

Abstract

Improving Coherence of Superconducting Qubits and Resonators

Kurtis Lee Geerlings

2013

Superconducting qubits and resonators with quality factors exceeding 10^7 are of great interest for quantum information processing applications. The improvement of present devices necessarily involves the consideration of participation ratios, which budget the influence of each physical component in the total energy decay rate. Experiments on compact resonators in which participation ratios were varied has demonstrated the validity of this method, yielding a two-fold improvement in quality factor. Similar experiments on compact transmon qubit devices led to a three-fold improvement over previous transmons, validating the method of participation ratios for qubits as well. Through the use of a 3D cavity, a further minimization of the participation of surface components combined with the removal of unnecessary components, produced an additional ten-fold increase in coherence times. Finally, the fluxonium qubit was redesigned in a similar minimalist environment with an improved superinductance, thus combining the advantages of the 3D architecture with the natural insensitivity to dissipation of the fluxonium, resulting in *another* ten-fold increase in relaxation times. This large increase in relaxation and coherence times enables experiments that were previously impossible, thus preparing the field of quantum information to advance on other fronts.

Improving Coherence of Superconducting Qubits and Resonators

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

by
Kurtis Lee Geerlings

Dissertation Director: Michel H. Devoret

August 19, 2013

Copyright © 2013 by Kurtis Lee Geerlings

All rights reserved.

Acknowledgements

This work was only possible through the effort and conscious guidance of many people and the years of prior research upon which the included work was built.

Firstly, and perhaps most directly involved and responsible, was my advisor, Michel Devoret. His long history in the field and infinite patience enabled him to explain even the most complicated topics from first principles, leading to discussions often many hours long. Brainstorming ideas with Michel was one of the best ways to spend time and somehow always left me more motivated and excited about my work. Michel cultivated a lab environment with loose controls that rewarded individual responsibility and ambition. For these reasons, I remain forever grateful to Michel for his influence on my work and my life.

The other professors at Yale that were each responsible for success through many discussions, ideas, and general levity are: Rob Schoelkopf, Leonid Glazman, Steve Girvin, Dan Prober and Mazyar Mirrahimi. This group, along with Michel, worked so closely on so many projects that the sum was surely greater than the parts. This close, supporting collaboration, unmatched in other Universities, created an environment of robust creativity, unparalleled enthusiasm and unbridled success.

In addition, the staff researcher, Luigi Frunzio, provided a sense of stability for the ever-changing group of graduate students. Luigi was a living history of past projects and fabrication recipes, and was a great resource for profitable discussion. Luigi's

booming, jubilant voice could easily be heard down the hallway, and was always a sign of exciting news.

Through my time in Qulab, I worked closely with three postdocs on various projects. Markus Brink, Shyam Shankar, and Ioan Pop were excellent partners, each with their own unique style of thinking and working. In working with them, I always had an excellent soundboard for ideas and a partner with whom to divide and conquer the necessary work.

Many other graduate students and postdocs, both in Michel and Rob's groups, contributed significantly to previous and in-parallel work, without which these experiments would not have been possible. This includes, in alphabetical order: Baleegh Abdo, Teresa Brecht, Lev Bishop, Jacob Blumoff, Gianluigi Catelani, Kevin Chou, Jerry Chow, Leonardo DiCarlo, Eustace Edwards, Andreas Fagner, Eran Ginossar, Michael Hatridge, Eric Holland, Blake Johnson, Archana Kamal, Gerhard Kirchmair, Zaki Leghtas, Yehan Liu, Vlad Manucharyan, Nick Masluk, Zlatko Minev, Anirudh Narla, Simon Nigg, Nissim Ofek, Hanhee Paik, Andrei Petrenko, Matt Reagor, Matthew Reed, Flavius Schackert, David Schuster, Adam Sears, Katrina Sliwa, Luyan Sun, Brian Vlastakis, Uri Vool, Chen Wang, and Terri Yu.

I would also like to thank my family for supporting me through the whole process. First and most importantly my wife, Heidi, who enthusiastically moved to New Haven and supported me financially throughout the whole process. Although my research meant long hours and delaying the decision to have children, her support was never in question. I would also like to thank my parents, whose guidance in my childhood put me on track to succeed as an adult. With my mom always teaching me subjects a few years before learning them in school and my dad showing by example how to never give up, I am forever indebted to them for my past and future achievements.

Lastly, I would like to thank my daughter Alina, who was born during the writing

of this thesis, for not crying during my dissertation defense.

Contents

Acknowledgements	iii
Contents	vi
List of Figures	x
List of Symbols	xvi
1 Introduction	1
1.1 Quantum Computing	3
1.1.1 What is a Qubit?	3
1.1.2 Quantum Resources	4
1.1.3 Qubit Coherence	6
1.1.4 Qubit Initialization	7
1.1.5 Circuit QED	8
1.2 Superconducting Qubit Species	9
1.2.1 Transmon	13
1.2.2 Fluxonium	16
1.3 Coherence Limitations	18
1.3.1 Relaxation Sources	18
1.3.2 Dephasing Sources	23

1.4	Thesis Overview	26
2	Resonator and Qubit Theory	30
2.1	Resonator Theory	30
2.1.1	2D Resonator Theory	30
2.1.2	3D Resonator Theory	32
2.2	Superconducting Artificial Atom Theory	34
2.2.1	Transmon Theory	34
2.2.2	Fluxonium Theory	43
3	Experimental Methods	57
3.1	Sample and Environment	57
3.1.1	Fabrication	57
3.1.2	Octobox Sample Holder	61
3.2	Qubit Readout	63
3.2.1	Dispersive Readout	64
3.2.2	High Power Readout	68
3.2.3	Heterodyne Measurement	71
3.3	Equipment	73
3.3.1	Dilution Refrigerators	73
3.3.2	Microwave Switches	76
3.3.3	Shields	80
4	Compact Resonators	90
4.1	Introduction	91
4.1.1	TLS Dissipation	91
4.1.2	Simulation	94

4.1.3	Design	99
4.2	Experimental Apparatus	101
4.3	Results	104
4.4	Analysis	111
4.5	Discussion	112
5	Compact Transmon	115
5.1	Sample Design	115
5.1.1	Simulation	117
5.2	Experimental Apparatus	120
5.3	Results	124
5.4	Analysis	132
5.5	Discussion	134
6	3D Transmon	136
6.1	Introduction	136
6.1.1	Origin of 3D Transmon	136
6.1.2	Existing Implementations	138
6.1.3	3D Transmon as “Calibration”	142
6.1.4	Reset Experiment	142
6.2	Experimental Apparatus	144
6.3	Results	144
6.3.1	Coherence Results	144
6.3.2	Qubit Reset	149
6.4	Analysis	153
6.5	Discussion	155

7 Fluxonium	157
7.1 Design	158
7.1.1 Superinductance	158
7.1.2 3D Cavity	159
7.1.3 Coupling	160
7.2 Experimental Apparatus	162
7.3 Results	164
7.4 Analysis	176
7.5 Discussion	185
8 Conclusions and Perspectives	189
A Resonator “Hanger” Equation	192
A.1 Derivation	193
A.2 Asymmetry Discussion	199
B Temperature Measurement Protocol	204
C Qubit Reset Theory	209
Bibliography	213

List of Figures

1.1	Bloch sphere representation of a qubit	4
1.2	Josephson tunnel junction schematic and circuit representation	11
1.3	Junction constitutive relation	12
1.4	Cooper Pair Box circuit schematic	12
1.5	Partial qubit hierarchy	13
1.6	Transmon circuit schematic	14
1.7	Transmon energy levels vs. gate charge	17
1.8	Fluxonium circuit schematic	18
1.9	Circuit schematic of qubit loss mechanisms	20
1.10	Qubit and cavity lifetimes summary over 15 years	29
2.1	Rectangular 3D cavity electric fields	34
2.2	Circuit schematic of transmon coupled to resonator	36
2.3	Qubit-cavity avoided crossing	40
2.4	Fluxonium parameters	44
2.5	Fluxonium wavefunctions and energy potential	47
2.6	Fluxonium transition spectrum	49
2.7	Fluxonium capacitance matrix elements	51
2.8	Fluxonium inductive matrix elements	52
2.9	Fluxonium quasiparticle matrix elements	54

2.10	Cartoon of relaxation due to quasiparticles in fluxonium	55
2.11	Symmetry of fluxonium quasiparticle matrix element	56
3.1	Octobox sample holder, open	62
3.2	Octobox sample holder, closed	62
3.3	Dispersive readout: amplitude and phase	66
3.4	Dispersive readout: IQ plane	67
3.5	High power readout bright state	69
3.6	High power readout contrast	70
3.7	High power readout histograms	71
3.8	Decay after high power readout	72
3.9	Heterodyne measurement experimental setup	73
3.10	Dilution refrigerator schematic	75
3.11	Pictures of Kelvinox25 and Kelvinox400 dilution refrigerators	77
3.12	Radial SP6T switch	79
3.13	Quinstar 8-12 GHz cryogenic circulator	81
3.14	K&L low-pass microwave filters	82
3.15	Eccosorb filters: box and pipe style	84
3.16	Measured response of Eccosorb filters	85
3.17	Copper powder filter	87
3.18	Infrared shield images	88
3.19	Coating on inside of infrared shield	89
4.1	TLS model	91
4.2	Participation of surface in CPW mode	95
4.3	Circuit schematic of inductor simulations	96
4.4	Inductor simulation results	96

4.5	Circuit schematic of capacitor simulations	97
4.6	Capacitor Simulation results	98
4.7	Compact resonator “Design A”	100
4.8	Compact resonator all designs	100
4.9	Experimental setup of compact resonator experiment (Kelvinox25) . .	102
4.10	Experimental setup of compact resonator experiment (Kelvinox400) .	103
4.11	Asymmetric resonator fit examples	104
4.12	Compact resonator Q_i vs. \bar{n}	105
4.13	Compact resonator frequency vs. temperature	106
4.14	Compact resonator, effect of different g_C	107
4.15	Dependence of Q_i on parameter values of compact resonators	108
4.16	Compact resonator Q_i vs. design	110
4.17	Compact resonator Q_i , fridge comparison	112
5.1	SEM image of CPW transmon	116
5.2	T_1 of CPW transmon	117
5.3	Circuit schematic of transmon coupled to resonator	118
5.4	Maxwell setup for capacitor simulation	119
5.5	Optical image of compact transmon device	121
5.6	Experimental setup of compact resonator experiment (Kelvinox400) .	122
5.7	Qubit temperature vs. magnetic field current	123
5.8	Basic transmon measurements	124
5.9	transmon measurement of T_1 , T_2^R , and T_2^E	126
5.10	Summary of coherence results for compact transmons	127
5.11	Coherence times vs. qubit frequency, compact transmon	129
5.12	Measured π -pulse power vs. frequency	131

5.13	Relaxation time vs. temperature for compact transmons	131
6.1	Cartoon of surface participation in 2D vs. 3D resonators	138
6.2	Image of physical 3D cavity	139
6.3	3D transmon device image	140
6.4	Experimental setup of 3D transmon experiment (Kelvinox400)	145
6.5	3D transmon coherence vs. frequency	146
6.6	Measurements of T_1 , T_2^R , and T_2^E for 3D transmon	147
6.7	Measurements of T_1 , T_2^R , and T_2^E for 3D transmon at low frequency	148
6.8	Relaxation time over 14 hours for 3D transmon	148
6.9	Effect of DDROP sequence vs. pulse duration	150
6.10	Expected vs. measured DDROP fidelity vs. \bar{n}	151
6.11	Expected vs. measured DDROP fidelity vs. Ω_R	152
6.12	Excited state measurements after DDROP sequence	154
7.1	Previous and updated fluxonium junction array designs	159
7.2	Circuit representation of fluxonium coupled to 3D cavity	161
7.3	Spectroscopy of 3D cavity and antenna mode vs. magnetic field	162
7.4	Images of fluxonium qubit and coupling	163
7.5	Measured fluxonium frequency vs. flux bias	165
7.6	Spectroscopy of Fluxonim qubit for different regions	167
7.7	χ measurement of fluxonium qubit	168
7.8	Coherence measurements of fluxonium at $\Phi_{ext}/\Phi_0 = 0.00$	169
7.9	Coherence measurements of fluxonium at $\Phi_{ext}/\Phi_0 = 0.03$	170
7.10	Single-exponential relaxation measurements of fluxonium at $\Phi_{ext}/\Phi_0 =$ 0.50	171

7.11 Double-exponential relaxation measurements of fluxonium at $\Phi_{ext}/\Phi_0 =$ 0.50	172
7.12 Log-scale relaxation measurement of fluxonium at $\Phi_{ext}/\Phi_0 = 0.50$. .	174
7.13 Measured T_1 as a function of applied flux for fluxonium	176
7.14 Coherence measurements of fluxonium at $\Phi_{ext}/\Phi_0 = 0.50$	177
7.15 Expected vs. measured relaxation times for capacitive loss (constant Q_{cap})	178
7.16 Expected vs. measured relaxation times for capacitive loss	179
7.17 Expected vs. measured relaxation times for inductive loss	179
7.18 Expected vs. measured relaxation times loss due to the Purcell effect	180
7.19 Expected vs. measured relaxation times for quasiparticle loss	181
7.20 Expected vs. measured relaxation times for quasiparticle loss near $\Phi_{ext}/\Phi_0 = 0.50$	182
7.21 Expected vs. measured relaxation times for combined loss	182
7.22 Predicted relaxation rates for possible explanation of double exponential	186
A.1 Resonance lineshapes of different measurements	193
A.2 Schematic of generic 2-port device	194
A.3 Circuit model of generic 2-port device	194
A.4 Theoretical hanger response curves for varying asymmetry	201
B.1 Sources of qubit and cavity temperature measurements	205
B.2 RPM qubit population measurement scheme	206
C.1 Level diagram of DDROP qubit reset procedure	211

List of Symbols

Constants

A	characteristic amplitude of various noise sources
$a_\nu(x)$	Mathieu's characteristic value
c	the speed of light in vacuum ($\approx 3 \times 10^8$ m/s)
e	electron charge ($\approx 1.602 \times 10^{-19}$ C)
h	Planck's constant ($\approx 6.626 \times 10^{-34}$ m ² kg/s)
\hbar	Reduced Planck's constant ($\equiv h/2\pi$)
i	“physicist's” imaginary unit, $i = +\sqrt{-1} = -j$
j	“engineer's” imaginary unit, $j = -\sqrt{-1} = -i$
k_B	Boltzmann constant ($\approx 1.381 \times 10^{-23}$ m ² kg/s ²)
R_Q	the resistance quantum ($\equiv \hbar/(2e)^2 \approx 1027\Omega$)
Φ_0	magnetic flux quantum ($\equiv h/2e \approx 2.067 \times 10^{-15}$ Wb)
ϕ_0	reduced magnetic flux quantum ($\equiv \hbar/2e = \Phi_0/2\pi$)
π	Ratio of circle circumference to diameter (≈ 3.141592)

Kets

$ g\rangle, 0\rangle$	quantum ground state
$ e\rangle, 1\rangle$	first quantum excited state
$ f\rangle$	second quantum excited state
$ \Psi\rangle$	arbitrary qubit state ($ \Psi\rangle = \alpha_g g\rangle + \alpha_e e\rangle$)

Abbreviations

AWG	Arbitrary Waveform Generator
CPB	Cooper Pair Box
CPMG	Carr-Purcell-Meiboom-Gill pulse sequence
CW	Continuous Wave excitation
DDROP	Double Drive Reset of Population
HEMT	High Electron Mobility Transistor
I-Q	In-phase and out-of-phase signal quadratures
IF	Intermediate Frequency
LC	short for oscillator made of inductor (L) and capacitor (C)
LO	Local Oscillator
MBE	Molecular Beam Epitaxy
NMR	Nuclear Magnetic Resonance
PCB	Printed Circuit Board
PMMA	Poly(Methyl MethAcrylate)
QED	Quantum ElectroDynamics
QND	Quantum Non-Demolition
RF	Radio Frequency

RMS	Root Mean Square
RPM	Rabi Population Measurement
SEM	Scanning Electron Microscope
SNR	Signal-to-Noise Ratio
SQUID	Superconducting QUantum Interference Device
TE	Transverse Electric field mode
TLS	Two Level System
TM	Transverse Magnetic field mode
UHV	Ultra-High Vacuum

Parameters and Variables

generic

C	generic capacitance
E_n	energy of n^{th} qubit or cavity level
f	frequency ($f = \omega/2\pi$)
f_0, ω_0	resonance frequency
H	Hamiltonian
I	current
I_0	Josephson tunnel junction critical current
L	generic inductance
M	mass
P	momentum
Q	charge or quality factor
T	temperature
T_c	critical temperature of a superconductor

t	time
$\tan\delta$	loss tangent
V	voltage
X	position coordinate
Y	admittance
Φ	branch flux or magnetic flux through a superconducting loop
ϕ	flux or phase (branch flux or superconducting phase)
ω	radial frequency ($\omega = 2\pi f$)
N, n, m, l	integer index of energy levels
$\varepsilon, \varepsilon', \varepsilon''$	generic permittivity ($\varepsilon = \varepsilon' - j\varepsilon''$)

qubits

C_J	capacitance of a Josephson tunnel junction
C_{ext}	external capacitance of transmon or fluxonium qubit
C_Σ	total capacitance of a qubit (for transmon, $C_\Sigma = C_J + C_{ext}$)
E_{ij}	energy of qubit $ i\rangle$ to $ j\rangle$ transition
E_C	charging energy ($\equiv e^2/2C_\Sigma$)
E_L	inductive energy ($\equiv \phi_0^2/2L_\Sigma$)
E_J	Josephson energy ($\equiv \phi_0 I_0$)
E_g, E_e, E_f	energy of qubit ground state, excited state, and second excited state
f_{ij}, ω_{ij}	qubit $ i\rangle$ to $ j\rangle$ transition frequency
n_g	offset charge
Q_1	quality factor of qubit mode ($Q_1 = \omega_Q T_1$)
T_1	qubit relaxation time
T_2	qubit decoherence time ($1/T_2 = 1/2T_1 + 1/T_\phi$)

T_2^E	Echo decay time, a measurement of T_2
T_2^R	Ramsey decay time, a measurement of T_2
T_ϕ	qubit dephasing time
V_g	gate voltage
α_g, α_e	ground and excited state probability amplitude
$\delta\omega$	fluctuations in qubit frequency
$\Gamma_\uparrow, \Gamma_\downarrow$	rate of qubit excitation and relaxation
Φ_{ext}	external flux applied through a superconducting loop
ϕ_B	Bloch sphere representation: azimuthal angle
θ_B	Bloch sphere representation: polar angle
ω_Q	qubit transition frequency

resonators

f_{mnl}	frequency of 3D cavity mode of index (m,n,l)
n_{th}	average number of photons in cavity due to thermal excitation
Q_0	total quality factor ($1/Q_0 = 1/Q_i + 1/Q_C$)
Q_i	internal quality factor
Q_C	coupling quality factor
TE_{mnl}, TM_{mnl}	frequency of TE or TM cavity mode of index (m,n,l)
ω_R, ω_0	resonator frequency

circuit QED

a, a^\dagger	cavity raising and lowering operators
E_{N_C, N_Q}	energy of qubit-cavity system with specified number of excitations
f_C^g, f_C^e, f_C^f	cavity frequency when qubit is in specified state

$f_{ge}^{N_C}$	qubit frequency specified for cavity excitation number
g	qubit-cavity coupling strength
H_{JC}	Jaynes-Cummings Hamiltonian
N_Q, N_C	number of excitations in qubit and cavity, respectively
\bar{n}	average number of photons in a cavity
α	qubit anharmonicity
Δ	detuning between qubit and cavity ($\Delta = \omega_Q - \omega_R$)
κ	linewidth of cavity mode
σ_z	spin 1/2 Pauli matrix
σ^-, σ^+	spin 1/2 (or two-level system) raising and lowering operators
χ	dispersive shift between qubit and cavity modes

circuit QED derivation

$A, A^\dagger, B, B^\dagger$	raising and lowering operators of dressed modes
$a, a^\dagger, b, b^\dagger$	raising and lowering operators of bare qubit and cavity modes
C_C	coupling capacitance between resonators
C_S	shunting capacitance of “qubit” mode
C_R	resonator capacitance
H_0	Hamiltonian of two coupled harmonic oscillators
H_1	anharmonicity perturbation Hamiltonian
L_R	resonator inductance
N_A, N_B	number of excitations in two dressed modes
Q_1, Q_2	charge of capacitor of either mode
Z_a, Z_b	characteristic impedance of bare qubit and cavity modes
ϕ_1, ϕ_2	phase across inductor of either mode
$\lambda_a, \lambda_b, \mu_a, \mu_b$	conversion matrix of (a,b) to (A,B)

χ_A, χ_B	self-Kerr of dressed modes
χ_{AB}	cross-Kerr between dressed modes
ω_A, ω_B	dressed mode resonant frequencies
ω'_A, ω'_B	final qubit and cavity frequencies with anharmonicity
ω_a, ω_b	qubit and cavity bare frequencies

participation ratios

C_E, L_E	environmental capacitance and inductance
G_E^{cap}, G_E^{ind}	conductance of environmental capacitance and inductance
G_{rad}	conductance due to radiation loss
G_J^{cap}, G_J^{ind}	conductance of junction capacitance and inductance
Q_{cap}	quality factor due to capacitance
p_{cap}^i	quality factor of i^{th} capacitive component ($\sum p_{cap}^i = 1$)
Q_{ind}	quality factor due to inductance
p_{ind}^i	quality factor of i^{th} inductive component ($\sum p_{ind}^i = 1$)
Q_{rad}	quality factor due to radiation

TLS loss

d	superconducting film thickness
E	RMS electric field amplitude
E_S	saturating electric field
F_{TLS}	filling factor for TLS loss
L_k	kinetic inductance
K_0, I_0	modified Bessel functions of first and second kind
α	kinetic inductance fraction
$\Delta\varepsilon(\omega, T)$	change in dielectric constant with frequency and temperature

$\Delta\sigma_2(\omega, T)$	change in imaginary conductivity with frequency and temperature
Ψ	complex digamma function
$\sigma, \sigma_1, \sigma_2$	complex conductivity of a superconductor ($\sigma = \sigma_1 - i\sigma_2$)
σ_N	normal-state conductivity

transmon

C_{ij}	element of capacitance matrix
C_X	particular combination of capacitance matrix elements
p_{surf}	surface participation ratio
t_{surf}	surface thickness
t_{mode}	distance between two conductors in resonant mode
X_{ij}	inverse of inverse capacitance matrix ($= 1/C_{ij}^{-1}$)

compact resonators

l_C, l_L	total capacitor or inductor length
g_C, g_L	gap between adjacent capacitor fingers or inductor meanders
g_R	gap between compact resonator and ground plane
w_C, w_L, w	width of capacitor traces, inductor traces, or both
Z_0	characteristic impedance of resonator or transmission line

fluxonium

C_g	capacitance to ground of each array island
C_{JA}	capacitance of an array junction
L_{JA}	inductance of an array junction
H_l	Hermite polynomials

N_A	number of junctions in array
Z_{JA}	impedance of an array junction ($\equiv \sqrt{L_{JA}/C_{JA}}$)
Z_J	impedance of phase-slip junction ($\equiv \sqrt{L_J/C_J}$)
ψ_l	wavefunction of l^{th} eigenstate
C_{pins}	capacitance between microwave lines and 3D cavity mode
C_{pad}	capacitance between antenna pads and 3D cavity mode
C_{self}	self-capacitance of antenna pads

fluxonium losses

\hat{C}	generic system operator that couples to noise source
\hat{X}	generic reservoir operator carries noise into system
S_{XX}	spectral density of operator \hat{X}
$\hat{\Phi}$	operator for phase across the phase-slip junction
$\hat{\varphi}$	operator for phase normalized by flux quantum
H_C, H_C^{cap}	coupling Hamiltonian for specific loss mechanism
H_C^{ind}, H_C^{qp}	
$\Gamma_C^\downarrow, \Gamma_C^\uparrow$	qubit relaxation and excitation rate due to generic component
$\Gamma^{cap}, \Gamma^{ind}, \Gamma^{qp}$	relaxation rate due to loss in specific component
$\hat{I}_{cap}, \hat{I}_{ind}$	operator of current in inductor or capacitor
Y_C, Y_{cap}, Y_{ind}	admittance of specific component
$Y_{qp}, Y_{purcell}$	
Q_{cap}, Q_{ind}	effective Q of capacitive or inductive components
x_{qp}	quasiparticle density normalized to Cooper pair density
G_t	conductance of a Josephson junction

appendix A

A, B, C, D	components of microwave ABCD-matrix
a, b, c, d, e	any complex number
f, g, h, k	
a_i, b_i, c_i, d_i	any complex number
$S_{11}, S_{12}, S_{21}, S_{22}$	components of microwave scattering matrix
x	scaled frequency, ($x = \frac{\omega - \omega_0}{\omega_0}$)
x_0	general Lorentzian resonance location
z_i, z_1, z_2, z_3	any impedance
ε	resonator internal loss rate
γ	total resonator loss rate
κ	combination of coupling loss and asymmetry ($\equiv \kappa_R + i\kappa_I$)
κ_R	resonator coupling loss rate
$\kappa_I, \delta\omega, \delta f$	hanger resonator asymmetry parameters
\tilde{Q}_c	true Q_c , taking asymmetry into account
\tilde{Q}_0	true Q_0 , taking asymmetry into account
$\tilde{\omega}_0$	true ω_0 , taking asymmetry into account

appendix B

A_g, A_e	amplitude of two Rabi oscillations
P_g, P_e, P_f	probability to find qubit in $ g\rangle$, $ e\rangle$, or $ f\rangle$ state
R_g, R_e, R_f	readout signal amplitude for $ g\rangle$, $ e\rangle$, and $ f\rangle$
$R_e^A, R_e^B, R_g^B, R_g^A$	extrema of two Rabi oscillations
θ	pulse rotation angle

appendix C

F	ground state preparation fidelity
-----	-----------------------------------

α_C	cavity coherence state amplitude ($\bar{n} = \alpha_C ^2$)
Ω_R	Rabi oscillation frequency

Introduction

It is now possible to fully control individual systems whose degrees of freedom obey the laws of quantum mechanics. Rydberg atoms [1, 2], ions held in traps [3], electron spins in quantum dots [4], nitrogen vacancies in diamond [5], and superconducting circuits [6, 7, 8, 9, 10, 11], are the main implementations. The creation and control of these devices with the goal of practical quantum computation comprises the field of “quantum engineering”. Perhaps the most fundamental and crucial requirement is that these systems must be dissipationless. Many experiments outlining the progress of the field are stimulated through the reduction of dissipation; thus the search for remaining dissipation is at the heart of quantum engineering.

The requirements to do useful computations with these systems were laid out by David DiVincenzo [12] in the year 2000, and the so-called DiVincenzo criteria are in summary: (1) a scalable physical system with well-characterized qubits, (2) the ability to initialize the qubits, (3) coherence times much longer than gate operation times, (4) a universal set of gates, and (5) a qubit-specific measurement. Additional requirements for quantum engineering systems like reproducible components, simple fabrication, and cost of operation would also be welcomed, but are not fundamentally necessary for basic physics experiments.

The world of superconducting qubits has evolved since the 1980s to become a major contender in the race to build a universal quantum computer. Superconducting qubits are simply electronic circuits made of mostly aluminum that express their quantum nature when they are cooled to low temperatures, as described below. Since these devices are created using common machine-shop equipment and wafer fabrication technologies, new designs with unique features can be quickly produced, while a scalable architecture remains a possibility. The major downside when comparing superconducting qubits to other systems has been concerning DiVincenzo criteria 3: the limited coherence times.

This thesis explores how to improve coherence of superconducting qubits and resonators for use in quantum information systems through the analysis technique of participation ratios. The method of participation ratios describes how different loss mechanisms participating in a resonant mode combine to yield an overall quality factor. An experiment on adjusting the geometry of compact resonators and compact transmons confirms the validity of the participation ratio argument through the display of improved quality factors and relaxation times up to $13 \mu\text{s}$. The idea of participation ratios is then pushed to its limit in the 3D transmon, with the drastically simplified mode structure and reduction of surface participation, resulting in coherence times of over $100 \mu\text{s}$. The notion of participation ratios is then extended with the fluxonium, which borrows the simplified geometry of the 3D transmon combined with a decreased sensitivity to dissipation, resulting in relaxation times of over 1 ms. In conclusion, the vast improvements of superconducting qubit and cavity coherence in the last five years enables experiments that were previously impossible, including an experiment described in this thesis, of a qubit reset on a 3D transmon. Thus the field of quantum information with superconducting qubits is now advancing on other fronts, a reflection that practical quantum computation is one step closer

to reality.

1.1 Quantum Computing

1.1.1 What is a Qubit?

All of the computing we do today with any type of modern-day computer is in fact *classical* computing. This computation scheme relies on operations involving bits, physical systems that can be in 2 states, labeled 0 and 1. In a computer, all storage and calculations (ROM, RAM, CPU, GPU, etc...) use registers of bits as the information medium. This computation scheme has limitations in how fast it can solve certain problems. Many problems require either that the size of the data to be processed or time needed to run the algorithm that scale exponentially. These problems become intractable and practically impossible for large enough sizes or long enough times.

In a quantum computer, the analog of the classical bit is called a “qubit”: defined as a pair of levels of any quantum system, called the ground state ($|g\rangle$ or $|0\rangle$) and the excited state ($|e\rangle$ or $|1\rangle$). Typically, the notations $|g\rangle$ and $|e\rangle$ are used to refer to a specific qubit design, and $|0\rangle$ and $|1\rangle$ are used in the abstract quantum information sense. These two levels may be two excited states of atoms, spin-up and spin-down of nuclear or electronic spins, two positions of a crystalline defect, two states of a quantum dot, or two energy levels of a superconducting circuit. Since the system is of a quantum nature, any combination of the two states are allowed, expressed as $|\Psi\rangle = \alpha|g\rangle + \beta|e\rangle$. Thus two complex numbers α_g and α_e define the qubit state, with normalization condition $|\alpha_g|^2 + |\alpha_e|^2 = 1$. This qubit state is commonly visualized using the Bloch sphere, as shown in Fig. 1.1, where $\alpha_g = \cos(\theta_B/2)$ and




Figure 1.1: Any pure state can be represented as a point on the surface of the Bloch sphere, with $|g\rangle$ and $|e\rangle$ as the poles. Mixed states, or states with incomplete knowledge, are represented as points inside the sphere.

$\alpha_e = e^{i\phi_B} \sin(\theta_B/2)$. This notation places $|g\rangle$ and $|e\rangle$ at the poles, while any other pure state can be represented by a point on the surface of the sphere. In addition, any point inside the sphere represents a mixed state, in which knowledge is incomplete, with the completely unknown qubit state in the center.

1.1.2 Quantum Resources

Quantum computing is a new paradigm of information processing that exploits the quantum nature of its components, namely the concepts of *superposition* and *entanglement*, to perform calculations. Quantum computation is thought to speed up the calculation of *some* classically “hard” problems, for example: prime factorization with Shor’s algorithm [13] and database searching with Grover’s algorithm [14, 15]. The effect of quantum speedup for Grover’s algorithm has been observed by Dewes

et al. [16] in 2012.

Superposition allows a degree of freedom to be in multiple energy eigenstates *at the same time*. For a qubit, this means that the qubit may be in both the ground state and excited state at the same time, with a notation corresponding to a sum $(|0\rangle + |1\rangle)/\sqrt{2}$. Superposition may manifest itself in physical systems as an atom being excited and not, the photon number of a resonator being zero and one, as the charge of a piece of metal being both positive and negative, or as current in a ring traveling both clockwise and counter-clockwise.

The other resource, entanglement, is also unlike anything in the classical realm. Entanglement, in the language of qubits, means that one bit of information can be shared by 2 or more different degrees of freedom, regardless of their physical separation. For example, the two qubit state $|00\rangle + |11\rangle$ is entangled, meaning that information is not stored in either qubit, but in the non-local correlation between the qubits. In this case, one non-local bit says that the qubits are in the same state and another non-local bit gives the sign of the entanglement (in this case, a “+”). A measurement on one qubit of result 0 instantly implies that the other qubit will be found in state 0. The qubits, once entangled, if they are left undisturbed, can be separated by any distance and their entanglement will remain. Entanglement allows for operations on one qubit to affect all entangled qubits, and thus is a resource that can be exploited.

Superposition and entanglement allow for systems to explore a wide configuration space all at once, a property often called “quantum parallelism”. Thus, a single operation on a system with n qubits is, in some sense, *like* doing operations on 2^n numbers in a classical computer; quantum parallelism is exponential in the number of qubits.

1.1.3 Qubit Coherence

While a qubit may start in a specific superposition of $|g\rangle$ and $|e\rangle$, it will not stay there forever. In classical computing, we rely on bits keeping their information, i.e. remaining in the state they have been set to, for long periods of time, for example in hard-drives or CD-ROMs. Qubits are not yet so robust, and in fact “lose” their state fairly quickly; the length of time the qubit retains its quantum information is called the coherence time. Since one of DiVincenzo’s criteria [12] is that coherence times must be much longer than gate operation times, it is important to understand and to raise this time as much as possible.

While coherence is a broad term often meant to include several types of information loss, the “coherence time” (denoted T_2) is the length of time for which the qubit maintains a given superposition of states. There are two processes that contribute to decoherence: relaxation processes (T_1) and dephasing processes (T_ϕ) [9, 11]. T_1 and T_ϕ combine to yield $1/T_2 = 1/2T_1 + 1/T_\phi$; if T_ϕ is made infinite, then $T_2 = 2T_1$. Relaxation processes prematurely bring the qubit from the excited state to the ground state (at rate Γ_\downarrow). If the qubit is in contact with a bath at nonzero temperature, there will also be excitation processes Γ_\uparrow , leading to a combined $T_1^{-1} = \Gamma_\downarrow + \Gamma_\uparrow$. In contrast, dephasing processes result in a change in phase. In the Bloch sphere notation, T_1 processes represent the state moving in the direction of the poles (diffusion along a meridian), and T_ϕ processes involve the state diffusing along a parallel of the sphere. In any physical system, there are many channels that may contribute each to T_1 and T_ϕ ; the purpose of this thesis is to identify and remove these channels in superconducting resonators and qubits.

Since relaxation processes involve the qubit decaying to the ground state, this implies a loss of energy. Thus, any microscopic or macroscopic effect that re-equilibrates

energy from the qubit causes relaxation. Dephasing, on the other hand, involves the qubit state diffusing along a parallel, thus not requiring a loss of energy. If the qubit frequency drifts, then the qubit state will rotate around the axis between the ground and excited state at the difference frequency. This rotation, when averaged over many instances, leads to dephasing; thus any effect that causes the qubit frequency to change over time is a source of dephasing.

While the coherence time is often loosely denoted as T_2 , there are two measurement methods, borrowed from the field of Nuclear Magnetic Resonance (NMR) that yield different results [17]. The first is the “Ramsey” time, T_2^R , which is measured by varying the time between two $\pi/2$ pulses. This measurement is sensitive to noise of all frequencies and is the more stringent of the two. The second is the “Echo” time, T_2^E , which is measured in the same way except a π -pulse is inserted half-way between the $\pi/2$ pulses. The π pulse reverses the direction of dephasing for the second half of the waiting time, thus it “echoes” out slow drifts in the qubit frequency. For this reason, the echo measurement is insensitive to low-frequency noise. Comparison between the two measurements in general gives some information about the spectral density of the dephasing noise source; more refined pulse methods (CPMG) can precisely access this spectral density [18, 19].

1.1.4 Qubit Initialization

Another one of DiVincenzo’s criteria is the ability to initialize qubits; another important feature of qubits to discuss. Initializing a qubit, sometimes also called resetting a qubit, means to put it into a known state regardless of its current state, with a fidelity defined by overlap between the created state and the desired state. In general, this target state can be anything, but it is often the ground state. While relaxation processes tend to prepare $|g\rangle$, there are also excitation processes that tend to pre-

pare $|e\rangle$. These excitation processes may be due to either thermal excitations or some other channel that gives energy to the qubit. In either case, these excitations, along with the given relaxation rate, will yield a steady-state excited state population (P_e), which can be expressed as an effective temperature.

This effective temperature T can be calculated using the Boltzmann factor, $P_e/(1-P_e) = \exp(hf_{ge}/k_B T)$, where f_{ge} is the qubit frequency. For temperatures high enough, not only the first excited state, but the second or third excited states may have considerable steady state population. Then the temperature can be calculated by fitting all of the level populations to that expected for a Boltzmann distribution.

Because of these possible excitations, it must be possible to reset the qubit to a known state before a given algorithm may begin. Our method for qubit reset, called DDROP is described theoretically in Appendix C and experimentally in Chapter 6.

1.1.5 Circuit QED

Nearly all current superconducting qubit systems utilize to some degree an analog of a system that is used to study light-matter interactions, called cavity quantum electrodynamics, or “cavity QED” [20, 21]. This experimental technique has been developed as a means of increasing the coupling between photons and atoms. Cavity QED involves passing atoms through a cavity between two mirrors, coupling the atoms to the quantum-mechanical electric field of the cavity. Due to the coupling, the atoms can get entangled with the photons in the cavity. Serge Haroche’s 2012 Nobel Prize in Physics was in part attributed for the development of Cavity QED which allowed his team to manipulate single Rydberg atoms interacting with a superconducting cavity. The Hamiltonian that describes the cavity QED system is the Jaynes-Cummings Hamiltonian [22], which depicts the atom (qubit) as a spin and the cavity as a quantum-mechanical electric field, as shown in Eq. 1.1.

$$H_{JC} = \hbar\omega_R \left(a^\dagger a + \frac{1}{2} \right) + \frac{\hbar\omega_Q}{2} \sigma_z + \hbar g (a^\dagger \sigma^- + a \sigma^+) \quad (1.1)$$

Circuit QED [8, 23], a circuit analog to cavity QED, has become one of the most common paradigms of superconducting qubit experiments. Circuit QED involves coupling artificial atoms to superconducting cavities as opposed to real atoms in cavity resonators, yet such systems can be analyzed and discussed in the same manner, with all of the same equations and terminology of cavity QED. Coupling the qubit only to a cavity and not directly to a microwave line protects the qubit from noise outside the resonator bandwidth. While coupling qubits to cavities has been done as early as the 1980s [24], utilizing photons in the cavity as a probe of the qubit state was developed around 2005 at Yale. This latter development is what has become known as circuit QED. Qubit readout is performed by passing photons through the cavity, which inherit a different transmission amplitude or phase depending on the qubit state. These artificial atoms have some advantages over natural atoms; perhaps the most important advantage is that the parameters the atom can be *engineered* by design. This flexibility is expressed in the variety of superconducting qubits that have been tested, a few of which are discussed in the following section.

1.2 Superconducting Qubit Species

Superconducting circuits present controllable quantum degrees of freedom, and therefore lend themselves to the design of artificial atoms. Since it is possible to design circuits that are described by a wide variety of Hamiltonians, superconducting circuits are very attractive as a quantum computing platform. The simplest circuit consists of a superconducting inductance and capacitance. This so-called LC oscillator is equivalent to a dissipationless spring-mass system, with the voltage on the

capacitor (or current in the inductor) varying sinusoidally in time, and a parabolic energy potential as a function of the flux through the inductor, $\phi = LI$, the analog of the “position” coordinate. In the spring-mass analogy, the mass is represented by the capacitance and the spring constant by the inverse of the inductance. The Hamiltonian of this oscillator along with that of the spring-mass system is shown in Eq. 1.2, showing that the canonically conjugate pair of flux (ϕ) and charge on the capacitor (Q) are analogous to position and momentum, respectively. When this oscillator is cooled to a temperature corresponding to an energy less than one photon at the resonant frequency ($T < \hbar\omega_0/k_B$, where $\omega_0 = 1/\sqrt{LC}$), it can express its quantum nature. The harmonic oscillator has equally spaced energy levels $E_n = \hbar\omega_0(n + 1/2)$, with minimum energy $E_0 = \hbar\omega_0/2$ and spacing $\hbar\omega_0$. Going from one energy level to the next corresponds to adding a quantum of energy in the oscillator. In these microwave LC oscillators, these quanta are called (standing) “photons” since they are quantized wavepackets of electromagnetic radiation even if they do not move through space.

$$\begin{aligned}
 (LC \text{ oscillator}) \quad H(\phi, Q) &= \frac{Q^2}{2C} + \frac{\phi^2}{2L} \\
 (spring - mass) \quad H(X, P) &= \frac{P^2}{2M} + \frac{kX^2}{2}
 \end{aligned}
 \tag{1.2}$$

The problem with using the first two levels of a harmonic oscillator as a qubit is that the levels are not individually addressable. It is impossible to move population between the two levels without also exciting many other levels, since they all share the same transition frequency. One solution is to modify the potential so that the energy levels are not equally spaced, and therefore any two levels may be addressed individually. To make an anharmonic oscillator, either the linear inductance or linear capacitance must be replaced by a nonlinear element. The only nonlinear circuit element without dissipation is the Josephson tunnel junction [25], a nonlinear induc-




Figure 1.2: (a) Simplified geometric aspect of a Josephson tunnel junction with two superconducting (S) electrodes separated by a thin insulator (I). (b) Circuit representation of a Josephson tunnel junction, and (c) equivalent circuit representation, consisting of a Josephson element and a capacitive element. The Josephson element is responsible for the nonlinear inductance. Each of these schematics doubles as a representation of the Cooper Pair Box qubit, which is composed of a superconducting island separated from ground by only a Josephson tunnel junction.

tance. A Josephson tunnel junction consists of two superconductors separated by a thin insulator as shown in Fig. 1.2; the most commonly used materials in superconducting qubits is Aluminum/Aluminum oxide/Aluminum. Since the Josephson tunnel junction resembles a parallel plate capacitor, the junction must be thought of as a Josephson element in parallel with a capacitor, in which the Josephson element models the tunneling of Cooper pairs.

To explain the effect of a Josephson tunnel junction in terms of its inductance, one must compare the constitutive equations relating the current through the inductor to the branch flux, a generalization of the magnetic flux that applies to all circuit elements. The branch flux of any given element is given by $\Phi(t) = \int_{-\infty}^t V(t_1) dt_1$, where $V(t)$ is the integral of the electric field along a current line inside the element. In the limit of lumped elements, any current line yields the same voltage. For the normal, linear inductor, the current and flux are linearly related:

$$I(t) = \frac{1}{L} \Phi(t) \quad (1.3)$$




Figure 1.3: Constitutive relation (Eq. 1.4) plotted as current I as a function of branch flux Φ . The dashed line is the expected relation for a linear inductor. Note that the sinusoidal dependence approximates a linear dependence in the limit of low Φ .




Figure 1.4: Circuit schematic of the Cooper Pair Box qubit, consisting of an isolated island electrode, highlighted in blue. The island is separated from the ground by a capacitance and a Josephson tunnel junction. The applied voltage is to charge bias the island as a means of suppressing the fluctuation of offset charges on the superconducting island.

In contrast, the Josephson element has a sinusoidal constitutive relation. Note that in the limit of low branch flux, the sinusoidal relationship approximates a linear inductor ($\sin(x) \sim x$), as shown in Fig. 1.3.

$$I(t) = I_0 \sin [2\pi\Phi(t)/\Phi_0] \tag{1.4}$$

Thus, a Josephson tunnel junction, being composed of a nonlinear inductor in parallel with a capacitor, is an anharmonic oscillator just by itself. This type of qubit

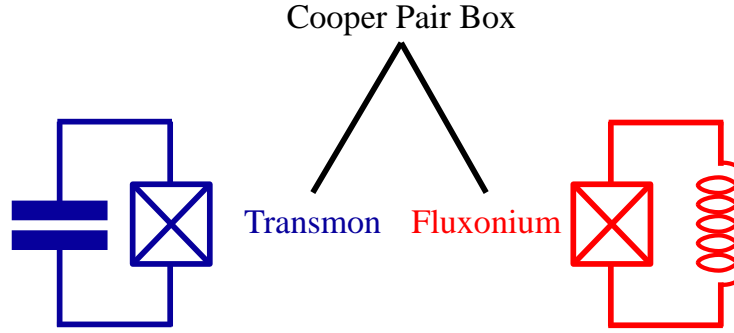


Figure 1.5: The fluxonium and transmon qubits are both “descendants” of the Cooper Pair Box. They each solve the problem of offset charge noise by shunting the Josephson junction by either an external capacitance or inductance.

is called the Cooper Pair Box (CPB) [6] when it is charge-biased. The capacitive energy of the qubit is given by a term proportional to $(\hat{n} - n_g)^2$, with \hat{n} , the integer number of charges on the superconducting island (highlighted by the box in Fig. 1.4) which determines the qubit state, and n_g , the offset charge induced by the voltage applied to the capacitance. One of the main limitations of this qubit was that when the offset charge drifts, the transition frequency changes significantly. Thus, qubit coherence was limited by these fluctuations of charge. There are two alternate qubit designs used in this thesis, based on the CPB, that were developed to eliminate the influence of charge fluctuations: the transmon and the fluxonium. Other superconducting qubit types include the Quantronium [26], the phase qubit [7], and the flux qubit [27, 28]. At the time of writing, these qubit designs are less popular and less coherent.

1.2.1 Transmon

The transmon is a superconducting qubit consisting of an anharmonic oscillator composed of a Josephson tunnel junction in parallel with a large external capacitance, as shown in Fig. 1.6. In order to distinguish the CPB and the transmon, it is im-

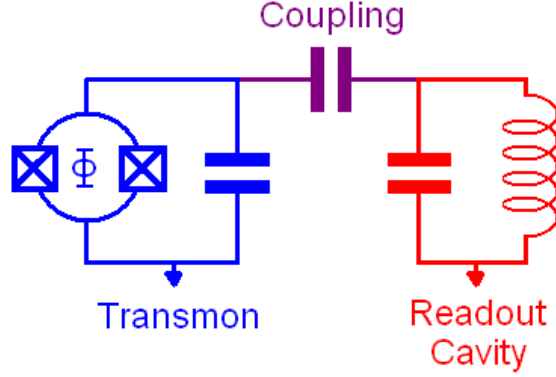


Figure 1.6: Circuit schematic of a transmon qubit capacitively coupled to a linear readout resonator. The only difference between the transmon and the CPB is the size of the capacitive energy relative to the junction energy.

portant to consider the Josephson energy E_J and the charging energy E_C . E_J is a parameter characterizing the junction inductance, and is defined as $E_J = \phi_0 I_0$, where $\phi_0 = \Phi_0/2\pi = \hbar/2e$ and I_0 is the critical current of the junction. E_C is the energy of the capacitance $E_C = e^2/2C_\Sigma$, where C_Σ is the sum of the junction capacitance C_J and the external capacitance C_{ext} . This figure also highlights the optional replacement of the Josephson tunnel junction with a SQUID loop consisting of two junctions in parallel, enabling a tunable qubit frequency via an applied magnetic flux Φ_{ext} , through $E_J(\Phi_{ext}) = E_J \cos(\pi\Phi_{ext}/\Phi_0)$. The transmon is distinguished from the CPB by the size of the capacitance; the CPB's capacitance comes solely from the junction. The purpose of the larger capacitance is to protect the qubit from fluctuations of offset charge of the superconducting island by flattening the transition frequency dependence on offset charge [29, 30, 31]. The transmon is a good example of how quantum circuit design can tailor a circuit to design a specific Hamiltonian to boost qubit performance.

The larger capacitance largely influences two qubit parameters, the anharmonicity and the charge dispersion. The anharmonicity is the difference in energy between

successive excitations of the system, denoted as $\alpha = (E_f - E_e) - (E_e - E_g)$; a harmonic oscillator has an anharmonicity of zero. While a large anharmonicity may not be fully exploitable, a small anharmonicity requires long qubit pulses in order to selectively excite only one pair of levels. Charge dispersion is the amount that the energy levels are dependent on the offset charge ($n_g = C_g V_g / 2e$) on the island making up the qubit. The CPB energy levels are so dependent on charge such that this charge must be kept constant using a gate voltage. The remarkable advantage of the transmon comes from the prediction that as the qubit parameter E_J/E_C is increased, the charge dispersion reduces *exponentially* and the anharmonicity decreases only with a power law. Because of this, with practical parameters, the charge dispersion can be made negligibly small while maintaining enough anharmonicity. The energy level of the m^{th} level [29] is shown in Eq. 1.5 where $a_\nu(n_g)$ is Mathieu's characteristic value and $k(m, n_g)$ is a function to sort the eigenvalues. From these energies, the anharmonicity and charge dispersion can be calculated. The charge dispersion is shown in Eq. 1.6, and in the limit of large E_J/E_C , $\alpha \sim E_C$.

$$E_m(n_g) = E_C a_{2[n_g + k(m, n_g)]}(-E_J/2E_C) \quad (1.5)$$

$$\varepsilon_m \simeq (-1)^m E_C \frac{2^{4m+5}}{m!} \sqrt{\frac{2}{\pi}} \left(\frac{E_J}{2E_C} \right)^{\frac{m}{2} + \frac{3}{4}} e^{-\sqrt{8E_J/E_C}} \quad (1.6)$$

The charge dispersion can be best visualized by plotting the transmon energy levels as a function of gate charge n_g . Fig. 1.7 shows that the first few energy levels of the CPB ($E_J/E_C \sim 1$) are highly dependent on charge. Transition energies (or frequencies) are represented by the differences between successive energy levels. Note that in this plot, the lowest CPB transition frequencies varies between roughly 3 and 10 GHz. There are sweet spots at values $n_g = 1/2 + n$ (where n is an integer),

where the first order charge noise goes to zero (see Section 1.3.2). As E_J/E_C is increased (via increasing the capacitance in parallel, thus reducing E_C), the energy levels become flatter. In the transmon limit, with $E_J/E_C \geq 50$, the energy spectrum of the first few energy levels are practically insensitive to charge. The remaining charge dependence is so small that reasonable parameters predict dephasing times due to charge noise on the order of seconds.

The transmon qubit was developed at Yale [29, 30, 31] between 2007 and 2009. Since then, the transmon has been adopted by many research groups around the world, and is, at the time of writing, the most commonly used species of superconducting qubit.

1.2.2 Fluxonium

The fluxonium qubit, on the other hand, solves the CPB's problem of charge noise sensitivity by *shunting* the Josephson tunnel junction by a large inductance, as shown in Fig. 1.8. A shunting inductor connects the two ends of the Josephson tunnel junction, thus connecting any previously isolated pieces of superconductor. The addition of this inductor destroys the charge quantization condition enforced on a superconducting island [32]. However, as this inductance approaches infinity, the properties of the circuit should match that as the CPB (which is effectively shunted by an infinite inductance). Thus, this so-called fluxonium qubit is *like* a CPB whose low frequency charge fluctuations are suppressed by the added inductor.

A derivation of the fluxonium energy level spectrum is given in Section 2.2.2. The fluxonium qubit was developed at Yale [33, 34, 32, 35] around the same time the transmon was developed. The original fluxonium qubits were fabricated and tested by Vladimir Manucharyan, with subsequent improvements made by Nicholas Masluk and Archana Kamal [36].




Figure 1.7: First few energy levels of qubit consisting of a Josephson tunnel junction in parallel to a capacitance plotted as a function of gate charge n_g . The dependence on gate charge differs significantly as the ratio of Josephson energy to charging energy (E_J/E_C) is tuned from the CPB regime ($E_J/E_C \sim 1$) to the transmon regime ($E_J/E_C \geq 50$). While the CPB suffers greatly from charge noise as a result of the large transition frequency dependence on charge, the transmon can be made arbitrarily insensitive to charge fluctuations.

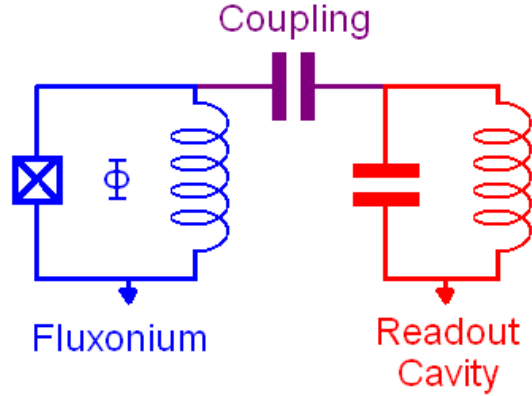


Figure 1.8: Circuit schematic of a fluxonium qubit capacitively coupled to a linear readout resonator. Note the lack of isolated islands in the fluxonium structure.

1.3 Coherence Limitations

There are many channels through which a qubit may relax or dephase, ultimately leading to decoherence. These channels each have varying strengths depending on the qubit species, parameters, and experimental setup. The most common sources of relaxation and dephasing for superconducting qubits will be introduced in the following sections.

1.3.1 Relaxation Sources

Participation Ratios

There are many sources of relaxation in every qubit; some of which may be orders of magnitude stronger than others. In general, any material through which an electric field is set or a current flows as part of the qubit mode, contributes to the relaxation rate of the qubit. The loss of each material is defined in terms of a “loss tangent” ($\tan\delta$), or inversely as a quality factor Q . The loss tangent of a dielectric insulating material with permittivity of $\varepsilon = \varepsilon' - j\varepsilon''$ is defined as $\tan\delta = \varepsilon''/\varepsilon'$. The quality factor for a given resonance is defined as 2π times the ratio of energy stored to energy

dissipated per cycle. While loss mechanisms act in parallel, these quality factors do not add directly; they add to the relaxation rate proportionally to their contribution in the qubit mode. This is better explained in terms of “participation ratios” [37].

The overall qubit mode has a quality factor $Q_1 = \omega_{ge}T_1$. This quality factor may be thought of as having contributions from electric field (capacitive loss), current (inductive loss), and radiative loss, which add to yield the overall Q_1 as shown in Eq. 1.7. Each of these groups themselves is made up of contributions of each material in that group. For example, capacitive loss is due to in part to each material in which electric field energy resides. Also, inductive loss is due in part to each material through which current travels. Since relaxation of the qubit implies energy loss, the energy that “leaves” the qubit by radiation to the outside instead of through dissipation of a material introduces a third term Q_{rad} ; in cases where the qubit is coupled to a cavity, this term is often called Purcell relaxation.

$$\frac{1}{\omega_{ge}T_1} = \frac{1}{Q_1} = \frac{1}{Q_{cap}} + \frac{1}{Q_{ind}} + \frac{1}{Q_{rad}} \quad (1.7)$$

Capacitive and inductive losses may be broken down into their constituent components via Eqs. 1.8. Here, the contribution from each material or component is weighted by the fraction of energy stored. This fraction of energy stored in each component is called the participation ratio, under the constraint $\sum p_{cap}^i = 1$ and $\sum p_{ind}^i = 1$ signifying that the energy of a given capacitance or inductance is stored somewhere. While some qubit or resonator designs may allow the participation ratios to be estimated from rough calculations, in practice full high frequency simulations are necessary to properly estimate participation ratios.

$$\begin{aligned} \frac{1}{Q_{cap}} &= \frac{p_{cap}^a}{Q_{cap}^a} + \frac{p_{cap}^b}{Q_{cap}^b} + \dots + \frac{p_{cap}^n}{Q_{cap}^n} \\ \frac{1}{Q_{ind}} &= \frac{p_{ind}^a}{Q_{ind}^a} + \frac{p_{ind}^b}{Q_{ind}^b} + \dots + \frac{p_{ind}^n}{Q_{ind}^n} \end{aligned} \quad (1.8)$$

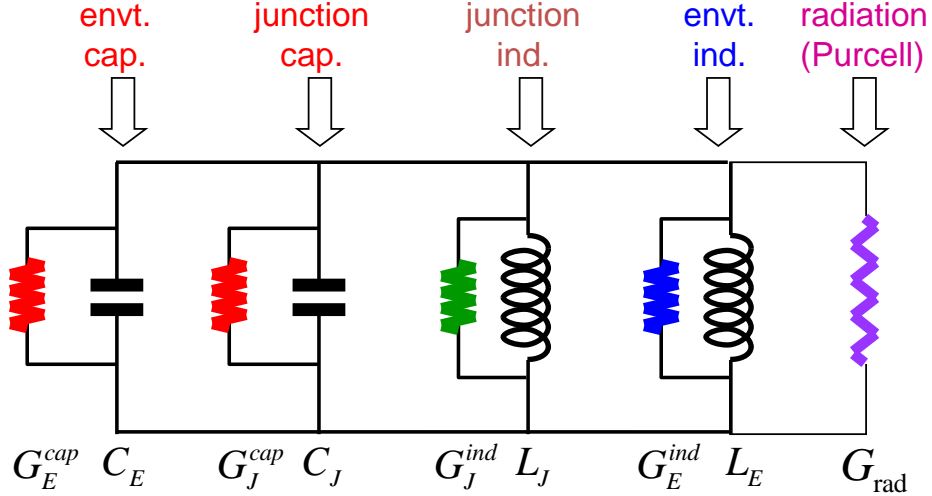


Figure 1.9: Circuit schematic for generic qubit, including junction capacitance (C_J), environment capacitance (C_E), junction inductance (L_J), environment inductance (L_E), and the Purcell effect. Each element comes with an attached conductance that represents the loss associated with that component.

Another way to think about these various loss mechanisms is to consider a resonant circuit representation in which each element comes with a resistor in parallel representing that component's loss. The resistors in parallel add to yield the total loss. The participation ratio, in this view, leads to different currents in different components; if there is a component with no current, its resistance does not count towards the total. Fig. 1.9 shows the circuit breakdown for a generic qubit, including the junction inductance and capacitance, the environment capacitance and inductance, and the Purcell effect. This circuit is used for the fluxonium qubit loss considerations in Chapter 7.

There are thus two ways to raise the overall quality factor of a given system: (1) raise the quality factor of the components (Q_{cap}^i , Q_{ind}^i , etc...) by changing materials or changing fabrication techniques or (2) reduce the participation ratios of the lowest quality components by changing the qubit design. However, it is often difficult to know which component or loss mechanism is limiting any qubit. Information from

the frequency dependence of T_1 or the measurements of several devices can be used to make educated guesses, as is done in Chapter 4 and Chapter 7.

In general, this formalism enforces a few rules that are important to remember. While it may be tempting to learn about loss by minimizing the participation of suspected components, it is the more difficult route to explore. If the suspected component were indeed limiting T_1 , then an increase in T_1 will be observed, but only until some other component is dominating. Since an observed increase may be small, it will be difficult to claim an improvement. Alternatively, one may purposefully increase the participation of each component in turn in order to accurately place bounds on the quality of each component. A measured value of T_1 implies that the quality factor of each component $Q_x > p_x Q_1$. For example, if the participation of a component is known to be 10% and the measured $Q_1 = 10,000$, that means that no matter what other components add to the loss, the Q of that component is at least 1,000, but possibly much higher. This means that if this component's Q were lower than 1,000, the measured Q_1 would be lower, *but* if it were higher than 1,000, the measured Q_1 would either increase (if Q_1 was limited by this component), or stay largely unchanged (if Q_1 was limited by some other component). In the limit of increasingly low participation, the bound becomes irrelevant. In the case where other losses are negligible, then the Q of that component is exactly 1,000. Thus, the higher the participation of a given component, the more stringent (and relevant) the bound will be; in the limit of 100% participation, the bound is more of a direct measurement.

Purcell Effect

The Purcell effect [38] has been well documented as a contribution to qubit relaxation [39, 40], and in the case of qubits coupled to resonators can be considered the main

contribution to the term Q_{rad} . This qubit relaxation mechanism arises when the qubit is coupled to a cavity that is coupled to a microwave line for readout and drive purposes. Since the qubit and cavity are coupled, with a given set of parameters, the dressed “qubit” mode partially inherits the properties of the undressed cavity mode, including its Q factor. In other words, since the cavity mode is coupled to a microwave line for measurement and stimulation, this coupling is passed on to the qubit. The strength of this coupling is controlled by the qubit-cavity coupling strength, the qubit-cavity detuning, and the cavity coupling quality factor. Assuming the qubit is only coupled to a single cavity mode, and assuming the dispersive regime of circuit QED (see Section 2.2.1), the Purcell limit of T_1 is calculated as $T_1 = \Delta^2/\kappa g^2$ where Δ is the detuning between bare cavity and qubit frequencies, and κ is the linewidth of the cavity mode. However, this simplified expression was shown not to be a good approximation [31], and influence from higher cavity modes cannot be ignored. The full expression for transmon qubits is shown in Eq. 1.9, where C_Σ is the capacitance of the transmon and $Y(\omega_{ge})$ is the admittance as seen by the qubit. Full 3D microwave simulations are the best way to obtain this admittance.

$$T_1^{purcell} = \frac{C_\Sigma}{\text{Re}[Y(\omega_{ge})]} \quad (1.9)$$

TLS Dissipation

An example of a capacitive material loss is dissipation to two level systems (TLS). Dielectric materials are often thought to have various defects both in the bulk and especially on their surface. Certain defects may have two states (thus, a TLS) with an difference in energy close to the qubit excitation energy. In this case, the qubit may relax while exciting a defect TLS. Eventually the TLS may dissipate the energy in the form of phonons into the substrate, thus preparing the TLS to absorb more energy

from the qubit. This loss mechanism and its smoking-gun features are discussed in detail in Chapter 4.

1.3.2 Dephasing Sources

In a manner similar to relaxation, there are many possible sources of dephasing. Dephasing from drifts in the qubit frequency over time result in the qubit state diffusing around the Bloch sphere instead of remaining steady. Any microscopic or macroscopic process causing the qubit frequency to drift or fluctuate is a source of dephasing; a few of these sources are discussed below. Possible sources that are not included here are dephasing due to quasiparticle tunneling [41] and dephasing due to E_C variations [29].

Photon Induced Dephasing

As derived in Section 2.2.1 and discussed in Section 3.2.1, coupling a qubit to a cavity introduces a term that makes the qubit frequency depend on the number of excitations in the cavity: $\omega_Q \rightarrow \omega_Q + \chi N_C$. Thus, it is clear that if the number of photons in the cavity fluctuates, then so does the qubit frequency. In the limit of χ larger than the width of the qubit peak itself and assuming an average number of photons $\bar{n} \gtrsim 0$, a change of one photon in the cavity with width κ completely dephases the qubit, resulting in a calculated dephasing rate [42] of:

$$T_2 = (\bar{n}\kappa)^{-1} \quad (1.10)$$

The full expression [43] for any χ and \bar{n} is shown below.

$$\Gamma = \frac{\kappa}{2} \text{Re} \left[\sqrt{\left(1 + \frac{2i\chi}{\kappa}\right)^2 + \left(\frac{8i\chi\bar{n}}{\kappa}\right)} - 1 \right] \quad (1.11)$$

Offset Charge Noise

Another contribution to dephasing comes from offset charge fluctuations. For example, fluctuations of the offset charge on the island of the CPB induce frequency fluctuations. The sensitivity of a given qubit design to charge noise is characterized by the slope $\partial E_{ge}/\partial n_g$ where $E_{ge} = \omega_{ge}/\hbar$ is the energy difference between the ground and excited state, and n_g is the “gate charge”, the background offset charge of a given superconducting island. A sweet spot for charge noise occurs when this slope goes to zero, meaning small charge fluctuations do not shift the qubit frequency to first order, thus leaving only second order fluctuations: $\delta\omega \propto (\partial^2\omega_{ge}/\partial n_g^2)\delta n_g^2$. Dephasing times can be calculated in general for a given slope and for the specific case of a transmon ($E_J/E_C \gg 1$) as shown in Eq. 1.12 [29]. Here A is a term denoting the amplitude of charge fluctuations, and may be of the order 10^{-3} or $10^{-4} e$ [44]. Section 1.2.1 discusses the difference in charge noise between the CPB and transmon, and is useful for understanding charge noise in general.

$$\begin{aligned}
 & \text{(general)} & T_2 & \sim \frac{\hbar}{A} \left| \frac{\partial E_{ge}}{\partial n_g} \right|^2 \\
 & \text{(transmon)} & T_2 & \sim \frac{\hbar}{A\pi} \left| -E_C 2^9 \sqrt{\frac{2}{\pi}} \left(\frac{E_J}{2E_C} \right)^{\frac{5}{4}} e^{-\sqrt{8E_J/E_C}} \right|^{-1}
 \end{aligned} \tag{1.12}$$

Flux Noise

Flux noise is another possible dephasing source, and results from flux fluctuations through a superconducting loop that influences the qubit frequency. The source of these fluctuations may be either from microscopic spins on the device surface or fluctuations of an externally applied field, for example. The sensitivity of a given qubit design to flux noise is characterized by the slope $\partial E_{ge}/\partial \Phi$ where Φ is the flux through the loop under consideration. A sweet spot for flux noise occurs when this

slope goes to zero, meaning small flux fluctuations do not shift the qubit frequency to first order, thus leaving only second-order fluctuations: $\delta\omega \sim \frac{1}{2}(\partial^2\omega_{ge}/\partial\Phi^2)\delta\Phi^2$. Dephasing times can be calculated in general for a given slope and for the specific case of a transmon as shown in Eq. 1.13 [29]. Here A is a term denoting the amplitude of flux fluctuations, and may be of the order of 10^{-5} or $10^{-6} \Phi_0$ [45].

$$\begin{aligned} \text{(general)} \quad T_2 &\sim \frac{\hbar}{A} \left| \frac{\partial E_{ge}}{\partial \Phi} \right|^{-1} \\ \text{(transmon)} \quad T_2 &\sim \frac{\hbar}{A} \frac{\Phi_0}{\pi} \left(2E_C E_J \left| \sin \frac{\pi\Phi}{\Phi_0} \tan \frac{\pi\Phi}{\Phi_0} \right| \right)^{-1/2} \end{aligned} \quad (1.13)$$

Critical Current Noise

Yet another possible source of dephasing is critical current noise. One of the primary parameters of a superconducting qubit is the Josephson energy E_J . This term is calculated using the critical current of the Josephson tunnel junction $E_J = I_0\hbar/2e$, and is physically derived from the thickness and area of the insulating layer. Fluctuations in I_0 can be converted into fluctuations in E_J , which in turn can be converted to fluctuations in qubit frequency, in the case of the transmon via $\omega_{ge} = \sqrt{8E_J E_C}$. The source of critical current *noise* in a Josephson tunnel junction is often attributed to a fluctuation in the number of transmission channels, possibly via defects fluctuating between two states. The sensitivity of a given qubit design to critical current noise is characterized by the slope $\partial E_{ge}/\partial I_0$, where I_0 is the critical current of the junction under consideration. Dephasing times can be calculated in general for a given slope and for the specific case of a transmon as shown in Eq. 1.14 [29]. Here A is a term denoting the amplitude of critical current fluctuations, and it may be of the order 10^{-6} or $10^{-7} I_0$ [46].

$$\begin{aligned}
\text{(general)} \quad T_2 &\sim \frac{\hbar}{A} \left| \frac{\partial E_{ge}}{\partial I_0} \right|^{-1} \\
\text{(transmon)} \quad T_2 &\sim \frac{2\hbar}{AE_{ge}}
\end{aligned}
\tag{1.14}$$

1.4 Thesis Overview

This dissertation discusses how the coherence times of superconducting resonators and qubits have been improved through the analysis technique of participation ratios. For superconducting qubits, the coherence time requirement of DiVincenzo’s criteria has been the main obstacle in the way of further experiments. The work described in this thesis shows that this obstacle has largely been removed, enabling a vast array of experiments that were impossible to attempt even a few years ago. All work described in this thesis was done in the group of Michel Devoret with strong additional influence from the groups of Rob Schoelkopf, Steve Girvin and Leonid Glazman.

On top of the introductions made in Chapter 1, Chapter 2 expands on the main theoretical concepts of the resonators and qubits involved in this work, including an alternative derivation of the Hamiltonian of a transmon coupled to a cavity, and the energy spectrum and loss calculations of the fluxonium qubit. Chapter 3 details the experimental techniques used for this dissertation, including a description of the sample fabrication, a discussion about sample boxes, how qubit readout is performed, and how various components and filters are hand-made.

There are four main experiments contained in this dissertation that led to an improvement in coherence times of four devices: planar resonators (Chapter 4), planar transmons (Chapter 5), 3D transmons (Chapter 6), fluxonium qubits (Chapter 7). For clarity, the evolution of transmon qubits will be described as three generations,

and can be identified by the type of resonator they are coupled to: coplanar waveguide (CPW) resonators (1st generation), compact resonators (2nd generation), and 3D cavity resonators (3rd generation).

Chapter 4 details an experiment (published as Ref. [47]) that measures the quality factor of compact resonators as a function of geometrical parameters. In the end, quality factors were improved by roughly a factor of 2 by optimizing the geometry, as a result of measuring quality factors as a function of each parameter individually. Quality factors were improved from roughly 70,000 for previous CPW resonators to a median of 160,000 for the initial compact resonators to a median of 380,000 for the optimized compact resonators.

Chapter 5 details an experiment that improved the coherence time of transmon qubits by implementing the same geometrical design improvements as compact resonators. Called compact transmons because they were coupled to compact resonators, T_1 times were improved by a factor of 3 over previous qubits, from a maximum of 4 μ s to 13 μ s. Coherence times were largely unchanged with the change to compact transmons, with a maximum T_2^R of about 3 μ s.

The improvement in compact resonators and compact transmon coherence is due to the impact of geometrical parameters on participation ratios. While the materials were unchanged, the impact of dissipative components was reduced. These experiments served as a confirmation of the validity of the participation ratios analysis technique. The idea of participation ratios is then pushed to its limit in the 3D transmon, described in Chapter 6, with its drastically simplified mode structure and further reduction of surface participation. This development originated at Yale and has spread within a couple of years to most superconducting qubit groups, and represented an immediate jump of a factor of 5 from the compact transmons and a factor of 15 from CPW transmons. This chapter also mentions how coherence times

were improved another factor of 2 from initial 3D transmons, leading in the end to approximately $100 \mu\text{s}$ for both T_1 and T_2 . This chapter also describes an experiment (published as Ref. [48]) that tested a new scheme to initialize a qubit to the ground state called the Double Driven Reset of Population (DDROP) method. This method is quick to calibrate and use, and does not rely on feedback, high fidelity readout, or precisely tuned pulses, making it easy to use and a big help towards achieving the initialization requirement of DiVincenzo’s criteria.

The notion of participation ratios is then extended with the fluxonium, described in Chapter 7, which borrows the simplified geometry of the 3D transmon combined with a decreased sensitivity to dissipation. The decreased sensitivity to loss is due to wavefunction symmetries and is calculated in Chapter 2. The improvements over previous fluxonium implementations include a redesign of the superinductance junction fabrication and cavity format. Relaxation times were improved by over a factor of 100 , from of order $10 \mu\text{s}$ to over 1 ms , and coherence times were improved by a factor of 3 from roughly $15 \mu\text{s}$ to roughly $50 \mu\text{s}$. Relaxation times as a function of frequency are fit to a predicted dependence including quasiparticle loss, dielectric loss, and loss due to the Purcell effect. In addition, the measurement of a peak in T_1 times at the half-flux symmetry point represents a smoking-gun measurement of coherent tunneling of quasiparticles across a Josephson tunnel junction.

In summary, the coherence times of each resonator and qubit type have been improved significantly through the application of the participation ratio analysis, indicating that coherence times are no longer the main limitation of superconducting quantum information systems. The overall trend of the improvement of qubit and cavity lifetimes since the first superconducting qubit experiments can be seen in the summary plot below; the contributions of this thesis on the right side of the plot are evident. A comparison between Moore’s law concerning the number and cost

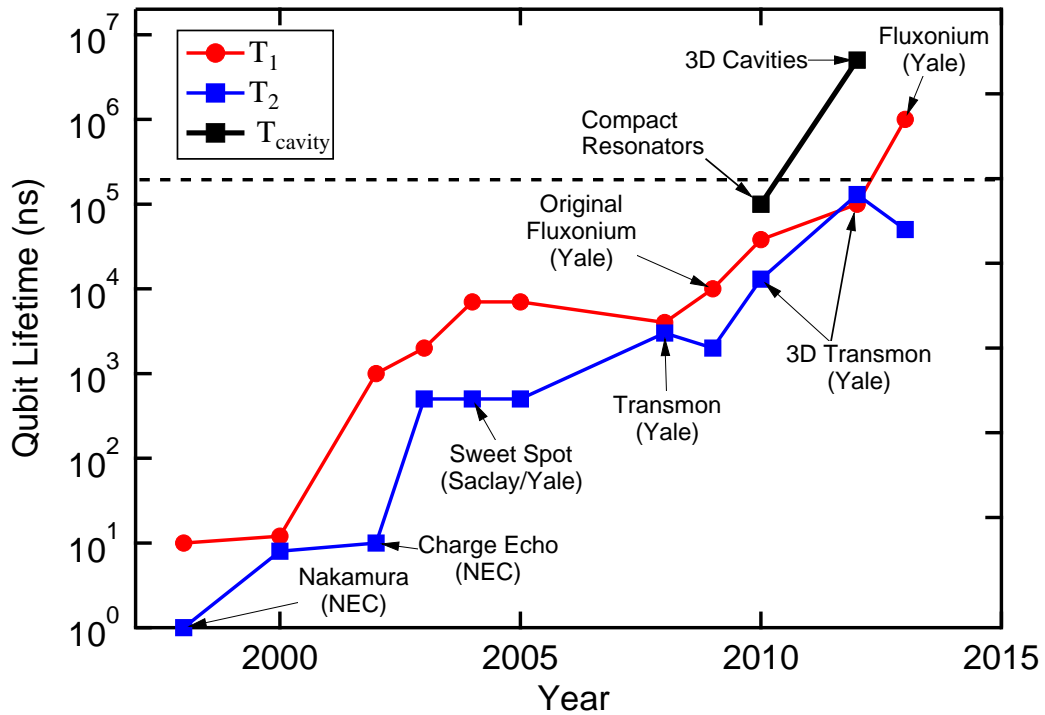


Figure 1.10: Qubit T_1 and T_2 , along with cavity lifetimes, are plotted as a function of year since the first superconducting qubit experiments. Typically, samples that produced new record values at the time of measurement were included. The results from compact resonators, 3D Transmon and fluxonium from this thesis are shown.

of transistors on integrated circuits can be made with superconducting qubits. As shown in this plot, “Schoelkopf’s law” observes that individual qubit relaxation and coherence times increase at a rate of approximately a factor of ten every three years.

Over the course of the next few years, the superconducting qubit community will continue performing experiments on multiple qubit algorithms, QND measurements, and error correction. If all goes well, qubit architectures will arise that support logical qubits: manifestations of a larger system that act as single qubits that live longer than any single physical qubit. From there, operations and algorithms on logical qubits will hopefully follow, which at some point in the future may lead to the holy grail of fault-tolerant quantum computation [49].

Resonator and Qubit Theory

2.1 Resonator Theory

2.1.1 2D Resonator Theory

An on-chip resonator is formed from thin-films deposited onto a substrate. The two commonly used substrates for superconducting qubits are silicon and “sapphire”. It should be noted that the sapphire substrate that we use is actually corundum, or pure Al_2O_3 , without the amount of impurities characterizing usual sapphire gems. Both silicon and sapphire (corundum) substrates can be ordered commercially with basically any specified thickness or diameters (up to 12”). The resonator structure is defined from the patterning of a thin metallic film which is deposited on the surface of the substrate. The metal can either be selectively deposited on the substrate, or deposited everywhere and selectively removed via etching. Selectivity is provided by a thin coating of a plastic material which is patterned by ultraviolet light or electron-beam lithography. Deposition is performed by evaporating the metal from a crucible onto the substrate; the thickness is determined from the rate of evaporation and the time.

In principle, any metal that can be evaporated can be used for a 2D resonator, but in the search for high quality factors, only superconducting materials are used, including: Al, Nb, TiN, NbTiN, and sometimes Re [50, 51, 52]. The thickness of the thin-film can vary from a few nanometers to a few microns depending on the application. The superconducting critical temperature (T_c) of each of these materials is between 1 K and 10 K, although the observed T_c , even for a given material, can vary depending on deposition parameters and film thickness. For experiments in this thesis, substrates of between 100 μm to 500 μm and film thicknesses between 10 nm and 200 nm were used.

There are several styles of planar resonators, each characterized by a unique electric and magnetic field pattern. One of the most commonly used on-chip resonator styles is the “coplanar waveguide” (CPW) resonator. The CPW resonator is characterized by two open strips cut into a metal plane, creating a central feedline and a plane on either side. The electric field travels in the two slots, with the voltage applied between the two ground planes and the central feedline. A CPW resonator is formed by interrupting a CPW trace with two impedance discontinuities: each either an inductance to ground or more commonly, a capacitor in the feedline. These interruptions cause the microwaves to reflect, defining resonant frequencies when standing waves are produced in the space between the two interruptions, in a manner equivalent to a Fabry-Perot cavity. The fundamental mode frequency (f_0) is defined as the frequency at which the length between the interruptions is exactly half of a wavelength. There are additional, equally spaced, modes at Nf_0 corresponding to a full wavelength, $3/2$ wavelengths, etc.

An alternative resonator design is the compact resonator, discussed at length in Chapter 4. Compact resonators, while still physically consisting of a thin-film deposited on a substrate, are distinct from CPW resonators because they consist of

a largely inductive region in parallel to a largely capacitive region. The inductive region is simply a thin length of wire, and the capacitor is two blocks of metal separated by some distance. The “lumped” approximation yields a resonance in which the voltage of the capacitor oscillates 90 degrees out of phase with the current in the inductor. The name “compact” indicates that these resonators are much smaller in size than CPW resonators while still not being fully lumped: the elements are not much smaller than a quarter wavelength. Compact resonators do not have equally spaced resonances; simulations predict a factor of 5 between the fundamental resonance and the second resonance. Truly lumped resonators will display much larger ratios.

2.1.2 3D Resonator Theory

While 2D resonators are quite flexible in design and small in size, they have some major limitations. Since vacuum is always by far the lowest loss material, it does not seem reasonable to concentrate the electromagnetic energy near the surface of the chip. While the debate over the exact loss tangent of sapphire or silicon at low power and low temperature is still going on, it is clear that a large participation of vacuum is better.

In much the same way a CPW resonator was formed from terminating a specific length of CPW transmission line, one can terminate a specific length of 3D rectangular waveguide in order to make a 3D resonator. The termination on either end of the resonator must be a short, otherwise the ends of the waveguide would radiate. Thus this 3D waveguide resonator is simply composed of a “cavity”, a region of vacuum bounded on all sides by metal. This simple geometry highlights the other main advantage over planar resonators: there is no Printed Circuit Board (PCB), wirebonds, or glue included in the resonating structure.

To find the fundamental mode of a rectangular cavity with side lengths $d > a > b$ (shown in Fig. 2.1), one must match the boundary conditions imposed by the metal surrounding the cavity. As with any rectangular waveguide, only TE and TM modes are supported. Derived in many textbooks [53], the frequencies of the TE_{mnl} and TM_{mnl} modes can be calculated via Eq. 2.1, with indices m , n , and l corresponding to the number of half-wavelengths in each direction (a , b , and d , respectively) of the cavity, and c the speed of light.

$$f_{mnl} = \frac{c}{2\pi} \sqrt{\left(\frac{m\pi}{a}\right)^2 + \left(\frac{n\pi}{b}\right)^2 + \left(\frac{l\pi}{d}\right)^2} \quad (2.1)$$

The fundamental mode of this cavity is the TE_{101} , where the length along the waveguide propagation direction is half a wavelength, in much the same way of a CPW resonator. Here, the shortest two sides form the cross-section of the waveguide and the longest side is the direction of propagation. The electric field of this mode is shown in Fig. 2.1. While the electric field everywhere in the cavity is pointing in the y direction, the shape of the magnitude dependence determines the mode structure. For the TE_{101} mode, as shown in the subsets of Fig. 2.1, the magnitude is constant in the y -direction while half an oscillation is observed in both the x and z -directions, hence the numbers 1-0-1.

The frequency of higher modes can be calculated from Eq. 2.1, and their electric field structure can be drawn easily by forcing the right number of half wavelengths in each direction. While there are some influences of the higher modes in calculations such as the Purcell effect (see Section 1.3.1), for the experiments in this thesis, only the TE_{101} mode is used intentionally.




Figure 2.1: Rectangular cavity with side lengths $d > a > b$ in the z , x , and y directions, respectively. Details of the TE_{101} (fundamental) mode of the cavity are shown. The electric field is entirely the y -direction and its dependence of magnitude along each axis is shown in the insets. For this mode, the field is constant in the y -direction, while half an oscillation is observed in both the x and z -directions.

2.2 Superconducting Artificial Atom Theory

2.2.1 Transmon Theory

Since the transmon is so widely used, there are many useful sources for transmon derivations and analyses [29, 23, 30, 31, 54, 55]. For this reason, a derivation of the transmon parameters will not be repeated here: equations for the energies (1.5) and charge dispersion (1.6) are given in Section 1.2.1. While the Jaynes-Cummings Hamiltonian has often been used in the past, its assumption of infinite anharmonicity is not well suited for the small anharmonicity of the transmon. A recent analysis by Nigg et al. [56] proposes a fully quantum, generalized multimode method for determining the Hamiltonian of a weakly anharmonic system. A semi-classical derivation of the Hamiltonian of a transmon coupled to a single resonator mode from the impedance seen across the junction has been first achieved by Shyam Shankar

and Michel Devoret.

Anharmonic Oscillator Derivation of transmon Equations

The Jaynes-Cummings Hamiltonian, shown in Eq. 2.2 is a general Hamiltonian for describing a two-level system coupled to a harmonic oscillator, and has been frequently used in many transmon studies [29]. Here ω_R and ω_Q are the resonator and qubit resonant frequencies, respectively, a and a^\dagger are the cavity raising and lowering operators, σ^- and σ^+ are the qubit raising and lowering operators, and g is the qubit-cavity coupling strength.

$$H_{JC} = \hbar\omega_R \left(a^\dagger a + \frac{1}{2} \right) + \frac{\hbar\omega_Q}{2} \sigma_z + \hbar g (a^\dagger \sigma^- + a \sigma^+) \quad (2.2)$$

However, with its reduced anharmonicity in the high E_J/E_C limit, a two-level system is a poor approximation for a transmon. The transmon can be more accurately represented as a slightly anharmonic resonator, treating the anharmonicity as a perturbation. The theory of an anharmonic oscillator coupled to a harmonic oscillator will now be derived as an alternative to the Jaynes-Cummings model.

Starting from the circuit model of a transmon coupled to a harmonic oscillator, shown in Fig. 2.2, the Hamiltonian can be written by summing the energy of each capacitor, inductor, and Josephson tunnel junction. The notation used here is such that the capacitances are: C_R , the resonator capacitance, C_S , the direct shunting capacitance of the transmon, and C_C , the effective total capacitance between the transmon and resonator. With details available in Chapter 2 of Ref. [57], the Hamiltonian of this circuit can be written as Eq. 2.3. Here, the terms in sequence are from: the cavity (mode 1), the qubit (mode 2), and the coupling; also the Q_i and ϕ_i represent the charge and phase of the respective capacitors and inductors.




Figure 2.2: Lumped element model of transmon qubit (dashed green box) coupled to resonator (dashed blue box). Particular shunting capacitances C_S and coupling capacitances C_C are needed in order to obtain a desired qubit frequency and coupling strength.

$$H = \left(\frac{Q_1^2}{2} \frac{1}{C_R + (C_C || C_S)} + \frac{\phi_1^2}{2L_R} \right) + \left(\frac{Q_2^2}{2} \frac{1}{C_S + (C_C || C_R)} - E_J \cos \left(\frac{\phi_2}{\phi_0} \right) \right) + Q_1 Q_2 \frac{C_C}{C_R C_S + C_S C_C + C_C C_R} \quad (2.3)$$

In order to approximate the qubit mode as a harmonic oscillator, one must expand the $\cos \left(\frac{\phi_2}{\phi_0} \right)$ term in the limit of small ϕ_2 , as shown in Eq. 2.4. Rewriting the above Hamiltonian and defining $L_J = \phi_0^2/E_J$, there are now two harmonic oscillator modes, along with a separate anharmonicity term of the qubit mode, shown in Eq. 2.5.

$$\cos \left(\frac{\phi_2}{\phi_0} \right) = 1 - \frac{1}{2!} \frac{\phi_2^2}{\phi_0^2} + \frac{1}{4!} \frac{\phi_2^4}{\phi_0^4} + \dots \quad (2.4)$$

$$H = \left(\frac{Q_1^2}{2} \frac{1}{C_R + (C_C || C_S)} + \frac{\phi_1^2}{2L_R} \right) + \left(\frac{Q_2^2}{2} \frac{1}{C_S + (C_C || C_R)} + \frac{\phi_2^2}{2L_J} \right) + Q_1 Q_2 \frac{C_C}{C_R C_S + C_S C_C + C_C C_R} - \frac{1}{4!} \frac{\phi_2^4}{L_J \phi_0^2} \quad (2.5)$$

Now, several substitutions are made to covert to convert these classical parameters into quantum operators. They introduce the resonant frequencies and impedances ($\omega_a, \omega_b, Z_a, Z_b$) and raising and lowering operators for each mode in terms of ϕ_i and Q_i : a and a^\dagger for the cavity, b and b^\dagger for the ‘‘harmonic’’ qubit mode. Also introduced at this point is the coupling strength g and qubit coupling energy E_C both expressed in terms of capacitances. The resulting Hamiltonian can be split into that of two coupled harmonic oscillators (H_0) and the anharmonicity perturbation (H_1), shown in Eq. 2.7.

$$\begin{aligned}
\phi_a &= (a + a^\dagger) \sqrt{\frac{\hbar Z_a}{2}} & \phi_b &= (b + b^\dagger) \sqrt{\frac{\hbar Z_b}{2}} \\
Q_a &= \frac{a - a^\dagger}{i} \sqrt{\frac{\hbar}{2Z_a}} & Q_b &= \frac{b - b^\dagger}{i} \sqrt{\frac{\hbar}{2Z_b}} \\
Z_a &= \sqrt{\frac{L_R}{(C_R + (C_C \| C_S))}} & Z_b &= \sqrt{\frac{L_J}{(C_S + (C_C \| C_R))}} \\
\omega_a &= \frac{1}{\sqrt{L_R(C_R + (C_C \| C_S))}} & \omega_b &= \frac{1}{\sqrt{L_J(C_S + (C_C \| C_R))}} \\
g &= \frac{1}{2} \sqrt{\omega_R \omega_Q} \frac{C_C}{\sqrt{(C_R + C_C)(C_S + C_C)}} \\
E_C &= \frac{e^2}{2(C_S + (C_C \| C_R))}
\end{aligned} \tag{2.6}$$

$$\begin{aligned}
H_0 &= \hbar \omega_a \left(a^\dagger a + \frac{1}{2} \right) + \hbar \omega_b \left(b^\dagger b + \frac{1}{2} \right) + \hbar g \left(a^\dagger b + ab^\dagger \right) + \hbar g \left(ab + a^\dagger b^\dagger \right) \\
H_1 &= -\frac{E_C}{12} \left(b^\dagger + b^\dagger \right)^4
\end{aligned} \tag{2.7}$$

In this simplified expression for H_0 , the term $\hbar g (ab + a^\dagger b^\dagger)$ can be ignored due to the rotating wave approximation in the case where $g \ll \omega_a, \omega_b$. The two modes at frequencies ω_a and ω_b are the bare modes, which are not eigenmodes of the coupled system. The coupling introduces a dressing which redefines these two resonant modes (this redefinition is also key to understanding the Purcell effect on relaxation, see Section 1.3.1). Ignoring the anharmonic term for now, one can diagonalize the Hamiltonian matrix. This yields two dressed modes, whose raising and lowering operators are given by Eq. 2.8, and dressed frequencies which corresponds to a linear combination of bare qubit and bare cavity frequencies. The dressed frequencies are given in Eq. 2.9, where Δ is the difference between the bare resonant frequencies $\omega_b - \omega_a$. The Hamiltonian in terms of the new raising and lowering operators A , A^\dagger , B , and B^\dagger is shown in Eq. 2.10.

$$\begin{aligned} \begin{pmatrix} a \\ b \end{pmatrix} &= \begin{pmatrix} \lambda_a & \mu_a \\ \lambda_b & \mu_b \end{pmatrix} \begin{pmatrix} A \\ B \end{pmatrix} \\ \begin{pmatrix} \lambda_a & \mu_a \\ \lambda_b & \mu_b \end{pmatrix} &= \begin{pmatrix} -g\sqrt{\frac{1+\frac{1}{4g^2}(\Delta+\sqrt{4g^2+\Delta^2})^2}{4g^2+\Delta^2}} & -\frac{\Delta-\sqrt{4g^2+\Delta^2}}{2}\sqrt{\frac{1+\frac{1}{4g^2}(\Delta+\sqrt{4g^2+\Delta^2})^2}{4g^2+\Delta^2}} \\ g\sqrt{\frac{1+\frac{1}{4g^2}(\Delta-\sqrt{4g^2+\Delta^2})^2}{4g^2+\Delta^2}} & \frac{\Delta+\sqrt{4g^2+\Delta^2}}{2}\sqrt{\frac{1+\frac{1}{4g^2}(\Delta-\sqrt{4g^2+\Delta^2})^2}{4g^2+\Delta^2}} \end{pmatrix} \end{aligned} \quad (2.8)$$

$$\begin{aligned} \omega_A &= \frac{1}{2} \left(\omega_a + \omega_b - \sqrt{4g^2 + \Delta^2} \right) \\ \omega_B &= \frac{1}{2} \left(\omega_a + \omega_b + \sqrt{4g^2 + \Delta^2} \right) \end{aligned} \quad (2.9)$$

$$H_0 = \hbar\omega_A \left(A^\dagger A + \frac{1}{2} \right) + \hbar\omega_B \left(B^\dagger B + \frac{1}{2} \right) \quad (2.10)$$

The energy levels of the dressed Hamiltonian can be calculated also using the number operators N_A and N_B quantifying the number of excitations in each oscillator.

$$E_{N_A, N_B} = \hbar\omega_A \left(N_A + \frac{1}{2} \right) + \hbar\omega_B \left(N_B + \frac{1}{2} \right) \quad (2.11)$$

Since ω_A and ω_B are the dressed frequencies, neither correspond directly to the qubit or cavity mode. However, in the limit of small g/Δ , the dressed frequencies approach the bare frequencies, $\omega_A \rightarrow \omega_R$ and $\omega_B \rightarrow \omega_Q$. The coupled resonant frequencies are plotted versus Δ in Fig. 2.3, showing that when $\Delta = 0$, the coupling creates an avoided crossing of $2g$. At this point, qubit and cavity are maximally coupled and each mode is half-qubit, half-cavity.

Returning to the anharmonic term H_1 , b and b^\dagger must be replaced with the dressed operators. The anharmonicity is characterized by the energy E_C in Eq. 2.7, and represents the difference in energy between the first and second transitions in the qubit

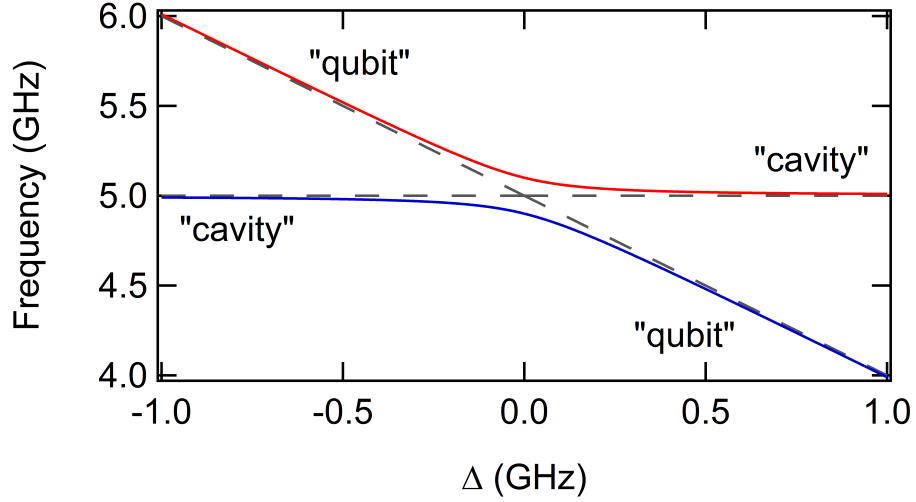


Figure 2.3: Dressed (red,blue) and bare (grey) frequencies as a function of detuning. Avoided crossing of frequency $2g$ is observed when $\Delta = 0$. Far from avoided crossing, qubit and cavity mode can be identified, but at the avoided crossing, each mode is part qubit and cavity.

spectrum. Typical transmon parameters yield $E_C/h = 300$ MHz; with $\omega_A, \omega_B \simeq 5$ - 10 GHz, E_C can be treated as a perturbation. Terms of order E_C^2 will thus be ignored. Collecting only the diagonal terms of $(b + b^\dagger)^4$ after the substitution of the dressed operators yields Eq. 2.12 in terms of the parameters characterizing the dressed operators in terms of the bare operators, λ_a , λ_b , μ_a , and μ_b . The resulting energy levels of the complete Hamiltonian $H = H_0 + H_1$ are given in Eq. 2.13.

$$H_1 \simeq -E_C (\lambda_b^2 A^\dagger A + \mu_b B^\dagger B) - \frac{E_C}{2} (\lambda_b^4 A^{\dagger 2} A^2 + \mu_b^2 B^{\dagger 2} B^2 + 4\lambda_b^2 \mu_b^2 A^\dagger A B^\dagger B) \quad (2.12)$$

$$E_{N_A, N_B} = \hbar\omega_A N_A + \hbar\omega_B N_B - E_C (\lambda_b^2 N_A + \mu_b^2 N_B) - \frac{E_C}{2} (\lambda_b^4 (N_A^2 - N_A) + \mu_b^4 (N_B^2 - N_B) + 4\lambda_b^2 \mu_b^2 N_A N_B) \quad (2.13)$$

These energy levels can be expressed in a much more understandable manner by making the following definitions. ω'_A and ω'_B are shifted resonant frequencies due to the anharmonicity. χ_A and χ_B are self-Kerr terms, and represent a shift in resonant frequency due to photons in the same mode. Non-zero self-Kerr means that each mode is now anharmonic, with transition frequencies changing with each additional photon. The redefinition of modes from a and b to A and B spread the anharmonicity from the qubit mode to the cavity mode. The term χ_{AB} , called the cross-Kerr is perhaps the most interesting, representing either a shift in ω'_A due to excitations in mode B or a shift in ω'_B due to excitations in mode A . This term is the basis of all qubit readout in circuit QED [23, 58] (see Section 3.2). The qubit-cavity Hamiltonian with these substitutions is given in Eq. 2.15, with energy levels given in Eq. 2.16.

$$\begin{aligned}
\omega'_A &= \omega_A - E_C \lambda_b^2 + \frac{E_C}{2} \lambda_b^4 \\
\omega'_B &= \omega_B - E_C \mu_b^2 + \frac{E_C}{2} \mu_b^4 \\
\chi_A &= -\frac{E_C}{2} \lambda_b^4 \\
\chi_B &= -\frac{E_C}{2} \mu_b^4 \\
\chi_{AB} &= -\frac{E_C}{2} (4\lambda_b^2 \mu_b^2)
\end{aligned} \tag{2.14}$$

$$H = \hbar\omega'_A \left(A^\dagger A + \frac{1}{2} \right) + \hbar\omega'_B \left(B^\dagger B + \frac{1}{2} \right) + \chi_A (A^\dagger A)^2 + \chi_B (B^\dagger B)^2 + \chi_{AB} (A^\dagger A B^\dagger B) \tag{2.15}$$

$$\frac{E_{N_A, N_B}}{\hbar} = \omega'_A N_A + \omega'_B N_B + \chi_A N_A^2 + \chi_B N_B^2 + \chi_{AB} N_A N_B \tag{2.16}$$

The full expressions for the self-Kerrs and the cross-Kerr in terms of g , E_C and Δ are, for reference, given in Eqs. 2.17. The often-quoted and highly useful results from the Jaynes-Cummings model can be derived by simplifying these expressions in the

limit where g/Δ is small. This limit corresponds to when the detuning between qubit and cavity modes is large enough compared to the coupling strength, that the qubit and cavity modes can almost be treated separately. The coupling can then be treated to first order in g/Δ . These simplified expressions are shown in Eqs. 2.18. Also in this limit $\omega'_A \rightarrow \omega_A$ and $\omega'_B \rightarrow \omega_B - E_C/2$. From these expressions, one can calculate the anharmonicities, showing that the cavity is now approximately harmonic, and the qubit has anharmonicity $-E_C$. In this dispersive limit, the Hamiltonian can be simplified as shown in Eq. 2.19, including a simple renaming of modes A and B to Q and C since the qubit and cavity are well-separated. To match common notation, anharmonicity χ_B is replaced with α ($\equiv 2\chi$) and dispersive shift χ_{AB} is simplified to χ .

$$\begin{aligned}
\chi_A &= -E_C \frac{g^4}{2} \left(\frac{1 + \frac{1}{4g^2} (\Delta - \sqrt{4g^2 + \Delta^2})}{4g^2 + \Delta^2} \right)^2 \\
\chi_B &= -E_C \frac{(\Delta + \sqrt{4g^2 + \Delta^2})^4}{32} \left(\frac{1 + \frac{1}{4g^2} (\Delta - \sqrt{4g^2 + \Delta^2})}{4g^2 + \Delta^2} \right)^2 \\
\chi_{AB} &= -E_C \frac{g^2 (\Delta + \sqrt{4g^2 + \Delta^2})^2}{2} \left(\frac{1 + \frac{1}{4g^2} (\Delta - \sqrt{4g^2 + \Delta^2})}{4g^2 + \Delta^2} \right)^2
\end{aligned} \tag{2.17}$$

$$\begin{aligned}
\chi_A &\rightarrow 0 \\
\chi_B &\rightarrow -E_C \frac{1}{2} \\
\chi_{AB} &\rightarrow -2E_C \frac{g^2}{\Delta^2}
\end{aligned} \tag{2.18}$$

$$\boxed{H = \hbar\omega_C \left(A^\dagger A + \frac{1}{2} \right) + \hbar\omega_Q \left(B^\dagger B + \frac{1}{2} \right) + \hbar\alpha (B^\dagger B)^2 + \hbar\chi (A^\dagger A B^\dagger B)} \tag{2.19}$$

In summary, this coupled harmonic oscillator approach to calculating the transition equations is a valid alternative to the Jaynes-Cummings model. This approach treats the qubit anharmonicity as a small parameter, as opposed to treating the

qubit as a two-level system. While the two-level approximation works well for CPBs, the anharmonicity for transmon qubits, by design, is quite low; thus this approach is more natural in the case of the transmon. Reassuringly, the results are equivalent when the dispersive limit is taken.

2.2.2 Fluxonium Theory

Superinductance

The fluxonium qubit, as introduced in Section 1.2.2, differs from the transmon in that it is composed of a Josephson tunnel junction shunted by a large inductance instead of a large capacitance. This inductance shunts charges across the junction islands and protects the circuit from low-frequency charge fluctuations. To create the energy spectrum of fluxonium, described below, this shunting inductance must be very high; so high in fact that its impedance must be higher than the resistance quantum. This regime is completely unreachable with geometric inductance due to the presence of self-resonant frequencies (associated with the self-capacitance of any geometric inductor). Utilizing Josephson tunnel junction kinetic inductance, an array of Josephson tunnel junctions is a good approximation to a linear inductor in the limit of low phase across each junction ($\sin\varphi \rightarrow \varphi$). This inductance, able to reach impedances much higher than the resistance quantum, is called a “superinductance”, and alone is responsible for many interesting experiments [59, 36].

There are four design considerations and constraints that should not be ignored when designing a superinductance for a fluxonium qubit [33]. The parameters of concern are shown in Fig. 2.4: L_{JA} , the inductance of a single array junction, C_{JA} , the capacitance of a single array junction, C_g , the capacitance of a junction island to ground, $Z_J = \sqrt{L_J/C_J}$, the impedance of the qubit junction, $Z_{JA} = \sqrt{L_{JA}/C_{JA}}$, the




Figure 2.4: Circuit schematic of the fluxonium qubit highlighting the various parameters. The qubit junction is shown on top, shunted by a series of array junctions. Each junction is composed of a capacitance and an inductance, C_J and L_J for the qubit junction and C_{JA} and L_{JA} for array junctions. The parameter N_A counts the number of array junctions. The parasitic capacitance to ground is shown by additional capacitances C_g from the islands between array junctions to the ground.

impedance of a single array junction, N_A , the number of junctions in the array, and $R_Q = \hbar/(2e)^2$, the resistance quantum. The first constraint is that $N_A L_{JA} \gg L_J$, where L_J is the inductance of the qubit junction; this allows the loop to support large flux variations instead of large charge variations. The second constraint is that in order to reduce the island offset charge, the size of the array junction must be chosen such that $e^{-8R_Q/Z_{JA}} \ll 1$. The lower this exponential, the lower the offset charge. The third constraint, $N_A e^{-8R_Q/Z_{JA}} \ll e^{-8R_Q/Z_J}$ prevents the array from being subject to phase slips. The last constraint limits the number of junctions in the array to $N_A < \sqrt{C_{JA}/C_g}$ as a means of ensuring self-resonant modes of the array are at a high enough frequency.

Energy Levels

While a detailed derivation of the fluxonium energy spectrum is included in the recent thesis of Nick Masluk [36], a summary will be provided here to facilitate basic understanding of the method. The energy levels of the fluxonium artificial atom can be calculated from the potential energy landscape as a function of the phase across the junction by diagonalizing the Hamiltonian matrix in the harmonic oscillator basis. The fluxonium Hamiltonian (without the cavity) in the phase basis is shown in Eq. 2.20; the terms are the energies of the capacitor (E_C), inductor (E_L), and junction (E_J), respectively. These energies are defined in Eqs. 2.21, where L_J is the inductance of the qubit (or “phase-slip”) junction, L_A is the shunting (junction array) inductance, and C_J is the qubit junction capacitance.

$$H = -4E_C \frac{\partial^2}{\partial \hat{\varphi}^2} + \frac{E_L}{2} \hat{\varphi}^2 + E_J \cos(\hat{\varphi}) \quad (2.20)$$

$$E_J = \frac{(\Phi_0/2\pi)^2}{L_J} \quad E_C = \frac{e^2}{2C_J} \quad E_L = \frac{(\Phi_0/2\pi)^2}{L_A} \quad (2.21)$$

The first two terms make up a harmonic LC oscillator whose wavefunctions (ψ_l) are known to be Hermite polynomials [60] (H_l), where the quantum number l represents the excitation number and ϕ_0 is the zero point motion.

$$\psi_l(\varphi) = \frac{1}{\sqrt{2^l l!} \sqrt{\pi} \phi_0} e^{-\frac{1}{2} \left(\frac{\varphi}{\phi_0}\right)^2} H_l \left(\frac{\varphi}{\phi_0}\right) \quad (2.22)$$

$$\phi_0 = \left(\frac{8E_C}{E_L}\right)^{1/4} \quad (2.23)$$

Due to the presence of the cosine term in the energy, these wavefunctions of the harmonic oscillator are not eigenfunctions of the fluxonium. The Hamiltonian matrix for the fluxonium can be completed by using the harmonic oscillator energies as diagonal elements and adding the off-diagonal elements from the cosine term as shown in Eq. 2.24, where the m, n are the matrix indices. The elements $\langle \psi_m(\varphi) | \sin \hat{\varphi} | \psi_l(\varphi) \rangle$ and $\langle \psi_m(\varphi) | \cos \hat{\varphi} | \psi_l(\varphi) \rangle$ evaluate to associated Laguerre polynomials.

$$\begin{aligned} & \langle \psi_m(\varphi) | E_J \cos \left(\hat{\varphi} - 2\pi \frac{\Phi_{ext}}{\Phi_0} \right) | \psi_l(\varphi) \rangle \\ &= \begin{cases} E_J \sin \left(2\pi \frac{\Phi_{ext}}{\Phi_0} \right) \langle \psi_m(\varphi) | \sin \hat{\varphi} | \psi_l(\varphi) \rangle, & \text{if } l + m \text{ is odd} \\ E_J \sin \left(2\pi \frac{\Phi_{ext}}{\Phi_0} \right) \langle \psi_m(\varphi) | \cos \hat{\varphi} | \psi_l(\varphi) \rangle, & \text{if } l + m \text{ is even} \end{cases} \end{aligned} \quad (2.24)$$

To calculate the fluxonium wavefunctions and energies, the Hamiltonian matrix is diagonalized, with the resulting eigenvectors representing the fluxonium eigenfunctions as linear superpositions of the harmonic oscillator wavefunctions, and the eigenenergies the fluxonium energy levels. In general, fluxonium potential can be

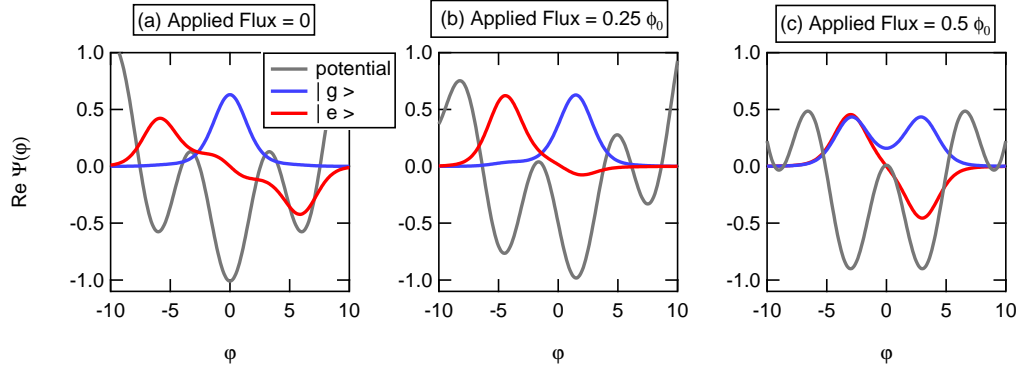


Figure 2.5: Wavefunctions and energy potential shown in the phase basis for first two eigenstates of the fluxonium artificial atom, with $E_J = 10.2$ GHz, $E_C = 3.6$ GHz, and $E_L = 0.46$ GHz. Ground (blue) and excited (red) wavefunctions change shape and character considerably as the knob of external flux is varied from zero to half-flux.

thought of as a sine wave on top of a parabola, with the ratio E_J/E_L controlling how many local minima are found in the potential. The parameter regime of fluxonium is when there are approximately 3 or 4 local minima, corresponding to $E_J/E_L \approx 20$. For parameters $E_J = 10.2$ GHz, $E_C = 3.6$ GHz, and $E_L = 0.46$ GHz (the parameters of one of the samples in chapter 7), the fluxonium wavefunctions in the phase basis can be seen at various applied flux values in Fig. 2.5. At zero applied flux, the ground state resides in the center well, and the excited state is anti-symmetrically split between the two side wells. At half flux, there are two equal center wells; the ground and excited states are symmetric and antisymmetric combinations of these two wells. This highly variable character is what makes the fluxonium a rich playground for interesting physics, yielding effects such as the insensitivity to quasiparticle tunneling at half-flux quantum, as shown in Section 7.5.

This strong variation in wavefunctions versus flux produces a transition spectrum whose lowest ($|g\rangle$ to $|e\rangle$) (qubit) transition varies over a wide range, typically from around 500 MHz to around 8 GHz. The transition from $|g\rangle$ to $|f\rangle$, denoted as frequency f_{gf} , on the other hand, varies much less as the flux is varied. The readout

strength can be thought of as derived from the difference in frequency from this transition to the cavity frequency. While the f_{ge} transition can be detuned far from the cavity, the coupling strength is maintained due to the proximity of the $|g\rangle$ to $|f\rangle$ transition. Details of the calculation of qubit-cavity coupling strength for the fluxonium can be found in Nick Masluk’s thesis [36]. Unlike the transmon, the fluxonium-cavity system is well represented by the Jaynes-Cummings Hamiltonian (see Section 2.2.1) since its anharmonicity is larger.

The three parameters E_J , E_C and E_L each control the transition spectrum in their own way. The parameter E_J/E_C determines the splittings between f_{ge} and f_{ef} at zero flux and between 0 and f_{ge} at half flux by roughly determining the number of excitations in each well of the sine wave potential. Low E_J/E_C leads to large half-flux frequencies. In the potential, E_L describes the height of the parabola compared to the sine wave period; large E_L leads to a “tall” parabola with few oscillations, while small E_L leads to a low parabola such that the overall potential is simply a sine wave. In the transition spectrum, E_L controls the slope of the linear portion between zero and half-flux. With the frequency at half-flux fixed, the slope is the dominant factor in determining the frequency at zero flux. For the qubit to be considered a fluxonium, the parameters must match approximately $2 < E_J/E_C < 5$ and $E_L \sim 0.5$ GHz.

Loss Calculations

Since the fluxonium qubit is tunable over a wide range of frequencies, this can be exploited as a means to learn about loss mechanisms. Each loss mechanism is predicted to have a unique dependence on frequency for the fluxonium; comparing the expected dependencies with measured relaxation times can help elucidate which loss mechanisms are dominating in any given qubit. Again, details are given in Nick




Figure 2.6: Transition spectrum of fluxonium qubit with $E_J = 10.2$ GHz, $E_C = 3.6$ GHz, and $E_L = 0.46$ GHz. Transitions shown are the first three transitions from the ground state, with frequencies denoted as f_{ge} , f_{gf} , and f_{gh} . The lowest (qubit) transition varies from 8 GHz to 500 MHz.

Masluk’s thesis [36], with basics described here.

As mentioned in Section 1.3.1, each component of a circuit can contribute to qubit relaxation with its own lossy component. These losses are combined in parallel, with rates that simply add to yield a total effective relaxation rate. The loss mechanisms considered for the fluxonium qubit are capacitive loss in both the qubit junction capacitance and the environmental capacitance, inductive loss in the qubit junction due to quasiparticles, inductive loss in the junction array, and loss due to the Purcell effect, i.e. radiative losses. These losses are called “capacitive”, “quasiparticle”, “inductive”, and “Purcell”, respectively.

Each loss mechanism is characterized by a coupling Hamiltonian, $H_C = \hat{C}\hat{X}$, where \hat{C} is the operator that couples to the noise of \hat{X} . Using Fermi’s golden rule, the qubit relaxation (Γ_C^\downarrow) and excitation (Γ_C^\uparrow) rates can be calculated from Eqs. 2.25, the observed qubit relaxation time can then be calculated as $1/T_1 = \Gamma_C^\downarrow + \Gamma_C^\uparrow$. The frequencies used are defined as $\hbar\omega_{ge} = E_e - E_g$ and $\hbar\omega_{eg} = E_g - E_e$. The term $S_{XX}[\omega]$ is the spectral density of fluctuations in \hat{X} given by Eq. 2.26, derived from the quantum fluctuation dissipation theorem [61]. The part of the spectral density that depends on the environment density of states is the real part of the admittance of the element under consideration, $\text{Re}[Y_C(\omega)]$. The matrix element of the coupling operator and this admittance are different for each loss mechanism and are responsible for the different frequency dependencies.

$$\begin{aligned}\Gamma_C^\downarrow &= \frac{1}{\hbar^2} \left| \langle e | \hat{C} | g \rangle \right|^2 S_{XX}(\omega_{ge}) \\ \Gamma_C^\uparrow &= \frac{1}{\hbar^2} \left| \langle e | \hat{C} | g \rangle \right|^2 S_{XX}(\omega_{eg})\end{aligned}\tag{2.25}$$

$$S_{XX}[\omega] = \hbar\omega \text{Re}[Y_C(\omega)] \left(\coth\left(\frac{\hbar\omega}{2k_B T}\right) + 1 \right)\tag{2.26}$$

For capacitive loss, here lumping both junction capacitance and environmental




Figure 2.7: Plots showing (a) $\hat{\varphi}$ matrix element and (b) ω_{ge}^2 , the factors largely responsible for the frequency dependence of capacitive loss, here plotted as a function of applied flux for a fluxonium qubit with the same parameters as Fig. 2.6. Plot (c) combines these factors to show that predicted T_1 depends only slightly on applied flux; the overall magnitude of T_1 is then controlled solely by the scale factor Q_{cap} .

capacitance into one effective capacitance C_Σ , although this capacitance is largely due to the junction capacitance, the coupling Hamiltonian and admittance are given in Eqs. 2.27, and the calculated relaxation rate is given in Eq. 2.28, where the amount of loss is controlled by the quality factor Q_{cap} . Fig. 2.7 shows the calculation of the matrix element $|\langle e | \hat{\varphi} | g \rangle|^2$ and ω_{ge}^2 as a function of flux. According to Eq. 2.28, these two factors are combined, along with basically a scaling factor, to give the relaxation rate, also shown in Fig. 2.7 for $Q_{cap} = 1,000,000$. Note that the matrix element and frequency terms have nearly equal, but opposite dependencies. This leads to a relaxation rate, and thus a T_1 , that is largely independent of frequency.

$$\begin{aligned}
 H_C^{cap} &= -\hat{\Phi} \hat{I}_{cap} = -\phi_0 \hat{\varphi} \hat{I}_{cap} \\
 \text{Re}[Y_{cap}(\omega)] &= \frac{\omega C_\Sigma}{Q_{cap}} \\
 C_\Sigma &= \frac{e^2}{2E_C}
 \end{aligned}
 \tag{2.27}$$

$$\Gamma^{cap} = \frac{\hbar}{8E_C Q_{cap}} \left(\coth\left(\frac{\hbar\omega_{ge}}{2k_B T}\right) + 1 \right) \omega_{ge}^2 |\langle e | \hat{\varphi} | g \rangle|^2
 \tag{2.28}$$




Figure 2.8: (a) Plot showing the $\hat{\varphi}$ matrix element, the only major factor responsible for the frequency dependence of inductive loss, here plotted as a function of applied flux for a fluxonium qubit with the same parameters as previously. Plot (b) shows that the predicted T_1 has the same, but inverted dependence as a function of applied flux; the overall magnitude of T_1 is then controlled solely by the scale factor Q_{ind} .

For inductive loss, here considering the total inductance L_Σ , which is almost entirely due to the array inductance L_A , the coupling Hamiltonian and admittance are given in Eqs. 2.29, and the calculated relaxation rate is given in Eq. 2.30, where the amount of loss is controlled by the quality factor Q_{ind} . Fig. 2.8 shows the calculation of the matrix element $|\langle e|\hat{\varphi}|g\rangle|^2$, the same as with capacitive loss. According to Eq. 2.30, this is the only factor with substantial frequency dependence, and thus the plot of expected T_1 for a constant Q_{ind} of 1,000,000, shown in Fig. 2.8 has the inverse dependence versus applied flux as the matrix element.

$$\begin{aligned}
 H_C^{ind} &= -\hat{\Phi}\hat{I}_{ind} = -\phi_0\hat{\varphi}\hat{I}_{ind} \\
 \text{Re}[Y_{ind}(\omega)] &= \frac{1}{\omega L_\Sigma Q_{ind}} \\
 L_\Sigma &= \frac{\phi_0^2}{2E_L}
 \end{aligned} \tag{2.29}$$

$$\Gamma^{ind} = \frac{E_L}{\hbar Q_{ind}} \left(\coth\left(\frac{\hbar\omega_{ge}}{2k_B T}\right) + 1 \right) |\langle e|\hat{\varphi}|g\rangle|^2 \tag{2.30}$$

For quasiparticle loss in the qubit junction (quasiparticle loss in the array would

be considered inductive loss), the coupling Hamiltonian and admittance are given in Eqs. 2.31 where G_t is the tunneling conductance of a given junction. The calculated relaxation rate is given in Eq. 2.32, where the amount of loss is controlled by the density of broken Cooper pairs x_{qp} . Note that the coupling operator is now $\sin(\hat{\varphi}/2)$ as opposed to simply $\hat{\varphi}$ for both capacitive and inductive losses. This change occurs because of the particular way in which the tunneling amplitude of quasiparticles depends on the phase across a junction [25, 62]. Fig. 2.9 shows the calculation of the matrix element $|\langle e | \sin(\hat{\varphi}/2) | g \rangle|^2$ and $\sqrt{1/\omega_{ge}}$ as a function of applied flux. According to Eq. 2.32, these two factors are combined, along with basically a scaling factor, to give the relaxation rate, also shown in Fig. 2.9 for $x_{qp} = 5 \times 10^{-6}$. Note that the contribution from $\sqrt{1/\omega_{ge}}$ varies only by a factor of 3 and the matrix element is the dominant contribution to the relaxation.

$$\begin{aligned}
H_C^{qp} &= -\frac{\hbar}{e} \sin(\hat{\varphi}/2) \\
\text{Re}[Y_{qp}(\omega)] &= \frac{1}{2} x_{qp} G_t \left(\frac{2\Delta}{\hbar\omega_{ge}} \right)^2 \\
G_t &= \frac{8E_J e^2}{\hbar\Delta}
\end{aligned} \tag{2.31}$$

$$\Gamma^{qp} = \frac{16x_{qp}E_L}{h} \sqrt{\frac{2\Delta}{\hbar}} \left(\coth\left(\frac{\hbar\omega_{ge}}{2k_B T}\right) + 1 \right) \sqrt{\frac{1}{\omega_{ge}}} |\langle e | \sin(\hat{\varphi}/2) | g \rangle|^2 \tag{2.32}$$

The predicted peak in T_1 for quasiparticle loss is particularly striking. A peak in T_1 means that the qubit becomes blind to the quasiparticle loss mechanism. Thus, at half-flux, the qubit junction becomes completely insensitive to quasiparticle tunneling. While explained in detail by Catelani et al. [62], there are two ways to think about this feature in T_1 . The first, highlighted in Fig. 2.10 and explained in further detail by D.N. Langenberg [63], considers that a quasiparticle may absorb energy from the qubit (and thus induce a relaxation process) and tunnel across a junction

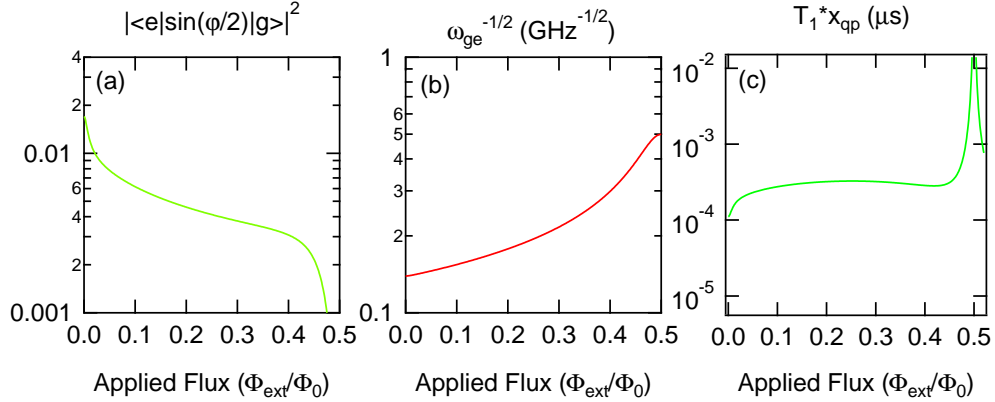


Figure 2.9: Plots showing (a) the $\sin(\hat{\varphi}/2)$ matrix element and (b) the only major factor responsible for the frequency dependence of inductive loss, here plotted as a function of applied flux for a fluxonium qubit with the same parameters as previously. Plot (c) combines these factors to show that the predicted T_1 is largely independent of flux except for a dramatic peak at an applied flux of a half flux quantum; the overall magnitude is controlled solely by the scale factor x_{qp} .

either as an electron or a hole. The probability amplitudes of these processes are $e^{+i\varphi/2}$ and $e^{-i\varphi/2}$, respectively. Since both are possible, the amplitudes are added to give $\cos(\varphi/2)$, which at $\varphi = \pi$ evaluates to zero. Thus the relaxation vanishes at half-flux due to quasiparticle interference. However, this way of thinking is not quite correct since φ has large quantum fluctuations.

The second way to think about the predicted peak in T_1 is to inspect the matrix element $|\langle e | \sin(\hat{\varphi}/2) | g \rangle|$ which contributes to the calculated dependence of Γ_{qp} . The matrix element is calculated by integrating the product of the operator with the ground and excited state wavefunctions over the entire range of φ . The various components, along with the energy potential, are shown in Fig. 2.11 for three different values of applied flux. The crucial feature is observed at half-flux where the symmetry of the three components is even-even-odd, meaning that the product is odd. Integrating an odd function yields identically zero; thus quasiparticle loss disappears at half-flux due to a symmetry in the wavefunctions, agreeing with the quantitative calculation of probability amplitudes above. This T_1 feature, along with




Figure 2.10: Quasiparticles contribute to qubit relaxation by absorbing energy from the qubit while tunneling across the junction. This tunneling may occur as an electron (with probability amplitude $e^{+i\varphi/2}$) or as a hole (with probability amplitude $e^{-i\varphi/2}$). These probability amplitudes add to yield $\cos(\varphi/2)$ which vanishes at $\varphi = \pi$, where $\varphi = 2\pi\Phi_{\text{ext}}/\Phi_0$.

relaxation measurements, is discussed further in Chapter 7.

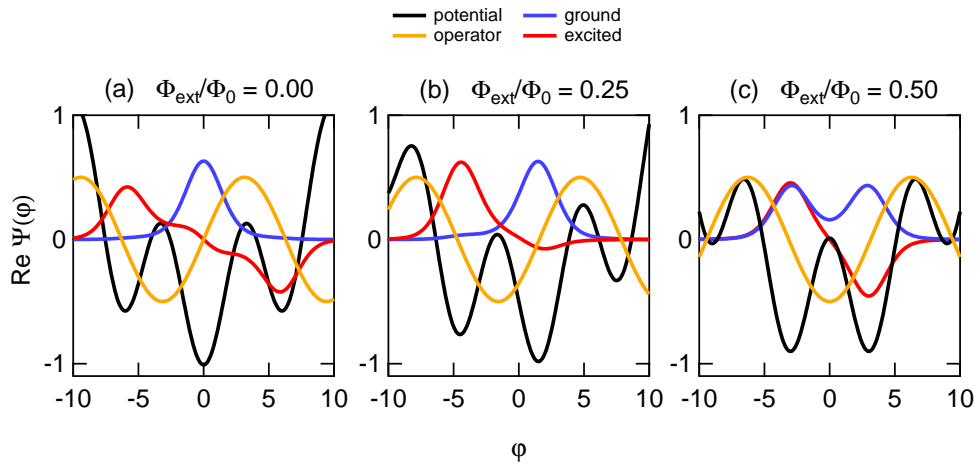


Figure 2.11: Components involved in calculation of quasiparticle matrix element $|\langle e | \sin(\hat{\varphi}/2) | g \rangle|$, along with the energy potential for three different values of applied flux. The matrix element is calculated by integrating over the product of the operator with the ground and excited state wavefunctions. Symmetry at half-flux yields a matrix element of zero, while evaluating to nonzero for every other flux.

Experimental Methods

3.1 Sample and Environment

3.1.1 Fabrication

Compact Resonators

The fabrication of compact resonators samples comprises of only one deposition step: photolithographic patterning of a Niobium film on a wafer of c-plane sapphire. The fabrication was done entirely at the wafer scale, utilizing 2 inch wafers throughout the process in order to produce many samples simultaneously. The experimental goal was to test the effect of different geometry, not different materials. Before metal deposition, the sapphire surface was prepared by a 60 s ion-milling using a 3 cm Kaufmann source that shoots 500 eV Argon ions at the wafer. The source operates at a flow rate of 4.25 sccm and a pressure of about 10 μ Torr, generating a current density of 0.67 mA/cm². A 200nm layer of Nb was then dc magnetron sputtered on the wafer. Photolithography was performed by patterning directly onto S1808 resist using a 365 nm laser. After development, the Nb was etched using a 1:2 mixture of Ar:SF₆ at 10 mTorr for 3 minutes. The wafer was then diced into individual chips

for measurement. Each sample consisted of a 2 mm x 7 mm chip with a coplanar waveguide (CPW) transmission line straight down the center with six resonators spaced along the feedline; each wafer contained 84 samples.

Compact Transmon

Compact transmons require two fabrication steps, corresponding to the resonator and qubit. The first step consists of photolithography to pattern Niobium on c-plane sapphire as described in the previous section. The second step utilizes e-beam lithography to produce the transmon capacitors and Al/AlO_x/Al. The junctions were fabricated using the Niemeyer-Dolan bridge technique [64, 65], which forms a bridge of resist, and double-angle evaporation to create an overlap between two layers of Aluminum, with an oxidation step between to create the junction. This technique is standard and has been used and described many times before.

The process for transmon fabrication is as follows. The wafer is first cleaned in NMP, acetone, and methanol, 3 minutes each in an ultrasonic bath. Microchem EL-13 copolymer is spun onto the wafer at 5000 RPM for 90 s and baked at 175 °C for 1 minute. The second resist layer, Microchem A-3 PMMA, is spun onto the wafer at 4000 RPM for 90 s and baked at 175 °C for 30 minutes. To prevent substrate charging during e-beam writing, a 13 nm Al discharge coating layer is then deposited. Depending on the sample, the e-beam was either written on an FEI XL40 SEM with a Nanometer Pattern Generation System (NPGS) add-on from J.C. Nability Lithography Systems using a 30 keV electron beam or a Vistec Electron Beam Pattern Generator (EBPG 5000+) using a 100 keV electron beam. The Al discharging layer is then stripped using MF312 for 90 s, and the resist is developed in MIBK:IPA (1:3) for 50 s and rinsed in IPA for 10 s. It was found that the temperature of the developer must be 25 ° C ± 1° C for reproducible results.

Junction deposition, done with a Plassys MEB550s, includes first a critical ion-beam cleaning step with an Anatech Argon ion gun operating at 250 V for 30 s. Aluminum deposition is done at 10^{-8} Torr at a rate of 1 nm/s; the first layer 25 nm at 0° followed by an oxidation at 15 Torr for 12 minutes. The second Aluminum layer is deposited at the same rate, but at an angle of 35° and to a thickness of 60 nm. A last oxidation is performed at 3 Torr for 10 minutes in order to oxidize the aluminum in a controlled environment as opposed to upon venting. The remaining resist and excess Aluminum liftoff is performed by soaking the wafer in acetone at 70°C for about an hour, making sure to place the wafer in a vertical orientation to prevent re-adhesion. After a final rinsing with water, sonication in acetone and rinsing with IPA, the wafer is diced into individual chips for measurement. As with compact resonators, each sample consisted of a 2 mm x 7 mm chip with a CPW feedline straight down the center, with six resonators spaced along the feedline. Each compact transmon device included one transmon placed inside one of the compact resonators. A total of 11 wafers were fabricated between the compact resonator and compact transmon experiments, each containing 84 samples.

3D Transmon

The fabrication of 3D transmon is simpler because the cavity is machined out of a block of copper or aluminum and is not lithographically defined. The fabrication of the junction of a 3D transmon is a single step of double-angle evaporated aluminum, with patterning defined by an electron-beam in much the same way as the compact transmon. In the time between the development of the compact transmon and the 3D transmon, a new e-beam machine and a new evaporator were installed, and both new machines are used for the 3D transmon.

The wafer used for the 3D transmon was $430\ \mu\text{m}$ c-plane sapphire. First the

substrate is cleaned via sonication in acetone for one minute and then 5 minutes in an oxygen plasma at 300 mBar and a power of 300 W. The wafer is then soaked in heated NMP at 90 °C for 10 minutes, sonicated in NMP and rinsed with acetone and methanol. Microchem EL-13 copolymer is spun onto the wafer at 2000 RPM for 100 s and baked at 200 °C for 5 minutes. The second resist layer, Microchem A-4 PMMA, is spun onto the wafer at 2000 RPM for 100 s and baked at 200 °C for 5 minutes. A gold coating was then deposited on the surface of the resist to aid in discharging of the substrate during e-beam writing (note that aluminum was used for the compact transmon). This gold layer was deposited using the Cressington Sputter Coater 108 for 45 seconds with Argon flow adjusted for 0.08 mBar and current at 30 mA, resulting in a 10 ± 1 nm gold layer.

The device is then written using the Vistec Electron Beam Pattern Generator (EBPG 5000+) using a 100 keV electron beam. The development is started by soaking in a potassium iodide/iodine solution for 10 seconds to remove the gold discharge layer. After being rinsed in water, the resist is developed in a 1:3 IPA:water mixture at 6 °C for 1 minute. Ultrasound is then turned on for 15 seconds, and the wafer is left in the mixture for 15 seconds after that.

Metal deposition is then performed in the Plassys UMS300 UHV e-beam multi-chamber e-beam evaporation system. First, the exposed surface is cleaned with an oxygen/argon plasma for 30 s to prepare a good surface for metal deposition. A titanium sweep is then performed to absorb all of the remaining oxygen, then the first layer of 30 nm of aluminum is deposited onto the substrate. The first layer is then oxidized with a 15% oxygen, 85% argon mixture for 10 minutes at 100 Torr. After that, the second layer of 50 nm aluminum is deposited. A final oxidation is then completed to ensure the aluminum does not oxidize upon removal from the vacuum system, for 10 minutes at 15 Torr.

Fluxonium

The materials and lithography of the fabrication recipe for the fluxonium artificial atom are exactly the same as the 3D transmon. The differences are solely in the design files for the e-beam patterning machine; in fact, the 3D transmon from Chapter 6 was made on the same wafer as preliminary fluxonium samples. The fabrication of junctions utilizes a bridge-free lithography technique [66, 67] that enables junctions to be of arbitrary size. This technique relies on a two-layer resist stack, with the top layer defining the structure and the undercut of the bottom layer providing selective removal of wires. By creating opposing undercuts on opposite sides of a junction, wires can be attached to both the top and bottom junction layer.

3.1.2 Octobox Sample Holder

A sample holder will, ideally, rigidly hold the device under test, allow microwaves to couple strongly to the device, shield the device from various types of radiation, all while not influencing the performance of the device. This, in practice, is very difficult and many sample holders are found to be flawed after many years of use. Developed by the group of Rob Schoelkopf, the “octobox”, shown in Figures 3.1 and 3.2, was a popular sample holder, and it was used for all experiments involving compact resonators and compact resonator transmons.

Performing measurements of on-chip devices requires methods to redirect microwave from coax lines to planar structures. The octobox uses coplanar waveguide (CPW) geometries to guide the microwaves on the PCB and on-chip. Therefore, we need a method for matching coax lines which are used at room temperature and in the dilution refrigerator to the CPW strips on-chip. Making this transition well is part of the challenge in designing a good sample holder. The octobox is built around




Figure 3.1: Octobox sample holder pieces, open. From left to right, the objects are: PCB with two devices mounted, base of octobox, shim of octobox, lid of octobox.



Figure 3.2: Octobox sample holder, closed. The eight holes of the top of the lid can be seen. The connectors on the PCB can be seen in four of the holes. The other four holes are plugged to prevent additional infrared radiation from reaching the sample.

the Rosenberger coax SMP-to-microstrip launcher, part number 19S102-40ML5.

Another main concern for the performance of a sample box is the influence of “box modes” on the device. Box modes are defined as extra, possibly unintended electromagnetic modes of the sample box. These modes are typically “cavity” modes, as explained in Section 2.1.2, modified by the presence of the PCB and sample chip. These box modes typically have low quality factors due to the participation of lossy materials and partially unsealed sample boxes. There are two methods to avoid detrimental coupling either resonators or qubits to box modes: 1) push the box modes to high frequencies, or 2) make the box modes higher quality. The latter is the route taken by the 3D transmon; the box modes are intentionally designed to be of higher quality and in fact are used for qubit readout. The former is the method chosen by planar devices using the octobox. The octobox includes a “shim”, shown in Fig. 3.1 that reduces the effective size of the box immediately around the chip and thus raises the frequency of box modes well above device frequencies. In fact, no box modes were observed below 12 GHz.

3.2 Qubit Readout

The efforts to improve relaxation and readout fidelity are often at odds. Relaxation is often improved by increasing the isolation between the qubit and an offending noise source, yet readout mechanisms consist of a strong coupling between the qubit mode and a measurement load. In circuit QED [23, 58], this coupling is performed through a cavity and readout is achieved by measuring cavity parameters that depend on the qubit state.

3.2.1 Dispersive Readout

Dispersive readout is the most common and original readout used for transmon and fluxonium qubits. Derived initially along with circuit QED, dispersive readout relies on a coupling that mixes the qubit and cavity states. The name dispersive is from the “dispersive” coupling regime, where the coupling strength is small compared to the frequency detuning is small, i.e. $g/\Delta \ll 1$. In this regime, the dressed qubit and cavity states are able to be separated into “mostly-qubit” and “mostly-cavity”, whereas at larger g/Δ the difference is less clear. The Hamiltonian of the qubit-cavity system in circuit QED, derived in Section 2.2.1 is shown in Eq. 3.1, with energy levels in Eq. 3.2, where N_C and N_Q represent the number of excitations in the cavity and qubit, respectively.

$$H = \hbar\omega_C \left(A^\dagger A + \frac{1}{2} \right) + \hbar\omega_Q \left(B^\dagger B + \frac{1}{2} \right) + \hbar\alpha (B^\dagger B)^2 + \hbar\chi (A^\dagger A B^\dagger B) \quad (3.1)$$

$$\frac{E_{N_C, N_Q}}{\hbar} = \omega_C N_C + \omega_Q N_Q + \alpha N_Q^2 + \chi N_C N_Q \quad (3.2)$$

Here, the linear cavity is characterized by the number of excitations N_C producing an energy $\hbar\omega_C N_C$, while the anharmonic qubit is represented by a number of excitations N_Q and an anharmonicity α . Without a coupling term, this Hamiltonian would simply consist of two separate oscillators, of little research interest; however, the addition of the coupling term $\chi N_C N_Q$ produces many interesting effects. The terms in the energy spectrum can be re-arranged in two ways, shown in Eq. 3.4 and 3.3. The first rearrangement shows that this coupling makes the qubit frequency dependent on the number of excitations in the cavity: $\omega_Q \rightarrow \omega_Q + \chi N_C$. This phenomenon is

called photon number splitting, and was studied in several papers [23, 68].

$$E_{N_C, N_Q} = \hbar(\omega_C) N_C + \hbar(\omega_Q + \alpha N_Q + \chi N_C) N_Q \quad (3.3)$$

$$E_{N_C, N_Q} = \hbar(\omega_C + \chi N_Q) N_C + \hbar(\omega_Q + \alpha N_Q) N_Q \quad (3.4)$$

The second re-arrangement highlights the use of the coupling as a readout since the cavity frequency becomes dependent on the number of excitations in the qubit: $\omega_C \rightarrow \omega_C + \chi N_Q$. Thus the cavity frequency when the qubit is in the ground state ω_C is different by from the cavity frequency when the qubit is in the excited state $\omega_C + \chi$, hence χ is typically called the “dispersive shift”. The shift is compounded for each additional qubit excitation. This effect is the same as in the Jaynes-Cummings Hamiltonian (see Section 2.2.1, with the added ability to access higher qubit states. This frequency shift allows for qubit readout via measuring cavity parameters that depend on this change. This cavity parameter is typically either microwave reflection off S_{11} or transmission through S_{21} the cavity, depending on the type of cavity and its coupling to the external microwave line. For the case of transmission through a cavity, Fig. 3.3 highlights the effect of the dispersive shift on cavity transmission parameters.

Qubit readout can be achieved at any frequency shown in this figure, but the contrast and the signal-to-noise ratio (SNR) will vary widely. Readout at either of the resonant frequencies, $f_C^g = \omega_C/2\pi$ or $f_C^e = (\omega_C - \chi)/2\pi$ results in a large amplitude difference and some phase difference. Maximum phase difference (with exactly no amplitude difference) is observed between the two resonant frequencies. While either amplitude or phase could strictly be used for qubit readout, one must use both in the form of the signal quadratures I and Q in order to be quantitative.




Figure 3.3: Amplitude and phase of cavity transmission S_{21} when qubit is in ground state (blue, $f_C^g = \omega_C/2\pi$) or excited state (red, $f_C^e = (\omega_C - \chi)/2\pi$). Readout can be performed at any frequency there is a difference of either transmitted amplitude or phase between ground and excited qubit states. Parameters for the curves shown are chosen to be typical: $f_C^g = 8.02$ GHz, $f_C^e = 7.98$ GHz, $\chi = 40$ MHz, $\kappa = 24$ MHz.




Figure 3.4: Expected readout signals plotted on the I - Q plane at three readout frequencies as noted, for ground (blue) and excited (red) qubit states. Readout parameters are the same as Fig. 3.3. Each circle represents a cloud of possible measurements consisting of a Gaussian centered on the circle. More averaging or lower noise temperatures decreases the width of these circles. Completely overlapping circles indicate the qubit states are not distinguishable, while far separated circles indicate high distinguishability.

Since the quadratures form an orthogonal basis and are linearly proportional to the measured voltage (see Section 3.2.3, measurement noise is Gaussian in the signal quadratures; noise is not treated properly in amplitude or phase. Qubit readout is easier to visualize when plotted as a cloud of measurements in the I - Q plane. Shown in Fig. 3.4 are readout signals for each of the three readout frequencies highlighted in Fig. 3.3: f_C^e , $(f_C^e + f_C^g)/2$, and f_C^g . Here, each circle represents a cloud of possible measurements consisting of a Gaussian centered on the circle. For these parameters, the readout SNR is higher the frequencies corresponding to the high amplitude peaks. While the phase difference is maximum for the case $(f_C^e + f_C^g)/2$, the lower amplitude reduces fidelity. However, in the case $\chi = \kappa$, the maximum readout fidelity is obtained at this middle frequency [69].

With optimized qubit-cavity systems and readout chains utilizing parametric amplifiers, high fidelity single-shot dispersive readout measurements have been performed by several groups [70, 71, 72, 69]. Even though the qubit state after a readout is different than the qubit state before readout, i.e. there exists a back-action, the

effect is calculable if the readout is of high enough fidelity [69], and thus dispersive readout is said to be QND.

3.2.2 High Power Readout

In addition to the well-defined and understood dispersive readout, the so-called “High Power Readout” [73] (occasionally called the “REED-out” in homage to the discoverer Matt Reed), allows a much higher readout fidelity at the expense of qubit demolition and time. Looking at the plot of cavity spectroscopy versus power, one sees that the cavity resonant frequency peak moves from 9.071 GHz at < -20 dBm to 9.055 GHz at > 5 dBm. At high powers, the cavity is observed at a frequency corresponding to the “bare” cavity frequency, i.e. the frequency not affected by the dressing of the qubit. Circuit QED theory predicts a small self-Kerr for the cavity, and thus a linear dependence of frequency on applied photon number (see Section 2.2.1). However, more detailed models show a diminishing cavity anharmonicity at high powers, and thus the frequency is stable above a critical power, matching qualitatively with what is observed [74]. The observed transmission amplitude of this “high-power” peak is much higher than that of the “low-power” peak and thus the cavity is said to be in the “bright state”. This readout works by mapping the qubit state onto easily-distinguishable classical states of the qubit-cavity system. All measurements shown in this Section are on the 3D transmon qubit discussed in Section 6.1.4.

The readout mechanism is derived from the observation that the critical power to reach the bright state depends on the initial qubit state; this effect is highlighted in Fig. 3.6. This figure shows how the transmission amplitude jumps from a low level (off-resonant) to a high level (on resonant, bright state) at a different power for the prepared ground and excited qubit states. Thus, any power in the region between

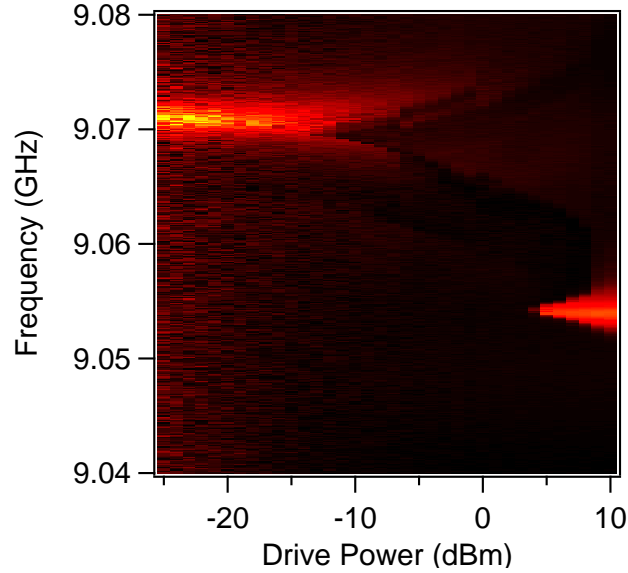


Figure 3.5: CW cavity spectroscopy versus drive power. At low powers, cavity response is a Lorentzian peak centered at 9.071 GHz, the dressed cavity frequency. At high power, this peak shifts to 9.055 GHz, corresponding to the bare cavity frequency.

the rises, where there is a large difference between ground and excited transmission amplitudes, is a possible choice of readout power. The readout procedure is thus the same as with dispersive readout, simply with a different frequency and power; no additional drives, tuning, or equipment of any kind is necessary.

Histograms of single-shot measurements, i.e. amplitudes from a single non-averaged measurement of the transmission amplitude, are shown in Fig. 3.7 at a power between the two critical powers for high power readout. This includes 25,000 alternating measurements of the qubit in the ground state and after a π -pulse. The main shape of each curve is composed of two Gaussians. In each case, one of the Gaussians is dominant, and the other is due to either poor state preparation due to high qubit temperatures (see Section 6.1.4) or decay during the measurement [75]. The two dominant Gaussians are separated by more than their width, indicating that the qubit state may be discriminated well from one measurement. Single-shot

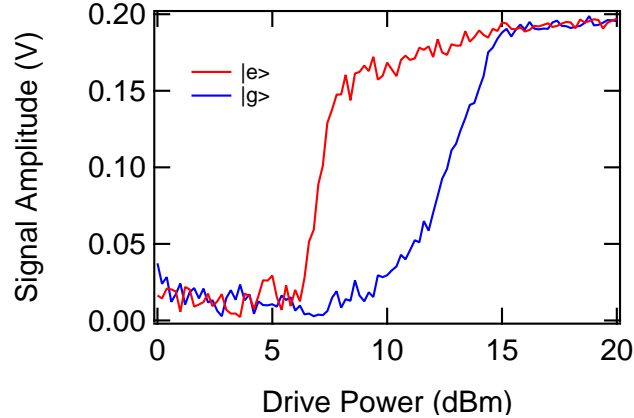


Figure 3.6: Normalized readout voltage as a function of drive power for prepared ground and excited states. The response switches to the “bright state”, i.e. higher voltage state, at a critical power that depends on qubit state. High power readouts can be performed at any power with a difference between ground and excited state transmission amplitude.

measurements of this qubit using the dispersive readout would yield Gaussians directly on top of each other, states could be distinguished only after many averages. Neither the readout chain nor qubit were optimized for readout, and much higher fidelities are possible.

High power readout is non-QND, meaning that the qubit state before readout is different from the qubit state after readout in a manner that erases information. In fact, the effect of high power readout is more complicated than simply projecting the qubit to the ground or excited state. Attempting to apply qubit rotation pulses shortly after high-power readout shows that after an excited state measurement (bright state), the qubit is unresponsive to rotation pulses. After a ground state readout, no such effect is observed. Thus, when the cavity enters the bright state, the qubit must leave the computational basis, only to return after the cavity decays from the bright state. Fig. 3.8 shows that after a measurement of the excited state, a following measurement at varying times shows that the qubit returns to the ground state with a time constant, in this case, of $100 \mu\text{s}$. Since T_1 for this qubit was




Figure 3.7: Histograms of single-shot measurements of thermal state and inverted state. Each plot is fit to two Gaussians, one centered on the voltage corresponding to the ground state, and one corresponding to the excited state. The two main peaks are separated by more than their widths.

40 μs , this shows that the effect of high power readout is long-lasting and must be considered when figuring an experimental repeat rate.

When comparing these two readout methods, there are situations when one is preferred over the other, although high single-shot fidelities are possible with both high power readout [73] and low power dispersive readout [69]. In the case where feedback is performed, it is preferable to use dispersive readout with its minimal (and possibly predictable) disruption of the qubit state. Alternatively, in the case where either a fast measurement is preferable or dispersive readout produces only low fidelity, the higher power readout may be desired.

3.2.3 Heterodyne Measurement

Both dispersive readout and high power readout rely on the same experimental setup. Each measurement relies on measuring the microwave transmission (or reflection) properties of a given system. Generating microwave pulses is easy via the use of commercial microwave generators, but accurately detecting a microwave signal's am-

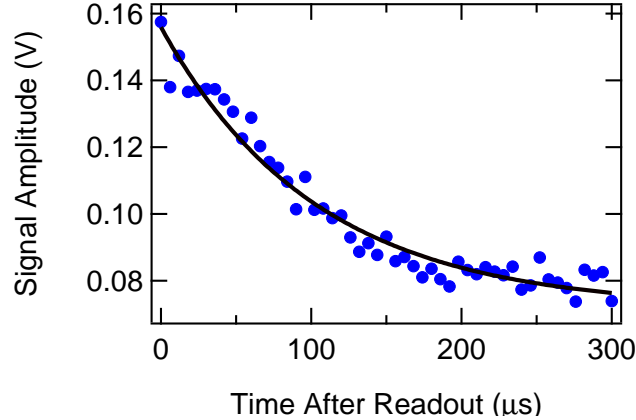


Figure 3.8: Measurement of a high power readout following a high power readout of the excited state (higher voltage). The qubit state decays back to the ground state with a time constant of $100 \mu\text{s}$, longer than the T_1 of $40 \mu\text{s}$.

plitude and phase is more difficult. While a vector network analyzer performs this duty in CW mode, and works great for resonator measurements, qubit measurements require certain pulse sequences that a network analyzer is unable to accomplish. For qubit measurements performed in this thesis, heterodyne measurements were used as a means of converting the signal into something that could be digitized with existing commercial digitizers.

The schematic of a heterodyne interferometric experiment is shown in Fig. 3.9. Three microwave waveform generators are shown, the qubit (ω_Q), cavity (ω_C), and LO ($\omega_{LO} = \omega_C + \omega_{IF}$) generators. The qubit generator produces microwaves at the qubit frequency, mixed with pulse envelopes from an arbitrary waveform generator (AWG) to create pulses to perform qubit rotations. The cavity and LO generator are used for the readout, with the cavity generator set to the interrogation frequency and the LO generator shifted by the IF frequency, $\omega_{IF} = 20 \text{ MHz}$. The cavity signal (either CW or pulsed) is split into two paths making the interferometer, one that goes through the device, and the other does not; each path ends in the RF input of a microwave mixer. The LO generator is split to feed the LO signal of both RF

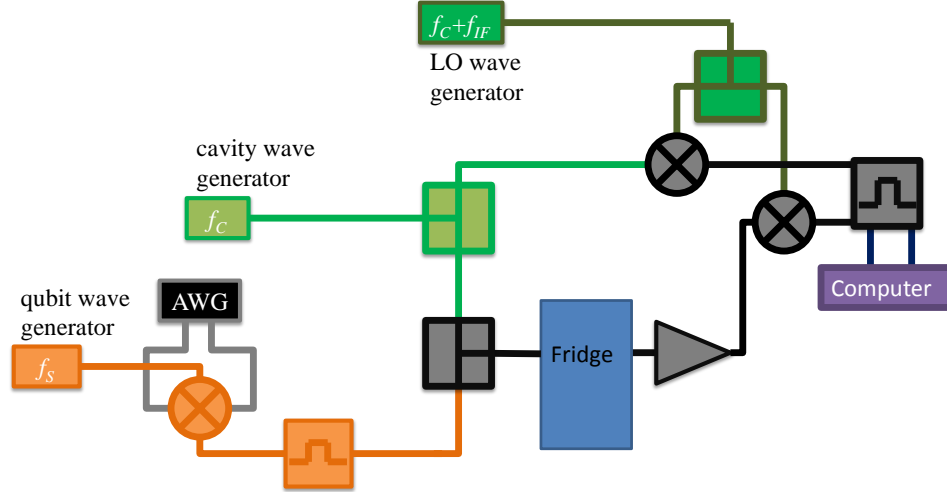


Figure 3.9: Cartoon schematic of experimental setup to perform heterodyne measurement, involving an interferometric measurement, which compares a microwave signal going through the device under test with a signal bypassing the device. Two microwave generators (cavity and LO) are mixed together after the two paths in order to produce a lower frequency tone at the difference frequency ω_{IF} that can be digitized in the computer. The additional (qubit) microwave generator can be used to stimulate the device and the effect on the cavity transmission can be measured.

mixers. This mixing operation produces signals at the sum and difference frequencies, $\omega_C + \omega_{LO}$ and $\omega_C - \omega_{LO}$. Additional filters remove the high-frequency component, and the ω_{IF} component is sent to a digitizer. The digitizer operates at 1 GS/s, and can effectively measure amplitude and phase of a signal at 20 MHz.

3.3 Equipment

3.3.1 Dilution Refrigerators

At room temperature, our resonators and qubits reduce to simple classical electrical circuits that can be decomposed into capacitances, inductances, and (unintentionally) resistances. In order to thermally prepare their quantum ground state, the devices must be cooled such that the thermal energy is below that of the transition

energies (ω_0) we wish to probe. For the case of superconducting qubits and resonators, this frequency is on the order of several GHz. In order for there to be an appreciable difference in ground state and excited state population, the system must be cooled to a temperature T such that the inequality $k_B T \leq \hbar \omega_0$ is well satisfied. This means that T must be less than roughly 500 mK for $\omega_0 = 10$ GHz, and 50 mK for $\omega_0/2\pi = 1$ GHz.

There are simple techniques for reaching temperatures of approximately 1.2 K: liquid Helium has a boiling temperature of 4.2 K, and thus can be used to cool equipment and samples down to this temperature. By pumping on a vessel of liquid Helium, one can reduce this temperature slightly to roughly 1.2 K by reducing the vapor pressure and forcing the liquid to boil. To reach temperatures below 1.2 K, more sophisticated equipment, a dilution refrigerator, is required.

Fig. 3.10 shows a generic schematic of a typical “wet” dilution refrigerator. This type of fridge is called “wet” because the dilution insert sits in a bath of liquid Helium as a means to cool the entire apparatus to 4.2 K. As this Helium evaporates, it must be re-filled, constituting the largest recurring cost of operating a dilution refrigerator. “Dry” dilution refrigerators use a closed circuit pulsed-tube that continually compresses and expands a fixed amount of Helium gas, thus cooling the screens around the dilution unit without the need for a liquid Helium bath. The experiments described in this thesis were performed in two “wet” refrigerators.

The cooling mechanism of the dilution refrigerator consists of a continually circulating circuit of ^3He and a stationary bath of ^4He . The ultimate cooling of the refrigerator is performed by the dilution of ^3He into ^4He in the so-called mixing chamber [76]. The rest of the circulation process is solely meant to facilitate this mixing. ^3He is removed from the mixing chamber by way of the still. The still is maintained at approximately 700 mK by balancing the heat loads on the various components.




Figure 3.10: Cartoon of the dilution unit of a dilution refrigerator showing the major components, from the Oxford Instruments Kelvinox operating instructions manual. ^3He is circulated through a closed loop to achieve cooling by continuously diluting into ^4He in the mixing chamber. A “wet” dilution refrigerator cools the dilution unit by immersion in a liquid Helium bath, and a “dry” dilution refrigerator cools the screens around the dilution unit with a pulse tube.

The vapor pressure of ^3He is much higher than ^4He at this pressure, so pumping on the still removes nearly pure ^3He . This ^3He is then re-liquified upon contact of an auxiliary refrigerator. This may be either a “1K pot” or a Joule-Thompson stage: a 1K pot is a small vessel of liquid Helium maintained at approximately 1.5 K by constant pumping. The ^3He continues to cool from contact in various efficient heat exchangers with helium from the opposite side of the cycle. The ^3He eventually reaches the mixing chamber and contributes again to the overall cooling by diluting from a nearly pure ^3He phase into the ^4He that filling the mixing chamber. Heat is needed for this dilution to occur, so forcing this process requires heat to be extracted from the environment, thus cooling whatever is in contact with the mixing chamber. The ^3He then gets removed from the still and begins the cycle again.

Commercial dilution refrigerators typically come with extra space below the mixing chamber on which we can mount our devices to measure. This space can vary from a few cubic inches to a few cubic feet. Two refrigerators were used in the experiments outlined in this thesis, both of the Kelvinox brand by Oxford. One is a Kelvinox25 with a nominal cooling power of $20 \mu\text{W}$ at 100 mK, the other is a Kelvinox400 with a cooling power of approximately $400 \mu\text{W}$ at 100 mK. The base temperatures according to the installed mixing chamber temperature sensors are 75 mK for the Kelvinox25 and 15 mK for the Kelvinox400. The dilution units of these two refrigerators are shown in Fig. 3.11.

3.3.2 Microwave Switches

The introduction of microwave switches was a helpful innovation that allows many samples to be measured in the same refrigerator at the same time. Without the use of a switch, each sample required its own input and output microwave lines, along with a complete set of filters, circulators, and amplifiers. This requires a lot of space





Figure 3.11: Images of dilution units of Kelvinox25 and Kelvinox400 dilution refrigerators. The Kelvinox25 refrigerator has a nominal cooling power of $20 \mu\text{W}$ at 100 mK and an operating base temperature of approximately 75 mK, while the Kelvinox400 refrigerator has a cooling power of approximately $400 \mu\text{W}$ at 100 mK and an operating base temperature of approximately 15 mK. The experimental space on the Kelvinox400 is much larger than the Kelvinox25.

inside the experimental area of the dilution refrigerator and prohibits the mounting and measurement of more than a couple of samples during a single cooldown. Since experiments may require the measurement of dozens of devices, this either requires many cooldowns to measure each sample, or a prohibitively expensive (and creatively packed) set of microwave lines to measure multiple devices. An alternative is to install a microwave switch that will allow the user to use the same input and output microwave lines to measure multiple devices. While the switch may only select one sample at a time, it can be used to switch between samples without warming up the refrigerator, and thus saves a lot of time.

The switch used is a SP6T electromechanical switch from Radiall, serial number R573423600, shown in Fig. 3.12. SP6T means that there is one input line that can be routed into any of the six outputs. To operate the switch, one sends a short current pulse, of order 200 mA to one of the six coils. This current pulse induces a magnetic field on one of the actuators, which either makes or breaks a microwave connection inside the switch. This mechanism allows more than one output to be selected simultaneously, however there will then be mismatches in impedance and severe reflections will be induced, thus only one output should be enabled at any given time. Since each coil and actuator is independent, the user is required to switch “off” outputs that are not intended to be used.

These switches are made by Radiall to work near room temperature and do not nominally work at low temperatures. We make a few minor modifications to the switch for low temperature usage. First, we remove the blue aluminum shield, as shown in Fig. 3.12. Then we remove the PCB which sits on top of the exposed connectors. Then we reinforce the exposed connectors with Stycast®(a type of epoxy glue) because they are very fragile. The coils in the switch to create the magnetic field are made from copper, which is resistive, and thus the switch dissipates heat when



Figure 3.12: Radiall SP6T switch, modified to work at low temperatures inside of a dilution refrigerator. On the left is the switch, with PCB removed and connectors glued in place. On the right is the body of the switch, removed to show the details inside.

switched. Since the resistance of copper drops significantly at low temperatures, this amount of heat is compatible with and does not disrupt the dilution refrigerator significantly. Depending on the efficiency of the thermal connection to the mixing chamber, the mixing chamber temperature may rise from 15 mK to 30 mK, requiring about 20 minutes to re-cool. This minor inconvenience is greatly preferred over having to warm up the entire fridge to switch samples.

3.3.3 Shields

During the time of this thesis, details in experimental setups have been blamed for limiting relaxation and coherence [77, 43]. Details about specific results and measurements are given in later chapters. Each time, adding some type of “shield” improved coherence times. Here, the term “shield” is used broadly, and serves to include both shields absorbing radiation from outside the microwave lines, and those absorbing radiation from inside the microwave lines. The latter are often called filters, as they absorb certain frequencies while allowing signal frequencies to pass through. Each type of shield used in this thesis will now be detailed.

Circulators

A circulator is a 3-port device that uses Faraday rotation and interference to create a non-reciprocal transmission matrix. Waves launched into ports 1, 2 or 3 will come out of ports 2, 3, or 1, respectively. Therefore, circulators can separate input waves from output waves and function as “diodes”. In this thesis, they are used solely with a 50Ω termination on one port, turning the circulator into an isolator which allows microwaves only from port 1 to 2, while absorbing waves that enter port 2. Alternatively, isolators can be purchased directly, but modifying circulators in this way produces a well-thermalized load.




Figure 3.13: Quinstar 8-12 GHz cryogenic circulator. In band, this device routes input microwaves to the next port in a clockwise fashion ($1 \rightarrow 2 \rightarrow 3$). This circulator can be turned into an isolator by placing a 50Ω termination on one port, then allowing microwaves to travel in one direction.

Isolators were used in every experiment in this thesis in order to impose a well-thermalized input state on the next element in the chain and to block noise from the HEMT amplifier from reaching the sample, backwards down the readout microwave line. This noise coming out the input of the HEMT is unknown, not specified by the producer, and likely varies from amplifier to amplifier. Since the device is active and it is known to dissipate energy, it is safest to prepare for the worst since devices can be heated or driven by noise at their resonant devices. Isolators are used to allow the signal to pass through the device to the HEMT, but to attenuate the noise coming back down from the HEMT. A typical cryogenic circulator from Quinstar (previously Pamtech) is shown in Fig. 3.13.

Quinstar circulators come with several possible bandwidths, Fig. 3.13 shows a 8-12 GHz circulator. The lower frequency circulators will be considerably larger, as the dimensions scale directly with the wavelength. Quinstar sells wider bandwidth isolators, but the dimensions is also considerably larger and limited by the lowest frequency in the band. While the circulator frequency is chosen to surround the main device frequencies, the performance out of band is not guaranteed and in fact attenuation in the reverse direction is observed to decrease to only a few dB at



Figure 3.14: K&L low-pass microwave reflection filters. They pass microwaves below a cutoff frequency (the filter shown here is 12 GHz) and are specified to reflect all other incoming frequencies up to 26 GHz.

frequencies in the 20 GHz range. Without equipment to measure far above these frequencies, it is really unknown what transmission is achieved, but it is not likely to be attenuated significantly. The noise out of band is unable to directly excite transitions in the circulator band, but can still interfere in the operation of qubit devices. The addition of K&L filters and Eccosorb filters is meant to help reduce the higher frequency noise.

K&L Low-pass filters

The purpose of K&L filters is to filter out unwanted microwave noise from outside the band of the circulators. The cutoff frequency of these low-pass reflective filters is chosen to coincide with the highest frequency of the chosen circulator. For 8-12 GHz circulators, a 12 GHz K&L filter is matched. Thus, above 12 GHz, where the circulator begins passing noise in the reverse direction, the low-pass filter will reject these frequencies. These devices are specified to have very low insertion loss (< 1 dB in band) as to not reduce the signal amplitude, yet guarantee rejection of at least 50 dB from 16 GHz to 26 GHz. At higher frequencies, their properties are not guaranteed, and this rejection must break down at some frequency. To help at even higher frequencies, Eccosorb filters are added.

Eccosorb Filters

Eccosorb filters are placed in a microwave line in order to prevent high frequency radiation from reaching the sample, in a manner similar to filters of Santavicca et al. [78]. The source of this radiation may be higher temperature components or HEMT amplifiers. In addition, radiation may enter the microwave line through loose or improperly soldered connectors. Eccosorb filters achieve a balance between passing signal frequencies and attenuating high frequencies by producing an attenuation that depends linearly with frequency. Eccosorb is optically opaque, so this attenuation continues at least to optical frequencies. The operating principle of this type of filter is to force the radiation to travel through a lossy material in a small section of the transmission line. Eccosorb is a proprietary material produced by Emerson & Cuming for the purpose of absorbing microwaves. Eccosorb CR-110 was chosen for this filter because it has the least attenuation below 18 GHz. This material has a roughly linear attenuation versus frequency curve; higher frequencies are attenuated more than lower frequencies.

The box-type Eccosorb filter consists of a short section of microstrip, in which the box serves as the ground plane and a thin metallic strip is used for the microstrip trace. Two SMA connectors are placed in the two opposing holes of the box, then a metallic strip is cut to the size corresponding to that of a 50 Ω microstrip trace. For the box size used here, the microstrip trace is cut out of 0.01" thick copper, aiming for a width of 0.090". The microstrip is then soldered in place onto the centerpins of the two SMA connectors. The Eccosorb CR-110 is then mixed (100 parts X with 12 parts Y by weight) and poured into the box. The lid is then secured and the whole filter is cured at an elevated temperature (70-90 °C depending on cure time); the filter is complete when curing is finished.




Figure 3.15: Different types of Eccosorb filters. (a) Box filter and (b) drawing of cross-section of box filter. The box filter is a short section of microstrip in which the dielectric is the absorbing material Eccosorb. (c) pipe filter and (d) drawing and cross-section of pipe filter. The pipe filter is a short section of coax in which the dielectric is Eccosorb.

The pipe-type Eccosorb filter consists of a short section of coax made out of a pipe and wire. Two SMA connectors are placed on opposite sides of a hollow pipe and a wire is soldered from centerpin to centerpin using conveniently drilled holes in the pipe. The pipe and wire diameter are chosen to match the 50Ω characteristic impedance of the rest of the microwave line. The pipe is then filled with Eccosorb and cured in the same way as the box filter. Both the box and pipe filter, along with cross-section drawings are shown in Fig. 3.15

While both filters work and have been used in various experiments, the pipe filter has a few advantages. First, the most important advantage is that it generally achieves lower reflections because it does not rely on changing from coax to microstrip. Second, subjectively, it is easier to make; it does not rely on any part being machined, one simply has to have the right pipe and wire diameter. Another advantage is that it is tunable. The attenuation can be varied by adjusting the length of the pipe. Lastly, it is much more compact in size and can fit in places where a box filter cannot.

Characterization data for two box filters and two pipe filters are shown in Fig. 3.16. The two box filters are nominally identical, while one pipe filter is significantly




Figure 3.16: Measured S_{21} and S_{11} of four Eccosorb filters: 2 box filters and 2 pipe filters. Each filter has an overall downward trend in transmission (S_{21}) as frequency increases. With the exception of the short pipe, this trend is approximately linear. The matching (S_{11}) shows how close the filter is to the desired 50Ω impedance. The frequency of relevance is the signal frequency of the qubit or cavity, so typically between 5 and 10 GHz.

longer than the other. The differences between the two box filter responses highlight the irreproducibility inherent in this design. Both the transmission and matching have a wide scatter when making many filters. The attenuation differences are likely due to differences in curing or mixing parameters, while the matching is due to slight differences in centerpin construction. The matching for the box filter can be very poor around 10 GHz as shown by ‘Box 1’, while the pipe filters more consistently have matching around -20 dBm at 10 GHz.

While minimizing reflections (low S_{11}) is always important for a filter, the transmission (S_{21}) at signal frequencies is more critical when used between the sample and the output amplifier. When the filter is placed on the input line, which usually contains 60 dB of intentional attenuation, a few dB more attenuation due to an Eccosorb filter is unimportant. However, when the filter is placed on the output of the device, any attenuation at the signal frequency degrades the measurement quality, so it is important to minimize the attenuation at the signal frequency.

Copper Powder Filters

Copper powder filters are an alternative to the Eccosorb filters, and also aim to attenuate high frequency noise from reaching the device. Copper powder filters have been made, in a style similar to the Eccosorb pipe filters. With further details in the thesis of Nick Masluk [36], these copper powder filters consist of a short section of coax transmission line where the dielectric is replaced by copper powder mixed with epoxy. This copper powder passes low frequencies well, but provides attenuation at frequencies well above the frequency range of interest. An example of a copper powder filters is provided in Fig. 3.17: a semi-rigid cable is stripped of its outer conductor, and the wider pipe is soldered on. The volume of the pipe is then filled with copper powder; the width of the pipe is chosen to maintain 50Ω impedance to minimize reflections.

Copper powder filters have some advantages and disadvantages. Similar to the Eccosorb pipe filters, the attenuation is adjustable by tuning the length of the copper powder section. The high frequency performance and matching is equally as easy as the Eccosorb pipe filters, due to the construction methods. A disadvantage is that due to the use of a single coax, the length of a copper powder filter is much longer than the attenuating length. Also, the attenuation for a given length is less than that of an Eccosorb filter, so longer lengths are typically needed. However, this can be an advantage if really small attenuations are desired.

Infrared Shield

In addition to filtering the noise inside the microwave lines, it is important to filter out infrared photons from external sources. These sources are likely hot components at other stages of the dilution refrigerator, or 300 K photons directly from room



Figure 3.17: Picture of copper powder filter showing typical length for a filter with 2 dB loss at 10 GHz.

temperature components. Studied extensively by Barends et al. [77], this stray infrared light can be avoided by building a sufficiently “light-tight” shield around the sample box area. The infrared shield detailed here was based loosely upon this work. Additional work at IBM [79] showed that this infrared light can significantly affect qubit performance and is thus critical, especially for qubit experiments.

The infrared shield is an additional copper box surrounding the experimental area. The infrared shield is composed of two pieces, a plate and a can. The plate is mounted rigidly onto the mixing chamber and holds connectors through which all input and output microwave signals will travel. There are 8 connectors in this infrared shield, thus allowing for up to four samples. The connectors used to penetrate this plate were hermetic room temperature SMA flange connectors. Being designed for room temperature, the connectors use a rubber O-ring for sealing; this was replaced with a bead of indium, which is commonly used for hermetic seals at cryogenic temperatures. The can is a cylinder with a flange that mates on one end to the plate. An indium seal is used between the plate and the can, and the can is screwed shut rigidly to achieve a hermetic seal. The entire structure, closed and open, can be observed in Fig. 3.18.

In addition to preventing infrared from reaching the sample box, this structure serves another purpose. A sheet of absorbing material is placed inside the copper can




Figure 3.18: (a) Picture showing bottom of infrared shield plate and attached samples. Four 3D cavities are shown here. The eight SMA connectors can clearly be seen, the hermetic seal is on the top side of the plate. (b) Picture showing closed infrared shield, completely enclosing the experimental area. In both images, the top half of the cryoperm shield is visible.



Figure 3.19: Image showing inside of infrared shield can. An infrared absorbent coating was applied to a thin copper sheet and placed on the walls and bottom of the can. The remaining indium seal between the can and the plate of the infrared shield can also be seen here.

in an effort to further reduce infrared radiation from reaching the sample box. Even if infrared photons reach inside the can, they will quickly be absorbed on the walls, rather than reflecting many times. Fig. 3.19 shows that this absorbent coating is applied to a thin copper sheet which is then placed along the walls and bottom of the can. Applying this coating in this way allows different coatings to be easily tested. For the experiments in this thesis, the coating used was a lab recipe of a mixture of Stycast 2850 and carbon powder, containing 7% carbon powder by weight. Before curing, this compound can be spread onto a surface easily; after curing it is very hard, but brittle.

Compact Resonators

This chapter describes an experiment aimed at improving the quality factor of on-chip lithographic resonators. The idea behind this experiment was to vary the resonator geometry so as to minimize the participation of elements that were suspected to contribute to losses. Low-loss on-chip resonators are important for several reasons: a) they can be used directly in quantum information experiments, b) they can be used as a proxy to study loss in on-chip qubits (see Chapter 5), and c) on-chip resonators offer some avenues for the scalability of quantum information. This work was performed in both dilution refrigerators described in Section 3.3.1, but mainly the Kelvinox25, and includes data from 24 experimental runs and measurements of over 150 resonators. In summary, our results indicated that surface loss is a dominant factor, and by tweaking the resonator's geometrical design parameters, the maximum quality factor at an approximate temperature of 200 mK, increased from from 210,000 to 500,000. The median quality factor increased similar 160,000 to 380,000.

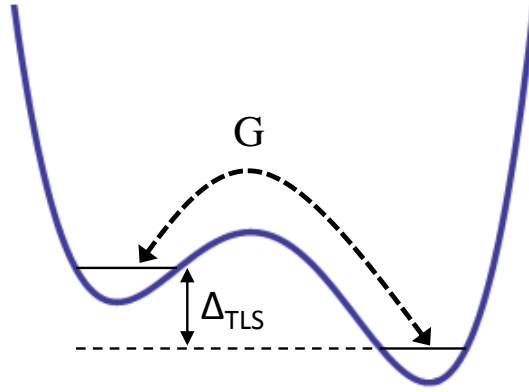


Figure 4.1: The model of a generic two level system in an amorphous solid. The system has two states, each corresponding to one of the two wells, with level asymmetry Δ_{TLS} and tunneling rate G .

4.1 Introduction

4.1.1 TLS Dissipation

Previous measurements have concluded that on-chip resonators are, in the low-power and low-temperature limit, frequently limited by a distribution of two-level systems (TLS)[80, 81, 51, 82]. These TLS have been theorized to originate from many sources, but in summary are attributed to some type of positional defect in amorphous materials [83, 84, 85]. A generic schematic of a TLS is given in Fig. 4.1, where a potential landscape has two separate wells, representing two states [86]. With an level asymmetry of Δ_{TLS} and a tunneling rate of G , the TLS energy is given by $\sqrt{\Delta_{TLS}^2 + G^2}$.

TLS loss is phenomenologically characterized by a few specific features. First, the dependence of the superconducting resonator frequency versus temperature deviates at low temperature [80, 81] from that expected from the dependence of the kinetic inductance on temperature, Mattis-Bardeen theory[87].

Mattis-Bardeen theory predicts a resonant frequency that asymptotically approaches a constant value at low temperature as the kinetic inductance (L_k) stabilizes far below T_c . The complex conductivity of a superconductor $\sigma = \sigma_1 - i\sigma_2$, where σ_1 characterizes the conductivity due to quasiparticles and σ_2 comes from the inertial response of the Cooper pair condensate, leading to a kinetic inductance $L_k \propto 1/d2\pi f\sigma_2$, where d is the film thickness [88]. The high frequency dependence of σ_2 leads to an increase in L_k as temperature rises, thus inducing a frequency reduction. The full form of the frequency dependence of σ_1 and σ_2 is given in Eqs. 4.1, where σ_N is the normal state conductivity [89].

$$\begin{aligned}\frac{\sigma_1(\omega, T)}{\sigma_N} &= \frac{4\Delta}{\hbar\omega} e^{-\Delta/k_B T} \sinh\left(\frac{\hbar\omega}{2k_B T}\right) K_0\left(\frac{\hbar\omega}{2k_B T}\right) \\ \frac{\sigma_2(\omega, T)}{\sigma_N} &= \frac{\pi\Delta}{\hbar\omega} [1 - 2e^{-\Delta/k_B T}] e^{-\hbar\omega/2k_B T} I_0\left(\frac{\hbar\omega}{2k_B T}\right)\end{aligned}\quad (4.1)$$

However, the observed data follows the frequency and temperature dependence given by Eq. 4.2, which contains both the effects of Mattis-Bardeen and TLSs. Empirically, the effect of TLSs is described through a dependence on the effective dielectric constant of participating materials due to a non-saturation of TLS energy levels only at the lowest temperatures, as shown in Eq. 4.3 [90]. In these two equations, ω_0 is the resonant frequency of the resonator, F_{TLS} is a fraction of electric field energy stored in TLS host material, ε is the effective dielectric constant, $\Delta\varepsilon$ is the change in effective dielectric constant as a function of frequency ω , α is the kinetic inductance fraction, σ is the imaginary part of the conductivity, and $\Delta\sigma(T)$ its change with temperature, Ψ is the complex digamma function, and δ characterizes the TLS loss tangent.

$$\frac{\Delta\omega_0(\omega_0, T)}{\omega_0} = \frac{\alpha}{2} \frac{\Delta\sigma(T)}{\sigma} - \frac{F}{2} \frac{\Delta\varepsilon(\omega_0, T)}{\varepsilon}\quad (4.2)$$

$$\frac{\Delta\varepsilon(\omega, T)}{\varepsilon} = -\frac{2\delta}{\pi} \left[\operatorname{Re}\Psi \left(\frac{1}{2} + \frac{1}{2\pi i} \frac{\hbar\omega}{k_B T} \right) - \log \frac{\hbar\omega}{k_B T} \right] \quad (4.3)$$

Another hallmark of TLS dissipation is shown in the dependence of Q_i on the circulating power in the resonator. Since Q_i is defined as the rate at which energy dissipates compared to the energy stored in the resonator, as more power is put into the resonator more power is dissipated. Since a TLS, by definition, has two energy levels, one with a finite lifetime, the relaxation rate of the TLS presents a maximum rate at which energy can leak from the resonator to an external bath through the TLS. The distribution of TLS is a bottleneck that becomes saturated at high resonator circulating powers. This bottleneck leads to an increase in Q_i at high circulating power since no more power can be dissipated through the TLSs. This effect is characterized by a dependence of Q_i on resonator electric field as shown in Eq. 4.4, where E is the RMS electric field value in the resonator, and E_s is the saturation electric field [51, 82], calculated from solving the Bloch equations for the TLS density matrix [86]. The $\tanh\left(\frac{\hbar f}{k_B T}\right)$ term accounts for partial saturation of TLS states due to temperature. For the purpose of this study, Q_i is measured at circulating powers corresponding to a single microwave photon, which is below the power at which Q_i begins rising noticeably. This is a typical power used for quantum information processes and thus it is important to characterize resonators at this power.

$$\frac{1}{Q_i} \propto \frac{\tanh\left(\frac{\hbar f}{k_B T}\right)}{\sqrt{1 + \frac{E^2}{E_s^2}}} \quad (4.4)$$

Measurements have revealed nonzero TLS filling factor F_{TLS} for CPW resonators, showing that they are coupled to some distribution of TLS. In addition, an experiment by Gao et al. [80] measured the variation of F_{TLS} with different CPW width

and gap parameters, thus helping to elucidate where the TLSs might be. One should note, however, that only the product of F_{TLS} and δ is found from a fit; changes in F_{TLS} are inferred by assuming a constant δ . As argued by Gao, if the TLSs are in the bulk dielectric, then F_{TLS} should be independent of width and gap, while if the TLSs were merely on the surface of the dielectric, then F_{TLS} should decrease with increasing width and gap. The latter was observed, leading to the conclusion that TLSs were primarily a surface distribution. The reason the F_{TLS} decreases as width and gap increases is that less energy is stored within the narrow surface layer; the electric field lines are spread out further into the vacuum and into the substrate, as shown in Fig. 4.2. Improving the quality factor of CPW resonators does not leave much flexibility beyond changing the width and gap parameters. The experiment described below is an extension of this study with compact resonators due to their highly flexible nature; descriptions of both CPW and compact resonators are given in Section 2.1.1.

4.1.2 Simulation

The idea behind compact resonators is to build a resonator from two connected subresonators, one with low impedance playing the role of a capacitance, and one with high impedance playing the role of an inductance [91, 92, 82], see Fig. 4.7. Because their size is smaller than the wavelength at the resonant frequency, these are often called lumped-element resonators. Alternatively, the compact resonator nomenclature does not presume that the attempted “lumped” nature is actually achieved, and simply names them from the fact that they are generally smaller than the more common CPW resonators. With thin-film lithography, a good way to make approximately lumped resonator elements is to use meanders for inductance and interdigitated fingers for capacitance.




Figure 4.2: Cartoon drawing of a cross-section of a CPW transmission line, highlighting different participation of surface (shown in red) with varying CPW gap. The gap is the distance between the centerpin and the ground planes of the CPW mode. In (a), the gap is small, and thus many electric fields pass through the surfaces, while in (b) the gap is large, and a smaller percentage of electric fields pass through the surface. One may note that in (b), the mode is less localized, and thus more likely to couple to other modes in the sample holder.

The inductor, composed of a thin meandering trace, has an inductance that scales with the length of the trace, without the gain in inductance from the coiling of a solenoid. The meander inductor is characterized by the parameters: l_L , the total length, g_L , the gap between adjacent meanders, and w_L , the width of the trace. To approximate the object as a lumped inductor, simulations performed in the 2.5D microwave simulation software package Sonnet were compared to that expected for a simple inductor. A 50Ω port was placed at each end of the inductor and the transmission was simulated.

Simulation results of inductance values of meander inductors inferred from the model shown in Fig. 4.3 are shown in Fig. 4.4. The addition of parasitic capacitances to ground is needed to properly infer an inductance from simulated transmission results. The parasitic capacitance is not surprising considering the geometry and the closeness of the conducting ground plane below the inductor. The simulations




Figure 4.3: Model for the extraction of “lumped” element inductor values from a meander inductor simulation. The parasitic capacitances to ground are needed in order to fit the simulation response, and their value shows the limit of the “lumped” element approximation.




Figure 4.4: Simulation results from Sonnet simulation software showing inferred lumped element inductance of meander inductors using the model of Fig. 4.3. (a) Inductance increases linearly with total meander length; different trace widths (w_L) produce different slopes. (b) Comparison of inductance per length as a function of w_L , showing that narrower traces have higher inductance, as expected.

show that inductance is indeed linear with total meander length. This confirms the idea that the inductance basically comes from the long length of the wire. Secondly, simulations show that the inductance of a given trace length varies inversely with the width of the trace.

Interdigitated capacitors are, so far, the best way to make high quality factor on-chip capacitors out of thin films. Alternatively, one may use overlap “parallel plate” capacitors, however these materials were at the time lower quality factor [93] and are just now yielding comparable quality factor [94]. While there are analytical




Figure 4.5: Model for the extraction of “lumped” element capacitor values from an interdigitated capacitor simulation. The parasitic inductances on each side of the capacitor are needed in order to fit the simulation response, and their value shows the limit of the “lumped” element approximation.

calculations of the capacitance of interdigitated capacitors [95], these formulas are generally only accurate in certain parameter regimes and may not take into account the effect of the distributed inductance that determines self-resonance frequencies. An interdigitated capacitor is defined by the parameters: l_C , the total length of adjacent fingers, g_C , the gap between adjacent fingers, and w_C , the width of the traces. To approximate the actual structure as a lumped capacitor, simulations are performed in Sonnet in the same manner as the inductor, with attaching 50Ω ports directly to either side of just the interdigitated capacitor.

Simulation results of capacitance values of meander inductors inferred from the model shown in Fig. 4.5 are shown in Fig. 4.6. The addition of the parasitic inductances on each arm of the capacitor is needed to infer a capacitance from simulated transmission response. The parasitic inductance is not surprising considering the potentially long length of the capacitor fingers. The simulations show that capacitance is indeed linear with total adjacent finger length, as expected. Next, simulations show that the capacitance for a given finger length varies logarithmically with the width of the trace, meaning that wider traces lead to higher capacitances. Also, capacitance is found to vary inversely with the gap between two capacitor fingers. None of these dependencies are surprising, but these simulations allow for the design of capacitors of a given capacitance with any desired parameters.

To make a compact resonator, one chooses the desired inductor and capacitor

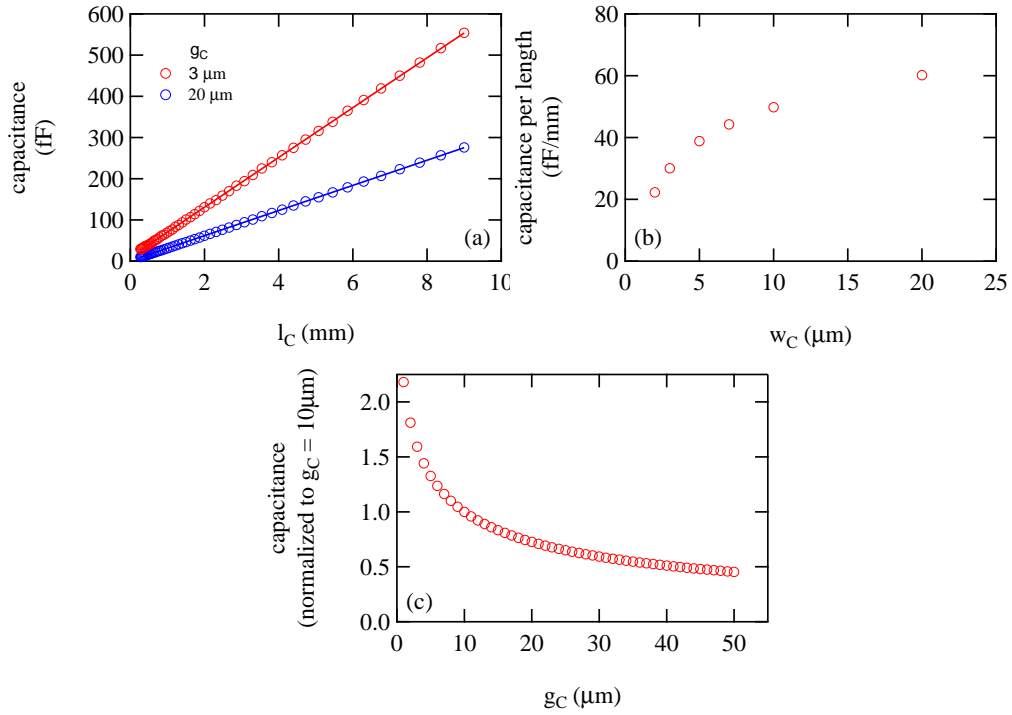


Figure 4.6: Simulation results from Sonnet showing inferred lumped element capacitance of interdigitated capacitors using the model of Fig. 4.5. (a) Capacitance increases linearly with total adjacent finger length; different gaps (g_C) produce different slopes. (b) Comparison of capacitance per length as a function of w_C , showing that wider traces have higher capacitance, as expected. (c) Comparison of capacitance per length as a function of g_C , showing that smaller gaps lead to higher capacitance, as expected.

from the frequency (ω_0) and impedance (Z_0) via $C = \frac{1}{\omega_0 Z_0}$ and $L = \frac{Z_0}{\omega_0}$. The frequency predicted from the ideal individual capacitance and inductance values is typically 30% higher than actual measured resonators. This difference is due to the parasitic capacitances and inductances shown in Figures 4.3 and 4.5, along with the additional capacitance from the ground plane, which was not included in the inductor and capacitor simulations. Simulations of full resonator structures are more accurate, with a typical 4% difference from measured values. For the purpose of measuring resonator quality factors, obtaining the exact frequency is not critical, so this amount of precision is adequate.

4.1.3 Design

Combining the meander inductor and interdigitated capacitor produces a resonator as shown in Fig. 4.7. These resonators were designed to be measured in the hanger-coupled scheme. For an in-depth analysis of this method of coupling, which allows for the extraction of both Q_c and Q_i , see Appendix A. The resonator is placed in a hole in the groundplane of the CPW feedline. The coupling between the resonator and the CPW feedline is achieved by placing a portion of the resonator inductance parallel to the feedline. The parameters that control the coupling strength are the width of the ground plane remaining between the resonator and the CPW centerpin (the shield width), and the length of the resonator inductor that is participating in the coupling. Narrower shields and longer coupling inductor lengths will produce stronger coupling, and therefore lower values of Q_c .

Three progressive resonator designs are used extensively in this study, called Designs A, B and C. Each of these designs is shown in Fig. 4.8. The different parameters and quality factors of these designs is discussed in Section 4.3.




Figure 4.7: Schematic showing design of “Design A” compact resonator and its coupling to the CPW feedline. The inset shows the design parameters g_L , the gap between meander inductors, g_C the gap between adjacent capacitor fingers, g_R the gap between the resonator and the ground plane, and w the width of both inductor and capacitor traces.




Figure 4.8: Schematic showing three resonator designs, without the CPW feedline. Resonator design is referred to as (a) Design A, (b) Design B, and (c) Design C. In this schematic, black represents metallic areas and white where the substrate is visible.

4.2 Experimental Apparatus

The measurements described in the next section were completed in both the Kelvinox25 and the Kelvinox400 dilution refrigerators, shown in Fig. 3.11. The Kelvinox25 refrigerator has an operational base temperature of about 80 mK with a cooling power of nominally 20 μ W at 100 mK. A schematic of the microwave setup inside the dilution refrigerator is shown in Fig. 4.9.

In this measurement setup, 4 chips were cooled at once, with each chip containing one feedline coupled to 6 independent resonators at frequencies between 5 and 8 GHz. Thus 24 resonators were tested in each cycle of the fridge. The chips were wire-bonded to a printed circuit board with Arlon dielectric and placed inside the octobox sample box (see Section 3.1.2). The octobox was mounted inside a magnetic shield (Amuneal A4K) and attenuators were installed totaling 50 dB on the input microwave line. All 4 chips were excited simultaneously using a passive 4-way microwave splitter. The output line consisted of two Pamtech 4-8 GHz isolators on the mixing chamber, a 12 GHz low-pass filter on the 700 mK stage, and a Caltech HEMT amplifier at the 4 K stage. The measurement line was switched between the 4 chips using a microwave switch (Radiall R573423600, see Section 3.3.2) mounted on the mixing chamber.

The Kelvinox400 refrigerator has an operational base temperature of about 15 mK with cooling power of approximately 400 μ W at 100 mK. A schematic of the microwave setup inside the dilution refrigerator is shown in Fig. 4.10. The only intentional change between the setups is the exchange of a 4-way splitter with a second switch.




Figure 4.9: Schematic of microwave lines in Kelvinox25 for compact resonator measurements and a picture of actual setup. The components on the schematic are labeled, not every component is visible in the included picture. Blue outlines indicate cryoperm shields.

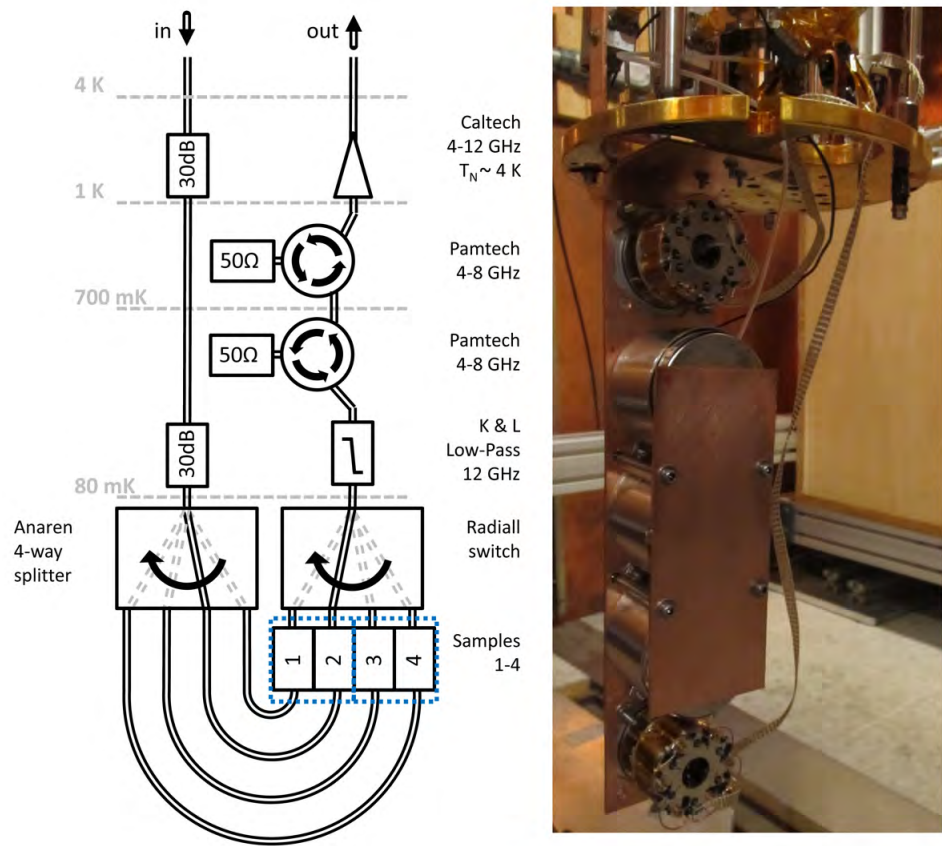


Figure 4.10: Schematic of microwave lines in Kelvinox400 for compact resonator measurements and a picture of actual setup. The components on the schematic are labeled, not every component is visible in the included picture. Blue outlines indicate cryoperm shields.




Figure 4.11: Two extreme examples of resonator response curves fit with Eq. 4.5. Responses typically fall between (a) symmetric and (b) strongly asymmetric about the resonant frequency.

4.3 Results

The results for this experiment are given in Ref. [47], and will be repeated here. While each individual resonator yields different results, there are certain traits which are common among them. Each measurement of resonant frequency and quality factor is extracted, as explained in Appendix A, from a transmission measurement through the feedline to which the resonator is coupled. The transmission yields information about the Q_c , Q_i and f_0 from a single measurement trace. Two example measurements are shown in Fig. 4.11, one of which is nearly symmetric, and the other quite asymmetric, illustrating the need to use the fully asymmetric fitting function derived in Appendix A, and shown in Eq. 4.5, where $\delta\omega$ is a parameter characterizing the extent of the asymmetry and $x = \frac{\omega - \omega_0}{\omega_0}$.




Figure 4.12: Q_i dependence on \bar{n} , the applied power expressed in terms of an average number of circulating photons in the resonator. The rise in Q_i is due to the saturation effect of TLS dissipation, explained in Eq. 4.4. Each color is a different resonator, all on the same chip and measured sequentially.

$$S_{21} = 1 - \frac{\frac{Q_0}{Q_c} - iQ_0 \frac{2\delta\omega}{\omega_0}}{1 + 2iQ_0 x} \quad (4.5)$$

A typical measurement of Q_i versus power is shown in Fig. 4.12, with power expressed in terms of an average number of circulating photons on the resonator, \bar{n} . The quality factor increases from a constant value at powers of order $\bar{n} = 1$ due to contributions from TLS dissipation in accordance with Eq. 4.4. Each color is a different resonator from the same chip, and measured sequentially during the same refrigerator cooldown. The maximum quality factor of approximately 1×10^6 at around 10^8 photons is typical; the maximum value of this peak was observed to be approximately 5×10^6 . The decrease in quality factor at highest powers is likely due to the finite critical currents and the breaking of Cooper pairs.

Measurements of the resonant frequency as a function of temperature also contain information about the contribution of TLS dissipation to resonator loss. An example measurement is shown in Fig. 4.13. The fit, using Eqs. 4.2 and 4.3 and the Mattis-Bardeen theory of kinetic inductance [87], is to the data for the 6.6 GHz resonator and yields values of $\alpha = 0.259$ and $F_{TLS} = 0.037$, where α is the kinetic inductance




Figure 4.13: Resonator resonant frequency dependence on temperature. This dependence is highlighted by Eq. 4.2, with two contributions; a decrease at low temperatures due to TLS dissipation and a decrease at high temperatures due to kinetic inductance. Each color is a different resonator on the same chip.

fraction and F_{TLS} is the TLS filling factor. This measurement confirms, along with the measurement of Q_i versus power, that TLS dissipation is contributing to the Q_i of these resonators. Note that the temperature used on the x -axis here is the mixing chamber temperature, which is known (see Fig. 4.17 and preceding discussion) to deviate from sample temperature in the low-temperature limit.

Measuring the resonant frequency dependence on temperature as a function of the parameter g_C , the gap between adjacent capacitor fingers, yields an interesting result. Fig. 4.14 shows the low-temperature portion of this dependence. In this range, the kinetic inductance contribution can be neglected and any remaining change in frequency is presumably due to TLS dissipation. Note that resonators with smaller g_C have a large positive slope while those with larger g_C have a negative slope. Slope changes monotonically between the smallest and highest value of g_C . Since positive slopes can be attributed to TLS dissipation, it is clear that as g_C is increased, the TLS filling factor is reduced and the effect of TLS dissipation on resonant frequency is dampened. The large negative slope at these temperatures cannot be explained by Mattis-Bardeen theory, but may be due to the fact that the x -axis is the mixing




Figure 4.14: Dependence of resonant frequency on temperature for resonators with varying g_C . Each color is a different resonator, all on the same chip, and measured sequentially during the same refrigerator cooldown. Lines through points are purely to guide the eye.

chamber temperature and not sample temperature.

With the goal of optimizing the design of compact resonators in mind, the method chosen was to vary independently each of the main parameters thought to control TLS loss. These parameters are: the gap g_C between two adjacent capacitor fingers, the distance g_L between two adjacent inductor meanders, the distance g_R between the resonator and the surrounding ground plane, and the width w of the resonator traces. The width of the inductor and capacitor traces are swept together and not varied independently in this study. In addition, the characteristic impedance Z_0 of the resonator was also varied.

As a means of reference, a benchmark set of resonators were measured, with parameters: $g_C = 10 \mu\text{m}$, $g_L = 20 \mu\text{m}$, $g_R = 10 \mu\text{m}$, $w = 5 \mu\text{m}$ and $Z_0 = 100 \Omega$. Resonators with this set of parameters are called “design A” resonators. A total of 25 design A resonators were measured with an average Q_i of 160,000 ($\pm 20,000$) and a maximum of 210,000 at single-photon power. Additionally, one chip with 6 resonators inexplicably had quality factors ranging from 40,000 to 70,000, much lower than the rest; this chip was not included in the benchmark.

g_C (μm)	3	5	10	20	30	40	
ΔQ_i	-31%	-14%	0%	16%	0%	18%	
g_L (μm)	2	5	10	20	40	60	
ΔQ_i	-6%	-10%	1%	0%	2%	2%	
g_R (μm)		5	10	20	50	100	200
ΔQ_i		28%	0%	-2%	-28%	-24%	-36%
w (μm)		3	5	7	10	20	
ΔQ_i		-16%	0%	17%	16%	60%	
Z_0 (Ω)		50	100	150	200	250	300
ΔQ_i		-12%	0%	26%	33%	19%	33%

Design A

Design B

Figure 4.15: Dependence of Q_i on parameter values. The changes in Q_i for a given mutant value are reported in reference to the average Q_i (160,000) of design A, with positive values representing improvements. The shaded column indicates the design A value of that parameter; ΔQ_i in this column is zero by definition. The square in each row with a bold border shows the value chosen for design B. Parameters for design C cannot easily be shown on this figure as explained in the main text.

To measure the effect of each parameter on the value of Q_i , 24 geometrical variants of design A were measured, with each “mutant” resonator having only one parameter value that is changed. For example, the mutant values of g_C were: 3, 5, 20, 30, and 40 μm . The results of the mutant resonators are shown in Fig. 4.15; percent changes in Q_i are given with respect to the design A resonator benchmark.

For g_C , small values lead to lower Q_i , and larger values lead to higher Q_i . The effect of changing g_L on Q_i is at least a factor of three smaller than for g_C . Thus, the gaps where electric fields are present (the capacitor and not the inductor), partially control Q_i , consistent with a surface loss mechanism coupled to the electric field. Similarly, Q_i increases for larger w , again consistent with surface loss, since wider traces lead to decreasing electric field strength at surfaces. Next, it is found that Q_i drops by roughly 25 % if $g_R \geq 50 \mu\text{m}$, suggesting that the ground plane prevents electric fields from reaching lossy materials such as the copper box or printed circuit

board (PCB) dielectric. Lastly, the trend indicating that larger values of Z_0 are beneficial to Q_i appears to contradict the usual hypothesis that dissipative mechanisms have a constant $\tan\delta$. The results for g_C , g_L , and w are all consistent with a loss dominated by surface electric field participation.

To combine the effect of each improved parameter, two new sets of parameter values were chosen using these results. Resonators with these parameters are called designs B and C resonators. Design B values were chosen to be relatively modest changes from design A while design C values were chosen to maximize Q_i . Design B chosen values were: $g_C = 20 \mu\text{m}$, $g_L = 5 \mu\text{m}$, $g_R = 10 \mu\text{m}$, $w = 10 \mu\text{m}$, $Z_0 = 200 \Omega$. Resonator size increases rapidly with g_L since the larger Z_0 requires twice the inductance. Therefore, to limit the overall size to roughly $700 \mu\text{m} \times 500 \mu\text{m}$, we reduced g_L to $5 \mu\text{m}$, despite the fact that this may lower Q_i by 10%. Design C chosen values were: $g_C = 80 \mu\text{m}$, $g_L = 10 \mu\text{m}$, $g_R = 10 \mu\text{m}$, $w_L = 10 \mu\text{m}$, $w_C = 40 \mu\text{m}$, $Z_0 = 300 \Omega$. Note that g_C was chosen beyond the range of tested mutant design A resonators. Also in design C, the trace width w was different for the capacitor (w_C) and inductor (w_L) halves in order to benefit from the larger capacitor width while keeping the resonator from being larger than $1000 \mu\text{m} \times 1000 \mu\text{m}$.

The results of all 49 design A, 73 design B, and 28 design C resonators are shown in Fig. 4.16. Designs B and C each show significantly higher Q_i than design A, with design C better on average than design B. While there exists a spread in Q_i for each design, we observed an overall increase in the range of measured Q_i . The maximum/median Q_i rose from 210,000/160,000 for design A to 370,000/280,000 for design B, and 500,000/380,000 for design C.

Resonators of type design A and B were also measured without the standard ion-milling cleaning step during fabrication. Ion-milling is meant to clean the bare sapphire surface after lithography but before deposition to make a clean substrate-

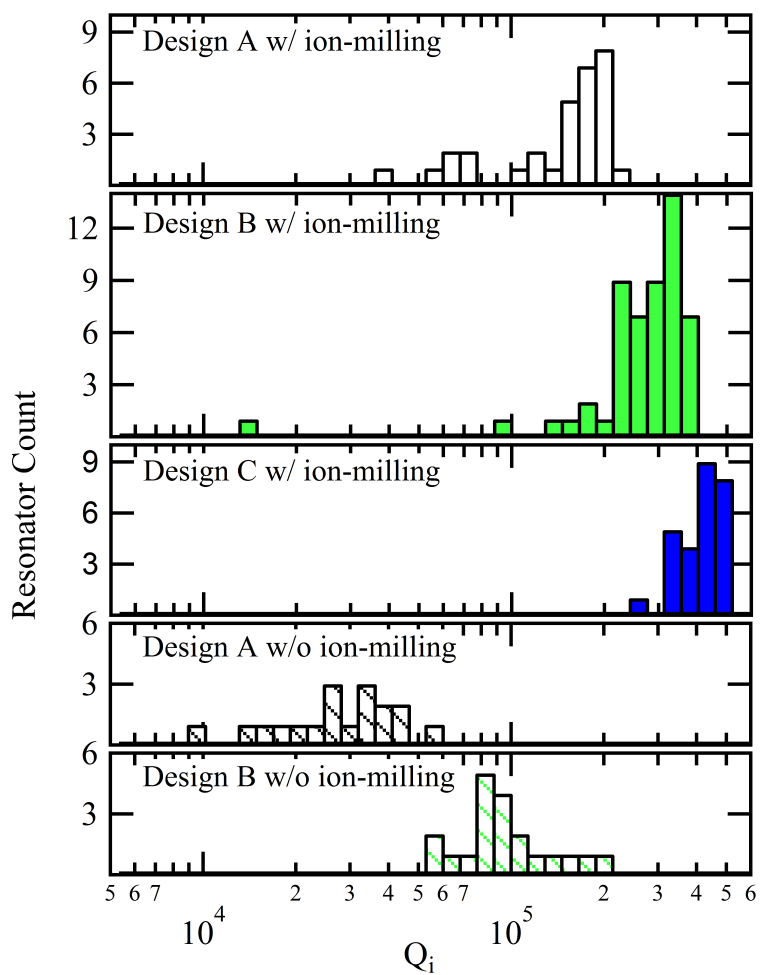


Figure 4.16: Histograms of single-photon internal quality factors for designs A, B, and C resonators with ion-milling and designs A and B without ion-milling. Q_i improves steadily from design A to B to C. For both designs A and B resonators, when ion-milling is not performed, Q_i is roughly a factor of 4 lower, maintaining the improvement between designs A and B.

metal interface. When ion-milling was not performed, the maximum/median Q_i was reduced to 50,000/30,000 for design A and 190,000/80,000 for design B. Design C was not measured without ion-milling. For both designs A and B, the median quality factor was reduced by roughly a factor of four when ion-milling was skipped.

4.4 Analysis

Repeat measurements were done in the Kelvinox400 dilution refrigerator, which showed lower quality factors than in the Kelvinox25, by roughly a factor of 2. Measuring the quality factor at different temperatures showed that the measurements agreed when the Kelvinox400 fridge was heated to approximately 200 mK, as shown in Fig. 4.17. This shows that while the temperature of the mixing chamber thermometer in the Kelvinox25 fridge said 80 mK, the actual temperature of the resonator sample was closer to 200 mK. TLS loss decreases at higher temperatures due to saturation in the same manner that higher power saturates TLS, so a higher fridge temperature leads to a higher resonator quality factor.

The strong effect of ion-milling of the sapphire surface (before metal deposition) on the observed values of Q_i help tell where losses are coming from. Since the ion-milling affects only the substrate-air interface and substrate-metal interface, we infer that these two surfaces participate strongly. The dominating participation of these surfaces has also been predicted by simulation [96]. This Q_i dependence on ion-milling also suggests that while the geometry controls the resonator sensitivity to the surface loss mechanism, the surface preparation determines the strength of the loss.

In summary, these results show that the Q_i of compact resonators depends strongly on geometrical factors controlling where electric fields are stored. This




Figure 4.17: Resonator quality factor measurements of a single resonator measured in both the Kelvinox25 and Kelvinox400 dilution refrigerators. The measurement in the Kelvinox25 was at the base temperature of 80 mK as quoted from the mixing chamber temperature sensor. Multiple measurements at different mixing chamber temperatures in the Kelvinox400 are shown.

surface loss mechanism is consistent with a distribution of TLS as shown by the dependence of Q_i on power, temperature, and by the resonant frequency dependence on temperature. In addition, substrate surface preparation prior to metal deposition is crucial for obtaining high quality surfaces. Using these results indicating that surface loss is dominant, quality factors have been increased, at an approximate temperature of 200 mK, from 210,000 to 500,000 by simply tweaking the resonator's geometrical design parameters.

4.5 Discussion

These results showed that for a given set of materials and processing steps, the quality factor could be improved significantly simply by design improvements. Any set of materials and processes that are still limited by surface losses should benefit from the same design improvements. This could be extended until the materials and pro-

cessing improve to such a level as to not be limited by surface losses. More recent results have achieved low-temperature (probably below 200 mK) single-photon quality factors of over 1×10^6 with aluminum on sapphire CPW resonators [97]. The improvements described were attributed to better fabrication, with best results utilizing MBE-grown aluminum heated to 850 °C in the presence of activated oxygen. These results still show an increase in quality factor with increased power, indicating that TLS dissipation still contributes to resonator loss. Thus, combining this improved fabrication recipe with improved compact resonator designs could yield quality factors of a few million and would be an interesting experiment to complete.

In hindsight, it is clear that the parameter g_R should have earned more attention. One of the reasons for the discrepancy between the measured frequency predicted by the individual inductor and capacitor values is that the presence of the ground plane adds a large amount of capacitance to the resonant mode. This capacitance lowers the resonant frequency, and thus participates significantly in the mode. Especially in design C resonators, where $g_C = 80 \mu\text{m}$, and $g_R = 10 \mu\text{m}$, the influence of g_R on both the frequency and the loss was likely dominant. In further studies, it would be prudent to treat g_R and g_C on an equal basis.

Another important point not to forget is the sample temperature. The sample temperature is often not equal to the mixing chamber temperature since there may not be enough thermal conductance between the sample and the mixing chamber. There may be a weak point at the substrate to sample-box connection, or possibly between the sample box and the mixing chamber, especially if there is cryoperm between the two. Every thermal link has to be strong in order to produce a low sample temperature. Even with good thermal conductivity, the microwave lines must also be properly filtered in order to prevent heating from thermal noise (see Section 3.3.3).

Concerning the possibility of a raised sample temperature, it is possible to measure the effective temperature of a qubit (see Appendix B), but it is very difficult to objectively measure the temperature of a harmonic resonator. One must always be cautious of “low-temperature” results for resonator quality factors, because an imperfect setup may produce higher quality factors.

In conclusion, measuring the quality factors of over 100 resonators of many different geometrical design variations has shown that participation ratios form the basis for the analysis of factors limiting relaxation. Geometrical adjustments decreased the participation ratio of surface loss (likely a TLS distribution), without changing the material components or processing techniques, resulting in an increase in resonator coherence. Quality factors have been increased, at an approximate temperature of 200 mK, from 210,000 to 500,000 through an increase in gaps and widths of the interdigitated capacitor, directly constituting a step towards the goal of significantly improving superconducting qubit and resonator coherence. In the next chapter, these improvements are applied to the transmon qubit.

Compact Transmon

5.1 Sample Design

The main motivation for improving the quality of compact resonators was that they constituted a proxy for superconducting qubits, the main enabler of quantum information experiments. Qubits with T_1 and T_2 on the order of milliseconds, without sacrificing readout fidelity or gate time, enable practical quantum computation. At the time of the start of this experiment, one of the most reproducible qubits with relatively long lifetimes was the transmon coupled to a CPW resonator. This 1st generation transmon qubit implementation will be called the “CPW transmon”, and can be seen in Fig. 5.1. The transmon is composed of the SQUID loop within the blue square and the effective capacitance in parallel, largely coming from the interdigitated capacitor, see Section 2.2.1 for an in-depth description of transmon theory.

Comparing the interdigitated capacitor of the CPW transmon with the interdigitated capacitor for compact resonators in Chapter 4, it is clear that the gap between adjacent capacitor fingers is much smaller. The gap between adjacent capacitor fingers in the capacitance directly across the junction (looks like “moose antlers”) is




Figure 5.1: SEM image of a CPW transmon device. Light colors are aluminum and the darker is the silicon substrate. The horizontal line in the middle is the centerpin of the CPW resonator, the transmon is placed in the gap between the centerpin and the ground plane. The large interdigitated capacitor of the transmon is visible, while the SQUID loop is within the blue square.

usually $3 \mu\text{m}$, while the gap between the CPW centerpin and one side of the qubit capacitor can be as small as $1 \mu\text{m}$ depending on alignment. Both of these capacitors contribute to C_Σ and therefore participate strongly in the qubit mode. Results from compact resonators have shown that while increasing g_C to $20 \mu\text{m}$ or further, quality factors are *still* determined by this gap. Gaps on the order of $1\text{-}3 \mu\text{m}$ should therefore be significantly worse, and are a likely culprit for limiting coherence times of CPW transmons.

The experiment described in this chapter entails designing and measuring a new implementation of the transmon qubit based around the results from the compact resonator study. The new implementation, called the “compact transmon”, or “ 2^{nd} generation” transmon, will include an interdigitated capacitor with large g_C and w in order to minimize coupling to surface loss. The compact transmon will be coupled to a compact resonator to benefit from the higher quality factor resonator as well. Results for typical CPW transmons are shown in Fig. 5.2 [39]. While this figure is directly from an experiment on the Purcell effect, it shows that when not

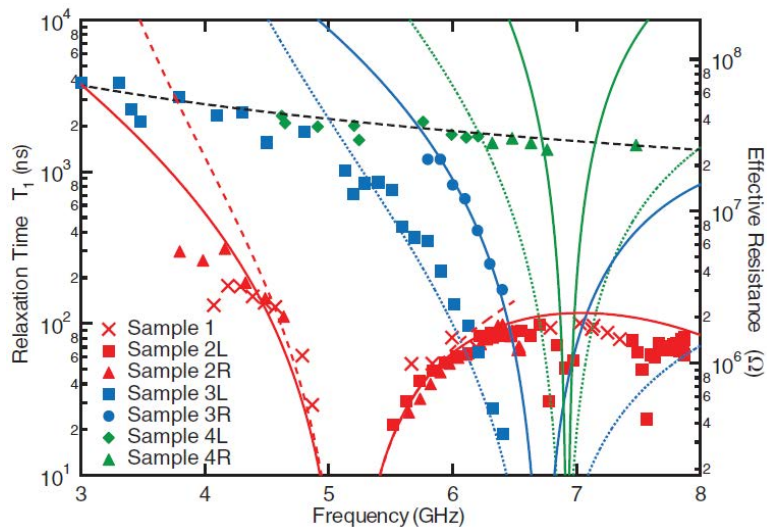


Figure 5.2: T_1 times for several CPW transmon devices, reprinted from Houck et al., 2008 [39]. For details on the solid lines and differences between qubits, see the reference. The nearly horizontal dashed line near the top of the graph corresponds to the lifetime predicted for a constant quality factor of 70,000. Note that the qubit lifetimes from several qubits seem to be bounded by this line.

Purcell limited, lifetime are limited by a quality factor about 70,000. At the lowest frequencies of 3 GHz, this corresponds to a lifetime of almost $4 \mu\text{s}$. Is this limited by the dielectric loss (TLS on surfaces) of the capacitor, and can compact transmons improve this lifetime?

5.1.1 Simulation

As described in Section 1.2.1 and 2.2.1 and shown in Fig. 1.6, the transmon qubit is composed of a Josephson tunnel junction (or SQUID for tunability) and a capacitance in parallel. In contrast to the CPB, the transmon's capacitance C_Σ largely comes from a capacitance separate from the junction. In both CPW and compact transmons, this capacitance is from an interdigitated capacitor patterned lithographically from the same material during the same fabrication step as junction fabrication (see Section 3.1.1 for fabrication details).




Figure 5.3: Lumped element model of transmon qubit (dashed green box) coupled to resonator (dashed blue box). Particular shunting capacitances C_S and coupling capacitances C_C are needed in order to obtain a desired qubit frequency and coupling strength.

When designing the compact transmon, the lumped element model in Fig. 5.3 is useful; it symmetrizes the coupling capacitor C_C in a way that more closely mimics the actual device. Simulations are used to predict values of C_S and C_C in order to obtain the desired transmon frequency and transmon-qubit coupling strength.

The software used was the fully 3D electromagnetic finite-element solver: Maxwell, from the company Ansoft. Maxwell is the DC brother of the high-frequency program HFSS, used for other simulations in this thesis (for example, of the Purcell effect in Section 7.4). The thin-film metals were simulated as perfect conducting infinitely thin conductors as an approximation since their thickness (200 nm) is much less than the closest distance between conductors (10 μm). Since Maxwell is a DC simulator, the resonator inductor was removed to prevent it from shorting the resonator. It is

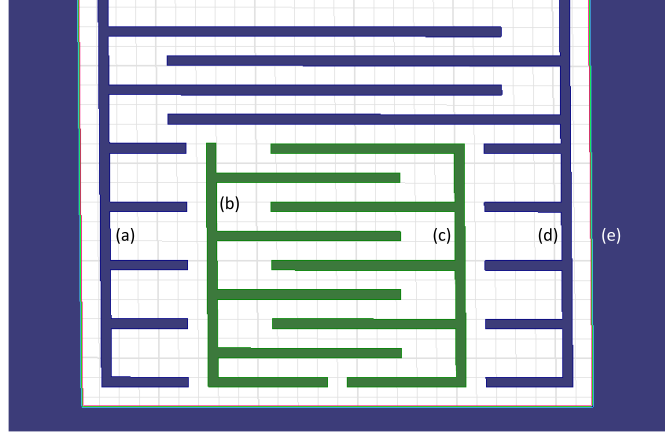


Figure 5.4: Screen shot from Maxwell simulation software to simulate capacitances in compact transmon mode. All colored shapes are metal, with (a) and (d) the two halves of the resonator capacitance, (b) and (c) the two halves of the transmon capacitance, and (e) the ground plane. The inductor of the capacitor is removed (not shown). The transmon junction is placed between parts (b) and (c) at the point of nearest approach at the bottom.

important that the ground planes below the chip and on the chip are included in the simulation, as they participate in the capacitance of the resonator and transmon. An example of a design simulation in Maxwell is shown in Fig. 5.4.

Maxwell produces a capacitance matrix, with each element given by $C_{ij} = \frac{\partial Q_i}{\partial V_j}$, from the selected conductors, in this case 5 metals comprising the two halves of the resonator capacitance, the two halves of the transmon capacitance, and the ground plane. The ground plane “belongs” to the CPW feedline to which the resonator is coupled, and in simulation is shorted to the ground plane of the sample box below the chip. In the capacitance matrix, Maxwell always includes another, artificial, ground at the outside of the simulation region. In order not to be influenced by this ground, one must make this ground far enough away that it can be approximated as infinity.

The capacitances in the model of Fig. 5.3 are merely proxies for the charging energy E_C and qubit-cavity coupling strength g . Actually, one obtains the total effective qubit (C_Q) and resonator (C_R) capacitances from Eqs. 5.1. In these equa-

tions, each component $X_{ij} = \frac{1}{C_{ij}^{-1}}$ is the $(i, j)^{th}$ element of the inverse of the inverse capacitance matrix, with the indices named to match Fig. 5.4. Then calculate the frequency directly from Eqs. 5.2. Equations 5.2 then shows how to calculate E_C from C_Q , then to calculate ω_Q from E_C and E_J , then how to calculate the coupling strength g from these parameters. Therefore, all of the relevant qubit and resonator parameters are known, since E_J , ω_R and κ are found using other means.

$$\begin{aligned}
C_q &= (X_{bb} + X_{cc} - 2X_{bc})^{-1} \\
C_r &= (X_{aa} + X_{dd} - 2X_{ad})^{-1} \\
C_x &= (X_{ba} + X_{cd} - X_{bd} - X_{ac})^{-1} \\
C_c &= \frac{C_Q C_R}{C_X}
\end{aligned} \tag{5.1}$$

$$\begin{aligned}
\frac{E_C}{h} &= \frac{e^2}{2C_Q} \\
\omega_Q &= \sqrt{8E_J E_C} \\
\frac{g}{2\pi} &= \frac{1}{2} \frac{\sqrt{\omega_Q \omega_R C_Q C_R}}{C_X}
\end{aligned} \tag{5.2}$$

An actual device whose results are shown in the next section can be seen in Fig. 5.5.

5.2 Experimental Apparatus

The measurement setup is largely identical to that for the compact resonators in Section 4.2. The first few compact transmons were measured in the Kelvinox25 dilution refrigerator, the rest were measured in the Kelvinox400. The experimental setup for the Kelvinox25 refrigerator is identical to that shown in Fig. 4.9, while the slightly modified setup in the Kelvinox400 refrigerator is shown in Fig. 4.10. Changes in this setup included the addition of a copper powder filter for infrared absorption on both the input and output microwave lines for *some* of the measurements, see




Figure 5.5: Optical image of compact transmon device. The compact resonator's inductance and capacitance are similar to those of compact resonators without qubits. Additional capacitor fingers were added to enable coupling to the qubit. At the bottom-center of this image, the interdigitated capacitor of lighter color is the shunting capacitance of the transmon, with small fingers coupling to the resonator. The junction can be seen between the shunting capacitance.

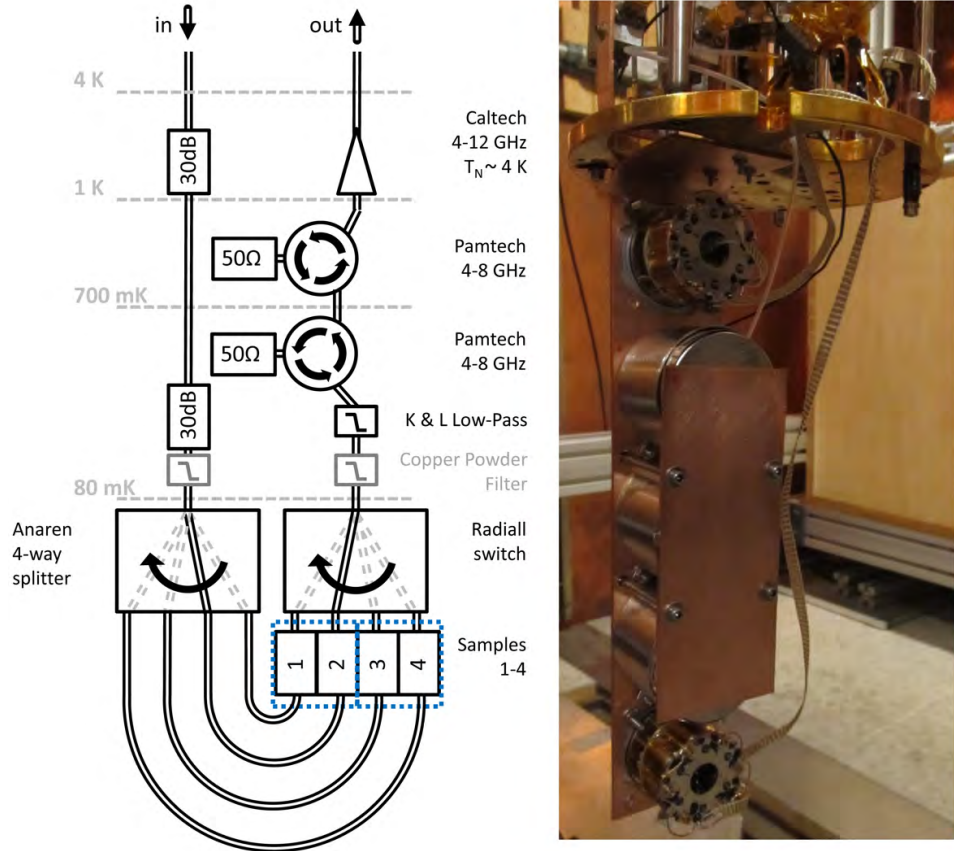


Figure 5.6: Schematic of microwave lines in Kelvinox400 for compact resonator measurements and a picture of the actual setup. The components on the schematic are labeled, but not every component is visible in the included picture. Blue outlines indicate cryoperm shields.

Section 3.3.3 for details on the filter.

Throughout the course of the measurements described in this experiment, it was discovered that many small modifications, together, reduce the effective temperature of the qubit sample. In general, poor thermalization combined with any heat sources can raise the “equilibrium” temperature of the qubit far above that of the mixing chamber. Fig. 5.13 shows T_1 before and after a list of modifications that will be described here. First, the connector for the magnetic field coil was discovered to have a few $m\Omega$ of resistance. This connector was first reduced in resistance, and then eliminated entirely, resulting in much less active heating. Second, the entire




Figure 5.7: Qubit temperature (measured with RPM) as a function of magnetic field coil current in several instances. Resistances are for a connector in the path of the applied current, 6 mΩ represents two connectors, 3 mΩ represents one connector, and 0 mΩ represents a direct connection with superconducting wire. A measurement with silver-backing and 3 mΩ resistance shows that silver-backing reduces the temperature for a given amount of heating. Note that when heating is removed, qubit temperature roughly matches sample box temperature.

backplane of the sapphire chip containing the device was coated in silver and was mounted to the octobox using silver paste instead of GE varnish, in order to obtain better thermal contact with the octobox metal. Measurements showing the heating effect of the magnetic field coil connector and the effect of silver-backing are shown in Fig. 5.7. In addition, thermal anchoring braids were added between the octobox and the cold finger of the mixing chamber, copper powder filters (3.3.3) were added on the input and output microwave lines and hand-formable braided cables were removed in favor of semi-rigid cables. These modifications were taken a step further by Barends et al. [77] and in this thesis work in Chapters 6 and 7.




Figure 5.8: (a) Cavity transmission amplitude as a function of frequency, showing how the cavity shifts due to qubit excitation. (b) Measurement of cavity transmission amplitude while sweeping additional qubit tone. Cavity transmission amplitude rises when qubit is excited, the two peaks shown are due to the f_{ge} and f_{ef} transitions. (c) Measurement of qubit frequencies as a function of applied flux, showing tunability of qubit frequencies.

5.3 Results

In the process of this experiment, 11 working compact transmons were measured. Each qubit had slightly different frequencies and coupling strengths, along with some important experimental setup changes along the way. While each transmon had its own coherence times, there are many similarities among the measurements.

First of all, each transmon spectroscopy looks quite similar. The spectroscopy is defined as the response of the cavity due to a frequency sweep of a probe (spectroscopic) tone. When this probe tone excites the qubit, the cavity shifts and the measured transmission on resonance is changed. For a more detailed description of dispersive readout, see Section 3.2.1. The basic points of spectroscopy can be seen in Fig. 5.8.

The first image shows a single trace of spectroscopy. The two peaks are the first two transition frequencies of the transmon, f_{ge} and f_{ef} . The difference between these two frequencies is a measure of qubit anharmonicity, which is often equal to E_C . For a sufficiently anharmonic qubit, E_C is kept around 300 MHz, and here it matched simulations well. Note that since there is a response at the f_{ef} with a single

spectroscopy tone proves that there is some $|e\rangle$ population to begin with. Since this is unexpected if the qubit were thermalized to the mixing chamber temperature, this effect is commonly called a “hot” qubit. While a temperature could be inferred from the height of these two peaks, this method is highly inaccurate; a better method to measure qubit temperature is described in Appendix B.

The second image is an array plot containing many slices of spectroscopy versus applied flux. Each of the 11 tested qubits had an inductive element made from a SQUID. This allows for tuning of the qubit transition frequency by applying a flux through this loop to influence the effective inductance. Applying flux through the loop increases the inductance and thus lowers the qubit frequency; the maximum frequency is at zero flux which may not correspond to zero applied flux. Both the f_{ge} and f_{ef} transition peaks can be seen to tune in this image.

The third image shows how the dispersive readout works. The trace is a measurement of the cavity transmission amplitude with and without a spectroscopic tone at the qubit frequency. The cavity shifts in frequency and allows for the two states to be distinguished by a change in transmission amplitude. The transmission shape is an inverted Lorentzian as expected for a resonator that is hanger-coupled to a feedline. Note that since this measurement uses qubit saturation instead of preparing the $|e\rangle$ state with a π -pulse, the average measured frequency difference is only half of the dispersive shift χ .

Besides these basic spectroscopy measurements, there are a few other measurements that are commonly done in order to characterize a transmon qubit. These include: cavity transmission as a function of power to find a good readout power, an AC stark shift measurement to calibrate the cavity photon number, and Rabi oscillations to calibrate π -pulses.

Once the qubit frequency is well localized and the π -pulse is well-tuned, one can




Figure 5.9: Measurement of T_1 , T_2^R , and T_2^E . Each fit is to an exponential decay, with T_2^R being measured on resonance. Each of these measurements was selected as the best non-outlier measurement of each coherence time. They are not from the same qubit or same experimental run.

measure the coherence times T_1 , T_2^R , and T_2^E . Fig. 5.9 shows the measurement traces of the highest coherence times for compact transmons. Note that the T_1 is considerably longer than the longest CPW transmon.

Fig. 5.10 shows the statistics and repeatability of compact transmons by including the maximum of each T_1 , T_2^R , and T_2^E for each of the 11 measured compact transmon devices plotted versus qubit transition frequency. Note that each qubit is tunable, and only the maximum value of each coherence time is quoted; the value for each coherence time is likely from a different qubit frequency. The dependence of coherence on qubit frequency is discussed next. While the maximum T_1 is 13 μs , there were four devices measured at or above 8 μs , showing some amount of repeatability. Each of the compact transmons measured were better than the best CPW transmon, with the minimum compact transmon being around 4 μs . The T_1 tends to increase at lower frequencies in a manner faster than a constant Q would predict. The two dashed lines show the T_1 expected for two values of Q , calculated from $Q = \omega_0 T_1$.

The values of T_2^R and T_2^E shown in Fig. 5.9 were chosen at the maximum qubit frequency, or “sweet-spot”. At this point, corresponding to zero flux through the transmon SQUID loop, the qubit frequency is first-order insensitive to flux noise. In the case of these qubits, this often leads to a maximum in T_2^R . As the qubit is tuned




Figure 5.10: Summary of coherence results for 11 compact transmon qubits; the best T_1 , T_2^R , and T_2^E for each qubit is plotted versus qubit frequency. Note that T_1 is in general higher than T_2^E , which is in general higher than T_2^R . There is a trend for higher T_1 values at lower frequencies, two dashed lines of $Q = 120,000$ and $Q = 200,000$ are added for reference.

down the slope, the sensitivity to flux noise increases and T_2^R decreases. T_2^E usually follows this trend, although if the flux noise is from low-frequencies, then the echo sequence diminishes this effect.

The dependence of each of the three coherence times on qubit frequency is an interesting method for learning what decoherence or dissipation methods are dominant. Different mechanisms of decoherence produce distinguishable dependences of T_1 , T_2^R , and T_2^E versus qubit frequency. While each qubit had its own peculiarities, many of them share common features. For example, Fig. 5.11 shows a typical coherence versus qubit frequency plot. The maximum qubit frequency for this qubit was 4.93 GHz with the ability to tune down to 3.22 GHz. While the qubit is expected to be tunable to a much lower frequency, measurements are often limited by either diminishing readout contrast or quickly increasing charge dispersion.

T_1 in this graph can be seen to increase significantly over the swept frequency range, rising as high as almost 8 μs at 3.22 GHz and dropping as low as almost 1 μs at 4.93 GHz. The solid red line shows the predicted Purcell limit, and shows the limit of T_1 due to the coupling to the cavity. To measure the effect of other loss mechanisms, it is best to increase the Purcell limit much higher than observed T_1 , whereas here it cannot be neglected. Here, the T_1 matches well to Purcell plus a constant Q between 120,000 and 200,000, as shown by the dashed lines.

T_2 as shown in this figure, appears to be limited by both low-frequency and high-frequency noise. The fact that T_2^E is everywhere larger than or equal to T_2^R shows that the qubit is subject to some low-frequency noise that the echo successfully cancels out. However, since T_2^E is only, in the best cases, equal to T_1 , shows that there is also some decoherence due to high-frequency noise, that the echo cannot remove. This is a common trait among the compact transmons; most devices show $T_1 \geq T_2^E \geq T_2^R$. In no device at any frequency was T_2^E above T_1 . Another observation from this graph




Figure 5.11: Example of coherence times versus qubit frequency for compact transmon. The solid red curve is the predicted single-mode Purcell limit estimate, while the dashed curves correspond to adding a further dissipation channel with a $Q = 120,000$ and $200,000$ independent of frequency.

is that T_2^R drops away from the flux sweet-spot maximum of just over $1 \mu\text{s}$. While the highest T_2^R observed at the sweet spot was $3.0 \mu\text{s}$, most devices were closer to $1 \mu\text{s}$.

Another interesting measurement is a method to indirectly measure the Purcell limit, not requiring any transmon theory. The Purcell limit attempts to calculate the coupling strength between the qubit and the dissipation of the cavity. If the cavity is under-coupled, this is represented by the various loss mechanisms of the cavity; however, if the cavity is over-coupled, this represents mainly the coupling of the cavity to the microwave line impedance. Therefore, in the case of an over-coupled cavity, the Purcell limit is determined by the coupling strength between the qubit and the microwave lines. Another way to directly measure this coupling strength is to measure the power required to complete a π -pulse given a fixed pulse shape and width. This measurement, in comparison to the Purcell limit is shown in Fig. 5.12. The Purcell limit is shown in log-scale because the π -pulse power is represented in dBm. Note that the π -pulse power changes by almost 10 dBm, corresponding to a factor of 10 change in power, and that the Purcell limit T_1 also changes by a factor of 10. While the overall shape is the same, the π -pulse measurement takes into account the actual impedances, resulting in the typical oscillations of slightly mismatched microwave lines. The qubit for which this was measured was not sufficiently Purcell limited in order to observe these oscillations in the qubit T_1 measurements.

A discussion in Chapter 4 highlighted results showing that the quality factor of compact resonators depended strongly on the *sample* temperature due to the dominance of TLS dissipation. The sample temperature was shown, in the case of the Kelvinox25 refrigerator to be considerably higher (200 mK) than the mixing chamber temperature (80 mK). This effect is also very important for qubit coherence times. Luckily, for the case of qubits, it is possible to directly measure the *temperature*




Figure 5.12: (a) Calibrated power to perform a π -pulse for a fixed qubit pulse shape. (b) Calculated Purcell limit in log-scale over the same frequency range. The coupling of the microwave line to the qubit, which is responsible for the Purcell limit of T_1 , can be measured via the power required to induce π -pulse. These graphs show that the expected shape of the Purcell limit is well-matched by the measured π -pulse power.




Figure 5.13: T_1 for a compact transmon device measured versus temperature. The red points in both graphs show T_1 plotted against the mixing chamber temperature, while the green points are plotted against the qubit temperature measured from RPM. Measurements are done (a) before modifications and (b) after modifications intended to produce lower qubit temperature.

by accurately quantifying the equilibrium amount of excited state population. Qubit temperature and methods to measure it are described in Appendix B. Using the Rabi Population Measurement, or RPM, method to measure qubit population, Fig. 5.13 shows how it is possible to be misled by certain measurements.

This graph shows how T_1 depends on temperature for a compact transmon, both before and after modifications intended to lower the qubit temperature. In the first graph, it looks like T_1 reaches a maximum at low temperatures when plotted against the mixing chamber temperature. However, when plotted against qubit temperature,

it is completely unclear whether the maximum is yet reached. In the second graph, the qubit temperature reaches lower values and confirms that the qubit has indeed reached a plateau in T_1 . While the second graph did not disprove the false conclusion from the first graph, it highlights the danger in assuming any sample temperature is equal to the mixing chamber. The modifications to the experimental setup are described in the previous section (Section 5.2).

5.4 Analysis

Comparing CPW and compact transmons, it is clear that compact transmons have significantly higher coherence times. Fig. 5.2 shows that relaxation times of CPW transmons seemed to be bounded by a quality factor of about 70,000, achieving almost $4 \mu\text{s}$ at around 3 GHz. It should be noted that this was in fact that longest lifetime of any CPW transmon; a typical lifetime was around $1 \mu\text{s}$ due to optimization for other parameters. Fig. 5.10 shows the summary of relaxation and coherence times for the 11 compact transmon devices; each of them individually shows a quality factor above 70,000. Most devices are between a quality factor of 120,000 and 200,000, while the highest quality factor is 250,000, with a lifetime of $13 \mu\text{s}$ at around 3 GHz. Thus, compact transmons typically achieved an improvement of about a factor of 3 (or 200% increase) in qubit lifetime.

While a small improvement was achieved, there were many specific improvements that could be responsible. The focus of this experiment was on the increased value of the width and gap of the transmon capacitance, other changes may have contributed to the increase in relaxation time. Other changes include: switching from silicon to sapphire substrates, switching from CPW to compact resonator as cavity, and many modifications to the experimental setup to improve shielding over-

all. While it has not been proven, sapphire has often been suggested to have lower loss than silicon, so this change may be responsible for some of the improvement. Experiments by Corcoles et al. [79] for qubits and Barends et al. [77] for resonators have shown that preventing stray infrared radiation from reaching the qubit or resonator results in drastic improvement to lifetimes. It has not been ruled out that experimental setup improvements are responsible for the relaxation time increase. While the change from CPW to compact resonator has resulted in using a higher cavity Q_i , the cavities are, in these experiments, always over-coupled. Therefore, this is unlikely to influence the qubit lifetime.

Although the improvement of relaxation time was the focus of this experiment, increases in coherence time are a long term necessity for quantum computation. CPW transmons typically had T_2^R and T_2^E values of around $1 \mu s$, showing that often, neither T_2^R nor T_2^E were limited by T_1 processes. Improving the relaxation time allows for higher coherence times, and in general the compact resonator coherence times were longer. As shown in Figures 5.10 and 5.11, T_2^E is almost always considerably longer than T_2^R , which is often limited at 1-2 μs at the flux sweet spot. This shows that low-frequency noise plays a role, and the fact that T_2^R decreases quickly away from the sweet spot shows that this is likely flux noise. Flux noise can either come from external noise, for example from the magnetic field coil current fluctuations, or from local noise, such as microscopic fluctuators on the surface of the SQUID loop. Experiments explicitly for the purpose of increasing the dephasing times, or to deduce their source, were not completed.

The importance of sample temperature cannot be overstated. Techniques such as RPM (see Appendix B) accurately measure the effective qubit temperature. This temperature is often drastically different from the mixing chamber temperature, depending on many details of the experimental setup. If this temperature is high

enough, it has been shown in some cases to be severely limiting qubit lifetime [79]. In the experiments here, qubit lifetime did not improve upon lowering qubit temperature, however it seems the temperature was *just* low enough, as shown in Fig. 5.13.

In the broad sense, it is not always true that the lowest temperature produces the best quality factors. If the temperature is high enough to increase the loss due to quasiparticles [62], then lower temperature will improve relaxation times. However, if the relaxation time at an elevated temperature is limited by TLS dielectric loss, then higher temperatures produce higher relaxation times. This was observed directly in Chapter 4, where measurements at 200 mK produced resonator quality factors roughly twice as high as measurements done at roughly 50 mK. Since the goal of this experiment was to apply lessons learned from compact resonators to transmons, the temperature is extremely important.

5.5 Discussion

The compact transmon design optimized for this experiment was used by other groups because of its good performance. In fact, the Berkeley group led by Irfan Siddiqi, openly borrowed an almost exact design of the compact transmon for their paper showing quantum jumps [98]. Similarly, flux qubits were also coupled to compact resonators for additional experiments [72].

While an increase in relaxation times was observed, possibly due to the increased width and gap of transmon capacitance, a thorough study of this effect was not completed. It would be interesting to measure the lifetimes of various qubits with different widths and gaps, much in the same way a variety of resonators were measured in Chapter 4. This study would require a lot of simulation, because one would

have to match capacitances with a wide variety of structures. Results from this experiment would shed light on whether the improvements observed were due to the increased width and gap or something else entirely.

Results from two recent experiments have done exactly as described, both by UCSB and IBM. These transmons have been tested with varying widths and gaps, in much the same way as the compact transmon described here. The UCSB group shows consistent lifetimes of 10-30 μs , with some measurements as high as 44 μs using their “Xmon” design, which includes larger width and gap parameters [99]. The Xmon is a transmon that uses a distributed capacitance to ground instead of an interdigitated capacitance. This capacitance has four legs, allowing for the coupling to four different components (resonators, feedlines, etc...) at once, thus the shape resembles an “X”, hence the name. IBM shows lifetimes of up to 60 μs by making the qubit capacitor out of TiN along with similar geometrical design changes [52]. The improvement observed in both of these groups is not surprising and is a natural extension to the work presented in this thesis.

In conclusion, the coherence improvements observed through the measurement of 11 compact transmon devices validate the analysis technique of participation ratios to superconducting qubits. Adjustments of capacitor width and gap that resulted in a decrease in surface participation, without changing the material components or processing techniques, led to an increase in coherence of both resonators and qubits. Compact transmons exhibit relaxation times as high as 13 μs , when compared to the maximum of 4 μs of CPW transmons. This increase in coherence represents an immediate advance towards the goal of improving superconducting qubit coherence. In the next chapter, now that the participation ratio method has been validated for small changes, the reduction in surface participation is taken to the extreme.

3D Transmon

6.1 Introduction

6.1.1 Origin of 3D Transmon

Chapters 4 and 5 describe experiments in which the geometry of resonators and qubits were slightly modified as a means of reducing the participation of lossy surfaces. The reduction in participation of surfaces resulted in an increase in relaxation times, validating the analysis method of participation ratios. With newfound confidence in the participation ratio method, the transmon was fully redesigned in order to force the participation of lossy materials to be much lower. To have quality factors of higher than 1,000,000, any material that has a 1% participation must have a Q of greater than 10,000, and a material with 0.1% participation must have a Q of greater than 1,000. In a system with lossy materials such as a PCB, GE varnish, and wirebonds, it is difficult to simulate properly and ensure that the participation to these offending materials is low enough. To guarantee that the transmon qubit is not limited by these lossy materials, they must simply be removed from the system.

In addition to removing dissipative materials, radiative loss must also be mini-

mized. The presence of a sample box is necessary to remove radiative losses to zeroth order and also to shield the chip from infrared radiation, yet the addition of a box usually adds spurious “box modes” which also contribute to dissipation. The octobox (see Section 3.1.2), used for experiments on CPW and compact transmons has box modes with very low quality factors due to the lack of attention to important details like how the box was closed and the wire-bonds and PCB inside the box. Realizing that the coupling to these modes cannot be reduced indefinitely, one must design a better sample box.

Thus, in order to remove dissipative materials and reduce loss through box modes, the transmon must be completely redesigned. The beauty of the “3rd generation” transmon is that both of these changes are implemented through the same change: the substitution of a 3D cavity resonator in place of the 2D compact resonator; the simple 3D cavity geometry, with no extra dissipative materials, provides both the readout resonator and a shield for infrared radiation. The radiation to box modes is removed, not by reducing the coupling between the transmon and box modes, but by making the box modes higher quality; in fact, one sample box mode doubles as the readout mode for the 3D transmon. The structure of the 3D cavity, as explained in Section 2.1.2, is composed of a machined chamber inside a solid block of metal. The 3D transmon is placed on a chip and mounted in the center of the cavity, held in place by friction between the two halves. This simple structure removes all lossy materials and is fully simulatable and possibly even calculable analytically.

The participation ratio benefit of the 3D transmon becomes clear when estimating the participation of surface layers. While a proper surface participation ratio can be calculated from simulation results, a simple estimation is enough to illustrate the difference between 3D and compact geometries. As shown in Fig. 6.1, a simple estimate compares the fraction of the volume filled by the surface: $p_{surf} = t_{surf}/t_{mode}$,




Figure 6.1: Rough estimate of surface participation in (a) 3D resonator or 3D transmon versus (b) compact resonator or compact transmon. This estimate calculates the fraction of mode volume filled with the offending surface. Assuming a surface thickness of approximately 5 nm, the surface participation of compact geometry is approximately 10^{-3} versus 10^{-6} in the 3D geometry.

where p_{surf} is the surface participation ratio, t_{surf} is the thickness of the surface layer in question, and t_{mode} is the distance between the two conductors in the mode. In the case of compact resonators or compact transmons, t_{mode} is equivalent to g_C , and supposing $t_{surf} = 5$ nm, a rough estimate of the surface participation is $5 \text{ nm}/5 \text{ μm} \sim 10^{-3}$. In a 3D cavity and a 3D transmon, the mode volume was increased significantly, thus increasing t_{mode} as well. A rough estimate of the surface participation in the 3D geometry is $5 \text{ nm}/5 \text{ mm} \sim 10^{-6}$. Therefore, by rough estimates, if a 3D cavity and compact resonator would both be limited by dissipation in a small surface layer (whatever that dissipation may be), the 3D cavity would have roughly a factor of 1,000 longer relaxation time. Since a 3D transmon necessarily (so far) includes the addition of planar structures inside the 3D cavity, they may not realize the full improvement as seen by 3D cavities.

6.1.2 Existing Implementations

The 3D transmon was first conceived and implemented at Yale University around 2010. After several different cavity types were tested, the cavity/sample box that was chosen for the first 3D transmon was a simple rectangular cavity. The modes




Figure 6.2: Picture of two halves of copper rectangular cavity. The cavity is simply milled out of a block of copper; the two halves combine to form one cavity. There is a ledge on the half on the right that allows one to place chips. Coupling to the cavity from microwave lines is achieved through one hole on either half of the cavity. Coupling strength is adjusted by changing the length of the pin sticking into the cavity from this hole.

of a rectangular cavity can be calculated analytically using formulas found in many books [53], and discussed in Section 2.1.2.

The most commonly used rectangular cavity used for the experiments at Yale is shown in Fig. 6.2. The cavity in this image is made out of copper, which limits the quality factor of all resonant modes to approximately 10,000 since copper remains a normal metal even at dilution refrigerator temperatures. Aluminum cavities of the same variety produce quality factors ranging from 1 to 50 million depending on various cleaning, machining and material properties. The discussion of 3D resonator quality factors is out of the scope of this thesis, but is detailed in a study by Reagor et al. [100].

Coupling qubits to 3D cavities is done in much the same way as a 2D cavity; the coupling is by capacitance between the two halves of the Josephson tunnel junction and the two voltage anti-nodes of the cavity. This is achieved in the 3D architecture by adding two large “pads” to the Josephson tunnel junction, one on either side, as




Figure 6.3: Picture of first 3D Transmon device from H. Paik et al. [101]. On the left, one half of the 3D cavity is shown. The design is similar to the updated cavity of Fig. 6.2, but this version is made of aluminum. On the right, a close up image of the antenna pads is shown. The junction is too small to be seen, but it connects the two pads in the center. The exact antenna pad geometry varies significantly among the 3D transmon devices made for different experiments.

shown in Fig. 6.3. The exact capacitor design varies from experiment to experiment, but is typically rectangular on the order of $100 \times 100 \mu\text{m}$. The total capacitance C_Σ of the transmon is highly affected by the cavity walls and therefore the cavity walls participate strongly in the qubit mode.

Coherence results of 3D transmons are extremely promising, with T_1 , T_2^R and T_2^E times each measured over $100 \mu\text{s}$. The original paper and results by Paik et al. [101] detailed T_1 values of between $20 \mu\text{s}$ and $60 \mu\text{s}$, with T_2^R typically in the range of $15\text{-}20 \mu\text{s}$. Echo typically increased coherence times by less than 20%. These results are approximately 4 times higher, both in T_1 and T_2^R than the compact transmons and 16 times higher than the best CPW transmon.

In the time since the original 3D transmon measurements, drastic improvements in coherence have been achieved. Work both at Yale and IBM [42, 43] showed that T_2^R was often limited by photon induced dephasing, described in Section 1.3.2. The

Kerr between the qubit and various cavity modes, combined with non-zero cavity populations, was found to be limiting T_2^R in the first 3D transmon experiments. A change of one photon in a cavity mode whose dispersive shift with the qubit is larger than the qubit linewidth completely dephases the qubit. The statistics of how often these photon jumps occurs, i.e. photon shot noise, determines the dephasing time of the qubit. For modes this strongly coupled, an T_2^R can be calculated from the steady state population \bar{n} and the cavity linewidth κ , via $1/T_2^R = \bar{n}\kappa$.

From these experiments on photon induced dephasing, it is clear that cavity populations must decrease in order to improve T_2^R . Decoupling the cavity from the environment and implementing better filtering have shown that the intrinsic dephasing times are at least several hundred microseconds [43, 42], resulting in coherence times of over 100 μs . A side effect of improved filtering and thermalization is that 3D transmon relaxation times have *also* increased to over 100 μs , indicating that either quasiparticles or infrared photons were previously limiting transmon relaxation.

These results spurred a revolution in the superconducting community, with many groups around the world quickly adopting 3D transmons. The fact that many groups duplicated the coherence results showed that the participation ratio idea underlying the 3D transmon is a very robust way to improve coherence, and does not rely on special fabrication methods. Their enhanced coherence times and simple fabrication allow for many experiments that may not previously have been possible. These experiments include initialization by measurement [70], initialization to an arbitrary point on the Bloch sphere by cavity-assisted bath engineering [102], stabilizing Rabi oscillations with feedback [103], measuring back-action of weak measurements [69], stabilizing a qubit trajectory [104], and measuring charge-parity fluctuation statistics [105].

6.1.3 3D Transmon as “Calibration”

As a means of confirming the experimental setup and as a comparison, under nearly identical conditions, to the fluxonium qubit results presented in Chapter 7, a 3D transmon was fabricated and measured for this thesis. Since it has been shown that the experimental environment is critical to maximizing the performance of transmon qubits, it is important to qualify any new experimental setup with existing designs before exploring new territory; this was a primary motivation for the measurement of this 3D transmon.

The measurements in this chapter were performed in a standard circuit QED setup on an aluminum transmon qubit, fabricated using a bridgeless double-angle evaporation technique [66, 67]. This fabrication technique is different from the Dolan bridge technique used for the original 3D transmon results. The bridge-free technique was chosen here because it was planned to be used for the fluxonium fabrication; its use in the 3D transmon measurements improves the future comparison.

To make further use of this experiment, this transmon was fabricated using a SQUID and measured inside a copper 3D cavity to allow an external coil to apply flux and control the qubit frequency (in the same manner as the compact transmons in Chapter 5).

6.1.4 Reset Experiment

With its extended coherence time, the 3D transmon qubit permits many experiments that were previously out of reach or impractical. One such experiment that takes advantage of the 3D transmon capabilities is a measurement of a new qubit reset mechanism. Qubit reset is an integral part of any quantum algorithm and is required at the beginning of and during algorithms as a means of initializing the qubit in a

known state. Any reset will be characterized on its speed and fidelity, with speed representing the time to initialize the state from any existing state and fidelity being the overlap with the expected initialization state. This reset mechanism works via redirecting energy of the qubit into the dissipation mechanism of the cavity. Named DDROP, for Double Drive Reset Of Population, this protocol is described in detail in Appendix C.

The existence of a proper qubit initialization mechanism is of great importance, not only for the eventual quantum algorithm, but also for the characterization of other quantum elements. For instance, a measurement of the fidelity of any gate will be limited by the fidelity of the initial state. Thus, to measure a 99.9% fidelity gate, one must start with a minimum ground state preparation fidelity of 99.9%. This means that a poor qubit initialization mechanism will limit every other aspect the quantum algorithm, and is thus one of the most important elements of a quantum system.

Qubit reset, in the case of a qubit with high equilibrium excited state population, also acts as a cooling mechanism. Initializing the qubit to the ground state requires (on average) removal of energy, thus the reset acts as a type of cooling. High temperature qubits are common [79], and can limit coherence, as discussed in Chapter 5. If qubit coherence times are limited by coupling to an environment at an elevated temperature, it is possible that one may cool this environment by continuously removing energy from the qubit. Results from this experiment will show whether this was achieved.

6.2 Experimental Apparatus

The experiment in this chapter was performed in the Kelvinox400 dilution refrigerator. The transmon sample was mounted inside a 3D copper cavity, which was thermally anchored to the mixing chamber of a dilution refrigerator with a base temperature of about 15 mK. The cavity was mounted inside a copper shield coated with infrared absorbing material on the inside. A high-frequency Eccosorb “box” filter, and a microwave 12 GHz low-pass filter were placed on each input and output microwave line. Two 8-12 GHz circulators were installed between the cavity and the HEMT amplifier. The microwave filters and infrared shield are described in more detail in Section 3.3.3.

6.3 Results

6.3.1 Coherence Results

The 3D transmon measurement protocol is exactly identical to that of the compact transmon, the differences in implementation may change parameters, but not the basic idea. Therefore, the basic results such as spectroscopy are very similar to that of the compact transmon. The qubit frequency still has a maximum frequency that doubles as a sweet spot for flux noise, while the Purcell effect still reduces the T_1 when the qubit frequency is close to the cavity frequency.

Measurements of the various coherence times as a function of flux are shown in Fig. 6.5. The maximum qubit frequency for this 3D transmon was 5.0 GHz; magnetic flux tuned the qubit frequency down to 4.4 GHz before the readout became too weak. Although it could be overlooked in the log-scale plot, T_1 increased as the

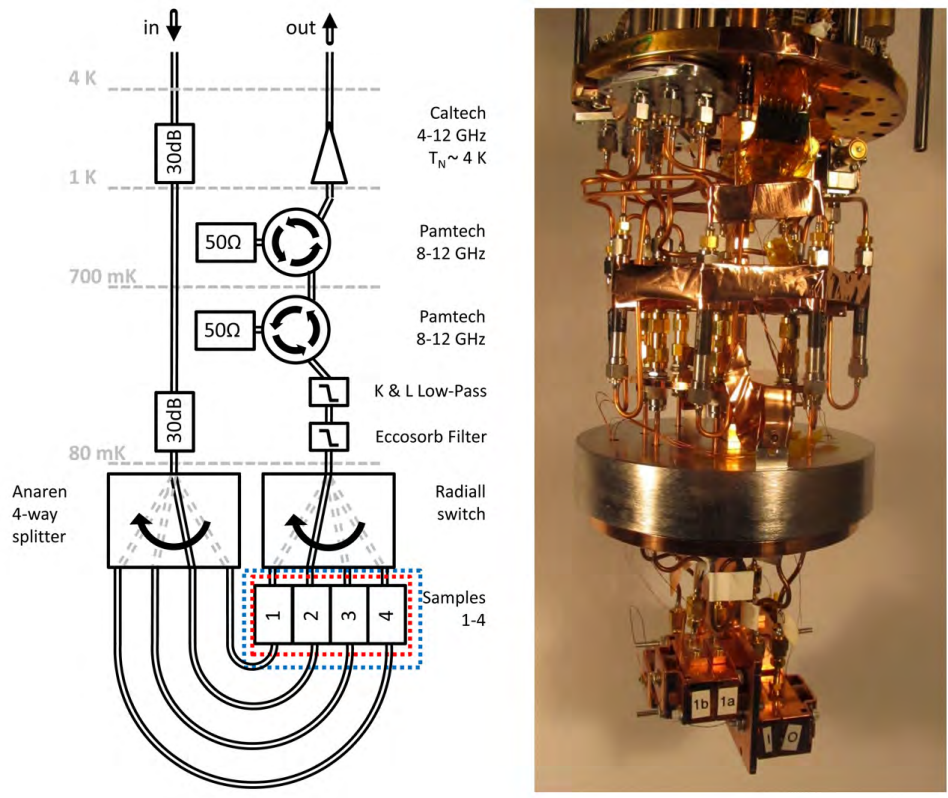


Figure 6.4: Schematic of microwave lines in Kelvinox400 for 3D transmon measurements and a picture of actual setup. The components on the schematic are labeled, not every component is visible in the included picture. The blue outline indicates the cryoperm shield and the red outline indicates the infrared shield. The bottom “can” of the infrared shield and cryoperm are not shown.




Figure 6.5: Measurements of T_1 and T_2^E versus qubit frequency for a tunable 3D transmon in a copper cavity. These measurements show a significant increase in T_1 as the frequency is reduced along with a drastic drop in T_2^E away from the maximum frequency sweet spot. T_2^R measurements are not included since the automated procedure failed to produce good results for this frequency sweep.

frequency was decreased from approximately $45 \mu\text{s}$ to approximately $65 \mu\text{s}$. Repeat measurements at the lowest frequency place this number at $70 \mu\text{s}$. This is a large increase in T_1 over a relatively short frequency span.

The Purcell effect calculation for 3D transmons is currently being investigated, but it is clear from this sample alone that the approximation in the single-mode strong-dispersive regime (described in Section 1.3.1) is insufficient. For the parameters at the maximum qubit frequency ($f_c^g = 9.1 \text{ GHz}$, $f_g^0 = 5.0 \text{ GHz}$, $\kappa/2\pi = 3 \text{ MHz}$, $\chi/2\pi = 7 \text{ MHz}$, $T_1 = 37 \mu\text{s}$), the estimated Purcell limit using this approximation yields $13 \mu\text{s}$. The observed value at this frequency is a factor of 3-4 higher than this estimate, highlighting the need for a more complete Purcell theory. At 4.3 GHz , this Purcell estimate rises to $21 \mu\text{s}$, still far from the measured T_1 . Note that the increase in T_1 from 5.0 to 4.3 GHz is faster than one would expect in the case of a constant Q ; $70 \mu\text{s}$ at 4.3 GHz corresponds to a Q of 1.9×10^6 while a T_1 of $40 \mu\text{s}$ at 5.0 GHz corresponds to a Q of 1.3×10^6 .

While the T_1 can be considered on par with typical (not the best) 3D transmon results, the T_2^E results are significantly lower away from the flux sweet spot. At the




Figure 6.6: Individual measurements of T_1 , T_2^R , and T_2^E at the maximum qubit frequency of 5.02 GHz. This frequency is first order insensitive to flux noise due to the zero derivative of frequency with applied flux. This results in the highest values of T_2^R and T_2^E for this qubit.

flux sweet spot, $T_2^E \approx T_1$, so the dephasing time $T_\phi \approx T_1$. This is likely due to flux noise as the T_2^E is observed to decrease very rapidly away from the sweet spot.

Individual measurements of each T_1 , T_2^R , and T_2^E are shown for qubit frequency 5.0 GHz in Fig. 6.6 and for 4.3 GHz in Fig. 6.7. Since T_2^R is considerably lower than T_2^E in both cases, it is clear that there is some high frequency noise, likely flux noise, influencing the qubit coherence. The fact that the echo procedure does not raise T_2^E as high as $2 T_1$ means that there is also some low-frequency noise. The value of T_2^R reduces drastically away from the sweet spot, dropping a factor of 40 from $20 \mu\text{s}$ to $0.5 \mu\text{s}$ while the T_2^E drops only a factor of 4 from $39 \mu\text{s}$ to $11 \mu\text{s}$.

On the scale of several hours, the relaxation time at either of these fluxes was noticed to fluctuate by almost a factor of 2. While not observed on every qubit device, this is a common occurrence and has been observed in compact transmons, this 3D transmon, and the fluxonium. These fluctuations may be due to subtle drifts in the amount of quasiparticles or accessible TLSs. An example measurement on the 3D transmon of T_1 variations over 14 hours is shown in Fig. 6.8.




Figure 6.7: Individual measurements of T_1 , T_2^R , and T_2^E at the qubit frequency of 4.3 GHz. This qubit at this frequency is much more sensitive to flux noise due to the non-zero derivative of frequency with applied flux. This results in much lower values of T_2^R and T_2^E . T_1 is higher at this frequency either due to the Purcell limit or due to lower loss at lower frequencies.




Figure 6.8: Measurement of relaxation time T_1 over the course of 14 hours for 3D transmon qubit. Each measurement here consisted of approximately 8 minutes of averaging. The fluctuation is nearly a factor of 2 from minimum to maximum, and there appears to be some structure in the fluctuations.

6.3.2 Qubit Reset

The DDROP reset mechanism, tested on this 3D transmon utilizes many of the 3D transmon advantages over previous transmon designs. Qubit and cavity parameters at the flux used for this experiment were measured to be $f_c^g = 9.1$ GHz, $f_g^0 = 5.0$ GHz, $\kappa/2\pi = 3$ MHz, $\chi/2\pi = 7$ MHz, $T_1 = 37 \mu\text{s}$, $T_2^R = 20 \mu\text{s}$, $T_2^E = 40 \mu\text{s}$, steady state $P_e = 9\%$, $\Gamma_\uparrow/2\pi \approx 400$ Hz.

As explained in Appendix C, for the DDROP protocol to improve the initialization of the ground state, there are two requirements: $\chi/\kappa > 2$ and $\kappa/\Gamma_\uparrow \gg 1$. For the parameters listed above, these requirements are met, with $\chi/\kappa = 2.3$ and $\kappa/\Gamma_\uparrow \simeq 8,000$, indicating that DDROP should prepare ground state better than the steady state 91%.

The effect of the DDROP protocol on this qubit is shown in Fig. 6.9, where the y -axis is the measured excited state population and the x -axis is the duration of the reset pulses (or delay time in the case without DDROP). Each data point is taken after waiting $1 \mu\text{s}$ ($20 \kappa^{-1}$) after the end of the DDROP pulse to allow the system adequate time to decay from $|g, \alpha\rangle$ to $|g, 0\rangle$. The two solid, nearly horizontal curves are the pre-reset ground and excited qubit states without DDROP pulse. The pre-reset is itself a $5 \mu\text{s}$ DDROP sequence done before all other pulses in order to suppress the initial excited state population. The slight downward trend in the excited state curve, due to the finite value of T_1 , is barely noticeable on this scale. The other two solid curves correspond to the same preparation, but show the effect of a DDROP pulse whose duration is varied across the x -axis. At short pulse duration, both initial populations tend towards 50% excitation, due to the Rabi drive. As the duration is increased, the population tends quickly towards the pre-reset ground state.

The four dashed curves represent an identical set of data taken without the pre-




Figure 6.9: Measured excited state population after reset pulse of varying duration, for four different initial preparations, measured after intervals of 40 ns. The solid lines include a pre-reset while the dashed lines begin with the steady state 9% excited population. The ‘w/ π pulse’ curve shows a slight downward trend due to the finite T_1 . The curves with DDROP show that, regardless of initial state, the qubit is driven to the ground state for pulse durations less than 2 μs . For this measurement, $\Omega_R \approx 0.8 \kappa$ and $\bar{n} = 8$.

reset, thus showing the effect of initial equilibrium population. Note that regardless of the initial state, DDROP forces the population to the ground state in less than 3 μs (including the 1 μs decay from $|g, \alpha\rangle$ to $|g, 0\rangle$). This is a factor 60 improvement over the standard protocol of waiting $5 T_1$, which would give a comparable reduction of excited state population in a *cold* qubit environment. With a significant initial excited state population, DDROP is not only faster, but better than simply waiting.

In order to determine the fidelity of the ground state preparation, it is important to characterize the resulting excited state population accurately. The measurement scheme presented in Appendix B allows for just that. Called the RPM (Rabi Population Method) for measuring temperature, this procedure accurately measures the excited state population.

To optimize this ground state initialization fidelity of DDROP, numerical simulations were performed of the expected fidelity F versus qubit drive amplitude and average cavity excitation, Ω_R and \bar{n} , respectively. The simulations were of the Lind-




Figure 6.10: (a) Contours of 90, 95 and 99% predicted ground state preparation fidelity from numerical simulations versus two Rabi drive amplitudes expressed as Ω_R/κ and \bar{n} . For fidelities greater than 99%, color of shaded area indicates reset time. (b) Measured excited state population from RPM method (crosses with error bars) compared to numerical simulation (solid line) versus \bar{n} for $\Omega_R/\kappa = 0.8$. This population decreases monotonically with \bar{n}

blad master equation obeyed by the qubit-cavity density operator, including the two drives and decoherence for both the qubit and the cavity. The initial state was chosen to be the cavity in vacuum and an equilibrium state for the qubit. The dependence of F on Ω_R for fixed \bar{n} was found to be weak and fidelities above 99% were found for $\Omega_R/\kappa > 0.3$. The numerical simulations show that F increases monotonically with \bar{n} for a fixed Ω_R and that with a higher Ω_R , higher \bar{n} is required to reach the same fidelity, as shown by the contours of constant fidelity in Fig. 6.10(a). The simulations did not account for self-Kerr effects that will reduce the fidelity at photon numbers much higher than the range shown.

With the guidance provided by these simulations and using RPM to experimentally quantify the fidelity, DDROP has been studied for a wide range of Ω_R and \bar{n} . The pulse duration was kept fixed at the value $5 \mu\text{s}$, chosen from simulation, to ensure DDROP has reached equilibrium in all conditions. Fidelities greater than 99% were achieved for Ω_R as low as 0.3κ and as high as 1.0κ , for $8 \leq \bar{n} \leq 50$.




Figure 6.11: Measurements and simulation of driven steady state qubit temperature (an alternate representation of P_e) versus Ω_R at fixed \bar{n} . Simulations with $\bar{n} = 4$ matches measurements well. Note that the dependence of temperature on Ω_R is weak, with only a 100% increase in qubit temperature (a much smaller increase in P_e) while increasing Ω_R/κ from 0.2 to 1.4.

For fixed $\Omega_R = 0.8 \kappa$, Fig. 6.10(b) shows measurement (markers) versus simulation (line) of remaining excited state population versus \bar{n} . Excited state population drops monotonically with \bar{n} , in good agreement with numerical simulation. On the other hand, above approximately $\bar{n} = 50$, the reset excited state population increased significantly. This is understood to be due to the breakdown of the dispersive approximation.

Measurements taken at fixed \bar{n} versus Ω_R also show decent agreement with theory, although the dependence is very weak. Fig. 6.11 shows measured qubit temperature versus Ω_R along with two predictions from simulation at $\bar{n} = 4$ and $\bar{n} = 8$. The measured data matches closely with the $\bar{n} = 4$ simulation, and the resultant qubit temperature changes very little with a wide change in Ω_R . The simulation with higher photon number predicts lower qubit temperature, as is shown in detail in Fig. 6.10(b). Overall, both drive amplitude parameters Ω_R and \bar{n} have a wide range for which DDROP works well, making it a reliable and stable protocol.

Note that the excited state population data in Fig. 6.11 is represented as a *effective* qubit temperature instead of an excited state percentage (P_e). These two

representations are equivalent, and conversions between the two can be made by interpreting P_e in terms of a Boltzmann distribution as shown in Eq. 6.1. It is natural to compare the effective qubit temperature representation with the mixing chamber temperature to see if an additional heating source is suspected. While the qubit temperature *assumes* a Boltzmann distribution and may not be the result of contact with a thermal bath, the effective temperature can give a more intuitive picture of the severity of the excited state population. In cases where the distribution is *proven* not to be thermal, the percentage excitation should be used.

$$\frac{P_e}{1 - P_e} = \exp\left(-\frac{hf_{ge}^0}{k_B T}\right) \quad (6.1)$$

While 99% fidelity is reached with a wide range of parameters, reset time is optimized when $\Omega_R \simeq \kappa$, yielding reset times comparable with those of two-qubit gates in the circuit QED architecture [106]. The reset time for the ground state population to reach 99% is shown by the colored pixels of Fig. 6.10(a).

As mentioned before, all of the DDROP characterization measurements included a 1 μ s ($20 \kappa^{-1}$) wait between drive pulses and the RPM measurement, to allow the cavity photons to decay. Therefore, the qubit excited state population begins returning to its equilibrium value as soon as the reset drives are turned off. This re-equilibration should occur on a timescale given by the mixing time T_1 , and this is what is found experimentally and shown in Fig. 6.12.

6.4 Analysis

A discussion of the nuance between cooling and reset is now in order. Qubit reset is ground state preparation, aiming for the highest fidelity in the shortest possible time, whereas qubit cooling reduces the excited state population below that




Figure 6.12: RPM excited state population measurements at various times after a DDROP pulse. At the shortest time of $1 \mu\text{s}$ after DDROP, P_e is indistinguishable from 0 and P_e rises quickly until it saturates at the steady state value of 9%. The curve is the simulation (not a fit) given the parameters of the DDROP pulse and steady state P_e , indicating great agreement between theory and experiment.

produced by contact with the external bath (steady state). As shown by these experiments, DDROP satisfies both definitions, yet it differs significantly from other dynamical cooling procedures. These methods, inherited from their counterpart in atomic physics [107], have been recently demonstrated in both nanomechanical systems [108, 109, 110] and superconducting qubits [111, 112, 36, 102].

It is important to note that although DDROP can be said to *cool* the qubit, it has been observed to have no effect on the qubit T_1 . While qubit T_1 certainly depends on qubit temperature, this type of cooling (state preparation) does not reduce the relaxation effects due to higher temperature (quasiparticles, for example). In fact, as shown in Fig. 6.12, it is clear that since the population re-equilibrates with a time constant equal to T_1 , it is clear that T_1 was unaffected by this cooling. It is possible that applying DDROP for a long time could physically cool the qubit environment; there was no evidence of this in these measurements.

The DDROP protocol for qubit reset has been experimentally demonstrated on a transmon in a 3D cavity to produce a fast, high-fidelity ground state preparation. This process satisfies the demand for qubit reset as part of an algorithm, and can be

used to improve the speed and fidelity of ground state preparation over that given by a return to equilibrium. The performance of the DDROP protocol has been evaluated by using the RPM method for quantifying the excited to ground state population ratio. The use of DDROP allowed experiments on this qubit to repeat at a rate 60 times faster than waiting $5 T_1$. Regardless of initial state, a ground state preparation fidelity of 99.5% was achieved in less than $3 \mu s$.

6.5 Discussion

As described earlier in this chapter, qubit reset is an integral and necessary element of any quantum algorithm. In addition, reset as a means of cooling may produce lower bath temperatures, although this has not yet been observed. Also, poor initialization fidelity limits the characterization of all other quantum gates, and thus initialization fidelity must be better than all other gate fidelities. DDROP performs all of these duties and the requirements and constraints of DDROP are fewer than other forms of reset; feedback, high fidelity readout and qubit tunability are each not necessary (this is discussed more in Appendix C). Therefore, DDROP is readily applicable and practically useful for most circuit QED systems.

The results and simulations from this experiment highlight an important next step in characterizing qubit reset. Simulations predict higher fidelities are possible; for example, simply reducing P_e from 9% to 1% and using $\bar{n} = 25$, simulations predict a fidelity of 99.99%. This parameter improvement has been achieved with other qubits, so in principle this experiment is can be completed relatively quickly.

In conclusion, the 3D transmon has proven itself as a robust qubit design with replicable fabrication, and as a result of pushing participation ratios to their limits, continued the trend of improved relaxation and coherence times of superconducting

qubits. The best compact transmons improved relaxation times by a factor of 3-4 over the CPW transmon and the best 3D transmon improved relaxation times even further by another factor of 10! Coherence times have also improved by a factor of 30 over compact transmons, yielding both relaxation and coherence times over 100 μ s. Coupling these minimized surface participation ratios with fluxonium's insensitivity to dissipation yield even further increases in relaxation times, as discussed in the next chapter.

Fluxonium

The fluxonium artificial atom, counterpart to the transmon, presents two advantages: possibility of probing different relaxation mechanisms over a wide frequency range and a natural insensitivity to dielectric dissipation. With typical parameters, the qubit transition of the fluxonium is tunable between 8 GHz and 500 MHz by applying an external magnetic field. This tuning, however, does more than change frequency. The energy potential landscape, a sine wave superimposed on a parabola, produces wildly different ground and excited state wavefunctions as the flux is varied. At some flux values, the ground state and excited state wavefunctions are non-overlapping in the φ representation, thus protecting the qubit from transitions induced by a varying electric field across the junction. For this reason, the fluxonium is expected to exhibit both shorter and longer relaxation times than the transmon depending on the loss mechanism and applied flux. This makes the fluxonium artificial atom a candidate of choice for exploring loss mechanisms and for maximizing relaxation time.

7.1 Design

7.1.1 Superinductance

The experiment described in this chapter is the measurement of a radical re-design of the fluxonium qubit to introduce both a new array design and a new sample holder. The junction array design for the superinductance was born out of the need to minimize capacitance to ground (see Section 2.2.2). This is important to eliminate any possible influence from self-resonant array modes. More precisely, the influence of the ground capacitance C_g limits the number of allowed junctions to $N_A < \sqrt{C_{JA}/C_g}$ [33], where C_{JA} is the junction capacitance of each array junction. The ground capacitance was found through simulations to be lowest for long skinny junctions with minimal metal between junctions. Fig. 7.1 shows the differences between the previous array design and the new array design. The new array design, utilizing the bridge-free fabrication technique [66, 67], has longer, narrower junctions and drastically less metal between junctions, with only a 100 nm wire between junctions. Lowering the ground capacitance as a means of increasing the number of junctions has the added benefit of allowing larger E_J/E_C ratio for array junctions while still maintaining a large total inductance. This reduction in array junction inductance prevents phase slips, which can limit coherence times [35].

The new superinductance design has been studied independently to verify its parameters. Compact resonators were made using superinductances and capacitors [113], coupled to a CPW feedline in the same way as other compact resonators. Internal quality factors were measured to be at least as high as 10^5 , comparable to non-superinductance compact resonators. The self-resonant modes were also measured using two-tone spectroscopy; the lowest observed mode was above 12 GHz,




Figure 7.1: Comparison between (a) first generation and (b) second generation junction array design, shown on the same length scale. The updated array design has longer, narrower junctions with lower ground capacitance. The previous array design was fabricated with the Dolan bridge technique, whereas the second generation design uses the bridge-free technique. Red rectangles highlight the size of the junction overlap.

sufficiently high not to interfere with the fluxonium qubit. Also with these superinductance resonators, phase slip rates were directly observed to be less than 1 mHz, with phase slips occurring on the scale of hours. Previous fluxonium qubit coherence was found to be limited by phase slip rates of 100 kHz [35], eight orders of magnitude more often.

7.1.2 3D Cavity

Besides the improvements in superinductance design, the other major improvement is the adoption of the 3D cavity (from the 3D transmon experiments, Chapter 6) as a contact-less sample holder. All previous fluxonium qubits were coupled to on-chip stripline resonators, thus necessitating the need for GE varnish, wire-bonds, and generally ignoring the environment of the sample holder, the same as with CPW and compact transmons as well. The fluxonium qubit is placed inside the 3D cavity, cou-

pling with antennae in the same way as the 3D transmon. Thus no glue, wire-bonds or other lossy materials are necessary. Without knowing exactly which component of 2D cavities and sample boxes are the problem, the use of 3D cavities eliminates many of them and has been shown to lead to higher relaxation and coherence times in 3D transmons.

However, the advantages of the 3D cavity are less clear for a fluxonium qubit than a transmon. While a quantitative analysis depends on qubit parameters, a large fraction of the capacitance of a transmon qubit comes from the self-capacitance of the antenna pads and the capacitance between the pads and the walls of the cavity. This means that a non-negligible amount of electric field energy is stored in the bulk dielectric and vacuum of the 3D cavity. For the fluxonium qubit, nearly 100% of the capacitance (that going into E_C) comes from the qubit junction capacitance. The small remaining amount comes from the small section of wire between the qubit junction and the first array junctions on either side, but with the coupling method described next, the antennae participate negligibly in the qubit mode. For this reason, the 3D cavity can be thought of as simply a contact-less sample holder for a 2D fluxonium qubit chip inside, whereas the transmon qubit itself is more “3D”.

7.1.3 Coupling

Coupling to the fluxonium with a 2D stripline resonator was achieved both capacitively with interdigitated capacitors to a voltage anti-node of the resonator and inductively with shared inductance at a current anti-node of the resonator [36]. For the 3D cavity, coupling capacitively while preventing E_C from decreasing out of the fluxonium range is difficult. For this reason, the coupling was chosen to be inductively via shared junctions in the superinductance. Antennae, attached across a few superinductance junctions couples the electric field inside the cavity to a current in

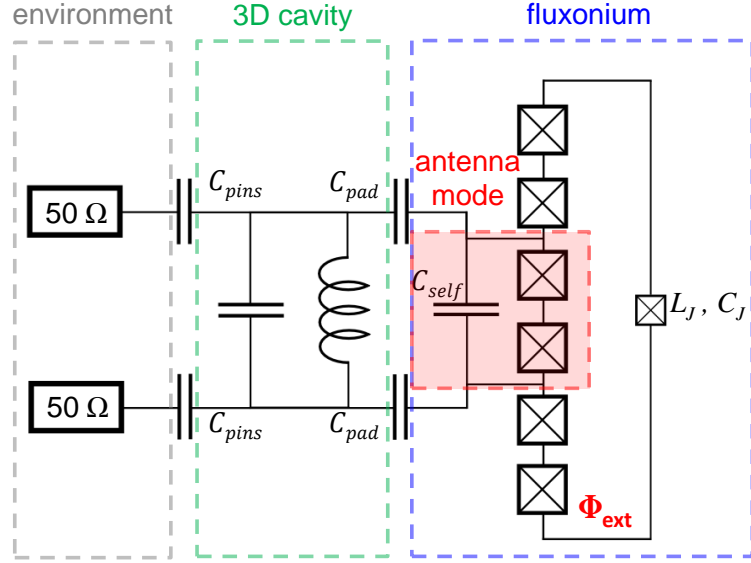


Figure 7.2: Approximate circuit diagram for fluxonium coupled inductively to 3D cavity. The fluxonium shares part of its superinductance with the antenna mode, which is capacitively coupled to the 3D cavity mode. The 3D cavity mode is itself capacitively coupled to the input and output 50Ω coax microwave lines. Qubit signals from the outside must pass through the 3D cavity and antenna mode “filters” to excite the qubit.

the superinductance, i.e. the fluxonium qubit mode. This type of coupling, drawn schematically in Fig. 7.2 introduces a new resonator mode: the antenna mode.

This mode, composed of the shared portion of the superinductance along with the capacitance from the antenna, is a separate mode from the fundamental mode of the 3D cavity. While the modes are coupled, excitations at the 3D cavity mode frequency will excite the antenna mode only weakly, and vice versa. Hence, to obtain strong coupling between the microwave line and the qubit mode (to allow for short qubit pulses), these two resonator modes must be relatively close in frequency. To make this aligning easier, the shared coupling junctions were replaced with coupling SQUIDs, thus allowing their inductance to be tuned with an external magnetic field. Since the fluxonium qubit is already tuned with a magnetic field, the SQUID loop area was made approximately 80 times smaller, allowing for full fluxonium tuning




Figure 7.3: Spectroscopy of 3D cavity mode as magnetic field is tuned. At low fields (low current), the antenna mode is at a higher frequency. At around 58 mA, the antenna mode passes through the cavity, and their avoided crossing is visible.

while making only small changes to the antenna mode frequency. This antenna mode tuning worked as planned and can be observed in Fig. 7.3.

An image showing the design of the fluxonium qubit coupling and superinductance is shown in Fig. 7.4. This image shows how the superinductance shunts the phase-slip junction, how the coupling SQUID junctions couple the antenna to the fluxonium, and how the antenna couples to the 3D cavity.

7.2 Experimental Apparatus

The experimental setup for the measurement of fluxonium qubits was equivalent to the setup for 3D transmons described in Section 6.2 and Fig. 6.4. All of the filters and infrared shielding was maintained between the two experiments. Since these changes were observed to have an effect on transmon coherence times, it is possible




Figure 7.4: Images showing the fluxonium qubit and its coupling to the antenna and 3D cavity modes. The top picture in an SEM image of the junction array shunting the phase slip junction. The coupling SQUID junctions join the antenna mode resonator with the fluxonium qubit loop. The antenna pads couple the antenna mode and fluxonium to the 3D cavity resonator, as shown in the bottom two images.

that the fluxonium may also be affected by these changes.

Due to the strong filtering of the 3D cavity mode at low frequencies, an additional room temperature amplifier is necessary to drive the fluxonium qubit at its lowest frequencies. For this purpose, the Mini-Circuits ZHL-20W-13+ amplifier, with a maximum output power of +43 dBm, was installed on the input line outside of the dilution refrigerator.

7.3 Results

The relaxation and coherence properties of two fluxonium qubits were measured. The qubits are nearly identical, with the only difference being in the coupling antennas. Fluxonium A had 3 coupling SQUIDs and a total antenna length of 1 mm, while fluxonium B had 4 coupling SQUIDs and a total antenna length of 2 mm. Most measurements described below come from fluxonium A, which was measured extensively and exhaustively. Fluxonium B was used to check the coherence repeatability, but was not measured in such detail.

Both qubits had visible spectroscopy peaks over the entire range of applied flux. This was an improvement over a few test samples with different antenna designs such that the qubit was only readable over a short range of frequencies. The improved coupler, described in the previous Section, while far from ideal, allows qubit readout over the entire range of frequencies. Qubit frequency is plotted as a function of applied flux in Fig. 7.5 for both fluxonium A and B. This frequency dependence is fit to that predicted from theory (Section 2.2.2) and shows good agreement. Fluxonium A shows some deviation from the fit, which so far cannot be explained. It has never been seen on any previous fluxonium samples, and is not repeated in fluxonium B, which shows better agreement. The fit parameters are the three energies that entirely




Figure 7.5: Measured qubit frequency as a function of applied flux over the entire tunable range. Fits of expected frequency dependence from theory match well with measured data and yield parameters as listed for each fluxonium sample.

characterize a fluxonium: E_J , the Josephson energy, E_C , the capacitive energy, and E_L , the inductive energy. From design and room temperature measurements, the parameters were expected to be roughly: $E_J = 12$ GHz, $E_C = 2.6$ GHz, and $E_L = 0.5$ GHz. Overall, the parameters are very close between fluxonium A and B, with only slight discrepancies between expected and fit parameters. The largest discrepancy is for E_C , which comes almost entirely from the junction capacitance. Here the estimate comes from the area of the junction and the commonly quoted $50 \text{ fF}/\mu\text{m}^2$, which is known to be an imprecise estimation method.

Fig. 7.6 shows spectroscopy traces for fluxonium A to highlight its features. Image (a) shows a plethora of spectroscopic lines that may make the fluxonium qubit look intimidating at first. While the qubit can be read out with either the 3D cavity mode or the antenna mode, here the 3D cavity mode is used; the two horizontal “blind spots” are the frequencies corresponding to these two modes. At

half-flux, there are many modes which enter this frequency range, and may be due to phonon modes of the array and/or multiple photon transitions of the fluxonium qubit. Regardless, the spectrum at zero flux is cleaner, as shown in Image (b). In previous fluxonium samples, these lines were each predicted and matched closely with theory [36]. Image (c) shows a close-up of the avoided crossing between f_{ge} and f_{ef} at zero flux. Notice that this frequency is a flux sweet spot, with frequency reaching a maximum, and thus first order insensitive to fluctuations in flux. The f_{ef} transition becomes invisible at exactly zero flux due to a symmetry in the Hamiltonian; this was also predicted and discussed with previous samples [36]. The last image, (d), shows the spectroscopy peak of f_{ge} at half flux, where the frequency reaches a minimum, another flux sweet spot. The spectroscopic line is visible over nearly the entire range, although χ does change signs a few times, leading to a few frequencies where $\chi = 0$.

A measurement of the dispersive shift χ (see Section 2.2.1) at an applied flux of $\Phi_{ext}/\Phi_0 = 0.03$, corresponding to $f_{ge} = 7.8$ GHz highlights one of the main limitations of this coupling design. While the qubit transition is coupled everywhere to the cavity, this coupling is small. At this frequency, Fig. 7.7 shows that the dispersive shift χ , the difference in cavity frequency when the qubit is in the excited state versus the ground state, is only 0.3 MHz. For the cavity linewidth of $\kappa = 5$ MHz, this means $\chi < \kappa$. The practical result of this small dispersive shift is that, especially without a parametric amplifier, many averages are needed to distinguish between the ground and excited states. This means that each measurement takes much longer than if this dispersive shift were higher. Although the detuning can be varied over a range of almost 8 GHz, the readout contrast (from χ/κ) is measured to range from 0 to 10 degrees, with the majority of measurements yielding close to 10 degrees. This highlights the fact that the coupling strength between the fluxonium and the cavity is related to the distance between the cavity and the f_{gf} transition, not only the f_{ge}

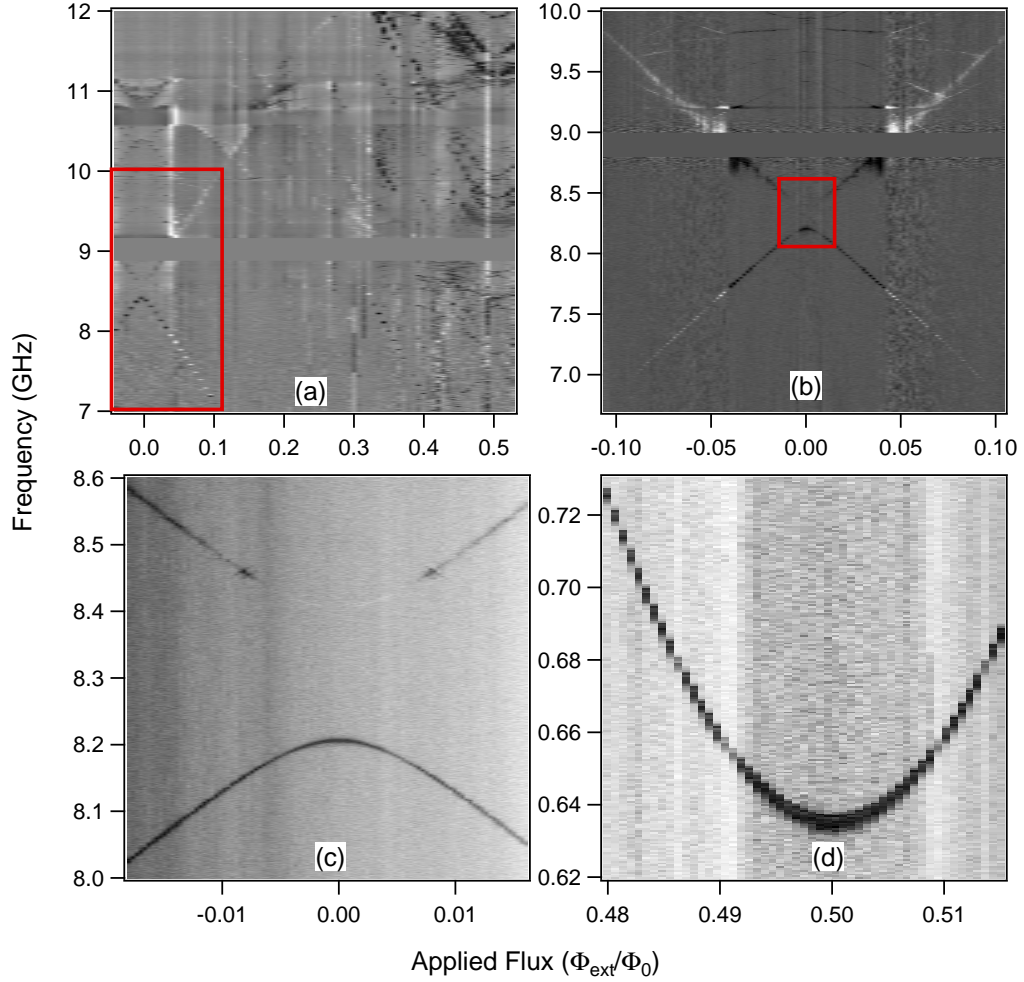


Figure 7.6: Spectroscopy shown over different regions of applied flux and frequency: (a) entire half-flux range, around cavity frequency, (b) closeup of qubit transitions near zero-flux, (c) further closeup showing avoided crossing of f_{ge} and f_{ef} at zero flux, and (d) f_{ge} around half-flux. Highlighted box in (a) shows region of (b) and highlighted box in (b) shows region in (c). In (a) and (b), the section around the cavity frequency is cut out to avoid confusion. In (a), many transitions at half-flux crowd the 10 GHz region.




Figure 7.7: Cavity transmission measurement as a function of frequency for prepared ground and excited qubit states, for $f_{ge} = 7.86 \text{ GHz}$. The cavity frequency shifts depending on the qubit state by $\chi = 0.3 \text{ MHz}$, a small fraction of the cavity linewidth $\kappa = 5 \text{ MHz}$.

transition.

Coherence measurements at the maximum frequency of 8.20 GHz are shown in Fig. 7.8. Since the qubit frequency has a maximum at zero flux and the cavity frequency (8.95 GHz) does not depend on flux, this is also the point of smallest detuning between the f_{ge} transition and the cavity; here $\Delta = 750 \text{ MHz}$. The T_1 is measured to be a respectable $43 \mu\text{s}$, longer than any previous fluxonium qubit, but lower than a typical 3D transmon. T_2^R is measured to be $15 \mu\text{s}$, definitely not limited by T_1 , indicating that there is some additional source of dephasing. Echo sequences are able to increase this dephasing time to $53 \mu\text{s}$. Since $T_2^E > T_2^R$ but still $T_2^E < 2T_1$, it is clear that while low frequencies dominate dephasing, there is also some high frequency noise inducing qubit dephasing.

Detuning away from the flux sweet spot at zero flux, but only to $\Phi_{ext}/\Phi_0 = 0.03$, f_{ge} decreases to 7.85 GHz, thus increasing the detuning to 1.1 GHz. For this small change in qubit transition frequency, there is a drastic change in relaxation and coherence parameters. Shown in Fig. 7.9, it is clear that the relaxation time has doubled to $T_1 = 95 \mu\text{s}$. This time is now comparable to the best 3D transmons at




Figure 7.8: Measured relaxation and coherence times at $\Phi_{ext}/\Phi_0 = 0.00$. (a) T_1 is measured by a decay from the excited state, yielding $T_1 = 43 \mu\text{s}$. (b) T_2^R is less than T_1 , while (c) T_2^E improves this time considerably.

this frequency [42]. This drastic change in relaxation time with such a small change in frequency is discussed in the next section. In stark contrast to the improvement in relaxation, the coherence time has plummeted. Without the protection of the flux sweet spot, T_2^R has decreased to 260 ns. At this flux, the qubit frequency depends strongly on the applied flux, with a slope of nearly $16 \text{ GHz}/\Phi_0$. The measurement of T_2^E yields a Gaussian shape instead of an exponential, but with a much longer time than T_2^R . All of these measurements say that the qubit coherence is limited by flux noise, either via microscopic local fluctuators or via noise on the external biasing coil. The Gaussian shape is a trademark of $1/f$ noise, which is the likely frequency dependence of flux noise. The exact coherence measurement values at this flux should not be considered too strongly. The pulses used in these sequences were each 160 ns long, a considerable fraction of T_2^R ; because of this, the echo sequence especially cannot be considered to be a true echo measurement since the echo cannot be approximated at “instantaneous”.

As the qubit is tuned towards lower frequencies, the relaxation time continues to increase. This trend continues until half flux, or $\Phi_{ext}/\Phi_0 = 0.50$, corresponding

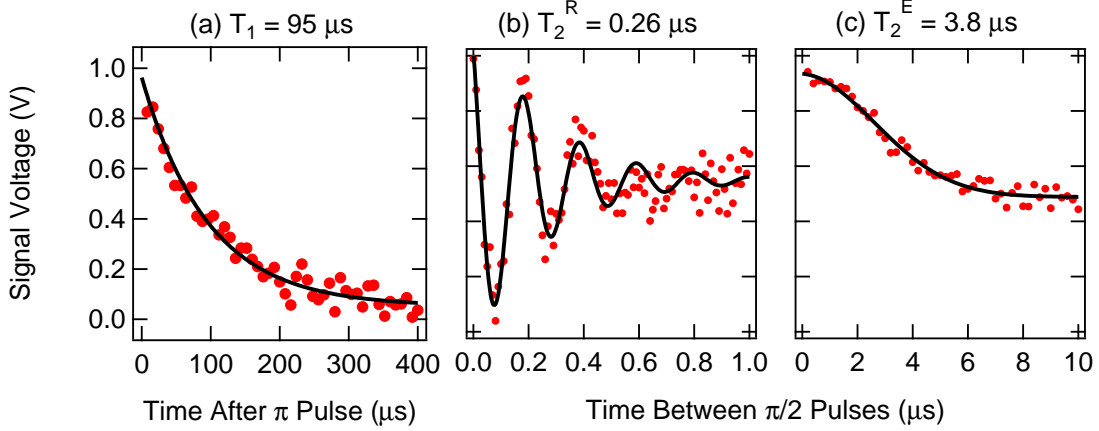


Figure 7.9: Measured relaxation and coherence times at $\Phi_{ext}/\Phi_0 = 0.03$. (a) T_1 is measured by a decay from the excited state, yielding $T_1 = 95 \mu\text{s}$. (b) T_2^R is much less than T_1 , indicating a significant dephasing source, (c) while T_2^E is a large improvement, it is still only about 4% of T_1 .

to the lowest qubit frequency. Around this flux, as shown in Fig. 7.10, T_1 of both fluxonium A and B is observed to be roughly 1 ms. The minimum frequencies for fluxonium A and B are 640 MHz and 480 MHz, respectively. At a frequency of 500 MHz, the mixing chamber temperature of 15 mK is comparable to the temperature $T = \hbar\omega_{ge}/k_B = 23 \text{ mK}$. This means that the qubit is nearly saturated in thermal equilibrium. Further, as seen by previous measurements (see Chapter 6), the qubit temperature can be significantly higher than this. Therefore, regardless of the exact qubit temperature, the effect of a π -pulse at this low frequency is to invert the population from being slightly more in the ground to being slightly more in the excited state. Interestingly, this still leads to a readout contrast of approximately 10 degrees, indicating that the dispersive shift may actually be much higher at this flux.

Repeating these relaxation measurements around half flux yields some interesting phenomena. While sometimes a single exponential decay is observed, as shown in Fig. 7.11, relaxation measurements *sometimes* exhibit a double exponential behavior.




Figure 7.10: Measured relaxation times *near* $\Phi_{ext}/\Phi_0 = 0.50$. Data is fit to a single exponential and reveals that for (a) fluxonium A at $f_{ge} = 640$ MHz and (b) fluxonium B at $f_{ge} = 750$ MHz, lifetimes are approximately 1 ms. The presence of single exponentials as shown here fluctuates in time, as shown by measurements at nearly the same flux shown in Fig. 7.11.

The two exponential times are very different, with one typically a factor of at least 10 greater than the other. While the fluxonium B measurement in this Fig. is at a slightly different qubit frequency as the single exponential measurement above, the fluxonium A measurement is at the exact same frequency as the above example. This shows that the presence of double exponentials varies both in flux and time.

These relaxation experiments measure qubit excitation following a π -pulse, and were measured simultaneously with Rabi oscillation measurements. The purpose of these Rabi oscillations is to calibrate the amplitude and offset of the exponential decays, which are then fixed parameters in the fits shown above. The single exponential measurements have only one free fitting parameter (T_1), while the double exponential measurements have three (T_1^a , T_1^b , and the amplitude fraction). This form of calibration is necessary to be sure the single exponential behavior is real, as it is easy to fit data with a slight double exponential to a single exponential if the amplitude and offset are allowed to vary. Also, when the short exponential is short enough, it is only represented as a point or two with the chosen linear spacing of




Figure 7.11: Measured relaxation times *near* $\Phi_{ext}/\Phi_0 = 0.50$ exhibiting double exponential behavior. Data is fit to a double exponential for (a) fluxonium A at $f_{ge} = 640$ MHz and (b) fluxonium B at $f_{ge} = 480$ MHz. The presence of double exponentials as shown here fluctuates in time, as shown by measurements at nearly the same flux shown in Fig. 7.10.

measurements points. Thus, without the calibration, these points could be mistaken for outliers and they would be neglected in the fit.

Since this measurement procedure requires a properly tuned π -pulse, combined with the long lifetimes, these measurements take a long time to tune and measure. With a repeat rate of 10 ms, averaging 2,000 times requires over 30 minutes. Many averages are required due to the small readout fidelity and lack of a parametric amplifier. Dispersive readout using approximately 1 photon and an averaging time of $30 \mu\text{s}$ (longer than some relaxation times!) was used here; the viability of the high-power readout has yet to be explored for the fluxonium qubit. Including the calibration of the π -pulse, a measurement at a single flux requires well over an hour.

An alternate relaxation measurement experiment, necessitated by these long times, is used to characterize the flux dependence of the relaxation time at low frequencies. First of all, switching from a π -pulse to a saturation pulse comes at a cost of half the readout amplitude, but yields many improvements. A saturation pulse is a moderate strength pulse of a length longer than T_2 that places the qubit

in an equal mixture of ground and excited state. Decay is measured from this state back to the thermal state, hence the readout contrast is decreased by half. Due to the simplicity of this pulse, no tuning or precise calibration is required, which cuts down on preparation time. Another improvement is that there is no need to wait for the qubit to decay fully to the ground state before repeating the experiment, as is necessary when π -pulses are used. A saturation pulse prepares the equal mixture regardless of initial state; this allows the experiment to proceed much faster, especially for the patterns with short wait times. Additionally, since an experiment with linear spacing of wait times (as shown above) requires pre-existing knowledge of the lifetime in order to choose the wait times, it is not possible to easily measure highly variable lifetimes with a single pulse sequence. A log-scale spacing of wait points, with many points at short wait times, and fewer points at longer wait times allows equally for any lifetime over a wide range to be measured without switching pulse sequences. Combined with the faster repeat time for short wait times, and the sparsely spaced points at longer wait times, this experiment runs at a speed around 5 times faster than the previous experiment. Taking into account the lack of pre-calibration, the ratio is far greater.

An example of this saturation log-scale relaxation measurement is shown in Fig. 7.12. In this Figure, the same data is displayed twice, once in the index basis, and the other in the time basis. A fit to this data assuming a single exponential yields a lifetime of 6.4 ms. A double exponential would look like a double step feature in the index-basis representation. The only observed anomaly for this type of measurement is the first point or two often seem to be outliers towards the ground state. The times corresponding to these points are too short to be the same short exponential from the double exponential (of order 100 μ s), and regardless they are outliers in the wrong direction for that to be the case. A hypothesis for these outliers is that




Figure 7.12: Measured relaxation time of fluxonium A at $\Phi_{ext}/\Phi_0 = 0.50$, corresponding to a qubit frequency of 640 MHz. The same data is expressed as a function of (a) time in log-scale (b) time in linear scale. Note that the measured points are equally spaced in (a).

the very fast repeat times used for these shortest wait times is interfering with the saturation pulse or readout pulse. This can be explored with further measurements. Ignoring these couple of points, the data fits well to a single exponential; a third representation (not shown here) shows that the data lies on a line in a log-log scale, further indicating that a single exponential is a good approximation. Interestingly, double exponentials were never observed with this type of measurement. Therefore, whatever causes the double exponential must be changed by the use of the saturation pulse or the faster readout times; this effect is not understood at this time.

While this result (6.4 ms) may seem much longer than the 2.3 ms from the long portion of the double exponential from Fig. 7.11(a), there are several reasons why this comparison is unwarranted. The first is that the 2.3 ms comes from a fit with a maximum measured wait time of 4 ms. This is not long enough to accurately measure such a long decay time; a factor of 3 between T_1 and longest measured wait time is a good minimum from experience. The double exponential plots shown have enforced limitations in an effort to show both exponentials, and do not accurately measure either the short or long lifetimes. Also, these measurements were not taken at the

same time, and lifetimes have been observed to fluctuate significantly over time. Third, sequential measurements were completed using both methods and found that the lifetimes always matched within 10%.

Nevertheless, the ease of these experiments allow them to be studied as a function of flux, and thus to complete the measured T_1 dependence for the region where T_1 becomes much longer than a few hundred microseconds. Many sweeps of T_1 versus flux were completed over the course of a month or so, all of them are included in Fig. 7.13. Data above roughly $\Phi_{ext}/\Phi_0 = 0.40$ is from the saturation/log method, whereas data at lower fluxes are from the traditional π -pulse/linear method. First, it is important to note that there is a large scatter at each flux value, with T_1 varying by almost a factor of 10 at each flux. The measurements were taken over the course of a month, but the scatter is evident even on the time scale of a few hours. Measuring T_1 over a small range usually reveals many suspected features, however these features are not reproducible. Plotting all measurements together blurs these features into an overall scatter; nevertheless there are some distinct features that stand out. There is a noticeable minimum in T_1 at the highest frequencies, with a drop from around 200 μ s to well below 100 μ s. There is also a very steady rising trend that is almost linear as a function of applied flux, over almost the entire range. Lastly, and perhaps most importantly, there is a noticeable *peak* in T_1 centered on $\Phi_{ext}/\Phi_0 = 0.50$. Although the scatter is evident, this peak is clearly defined and represents an increase in T_1 by almost an order of magnitude over a change in flux from applied flux from 0.47 to 0.50 Φ_0 . Relaxation around this peak was measured many times over the course of a month, which highlights its stability and repeatability.

Since T_1 has a maximum at half flux, it is important to see how this affects the coherence time. Shown in Fig. 7.14 are coherence measurements for fluxonium B at half flux. Results from fluxonium A are nearly identical. The minimum frequency




Figure 7.13: Measured T_1 as a function of applied flux for fluxonium A, which tunes the qubit frequency roughly from 8.2 GHz at $0.0 \Phi_0$ to 640 MHz at $0.5 \Phi_0$. T_1 has a minimum below $100 \mu\text{s}$ at the maximum frequency and a maximum of nearly 10 ms at the lowest frequency.

is a flux sweet spot in the same manner as the maximum frequency. Therefore, it is not surprising that the coherence time is longer than anywhere at $\Phi_{\text{ext}}/\Phi_0 = 0.03$. The coherence times, in fact, are very similar to those at zero flux.

7.4 Analysis

From the frequency dependence of T_1 , it is possible to deduce which of several loss mechanisms is dominating since each loss mechanism has a different frequency dependence. The method for calculating the expected relaxation rate for a given loss mechanism is described in Section 2.2.2. The loss mechanisms considered here are capacitive loss (both junction and environmental capacitance lumped together), inductive loss (loss in the superinductance), quasiparticle loss (in the qubit junction), and Purcell (loss due to coupling with the cavity). It is important to note that inductive loss may also be determined by quasiparticles in the inductor array, but here the term “quasiparticle loss” will mean loss due to quasiparticles tunneling across




Figure 7.14: Measured coherence times of fluxonium B at $\Phi_{ext}/\Phi_0 = 0.50$. T_1 at this flux is shown in Figures 7.10 and 7.11. (a) T_2^R is much less than T_1 , indicating a significant dephasing source, (c) T_2^E is a small improvement, indicating a high frequency noise source.

the qubit junction. Each of these loss mechanisms predicts a unique dependence on frequency.

Capacitive loss, characterized by a constant quality factor Q_{cap} , predicts only a small frequency dependence. As shown in Fig. 7.15, this flat dependence does not well match the data. For a given curve, any measured data above the curve is inconsistent, and thus that curve is invalid. This is because additional sources of loss can only reduce T_1 since the losses add like resistors in parallel. For example, the curve corresponding to $Q_{cap} = 1 \text{ M}$ matches most of the data quite well, except around half flux where the T_1 increases far above the curve. This means that Q_{cap} is greater than 1 M. From the Figure, to be consistent with the data, a lower bound of a constant Q_{cap} is 15,000,000.

In an effort to better model the loss observed by the fluxonium, Fig. 7.16 uses a Q_{cap} that depends on frequency as $(f_x/f_{ge})^{0.7}$. This dependence was measured at low temperature for the sapphire dielectric loss tangent [114]. This was chosen since both junction and environmental capacitance use aluminum oxide as dielectric. The Q_{cap} values given in the Fig. are for the frequency $f_x = 5 \text{ GHz}$. These curves capture




Figure 7.15: Measured T_1 data superimposed upon predicted relaxation times for capacitive loss, assuming a constant Q_{cap} .

the overall trend of increasing T_1 at lower frequencies very well. The slope of the curves matches the slope of the data. To remain consistent with the data, Q_{cap} at 5 GHz must be at least 3,000,000.

In contrast to the capacitive loss, inductive loss characterized by a constant Q_{ind} exhibits the opposite dependence on frequency. As shown in Fig. 7.17, the curves have a minimum at half flux, where the frequency is lowest. The only curve that is *consistent* with the data is that of $Q_{ind} = 500,000,000$. This curve shows that if inductive loss dominates the dissipation at half flux, it is basically irrelevant everywhere else.

A calculation of the Purcell effect is basically the same as any other loss mechanism, as detailed in Section 2.2.2, except that the term $\text{Re}[Y_{purcell}]$ is simulated using HFSS. This simulation results in the estimate of Purcell contribution to T_1 as shown in Fig. 7.18. From the limited portion of the data that is close to this Purcell limit, it is not possible to confirm whether this is an accurate estimate or not. Regardless, if accurate, then Purcell only dominates relaxation over a narrow range of applied flux,




Figure 7.16: Measured T_1 data superimposed upon predicted relaxation times for capacitive loss, assuming a $Q_{cap} \propto (f_x/f_{ge})^{0.7}$.




Figure 7.17: Measured T_1 data superimposed upon predicted relaxation times for inductive loss, assuming a constant Q_{ind} .




Figure 7.18: Measured T_1 data superimposed upon predicted relaxation times loss due to the Purcell effect. The Purcell effect is due to coupling with the 3D cavity mode at 8.95 GHz with Q of approximately 1,500.

where the qubit frequency is less than 1 GHz detuned from the cavity. The Purcell effect may be responsible for the doubling in T_1 observed between $\Phi_{ext}/\Phi_0 = 0.00$ and $\Phi_{ext}/\Phi_0 = 0.03$.

Lastly quasiparticle loss has a dramatically different predicted T_1 as a function of applied flux. Fig. 7.19 highlights its uniqueness well by showing the strong peak observed at half flux. Note that in addition to T_1 from quasiparticle loss, a constant loss (equivalent to a T_1 of 8 ms) was added to represent a limitation from some other loss mechanism. In truth, the quasiparticle contribution to T_1 goes to infinity as the dissipation is reduced to zero because of symmetry. This mechanism is explained in Section 2.2.2.

A close-up around half-flux of the predicted T_1 from quasiparticle loss, characterized by curves of constant broken Cooper pair fractions, x_{qp} , is shown in Fig. 7.20. The parameter x_{qp} is defined as the ratio of quasiparticles to Cooper pairs, assuming a constant density. This figure shows that these curves really capture the




Figure 7.19: Measured T_1 data superimposed upon predicted relaxation times for quasiparticle loss, assuming a constant fraction of broken Cooper pairs, x_{qp} .

trend measured by the data. Neither capacitive, inductive, or Purcell predicted any feature similar to this peak, yet quasiparticle loss matches very well. Although there is a large scatter in the data, this is also captured well by the curves. If the value of x_{qp} varies between 1×10^{-7} and 2.5×10^{-6} , the reduction in scatter at half flux predicted by the most extreme curves is well matched to the range of scatter in the data. Note that scatter reduces as the flux is tuned exactly at $0.5 \Phi_0$ as other loss mechanisms become important.

In order to have a possible explanation for loss everywhere, Fig. 7.21 shows one scenario that combines quasiparticle loss, capacitive loss, and loss due to the Purcell effect to match roughly with the upper bound of the scatter on the data. The black line in this Fig. matches the dip at zero flux, the peak at half flux, and the slow upward trend between. The capacitive loss contributes both to the overall upward trend at low frequencies, and as the limiting factor at exactly half flux, where quasiparticle loss disappears.

The exact participation ratios of the junction capacitance and capacitance from




Figure 7.20: Measured T_1 data superimposed upon predicted relaxation times for quasiparticle loss, assuming a constant fraction of broken Cooper pairs, x_{qp} ; close-up around half-flux




Figure 7.21: Measured T_1 data superimposed upon predicted relaxation times for capacitive loss for $Q_{cap} \propto (f_0/f_{ge})^{0.7}$, with a value of 3,000,000 at $f_x = 5$ GHz, the loss due to the Purcell effect, and the loss due to quasiparticles with constant $x_{qp} = 2 \times 10^{-7}$. These contributions are added to yield the effective relaxation time, which matches roughly with the upper end of the data scatter for the entire range of applied flux.

the leads is unknown, simulations show that for the chosen configuration of junctions and wires, the wires should add less than 1% to the capacitance. Therefore a value of $Q_{cap} \geq 3,000,000$ may say that either the junction dielectric has $Q \geq 3,000,000$ or that the bulk sapphire dielectric has $Q \geq 30,000$. Further experiments would be necessary to distinguish the two capacitances.

Since relaxation rates from different mechanisms add, it is easy to achieve a dip in T_1 by having one contribution dominate, a *peak* in T_1 requires contributions from *all* other mechanisms to be considerably lower than, in this case, quasiparticle loss. Thus, near half flux, quasiparticle loss must dominate the loss, with all other contributions being small. Then, approaching half flux, the quasiparticle loss relaxation rate decreases and the T_1 rises. Then at exactly half flux, quasiparticle loss is eliminated, and the second strongest contribution now limits relaxation.

The observation of this *peak* in T_1 at half flux is evidence for the observation of coherent quasiparticle interference. This effect, described in more detail in Section 2.2.2 was predicted as the fourth Josephson term in Josephson's original paper [25] and described more recently in this context by G. Catelani et al. [62]. While many experiments have attempted to measure this term [115, 116, 117, 118, 119, 120], the available measurements do not agree and are not conclusive [121]; the present experiment represents a verification of a long-standing prediction, 50 years after its description.

The presence of double exponentials may also be due to quasiparticles. Since the measurement of relaxation time takes many averages and many minutes to acquire these averages, it is possible that the number of quasiparticles changes significantly, perhaps many times. If some measurements are during a particular quasiparticle distribution and other measurements during a different quasiparticle distribution, then one would expect a double exponential since the results are simply averaged

together. This hypothesis is that the qubit is like “Dr. Jekyll and Mr. Hyde”, where the same qubit is sometimes bad (low T_1) and sometimes good (high T_1). A possible scenario is that during the high T_1 (Dr. Jekyll) times, the qubit is limited by a quasiparticle density of roughly 2×10^{-7} , with quasiparticles in the array being low enough such that $Q_{ind} > 500,000,000$ and the low T_1 (Mr. Hyde) times corresponding to a sudden influx of quasiparticles in the array leading to $Q_{ind} \sim 10,000,000$. In this latter case, the loss of the array would then dominate the loss, and the T_1 would be drastically lower. The quasiparticle density to explain this increased loss is roughly one quasiparticle per array junction into the array, perhaps generated via bombardment from a radiative decay, a cosmic ray, or any other high energy particle. In contrast, the actual number of quasiparticles in the high T_1 case, for $x_{qp} \sim 10^{-7}$ is of order one, which may explain the large scatter observed as the fluctuations are of the same order.

On another note, from decoherence measurements, it is clear from the lower T_2 on the slope of the spectrum as opposed to at the sweet spots, that flux noise is the limiting factor in dephasing along the slope. At the maximum frequency, photon induced dephasing is a real possibility. At zero flux, with a cavity linewidth $\kappa = 5$ MHz, and a measured dispersive shift $\chi = 0.3$ MHz (measured slightly away from zero flux, so this is an approximation), the photon induced dephasing rate can be calculated via Eq. 1.11 [43, 42]. The dephasing expected for a thermal average number photons $n_{th} = 0.5$ in the cavity is $20 \mu\text{s}$, in the range of what is observed. For even a slightly higher χ of 0.4 MHz, this drops to $12 \mu\text{s}$. While 0.5 photons seems to be high for a copper cavity, it is not impossible. While the cavity photon population was not measured directly (in fact it is very difficult to measure when $\chi < \kappa$), this effect could be tested by either reducing the coupling, decreasing κ or improving shielding.

7.5 Discussion

Relaxation times from saturation measurements of longer than 1 ms make the fluxonium the longest lived superconducting qubit to our knowledge at the time of this writing. The previous record was held by a Cooper pair box from the group at the Laboratory for Physical Sciences (LPS) [122] with a relaxation time of approximately 200 μs . The longest transmon relaxation and coherence times are between 100 μs and 200 μs , with similar results from Yale and IBM. The improvements of fluxonium come from the combination of the reduced surface participation and the expected insensitivity to loss sources. Further experiments to improve our understanding of the fluxonium qubit that may lead to improved relaxation and coherence times, including: waveguide coupling, reducing flux noise, reducing possible cavity occupation, and investigating quasiparticle dynamics.

The influence of flux noise on fluxonium coherence can be examined further by measuring a redesigned fluxonium qubit that is engineered to be insensitive to flux noise: the gradiometric fluxonium. This design, built from two superconducting loops instead of one, would allow for the fluxonium qubit to be locked to either the maximum or minimum frequency, flattening the frequency dependence on applied flux to remove that source of dephasing. In addition, it is possible to add a superconducting shield to improve the magnetic field shielding. This experiment will reveal whether the fluxonium was limited by flux noise and whether that flux noise is environmental or microscopic on the surface of the sample.

A further experiment to check the validity of the quasiparticle dynamics theory (Dr. Jekyll and Mr. Hyde) for the double exponential is to measure the two exponential times as a function of flux. A possible flux dependence is shown in Fig. 7.22, where the flux dependence of the Dr. Jekyll and Mr. Hyde cases are shown.




Figure 7.22: Predicted flux dependence of T_1 for quasiparticle loss with density $x_{qp} = 2 \times 10^{-7}$ (Dr. Jekyll) and for inductive loss with $Q_{ind} = 10$ M (Mr. Hyde). These two cases may represent the dominating factors in the observed double exponentials. The arrow highlights that at a given flux, quasiparticle dynamics may cause the relaxation time to jump between the two curves, leading to a short and long T_1 with approximately the values observed.

Measuring this dependence will confirm or deny this hypothesis. One way to measure this is to use a parametric amplifier to enable higher fidelity readout, and thus measurements will require fewer averages. Then, the calibration and measurement of double exponentials as a function of flux is no longer impractical. In addition, if the switching time between the two cases is longer than a few seconds, perhaps T_1 can be measured in a time where no switches occur, thus leaving either simply a short or long T_1 . Another method is to use the same parametric amplifier to observe quantum jumps and infer T_1 from the jump statistics. One should observe periods of fast jumps and periods of slow jumps corresponding to the two different relaxation rates.

The 3D cavity resonator may be replaced with a section of waveguide in order to side-step the problem of the different cavity and antenna mode frequencies and

may improve qubit-cavity coupling. The weak coupling in the current design was a limitation for both the readout and the practicality of short pulses. The filtering aspect of the 3D cavity may intentionally be minimized by engineering a very wide bandwidth. A “cavity” with the widest bandwidth possible is that whose discontinuities are replaced with matched connections, thus reducing the cavity to a section of waveguide. Since only one mode is required for qubit readout, the antenna mode remains sufficient for that purpose. This waveguide coupling method is similar to the “hanger” coupling method of compact resonators to a CPW feedline, and may produce more strongly coupled fluxonium qubits.

Ground state preparation is another procedure that will help increase the readout contrast and will allow for faster repeat times, which are now necessary in light of the increased relaxation times. Sideband cooling as a means of preparing the ground state has been achieved with previous fluxonium qubits, and there is no reason it would not be possible in the new fluxonium design [36]. As an alternative, the DDROP sequence (See Chapter 6 and Section 6.1.4) may be implemented for the fluxonium and used as a form of reset.

The possible effect of cavity population on fluxonium dephasing times may be studied by adjusting the experimental apparatus to produce lower cavity populations. The cavity population is impossible to measure with existing coupling strengths, yet it can be measured through the use of stronger coupled fluxonium samples. With or without measuring the cavity population, improved filtering or infrared shielding may reduce the cavity population, and any change in dephasing rates can be measured.

Another way to explore the impact of quasiparticle loss on fluxonium relaxation and decoherence is to attempt to remove quasiparticles from the system. This could possibly be achieved through the addition of normal metal quasiparticle traps. Depending on where these traps are placed, they may be used to remove quasiparticles

from the superconducting islands of the qubit and prevent them from causing dissipation.

In conclusion, the fluxonium qubit has shown remarkable improvements in relaxation time over even the 3D transmon. These improvements are due to matching the minimized surface participation ratios of the 3D transmon with the insensitivity to dissipation of the fluxonium. While these initial fluxonium relaxation measurements of longer than 1 ms represent a drastic improvement in superconducting qubit coherence, the above experiments detail a path for further investigation into the limits of the fluxonium qubit.

Conclusions and Perspectives

In conclusion, it is clear that the coherence of superconducting qubits have improved drastically over the last five years through the application of the participation ratio analysis technique. With validation provided by observed improvements in the quality factor of compact resonators from 160,00 to 380,000 and improvements in 2D transmon relaxation times from a maximum of 4 μs to a maximum of 13 μs , the idea of participation ratios was taken to the extreme with a 3D architecture. The introduction of a 3D cavity for transmon qubits improved coherence times to around 100 μs . Finally, the simple 3D cavity design was combined with the insensitivity to dissipation of the fluxonium qubit to produce remarkable relaxation times of longer than 1 ms. Altogether, these improvements have helped to remove the label of “limiting factor” from superconducting qubit coherence. From here, there are many experiments that can be done that may help continue these trends.

The next step for compact resonators has already been taken, in the combination of geometrical improvements and better materials. Results from UCSB have achieved quality factors of over 1×10^6 with aluminum on sapphire CPW resonators [97]. While these are not compact resonators, they were designed with wide centerpins and wide gaps. Since there is still an observed increase of quality factor at high

power, it is likely that the quality factor is still controlled in part by TLS. Thus, further improvements from geometry should be possible.

The compact transmon measurements in this paper have already inspired others to go a step farther. Both UCSB and IBM have since measured transmon qubits with wide gap and width features combined with more optimized materials for low loss. These experiments have resulted in transmons with typical T_1 of up to $30 \mu\text{s}$ (with some measurements up to $44 \mu\text{s}$) from UCSB and up to $60 \mu\text{s}$ from IBM, both of which attribute geometrical improvements as a key to the higher relaxation times.

3D transmon coherence is currently being heavily explored by several groups including Yale and IBM. These experiments are focusing on new materials and fabrication techniques. The reset experiment detailed in this thesis can be extended quite easily in a proposed experiment. The 3D transmon used for the reset experiment had moderate parameters, with $T_1 = 40 \mu\text{s}$ and a P_e of 9%. Using transmons that exist now, the DDROP ground state preparation fidelity is expected to increase from the current 99.5% to 99.99% simply by reducing P_e to 1% and increasing \bar{n} to 25 photons. Any increase in T_1 will improve fidelity further. Since these parameters have been achieved in existing qubits, this experiment could be completed relatively quickly.

The fluxonium results shown here barely scratch the surface of what is possible with this artificial atom. Keep in mind that the measurements in this thesis are simply from the first two redesigned fluxonium samples, while surely over 100 transmon samples have been measured in total from all research groups. Our understanding of various loss mechanisms including quasiparticle tunneling can be extended by measuring more qubits, and specifically conducting several experiments. First, with the existing samples, the double exponential effect can be explored in detail with the addition of a low noise amplifier that allows measurement of quantum jumps. The

exact statistics of the quantum jumps should give as much information as possible concerning this effect. In addition, the amplifier can be used to prepare the ground state better than thermal equilibrium at half-flux, thus allowing experiments to take advantage of the full signal contrast. Also, the effect of quasiparticles on relaxation time can be measured by injecting quasiparticles into the system with high power microwave pulses.

Further experiments include the measurement of a more strongly coupled fluxonium, possibly inside a waveguide section to avoid the problem of two non-aligned filters. Also, simply measuring more fluxonium samples will reveal more statistics about the variation in loss mechanisms. The influence of flux noise on fluxonium coherence can be examined by measuring a redesigned gradiometric fluxonium qubit that is engineered to be insensitive to flux noise, thus revealing more information about the flux noise dephasing limitation.

Resonator “Hanger” Equation

A common goal of resonator research is to measure the internal quality factor (Q_i) as a means of quantifying certain loss mechanisms. Coupling to the resonator always introduces another quality factor, the coupling quality factor Q_c . Thus measuring the total quality factor Q_0 (defined by $1/Q_0 = 1/Q_i + 1/Q_c$) only yields information about Q_i in the limit that $Q_c \gg Q_i$. Depending on the coupling mechanism, relying on being in this limit usually forces weak measurement and low SNR.

Another way to measure the Q_i directly is to use the “hanger” coupling method, which is accomplished by coupling the resonator between an unbroken microwave transmission line (feedline) and ground. This type of coupling enables one to actually measure both Q_c and Q_i from a single measurement. While the width of the observed Lorentzian lineshape from a transmission resonator measurement reveals Q_0 , the reduction from full transmission on resonance is required to extract Q_i , so one must separately calibrate the full transmission amplitude. For hanger measurements, the Lorentzian is inverted, and thus the deviation from zero transmission is required to extract Q_i . Clearly there is no calibration required for zero transmission, and therein lies the beauty of the hanger measurement scheme. Additionally, in the hanger measurement scheme, the full transmission amplitude is also obtained for




Figure A.1: Expected amplitude lineshape of resonance measured in transmission (red) or as a hanger (blue). In either case, the lineshape is a Lorentzian, with the hanger measurement showing the inverse of the transmission measurement. Transmission measurements require calibration to full amplitude, while hanger-coupling gives this information from the amplitude off-resonance.

free, by measuring the transmission away from resonance. Plots of an ideal (infinite Q_i) resonance measured in transmission and in the hanger-coupled scheme is shown in Fig. A.1.

A.1 Derivation

A derivation of the expected resonator lineshape is now in order. This derivation will make only minimal assumptions. The first assumption is that measurement is of transmission through a microwave feedline to which the resonator is somehow coupled. This implies that we are considering a 2-port device. The most generic 2-port device is shown in Fig. A.2, or alternatively as a lumped element circuit model in Fig. A.3. For now, the three impedances can be anything. The resonator is considered to be lossless at first.

The second assumption is that there is only one resonator, implying that at most, each of the impedances in Fig. A.2 can have one pole. Thus the form of the impedances can be written as shown in Eq. A.1. In these equations, each coefficient




Figure A.2: Schematic of the most generic 2-port device. Each impedance Z_1 , Z_2 , and Z_3 can have any value and any frequency dependence.




Figure A.3: Circuit model for derivation of “hanger” resonator response calculation, where the transmission lines of characteristic impedance Z_0 are replaced by resistors of the same impedance.

(a_i, b_i, \dots) can be any complex number. This yields a total of 24 real variables.

$$z_i = \frac{a_i + b_i i \omega}{c_i + d_i i \omega} \quad (\text{A.1})$$

Since measurements typically involve measuring the scattering matrix (S-matrix) and not the impedances directly, it is convenient to calculate the S-matrix from this impedance structure. First calculating the transmission (ABCD) matrix and then the S-matrix yields Eqs. A.2 and A.3, where Z_0 is the characteristic impedance of the feedline (see chapter 4 of Ref. [53]).

$$\begin{aligned}
A &= 1 + \frac{z_1}{z_2} \\
B &= z_1 + z_3 + \frac{z_1 z_3}{z_2} \\
C &= \frac{1}{z_2} \\
D &= 1 + \frac{z_3}{z_2}
\end{aligned} \tag{A.2}$$

$$\begin{aligned}
S_{11} &= \frac{A + \frac{B}{Z_0} - CZ_0 - D}{A + \frac{B}{Z_0} + CZ_0 + D} \\
S_{12} &= \frac{2}{A + \frac{B}{Z_0} + CZ_0 + D} \\
S_{21} &= \frac{2}{A + \frac{B}{Z_0} + CZ_0 + D} \\
S_{22} &= \frac{-A + \frac{B}{Z_0} - CZ_0 + D}{A + \frac{B}{Z_0} + CZ_0 + D}
\end{aligned} \tag{A.3}$$

The current assumptions allow for elimination of some variables. Adhering to the restriction that there is only one pole in frequency, that $S_{12} = S_{21}$, and noticing that the denominators of each S-matrix element are the same, one can simplify Eqs. A.3 to form Eqs. A.4, where each of the coefficients here (a, b, ...) are free to be any complex number.

$$\begin{aligned}
S_{11} &= \frac{a + bi\omega}{c + di\omega} \\
S_{12} &= \frac{g + hi\omega}{c + di\omega} \\
S_{21} &= \frac{g + hi\omega}{c + di\omega} \\
S_{22} &= \frac{e + fi\omega}{c + di\omega}
\end{aligned} \tag{A.4}$$

The only measurement needed to characterize the resonator is S_{21} , on which we can apply one more assumption. The third assumption is that when the resonator is lossless, there is zero transmission on resonance. Transmission is exactly zero for hanger-coupled resonators that have no internal loss (infinite Q_i , the current assumption), since all power from the feedline is shunted into the resonator. Defining the resonant frequency to be ω_0 , this forces S_{21} into the form shown in Eq. A.5. This

expression highlights the fact that through the above assumptions, S_{21} has only one pole and one zero.

$$S_{21} = \frac{hi(\omega - \omega_0)}{c + di\omega} \quad (\text{A.5})$$

As an aside, an alternate derivation, which calculates the transmission directly using the circuit model of Fig. A.3, via $S_{21} = 2V_{out}/V_{in}$, results in an equivalent expression. Ohm's law and the equations for current division can be used to express V_{out} in terms of V_{in} , with the result shown in Eq. A.6, where $\tilde{z}_1 = z_1 + Z_0$ and $\tilde{z}_3 = z_3 + Z_0$.

$$S_{21} = 2 \frac{z_2 \frac{Z_0}{\tilde{z}_1 + \tilde{z}_3}}{z_2 + \frac{\tilde{z}_1 \tilde{z}_3}{\tilde{z}_1 + \tilde{z}_3}} \quad (\text{A.6})$$

Continuing the derivation, additional simplification of Eq. A.5 is achieved by declaring that the overall magnitude and phase of S_{21} is unimportant. Outside measurement lines will always add attenuation and change the length unless a thorough calibration is performed. In this case, the overall magnitude and phase are completely irrelevant, so no calibration is necessary. This is accomplished by dividing both the numerator and denominator by h , thus eliminating $\frac{d}{h}$ by choice. This leaves Eq. A.7. Note that $k = \frac{c}{h}$.

$$S_{21} = \frac{i(\omega - \omega_0)}{k + i\omega} \quad (\text{A.7})$$

Replacing k with $\kappa = k + i\omega_0$ is useful for a reason that will become clear later, thus yielding Eq. A.8. A simple rearrangement yields Eq. A.9. Note that the second term here is a Lorentzian and the overall form is already reminiscent of Fig. A.4, the inverted Lorentzian.

$$S_{21} = \frac{i(\omega - \omega_0)}{\kappa + i(\omega - \omega_0)} \quad (\text{A.8})$$

$$S_{21} = 1 - \frac{\kappa}{\kappa + i(\omega - \omega_0)} \quad (\text{A.9})$$

The familiar Lorentzian form exposes these placeholder variables as familiar parameters of a resonance. Splitting the complex κ into its real and imaginary components, κ_R and κ_I aids this process. Before proceeding, dissipation can also be introduced into the resonator; this dissipation parameter will evolve into Q_i . Introducing dissipation will add an imaginary component to the resonance frequency, thus $\omega_0 \rightarrow \omega_0 + i\varepsilon$. These substitutions result in Eq. [A.12](#).

$$\frac{\kappa}{\kappa + i(\omega - \omega_0)} = \frac{\kappa_R + i\kappa_I}{\kappa_R + i\kappa_I + i(\omega - \omega_0)} \quad (\text{A.10})$$

$$= \frac{\kappa_R + i\kappa_I}{\kappa_R + i\kappa_I + i(\omega - (\omega_0 + i\varepsilon))} \quad (\text{A.11})$$

$$= \frac{\kappa_R + i\kappa_I}{(\kappa_R + \varepsilon) + i(\omega - (\omega_0 - \kappa_I))} \quad (\text{A.12})$$

A Lorentzian is commonly expressed in the form of Eq. [A.13](#), where x_0 is the resonance location and γ is the full-width at half maximum, or FWHM. From the definition of total quality factor, it is clear that $Q_0 = \frac{x_0}{\gamma}$. By looking at Equation [A.12](#), one can identify the resonant frequency as $\tilde{\omega}_0 = \omega_0 - \kappa_I$, thus observing that κ_I is effectively a frequency shift of the resonance. Replacing κ_I by $\delta\omega$ using $\kappa_I = -\delta\omega$, thus defining $\tilde{\omega}_0 = \omega_0 + \delta\omega$, simplifies the Lorentzian to Equation [A.14](#).

$$\frac{1}{\frac{1}{2}\gamma + i(x - x_0)} \quad (\text{A.13})$$

$$= \frac{\kappa_R - i\delta\omega}{(\kappa_R + \varepsilon) + i(\omega - \tilde{\omega}_0)} \quad (\text{A.14})$$

The real part of the denominator can be replaced by introducing the total quality factor Q_0 . Matching for the forms of Equations A.13 and A.14 shows that $\kappa_R + \varepsilon = \frac{1}{2}\gamma$, and by definition $Q_0 = \frac{\tilde{\omega}_0}{2(\kappa_R + \varepsilon)}$. Remember that Q_0 can be decomposed into two components, an internal quality factor Q_i and a coupling (external) quality factor Q_c , and note that the system is assumed to be dissipation-less except for the loss added explicitly by ε . Thus it must be that $\frac{2\varepsilon}{\tilde{\omega}_0} = \frac{1}{Q_i}$ and $\frac{2\kappa_R}{\tilde{\omega}_0} = \frac{1}{Q_c}$. Rewriting Equation A.14 in terms of Q_i and Q_c , yields Equation A.16.

$$\frac{1}{Q_0} = \frac{2\kappa_R}{\tilde{\omega}_0} + \frac{2\varepsilon}{\tilde{\omega}_0} \quad (\text{A.15})$$

$$\frac{\frac{\tilde{\omega}_0}{2Q_c} - i\delta\omega}{\left(\frac{\tilde{\omega}_0}{2Q_c} + \frac{\tilde{\omega}_0}{2Q_i}\right) + i(\omega - \tilde{\omega}_0)} \quad (\text{A.16})$$

Taking all of these replacements and writing the full expression for S_{21} yields Eq. A.17. Two simplified equations are also shown, with $x = \frac{\omega - \tilde{\omega}_0}{\tilde{\omega}_0}$. Note that this derived equation is equivalent to another derivation that assumes a lumped element circuit model [123]. The fact that this simple derivation with minimal assumptions derives the full expression from much more complicated models is quite beautiful and highlights the simplistic and fundamental nature of the hanger-coupled resonator.

$$S_{21} = 1 - \frac{\frac{\tilde{\omega}_0}{2Q_c} - i\delta\omega}{\left(\frac{\tilde{\omega}_0}{2Q_c} + \frac{\tilde{\omega}_0}{2Q_i}\right) + i(\omega - \tilde{\omega}_0)} \quad (\text{A.17})$$

$$\boxed{S_{21} = 1 - \frac{\frac{Q_0}{Q_c} - iQ_0\frac{2\delta\omega}{\omega_0}}{1 + 2iQ_0x} = \frac{Q_c + iQ_cQ_i\left(2x + \frac{2\delta\omega}{\omega_0}\right)}{(Q_i + Q_c) + 2iQ_cQ_ix}} \quad (\text{A.18})$$

For comparison, the resonant response with $\delta\omega = 0$, the so-called “symmetric” response, derived in the thesis of J. Gao [124] and used in many earlier papers [91, 82], is shown in Eq. A.19. The “asymmetric” equation can be written in a form that matches the symmetric equation, except with a complex coupling quality factor, \tilde{Q}_c (which also makes a complex total quality factor \tilde{Q}_0); this form is shown in Eq. A.20. The conversion between \tilde{Q}_c and the usual Q_c and $\delta\omega$ are shown in Eqs. A.21 and A.22.

$$S_{21} = 1 - \frac{\frac{Q_0}{Q_c}}{1 + 2iQ_0x} = \frac{Q_c + 2iQ_cQ_ix}{(Q_i + Q_c) + 2iQ_cQ_ix} \quad (\text{A.19})$$

$$S_{21} = 1 - \frac{\frac{\tilde{Q}_0}{\tilde{Q}_c}}{1 + 2i\tilde{Q}_0x} = \frac{\tilde{Q}_c + 2i\tilde{Q}_cQ_ix}{(Q_i + \tilde{Q}_c) + 2i\tilde{Q}_cQ_ix} \quad (\text{A.20})$$

$$\frac{1}{Q_c} = \text{Re} \left\{ \frac{1}{\tilde{Q}_c} \right\} = \frac{|\tilde{Q}_c|^2}{\text{Re}[\tilde{Q}_c]} \quad (\text{A.21})$$

$$\delta\omega \sim \text{Im} \left\{ \frac{1}{\tilde{Q}_c} \right\} \quad (\text{A.22})$$

A.2 Asymmetry Discussion

While actual measurements do not always exhibit a large asymmetry, it is important to use the asymmetric fitting function. Forcing the data to fit a symmetric response results in exaggerated values of Q_i that get more incorrect as the asymmetry parameter $\delta\omega$ increases. A study by Khalil et al. showed that as the asymmetry increases,

one may overestimate Q_i by over an order of magnitude by fitting with the symmetric equation [123]. Therefore, failure to take the asymmetry into account may lead to drastically incorrect results.

In order to better understand the response equation and fitting procedure, one should look at the possible theoretical response curves shown in Fig. A.4. The only difference in the three panels of this figure is the asymmetry parameter $\delta\omega$; the quality factors and resonant frequency are held constant. In the symmetric case, one can gain an intuition for fitting parameters. For the magnitude case, the width determines Q_0 , while the depth determines ratio between Q_i and Q_c . In the Smith-chart representation, essentially a polar plot of the complex S_{21} , the response is a circle whose diameter determines the ratio between Q_i and Q_c and the closest approach to origin determines Q_i . Khalil et al. [123] detail how, in the asymmetric case, these perspectives are skewed by the parameter $\delta\omega$.

Even with agreement on the asymmetric model in the superconducting resonator field, there are debates about how the fitting should be done. In this thesis, all fitting was done in magnitude only, and no noticeable difference was observed when fitting the full Smith chart. Khalil et al. advocate the fitting of the full Smith chart [123], while Megrant et al. argue that one should use fit the data in the form of S_{21}^{-1} [97]; also represented by a circle. The argument for the inverse transmission plot is that the fit parameters are expressed in a more direct way. Regardless of fitting methods, it is clear that the asymmetric model is necessary to accurately obtain quality factor information from hanger-coupled resonators.

The source of the asymmetry in the observed lineshape is often attributed to reflections in the feedline to which the resonator is coupled [123, 97]. A complete derivation of the transmission equation of a hanger-coupled resonator with a generic 4-port coupler yields the symmetric solution, therefore the coupling itself cannot be




Figure A.4: Theoretical hanger response curves for resonator with $Q_i = 10,000$, $Q_c = 1,000$, $f_0 = 10$ GHz, and variable $\delta\omega$ (or δf). The parameter $\delta f = \delta\omega/2\pi$ is shown for three values, (a) 0 Hz (b) 1 MHz and (c) 10 MHz. For the chosen quality factors, these δf values correspond to no asymmetry, slight asymmetry, and extreme asymmetry. In each case, the full response Smith-chart is plotted, along with the magnitude versus frequency. The markers are simulated noise, while the black are equations drawn with correct parameters.

responsible for the asymmetry. However, taking the symmetric equations and adding reflections yields exactly the asymmetric equation.

A better understanding can be found by relating the impedances directly to the asymmetry. While this can be done using the previous derivation, such a calculation yields a large and unhelpful expression. A derivation made with further, simplifying assumptions allows for a much more useful relation to be found. Below, the derivation of S_{21} is branched from Eq. A.6. Before applying the assumptions, terms in Eq. A.6 can be rearranged to yield Eq. A.23 and a further rearrangement yields Eq. A.24, where for simplicity $1/\tilde{z} = (\tilde{z}_1 + \tilde{z}_3)/\tilde{z}_1\tilde{z}_3$.

$$S_{21} = 2 \frac{\frac{z_2 Z_0}{z_1 z_3}}{z_2 \frac{\tilde{z}_1 + \tilde{z}_3}{\tilde{z}_1 \tilde{z}_3} + 1} \quad (\text{A.23})$$

$$S_{21} = \frac{2Z_0}{\tilde{z}_1 + \tilde{z}_3} \left(1 - \frac{1}{z_2/\tilde{z} + 1} \right) \quad (\text{A.24})$$

Eq. A.24 begins to resemble the equation for a circle in the complex plane. The initial scale factor can be ignored to yield a scaled transmission \widetilde{S}_{21} as a means of studying the asymmetry.

Until now, this derivation is equivalent and parallel to the previous derivation. Now, some simplifying assumptions will be made, starting with the assertion that z_2 has a resonance at a frequency ω_0 , leading to an impedance of the form $z_2 \sim x_2 \left(\frac{\omega - \omega_0}{\omega_0} \right)$, where x_2 is purely imaginary and much greater than Z_0 . Also, since there is only one resonant structure under consideration, both z_1 and z_3 are *nonresonant* around ω_0 . Since it must be possible to measure transmission through this device, both z_1 and z_3 must not be too large compared to the characteristic impedance Z_0 . For simplicity, only the case of infinite internal Q is considered since it is not necessary include loss to understand the asymmetry.

These assumptions allow the re-parameterization shown in Eq. A.25, where r is a real number, thus yielding the simplified Eq. A.26. It is important to note that this is an equivalent re-parameterization of Eq. A.20.

$$\frac{z_2}{\tilde{z}} \equiv r e^{i\theta} \left(\frac{\omega - \omega_0}{\omega_0} \right) \quad (\text{A.25})$$

$$\tilde{S}_{21} = 1 - \frac{1}{r e^{i\theta} \left(\frac{\omega - \omega_0}{\omega_0} \right) + 1} \quad (\text{A.26})$$

This is the equation for a circle in the complex plane, whose value far from ω_0 is equal to $\text{Re}[S_{21}] = 1$. The parameter r controls the quality factor of the resonance, which in terms of the circle describes how quickly the lineshape returns to $\text{Re}[S_{21}] = 1$ away from resonance. The parameter θ (a re-parameterization of $\delta\omega$) characterizes the asymmetry by the amount of deviation of θ from $\pi/2$ ($\delta\omega \propto \theta - \pi/2$), with a value of exactly $\pi/2$ corresponding to no asymmetry. The angle θ is the angle of the line tangent to the circle at the resonant frequency, and is thus vertical for $\theta = \pi/2$, leading to a resonance at the origin and a symmetric lineshape. Views of this circle for various θ can be seen below in Fig. A.4, in which the asymmetric shape is shown both as a rotation around the point $\text{Re}[S_{21}] = 1$ and as a scaling of the diameter as explained above.

Temperature Measurement Protocol

Concerning the case of a qubit coupled to a resonator, there are fundamentally two ways to measure the temperature of the cavity, and two ways to measure the temperature of a qubit (see Section 1.1.4 for a qubit temperature definition); these are shown in Fig. B.1. It is very difficult to measure the temperature of a harmonic oscillation from its level population, since it is not possible to address individual levels. However, a cavity coupled to a qubit will inherit some anharmonicity from the qubit, and therefore temperature measurements may be possible. Depending on the qubit and cavity parameters for a particular device, some of these schemes may be impossible; however they are each possible in the limit of vanishing κ and γ .

While method (b) in Fig. B.1 to measure cavity population has already been explored[68], the method described below is a variant of method (c) to measure qubit population. The commonly used method for evaluating P_e in the past compares the heights of the two spectroscopic peaks corresponding to the $|g\rangle$ to $|e\rangle$ (f_{ge}^0) and $|e\rangle$ to $|f\rangle$ (f_{ef}^0) qubit transitions. By measuring the heights of these two peaks, one measures the ratio of population between $|g\rangle$ and $|e\rangle$. While this is generally true, this method does not take into account the readout efficiency variation with qubit state, and is therefore not quantitative without further corrections or calibrations.

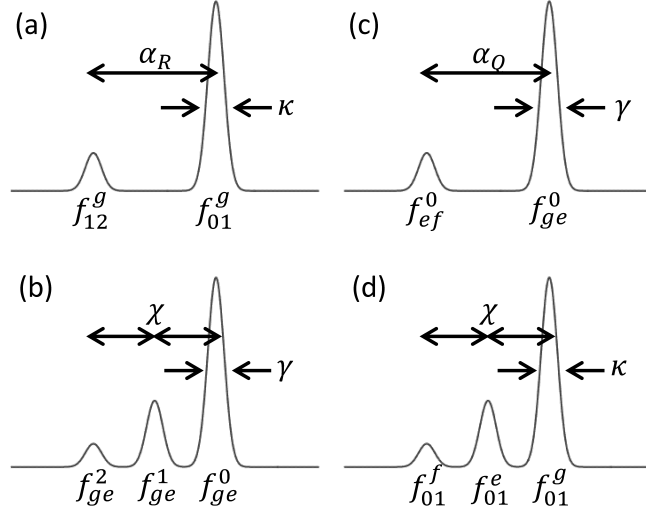


Figure B.1: Four fundamental ways to measure temperature: two for a qubit, and two for cavity, each from various spectroscopic peak heights. One can infer cavity temperature by measuring (a) first two cavity transitions, or (b) multiple qubit lines associated with individual photon numbers. One can infer qubit temperature by measuring (c) first two qubit transitions, or (c) multiple cavity transitions associated with individual qubit states. Note that anharmonicities (α_R and α_Q) in (a) and (c) are typically very different, but the dispersive shift χ is equal in (b) and (d).

This method should be kept for quick checks of whether the qubit is “hot” or not, but should never be used to quote a quantitative P_e or temperature.

While for egregiously “hot” qubits, this qualitative assessment is sufficient, in cases where it is important to quantitatively measure the qubit temperature, it is important to accurately measure the steady-state $|e\rangle$ population, P_e . An improved method for this purpose was developed at Yale, and first written about in Ref. [48], is an adapted version of a measurement of cavity population[72]. This method is called the “Rabi Population Measurement”, or RPM for short. The basic idea of RPM is to measure two Rabi oscillations whose amplitude ratio corresponds directly to the ratio of initial excited state (P_e) to ground state population (P_g). This method is similar to, but different from, techniques previously used in phase qubits [125, 126]. Note that, as mentioned, for the low-temperature limit, it is assumed that $P_g + P_e = 1$.

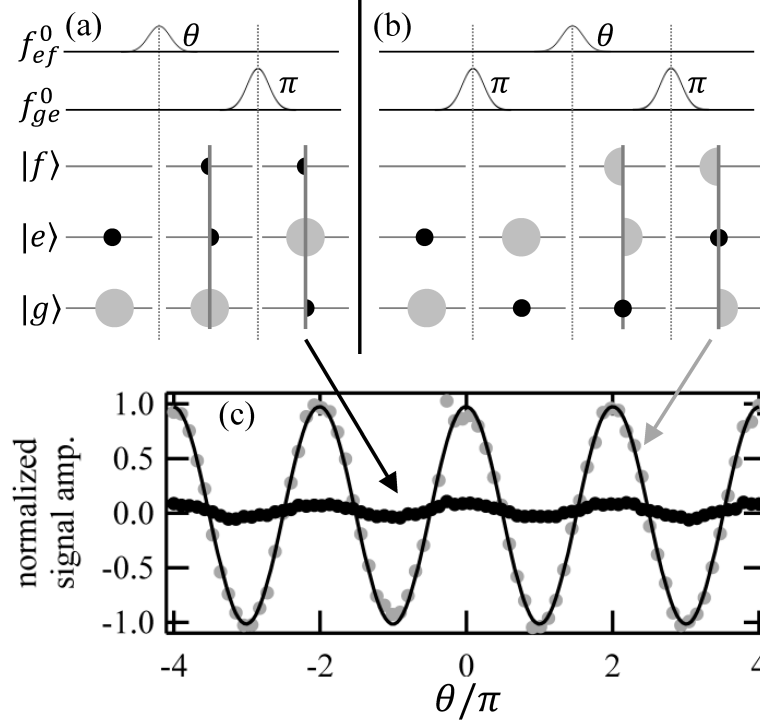


Figure B.2: Upper panel: pulse sequences used to perform RPM qubit population measurement, each producing an oscillation whose amplitude is proportional to initial (a) excited and (b) ground state population. Circle radii indicate population in each state, vertical bars separate the two extrema in Rabi oscillations. Lower panel (c): example normalized data for measurement of 7% excited state population.

The RPM is performed by applying two sequences of qubit pulses as shown in Fig. B.2. The first sequence consists of a pulse performing a rotation around X on the $|e\rangle$ to $|f\rangle$ transition with varying angle $\theta \in [0, 2\pi]$, followed by a π -pulse on the $|g\rangle$ to $|e\rangle$ transition. Measuring the population of the $|g\rangle$ state results in a Rabi oscillation $A_e \cos(\theta)$ with an amplitude A_e proportional to P_e . The second sequence differs only by the insertion of a π -pulse to first invert the population of the $|g\rangle$ and $|e\rangle$ states, yielding a Rabi oscillation $A_g \cos(\theta)$ with an amplitude A_g proportional to P_g . The proportionality constants between the Rabi oscillation amplitudes and the corresponding populations are equal since *the same transition* is used in both sequences, thus avoiding problems from differing readout efficiencies.

From the two oscillation amplitudes, an estimate of the population and its asso-

ciated standard deviation can be calculated from $P_e = A_e / (A_e + A_g)$. The RPM protocol is self-calibrating and accesses smaller amplitudes than crude population measurements since it relies on the amplitude of an oscillation in a lock-in fashion, instead of just one value. The minimum measurable value of P_e is limited by the SNR of the qubit readout, corresponding to the smallest measurable oscillation amplitude. In the reset experiment described in Chapter 6, the minimum distinguishable value was P_e , limited for technical reasons by the characteristics of the readout amplifier chain.

There are a few variants of RPM, each of which use the same principle. Each variant requires two choices: one must choose a pair of levels for the Rabi oscillations, and a readout frequency with optimum readout efficiency. The choice of levels for the Rabi oscillation affects the frequency of the Rabi pulses and the optional π -pulse to move the population to the correct state beforehand. The choice of readout frequency, between f_{01}^g , f_{01}^e , or f_{01}^f determines whether or not a π -pulse is needed after the Rabi pulse. For the variant described above, Rabi oscillations are performed between $|e\rangle$ and $|f\rangle$, and readout is performed at f_{01}^g .

One major limitation of the RPM scheme is that it only works in the low-temperature limit where population of the $|f\rangle$ state or higher is negligible. This limit can be illustrated by deriving the amplitude of the Rabi oscillations. This is done by following the population as done in Fig. B.2, parts (a) and (b), while including an initial $|f\rangle$ state population as well. Calling the initial state populations P_g , P_e , and P_f and readout amplitudes (at a single readout frequency) of R_g , R_e , and R_f , the extrema of the the Fig. B.2(a) Rabi oscillation are shown in Eqs B.1.

$$\begin{aligned} P_f R_g + P_g R_e + P_e R_f &= R_e^A \\ P_e R_g + P_g R_e + P_f R_f &= R_e^B \end{aligned} \tag{B.1}$$

The difference of these two extrema values yields the Rabi amplitude, as shown in Eq. B.2. The Rabi oscillation amplitude is proportional to $P_e - P_f$, with readout sensitivity (scale factor) $R_g - R_f$. Thus in the limit of low temperatures, with P_f negligible, this Rabi oscillation amplitude is proportional to P_e as described above. However, when there is a significant $|f\rangle$ state population, the Rabi oscillation amplitude is affected and RPM no longer effectively estimates the temperature.

$$A_e = R_e^B - R_e^A = (R_g - R_f) (P_e - P_f) \quad (\text{B.2})$$

The Rabi oscillation from Fig. B.2(b) is similarly considered below, with extrema in Eqs B.3 and Rabi amplitude in Eq. B.4. This Rabi oscillation is proportional to $P_g - P_f$, with the same readout sensitivity. The fact that the readout sensitivity (scale factor) is the same between the two Rabi oscillations is the reason RPM works in the first place, as described earlier. When there is a significant $|f\rangle$ state population, this Rabi oscillation amplitude is also affected, although likely to a lesser degree.

$$\begin{aligned} P_f R_g + P_e R_e + P_g R_f &= R_g^A \\ P_g R_g + P_g R_e + P_f R_f &= R_g^B \end{aligned} \quad (\text{B.3})$$

$$A_g = R_g^B - R_g^A = (R_g - R_f) (P_g - P_f) \quad (\text{B.4})$$

Therefore, at higher temperatures, the ratio of the two oscillation amplitudes yields the fairly meaningless (and possibly misleading) $(P_g - P_f)/(P_e - P_f)$. In the low-temperature limit, this ratio simplifies to P_g/P_e , a useful metric for measuring qubit temperature.

Qubit Reset Theory

A method for qubit initialization is one of the fundamental requirements of quantum information processing laid out by DiVincenzo [12]. Due to recent advancements in extending superconducting qubit relaxation times to the 100 μs range [101], active ground state preparation (qubit reset), other than by passively waiting for equilibration with a cold bath, is becoming a necessity. The main use for a fast, high-fidelity reset is to place the qubit into a known pure state either before or during an algorithm. Active reset is preferred over passive reset when a) the qubit thermal environment is hot on the scale of the transition frequency, and b) rapid evacuation of entropy from the system is necessary, as in implementations of quantum error correction [127, 128].

The ancestor of active qubit reset is dynamical cooling of nuclear spins using paramagnetic impurities [129]. Superconducting qubits are analogous to single spins in a controlled environment, and it is therefore possible to design similar dynamical cooling methods to achieve reset times much faster than the relaxation times T_1 . While several methods [111, 112, 40, 130, 71, 72, 70, 104] for reset and dynamical cooling have been demonstrated in superconducting qubits, they each require either qubit tunability or some form of feedback and high-fidelity readout. Described here

is a practical dynamical cooling protocol without these requirements. This protocol is related to dissipation engineering [131], as it uses the dissipation through the cavity to stabilize the qubit ground state. Nicknamed DDROP (Double Drive Rest of Population), this protocol is tested on a transmon qubit in a 3D cavity, but can be applied to any circuit QED system.

DDROP consists of a pulse sequence that manipulates the transition landscape of the qubit-cavity system in order to quickly drive the qubit to the ground state, i.e. removing the entropy from the qubit through the decay of the cavity. The protocol relies on the number splitting property of the strong dispersive regime [23] of circuit QED, where the dispersive shift χ of the cavity due to a qubit excitation is larger than twice the cavity linewidth κ and qubit linewidth $1/T_2$. Thus the cavity frequency depends on the state of excitation of the qubit, and the qubit frequency depends on the number of excitations in the cavity. Another requirement is needed: κ must be much larger than $\Gamma_{\uparrow} = P_e/T_1$, where P_e is the equilibrium excited state population. This condition is easy to satisfy with the recent advances in extending T_1 . Apart from special cases where it is desirable to have small κ , most transmons and other qubits read by a superconducting cavity are candidates for this type of reset.

In the DDROP protocol, shown graphically in Fig. C.1, two microwave drives are applied simultaneously for a duration of order $10 \kappa^{-1}$ in order to reach a steady state. The first drive frequency, f_{ge}^0 , is chosen in order to Rabi drive the qubit if the cavity has zero photons. The amplitude of this drive is quantified by Ω_R , the Rabi frequency. The second frequency, f_c^g , is chosen as to populate the cavity with photons if and only if the qubit is in the ground state. The role of the cavity drive is to lift the population of $|g, 0\rangle$ to the coherent state $|g, \alpha\rangle$, where $|\alpha|^2 = \bar{n}$, the steady state average photon number in the cavity. Due to number splitting, the




Figure C.1: Level structure of the transmon qubit coupled dispersively to a single resonator mode. The qubit excitations are spanned vertically while the resonator photon numbers are spanned horizontally. The arrows show the transitions involved in the DDROP procedure along with their rates, with $\Gamma_{\uparrow} \ll \kappa \approx \Omega_R < \chi/2$. The double arrows are driven transitions, while single arrows are spontaneous. Qubit transitions are represented by straight lines while cavity transitions are wavy lines. The steady-state equilibrium qubit/cavity state is the coherent state $|g, \alpha\rangle$. For visualization, the state $|g, m\rangle$ is highlighted, where m is the closest integer to the steady-state average number of photons in the cavity.

qubit transition frequencies when the cavity is in state $|\alpha\rangle$ differ sufficiently from f_{ge}^0 that the Rabi drive does not excite $|g, \alpha\rangle$. The only way for the system to leave $|g, \alpha\rangle$ is through a spontaneous excitation happening at a rate Γ_{\uparrow} , which is slow compared to all other rates in the system. Once in $|e, \alpha\rangle$, the system rapidly falls back to $|e, 0\rangle$ in a time of order κ^{-1} . The role of the Rabi drive, with Rabi frequency of order κ , is to speed up the transition between $|e, 0\rangle$ and $|g, 0\rangle$, thus allowing for a fast return to $|g, \alpha\rangle$. With both drives on, the system will be driven to $|g, \alpha\rangle$ at a rate of order κ regardless of initial state, while the rate Γ_{\uparrow} away from this state is slow. Eventually, to prepare $|g, 0\rangle$ instead of $|g, \alpha\rangle$, one must turn off the drives and wait for the photons to decay in a time of several κ^{-1} . Since the cavity is in a coherent state, this waiting time could be avoided by using a displacement pulse, which is easier to calibrate for cavities with higher quality factor. The ratio κ/Γ_{\uparrow} determines the fidelity of the ground state preparation, and must therefore be much greater than

1.

Results showing measurements of the DDROP protocol performed on a 3D transmon qubit are shown in Section 6.3.2. Regardless of initial state, a ground state preparation fidelity of 99.5% was achieved in less than 3 μ s. Simulation predicts higher fidelities are possible; for reasons discussed in Chapter 6.

DDROP is not the first demonstrated qubit reset mechanism to work on superconducting qubits; several distinct methods have been shown previously, including: sideband cooling through higher energy levels [111], sweeping the qubit frequency into resonance with a low- Q cavity [40, 130], a feedback loop with conditional coherent driving [71], and strong projective measurements [72, 70, 104]. However, DDROP has many advantages when compared to each of these processes. First, there is no need to tune in real time the qubit frequency, which means DDROP will still work with fixed-frequency qubits. There is no need for fast external feedback of any kind, thus simplifying the required setup. There is also no need for high-fidelity, single-shot readouts or in fact a low-noise amplifier at all. Finally, the decisive qualitative advantage is that sensitivity to drive amplitudes is low, and that there is no need for accurate pulse timing or shapes; DDROP can be quickly tuned to near optimum parameters.

Bibliography

- [1] D. Jaksch, J. I. Cirac, P. Zoller, S. L. Rolston, R. Côté, and M. D. Lukin. Fast quantum gates for neutral atoms. *Phys. Rev. Lett.*, 85:2208–2211, 2000. [1](#)
- [2] M. Saffman, T. G. Walker, and K. Mølmer. Quantum information with Rydberg atoms. *Rev. Mod. Phys.*, 82:2313–2363, 2010. [1](#)
- [3] C. Monroe, D. M. Meekhof, B. E. King, W. M. Itano, and D. J. Wineland. Demonstration of a fundamental quantum logic gate. *Phys. Rev. Lett.*, 75:4714–4717, 1995. [1](#)
- [4] D. Loss and D. P. DiVincenzo. Quantum computation with quantum dots. *Phys. Rev. A*, 57:120–126, 1998. [1](#)
- [5] J. R. Weber, W. F. Koehl, J. B. Varley, A. Janotti, B. B. Buckley, C. G. Van de Walle, and D. D. Awschalom. Quantum computing with defects. *Proceedings of the National Academy of Sciences*, 107:8513–8518, 2010. [1](#)
- [6] V. Bouchiat, D. Vion, P. Joyez, D. Esteve, and M. H. Devoret. Quantum coherence with a single Cooper pair. *Physica Scripta*, 1998:165, 1998. [1](#), [13](#)
- [7] J. M. Martinis, S. Nam, J. Aumentado, and C. Urbina. Rabi oscillations in a large Josephson-junction qubit. *Phys. Rev. Lett.*, 89:117901, 2002. [1](#), [13](#)

- [8] A. Wallraff, D. I. Schuster, A. Blais, L. Frunzio, R.-S. Huang, J. Majer, S. Kumar, S. M. Girvin, and R. J. Schoelkopf. Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature*, 431:162–167, 2004. [1](#), [9](#)
- [9] M. H. Devoret, A. Wallraff, and J. M. Martinis. Superconducting qubits: A short review. *ArXiv Condensed Matter e-prints*, 2004. [1](#), [6](#)
- [10] I. Chiorescu, P. Bertet, K. Semba, Y. Nakamura, C. J. P. M. Harmans, and J. E. Mooij. Coherent dynamics of a flux qubit coupled to a harmonic oscillator. *Nature*, 431:159–162, 2004. [1](#)
- [11] J. Clarke and F. K. Wilhelm. Superconducting quantum bits. *Nature*, 453:1031–1042, 2008. [1](#), [6](#)
- [12] D. P. DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik*, 48:771–783, 2000. [1](#), [6](#), [209](#)
- [13] P.W. Shor. Algorithms for quantum computation: discrete logarithms and factoring. In *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*, pages 124–134, 1994. [4](#)
- [14] L. K. Grover. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 212–219, Philadelphia, Pennsylvania, USA, 1996. ACM. [4](#)
- [15] L. K. Grover. Quantum mechanics helps in searching for a needle in a haystack. *Phys. Rev. Lett.*, 79:325–328, 1997. [4](#)
- [16] A. Dewes, R. Lauro, F. R. Ong, V. Schmitt, P. Milman, P. Bertet, D. Vion,

- and D. Esteve. Quantum speeding-up of computation demonstrated in a superconducting two-qubit processor. *Phys. Rev. B*, 85:140503, 2012. 5
- [17] E. L. Hahn. Spin echoes. *Phys. Rev.*, 80:580–594, 1950. 7
- [18] H. Y. Carr and E. M. Purcell. Effects of diffusion on free precession in nuclear magnetic resonance experiments. *Phys. Rev.*, 94:630–638, 1954. 7
- [19] S. Meiboom and D. Gill. Modified spin-echo method for measuring nuclear relaxation times. *Rev. Sci. Instrum.*, 29:688–691, 1958. 7
- [20] M. Brune, E. Hagley, X. Maitre, G. Nogues, C. Wunderlich, J-M. Raimond, and S. Haroche. Manipulating entanglement with atoms and photons in a cavity. In *Quantum Electronics Conference, 1998. IQEC 98. Technical Digest. Summaries of papers presented at the International*, pages 145–, 1998. 8
- [21] H. Mabuchi and A. C. Doherty. Cavity quantum electrodynamics: Coherence in context. *Science*, 298:1372–1377, 2002. 8
- [22] E. T. Jaynes and F. W. Cummings. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proceedings of the IEEE*, 51:89–109, 1963. 8
- [23] D. I. Schuster, A. A. Houck, J. A. Schreier, A. Wallraff, J. M. Gambetta, A. Blais, L. Frunzio, J. Majer, B. Johnson, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Resolving photon number states in a superconducting circuit. *Nature*, 445:515–518, 2007. 9, 34, 41, 63, 65, 210
- [24] M. H Devoret, D. Esteve, J. M Martinis, and C. Urbina. Effect of an adjustable admittance on the macroscopic energy levels of a current biased josephson junction. *Physica Scripta*, 1989:118, 1989. 9

- [25] B.D. Josephson. Possible new effects in superconductive tunnelling. *Physics Letters*, 1:251–253, 1962. [10](#), [53](#), [183](#)
- [26] D. Vion, A. Aassime, A. Cottet, P. Joyez, H. Pothier, C. Urbina, D. Esteve, and M. H. Devoret. Manipulating the quantum state of an electrical circuit. *Science*, 296:886–889, 2002. [13](#)
- [27] J. R. Friedman, V. Patel, W. Chen, S. K. Tolpygo, and J. E. Lukens. Quantum superposition of distinct macroscopic states. *Nature*, 406:43–46, 2000. [13](#)
- [28] C. H. van der Wal, A. C. J. ter Haar, F. K. Wilhelm, R. N. Schouten, C. J. P. M. Harmans, T. P. Orlando, S. Lloyd, and J. E. Mooij. Quantum superposition of macroscopic persistent-current states. *Science*, 290:773–777, 2000. [13](#)
- [29] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Charge-insensitive qubit design derived from the Cooper pair box. *Phys. Rev. A*, 76:042319, 2007. [14](#), [15](#), [16](#), [23](#), [24](#), [25](#), [34](#), [35](#)
- [30] J. A. Schreier, A. A. Houck, Jens Koch, D. I. Schuster, B. R. Johnson, J. M. Chow, J. M. Gambetta, J. Majer, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Suppressing charge noise decoherence in superconducting charge qubits. *Phys. Rev. B*, 77:180502, 2008. [14](#), [16](#), [34](#)
- [31] A. A. Houck, J. Koch, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Life after charge noise: recent results with transmon qubits. 8:105–115, 2009. [14](#), [16](#), [22](#), [34](#)
- [32] J. Koch, V. Manucharyan, M. H. Devoret, and L. I. Glazman. Charging effects in the inductively shunted Josephson junction. *Phys. Rev. Lett.*, 103:217004, 2009. [16](#)

- [33] V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret. Fluxonium: Single Cooper-pair circuit free of charge offsets. *Science*, 326:113–116, 2009. [16](#), [43](#), [158](#)
- [34] V. E. Manucharyan, J. Koch, M. Brink, L. I. Glazman, and M. H. Devoret. Coherent oscillations between classically separable quantum states of a superconducting loop. *arXiv e-prints*, 0910:3039, 2009. [16](#)
- [35] V. E. Manucharyan, N. A. Masluk, A. Kamal, J. Koch, L. I. Glazman, and M. H. Devoret. Evidence for coherent quantum phase slips across a Josephson junction array. *Phys. Rev. B*, 85:024521, 2012. [16](#), [158](#), [159](#)
- [36] N. Masluk. *Reducing the Losses of the Fluxonium Artificial Atom*. PhD thesis, Yale University, 2012. [16](#), [43](#), [45](#), [48](#), [50](#), [86](#), [154](#), [160](#), [166](#), [187](#)
- [37] L. Frunzio, G. Catelani, A. Kamal, L.I. Glazman, M.H. Devoret, and R.J. Schoelkopf. Analyzing energy relaxation of superconducting qubits. *in Preparation*, 2013. [19](#)
- [38] E. M. Purcell. Spontaneous emission probabilities at radio frequencies. *Physical Review*, 69:681, 1946. [21](#)
- [39] A. A. Houck, J. A. Schreier, B. R. Johnson, J. M. Chow, Jens Koch, J. M. Gambetta, D. I. Schuster, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf. Controlling the spontaneous emission of a superconducting transmon qubit. *Phys. Rev. Lett.*, 101:080502, 2008. [21](#), [116](#), [117](#)
- [40] M. D. Reed, B. R. Johnson, A. A. Houck, L. DiCarlo, J. M. Chow, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf. Fast reset and suppressing spontaneous emission of a superconducting qubit. *Appl. Phys. Lett.*, 96:203110–3, 2010. [21](#), [209](#), [212](#)

- [41] R. M. Lutchyn, L. I. Glazman, and A. I. Larkin. Kinetics of the superconducting charge qubit in the presence of a quasiparticle. *Phys. Rev. B*, 74:064515, 2006. [23](#)
- [42] A. P. Sears, A. Petrenko, G. Catelani, L. Sun, H. Paik, G. Kirchmair, L. Frunzio, L. I. Glazman, S. M. Girvin, and R. J. Schoelkopf. Photon shot noise dephasing in the strong-dispersive limit of circuit QED. *Phys. Rev. B*, 86:180504, 2012. [23](#), [140](#), [141](#), [169](#), [184](#)
- [43] C. Rigetti, J. M. Gambetta, S. Poletto, B. L. T. Plourde, J. M. Chow, A. D. Corcoles, J. A. Smolin, S. T. Merkel, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen. Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms. *Phys. Rev. B*, 86:100506, 2012. [23](#), [80](#), [140](#), [141](#), [184](#)
- [44] A. B. Zorin, F.-J. Ahlers, J. Niemeyer, T. Weimann, H. Wolf, V. A. Kru-
penin, and S. V. Lotkhov. Background charge noise in metallic single-electron
tunneling devices. *Phys. Rev. B*, 53:13682–13687, 1996. [24](#)
- [45] F. C. Wellstood, C. Urbina, and J. Clarke. Low-frequency noise in dc supercon-
ducting quantum interference devices below 1 K. *Appl. Phys. Lett.*, 50:772–774,
1987. [25](#)
- [46] D. J. Van Harlingen, T. L. Robertson, B. L. T. Plourde, P. A. Reichardt,
T. A. Crane, and J. Clarke. Decoherence in Josephson-junction qubits due to
critical-current fluctuations. *Phys. Rev. B*, 70:064517, 2004. [25](#)
- [47] K. Geerlings, S. Shankar, E. Edwards, L. Frunzio, R. J. Schoelkopf, and M. H.
Devoret. Improving the quality factor of microwave compact resonators by op-

- timizing their geometrical parameters. *Appl. Phys. Lett.*, 100:192601–3, 2012. [27](#), [104](#)
- [48] K. Geerlings, Z. Leghtas, I. M. Pop, S. Shankar, L. Frunzio, R. J. Schoelkopf, M. Mirrahimi, and M. H. Devoret. Demonstrating a driven reset protocol for a superconducting qubit. *Phys. Rev. Lett.*, 110:120501, 2013. [28](#), [205](#)
- [49] M. H. Devoret and R. J. Schoelkopf. Superconducting circuits for quantum information: An outlook. *Science*, 339:1169–1174, 2013. [29](#)
- [50] H. Wang, M. Hofheinz, J. Wenner, M. Ansmann, R. C. Bialczak, M. Lenander, E. Lucero, M. Neeley, A. D. O’Connell, D. Sank, M. Weides, A. N. Cleland, and J. M. Martinis. Improving the coherence time of superconducting coplanar resonators. *Appl. Phys. Lett.*, 95:233508, 2009. [31](#)
- [51] R. Barends, N. Vercruyssen, A. Endo, P. J. de Visser, T. Zijlstra, T. M. Klapwijk, P. Diener, S. J. C. Yates, and J. J. A. Baselmans. Minimal resonator loss for circuit quantum electrodynamics. *Appl. Phys. Lett.*, 97:023508, 2010. [31](#), [91](#), [93](#)
- [52] J. Chang, M. R. Vissers, A. D. Corcoles, M. Sandberg, J. Gao, David W. Abraham, Jerry M. Chow, Jay M. Gambetta, M. B. Rothwell, G. A. Keefe, M. Steffen, and D. P. Pappas. Improved superconducting qubit coherence using titanium nitride. *ArXiv e-prints*, 1303:4071, 2013. [31](#), [135](#)
- [53] D. M. Pozar. *Microwave Engineering*. Wiley, 2004. [33](#), [139](#), [194](#)
- [54] L. Bishop. *Circuit Quantum Electrodynamics*. PhD thesis, Yale University, 2010. [34](#)

- [55] J. Chow. *Quantum Information Processing with Superconducting Qubits*. PhD thesis, Yale University, 2010. 34
- [56] S. E. Nigg, H. Paik, B. Vlastakis, G. Kirchmair, S. Shankar, L. Frunzio, M. H. Devoret, R. J. Schoelkopf, and S. M. Girvin. Black-box superconducting circuit quantization. *Phys. Rev. Lett.*, 108:240502, 2012. 34
- [57] M. H. Devoret, E. Giacobino, and J. Zinn-Justin. Quantum fluctuations in electrical circuits. In S. Reynaud, editor, *Fluctuations Quantiques/Quantum Fluctuations*, page 351, 1997. 35
- [58] S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. Circuit QED and engineering charge-based superconducting qubits. *Physica Scripta*, 2009:014012, 2009. 41, 63
- [59] V. Manucharyan. *Superinductance*. PhD thesis, Yale University, 2012. 43
- [60] R. Shankar. *Principles of Quantum Mechanics*. Springer, 1994. 46
- [61] R. J. Schoelkopf, A. A. Clerk, S. M. Girvin, K. W. Lehnert, and M. H. Devoret. Qubits as spectrometers of quantum noise. *arXiv*, page 10247, 2002. 50
- [62] G. Catelani, R. J. Schoelkopf, M. H. Devoret, and L. I. Glazman. Relaxation and frequency shifts induced by quasiparticles in superconducting qubits. *Phys. Rev. B*, 84:064517, 2011. 53, 134, 183
- [63] D. N. Langenberg. Physical interpretation of the $\cos\varphi$ term and implications for detectors. *Rev. Phys. Appl. (Paris)*, 9:35–40, 1974. 53
- [64] J. Niemeyer and V. Kose. Observation of large dc supercurrents at nonzero voltages in Josephson tunnel junctions. *Appl. Phys. Lett.*, 29:380–382, 1976. 58

- [65] G. J. Dolan. Offset masks for lift-off photoprocessing. *Appl. Phys. Lett.*, 31:337–339, 1977. [58](#)
- [66] C. Rigetti. *Quantum Gates For Superconducting Qubits*. PhD thesis, Yale University, 2009. [61](#), [142](#), [158](#)
- [67] F. Lecocq, I. M Pop, Z. Peng, I. Matei, T. Crozes, T. Fournier, C. Naud, W. Guichard, and O. Buisson. Junction fabrication by shadow evaporation without a suspended bridge. *Nanotechnology*, 22:315302, 2011. [61](#), [142](#), [158](#)
- [68] B. R. Johnson, M. D. Reed, A. A. Houck, D. I. Schuster, Lev S. Bishop, E. Ginossar, J. M. Gambetta, L. DiCarlo, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. Quantum non-demolition detection of single microwave photons in a circuit. *Nat Phys*, 6:663–667, 2010. [65](#), [204](#)
- [69] M. Hatridge, S. Shankar, M. Mirrahimi, F. Schackert, K. Geerlings, T. Brecht, K. M. Sliwa, B. Abdo, L. Frunzio, S. M. Girvin, R. J. Schoelkopf, and M. H. Devoret. Quantum back-action of an individual variable-strength measurement. *Science*, 339:178–181, 2013. [67](#), [68](#), [71](#), [141](#)
- [70] D. Ristè, J. G. van Leeuwen, H.-S. Ku, K. W. Lehnert, and L. DiCarlo. Initialization by measurement of a superconducting quantum bit circuit. *Phys. Rev. Lett.*, 109:050507, 2012. [67](#), [141](#), [209](#), [212](#)
- [71] D. Ristè, C. C. Bultink, K. W. Lehnert, and L. DiCarlo. Feedback control of a solid-state qubit using high-fidelity projective measurement. *arXiv e-prints*, 2012. [67](#), [209](#), [212](#)
- [72] J. E. Johnson, C. Macklin, D. H. Slichter, R. Vijay, E. B. Weingarten, John Clarke, and I. Siddiqi. Heralded state preparation in a superconducting qubit. *Phys. Rev. Lett.*, 109:050506, 2012. [67](#), [134](#), [205](#), [209](#), [212](#)

- [73] M. D. Reed, L. DiCarlo, B. R. Johnson, L. Sun, D. I. Schuster, L. Frunzio, and R. J. Schoelkopf. High-fidelity readout in circuit quantum electrodynamics using the Jaynes-Cummings nonlinearity. *Physical Review Letters*, 105:173601, 2010. [68](#), [71](#)
- [74] L. S. Bishop, E. Ginossar, and S. M. Girvin. Response of the strongly driven Jaynes-Cummings oscillator. *Phys. Rev. Lett.*, 105:100505, 2010. [68](#)
- [75] Jay Gambetta, W. A. Braff, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf. Protocols for optimal readout of qubits using a continuous quantum nondemolition measurement. *Phys. Rev. A*, 76(1):012325–, July 2007. [69](#)
- [76] F. Pobell. *Matter and Methods at Low Temperatures*. Springer-Verlag, 1996. [74](#)
- [77] R. Barends, J. Wenner, M. Lenander, Y. Chen, R. C. Bialczak, J. Kelly, E. Lucero, P. O’Malley, M. Mariantoni, D. Sank, H. Wang, T. C. White, Y. Yin, J. Zhao, A. N. Cleland, J. M. Martinis, and J. J. A. Baselmans. Loss and decoherence due to stray infrared light in superconducting quantum circuits. *ArXiv e-prints*, 2011. [80](#), [87](#), [123](#), [133](#)
- [78] D. F. Santavicca and D. E. Prober. Impedance-matched low-pass stripline filters. *Measurement Science and Technology*, 19:087001, 2008. [83](#)
- [79] A. D. Corcoles, J. M. Chow, J. M. Gambetta, C. Rigetti, J. R. Rozen, G. A. Keefe, M. B. Rothwell, M. B. Ketchen, and M. Steffen. Protecting superconducting qubits from radiation. *Appl. Phys. Lett.*, 99:181906–3, 2011. [87](#), [133](#), [134](#), [143](#)
- [80] J. S. Gao, M. Daal, A. Vayonakis, S. Kumar, J. Zmuidzinas, B. Sadoulet, B. A. Mazin, P. K. Day, and H. G. Leduc. Experimental evidence for a surface

- distribution of two-level systems in superconducting lithographed microwave resonators. *Appl. Phys. Lett.*, 92:152505, 2008. [91](#), [93](#)
- [81] R. Barends, H. L. Hortensius, T. Zijlstra, J. J. A. Baselmans, S. J. C. Yates, J. R. Gao, and T. M. Klapwijk. Contribution of dielectrics to frequency and noise of NbTiN superconducting resonators. *Appl. Phys. Lett.*, 92:223502–3, 2008. [91](#)
- [82] M. S. Khalil, F. C. Wellstood, and K. D. Osborn. Loss dependence on geometry and applied power in superconducting coplanar resonators. *Applied Superconductivity, IEEE Transactions on*, 21:879–882, 2011. [91](#), [93](#), [94](#), [199](#)
- [83] W.A. Phillips. *Amorphous solids: low-temperature properties*. Springer-Verlag, 1981. [91](#)
- [84] J. Classen, C. Enss, C. Bechinger, G. Weiss, and S. Hunklinger. Low frequency acoustic and dielectric measurements on glasses. *Ann. Phys.*, 506:315–335, 1994. [91](#)
- [85] P. Esquinazi. *Tunneling Systems in Amorphous and Crystalline Solids*. Springer, 1998. [91](#)
- [86] A. L. Burin, M. S. Khalil, and Kevin D. Osborn. Universal dielectric loss in glass from simultaneous bias and microwave fields. *Phys. Rev. Lett.*, 110:157002, 2013. [91](#), [93](#)
- [87] D. C. Mattis and J. Bardeen. Theory of the anomalous skin effect in normal and superconducting metals. *Phys. Rev.*, 111:412–417, 1958. [91](#), [105](#)
- [88] M. Tinkham. *Introduction to superconductivity*. Dover Publications, Incorporated, 1996. [92](#)

- [89] R. Barends. *Photon-detecting superconducting resonators*. PhD thesis, Delft University of Technology, 2009. 92
- [90] M. von Schickfus and S. Hunklinger. The dielectric coupling of low-energy excitations in vitreous silica to electromagnetic waves. *Journal of Physics C: Solid State Physics*, 9:L439, 1976. 92
- [91] T. Lindstroem, J. E. Healey, M. S. Colclough, C. M. Muirhead, and A. Y. Tzalenchuk. Properties of superconducting planar resonators at millikelvin temperatures. *Phys. Rev. B*, 80:132501, 2009. 94, 199
- [92] H. G. Leduc, B. Bumble, P. K. Day, B. H. Eom, J. Gao, S. Golwala, B. A. Mazin, S. McHugh, A. Merrill, D. C. Moore, O. Noroozian, A. D. Turner, and J. Zmuidzinas. Titanium nitride films for ultrasensitive microresonator detectors. *Appl. Phys. Lett.*, 97:102509–3, 2010. 94
- [93] A. D. O’Connell, M. Ansmann, R. C. Bialczak, M. Hofheinz, N. Katz, E. Lucero, C. McKenney, M. Neeley, H. Wang, E. M. Weig, A. N. Cleland, and J. M. Martinis. Microwave dielectric loss at single photon energies and millikelvin temperatures. *Appl. Phys. Lett.*, 92:112903–3, 2008. 96
- [94] S. J. Weber, K. W. Murch, D. H. Slichter, R. Vijay, and I. Siddiqi. Single crystal silicon capacitors with low microwave loss in the single photon regime. *Appl. Phys. Lett.*, 98:172510–3, 2011. 96
- [95] P. Benedek. Capacitances of a planar multiconductor configuration on a dielectric substrate by a mixed order finite-element method. *Circuits and Systems, IEEE Transactions on*, 23:279–284, 1976. 97
- [96] J. Wenner, R. Barends, R. C. Bialczak, Yu Chen, J. Kelly, E. Lucero, M. Mariantoni, A. Megrant, P. J. J. O’Malley, D. Sank, A. Vainsencher,

- H. Wang, T. C. White, Y. Yin, J. Zhao, A. N. Cleland, and J. M. Martinis. Surface loss simulations of superconducting coplanar waveguide resonators. *Appl. Phys. Lett.*, 99:113513–3, 2011. [111](#)
- [97] A. Megrant, C. Neill, R. Barends, B. Chiaro, Yu Chen, L. Feigl, J. Kelly, E. Lucero, M. Mariantoni, P. J. J. O’Malley, D. Sank, A. Vainsencher, J. Wenner, T. C. White, Y. Yin, J. Zhao, C. J. Palmstrom, J. M. Martinis, and A. N. Cleland. Planar superconducting resonators with internal quality factors above one million. *Appl. Phys. Lett.*, 100:113510–4, 2012. [113](#), [189](#), [200](#)
- [98] R. Vijay, D. H. Slichter, and I. Siddiqi. Observation of quantum jumps in a superconducting artificial atom. *Phys. Rev. Lett.*, 106:110502, 2011. [134](#)
- [99] R. Barends, J. Kelly, A. Megrant, D. Sank, E. Jeffrey, Y. Chen, Y. Yin, B. Chiaro, J. Mutus, C. Neill, P. O’Malley, P. Roushan, J. Wenner, T. C. White, A. N. Cleland, and J. M. Martinis. Coherent Josephson qubit suitable for scalable quantum integrated circuits. *ArXiv e-prints*, 1304:2322, 2013. [135](#)
- [100] M. Reagor, H. Paik, G. Catelani, L. Sun, C. Axline, E. Holland, I. M. Pop, N. A. Masluk, T. Brecht, L. Frunzio, M. H. Devoret, L. I. Glazman, and R. J. Schoelkopf. Ten milliseconds for aluminum cavities in the quantum regime. *ArXiv e-prints*, 2013. [139](#)
- [101] H. Paik, D. I. Schuster, L. S. Bishop, G. Kirchmair, G. Catelani, A. P. Sears, B. R. Johnson, M. J. Reagor, L. Frunzio, L. I. Glazman, S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf. Observation of high coherence in Josephson junction qubits measured in a three-dimensional circuit QED architecture. *Phys. Rev. Lett.*, 107:240501, 2011. [140](#), [209](#)
- [102] K. W. Murch, U. Vool, D. Zhou, S. J. Weber, S. M. Girvin, and I. Siddiqi.

- Cavity-assisted quantum bath engineering. *Phys. Rev. Lett.*, 109:183602, 2012. [141](#), [154](#)
- [103] R. Vijay, C. Macklin, D. H. Slichter, S. J. Weber, K. W. Murch, R. Naik, A. N. Korotkov, and I. Siddiqi. Stabilizing Rabi oscillations in a superconducting qubit using quantum feedback. *Nature*, 490:77–80, 2012. [141](#)
- [104] P. Campagne-Ibarcq, E. Flurin, N. Roch, D. Darson, P. Morfin, M. Mirrahimi, M. H. Devoret, F. Mallet, and B. Huard. Stabilizing the trajectory of a superconducting qubit by projective measurement feedback. *ArXiv e-prints*, 1301:6095, 2013. [141](#), [209](#), [212](#)
- [105] D. Ristè, C. C. Bultink, M. J. Tiggelman, R. N. Schouten, K. W. Lehnert, and L. DiCarlo. Millisecond charge-parity fluctuations and induced decoherence in a superconducting qubit. *ArXiv e-prints*, 1212:5459, 2012. [141](#)
- [106] H. Paik and R. J. Schoelkopf. *in Preparation*, 2012. [153](#)
- [107] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. *Rev. Mod. Phys.*, 75:281–324, 2003. [154](#)
- [108] J. Chan, T. P. Mayer Alegre, A. H. Safavi-Naeini, J. T. Hill, A. Krause, S. Gröblacher, M. Aspelmeyer, and O. Painter. Laser cooling of a nanomechanical oscillator into its quantum ground state. *Nature*, 478:89–92, 2011. [154](#)
- [109] R. Rivière, S. Deléglise, S. Weis, E. Gavartin, O. Arcizet, A. Schliesser, and T. J. Kippenberg. Optomechanical sideband cooling of a micromechanical oscillator close to the quantum ground state. *Phys. Rev. A*, 83:063835, 2011. [154](#)

- [110] J. D. Teufel, T. Donner, Dale Li, J. W. Harlow, M. S. Allman, K. Cicak, A. J. Sirois, J. D. Whittaker, K. W. Lehnert, and R. W. Simmonds. Sideband cooling of micromechanical motion to the quantum ground state. *Nature*, 475:359–363, 2011. [154](#)
- [111] S. O. Valenzuela, W. D. Oliver, D. M. Berns, K. K. Berggren, L. S. Levitov, and T. P. Orlando. Microwave-induced cooling of a superconducting qubit. *Science*, 314:1589–1592, 2006. [154](#), [209](#), [212](#)
- [112] M. Grajcar, S. H. W. van der Ploeg, A. Izmailkov, E. Il'ichev, H.-G. Meyer, A. Fedorov, A. Shnirman, and Gerd Schön. Sisyphus cooling and amplification by a superconducting qubit. *Nat Phys*, 4:612–616, 2008. [154](#), [209](#)
- [113] N. A. Masluk, I. M. Pop, A. Kamal, Z. K. Mineev, and M. H. Devoret. Microwave characterization of Josephson junction arrays: Implementing a low loss superinductance. *Phys. Rev. Lett.*, 109:137002, 2012. [158](#)
- [114] V. B. Braginsky, V. S. Ilchenko, and Kh. S. Bagdassarov. Experimental observation of fundamental microwave absorption in high-quality dielectric crystals. *Physics Letters A*, 120:300–305, 1987. [177](#)
- [115] N. F. Pedersen, T. F. Finnegan, and D. N. Langenberg. Magnetic field dependence and Q of the Josephson plasma resonance. *Phys. Rev. B*, 6:4151–4159, 1972. [183](#)
- [116] C. M. Falco, W. H. Parker, and S. E. Trullinger. Observation of a phase-modulated quasiparticle current in superconducting weak links. *Phys. Rev. Lett.*, 31:933–936, 1973. [183](#)
- [117] D. A. Vincent and Jr. Deaver, B. S. Observation of a phase-dependent conduc-

- tivity in superconducting point contacts. *Phys. Rev. Lett.*, 32:212–215, 1974. [183](#)
- [118] M. Nisenoff and S. Wolf. Observation of a $\cos\varphi$ term in the current-phase relation for "dayem"-type weak link contained in an rf-biased superconducting quantum interference device. *Phys. Rev. B*, 12:1712–1714, 1975. [183](#)
- [119] O. H. Soerensen, J. Mygind, and N. F. Pedersen. Measured temperature dependence of the $\cos\varphi$ conductance in Josephson tunnel junctions. *Phys. Rev. Lett.*, 39:1018–1021, 1977. [183](#)
- [120] F. Giazotto and M. J. Martinez-Perez. The josephson heat interferometer. *Nature*, 492:401–405, 2012. [183](#)
- [121] J. Leppäkangas, M. Marthaler, and G. Schön. Phase-dependent quasiparticle tunneling in Josephson junctions: Measuring the $\cos\varphi$ term with a superconducting charge qubit. *Phys. Rev. B*, 84:060505, 2011. [183](#)
- [122] Z. Kim, B. Suri, V. Zaretsky, S. Novikov, K. D. Osborn, A. Mizel, F. C. Wellstood, and B. S. Palmer. Decoupling a cooper-pair box to enhance the lifetime to 0.2ms. *Phys. Rev. Lett.*, 106:120501, 2011. [185](#)
- [123] M. S. Khalil, M. J. A. Stoutimore, F. C. Wellstood, and K. D. Osborn. An analysis method for asymmetric resonator transmission applied to superconducting devices. *J. Appl. Phys.*, 111:054510–6, 2012. [198](#), [200](#)
- [124] J. Gao. *The Physics of Superconducting Microwave Resonators*. PhD thesis, California Institute of Technology, 2008. [199](#)
- [125] E. Lucero, M. Hofheinz, M. Ansmann, R. C. Bialczak, N. Katz, M. Neeley,

- A. D. O’Connell, H. Wang, A. N. Cleland, and J. M. Martinis. High-fidelity gates in a single Josephson qubit. *Phys. Rev. Lett.*, 100:247001, 2008. [205](#)
- [126] M. Neeley, M. Ansmann, R. C. Bialczak, M. Hofheinz, E. Lucero, A. D. O’Connell, D. Sank, H. Wang, J. Wenner, A. N. Cleland, M. R. Geller, and J. M. Martinis. Emulation of a quantum spin with a superconducting phase qudit. *Science*, 325:722–725, 2009. [205](#)
- [127] P. Schindler, J. T. Barreiro, T. Monz, V. Nebendahl, D. Nigg, M. Chwalla, M. Hennrich, and R. Blatt. Experimental repetitive quantum error correction. *Science*, 332:1059–1061, 2011. [209](#)
- [128] M. D. Reed, L. DiCarlo, S. E. Nigg, L. Sun, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf. Realization of three-qubit quantum error correction with superconducting circuits. *Nature*, 482:382–385, 2012. [209](#)
- [129] A. Abragam. *The principles of nuclear magnetism*. Clarendon Press, Oxford, 1961. 25 cm. The International series of monographs on physics. [209](#)
- [130] M. Mariantoni, H. Wang, T. Yamamoto, M. Neeley, R. C. Bialczak, Y. Chen, M. Lenander, E. Lucero, A. D. O’Connell, D. Sank, M. Weides, J. Wenner, Y. Yin, J. Zhao, A. N. Korotkov, A. N. Cleland, and J. M. Martinis. Implementing the quantum von Neumann architecture with superconducting circuits. *Science*, 334:61–65, 2011. [209](#), [212](#)
- [131] J. F. Poyatos, J. I. Cirac, and P. Zoller. Quantum reservoir engineering with laser cooled trapped ions. *Phys. Rev. Lett.*, 77:4728–4731, 1996. [210](#)