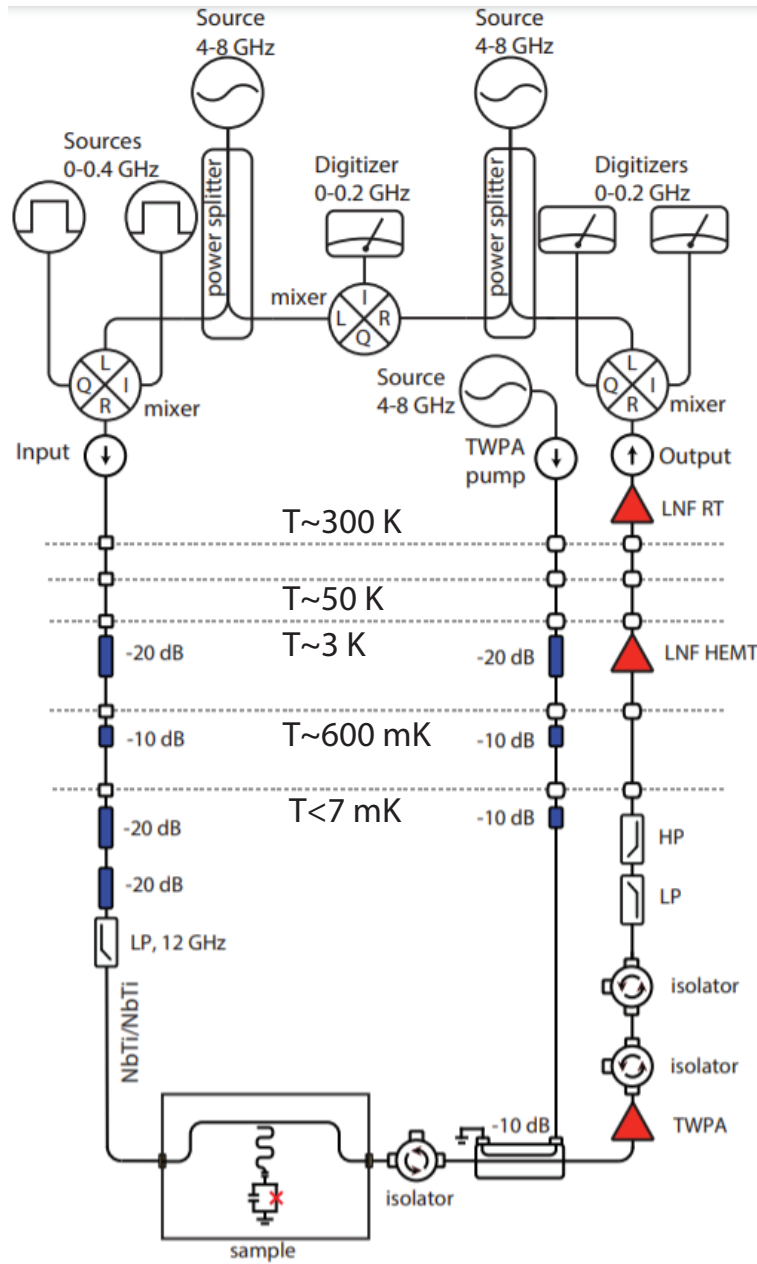


Figure 5: The room-temperature hardware send signals down lines of the dilution refrigerator, which interact with our sample at the coldest stage before being returned to the room temperature data-acquisition hardware.

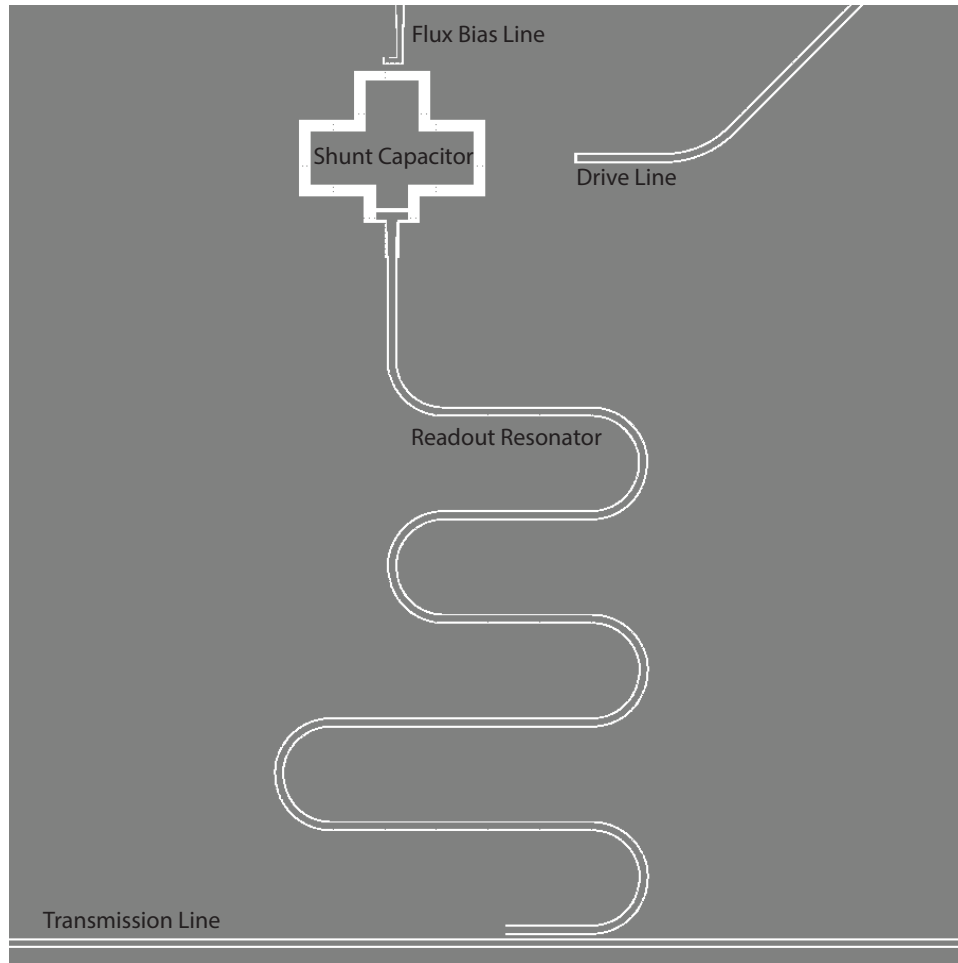


Figure 6: This image provides a reminder of the device geometry; the signal I described above travels through the transmission line shown here, thus interacting with the rest of the on-chip circuit.

where it is sent to a computer and turned into useful data plots.

On the outgoing line, there are filters meant to screen out noise and thermal radiation traveling backwards down that line, and there is a tertiary line for operating the TWPA, which is a cryogenic, nearly quantum-limited amplifier that enhances our Signal-to-Noise ratio (SNR) while distorting the outgoing signal to nearly the minimum allowed by quantum mechanics [24].




Figure 7: Data plot of single-tone spectroscopy performed on a readout resonator through the transmission-line – the center of the Lorentzian feature is the fundamental frequency of the readout resonator (f_R)

3.2 Spectroscopy for Qubit State Read-Out

The first task I typically perform is single-tone spectroscopy to locate each resonator in frequency-space. The easiest way is to measure the S_{21} —the ratio of output and input microwave energy—of the transmission line and search for sharp dips:

In Fig. 7, the Y-axis is a measure in volts of the signal amplitude through the transmission line. The frequency at which this sharp dip in transmission appears is the frequency at which the signal destructively interferes with the signal reflection off the resonator at its resonance frequency (f_R).

During this measurement, since the device is at a base-temperature low enough to suppress the thermal occupation of qubit excited states to near zero percent (residual thermal occupation is extremely low at mK temperatures given operating frequencies of devices like ours) [25], we can safely assume that the resonator frequency we observe during this measurement is indeed the resonator’s frequency while the qubit is in its ground state. At least, that is assuming that the power to which we’ve set our tone

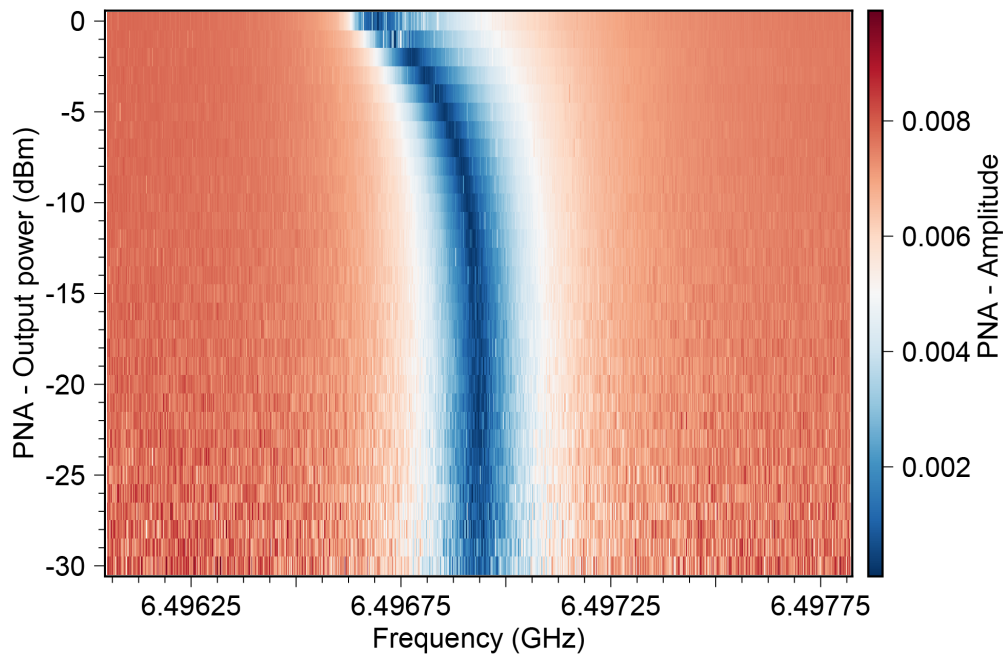


Figure 8: This is a power scan of a resonator in spectroscopy; at low powers the resonator’s fundamental frequency does not vary as a function of power, but at higher powers we see an abrupt change in the resonance.

is low enough not to populate the resonator with too many photons and accidentally excite our qubit. We can test this directly, as a matter of fact, by scanning that power as shown in Fig. 8.

When the resonator frequency starts to rapidly veer off as a function of increasing power, we can safely conclude that power is high enough to accidentally excite the qubit at least some of the time. When choosing the power of our resonator spectroscopy tone, we make sure the power is high enough that our signal-to-noise ratio (the difference in signal amplitude between when we’re on resonance with the resonator and when we aren’t) is quite good, but low enough that we haven’t begun unintentionally exciting our qubit.

Importantly, the behavior of the resonator as a function of this unintentional excitation demonstrates something crucial to our measurement: when the qubit is in the first excited state, the resonator’s position will be different than when the qubit is in its ground state [26].

In Fig. 9 is an example of the resonator being pulled from one frequency (f_R) to

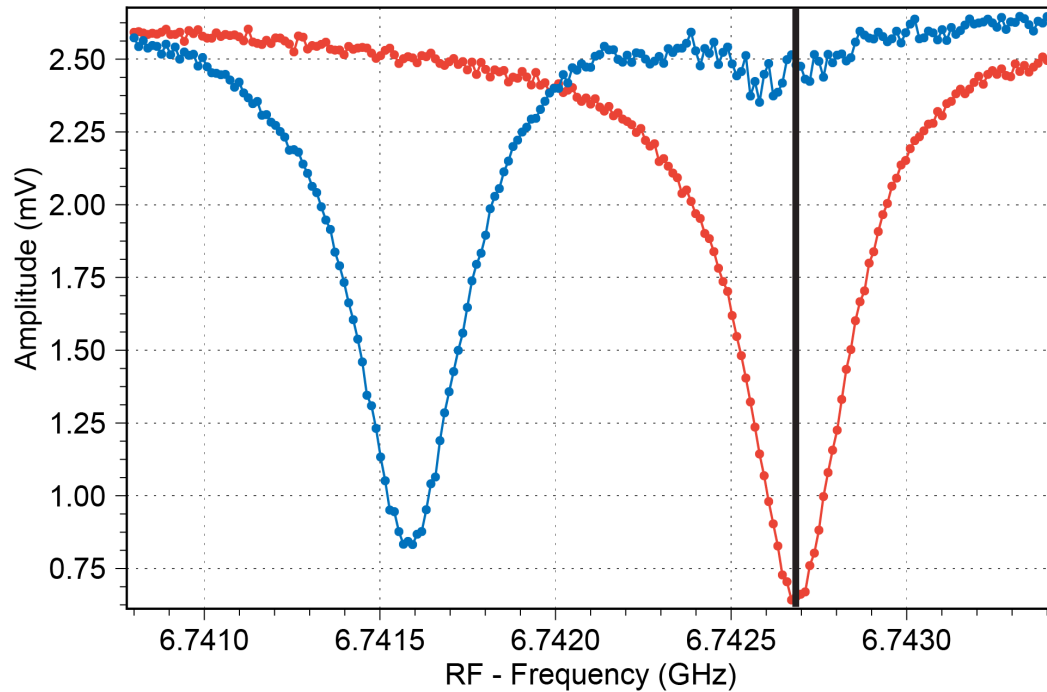


Figure 9: Two scans of the readout resonator feature, the red curve is taken when the qubit coupled to it is in the ground state, and the blue curve when the qubit is in the first excited state. The black line is the marker which represents the frequency (f_R) where we would then choose to position our readout-resonator-probe-tone for future two-tone-spectroscopy. The difference in the frequencies between the center of these two features is Δ_{QR} .

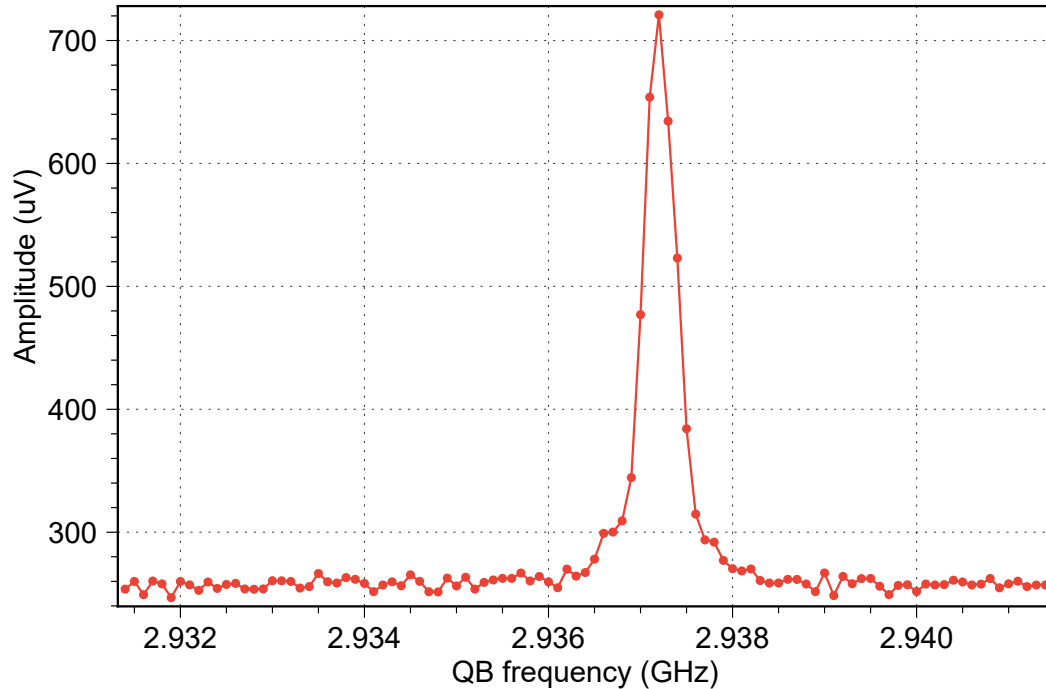


Figure 10: Two-tone spectroscopy performed by reading out the amplitude of the tone at the resonator’s un-pulled resonance frequency (f_R), while sourcing a second tone at (f_{spec}) trying to excite the qubit.

another (f_{RP}) due to the excitation of the qubit from the ground state to the first excited state. This gives us a good way to set up the readout of our computational basis. If we’re always looking at the frequency of the resonator when its qubit is in the ground state as implied by the black marker on Fig. 9, we’d read an amplitude of just under 0.75 mV as shown where the marker meets the red curve; this corresponds to the ground state of the qubit. If when we send in our resonator-tone, while our qubit is in the first excited state, we would read an amplitude of about 2.5 mV, as shown where the solid black line meets the blue curve.

3.3 Two-Tone Qubit Spectroscopy

At this point, while probing the readout frequency (f_R), we send in a second tone in an attempt to excite the qubit, and see the readout amplitude shift from low to high.

Each of the data points on the plot shown in Fig: 10 (in this particular case) is 1000

measurements of the resonator amplitude being averaged together. Before our second tone arrives near the qubit resonance (f_{01}), we read nearly all 1000 measurements at the point of lowest amplitude with the qubit in the ground state. As we approach near to the qubit resonance, the qubit occasionally excites into the first excited state, therefore, some of the 1000 samples appear at the point of high amplitude, and some appear at the point of low amplitude; these average to some middle value. At the very center of the feature, the point of highest apparent amplitude over almost all 1000 samples, is where the qubit is the most excited on average. It is this point that represents where the population was highest in the first excited state, and is approximately the actual resonance frequency of the qubit; in other words, the transition energy f_{01} between the ground state and first excited state of our computational basis.

3.4 Qubit State Manipulation

At this point, we have established readout of the computational basis, and located the transition frequency between the two eigenstates in that basis. Next, I'll tune-up a π -pulse—the pulse which fully excites the qubit from the ground state to the first excited state.

I will choose the qubit frequency to be the center of the spectroscopy feature like the one shown in Fig: 10. This isn't exact, but it's close enough for now and can be optimized later.

We set the first microwave tone to the central frequency of the un-pulled resonator feature (f_r), and the second tone to the transition frequency we found for the qubit (f_{01}), and pick some nominal duration. Typically, I choose around 50-100ns in total duration. Short enough to be very short in comparison to relevant lifetimes of the qubit, but long enough to be relatively gentle – low in total power compared to relevant energy scales. This varies, and may be very different depending on the system in question.

I then scan the amplitude of the pulse, and receive a plot similar to the following:

At zero amplitude, of course the qubit is not being excited at all. As I increase the amplitude of the pulse, the average state population in the first excited state increases until some maximum (ideally 100 percent), before decreasing again. When the qubit is in its ground state, a tone incident upon the device may or may not get




Figure 11: This is a data plot of coherent oscillations of the qubit as a function of drive power. A pulse of constant duration is sourced at the qubit transition frequency (f_{01}) and its amplitude is scanned, producing this oscillation between the two eigenstates.




Figure 12: This is a data plot of coherent oscillations of the qubit as a function of drive duration. A pulse of constant amplitude is sourced at the qubit transition frequency (f_{01}) and its duration is scanned, producing this oscillation between the two eigenstates.

absorbed and cause an excitation to the first excited state. The chances of that tone in-fact exciting the qubit into the first excited state is a function of how close that tone is to the transition frequency, and how close it is to the right amplitude. That is, when the total power applied for some duration at the set frequency f_{01} is equal to the energy of a single photon of frequency f_{01} , the chances of exciting the qubit are optimally high. If the total energy applied is more than that, there is a chance that the qubit gets driven into the first excited-state and then subsequently driven back into the ground state during the same pulse, which causes the oscillation as a function of drive-power. This is a type of “Rabi oscillation”, specifically a Rabi oscillation as a function of drive amplitude, or a “Rabi amp”, colloquially.

An important measure of the apparent strength of the drive seen by the qubit is a Rabi Oscillation as a function of duration as shown in Fig. 12.

The frequency of this oscillation is often written as Ω and is representative of the

energy scale of a resonant drive tone, or the apparent amplitude of a resonant drive which the qubit sees. More aptly, Ω is a measure of the rate at which the qubit occupation in the computational basis evolves in the presence of such a drive-tone of a particular amplitude. That is, the Hamiltonian of the qubit has a time independent portion and a time dependent portion, the latter of which is modified by apparent drives that change the time-dynamics of the system. This relationship is defined as [2]:

$$H = -\frac{\omega_{01}}{2}\sigma_z + \Omega V_d(t)\sigma_y \quad (3.1)$$

The latter of these two terms is the time-dependent portion which varies as a function of the drive amplitude Ω , and the associated normalized time-dependent function $V_d(t)$. This determines how the qubit responds to an apparent RF-drive. This measurement as described is only accurate if the frequency of the tone and the transition frequency of the qubit are the same (if the tone is on resonance with the qubit), otherwise there is yet an additional term which increases the apparent Rabi frequency due to additional azimuthal rotation as a function of detuning. In fact, one of the ways I might choose to better tune up the frequency of my drive-tone is to take multiple of the above scans at different values of frequency detuning (Δ_f) and minimize the Rabi frequency at a given drive-amplitude. The drive-tone-frequency at which the Rabi frequency is minimized is the computational basis gap energy (the qubit frequency). An example of this kind of measurement is as follows where each vertical cut through the colormap in Fig. 13 is a single Rabi oscillation plot.

3.5 Qubit Characteristic Lifetime and Decoherence time

So at this point, we have successfully identified the transition frequency of the qubit computational basis and tuned up a reasonable π -pulse, which brings the qubit from the ground state to the first excited state. Depending on the pulse-duration chosen until this point, and the resolution of the plot (like Figure 13) used to examine the exact qubit frequency, this π -pulse may be pretty low fidelity, but even in that case is good enough to measure rates of decay or decoherence to sufficient precision.

Now we can make our first measurement of a characteristic-time of the qubit. The plan here is to pulse the qubit from the ground state into the first excited state, and

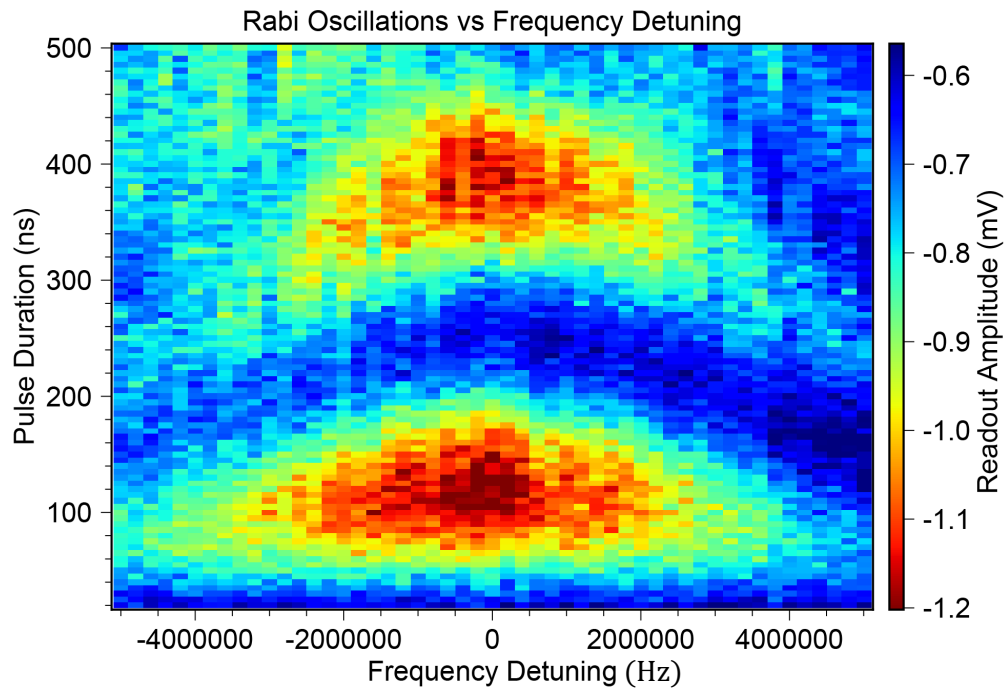


Figure 13: This is a colormap of coherent oscillations over time made at different frequency detuning (Δ_f). The oscillations are slowest when the tone frequency matches the qubit frequency. Here, the colorscale represents the measured probe amplitude, which is the established proxy for qubit-state.

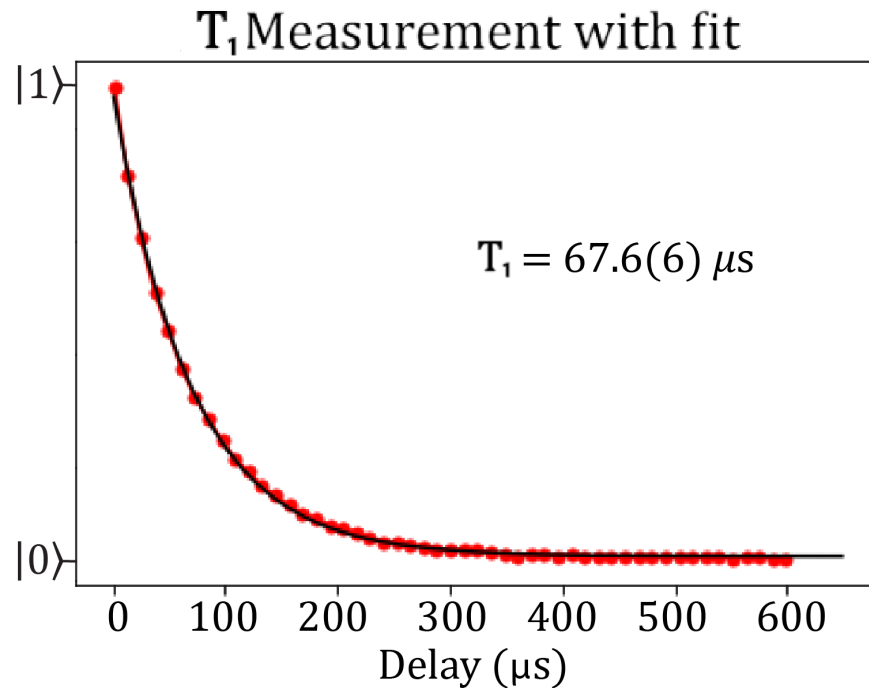


Figure 14: A scan of the decay time T_1 , taken 36 times and averaged. Each datapoint on this plot represents 36,000 measurements of the qubit state at each time after preparation in the first excited state all averaged together.

then wait to see how long it takes for the qubit to decay. This is a measurement of the qubit's relaxation, or decay time. The average of many such scans might look as shown in Figure 14.

The scan in Figure 14 scan is the average of 36 individual scans; a single scan will almost never look quite this smooth. Within a single scan, the first datapoint is a measurement of the qubit immediately after we finish sending in the π -pulse. That datapoint is 1000 measurements of the qubit-state averaged together, almost all of which should be in the first excited state at this point assuming we've tuned up our pulse well. After that point, each datapoint in the scan waits a little longer before measuring the qubit state; this shows us an exponential decay that corresponds to the rate at which our qubit decays out of the first excited state back to the ground state. This plot, when fitted to an exponential decay, has a decay constant in units

of inverse-seconds (Γ_1), the inverse of which we call “ T_1 ”, the inverse characteristic decay constant that defines the rate at which the system decays. This value is the amount of time that must pass before the system decays from 100 percent first excited state population to $\frac{1}{e}$ of that. This is a measure of how long the qubit can live in the first excited state before spontaneously emitting a photon into its environment and relaxing into the ground state.

This value can be limited by a whole host of noise-sources usually called “transverse noise”, including thermal noise and quantum Nyquist-noise among others [2]. It can be roughly understood as how many other parasitic modes are nearby that can accept a photon of the qubit frequency and what their coupling strength is to the qubit. This can include intentionally coupled devices, like the resonators we use for read-out (Purcell decay into the resonator), two-level-systems (TLS) in the environment – most typically atomic-scale defects in native oxides formed on superconducting metals that are near to the qubit, or relaxation into the vacuum to name a few. Relaxation into this overall environment can also be driven to become more likely by noise in the electromagnetic field in the qubit’s environment (think stimulated emission).

The next step is to cut the π -pulse’s amplitude in half to make a rough $\frac{\pi}{2}$ pulse, which gates the qubit from the ground state to the equator of the Bloch Sphere, a state of $|+x\rangle$. This is a coarse method that provides a pretty low fidelity $\frac{\pi}{2}$ pulse. It can be tuned up to high fidelity of course, but this rough preparation is often sufficient for a basic T_2 measurement.

Figure 15 shows the Bloch Sphere, which is a representation of the superposition of states that a given qubit is in. The North and South poles are the ground state and first excited state respectively (although the choice of which is arbitrary and merely a matter of convention) and each point on the surface of the sphere represents a different superposition of these two states. Starting at one of the poles, the π -pulse brings us directly from 0 to 1 or 1 to 0. More generally, it effectively flips the state-population of any given superposition (this can perhaps be thought of as a NOT gate, in traditional digital logic).

The value along the vertical Z axis in the Bloch sphere basis is representative of the state occupation – whether the qubit is in the ground state or the first excited state. The X and Y axes are representative of the real and imaginary relative phase between the drive-tone and the qubit; also understood, respectively, as the real and




Figure 15: The conventional representation of superposition states of the qubit; the Bloch sphere.

imaginary superpositions of the two eigenstates of the computational basis. What I mean here by "relative phase" is that the qubit is at some frequency, and in the case of a traditional π -pulse, we can say that the drive-tone is set at the exact same frequency. The relative phase between the drive-tone and the qubit wavefunction corresponds to a vector on this sphere; the key is that this phase is a complex one. The phase can be real, imaginary, or some combination of both, which is why the Bloch-sphere has 3 axes; the Z axis to define the energy difference between the two eigenstates, and the X and Y axes to define the real and imaginary components of the relative phase between the qubit wavefunction and the drive-tone's. These superpositions can be written with the Bloch equation in the following form:

$$|\Psi\rangle = \cos(\theta/2) |0\rangle + e^{i\varphi} \sin(\theta/2) |1\rangle. \quad (3.2)$$

Rotations about axes on the Bloch sphere are evolutions of θ or φ in Eq. 3.1. The $\frac{\pi}{2}$ pulse is a gate which gets us a quarter turn around the sphere, bringing us, for example, from 0 to the equator, then from that point on the equator to 1, and so on. If the $\frac{\pi}{2}$ pulse is about the y-axis and we begin at the north pole, it will rotate the Bloch vector from 0 to +x, then +x to 1, then 1 to -x, then -x to 0, and so on. These gates can be rotated in phase, such that the user can choose an x $\frac{\pi}{2}$ or a y $\frac{\pi}{2}$, or some combination. Combining and changing this phase intentionally and dynamically is a kind of virtual Z-gate, and it can allow the user to reduce the total number of gates required to perform a given operation. Similarly, anything that allows phase evolution between the qubit and the drive-tone can be a Z-gate. This would include establishing (Δ_f) between the qubit frequency and the drive-tone.

Here, we just plan to choose a single $\frac{\pi}{2}$ pulse of some phase we can call γ . So, our trajectory will be 0 to +x, +x to 1, 1 to -x and -x to 0.

With this $\frac{\pi}{2}$ pulse we can perform a Ramsey measurement. Here, we send in a single $\frac{\pi}{2}$ pulse, then wait some period of time and send in another. Crucially, the pulses we program for this measurement are intentionally detuned to a frequency that is different from the qubit-transition frequency. This difference between the gate-pulse and the qubit transition frequencies lead to a relative phase accumulation between them which produces an oscillation at the frequency of that (Δ_f) in our qubit superposition. That is to say, if the qubit frequency is 4.05 GHz, and we send

a gate-pulse intentionally set to 4.052 GHz, there will be a 2 MHz oscillation in the scan. This is the same as Z-gating, or rotating about the Z axis, at a rate of 2 MHz. In order to see this oscillation, all we have to do is sweep the duration of the wait-time between the first and second $\frac{\pi}{2}$ pulses; one such scan is shown in Fig. 16.

This scan will oscillate from near the first excited state to near the ground state as shown, and will decay over time as the qubit “dephases”. The phase accumulation described can also be understood as a dynamic oscillation around the Bloch Sphere. When the first $\frac{\pi}{2}$ pulse is seen by the qubit, the Bloch vector rotates from 0 to +x, at which point it will begin to precess, or rotate around the Z-axis before being gated again by the second $\frac{\pi}{2}$ pulse. The distance covered during this precession is set by the (Δ_f) chosen and the wait time in between pulses. When the qubit sees the second $\frac{\pi}{2}$ pulse, it will have precessed some amount and will be rotated to a new position before being read-out. The change in the wait-time provides the characteristic oscillation.

The decay of this waveform is due to decoherence, the majority of which in this case is the result of pure dephasing. There is an uncertainty in exactly what the phase between the qubit and the drive-tone is, and it is this uncertainty getting larger and larger as a function of time after any given state-preparation which gives rise to the decay of the envelope of this oscillation. We call this increase in the overall uncertainty in the phase as a function of time “decoherence”. The rate of decoherence is (Γ_2) the inverse of which is a characteristic time we call T_2 , which is the time any given superposition will survive before being a factor of $\frac{1}{e}$ closer to indistinguishable from the flip of a coin. This occurs due to the increasing uncertainty in the relative phase between the qubit and the drive-tone being used to gate it. One useful visualization of this concept is that when a superposition is initially selected, the Bloch sphere diagram might start with a well defined Bloch vector, and evolve over some time to a less well-defined superposition, as is shown in Fig. 17

The superposition becomes less and less well defined, and eventually the area that represents the superposition(s) will (absent any interactions which would otherwise change the system) encompass the entire Bloch sphere, such that the uncertainty is maximum and any measurement of the state of the qubit results in a perfect 50/50 coin-flip in response. This can also be discussed as a function of the Bloch-vector magnitude. A Bloch-vector amplitude of 1 is a pure state with perfect certainty,

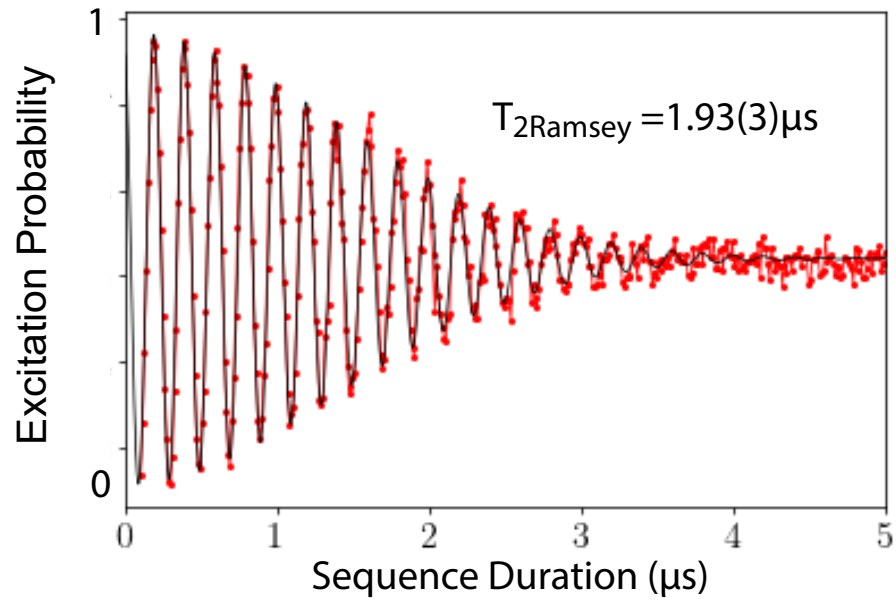


Figure 16: A measurement of decoherence on an intentionally low-coherence device performed with a Ramsey experiment. The noise responsible for the dominant proportion of the decoherence is pure-dephasing from flux-frequency noise on a flux-sensitive transmon.




Figure 17: A demonstration of a way to visualize decoherence on the Bloch sphere. The superposition starts well-defined, and progresses from blue to orange to red, smearing out as a function of time.

and a Bloch-vector amplitude of 0 is a perfectly mixed state of all superpositions. A preparation which has smeared out to encompass the entire Bloch sphere is analogous to a Bloch vector that has decayed to zero length, and a preparation which is of infinitesimal area is analogous to a Bloch vector with a length of 1.

This dephasing can be made faster by any uncontrolled or unknown noise or changes in either the phase of the qubit or the phase of the drive-tone itself, since the important quantity is the relative phase between them. That is, any physical process which changes the frequency of the qubit during the measurement or operation will deleteriously impact the T_2 , but similarly any physical process which changes the frequency or phase of the drive tone being used will also contribute to dephasing in the same way. Decoherence as a whole can be represented as follows, where the noise-related dephasing I've been talking about is represented by $\Gamma_{\text{pure-dephasing}}$ [27].

$$T_2 = \left(\frac{1}{2T_1} + \Gamma_{\text{pure-dephasing}} \right)^{-1} \quad (3.3)$$

Another method of measuring the decoherence of a system is a Hahn-Echo experiment. This is essentially a Ramsey experiment with a π -pulse of orthogonal phase exactly in between the two $\frac{\pi}{2}$ pulses in time. This additional pulse effectively unwinds noise from the first half of the sequence that is present and the same during the second half of the sequence. Any noise which is correlated with time, or constant, or otherwise non-randomly distributed, will have a tendency to be unwound by a Hahn-Echo pulse sequence. This typically results in higher T_2 times assuming there are some non-random noise sources in the system and the system isn't already T_1 limited; the average of many such scans might look like Fig. 18 (a single scan will not be so smooth):

In the case of a Hahn-Echo experiment, every pulse sequence is attempting to prepare the first excited state, and the degree to which this preparation fails is the measure of decoherence in this case.

The difference between the $T_{2\text{Ramsey}}$ and the $T_{2\text{Echo}}$ can be thought of as a measure of non-random noise in the system. The better $T_{2\text{Echo}}$ is when compared to $T_{2\text{Ramsey}}$, the more noise was unwound by the phase-evolution-inversion of the π -pulse in the middle of the pulse-sequence, which can inform the user about the type of noise impacting their measurements.

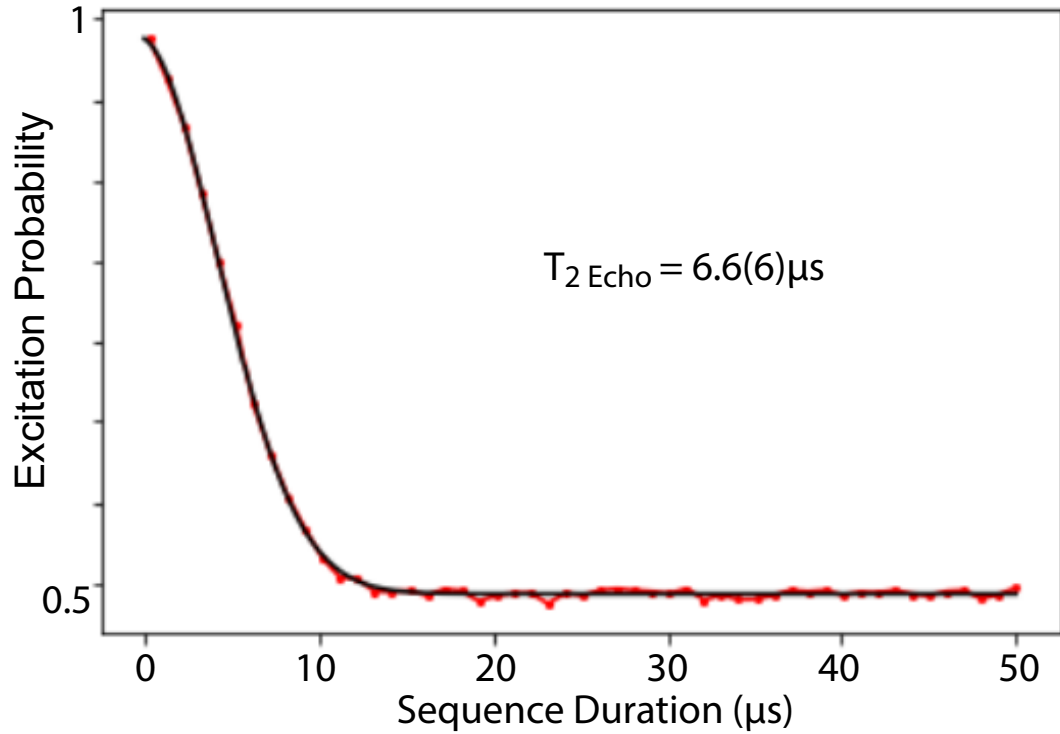


Figure 18: A measurement of decoherence on an intentionally low-coherence device performed with a Hahn-Echo experiment. The noise responsible for the dominant proportion of the decoherence is pure-dephasing from flux-frequency noise on a flux-sensitive transmon. Some of it is mitigated by this measurement technique.

3.6 Second Excited State and Anharmonicity

The only remaining detail typically needed for operating superconducting transmon qubits is the anharmonicity of the system. Chances are, estimates of this value extracted via capacitance simulation are pretty close to accurate, however measuring the value is easy.

Since we already have our π -pulse, all we need to do is prepare a gate-sequence which begins with a π -pulse, and concludes immediately after with a spectroscopy pulse of some reasonable duration and amplitude, which we scan in frequency in order to find where the 1 to 2 transition is. This scan might look as shown in Fig. 19

In the case shown, the center of the feature which I've marked with a black line, shows the frequency detuning between the computational basis frequency and the frequency of this transition. This is a direct measure of α . It pays to mention, however, that pending additional scans to understand how the readout resonator behaves when the qubit is in the second excited state, the readout may not look so qualitatively similar to the computational basis spectroscopy. Extracting this information may require looking for the resonator position in the ground state, first excited state, and second-excited-state, and reading out differently. Without this, the readout here might have very low signal-to-noise-ratio (SNR), or might have a very different shape. This is because the dispersive interaction between the resonator and the computational basis eigenstates is different than the dispersive interaction between the resonator and the second-excited state of the transmon.

We now have all of the information required to start trying continuous decoupling with the intent of performing universal gates on the decoupled splitting in the rotating frame.

Before we actually begin doing any continuous decoupling, it would be extremely useful to measure the relationship between the drive-tone amplitude we choose in our control software and the actual Rabi frequency (Ω) at which the qubit is driven. Whether we're setting the power output level in dBm of some generator or choosing the bias voltage of some mixer or amplifier, we need to know the relationship between whatever knob we want to turn and the actual Rabi frequency (Ω) of the qubit. We need a 3-dimensional series of scans which takes some pulse of some particular amplitude and sweeps its duration, and then steps the amplitude and repeats the

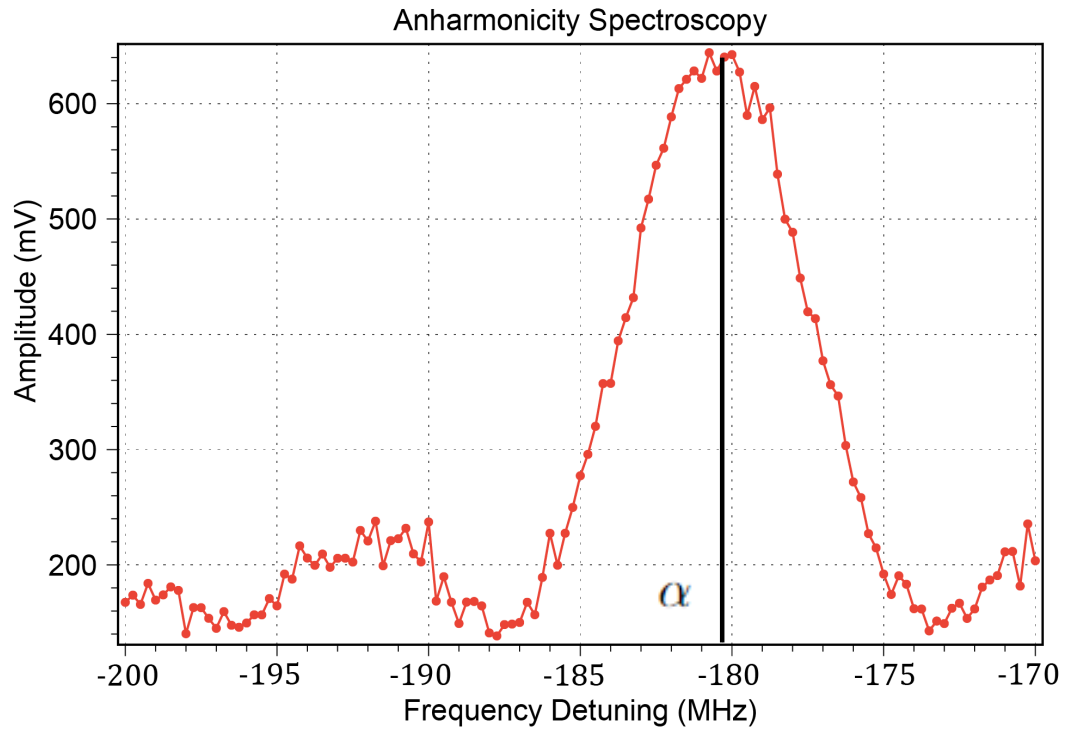


Figure 19: A measurement of the transition between the first excited state and the second-excited state made by exciting the qubit into the first excited state before scanning a pulse through frequency-space searching for the second-excited state transition.

sweep for some range of amplitudes. One of these sweeps would look like Fig. 12, with each one having faster oscillations than the last through the entire series of scans.

The higher the amplitude of the drive-tone selected, the faster the Rabi oscillations. When continuously decoupling a qubit, the splitting gap-energy we generate will be nearly equal to the Rabi frequency (Ω) of the decoupling-drive-tone, so knowing the relationship between room-temperature-hardware-settings and actual Rabi frequency (Ω) in advance saves a lot of headache.

Chapter 4

Continuous Decoupling and Fast, High-Fidelity Universal gates

Longitudinal noise is a crucial limiting factor for operating qubits in large-scale qubit platforms. Protecting qubits from this ubiquitous dephasing is especially critical when interfacing computational devices with noise-dense environments to facilitate quantum networking, quantum transduction, and quantum computation. Further, many qubit species have constrained operational ranges, such that only a particular parameter-range allows for high-coherence operation and relaxing these constraints may broaden potential applications. Towards these ends, I have performed fast, high fidelity, universal single-qubit gates on a flux-sensitive transmon qubit far from its flux-insensitive sweet-spot while continuously decoupled from longitudinal noise. Using a symmetric flux-tunable transmon qubit intentionally tuned far from its flux-insensitive upper-sweet-spot (USS), I drove a 23 MHz splitting between the otherwise degenerate $+x$ and $-x$ superpositions. I operated this splitting as a qubit, improving coherence over traditional operation by an order of magnitude. I gated this splitting with 99.44 percent fidelity as measured via randomized benchmarking. Using the same technique, I demonstrated fast universal gates on a high coherence fixed-frequency device, measuring a fidelity of 99.85 percent. This method vastly reduces a qubit's sensitivity to longitudinal noise allowing high-fidelity qubit operation in regimes previously prohibited by sources of decoherence, thus a potentially useful tool for operating qubits in poor environments, or operating noise-sensitive devices.

In the remainder of this chapter, I present the central results of my thesis work;

the first performance of fast, universal, single qubit gates on a continuously decoupled qubit.

4.1 Longitudinal Noise and Continuous Decoupling

In order to perform useful algorithms on a quantum system, the quantum information sciences community seeks to maximize the fidelity and number of operations that can be performed within a single qubit-lifetime. This means that we want long lifetimes, the ability to implement rapid, high-fidelity universal commands in arbitrary sequences, and of course, we want to be able to reliably read-out the device.

As outlined chiefly in Chapter 3, coherence of a quantum system sets the upper limit of the total amount of time any single calculation or algorithm can take before the result is no different than flipping a coin. Gating, or manipulating a qubit from one superposition state to another must be done quickly and correctly, else an algorithm that requires several hundreds of manipulations will fail to complete before the prepared superposition dephases or accumulates a prohibitive amount of error; this demands high fidelity gates that are short in duration when compared to the qubit coherence times (e.g. $T_{2\text{Ramsey}}$ or $T_{2\text{Echo}}$).

Within experimental quantum information communities, the focus for many modern research efforts has been departing qubit regimes dominated by longitudinal noise—that is, dither in the computational-basis' gap-energy which reduces qubit lifetime. Unfortunately, creating devices that are less sensitive to sources of noise in the qubit environment may not alone be sufficient.

The reality of qubit technologies is that in order to build a quantum computer of any reasonable scale, capable of actually performing algorithms and demonstrating quantum supremacy [28], qubits are often by necessity placed in an environment that is full of longitudinal noise – an environment antithetical to long coherence. Real quantum processors will be replete with the following non-idealities: charge noise, flux noise, resonator crosstalk, single qubit gate crosstalk, ZZ-coupling, direct parasitic capacitance between qubits, local TLS noise, "1/f" noise broadly and others [2, 29, 30].

Introducing new control parameters or increasing physical coupling between a qubit and classical control-lines tends to increase total decoherence. As described

earlier, this is the result of either adding new channels for decay and dephasing or increasing the rate at which decay and decoherence events occur in channels already present; the same is true for enhancing the qubit's ability to see other quantum systems by increasing the physical coupling parameters between them.

These issues are likely to become an especially important challenge as quantum information systems of different modalities are interfaced with one another, i.e. hybrid quantum systems, for quantum transduction or other advanced functionalities that rely upon the respective favorable attributes of multiple types of quantum hardware components [5, 31, 32, 33, 34, 35]. Generally speaking, optimization of the performance of each quantum modality with respect to material properties, wavelengths, and geometries is unique and can be non-ideal for the other modalities. In such hybrid applications, sufficient coupling between modalities is a challenge that must be overcome, potentially increasing sensitivity to the very worst of the many sources of longitudinal noise. One important example of a challenge that will need to be addressed is the deleterious impact of quasiparticles on superconducting circuits. This is because the wavelengths of light (infrared, visible, and uv wavelength) that are traditionally utilized for quantum information processing in quantum photonic and trapped-ions platforms are sufficiently high energy to break Cooper-pairs and generate quasiparticle excitations in large numbers when they are absorbed by the superconductor [36, 37].

Noise-decoupling control techniques can preserve coherence in operating regimes dominated by longitudinal noise. From nitrogen-vacancy center spin qubits to ultra-cold gasses, continuous decoupling has been studied extensively, and has produced results that show a formidable difference in the coherence of such qubits during free vs driven evolution [11, 12]. In the field of NMR and in superconducting qubits, a similar driving technique has been leveraged as a metrology tool these communities have termed "Spin-locking", which allows for examination of long timescale noise and noise sources at small detunings from the gap energy of a qubit [39, 38, 40, 41, 42].

Qubits under strong drives of the correct form are continuously decoupled from longitudinal noise. When applied correctly, a continuous drive can be used to transform the qubit basis into a rotating frame with very different properties from the rest-frame device. The computational basis transition-energy from the ground-state of this new rotating-frame qubit to its first-excited state has a gradual curvature as




Figure 20: The Bloch sphere of the bare transmon overlaid with a drive along the x-axis. The following demonstration of continuous decoupling is prepared with the continuous drive and Bloch-vector (red arrow) colinear along this axis.

a function of frequency-detuning (Δ_f), and crucially, a point at which the slope goes to zero as shown in Fig. 21.

This is the defining feature which gives rise to coherence enhancement. What I mean is, the primary source of dephasing in traditional qubits is frequency-dither i.e. the actual frequency of the computational basis – or the microwave-tone used to gate it – changing as a function of time due to random noise in the environment or hardware imperfection respectively. When these frequencies change, for a typical qubit, it directly results in a random change in the detuning (Δ_f) between the qubit basis and the microwave tone, which results in an unknown change in the phase – this is dephasing. When in the rotating frame, operating this new splitting as the computational basis, if the system is prepared at the point of zero slope on the above graph, small changes in (Δ_f) don't change your actual gap energy, and thus no longer result in frequency-dither. To first order, this system is insensitive to changes in the frequency of either the qubit itself or the microwave signals being sourced.

The problem with operating these systems as qubits is they are difficult to gate due to their small gaps, and as such have never before been operated with universal fast gates. The main difficulties gating such systems are as follows:

1. Unlike the rest-frame qubit transition (ω), the transition energy of the decoupled computational basis (Ω) is small such that typical Rabi frequencies for gate pulses (Ω_D) are on its same order. They are small gap qubits in the rotating frame of the decoupling drive. This leads to higher order counter-rotating terms in the Hamiltonian becoming non-negligible, thus leading to a failure of the Rabi approximation for any applied fast gates.

2. The decoupling drive with amplitude Ω and gate pulses with amplitude Ω_D meant to address the splitting must have identical frequency so that they can be phase-locked at 90 degrees from one another, but the splitting (Ω) has a frequency-noise insensitive “sweet spot” (f_p) detuned from the frequency that the gate-pulse must be applied – with typical gate-pulses this would make transitions difficult to excite.

In the following text, I describe solutions to these challenges. I alleviate the apparent dichotomy between needing high coherence simultaneously with stronger, more numerous, or noisier coupling channels which carry limiting levels of longitudinal

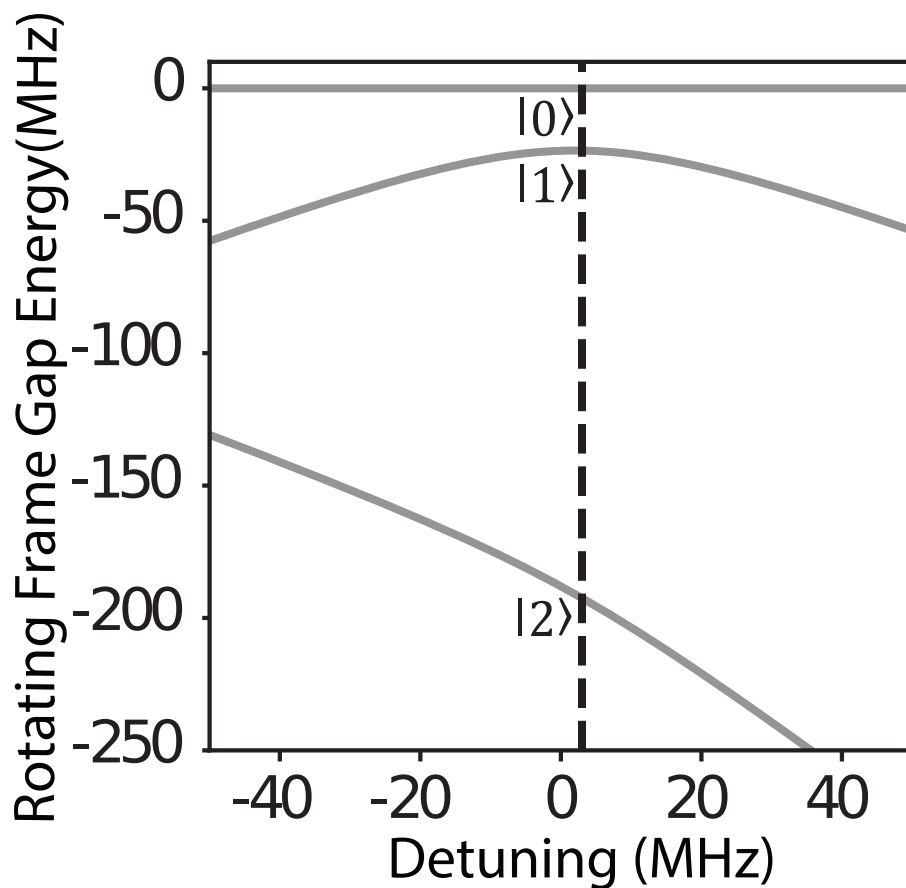


Figure 21: These are the numerically-simulated energy levels of the rotating-frame splitting when driven at a gap (Ω) of about 23 MHz. The frequency-detuning-insensitivity-sweet-spot is the point of minimum gap-energy (Ω) and is detuned from the frequency of the rest-frame-qubit. This detuning Δ_ρ is marked with the vertical dashed line.

noise, by demonstrating fast, universal gates on a system continuously decoupled from coherence-limiting longitudinal noise.

4.2 Mitigating the failure of the Rabi Approximation

Addressing the first challenge, when Rabi-gating a qubit one must select a drive-power and pulse-duration that together creates a π -pulse which rotates the qubit from the ground-state to the first excited-state. Since high-fidelity gates are a typical goal, gates must be fast, else dephasing during the gate reduces its fidelity. Since gates must be fast, the power of the drive-tone being used to gate the qubit must be relatively high in order for the total energy input to be equivalent to a single photon of the transition frequency of the qubit.

In the Hamiltonian of a transmon (and indeed any qubit system that can be driven), the typically negligible counter-rotating terms scale as a function of the Rabi frequency (Ω) – the measure of the drive-tone-amplitude apparent to the qubit. Consequently the Rabi-approximation fails when the qubit’s undriven gap-energy is similar in scale to the Rabi frequency characteristic of the applied drive: the actual gap energy of the system changes strongly as a function of pulse amplitude [44, 43]. This means the qubit gap energy is different when the gate pulse is on, as opposed to when the gate pulse is off. Since gate pulses are continuous operations that need to be ramped on and ramped off, a qubit which is being gated by a pulse of similar scale to its natural gap energy can be expected to chirp around in frequency-space while being gated, which presents a serious problem. If at any point during the pulse, the frequency of the qubit and the frequency of the gate pulse become different, this generates what is known as an off-resonance rotational (ORR) error in the Bloch-sphere trajectory of the qubit [43, 44, 45]. That is, much in the same way as we intentionally generate continuous azimuthal revolutions around the Bloch-sphere by detuning pulses to perform a Ramsey experiment, any detuning present between the qubit and the gate pulse will result in azimuthal rotation. As a reminder, azimuthal rotation about the Bloch sphere is physically identical to dynamic phase evolution as a result of the gate-pulse being a different frequency than the qubit transition. If the qubit frequency chirps during the gate pulse, and the gate pulse isn’t adjusted accordingly to follow it, the resulting accidental azimuthal rotation will cause the

gate to miss. This produces Z errors that accumulate dynamically during the gate, which translate into X/Y errors in every gate. In a traditional Rabi gating scheme, each new gate might have a new pulse-amplitude, so the exact chirp-trajectory of the qubit would have to be compensated for individually with each gate-pulse. Of course, it is untenable to use an arbitrary gate-set for which each of N gates must be individually optimized to high fidelity, so we must find another solution.

Fortunately, the phase of the continuously-decoupled splitting Z-evolves continuously at a rate equal to its gap energy Ω .

$$Z_\phi = e^{\frac{-(i\Omega\sigma_z t)}{2}} \quad (4.1)$$

This continuous Z evolution where σ_z is just the Pauli gate about the Z-axis and t is time, provides us three crucial capabilities: the ability to unwind these erroneous ORRs, the freedom to Z-gate merely by waiting, and consequently as suggested in Fig. 22, the capability to build arbitrary single qubit gates out of only 1 explicit gate-pulse. Given this continuous Z evolution, the only feature missing to achieve a universal single-qubit gate-set is a single $\pi/2$ pulse. With these features together, our problem goes from manually tuning up N pulses to tuning up a single $\pi/2$ pulse.

To be clear, the ORR still occurs, and we still have to fix it. The good news is though, because we can build a universal single qubit gate-set out of a single $\frac{\pi}{2}$ and we can unwind these mistakes by just waiting, it means we have to tune up the pulse amplitude, shape, and the wait-time before and after the pulse only once, and then the problem is solved for our entire gate-set.

4.3 Fast Universal Gates

One gate scheme ideal for generating Z-evolution to unwind ORR errors was recently demonstrated in Ref. [44]. Crucially, this gate scheme has been demonstrated to work on small gap systems for that very reason. This unique gate scheme can produce fast, high amplitude gates that can successfully operate qubits with transitions on the same order as the Rabi frequency characteristic of the pulse amplitude used. Moreover, since the only X or Y gate-pulse this scheme uses is one $\pi/2$ pulse of a single chosen phase, the user need only optimize the ORR correction for a single pulse shape [44].




Figure 22: X or Y axis $\pi/2$ pulses are sufficient when combined with a single continuously variable gate-type (such as continuous phase evolution when the relative phase between the qubit and its drive evolves as a function of time), to provide universal access to the entire space.

The ideal tuning of these pulses for a given system depends on the transition energy (Ω) and anharmonicity (α_ρ) of the splitting. Our gate pulses have an interesting Fourier profile which make them uniquely suited to this particular application:

With some forethought the exact splitting and frequency (Δ_f) can be chosen such that the frequency of the computational transition falls near to one of the primary Fourier-maxima, while the frequency of each the ground-state to second-excited-state transition and the first to second excited state transition each fall on a Fourier-minimum. This minimizes leakage into higher states in the manifold, optimizing our confinement to our computational basis. That is, the maxima of our Fourier transform (FT) are the frequencies at which excitations will be most encouraged, and the minima of our FT are the frequencies at which there will be minimal excitation.

4.4 Designing gates despite frequency constraints

Finally difficulty 2 must be addressed; the frequency detuning of this gate-pulse and the frequency detuning of our decoupling drive-tone must be identical in order to lock their relative phase at 90 degrees. The decoupling drive, however, is constrained to a specific frequency detuning ($\Delta_\rho = f_{01} - f_\rho$) away from the rest-frame qubit gap –

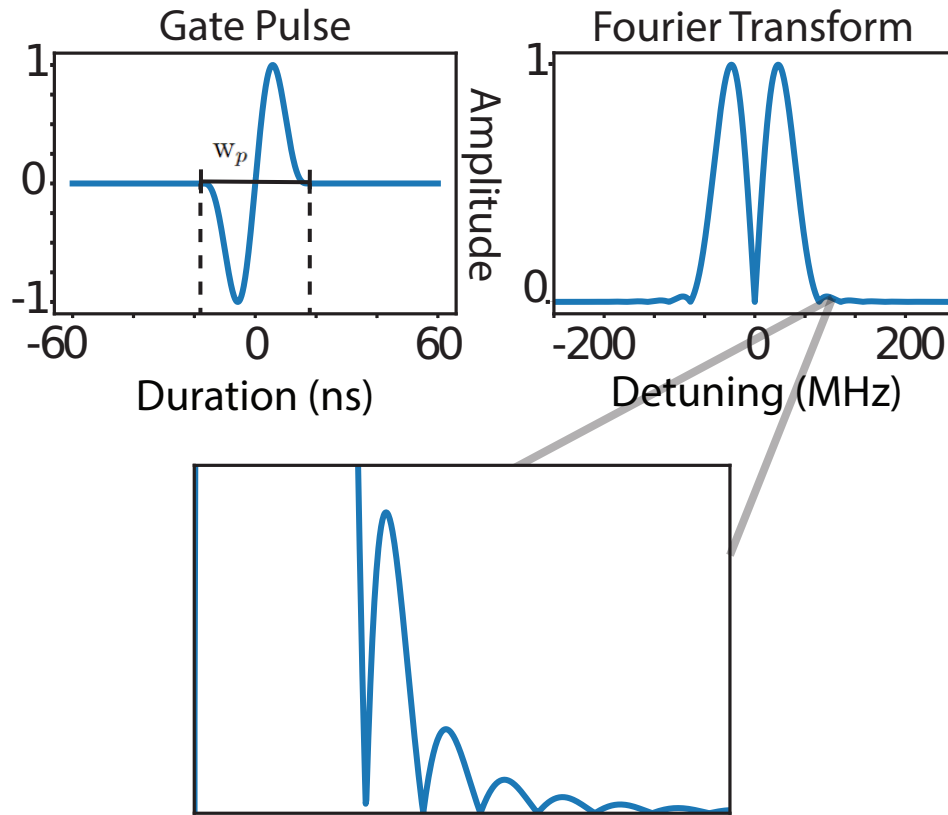


Figure 23: An example of a $\pi/2$ pulse and its Fourier transform (FT); crucially, the FT of this pulse has sinc-like features which present an opportunity to create regions of intentionally low excitation: wherever the FT goes to zero, interaction with modes at the same location in frequency-space are minimized. This gives us the freedom to minimize leakage into the second-excited state. The time domain pulse-width w_p in nanoseconds determines the profile of the FT.

the frequency detuning-insensitive sweet-spot, where f_ρ is the absolute frequency of that sweet-spot. As discussed earlier in this section, this is the detuning at which frequency-noise has little to no impact on the dephasing of the system. This fact constrains the frequency detuning of the gate-pulses as well, so we must design the Fourier profile of our pulse so that we can gate our computational basis with strong confinement as outlined prior – i.e. rather than changing our frequency-detuning to minimize excitation out of the computational basis, we are constrained to a single detuning (the insensitive sweet-spot), and can only change the pulse-width (w_p) to achieve this optimization while also optimizing our excitation of the computational basis – this is because as we change the time-domain width of the gate-pulse, we inversely change its frequency-domain width or the width of the FT. That is, the shorter the pulse in time, the wider its profile in frequency, and *visa versa*.

The frequency at which our gate-pulse Fourier maxima and minima show up is a function of the width of our gate pulse (w_p). The shorter in duration the pulse, the further apart our Fourier-minima are spaced. I choose the gate-pulse-width (w_p) to target the computational transition near a primary Fourier-maxima and the 2nd-excited-state transitions on Fourier-minima to minimize second-excited-state leakage.

To set up the gates, I select the detuning of the decoupling drive to be the frequency-insensitive sweet-spot (Δ_ρ) for a given Ω . That is, we want Δ_f for all gate-pulses and drive tones to be equal to Δ_ρ . As discussed earlier, this maximizes our coherence enhancement by minimizing the influence that small dithers in effective Δ_f have on the actual gap-energy of the system we’re now operating. I then set a gate-pulse-width (w_p) which places the Fourier-spectrum maximum near the splitting transition, while similarly placing Fourier-spectrum minima on each of the 2nd-excited-state transitions. This successfully minimizes all sources of dephasing which would have impacted the rest-frame qubit, while keeping us optimally confined to our rotating frame computational basis. This does restrict us to only certain combinations of transmon-anharmonicities (α) and splitting energies (Ω), which needs to be simulated/calculated ahead of time. The original transmon α and splitting Ω I choose determines the splitting’s anharmonicity we might call α_ρ , which sets the target gate-pulse-width; if we choose the splitting poorly for any given anharmonicity, it will be difficult to line up Fourier-minima with all 2nd-excited state transition frequencies as shown in Fig. 24.

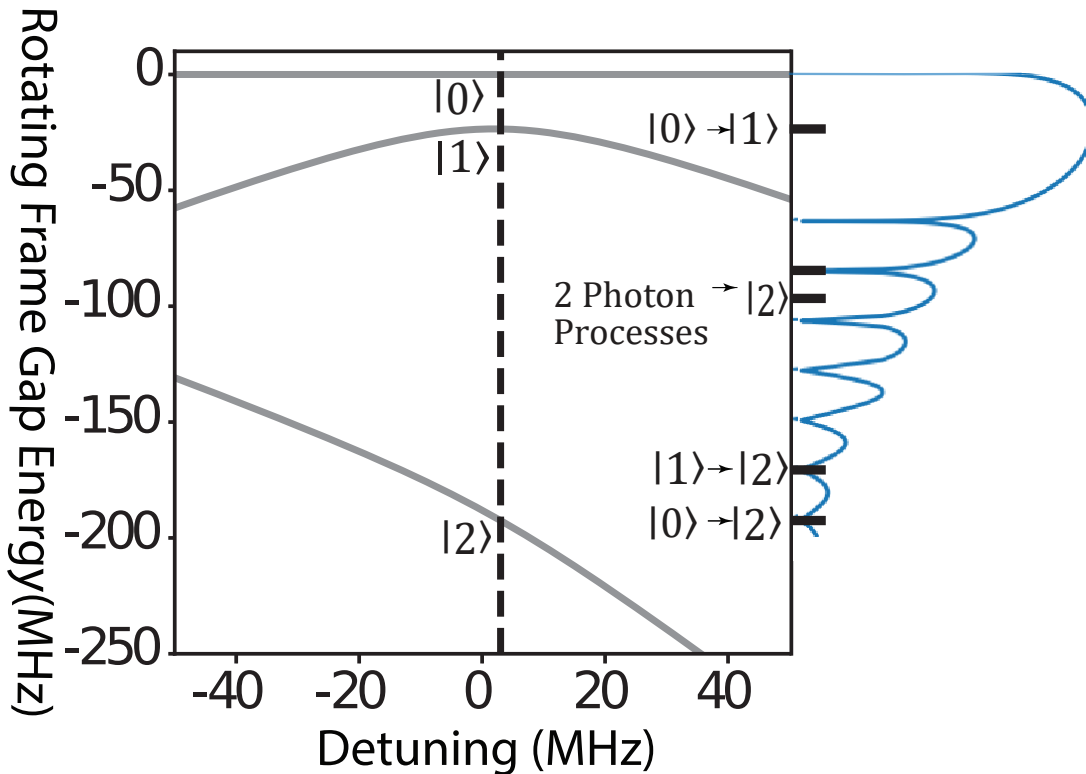


Figure 24: An example of a $\pi/2$ pulse in the frequency-domain overlaid (on the right-axis) atop the level manifold of the continuously decoupled transmon. The maximum of the FT is placed at the frequency of the rotating frame computational basis ($f_{\rho 01}$), which is the continuously decoupled qubit. The frequencies at which 2nd-excited-state transitions are present are to the best of our ability positioned at minima of this FT to optimize computational basis confinement.

When these gates are performed during the decoupling-drive tone, the traces which represent the RF signals being sent to the sample look as shown in Fig. 25.

4.5 State Preparation and Tuneup

Lastly, we need to minimize state-preparation and measurement error (SPAM) by carefully designing our state-preparation. In many cases, such as in the spin-locking projects cited in Chapter 1, the rest-frame device is gated to the target Bloch sphere position with a pulse, then immediately after, the continuous drive is turned on.

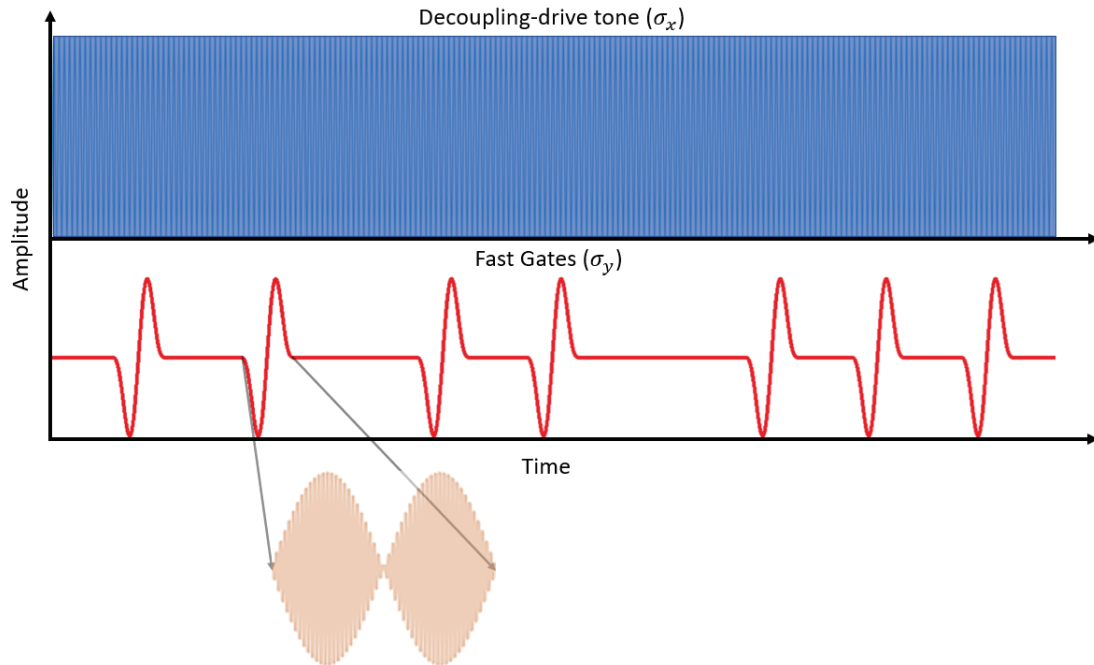


Figure 25: An example of the decoupling-drive tone (shown in blue) while a random selection of fast gates (amplitude envelope in red, tone example in orange) is performed on the rotating frame qubit. The decoupling-drive tone is a continuous tone of constant amplitude with a phase that makes it a rotation about the X-axis of the rest-frame basis, while the fast gates have a chosen phase which would make them a rotation about the orthogonal Y-axis in the rest-frame.

In our case, instead of fast-pulsing the rest-frame device to the target state, we chose a Gaussian-ramp for the decoupling drive-tone at a rate which is adiabatic with respect to the rest-frame qubit transition f_{01} . This preparation slowly drags the Bloch vector to its intended Bloch sphere orientation rather than fast-pulsing it there. I used a $1\mu s$ ramp on/off, but it may be possible to optimize this to a shorter duration or choose a different, less adiabatic preparation that doesn't introduce additional SPAM—this was not a focus of our experiment but might wind up being an important area of investigation if fast preparation is needed.

In order to maintain adiabaticity, I often needed to chirp the frequency of the drive tone such that the amplitude ramp occurs while the tone is deliberately off-resonant with the rest-frame qubit before being chirped to resonance with the splitting. This helps maintain adiabaticity as the rate of change of the effective Rabi frequency is low

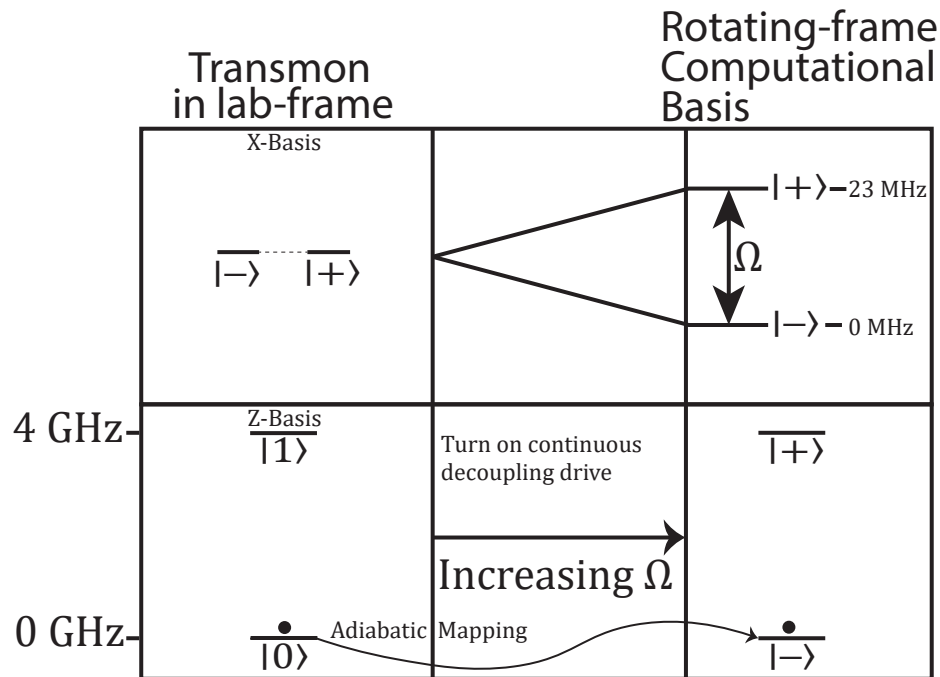


Figure 26: Continuous Decoupling preparation and mapping to and from the rest-frame and the rotating frame. Under typical transmon operation, the qubit is just the ground and first excited states of the transmon in the rest frame. Under continuously decoupled operation, the qubit is now the previously degenerate \pm - X states in the rotating frame, split by the Rabi frequency of the drive along the X axis. Readout is performed in the rest-frame.

during the amplitude ramp thanks to the relative detuning before being brought into resonance. Once we've determined how we'll perform state-preparation, we can try preparing the splitting and begin to probe it in spectroscopy. Typically, the first scan I perform when preparing this technique would be a sort of noise-spectroscopy [40].

There are going to be drive-amplitudes that result in gaps at which the system decays more rapidly than others. My working hypothesis to explain this is just the presence of TLSs near these transitions in frequency-space, or perhaps higher spectral density of noise from other sources which drives decay at these frequencies. Work done

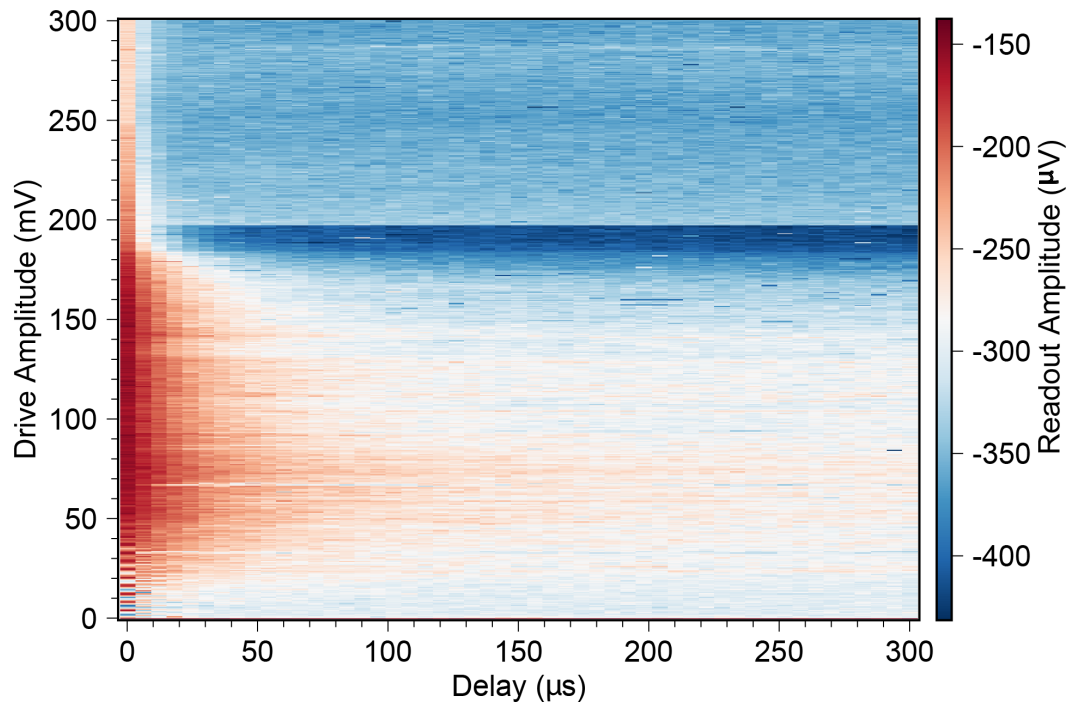


Figure 27: A series of scans typically used to determine $T_{1\rho}$ vs. Continuous-decoupling drive amplitude, searching for regions of drive amplitude which have suitably long $T_{1\rho}$ for coherent operation.

by Fei Yan with the Will Oliver Group in 2013 suggests that noise features exposed by $T_{1\rho}$ -spectroscopy in their case were "sets of coherent TLSs" [40]. Since I'm merely using this scan of $T_{1\rho}$ (the T_1 of the rotating-frame-computational-basis, denoted with subscript ρ) as a function of drive-amplitude to find broad regions of amplitude-space have long enough decay lifetime for coherent operation, it's beyond scope to attempt to dive any deeper and investigate the cause. In my case, this is just an engineering technique to get a reasonable $T_{1\rho}$.

To be entirely clear, this is a scan typically used to measure $T_{1\rho}$ which does not require that we gate the qubit. In the rest-frame computational basis for a typical transmon under nominal conditions, the T_1 of the ground-state is effectively infinite. While there is residual first excited state population due to non-zero effective charge-carrier temperature and sources of noise which carry non-zero energy near to the transition frequency, the rate at which the ground-state spontaneously excites to the first-excited state without intervention is exponentially suppressed, if we've done our

jobs properly and our effective temperatures are low. In the rotating frame, both the ground-state and the first-excited state decay. Since it's a small-gap qubit, they tend to mix to something near to 50/50 state population; as a result, we can measure $T_{1\rho}$ without designing a π -pulse. All we need to do is turn on the drive-tone and vary the duration of that tone before ramping it off and reading-out.

This scan helps us identify regions to avoid, but it's important to remember that these measurements are often slow and convolve time-dependent noise-features in the y-axis, For example, there seems to be an abrupt change that occurs just below 200 mV; this is probably due to an event that occurred at the time that trace was being taken, such as a flux-jump or a quasi-particle generating event e.g. cosmic ray collision, muon interaction, or radioactive decay of some nearby naturally occurring isotope in the fridge. The region above 200 mV on the y-axis that looks like it has a terrible $T_{1\rho}$ might actually be as bad as it looks, or might need to be tuned up and scanned again to reveal the actual $T_{1\rho}$ at the Rabi frequencies that correspond to those drive amplitudes.

I should mention that our target detuning Δ_ρ changes as a function of drive amplitude, so during that scan I need to follow that frequency value as I scan amplitude. To facilitate that, we should simulate this so that we aren't blindly tuning through huge parameter ranges just to find our optimum frequency, and build an intuition for how fast that sweet-spot detuning changes as a function of drive-amplitude.

For this purpose, we built Python code that uses the Qutip package [46] to simulate a transmon of a particular junction energy and anharmonicity under the influence of a decoupling drive-tone (written in conjunction by Dan Campbell and myself).

This code takes the diagonalized Hamiltonian later detailed in Chapter 5, and simulates the eigenstate manifold and the impacts of gate-pulses on that system.

4.5.1 Frequency and Time Domain measurements

Once we know roughly where that sweet-spot should be in the frequency domain, we can find it and check whether or not we are operating at it for any given Ω . To find it, we might try spectroscopy. Unfortunately, doing spectroscopy on a continuously decoupled state is a bit of a problem. In the rest-frame, I can do rapid spectroscopy on a qubit by using high amplitude, short duration pulses with a large Fourier-profile. That is, a 50 ns Gaussian pulse has a Fourier width of 20 MHz. This means when doing

spectroscopy, taking 5 MHz steps in frequency is more than sufficient to guarantee that the qubit response will be visible; looking for the qubit between 3.95 GHz and 4.05 GHz with steps of 5 MHz can be quite fast.

In the case of a continuously decoupled state there are several problems that each make it impossible to do spectroscopy with high amplitude pulses. If the spectroscopy pulse being used isn't very low in amplitude when compared to the decoupling-drive, the scan will be replete with erroneous interference features that make discerning the qubit from some erroneous feature very difficult, and worse, the splitting itself will be completely different while the spectroscopy pulse is on, and as a result won't even appear at the correct frequency—the frequency at which it sits when the spectroscopy pulse is off. These are literally the very problems I avoid by using my chosen gating-scheme, but when performing spectroscopy we are somewhat constrained to using more traditional Gaussian-like pulses.

The solution to this (while still being able to do traditional spectroscopy with a basic Gaussian pulse) is to use very low amplitude, very long-duration pulses. After all, we aren't trying to do a high-fidelity gate or a fast-pulse right now, we're merely trying to see our qubit-response in spectroscopy. The problem with this, of course, is that if the spectroscopy pulse is very long in duration, its Fourier presence is very narrow. A 10 μ s long pulse has a Fourier width of 100 kHz, so scanning the same range of frequency as before now takes a factor of 50 more data points, i.e. a factor of 50 longer time for the scan to complete.

The second type of scan would just be a series of Ramsey measurements at different detuning like those discussed in Chapter 3. For each new detuning, a new $\pi/2$ pulse may need to be tuned up, but this can often still be faster than the spectroscopy method. Here, we would run Ramsey oscillations at different values of detuning, and observe that the Ramsey-frequency changes as a function of detuning, Unlike typical Ramsey scans, where the Ramsey frequency is equal to the detuning away from the qubit frequency, this Ramsey frequency is equal to $(\Omega + (\Delta_f - \Delta_\rho))$, so at whatever chosen frequency the Ramsey frequency is minimized is where Δ_f and Δ_ρ are equal, and this frequency is the sweet-spot.

Whichever method we employ, we have now located our frequency-noise insensitive sweet-spot for a particular drive-amplitude. This is the frequency at which we will set our decoupling-drive-tone.

At this point, I would finish by measuring the frequency of this sweet-spot as a function of drive-amplitude within a range of interest so I can hop around to different gap energies with relative freedom within that range and only need to do very quick, narrow scans to make small corrections to the Δ_f . Now that I know the target sweet-spot, I would rescan $T_{1\rho}$ from above in some range of interest such that each sweep is at its corresponding sweet-spot, as the original scan was probably not done at the exact sweet-spot for each drive-amplitude.

4.5.2 Designing Gates

Now in order to go any further, we need to build gates. Crucially, we want to design fast, high fidelity gates, which again means we need to carefully choose the splitting we use to optimize several features of the system. We need to operate in a region of drive-amplitude that provides high $T_{1\rho}$, we need to operate at the frequency-noise-insensitive sweet-spot characteristic of the splitting at the chosen drive-amplitude to ensure high $T_{2\rho}$, and the gates that we use must also share the same frequency as the decoupling-drive-tone in order to remain phase-locked. Together, these attributes lead to an elaborate optimization problem. At this point, we will have a device-anharmonicity and transition frequency which allows us to perform a Hamiltonian simulation of that device under a drive tone, which produces a rotating frame manifold diagram like the one shown in Fig. 21 in this chapter. Given this manifold diagram, we can calculate the Fourier transform of our gate-pulses as a function of pulse-duration, and choose a duration that gives high excitation at the computational transition with minimal excitation at the 2nd-excited-state transitions. Once we've done this, we now have chosen a pulse duration and pulse detuning.

The first step to taking that information and using it to make a gate-set is to optimize a high-fidelity $\pi/2$ gate, which I'll call the X-pulse. Fidelity is, colloquially, how correct any preparation is. Mathematically, it can be written as follows [47, 48, 49]:

$$F = |\langle \psi_{actual} | \psi_{target} \rangle|^2 \quad (4.2)$$

Where ψ_{target} is the target state preparation and ψ_{actual} is the state in-fact prepared by the manipulations performed.

In order to measure this, we need to perform a standard measurement called randomized benchmarking. This is the performance of a statistically stochastic distribution of combinations of transformations on the Bloch sphere picked from the single-qubit Clifford set followed by a single recovery gate meant to undo everything performed by the sequence prior. The single-qubit Clifford set is the set of all gates that transform the Bloch vector from any one of the 6 cardinal points on the surface of the Bloch sphere to any of those 6 points. How well this totally pseudorandom selection of gate-sequences can be repeatably performed is the measure of how good any given single qubit is at single-qubit operations.

The result of this experiment, combined with the dimensionality of Clifford gates performed in standard randomized benchmarking provides the following equation for the measurement of the fidelity of single-qubit Clifford gates [47, 48, 49]:

$$F_C = 1 - \frac{1 - \left(\frac{1}{e}\right)^{\left(\frac{1}{N}\right)}}{2} \quad (4.3)$$

Where N is the number of Clifford gates performed before the proportion of measurements which result in the desired target after the recovery gate is equal to $\frac{1}{e}$.

This is a conveniently quick calculation for the per-gate fidelity which falls out of the recovery probability equation

$$P_{Recovery} = P_C^N, \quad (4.4)$$

where P_C is the probability that any single Clifford gate is correct and $P_{Recovery}$ is the probability that the recovery gate at the end of a sequence of N number of single Clifford gates will produce the expected result.

This is the standardized experimental technique for measuring any qubit's ability to perform gates, and the measurement we are leading up to. In order to perform this, we need to make sure all of our individual operations which will make up our Clifford gate-set are tuned-up properly.

By now, we've already tuned up a low to medium fidelity $\pi/2$ gate, but that isn't good enough. Using typical amplitude scans of a single pulse, the best fidelity I might hope to achieve is perhaps 95 percent or so. This is because the difference between the

readout of a full 100 percent state population in the first-excited state and 95 percent in the first-excited state and 5 percent in the ground-state is almost imperceptible – that is, very low SNR – so taking better data to aid in tuneup is crucial. The solution is to prepare scans with more pulses. That is, a pulse with 95 percent fidelity, if used in a 10-pulse sequence, will have a total sequence fidelity of just under 60 percent, which produces a much more pronounced effect on the output values produced by readout, and thus increases the precision of our tuneup. In order to prepare such a sequence, we would ordinarily need to optimize three dependent variables. One of the three, the duration of the X-pulse, is already set for us by our need to optimize computational basis confinement as discussed in Section 4.4. The remaining two are the amplitude of the X-pulse and the spacing required between X-pulses which is the requisite Z-evolution to make each of these pulses actually a $\pi/2$ X-gate and not some other phase. To be clear, it's not actually important that the chosen gate be a $\pi/2$ X gate, but it must be a single chosen phase; my choice of X is somewhat arbitrary.

Remember, now we are building a gate-scheme which will have a single explicit $\pi/2$ gate, and a continuous Z-evolution for all time, so any wait-time including the set duration of the X-gate itself will also be a Z-gate of some arbitrary amount. The best way to tune this up is to first take the previously established low-fidelity $\pi/2$ gate and prepare a sequence of 2 of them with some delay in between to make a π gate of the same phase. The first parameter scan will be the delay between the pulses. Start with a sequence where the second pulse begins immediately after the first pulse ends, and scan nanosecond by nanosecond of increased delay. The delay which provides the highest readout amplitude is closest to a perfect π sequence. After this, I sweep the amplitude with a constant delay, and pick the best preparation, repeat the cycle two or three times.

Once this is done, I now have a low-fidelity π -sequence, for which I have selected roughly what the amplitude of any given $\pi/2$ X-gate should be, and roughly what the delay time between each pulse should be. Next, I extend the pulse-sequence.

Now I prepare a set of sequences. My personal favorite method is to prepare 4 sequences, each with 11, 12, 13, and 14 X-pulses respectively, and scan amplitude. This should produce a scan that looks like Fig. 28.

Each of these traces is a different number of pulses, so each one has a different target. 10 X pulses should prepare the first-excited state, 11 should prepare a 50/50

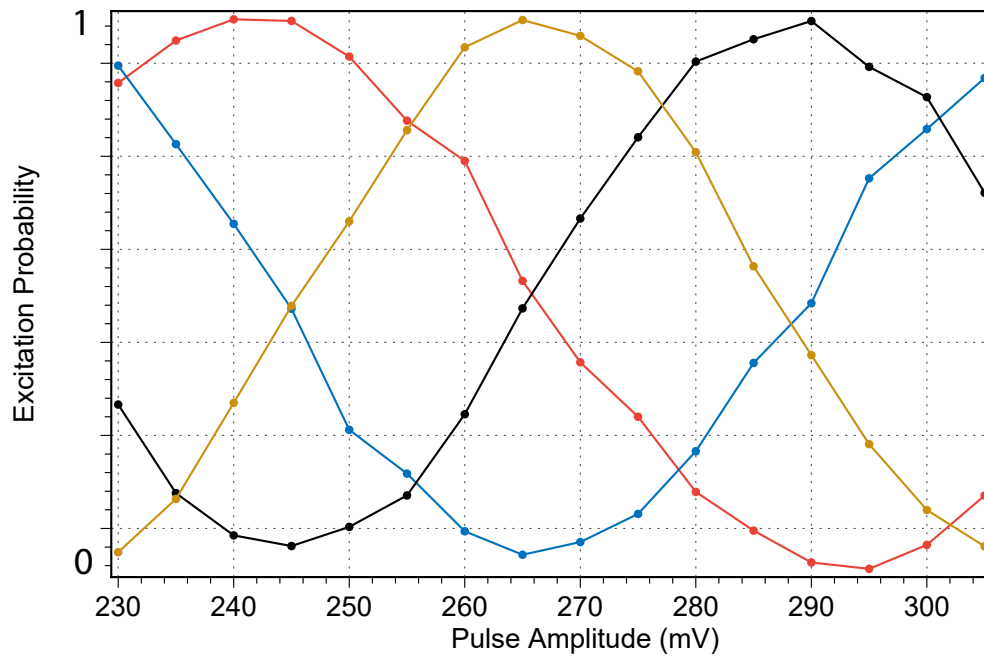


Figure 28: A series of sequences each with a different number of $\pi/2$ X-gates, allowing us to tune-up a high-fidelity $\pi/2$ X-gate.

superposition, 12 should prepare the ground-state, and 13 should prepare a 50/50 superposition again. Using this information, you can tell to relatively high precision what amplitude is the correct one to choose.

Then, do the same type of scan for a small range of delay at the chosen amplitude. Repeat the cycle perhaps twice.

Next, scale up to 20 or 30-something pulses. Rinse and repeat with higher and higher numbers of pulses. Note that at some point, the total dephasing that will occur over the entire sequence will make preparing tuneup plots that can be read and deciphered to high-precision more difficult, so there is a maximum number of pulses that is appropriate for this tuneup. For my purposes, I typically didn't go higher than 100 pulses or so. If any one sequence can be tuned-up reliably to a fidelity of 95 percent as I suggested earlier, my theoretical preparation limit at 100 pulses would be about 99.95 percent individual X-gate fidelity. For higher fidelity tuned to higher precision, you would need more pulses, but if your actual X-pulse fidelity is less than this, due merely to the proportion of the qubit coherence occupied by a single pulse-duration, more pulses will result in messy data plots that don't have clear or reliable

optima.

Once this process is complete, the result is a high fidelity X-gate with a given amplitude, duration, and delay between pulses.

4.5.3 Tuning up time-evolution Z-gates

The next step is to tune-up the optimal Z-gate wait-time for any given target Z-rotation. The center-to-center wait time between X-pulses that we already tuned up and the actual 2π Z-gate duration are different. That is to say, the Z-evolution is different as a function of time during X-gates, and so having tuned up the correct X-gate-associated delay is not the same as having tuned up the correct 2π Z-gate. Typically for this, the best strategy is merely to execute a Ramsey measurement with a moderate number of cycles. For a typical Ramsey measurement, we intentionally detune our gates to generate azimuthal rotations that produce the characteristic Ramsey oscillations. In this case, since our system is continuously Z-evolving as a function of time, we don't need to do that. A single $\pi/2$ of our previously chosen parameters, followed by a wait time, followed by another single $\pi/2$ as usual will produce Ramsey oscillation, but these will be at the frequency of our computational basis' transition frequency exactly – assuming the preparation is correct. This Ramsey sequence can, as it happens, also be thought of as a single X pulse followed by a sequence of arbitrary Z gates followed by a single X pulse. As before, the longer this sequence, the higher our precision in tuning up the fidelity, to a point. The more Ramsey oscillations the better, until some maximum at which point dephasing or slow frequency drifting can cause your tuneup process to become worse as a function of increasing Z-duration rather than better. I typically stopped at around 10 cycles, but I have not determined whether this is in fact the optimal number of cycles for this tuneup. When the Ramsey frequency has been determined to relatively high precision, we now know the 2π Z-gate duration.

All of these parameters can drift as a function of time; this is usually the result of slow frequency drift of any of the operating bases, whether the rest-frame qubit device or the microwave generator, or due to noise-sources in the qubit-environment changing as a function of time. So, retuning may be required. If retuning is required, it's important to retune all of the above parameters and not just one of them. If the Z-gate duration changes, rest-assured that the X-gate-delay must also change, and

may perhaps mean the X-gate-amplitude requires tweaking as well. Alternatively, it may be sufficient to retune the rest-frame-qubit-frequency instead, and then recheck the tuning of the rotating-basis gates, and the frequency-domain-position relative to the target detuning-insensitive-sweet-spot.

Once I have performed all of these tuneups, I have successfully prepared fast, high-fidelity universal single qubit gates for a continuously decoupled rotating-frame qubit.

4.6 Data and Results

After having tuned-up the measurement circuitry as described in the previous subsections, what follows here are fits to my data of the T_2 , $T_{2\rho}$, and randomized benchmarking shown in Figures 29 and 30.

The device on which these measurements were performed is a symmetric flux-tunable transmon with flux-insensitive-upper-sweet-spot (USS) at 4.64 GHz, and a readout resonator at 5.86 GHz with a relatively nominal coupling g (about 100 MHz). This frequency-proximity and coupling makes for an overcoupled, Purcell limited qubit with a T_1 of roughly 20 us, and a T_2 of about 40 us as measured with decay and Hahn-Echo measurements not shown here. In order to reduce the Purcell limit, we flux-tune the device to gap-energy of 2.9 GHz. Here, the T_1 has improved to a much more respectable 69.25 us (not shown), but the device has suffered vastly in its coherence with a $T_{2Ramsey}$ of 1.93 us (not shown) and a T_{2Echo} of 6.51 us (shown in the blue curve in Fig. 29, limited by the flux-frequency noise being 2 GHz away from the relevant flux-noise insensitive sweet-spot, thus having high flux-frequency slope [50]).

In order to remedy this problem, I continuously decouple the qubit by inputting a continuous drive with an amplitude corresponding to an effective Rabi frequency Ω of about 23 MHz with a phase of +X on the Bloch sphere. At this point, I have split a 23 MHz non-degeneracy between the +x and -x states of the rest-frame qubit. This rotating-frame splitting has a decay time $T_{1\rho}$ of 45 us – notably, slightly worse than the rest-frame device – and dephasing times $T_{2\rho Ramsey}=39$ us and $T_{2\rho Echo}=61$ us (shown in the red curve of Fig. 29; an order of magnitude better than the rest-frame

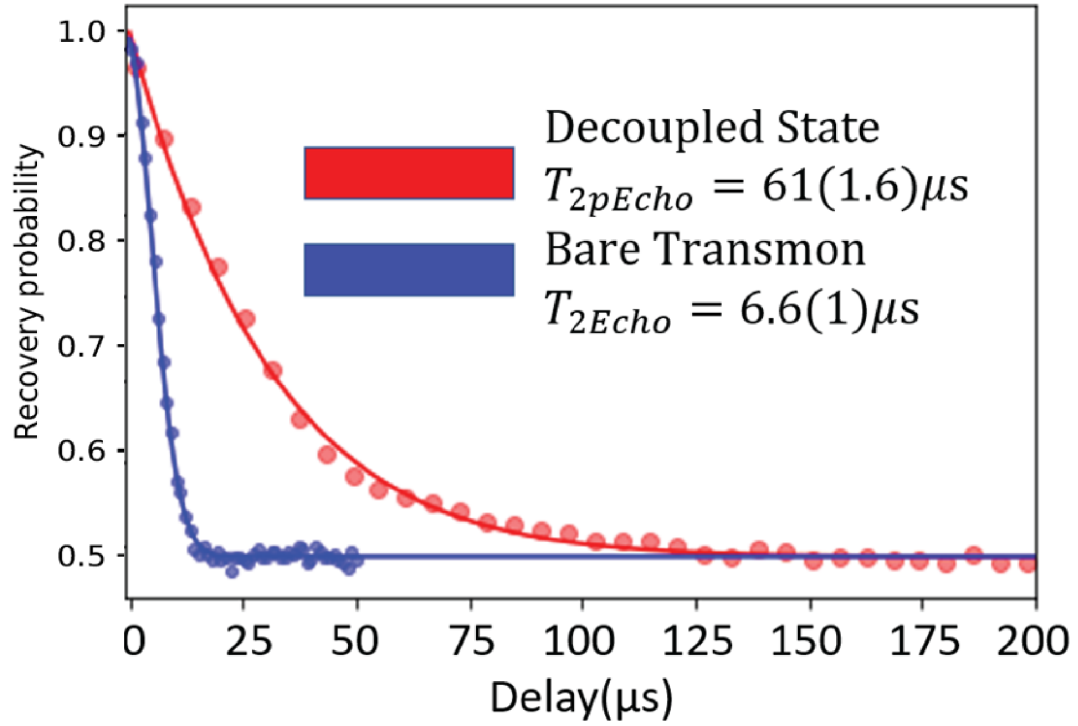


Figure 29: Measurements of coherence performed with Hahn-Echo experiments of each the rest-frame transmon operated traditionally (i.e. bare transmon), and the continuously-decoupled state in the rotating-frame gated with our novel technique, along with exponential fit-curves shown in the solid blue and red lines, respectively.

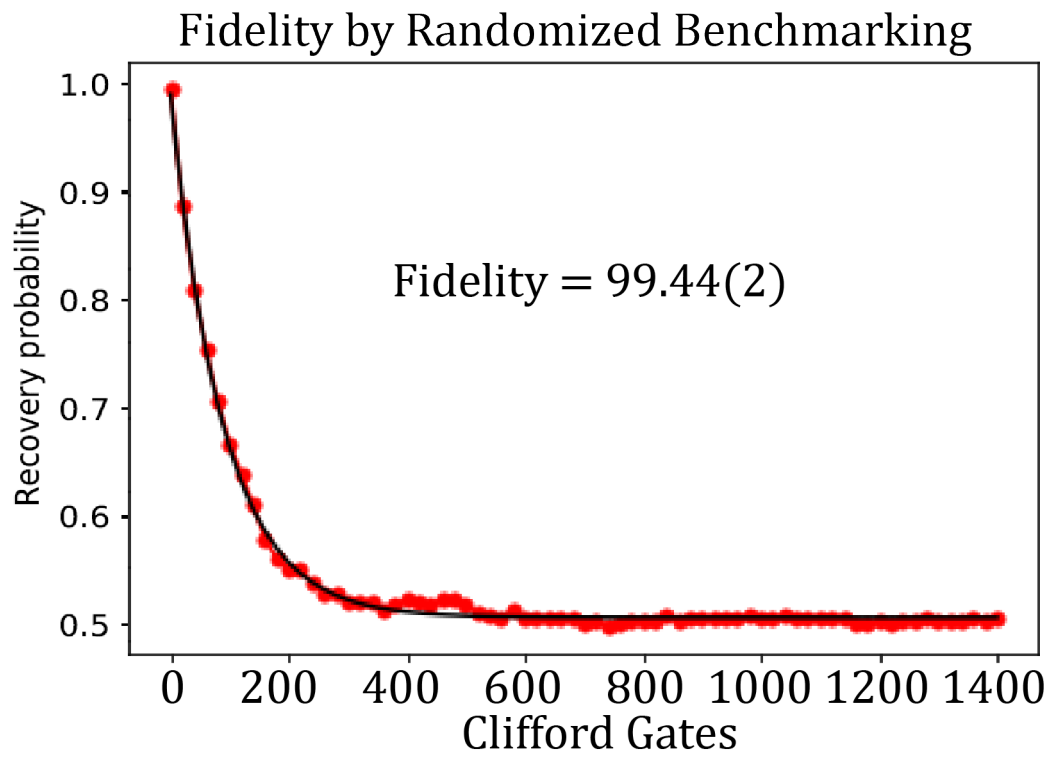


Figure 30: Measurement of Fidelity of Clifford gates measured via randomized benchmarking, along with exponential fit-curve shown.

qubit operated conventionally at the same position in flux-space. This enhancement is shown directly in Figure 29.

Just for clarity, I want to tease apart the distinction between the typical flux-noise insensitive sweet-spot and the detuning-insensitive sweet-spot. As discussed in Chapter 2, in the case of a flux-tunable transmon, the frequency of the device can be tuned as a function of the magnetic flux threaded through the SQUID-loop in which the transmon's Josephson junctions reside. When that magnetic flux is tuned, so too is the frequency, between some maximum frequency and some minimum frequency. At the top of the tuning range, it reaches its maximum frequency and turns around. This curvature leads to a position on the flux vs frequency plot where the slope is zero, which provides flux-noise insensitivity when tuned to that flux. An example of this is shown in Fig. 31.

That is, small changes in the magnetic flux don't change the frequency of the qubit at this position if they are sufficiently small, because the slope is zero. In the case of the rotating frame-splitting's detuning-insensitivity sweet-spot, we are now no longer speaking of flux and magnetic fields, but rather the detuning between the rest-frame qubit basis and the gate-pulse's basis as shown in 24. In this new system, the transmon itself is still at what would be a flux-sensitive position in the rest-frame – i.e. the flux vs frequency slope is explicitly not zero were the transmon operated at that flux under nominal conditions – in fact it's quite large as shown at the orange marker in Fig. 31. In operating the device with this continuous decoupling technique, we make the new splitting insensitive to the frequency dither produced by that flux-noise sensitivity by generating that point of zero slope between the splitting-gap frequency and the gate-pulse detuning, which is shown in the top of the curvature between the two first energy eigenstates shown in Fig. 24. We are still in a flux-noise-sensitive regime, but now we're insensitive to the frequency dither that would ordinarily limit our coherence, and so our $T_{2\rho Ramsey}$ sees enhancement of a factor of 20 over our rest-frame T_2 .

When operated at the flux-insensitive sweet-spot, the qubit coherence is unharmed by small changes in magnetic flux, but still nominally sensitive to small changes in frequency from other sources. When operated at the frequency-noise insensitive sweet-spot, the rotating-frame qubit is unharmed by small changes in the frequency of the rest-frame qubit or the gate-pulses used to manipulate it irrespective of the source of




Figure 31: Typically, symmetric flux tunable transmons are operated at the blue points, where the frequency is maximum, called the Upper Sweet Spot (USS). In the case of my experiments in continuous decoupling, the device was prepared at the spot denoted by the orange dot, at a point of high-slope, and thus, low coherence.

that dither.

The way this actually looks when understood in terms of flux is as shown in Fig. 32, where I have re-parameterized the rotating-frame Hamiltonian in terms of flux rather than detuning. Here we see the new splitting in the midst of the flux tuning, and the corresponding point of zero-slope.

After establishing this new basis and confirming the coherence enhancement, I then tuned up X and Z gates, and measured our Clifford-gate fidelity via randomized benchmarking as shown in Figure 30. With average Clifford gate-durations of about 125 ns, I measured an average Clifford gate fidelity of 99.44 percent limited by a combination of dephasing and coherent error from second-excite-state leakage and slow frequency drift.

In summary, I took a flux-tunable device that was limited by decay by frequency-proximity to its readout resonator, traded precious coherence for additional decay-lifetime, decoupled it from noise to enhance its coherence even above its USS maximum, and finally demonstrated universal high-fidelity single-qubit gates never before performed on a continuously decoupled system.

I want to note that separately, on a fixed-frequency, high-coherence device, I verified the same technique works. While continuously decoupled (with nominally identical parameters between the two experiments performed on each device), this device had a similar $T_{2\rho Echo} = 54$ μs . On this device, I was able to achieve an average Clifford gate fidelity of 99.85 percent. This implies that with some intelligent correction of the slow frequency drift similar fidelity can be reached on the flux-tunable device – that is, yet higher fidelities can be achieved on systems which require this kind of coherence enhancement for high-coherence operation.

For a rough estimate of the error budget on the key result, I presume that the effective decoherence time during randomized benchmarking is about 50 μs , about halfway or so between the Ramsey and Echo measurements, as randomized benchmarking tends to unwind some dephasing in the same way that the Hahn-Echo measurement does, but to a lesser degree. Given a Clifford-gate duration of 125 ns, the best possible Clifford-gate fidelity we could achieve would be 99.86 percent with 100 percent of our infidelity contributed by decoherence alone. This is roughly 25 percent of the measured infidelity. Of the remaining infidelity, the slow DC frequency drift would be responsible for the difference between 99.85 percent (the fidelity measured

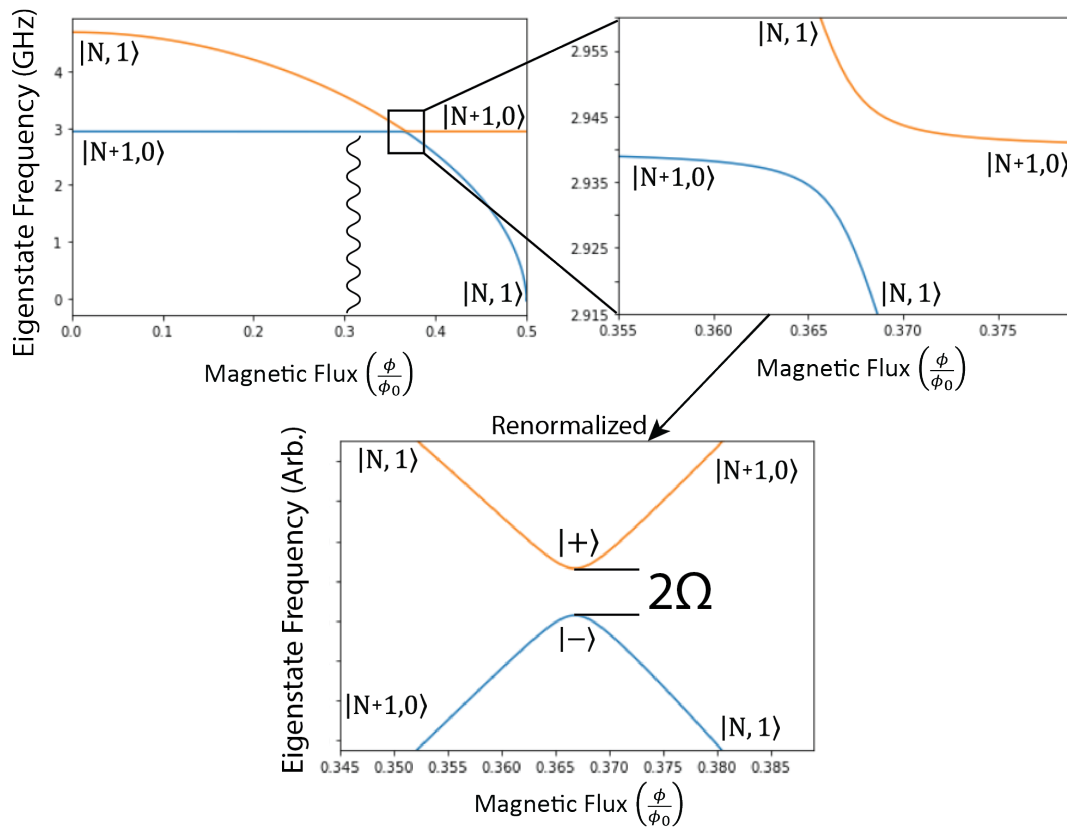


Figure 32: Here, the rotating frame Hamiltonian has been reparameterized as a function of magnetic flux. For proper bookkeeping I am working in the number-state basis N for the drive field and the rotating frame energy eigenbasis for the qubit states. The top left plot shows the full flux-tuning from the USS in Frequency vs Flux to zero frequency. Due to the decoupling-drive, the eigenstates would cross at the chosen frequency of the decoupling-drive tone, and it is at this point where we see our avoided-level crossing representative of the rotating-frame splitting which is our new qubit in the top right plot. The bottom plot is merely the top right plot renormalized to show clearly the point of closest approach, the frequency-insensitive sweet-spot at which we operate the system, such that our new computational basis is made up of the two states shown, $|+\rangle$ and $|-\rangle$.

on the fixed frequency device) and 99.44 percent, which makes up about 73.2 percent of the measured infidelity. The remainder would likely be the result of leakage, making up about 1.8 percent of the measured infidelity. Here I am tacitly assuming that both the dephasing contribution and the leakage contribution is about the same between the fixed-frequency and frequency-tunable devices, as the two devices' anharmonicities and Clifford-gate duration to T_2 ratio were about the same, and they were both operated with nearly the same decoupling-drive amplitude. The portion of the error caused by slow drift can be mitigated by interleaving a scan which re-measures and sets the target frequency in between each set of scans performed. The portion of the error caused by dephasing can be mitigated by increasing gate-speed. The portion of the error caused by leakage might be further optimized to lower values by modifying the shape or detuning of the pulses, or by choosing a better combination of transmon anharmonicity, decoupling-drive amplitude and gate-pulse width.

If instead of using continuous decoupling, I measured the fidelity of Clifford gates performed on the rest-frame device at the same position in parameter-space, an estimate of the fidelity follows.

The coherence time would be roughly $3 \mu s$, between the Ramsey and Hahn-echo measured coherence, and a standard Clifford gate duration for standard qubit operation might be about 40 ns. This would mean that if the tuneup was ideal and all infidelity was only caused by dephasing, I would expect about 75 gates before decaying to a recovery probability of $\frac{1}{e}$, which corresponds to a fidelity of about 99.3 percent.

4.6.1 Paths for improvement

In order to shorten gate-duration and boost fidelity, one method could be to increase the decoupled splitting transition energy by increasing our drive-tone amplitude. This pushes the Z-evolution rate higher and allows for faster average Clifford gates. This would require a device with higher absolute anharmonicity α ; when I tried to push to higher decoupling-drive-amplitudes on these devices, I found the system limited by leakage into the second-excited state and other general misbehaviors that seem to occur where Ω is approaching α in magnitude.

Of course, once we progress to faster gates, implementing Derivative Reduction by Adiabatic Gates (DRAG) [51] on this system would become progressively more important despite the leakage being mitigated to first-order by the FT of the pulses we use. That is, eventually, despite our leakage being largely limited by the fact that we've strategically placed those FT minima on our second-excited-state transitions, if our gate-pulses become sufficiently short and thus our pulse-amplitudes sufficiently large, the relative positions of these minima and maxima will chirp in frequency during the pulse. DRAG is often implemented by pulsing the frequency of the drive-tone to follow the predicted chirp of the qubit-basis resulting from the gate-pulse. If pushing to sufficiently short, thus sufficiently strong pulses, this would become important even in our case.

The measurements described above were not dynamically corrected for long-timescale frequency drift, and so average fidelity could yet be further improved by implementing an interleaved scan which measures the frequency of the system and re-tunes pulse spacing to compensate for that long-timescale drift. Additionally, modifications on the gate-pulses (such as multiple pulses per $\pi/2$ gate) may help to optimize pulse-amplitude and thus confinement.

4.6.2 Conclusion

I have demonstrated fast, universal, high-fidelity single qubit gates on a continuously decoupled splitting prepared on a low-coherence flux-noise-sensitive transmon. The implication is that any qubit in a coherence-limiting environment primarily under the effects of longitudinal-noise might be better operated with a technique like this. While there are huge areas of parameter space in device-modality as well as operation parameters that need to be explored, this technique shows promise as a boon for any application in which qubit coherence is limited by longitudinal-noise in the qubit environment. This technique can broaden operational ranges of flux-tunable transmons from just regions around their sweet-spots to almost their entire frequency range; may make useful qubit species that have since been passed-over for being low coherence, such as the Cooper pair box, and potentially allow high-coherence operation of qubits within high order quantum networks otherwise replete with coherence-limiting noise.

Chapter 5

Hamiltonian Design

The device I used as the test-platform for the work showcased in this thesis was a transmon, which when continuously decoupled, has the Hamiltonian in Eq. 5.1. Here, ω_{10} is the transition energy from the ground-state to the first-excited state, ω_{20} is the transition from the ground-state to the second-excited state, ω_d is the continuous decoupling drive-tone-frequency, Ω is the continuous decoupling drive-tone-amplitude and α is the anharmonicity of the transmon, which while not present explicitly below, $\alpha = \frac{(\omega_{20}-2\omega_{10})}{2\pi}$.

$$H = \hbar \begin{pmatrix} 0 & 2\Omega \cos(\omega_d t) & 0 \\ 2\Omega \cos(\omega_d t) & \omega_{10} & 2\sqrt{2}\Omega \cos(\omega_d t) \\ 0 & 2\sqrt{2}\Omega \cos(\omega_d t) & \omega_{20} \end{pmatrix} \quad (5.1)$$

After performing the rotating-frame transformation as follows:

$$i\hbar\delta_t\Psi = H\Psi \quad (5.2)$$

$$\Psi = U\Psi' \quad (5.3)$$

$$H = UH'U^\dagger \quad (5.4)$$

$$U^\dagger i\hbar\partial_t(U\Psi') = U^\dagger UH'U^\dagger U\Psi' = H'\Psi' \quad (5.5)$$

$$i\hbar\partial_t\Psi' = (H' - i\hbar U^\dagger\partial_t U)\Psi' \quad (5.6)$$

$$U = \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{i\omega_d t} & 0 \\ 0 & 0 & e^{i2\omega_d t} \end{pmatrix}, \quad U^\dagger = \begin{pmatrix} 1 & 0 & 0 \\ 0 & e^{-i\omega_d t} & 0 \\ 0 & 0 & e^{-i2\omega_d t} \end{pmatrix} \quad (5.7)$$

we obtain Eq. 5.8.

$$H_{rot} = \hbar \begin{pmatrix} 0 & \Omega(e^{i2\omega_d t} + 1) & 0 \\ \Omega(e^{i2\omega_d t} + 1) & \omega_{10} - \omega_d & \sqrt{2}\Omega(e^{-i2\omega_d t} + 1) \\ 0 & \sqrt{2}\Omega(e^{-i2\omega_d t} + 1) & \omega_{20} - 2\omega_d \end{pmatrix} \quad (5.8)$$

When prepared resonantly, the oscillating terms in the off-diagonals are fast-oscillating functions that average to zero on the timescales set by the eigenbasis of the Hamiltonian in rotating frame, so we can neglect them, performing the rotating-wave approximation to arrive at Eq. 5.9.

$$H_{rot} \approx \hbar \begin{pmatrix} 0 & \Omega & 0 \\ \Omega & \omega_{10} - \omega_d & \sqrt{2}\Omega \\ 0 & \sqrt{2}\Omega & \omega_{20} - 2\omega_d \end{pmatrix} \quad (5.9)$$

We diagonalize this, and using Qutip, make figure 21 which shows the first three lowest energy eigenstates of the continuously decoupled splitting, in the rotating-frame basis, the first two lowest of which make our computational basis, and the 3rd is our primary source of leakage related error that we must carefully avoid.

It is this simulation that we use to choose our time-domain design for our qubit gates, that is, our pulse durations, at least roughly, before doing direct measurements on the frequency-locations of these transitions and tuning up to the real physical system.

Chapter 6

Departing Remarks and Future Work

In essence, I have demonstrated that a qubit operated under the continuous decoupling technique can be universally gated with high-fidelity single-qubit gates, and thus that such a system can be robustly operated as a qubit. Consequently, the quantum-information community can take full advantage of the coherence-enhancing effects of this technique, and retain full operational capability. This technique can expand operational regimes of qubits, allowing us to operate flux-tunable qubits away from their flux-insensitive sweet-spots, or operate high-pure-dephasing devices as high-coherence qubits.

Although my work was performed on the superconducting qubit platform, this work should extend to any and all qubit modalities which can be driven and operated in the way I described. I suspect, indeed, that this method should apply to all deterministic qubit technologies that can be driven.

I imagine that the ideal applications will be modalities that have much higher decay-lifetimes than coherence lifetimes which have their overall efficacy dominantly limited by pure dephasing.

Interestingly, within superconducting qubits, transmons are perhaps a poor example of this. State-of-the-art transmons often have T_1 limited coherence, with a very long coherence time and very low pure dephasing. That is, the very best single, isolated transmons have coherence limited almost entirely by the system's decay-lifetime, which is defined by transverse noise rather than longitudinal noise [2].

Poorer charge qubits such as CPBs are indeed overwhelmingly limited by longitudinal noise [15], and typically have much longer decay lifetimes than coherence lifetimes [52]. As well, other superconducting qubits might have higher anharmonicities and higher decay lifetimes, both of which may make other superconducting varieties better showcases of this technique.

Similarly, many qubit modalities have quite low coherence lifetimes compared to their decay lifetimes; I suspect these modalities will benefit most from a technique like this. Personally I am interested in investigating the application of this technique on capacitively-shunted-flux-qubits (CSFQs) [27], fluxonium qubits [53], CPBs, and other devices with high anharmonicities and potentially lower T_2 than their T_1 limit, or constrained operational ranges akin to the flux-tunable transmon which can typically only perform high-coherence, high-fidelity operation at the USS.

Beyond this, I suspect the next crucial step would be to investigate the impacts of this technique on multi-qubit-gates, obviously starting with 2-qubit-gates. I'd be interested in answers to questions like: how do SWAP operations on continuously decoupled qubits work? Are there parasitics that exist in these systems that don't exist in other systems, or do certain parasitics typically present disappear? What happens if you try a 2-qubit-gate between one qubit in the rotating frame and one in the rest frame? 2 qubits in the rotating frame? Is ZZ-coupling in these systems any different? In the end, if this is meant to be a technique we leverage to improve the scalability problem, we need to at least understand how to operate multi-qubit-gates with it, so I suspect this is the next important path of investigation.

Bibliography

- [1] Morten Kjaergaard et al. “Superconducting Qubits: Current State of Play”. In: *Annual Review of Condensed Matter Physics* 11.1 (2020), pp. 369–395. URL: <https://doi.org/10.1146/annurev-conmatphys-031119-050605>.
- [2] P. Krantz et al. “A quantum engineer’s guide to superconducting qubits”. In: *Applied Physics Reviews* 6.2 (June 2019). Publisher: American Institute of Physics, p. 021318. DOI: [10.1063/1.5089550](https://doi.org/10.1063/1.5089550). URL: <https://aip.scitation.org/doi/10.1063/1.5089550>.
- [3] David Isaac Schuster. “Circuit Quantum Electrodynamics”. en. PhD thesis. 2007.
- [4] A. A. Clerk et al. “Hybrid quantum systems with circuit quantum electrodynamics”. en. In: *Nature Physics* 16.3 (Mar. 2020). Number: 3 Publisher: Nature Publishing Group, pp. 257–267. ISSN: 1745-2481. DOI: [10.1038/s41567-020-0797-9](https://doi.org/10.1038/s41567-020-0797-9). URL: <https://www.nature.com/articles/s41567-020-0797-9>.
- [5] Ze-Liang Xiang et al. “Hybrid quantum circuits: Superconducting circuits interacting with other quantum systems”. In: *Reviews of Modern Physics* 85.2 (Apr. 2013). Publisher: American Physical Society, pp. 623–653. DOI: [10.1103/RevModPhys.85.623](https://doi.org/10.1103/RevModPhys.85.623). URL: <https://link.aps.org/doi/10.1103/RevModPhys.85.623>.
- [6] Jonas Bylander et al. “Dynamical decoupling and noise spectroscopy with a superconducting flux qubit”. In: *Nature Physics* 7.7 (July 2011). arXiv:1101.4707 [cond-mat, physics:quant-ph], pp. 565–570. ISSN: 1745-2473, 1745-2481. DOI: [10.1038/nphys1994](https://doi.org/10.1038/nphys1994). URL: <http://arxiv.org/abs/1101.4707>.

- [7] Leonid V. Abdurakhimov et al. “Driven-state relaxation of a coupled qubit-defect system in spin-locking measurements”. In: *Physical Review B* 102.10 (Sept. 2020). arXiv: 2006.05820, p. 100502. ISSN: 2469-9950, 2469-9969. DOI: [10.1103/PhysRevB.102.100502](https://doi.org/10.1103/PhysRevB.102.100502). URL: <http://arxiv.org/abs/2006.05820>.
- [8] S. Gustavsson et al. “Driven Dynamics and Rotary Echo of a Qubit Tunably Coupled to a Harmonic Oscillator”. In: *Physical Review Letters* 108.17 (Apr. 2012). Publisher: American Physical Society, p. 170503. DOI: [10.1103/PhysRevLett.108.170503](https://doi.org/10.1103/PhysRevLett.108.170503). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.108.170503>.
- [9] W. K. Wootters and W. H. Zurek. “A single quantum cannot be cloned”. en. In: *Nature* 299.5886 (Oct. 1982). Number: 5886 Publisher: Nature Publishing Group, pp. 802–803. ISSN: 1476-4687. DOI: [10.1038/299802a0](https://doi.org/10.1038/299802a0). URL: <https://www.nature.com/articles/299802a0>.
- [10] Anatoly Yu. Smirnov. “Decoherence and relaxation of a quantum bit in the presence of Rabi oscillations”. In: *Physical Review B* 67.15 (Apr. 2003). Publisher: American Physical Society, p. 155104. DOI: [10.1103/PhysRevB.67.155104](https://doi.org/10.1103/PhysRevB.67.155104). URL: <https://link.aps.org/doi/10.1103/PhysRevB.67.155104>.
- [11] Kevin C. Miao et al. “Universal coherence protection in a solid-state spin qubit”. In: *Science* 369.6510 (Sept. 2020). Publisher: American Association for the Advancement of Science, pp. 1493–1497. DOI: [10.1126/science.abc5186](https://doi.org/10.1126/science.abc5186). URL: <https://www.science.org/doi/10.1126/science.abc5186>.
- [12] Chen Avinadav et al. “Time-Optimal Universal Control of Two-Level Systems under Strong Driving”. In: *Physical Review B* 89.24 (June 2014). arXiv: 1402.4234, p. 245311. ISSN: 1098-0121, 1550-235X. DOI: [10.1103/PhysRevB.89.245311](https://doi.org/10.1103/PhysRevB.89.245311). URL: <http://arxiv.org/abs/1402.4234>.
- [13] Shi Yu et al. “Suppressing phase decoherence of a single atom qubit with Carr-Purcell-Meiboom-Gill sequence”. EN. In: *Optics Express* 21.26 (Dec. 2013). Publisher: Optica Publishing Group, pp. 32130–32140. ISSN: 1094-4087. DOI: [10.1364/OE.21.032130](https://doi.org/10.1364/OE.21.032130). URL: <https://opg.optica.org/oe/abstract.cfm?uri=oe-21-26-32130>.

- [14] Y. Nakamura, Yu A. Pashkin, and J. S. Tsai. “Coherent control of macroscopic quantum states in a single-Cooper-pair box”. en. In: *Nature* 398.6730 (Apr. 1999). Number: 6730 Publisher: Nature Publishing Group, pp. 786–788. ISSN: 1476-4687. DOI: [10.1038/19718](https://doi.org/10.1038/19718). URL: <https://www.nature.com/articles/19718>.
- [15] Jens Koch et al. “Charge-insensitive qubit design derived from the Cooper pair box”. In: *Physical Review A* 76.4 (Oct. 2007). Publisher: American Physical Society, p. 042319. DOI: [10.1103/PhysRevA.76.042319](https://doi.org/10.1103/PhysRevA.76.042319). URL: <https://link.aps.org/doi/10.1103/PhysRevA.76.042319>.
- [16] Aaron Somoroff et al. *Millisecond coherence in a superconducting qubit*. Number: arXiv:2103.08578 arXiv:2103.08578 [cond-mat, physics:quant-ph]. Mar. 2021. DOI: [10.48550/arXiv.2103.08578](https://doi.org/10.48550/arXiv.2103.08578). URL: <http://arxiv.org/abs/2103.08578>.
- [17] Crispin Gardiner, P. Zoller, and Peter Zoller. *Quantum Noise: A Handbook of Markovian and Non-Markovian Quantum Stochastic Methods with Applications to Quantum Optics*. en. Google-Books-ID: a_xsT8oGhdgC. Springer Science & Business Media, Aug. 2004. ISBN: 978-3-540-22301-6.
- [18] Markus Brink et al. “Device challenges for near term superconducting quantum processors: frequency collisions”. In: *2018 IEEE International Electron Devices Meeting (IEDM)*. ISSN: 2156-017X. Dec. 2018, pp. 6.1.1–6.1.3. DOI: [10.1109/IEDM.2018.8614500](https://doi.org/10.1109/IEDM.2018.8614500).
- [19] M. Göppl et al. “Coplanar Waveguide Resonators for Circuit Quantum Electrodynamics”. In: *Journal of Applied Physics* 104.11 (Dec. 2008). arXiv:0807.4094 [cond-mat], p. 113904. ISSN: 0021-8979, 1089-7550. DOI: [10.1063/1.3010859](https://doi.org/10.1063/1.3010859). URL: <http://arxiv.org/abs/0807.4094>.
- [20] John M. Martinis and Kevin Osborne. *Superconducting Qubits and the Physics of Josephson Junctions*. Feb. 2004. URL: <http://arxiv.org/abs/cond-mat/0402415>.
- [21] Michael Tinkham. *Introduction to Superconductivity*. 2004. URL: <https://www.bibsonomy.org/bibtex/a3ee4e22d9ab2f1dce3bf1a888cede14>.
- [22] Julian S Kelly. “Fault-tolerant superconducting qubits”. en. In: (2015), p. 207.

- [23] A. A. Houck et al. “Controlling the Spontaneous Emission of a Superconducting Transmon Qubit”. In: *Physical Review Letters* 101.8 (Aug. 2008). Publisher: American Physical Society, p. 080502. DOI: [10.1103/PhysRevLett.101.080502](https://doi.org/10.1103/PhysRevLett.101.080502). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.101.080502>.
- [24] C. Macklin et al. “A near-quantum-limited Josephson traveling-wave parametric amplifier”. In: *Science* 350.6258 (Oct. 2015). Publisher: American Association for the Advancement of Science, pp. 307–310. DOI: [10.1126/science.aaa8525](https://doi.org/10.1126/science.aaa8525). URL: <https://www.science.org/doi/10.1126/science.aaa8525>.
- [25] X. Y. Jin et al. “Thermal and Residual Excited-State Population in a 3D Transmon Qubit”. In: *Physical Review Letters* 114.24 (June 2015). Publisher: American Physical Society, p. 240501. DOI: [10.1103/PhysRevLett.114.240501](https://doi.org/10.1103/PhysRevLett.114.240501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.114.240501>.
- [26] Jay Gambetta et al. “Qubit-photon interactions in a cavity: Measurement-induced dephasing and number splitting”. en. In: *Physical Review A* 74.4 (Oct. 2006), p. 042318. ISSN: 1050-2947, 1094-1622. DOI: [10.1103/PhysRevA.74.042318](https://doi.org/10.1103/PhysRevA.74.042318). URL: <https://link.aps.org/doi/10.1103/PhysRevA.74.042318>.
- [27] Fei Yan et al. “The flux qubit revisited to enhance coherence and reproducibility”. en. In: *Nature Communications* 7.1 (Nov. 2016). Number: 1 Publisher: Nature Publishing Group, p. 12964. ISSN: 2041-1723. DOI: [10.1038/ncomms12964](https://doi.org/10.1038/ncomms12964). URL: <https://www.nature.com/articles/ncomms12964>.
- [28] Frank Arute et al. “Quantum supremacy using a programmable superconducting processor”. en. In: *Nature* 574.7779 (Oct. 2019). Number: 7779 Publisher: Nature Publishing Group, pp. 505–510. ISSN: 1476-4687. DOI: [10.1038/s41586-019-1666-5](https://doi.org/10.1038/s41586-019-1666-5). URL: <https://www.nature.com/articles/s41586-019-1666-5>.
- [29] E. Paladino et al. “ $\{1\}$ noise: Implications for solid-state quantum information”. In: *Reviews of Modern Physics* 86.2 (Apr. 2014). Publisher: American Physical Society, pp. 361–418. DOI: [10.1103/RevModPhys.86.361](https://doi.org/10.1103/RevModPhys.86.361). URL: <https://link.aps.org/doi/10.1103/RevModPhys.86.361>.

- [30] John Preskill. “Quantum Computing in the NISQ era and beyond”. en-GB. In: *Quantum* 2 (Aug. 2018). Publisher: Verein zur Förderung des Open Access Publizierens in den Quantenwissenschaften, p. 79. DOI: [10.22331/q-2018-08-06-79](https://doi.org/10.22331/q-2018-08-06-79). URL: <https://quantum-journal.org/papers/q-2018-08-06-79/>.
- [31] David Awschalom et al. “Development of Quantum Interconnects (QuICs) for Next-Generation Information Technologies”. In: *PRX Quantum* 2.1 (Feb. 2021). Publisher: American Physical Society, p. 017002. DOI: [10.1103/PRXQuantum.2.017002](https://doi.org/10.1103/PRXQuantum.2.017002). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.2.017002>.
- [32] Xianjing Zhou et al. “Single electrons on solid neon as a solid-state qubit platform”. en. In: *Nature* 605.7908 (May 2022). Number: 7908 Publisher: Nature Publishing Group, pp. 46–50. ISSN: 1476-4687. DOI: [10.1038/s41586-022-04539-x](https://doi.org/10.1038/s41586-022-04539-x). URL: <https://www.nature.com/articles/s41586-022-04539-x>.
- [33] Gershon Kurizki et al. “Quantum technologies with hybrid systems”. In: *Proceedings of the National Academy of Sciences* 112.13 (Mar. 2015). Publisher: Proceedings of the National Academy of Sciences, pp. 3866–3873. DOI: [10.1073/pnas.1419326112](https://doi.org/10.1073/pnas.1419326112). URL: <https://www.pnas.org/doi/10.1073/pnas.1419326112>.
- [34] Xu Han et al. “Microwave-optical quantum frequency conversion”. EN. In: *Optica* 8.8 (Aug. 2021). Publisher: Optica Publishing Group, pp. 1050–1064. ISSN: 2334-2536. DOI: [10.1364/OPTICA.425414](https://doi.org/10.1364/OPTICA.425414). URL: <https://opg.optica.org/optica/abstract.cfm?uri=optica-8-8-1050>.
- [35] D. De Motte et al. “Experimental system design for the integration of trapped-ion and superconducting qubit systems”. en. In: *Quantum Information Processing* 15.12 (Dec. 2016), pp. 5385–5414. ISSN: 1570-0755, 1573-1332. DOI: [10.1007/s11128-016-1368-y](https://doi.org/10.1007/s11128-016-1368-y). URL: <http://link.springer.com/10.1007/s11128-016-1368-y>.
- [36] G. Catelani et al. “Decoherence of superconducting qubits caused by quasiparticle tunneling”. In: *Physical Review B* 86.18 (Nov. 2012). Publisher: American Physical Society, p. 184514. DOI: [10.1103/PhysRevB.86.184514](https://doi.org/10.1103/PhysRevB.86.184514). URL: <https://link.aps.org/doi/10.1103/PhysRevB.86.184514>.

- [37] Simon Gustavsson et al. “Suppressing relaxation in superconducting qubits by quasiparticle pumping”. In: *Science* 354.6319 (Dec. 2016). Publisher: American Association for the Advancement of Science, pp. 1573–1577. DOI: [10.1126/science.aah5844](https://doi.org/10.1126/science.aah5844). URL: <https://www.science.org/doi/10.1126/science.aah5844>.
- [38] Uwe von Lüpke et al. “Two-Qubit Spectroscopy of Spatiotemporally Correlated Quantum Noise in Superconducting Qubits”. In: *PRX Quantum* 1.1 (Sept. 2020), p. 010305. DOI: [10.1103/PRXQuantum.1.010305](https://doi.org/10.1103/PRXQuantum.1.010305). URL: <https://link.aps.org/doi/10.1103/PRXQuantum.1.010305>.
- [39] Fei Yan et al. “Distinguishing Coherent and Thermal Photon Noise in a Circuit Quantum Electrodynamical System”. In: *Physical Review Letters* 120.26 (June 2018). Publisher: American Physical Society, p. 260504. DOI: [10.1103/PhysRevLett.120.260504](https://doi.org/10.1103/PhysRevLett.120.260504). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.120.260504>.
- [40] Fei Yan et al. “Rotating-frame relaxation as a noise spectrum analyser of a superconducting qubit undergoing driven evolution”. en. In: *Nature Communications* 4.1 (Dec. 2013), p. 2337. ISSN: 2041-1723. DOI: [10.1038/ncomms3337](https://doi.org/10.1038/ncomms3337). URL: <http://www.nature.com/articles/ncomms3337>.
- [41] Youngkyu Sung et al. “Multi-level quantum noise spectroscopy”. en. In: *Nature Communications* 12.1 (Feb. 2021). Number: 1 Publisher: Nature Publishing Group, p. 967. ISSN: 2041-1723. DOI: [10.1038/s41467-021-21098-3](https://doi.org/10.1038/s41467-021-21098-3). URL: <https://www.nature.com/articles/s41467-021-21098-3>.
- [42] Fumiki Yoshihara et al. “Flux qubit noise spectroscopy using Rabi oscillations under strong driving conditions”. In: *Physical Review B* 89.2 (Jan. 2014). Publisher: American Physical Society, p. 020503. DOI: [10.1103/PhysRevB.89.020503](https://doi.org/10.1103/PhysRevB.89.020503). URL: <https://link.aps.org/doi/10.1103/PhysRevB.89.020503>.
- [43] P. Forn-Díaz et al. “Broken selection rule in the quantum Rabi model”. en. In: *Scientific Reports* 6.1 (June 2016). Number: 1 Publisher: Nature Publishing Group, p. 26720. ISSN: 2045-2322. DOI: [10.1038/srep26720](https://doi.org/10.1038/srep26720). URL: <https://www.nature.com/articles/srep26720>.

- [44] Daniel L. Campbell et al. “Universal Nonadiabatic Control of Small-Gap Superconducting Qubits”. In: *Physical Review X* 10.4 (Dec. 2020). Publisher: American Physical Society, p. 041051. DOI: [10.1103/PhysRevX.10.041051](https://doi.org/10.1103/PhysRevX.10.041051). URL: <https://link.aps.org/doi/10.1103/PhysRevX.10.041051>.
- [45] David C. McKay et al. “Efficient Z-Gates for Quantum Computing”. In: *Physical Review A* 96.2 (Aug. 2017). arXiv:1612.00858 [quant-ph], p. 022330. ISSN: 2469-9926, 2469-9934. DOI: [10.1103/PhysRevA.96.022330](https://doi.org/10.1103/PhysRevA.96.022330). URL: <http://arxiv.org/abs/1612.00858>.
- [46] J. R. Johansson, P. D. Nation, and Franco Nori. “QuTiP 2: A Python framework for the dynamics of open quantum systems”. en. In: *Computer Physics Communications* 184.4 (Apr. 2013), pp. 1234–1240. ISSN: 0010-4655. DOI: [10.1016/j.cpc.2012.11.019](https://doi.org/10.1016/j.cpc.2012.11.019). URL: <https://www.sciencedirect.com/science/article/pii/S0010465512003955> (visited on 11/04/2022).
- [47] Easwar Magesan, Jay M. Gambetta, and Joseph Emerson. “Characterizing quantum gates via randomized benchmarking”. In: *Physical Review A* 85.4 (Apr. 2012). Publisher: American Physical Society, p. 042311. DOI: [10.1103/PhysRevA.85.042311](https://doi.org/10.1103/PhysRevA.85.042311). URL: <https://link.aps.org/doi/10.1103/PhysRevA.85.042311>.
- [48] R. Barends et al. “Superconducting quantum circuits at the surface code threshold for fault tolerance”. en. In: *Nature* 508.7497 (Apr. 2014). Number: 7497 Publisher: Nature Publishing Group, pp. 500–503. ISSN: 1476-4687. DOI: [10.1038/nature13171](https://doi.org/10.1038/nature13171). URL: <https://www.nature.com/articles/nature13171>.
- [49] E. Knill et al. “Randomized benchmarking of quantum gates”. In: *Physical Review A* 77.1 (Jan. 2008). Publisher: American Physical Society, p. 012307. DOI: [10.1103/PhysRevA.77.012307](https://doi.org/10.1103/PhysRevA.77.012307). URL: <https://link.aps.org/doi/10.1103/PhysRevA.77.012307>.
- [50] M. D. Hutchings et al. “Tunable Superconducting Qubits with Flux-Independent Coherence”. In: *Physical Review Applied* 8.4 (Oct. 2017), p. 044003. ISSN: 2331-7019. DOI: [10.1103/PhysRevApplied.8.044003](https://doi.org/10.1103/PhysRevApplied.8.044003). URL: <http://arxiv.org/abs/1702.02253>.

- [51] F. Motzoi et al. “Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits”. In: *Physical Review Letters* 103.11 (Sept. 2009). Publisher: American Physical Society, p. 110501. DOI: [10.1103/PhysRevLett.103.110501](https://doi.org/10.1103/PhysRevLett.103.110501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.103.110501>.
- [52] Z. Kim et al. “Decoupling a Cooper-Pair Box to Enhance the Lifetime to 0.2 ms”. In: *Physical Review Letters* 106.12 (Mar. 2011). Publisher: American Physical Society, p. 120501. DOI: [10.1103/PhysRevLett.106.120501](https://doi.org/10.1103/PhysRevLett.106.120501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.106.120501>.
- [53] Long B. Nguyen et al. “The high-coherence fluxonium qubit”. In: *Physical Review X* 9.4 (Nov. 2019), p. 041041. ISSN: 2160-3308. DOI: [10.1103/PhysRevX.9.041041](https://doi.org/10.1103/PhysRevX.9.041041). URL: <http://arxiv.org/abs/1810.11006>.
- [54] Junling Long et al. *A universal quantum gate set for transmon qubits with strong ZZ interactions*. arXiv:2103.12305 [quant-ph]. Mar. 2021. URL: <http://arxiv.org/abs/2103.12305> (visited on 11/21/2022).
- [55] Ray Freeman. “Shaped radiofrequency pulses in high resolution NMR”. en. In: *Progress in Nuclear Magnetic Resonance Spectroscopy* 32.1 (Feb. 1998), pp. 59–106. ISSN: 0079-6565. DOI: [10.1016/S0079-6565\(97\)00024-1](https://doi.org/10.1016/S0079-6565(97)00024-1). URL: <https://www.sciencedirect.com/science/article/pii/S0079656597000241> (visited on 11/17/2022).
- [56] Robin Harper and Steven T. Flammia. “Estimating the fidelity of T gates using standard interleaved randomized benchmarking”. In: *Quantum Science and Technology* 2.1 (Mar. 2017). arXiv: 1608.02943, p. 015008. ISSN: 2058-9565. DOI: [10.1088/2058-9565/aa5f8d](https://doi.org/10.1088/2058-9565/aa5f8d). URL: <http://arxiv.org/abs/1608.02943>.
- [57] G. Ithier et al. “Decoherence in a superconducting quantum bit circuit”. In: *Physical Review B* 72.13 (Oct. 2005). Publisher: American Physical Society, p. 134519. DOI: [10.1103/PhysRevB.72.134519](https://doi.org/10.1103/PhysRevB.72.134519). URL: <https://link.aps.org/doi/10.1103/PhysRevB.72.134519>.
- [58] *Building scalable, innovative quantum systems*. en. URL: <https://www.rigetti.com/what-we-build>.
- [59] *A near-quantum-limited Josephson traveling-wave parametric amplifier* / *Science*. URL: <https://www.science.org/doi/10.1126/science.aaa8525>.

- [60] X. Y. Jin et al. “Thermal and Residual Excited-State Population in a 3D Transmon Qubit”. en. In: *Physical Review Letters* 114.24 (June 2015), p. 240501. ISSN: 0031-9007, 1079-7114. DOI: [10.1103/PhysRevLett.114.240501](https://doi.org/10.1103/PhysRevLett.114.240501). URL: <https://link.aps.org/doi/10.1103/PhysRevLett.114.240501>.
- [61] Harrison Ball, William D. Oliver, and Michael J. Biercuk. “The role of master clock stability in quantum information processing”. en. In: *npj Quantum Information* 2.1 (Nov. 2016). Number: 1 Publisher: Nature Publishing Group, pp. 1–8. ISSN: 2056-6387. DOI: [10.1038/npjqi.2016.33](https://doi.org/10.1038/npjqi.2016.33). URL: <https://www.nature.com/articles/npjqi201633>.
- [62] F. Yan et al. “The Flux Qubit Revisited to Enhance Coherence and Reproducibility”. In: *Nature Communications* 7.1 (Dec. 2016). arXiv:1508.06299 [quant-ph], p. 12964. ISSN: 2041-1723. DOI: [10.1038/ncomms12964](https://doi.org/10.1038/ncomms12964). URL: <http://arxiv.org/abs/1508.06299>.
- [63] Jonas Bylander et al. “Noise spectroscopy through dynamical decoupling with a superconducting flux qubit”. en. In: *Nature Physics* 7.7 (July 2011). Number: 7 Publisher: Nature Publishing Group, pp. 565–570. ISSN: 1745-2481. DOI: [10.1038/nphys1994](https://doi.org/10.1038/nphys1994). URL: <https://www.nature.com/articles/nphys1994>.
- [64] Daniel L. Campbell et al. *Modular tunable coupler for superconducting qubits*. Number: arXiv:2207.06607 arXiv:2207.06607 [quant-ph]. July 2022. URL: <http://arxiv.org/abs/2207.06607>.
- [65] John M Martinis and Kevin Osborne. “Superconducting Qubits and the Physics of Josephson Junctions”. en. In: (), p. 26.
- [66] He-Liang Huang et al. “Superconducting Quantum Computing: A Review”. In: *Science China Information Sciences* 63 (June 2020). DOI: [10.1007/s11432-020-2881-9](https://doi.org/10.1007/s11432-020-2881-9).

Curriculum Vitae

Personal Details

Name: Michael Senatore
Place of Birth: Staten Island, New York
Date of Birth: January, 20, 1992
Email: Michael.Senatore.Research@gmail.com

Education

Syracuse University

Doctor of Philosophy in Physics
Advisor: Dr. Britton Plourde
Expected graduation date: Dec 2022

Syracuse University

Master's in Physics
Graduation date: December 2017

Colgate University

Bachelor's of Art in Physics
Bachelor's of Art in Philosophy
Graduation date: May 2014

Publications

- In preparation: Michael Senatore et. al. Fast Universal Gating of a Coherence-enhanced Flux Sensitive Transmon

- Campbell, D. L., Kamal, A., Ranzani, L., Senatore, M., LaHaye, M. (2022). Modular tunable coupler for superconducting qubits (arXiv:2207.06607). arXiv. <http://arxiv.org/abs/2207.06607>

Presentations and Talks

- "Demonstration of high-fidelity universal gates on a continuously noise-decoupled qubit" Michael Senatore, Daniel L Campbell, Oleksiy Redko, Matthew LaHaye—Bulletin of the American Physical Society, 2022
- "Modular coupling approach using ancilla transmons with flux-tunable hybridization" Daniel L Campbell, Michael Senatore, Archana Kamal, Leonardo M Ranzani, Matthew LaHaye—Bulletin of the American Physical Society, 2022
- "Gate sequencing approach to reduce calibration and equipment overhead in qubit control" Oleksiy Redko, Michael Senatore, Daniel L Campbell, Matthew LaHaye—Bulletin of the American Physical Society, 2022
- "Characterization and control of topologically protected charge-parity qubits: Part 1" Kenneth Dodge, Yebin Liu, Michael Senatore, FNU Naveen, Shaojiang Zhu, Abigail Shearrow, Andrey Klots, Lara Faoro, Lev Ioffe, Robert McDermott, Britton Plourde—Bulletin of the American Physical Society, 2020
- "Characterization of topologically protected charge-parity qubits: Part 2" FNU Naveen, Abigail Shearrow, Shaojiang Zhu, Kenneth Dodge, Yebin Liu, Michael Senatore, Andrey Klots, Lara Faoro, Lev Ioffe, Britton Plourde, Robert McDermott— Bulletin of the American Physical Society, 2020
- "Implementation of π -periodic Josephson Elements for Topologically Protected Charge-Parity Qubits", Yebin Liu, Kenneth Dodge, Michael Anthony Senatore, Shaojiang Zhu, FNU Naveen, Abigail J Shearrow, Francisco Schlenker, Andrey Klotz, Lara Faoro, Lev B Ioffe, Robert F McDermott, B.L.T. Plourde—March Meeting of the American Physical Society Bulletin, 2019, Boston, MA
- "Fabrication of Superconducting Devices for Quantum Information Science", B.L.T. Plourde, Andrew Ballard, Jaesung Ku, Yebin Liu, Michael Senatore,

Physics and Nanostructure Physics, Research Accomplishments 2017-2018, CNF at Cornell University

- “Microwave Response of Vortices in Superconducting Resonators with High Kinetic Inductance”, Kenneth Dodge, JJ Nelson, Michael Senatore, Peng Xu, Kevin Osborn, David Pappas, Britton Plourde–March Meeting of the American Physical Society Bulletin, 2018, Los Angeles, CA
- "Micro-rotation Through Indirect Manipulation of Absorbing Objects," C. Herne and M. A. Senatore, Meetings, OSA Meeting Technical Digest, 2014.
- “Production of two-photon cluster states in polarization and spatial modes,” E.J. Galvez, W.H. Schubert and M. Senatore. Contributed talk at the conference of the SPIE Photonics West, San Francisco, February 6, 2013.
- “Beyond qubits with spatial modes of light,” E.J. Galvez, M. Senatore, X. Cheng, and B.L. Rojec. Contributed abstract at the conference of the SPIE Defense and Security, Baltimore, May 3, 2013.