

Tileable Superconducting Quantum Circuits with Magnetic Flux Control

Simone Diego Fasciati

Lincoln College

University of Oxford



A thesis submitted for the degree of

Doctor of Philosophy

in Hilary Term 2023

Supervised by Dr. Peter J. Leek

© 2023 Simone Diego Fasciati
All rights reserved.

Contents

Abstract	iv
Acknowledgements	v
1. General introduction and outline	1
2. Theoretical aspects of superconducting quantum circuits	4
2.1. Introduction	4
2.2. Basic superconducting circuits	5
2.2.1. The Josephson junction and the Cooper pair box	5
2.2.2. The transmon qubit	7
2.2.3. Adding flux tunability	8
2.3. The RF-SQUID circuit	10
2.3.1. Superinductance regime ($E_L \ll E_J$)	12
2.3.2. Single well regime ($E_L \geq E_J$)	13
2.3.3. A special case: the quarton ($E_L = E_J$)	20
2.4. Coupled qubits	21
2.4.1. Multi-species entanglement	21
2.4.2. Direct qubit coupling and the static ZZ interaction	22
2.4.3. ZZ suppression techniques	24
2.5. Summary	27
3. Flux control in superconducting quantum circuits	29
3.1. Introduction	29
3.2. Approaches to flux control	30
3.2.1. On-chip flux bias lines	30
3.2.2. Off-chip flux delivery	31
3.3. Flux noise	33
3.3.1. Common sources of noise	33
3.3.2. Experimental estimation of $1/f$ flux noise amplitude	36
3.3.3. Influence of SQUID geometry on flux noise	37
3.3.4. Radiative loss through the FBL	40
3.4. Summary	41

4. Realizing a tileable off-chip flux control architecture	42
4.1. Introduction	42
4.2. The coaxial qubit architecture	42
4.2.1. Adding flux control	43
4.2.2. Gradiometric transmon design	46
4.3. Design and construction of flux-tuning sample enclosure	49
4.3.1. Design and simulation	49
4.3.2. Manufacturing of dFBLs	53
4.4. Cryogenic characterization	56
4.4.1. Continuous-wave spectroscopy	57
4.4.2. Coherence and flux noise	61
4.4.3. Flux crosstalk	65
4.4.4. Rabi drive calibration	69
4.4.5. T_1 spectroscopy	71
4.5. Readout multiplexing	74
4.5.1. Qubit drive calibration with multiplexer	77
4.5.2. T_1 spectroscopy with multiplexer	79
4.6. Summary	81
5. The gradiometric inductively shunted transmon	83
5.1. Introduction	83
5.2. IST design	84
5.2.1. Choosing the inductance range	84
5.2.2. IST geometry and simulation	84
5.2.3. Flux tunability and flux noise	88
5.2.4. Parasitic inductor modes	91
5.2.5. Future design improvements	94
5.3. Experiment overview	95
5.3.1. Fabrication	95
5.3.2. Experiment design	96
5.3.3. Frequency selection	98
5.4. Spectroscopic measurements	99
5.4.1. Resonator spectroscopy	99
5.4.2. Qubit spectroscopy	101

5.5. Time-domain measurements	109
5.5.1. Lifetime and coherence at the flux sweet spot	110
5.5.2. Investigation of flux noise	113
5.6. Summary	115
6. Interactions in a multi-species qubit system	117
6.1. Introduction	117
6.2. Coupled device overview	117
6.3. Static ZZ interaction	120
6.3.1. ZZ interaction vs. flux	122
6.3.2. Simulation of ZZ landscape	124
6.3.3. Implications for scaling	127
6.4. Sideband transitions	129
6.4.1. Concept	129
6.4.2. Device metrics at finite flux bias	130
6.4.3. Driving the transitions	131
6.5. Demonstration of a sideband CZ gate	134
6.5.1. Active ZZ cancellation	134
6.5.2. Single qubit gate benchmarking	136
6.5.3. Sideband pulse calibration	139
6.5.4. CZ gate benchmarking	143
6.6. Summary	146
7. Conclusion and Outlook	147
A. Fabrication	152
A.1. Process overview	152
A.2. Detailed fabrication recipe	154
B. Experimental setup	158
C. Additional datasets	163
Bibliography	167

Abstract

Tileable Superconducting Quantum Circuits with Magnetic Flux Control

Simone Diego Fasciati, Lincoln College, University of Oxford

A thesis submitted for the degree of Doctor of Philosophy

Hilary Term 2023

Superconducting circuits are a leading platform for quantum information processing, partly due to the great freedom of tailoring circuit parameters which enables the implementation of a wide variety of Hamiltonians. Many simple superconducting circuits commonly employed as qubits contain a variation of superconducting quantum interference device (SQUID) that adds in-situ magnetic flux tunability to the system and further increases its flexibility. At the same time, this control parameter requires additional dedicated circuitry and can introduce flux noise which is detrimental to qubit performance. This increases hardware complexity and hinders the scaling to large qubit numbers. In this thesis, we develop and test a simple 3D-integrated architecture for individual flux control of tileable, coaxial, gradiometric superconducting qubits, achieving highly selective flux bias (low crosstalk) and incorporating both charge and flux control into a single off-chip element. The addition of flux tunability fully retains the simplicity of the fabrication and packaging process of the original, fixed-frequency coaxial architecture. We then use this experimental platform to study the inductively shunted transmon (IST), an interesting qubit species based on a radio-frequency (RF) SQUID loop with a small linear shunting inductance. Its qualitative behavior is similar to a transmon qubit but with a positive instead of negative anharmonicity. We design, simulate, fabricate and experimentally characterize gradiometric IST qubits, showing that they can be easily integrated into an existing transmon-based quantum processor architecture. Further, by directly coupling an IST to a transmon via a mutual capacitance, we demonstrate how the opposite signs of anharmonicity can be exploited to effectively reduce the undesired static longitudinal coupling (ZZ interaction) to zero. We also investigate microwave sideband transitions in this two-qubit system and benchmark a controlled-Z (CZ) entangling gate. This work paves the way towards hardware-efficient, crosstalk-suppressed superconducting quantum processors based on multi-species qubit lattices.

Acknowledgements

Completing this thesis would naturally not have been possible without the help and support of many individuals. My sincere gratitude and recognition therefore go to the following people, in a somewhat meaningful order (but please do not use this against me in the future):

First and foremost to Peter, for giving me a shot at doing research and always providing me with deep insights into the mysterious art of experimental science. Your righteousness, integrity and discipline, both as a scientist and as a person, are truly inspirational. The supportive atmosphere you have fostered in the group has enabled me and other students to develop ourselves freely and without pressure.

To my fellow ninja turtles, Giulio, James and Shuxiang, for making what could have been an unbearably long and lonely existence as a doctoral student into a hilarious experience that will not be easily forgotten. Whenever one of us was ready to throw in the towel, there was always someone else ready to pick them up again. To Brian, for providing endless wisdom and ideas, which were sorely needed especially early on in my degree. To Joseph, for teaching me the ropes and introducing me to the fabulous world of flux tuning, which ended up becoming an integral part of this work. To Peter S. and Takahiro, for providing both mentorship and assistance, but mainly for all the comedy. To Salha, Kitti and Martina, for being patient and understanding with a student who didn't quite know yet what they were doing. To Boris and Vivek, for being great colleagues to work with, and great friends to hang out with. To Mustafa, for bringing in a fresh perspective and always keeping us physicists on our toes. To Michele, Mohammed and Sajjad, for putting up with a stressed student trying to wrap up.

Outside the lab, to JJ, Freya, Baptiste and others, for accepting me and deciding to tolerate me in your daily lives. To everyone involved at Lincoln College Boat Club and the various crews I have been part of, for involuntarily convincing me that one of the most boring sports in the world could actually be a lot of fun.

To Angeliki, for being so caring and supportive through the highs and lows of the last years, for keeping me sane and for being understanding of my thesis shenanigans. I would not have completed this without you, so this work belongs to you as well. Further, to all my family and friends who have helped me or encouraged me over the years in one way or another, leading up to this milestone. And finally, to my parents for your unconditional and unwavering support. I don't think you have ever told me once what I should pursue, but instead always encouraged me to do what I wanted. And that is why I dedicate this thesis to you. Grazie.

Ai miei genitori

1.

General introduction and outline

Quantum computation is the manipulation of information using the unitary evolution of a quantum mechanical system [1]. This now established field of research, at the boundary between physics and computer science, started out as a theoretical proposal in the early 1980's to use atoms themselves as units of information to more efficiently simulate quantum physics [2, 3]¹. The development of formal quantum-based algorithms for general computational purposes, predicted to offer a remarkable speedup compared to their classical counterparts, lent significant credibility to this idea [5–7]. This early theoretical groundwork spawned an incredibly active field of experimental research in the decades that followed, focused on performing quantum information processing in a variety of physical platforms [8, 9]. The field is currently at the threshold of *quantum advantage*, where first examples of quantum algorithms at or beyond the reach of classical simulation resources are being reported [10, 11]. Note that this does not necessarily imply that these algorithms are of practical significance just yet.

There are important distinctions that should be made between different approaches to quantum information processing. When talking informally about a *quantum computer*, we often refer to universal gate-based computing using discrete sets of digital operations (gates) on a subset of the quantum bits (qubits) making up the quantum processor [1, 12]. While a lot of platforms fall under this broad umbrella, including the experiments carried out in the present work, fundamentally different approaches such as quantum annealing [13, 14] or analog quantum simulation [15, 16] are also being pursued for specific applications, showing great promise.

It is generally accepted that performing a practically useful quantum computation will require quantum processors orders of magnitude larger and much more robust to errors than current small-to-intermediate scale systems [17]. Ultimately, quantum error correction (QEC) of sufficiently large and high-quality physical systems will

¹Yes, this is yet another thesis in quantum computing that begins with the certainly overused reference to Richard Feynman. See also a related perspective in Ref. 4.

hopefully lead to *fault-tolerant quantum computation*, unlocking the full potential of this technology [18–20].

In this thesis, we narrow our scope and focus on one particular platform for quantum information processing: superconducting circuits [21, 22]. This approach uses dissipationless electrical circuits that exhibit coherent dynamics in the quantum regime, and benefits from established planar nanofabrication techniques to easily build and scale two-dimensional circuits. Currently one of the most promising platforms for quantum computing, superconducting quantum processors hosting on the order of tens to more than hundred qubits have been successfully realized [10, 23–25]. First viable demonstrations of QEC [26, 27] give an indication of the rapid development in this field, while important obstacles to further development lie ahead [28].

Over the course of this work, we will touch upon a few of the main practical challenges of scaling superconducting quantum circuits and provide a perspective on existing approaches and solutions to tackling them. We will then set out to address some of these challenges in our own experimental architecture, developing it further and using it to explore novel types of superconducting quantum circuits.

Thesis outline

This thesis is structured as follows. In [Chapter 2](#) we provide an overview of simple superconducting electrical circuits commonly used as qubits in quantum information processing. We analyze one such circuit in more detail, the inductively shunted transmon (IST). We also touch upon two-qubit dynamics and the issue of static ZZ-type interaction that can limit quantum processor performance, as well as methods for its suppression.

We continue with a discussion of the practical aspects of magnetic flux control of superconducting qubits in [Chapter 3](#), an important ingredient in many circuit realizations. This includes design considerations of the flux biasing elements and of the qubits themselves, and explains the issue of magnetic flux noise and how it can be mitigated.

Moving on to the experimental part of this work, we address some of the challenges of flux control with the development of a novel 3D-integrated flux tuning architecture for circuit QED in [Chapter 4](#). Its central element is the differential flux bias line (dFBL), which can provide full individual XYZ control of coaxial gradiometric qubits. Design, numerical simulations and cryogenic characterization are carried out, using a test device with four tuneable transmons.

We then proceed to the design, fabrication and measurements of IST qubits in our dFBL platform to explore their physics. This is reported in [Chapter 5](#), showing that these circuits behave as expected from theory and are fully compatible with our existing circuit QED architecture.

Building upon these results, we study the two-qubit dynamics of a coupled transmon-IST circuit in [Chapter 6](#). The system once again behaves as expected from theory, showing an interesting landscape of static ZZ interaction which includes an operating point with full ZZ suppression. We also identify sideband transitions in this system, and use one such transition to tune up and benchmark an entangling gate between the two qubits.

To conclude, our findings are summarized in [Chapter 7](#) and we provide an outlook on further exploratory studies in this direction.

2.

Theoretical aspects of superconducting quantum circuits

2.1. Introduction

We begin with an important disclaimer: the aim of this chapter is *not* to offer a general and complete introduction to quantum computing with superconducting circuits. The reason is twofold: firstly, it would not be possible to describe in this limited space all the important milestones achieved since coherent quantum oscillations were first observed in a superconducting Cooper-pair box (CPB) circuit almost 25 years ago [29], which is often viewed as the starting point for this experimental field of research. Secondly, a number of accessible and detailed reviews have already been written on the subject (or parts of it) from different perspectives, ranging from the purely theoretical to the explicitly practical, which should satisfy readers with a variety of backgrounds [21, 22, 30–33]. What this chapter *does* attempt to do is instead to provide a few useful concepts and tools that will be required to understand the ideas developed later in this work.

In the following we describe the prototypical superconducting charge-based qubit, i.e. the CPB mentioned above. This is but one species in an ever-expanding *qubit zoo*, made possible by the ease with which the Hamiltonian of a superconducting circuit can be tailored to fit different needs and realize a variety of parameter regimes. The fundamental building block of all these circuits is the Josephson junction, a tunneling barrier between two superconductors which acts as a non-linear inductor and gives rise to many interesting phenomena. Simple modifications to the Cooper-pair box can make the non-linear inductance tunable by magnetic flux (using a SQUID, or superconducting quantum interference device), or introduce various degrees of protection from noise sources such as fluctuations of the charge degree of freedom.

As a further important example we discuss in more detail the RF-SQUID (radio-frequency SQUID) circuit consisting of a Josephson junction connected in parallel with a linear inductor. This can be used to implement a variety of flux-based qubits

with fundamentally different properties to charge qubits. We perform perturbative and numerical analysis of the RF-SQUID in the low-inductance regime, i.e. where the ratio of inductive energy E_L to Josephson energy E_J is larger than one, to obtain a relatively unexplored parameter regime that is simple to understand and is potentially useful as a qubit.

Next, we present a few basic concepts of two-qubit coupling schemes and entangling gates, which are all important considerations when trying to build an actual quantum processor out of individual qubits. We focus in particular on the aspect of unwanted coupling, or crosstalk, in the form of longitudinal ZZ-type coupling between qubits, which is often a limiting factor for the fidelity of two-qubit operations in superconducting circuits. Several techniques to mitigate or suppress ZZ crosstalk have been developed in a variety of two-qubit coupling schemes. This provides the context for some of the experimental studies carried out in this work, which will investigate one such method for ZZ suppression with minimal impact on the hardware and software complexity of the quantum circuit.

2.2. Basic superconducting circuits

2.2.1. The Josephson junction and the Cooper pair box

In superconducting quantum circuits, one particular phenomenon lies at the heart of much of the interesting physics that we observe: the *Josephson effect* [34, 35], which describes the behavior of two superconductors separated by a thin insulating barrier. The superconductors are each associated with a macroscopic quantum phase ϕ_1 and ϕ_2 respectively, and their difference is $\varphi = \phi_1 - \phi_2$. Such an element is called a *Josephson junction* (JJ), and the current I and voltage V across it evolve according to the Josephson relations [36]

$$I(t) = I_c \sin(\varphi), \quad V(t) = \frac{\hbar}{2e} \frac{\partial \varphi}{\partial t} = \frac{\Phi_0}{2\pi} \frac{\partial \varphi}{\partial t}, \quad (2.1)$$

where $\Phi_0 = h/2e \approx 2.068 \times 10^{-15}$ Wb is the magnetic flux quantum and I_c is the *critical current* of the JJ, which is related to its physical properties such as dimensions and materials. The energy scale associated with this is the Josephson energy $E_J = (\Phi_0/2\pi)I_c$. Furthermore, a corresponding inductance scale can be deduced by rewriting and combining the Josephson relation to obtain

$$V = L(\varphi) \frac{\partial I}{\partial t}, \quad L(\varphi) = \frac{L_J}{\cos(\varphi)}. \quad (2.2)$$




Figure 2.1: Basic superconducting circuits. (a) Josephson junction (JJ) element, characterized by the Josephson energy E_J and intrinsic charging energy E_{C_J} , usually written as a combined element. (b) Simple Cooper pair box (CPB) circuit, consisting of a superconducting island weakly linked to a reservoir via a JJ. The charge on the island can be controlled via an external gate voltage. (c) Standard transmon circuit, where a JJ is shunted by a large capacitance. (d) Replacing the JJ by a DC-SQUID makes the Josephson energy tunable via an external magnetic flux Φ_{ext} .

This behavior mimics that of an inductor, but with an effective inductance that depends non-linearly on the phase φ , and with a characteristic *Josephson inductance* $L_J = (\Phi_0/2\pi)/I_c$. This is a kinetic inductance originating from the kinetic energy of Cooper pairs in the superconductor. The JJ is therefore often referred to as a non-linear inductive element and is usually represented in circuits with a cross symbol, see Fig. 2.1 (a). In practical realizations, the proximity of the two superconductors will also lead to a small capacitance C_J in parallel with L_J , with corresponding energy scale $E_{C_J} = e^2/(2C_J)$. The combined element is represented as a cross inside a square.

A single JJ is already enough to build a useful qubit, i.e. a system with a pair of energy levels that can be addressed individually. An isolated superconducting island is weakly connected to a large superconducting bath via a JJ. This is called a *Cooper pair box* (CPB) and is depicted in Fig. 2.1 (b). In the typical case $E_C \gg E_J$, adding an extra Cooper pair on the island by tunneling it across the JJ costs significant charging energy, and therefore the number N of excess Cooper pairs in the CPB at low temperatures is usually close to zero. The charging energy can be overcome by applying an offset DC voltage V_g , which will lead to a tunable amount of offset charge $N_g = C_g V_g / 2e$ on the island. The quantum Hamiltonian of the CPB is the sum of the charge and Josephson term and reads

$$\mathcal{H} = 4E_C(\hat{N} - N_g)^2 - E_J \cos(\hat{\varphi}), \quad (2.3)$$

with the charge and phase operators obeying the canonical commutation relation $[\hat{\varphi}, \hat{N}] = i$. Here the charging energy includes all capacitive contributions, $C = C_J + C_g$. The eigenstates are commonly represented in the charge basis, i.e. $|0\rangle, |1\rangle, \dots$ corresponding to the number of discrete excess Cooper pairs on the

island. One finds that at a half integer value of the offset charge, $N_g = 0.5$, the states $|0\rangle$ and $|1\rangle$ are degenerate. The small finite tunneling probability through the JJ couples these states and leads to a splitting of magnitude E_J . Since all other charge states are separated by at least $E_C \gg E_J$, we have thus created a single addressable transition – a qubit.

2.2.2. The transmon qubit

The CPB qubit is one of the simplest superconducting qubits and indeed one of the earliest to be implemented experimentally [29, 37], and was utilized in several important early experiments in the field [38, 39]. Unfortunately, encoding the qubit in the charge degree of freedom makes it highly sensitive to the electrostatic environment, which can fluctuate in time due to the presence of charged defects and therefore adds considerable noise to the qubit mode, limiting its coherence. This prompted the shift towards superconducting qubits based instead on flux encoding (more on this in Sec. 2.3), which are exposed to magnetic flux noise but not to charge noise. On the other hand, a simple modification of the CPB also alleviates the charge noise problem. Adding a large shunting capacitance C in parallel with a JJ, as depicted in Fig. 2.1 (c), makes the total charging energy E_C very small and achieves the regime $E_J \gg E_C$. The enhanced tunneling of Cooper pairs smears the eigenstates across many charge states simultaneously, therefore making their eigenenergies approximately independent of the changes in offset charge N_g . This is the basic operating principle of the *transmon* qubit [40]. The Hamiltonian is the same as the CPB in Eq. (2.3), but with the offset charge control N_g removed as it is not anymore a useful tuning knob for the qubit energy. Instead, the large shunting capacitance makes it easy to couple the macroscopic plasma oscillations of the transmon mode to linear microwave cavities and external electromagnetic fields (usually in the range of a few GHz). This is used to good effect for both qubit control and readout, as well as qubit-qubit coupling, in the now very popular framework of circuit quantum electrodynamics (circuit QED) [39, 41, 42].

The basic physics of the transmon [40] can be understood intuitively by making an analogy to a mechanical oscillator, where a particle of mass m is trapped in a potential (e.g. due to a spring with constant k). In the electrical oscillator version, i.e. the LC circuit, the mass by convention corresponds to the capacitance C , and the spring constant to the inductance L . The transmon is therefore referred to as a *heavy* charge mode due to its very large value of C . This heavy particle oscillates around the minimum of the cosinusoidal phase potential provided by the inductive

term $U(\hat{\varphi}) = -E_J \cos(\hat{\varphi})$ of the Josephson junction. Expanding the potential for small phase excursions around zero gives us

$$U(\hat{\varphi}) \approx -E_J + \frac{E_J}{2}\hat{\varphi}^2 - \frac{E_J}{24}\hat{\varphi}^4, \quad (2.4)$$

which means that the heavy charge particle sees an approximately quadratic (harmonic) potential with a negative quartic correction. The transmon Hamiltonian is therefore often approximated as that of a quantum non-linear (Duffing) oscillator, written in second quantization form as [40, 43]

$$\mathcal{H}_{tr} \approx \hbar\omega_0 \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) + \frac{\alpha}{2} \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a}, \quad (2.5)$$

with the creation and annihilation operators \hat{a}^\dagger and \hat{a} that either add or remove an excitation in the system, $\omega_0 = \sqrt{8E_C E_J}$ is the plasma frequency, and $\alpha \approx -E_C$ is the non-linearity or *anharmonicity* of the circuit. Thus, the transmon is a weakly anharmonic oscillator with a first transition frequency on the order of ω_0 , spectrally separated from the nearest transition by a modest anharmonicity $\alpha < \omega_0$. In addition to the protection from charge noise mentioned above, the simple transmon also does not suffer from magnetic flux noise as it does not couple to external magnetic fields.

It is difficult to overstate how impactful the invention of the transmon qubit has been. Its simplicity and resilience to noise has led to the rapid scaling of quantum processors to a few tens of qubits [10, 23, 26] and recently even to over 100 qubits [25]. Refinements in materials and in circuit design have led to transmons achieving relaxation times T_1 and coherence times T_2 of a few 100 μs [44–48].

One of the main drawbacks of the transmon is its anharmonicity α , which is necessarily small in the charge-insensitive regime $E_J \gg E_C$ [40]. This leads to higher excited states lying close in frequency to the qubit transition, causing leakage of information into these states. While this can be elegantly used to execute two-qubit entangling gates that leverage coupling to higher states [49, 50], it is generally considered a disadvantage as it limits the speed of qubit manipulations and requires careful pulse shaping to minimize leakage [51, 52].

2.2.3. Adding flux tunability

Although both CPB and transmon can be realized using a single Josephson junction, a simple modification makes them much more flexible and controllable. This consists

of replacing the JJ by a SQUID (superconducting quantum interference device), a closed superconducting loop interrupted by two JJs with energy E_{J_1} and E_{J_2} respectively, as shown in Fig. 2.1 (d). The two junctions are not necessarily identical. Assuming $E_{J_1} \geq E_{J_2}$, we define the asymmetry $d = (E_{J_1} - E_{J_2})/E_{J,tot}$ with $E_{J,tot} = E_{J_1} + E_{J_2}$. The Hamiltonian now has two separate inductive contributions,

$$\mathcal{H} = 4E_C \hat{N}^2 - E_{J_1} \cos(\hat{\varphi}_1) - E_{J_2} \cos(\hat{\varphi}_2), \quad (2.6)$$

where we denote the phase operator across each junction as $\hat{\varphi}_1$ and $\hat{\varphi}_2$. The presence of a closed loop allows a finite amount of magnetic flux Φ_{ext} to be threaded through it by applying an external magnetic field perpendicular to the SQUID plane. Due to the fundamental requirement that flux through a superconducting ring must be quantized to a multiple of Φ_0 , the external flux will be compensated by an induced supercurrent that will in turn current-bias the two JJs [36]. The flux quantization condition reads

$$\hat{\varphi}_1 - \hat{\varphi}_2 + \varphi_{\text{ext}} = 2\pi n, \quad \varphi_{\text{ext}} = 2\pi \frac{\Phi_{\text{ext}}}{\Phi_0}, \quad (2.7)$$

with $n \in \mathbb{N}$, and we have introduced the reduced external flux φ_{ext} for convenience. This constraint allows us to eliminate one of the two phase variables, introducing the average phase drop $\hat{\varphi} = (\hat{\varphi}_1 + \hat{\varphi}_2)/2$ and rewriting Eq. (2.6) as [40]

$$\mathcal{H} = 4E_C \hat{N}^2 - E_J(\varphi_{\text{ext}}) \cos(\hat{\varphi}), \quad (2.8)$$

which is reduced to the same functional Hamiltonian as the single-junction case, but with a flux-dependent Josephson energy

$$E_J(\varphi_{\text{ext}}) = E_{J,tot} \times \left| \cos\left(\frac{\varphi_{\text{ext}}}{2}\right) \right| \times \sqrt{1 + d^2 \tan^2\left(\frac{\varphi_{\text{ext}}}{2}\right)}. \quad (2.9)$$

This makes the transmon (or CPB) transition frequency tunable in-situ by changing the amount of applied external flux, without affecting its underlying physics. This is in fact a very common technique used experimentally [29, 53, 54] to add more freedom and flexibility but also to make up for shortcomings in fabrication, which limit the accuracy and precision with which E_J of an individual junction can be set. The major drawback is that this makes the qubit sensitive to noise in the flux control parameter, similar to the case of flux-based qubits discussed later. Fluctuations in flux will modify the qubit frequency, leading to dephasing. The

practicalities of flux tuning and flux noise will be discussed in greater detail in [Chapter 3](#).

The inductive energy in [Eq. \(2.9\)](#) is periodic in applied flux, as expected from flux quantization. In case of symmetric junctions with $d = 0$, the expression simplifies to a simple cosine dependence $E_J(\varphi_{\text{ext}}) = E_{J,\text{tot}} |\cos(\varphi_{\text{ext}}/2)|$, meaning that the Josephson energy can be brought to zero at half-integer values of flux bias. Instead, for finite asymmetry the square-root term reduces the total tunability range to between $E_{J,\text{tot}}$ and $E_{J1} - E_{J2} > 0$. This is sometimes done intentionally to provide only the tunability necessary for a particular experiment while limiting the amount of dephasing due to flux noise to acceptable levels [[55](#), [56](#)].

2.3. The RF-SQUID circuit

The RF-SQUID (radio-frequency SQUID) is a circuit consisting of a closed superconducting loop with self-inductance L interrupted at one location by a Josephson junction. Three energy scales define the dynamics of this circuit: the inductive energy $E_L = (\Phi_0/2\pi)/L$ coming from the linear inductance L of the ring, the Josephson energy E_J of the JJ, and a charging energy $E_C = e^2/(2C)$ associated with the total capacitive contribution C in the circuit, such as the JJ capacitance or any additional shunting capacitance. The equivalent circuit representation is shown in [Fig. 2.2](#) (a). The RF-SQUID thus incorporates the three basic elements of a dissipationless circuit: linear inductance, non-linear Josephson inductance, and capacitance. The quantum Hamiltonian that describes this circuit is given by

$$\mathcal{H}_0 = 4E_C \hat{N}^2 - E_J \cos(\hat{\varphi} + \varphi_{\text{ext}}) + \frac{1}{2} E_L \hat{\varphi}^2, \quad (2.10)$$

where \hat{N} and $\hat{\varphi}$ are the conjugate operators for charge number and phase, and φ_{ext} is the reduced external flux through the loop, as before. Most superconducting qubits can be treated as a special case or slight modification of the generic RF-SQUID circuit¹ [[21](#), [58](#)]. Even the physics of the CPB described above, which is topologically distinct due to the absence of a closed loop, can be recovered in the limit of an infinitely large inductor $L \rightarrow \infty$ [[59](#)].

One feature of the RF-SQUID loop is that there is no isolated superconducting island that can hold a discrete number of Cooper pairs, as in the CPB case. The linear inductive element has zero DC impedance, and any excess charge that might

¹This statement is limited to circuits that host a single electromagnetic mode. Multi-mode qubits are a separate topic offering very rich physics [[57](#)].




Figure 2.2: RF-SQUID and related circuits. (a) General RF-SQUID circuit consisting of a linear inductance, Josephson junction and capacitance in parallel. (b) Conventional three-junction flux qubit. A small JJ is shunted by two larger JJs, with an energy ratio $\alpha = E_J/E'_J < 1$. (c) Fluxonium circuit, where a large number N of JJs forms an effectively linear inductive shunt with $E_L \approx E'_J/N \ll E_J$.

form on one side of the JJ is easily equalized via this link. This automatically removes any static charge offsets, and modifies the charge noise spectrum to become linear in frequency, as opposed to typical $1/f$ -type charge noise [59]. This modified noise spectrum leads to a strongly reduced impact on qubit coherence. Note that this is fundamentally different from the approach adopted in the transmon [40], where DC offset charges are still very much present but the qubit mode has been made insensitive to them by flattening its charge dispersion bands at the expense of decreased circuit anharmonicity.

In practical implementations, the linear inductance L is often approximated by a discrete number of JJs in series. While this breaks the loop into discrete sections of superconducting islands and adds independent phase variables to the Hamiltonian, the effective low-energy behavior is still similar to that of a linear inductive shunt. As discussed in Sec. 2.2.3, two JJs in parallel form a DC-SQUID but do not fundamentally modify the Hamiltonian of the circuit. The smallest number of junctions required in a loop to introduce new non-trivial effects is therefore three, as in the circuit depicted in Fig. 2.2 (b). Typically, a small junction with low E_J is connected in parallel with two larger JJs. The ratio $\alpha = E_J/E'_J < 1$ (not to be confused with the anharmonicity) is a crucial parameter of the system. This circuit realizes the conventional *flux qubit* [60–62], an important early superconducting qubit implementation. The closed loop can be biased to induce supercurrents in either direction, clockwise or counter-clockwise. Just as two charge states are made degenerate in the CPB by biasing its island to a half charge offset, in the flux qubit two opposite current (or flux) states are made degenerate by biasing the loop to half a flux quantum, with a splitting on the order of $\Delta \sim E_J$.

The development of the flux qubit was motivated at least partly by the realization that charge-based qubits inevitably suffered from large fluctuations of their electrostatic environments over time, while the flux qubit encodes information in

supercurrents flowing around its loop and is therefore only exposed to the magnetic environment. It turns out that magnetic flux noise also represents a significant obstacle (see [Chapter 3](#)), and the advent of the transmon shifted the focus once more towards charge qubits. Nevertheless, derivatives of the original flux qubit have shown great promise in recent years and are actively being pursued, such as the capacitively shunted flux qubit (CSFQ) [63] and the fluxonium [64, 65].

2.3.1. Superinductance regime ($E_L \ll E_J$)

An important evolution of the flux qubit is the addition of an increasing number N of large-area Josephson junctions in series to form a long inductive array, depicted in [Fig. 2.2 \(c\)](#). Values of $N \sim 100$ or larger are not uncommon. Although each junction is associated with an individual phase variable $\hat{\varphi}_i$, a one-dimensional approximation can reduce this to an effective RF-SQUID Hamiltonian as in [Eq. \(2.10\)](#) with an equivalent linear inductive energy $E_L = E'_J/N \ll E_J$ much smaller than the energy of the small JJ. The array is referred to as a *superinductance* due to its very large kinetic inductance, with a characteristic impedance that typically surpasses the resistance quantum $R_Q = \hbar/(2e)^2$. This type of circuit is known as the *fluxonium* [64], which was created in an attempt to solve both the charge noise problem of the CPB and the flux noise problem of the conventional flux qubit. Although the loop is broken up into discrete superconducting sections between each JJ, the large E'_J makes the junctions fairly transparent to Cooper pairs and allow for an efficient equalization of charge offsets, which takes care of reducing charge noise.

While a detailed description of the rich fluxonium physics is beyond the scope of this work and can be found elsewhere [65–70], some intuition for its behavior can be gained by observing that in the regime $E_L \ll E_J$ the quadratic phase potential term in [Eq. \(2.10\)](#) is very weak compared to the cosine term of the small JJ. The effective potential $U(\hat{\varphi})$ therefore looks like a shallow corrugated parabola with many local minima. Depending on the details of the circuit and on the choice of external flux bias, the two lowest eigenstates will either be located in the same potential well (*plasmon*-like excitation, as in the transmon), or in two adjacent wells (*fluxon*-like excitation, as in the flux qubit). Which and how many of the potential wells are populated will depend, among other things, on the ratio E_J/E_C that defines how light or heavy the circuit mode is, and therefore how easily it can explore the potential landscape, analogously to the transmon discussed in [Sec. 2.2.2](#). The different energy scales associated with fluxon and plasmon transitions can lead

to an extremely high anharmonicity, several times larger than the qubit frequency transition itself, making the fluxonium a prime candidate as a replacement for the weakly anharmonic transmon. It has also been shown to reach relaxation and coherence times beyond 1 millisecond [71], pushing the state of the art for JJ-based superconducting qubits.

Recently, the engineering of a long suspended JJ chain with a reduced parasitic capacitance contribution has enabled the realization of a circuit where both $E_L \ll E_J$ and $E_L \ll E_C$, opening up a regime beyond the fluxonium called the quasicharge qubit (or *Blochonium*) [72]. Separately, a different substrate engineering technique has been used to fabricate qubits in the superinductance regime using a purely geometric inductor [73, 74], instead of requiring a kinetic inductance chain as described above.

2.3.2. Single well regime ($E_L \geq E_J$)

A very different scenario presents itself if we make E_L comparable to or larger than E_J , meaning the shunting inductance becomes small. If $E_L \gg E_J$, the phase potential is characterized by a strong harmonic confinement $\sim E_L \hat{\varphi}^2/2$, with only a small perturbation due to the Josephson cosine term. This leads to an approximately linear oscillator with negligible anharmonicity, which is not relevant for qubit encoding. Still, this limit might be of interest for applications such as the study of strongly driven systems and Floquet dynamics [75, 76], as the parabolic confinement of the phase prevents the circuit from escaping into unconfined states like in the transmon case [77, 78].

Since we are primarily interested in using this circuit as a novel type of qubit, requiring a finite anharmonicity, we focus here on the case $E_L \gtrsim E_J$, when neither of the potential energy terms can be regarded as a small perturbation to the other. We begin by studying the physics of this low-inductance RF-SQUID at the half flux bias point $\varphi_{\text{ext}} = \pi$. We expand the phase potential around zero and obtain

$$U(\hat{\varphi}) = +E_J + \frac{E_L - E_J}{2} \hat{\varphi}^2 + \frac{E_J}{24} \hat{\varphi}^4 + \mathcal{O}(\hat{\varphi}^6) \approx (\beta - 1) \frac{E_J}{2} \hat{\varphi}^2 + \frac{E_J}{24} \hat{\varphi}^4, \quad (2.11)$$

where $\mathcal{O}(\hat{\varphi}^6)$ collects all terms of order 6 and higher in $\hat{\varphi}$, and we drop the constant term $+E_J$. We define the ratio of inductances $\beta = E_L/E_J$ for convenience. This expression has the same Duffing oscillator form as the transmon in Eq. (2.4) but with the quadratic term rescaled by a prefactor $\beta - 1$ and the quartic term with




Figure 2.3: Phase potential of the low-inductance RF-SQUID. Phase potential $U(\varphi)$ at $\varphi_{\text{ext}}/2\pi = 0.5$ for three different inductive energy ratios $\beta = E_L/E_J$: (a) weakly anharmonic regime $\beta = 2$, (b) quartic regime $\beta = 1$, and (c) double-well regime $\beta = 0.5$.

positive instead of negative sign². The curvature of this potential is strictly positive for $\beta > 1$, leading to a single potential well and to plasmonic excitations as in the transmon. We conventionally define a qubit as the subsystem comprising the two lowest eigenstates of this system.

In more detail, we distinguish between the following cases:

- ($\beta > 2$) A large quadratic term weakens the non-linearity of the potential compared to the transmon, resulting in smaller α .
- ($\beta = 2$) This case is qualitatively equivalent to the transmon potential, but with positive non-linearity.
- ($2 > \beta > 1$) As the relative contribution of the quadratic term is weakened, the non-linearity of the circuit increases, which results in larger α .
- ($\beta = 1$) The quadratic term disappears and the problem is reduced to that of a particle in a quartic potential. This is the so-called *quarton* regime [58], which will be treated separately in Sec. 2.3.3.
- ($\beta < 1$) The quadratic term is now negative, resulting in a double-well phase potential. This is analogous to the persistent-current flux qubit [60, 61] discussed earlier, where the lowest eigenstates are fluxon states localized in either of the two potential wells, associated with currents flowing in opposite directions around the circuit loop.

We visualize the three qualitatively different flavours of phase potential $U(\varphi)$ in Fig. 2.3. While the double-well potential is a very interesting regime in its own

²One can already predict that the opposite sign of the oscillator non-linearity flips the sign of anharmonicity, which will be confirmed later.

right, we are going to limit ourselves to the description of the plasmonic single-well physics in the rest of this section.

Perturbation theory

Approximate analytical expressions for the lowest transition frequency ω_{01} and anharmonicity α can be derived for $\beta > 1$ via perturbation theory starting from the potential in Eq. (2.11), following the same approach as for the transmon outlined in Refs. 40, 43. In this case, the quadratic term is positive and we can solve the harmonic Hamiltonian

$$\mathcal{H}_{\text{HO}} = 4E_C \hat{N}^2 + (\beta - 1) \frac{E_J}{2} \hat{\varphi}^2 \quad (2.12)$$

in the conventional way by defining bosonic ladder operators \hat{a}, \hat{a}^\dagger satisfying $[\hat{a}, \hat{a}^\dagger] = 1$ via the relations

$$\hat{\varphi} = \frac{\varphi_{\text{zpf}}}{\sqrt{2}} (\hat{a}^\dagger + \hat{a}), \quad \hat{N} = \frac{i}{\sqrt{2}\varphi_{\text{zpf}}} (\hat{a}^\dagger - \hat{a}), \quad (2.13)$$

where $\varphi_{\text{zpf}} > 0$ is a normalization factor that is related to the zero-point fluctuations of the phase³. By setting $\varphi_{\text{zpf}} = [8E_C/(\beta - 1)E_J]^{1/4}$ and expressing Eq. (2.12) in terms of ladder operators we obtain

$$\mathcal{H}_{\text{HO}} = \hbar\omega_0 \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right), \quad (2.14)$$

where we defined the plasma frequency $\omega_0 = \sqrt{8(\beta - 1)E_J E_C}$. The eigenstates are the Fock states $|n\rangle$ for $n = 0, 1, 2, \dots$ with eigenenergies $E_n^{(0)} = \hbar\omega_0 (n + 1/2)$.

We now treat the quartic term of $U(\hat{\varphi})$ as a perturbation

$$\mathcal{H}_p = \frac{E_C}{12(\beta - 1)} (\hat{a}^\dagger + \hat{a})^4 \quad (2.15)$$

to \mathcal{H}_{HO} and calculate the first-order perturbative correction $E_n^{(1)}$ to the eigenenergies of the system as [40]

$$E_n^{(1)} = \langle n | \mathcal{H}_p | n \rangle = \frac{E_C}{12(\beta - 1)} (6n^2 + 6n + 3). \quad (2.16)$$

³Equivalently, we can also define the zero-point fluctuations of charge as $N_{\text{zpf}} = 1/\varphi_{\text{zpf}}$ via the canonical commutation relation $[\hat{\varphi}, \hat{N}] = i$.

The new eigenenergies $E_n \approx E_n^{(0)} + E_n^{(1)}$ allow us to obtain perturbative expressions for the lowest transition frequency ω_{01} and anharmonicity α of the non-linear oscillator, defined as follows:

$$\omega_{01} = E_1 - E_0 \approx \sqrt{8(\beta - 1)E_C E_J} + \frac{E_C}{\beta - 1}, \quad \alpha = (E_2 - E_1) - (E_1 - E_0) \approx +\frac{E_C}{\beta - 1}. \quad (2.17)$$

This is a simple result that nevertheless captures many important concepts. In Fig. 2.4 (a) and (b) we plot these expressions for ω_{01} and α . When $\beta = 2$ we recover the transmon-like case with $\omega_{01} = \sqrt{8E_C E_J} + \alpha$ and $|\alpha| = E_C$, but with positive instead of negative anharmonicity. The anharmonicity decreases as β increases, reaching the harmonic oscillator limit when $\beta \rightarrow \infty$. Most interestingly, we can achieve values of α significantly larger than E_C for $1 < \beta < 2$, going beyond the limitations of the transmon.

While Eq. (2.17) suggests that both ω_{01} and α diverge for $\beta \rightarrow 1$, it is important to keep in mind that the perturbation theory outlined above breaks down when the quadratic term becomes comparable to or smaller than the quartic term. Another way to realize this is to observe that the magnitude of phase fluctuations φ_{zpf} diverges when $\beta \rightarrow 1$, meaning that the Taylor expansion around small phase values in Eq. (2.11) is not justified. Values of β approaching 1 therefore require either the inclusion of many higher order terms [43] or a numerical treatment in order to obtain reliable values for ω_{01} and α , as discussed further below.

An important figure of merit in a system used for qubit encoding is the relative spectral separation between the qubit transition of interest and the next higher transition of the system, quantified by the relative anharmonicity $\alpha_r = |\alpha/\omega_{01}|$. This ratio tells us how fast the qubit can be driven relative to its own frequency without exciting unwanted transitions. A large ratio is beneficial for performing fast operations. From Eq. (2.17) we directly obtain

$$\alpha_r \approx \left(\sqrt{8 \frac{E_J}{E_C} (\beta - 1)^3 + 1} \right)^{-1}, \quad (2.18)$$

which we plot in Fig. 2.4 (c) for a fixed ratio $E_J/E_C = 60$. We also include the value $\alpha_r \sim 5\%$ of an equivalent transmon for comparison. Perturbation theory significantly overshoots the more accurate numerical solution (see details further below) when $\beta \rightarrow 1$, but it correctly predicts that α_r increases monotonically and surpasses the transmon.

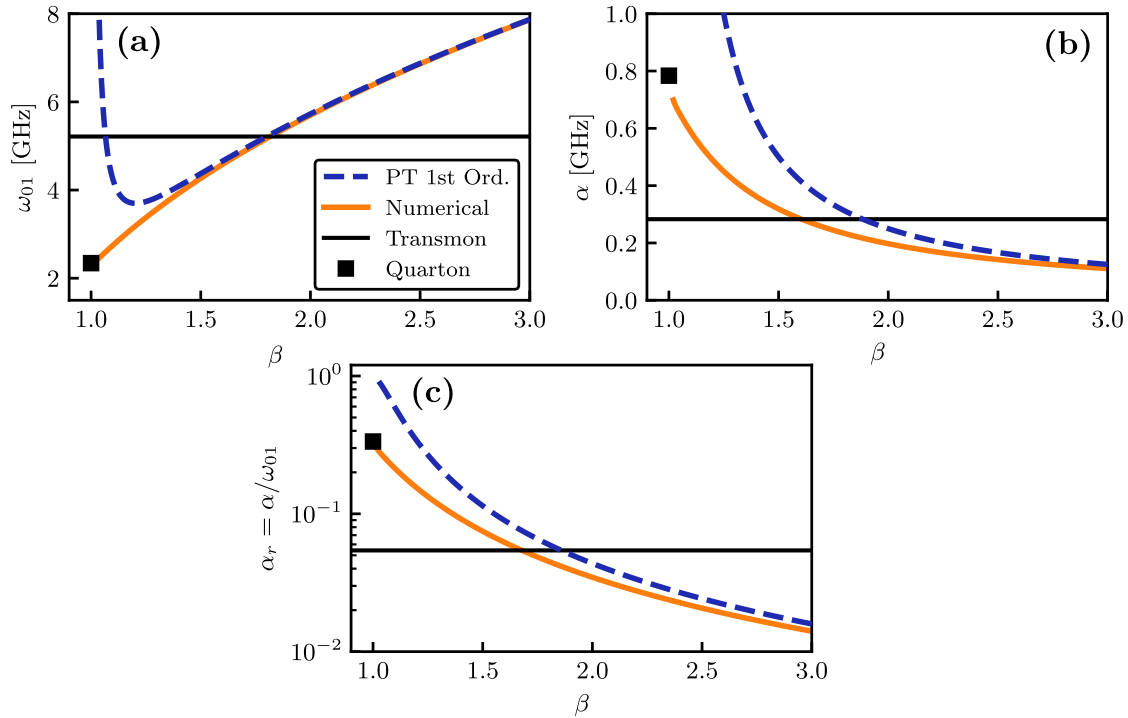


Figure 2.4: IST qubit frequency and anharmonicity at half flux bias. Theoretical predictions for (a) qubit frequency, (b) anharmonicity, and (c) relative anharmonicity of an inductively shunted transmon (IST) circuit biased at $\varphi_{\text{ext}} = 0.5$. Perturbation theory and numerical diagonalization results are shown in dashed blue and solid orange lines, respectively. We fix $E_C = 250$ MHz and $E_J = 60 \times E_C$, and sweep the inductance ratio β . The numerical result approaches the predicted quarton values (squares) when $\beta \rightarrow 1$. The reference values for a transmon with the same E_J and E_C , also obtained numerically, are added as black lines for comparison.

Due to the similarity of the perturbation formulas in Eq. (2.17) to the transmon case, as well as a relevant range of values for E_J and E_C similar to existing transmons, we refer to the RF-SQUID circuit with $\beta > 1$ presented here as the *inductively shunted transmon* (IST), which is the name given to a recent experimental realization of this type of circuit [79]⁴. Broadly speaking, the properties of the IST circuit are also reminiscent of the C-shunted flux qubit (CSFQ) [63, 81], which is a flux qubit with a small E_C and a typical junction ratio $\alpha < 0.5$ that implements a single-well potential and a plasmonic excitation spectrum with pronounced positive anharmonicity. This is very similar to our circuit with $1 < \beta < 2$. On the other hand, we should note that this similarity hides a fundamental difference between these circuits – namely that the CSFQ, like the transmon, is an island-based circuit

⁴This is not to be confused with another recent realization of RF-SQUID circuit that is also called IST, but which employs a large geometric superinductance to achieve the regime $\beta \ll 1$ [80].

and will therefore be governed by different charge dynamics compared to a true RF-SQUID loop .

For the sake of completeness, we repeat the above perturbative derivation for the IST circuit now biased at $\varphi_{\text{ext}} = 0$ and summarize the result here. The potential is

$$U(\hat{\varphi}) \approx (\beta + 1) \frac{E_J}{2} \hat{\varphi}^2 - \frac{E_J}{24} \hat{\varphi}^4, \quad (2.19)$$

leading to the perturbative expressions

$$\omega_{01} \approx \sqrt{(\beta + 1) 8E_J E_C} - \frac{E_C}{\beta + 1}, \quad \alpha \approx -\frac{E_C}{\beta + 1}. \quad (2.20)$$

The two main differences to Eq. (2.17) are the new rescaling factor $\beta + 1$ and the flipped sign of anharmonicity, which is now negative and always smaller than E_C for arbitrary values of β . One interpretation is that at zero applied flux the two inductive terms are added instead of subtracted, giving an effective linear inductance

$$L_{\text{eff}} = \frac{\Phi_0/2\pi}{(\beta + 1)E_J} = \frac{\Phi_0/2\pi}{E_L + E_J} = \left(\frac{1}{L} + \frac{1}{L_J} \right)^{-1}, \quad (2.21)$$

i.e. the Josephson inductance is diluted by the linear inductance. This leads to an enhanced quadratic term irrespective of the value of β and therefore to a reduced anharmonicity compared to a transmon with the same E_J and E_C . This makes the $\varphi_{\text{ext}} = 0$ working point relatively uninteresting from a practical perspective.

Lastly, by taking the ratio of the two plasma frequencies (zeroth order terms of ω_{01}) in Eq. (2.17) and Eq. (2.20) we conclude that the frequency at zero applied flux should be larger than the value at half flux, since $\sqrt{(\beta + 1)/(\beta - 1)} > 1$.

Numerical treatment

Having built some intuition for the physics of the IST circuit we now turn to numerical solutions and compare these to the perturbation theory above. We use the Python library *QuTiP* [82] to numerically diagonalize the time-independent RF-SQUID Hamiltonian in Eq. (2.10) and calculate the eigenspectrum $\{E_n\}$ as a function of E_C, E_L, E_J and φ_{ext} , from which we obtain the qubit frequency ω_{01} and anharmonicity α in terms of the eigenenergies $\{E_n\}$ as defined in Eq. (2.17).

For $\varphi_{\text{ext}} = 0.5$, we add the numerical result to Fig. 2.4 and observe that there is reasonable agreement with the perturbation theory for weak non-linearity (large β), as expected. While the PT expressions for ω_{01} and α both diverge when β approaches 1, giving an unphysical result, the numerics provide finite values and




Figure 2.5: IST qubit frequency and anharmonicity versus flux. Numerical prediction for qubit frequency (left panel) and anharmonicity (right panel) versus applied reduced flux φ_{ext} . We fix $E_J = 15$ GHz and show three different combinations of β and E_J/E_C . The anharmonicity undergoes a zero crossing in proximity of the half flux bias point.

confirm that the relative anharmonicity α indeed increases substantially beyond the transmon. Moreover, the values for $\beta \rightarrow 1$ tend asymptotically to the predictions for the quarton regime obtained from independent formulas (see [Sec. 2.3.3](#) below), further supporting the analysis presented here.

Given that φ_{ext} is a parameter of the Hamiltonian, we can easily obtain the flux dependence of the qubit spectrum as shown in [Fig. 2.5](#). Qualitatively, the IST spectrum shows two extrema as a function of applied flux, which we call *flux sweet spots* due to the disappearing flux sensitivity of the frequency, $\partial\omega_{01}/\partial\varphi_{\text{ext}} = 0$. The frequency minimum is achieved at half flux bias, and the positive anharmonicity quickly decreases and crosses zero when moving away from this point, analogously to the CSFQ [[63](#), [83](#)]. We fix E_J and choose three values of β to represent a high, middle, and low anharmonicity qubit, and then tune E_C such that the frequency at half flux lies within the range 4-5 GHz. As demonstrated by the three examples plotted in [Fig. 2.5](#), the behavior is qualitatively similar over a wide range of E_J/E_C and β . Note that we can freely reduce E_J/E_C to achieve the desired frequency since the circuit is immune to low-frequency charge noise, unlike the transmon. In fact, the high anharmonicity case $\beta \rightarrow 1$ tends to require higher E_C to remain in the 4-5 GHz region. One direct practical consequence is that the shunting capacitance C could be made smaller than in a typical transmon, reducing the device footprint.

2.3.3. A special case: the quarton ($E_L = E_J$)

When inductive and Josephson energies match exactly ($\beta = 1$), the quadratic term in $U(\hat{\varphi})$ disappears, leaving us with the purely quartic potential $U(\hat{\varphi}) \approx E_J \hat{\varphi}^4$. In this particular case the eigenenergies take the form [58, 84]

$$E_n = \lambda_n \left(\frac{2}{3} E_J E_C^2 \right)^{1/3}, \quad (2.22)$$

with the dimensionless prefactors λ_n to be computed numerically. The first three values are $\lambda_{0,1,2} \approx 1.06, 3.80, 7.46$. Note how the frequency scaling differs fundamentally from the weakly anharmonic case with $\omega \sim (E_J E_C)^{1/2}$. The stronger (weaker) dependence on E_C (E_J) should make the qubit properties more repeatable in practice, since E_J typically has larger associated fabrication errors than E_C [85].

Interestingly, the relative anharmonicity α_r assumes the constant value

$$\alpha_r = \frac{(\lambda_2 - \lambda_1) - (\lambda_1 - \lambda_0)}{\lambda_1 - \lambda_0} \approx 33.5\% \quad (2.23)$$

independently of E_J and E_C , meaning that the (positive) quarton anharmonicity is always approximately one third of its frequency. This can be regarded as the maximal non-linearity achievable in a single-well potential, since the linear (harmonic) contribution has been completely removed⁵. As mentioned earlier, going to $\beta < 1$ will cause the formation of a double well and the appearance of fluxon states, leading to flux qubit physics. Of course, the presence of multiple potential wells for $\beta \ll 1$ is what enables many interesting properties in the first place, such as decay-protected states and extremely large anharmonicities $\alpha_r \gg 1$ in the fluxonium and other circuits, as discussed in Sec. 2.3.1, but this goes beyond the scope of the present work.

Quarton qubits have been realized experimentally in CSFQ-like circuits using few-junction shunting arrays with $N \lesssim 10$ [58], with anharmonicities around 1 GHz and very promising coherence properties. More recently, a circuit called *unimon* [86], consisting of a single Josephson junction embedded in a shorted co-planar waveguide (CPW) resonator, has been shown to approach the ideal quarton case, with $\beta > 1.07$ and α_r up to 21 %. Of interest is the fact that the energy scales E_C and E_L in the unimon are provided mainly by the distributed capacitance and inductance of the CPW modes, in contrast to many other junction-based

⁵Hypothetically, the non-linearity could be further enhanced by engineering a potential where the quartic term itself is cancelled, leading to an extremely steep confinement $\sim \hat{\varphi}^6$. Ideas on how to implement such a potential are very welcome.

superconducting qubit implementations – including those presented in this work – that rely on effectively lumped circuits.

2.4. Coupled qubits

For a quantum computer to be of any practical use, a universal set of quantum gates has to be provided with it [12]. *Universal* here means that any arbitrary unitary operation on all qubits can be represented as a combination of gates from this set. Fortunately, it is enough to provide full single-qubit control (e.g. using rotations along the qubit X and Z axes) along with two-qubit entanglement between nearest-neighbor qubits for this to be fulfilled [1].

The literature on two-qubit entangling gates in superconducting circuits is vast and lies beyond the scope of this short overview. In particular, transmon-transmon gates have been studied extensively for well over a decade due to the simplicity and popularity of the transmon qubit, leading to an immense variety of implementations [49, 50, 56, 87–92]. More recently, increased interest in the fluxonium qubit as an alternative to the transmon has led to rapid progress in the theory and implementation of two-qubit fluxonium gates [93–97]. Here we are going to restrict ourselves to a small niche within this large field: coupling two different flavours of superconducting qubits.

2.4.1. Multi-species entanglement

Most of the research is focused on entangling qubits of the same type with each other, which is the simplest and perhaps most intuitive approach. On the other hand, it is also interesting to ask what happens when we try to couple different types of qubits, forming a *heterogeneous* or *multi-species* system⁶. Combining building blocks with fundamentally different properties could lead to interesting features, such as new interaction schemes or the reduction of crosstalk and decoherence in the overall system. There are a few such examples in the literature. The one which is most relevant to this work, and the only one demonstrated experimentally so far, is the CSFQ-transmon system [83, 100, 101], which is interesting from the point of view of static ZZ cancellation and is going to be discussed further in [Sec. 2.4.3](#) below. Recent theoretical proposals also look at coupling a transmon either to a fluxonium [102] or to a Cooper pair parity-protected qubit [103]. This is to strike a

⁶We intentionally avoid using the term *hybrid*, which is often used in the context of superconducting/semiconducting device physics [98, 99].

balance between the ease of manipulation and readout of transmons on one hand, and the higher degree of protection from the environment enjoyed by other qubit species on the other hand.

Slightly blurring the definition of what constitutes an individual qubit, we should also briefly mention the very interesting topic of multi-mode qubits [57]. While these are treated operationally as a single *logical* qubit, they show how two different *physical* modes with complementary properties can be combined in one circuit to obtain new and useful features. For example, a simple two-junction circuit can host two strongly coupled transmon-like modes with different field symmetries, leading to selective coupling of one mode to neighboring qubits and readout circuitry while isolating the other [104, 105]. More complex two-mode circuits utilize the properties of fluxon and plasmon modes in a clever way to encode a qubit state that is fully (i.e. exponentially) protected from relaxation and decoherence, at the expense of potentially more challenging device fabrication, qubit control and readout. The most prominent example of this approach is the *zero- π qubit* [57, 106].

2.4.2. Direct qubit coupling and the static ZZ interaction

A common method of physically coupling two superconducting qubits A and B together consists of introducing a mutual capacitance C_c between the superconducting islands defining each qubit, which couples their charge degrees of freedom \hat{N}_A, \hat{N}_B . This gives rise to a coupling Hamiltonian $\mathcal{H}_{int} = g_c \hat{N}_A \hat{N}_B$ [33], where the fixed coupling strength g_c is purely a function of the capacitive network of the system. Given the single-qubit Hamiltonians \mathcal{H}_A and \mathcal{H}_B , and expressing the charge operators in \mathcal{H}_{int} in terms of creation and annihilation operators as in Eq. (2.13), the total system Hamiltonian reads

$$\mathcal{H}_{2Q} = \mathcal{H}_A + \mathcal{H}_B + \mathcal{H}_{int} = \mathcal{H}_A + \mathcal{H}_B + \hbar J (\hat{a}^\dagger - \hat{a})(\hat{b}^\dagger - \hat{b}). \quad (2.24)$$

We denote the annihilation operators of qubits A and B as \hat{a} and \hat{b} respectively, and we introduced the *exchange coupling* strength

$$J = N_{zpf,A} N_{zpf,B} \times g_c, \quad (2.25)$$

which incorporates the charge zero-point fluctuations of the individual qubits. \mathcal{H}_{int} contains terms of type $J \hat{a}^\dagger \hat{b}$, which transfer a photon from one qubit to the other at a rate J . Non-energy-conserving terms of type $\hat{a}^\dagger \hat{b}^\dagger$ and $\hat{a} \hat{b}$ (creation or annihilation of two photons) are usually neglected for simplicity in the so-called rotating-wave

approximation (RWA), leaving us with the simplified and often-used interaction term $\hbar J(\hat{a}^\dagger \hat{b} + \hat{b}^\dagger \hat{a})$. This is the simple underlying XY-type interaction that enables the execution of common two-qubit unitary operations such as the *iSWAP*, where the two qubit transitions are brought on resonance to activate the exchange at rate J [42, 107], or the slightly more sophisticated *controlled phase (CPHASE)* gate, where exchange interactions with non-computational states introduce a phase on one qubit conditional on the state of the other [49]. The exchange term J is directly observable since it leads to the hybridization of the two coupled qubits and therefore gives rise to an avoided level crossing with a splitting of order $2J$, measurable e.g. via qubit spectroscopy.

If the qubit-qubit detuning $\Delta = \omega_A - \omega_B$ is much larger than the coupling strength, $J/|\Delta| \ll 1$, the resonant exchange is strongly suppressed and the qubits only weakly feel each other's presence. In this scenario, the transverse interaction in Eq. (2.24) can be reduced to an effective *longitudinal* term (ZZ-type) via a Schrieffer-Wolff transformation [33]. This is called the *dispersive regime* and is analogous to the usual interaction regime between a qubit and a far-detuned linear oscillator commonly used for qubit readout in circuit QED [39], where a transversal coupling g leads to a dispersive shift $\chi \approx g^2 \alpha / \Delta^2$ of the qubit frequency ω_q dependent on the state of the resonator with frequency $\omega_r = \omega_q + \Delta$, and vice versa⁷. In a similar fashion, consider a two-qubit system with single-qubit eigenstates $|i\rangle_A$ and $|j\rangle_B$, where $i, j \in \mathbb{N}$ are the number of excitations in qubits A and B. In the absence of interactions, $J = 0$, the bare product states $|ij\rangle = |i\rangle_A \otimes |j\rangle_B$ are also eigenstates of \mathcal{H}_{2Q} , with eigenenergies $E_{ij} = E_i + E_j$. Turning on the interaction in the dispersive regime does not significantly hybridize these states but only weakly perturbs them, leading to dressed eigenstates $|\tilde{ij}\rangle$ with slightly shifted energies \tilde{E}_{ij} . The state-dependent dispersive shift ζ can be expressed in terms of these eigenenergies,

$$\hbar\zeta = (\tilde{E}_{11} - \tilde{E}_{10}) - (\tilde{E}_{01} - \tilde{E}_{00}). \quad (2.26)$$

We can also rewrite it in terms of transition frequencies relative to the ground state as $\zeta = \omega_{11} - (\omega_{01} + \omega_{10})$, where $\hbar\omega_{ij} = \tilde{E}_{ij} - \tilde{E}_{00}$, which makes it obvious that $\zeta = 0$ if the qubits are not coupled. The nature of this shift is related to the fact that our superconducting qubits are not ideal two-level systems, but are often weakly non-linear oscillators with large manifolds of higher excited states that are not necessarily spectrally well-separated. All of these exert noticeable dispersive shifts on each

⁷An alternative, and promising, approach is to fundamentally replace the transverse coupling by a longitudinal one, leading to new interaction and readout schemes [108, 109].

other, most prominently within the two-excitation manifold $\{|\widetilde{11}\rangle, |\widetilde{02}\rangle, |\widetilde{20}\rangle\}$, where the two non-computational states $|\widetilde{02}\rangle$ and $|\widetilde{20}\rangle$ can shift $|\widetilde{11}\rangle$ appreciably and thus affect qubit operation. This type of interaction is commonly referred to as a longitudinal or ZZ interaction, as it takes the simple form $(\hbar\zeta/2)\hat{\sigma}_{z,A}\hat{\sigma}_{z,B}$ in the two-qubit subspace, with $\hat{\sigma}_{z,i}$ the Pauli Z operators for each qubit. This is in contrast to the resonant case, where the bare exchange interaction takes the form of an XY term. In the case of qubits with small anharmonicities (Duffing nonlinearities) α_A and α_B , such as in the transmon case, a perturbative expansion leads to the approximate expression [101]

$$\zeta \approx \frac{2J^2(\alpha_A + \alpha_B)}{(\Delta + \alpha_A)(\Delta - \alpha_B)}. \quad (2.27)$$

Although small for $|J/\Delta| \ll 1$, the ZZ shift ζ is always present and therefore represents a source of error in multi-qubit systems: if a qubit's frequency is dependent on the state of its neighbors, it will constantly shift around during the application of an arbitrary series of gates forming an algorithm, leading to the accumulation of unwanted conditional phases. These errors are not random, such as stochastic relaxation and dephasing processes, but arise due to the purely unitary evolution of the entire multi-qubit system and are therefore *coherent* errors that cause unwanted entanglement. These errors can represent a limiting factor for the fidelity of two-qubit gates when the effect of incoherent processes (qubit dephasing and decay) is small enough, which is a case that is often realized in contemporary experiments when the gate implementation does not require flux tuning [91, 100, 110]. Furthermore, the malicious effect of the ZZ interaction extends to next-nearest neighbor qubits, so-called *spectators*, which are not directly involved in the two-qubit gate but experience error accumulation due to the always-on nature of the ZZ-term [111, 112]. This will only make matters worse as circuits are scaled up to larger numbers of qubits.

2.4.3. ZZ suppression techniques

The detrimental effect of the static ZZ interaction and its ubiquity in coupled superconducting qubits has led to the development of several solutions to mitigate it or suppress it. In the following we offer a brief overview of the most common techniques.

Looking at Eq. (2.27), one can decrease the ZZ term simply by making the detuning Δ very large. This is commonly employed in conventional flux-tunable

gates [49, 90, 113], where the frequency tunability of the qubits allows them to stay far detuned when idling and are pulsed to (or close to) a resonance condition to carry out the gate. While simple and effective, this method is quite limited as it can never cancel ζ completely, since Δ cannot be made arbitrarily large, and it might run into frequency crowding issues as the number of coupled qubits increases. Tuning a qubit across a wide frequency range will inevitably cause brief collisions with unwanted transitions, accumulating coherent errors.

Another option is to make J an in-situ tunable parameter instead of a fixed quantity, such that it can be turned on only when needed for performing a gate, and can be otherwise turned off (at least approximately) to reduce the static ZZ interaction. This cannot be achieved with fixed capacitive coupling as described above, but instead by adding a coupling element that is itself tunable. Several realizations of this idea have been demonstrated, such as using a tunable inductance [114, 115] or an additional tunable qubit mode (often called a tunable *bus*) [116–118]. Going one step further, by intentionally coupling two qubits both directly via a capacitance and indirectly via a tunable mode, more sophisticated *multi-path* coupling can be engineered [96, 119–123]. The underlying idea is that couplings along different paths can interfere constructively or destructively, allowing for large on/off ratios. While generally implemented using at least one tunable mode, a version with only direct coupling and a fixed linear bus mode has also been demonstrated [91]. This has the advantage that no flux tunability is required, reducing the number of control parameters and loss channels.

While there is no doubt that the tunable coupling approach has proven to be successful experimentally [10], the associated hardware costs might be undesirable. The couplers take up valuable real estate on the device substrate, and they require additional control lines and electronics. One can instead choose a fundamentally different approach, resorting to purely software-based solutions that rely on optimal system control. An example is the technique of dynamical decoupling (DD), borrowed from the field of nuclear magnetic resonance, whereby the insertion of echo pulses during system evolution can cancel out unwanted coherent interactions [124–127]. A drawback is that determining the correct DD sequences might add significant computational costs as systems increase in size. More recently, the suppression of static ZZ error has also been demonstrated through the application of continuous, off-resonant microwave drives to individual qubits, causing AC-Stark shifts of the energy levels $|i\tilde{j}\rangle$ which negates the static dispersive shifts [83, 110, 128, 129]. This is an attractive approach, since the AC-stark shift tones can be applied via the

same control lines used for standard qubit drives, and the successful realization in both transmon and fluxonium systems demonstrates its generality. While the hardware and software overhead is arguably smaller than for other methods, there is still the requirement of additional electronic components at room temperature and more parameters to tune in software.

Intrinsic ZZ suppression

Broadly speaking, the ZZ suppression methods discussed so far try to externally solve an issue that is assumed to be intrinsic to all practical superconducting circuits, because it is inevitably tied to the use of non-ideal multi-level systems with often non-trivial dispersive interactions. Let us now move away from this assumption and ask whether we can instead modify or replace the underlying circuits to remove the issue of static ZZ interaction at a more fundamental level. As mentioned earlier, the spectral proximity of levels such as $|\widetilde{02}\rangle$, $|\widetilde{20}\rangle$ and $|\widetilde{11}\rangle$ is the main source of the unwanted ZZ. Therefore, pushing these levels further apart will naturally reduce this effect, which is equivalent to increasing the anharmonicity of the circuit and make it appear more similar to an ideal two-level system. This has in fact been a major driver in the quest to find a replacement for the transmon, which by construction has a small relative anharmonicity limited to a few % of its frequency [40]. The most prominent candidate is undoubtedly the fluxonium [64], which has gained considerable popularity in recent years. Two-qubit gate implementations with statically coupled fluxoniums have been realized which indeed suffer less from static ZZ than their transmon counterparts [93–95]. Furthermore, fluxoniums are generally compatible with extrinsic ZZ suppression methods that were initially developed for the transmon, as shown in some of the literature referenced above.

An alternative route that does not require the move to a qubit such as the fluxonium, which can be challenging to fabricate due to the use of long junction arrays, can be intuitively derived by once again looking at the expression for ζ in Eq. (2.27). The sum of anharmonicities in the numerator is always non-zero when using identical types of qubits. Going back to our discussion of multi-species qubit systems in Sec. 2.4.1, let us instead choose to couple a qubit with positive and one with negative anharmonicity – this allows us to reduce the ZZ shift, or indeed cancel it completely for $\alpha_A = -\alpha_B$. This idea has previously been treated theoretically [101, 130]. The intuition is that the $|\widetilde{02}\rangle$ and $|\widetilde{20}\rangle$ states now exert opposite *pressure* on the computational $|\widetilde{11}\rangle$ state instead of pushing it into the same direction, leading to exactly $\zeta = 0$ when the system is in balance. While the

perturbative expression in Eq. (2.27) only provides a qualitative guideline, as it does not take all dispersive shifts from higher states into account, this simple idea holds in practice and has in fact been demonstrated experimentally by coupling a transmon to a CSFQ [100]. The CSFQ typically has a positive anharmonicity of a few 100 MHz when biased at its half flux sweet spot [63] and is therefore an ideal companion to the transmon. In the experiment of Ref. 100, a scan of the CSFQ flux bias identifies the optimal point of ZZ cancellation, and the benchmarking of a cross-resonance entangling gate shows how this point leads to a higher gate fidelity. While a linear bus mode is used there as the coupler between the two qubits, the ZZ cancellation by opposite anharmonicities is equally possible in directly coupled qubits [101]. This provides us with a minimal working ZZ-suppressed system without the need for any additional hardware or software elements. Crucially, the ZZ suppression here is completely *passive*.

While the CSFQ is a relatively easy qubit design to implement, with control and readout schemes very similar to the transmon, its nature as a multi-junction device makes it more sensitive to fabrication imperfections than a single-junction transmon. Luckily, we can further simplify this system by replacing the CSFQ with another plasmonic qubit with positive anharmonicity and essentially the same physics – the IST discussed in Sec. 2.3.2. This has been demonstrated experimentally in a preliminary study [79] where an IST, realized using a meandering wire as the shunting inductance, is statically coupled to a transmon and shows extremely low levels of ZZ shift. Like the transmon, the IST contains only one Josephson junction, and this is sufficient to create a non-trivial interaction. This is a surprisingly simple and at the same time powerful approach that merits further investigation, and we will therefore devote considerable effort to explore this direction over the course of this work, namely in Chapter 5 and Chapter 6.

2.5. Summary

We began by offering a brief overview of the basics of superconducting quantum circuits, which owe much of their interesting physics to the Josephson junction and its associated non-linear inductance. The CPB and the transmon are simple realizations of superconducting qubits that can be understood in terms of their charging energy E_C and Josephson energy E_J . The transmon in particular has become the dominant type of superconducting qubit in large-scale quantum circuits

over the last several years, thanks to its simple design and increasingly high coherence properties.

We then looked more closely at the RF-SQUID, which is a general circuit from which many types of technologically relevant superconducting qubits can be derived by simply choosing different relative scales of energies E_J , E_L , and E_C . Although there has been considerable interest in superinductance-based circuits where typically $\beta = E_L/E_J \ll 1$, we instead choose to explore the opposite regime of $\beta > 1$, which is treated both perturbatively and numerically to extract the basic physics of this type of circuit. We refer to it as the inductively shunted transmon (IST) due to its qualitative similarity to the conventional transmon qubit. Both the transmon and the IST show a simple plasmonic excitation spectrum with relatively weak nonlinearity, and both realize useful qubit regimes with similar values of E_J and E_C . The IST has the additional feature of a strongly flux-dependent anharmonicity, which is positive (and generally larger) at the half-flux bias point of the RF-SQUID loop, while it crosses zero and becomes negative as the flux bias is swept. Further, the interesting limiting case $\beta = 1$ allows this circuit to reach the quarton regime with a maximal nonlinearity equivalent to roughly 33 % of the qubit frequency, much larger than the transmon.

The remainder of this chapter focused on the description of basic two-qubit dynamics in superconducting circuits, in particular the ubiquitous but unwanted static longitudinal (ZZ-type) interaction that plagues many circuit realizations. We discussed the most common techniques to suppress the ZZ interaction between pairs of qubits, ranging from additional hardware components such as tunable couplers, to control schemes that cancel ZZ via active pulses or continuous drives. A fundamentally different approach consists of coupling two different species of qubits together. When their anharmonicities are opposite in sign, the ZZ term can be passively cancelled even in a pair of statically coupled qubits without any additional hardware or software requirements. This simple but elegant idea will be explored further over the course of this work.

3.

Flux control in superconducting quantum circuits

3.1. Introduction

So far we have described a few simple superconducting circuits commonly used in quantum computing applications. With the notable exception of the fixed-frequency transmon, the majority of technologically relevant superconducting qubit species contain a closed loop and implement some variation of either a DC-SQUID or RF-SQUID, therefore relying on magnetic flux control for their operation. This ranges from a simple DC (static) flux bias for choosing the operating point of the qubit – usually a flux sweet spot that maximizes the qubit coherence time – to finely calibrated time-domain flux pulses for implementing high-fidelity two-qubit entangling gates. Even in the case of fixed-frequency transmons, an additional flux-tunable circuit element is sometimes added as a tunable coupler to provide fast gates while minimizing unwanted coupling and crosstalk. The ubiquity of flux tuning in superconducting quantum circuits therefore motivates a closer inspection of this particular technological aspect.

Flux tunability was an integral part of the theoretical discussion in [Chapter 2](#), with the external flux variable Φ_{ext} often entering the circuit Hamiltonians, but the practicalities of this control parameter were not further explained. While this is not necessary in order to understand the basic physics of circuits such as the tunable transmon, fluxonium or inductively shunted transmon (IST), it will affect both the practical design and the performance of such circuits. The additional external degree of freedom enables a wide variety of interesting circuits and operation schemes, but it comes with a new noise channel that needs to be kept under control.

Here, we aim to offer a short overview of different magnetic flux bias approaches in superconducting quantum circuits, each with their advantages and drawbacks. We then look more closely at the effect of flux noise on the performance of flux-tunable circuits, and distinguish between different sources of noise. In particular, the existence of microscopic magnetic defects causing local flux noise in the SQUID

environment has been studied extensively in the field. We discuss how the sensitivity to this local flux noise changes as a function of parameters such as the geometry of the SQUID loop. This will provide a set of tools to understand and quantify flux noise sources in experiments, and hopefully mitigate them.

3.2. Approaches to flux control

Let us assume that we have a superconducting qubit with a SQUID loop that we want to flux bias. The simplest way to achieve this is to generate a magnetic field by running a current I through a wire and placing it close to the SQUID in the correct orientation, causing a net flux Φ_{ext} to be enclosed within the SQUID loop. This inductive coupling can be quantified by the mutual inductance $M = \Phi_{\text{ext}}/I$ between the two elements. Broadly speaking, we can separate practical flux bias designs into two categories: whether the bias circuit is located on the same substrate hosting the SQUID or not.

3.2.1. On-chip flux bias lines

A common flux delivery method consists of using a small inductive loop patterned lithographically in proximity of the SQUID [49], usually in the same fabrication steps that define the rest of the circuit. This on-chip *flux bias line* (FBL) is sketched in Fig. 3.1 (a). The resolution and repeatability afforded by lithography means that the FBL and SQUID can be positioned very closely to each other and can thus be made small, usually in the 10's of μm . This is advantageous from the point of view of selectivity and reduction of crosstalk in multi-qubit circuits, as each SQUID can be biased by its own FBL without causing significant stray fields that will bias other SQUIDs. One end of the FBL is typically shorted to a 2D ground plane on the substrate, meaning that the injected current I is dispersed into this plane and then finds its way back out via wirebonds at the edge of the substrate or similar connections that eventually lead to the cryostat reference ground.

On-chip FBLs are used routinely in state-of-the-art experiments with planar multi-qubit circuits [23, 26, 131], thanks to their simple design and compact size. The small self-inductance of the FBL translates into a fast response time, which can be used to apply short (10's of ns) time-varying flux pulses to the tunable circuits, a technique employed in many common entangling gate schemes [49, 56, 90]. On the other hand, planar FBLs are not without drawbacks. For example, the uncontrolled paths of the return current on the substrate ground plane can cause unwanted stray

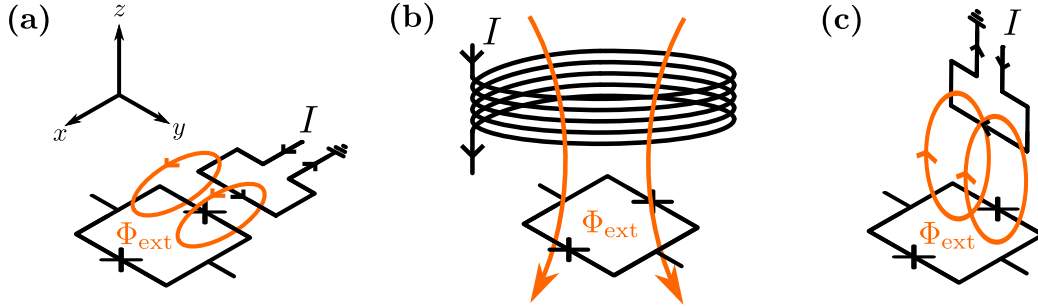


Figure 3.1: Simple magnetic flux control of SQUIDs. (a) Conventional on-chip flux bias line (FBL) formed by a small inductive loop patterned in proximity of the DC-SQUID loop to be tuned. A current I generates a flux Φ_{ext} . (b) Large external solenoid used to apply flux remotely. (c) Modified FBL design where the current loop is located outside of the SQUID device plane.

magnetic fields and lead to non-local crosstalk between qubits [132]. Flux crosstalk is usually expressed in terms of the unwanted amount of flux bias on a neighboring qubit relative to the bias applied intentionally to the target qubit. The typical flux crosstalk in these architectures is in the range of 1-10 % [132–136], although smaller values on the order of 0.1 % have been reported recently in a 17-transmon device [26]. Furthermore, the presence of possibly resistive components (wirebonds, solder joints) can lead to substrate heating, affecting the device performance and adding heat load to the cryogenic environment [137]. These galvanic connections might also degrade or fail upon repeated thermal cycling. Each FBL requires a dedicated input line, increasing the density of wiring both on the substrate and in the cryostat. Line crowding on the substrate can be alleviated with air-bridges [138] or by moving to 3D integration solutions such as through-silicon vias [139], but these increase the fabrication and design complexity, potentially degrading qubit performance.

3.2.2. Off-chip flux delivery

A different approach consists of removing the flux bias circuit from the substrate hosting the quantum circuit to be tuned. This has the advantage that no galvanic connection to the substrate is needed, simplifying the qubit circuit itself and removing the problem of resistive substrate heating. The simplest solution is to use a large solenoid, located above or around the substrate, as depicted in Fig. 3.1 (b). The solenoid is often attached to the outside of the sample enclosure containing the substrate, and superconducting wire is typically used in order to avoid resistive heating in the coil. While this allows for the remote application of flux, it makes

the bias much less selective and requires larger SQUID loop areas. Additionally, the large inductance of the solenoid does not support the delivery of fast flux signals. External solenoids are therefore only used in small-scale experiments when merely static (DC) flux tuning of one or two SQUIDs is needed without the requirement of selectivity [29, 53]. This method is also not compatible with the use of superconducting sample enclosures (e.g. aluminium), because these will shield any external field due to the Meissner effect. This particular issue can be solved by using a so-called magnetic hose, consisting of alternating concentric layers of superconductor and high-permeability material [140]. This method has been used successfully to externally flux-tune transmons located inside an aluminium 3D cavity [141]. Unfortunately, the bulky nature of the magnetic hose still does not make it practical for selective flux tuning of individual SQUIDs.

To regain flux bias selectivity, one has to again shrink the bias circuit and bring it closer to the SQUID, as sketched in Fig. 3.1 (c). This technical challenge has been solved by lithographically patterning a planar FBL on a separate substrate, and then bringing it into proximity with the flux-tunable circuit [142, 143]. In Ref. 142, the transmon and FBL are located on individual sapphire chips and placed next to each other inside a 3D cavity architecture. The FBL also incorporates low-pass filtering to reduce Purcell decay of the qubit excitation via the FBL line. A more compact and scalable solution has been developed in Ref. 143, which uses a flip-chip technique [144] to accurately position the FBLs above the qubits with a separation on the order of 25 μm . Furthermore, the design is chosen to simultaneously provide flux tuning and standard microwave control through the same element, implementing a compact *XYZ line* for full qubit control¹. This is often not possible using planar on-chip FBLs, since the capacitive coupling of qubit to FBL is enhanced by the large relative permittivity $\epsilon_r \sim 10$ of typical substrates, while the magnetic permeability $\mu_r \sim 1$ is the same as in vacuum. This can lead to significant Purcell decay of the qubit via the FBL when targeting design parameters that provide a useful amount of flux tunability [142, 143, 145]. The solution is to heavily low-pass filter the FBL input to block the qubit excitation from escaping into the environment, but this also makes it difficult to drive the qubit with microwave signals. Moving the FBL away from the substrate and using only vacuum as the dielectric medium balances out the strength of capacitive and inductive coupling, allowing for simultaneous XYZ control.

¹Other implementations of XYZ lines have been recently reported, but without insights into their specific design [23, 96].

The off-chip FBL solution of Ref. 143 is promising, yet it is not immediately clear whether flux crosstalk will be improved compared to standard planar designs². This is because the individual FBLs are still shorted to a common ground plane on the auxiliary chip, which can be expected to lead to similar crosstalk behavior as in conventional designs [132]. For the purpose of scalability and achieving reliable high-fidelity qubit operations, it would be advantageous to intentionally engineer low levels of flux crosstalk. In addition, due to the flip-chip technology involving the bonding of the FBL chip to the qubit chip, a new FBL substrate has to be fabricated for each new device to be measured. Making the FBLs reusable would simplify the process and make device change-overs faster. We will attempt to address these challenges through the development of a novel, fully 3D-integrated flux tuning architecture in Chapter 4.

3.3. Flux noise

Coupling a SQUID to an external flux control parameter Φ_{ext} enables interesting functionality, but this comes at a cost. Any noise on this control parameter will also couple into the SQUID and affect its operation. In this section we provide an overview of different sources of noise and how they impact the performance and operation of tunable qubits.

Any source of magnetic field, whether intentionally provided or not, can in principle couple to the SQUID loop. If this source is perfectly static, it can be compensated simply by applying a corresponding DC bias on the FBL of the SQUID. But if the source is noisy and fluctuates in time, it will cause the applied flux to become a time-dependent quantity $\Phi_{\text{ext}}(t)$. These random fluctuations modify the effective Josephson energy of the loop, such as in Eq. (2.9) for a DC SQUID, which in turn causes fluctuations in the qubit frequency. A randomly time-varying qubit frequency is what leads to the loss of information about the phase of the qubit wave-function, a process referred to as *dephasing* [32]. This is usually quantified by a relevant dephasing timescale T_ϕ and dephasing rate $\Gamma_\phi = 1/T_\phi$ that depend on the details of the underlying physical dephasing processes.

3.3.1. Common sources of noise

In Fig. 3.2 we depict a few common sources of flux noise relevant to superconducting circuits. The intentionally designed flux bias circuit, such as a large inductive coil

²No crosstalk metrics are reported in that study.




Figure 3.2: Sources of magnetic flux noise in SQUIDs. Cartoon showing common sources of unwanted magnetic flux noise in superconducting circuits. The current I provided by the flux-bias line contains a noise component $\delta I(t)$ that is coupled into the SQUID via the mutual inductance M . External sources of magnetic fields, e.g. electronic instruments or magnetic materials, add an environmental background $B_{\text{env}}(t)$. Independent local magnetic fluctuators with moment m and surface density σ provide a local $1/f$ noise component.

or an on-chip FBL, is usually driven by an external current source, which will have an instrument-specific noise output $\delta I(t)$ that is added on top of the desired signal I . The noise on the supplied current can be quantified by its power spectral density (PSD) $S_I(\omega)$, describing the frequency spectrum of the noise. This will be converted to magnetic flux noise felt by the SQUID via the mutual inductance M as

$$S_{\Phi}(\omega) = \left| \frac{\partial \Phi}{\partial I} \right|^2 S_I(\omega) = M^2 S_I(\omega). \quad (3.1)$$

This requires the use of high-quality, low-noise instrumentation, as well as careful filtering and attenuation of the cryogenic input lines carrying the current. The filtering also helps to reduce thermal noise that enters the cryostat via the cables at room temperature.

In large-scale systems, the combined DC currents required for all tunable components can lead to significant heat load on the cryogenic environment [137]. Therefore, large mutual inductances are desired to minimize the amount of required current. On the other hand, Eq. (3.1) cautions us against coupling too much current noise from the FBL into the SQUID, leading to a delicate trade-off between tunability and noise. This is a typical problem faced when trying to couple a quantum system to the external world to allow us to control it, while also protecting it from decay and dephasing channels added by the environment.

Magnetic fields external to the experiment are always present, generated by other instruments in the lab environment, computers, monitors, ferromagnetic materials and components (iron, nickel, etc.), and even earth's own magnetic field. This will

lead to a varying background magnetic field $B_{\text{env}}(t)$. Since most commercial cryostat radiation shields are made of non-magnetic materials such as aluminium or copper with relative magnetic permeability $\mu_r \approx 1$ similar to vacuum, magnetic fields will easily penetrate all the way to the cryogenic experiment. The most common method for reducing this effect is the use of high-permeability metal shields (typically made of *mu-metal*, a ferromagnetic nickel-iron alloy) placed around the device inside the cryostat. This will attenuate external fields and reduce their influence on SQUID operation. One can also decrease the intrinsic circuit sensitivity to such external fields by using a gradiometric SQUID design [146–148], composed of two independent loops of equal size. Such a gradiometer is insensitive to homogeneous fields, and is only tuned by fields that have a significant gradient between the two SQUID halves, such as an appropriately located FBL. This approach will also be used in our tunable transmon design in Chapter 4.

Even when external fields are shielded, instrumentation noise on the FBL has been filtered out, and no magnetic materials and components are used in proximity of the device under investigation, one still has to worry about the presence of microscopic magnetic impurities on the device substrate itself. It has been found repeatedly in various experiments with SQUIDs of vastly different sizes that the dominant flux noise spectrum is often of $1/f$ type [149–152] and that it scales with SQUID geometry and size in a way that is not consistent with the assumption of a distant external field source. Instead, it is compatible with an ensemble of independent microscopic fluctuators with material-dependent magnetic moments m and surface density σ located in the surfaces and interfaces surrounding the SQUID itself [152, 153], as depicted qualitatively in Fig. 3.2. This model will be discussed in more detail in Sec. 3.3.3. The $1/f$ noise spectrum is characterized by a power spectral density (PSD) of the form $S_{\Phi,1/f}(\omega) = A_{\Phi}/|\omega|$, with the *flux noise amplitude* $\sqrt{A_{\Phi}}$ typically in the range $\sqrt{A_{\Phi}} \sim 1 - 10 \mu\Phi_0$ in these experiments³. Adding this noise component to the one from Eq. (3.1), we obtain a combined noise PSD

$$S_{\Phi}(\omega) = M^2 S_I(\omega) + A_{\Phi}/|\omega|. \quad (3.2)$$

The fact that this finding is reproducible using different SQUID materials and across orders of magnitudes in SQUID loop areas suggests a common origin for this microscopic noise, which is often simply referred to as *universal flux noise*. Several candidates for these small magnetic impurities have been suggested, such as adsorbed gas molecules on the SQUID surface [154–156], or paramagnetic defect

³The unit $\mu\Phi_0 = 10^{-6}\Phi_0$ is commonly used to quantify this type of noise.

centers embedded in surface oxides [157]. This is an active field of research and a conclusive answer has not yet been reached. The importance of understanding and eventually mitigating universal flux noise cannot be understated, as it would remove the main drawback of many flux-tunable qubit designs and gate implementations, whose performance is often limited by decoherence due to flux sensitivity of circuit parameters. A noiseless flux environment would also render the use of flux sweet spots superfluous, vastly increasing the practical parameter space and allowing for arbitrary precision in setting frequencies and detunings without requiring stringent fabrication tolerances.

3.3.2. Experimental estimation of 1/f flux noise amplitude

Here we briefly describe how to experimentally quantify the level of 1/f flux noise present in the environment of a tunable qubit. Assuming that other noise components (e.g. white noise) are negligible, qubit dephasing is dominated by flux noise with a PSD $S_\Phi(\omega) = A_\Phi/\omega$. In this case, it is possible to easily estimate the amplitude A_Φ by performing coherence measurements of the qubit as a function of applied flux [150, 152, 158]. When subjected to 1/f noise, one can derive that the qubit state decoheres in time according to a Gaussian decay function $\propto \exp[-(\Gamma_\phi^G t)^2]$, with the Gaussian pure dephasing rate Γ_ϕ^G , as opposed to a simple exponential decay function (such as when limited by stochastic white noise). The choice of coherence measurement sequence (bare Ramsey sequence, single echo pulse, multiple echo pulses) acts as an effective filter function that will expose the qubit to different regions of the noise spectrum [158] and therefore lead to different values of dephasing rate. For a single echo pulse, i.e. a standard Hahn-echo sequence, the following expression for the relevant dephasing time $\Gamma_{\phi,E}^G$ is derived [150]:

$$\Gamma_{\phi,E}^G = \sqrt{A_\Phi \ln 2} \times \left| \frac{\partial \omega_{01}}{\partial \Phi} \right|, \quad (3.3)$$

from which the flux noise amplitude $\sqrt{A_\Phi}$ can be extracted. A different coherence measurement sequence can equivalently be used, leading to the same scaling but with a different numerical prefactor in Eq. (3.3) – for example, a Ramsey sequence leads to a larger dephasing rate $\Gamma_{\phi,R}^G \approx 5 \times \Gamma_{\phi,E}^G$, which highlights the degree of noise protection afforded by the echo pulse.

According to the model presented here, dephasing due to flux noise appears only when the qubit is operated at a point of finite frequency slope $|\partial \omega_{01}/\partial \Phi|$, where a fluctuation in flux will lead to a first-order correction of the qubit frequency. In

contrast, this dephasing rate disappears at a flux sweet spot where $|\partial\omega_{01}/\partial\Phi| = 0$. This dependence has led to the development of flux-tunable qubits which have a relatively flat frequency-flux dispersion across the whole tunability range and are therefore less exposed to noise, such as highly asymmetric DC SQUIDs [55] or superinductance qubits in a particular parameter regime [80]. This approach can be thought of as Hamiltonian engineering, where system properties are tailored to make the circuit more robust to external parameters. For completeness we mention here that higher-order dephasing contributions, proportional to the curvature $\partial^2\omega_{01}/\partial\Phi^2$ and further terms, can be derived similarly but are predicted to be negligible in practice compared to other dephasing mechanisms [40].

In a general scenario, multiple dephasing sources will play a role in qubit dynamics. In the presence of both stochastic and $1/f$ noise, with associated dephasing rates Γ_ϕ and Γ_ϕ^G respectively, the total decoherence function of the qubit (in the off-diagonal matrix elements of the density matrix) can be expressed as [32]

$$\chi(t) = e^{-\frac{\Gamma_1}{2}t} e^{-\Gamma_\phi t} e^{-(\Gamma_\phi^G t)^2}, \quad (3.4)$$

which should be used when fitting experimental decoherence curves to extract the various decay rates correctly. Here $\Gamma_1 = 1/T_1$ denotes the conventional energy relaxation rate of the qubit. In the absence of $1/f$ noise, the expression reduces to a pure exponential decay $\exp(-t/T_2)$ with the familiar definition of the coherence time $T_2 = (\Gamma_1/2 + \Gamma_\phi)^{-1}$. Further, if dephasing is negligible compared to the rate of energy relaxation, we obtain the upper limit $T_2 = 2/\Gamma_1 = 2T_1$ for relaxation-limited coherence of a qubit.

3.3.3. Influence of SQUID geometry on flux noise

In [Sec. 3.3.1](#) we mentioned studies of local flux noise as a function of SQUID size and geometry, and here we want to describe this particular aspect in more detail. This section is mainly based on Ref. [152], which expands on a model for local flux noise developed earlier [151, 153, 159]. The main finding is that the aspect ratio of the SQUID loop, i.e. the ratio of loop perimeter over the width of the superconducting wire, is the main parameter that affects how much flux noise will be picked up by the SQUID, quantified by the magnitude of the flux noise amplitude A_Φ . This is fundamentally different from the naive picture of a global external flux noise, which should simply scale proportionally with the total SQUID

area. It is therefore more important to optimize the relative dimensions of the SQUID rather than its absolute size.

We start with the assumption that all surfaces surrounding the SQUID contain a homogeneous density σ of independently and randomly flipping spins, each with magnetic moment m . Each of these spins generate a local field that will couple a small amount of flux into the SQUID loop. Assuming the fluctuator ensemble to be isotropic, the total average captured flux will be $\langle \Phi \rangle = 0$ but its variance $\langle \Phi^2 \rangle$ will be finite. The latter can be calculated as (see supplement of Ref. 152):

$$\langle \Phi^2 \rangle = \frac{1}{3} m^2 \sigma \int dA \left(\frac{B}{I} \right)^2,$$

where we integrate over all surfaces of interest. B is the magnetic field generated at each point due to a test current I through the SQUID loop. It is predicted that the largest contribution to flux noise comes from fluctuators residing in the immediate proximity of the SQUID current, i.e. on the surface of the superconducting metal itself [153]. For a toroidal loop of perimeter P made from a wire with circular cross-section of diameter W , Ampere's law gives $B = \mu_0 I / (\pi W)$, a constant field along the whole surface of the wire, so we simply need to integrate the surface area of the torus $A = \pi W P$, leading to

$$\langle \Phi^2 \rangle = \frac{1}{3} m^2 \sigma \left(\frac{\mu_0}{\pi W} \right)^2 \int dA = \frac{\mu_0^2}{3\pi} m^2 \sigma \times \frac{P}{W}. \quad (3.5)$$

This expression neatly separates the two fundamental approaches to the targeted mitigation of flux noise. On one hand, the term $m^2 \sigma$ contains all the information about the properties of the materials and processes used to build the SQUID. This is the aspect that is tackled in studies of surface magnetic defects and their removal [156, 157]. On the other hand, the term P/W is purely dependent on the SQUID geometry, and can therefore be optimized by careful design choices.

A correction to the formula in Eq. (3.5) can be derived [152] when assuming a more realistic wire cross-section, i.e. a thin rectangular strip of width W and thickness $b \ll W$ instead of a circular wire. This results in a weak logarithmic correction that is typically less than a factor of 2 over more than one order of magnitude variation in W for typical thin films with $b \sim 100$ nm. This means that for a qualitative treatment of the physics we can neglect this logarithmic correction and focus for now on the simple linear dependence on the ratio P/W .

The missing step is now to connect the flux variance $\langle \Phi^2 \rangle$ from this microscopic model to a measurable quantity in experiments, such as flux-induced Gaussian

dephasing of the qubit at a rate $\Gamma_{\phi,E}^G$. The total flux variance (or flux noise power) is obtained by integrating the noise PSD $S_{\Phi}(\omega)$ over all frequencies. In practice, the particular choice of experimental sequence to measure qubit coherence will be most sensitive to certain regions of the noise spectrum, introducing an effective filter function $g(\omega)$. The flux variance is then obtained as $\langle \Phi^2 \rangle = \int S_{\Phi}(\omega)g(\omega)d\omega$. Assuming a 1/f noise PSD as described earlier, and the filter function for a Hahn-echo sequence [32], we obtain $\langle \Phi^2 \rangle = 2A_{\Phi} \ln 2$. This finally allows us to rewrite the Gaussian dephasing rate from Eq. (3.3) as

$$\left(\Gamma_{\phi,E}^G\right)^2 = \frac{\langle \Phi^2 \rangle}{2} \times \left| \frac{\partial \omega_{01}}{\partial \Phi} \right|^2 = \frac{\mu_0^2}{6\pi} m^2 \sigma \times \frac{P}{W} \times \left| \frac{\partial \omega_{01}}{\partial \Phi} \right|^2. \quad (3.6)$$

This is the main result of this section – the flux-limited dephasing rate of a qubit is a combination of three main ingredients: material/fabrication properties, SQUID geometry, and Hamiltonian engineering. The simple geometric scaling $\Gamma_{\phi,E}^G \sim \sqrt{P/W}$ has been reproduced experimentally [152], underscoring the importance of SQUID design optimization.

Extension to general SQUID geometries

In order to apply this model to more general SQUID geometries such as the gradiometric coaxial qubits studied in this work, we need to take into account the fact that in a realistic device the SQUID loop is made out of very narrow sections (the junction leads and, in the case of the inductively shunted transmon presented in Chapter 5, the meandering inductor) and wider sections. We cannot integrate along the whole perimeter assuming a fixed width W , but instead need to break it up into sections P_i with $\sum_i P_i = P$, each having a different width W_i . By integrating over the different sections individually along the entire closed perimeter, the flux noise power is estimated as

$$\langle \Phi^2 \rangle = \frac{\mu_0^2}{3\pi} m^2 \sigma \sum_i \frac{P_i}{W_i}. \quad (3.7)$$

The figure of merit to be used for comparison of SQUID designs is therefore the purely geometric quantity

$$\rho := \sum_i \frac{\Delta P_i}{W_i}, \quad (3.8)$$

which we refer to as a *generalized aspect ratio*. This gives us an intuitive dimensionless parameter to compare different qubit geometries. Note that from Eq. (3.6)

we can deduce $\Gamma_\phi^G \sim \sqrt{\langle \Phi^2 \rangle} \sim \sqrt{\rho}$, so we are interested in the square root of ρ if we want to quantify the effect of geometry changes on the dephasing rate.

For the optimal geometry chosen in Ref. [152] ($W = 2 \mu\text{m}$, $P = 32 \mu\text{m}$) the simple value $\rho = P/W = 16$ is obtained. Values much smaller than this are not going to be practically useful. If the inner perimeter of the loop goes towards zero, ρ asymptotically approaches 4, but the SQUID has now lost the ability to be tuned effectively. Note that the authors of Ref. 152 do not include the thin JJ leads in their definition of aspect ratio, which likely leads to an underestimate of the noise sensitivity discussed here. We instead choose to include the entire SQUID perimeter in the definition of Eq. (3.8). This adds a value of $\sim 30 - 40$ to the total value of ρ for each JJ added to the SQUID, based on our JJ design presented later (see Fig. 4.2). Our intuitive reasoning is as follows: given predicted surface spin densities on the order of $\sigma \sim 10^{17} \text{ m}^{-2} = 0.1 \text{ nm}^{-2}$ [152, 153, 160] compatible with experiments, a section of narrow strip of metal on the order of $W \sim 100 \text{ nm}$ or more, typical for aluminium JJ leads, is still large enough to host on average a considerable number of defects which will affect the flowing supercurrents. On the other hand, we note that the previously made assumptions of a large homogeneous spin density likely breaks down in this mesoscopic regime, requiring a more careful treatment.

3.3.4. Radiative loss through the FBL

In this section, we have mostly been concerned with flux noise that is added to our SQUID circuit from different sources, and how this affects the qubit dephasing time T_ϕ . But, as mentioned earlier in Sec. 3.2, the mere presence of a FBL creates a new channel for the qubit excitation to be dissipated into its environment (radiative Purcell decay) and can therefore limit the qubit lifetime T_1 as well. This radiative loss can happen both through inductive and capacitive coupling between qubit and FBL. The decay rate due to the inductive component is found to be proportional to M^2 [40], and is non-zero only in the case of finite SQUID asymmetry d . Fortunately, for realistic parameters the radiative lifetime is at least on the order of 10 ms even for very large asymmetries [55], and is not going to be a limiting factor in the near future. On the other hand, the decay rate due to a finite coupling capacitance C_c is proportional to its square C_c^2 and can severely limit the qubit T_1 in typical circuits [145]. This often requires the use of low-pass filtering on the FBL input to protect the qubit from decay [142]. We will touch upon this topic again in Sec. 4.3

when numerically simulating the radiative decay through a 3D-integrated FBL in our device architecture.

3.4. Summary

In this chapter we have seen how implementing flux tunability in superconducting qubits comes with its own set of practical challenges. The method of flux delivery, either on-chip or off-chip, determines how effectively a SQUID can be tuned and how selective the applied flux bias is. Compared to fixed-frequency architectures, flux tuning requires additional physical space on the substrate itself, additional input lines and filtering in the cryostat, and additional control electronics at room temperature. It also introduces new potential channels for energy relaxation and decoherence. On the other hand, it is difficult to ignore the rich physics enabled by flux-tunable qubits, some of which were presented in [Chapter 2](#). It is therefore worthwhile to tackle these practical challenges and reduce the circuit complexity while retaining or improving performance. This is the main motivation behind the development of a novel flux tuning architecture that will be presented in [Chapter 4](#).

The second half of the chapter was devoted to the general description of flux noise, an important limiting factor for qubit coherence in tunable superconducting circuits. Adding an external control parameter inevitably brings unwanted noise with it, which can take several forms. Externally added noise, such as current noise on the FBL itself or from stray magnetic fields in the laboratory environment, needs to be filtered and shielded properly. Universal $1/f$ flux noise, due to microscopic magnetic impurities residing on the substrate itself, has been shown to be ubiquitous in SQUIDs of various geometries and material compositions. The precise source of this noise is the subject of ongoing studies. Its suppression would have an enormous impact on the performance and flexibility of tunable superconducting qubits. In the meantime, SQUID design optimization, as well as engineering of the qubit Hamiltonian to reduce noise sensitivity, can help mitigate the effect of this universal flux noise.

4.

Realizing a tileable off-chip flux control architecture

4.1. Introduction

The present chapter introduces an architecture for 3D-integrated flux control of superconducting qubits. This forms an extension to the previously implemented coaxial circuit quantum electrodynamics (QED) architecture [46, 161], which was primarily developed for fixed-frequency transmon qubits. The main idea consists of utilizing a gradiometric SQUID which can be flux-tuned selectively by applying a spatially inhomogeneous magnetic field using an out-of-plane inductive loop. We refer to this as a *differential flux bias line* (dFBL). This approach keeps the flux control wiring away from the substrate while ensuring individual addressability of each qubit. Crucially, it also preserves the tileability of the coaxial architecture, making it an attractive solution for larger scale circuits.

We design and build a superconducting sample enclosure incorporating four individual dFBL elements, and test its performance cryogenically by measuring a device hosting four coaxial gradiometric transmons. Beyond providing individual flux control, the dFBL can also be used to perform charge control of the transmons by applying microwave drives, providing an all-in-one control element for full qubit manipulation. We carry out measurements both with individual readout lines for each qubit as well as with a 3D-integrated readout multiplexer. Furthermore, we use the qubit itself as a spectrum analyzer to characterize the transfer function of both control and readout lines, leading to a better understanding of couplings and radiative decay contributions.

4.2. The coaxial qubit architecture

We begin by offering a brief overview of the *coaxmon* architecture for circuit QED, which has been developed in our research group over several years [46, 105, 161, 162]. The underlying idea is to use concentric transmons [148] that have no electric dipole

moment, and to couple them capacitively to an out-of-plane control line formed by a terminated coaxial cable, as depicted in Fig. 4.1 (a). Readout of qubits is performed via standard dispersive coupling to a linear resonator [39], here implemented as a coaxial LC spiral resonator patterned on the opposing side of the substrate hosting the transmon. Circuits are fabricated entirely out of aluminium thin films on either sapphire or high-resistivity undoped silicon as the substrate material, providing a clean dielectric medium through which the transmon and resonator electrodes can couple capacitively. The readout resonator is probed through a separate coaxial line. The arrangement of all these components along the Z axis (perpendicular to the substrate surface) makes this single-qubit unit cell easily tileable along the XY plane, offering a simple blueprint for building large-scale two-dimensional square lattices of coupled qubits while maintaining addressability of the individual qubits and avoiding crowding of control lines on the substrate. Crosstalk between neighboring qubits and control lines is suppressed by design thanks to the absence of a dipole moment in the concentric transmons which weakens coupling to distant electric fields.

The coaxmon architecture has so far been implemented experimentally at a scale of up to 16 qubits [163]. The substrate hosting the quantum circuit is packaged in a micro-machined aluminium enclosure with inductive shunts for the removal of spurious cavity modes [164]. This ensures a clean microwave environment with low levels of unwanted qubit-qubit crosstalk at arbitrary scales. The coaxial control and readout lines are embedded in the enclosure. Since no galvanic contact is required between the qubit chip and the lines, no wire-bonding or similar steps are required and the chip can be easily removed from the enclosure and re-utilized later.

4.2.1. Adding flux control

The initial implementation of the coaxmon architecture was based on single-junction, fixed-frequency transmons that do not require any flux-tuning to be operated. As discussed in Chapter 2 and Chapter 3, magnetic flux tunability is a powerful resource in superconducting circuits that opens up a wide variety of possible qubit designs and interactions schemes. It is therefore desirable to add this feature to the coaxmon. Important exploratory work in this direction has been carried out in our research group by J. Rahamim [165], which is summarized in the following.

To retain the tileability of the architecture, the flux control must be implemented outside of the device plane. Ideally, magnetic flux should also be applied without requiring an electric current flowing on the qubit device itself, avoiding the addition




Figure 4.1: Overview of the coaxmon architecture with extension for flux control. (a) Schematic drawing of a single unit cell of the coaxial circuit QED architecture. Control and readout ports consist of exposed inner conductor of an open-ended coaxial cable inserted into a micromachined aluminium enclosure hosting the device substrate. Qubit and readout resonator are patterned on opposing sides of the substrate, and are coupled capacitively through it. (b) Introduction of flux tunability in the coaxmon, developed in Ref. 165. A gradiometric SQUID transmon is controlled via a differential flux bias line (dFBL) fabricated out of a section of inner conductor of coaxial cable shorted to its outer conductor. (c) The dFBL design variation studied in this work. A short section of inner conductor is exposed and bent into a U shape to create an inductive loop. The bias current flows exclusively in the inner conductor, and both ends are connected to an isolated external DC source.

of galvanic contacts. As already discussed in Sec. 3.2, beyond using a large external magnetic coil (which denies individual addressability of qubits), solutions exist for applying magnetic fields remotely using small inductive loops placed at a fixed distance above the substrate [142, 143]. These off-chip flux bias lines (FBLs) are patterned lithographically on a separate substrate, which is then brought into proximity of the qubit(s) to be tuned. Given that the coaxmon architecture already uses coaxial lines for control and readout, here we choose to fabricate FBLs simply by bending a section of exposed inner conductor and soldering it to its own outer conductor. This is shown schematically in Fig. 4.1 (b). The current is applied to the inner conductor and then flows through the loop and returns via the outer conductor, reaching ground via the outer housing of the SMA connector at the end of the coaxial cable. This loop replaces the previous coaxial control port inside the device enclosure. To maximize the mutual inductance between qubit and FBL, we take advantage of the particular symmetry of the generated magnetic field by using a gradiometric SQUID [146–148], which is tuned by a magnetic field gradient

rather than a homogeneous (global) field. We call this loop design a *differential flux bias line* (dFBL), owing to the fact that only a differential but not a global field is applied on the device surface (see also field simulations in [Sec. 4.3](#)).

A first dFBL implementation [165] has been shown to successfully flux-bias gradiometric coaxial transmons with mutual inductances up to $M \approx 20$ pH and with low levels of DC flux crosstalk between neighboring qubits and dFBLs, on the order of $1 - 3 \times 10^{-3}$. Furthermore, the dFBLs can also be used for fast flux tuning, with a line response time on the order of 1 ns, sufficiently fast for the operation of fast flux-based gates.

In this work, we want to further develop this flux control scheme. One drawback of the original implementation is the fact that bending the loops into shape and soldering them is a difficult manual task that does not lead to reproducible shapes of the dFBLs. A more reliable method for shaping the loops and aligning them inside the sample enclosure is required. Furthermore, preliminary coherence measurements on a tunable transmon in this architecture give a Ramsey coherence time $T_2^* \approx 2$ μ s at the flux-insensitive point (flux sweet spot) of the transmon and a rather large flux noise amplitude $\sqrt{A_\Phi} \approx 270$ $\mu\Phi_0$ when biased away from the sweet spot. This suggests that mitigation of flux noise will be necessary in order to achieve acceptable coherence times, especially if implementing gate schemes where the qubits are required to be temporarily kept at finite flux bias.

We propose to address the former issue by modifying the dFBL design as shown in [Fig. 4.1](#) (c). By removing the outer conductor and dielectric from a short section of coaxial cable and bending the exposed inner conductor into a U shape, a simple loop can be formed. This can then be connectorized at both ends using standard SMA parts or similar components. This creates a well-defined path for the current to flow, with no potential resistive elements (such as a solder joint) in the immediate proximity of the qubit. We will show that dFBLs can be fabricated reproducibly with this method, leading to a small variation of the mutual inductances to qubits.

The issue of mitigating flux noise is less straightforward, as it can originate from various sources (see [Sec. 3.3](#)) – microscopic magnetic impurities on the chip, insufficient shielding from the magnetic environment of the experiment, or from current noise on the FBL itself. From the point of view of the dFBL design, we choose to increase the vertical separation h_{dFBL} between qubit and loop to reduce M to about 1 – 2 pH. This has two main consequences: it reduces the coupling strength of current noise from the dFBL into the qubit SQUID, while increasing the amount of current required to thread one flux quantum through the SQUID.

We believe this to be a reasonable trade-off, as the current travels mostly along a superconducting DC wire inside the cryostat, which should lead to a negligible increase in heat load with larger currents. Furthermore, in our dFBL design the two ends of the loop can be connected to a twisted pair of superconducting DC wires in the cryostat, providing built-in low-pass filtering of the current signal, potentially aiding in reducing thermal and instrumentation-induced flux noise.

4.2.2. Gradiometric transmon design

Before presenting the design and manufacturing process of the improved dFBL architecture, we briefly discuss a few design aspects of the gradiometric coaxial transmons fabricated and measured in this work, shown schematically in Fig. 4.2. This represents a slight modification of the original gradiometric coaxmon [165]. We choose the radii $r_{\text{int}} = 150 \text{ }\mu\text{m}$, $r_{\text{out1}} = 300 \text{ }\mu\text{m}$, $r_{\text{out2}} = 450 \text{ }\mu\text{m}$ defining the coaxial electrodes. While the previous design had very thin leads connecting the Josephson junctions to the electrodes ($w_l = 2 \text{ }\mu\text{m}$), we widen the leads to $w_l = 20 \text{ }\mu\text{m}$ and add a taper towards the junction. The reason for this is to minimize the generalized aspect ratio ρ of the SQUID, defined in Eq. (3.8). As discussed in Sec. 3.3.3, the aspect ratio of a SQUID is found both theoretically and experimentally to correlate with its sensitivity to microscopic magnetic fluctuators. With this simple modification, we can reduce ρ roughly from 320 to 100, making the qubit less sensitive to local flux noise. Separately, there is experimental evidence that wide-tapered junction leads might be an effective way of decreasing the quasiparticle tunneling rate through the JJ, reducing associated decay and dephasing events [166].

Furthermore, we incorporate the widened leads into the photolithography pattern (orange color in Fig. 4.2) to reduce the surface area of substrate that is directly exposed to argon ion milling during the subsequent Josephson junction fabrication step. This is motivated by findings that show how ion milling increases the surface roughness of silicon, which in turn can negatively impact quality factors of superconducting resonators and qubits [45, 167]. Thanks to this design change, the pattern defined by electron-beam lithography (EBL, blue color in Fig. 4.2) is now mostly deposited on top of existing metal, except for the Josephson junctions themselves and a short section of the leads. More details on the fabrication process can be found in Appendix A, and example images of circuits fabricated with this method are shown in Fig. 4.3. We also design a variation of the transmon in which the gradiometer orientation is horizontal rather than vertical (i.e. perpendicular to the orientation of the JJ double-angle evaporation), shown in Fig. C.1 in

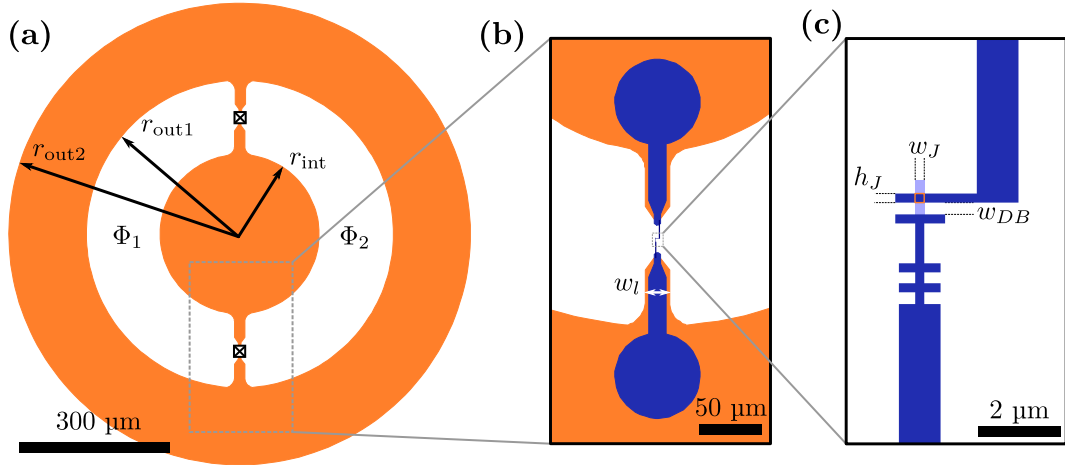


Figure 4.2: Gradiometric transmon design. (a) Schematic drawing of the coaxial gradiometric transmon fabricated and measured in this work, derived from previous implementations [148, 165]. The inner and outer concentric electrodes are connected via two Josephson junctions, forming a SQUID gradiometer with two loops of equal size threaded by external flux Φ_1 and Φ_2 , respectively. (b) Josephson junction patterns (blue) are defined by electron-beam lithography (EBL) and deposited on top of the transmon electrodes defined by photolithography (orange). The large overlap area ensures good electrical contact. Rounded edges and tapering are implemented where possible to minimize electric field concentration. (c) Detailed drawing of the Josephson junction EBL pattern (dark blue), showing the critical junction dimensions h_J and w_J , as well as the Dolan bridge width w_{DB} . Shadow evaporation creates an extension of the lower electrode (light blue), forming the junction overlap area (orange box).

the appendix. This is to allow us to fabricate transmons with different relative gradiometer orientations on the same substrate in the same junction fabrication step.

The effective transmon shunting capacitance C to ground can be estimated using a 3D finite-element electromagnetic solver such as ANSYS Maxwell. To obtain a reliable prediction, we embed the transmon and silicon substrate inside an aluminium box resembling the actual experimental enclosure (see next section). This is because the large coaxial electrodes have a significant capacitance to the grounded enclosure. We remove the Josephson junctions and only simulate the large-scale features of the transmon, i.e. those defined by photolithography. We define the three circuit nodes i , o , and g (inner and outer electrodes, and ground), and calculate the total capacitance matrix. The effective capacitance C of a single reduced node to ground can be obtained from the full matrix as

$$C = \frac{C_{ig}C_{og} + C_{ig}C_{io} + C_{og}C_{io}}{C_{ig} + C_{og}}, \quad (4.1)$$




Figure 4.3: Fabricated coaxial resonators and gradiometric transmons. Optical images of (a) LC spiral resonator, (b) gradiometric coaxial transmon, and (c) detail of Josephson junction. Resonators and qubits are fabricated out of aluminium (bright grey) on opposing sides of a high-resistivity silicon substrate (dark grey). A single step of photolithography defines the entire spiral resonators, while the combination of photolithography with EBL and ion milling defines the transmons.

where C_{ig} denotes the capacitance between nodes i and g , and so forth. For our chosen transmon design we obtain an effective capacitance $C \approx 92$ fF, or equivalently a charging energy $E_C \approx 210$ MHz. Note that the wide junction leads add a small but noticeable contribution of around 5 fF compared to the bare coaxial electrodes, which will cause concentration of electric field between the lead tips. This could be minimized in the future by using a wider and shallower taper design. The modified leads of the horizontal gradiometer design mentioned earlier increases the total capacitance only slightly by around 1–2 fF compared to the reference design, which we do not correct for.

The value of C determined here is only an approximation, as the Josephson junctions themselves provide a small additional capacitance, and we neglected the presence of the readout resonator which has a significant capacitive coupling to the transmon through the substrate. Nevertheless, this simple estimate is found to be within less than 5 % of experimentally measured values (see [Table 4.1](#)).

4.3. Design and construction of flux-tuning sample enclosure

4.3.1. Design and simulation

In the following, we describe the design and manufacturing process of a sample enclosure with dFBLs that can be used to measure a 2×2 qubit device cryogenically. The starting point is the four-qubit micro-machined aluminium enclosure developed and built in Ref. 165. We choose the same 1.5 mm pitch between neighboring qubits, and use the same 5×5 mm² square recess to host the substrate. The only modification required is the lid of the sample enclosure which hosts the dFBLs. A partial schematic drawing with relevant dimensions is shown in Fig. 4.4 (a). The qubits are located at a vertical separation $h_{sep} = 0.2$ mm from the surface of the lid, centered below a cylindrical aperture of diameter $d_p = 1.2$ mm which constitutes the dFBL port. The value of d_p remains unchanged from previous coaxmon sample enclosure designs. The dFBL themselves are aligned inside these ports by drilling two adjacent holes of diameter matching the coaxial cable chosen here (type UT-020-CSP). The vertical separation h_{dFBL} between transmon and the center of the dFBL inner conductor has to be chosen carefully, as it determines the mutual inductance M between the gradiometric SQUID loop and the dFBL. As in the initial dFBL realization [165], neighboring qubits and dFBLs are rotated by 90° to intentionally mismatch the dFBL field orientation and gradiometer alignment, leading to reduced flux crosstalk.

To gain a better understanding of how the differential flux biasing works, we model the magnetostatics of this system by using ANSYS Maxwell, a finite-element (FE) solver. Fig. 4.4 (b) plots the magnitude of magnetic field $|\vec{B}|$ on the substrate surface induced by a constant current of 1 mA flowing through the inner conductor. The transmon is modelled by shorting the two electrodes together, creating a continuous gradiometric ring. The aluminium lid is not shown here but is included in the simulation. The whole enclosure is modelled as a superconductor by setting its relative magnetic permeability close to zero, $\mu_r \approx 0$, leading to full rejection of magnetic field. We see that the field has a figure-eight shape and is concentrated in the vicinity of the transmon. Note that the cylindrical superconducting wall around the dFBL constrains the field laterally and channels it towards the substrate, leading to selective bias of a single qubit. This is similar to the principle of a magnetic hose [141]. The two halves of the gradiometer enclose a total flux Φ_1 and Φ_2 , respectively. By plotting the Z component B_z of the magnetic field in

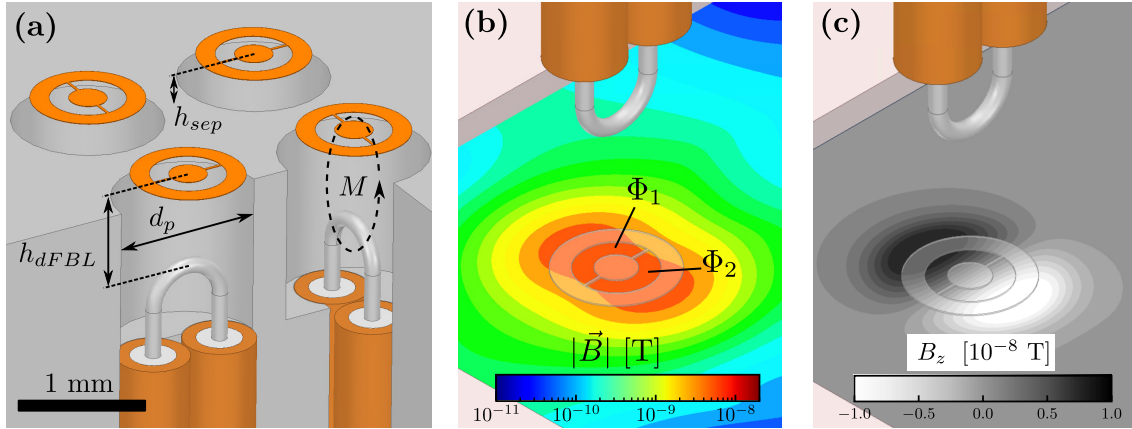


Figure 4.4: Design and simulation of the dFBL. (a) 3D model of the coaxmon dFBL architecture created in Ansys Maxwell for magnetostatic simulations. The qubits are placed a distance $h_{sep} = 0.2$ mm away from the enclosure "lid" hosting the dFBLs, which reside in circular ports with diameter $d_p = 1.2$ mm. The dFBL separation h_{dFBL} is varied to find the optimal value of mutual inductance M . (b) Magnetostatic FE simulation in Maxwell showing the absolute value of magnetic field on the substrate surface, induced by a 1 mA current flowing in the dFBL. The total differential flux threaded through the gradiometric SQUID is obtained as $\Phi_{\Delta} = \Phi_1 - \Phi_2$. (c) Same simulation as (b) but we plot the z component of the field B_z on a linear scale, showing the opposite signs in the two regions of the gradiometer.

Fig. 4.4 (c), we can identify that the field has the same magnitude but opposite sign inside the two halves, leading to the absence of net global flux $\Phi_{\Sigma} = \Phi_1 + \Phi_2 = 0$ but a finite differential flux $\Phi_{\Delta} = \Phi_1 - \Phi_2$. The latter is the relevant quantity that determines the flux bias of the gradiometer. This geometry is therefore capable of efficiently biasing a SQUID without introducing any net flux on the device surface.

Repeating the magnetostatic simulation while sweeping h_{dFBL} allows us to determine how the amount of threaded differential flux Φ_{Δ} varies with separation, see Fig. 4.5 (a)–(b). The value of Φ_{Δ} is obtained by numerically integrating the simulated magnetic field in the two halves of the gradiometer loop. The mutual inductance is then simply calculated as $M = \Phi_{\Delta}/I$, where we fix $I = 1$ mA in the simulation. We can observe an exponential dependence of M with distance, which is analogous to the scaling of capacitive coupling to coaxial ports in the original coaxmon architecture [168]. This highlights the importance of accurately setting h_{dFBL} when assembling the enclosure.

Due to the fact that the transmon electrodes are superconducting but the film thickness of 100 nm of the real fabricated devices is not much larger than the London penetration depth of aluminium [170, 171], we separately consider two limiting cases of the SQUID metal permeability $\mu_{r,Q}$. For $\mu_{r,Q} = 1$, the qubit electrodes are

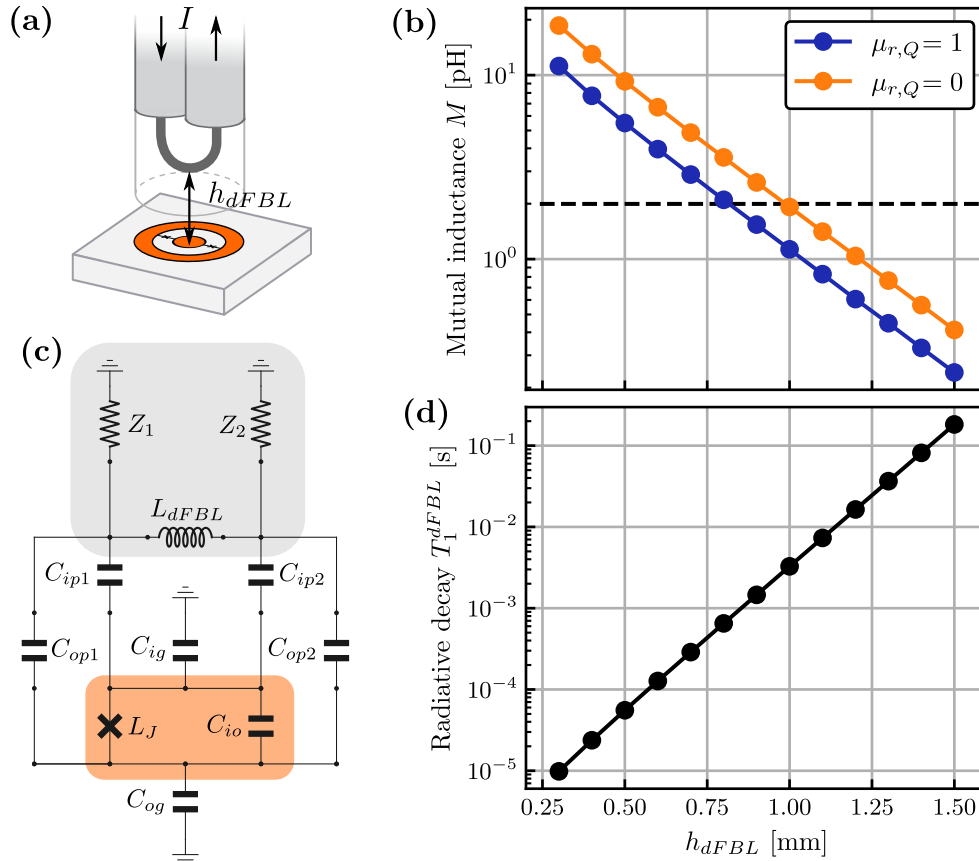


Figure 4.5: Simulation of dFBL mutual inductance and radiative decay. (a) Schematic drawing of the studied system where the vertical qubit-dFBL separation h_{dFBL} is varied. (b) Mutual inductance M between qubit and dFBL as a function of h_{dFBL} , obtained directly from FE magnetostatic simulation in Ansys Maxwell of the 3D model shown in Fig. 4.4. Different results are obtained when assuming screening ($\mu_{r,q} = 0$) or transparent ($\mu_{r,q} = 1$) transmon electrodes. The horizontal dashed line indicates the average value of M in the measured transmon device, see Sec. 4.4. (c) Circuit model of the qubit-dFBL system generated using the QuCAT Python library [169], which allows for the numerical estimation of various quantum circuit parameters, including the finite qubit lifetime due to dissipation in the line impedances $Z_1 = Z_2 = 50 \Omega$. The capacitances are obtained from an electrostatic simulation in Maxwell, and we set $L_{dFBL} = 1$ nH and $L_J = 10$ nH. The resulting transmon has $\omega_{01} \approx 5$ GHz and $\alpha \approx 215$ MHz. (d) Radiative qubit lifetime T_1^{dFBL} versus h_{dFBL} obtained from the QuCAT circuit in (c), assuming no additional internal or external loss mechanisms.

transparent to magnetic field and they do not affect the field lines, whereas for $\mu_{r,Q} = 0$ the field is completely expelled from the electrodes and diverted either into the gradiometer or outside of it. This leads to a focussing effect that increases the total enclosed flux and therefore M in the case $\mu_{r,Q} = 0$, see Fig. 4.5 (b). We note that these are only theoretical limiting cases, and that the real behavior of our circuits is likely to lie in between.

Another important aspect of FBL design, discussed in [Sec. 3.3.4](#), is radiative (Purcell) decay of the qubit excitation through the FBL port. The inevitable presence of a capacitive as well as inductive coupling to the dFBL leads to a possible decay channel that will limit the qubit lifetime to an upper bound T_1^{dFBL} . We can estimate the importance of this decay mechanism from a circuit model implemented using the Python library QuCAT [169], drawn in [Fig. 4.5 \(c\)](#). We approximately model the dFBL loop by a lumped inductance $L_{dFBL} = 1$ nH, with either end symmetrically coupled capacitively to the transmon electrodes. The dFBL is connected to two impedances $Z_1 = Z_2 = 50\Omega$ representing the coaxial lines, which is where the qubit excitation can be dissipated. The capacitive network is calculated as before from an electrostatic simulation in Maxwell, adding a new node p for the inner conductor of the dFBL. The symmetrized capacitances are simply defined as $C_{ip1} = C_{ip2} = C_{ip}/2$ and $C_{op1} = C_{op2} = C_{op}/2$, where C_{ip} and C_{op} are the values obtained from Maxwell. Given all circuit parameters, QuCAT solves for the eigenmodes of the circuit, in this case a single non-linear mode with frequency $\omega \approx 5$ GHz representing the transmon. The solution includes the dissipation rate κ from which we obtain $T_1^{dFBL} = \kappa^{-1}$, plotted in [Fig. 4.5 \(d\)](#). As in the case of M , we observe an exponential dependence of lifetime on the dFBL separation. The lifetime is found to scale roughly as $T_1^{dFBL} \sim 1/C_{ip}^2$, in agreement with analytical treatment of a simple circuit model [145]. Fortunately, it is possible to find a value h_{dFBL} for which T_1^{dFBL} exceeds 1 ms, which is among the highest values currently seen in state-of-the-art superconducting qubits [71], while at the same time offering a practical¹ mutual inductance M around 1–2 pH.

This interplay of inductive and capacitive coupling has been discussed already in the context of a different implementation of off-chip FBL [143], and can be intuitively understood as the balance between the dielectric constant ϵ_0 and magnetic permeability μ_0 of vacuum, which lead simultaneously to desirable levels of inductive and capacitive coupling. In comparison, on-chip FBL designs suffer from the higher dielectric constants of typical substrate materials (sapphire $\epsilon_r \approx 10$, silicon $\epsilon_r \approx 11$, both with $\mu_r = 1$), leading to stronger capacitive coupling with no advantage in terms of inductive coupling.

It is important to point out that our circuit model is only a crude approximation of the real system, and that the estimated radiative lifetime T_1^{dFBL} should therefore be interpreted as a semi-quantitative result. Still, it is useful to understand the scaling with h_{dFBL} . Once superconducting circuits consistently reach and surpass

¹Practical here means to apply a full flux quantum $\Phi_0 = h/2e \approx 2.06 \times 10^{-15}$ Wb with a typical current range of \sim mA, resulting in mutual inductances in the range of pH.

lifetimes of 1 ms (limited by other mechanisms), it will be necessary to decrease the coupling to external lines, which in our case can be easily achieved by only slightly moving the dFBLs further away and without having to modify the qubit circuit and layout. This also means decreasing M and taking larger bias currents into account, which will be acceptable only if resistive heating in the dFBL lines can be kept at a minimum.

4.3.2. Manufacturing of dFBLs

We now turn to the manufacturing and assembly of the dFBL sample enclosure. As the starting material for the loops we choose semi-rigid coaxial cable of type UT-020-CSP, which has both inner and outer conductors made out of silver-plated copper (avoiding any magnetic materials), with PTFE as dielectric and an outer diameter of 0.020 inches (~ 0.51 mm). As can be seen in Fig. 4.4 (a), drawn at scale, the diameter is small enough such that two cables can be placed parallel to each other without increasing the footprint of a single coaxmon unit cell with 1.5 mm pitch. We also expect this dFBL design to be fully compatible with a larger 2 mm pitch and the use of inductive shunts in the enclosure [46, 164], which will be required to suppress cavity modes in larger-scale circuits.

The main steps of the procedure are depicted in Fig. 4.6. A short section of coaxial cable of predefined length is scribed by using a caliper and a clean razor blade under an optical stereo microscope. The outer conductor and dielectric are then removed by carefully cutting them open with razor blade and tweezers. Since the inner conductor is made of copper, direct contact with the steel razor blade should be avoided as it will be damaged very quickly, compromising the silver plating and the conductive properties of the wire. This is by far the most difficult step, which requires a steady hand and a good dose of patience. To produce a large number of dFBLs reliably, this step should definitely be carried out in an automated fashion with purpose-built mechanical tools. Once the inner conductor is exposed, it is bent into a U-shape around a copper wire of the correct dimensions by gently pulling both ends of the loop simultaneously using tweezers or pliers. It is relatively easy to reproduce the same shape and curvature of the loop every time, which should lead to a consistent magnetic field bias. Four dFBLs are inserted into a square plate of oxygen-free copper and are soldered to it to fix them into position. Here it is crucial that the loops are all set at the correct height. We aim for a height that will result in a separation of about $h_{dFBL} \approx 900$ μm in the final enclosure, based on the simulation results shown in Fig. 4.4. The dFBL




Figure 4.6: dFBL manufacturing and sample enclosure assembly. (a) A section of predefined length of coaxial cable (type UT-020-CSP) is gently scribed with a razor blade. (b) The outer conductor is removed, revealing the PTFE dielectric and inner conductor. (c)–(d) After carefully removing the PTFE, the bare inner conductor is formed into a reproducible loop shape by bending it around a copper wire. (e) Four dFBL loops are inserted into a copper holding plate and soldered to it. (f) The soldered dFBL assembly is inserted into the lid of the Al sample enclosure. Dowel pins facilitate alignment and prevent the loops from being bent or damaged. (g) dFBLs are connectorized with standard SMA parts and bulkhead mounts. (h) Device-side view of the dFBLs located in their individual ports. (i) The sample enclosure lid with dFBLs (left) is combined with a standard coaxial qubit enclosure as described previously [165], containing the qubit device in a 5×5 (mm)² pocket and four independent coaxial readout lines.

assembly is then carefully inserted into the aluminium enclosure lid. Female SMA connectors are soldered to the ends of the coaxial cables, which are connected to a bulkhead SMA mount attached to the sample enclosure for ease of use. The final result is an enclosure lid with four dFBLs and eight SMA connectors which can be combined with the bottom half of a previously designed and tested four-qubit coaxmon enclosure. This features a square pocket for holding the substrate, a groove for indium sealing, and individual coaxial readout ports for each resonator.

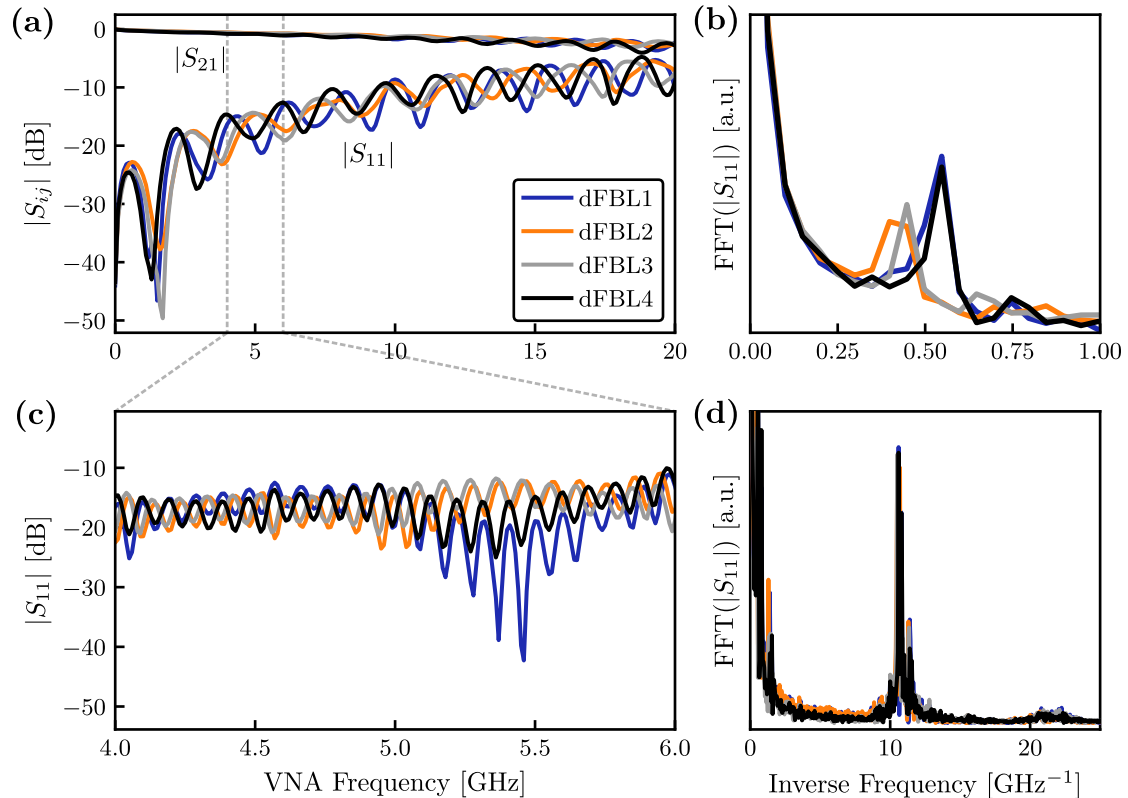


Figure 4.7: dFBL frequency response at room temperature. (a) Wide-range VNA scan of transmission and reflection parameters $|S_{21}|$ and $|S_{11}|$ through the four different dFBLs, measured directly at the SMA connectors visible in Fig. 4.6 (g). (b) Fast Fourier transform (FFT) of $|S_{11}|$ in (a), showing the period of the oscillations in the frequency response. (c) Higher-resolution scan of $|S_{11}|$ in a narrower frequency range and (d) its FFT, showing a fast oscillating component.

An advantage of this dFBL implementation is that both of its ports are accessible, instead of one being shorted to ground, meaning that we can fully characterize its microwave properties by connecting both ends of the dFBL to a virtual network analyzer (VNA). We do this at room-temperature with a fully assembled sample enclosure but without any device substrate present. The magnitude of the transmission coefficient $|S_{21}|$ and reflection coefficient $|S_{11}|$ up to 20 GHz are shown in Fig. 4.7 (a). Even with the impedance mismatch created by the section of missing

dielectric and outer conductor, the dFBLs have insertion loss smaller than -5 dB up to 20 GHz, demonstrating that they can support microwave signals in the range required to drive qubit excitations. This motivates us to study direct qubit Rabi driving via the dFBL in a cryogenic experiment, see next section.

The impedance mismatch is best seen in the oscillation of the reflected signal $|S_{11}|$. By taking the fast Fourier transform (FFT) of each trace in Fig. 4.7 (b), we can observe a distinct periodicity of these oscillations. We attribute this to destructive interference between the input signal and the signal reflected at the loop location. The period Δf of these frequency oscillations can be converted to a physical length scale $l = \kappa c / \Delta f$, where c is the speed of light in vacuum and $\kappa \approx 0.7$ is the typical velocity factor for a coaxial line. For $\Delta f \sim 2$ GHz we obtain a length scale on the order of $l \sim 10$ cm, compatible with the length of coaxial cable used for the dFBLs. The oscillations in the response of dFBL1 and dFBL4 are faster (smaller Δf) than those of dFBL2 and dFBL3, consistent with the fact that the former have longer coaxial lines than the latter, as visible in Fig. 4.6 (g). Zooming into the frequency range 4–6 GHz typical for transmon qubits reveals a much faster oscillating component² in $|S_{11}|$, see data and corresponding FFT in Fig. 4.7 (c)–(d). This can be converted to a length scale of 1–2 m, which leads us to attribute this to mismatches from connectors in the microwave cables used for the VNA measurement.

4.4. Cryogenic characterization

The ultimate test of the dFBL sample enclosure described so far in this chapter consists of carrying out a cryogenic measurement of a superconducting quantum device to verify its performance. We choose a device with four individual (i.e. not coupled) gradiometric coaxial transmon qubits with corresponding readout resonators, as described and shown in Sec. 4.2.2. We label qubits as Q1–Q4 and resonators as R1–R4 matching the labelling convention of dFBL1–4 shown in Fig. 4.6. The Josephson energy E_J of the transmons is defined by the lateral dimensions of the Josephson junction during fabrication, and can be estimated from a room-temperature resistance measurement. The device investigated here is designed to result in qubit frequencies around 6 GHz. Readout resonators

²These oscillations are not visible in Fig. 4.7 (a) due to lower frequency resolution and the use of a larger IF frequency which filters out high-frequency features.

are designed to be clustered around 10 GHz, with 150 MHz between them. See [Appendix A](#) for more details on fabrication.

The chosen device is packaged inside the dFBL sample enclosure and mounted on the lowest-temperature stage of a cryogen-free dilution refrigerator at around 20 mK. A description of the experimental setup, including wiring of the control lines and the measurement chain, can be found in [Appendix B](#). Here we only point out a few details about the dFBL wiring and operation. All dFBLs receive their DC bias via a pair of twisted wires from a superconducting wire loom inside the cryostat, connected to two cryogenic RF bias tees located at base temperature that allow us to add a high-frequency signal to the dFBL. The DC bias current is provided by low-noise DC voltage sources at room temperature. The use of a low-pass RLC filter (cutoff 160 Hz) at the 4K stage of the cryostat adds a $\sim 10\text{ k}\Omega$ resistance to each DC line, effectively converting a voltage source with range $\pm 20\text{ V}$ into a current source with range $\pm 1\text{ mA}$. To test different configurations of microwave (MW) input to the dFBL, we attach both ports of dFBL1 to MW input coaxial lines with small amount of attenuation ($\sim 20\text{ dB}$), adding 1 GHz low-pass filters to protect the qubit from radiation at its frequency³. The other three dFBLs are instead connected at one end to a MW line with 60 dB attenuation, while the other end is terminated with a cryogenic $50\ \Omega$ load on the RF end of the bias tee. Both configurations are shown in [Fig. B.1](#) in [Appendix B](#). Given that the dFBLs themselves are capable of carrying microwave signals with low insertion loss (see [Fig. 4.7](#)), we expect that it should be possible to drive qubit excitations using either of these wiring methods.

4.4.1. Continuous-wave spectroscopy

We begin by performing standard continuous-wave (CW) spectroscopy of resonators and qubits. We use a VNA to send a microwave signal to a coaxial readout port in the sample enclosure, which will inject photons into the resonator when driving at the appropriate resonant frequency. The signal reflected from the readout port is amplified in multiple stages and is then received by the VNA. By observing changes in the measured complex signal S_{21} we can detect the presence of a resonator. Moreover, due to the dispersive qubit-resonator coupling, changing the flux bias on the qubit dFBL will cause both the qubit and the resonator to shift in frequency. Such a flux-dependent resonator measurement is shown in [Fig. 4.8](#) (a) for R3

³We use reflective low-pass filters with insertion loss of around -40 dB to -60 dB at typical qubit frequencies.

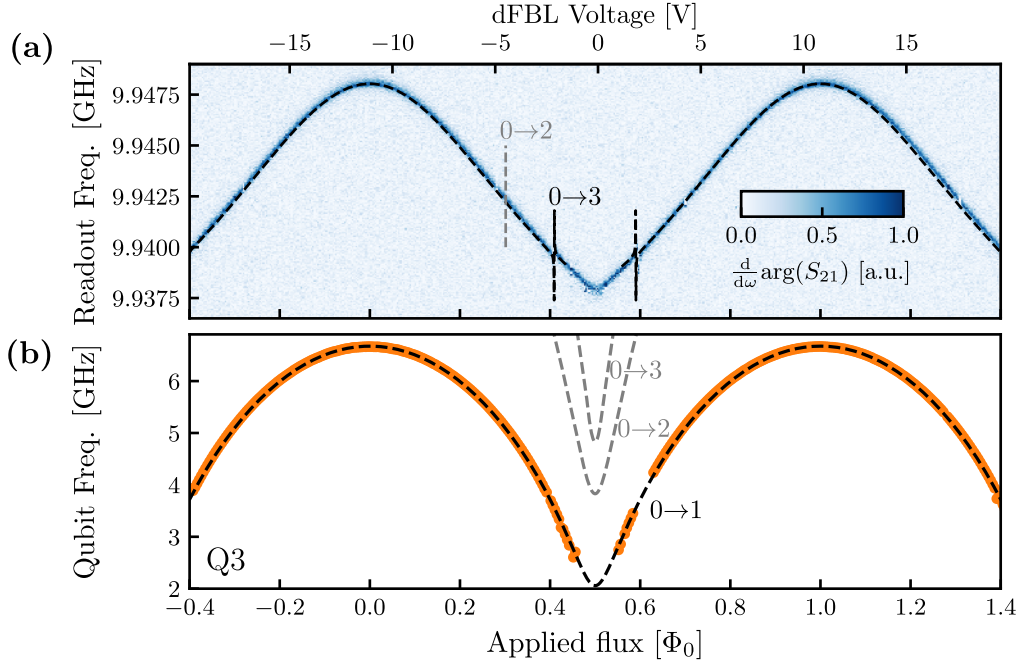


Figure 4.8: Flux-dependent spectroscopy of resonator and transmon. Continuous wave spectroscopy of (a) resonator R3 and (b) qubit Q3 as a function of applied external flux on dFBL3. The flux is controlled by an external voltage source (top axis), which can be converted to a flux scale (bottom scale) by observing the periodicity of the spectrum. Dashed lines are a fit of the full numerical qubit-resonator Hamiltonian to the data, with labelled qubit transitions. Lines and labels shown in grey are not observed directly in the experiment but are added for completion.

when changing the voltage applied to dFBL3. We plot the phase derivative of S_{21} , showing a single clear spectral feature which is periodic in applied voltage (i.e. flux).

Since we know that the spectrum of a tunable transmon has a period of one flux quantum $\Phi_0 = h/2e$, see [Sec. 2.2.3](#), we can convert the voltage scale to a flux scale and add it to the plot (bottom axis). A maximum in resonator (and qubit) frequency corresponds to an integer number of flux quanta threading the SQUID loop. The caveat is that we can only determine the value of applied flux bias up to multiples of Φ_0 . The zero of the flux axis is therefore an arbitrary choice. The conversion between applied bias current I (or voltage V , given the known resistance R in the DC lines mentioned earlier) and threaded flux Φ directly gives us the important quantity M , which was discussed during the design of the dFBL architecture in [Sec. 4.3.1](#), via the relation $\Phi = M \times I = M \times V/R$. In the case of R3 and Q3 we obtain $M \approx 1.9$ pH, which lies within the desired range.

Once the resonator frequency versus flux has been calibrated, we can perform qubit spectroscopy by applying a constant drive to the resonator at its correct

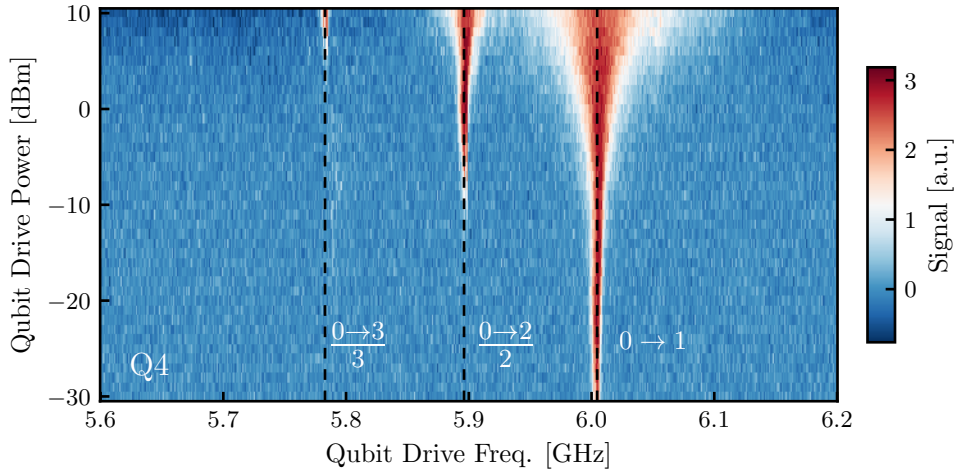


Figure 4.9: Power-dependent qubit spectroscopy. Continuous-wave spectroscopy of Q4 at $\varphi_{\text{ext}} = 0$ as a function of the qubit drive power at room temperature. The dashed lines and labels show the predicted frequencies of the three lowest qubit transitions from the numerical fit to the full flux-dependent qubit spectrum.

frequency while driving a separate tone on the qubit. We do this using an additional RF generator port on our 4-port VNA and applying the signal to the qubit dFBL via the corresponding cryostat microwave line, which also acts as a test of the high-frequency capabilities of the dFBL. At low drive powers, a single spectral feature can usually be distinguished, which we associate with driving the lowest transition $|0\rangle \rightarrow |1\rangle$ of the transmon. The frequency ω_{01} of the identified peak as a function of applied flux for Q3 is shown in Fig. 4.8 (b). This particular qubit has a maximum frequency at its flux sweet spot of 6.674 GHz and can be tuned below 3 GHz before the signal becomes too weak to be read out reliably.

The multi-level nature of the transmon can be probed by performing spectroscopy at a fixed value of applied flux while changing the applied drive power. Such a measurement is shown for Q4 in Fig. 4.9. As the drive power is increased, the main transition frequency (here close to 6 GHz) experiences power broadening, while additional sharper transitions appear at lower frequencies. The narrower spectral width is associated with the driving of multiple photons to excite the transition. The most prominent one is attributed to the two-photon driving of the second transmon excited level $|0\rangle \rightarrow |2\rangle$ at frequency $\omega_{02}/2$, from which we can extract the transmon anharmonicity as $\alpha = \omega_{02} - 2\omega_{01}$ (here -217 MHz).

We perform resonator and qubit spectroscopy for all transmons on the device, and we report the main measured quantities in the top section of Table 4.1. Unless noted otherwise, these correspond to the values measured at zero applied flux $\varphi_{\text{ext}} = 0$ for each qubit. The resonator linewidths κ are extracted from a Lorentzian

Table 4.1: Basic characterization of gradiometric transmon device. Overview of measured and extracted quantities of the four-qubit device characterized using the dFBL sample enclosure. Basic spectroscopic quantities and coherence statistics are obtained at each qubit’s flux sweet spot ($\varphi_{\text{ext}} = 0$). Circuit parameters E_J , E_C , d and g_r are obtained from a numerical fit of the full Hamiltonian to the flux-dependent qubit and resonator spectra. Numbers in brackets denote measurement uncertainty for measured quantities, fit uncertainty for numerically extracted parameters, and statistical variation for the repeated coherence measurements.

Quantity	Unit	Q1	Q2	Q3	Q4
Qubit frequency $\omega_{01}/2\pi$	(GHz)	6.742	5.915	6.674	6.009
Anharmonicity $\alpha/2\pi$	(MHz)	-216(1)	-216(1)	-216(1)	-217(1)
Resonator frequency $\omega_r/2\pi$	(GHz)	9.686	9.814	9.948	10.080
Resonator linewidth $\kappa/2\pi$	(MHz)	0.41(2)	0.92(5)	0.36(1)	0.25(1)
Dispersive shift $2\chi/2\pi$	(MHz)	-2.6(1)	-1.48(1)	-2.29(1)	-1.45(1)
Mutual inductance M	(pH)	2.10	2.05	1.9	1.93
Charging energy E_C/h	(GHz)	0.203(1)	0.201(1)	0.203(1)	0.201(1)
Tot. Josephson En. E_J/h	(GHz)	29.8(1)	23.5(1)	29.3(1)	24.1(1)
Ratio E_J/E_C		147	117	144	120
SQUID asymmetry d		0.15(1)	0.29(1)	0.15(1)	0.28(1)
Resonator coupling $g_r/2\pi$	(MHz)	140(1)	140(1)	142(1)	142(1)
Relaxation time T_1	(μs)	49(6)	36(5)	37(4)	42(9)
Ramsey coherence time T_2^*	(μs)	15(4)	30(4)	13(3)	49(13)
Echo coherence time T_{2E}	(μs)	54(6)	39(5)	30(4)	52(11)
Frequency variation σ_f	(kHz)	12	4.5	12	5
# of repetitions		217	77	219	152
Total run time	(h)	16.5	5.8	16.6	11.5
Flux noise amplitude A_Φ	($\mu\Phi_0$)	6.2(1)	6.7(2)	6.0(1)	6.6(1)

fit of the resonator phase gradient data at low drive power. The dispersive shift χ is a measure of the dispersive coupling strength between qubit and resonator, and is obtained from the shift of resonator frequency when the qubit is in the ground state versus the excited state⁴.

The most important result from the point of view of dFBL design is that all four measured transmons can be successfully flux-biased and that the mutual inductances M all lie very close to 2 pH, in accordance with our design target range. Moreover, the very small standard deviation $\sigma_M \approx 0.1$ pH proves the reproducibility of our chosen dFBL design and manufacturing process. If we assume σ_M to stem

⁴ χ is estimated here from highly averaged resonator spectroscopy due to a small residual excited state population of the qubit at thermal equilibrium, but a more accurate measure is obtained by driving a X_π pulse on the qubit in the time domain.

exclusively from inaccuracy in setting h_{dFBL} during assembly of the dFBL enclosure, we can convert this to a variation in loop heights of $\sigma_h \approx 20 \text{ }\mu\text{m}$ thanks to the simulated scaling in Fig. 4.5. This is to be interpreted as an upper bound, as other factors such as dFBL loop shape variation and tilt of the device substrate inside the sample enclosure pocket can also lead to variations in measured M .

From the resonator and qubit flux-dependent frequency datasets, it is possible to extract all circuit parameters of interest by numerically calculating and fitting the Hamiltonian of the coupled qubit-resonator system [39] to the experimental spectra. We obtain fitted values for E_J , E_C , the SQUID junction asymmetry d and the qubit-resonator static coupling strength g_r , which we add to Table 4.1. The numerical flux-dependent spectra for R3 and Q3 are overlaid to the data in Fig. 4.8, showing excellent agreement. We also add the expected frequencies of higher-order transitions $\omega_{02}/2$ and $\omega_{03}/3$ for Q4 in Fig. 4.9, correctly predicting the experimental values (which themselves are not an input to our fitting routine) and therefore demonstrating the consistency of the model.

We observe that all qubits lie in the strong transmon limit, $E_J \gg E_C$. Here, $E_J = E_{J1} + E_{J2}$ is understood as the total inductive energy of the transmon SQUID. All qubits have a significant SQUID asymmetry, even though both SQUID junctions are designed to be identical during fabrication. While variability in fabrication could explain this, we note that Q2 and Q4 have a horizontally oriented gradiometer (see Fig. C.1) while Q1 and Q3 have a vertical gradiometer (see Fig. 4.3), which seems to correlate with the systematic differences in d between these two pairs. There is also a systematic shift in qubit frequency between the two orientations. From the numerical fits we can see that this is not due to a different capacitive contribution of the gradiometer arms, since E_C (and therefore α) is virtually identical across all qubits, but rather due to large variations in Josephson energies E_J that cannot be explained from the EBL junction designs. This issue should be investigated further by testing different Josephson junction patterns and gradiometer designs.

4.4.2. Coherence and flux noise

After having obtained a general picture of the transmon device and its circuit parameters from frequency domain data, we move to characterization in the time domain. We utilize the electronic control setup for up- and down-conversion of microwave signals described in Appendix B, which allows us to create arbitrary sequences of shaped pulses and send them to the qubits and resonators. For

readout, we generally use a simple square pulse shape of width 5–10 μs at the resting resonator frequency, while for qubits we typically use Gaussian-shaped pulse envelopes with width σ between 50 ns and 100 ns, with truncation of the pulse at 3σ . Given that the qubit spectroscopy presented in the previous section could successfully be carried out by driving the qubit through its dFBL, we also choose to apply time-domain qubit pulses via the same route. We calibrate a full qubit state flip, which implements a X_π rotation in the qubit Bloch sphere, by measuring Rabi oscillations as a function of the qubit pulse amplitude and picking the value of amplitude that maximizes the qubit population inversion. A half rotation, or $X_{\pi/2}$ pulse, is implemented by setting half the X_π pulse amplitude.

With calibrated qubit pulses we can now determine basic quantities such as the qubit relaxation time (or decay time) T_1 , as well as the coherence time T_2 . The latter is quantified both via a Ramsey sequence, resulting in the bare coherence time T_2^* , and via a Hahn-Echo sequence, resulting in the echoed coherence time T_2^E which is insensitive to low-frequency dephasing sources. Example data for all of these three sequences, measured on qubit Q2, are shown in Fig. 4.10 along with their respective pulse schemes. In each case, the total sequence duration is swept by changing the delay time τ as defined in the pulse scheme diagrams. Standard exponential decay curves are fitted to extract relaxation and coherence times. The exponential envelope is found to be a good fit for transmons at the sweet spot, where they are insensitive to flux noise and should therefore not exhibit any Gaussian decay component (see Sec. 3.3.2). In the case of the Ramsey sequence, an intentional detuning of the qubit drive frequency leads to the appearance of sinusoidal Ramsey fringes, which can be used to very accurately determine the qubit frequency ω_{01} .

In the following, data is always collected with the qubits at their flux sweet spots. Individual measurement sequences for T_1 , T_2^* and $T_{2,E}$ are repeated in an interleaved fashion many times for each qubit in order to obtain a statistical distribution for these quantities. In Table 4.1 we report the mean values with standard deviations, as well as the number of individual repetitions for each sequence type and the total duration of the measurement protocol. Measurement runs were performed individually for each qubit, with all other qubits at rest. From the Ramsey fringes we can also determine how much the transmon frequency fluctuates over the course of this measurement run, quantified by the frequency standard deviation σ_f .

The relaxation and coherence times measured in this device compare favourably to the previous dFBL implementation [165], with similar or better T_1 and with

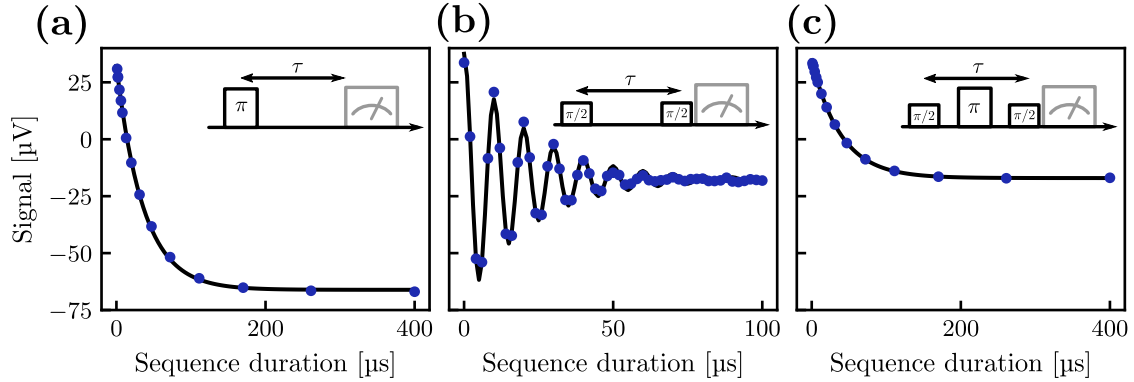


Figure 4.10: Measurement of relaxation and coherence times. Standard measurement sequences for determination of (a) relaxation time T_1 , (b) Ramsey coherence time T_2^* , and (c) Hahn-echo coherence time T_2^E . The insets show the respective pulse sequences used to obtain the data. The sequence duration τ is swept in each experiment to observe the decay of the readout signal. Shown is data obtained for Q2 at the flux sweet spot (blue points) which fits well to exponential decay envelopes (black lines). In the Ramsey experiment, an intentional detuning is applied to the qubit drive which causes the appearance of sinusoidal oscillations. These allow for accurate determination of the qubit frequency. In the Hahn-Echo sequence, an additional X_π echo pulse is inserted at a time $\tau/2$ after the start of the sequence.

coherence times T_2^* and T_2^E improved by one order of magnitude. The increase in coherence time could stem from the reduced mutual inductance M to the dFBL, protecting the qubits from externally injected flux noise, although we should point out that the qubit device presented here is fabricated using a vastly different process and with a modified transmon design, making a fair comparison of dFBL performance rather difficult. Nevertheless, the general improvement in coherence times and especially the achievement of values $T_2^E \sim T_1$ is a promising sign for this architecture, providing a platform for qubits with good coherence properties for a wide range of circuit QED experiments.

Determination of flux noise amplitude

So far we have mainly characterized qubit performance at the flux sweet spot. Since a variety of gates and interaction schemes involve tuning the qubits away from this point for short periods of time, it is important to quantify how much decoherence they will experience. As discussed in [Sec. 3.3](#), operation at a finite qubit frequency slope $\partial\omega_{01}/\partial\Phi$ will expose the transmons to noise in the applied external flux (usually $1/f$ -type noise), leading to fluctuations in qubit frequency and therefore to enhanced dephasing of the qubit. A $1/f$ noise power spectrum $S_\Phi(\omega) = A_\Phi/\omega$ leads to a Gaussian envelope of the qubit coherence signal, which combined with




Figure 4.11: Determination of $1/f$ flux noise amplitude on Q3. (a) Measurement of T_1 as a function of flux around the upper sweet spot. (b) Measurement of Gaussian pure dephasing time $T_{\phi,G}^E$, extracted from a Hahn echo sequence (see main text for details). (c) When plotted as a dephasing rate against the slope of the qubit frequency curve with respect to flux, the Gaussian dephasing follows a simple linear model (black dashed line) due to $1/f$ noise spectrum according to Eq. (3.3), from which the flux noise amplitude A_Φ can be extracted.

the exponential decay from relaxation (at rate $\Gamma_1 = T_1^{-1}$) and flux-independent pure dephasing due to stochastic noise (rate $\Gamma_\phi = T_\phi^{-1}$) results in the combined decay function in Eq. (3.4). By performing Hahn-echo sequences as a function of flux, and dividing out the contributions due to Γ_1 and Γ_ϕ from the measured decay curves, we can extract the echoed Gaussian dephasing rate $\Gamma_{\phi,G}^E$, which we here attribute exclusively to flux noise. Furthermore, we can calculate the slope $\partial\omega_{01}/\partial\Phi$ from the full numerical spectrum fit obtained previously, as shown in Fig. 4.8. This allows us to then finally extract the amplitude $\sqrt{A_\Phi}$ of the $1/f$ flux noise spectrum via Eq. (3.3).

We perform Hahn-echo sequences as a function of flux for each qubit in the region between $\varphi_{\text{ext}} = \pm 0.25$. At each flux value, resonator frequency, qubit frequency, and X_π pulse amplitude have to be recalibrated. We also carry out T_1 sequences at each flux point instead of just relying on the value measured at the flux sweet spot, in order to take into account possible variations of relaxation time as a function of qubit frequency. In fact, we see in Fig. 4.11 (a) that T_1 can fluctuate significantly⁵ across the flux range. The stochastic pure dephasing rate Γ_ϕ is instead fixed to a single value obtained from the sweet spot T_1 and T_2^E values in Table 4.1 as $\Gamma_\phi = (T_2^E)^{-1} - (2T_1)^{-1}$ assuming a strictly exponential decoherence function. We are therefore treating white-noise dephasing as a constant background term independent of flux. As explained above, this then allows us to extract the

⁵But not randomly – this will be discussed in more detail when performing detailed T_1 spectroscopy in Sec. 4.4.5.

pure Gaussian dephasing time $T_{\phi,G}^E = 1/\Gamma_{\phi,G}^E$ at each flux value, which we plot in Fig. 4.11 (b). The dephasing time steadily and rapidly decreases away from the flux sweet spot, as the qubit frequency slope and therefore its sensitivity to flux noise increases. Equivalently, plotting the same data points as $\Gamma_{\phi,G}^E$ versus $\partial\omega_{01}/\partial\Phi$ in Fig. 4.11 (c) then manifests the linear behavior predicted by Eq. (3.3). We infer that our experimental data is therefore compatible with the assumption of 1/f flux noise being the dominant mechanism for dephasing at finite flux bias. The flux noise amplitudes $\sqrt{A_\Phi}$ extracted from a linear fit are reported at the bottom of Table 4.1 and lie in the range 6 – 7 $\mu\Phi_0$ for all four qubits. These are somewhat higher than the values of *universal* flux noise $\sqrt{A_\Phi} \sim 1 - 2 \mu\Phi_0$ found in the literature [150, 152], suggesting that improvements in either magnetic shielding, noise filtering, SQUID design, or materials and fabrication might be required. On the other hand, this compares well to the much higher value on the order of $\sim 300 \mu\Phi_0$ measured in the first dFBL implementation [165], which can be at least partially explained by the much larger value $M \approx 15$ pH in that experiment if current noise on the dFBL is assumed to be a significant contribution. A more systematic study of instrumentation noise, DC line filtering and dependence of $\sqrt{A_\Phi}$ on dFBL distance should be carried out to hopefully further decrease noise amplitudes and improve qubit coherence away from the sweet spot.

4.4.3. Flux crosstalk

When operating multiple tunable qubits simultaneously, it is important to be able to flux-bias each qubit individually without disturbing those around it, which could lead to unwanted dephasing and to coherent control errors accumulated on idling qubits. We therefore want to quantify the *DC flux crosstalk* in our experiment. We define this as a (positive) ratio of external fluxes Φ_{ii} and Φ_{ij} threaded through qubit i when applying an identical bias current I either to its own dFBL i or to a neighboring dFBL j ,

$$\chi_{ij} = \left| \frac{\Phi_{ij}}{\Phi_{ii}} \right| = \left| \frac{M_{ij}I}{M_{ii}I} \right| = \left| \frac{M_{ij}}{M_{ii}} \right|, \quad (4.2)$$

as shown schematically in Fig. 4.12 (a). χ_{ij} is therefore a ratio of mutual inductances M_{ij} between different elements in our system – we are essentially asking the question of how strongly the *wrong* dFBL can couple to a given qubit. Since we have four qubits and four dFBLs, the individual crosstalk values χ_{ij} can be represented as a 4×4 square matrix $(\chi)_{ij}$, with real positive values and diagonal elements $\chi_{ii} = 1$ by definition. It has already been shown that the coaxial dFBL architecture can

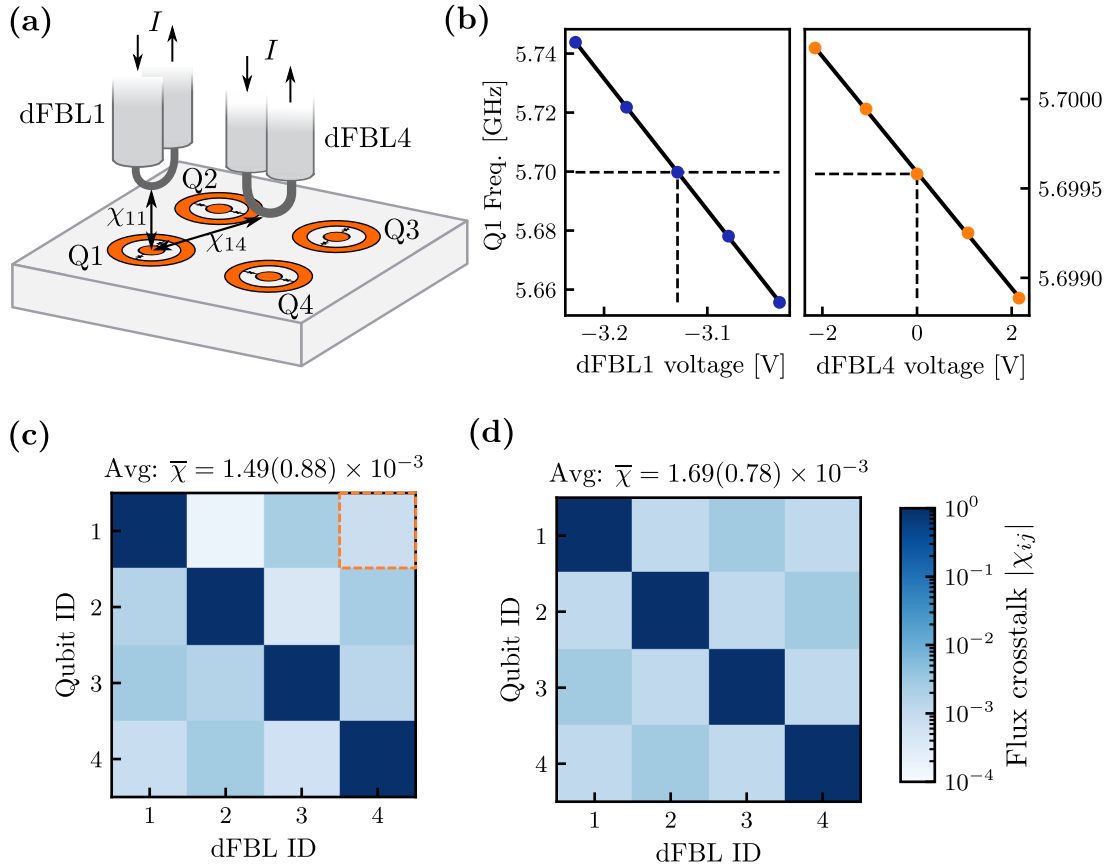


Figure 4.12: Measurement and simulation of DC flux crosstalk. (a) DC flux crosstalk is determined by measuring shifts in qubit frequencies as a function of flux applied via each dFBL. (b) Q1 frequency versus voltage of dFBL1 (left) and dFBL4 (right) measured via Ramsey sequence around the bias point $\varphi_{\text{ext}} = 0.25$ to enhance flux sensitivity. Error bars are smaller than the data points. The ratio of the two slopes obtained from linear fits (black solid lines) is used to extract the flux crosstalk element $\chi_{14} \approx 8 \times 10^{-4}$. (c) Color map of measured crosstalk matrix $(\chi)_{ij}$ between each pair of qubit and dFBL across the whole device. The quoted average value $\bar{\chi}$ does not include the diagonal elements, which are equal to 1 by definition. The matrix element χ_{14} obtained from the data in (b) is highlighted with a dashed box. (d) Simulated flux crosstalk matrix, obtained from magnetostatic simulation in Maxwell. The same color scale is used for both (c) and (d).

achieve individual crosstalk values on the order of $\chi_{ij} \sim 1 - 3 \times 10^{-3}$ [165]. This is attributed to the optimal matching of the dFBL and SQUID geometries, in addition to the intentionally mismatched orientation of neighboring SQUIDs. Here we want to perform a more systematic characterization of the full crosstalk matrix $(\chi)_{ij}$ in our device and sample enclosure, including a direct comparison between experiment and magnetostatic simulations.

The strong dependence of qubit frequency on flux when biased away from the sweet spot can be employed as a very sensitive magnetometer. One can either

use shifts in resonator frequency as a proxy for the qubit, or measure the qubit frequency directly via Ramsey fringes [132]. Here we choose the latter method due to its higher accuracy in determining extremely small shifts. We bias an individual qubit i at a quarter flux bias $\varphi_{\text{ext}} = 0.25$ while keeping all others at their sweet spots, and perform a calibration measurement of its frequency versus its own dFBL bias as shown in the left panel of Fig. 4.12 (b). In a small enough region, the flux dependence of the qubit is approximately linear - the data shown here for Q1 fits well to a linear model, with a coefficient of determination $r^2 \approx 0.99998$. We denote by k_{ii} the slope of the calibration line of qubit i biased through dFBL i . We now repeat the measurement but by keeping dFBL i fixed at the bias point $\varphi_{\text{ext}} = 0.25$ and instead sweeping the bias of a different dFBL j symmetrically around the sweet spot of qubit j , resulting in a different slope k_{ij} of the frequency of qubit i , as shown in the right panel of Fig. 4.12 (b). Notice the larger range of applied voltage and the much smaller measured frequency shift. Since the frequency response is linear in flux in both cases, and the flux itself is linear in applied external voltage, the ratio $|k_{ij}/k_{ii}|$ is simply the ratio of mutual inductances between qubit i and different dFBLs, and it is therefore equal to the definition of χ_{ij} in Eq. (4.2). This is how we can experimentally extract the desired values of crosstalk. For example, from the two slopes in Fig. 4.12 (b) we obtain the value $\chi_{14} = 0.80(1) \times 10^{-3}$.

One should be careful when biasing neighboring qubits, since small residual dispersive shifts between qubits (ZZ crosstalk, as discussed in Sec. 2.4.2) might distort the frequency dependence that we use to extract the linear slope k_{ij} . This is because when biasing dFBL j , not only will qubit i experience a change in applied flux (classical crosstalk), but also a dispersive shift due to the change in qubit j 's frequency (quantum crosstalk). Since the transmons in our device have been designed with no intentional capacitive coupling g_c between them, we can expect this residual coupling to be extremely small thanks to the geometry of the coaxmon [46]. Furthermore, biasing qubit j symmetrically around its sweet spot would lead to the same sign of frequency shift of qubit i for both signs of bias, which would be detectable as a quadratic distortion of the linear dependence in Fig. 4.12 (b). Given that we obtain a good linear fit with $r^2 \approx 0.9998$ when biasing Q1 from dFBL4, we can assume this effect to be small enough to not invalidate the present analysis.

Repeating the same procedure for each combination of qubit and dFBL, we obtain the full experimental flux crosstalk matrix $(\chi)_{ij}$ shown in Fig. 4.12 (c). The average crosstalk across the whole device is found to be $\bar{\chi} = 1.5(9) \times 10^{-3}$

(excluding the diagonal matrix elements) and all values are smaller than 3×10^{-3} , consistent with the preliminary findings from Ref. 165. To better understand these results, we compare them to a predicted crosstalk matrix obtained from the same magnetostatic simulations as performed in Sec. 4.3.1 to obtain the mutual inductance M , see also Fig. 4.4. We simply move and rotate the qubit to one of the neighboring lattice sites and compute the new threaded differential flux. The result is plotted in Fig. 4.12 (d). Due to symmetry considerations, there are only three distinct values of crosstalk, corresponding to the three different types of neighbors that each qubit has: adjacent in X, adjacent in Y, and diagonal. The agreement between simulation and experiment is striking - not only do we obtain a very similar average crosstalk, but the structure of the matrix is also similar. In the experimental data, all of the four diagonal crosstalk elements (i.e. between pairs $1 \leftrightarrow 3$ and $2 \leftrightarrow 4$) are larger than all the remaining elements. This is correctly predicted by the simulation, and can be explained by the fact that the gradiometer orientation of diagonal (next-nearest) neighbours does not match the direction of magnetic field gradient, therefore resulting in a finite differential flux.

In the simulations above, we fixed $h_{dFBL} = 1$ mm and used transparent qubit metal electrodes with $\mu_{r,q} = 1$. We test the dependence of the DC crosstalk on these parameters by repeating the simulations while sweeping h_{dFBL} for both $\mu_{r,q} = 0, 1$ and calculating the average crosstalk value $\bar{\chi}$ in each case. The result is plotted in Fig. C.2 in the appendix and shows, encouragingly and perhaps surprisingly, that flux crosstalk is only weakly sensitive to these parameters. This is in stark contrast with the exponential dependence of the absolute mutual inductance M with separation seen in Fig. 4.5, meaning that in this architecture we can freely optimize for the desired M and T_1^{dFBL} while retaining consistent crosstalk performance.

We note that the localized nature of the dFBL currents in space means that qubits located further away from a given dFBL will necessarily experience weaker magnetic fields, ensuring a well-behaved crosstalk environment in larger-scale devices. This can be seen from the smooth decay of the magnetic field magnitude simulated in Fig. 4.4. The excellent agreement between experiment and simulation in our study suggests that the magnetostatics are well understood and that the performance of future larger-scale systems could be reliably simulated and predicted. In contrast, on-chip FBL designs usually inject the bias current into a device-wide ground plane, leading to complicated return paths for the current across the entire device and therefore to a non-local crosstalk landscape which can be difficult to predict and mitigate [132]. The only published demonstration of individual off-chip flux control

in a multi-qubit device, reported in Ref. [143], similarly utilizes grounded FBLs on a two-dimensional substrate and does not include flux crosstalk characterization, so it is unclear how the issue of return currents extends to off-chip FBL designs. Our present work is thus the first systematic study of flux crosstalk in a fully 3D-integrated flux control scheme for superconducting qubits.

An outstanding question in this dFBL architecture is how the crosstalk performance extends to fast AC flux pulses, although previous studies in 2D circuits suggest that AC flux crosstalk is comparable to or weaker than typical DC crosstalk [26, 132, 136] and should therefore not pose a fundamental obstacle.

4.4.4. Rabi drive calibration

Techniques exist to calibrate the frequency response of cryogenic microwave lines that carry signals to the qubit [172, 173]. This allows for the correction of line distortions that arise from an imperfect transfer function of the different components along the control lines, such as cables, filters, circulators, connectors and solder joints.

In the preceding sections, we have already demonstrated that our dFBLs have sufficient bandwidth to carry qubit-frequency voltage signals and thus drive Rabi oscillations on the qubits, providing charge control in addition to flux control. Here, we want to investigate the frequency response of the dFBLs at cryogenic temperatures more systematically and compare it to the reference VNA data in Fig. 4.7. Knowing the absolute value of a voltage oscillation at the dFBL requires careful calibration of the attenuation of input lines at cryogenic temperatures, which we did not carry out here. Instead, as a combined proxy for the dFBL transmission and its coupling strength to the qubit we simply use the AWG output amplitude A_π (in V) of a Gaussian-envelope pulse of fixed length (usually $\sigma = 50$ ns) required to drive a X_π rotation on the qubit, effectively using the qubit itself as a detector. By performing Rabi drive calibration as a function of qubit frequency ω_{01} we then obtain a calibration curve $A_\pi(\omega_{01})$ such as the one shown in Fig. 4.13 (a) for Q2, which gives us a relative measure of how strongly the qubit is coupled to its dFBL. A larger value of A_π can be associated with higher losses in the dFBL wiring or with weaker dFBL-qubit coupling, since a stronger signal is required to induce the same effect (a X_π rotation), and vice-versa. We also add the previously obtained room-temperature VNA trace of $|S_{11}|$ for a qualitative comparison of the frequency response. We can see that the rough trend of the two curves is similar, but the oscillations of the room-temperature data are roughly twice as fast as in

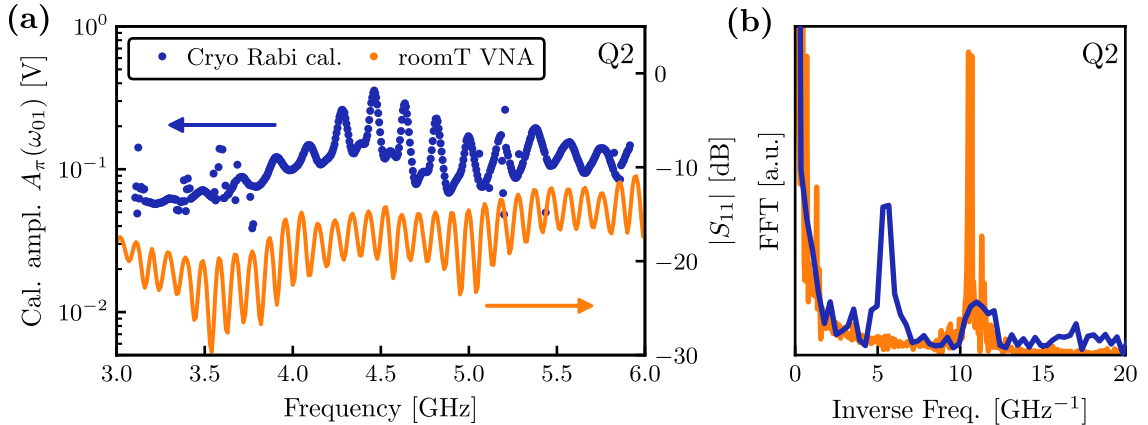


Figure 4.13: Calibration of Rabi drive through dFBL. AWG drive amplitude required to perform a X_π pulse on Q2 as a function of its frequency (blue points). A larger value corresponds to weaker coupling to the qubit. The reference VNA data from Fig. 4.7 is added for comparison (orange). (b) FFT of the two traces in (a), revealing the period of frequency oscillations.

the cryogenic measurement. We can see this more clearly by plotting the FFT of both traces in Fig. 4.13 (b). The cause of these different oscillations is not precisely known, but as discussed in Sec. 4.3.2 they could be due to impedance mismatches and reflections in coaxial cable sections of length on the order of a meter. A numerical simulation of such a system, including the impedance mismatches of the dFBL, could possibly allow us to better understand the difference between these datasets.

We repeat this measurement and analysis for the other three qubits, which can be found in Fig. C.3 in the appendix. There is generally no clear agreement between the cryogenic and room-temperature data. The FFT peaks are located at different inverse frequencies from qubit to qubit, which might again be related to varying coaxial cable lengths of our dFBL wiring.

In a similar fashion, we can also drive Rabi oscillations via the coaxial readout ports, which have a small residual coupling to the qubits through their readout resonators and the substrate. This method has already been used to drive qubit transitions in previous experiments within our coaxmon architecture [165]. Qubit and readout drives are simply added together at room temperature using a power combiner. The only caveat is that the readout lines contain microwave components such as filters and circulators that are tailored to the 8–12 GHz band, meaning that signals at qubit frequencies are located out of band and will experience increased distortion and losses. The resulting Rabi calibrations for Q2, Q3 and Q4 are found




Figure 4.14: Calibration of Rabi drive through readout port. AWG drive amplitude required to perform a X_π pulse on qubits Q2, Q3 and Q4 when driven through the coaxial readout (RO) lines of their respective resonators.

in Fig. 4.14, again showing a clear oscillatory behavior with varying periodicities. No data was collected for Q1 in this measurement.

We note that in general, different pulse widths and RF generator powers have to be chosen depending on the drive port used. This is due to the limited amplitude range of the AWG and the wide variations in coupling strengths of the different coaxial lines. As a result, the absolute values of $A_\pi(\omega_{01})$ cannot be compared directly between different qubits, or between dFBL and readout port for the same qubit.

4.4.5. T_1 spectroscopy

With a detailed calibration of the X_π pulse at each qubit frequency, it is now straightforward to obtain the energy relaxation T_1 as a function of ω_{01} , similar as was done in Fig. 4.11 (a). Such measurements of frequency-dependent T_1 can provide useful information on the loss and decay mechanisms that affect the qubits in a given experiment [54, 63]. We collect data individually for each qubit while leaving the others at rest. Each dataset consists of between 1000 and 2000 individual T_1 sequences measured over durations between 16–24 hours. We collect these sequences in randomized qubit frequency order to avoid distorting the obtained T_1 spectrum systematically in the case of any slow temporal drifts of the qubit parameters or the electronic setup. The result is plotted in Fig. 4.15. We see values fluctuating significantly in the range 20–100 μs for most qubits, with a few values beyond 100 μs at frequencies below 4 GHz, but the fluctuations are not random –

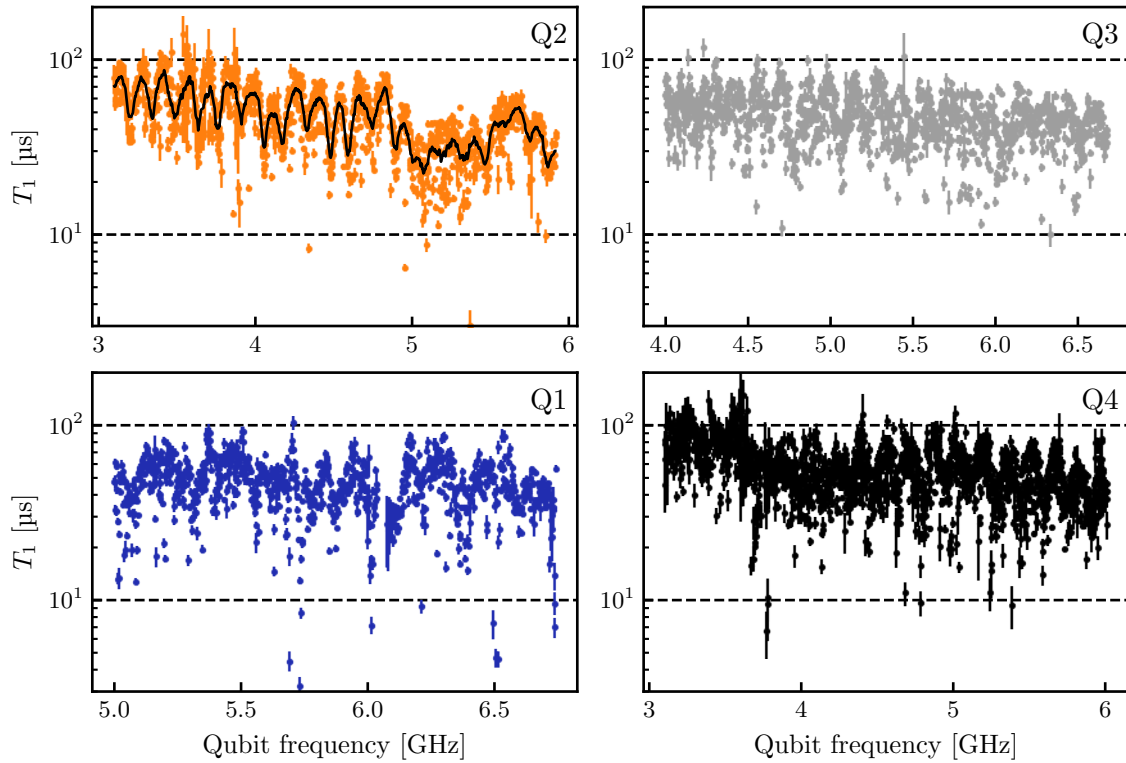


Figure 4.15: Frequency dependence of T_1 . Measurements of energy relaxation time T_1 versus qubit frequency performed individually on each qubit. The black line overlaid to the Q2 data is a running average obtained by applying a square box filter with width $n = 10$ to the raw data points.

there is a visible structure to the data. This can be better visualized by overlaying a running average trend to the data, as shown in the case of Q2. Since temporal correlations can be ruled out due to random sampling along the frequency axis, we consider this to be a stable spectrum of qubit energy relaxation, telling us at which frequencies the qubit is more or less protected from energy decay channels.

Given the oscillations visible in the T_1 datasets, it is natural to compare these to the previously obtained Rabi amplitude calibration curves from Sec. 4.4.4, which show a similar pattern that we attributed to the transfer function of the microwave lines used to drive the Rabi pulse. We therefore plot in Fig. 4.16 the FFT of the following datasets for each qubit: Rabi calibration through dFBL (Fig. 4.13 and Fig. C.3), Rabi calibration through readout port (Fig. 4.14, except Q1), and T_1 data (Fig. 4.15). Strikingly, we observe a clear correlation between the readout port Rabi calibration and the T_1 data, since their FFTs show a single prominent peak at the same frequency period, while the dFBL calibration has a clearly distinct component. Furthermore, one can see that a local maximum in T_1 corresponds to a maximum

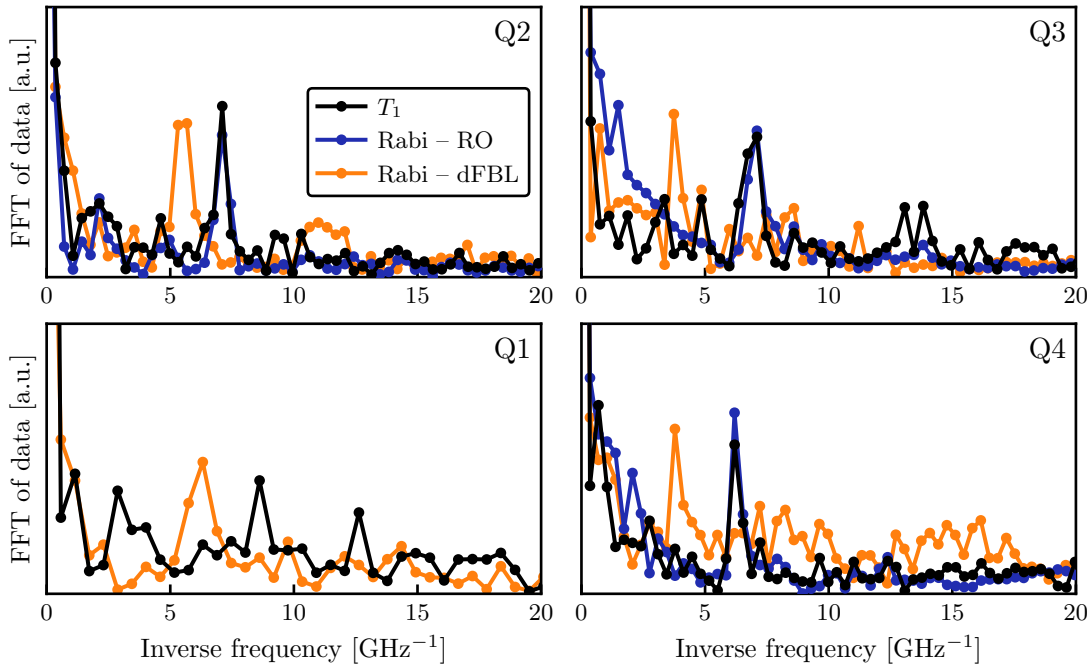


Figure 4.16: Comparison of T_1 and Rabi calibration data. FFT of T_1 versus frequency (black), Rabi calibration curves $A_\pi(\omega_{01})$ through readout port (blue) and through dFBL (orange) for each qubit. No readout port calibration was obtained for Q1.

in readout port Rabi pulse amplitude by suggestively overlaying the full datasets on top of each other, see Fig. C.4 in the appendix. There is excellent agreement between these two datasets for all three qubits for which data is available. Since the dFBLs and the coaxial readout lines in our experiment are passive reciprocal elements, i.e. we expect transmission of signals to be the same in both directions, we conclude that qubit relaxation in our experiment is strongly affected by radiative decay into the coaxial readout port. This is what we commonly refer to as Purcell decay [54, 174]. When transmission is higher, a weaker external signal can excite the qubit – conversely, an excitation in the qubit can more easily decay when the coupling to the environment is enhanced. This offers a relatively clear path forward for increasing qubit lifetimes in our architecture: Purcell decay has to be limited either by reducing the external resonator coupling κ or the dispersive shift χ , at the expense of readout efficiency, or the qubit mode needs to be additionally protected through the inclusion of a Purcell filter [175, 176]. First steps have already been taken in this direction by making the qubit readout circuitry more controlled and reproducible via the use of a micro-machined multiplexing cavity, as will be presented in Sec. 4.5.

On the other hand, the absence of correlations between T_1 and the dFBL Rabi curve for any of the qubits suggests that Purcell decay through the dFBL is not a significant contribution to T_1 at the current level, as predicted by the radiative lifetime simulation T_1^{dFBL} in Fig. 4.5. This is corroborated by the fact that, especially in the case of Q2 and Q4, the dFBL transmission tends to be stronger at lower frequencies (see Fig. 4.13 and Fig. C.3), where higher qubit T_1 values are usually measured. This makes enhanced decay through the dFBL an unlikely mechanism.

The important conclusion to draw here is that our dFBL design successfully realizes a 3D-integrated charge and flux coupling element for coaxial superconducting qubits that does not limit their lifetime, at least not at the $T_1 \sim 50 - 100 \mu\text{s}$ level observed in our experiment. Going to a radiative limit of $T_1^{dFBL} \sim 1 \text{ ms}$ and beyond might require trading off flux and charge tunability for increased lifetime by increasing the separation of the dFBL or modifying its geometry, as was already discussed in Sec. 4.3. On the other hand, this will lead to larger required signal amplitudes travelling through the microwave lines in the cryostat, adding further thermal load. This is ultimately still an unsolved challenge in our field, and there is ongoing work on trying to break this trade-off between lifetime and manipulability of superconducting qubits [177, 178].

4.5. Readout multiplexing

In the experiment presented so far, each readout resonator is probed through independent coaxial lines. This becomes cumbersome very quickly when increasing the qubit count. The additional input and output lines take up invaluable physical real estate in the cryogenic setup and add unwanted thermal load [137]. It is therefore desirable to reduce the number of these lines to improve scalability. Fortunately, device-level multiplexed dispersive readout of qubits is common practice in superconducting quantum circuits. Early demonstrations of two-qubit interactions in circuit QED already involved qubits coupled to a shared co-planar waveguide (CPW) transmission line for dispersive readout [42, 179]. In more recent implementations, each qubit usually has its own dedicated readout resonator, but several of these resonators are coupled to a common feedline to allow simultaneous measurement via frequency-multiplexing of the readout tone [180, 181]. This can be combined with the usage of Purcell filters to protect qubits from radiative decay through the readout line while enabling fast measurement.

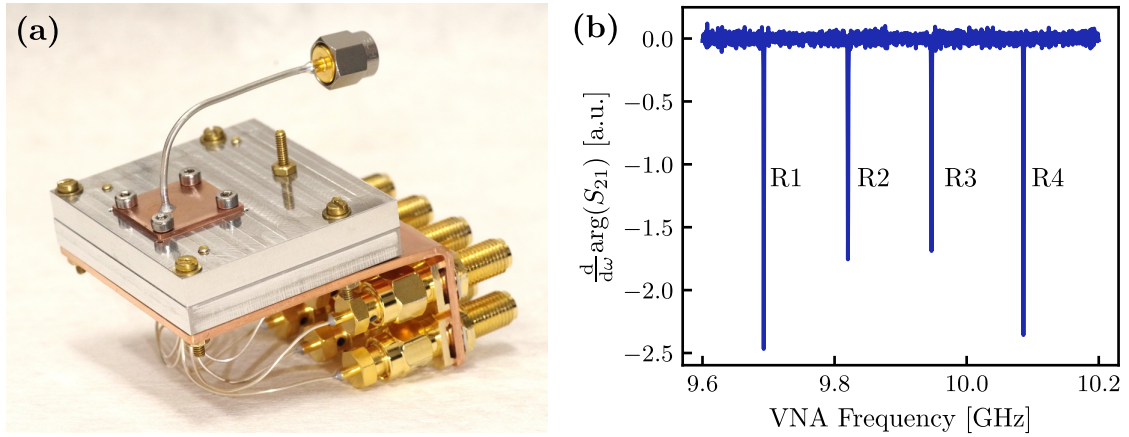


Figure 4.17: 3D-integrated multiplexed readout. (a) Picture of modified sample enclosure incorporating a 3D-integrated readout multiplexer cavity, located underneath the square copper plate. A single coaxial line allows for simultaneous readout of a four-qubit device. The dFBL enclosure lid (bottom side) is the same as presented in Fig. 4.6 and related text. (b) Continuous-wave resonator spectroscopy at cryogenic temperatures, showing all four readout resonators.

A solution for incorporating readout multiplexing into the coaxmon architecture while maintaining full 3D-integration and scalability has recently been developed in our research group by Dr. Mustafa Bakr, with details to be published separately in future work. In summary, we use a micro-machined three-dimensional cavity with a fundamental mode centered around ~ 10 GHz to capacitively couple four readout resonators on the device to the same external coaxial line for readout. A picture of the modified sample enclosure is shown in Fig. 4.17 (a). The cavity is located underneath the square copper plate, with a single coaxial line that can be connected to the same readout signal chain as before. We re-utilize the same qubit control side as before, i.e. the sample enclosure lid with dFBLs that was presented in Sec. 4.3 and used in the previous cryogenic experiment of Sec. 4.4.

In the following, we want to apply some of the experimental techniques discussed so far in this chapter to study and characterize the readout multiplexer assembly. The same four-qubit device measured previously is mounted in this new enclosure and is cooled down again in the same cryostat as before. Since all four readout resonators are now coupled to the same readout line, we can easily measure them in a single VNA trace as shown in Fig. 4.17 (b). This allows us to use a single AWG waveform with a frequency-multiplexed square readout pulse to populate all resonators simultaneously, greatly improving the hardware efficiency of our experiment and also simplifying the wiring inside the cryostat, as cryogenic

Table 4.2: Basic characterization of gradiometric transmon device with multiplexed readout. Overview of measured quantities from the four-qubit device characterized using a modified dFBL sample enclosure with multiplexed readout. Basic spectroscopic quantities and coherence statistics are obtained at the flux sweet spot ($\varphi_{\text{ext}} = 0$) of each qubit. Numbers in brackets denote measurement uncertainty for measured quantities and statistical variation for the repeated coherence measurements.

Quantity	Unit	Q1	Q2	Q3	Q4
Qubit frequency $\omega_{01}/2\pi$	(GHz)	6.730	5.908	6.689	6.008
Anharmonicity $\alpha/2\pi$	(MHz)	-215(1)	-215(1)	-216(1)	-216(1)
Resonator frequency $\omega_r/2\pi$	(GHz)	9.693	9.821	9.953	10.086
Resonator linewidth $\kappa/2\pi$	(MHz)	0.26(1)	0.50(2)	0.50(2)	0.29(1)
Dispersive shift $2\chi/2\pi$	(MHz)	-2.6(1)	-1.4(1)	-2.3(1)	-1.5(1)
Mutual inductance M	(pH)	1.92	1.76	1.58	1.73
Relaxation time T_1	(μs)	20(4)	30(9)	10(1)	29(5)
Ramsey coherence time T_2^*	(μs)	18(3)	24(6)	13(1)	27(5)
Echo coherence time T_{2E}	(μs)	29(5)	32(8)	17(1)	36(7)
Frequency variation σ_f	(kHz)	2	6.8	2.3	5
# of repetitions		156	160	95	292
Total run time	(h)	6.8	8.2	4.9	15

switches are not required anymore. This new modified wiring configuration is shown in Fig. B.2 in Appendix B.

It is interesting to investigate how basic device properties might have been affected by the thermal cycle and by the change in sample enclosure and readout mechanism. The main quantities from continuous-wave spectroscopy are summarized in Table 4.2. Comparing these with the previous dataset in Table 4.1, we can notice a few important differences. Firstly, the external resonator linewidths κ are different, as expected from the drastic modification of the readout circuitry. Moreover, the values are very similar between the pairs (R1, R4) and (R2, R3). Looking at the symmetry of the phase derivative response peaks in Fig. 4.17 (b), we speculate that this could be due to the multiplexer cavity frequency being centered around ~ 9.9 GHz, leading to stronger external decay of R2 and R3 compared to R1 and R4 as they are closer to the center frequency of the common mode.

Qubit frequencies have shifted by up to 15 MHz, which is not an uncommon observation when thermally cycling transmon qubits. More interestingly, resonator frequencies have all systematically increased by 5–7 MHz – this could be explained by a slightly lower capacitive loading of the resonators due to the new sample enclosure. In a similar fashion, the now different values of mutual inductance M between qubits and dFBLs might be caused by changes in vertical qubit-dFBL

separation h_{dFBL} . The decrease in average M from 2.00(8) pH to 1.75(12) pH between experimental runs could be explained by an increase in average separation by $\sim 50 \mu\text{m}$ given the simulated scaling in Fig. 4.5, compatible with the tolerances of machined parts. The larger spread of values could instead be explained by a larger unintentional tilt of the sample substrate inside the enclosure. Overall, these changes do not affect the operation of the device, and instead suggest that similar flux tuning performance can be achieved between runs after disassembling and even drastically modifying the sample enclosure.

Repeated measurements of relaxation and coherence times are also carried out to verify the time-domain characteristics of the device in the new enclosure. The data is added to Table 4.2. Compared to the data collected in the first cooldown and presented in Sec. 4.4.2, relaxation times T_1 seem to have generally decreased. Given the observation made in Sec. 4.4.5 that T_1 is strongly affected by coupling to the readout port, the change in relaxation time is likely due to the newly installed readout multiplexer. This will be further investigated in Sec. 4.5.2. Meanwhile, coherence times also seem to have either remained similar or decreased, albeit without such a clear trend. The Ramsey coherence time T_2^* has decreased significantly (i.e. beyond the widths of the statistical distributions) in only one qubit, and the echoed coherence time T_2^E in the case of two qubits. This is largely accounted for by the aforementioned decrease in T_1 via the standard relation $1/T_2 = 1/2T_1 + 1/T_\phi$, without requiring the assumption of a significant change in pure dephasing time T_ϕ .

4.5.1. Qubit drive calibration with multiplexer

Although the readout multiplexer cavity acts as a bandpass filter, it might still be possible to transmit qubit-frequency signals through the readout port and to the qubits. In the same spirit of the Rabi drive calibration via readout ports performed in Sec. 4.4.4, we therefore calibrate the X_π pulse amplitude for different qubits when driving them via the single multiplexed readout port to observe its frequency response. The result is plotted in Fig. 4.18. Note that while we obtain each dataset independently by sweeping one qubit frequency at a time and keeping all others at their sweet spots, the multiplexer in principle couples to all qubits simultaneously and is therefore not a viable solution for individual qubit control. As in previous datasets, we observe oscillations that we attribute to reflections and impedance mismatches in the line. The general trend is that transmission is weaker at lower frequency, requiring larger voltages to drive the same qubit excitation, consistent

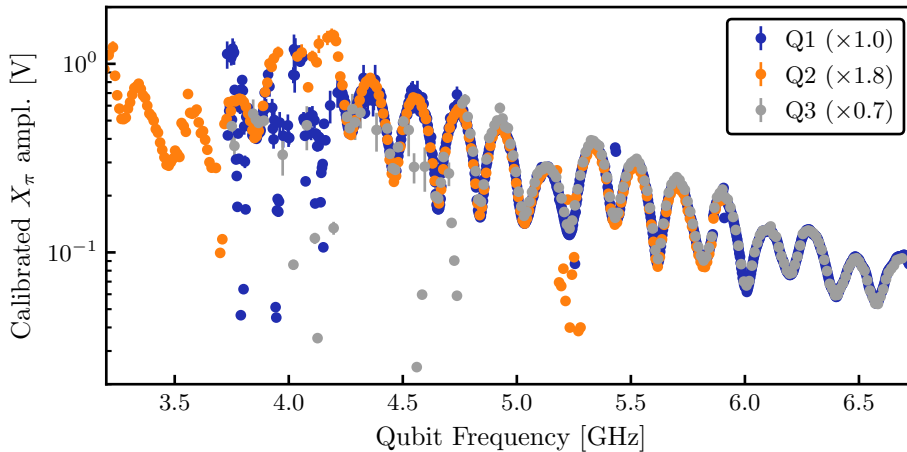


Figure 4.18: Qubit Rabi drive calibration through multiplexer port. AWG drive amplitude required to perform a X_π pulse on qubits Q1, Q2 and Q3 by driving them through the common multiplexed readout port. Note that the individual datasets are multiplied with constant factors as in the plot legend. No data for Q4 was collected in this measurement.

with the fact that the multiplexer cavity centered around 10 GHz acts as a wide bandpass filter. The most striking observation is that the datasets closely resemble each other up to a constant factor. This confirms that we are indeed measuring the transfer function of the same line multiple times, with the multiplicative factors representing the relative effective coupling strength from the readout port to the different qubits. We contrast this with the calibration of individual readout ports in Fig. 4.14, where the oscillation periods do not match.

Since we are using the same dFBL enclosure lid as in the previous experiment, we can also check for the robustness of the dFBL Rabi calibrations shown in Fig. 4.13 and Fig. C.3 by carrying out the same measurement here. The result for Q2 is displayed in Fig. 4.19, directly comparing the two datasets without any scaling or conversion. Perhaps unsurprisingly, the two calibration curves are very similar, as the dFBL side of the sample enclosure was not modified. Still, it is encouraging to observe that the calibration is repeatable to such an extent when taking into account possible changes in chip alignment after loading it into the new enclosure. Since the dFBLs are part of the sample enclosure and can therefore be reused many times, having a pre-existing line calibration could speed up subsequent measurement runs with similar qubit devices. This is in contrast to on-chip dFBLs, which are wire-bonded to a sample carrier or PCB and will therefore have a slightly different transfer function on every new device.

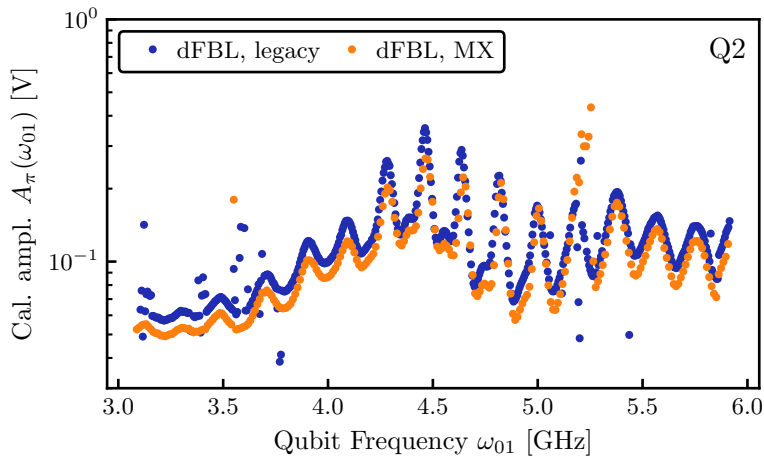


Figure 4.19: Comparison of dFBL Rabi drive calibration between cooldowns. Rabi calibration curves for qubit Q2 driven through its dFBL. Blue data points correspond to first cooldown with a sample enclosure with individual readout lines (legacy), while orange data points belong to second cooldown using the sample enclosure with multiplexer (MX).

We repeat the dFBL calibration comparison for all qubits, with the data reported in Fig. C.5 in the appendix. The calibration curves for Q3 and especially Q1 differ more significantly between cooldowns, but we note that this coincides with the removal of infrared (Eccosorb) filters on the control lines of these two qubits between cooldowns, which could confound the measurement. On the other hand, data for Q2 and Q4 shows good reproducibility. This highlights that the entirety of the control lines, including filters and other components, matters for the calibration.

4.5.2. T_1 spectroscopy with multiplexer

To complete the picture, we perform a wide-frequency scan of energy relaxation T_1 in the new enclosure. We observe in Fig. 4.20 that the T_1 spectra have changed considerably on a qualitative level compared with previous data in Fig. 4.15. This is consistent with our finding that in the original sample enclosure the qubit lifetime was predominantly limited by radiative decay through the readout coaxial ports. The modification of precisely this part of the qubit environment is therefore expected to cause a change in the T_1 spectrum. Given that the T_1 values measured here generally lie in a similar range between 10–100 μs as before, we conclude that the same radiative limitation is present here. It is unlikely that an additional source of energy relaxation was introduced between the two experiments, since the

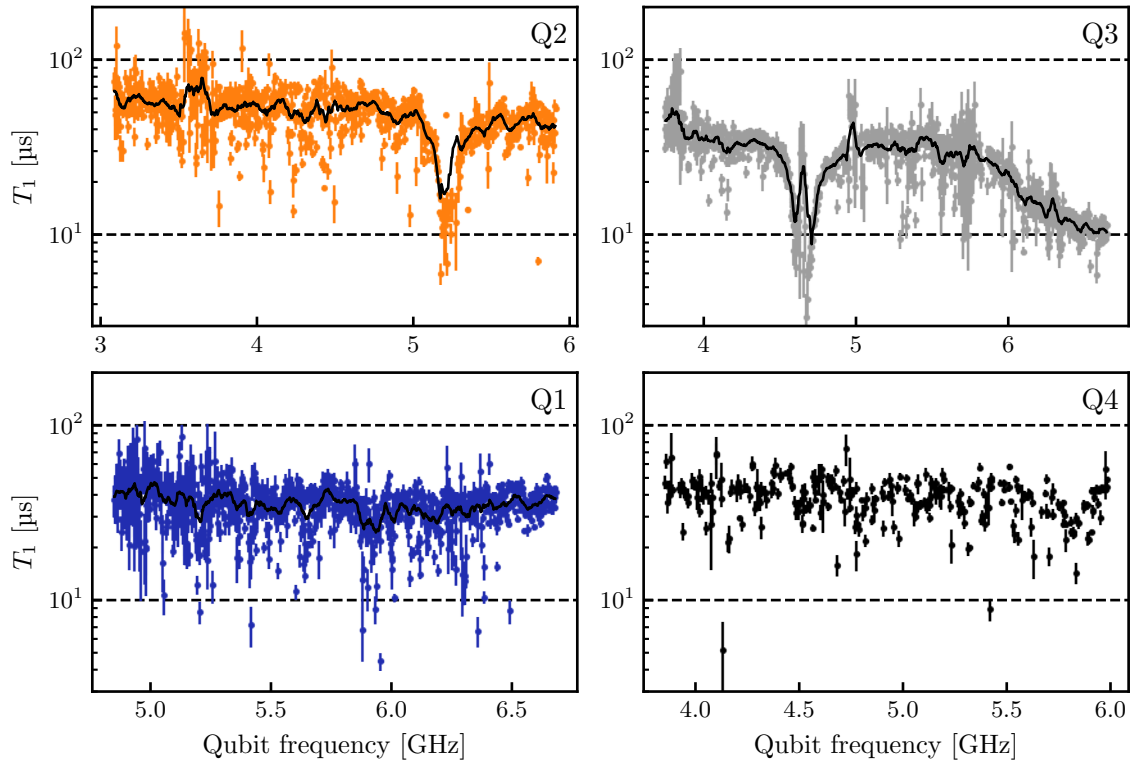


Figure 4.20: Frequency dependence of T_1 in enclosure with multiplexer. Measurements of energy relaxation time T_1 versus qubit frequency performed individually on each qubit. The black line overlaid to the data is a running average obtained by applying a square box filter with width $n = 20$ to the raw data points.

transmon device, dFBL enclosure lid, cryostat and wiring setup have remained the same.

Strikingly, we do not observe clear oscillations of T_1 versus frequency as we did previously. This is somewhat unexpected, as the Rabi calibration through the multiplexer shown in Fig. 4.18 is qualitatively similar to what was observed for the readout lines in the original enclosure, shown in Fig. 4.14. The similarity suggests comparable levels of impedance mismatch between the qubit and the coaxial readout lines in both cases, and should therefore lead to visible oscillations in radiatively limited T_1 in both cases. The main difference is the presence of the multiplexer cavity lying between the qubit and its environment, which might act as an impedance transformer from the point of view of the qubit, but we recognize the need for further work to understand this more thoroughly.

Qubits Q2 and Q3 show prominent dips in their spectrum, which we can match with large avoided crossings in their respective flux-dependent spectroscopy (not shown). It is currently unclear whether these modes are part of the multiplexer

cavity assembly, or whether these are local defects that interact with the qubit. The latter could potentially be removed by thermally cycling the device and retaking the same spectra, which was not done in the present study.

From a wider perspective, the type of measurements presented here and in the preceding section allows us to utilize the qubit not as a computing resource, but rather as a very sensitive detector of its environment. This can help us understand and optimize our experimental design and improve future iterations of the architecture.

4.6. Summary

In this chapter we presented the design, construction and experimental characterization of a cryogenic sample enclosure for off-chip flux control of 3D-integrated superconducting circuits. The core element of this architecture, the dFBL, allows for highly selective and well-controlled remote flux biasing of gradiometric SQUID loops without requiring galvanic contact to the substrate hosting the SQUIDs themselves. It is seamlessly integrated into our existing coaxial circuit QED architecture without increasing the footprint of a single coaxmon unit cell, i.e. it retains the tileable nature of the design and the simplicity of the device fabrication and packaging process.

Crucially, the dFBL can simultaneously provide both flux and charge control of transmon qubits and is therefore a single, fully integrated XYZ tuning element, removing the need for separate flux and charge bias lines. This achieves a similar goal as the approach presented in Ref [143], although our solution does not require the lithographic patterning of an additional substrate to create the flux bias lines, nor does it involve bonding to the qubit substrate. The dFBLs are incorporated inside the sample enclosure and are therefore fully reusable.

The spatially well-defined paths for the bias currents lead to a simple magnetic field distribution in three dimensions, which in turn allows for the reliable prediction of flux tuning performance via finite-element magnetostatic simulations. We use this method to estimate the mutual inductance M between dFBL and corresponding gradiometric SQUID, as well as the full DC flux crosstalk matrix $(\chi)_{ij}$ in our dFBL enclosure, finding excellent agreement with experimentally measured values. Our design achieves at the same time a sufficiently large $M \sim 2$ pH for efficient flux biasing and also a large numerically predicted energy relaxation limit $T_1^{dFBL} > 500$ μ s which is not currently a limiting T_1 factor in our experiment. Flux crosstalk

is found to be $1.5(9) \times 10^{-3}$ on average across the measured four-transmon device, demonstrating the capability for highly selective flux bias. We attribute this to the optimal matching of the magnetic field generated by the dFBLs with the geometry and orientation of the gradiometric SQUID loops. According to simulations, M and T_1^{dFBL} can be tuned over several orders of magnitude by adjusting the vertical separation between transmon and dFBL, with only a modest effect on flux crosstalk performance. The architecture can therefore be easily adapted to always provide the optimal level of protection of qubits from their external environment.

Further, we use the tunability of our transmons over several GHz to characterize the frequency response of both the dFBLs and the readout lines in our cryogenic experiment. This is carried out both in a sample enclosure with individual coaxial readout ports, as well as with a four-way 3D readout multiplexer. By comparing the AWG voltage amplitude required to excite a qubit via different ports to the behavior of qubit energy relaxation T_1 as a function of frequency, we can identify coupling to the readout lines as a likely limiting factor for T_1 in our experiment due to enhanced radiative (Purcell) decay. This provides a clear path towards improving qubit lifetimes in our architecture. In contrast, decay through the dFBL is unlikely to contribute significantly to T_1 at current levels, as predicted by simulations.

While the dFBL architecture has been tested here specifically with transmon qubits due to their ubiquity and simplicity, we are by no means limited to only this type of circuit. It is relatively straightforward to modify the gradiometric SQUID design to include arrays of Josephson junctions or even linear inductors, giving access to many of the superconducting qubit species discussed in [Chapter 2](#) such as the fluxonium, flux qubit and the inductively shunted transmon (IST). This makes our design *qubit-agnostic*, providing a general flux control approach for a variety of circuit QED implementations. The following chapters will provide support to this claim by investigating ISTs in this architecture.

5.

The gradiometric inductively shunted transmon

5.1. Introduction

We turn our attention to an important obstacle to scaling quantum circuits that was discussed in [Sec. 2.4](#), i.e. quantum crosstalk in coupled systems. More specifically, we focus on static ZZ -type interactions between qubits, which can be an important source of coherent error in state-of-the-art entangling gate implementations, especially between qubits with fixed coupling [[91](#), [111](#)].

One elegant method for suppressing the static ZZ term is to use qubits with opposite anharmonicity. Previous work on this native ZZ suppression has focused mainly on the capacitively shunted flux qubit (CSFQ) [[100](#), [101](#), [130](#)], while a recent preliminary study showed how an RF-SQUID circuit with a small geometric inductance ($E_L \approx 2E_J$) can achieve a similar outcome [[79](#)]. The latter is an experimental implementation of the inductively shunted transmon (IST), which we treated theoretically in [Sec. 2.3](#). Here we choose to focus on this approach, motivated mainly by the simplicity of the IST – a single-junction circuit with a linear shunting inductor. Recent work on superinductance qubits employing large geometric inductances and substrate engineering to realize a variety of qubit regimes [[74](#), [80](#)], as well as a junction shunted by a transmission line [[86](#)], further motivates the study of this simple but relatively unexplored approach to implementing RF-SQUID-type qubits.

This chapter presents the design, fabrication, and cryogenic characterization of gradiometric IST qubits embedded in our coaxial architecture with differential flux bias control scheme, developed in [Chapter 4](#). The aim is to understand the basic practical properties of the IST, from circuit design and simulation to cryogenic measurements in frequency and time domain. This work lays the foundations for the study of IST-transmon dynamics and ZZ suppression, which will be the subject of [Chapter 6](#).

5.2. IST design

5.2.1. Choosing the inductance range

The RF-SQUID is a simple flux-tunable circuit consisting of a superconducting ring with inductance L interrupted by a Josephson junction (JJ), to which we can assign an inductance scale $L_J = (\Phi_0/2\pi)^2/E_J$. Based on our treatment of the RF-SQUID in [Sec. 2.3](#), we aim to realize circuits with a ratio of Josephson to geometric inductance $\beta = L_J/L = E_L/E_J$ in the range $1 < \beta \lesssim 2$. We also recreate the heavy-charge regime $E_J/E_C \gg 1$ as in the transmon by adding a large shunting capacitance to the circuit¹. These parameters lead to a single-well phase potential and a plasmonic energy level structure with sufficient anharmonicity for the operation as a qubit.

Given typical values of Josephson inductance on the order of ~ 10 nH, we aim to design linear inductors in the range 5 – 10 nH. This is a very modest target when compared to other realizations of RF-SQUID qubits such as the fluxonium [[64](#), [65](#)] (several 100 nH) and the quasicharge qubit [[72](#)] (1 μ H and beyond). These large values are provided either by long JJ arrays, by high kinetic inductance materials [[182](#), [183](#)], or by geometric superinductors with substrate engineering [[73](#)]. For our purposes, it is sufficient to instead implement simple geometric inductors formed by an aluminium wire on high-resistivity silicon, similar to the spiral inductors that form part of our coaxial readout resonators. This avoids adding any fabrication complexity and/or new materials to our process, and also allows us to model the inductance classically using finite-element (FE) numerical methods.

5.2.2. IST geometry and simulation

We build a gradiometric RF-SQUID by starting from the tunable coaxial transmon presented in [Chapter 4](#) and replacing one of the two JJs by a meandering wire with N_t turns, as shown schematically in [Fig. 5.1](#). This is similar to the design used in [Ref. 79](#). We use concentric sections of wire to efficiently fill out the space between the inner and outer qubit electrodes. The exclusion angle θ_{excl} defines how much empty space is left between the inductor and the JJ leads. The geometric inductance L is roughly proportional to the total length of the wire, although we note that the tight packing of the meander inevitably reduces L when compared to a straight wire of equivalent length [[184](#)]. The gradiometric topology is maintained,

¹This is not required on a fundamental level, but is chosen for practical reasons to keep the capacitive network of the device and its coupling to the environment similar to our transmon.

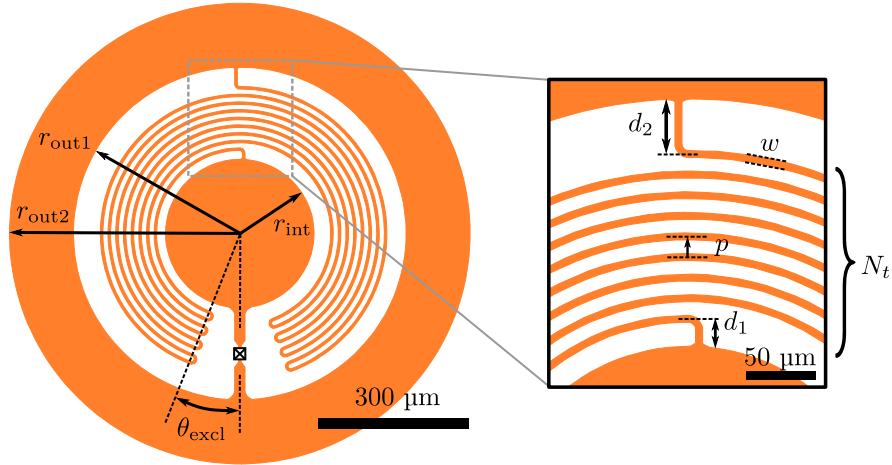


Figure 5.1: Gradiometric IST inductor design. The gradiometric IST consists of a single-JJ coaxial transmon qubit shunted by a meandering inductor wire of width w and pitch p with a total of N_t turns. In this drawing, $N_t = 8$.

but the convoluted shape of the two gradiometer loops has an impact on the flux tunability of this circuit, which we will touch upon in [Sec. 5.2.3](#).

The superconductor will also add some kinetic inductance L_K . For a wire of length l , width w and thickness t , we can estimate it via the relation [\[182\]](#)

$$L_K = \mu_0 \lambda^2(0) \left(\frac{l}{wd} \right), \quad (5.1)$$

where μ_0 is the vacuum permeability and $\lambda(0)$ is the penetration depth of the superconductor at zero temperature. The value of $\lambda(0)$ is roughly 45 nm for bulk aluminium, but increases in thin films [\[170, 171\]](#). Assuming $\lambda(0) \approx 60$ nm for our typical films with $t = 100$ nm, we find that the kinetic inductance of a 10 mm long and 5 μm wide wire is $L_K \approx 0.1$ nH. This constitutes at most a few % of the target values of geometric inductance mentioned above and is therefore going to be neglected in the following discussion.

Our design philosophy is to minimally modify our existing gradiometric transmon design to implement the IST. The following design steps are undertaken:

- The inductor wire width w is chosen to be 5 μm , the same as the inductor wires used for our spiral resonators. This allows us to keep using the same photolithography and etching process as before, to maximize the quality and reproducibility of the inductor.
- The same values of shunting capacitance C of around 90-95 fF as in the tunable transmon are targeted. This is motivated by the notion that equal

but opposite anharmonicities (and therefore charging energies²) are required to suppress the static ZZ interaction between a transmon and a positive-anharmonicity qubit [185]. While numerical treatment of the problem reveals a slightly more complicated picture (see Sec. 6.3.2), this choice of C is still very reasonable and leads to a useful parameter regime.

- Four variations of the meandering inductor are implemented, covering the range between 3.5 nH and 8 nH. Given the constraints of inductor wire thickness and of qubit capacitance mentioned above, there is not enough space between the qubit electrodes to fit much longer inductors, but this range is sufficient for the present study.
- The silicon substrate, sample enclosure, dFBL, spiral resonator and readout circuitry remain the same as in the experiments of Chapter 4. The intuition here is that the largely unaltered size and geometry of the coaxial qubit electrodes should lead to a similar capacitive network, and therefore to similar coupling of the qubit excitation (which is still of plasmonic nature) to its environment.

FE simulation

In order to predict the IST circuit parameters, two separate FE simulations are performed. Firstly, we obtain the full capacitance network of the IST in ANSYS Maxwell, and then calculate the equivalent capacitance C . For simplicity, we exclude the readout resonator to make the problem easier to treat. To model the capacitance sub-network, we remove the JJ and the two leads connecting the meander to the qubit electrodes, as shown in Fig. 5.2 (a), leaving us with four nodes: inner (i) and outer (o) qubit electrode, floating meander (m), and ground (g). The equivalent network, represented schematically in Fig. 5.2 (b), can now be reduced to a single node with effective capacitance C to ground. The meander node is eliminated with a Y-delta transformation, and the remaining network is simplified as was done for the transmon in Sec. 4.2.2, leaving us with the exact expression

$$C = C_{io} + C_b + \frac{(C_{og} + C_a)(C_{ig} + C_c)}{C_{ig} + C_c + C_{og} + C_a}, \quad (5.2)$$

²To lowest order, $|\alpha| \approx E_C$ is valid for both the transmon and the IST at half flux bias, see perturbation theory in Sec. 2.3.2.

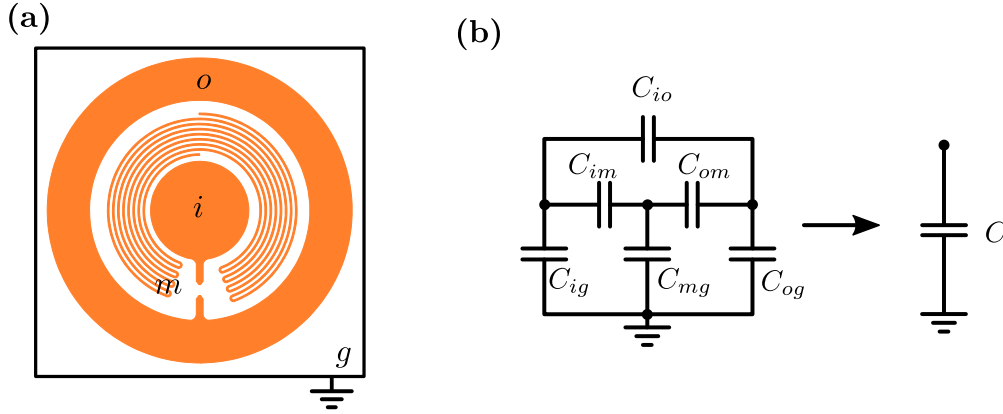


Figure 5.2: IST Capacitance network. a) Simplified IST circuit for capacitance calculation. The device enclosure provides the ground terminal. b) Equivalent capacitance network, which can be reduced to a single node with capacitance C to ground (see main text).

where

$$C_a = \frac{C_{mg}C_{om}}{C^*}, \quad C_b = \frac{C_{im}C_{om}}{C^*}, \quad C_c = \frac{C_{im}C_{mg}}{C^*},$$

$$C^* = C_{im} + C_{om} + C_{mg}.$$

Note that this is a relatively crude approximation, since we are neglecting the distributed nature of the capacitance along the length of the meandering inductor.

Secondly, an HFSS Eigenmode simulation of the IST with the JJ removed gives us the fundamental frequency ω of the linear LC circuit consisting only of the qubit electrodes and the meandering inductor. The inductance L is then extracted directly from $\omega = 1/\sqrt{LC}$ using the capacitance obtained from Eq. (5.2). In Fig. 5.3 we show simulated values of L versus the inductor dimensions N_t and θ_{excl} for a representative IST design. The value of N_t sets the approximate range of L , while the exclusion angle θ_{excl} can be used to continuously fine-tune L . Adding more meander turns leads to ever smaller gains in L , due to the tighter packing of the meander in the same space and the increased (detrimental) mutual inductance between neighboring wires.

Final choice of parameters

Based on these simulations, we choose four designs to evenly cover the range of L between 8 and 3.5 nH, labelled IST1 to IST4. The physical dimensions and simulated values of L and C for these four designs are reported in Table 5.1. By slightly adjusting the radii of the coaxial electrodes we achieve a roughly constant C across all variations.




Figure 5.3: Simulated IST inductance. Inductance L extracted from FE simulation of capacitance C and eigenfrequency ω as a function of inductor parameters N_t and θ_{excl} for a representative IST design with $(r_{in}, r_{out1}, r_{out2}, w, p) = (135, 350, 450, 5, 15)$ μm and symmetric gaps $d_1 = d_2$. As expected, a longer inductor (larger N_t , smaller θ_{excl}) leads to higher L .

In Sec. 5.3 we will provide an overview of devices to be fabricated and discuss the choice of E_J and frequency targets. Before that, we want to make predictions about some important practical properties of the IST, including flux noise properties and higher-order modes of the circuit.

Table 5.1: IST design parameters. Physical dimensions of the four implemented IST designs, labelled as in Fig. 5.1 (a), along with simulated values of L and C . The following quantities are the same in all designs and are omitted from the table: $w = 5$ μm , $d_1 = 20$ μm , $r_{out2} = 450$ μm .

	<i>Physical dimensions</i>					<i>Simulated parameters</i>		
	N_t	θ_{excl} (deg)	r_{in} (μm)	r_{out1} (μm)	d_2 (μm)	p (μm)	L (nH)	C (fF)
IST1	12	21	135	350	20	14.6	7.9	91.4
IST2	10	21	145	350	40	14.5	6.3	94.0
IST3	8	21	145	325	40	15.0	4.9	93.5
IST4	8	45	145	300	40	11.9	3.5	92.6

5.2.3. Flux tunability and flux noise

We want to understand how the presence of the meandering inductor will affect the flux tunability via the dFBL, as well as the sensitivity of the RF-SQUID loop to flux noise.

It is easy to see from Fig. 5.1 (a) that the two interdigitated areas of the gradiometer loops occupy regions of both negative and positive differential field generated by the dFBL, meaning that the total integrated flux threading each loop will be reduced when compared to a standard gradiometric transmon of similar dimensions. This in turn will lead to a smaller net flux bias Φ_Δ of the SQUID. In order to ensure that we can still suitably bias this qubit with our existing dFBL, we repeat the same Maxwell magnetostatic simulations as in Sec. 4.3 to extract the mutual inductance M between the IST and a dFBL centered above it at 1 mm distance from the substrate, as well as the elements of the DC flux crosstalk matrix $(\chi)_{ij}$ for a 2×2 qubit array. As before, we perform separate simulations assuming either complete field exclusion (permeability $\mu_{rel} = 0$) or transparency ($\mu_{rel} = 1$) of the qubit metal. The result is reported in Table 5.2, where we also include the previously obtained numbers for the gradiometric transmon. All IST designs have significantly decreased values of M compared to the gradiometric transmon, owing to the the partial cancellation of differential flux mentioned above. Most of the flux crosstalk values are on the order of $1 - 3 \times 10^{-3}$, not significantly different from the transmon case, showing that the selectivity of the dFBL architecture is robust with respect to changes in qubit geometry. The notable outlier is χ_y for $\mu_{rel} = 0$, which has anomalously large values for three out of four IST designs. This is not currently understood, but could be originating from an artefact of the FE simulation.

Table 5.2: IST flux tunability and crosstalk. Mutual inductance M between IST and dFBL, as well as DC flux crosstalk matrix elements χ_x , χ_y , χ_d obtained from magnetostatic simulation of a 2×2 IST array assuming either superconducting ($\mu_r = 0$) or normal ($\mu_r = 1$) permeability of the qubit metal. The distance between qubit and dFBL is set to 1 mm. See Sec. 4.3 for simulation methods and definitions. The corresponding transmon results are included for comparison.

	$\mu_r = 0$				$\mu_r = 1$			
	M (pH)	χ_x (10^{-3})	χ_y (10^{-3})	χ_d (10^{-3})	M (pH)	χ_x (10^{-3})	χ_y (10^{-3})	χ_d (10^{-3})
IST1	0.93	1.5	1.6	3.4	0.42	2.1	2.4	3.0
IST2	1.24	1.2	6.8	2.5	0.69	1.7	0.6	3.0
IST3	1.22	1.8	7.8	2.7	0.61	1.5	0.4	2.9
IST4	1.26	1.3	7.6	3.3	0.60	1.7	0.3	2.9
Tr.	1.92	1.2	1.2	2.7	1.13	1.1	1.1	2.8

Table 5.3: IST SQUID aspect ratio and flux noise amplitude. Length of meandering inductor for the four implemented IST designs, together with the corresponding SQUID aspect ratio ρ calculated using Eq. (3.8). We also include the expected relative change in microscopic flux noise amplitude $\sqrt{A_\Phi} \propto \sqrt{\rho}$ compared to the transmon design with $\rho \approx 100$ from Chapter 4.

	Inductor length (mm)	Aspect ratio ρ (10^3)	Relative $\sqrt{A_\Phi}$ ($\sqrt{A_\Phi(\rho=100)}$)
IST1	16.4	3.3	5.8
IST2	13.4	2.7	5.2
IST3	10.1	2.1	4.5
IST4	8.2	1.7	4.1
Transmon	/	0.1	1

Fortunately, even the lower bounds on M for the case $\mu_r = 1$ are still sufficient to apply a full flux period in the experiment³. We are helped by the fact that the magnetic field of the dFBL is highly inhomogeneous and is strongest close to the outer qubit ring (see Fig. 4.4), where there is sufficient empty space belonging to only one of the two gradiometer loops.

Another important consequence of adding the inductor is that it creates a macroscopically long ($\sim 10 - 15$ mm) superconducting loop with a high aspect ratio ρ . Going back to Sec. 3.3.3, there is experimental evidence that the sensitivity of a SQUID loop to microscopic flux noise increases with its aspect ratio [152, 159], which is what led us to design our gradiometric transmon with relatively wide JJ leads. In Table 5.3 we report the inductor length of each IST design, as well as the generalized aspect ratio ρ calculated using Eq. (3.8). To better understand the practical effect of different aspect ratios on noise properties, we include the relative increase in flux noise amplitude $\sqrt{A_\Phi} \propto \sqrt{\rho}$ in the IST compared to the gradiometric transmon with $\rho \approx 100$. We conclude that the IST should experience significantly higher flux noise amplitudes by about a factor of 4 to 6, if microscopic local fluctuators are in fact the dominant contribution to flux noise in the experiment⁴.

The inevitable large aspect ratio could represent a fundamental drawback of the geometric inductive shunt design, at least as long as the sources of local flux noise in SQUIDs cannot be strongly suppressed below current levels. Interestingly, the authors of Refs. 74, 80 report large flux noise amplitudes $\sqrt{A_\Phi} \sim 100 - 800 \mu\Phi_0$

³A welcome side effect is that a smaller M decreases the qubit sensitivity to current noise on the dFBL according to Eq. (3.1).

⁴As opposed to other noise sources, e.g. current noise on the dFBL or stray magnetic fields in the experimental environment.

measured in various realizations of geometric superinductance qubits, which they attribute to the long perimeter of their inductor coils. Based on the published images and design parameters of those devices, we indeed estimate their aspect ratios to be extremely high, around $\rho \sim 10^4 - 10^5$. We shall return to this topic when analyzing experimental coherence data in [Sec. 5.5.2](#).

5.2.4. Parasitic inductor modes

The RF-SQUID Hamiltonian in [Eq. \(2.10\)](#) assumes that the entire superconducting loop can be described by a single macroscopic superconducting phase variable $\hat{\varphi}$. In the case of shunting inductances made from JJ arrays, where many individual superconducting islands exist, this is merely an approximation of a more complicated, multi-dimensional Hamiltonian [\[63, 64\]](#). In contrast, the use of a geometric inductor guarantees an uninterrupted superconducting ring, for which [Eq. \(2.10\)](#) is an accurate description.

In practice, this is only true as long as our circuit can be described by ideal lumped elements – or equivalently, as long as we operate at frequencies well below any self-resonances of the individual circuit components. In the case of a meandering inductor, the presence of a distributed capacitance creates an effective transmission line with a resonant frequency that depends on the inductor geometry [\[73\]](#). The same effect also appears in JJ arrays due to the finite capacitance of the JJs themselves [\[64, 68\]](#). We refer to these spurious resonances and their higher harmonics as *parasitic modes* of the inductor, as they are usually undesired and can negatively affect the lumped mode that we want to utilize as the qubit. We denote their frequencies by $\omega_{p,n}$, where $n = 0, 1, \dots$ labels the modes in ascending frequency order.

The same eigenmode simulations performed in [Sec. 5.2.2](#) to obtain the fundamental mode of the IST can also help us to understand the parasitic mode structure. [Fig. 5.4](#) (a) and (b) visualize the lumped LC mode (here at $\omega = 7.9$ GHz) and the first parasitic inductor mode ($\omega_{p,0} = 11.8$ GHz) of design IST2 by plotting the electric field magnitude associated with each of these modes. The JJ is approximated by a 10 nH lumped linear inductor. While the lumped mode is distributed over the entire circuit, including the JJ and the outer ring electrode, the field of the parasitic mode is strongly localized in the inductor region. Higher harmonics of the inductor mode (not shown) have a similar field distribution. In [Fig. 5.4](#) (c) we then plot the frequencies of the five lowest parasitic modes for each IST design, showing that these modes exist at frequencies down to around

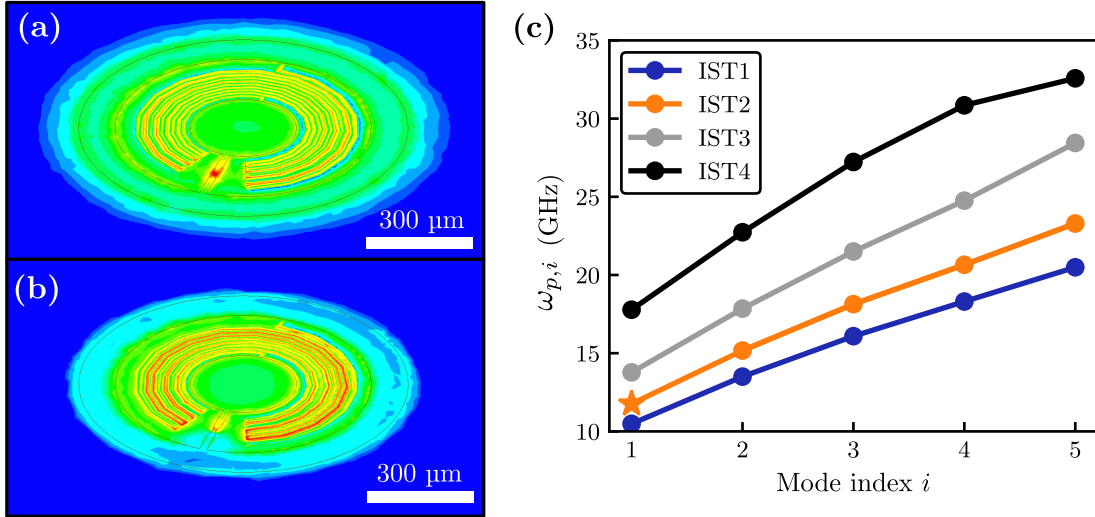


Figure 5.4: Parasitic modes of the inductor. (a)-(b) Magnitude of electric field associated with the lumped mode (top) and first parasitic inductor mode (bottom) of IST2, obtained from HFSS. Scale bar and logarithmic color map are identical for both plots, with large (small) electric field values shown in red (blue). While the lumped mode extends over the entirety of the coaxial electrodes, the parasitic mode is strongly localized in the inductor region. (c) Simulated frequency of the five lowest parasitic inductor modes for each IST design. The mode plotted in (b) is denoted by a star.

10 GHz in our chosen designs. It is therefore reasonable to expect that they will couple to the lumped qubit mode and modify its spectrum, as in other RF-SQUID implementations [64, 74]. This is treated phenomenologically by considering each of the (infinitely many) parasitic modes as an independent linear oscillator coupled to the qubit, analogously to the case of a qubit coupled capacitively to the modes of a transmission line, resulting in the Hamiltonian [186]

$$\mathcal{H} = \mathcal{H}_0 + \sum_n \left[\hbar\omega_{p,n}\hat{c}_n^\dagger\hat{c}_n + \hbar g_{p,n}\hat{N}(\hat{c}_n^\dagger + \hat{c}_n) \right], \quad (5.3)$$

with the RF-SQUID Hamiltonian \mathcal{H}_0 from Eq. (2.10), \hat{c}_n^\dagger the annihilation operator for parasitic mode n , and $g_{p,n}$ the coupling strength between qubit and mode n . The latter scales as $g_{p,n} \sim g_{p,0}\sqrt{n+1}$ for small n [186]; this means that higher modes, although further detuned from the qubit, might still lead to a noticeable Lamb shift $\sim g_{p,0}^2(n+1)/\Delta$ of the dressed qubit frequency, complicating the analysis. In the case of a geometric inductor, the magnitude of $g_{p,0}$ can be estimated classically by fitting the admittance $Y(\omega)$ of the circuit at the Josephson junction port, obtained from FE simulations, to a lumped circuit model that includes the self-resonance of the inductor [73].




Figure 5.5: Influence of parasitic modes on IST spectrum. (a) Numerically calculated IST spectrum in the presence of coupling to N parasitic modes, with N between 0 and 4. IST parameters E_L , E_C are chosen to match IST2, with $\beta = 1.5$, and the parasitic mode frequencies are taken from FE simulations. Coupling to the lowest mode is set to $g_{p,0} = 500$ MHz, and higher mode couplings are scaled by $\sqrt{n+1}$. (b) Anharmonicity of IST for $N = 0$ (black) and $N = 4$ (orange). The IST second transition ω_{02} crosses the lowest parasitic mode, leading to an avoided crossing.

It is important to point out that the infinite number of modes in Eq. (5.3) lead to divergences when trying to derive system quantities naively [54], but a more careful treatment of the problem results in a natural ultraviolet cutoff that removes these divergences [186, 187]. For practical purposes, only a finite number of modes are usually considered when fitting experimental data to this model.

We briefly investigate the effect of parasitic modes on the IST spectrum by numerically simulating the system in Eq. (5.3) for a finite number N of modes included in the sum. We choose E_L and E_C corresponding to design IST2, as well as $\beta = 1.5$ and $g_{p,0} = 500$ MHz, which are reasonable values based on experimental results (see Sec. 5.4). We assume the scaling $g_{p,n} \sim g_{p,0}\sqrt{n+1}$. The frequencies $\omega_{p,n}$ are as in the FE simulations of Fig. 5.4. We plot the qubit frequency ω_{01} , second transition frequency ω_{02} , and anharmonicity α versus flux in Fig. 5.5 for N up to 4, and observe how the spectrum is distorted by the addition of each parasitic mode, especially at the upper sweet spot where the qubit-mode detunings are smaller. The second transition frequency crosses the lowest parasitic mode, leading to an avoided crossing in the anharmonicity. Note that in the simulation, the avoided crossing makes the labelling of modes ambiguous – the second IST level is identified as either E_2 , or E_3 in the numerical spectrum, depending on where it lies with respect to the parasitic modes.

It is worrying that the frequency shift caused by each of the modes is roughly constant and does not decrease with higher mode index. We attribute this to the scaling $\sqrt{n+1}$ we assumed for the mode couplings, which is likely to be an oversimplification of the real system. In reality, the effective coupling of each parasitic mode to the qubit will depend on the spatial extent of the distributed parasitic mode excitations and their overlap with the qubit mode. Still, due the qualitative effect shown here we conclude that the inclusion of parasitic modes is likely to be necessary for the correct quantitative fitting and extraction of circuit parameters from experimental spectra, similar to other RF-SQUID type circuits [64, 74, 80]. The cutoff N is in principle arbitrary, but practical considerations such as the computational time needed to simulate many modes will be required.

More advanced methods to simulate the system of interest are available and could be implemented in a future study. For example, the black box quantization (BBQ) method [188–190] allows one to obtain accurate parameters of the quantum system such as frequencies and non-linear couplings from the classical impedance response $Z(\omega)$ of the system. More recently, the energy participation ratio (EPR) framework [191] has been proposed as an intuitive and computationally efficient approach that relies only on the result from a single FE eigenmode simulation. Furthermore, EPR is derived ab initio and is valid for arbitrary non-linear systems, while BBQ uses a perturbative approach and is therefore only suited for weakly non-linear circuits.

5.2.5. Future design improvements

With an eye towards a future iteration of the IST, we have come across several design aspects that would benefit from a miniaturization of the meandering inductor:

- The achievable value of L is limited by how much inductor length can be fitted inside the coaxial electrodes. We don't foresee requiring inductances much larger than 10 nH for IST qubits in the single-well regime, which should not require fundamental design changes. If the aim is to keep the IST compatible with our existing architecture, the overall footprint of the qubit should remain similar, and so the only practical way to increase L is to decrease w and fit more meander turns into the same space.
- The large footprint of the inductor reduces the tunability of the gradiometric SQUID via the dFBL. A thin enough inductor wire might achieve sufficient L

without having to fill out most of the SQUID loop area, allowing for a larger net differential flux to be threaded.

- Furthermore, the parasitic modes of the inductor could be moved to higher frequency by decreasing the self-capacitance of the meander. This can again be achieved by a smaller overall inductor footprint [73], which might also coincidentally reduce the coupling strength $g_{p,i}$ to the qubit by reducing the spatial extent of the parasitic modes.

We predict that an inductor with $w \approx 1 - 2 \mu\text{m}$ should be achievable in practice without drastic changes to our photolithography process. One drawback of thinner inductors is that given a fixed error margin on lateral feature sizes of, say, $\pm 100 \text{ nm}$, the relative error on the quantity w becomes larger. This will affect the reproducibility of L and might partially remove the benefits of using a geometric inductor versus a junction array. We do not see this as a fundamental issue, but rather as an engineering challenge that can be overcome. State-of-the-art photolithography in the semiconductor industry is already capable of feature sizes orders of magnitude smaller than what we are aiming for here. Of course, features down to 100 nm and below can also be easily obtained by switching to electron-beam lithography, which is used in other recent implementations of geometric inductance qubits [73, 79]. One question that remains to be answered is which type of process will lead to better reproducibility and to the highest quality of the inductor and surrounding interfaces.

5.3. Experiment overview

5.3.1. Fabrication

After having presented the IST design and its predicted properties, we now turn to implementing it experimentally and carrying out a proof of principle of this new circuit.

We fabricate $5 \times 5 \text{ mm}^2$ samples through our established wafer-scale process on a 3-inch high-resistivity silicon wafer and subsequent dicing. Details of the fabrication recipe can be found in [Appendix A](#). The IST devices presented in the following are, in fact, fabricated on the exact same wafer as the gradiometric transmon device from [Chapter 4](#). This underscores how simple it is to practically implement geometric inductive shunts with $\beta = E_L/E_J \gtrsim 1$ in an existing transmon-based




Figure 5.6: Fabricated IST circuits. (a) Optical image of a partially fabricated IST qubit, after a single step of photolithography and aluminium wet etching. The coaxial electrodes and the meandering inductor are fabricated at the same time. (b) Detail of a meandering inductor with $w = 5 \mu\text{m}$.

cQED architecture, which is a point that will become even more apparent when presenting experimental results.

The meandering inductors are fabricated in the same step as the qubit electrodes using photolithography and wet etching of aluminium. Fig. 5.6 shows example optical images of features created after this fabrication step. The whole circuit, with the exception of the Josephson junctions added later by EBL, consists of a single aluminium layer deposited on a pristine Si substrate that has been BOE-treated. This ensures that we have a clean substrate-metal interface. While we have not carried out a systematic study of IST inductors made using different fabrication techniques, we believe that the cleanliness and simplicity of this process might play an important role in ensuring good performance and reproducibility of the IST inductor. As an example, the removal of native substrate oxide prior to deposition of the inductor metal might help in removing defect sites that could provide a source of magnetic flux noise in SQUIDs [157].

5.3.2. Experiment design

The first experiment we design (device A) consists of four uncoupled qubits with each of the IST design variations listed in Table 5.1 represented once, as shown schematically in Fig. 5.7 (a). The aim of this experiment is to test individual IST qubits for their performance and gain a better understanding of how their properties change as a function of L . Ideally, we want to extract the parameters L and C from experimental data and compare them to predicted values. Coherence

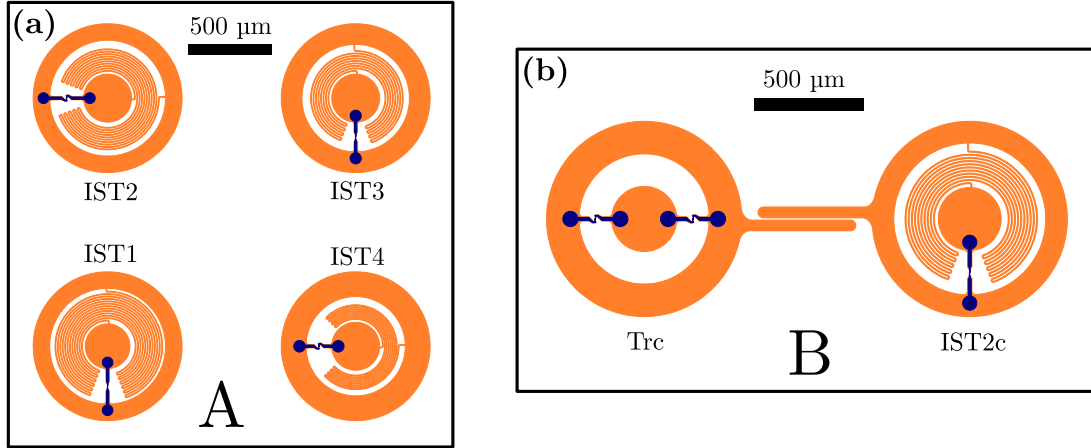


Figure 5.7: Overview of measured IST devices. (a) Schematic drawing of device A consisting of four single IST qubits, representing each of the four chosen inductor designs. (b) Device B, consisting of a coupled transmon-IST circuit for the investigation of two-qubit dynamics and ZZ suppression. We use simple coupling arms extending from the outer electrodes to implement direct capacitive coupling between qubits.

properties will also be studied, to further assess the viability of these circuits as qubits.

To investigate the behavior of two-qubit dynamics involving the IST, we also design devices with coupled IST-transmon pairs. Fig. 5.7 (b) shows such a coupled pair (device B), where we label the qubits as Trc and IST2c respectively. While we choose to use a tunable gradiometric transmon in this particular case to allow for more experimental flexibility, we note that this is not a requirement for achieving static ZZ suppression between coupled qubits, for which a fixed-frequency transmon would be sufficient (See Sec. 2.4.2).

As in the previous transmon device, the gradiometer loops of neighboring qubits are rotated by 90 degrees to match the orientation of the dFBLs in the sample enclosure. The readout spiral resonators fabricated on the opposite side of the substrate (not shown) also remain the same, with target frequencies between 9.7 GHz and 10.2 GHz.

We perform cryogenic microwave measurements on device A and device B in separate cooldowns, using the same experimental setup and techniques as in Chapter 4. Details of the cryogenic and room-temperature measurement setup can be found in Appendix B. In particular, we use the original version of our sample enclosure with individual coaxial readout ports for measuring device A, while we use the modified version which employs a 3D-integrated, four-way readout multiplexer, presented in Sec. 4.5, to measure device B. DC currents are applied to the dFBLs via isolated battery-operated voltage sources with output range ± 20 V.

The total line resistance of 20.08(1) k Ω , coming mainly from the resistors in the RLC filters installed at the 4K stage of the cryostat, converts the voltages to an approximate current range of ± 1 mA. An exception is made for IST1, which is instead biased using a low-noise current source due to its smaller mutual inductance to the dFBL (as predicted in Table 5.2), requiring larger currents in order to thread a full flux quantum through the IST loop.

In the remainder of this chapter we will focus on frequency-domain and time-domain characterization of individual qubits IST1-4 on device A, as well as qubit IST2c on device B. We defer any measurements and analysis related to qubit Trc and to the two-qubit dynamics of device B to Chapter 6.

5.3.3. Frequency selection

The spectrum of the RF-SQUID Hamiltonian is defined by the three energy scales E_C , E_L and E_J . The first two have already been fixed by the choice of IST geometry in Sec. 5.2, making E_J the only parameter left to determine. The most interesting operation point for the IST is at half flux bias, $\varphi_{\text{ext}} = 0.5$, where the anharmonicity is positive and the qubit is first-order insensitive to flux noise. We therefore aim for an E_J that results in $\omega_{01}(\varphi_{\text{ext}} = 0.5) \approx 4 - 5$ GHz, a typical range for our transmon qubits. These values of E_J are then targeted during the junction fabrication step.

We determine the target E_J required for 4 GHz and 5 GHz numerically⁵ and report these in Table 5.4 for each IST design, along with the ratio β . We also include the resulting anharmonicity, which varies significantly as a function of E_J . All IST designs have similar $E_C \sim 200$ MHz, resulting in $E_J/E_C \sim 30 - 200$. Interestingly, the largest values of E_J/E_C here also correspond to the largest relative anharmonicities. This is due to $\beta \propto 1/E_J$ becoming smaller and tending towards the highly anharmonic quarton limit (Sec. 2.3.3). We note that this is a striking departure from the well-known transmon scaling.

The dependence of α on E_J and E_L is likely to make it more challenging to achieve a small relative spread in α between identical devices compared to the transmon, which we determined to be within 1% in our four-transmon device (Table 4.1). On the other hand, this dependence also allows us to freely adjust α without having to modify E_C and the capacitive network of the system. This should ensure that capacitive couplings to external lines, readout resonators, and other

⁵The 1st order perturbation formulas in Eq. (2.17) are not reliable enough, since they fail to predict any meaningful value of E_J for small β due to the divergence visible in Fig. 2.4 (a).

Table 5.4: IST frequency targets. Required values of E_J to achieve target frequencies of 4 GHz and 5 GHz in the four implemented IST designs, predicted numerically using E_L and E_C converted from L and C in Table 5.1. The resulting ratio β and anharmonicity α are included.

	Target: 4 GHz			Target: 5 GHz		
	E_J (GHz)	β	α (MHz)	E_J (GHz)	β	α (MHz)
IST1	12.5	1.66	226	6.5	3.20	83
IST2	17.9	1.45	298	11.7	2.22	138
IST3	26.2	1.28	418	19.8	1.68	229
IST4	40.9	1.14	615	34.4	1.36	383

qubits can be kept similar for a wide range of IST frequencies and anharmonicities without requiring any design adjustments.

Another interesting feature of the IST is that the frequency at half flux bias increases when E_J decreases, which can be deduced from the negative sign in the plasma frequency $\omega_0 = (8(E_L - E_J)E_C)^{1/2}$ (see perturbation theory in Sec. 2.3.2). This is yet another fundamental difference to the transmon, where instead we have $\omega_0 = \sqrt{8E_JE_C}$. This means that a change in E_J leads to opposite shifts in frequency – something that has to be considered carefully when targeting frequencies and detunings of coupled IST-transmon systems. This will be further discussed in Chapter 6.

5.4. Spectroscopic measurements

5.4.1. Resonator spectroscopy

We begin by carrying out standard continuous-wave (CW) spectroscopy of readout resonators. If the IST qubits are functional, we should be able to observe flux dependence of the resonator frequency due to the dispersive qubit-resonator coupling. Just like the transmon, the IST qubit is expected to be periodic in full flux quanta Φ_0 due to the term $E_J \cos(\hat{\varphi} - \varphi_{\text{ext}})$ in its Hamiltonian.

Fig. 5.8 shows a low-power spectroscopy of the readout resonator belonging to IST1 as a function of applied dFBL current. As the qubit is tuned in frequency, the resonator is Lamb-shifted in tandem by an amount $\sim g_r^2/\Delta$. As will later be confirmed directly by qubit spectroscopy, the resonator frequency is always higher than the qubit in the devices presented here, $\Delta = \omega_r - \omega_{01} > 0$, such that the Lamb




Figure 5.8: IST resonator spectroscopy. Low power spectroscopy of IST1 readout resonator versus applied dFBL bias current. Shown is the derivative of the phase response of the VNA readout signal S_{21} . The resonator frequency is periodic in applied flux, which allows us to convert the current axis to a flux axis. The dashed black line is a simultaneous numerical fit to resonator and IST spectra (see main text for details). The avoided crossing is attributed to the IST transition $0 \rightarrow 2$.

shift is always positive. A maximum in ω_r therefore corresponds to a maximum in ω_{01} and to an integer number of flux quanta threaded by the loop, equivalent to $\varphi_{\text{ext}} = 0$, while a minimum corresponds to $\varphi_{\text{ext}} = 0.5$. The current axis can thus be easily converted to flux, and we obtain a mutual inductance $M = 0.97$ pH as the conversion factor in the case of IST1. The zero of the flux axis in Fig. 5.8 is chosen arbitrarily, since it is impossible to know in absolute terms how many multiples of Φ_0 are being applied to the SQUID loop. A fixed flux offset has to be removed, which is roughly constant for each qubit within one cooldown (subject to small fluctuations), but varies randomly from qubit to qubit and also between cooldowns for the same qubit. We observe the same in our transmon devices. We attribute this effect to a variable amount of flux trapped inside the SQUID loop each time a circuit is cooled down through its critical temperature, coming from stray fields in the environment of the experiment.

The resonator also shows a prominent avoided crossing – as we shall see below in Sec. 5.4.2, this is an anti-crossing with the second IST transition $0 \rightarrow 2$. While this does not affect the $0 \rightarrow 1$ transition directly, it limits our ability to read out the qubit in this region and also affects its anharmonicity. It is therefore advisable to avoid operating the qubit close to this region. In a future iteration, different qubit and resonator frequencies could be targeted to remove this avoided crossing entirely.

Table 5.5: Measured IST quantities. Resonator frequency ω_r , resonator linewidth κ , qubit frequency ω_{01} , and qubit anharmonicity α obtained from low-power spectroscopic measurements at the half flux bias point $\varphi_{\text{ext}} = 0.5$ of each IST qubit. The mutual inductance M between qubit and dFBL is determined from the flux periodicity of the resonator frequency. Numbers in brackets are the fit uncertainties on the last digit.

Qubit	$\omega_r/2\pi$ (GHz)	$\kappa/2\pi$ (MHz)	$\omega_{01}/2\pi$ (GHz)	$\alpha/2\pi$ (MHz)	M (pH)
IST1	9.730	0.76(6)	3.659	+238	0.97
IST2	9.874	0.78(9)	4.571	+150	1.47
IST3	9.992	0.27(8)	~ 4.7	–	1.27
IST4	10.132	0.23(2)	5.621	+280	1.32
IST2c	9.969	0.53(1)	3.785	+281	1.10

We repeat this spectroscopy for all resonators and report the measured resonator frequencies ω_r and linewidths κ at $\varphi_{\text{ext}} = 0.5$, as well as the mutual inductances M obtained from the current or voltage periodicity, in Table 5.5. The values of M found here are compatible with those estimated earlier from magnetostatic FE simulations, see Table 5.2. This shows that our relatively simple method of predicting M is useful even in the case of the IST where the SQUID loop geometry is rather convoluted. Although the experimental values closely match the case of complete field expulsion from the SQUID loop metal ($\mu_r = 0$), the effect of different μ_r cannot be distinguished from the effect of deviations in dFBL–substrate separation based solely on our experimental data, as discussed in Sec. 4.3.

5.4.2. Qubit spectroscopy

Next, we perform CW qubit spectroscopy by driving the resonator at the correct flux-dependent frequency, determined from the previous resonator spectroscopy, while scanning the frequency of an additional CW drive applied to the qubit via its corresponding dFBL. To better distinguish the fundamental qubit transition from additional features, we also sweep the power of the qubit drive. An example is shown in Fig. 5.9 for qubit IST1 biased at the half flux point. Only one transition at around 3.65 GHz is visible at low power, which we associate with the fundamental qubit transition $0 \rightarrow 1$. A second feature appears at a higher frequency when the power is increased; since we predict the IST to have positive anharmonicity at this bias point, we identify this as the two-photon driving of the $0 \rightarrow 2$ transition, yielding an anharmonicity of $\alpha/2\pi = 238$ MHz corresponding to twice the separation

between ω_{01} and $\omega_{02}/2$. More features, possibly three-photon transitions, appear at even higher powers.




Figure 5.9: IST qubit power spectroscopy. Spectroscopy of IST1 at $\varphi_{\text{ext}} = 0.5$ as a function of the qubit drive power at room temperature. The white dashed lines show the predicted frequencies of the three lowest qubit transitions from the numerical fit to the full flux-dependent qubit spectrum.

We repeat this measurement for all other IST qubits, and report the obtained qubit frequencies and anharmonicities in [Table 5.5](#). Broadly speaking, we observe frequencies and anharmonicities comparable to typical transmons. This is an important result, because it confirms that we have implemented a novel type of qubit that can be controlled with the same microwave electronics and with the same measurement schemes as in other existing cQED implementations. Moreover, the IST anharmonicity is positive at the half flux bias point, confirming our theoretical predictions.

It was not possible to obtain accurate spectroscopy data for IST3, because the avoided crossing between resonator and second IST transition is located very close to the half flux bias point, making the spectrum difficult to interpret. A broad transition around 4.7 GHz can be identified at low powers, which we cautiously interpret as the qubit transition.

To gain a complete picture of the IST spectrum, we perform qubit spectroscopy over the whole flux range. The resonator probing frequency at each flux step is fixed to the value obtained from the flux-dependent resonator spectroscopy in [Sec. 5.4.1](#). A single wide-range scan at low power and high resolution for IST1 is shown in [Fig. 5.10](#) (a). Only the main qubit transition is visible across the entire range. The transition becomes weaker and disappears in the frequency range 4.5 – 5 GHz, which we attribute to lower signal-to-noise ratios in proximity of the




Figure 5.10: Full spectrum of IST qubit. (a) Low-power, high-resolution qubit spectroscopy of IST1 versus applied external flux. Only one transition is visible at low powers. Dashed line is a simultaneous numerical fit of resonator and qubit spectra, showing excellent agreement with the data. (b) High-power spectroscopy in the region indicated by the rectangle in (a). The middle feature matches the two-photon transition ($0 \rightarrow 2$)/2 when using the same fit as in (a), while the top feature is interpreted as a sideband transition (SB, red dashed line), see main text for details.

resonator avoided crossing mentioned before and shown in Fig. 5.8. To observe higher-order transitions, we repeat the measurement in a small region around the half flux bias point but at a higher qubit drive power, resulting in the spectrum shown in Fig. 5.10 (b). Comparing this to the power-dependent spectroscopy of Fig. 5.9, we can easily identify the same two-photon transition. Crucially, the two frequencies ω_{01} and $\omega_{02}/2$ converge when moving away from the symmetry point, becoming almost indistinguishable – this is in perfect agreement with theory (see e.g. Fig. 2.5), which predicts a zero crossing of the IST anharmonicity. A third transition is also visible, but it is not immediately clear what this corresponds to. It has the same flux curvature as ω_{01} , but is shifted by about +270 MHz. We interpret this as being a sideband transition involving the qubit transition itself and at least one additional mode. This will be further discussed below.

Similar wide-range scans are also carried out for IST2 and IST2c. This measurement is not performed on IST3, given the unfortunate location of the resonator avoided crossing. IST4 suffers from a different issue: while the frequency at half flux bias is safely above $\omega_r/2$ and therefore no avoided crossing is observed in resonator spectroscopy, the resonator response becomes very weak and its frequency shifts by several 100 MHz (as opposed to $\sim 10 - 40$ MHz for the other resonators) as the qubit is tuned to the zero flux bias point, suggesting that the qubit and

resonator detuning becomes small and the two modes hybridize strongly. While this is inconvenient, we note that there is nothing in our data which suggests that IST3 and IST4 are clearly malfunctioning or misbehaving. Thanks to the insights gained here, it should be possible to avoid these issues in a second generation of IST devices by using more appropriate targets for IST and resonator frequencies.

Numerical fit of IST spectra

The experimentally determined flux-dependent spectra of qubits and resonators, such as those shown in Fig. 5.8 and Fig. 5.10 for IST1, provide a good dataset against which the Hamiltonian of Eq. (5.3) can be tested. As discussed in Sec. 5.2.4, parasitic inductor modes are expected to significantly distort the qubit spectrum, making their inclusion necessary for quantitative extraction of circuit parameters such as E_L , E_J and E_C . Given that the computational times on the desktop computer used here become unreasonably long when adding even just 4 or 5 modes in addition to the qubit, we limit ourselves to inclusion of the readout resonator, with the free parameters being the bare resonator frequency ω_r and resonator-qubit coupling strength g_r , and the two lowest frequency parasitic modes. To limit the number of free parameters and avoid over-fitting, we fix the parasitic mode frequencies to the values predicted from FE simulations, and also enforce that $g_{p,1} = \sqrt{2}g_{p,0}$ [186], leaving $g_{p,0}$ as the only free parameter associated with the parasitic modes.

We perform a simultaneous least-squares fit of the model Hamiltonian to the experimental data of both the qubit frequency and resonator frequency across the whole flux range, in addition to the second qubit transition where available. This provides us with estimates for the circuit quantities E_L , E_J , E_C , g_r and $g_{p,0}$, along with their fit uncertainties, which we report in Table 5.6 for qubits IST1, IST2, and IST2c. Using the obtained fit parameters, we add the numerical prediction of qubit and resonator frequencies to the spectroscopy plots of IST1 in Fig. 5.8, Fig. 5.9 and Fig. 5.10. We find excellent agreement with the data, indicating that the chosen Hamiltonian is suitable for describing our circuit. The model is also able to match the third transition frequency ($0 \rightarrow 3$)/3 with a feature visible in the power-dependent spectroscopy of Fig. 5.9, providing a self-consistency check of our method.

In the following we discuss the obtained fit parameters. Starting with E_C , we see that the values are similar for all three qubits, consistent with the fact that their qubit electrodes are overall comparable in size. On the other hand, the equivalent

Table 5.6: Extraction of IST circuit parameters. The full flux-dependent spectra of qubits IST1, IST2 and IST2c and corresponding resonators are fitted numerically with a Hamiltonian of type Eq. (5.3) to extract the RF-SQUID circuit parameters E_L , E_J , and E_C , as well as the coupling strengths g_r and $g_{p,0}$ to the resonator and to the lowest parasitic inductor mode, respectively. We also convert E_L to L for direct comparison with simulations, and report the ratios β and E_J/E_C . The numbers in brackets are the fit uncertainties on the last digit.

Qubit	E_C/h (GHz)	E_J/h (GHz)	E_L/h (GHz)	g_r (GHz)	$g_{p,0}$ (GHz)	L (nH)	β	E_J/E_C
IST1	0.245(4)	14.9(2)	22.2(3)	0.18(1)	0.90(1)	7.4(1)	1.5	61
IST2	0.242(2)	13.7(1)	26.7(2)	0.174(9)	0.95(1)	6.12(5)	1.9	57
IST2c	0.238(1)	19.4(1)	27.2(1)	0.20(1)	0.906(7)	6.01(3)	1.4	82

values of capacitance $C \sim 80$ fF are somewhat lower than the 92 – 94 fF predicted from FE simulations in Sec. 5.2.2. This is likely due to the simplified treatment of the meandering inductor for the calculation of the total capacitance, as shown in Fig. 5.2, which ignores its distributed nature. A more sophisticated model will be required to make better predictions of the IST capacitance and achieve a similar accuracy to the transmon case.

We calculate the equivalent linear inductance values L from E_L and include them in Table 5.6. These are slightly lower than the simulated values of 7.9 nH and 6.3 nH for designs IST1 and IST2, but it is encouraging that we are able to estimate L to within less than 10 % based only on classical FE simulations. Moreover, the inductances of qubits IST2 and IST2c are very close to each other, as expected from the fact that their inductor dimensions are identical. This provides a first indication that fabricating IST qubits with repeatable values of L should be feasible, but a much larger dataset will be required to obtain significant statistics on reproducibility. Note that the method of predicting L in Sec. 5.2.2 is based on the simulated capacitance C , which as just mentioned above deviates from the fitted value by around 15 %. Using the fitted C instead would give a higher prediction for L , leading to a larger discrepancy from the experimental result.

The ratios β and E_J/E_C are obtained from the fitted quantities and added to Table 5.6. We find that our IST qubits lie in the transmon regime $E_J/E_C \gg 1$, as expected. The inductive ratios β lie within the range of interest for our experiment, leading to anharmonicities that are comparable to our transmon qubits.

Lastly, we comment on the coupling strengths g_r and $g_{p,0}$. The resonator couplings are found to be comparable to our transmon devices (see Table 4.1), indicating that the ability to effectively read out the qubit state is retained in the IST. The

coupling $g_{p,0}$ to the lowest parasitic inductor mode is much larger – we attribute this to the large field overlap of the IST lumped mode and the inductor parasitic mode, as shown in Fig. 5.4.

Direct observation of parasitic inductor modes

So far, the existence and frequency of the parasitic inductor modes have been assumed from the FE simulation results. While their inclusion in the model allows us to achieve a good fit of the IST spectrum⁶, it could be argued that the parameter $g_{p,0}$ simply acts as an ad-hoc correction factor and does not necessarily have any physical meaning. We therefore want to investigate the parasitic modes experimentally and discuss whether our model is justified.

Given that we expect several parasitic modes in the range 10 – 20 GHz from FE simulations and that they should be strongly coupled to the IST mode according to our numerical fit, we can attempt to observe them directly in spectroscopy. We perform high-power, wide-range spectroscopy of qubit IST2c up to 20 GHz, shown in Fig. 5.11 (a). This utilizes the full frequency range of the 4-port virtual network analyzer (VNA) used for these experiments, although higher frequencies could be achieved by using a separate RF source as the qubit drive if required. At first glance, we find a rich spectrum with a large number of features. We start by identifying the qubit 01 and 02 transitions, based on separate spectroscopy data as discussed earlier in this section. Two prominent avoided crossings of the 02 transition with flux-independent modes at around 11.4 GHz and 14.8 GHz can also be seen, in relatively good agreement with the frequencies 11.8 and 15.2 GHz of the two lowest inductor parasitic modes of IST2 predicted from FE simulations (Sec. 5.2.4). We therefore assign the observed features to these parasitic modes. Taking these corrected frequencies into account, the spectrum of IST2c is numerically fitted again⁷ starting from the parameters in Table 5.6 and overlaid to the data as black dashed lines, as shown in Fig. 5.11 (b). The readout resonator is also added, although we note that it is virtually invisible in this spectroscopy data – this is expected, since the qubit drive is applied to the dFBL which is located far away from the resonator and shielded by the IST. It is interesting that the parasitic modes are detectable not only indirectly via their avoided crossings, but are also faintly visible when they are detuned from any of the qubit modes. This can be

⁶We have attempted to fit the data using the simple RF-SQUID Hamiltonian \mathcal{H}_0 , but found that we could never accurately reproduce the shape of the qubit spectrum.

⁷The qubit fit parameters E_J , E_L and E_J deviate only by a few % at most from the previously reported values, showing that the fitting method is robust.

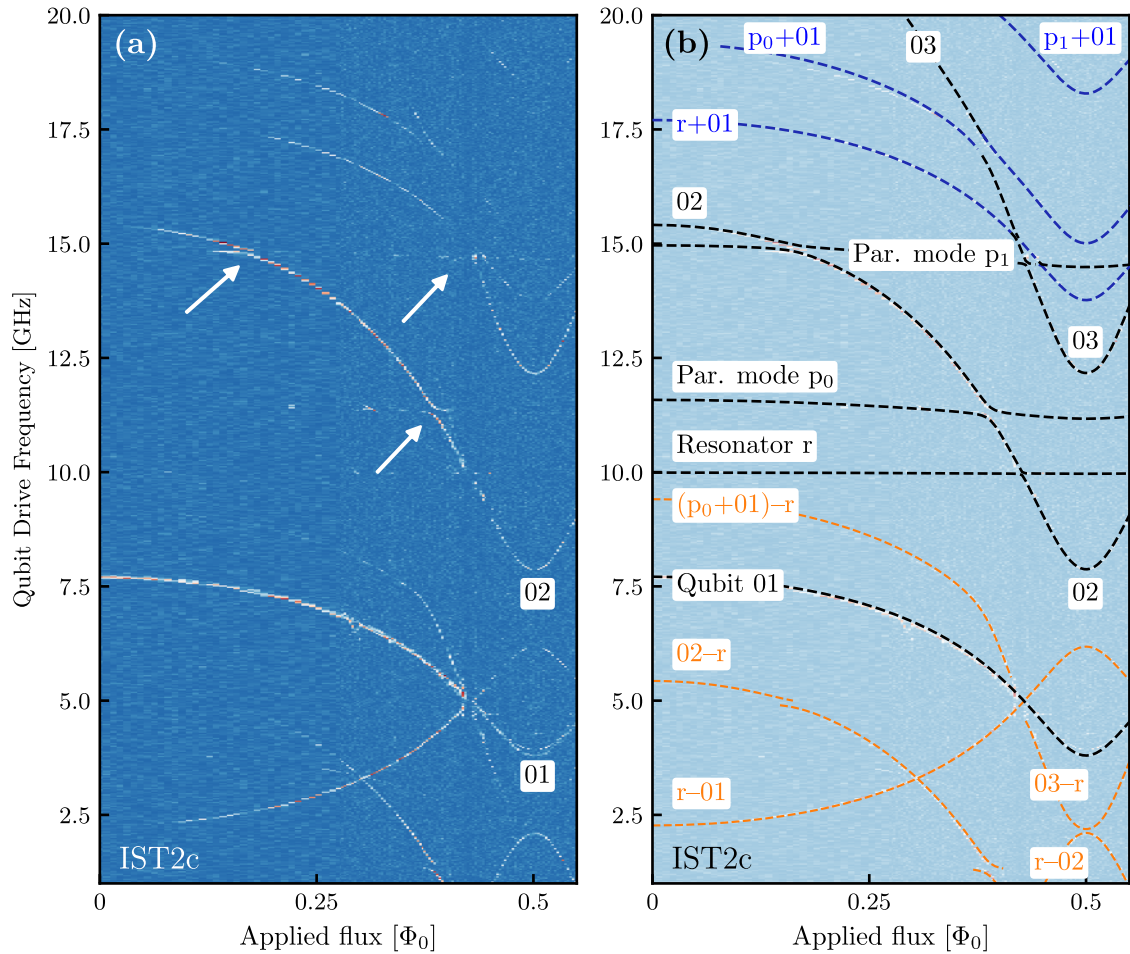


Figure 5.11: Wide-range IST spectroscopy. High-power spectroscopy of qubit IST2c as a function of flux up to 20 GHz. (a) Qubit transitions 01 and 02 are identified based on separate measurements. Horizontal features and avoided crossings are visible, indicated by the white arrows. (b) Same data as in (a), overlaid with numerical fits to all visible transitions. The horizontal features are assigned to parasitic modes. In addition to the main modes (black lines), we also observe blue and red sideband transitions (blue and orange lines). The labels indicate which modes are involved in the sidebands.

explained by the presence of a small direct coupling between inductor modes and dFBL, or alternatively by a weak hybridisation of the qubit and parasitic mode due to the large couplings strengths g_p .

Many additional features in the spectrum can be understood by considering that when any two modes at frequencies ω_i and ω_j are coupled to each other it is possible to drive *sideband* transitions between them [192, 193], either at their sum frequency $\omega_+ = \omega_i + \omega_j$ (blue sideband) or difference frequency $\omega_- = |\omega_i - \omega_j|$ (red sideband). These correspond to exciting both modes simultaneously in the former case, and exchanging an excitation between the two in the latter. With the knowledge of

qubit, resonator and parasitic mode frequencies, one can then try out different combinations of sum and difference frequencies to match the observed spectral features. This allows us to successfully label most of the remaining features of the spectrum in Fig. 5.11 (b), where blue and red sideband transitions are labelled in blue and orange, respectively. Note that all of these sideband transitions are already fully constrained by the fit performed on the main qubit and resonator data, so there are no free parameters. The overall agreement between data and numerical predictions across such a wide range of frequency and flux is striking. The fact that the frequency of the 03 transition is predicted correctly provides a further self-consistency check of the model.

The additional transition seen in the high-power spectrum of IST1 in Fig. 5.10 can also be interpreted as a sideband transition, namely the sideband $(p_0 + 01) - r$, if we assume the detuning between parasitic mode and resonator to be 270 MHz, placing the mode at around 10 GHz. This is consistent with the FE prediction of 10.5 GHz and the fact that parasitic modes in IST2c are found experimentally at lower frequencies than simulated.

We now look more closely at the parasitic modes themselves. We show detailed views of the spectroscopy of Fig. 5.11 in the two regions of interest in Fig. 5.12 (a) and (b). The contrast of these modes is very poor away from the avoided crossings, but it is possible to identify their maximum and minimum frequency. The same numerical fit as before is overlaid to the data, showing that while the qualitative behavior is captured correctly, the fitting routine grossly overestimates the coupling strength $g_{p,0} \approx 900$ MHz. By manual adjustment we find that a value around 500 – 600 MHz provides a better fit, but this in turn leads to poor fitting of the remaining qubit spectrum. One possible explanation for this is that a larger number of modes should be included in the model as shown numerically in Fig. 5.5, which would help to compensate for the smaller values of $g_{p,i}$. The FE simulations predict a third mode at 18.1 GHz, and we can indeed find an extremely weak feature at 17.8 GHz in the spectrum of IST2c, but without a quantifiable avoided crossing. It is therefore difficult to predict how this and higher modes will affect the spectrum, if at all.

To conclude, we have identified features in the IST spectrum that match the predicted frequency and behavior of inductor parasitic modes, but note that the coupling strengths and their influence on the qubit spectrum are not fully understood quantitatively. Regardless, we expect that these effects will play a lesser role if the IST inductor is miniaturized in a future design iteration, pushing




Figure 5.12: Spectroscopy of parasitic inductor modes. Details of the IST2c spectroscopy data from Fig. 5.11 showing additional modes of the system. These are interpreted as the (a) first and (b) second lowest parasitic inductor modes of IST2c. The same numerical fit curves as in Fig. 5.11 are added here as white dashed lines. The maximum and minimum frequencies of the modes are indicated by white arrows.

the modes higher up in frequency. Furthermore, we will see in the remainder of this chapter and in the next that these modes do not by themselves preclude the operation of IST qubits.

5.5. Time-domain measurements

The insights gained from the spectroscopic characterization of IST qubits allow us to now proceed to time-domain (pulsed) measurements and to obtain preliminary coherence data. From a practical point of view, once the qubit spectrum has been understood and the flux sweet spot and qubit frequency have been identified, tuning up and characterizing IST qubits is virtually identical to the procedure used for transmons. All of the same measurement techniques and pulse sequences used before can be applied here.

In the following we focus on IST1 and IST2 on device A, as well as IST2C on device B. We perform Rabi drive calibration and tune up X_π pulses, typically with Gaussian envelopes and 50 ns full width at half maximum (FWHM). Qubit readout is performed with a 5 – 10 μ s wide, low-power square readout pulse at the ground-state resonator frequency determined by spectroscopy. To obtain a suitable signal-to-noise ratio, each identical measurement is usually repeated and averaged 500 – 1000 times.

5.5.1. Lifetime and coherence at the flux sweet spot

We carry out standard measurement sequences of energy relaxation time T_1 , Ramsey coherence time T_2^* , and Hahn-echo coherence time $T_{2,E}$ on IST qubits at their half-flux bias point. To obtain statistical information on these quantities, each sequence is repeated N times in an interleaved fashion. In Fig. 5.13 we show data obtained for IST1, averaged across $N = 10$ individual repetitions. The data fits well to the standard exponential decay envelopes, and the Ramsey oscillations do not show signs of beating (i.e. multiple frequency components). We report the extracted mean values for T_1 , T_2^* and $T_{2,E}$ over N repetitions and their statistical spread (one standard deviation) for qubits IST1, IST2 and IST2c in Table 5.7. From the Ramsey data we can also extract the qubit frequency, and we report the standard deviation σ_f of measured values from their mean to quantify the frequency stability of the IST qubits.

Broadly speaking, we find that the IST qubits show promising coherence metrics that are roughly comparable to data from tunable transmon qubits from the same fabricated wafer and measured in the same cryogenic setup (see Table 4.1). This preliminary dataset, while limited in scope, does not point to the existence of a strong limiting factor for relaxation and decoherence unique to the IST circuit implemented here. Instead, we interpret this as a first indication that the gradiometric IST is a viable candidate for multi-qubit cQED experiments and should be pursued further. To our knowledge, the coherence times measured on qubit IST1 are higher than reported in any previous implementation of RF-SQUID qubits with a geometric inductive shunt, including the original transmon-like IST [79], the geometric superinductance qubits [74, 80] and the unimon [86].




Figure 5.13: IST energy relaxation and coherence time measurement. Averaged data over 10 repetitions of T_1 (left), T_2^* (middle) and $T_{2,E}$ (right) measurement sequences for qubit IST1 at the half-flux bias point. Solid black lines are a fit to the expected decay functions for each sequence.

Table 5.7: IST coherence statistics. Summary of energy relaxation time T_1 , Ramsey coherence time T_2^* and Hahn-echo coherence time $T_{2,E}$ for qubits IST1, IST2 and IST2c at the half-flux bias point. Tabulated values are the mean with standard deviation in brackets. Total number of datapoints N for each quantity and the total measurement run time are included. We also calculate the standard deviation of qubit frequency over the entire measurement run, denoted as σ_f , obtained from Ramsey data. Two datasets for IST2c are reported, one for each cooldown (CD) of device B.

Qubit	N	Meas. time (h)	T_1 (μs)	T_2^* (μs)	$T_{2,E}$ (μs)	σ_f (kHz)
IST1	10	0.5	37(3)	39(4)	46(3)	0.25
IST2	10	0.5	26(1)	30(2)	34(3)	2.8
IST2c (CD 1)	246	12.9	47(4)	28(3)	38(2)	4.5
IST2c (CD 2)	96	8.8	47(5)	12(2)	23(2)	7.2

Relaxation and coherence of qubit IST2c was measured repeatedly over the course of two cooldowns of device B; the two sets of values reported in the table show representative values for each cooldown. T_1 is virtually identical between cooldowns; this is compatible with the assumption that energy relaxation of our qubits is largely radiatively limited as suggested by the transmon data in [Sec. 4.5](#), although this is not enough to prove it. On the other hand, coherence times T_2^* and $T_{2,E}$ decrease significantly, which given the unchanged T_1 has to originate from a reduction of the pure dephasing time T_ϕ . At the flux sweet spot we do not expect dephasing due to flux noise to play an important role, and in [Sec. 5.5.2](#) we shall see that flux noise amplitudes do not change drastically in any case. Higher residual thermal photons in the readout resonator could contribute to increased dephasing, but we do not currently have enough data to reach a conclusive answer.

Temporal fluctuations

What the data presented so far does not capture is the nature of temporal fluctuation of relaxation and coherence times and their stability over timescales relevant for typical experiments (minutes, hours, days), which is a common issue in superconducting qubits and has been studied in detail using both fixed-frequency and tunable transmons [[194–196](#)]. For qubit IST2c, we plot T_1 , T_2^* and $T_{2,E}$ versus time in [Fig. 5.14](#) (a) for the dataset corresponding to the third row in [Table 5.7](#), collected over a period of around 13 hours. We see that during this time all three quantities fluctuate but do not show any obvious structure such as peaks, dips, or slow systematic drifts. We plot histograms of all quantities in [Fig. 5.14](#) (c) to better visualize their distribution. The Ramsey time distribution vaguely resembles a

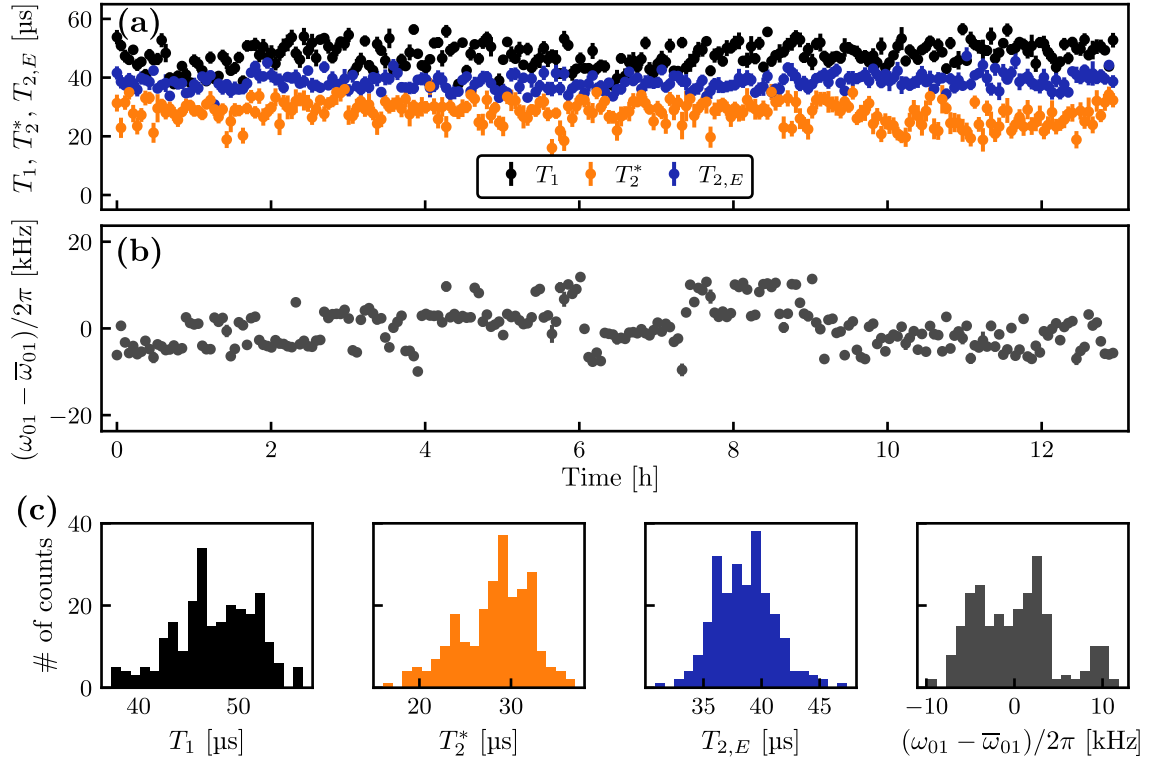


Figure 5.14: Time variation of IST relaxation, coherence and frequency. (a) Extracted T_1 , T_2^* and $T_{2,E}$ from $N = 246$ repeated measurements on qubit IST2c over the course of 12.9 hours during cooldown 1 of device B. (b) Qubit frequency variation from its average during the same time period, obtained from Ramsey oscillations. (c) Histograms of the datasets in (a) and (b).

double-peak structure, but more data would be required to make a significant statement. Overall, it is encouraging to observe such stable parameters in a flux-tunable qubit, but we note that this is only one example dataset. Longer measurements on multiple qubits would be required to study long-term device performance more systematically.

We also extract the qubit frequency over the course of the measurement from the Ramsey fringes and plot the difference from its mean value in Fig. 5.14 (b). The ensemble of measured values has a standard deviation of $\sigma_f \approx 5$ kHz over the course of the experiment, but the histogram in Fig. 5.14 (c) shows that the distribution is clearly non-Gaussian. Instead the frequency shows signs of switching behavior between discrete values. The fact that the frequency occasionally jumps suddenly in between periods of relative stability indicates that the characteristic timescale of switching is longer than the acquisition time of one Ramsey trace (here 1 minute). Quasiparticle dynamics take place on much faster timescales not resolvable here [194] and therefore cannot explain our observations. Charge

dispersion due to slowly fluctuating offset charges, typical for island-based qubits, can also be ruled out due to the presence of the inductive shunt [59]. Instead, the frequency jumps could originate from state-dependent dispersive shifts caused by weakly coupled modes in the qubit environment. It was found earlier that the parasitic inductor modes of the IST can lead to frequency shifts, but these are expected to be much larger, and it is unlikely for the mode lifetimes to be on the order of minutes or longer. Instead, microscopic charged defects in the proximity of the qubit have been suggested as a candidate for such weakly-coupled, slow-switching modes and could provide a possible explanation [195, 197]. Lastly, changes in the applied DC flux bias could also shift the qubit frequency. Starting from the flux sweet spot, a change in flux by $1.5 \times 10^{-4} \Phi_0$ (or 6 mV on our voltage source, larger than the 1 mV instrument resolution) is required to shift the frequency by 10 kHz. The long-term stability of the DC source output should be investigated to conclude whether this is a potential factor or not. Strong switching electromagnetic fields in the laboratory environment could also interfere with the DC sources and the attached cables.

It is worth pointing out that transmon qubit Trc is coupled to the IST and can obviously influence the frequency via the dispersive ZZ shift, but we note that Trc is kept in its equilibrium state for the whole duration of the measurements shown here, and that the ZZ shift of roughly 30 kHz when both qubits are at their flux sweet spots (see Sec. 6.3) is larger than the observed range of frequency jumps.

5.5.2. Investigation of flux noise

Since the IST is flux-tunable, it is important to understand how qubit coherence is affected by flux noise present in the system. This has been discussed at length in Chapter 3, and in Sec. 5.2.3 we have made the prediction that based on the IST loop aspect ratio ρ , the sensitivity to microscopic $1/f$ flux noise should be higher than in typical transmons with much smaller ρ . We therefore repeat the measurement of relaxation and coherence times versus flux that was carried out on the transmon device in Sec. 4.4.2 to extract the amplitude A_Φ of $1/f$ flux noise in the case of qubit IST2c. The result of such a measurement is plotted in Fig. 5.15 (a), showing that the IST dephasing as a function of flux derivative of the qubit frequency agrees well with the simple model of $1/f$ noise. This measurement is carried out four times over the course of two cooldowns, as shown in Fig. 5.15 (b). We notice that the flux noise amplitude increases slightly over time, both between the two cooldowns

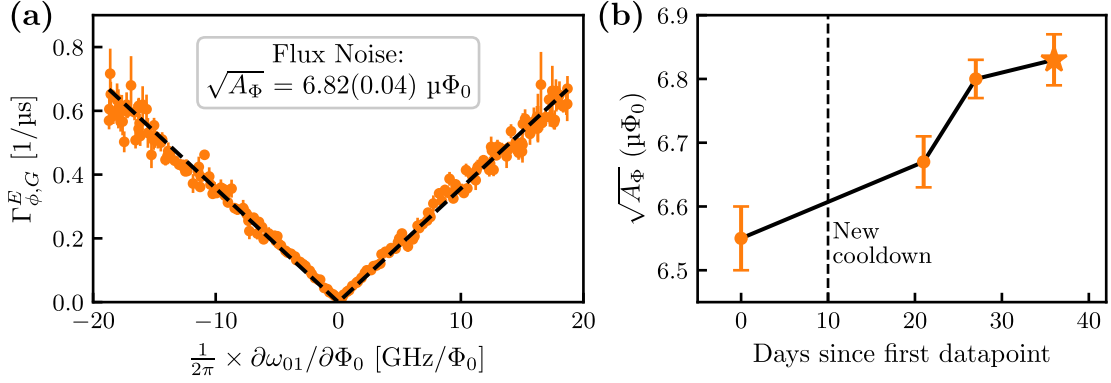


Figure 5.15: IST flux noise amplitude. (a) Measurement of echoed Gaussian dephasing rate of qubit IST2c as a function of flux, plotted against the qubit frequency flux dispersion. The data (orange points) fits well to a model assuming $1/f$ flux noise (dashed lines), from which $\sqrt{A_\Phi}$ can be extracted. (b) The same measurement is carried out four times across two cooldowns of the same device. The star denotes the datapoint obtained from (a).

and within the same cooldown, suggesting a slow drift in the magnetic environment of the qubit.

The values of $\sqrt{A_\Phi}$ observed for qubit IST2c are only slightly higher than the average of $6.3 \mu\Phi_0$ in the transmon device (see Table 4.1). This is a striking result if we consider the prediction made in Table 5.3 that the IST should be subjected to roughly five times larger microscopic flux noise amplitudes than our transmons. If the model is to be relied upon, we have to conclude that flux noise in our system is dominated by other sources, e.g. current noise on the dFBL. This could be mitigated by improved filtering on the DC bias lines. We expect that due to the gradiometric topology, our qubits should be relatively insensitive to global magnetic fields originating far away from the qubits themselves, as was shown for gradiometric fluxonium qubits in strong magnetic fields [198].

If flux noise predominantly couples to the SQUID loops via the dFBL, we then expect the flux noise amplitude to scale with the mutual inductance M according to Eq. (3.1). IST2c has about half the value of M of our measured transmons, meaning that all else being equal we should in fact observe a *smaller* value of $\sqrt{A_\Phi}$ in the experiment, which is clearly not the case. In general, we should assume that multiple noise sources are present at any given time and whose power spectral densities are added together. Assuming that both the current noise and the local universal flux noise have a $1/f$ spectrum, the total $1/f$ noise amplitude measured in the experiment will be $\sqrt{A_\Phi} = \sqrt{M^2 A_I + A_{\Phi,0}}$. We can reconcile this model with the experimental data by setting $\sqrt{A_{\Phi,0}} \approx 1.2 \mu\Phi_0$, compatible with the typical

strength of universal flux noise measured in a variety of experiments [152, 153], and $\sqrt{A_I} \approx 6.5$ nA for the added current noise. While this seemingly explains our findings, we note that our DC voltage sources have a white noise floor with total noise around 0.5 nA in a 1 kHz bandwidth according to the manufacturer specifications, which invalidates the assumption of $1/f$ noise being injected by the DC source. Further studies are required to pinpoint the exact mechanisms leading to the observed flux noise in our experiments. More broadly, the IST could be employed in the future to systematically study sources of flux noise and the validity of existing flux noise models, due to its particular geometry that unlocks a much wider range of SQUID parameters compared to what has been achieved in transmons [152].

The observed flux noise amplitudes for IST2c in the range $6.5\text{--}7 \mu\Phi_0$ compare favourably to values reported for the unimon qubit ($6\text{--}15 \mu\Phi_0$) [86] and geometric superinductance qubits with large aspect ratios ($100\text{--}800 \mu\Phi_0$) [74, 80].

5.6. Summary

Over the course of this chapter we have presented the design, simulation, and experimental characterization of the gradiometric IST, a realization of the RF-SQUID qubit in the low-inductance regime $\beta \gtrsim 1$. The IST is a promising qubit candidate due to its simplicity as a single-junction circuit, its insensitivity to charge noise, and its positive anharmonicity that complements the transmon, with the possibility of also drastically increasing the anharmonicity towards the quarton regime.

Simple FE simulations allow us to roughly predict the IST circuit properties, such as the circuit lumped inductance L , capacitance C , mutual inductance M to the dFBL, and the frequency of parasitic inductor modes. In conjunction with the experimental data obtained on the first-generation IST devices measured here, these simulations should enable us to make future design iterations even more predictable. An important goal will be to miniaturize the shunting inductance by using a thinner inductor wire. This will be beneficial to most relevant properties of the IST, from increased L and M to weaker influence of parasitic inductor modes.

From the point of view of fabrication, our IST does not require the introduction of any new materials or additional fabrication steps. This makes the IST easily accessible to every research team or institution that already has the expertise required to fabricate single-junction transmon qubits. While the IST design

presented here is tailored to our coaxial architecture, we predict that adding an appropriately designed meandering inductor should be possible in a large variety of physical circuit implementations, making the IST a drop-in building block that instantly expands the cQED toolbox at our disposal.

We successfully measured IST qubits in our existing coaxial architecture with off-chip flux control, showing that these novel qubits are readily compatible with existing transmon-oriented platforms and techniques. The IST shows similar levels of coupling to external lines and to readout resonators as our transmons, making their operation straightforward. We attribute this both to their essentially equivalent low-energy physics (both behave as weakly anharmonic, plasmonic Duffing oscillators), and to their similar geometry which leads to a mostly equivalent network of capacitive couplings. The frequency spectrum of the IST can be understood both qualitatively and quantitatively. Coherence properties and flux noise amplitudes are comparable to the tunable transmons presented in [Chapter 4](#), showing that the IST is a promising candidate for further cQED experiments. The similar coherence metrics of our IST and tunable transmon further suggests that the current limitations to qubit relaxation and dephasing might be similar in these two cases and that ongoing efforts in the research field to improve transmon performance could also benefit the IST.

6.

Interactions in a multi-species qubit system

6.1. Introduction

After having demonstrated a fully working implementation of the gradiometric IST at the single-qubit level, we now pose ourselves the question of how to realize qubit-qubit entanglement.

Broadly speaking, the IST is a weakly anharmonic, plasmonic, flux-tunable qubit that can be treated qualitatively as a Duffing oscillator in its low-energy physics. We therefore expect that ISTs can be entangled with each other using similar protocols as routinely used in all-transmon systems. While the higher anharmonicity allowed in the IST should be beneficial for gate fidelities and suppression of leakage, coupling two ISTs does not seem to offer a fundamental advantage at first sight. Therefore, we go one step further and choose to study the behavior of a multi-species system consisting of an IST coupled to a transmon. This allows us to not only show that the IST can be entangled with a different type of qubit, but also to investigate static ZZ suppression due to opposite anharmonicities as discussed in [Sec. 2.4.2](#).

In the following, we present experimental results on a coupled transmon-IST device. After performing basic characterization, we map out the ZZ interaction landscape across the IST flux range and compare it to numerical simulations. We then investigate interactions and transitions available in this system with the goal of implementing and benchmarking a two-qubit entangling gate.

6.2. Coupled device overview

Here we study IST device B presented in [Sec. 5.3.2](#) (see also the schematic drawing in [Fig. 5.7](#)). An optical image of the two coupled qubits is shown in [Fig. 6.1](#) (a). The qubits are coupled via simple capacitive arms that extend from their respective outer electrodes, with a lateral gap of 10 μm between the arms. We employ a gradiometric transmon with the same design as in [Chapter 4](#), and the IST




Figure 6.1: Coupled IST-transmon device. (a) Optical image of IST device B, consisting of a tunable transmon and an IST coupled directly to each other via capacitive arms extending from their outer electrodes. (b) Transmon spectroscopy versus applied flux on the IST, showing an avoided crossing with splitting $2J/2\pi = 17.15(5)$ MHz. The transmon is biased at its flux sweet spot, $\phi_{ext} = 0$.

corresponds to design variation IST2 with $L \approx 6$ nH. This device was measured during two separate cooldowns as discussed in [Sec. 5.5](#), but all the data presented in this chapter was collected during the second cooldown.

We carry out a general characterization of this device with both qubits biased at their respective flux sweet spots, and summarize the results in [Table 6.1](#). Basic single-qubit characterization of the IST was already presented in [Chapter 5](#) (where it was labelled as IST2c) and is reported here for convenience. As expected, the two qubits have opposite signs of anharmonicity, which also lead to opposite signs of dispersive shift χ of their readout resonators. Circuit parameters such as E_J , E_C , E_L are extracted from a numerical fit to the flux-dependent qubit spectra, as described previously for both transmons and ISTs.

Relaxation and coherence statistics are broadly in agreement with our previous data on individual transmons and ISTs ([Table 4.1](#) and [Table 5.7](#)). Note that coherence was actually higher for both qubits during the first device cooldown (average T_2^* and $T_{2,E}$ of 30(3) and 35(4) for the transmon, and 28(3), 38(2) for the IST respectively), but we report the lower numbers here for consistency with the rest of the dataset in [Table 6.1](#). The reason for the lower coherence is currently not understood, but in our experience it is not uncommon to see fluctuations in these metrics between cooldowns. Overall, these numbers give a first indication that the addition of the coupling arms does not significantly reduce coherence compared to isolated qubits.

Table 6.1: Main parameters of coupled IST-transmon device. All quantities valid for both qubits biased at their flux sweet spots ($\varphi_{ext} = 0$ for transmon, $\varphi_{ext} = 0.5$ for IST) except for exchange couplings J . Numbers in brackets denote measurement uncertainty for measured quantities, fit uncertainty for numerically extracted parameters, and statistical variation for the repeated coherence measurements.

	Quantity	Unit	Transmon	IST
<i>Basic quantities</i>	Qubit frequency $\omega_{01}/2\pi$	(GHz)	5.644	3.785
	Anharmonicity $\alpha/2\pi$	(MHz)	-220	+281
	Resonator frequency $\omega_r/2\pi$	(GHz)	9.843	9.969
	Resonator linewidth $\kappa/2\pi$	(MHz)	0.52(5)	0.53(1)
	Dispersive shift $2\chi/2\pi$	(MHz)	-1.5(1)	+1.2(1)
	Mutual inductance M	(pH)	1.71	1.10
<i>Numerical fit</i>	Charging energy E_C/h	(GHz)	0.203(2)	0.238(1)
	Total Josephson energy E_J/h	(GHz)	21.2(2)	19.4(1)
	Inductive energy E_L/h	(GHz)	-	27.2(1)
	SQUID asymmetry d		0.20(2)	-
	Ratio E_J/E_C		104	82
	Ratio $\beta = E_L/E_J$		-	1.4
	Resonator coupling $g_r/2\pi$	(MHz)	154(1)	201(10)
<i>Coherence stats.</i>	Relaxation time T_1	(μ s)	38(9)	47(5)
	Ramsey coherence time T_2^*	(μ s)	18(2)	12(2)
	Hahn-echo coherence time $T_{2,E}$	(μ s)	26(3)	23(2)
	Frequency variation σ_f	(kHz)	5.5	7.2
	# of repetitions		130	96
	Total run time	(h)	11.2	8.8
<i>Coupling</i>	Exchange coupling $J/2\pi$	(MHz)	10.93(4) at ~ 3.79 GHz 17.15(5) at ~ 5.64 GHz	
	Fixed cap. coupling $g_c/2\pi$	(MHz)	9.7(1)	
	Static ZZ shift $\zeta/2\pi$	(kHz)	27(1.8)	

Measurement of exchange coupling

Thanks to the particular frequency configuration of these two qubits, the IST can be tuned to the transmon sweet spot frequency and vice versa. It is therefore easy to directly obtain the magnitude of the exchange coupling J between the qubits via spectroscopy, as shown in Fig. 6.1 (b). Here the IST flux bias is tuned until the IST frequency crosses the transmon, resulting in a clear avoided crossing in the transmon signal. The two lines, separated by $2J$, correspond to the hybridised states of the two-qubit system. We fit the two branches [199] and obtain $J/2\pi = 17.15(5)$ MHz. This coupling is slightly larger than in typical experiments with statically coupled

transmons [110, 111, 129, 200], where the tendency is to have smaller $J < 10$ MHz to avoid excessive static ZZ due to the unfavourable scaling $\zeta \sim J^2$. The large J is in fact desired in our case, as it helps to measure ζ more easily when qubits are far detuned and enables a more stringent test of the theory, in addition to ultimately allowing for faster gates.

The measurement of J can be repeated by swapping the roles of the two qubits. This leads to a different value $J/2\pi = 10.93(4)$ MHz measured at the IST sweet spot frequency. This difference is due to the frequency-dependent nature of J , which in general can be expressed as $J = N_{\text{zpf},A}N_{\text{zpf},B} \times g_c$ (see also Eq. (2.25)), with $N_{\text{zpf},i}$ the charge zero-point fluctuations of qubit $i = 1, 2$. Their product is roughly proportional to the average qubit frequency ($N_{\text{zpf}} \propto E_J^{1/4} \propto \omega_{01}^{1/2}$ for transmons). The parameter g_c is the fixed capacitive coupling strength between the two qubits, which is purely a function of the capacitive network of the circuit. From our device parameters and the measured J values, we can extract $g_c = 9.7(1)$ MHz.

6.3. Static ZZ interaction

We now turn to one of the main topics of interest in our coupled device, namely the investigation of static ZZ interaction. As discussed in Sec. 2.4.2, when the qubit-qubit detuning Δ is much larger than the coupling strength J , the qubits are effectively dispersively coupled (much like the usual regime of qubit-resonator operation), and the only noticeable effect is a renormalization of the bare states $|ij\rangle = |i\rangle_{\text{Tr}} \otimes |j\rangle_{\text{IST}}$ with eigenenergies $E_{ij} = E_i + E_j$ to dressed states $|\tilde{i}\tilde{j}\rangle$ with slightly shifted energies \tilde{E}_{ij} . In this notation, i and j indicate the number of excitations in the transmon and IST, respectively. Describing the new eigenstates in terms of individual excitations is a reasonable choice as long as $J/|\Delta| \ll 1$ and the states do not hybridize significantly. The dispersive shifts are what cause qubit transition frequencies to be dependent on the state of the other qubit, as depicted qualitatively in Fig. 6.2 (a). Eigenstates in the computational subspace spanned by $\{|\tilde{00}\rangle, |\tilde{01}\rangle, |\tilde{10}\rangle, |\tilde{11}\rangle\}$ are affected by dispersive interactions with higher states, most prominently within the two-excitation manifold. The frequency of an individual qubit transition will now be different depending on the starting state, as denoted by the blue and orange arrows. This difference is what we commonly define as the static ZZ shift ζ , see Eq. (2.26).

To determine ζ we use the simple conditional Ramsey sequence shown in Fig. 6.2 (b). One qubit (in this case the transmon) acts as control, starting

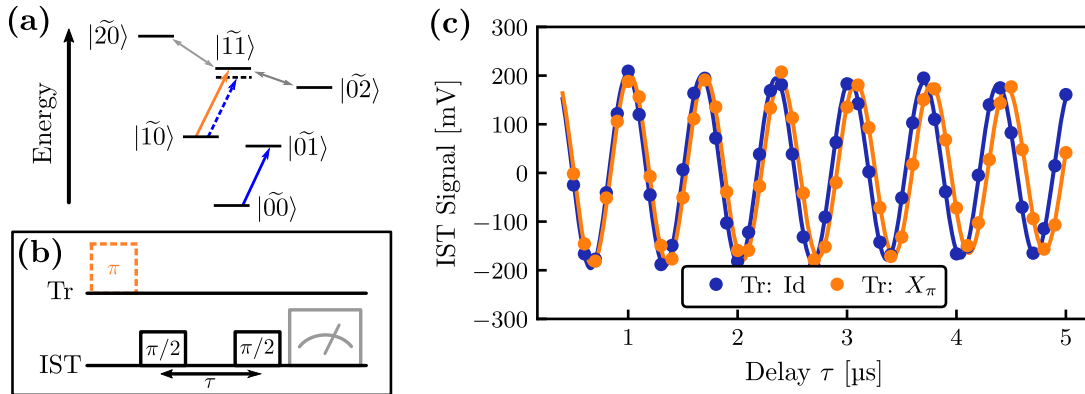


Figure 6.2: Static ZZ interaction between transmon and IST. (a) Qualitative energy level diagram of the coupled transmon-IST system showing dressed states $|\tilde{i}\tilde{j}\rangle$ up to the two-excitation manifold. The ZZ shift ζ is equal to the frequency difference between the transitions denoted in orange and blue. (b) Conditional Ramsey pulse sequence for determination of ZZ shift, repeated with/without an X_π pulse on the transmon (orange box). (c) An example of the measured conditional Ramsey oscillations on the IST. The fit to the two Ramsey traces yields a frequency shift of $\zeta/2\pi = 27.0(1.8)$ kHz.

in either the ground or excited state. The other qubit (here the IST) undergoes a standard Ramsey sequence, from which the oscillation frequency is fitted. ζ corresponds to the difference in oscillation frequency between the two cases. The choice of control qubit is arbitrary, as the definition of ζ in Eq. (2.26) is symmetric with respect to qubit labelling. Fig. 6.2 (c) shows the result of such a measurement for both qubits at the sweet spot. Even for short sequence durations up to 5 μs it is already possible to detect a clear detuning between the two traces, in this case $\zeta/2\pi = 27.0(1.8)$ kHz, where the error bound originates from the measurement uncertainty of the individual Ramsey frequencies. Although a longer sequence would determine the frequency more precisely, one should not forget that due to the finite T_1 of the control qubit the Ramsey signal will increasingly contain both frequency components, complicating the analysis. We choose $\tau \leq 5 \mu\text{s}$ as a reasonable middle ground.

The measured ZZ shift at the sweet spot is very small compared to the bare exchange coupling J , which is to be expected from the very large detuning Δ . To perform a simple consistency check with the theory, consider the perturbative expression for ζ in Eq. (2.27), valid for two weakly anharmonic oscillators. This predicts a shift of ~ 11 kHz for our device, while for a hypothetical device where the IST is replaced by a transmon with the same frequency but opposite anharmonicity we obtain ~ 65 kHz. The measured value lies between these two, meaning that we cannot reach a clear conclusion on whether our device has smaller ZZ than

an equivalent transmon-transmon counterpart and is therefore achieving built-in ZZ suppression. We believe that this is simply due to the fact that a first-order perturbative approach is ill-suited for making any quantitative predictions about an IST with small inductance ratio β (and therefore large non-linearity), as was found in Sec. 2.3.2. Instead, we can obtain a numerical prediction for ζ directly from its general definition in Eq. (2.26) by calculating the spectrum of the coupled IST-transmon system. The total Hamiltonian $\mathcal{H}_{tot} = \mathcal{H}_{Tr} + \mathcal{H}_{IST} + \mathcal{H}_c$ includes the two single-qubit Hamiltonians from Eq. (2.8) and Eq. (5.3), as well as a transverse coupling term $\mathcal{H}_c = \hbar J(\hat{a}_{Tr}^\dagger - \hat{a}_{Tr})(\hat{a}_{IST}^\dagger - \hat{a}_{IST})$. We include the two lowest inductor parasitic modes in \mathcal{H}_{IST} to closely reproduce the experiment. All parameters in the Hamiltonian are determined independently from spectroscopic measurements, and the simulation is therefore fully constrained with no free fitting parameters. The result is $\zeta \approx 35$ kHz, closer to the measured value. This highlights once again that the nonlinearity of the IST requires a more careful analytical approach or a full numerical treatment in order to make meaningful quantitative predictions.

6.3.1. ZZ interaction vs. flux

We now move to the exploration of static ZZ interaction as a function of applied flux on the IST. As in the CSFQ-transmon experiment of Ref. 100, tuning the IST away from the half-flux sweet spot increases its frequency while at the same time lowering its anharmonicity until it becomes negative. Since $|\alpha_{IST}| > |\alpha_{tr}|$ at the sweet spot, we can tune α_{IST} through a "resonance" with $|\alpha_{tr}|$, hopefully eliminating the dispersive shift that gives rise to a finite ζ term [101].

To confirm this, we repeat the ZZ measurement presented in Fig. 6.2 while sweeping the IST flux around its sweet spot. We also swap the roles of transmon and IST, given that the IST coherence deteriorates very quickly away from the sweet spot making it unsuitable for a Ramsey experiment, while its T_1 remains largely unaffected. At each flux point, we re-calibrate the X_π pulse on the IST via Rabi amplitude calibration. The result is plotted in Fig. 6.3 (a), showing that ζ decreases away from the sweet spot and crosses zero near $\varphi_{ext} = 0.47$. Furthermore, by simulating the full Hamiltonian \mathcal{H}_{tot} as described earlier, we can add the numerical prediction to the data (black line) and observe that we are able to approximately match the location of the ZZ suppression point, even though the numerics seem to overestimate the steepness of ζ as a function of flux. This could be due to a slight error in the extracted fixed capacitive coupling g_c , which strongly affects ζ via the $\zeta \sim J^2$ scaling.




Figure 6.3: Static ZZ interaction versus IST flux. (a) ZZ shift ζ measured in proximity of the IST flux sweet spot, showing ZZ suppression near $\varphi_{\text{ext}} = 0.47$. The black solid line is the numerical prediction from the full Hamiltonian \mathcal{H}_{tot} of the system, with all parameters determined independently. (b) Wide-range measurement of ζ across the whole IST flux range. We plot the absolute value $|\zeta|$ on a logarithmic scale, and the x-axis is chosen to be linear in the IST frequency (top axis) instead of flux (bottom axis). The model agrees well with the data over many orders of magnitude in ζ .

This is one of the main results of this chapter – we can suppress the static ZZ interaction between a directly coupled transmon-IST pair, without the need for tunable couplers or external drives. Furthermore, its location in flux bias can be predicted numerically using independently obtained parameters. While the ZZ suppression point is not located at the IST sweet spot in this particular device, a slightly different choice of device parameters leads to the optimal case where ZZ can be minimized at the sweet spot itself, preserving coherence. This will be discussed in more detail in the next subsection.

To further explore the ZZ interaction, we collect data over the entire available IST flux range, save for regions where a clear signal cannot be obtained e.g. due to an avoided crossing with the readout resonator or with parasitic inductor modes of the IST. The whole dataset is shown in Fig. 6.3 (b). We plot the absolute value $|\zeta|$ on a logarithmic scale, since the ZZ shift varies by several orders of magnitude as the IST is tuned over more than 3 GHz of frequency, crossing the transmon ($\Delta = 0$) at around $\varphi_{\text{ext}} \approx 0.39$. We convert the x-axis to be linear in IST frequency instead of flux for better display. As expected, the ZZ shift generally increases as the qubits are brought closer in frequency. The maximum of ζ is reached when the IST frequency is lower than the transmon's by an amount $\Delta \approx \alpha_{Tr}$, which is similar in principle to a transmon-transmon system [201]. This corresponds to the

avoided crossing between $|\widetilde{11}\rangle$ and $|\widetilde{20}\rangle$. Note that the concept of ZZ (or dispersive) shift does not make sense when qubit transitions are near-resonant and hybridize strongly. While ζ itself is still a measurable physical quantity, in this case the definition in Eq. (2.26) gives us half of the direct exchange splitting between $|\widetilde{11}\rangle$ and $|\widetilde{20}\rangle$ [101].

We perform the same numerical calculation of ζ as before and add it to the plot in Fig. 6.3 (b). The agreement with the data over more than four orders of magnitude in ζ is remarkable, especially given that no fitting parameters are used here. The numerical model correctly predicts the ZZ behavior qualitatively and, for the most part, quantitatively. This gives us a strong indication that we are correctly modelling the basic physics of the coupled IST-transmon system. The dip in ζ very close to $\Delta = 0$, present in both data and simulation, is not inherent to the IST-transmon physics but is due to the avoided crossing of $|\widetilde{02}\rangle$ with the lowest parasitic mode of the IST (see Sec. 5.4.2). This can be verified by excluding the parasitic mode from the numerical Hamiltonian and repeating the simulation, as shown in Fig. C.6 in the appendix. These parasitic modes could be moved to higher frequencies in a future IST design iteration as discussed in Sec. 5.2.5, removing these effects.

6.3.2. Simulation of ZZ landscape

The next question to ask is whether the ZZ suppression in the measured device is merely a fortunate coincidence or whether it is a general property that can be achieved for a wide variety of circuit parameters. Given the relatively low computational requirements to simulate our two-qubit system (an advantage of using statically coupled qubits without any additional coupling modes), and the fact that the model seems to reliably predict the experiment as seen in Fig. 6.3, we can explore the ZZ landscape as a function of various system parameters to try and answer this question. We choose to exclude the parasitic inductor modes from the IST Hamiltonian in order to further speed up the computation, while keeping all other device parameters as in Table 6.1. While this causes the IST energy levels to shift significantly and deviate from the measured device, it does not change the qualitative picture of the coupled IST-transmon system.

We perform numerical calculations of ZZ shift ζ at the IST sweet spot starting from the same circuit parameters as in our experimental device and varying the IST Josephson inductance $L_{J,\text{IST}}$ and capacitance C_{IST} , resulting in the two-dimensional landscape of ζ shown in Fig. 6.4. We choose to plot the data against




Figure 6.4: Simulation of static ZZ interaction versus IST circuit parameters. Two-dimensional map of $|\zeta|$ at the flux sweet spot obtained numerically by varying the IST junction inductance $L_{J,IST}$ and capacitance C_{IST} . All other parameters are kept the same as in the experimental device, see Table 6.1. Black and white contours indicate constant values of qubit-qubit detuning Δ and of IST anharmonicity α_{IST} , respectively. The $\Delta = 0$ line is calculated for a hypothetical uncoupled system with $J = 0$. The blue star denotes the values of $L_{J,IST}$ and C_{IST} in the measured device.

the dimensionless quantities $\beta = L_{J,IST}/L_{IST}$ and C_{IST}/C_{Tr} . We also add a blue star to indicate the parameters corresponding to our experimental device. Two main features are visible: an almost horizontal line where $\zeta = 0$, and a region of maximal ζ extending diagonally. It is encouraging to observe that ZZ suppression can be achieved not only at isolated points, but across a continuous one-dimensional manifold in parameter space. The suppression happens at an almost constant value of β , where the opposite non-linearities of IST and transmon are balanced. Although not visible when plotting the absolute value, we have $\zeta < 0$ (transmon-like) for large β , and $\zeta > 0$ for small β . We also find that our experimental device has a β which is too small to achieve ZZ suppression at the flux sweet spot. This is partly due to the Josephson inductance $L_{J,IST}$ being smaller than the target value, caused by fabrication imperfections.

To better visualize how the IST properties change across this space, we overlay contour lines (isolines) of the IST-transmon detuning Δ in black and of the IST anharmonicity α_{IST} in white. Note that Δ here is relative to a fixed transmon frequency of 5.644 GHz. An important result is that the ZZ suppression manifold extends over arbitrary values of Δ – this means that it is possible to create ZZ-free




Figure 6.5: Flux dependence of ZZ landscape. The same simulation as in Fig. 6.4 is repeated for different values of external flux φ_{ext} on the IST. All axes and color scales are kept the same as before. The $\zeta = 0$ manifold shifts until it crosses the point corresponding to our experimental device (blue star) near $\varphi_{\text{ext}} \approx 0.47$, as in the experiment. Beyond this value of flux, ζ is finite across the whole parameter space.

IST-transmon pairs independently of the choice of detuning. This method of ZZ suppression therefore does not impose any constraints on frequency allocation between qubits, which is an important consideration when choosing connectivity and gate implementations in larger lattices [202]. Also, the ZZ suppression line is not parallel to the anharmonicity isolines. The simple condition $\alpha_{\text{IST}} = -\alpha_{\text{Tr}}$ suggested by the perturbative expression in Eq. (2.27) is therefore not sufficient to guarantee suppression of ZZ in practice. Rather, it is more convenient from a design perspective to think directly in terms of physical circuit parameters such as inductances and capacitances, which is the approach chosen in Fig. 6.4.

The region of maximal $|\zeta|$ is where the dressed $|\widetilde{11}\rangle$ state crosses either of the two non-computational states $|\widetilde{20}\rangle$ and $|\widetilde{02}\rangle$. Since $\alpha_{\text{IST}} > 0$ but $\alpha_{\text{Tr}} < 0$, both of these avoided crossings happen when the IST is below the transmon, i.e. $\Delta < 0$. This is in contrast to the transmon-transmon case, where the two avoided crossings happen on either side of the $\Delta = 0$ boundary. One consequence is that here the $\Delta > 0$ half-plane is well-behaved, making it easy to target a region of low ZZ without the risk of frequency collisions with non-computational states. An interesting phenomenon manifests itself when the anharmonicities match, $\alpha_{\text{IST}} = -\alpha_{\text{Tr}}$: the three states $|\widetilde{11}\rangle$, $|\widetilde{20}\rangle$ and $|\widetilde{02}\rangle$ then form a triple degeneracy point at detuning $\Delta = -\alpha_{\text{Tr}}$ [101], which is where the ZZ suppression manifold crosses the barrier of maximal $|\zeta|$, creating a saddle point. The hybridisation of $|\widetilde{11}\rangle$ with two non-computational states at the same time leads to an exchange coupling larger by a factor of $\sqrt{2}$ compared to the transmon-transmon case, offering the possibility of performing a very fast conditional phase gate by non-adiabatically flux-tuning in and out of this condition [50, 90, 101].

The only parameter that can be tuned in-situ during an experiment is the external flux, so it is illustrative to vary this and observe the changes in the ZZ landscape.

We therefore repeat the same two-dimensional ZZ map for several values of flux, as shown in Fig. 6.5. The $\zeta = 0$ manifold shifts down towards smaller values of β , morphing into an arc that traverses the location of our measured device (blue star), corresponding to the zero-crossing of ZZ obtained experimentally in Fig. 6.3. We explain the downward shift of this line with the rapid reduction of α_{IST} away from the sweet spot, which can be compensated up to a point by choosing a larger IST nonlinearity β . Eventually, this is not sufficient since the IST anharmonicity always tends to zero away from the sweet spot (see e.g. Fig. 2.5), and the ZZ suppression manifold then quickly vanishes from the parameter space, leaving a trivial landscape where the negative transmon anharmonicity dominates and $|\zeta|$ is always finite. Going further, the IST frequency keeps increasing and crosses the transmon, ending up on the other side of the high- ζ region.

This sequence of plots conveniently visualizes what happens to the ZZ interaction as the IST moves from positive to negative anharmonicity, creating an interesting landscape. Generally speaking, it should be possible to reach a $\zeta = 0$ point as a function of flux for an arbitrary IST-transmon pair, as long as the IST anharmonicity is large enough to begin with when biased at the flux sweet spot.

6.3.3. Implications for scaling

Is the static ZZ suppression method presented in this section suitable for scaling to larger systems? As mentioned previously, we can achieve the ideal $\zeta = 0$ case for arbitrary IST-transmon detunings. From a different perspective, if we fix a desired detuning Δ then Fig. 6.4 tells us that there is only one point in parameter space where we can suppress ZZ. We therefore need to estimate how close we can come to that point, and what amount of residual ZZ is acceptable for our system. It is difficult to make a general statement about this, since the effect of static ZZ on e.g. the fidelity of a two-qubit gate will depend on the particular gate implementation and its duration. What complicates things is that the gate itself can add a new, dynamical contribution to ZZ, such as in the case of the cross-resonance gate [130, 203].

Generally, the amount of residual static ZZ limits the on-off ratio of a two-qubit gate since it constitutes an always-on interaction even when no gate is activated. If we can perform the gate at a rate J , the on-off ratio will be $|J/\zeta|$. If we want a ratio of 10^3 , assuming $J \sim 10$ MHz similar to our device, we need the ZZ shift to be on the order of 10 kHz or less. Looking at Fig. 6.4, this puts a strict requirement on the accuracy of setting the IST Josephson inductance $L_{J,\text{IST}}$, while C_{IST} does

not play a significant role. Of course, one can also choose to operate slower gates with a lower value of J , which causes the ZZ interaction to decrease quickly due to the scaling $\zeta \sim J^2$ – but this is only advantageous as long as qubit coherence does not become the limiting factor.

The issue outlined here is analogous to the well-known problem of targeting qubit frequencies in fixed-frequency transmon devices, which is challenging due to the sensitivity of Josephson junction inductances to fabrication imperfections. Here we are ignoring the additional parameter L_{IST} , which we assume to be more robust than $L_{\text{J,IST}}$ since it is defined by a simple top-down lithography step, same as the capacitance C_{IST} .

Recently, a technique for post-fabrication laser-annealing of individual junctions has shown great promise in reducing the spread of room-temperature junction resistances, and therefore transmon frequencies, in large-scale processors with up to 65 qubits [85, 204]. This method can tune the room-temperature resistance of single Josephson junctions (which is proportional to the inductance L_J) to within $\pm 0.3\%$ of a given target. This accuracy would allow us to guarantee $|\zeta| \lesssim 10$ kHz at detunings $\Delta \geq +100$ MHz or $\Delta \leq -500$ MHz in the case of the device simulated in Fig. 6.4. If $\zeta \sim 50$ kHz is acceptable, virtually all practical detunings (at least $5 - 10 \times J$ away from the anti-crossings between $|\widetilde{11}\rangle$, $|\widetilde{20}\rangle$ and $|\widetilde{02}\rangle$) become accessible.

We believe that the laser annealing technique is applicable to the IST, since it only contains a single junction¹. We see this as a fundamental advantage of the IST compared to RF-SQUID qubit designs that rely on tightly packed junction arrays. The only caveat is that the junction resistance cannot be measured by a simple DC probe due to the presence of the inductive shunt, which means that the IST junction would have to be annealed blindly based on previous calibrations, reducing the accuracy of the technique. We note that a more detailed analysis beyond the scope of the present discussion is required to make a final conclusion; nevertheless, we believe this technique to be a promising avenue to enable larger lattices of coupled ISTs and transmons with reproducible ZZ suppression.

One additional point to consider is that the use of built-in ZZ suppression, while elegant on its own, does not preclude the use of *additional* methods for ZZ suppression. In fact, combining two independent techniques might be necessary to ensure good performance at larger scales in the long run. Recent examples are the combination of multipath coupling and ac-Stark based active ZZ cancellation in

¹Even our gradiometric transmon could be tuned in this way, since the two junctions forming the SQUID are separated by hundreds of μm , allowing the laser to target each one individually.

a seven-qubit device [129], as well as the combination of fluxonium qubits (high anharmonicity) with a tunable coupler [96].

6.4. Sideband transitions

6.4.1. Concept

So far we have only investigated the system at rest, i.e. when no external drives are applied. We now want to explore how we can create entanglement between the IST and transmon. The large detuning $\Delta \approx 1.9$ GHz makes it unfeasible to drive a gate (e.g. a conventional cross-resonance gate) efficiently when both qubits are parked at their sweet spots. We could leverage the flux tunability of either of our qubits to implement a flux-based gate, either using baseband flux pulses [49, 50] or with a parametric drive [56, 92]. Instead, here we want to make use of a property specific to the RF-SQUID Hamiltonian, namely the fact that biasing the RF-SQUID at a finite flux (i.e. away from either of the sweet spots) introduces odd terms in the phase potential, leading to a three-wave mixing interaction and to modified selection rules for transitions between energy levels. This was explored in the transmon-CSFQ experiment of Ref. 83, and it is also the same type of interaction that enables three-wave mixing in SNAIL-based parametric amplifiers [205, 206].

When $\varphi_{\text{ext}} \neq 0, 0.5$ the symmetry of the RF-SQUID Hamiltonian Eq. (2.10) with respect to parity $\hat{\varphi} \rightarrow -\hat{\varphi}$ is broken, and the expansion of the Hamiltonian reads [83]

$$\mathcal{H}_0/\hbar \approx \omega_0 \hat{a}^\dagger \hat{a} + \beta_c (\hat{a}^\dagger \hat{a}^\dagger \hat{a} + \hat{a}^\dagger \hat{a} \hat{a}) + \frac{\alpha}{2} \hat{a}^\dagger \hat{a}^\dagger \hat{a} \hat{a}, \quad (6.1)$$

with a new cubic non-linearity β_c in addition to the usual Duffing fourth order non-linearity α . Third-order terms of type $\hat{a}^\dagger \hat{a}^\dagger \hat{a}$ lead to non-zero transition matrix elements between eigenstates of the same parity, such as $|0\rangle \rightarrow |2\rangle$, which are not allowed as single-photon transitions in the transmon [40]. In the two-qubit system consisting of transmon and IST, this cubic nonlinearity therefore allows us to directly drive sideband transitions by applying a tone at frequency ω_{sb} to the IST either at the difference (red sideband) or sum (blue sideband) of qubit transition frequencies, such as $\omega_{\text{sb}} = \omega_{01} \pm \omega_{10}$, manifesting the three-wave mixing interaction. This is the same mechanism behind the IST-resonator or IST-parasitic mode sidebands that we observed earlier in IST spectroscopy (see Fig. 5.11 and related discussion). We emphasize that, in the absence of the cubic nonlinearity, these sidebands are only possible via two-photon processes [87, 193, 207], requiring much

higher drive powers, or via additional tunable coupling elements [117], requiring more complex circuitry.

6.4.2. Device metrics at finite flux bias

To turn on the cubic nonlinearity we bias the IST at $\varphi_{\text{ext}} \approx 0.37$, far away from either of its flux sweet spots, while keeping the transmon at $\varphi_{\text{ext}} = 0$ as before. The main device parameters at this operation point are summarized in Table 6.2. In this regime the IST has a (small) negative anharmonicity due to the interplay of the Josephson and inductive potential terms as discussed in Sec. 2.3.2. This particular choice of bias leads to a relatively small qubit-qubit detuning $\Delta \approx +250$ MHz that provides sufficient interaction strength while avoiding hybridization of the qubits via the direct exchange J . The drawback is that the static ZZ interaction is now much stronger, as expected from the data in Fig. 6.3 (b).

Table 6.2: Device parameters at the sideband operation point. Basic properties of the IST-transmon device when the IST is biased at $\varphi_{\text{ext}} \approx 0.37$. The continuous-wave (CW) tone for active ZZ cancellation is presented in Sec. 6.5.1.

	Quantity	Unit	Transmon	IST
	Qubit frequency $\omega_{01}/2\pi$	(GHz)	5.642	5.898(1)
	Anharmonicity $\alpha/2\pi$	(MHz)	-220	-103(1)
	Dispersive shift $2\chi/2\pi$	(MHz)	-1.5(1)	-2.5(1)
	Relaxation time T_1	(μs)	46(7)	25(2)
	Ramsey coherence time T_2^*	(μs)	20(3)	0.39(3)
	Hahn-echo coherence time $T_{2,E}$	(μs)	30(4)	1.2(2)
	Frequency variation σ_f	(kHz)	2.8	450
	# of repetitions			69
	Total run time	(h)		9.3
	ZZ shift $ \zeta /2\pi$	(kHz)	2200(100)	
	ZZ shift $ \zeta /2\pi$ (with CW tone)	(kHz)		< 4

The IST-resonator dispersive shift 2χ is now negative instead of positive, as expected from the flipped sign of α , but its value is unusually large compared to the transmon which has more than twice the amount of anharmonicity. One would naively expect a dispersive shift on the order of -1 MHz from the familiar expression $2\chi \approx 4g_r^2\alpha/\Delta_r(\Delta_r + \alpha)$ for a weakly anharmonic qubit [33], where Δ_r is the qubit-resonator detuning. Note that this formula does correctly predict -1.5 MHz for our transmon. This discrepancy is due to the cubic nonlinearity of the IST at finite bias, see Eq. (6.1), which invalidates the usual treatment of the

qubit as a Duffing oscillator. Indeed, the numerical fit of the exact Hamiltonian to the IST-resonator spectrum as presented in [Sec. 5.4.2](#) predicts a much larger dispersive shift of -2.7 MHz, very close to the measured value.

While the transmon decay and coherence properties are largely unaffected, the IST coherence is reduced by more than one order of magnitude compared to the sweet spot. This is expected from the presence of a finite flux noise amplitude A_Φ , as determined in [Sec. 5.5.2](#), and the steep slope of frequency versus flux at this operation point. While it is still sufficient to investigate two-qubit dynamics, we predict that the short coherence times will be a severe limiting factor for the fidelity of qubit operations.

6.4.3. Driving the transitions

The main sideband transitions of interest are shown schematically in [Fig. 6.6](#) (a). The labels denote the two-qubit unitary operations that can be implemented by using the corresponding sideband transition as a primitive. The simplest case is the red sideband at the qubit detuning $\omega_{\text{sb}} = \Delta$, which implements a direct excitation swap $|\widetilde{01}\rangle \rightarrow |\widetilde{10}\rangle$ and is therefore a primitive for an iSWAP unitary. In the two-excitation manifold, driving at either $\omega_{\text{sb}} = \Delta + \alpha_{\text{Tr}}$ or $\omega_{\text{sb}} = \Delta - \alpha_{\text{IST}}$ brings $|\widetilde{11}\rangle$ out of the computational space and into either $|\widetilde{20}\rangle$ or $|\widetilde{02}\rangle$, but driving for a full Rabi period returns the state to $|\widetilde{11}\rangle$ with an added phase ϕ , useful for implementing a CPHASE(ϕ) gate. The iSWAP and CPHASE sideband gates have been successfully demonstrated in the transmon-CSFQ system of [Ref 83](#). Furthermore, we can directly Rabi-drive the blue sideband $|\widetilde{00}\rangle \rightarrow |\widetilde{11}\rangle$ starting from the ground state and applying a tone at the sum frequency $\omega_{\text{sb}} = \omega_{01} + \omega_{10}$. This transition adds two excitations to the system at once and can be used to implement a gate called bSWAP, which has been demonstrated before either via two-photon drive in a statically coupled transmon-transmon system [[87](#)] or by flux-driving a tunable coupler between two transmons [[117](#)].

In our experimental setup, blue sideband drives are generated with the same up-conversion signal chain as regular qubit microwave drives. Red sideband drives on the order of a few 100 MHz are instead directly generated as baseband signals from a single AWG channel with sampling rate of 2 GSa/s and then combined with the microwave drive line at room temperature, before being sent to the device in the cryostat. See [Appendix B](#) and specifically [Fig. B.2](#) for more details. We use a soft square pulse shape, with 4.5 ns Gaussian rising and falling edges ($\sigma = 3$ ns, truncation after 3σ), and a variable square plateau time τ .




Figure 6.6: Sideband transitions between IST and transmon. (a) Qualitative energy diagram of the coupled IST-transmon system, showing red and blue sideband transitions allowed by the IST cubic non-linearity. The labels indicate the unitary two-qubit operations that can be implemented based on these transitions. (b) To probe the red sideband $|\widetilde{01}\rangle \rightarrow |\widetilde{10}\rangle$, the transmon is excited and a sideband tone of variable frequency ω_{sb} and duration τ is then applied to the IST. The resulting transmon and IST readout signals are shown in (c) and (d), respectively. The colormap is chosen such that blue (red) corresponds to the qubit in the ground (excited) state.

As a first example, we probe the $|\widetilde{01}\rangle \rightarrow |\widetilde{10}\rangle$ transition by preparing the system in $|\widetilde{10}\rangle$ with a X_π pulse on the transmon and then applying a sideband drive of duration τ and frequency $\omega_{\text{sb}} \approx \Delta$ to the IST, as shown in the pulse sequence in Fig. 6.6 (b). By varying ω_{sb} and τ we can observe the typical chevron-like pattern of coherent Rabi oscillations in the transmon and IST signals, shown in Fig. 6.6 (c)-(d). The excitation first present in the transmon is swapped back and forth between the two qubits at a rate $\Omega_{01 \rightarrow 10}/2\pi \approx 15$ MHz when driving on resonance. This allows us to fully swap the qubit excitation when driving for half a period, around $t \approx 33$ ns. This rate is dependent on several factors such as the voltage drive amplitude, the bare exchange interaction J , and the strength of the IST cubic nonlinearity β_c enabling this transition [83]. Faster rates should be possible in principle, but in our experiment we are limited by the AWG output amplitude and the amount of attenuation in the drive lines.

One can obtain similar chevron patterns when driving one of the red sidebands in the two-excitation manifold, $|\widetilde{11}\rangle \rightarrow |\widetilde{02}\rangle$ at $\omega_{\text{sb}} = \Delta + \alpha_{\text{IST}} \approx 150$ MHz or $|\widetilde{11}\rangle \rightarrow |\widetilde{20}\rangle$ at $\omega_{\text{sb}} = \Delta - \alpha_{\text{Tr}} \approx 460$ MHz. This requires both qubits to be excited with a X_π pulse to create the starting state $|\widetilde{11}\rangle$. For the sideband involving the second transmon excited state $|\widetilde{20}\rangle$, we obtain a Rabi rate $\Omega_{11 \rightarrow 20}/2\pi \approx 13$ MHz with the same pulse amplitude as used in Fig. 6.6. A more detailed discussion will be offered in Sec. 6.5.3 when calibrating a CZ gate with this sideband.

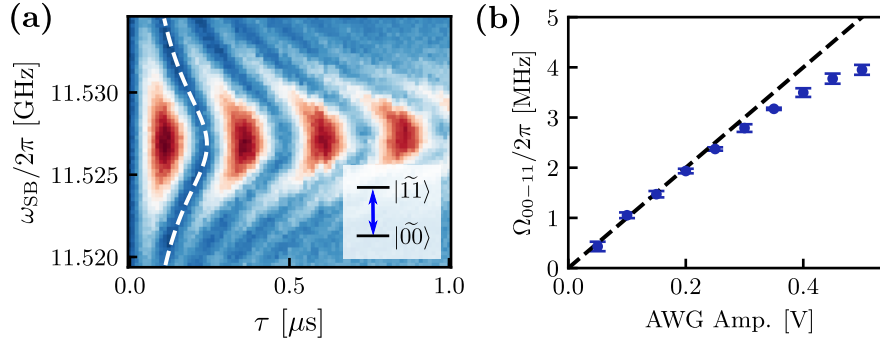


Figure 6.7: Blue sideband transition. (a) Oscillation pattern measured on the IST while driving the $|\widetilde{00}\rangle \rightarrow |\widetilde{11}\rangle$ blue sideband transition around the sum frequency of the two qubits. The white dashed line is a fit showing the full oscillation period at each drive frequency, obtained from fitting the Rabi rate $\Omega_{00 \rightarrow 11}$ (see main text). (b) Rabi rate versus sideband drive amplitude at the AWG output, obtained by repeating the measurement in (a). The black dashed line is a guide to the eye, showing the deviation of the data from a linear relation.

We now turn to the blue sideband $|\widetilde{00}\rangle \rightarrow |\widetilde{11}\rangle$, which does not exchange excitations between qubits, but instead adds and removes two excitations at once. Repeating the same scan of ω_{sb} and τ as before, this time centered around $\omega_{\text{sb}} \approx \omega_{01} + \omega_{10}$, we again observe coherent oscillations, shown in Fig. 6.7 (a) for the IST. The $|\widetilde{11}\rangle$ state is reached after $\tau \approx 120$ ns. If instead we drive for half that amount of time, we end up in a coherent superposition $(|\widetilde{00}\rangle + |\widetilde{11}\rangle)/\sqrt{2}$, i.e. we have generated a Bell state [87]. This means that the blue sideband transition offers a simple method of creating a maximally entangled Bell state with only a single 60 ns RF pulse starting from the ground state. To see whether this can be sped up further, we investigate the dependence of the Rabi rate $\Omega_{00 \rightarrow 11}$ on the AWG output amplitude of the sideband drive by repeating the chevron pattern measurement for different drive amplitudes. We fit the oscillation rate Ω at each value of ω_{sb} to the Rabi rate equation $\Omega^2 = \Omega_{00 \rightarrow 11}^2 + \Delta_{\text{sb}}^2$ to obtain the true Rabi rate $\Omega_{00 \rightarrow 11}$, where Δ_{sb} is the detuning of the drive from the true sideband frequency. The result is plotted in Fig. 6.6 (b), showing that the Rabi rate initially increases roughly linearly with amplitude (as expected from a single-photon transition) but then begins to saturate as the drive strength increases further. This indicates that we might already be approaching the limit of the effective interaction rate possible for this particular transition at the chosen operation point. The largest value measured here corresponds to the plot in (a).

6.5. Demonstration of a sideband CZ gate

The variety of sideband transitions we have explored offers an interesting playground for implementing various types of entangling gates with the IST qubit. In this section, we show as an example the calibration of a CZ gate based on the red sideband $|\widetilde{11}\rangle \rightarrow |\widetilde{20}\rangle$, largely following the procedure outlined in Ref. 83 for sideband gates in a transmon-CSFQ system. To our knowledge, this is the first demonstration of a two-qubit entangling gate employing a geometric inductively shunted qubit.

The IST is biased at $\varphi_{\text{ext}} \approx 0.37$, same as for the sideband experiments of Sec. 6.4. We start by removing the (large) static ZZ interaction with an ac-Stark shift drive, then calibrate and benchmark single-qubit gates, before moving on to the characterization of the CZ sideband gate itself.

6.5.1. Active ZZ cancellation

The relatively small detuning and large coupling J between the qubits leads to a large ZZ shift $\zeta > 2$ MHz at the chosen sideband operation point, which can degrade the fidelity of qubit operations by introducing coherent (unitary) errors. Fortunately, the same three-wave mixing term that enables single-photon sideband interactions can also be used to actively suppress the ZZ shift, as explained in Ref. 83. A continuous wave (CW) tone is applied to the IST at a frequency ω_{CW} slightly detuned from the $|\widetilde{11}\rangle \rightarrow |\widetilde{02}\rangle$ sideband. This off-resonant tone does not drive the transition itself but causes an ac-Stark shift of the involved energy levels and leads to complete cancellation of ζ when the frequency and amplitude of the CW tone is chosen correctly. The principle is similar to other implementations of ac-Stark shifts for ZZ cancellation in transmon [110, 129] and fluxonium [208] systems, with the crucial difference that these rely on primary transition frequencies, usually on the order of GHz or more. Instead, the IST three-wave mixing term allows us to use much lower frequency drives (a few 100 MHz) as for the red sideband transitions discussed here, which are easily implemented with baseband AWG pulses without requiring dedicated microwave sources. Furthermore, the CW drive in our case is added in software to the same waveform output as the other sideband drives, resulting in zero hardware overheads.

We observe the effect of the CW tone by performing spectroscopy on the transmon as a function of CW tone amplitude with a fixed value $\omega_{\text{CW}} = 180$ MHz. The spectroscopy is carried out twice, once with the IST in the ground state and once

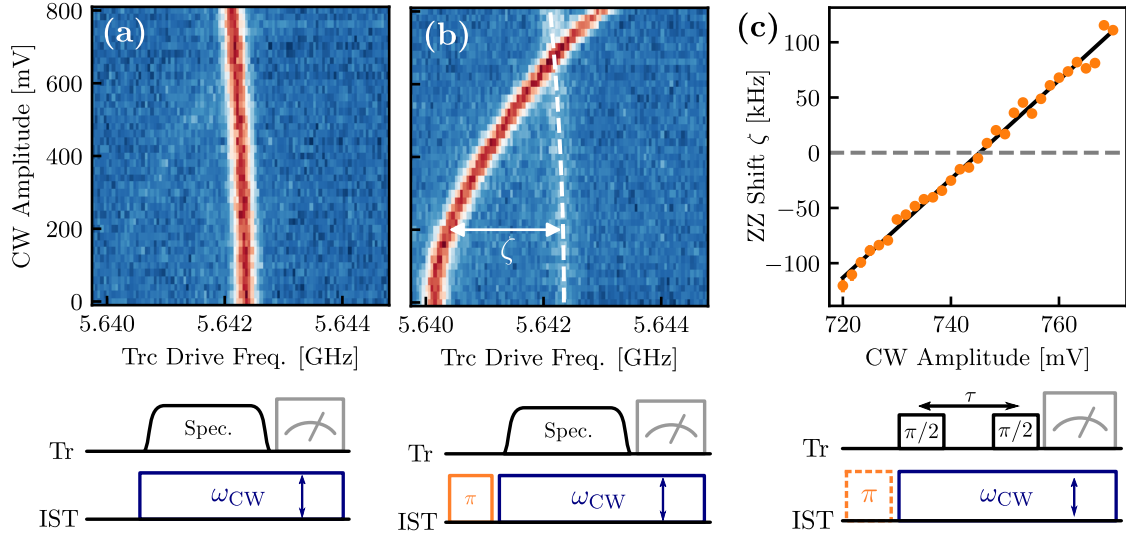


Figure 6.8: Calibration of CW tone for ZZ cancellation. Transmon spectroscopy at the CZ operation point versus amplitude of a 180 MHz CW drive on the IST, (a) without and (b) with an additional X_π pulse applied to the IST. Dashed line in (b) shows a quadratic fit to the transmon frequency from (a). The difference between the two curves is simply the ZZ shift ζ . (c) Fine-tuning of CW tone amplitude by Ramsey-based ZZ measurement. The optimal amplitude for ZZ cancellation is obtained from a linear fit. The pulse sequences for (a)–(c) are shown below the respective plots.

with a X_π pulse applied to it to induce a ZZ shift on the transmon, and the result is shown in Fig. 6.8 (a)–(b). As the CW strength is increased, the ac-Stark shifted transmon frequencies converge to the same value, making the transmon completely insensitive to the state of the IST – in other words, ζ has been suppressed. Note that the CW tone is applied to the IST, but the effect is here observed in the transmon thanks to their dispersive coupling.

The cancellation of ζ can be determined more accurately via conditional Ramsey on the transmon, see Fig. 6.8 (c). A linear fit to the data gives us an optimal CW amplitude of 748(1) mV in this case. While the trend seems robust, the residuals (difference between fit line and data points) are 4 kHz on average, larger than the errors observed for ZZ measurements at the sweet spot. We attribute this to the larger fluctuations $\sigma_f \sim 450$ kHz of the IST frequency at finite flux bias, obtained from repeated Ramsey data (Table 6.2), which affect the qubit-qubit detuning and the bare value of ζ , and thus the optimal CW tone parameters. We can therefore only give an upper bound $|\zeta| \lesssim 4$ kHz for the active ZZ suppression method once the CW tone has been fixed, corresponding to a still significant reduction in static ZZ strength by a factor of more than 500 compared to the bare value.

We also investigate whether the presence of ac-Stark shifts due to the CW tone has an adverse effect on qubit relaxation and coherence at the ZZ cancellation point. We obtain the following values of T_1 , T_2^* and $T_{2,E}$: 37(6), 18(1) and 25(3) μs for the transmon, and 22(1), 0.37(3) and 1.2(1) μs for the IST. Compared with the data in Table 6.2, the IST is not affected while the transmon seems to degrade slightly, but we note that the data here agrees well with the transmon reference values from Table 6.1, so we attribute these differences to random temporal fluctuations of qubit parameters rather than any effect of the CW tone itself.

6.5.2. Single qubit gate benchmarking

So far, single-qubit gates have been calibrated only approximately via a simple Rabi amplitude scan to obtain the pulse amplitude that optimally inverts the qubit population (in the case of a X_π gate). While this is sufficient for the rather simple pulse sequences presented until now, benchmarking a two-qubit gate requires well-calibrated single-qubit gates to minimize error contributions. We use a procedure similar to the one presented in Ref. 168. We choose Gaussian pulse envelopes with $\sigma = 15$ ns and truncation at 3σ , with an added 5 ns buffer after each pulse. This results in a total gate duration $t_g = 50$ ns. After a rough amplitude calibration and a Ramsey sequence to determine the qubit frequency precisely, we carry out a pulse train sequence which amplifies errors in amplitude [209]. We separately calibrate X_π and $X_{\pi/2}$ amplitudes. Subsequently, we minimize leakage contributions to error via derivative removal by adiabatic gate (DRAG) [51]. This technique is especially important when qubit anharmonicity is small, such as for the IST at this operating point, since any finite-length pulse will contain spectral components that can drive transitions outside of the computational space. Since the DRAG pulse shape optimization slightly modifies the total pulse amplitude, we then repeat the pulse train sequence. One can iterate back and forth until stable pulse parameters are reached – we found that two DRAG rounds with a subsequent pulse train recalibration are sufficient for our purposes. All of these calibration routines are performed individually on each qubit, in the presence of the CW tone for ZZ suppression discussed earlier.

We now characterize the single qubit gates via the standard technique of *randomized benchmarking* (RB) [210, 211], which is a resource-efficient method to obtain an estimate of gate errors. Briefly, standard RB pulse sequences consist of a variable number m of randomly chosen single-qubit gates from the Clifford group with 24 elements. At the end of the sequence, the inverse of the product of all

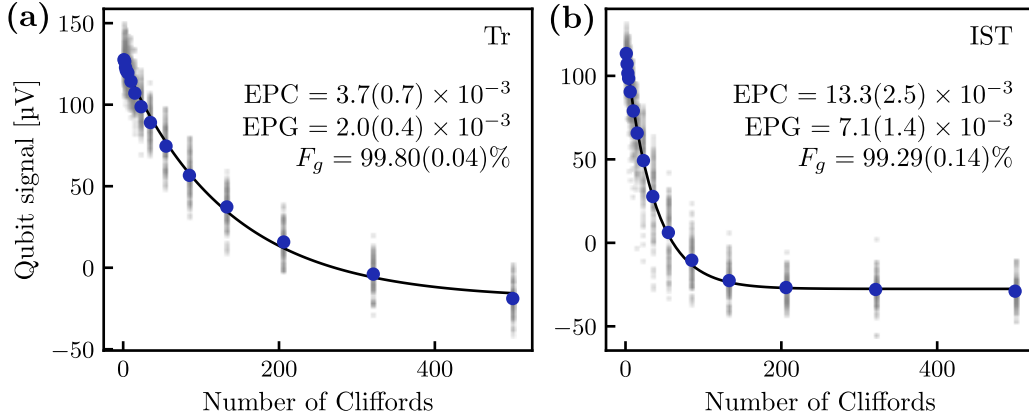


Figure 6.9: Single-qubit randomized benchmarking. Single-qubit gate randomized benchmarking (RB) performed individually on (a) transmon and (b) IST after gate calibration routine as explained in the main text. Raw datapoints for 100 different Clifford group samples at each Clifford sequence length are shown as grey points. The standard RB decay curve is fitted (black line) to the mean values at each sequence length (blue dots). Extracted error per Clifford (EPC), error per physical gate (EPG) and gate fidelity $F_G = 1 - \text{EPG}$ are reported.

these gates is applied in order to recover the initial state (usually $|0\rangle$). The longer the sequence, the larger the accumulated error, and the smaller the probability of recovering the initial state. For many common types of error, this probability can be fitted with a simple exponential decay

$$P_{|0\rangle} = A \cdot p^m + B, \quad (6.2)$$

where A , B are scaling parameters and p is known as the depolarizing parameter, with possible values $0 < p < 1$. The average error per applied Clifford (EPC) is then obtained as $\text{EPC} = (1 - p)/2$. In a real circuit implementation, this number will depend on how efficiently the Clifford operations are decomposed into actual native physical gates, such as X_π and $X_{\pi/2}$. A more meaningful measure to characterize a physical system might then be to instead quote the average error per applied physical gate (EPG), defined as $\text{EPG} = (1 - p^{1/N_g})/2$, where N_g is the average number of physical gates required per Clifford. We use the decomposition $\{I, X_\pi, X_{\pi/2}, Y_\pi, Y_{\pi/2}\}$, which results in $N_g = 45/24$. We note that a smaller N_g can be achieved by implementing virtual (i.e. zero-duration, software-based) Z gates [212] and therefore using fewer physical gates, but this only affects EPC and not the underlying EPG.

We perform RB individually on each qubit, sampling 100 random Clifford sequences at each sequence length m to ensure sufficient randomization. The

resulting measurements of RB decay are shown in Fig. 6.9 together with the exponential fit. The average gate fidelity is simply defined as $F_g = 1 - \text{EPG}$. We obtain 99.80(4) % for the transmon and 99.29(14) % for the IST. The IST has lower fidelity than the transmon, which is not surprising given its much shorter coherence time. A simple estimation of the error expected from decoherence [213] provides the upper bounds 99.9 % and 96.0 % for transmon and IST. The measured transmon fidelity is slightly lower than this estimate, suggesting the presence of coherent errors such as pulse miscalibration. On the other hand, the IST gate fidelity is much *higher* than the predicted coherence limit, which is an unphysical result. In fact, a fidelity on the order of 99 % or more seems implausible when the IST decoheres on average after only $T_2^*/t_g \approx 8$ gate applications. The solution lies in the fact that T_2^* only probes the "worst-case scenario" of bare decoherence when the qubit is idling on the equator of the Bloch sphere, while during the RB sequence the qubit is being driven most of the time. One can derive [213] that if slow 1/f-type noise is the dominant decoherence source, as in the case of the IST at finite flux bias, the RB sequence partially protects the qubit from this noise and extends its coherence under drive, leading to a much higher gate fidelity compared to the case of stochastic (white) noise. The corresponding formula $F_g = (2 + \exp(-(t_g/T_2^*)^2))/3 \approx 99.5$ % is indeed in good agreement with our experimental result, assuming a small contribution from coherent error as for the transmon.

One can also perform the RB protocol simultaneously on both qubits [214], which in the presence of qubit-qubit crosstalk will generally lead to larger errors than individual RB. Indeed, we obtain the gate fidelities 99.39(6) % and 99.07(11) % for transmon and IST, respectively, when the RB is performed simultaneously. This indicates that despite the active ZZ cancellation a finite contribution from crosstalk remains. This could either come from fluctuations of the optimal ZZ suppression point in time as mentioned previously, or from classical crosstalk between the qubits, i.e. the drive signals applied to the control lines can weakly drive neighboring qubits. It is possible to carry out calibration sequences to identify such sources of error [46].

As a side note, the gate fidelity of approximately 99.8 % measured for the transmon, while nowhere near the current state-of-the-art for superconducting circuits [46, 71], is an encouraging result from the point of view of our dFBL design. It confirms that our single control element is capable both of highly selective flux biasing, as characterized in Sec. 4.4, and also high-fidelity qubit operations. We

do not believe that the transmission properties of the dFBL are a limiting factor to fidelity at the current level. Future experiments with higher-coherence qubits and more careful pulse shaping will be required to perform a stricter test of dFBL performance.

6.5.3. Sideband pulse calibration

In principle, either of the red sidebands $|\widetilde{11}\rangle \leftrightarrow |\widetilde{02}\rangle$ and $|\widetilde{11}\rangle \leftrightarrow |\widetilde{20}\rangle$ can be used to implement a CPHASE(ϕ) gate, where ϕ is the phase accumulated on the $|\widetilde{11}\rangle$ state relative to the rest of the computational subspace after performing one full Rabi oscillation period. Here we pick $|\widetilde{11}\rangle \leftrightarrow |\widetilde{20}\rangle$, located around 460 MHz, since it is well-separated from the other sidebands and from the ZZ cancellation tone at 180 MHz.

We begin by identifying the optimal sideband drive parameters, such as the frequency ω_{sb} and drive plateau duration τ to implement a full 2π rotation in the $\{|\widetilde{11}\rangle, |\widetilde{20}\rangle\}$ subspace. The drive amplitude at room temperature is fixed at 0.5 V (1 V peak-to-peak), which is the maximum output allowed by the AWG channel. We predict that we could achieve a faster gate by removing some attenuation from the control line in the cryostat or using a DAC channel with higher output voltage. The pulse sequence and resulting measured chevron pattern are shown in Fig. 6.10 (a)-(b). We observe the oscillations most clearly in the IST, since the transmon never returns to its ground state during the $|\widetilde{11}\rangle \leftrightarrow |\widetilde{20}\rangle$ Rabi flop while the IST does, and our readout calibration is most sensitive to changes in the

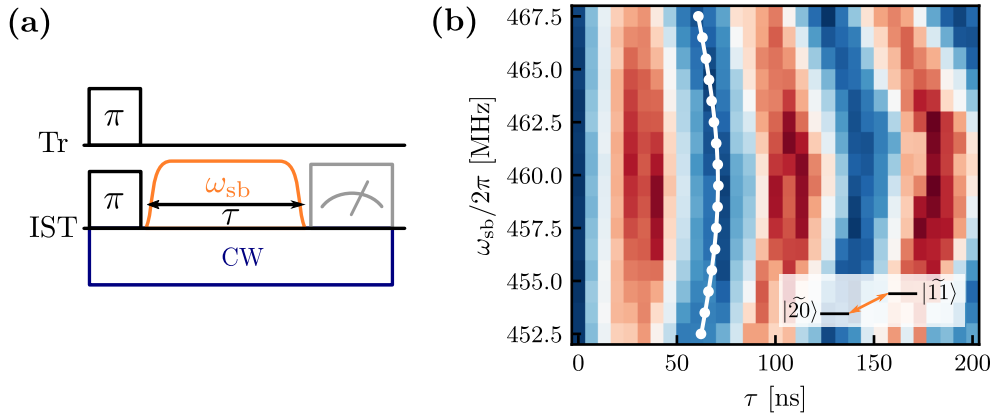


Figure 6.10: Calibration of red sideband transition $|\widetilde{11}\rangle \leftrightarrow |\widetilde{20}\rangle$. (a) Pulse sequence to implement Rabi driving of the $|\widetilde{11}\rangle \leftrightarrow |\widetilde{20}\rangle$ sideband. The system is initialized in the $|\widetilde{11}\rangle$ state before applying a sideband drive of variable duration and frequency. (b) Rabi oscillations observed in the IST signal. By fitting the Rabi rate at each frequency, we obtain the optimal drive time for a full oscillation period (white dots and line).

ground state qubit population $P_{|0\rangle}$. As done for the blue sideband data in Fig. 6.7, we fit the oscillations to the Rabi rate equation $\Omega^2 = \Omega_{11 \rightarrow 20}^2 + \Delta_{\text{sb}}^2$ and extract $\Omega_{11 \rightarrow 20} \approx 13$ MHz. The fit is added to Fig. 6.10 (b). This also gives us the optimal gate time duration versus drive frequency $t_g(\omega_{\text{sb}})$ that performs a full rotation and minimizes leakage into the $|\widetilde{20}\rangle$ state. For now we choose to drive the sideband on resonance, but this will be slightly modified later to achieve exactly $\phi = \pi$.

In addition to the desired conditional phase ϕ , the sideband drive causes small ac-Stark shifts on all system levels. This is the same mechanism that allowed us to cancel out the static ZZ interaction in Sec. 6.5.1, but it is unwanted here since it means that additional local phases will be accumulated on both qubits as they will be temporarily detuned from the lab frame of reference (defined by the frequency and phase of their microwave source). Luckily, this phase is easy to calibrate out using Ramsey interferometry as shown in Fig. 6.11 (a). The sideband pulse as determined above is inserted into a fixed-length Ramsey sequence, here shown for the IST. With the qubit initialized on the equator of the Bloch sphere, the phase imparted during the sideband will make the qubit rotate to a new point on the equator. We then add an intentional VZ(θ) rotation by an angle θ (implemented in software by rotating the lab frame for all subsequent pulses by $-\theta$) before projecting the qubit back and measuring it. This maps the total accumulated phase onto the qubit Z axis. By sweeping θ , with and without the sideband drive, we obtain the two curves plotted in Fig. 6.11 (b). The offset between the two corresponds to the phase added solely by the sideband, which is deterministic and can therefore be compensated in software by adding a VZ gate with the opposite angle to the definition of the CPHASE gate. Repeating this experiment should now yield two perfectly overlapping curves. The same procedure can then be applied to the other qubit to remove its own local accumulated phase. Note that the sideband and CW cancellation tones are always applied to the IST, regardless of which qubit the Ramsey interferometry is performed on.

Now that single-qubit phases have been calibrated out, we can observe the conditional phase induced by the sideband via a modified sequence [83] where the sideband drive is always present, and a conditional pulse is added to the control qubit as shown in Fig. 6.11 (c). When the control is in the ground state, nothing changes compared to before. If the control is excited, the target qubit now oscillates out of phase due to the accumulated phase $\phi \approx \pi$ during the sideband drive, see data in Fig. 6.11 (d). We note that in our experiment, the contrast of the oscillations is usually reduced by around 20-30 % when the control qubit activates

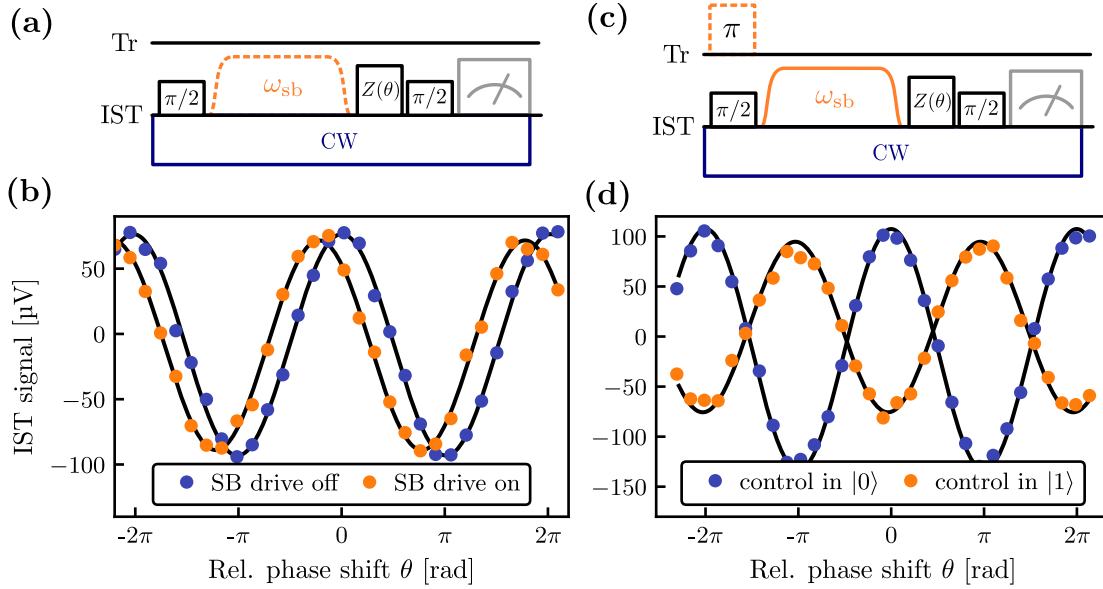


Figure 6.11: Calibration of single and conditional phases of the CPHASE sideband gate. Ramsey interferometry experiments to detect phases added by the presence of the CPHASE sideband drive. (a)-(b) Single-qubit phases can be detected by sweeping the phase θ of a virtual Z gate prior to projecting the qubit back onto its measurement axis, with or without the sideband drive present. The phase offset is equal to the unwanted single-qubit rotation added by the sideband drive, which can now be calibrated out. (c)-(d) The conditional phase ϕ can be detected by repeating the Ramsey interferometry as a function of state of the control qubit when driving the sideband. A phase difference close to π is observed here. Note that in all these experiments, a CW tone for ZZ cancellation is always applied to the IST.

the conditional phase. We do not currently understand the origin of this effect, but we also do not think that this is a fundamental issue with the sideband drive, as the chevron pattern in Fig. 6.10 does not show such an obvious loss of contrast even after two full oscillation periods. This rules out loss of contrast due to leakage into $|\tilde{2}0\rangle$. One possibility is that instead dephasing is enhanced while the higher state $|\tilde{2}0\rangle$ is populated during the rotation, which could cause the Ramsey interferometry signal to weaken. It could also be related to readout, due to unwanted coupling of one qubit to the wrong readout resonator [215], but more careful characterization will have to be performed in order to confirm this. In any case we find that the sideband gate still performs reasonably well in randomized benchmarking, as we will see in Sec. 6.5.4.

The phase offset between the two curves in Fig. 6.11 (d) is exactly the conditional phase ϕ . We find that it is not exactly equal to π for our initial guess of drive frequency, but it can be tuned freely by moving along the optimal one-dimensional manifold in Fig. 6.10 (b). This is analogous to the flux-based parametric CZ gate

in transmon qubits [200]. To confirm this, we repeat the conditional interferometry of Fig. 6.11 (c)–(d) to extract ϕ while varying the sideband drive frequency ω_{sb} and following the optimal $t_g(\omega_{\text{sb}})$ curve. Further, we repeat this by swapping the role of control and target qubit as a consistency check of the method. The result is plotted in Fig. 6.12, showing that indeed the conditional phase can be tuned continuously and that it is consistent when swapping the roles of the qubits, as the two lines have opposite slope and intersect exactly at a value of $\phi = \pi$. We therefore choose this operation point, which is slightly detuned from the true transition frequency, and then recalibrate the single qubit Z corrections as before.

The procedure outlined in this section provides us with a calibrated $\text{CPHASE}(\pi) = \text{CZ}$ with a total length $t_g = 83$ ns including rise and fall edges and a 5 ns buffer inserted after the pulse. Overall, we find that the sideband CZ gate is relatively easy to tune up and its behavior in parameter space is fairly predictable, as demonstrated by the tuning procedure outlined here. Although not further investigated in this study, the access to a continuous set of $\text{CPHASE}(\phi)$ gates can be a useful resource to compile circuits more efficiently than with only one discrete two-qubit unitary [216].




Figure 6.12: Varying the conditional phase of the sideband gate. The conditional phase ϕ induced by the $\text{CPHASE}(\phi)$ gate is measured by Ramsey interferometry as a function of the sideband drive frequency ω_{sb} , adjusting the optimal gate time as in the fit from Fig. 6.10. The experiment is carried out twice, with each of the qubits acting as control. Error bars on ϕ are present but are smaller than the data points. Linear fits to the two datasets give an optimal drive frequency $\omega_{\text{sb}} = 463.12(5)$ MHz that implements a $\text{CPHASE}(\pi) = \text{CZ}$ gate in this particular calibration run.

6.5.4. CZ gate benchmarking

The CZ gate fidelity is determined using RB, as for the single-qubit gates. The extension of RB to the two-qubit (2Q) case is straightforward and the principle remains the same [217, 218]. The 2Q RB protocol attempts to restore the starting state $|\widetilde{00}\rangle$ after applying a random sequence of Cliffords. The main difference is that the full 2Q Clifford group with 11'520 elements is now significantly larger than in the single-qubit case, meaning that the group has to be sampled more times in order to obtain a reliable estimate of the error. In general, given a native gate such as the CZ implemented here, a general 2Q Clifford will consist of several CZ gates and single-qubit operations. The EPC will therefore include a mixture of single- and two-qubit errors. To obtain the error due to the actual CZ gate, one can make use of a refined protocol called *interleaved* RB [219]. Here, a copy of the particular gate of interest G is inserted between each random Clifford of a standard (reference) RB sequence. The additional error accumulated in interleaved RB compared to reference RB is purely due to G . The individual gate error EPG_G is then obtained from the depolarizing parameters p and p'_G of the reference and interleaved RB, respectively, as $\text{EPG}_G = 3(1 - p'_G/p)/4$.

We perform reference and interleaved 2Q RB for the CZ gate on our device by randomly sampling the Clifford group 150 times at each sequence length. While this number is too small to sufficiently represent the group, this experiment already takes ~ 12 hours to complete. Much longer sequences would likely require intermediate recalibration of pulse parameters due to slow drifts in the experimental setup. We are limited by: the loading of the long required digital waveforms on our AWG, by the data acquisition time (each unique sequence is averaged over 5'000 repetitions to achieve a sufficient signal-to-noise ratio), and by the time required for the computation of the final inverse 2Q Clifford. Improving the readout setup to achieve higher signal-to-noise ratio, as well as pre-computing the 2Q Clifford inverse for each sequence and storing it as a list, would improve the efficiency of this benchmark. Note that this is not a fundamental issue, as Clifford gates can be efficiently computed classically [1], but simply a limitation of our particular hardware and software used.

As a positive side effect, this long experiment shows us that our system is relatively stable over long periods of time. No intermediate pulse recalibrations were performed during this 12 h period.

The measured data for both standard and interleaved RB is plotted in Fig. 6.13 along with the RB fit to each sequence. The CZ gate fidelity $F_{\text{CZ}} = 1 - \text{EPG}_{\text{CZ}}$ is

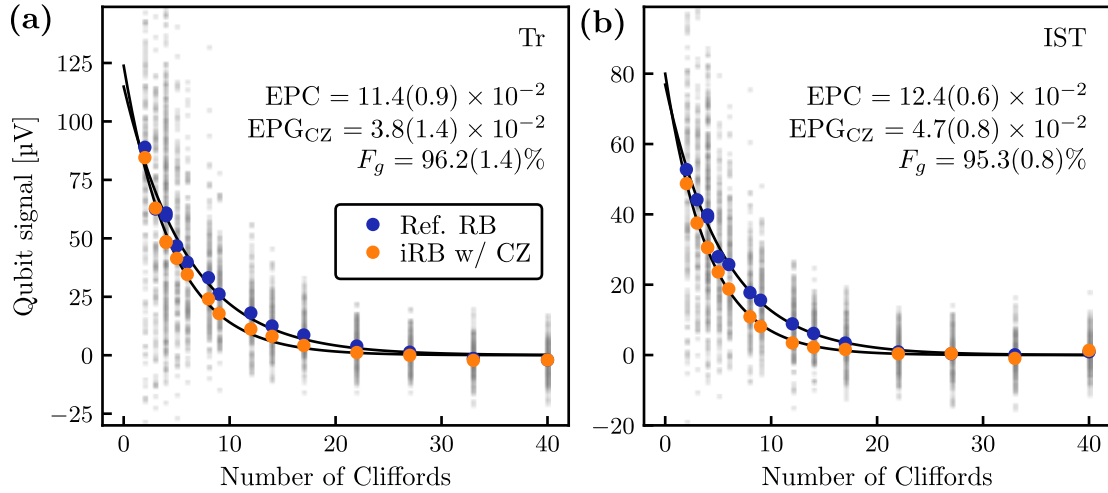


Figure 6.13: Sideband CZ gate randomized benchmarking. Standard (blue) and interleaved (orange) two-qubit randomized benchmarking experiment, showing simultaneously recorded data on the (a) transmon and (b) IST. The raw data points (grey) are shown for the reference RB sequence only. The fitted decay curves (black lines) provide an estimate for the average error per Clifford (EPC) and the error per physical CZ gate (EPG_{CZ}). The quoted gate fidelity is $F_{\text{CZ}} = 1 - \text{EPG}_{\text{CZ}}$.

obtained from the combined standard and interleaved RB data as explained above. We separately plot and fit the data for each qubit, recorded simultaneously. One would expect both qubits to provide the same error estimate [89]. We find that the extracted fidelity F_{CZ} is 96.2(1.4) % for the transmon and 95.3(0.8) % for the IST. The estimated fidelities are slightly different but are still within fit uncertainty of each other. Looking at the raw data for reference RB in Fig. 6.13 we can see that especially at small Clifford numbers the spread of recorded values is extremely large. We attribute this to the effect of finite sampling, which is the likely cause for the large uncertainties on F_g , although we note that longer experiments with more sampling would be required to confirm whether this is the case or whether there is a more fundamental issue in our RB implementation. For practical purposes, here we choose to take the average of the fidelities obtained from the two qubits as our final result, which is $F_{\text{CZ}} = 95.8(1.6)$ %.

We briefly discuss the limitations of gate fidelity. While coherent errors due to e.g. residual ZZ crosstalk or miscalibration of the single-qubit and conditional phases as discussed in Sec. 6.5.3 can play a role, we expect the gate error to be mostly coherence-limited based on the very low coherence of the IST at finite flux bias. We are not aware of any theoretical work that analyses two-qubit gate error limits in the presence of $1/f$ flux noise, as was discussed for the single-qubit case [213] (see also details in Sec. 6.5.2). This is further complicated by the fact that the transmon

and IST are subjected to different noise spectra in our experiment, creating an asymmetry in the dynamics. The authors of Ref. 100 observed experimentally that the fidelity of a CR gate on a transmon-CSFQ device is higher than expected from a standard coherence limit, which they also attribute to the flux noise limited coherence of their CSFQ.

Here we offer an intuitive picture for the case of our CZ gate. During gate operation, the sideband transition is only being driven when the two-qubit system has a non-zero probability of being in the $|\widetilde{11}\rangle$ eigenstate, which on average is 1/4 for a sufficiently randomized state. All remaining eigenstates are idling during this time, accumulating only single-qubit phases, making the IST more sensitive to the bare 1/f noise than during single-qubit benchmarking, which mostly consists of actively driving the qubit. We therefore expect the 2Q RB to be less forgiving in this regard, due to our specific choice of native two-qubit gate. Assuming all error to come from IST decoherence, the naive coherence limit for our $t_g = 83$ ns gate is 93.6 %, while the single-qubit 1/f noise scenario gives 98.5 %. Our actual gate fidelity sits between these two limiting cases, supporting the intuition developed here. As an additional comparison, the original sideband CZ gate implementation in Ref. [83] achieves a slightly higher fidelity of 97 % with a slightly higher CSFQ coherence of $T_R^* \approx 0.6$ μ s and shorter gate time $t_g = 50$ ns, suggesting that both these gate implementations are limited by the same mechanism.

Ultimately, the short coherence makes this gate a relatively unattractive choice unless gate times can be decreased significantly with larger J or flux noise in the system can be mitigated. One solution to partially reduce the effect of decoherence would be to keep the IST at its sweet spot when idling or when performing single qubit operations, and quickly tune it to the sideband operation frequency before applying the CZ drive. This is analogous to conventional flux-based gates with fast baseband pulses [49], which should be possible in our architecture thanks to the large bandwidth provided by the dFBL. While this does not remove the issue of low coherence during CZ operation, it improves single-qubit fidelities and the time-averaged coherence time of the IST. Additionally, one could also split the sideband drive into two halves with an interleaved echo pulse, such as for the echoed version of the cross-resonance gate [217], to additionally remove slow noise contributions during gate operation.

6.6. Summary

The study of a coupled device consisting of a transmon and IST presented in this chapter has led to several interesting results. Firstly, the device is found to operate as expected, with stable parameters and acceptable coherence metrics. The static ZZ interaction between the two qubits can be suppressed by choosing the optimal flux working point of the IST such that the opposite anharmonicities of the two qubits cancel out unwanted dispersive shifts in the computational subspace. Crucially, a numerical simulation of the full system can closely reproduce the behavior of ZZ interaction across the entire IST flux range without any fitting parameters. This leads us to conclude that the properties of future optimized devices could be reliably predicted using this numerical model, increasing the chances of achieving sufficient ZZ suppression. While challenging, the reliable targeting of circuit parameters in practice should become feasible in the near future, thanks to techniques such as laser annealing of Josephson junctions.

Although not explored in this device, we believe that conventional two-qubit gates such as the cross-resonance gate and parametric flux gates should be possible in this system. This opens up the possibility of both suppressing ZZ contributions and performing fast high-fidelity gates in a simple statically-coupled device, without the need for additional coupling modes and related control circuitry.

Further, single-photon sideband transitions enabled by odd terms in the IST Hamiltonian when biased away from the flux sweet spot have been explored, showing a variety of two-qubit interactions that can be used to implement two-qubit entangling gates. This is a resource that is not available in conventional all-transmon systems. As an example, we carefully tune up and benchmark single-qubit gates as well as a two-qubit CZ gate based on a red sideband transition. To our knowledge, this is the first experimental implementation of a two-qubit gate using a geometrically inductively shunted qubit, and demonstrates the feasibility of incorporating this qubit species in existing transmon-based architectures at virtually no cost.

7.

Conclusion and Outlook

We began this thesis by offering an overview of simple superconducting circuits often used for quantum information processing, starting from the description of the basic building block, the Josephson junction. We described one particular circuit in more detail, the RF SQUID, which is tunable in-situ via the application of a magnetic flux through its closed loop. The three main energy scales of this circuit are the charging energy E_C , the Josephson energy E_J , and the inductive energy E_L . Different choices of parameters can lead to vastly different physics. One example is the IST, a weakly anharmonic oscillator which bears striking resemblance to the popular transmon qubit despite a fundamentally different topology. The IST lies in the easily realizable $E_L > E_J$ regime, and is therefore a simple alternative to other more involved designs.

We then discussed the practical challenges faced when implementing flux control in superconducting circuits, such as the additional hardware requirements and the issue of 3D integration and scalability. Performance in flux-tunable circuits can be impacted by factors such as resistive heating, flux crosstalk, qubit dephasing induced by flux noise, or qubit Purcell decay through flux bias lines (FBL).

With these issues in mind, we set out to develop a novel 3D-integrated architecture for flux control of superconducting circuits, based on an extension of our existing coaxial circuit QED design developed previously in our research group. The main ingredient of this design is the differential FBL (dFBL), a control element fabricated out of coaxial wire to form an inductive loop and embedded inside a superconducting sample enclosure. This approach removes the FBL from the qubit substrate and does not require any sort of wirebonding, bump-bonding or other involved packaging technique. This allows us to retain the simplicity and tileability of the original coaxial architecture, ensuring that circuits can be increased in size without intrinsic constraints. We tested the dFBL design at the four-qubit level via cryogenic measurement of four tunable gradiometric coaxial transmons. We can control both the flux (Z) and charge (XY) degree of freedom of the transmons via the dFBL, proving that it is capable of full XYZ control. This is in contrast to many conventional architectures that utilize two independent control lines per qubit,

and is therefore an important step towards the reduction of hardware overheads. Moreover, the dFBLs showed very high flux bias selectivity, with DC flux crosstalk on the order of 10^{-3} across the entire enclosure. We attributed this to the use of gradiometric SQUID loops and the optimal matching of the dFBL geometry and orientation with its corresponding SQUID. The flux crosstalk behavior is easily reproduced via finite-element magnetostatic simulations, and we expect this to scale favorably as circuit size increases.

To carefully characterize our architecture, we used our tunable transmons to measure the transmission through the dFBLs and the readout ports. We also repeated the measurements using a modified enclosure, which includes a 3D-integrated readout multiplexer for simultaneous readout of all four qubits via a single microwave line. By comparing this data to qubit T_1 as a function of frequency, we found strong indications that our qubit lifetimes are currently limited by radiative Purcell decay through the readout port, and not by the dFBLs.

The realization of the dFBL platform then enabled us to proceed to the study of novel flux-tunable circuits. We focused on the design, simulation and realization of gradiometric IST qubits to investigate their relatively unexplored physics. We showed how these qubits can be easily integrated into our existing transmon-based architecture without requiring any modifications of the experimental setup or the control and measurement scheme. The IST qubits were found to have a simple flux spectrum and to behave in accordance with numerical predictions from the full Hamiltonian. In particular, the IST anharmonicity is positive at the half flux bias point and is easily tuned through zero to become negative. The overall energy relaxation and coherence properties of our ISTs are on par with our tunable transmons fabricated using the exact same process, showing that the IST is a viable candidate for further interesting experiments in circuit QED.

The first step in this direction was taken by fabricating and characterizing a device consisting of a transmon coupled capacitively to an IST. This allowed us to study the two-qubit dynamics of a system where the opposite signs and comparable magnitudes of the two qubit anharmonicities can lead to exact (within measurement uncertainty) suppression of the unwanted longitudinal ZZ coupling. This behavior could be accurately reproduced numerically across many orders of magnitude in ZZ coupling strength, highlighting the simplicity of this two-qubit Hamiltonian. The ZZ term is commonly found to be a limiting factor in circuit QED experiments, and several solutions involving either additional hardware or software requirements have been developed. Here we showed that this is possible in a minimalistic two-

qubit system without the need for external elements, confirming a similar study that had been previously carried out using a transmon coupled to a C-shunted flux qubit [100]. Further, using the parity-breaking of the IST Hamiltonian at finite flux bias, we activated single-photon red and blue sideband transitions which are not possible in an all-transmon system. These can be employed to perform entangling operations, and as an example we tuned up a controlled-Z (CZ) gate and benchmarked it via randomized benchmarking, achieving a CZ gate fidelity of 96 %. To our knowledge, this is the first demonstration of a two-qubit gate involving a qubit species with a geometric inductive shunt.

Outlook

This work realized a proof of principle of fully 3D-integrated flux-tunable superconducting circuits, compatible with multi-qubit operations while also providing low crosstalk and hardware-efficient XYZ qubit control. Going forward, we would like to further scale up the dFBL architecture and show how its useful properties translate to larger circuits, e.g. on the scale of 16 qubits, 64 qubits, and beyond. The combination of dFBLs with a larger unit cell pitch of 2 mm and the inductive shunting of the enclosure to suppress spurious modes, demonstrated previously [46, 168], will maintain the full tileability of our approach. The integration of readout multiplexing, already demonstrated at the 4-to-1 level, will be further pursued in separate work. The main remaining challenges will then be found outside of the sample enclosure, for example due to the limited physical space and cooling capacity of cryogenic setups [137] and the sheer amount of classical control electronics required for large number of qubits. We note that these issues are not unique to our experiments – the research field as a whole is trying to overcome these obstacles, and we are hopeful that significant progress will be made in the coming years through collective effort.

With the simplicity and flexibility provided by the dFBL design, it is then possible to focus on the perhaps more interesting aspect of Hamiltonian design and the exploration of large-scale entangled quantum systems. After the successful demonstration of intrinsic ZZ suppression in a single transmon-IST pair, it is natural to wonder how this will translate to larger lattices of multi-species systems, such as the one sketched in Fig. 7.1. Simultaneously hitting the optimal ZZ operation point between many pairs of qubits will require stringent fabrication requirements, similar to the well-known problem of avoiding frequency collisions in fixed-frequency transmon circuits [204]. Spectator effects on neighboring qubits might limit gate




Figure 7.1: Extension of coaxial transmon-IST circuits. Conceptual drawing of a large multi-species 2D lattice consisting of alternating tunable transmons (blue) and ISTs (orange). Individual full XYZ control is provided through one dFBL per qubit, and nearest-neighbor coupling is achieved via simple capacitive arms.

fidelities even if every nearest-neighbor ZZ coupling is itself suppressed [111]. It is possible that robust ZZ crosstalk suppression might only be realizable in large circuits by combining multiple techniques, as discussed in [Sec. 2.4.3](#).

While the most hardware-efficient solution is to treat all transmons and ISTs in the lattice as physical data qubits, with only static capacitive couplings between them, one could also easily convert this into a tunable coupling architecture [119, 220] by redefining every second qubit as a coupler. This could even be done in-situ using the same device, to test different topologies and configurations that might deliver the best trade-off between qubit count and overall levels of crosstalk and gate fidelities.

The IST qubit itself offers further avenues of study. Realizing the highly anharmonic quarton regime (see [Sec. 2.3.3](#)) is fully within the scope of the present IST design and will provide modest but immediate benefits compared to conventional transmons in terms of addressability of the qubit transition and single-qubit gate speeds. On a more fundamental level, the highly confined phase potential of the IST Hamiltonian could be useful for experimentally investigating the physics of strongly driven systems, comparing it to the chaotic behavior observed in transmons [77, 221]. Meanwhile, the tunable nature of the IST anharmonicity could open up the possibility of simulating different types of 2D lattices in-situ by going from positive to zero to negative anharmonicity. This implements different regimes of the two-dimensional Bose-Hubbard model, where the anharmonicity plays the role of an on-site interaction [23, 136, 222]. All of the possibilities mentioned

here provide abundant opportunities to explore the dynamics of intermediate-scale quantum systems, be it for the purpose of gate-based quantum computation or for the simulation of fundamental states of matter, hopefully leading to new insights in the coming years.

A.

Fabrication

A.1. Process overview

In this section we provide a brief overview of the nanofabrication process employed to create the superconducting circuits measured over the course of this work. This process has been developed and refined over several years within our research group in Oxford, thanks to the contribution of many current and former group members. The original coaxmon qubit [161] was implemented on a sapphire substrate with qubits and resonators made entirely out of aluminium, emulating the materials of a typical 3D transmon [44]. More recently, the fabrication was moved to aluminium circuits on silicon to take advantage of the ease of processing that this substrate allows. A fairly recent and up-to-date version of our process, including Josephson junction design and statistics, is described with great detail in the doctoral thesis of Peter Spring [168]. The aim of the present section is therefore not to repeat the whole process in detail, but rather to give a general overview and to point out any relevant changes between the previously published version and the one used here. In addition, in [Sec. A.2](#) we provide for reference the step-by-step fabrication recipe used for the devices in this work.

Fabrication is carried out on double-side polished, 3-inch intrinsic silicon wafers with a thickness of around 500 μm . Since the coaxmon architecture involves the patterning of both sides of the same substrate (see [Fig. 4.1](#)), care has to be taken not to damage or contaminate the back side of the wafer while the other side is being processed. This requires a protective resist layer on the back side and the use of special tools and holders that avoid contact with any potentially contaminated surface, e.g. glassware, spincoater vacuum chuck, EBL wafer carrier. The process starts with a bare silicon wafer undergoing a standard buffered oxide etch (BOE) to remove silicon oxide, after which it is loaded as quickly as possible (less than 10 minutes) into the load-lock of a dedicated Plassys MEB550 S2 ultra-high vacuum (UHV) electron-beam evaporation tool, to limit the oxide regrowth. After depositing a single 100 nm layer of aluminium, the wafer is unloaded, flipped, and quickly loaded again for metallization of the second side. Subsequently, a single step of

standard photolithography and wet etching is carried out sequentially on each side to fully define the spiral resonators and the transmon electrodes, respectively. We use an EVG 620 mask aligner system that crucially allows us to align features on both sides of the substrate to each other, even for optically opaque materials such as silicon. This ensures good lateral alignment of the resonators and transmons. The last major step involves the fabrication of Josephson junctions on the qubit side of the wafer, performed with a single EBL step and double-angle aluminium evaporation using the Dolan bridge technique¹ [223]. An in-situ argon ion mill is performed inside the evaporation tool immediately prior to metal deposition in order to remove aluminium oxide and ensure good electrical contact between the junction metal leads and the large transmon electrodes already present on the substrate. After metal lift-off and resist strip, the wafer is then diced into $5 \times 5 \text{ mm}^2$ dies, which can now be packaged into our sample enclosure (presented in Sec. 4.3) and characterized cryogenically.

Updates to the fabrication process

In the following, we describe the main modifications that have been made to our previously described process [168] in the context of the fabrication carried out for this work:

- A pre-bake of the wafer for 5 minutes at $120 \text{ }^\circ\text{C}$ to remove water is introduced before spincoating the photoresist during photolithography steps.
- Photoresist exposure dose is changed from 20 mJ/cm^2 to 40 mJ/cm^2 to lie closer to the typical range for AZ 1514 H photoresist. Development time is then shortened from 80 s to 40 s to avoid over-development. We split the development time between two separate beakers of AZ 726 MIF developer, 30 s in the first and 10 s in the second, to promote a constant development rate and reduce the amount of dissolved resist being carried into the rinse steps.
- The temperature of both the developer bath and the etchant bath is changed from $23 \text{ }^\circ\text{C}$ to $25 \text{ }^\circ\text{C}$ for simplicity.
- Three (previously two) steps of DI water rinse are carried out after both the etching bath and the resist strip in DMSO, to ensure sufficient removal of

¹We refer the reader back to the previously mentioned thesis of Peter Spring [168] for a detailed discussion of our Josephson junction design, which is less sensitive to edge effects and resist thickness variation than the traditional Dolan junction.

solvent traces and contaminations. Due to the viscosity of DMSO at room temperature, it is recommended to warm up the first two DI water rinse beakers to aid with solvent removal.

- EBL is carried out with a 30 keV electron beam for both small and large design features, instead of switching to 20 keV for larger features. This reduces the size of the undercut profile and makes the developed PMMA mask more stable.
- Development of EBL resist is kept the same, but the post-development rinse is simplified to a single 30 s rinse in IPA before gentle N₂ blow-dry.
- The plasma ashing step (often referred to as *de-scum*) after development and before metal deposition is omitted. This is because the in-situ ion milling step, in addition to removing metal oxide, is also likely to remove organic residues from the exposed surfaces, rendering a dedicated de-scum step unnecessary. Omitting this step reduces the broadening of PMMA mask features, which should have a positive impact on reproducibility.
- Lift-off and EBL resist strip after junction deposition is carried out in DMSO at 80 °C instead of acetone. DMSO is a much safer solvent (not flammable; higher boiling point) and does not lead to the formation of visible organic residue often seen after lift-off in acetone.

A.2. Detailed fabrication recipe

We divide the entire fabrication process into four main sections: initial wafer clean and metallization, photolithography and etching (double-sided), EBL (one-sided), and final dicing and clean.

Initial wafer cleaning and metallization

Begin with a fresh 3-inch silicon wafer.

- Remove native silicon oxide in BOE (volume ratio 10:1 of NH₄F to HF) for 2 minutes, followed by two DI water rinse steps of 30 s each. Blow-dry with N₂.
- Quickly transfer to vacuum chamber of evaporation tool.
- Bake wafer in UHV chamber above 100 °C for 6 hours. Let cool down to room temperature.

- Deposit 100 nm of aluminium at rate 1 nm/s.
- Unload wafer from chamber, flip it, load again. deposit 100 nm of aluminium on the second side with the same procedure.

The wafer is now fully metallized and ready for lithography. The side that was metallized first is defined as the qubit side, the other as the resonator side.

Photolithography and etching

Positive photoresist AZ 1514 H is used throughout the process. Standard spincoating parameters: 10 s spin at 500 RPM, followed by 50 s spin at 4000 RPM.

- Pre-bake wafer on hot plate for 5 min at 120 °C. Cool down for 5 min.
- Spincoat qubit side using non-contact spincoater chuck to prevent contamination of backside. Soft-bake for 50 s at 100 °C. Cool down for 5 min.
- Spincoat resonator side using standard chuck. Soft-bake for 50 s at 100 °C. Cool down for 5 min.
- Expose resonator side using chrome/glass photomask. Dose 40 mJ/cm².
- Develop resist using two beakers of AZ 726 MIF developer at 25 °C, for 30 s in one and then 10 s in the other. Rinse in two beakers of DI water for 30 s each. Gentle N₂ blow-dry.
- Etch aluminium in etchant bath (Alfa Aesar 44581 aluminium etchant) at 25 °C. Total etch time is highly temperature dependent, usually 2–3 minutes for 100 nm. Etch until substrate is cleared, then wait an additional 30 s before removing. Rinse in three consecutive beakers of DI water, for 10 s, 30 s, and 30 s respectively. N₂ blow-dry.
- Strip resist in DMSO at 80 °C for 30 min. Rinse in three beakers of DI water, for 30 s each. Heat the first two DI water beakers to facilitate DMSO removal from the substrate. Then N₂ blow-dry.

Now, repeat the above steps by swapping the roles of qubit and resonator sides. Protect the already patterned resonator side with a layer of resist, then process the qubit side. The only difference is that before exposure of the resist, the new pattern needs to be aligned to the existing pattern using the bottom optics of the EVG 620 mask aligner.

EBL and junction fabrication

We use a standard bilayer EBL resist stack consisting of MA/MMA copolymer (Allresist AR-P 617.08) and PMMA 950K (Allresist AR-P 672.045). Both are spincoated as before, 10 s at 500 RPM followed by 50 s at 4000 RPM.

- Pre-bake wafer on hot plate for 5 min at 120 °C. Cool down for 5 min.
- Spincoat a protection layer of copolymer resist on resonator side using non-vacuum spincoater chuck. Bake at 180 °C for 5 min, cool down for 5 min.
- Spincoat copolymer resist on qubit side using standard chuck. Bake at 180 °C for 5 min, cool down for 5 min.
- Spincoat PMMA resist on qubit side using standard chuck. Bake at 180 °C for 5 min, cool down for 5 min.
- EBL is carried out in a Raith EBL system at 30 keV, using aperture size 30 μm^2 for small features and 120 μm^2 for large features. Doses are typically between 500-700 $\mu\text{C}/\text{cm}^2$.
- Develop EBL resist in mixture of MIBK and IPA (volume ratio 1:3) at 25 °C for 45 s. Stop development in a beaker of IPA for 30 s. Gentle N₂ blow-dry.
- Load wafer into Plassys MEB550S2 evaporation tool. Perform Ar ion milling for 1 minute (voltage 400 V, acc. voltage 90 V, current 15 mA). Deposit 60 nm of Al at angle of 60° from normal, rate 0.5 nm/s. Perform static oxidation, typically for 5-10 min at 5-10 mbar, depending on target junction resistance. Pump back down to UHV, deposit 70 nm of Al at angle of 0°, rate 0.5 nm/s.
- Lift-off and resist strip in DMSO at 80 °C, typically for 1 hour or more. Gently blow with disposable pipette to aid lift-off mechanically. After lift-off is completed, rinse in a clean beaker of DMSO at 80 °C, then in three beakers of DI water as described for the photolithography above.

Dicing and cleaning

- Spincoat protective layer of AZ 1514 H photoresist on both sides of wafer, using same spincoating and soft-baking parameters as for photolithography above. Pre-bake of wafer before spincoating is not necessary.

- Dice the entire silicon wafer into $5 \times 5 \text{ mm}^2$ dies. For the devices fabricated in this work, the dicing was carried out by collaborators at the University of Southampton, UK.
- The individual dies are cleaned in DMSO at $80 \text{ }^\circ\text{C}$ for 30 min, and then rinsed in three DI water beakers. The procedure is the same as in the photolithography above.

The individual dies are then inspected optically for cleanliness and defects, and the Josephson junctions are probed via a DC resistance measurement at room temperature.

B.

Experimental setup

This chapter provides details on the experimental setup and components used for the cryogenic measurements of superconducting circuits presented in this work. The overall wiring diagram for our cryogenic setup is shown in [Fig. B.1](#).

Our device is cooled down inside a dilution cryostat (Oxford Instruments Triton 200) to a temperature of around 20 mK. The quantum device, packaged inside the superconducting enclosure with differential flux bias lines (dFBLs) presented in [Chapter 4](#), is housed within a high-permeability magnetic shield with an indium-sealed lid made out of oxygen-free copper. The individual SMA feed-through connectors on the lid are also indium-sealed. This provides protection from stray radiation present inside the cryostat, coming from higher-temperature surfaces. This radiation shielding solution was designed and realized in our group by Brian Vlastakis and James Wills. Semi-rigid coaxial cables are used outside of this shield to minimize leakage of radiation into the coaxial lines, while hand-formable braided coaxial cables are used inside the shield for convenience.

Control and readout signals are generated at room temperature and sent to the device via attenuated microwave lines inside the cryostat. All electronic instruments are synchronized by a common 10 MHz reference clock. We perform readout of our resonators in a reflection configuration, applying a signal to a readout input port (R_{in}) which then bounces off the coaxial readout port inside the device enclosure, and extracting the reflected signal via circulators and an amplified output line (R_{out}). For continuous-wave spectroscopic measurements, we simply connect ports 1 and 2 of a 4-port virtual network analyzer (VNA; Rohde & Schwarz ZNB 20) to R_{in} and R_{out} respectively. A qubit spectroscopy tone is then generated via an additional port on the VNA and sent to the qubit via its dFBL ($QDrive$). For time-domain experiments, we create pulsed waveforms via up-conversion of intermediate frequency (IF) signals from arbitrary waveform generators (AWG; we use one Tektronix AWG5014C and one Tabor WX2184C). The local oscillator (LO) tones are provided by Rohde & Schwarz SGS100A sources. For qubit control up-conversion we use the built-in IQ mixers of the SGS100A. In the case of the readout signal chain, mixing is instead performed using external IQ mixers. The

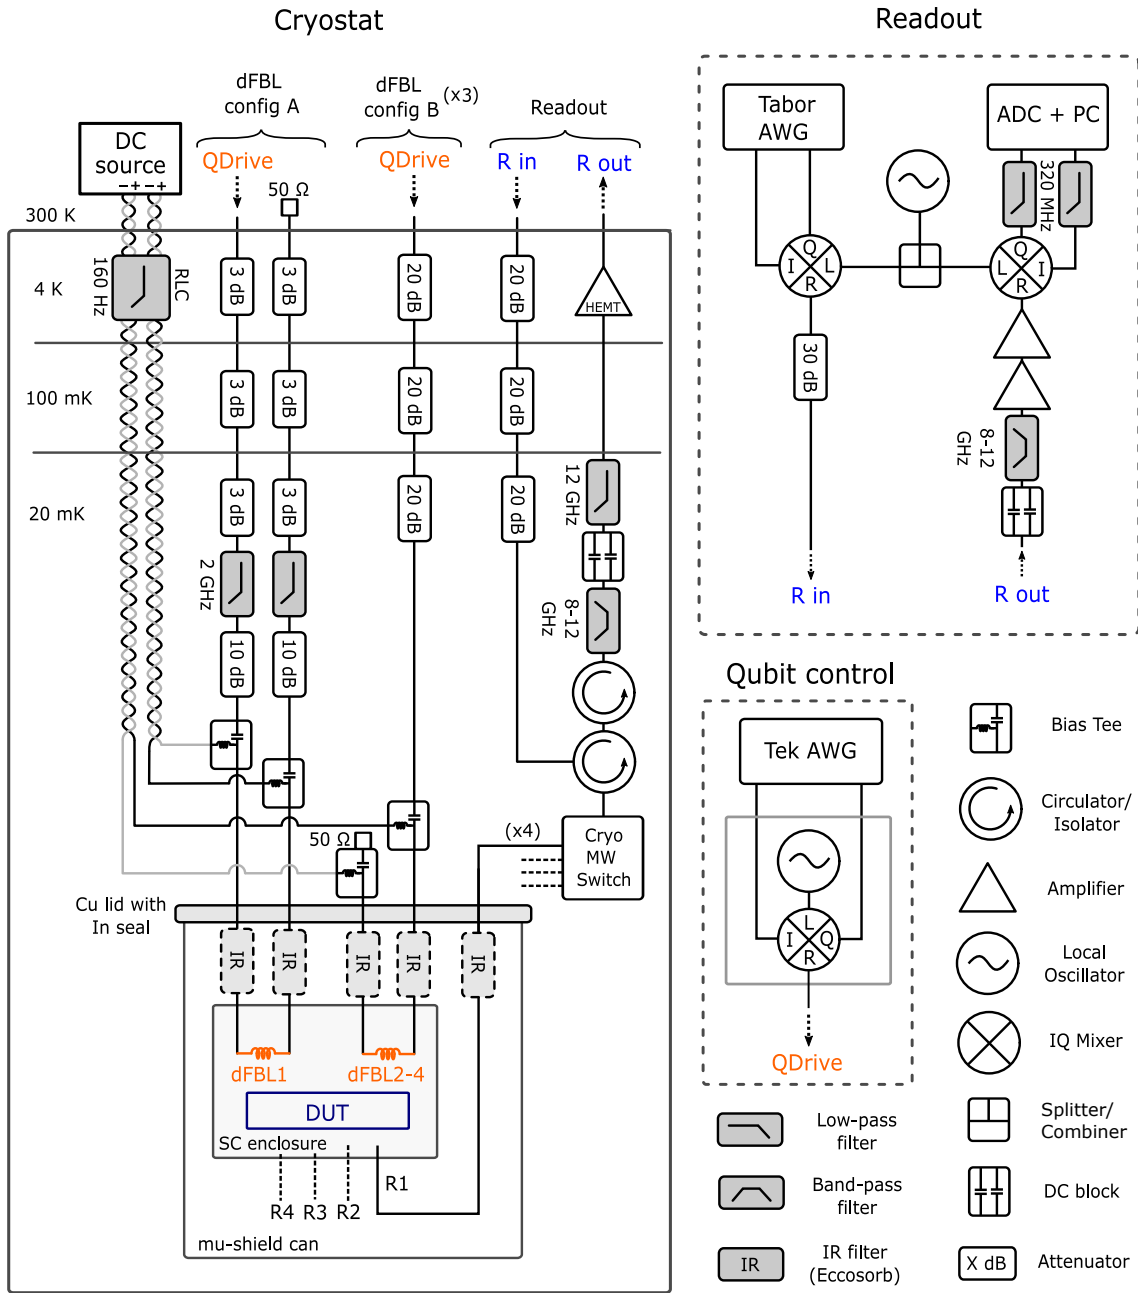


Figure B.1: Experimental setup – Standard configuration. Simplified wiring diagram of the cryogenic experimental setup used in this work. The superconducting qubit device under test (DUT) is located at the lowest temperature stage of an Oxford Instrument Triton 200 dilution cryostat. Shown here is the wiring for individual, non-multiplexed qubit readout using a cryogenic switch to address each readout resonator. Note the two wiring configurations for the differential flux bias lines (dFBLs) employed in this work. Config A is used for dFBL 1, while config B is used for the remaining three dFBLs 2-4 (only one set is shown here for clarity). Infrared (Eccosorb) filters were added inside the magnetic shield on some of the microwave lines during a subset of the cooldowns, see main text for details.

readout output signal carrying information about the experiment is down-converted to IF again, before being recorded by a digitizer (AlazarTech ATS9360) connected to a PC. The LO for both up- and down-conversion of readout signals is provided by the same source via a power splitter.

Wiring up the dFBLs inside the cryostat deserves careful consideration. We choose to keep DC and AC components of the input signal separated for as long as possible, to allow dedicated filtering to be applied to each, and combine them at base temperature using cryogenic bias tees (Anritsu K251). This is different from other common wiring schemes [56, 137] where the bias tee is located outside the cryostat and DC signals therefore travel through attenuated microwave lines, causing resistive heating. Each of our dFBLs is supplied with a DC bias from an isolated, battery-powered voltage source¹ (Stanford Research Systems SIM928), applied via a pair of twisted wires of a superconducting loom inside the cryostat. These signals are low-pass filtered at the 4K stage to reduce high-frequency noise – in addition, the capacitance between the wires of each twisted pair offers a natural low-pass filtering behavior as well.

For the RF component of the input signal, we implement two different microwave line configurations. Configuration A (used for dFBL1) combines low total attenuation (~ 20 dB) with a reflective low-pass filter (LPF) to shield the qubits from thermal radiation at their frequency coming from the higher-temperature stages. While not specifically used in this work, this configuration has the advantage that low-frequency signals below the LPF cutoff (here 2 GHz) are only weakly attenuated, which will allow for the application of strong, fast AC flux modulations to the tunable qubit. This is useful for many flux-based interaction schemes employing either baseband flux pulses [49, 50, 90] or finite-frequency parametric drives [56, 92]. Both dFBL ports are connected to individual input lines, but within the scope of this work we only apply signals to one input line and terminate the other with a $50\ \Omega$ load. Meanwhile, configuration B (used for dFBL2-4) utilizes a more conventional high-attenuation line (60 dB) to provide qubit drive pulses. As a consequence, only static flux bias can be applied in this configuration, which is sufficient for all experiments carried out here. For simplicity, only one dFBL port is connected to a drive line, while the other is terminated with a $50\ \Omega$ cryogenic load after the DC component has been removed by the bias tee.

¹An exception is made for qubit IST1 measured in [Chapter 5](#), which requires larger currents that can be supplied by the SRS sources and is instead biased with a Keithley 6221 current source.

The addition of lossy infrared filters (Eccosorb) in proximity of the device enclosure to filter out high-energy radiation was trialled during initial cooldowns. The use of IR filters has been shown to reduce the rate of generation of quasiparticles which might negatively affect the performance of superconducting qubits [224, 225]; In particular, placing the IR filters inside a light-tight radiation shield and as close as possible to the quantum device was found to be crucial. From the admittedly limited testing performed in our lab over the course of several cooldowns, we have not been able to detect a significant effect of the presence of IR filters on e.g. qubit lifetime and coherence. A more systematic investigation will be necessary to reach a conclusive answer. For simplicity, the IR filters were removed in later cooldowns.

In the following, we further mention two important modifications to the wiring and control diagram that were implemented for a subset of the experiments. The first concerns the addition of a 3D-integrated readout multiplexer in our sample enclosure, as presented in Sec. 4.5. This allows us to simplify the wiring inside the cryostat by removing the cryogenic microwave switch and reducing the number of coaxial cables going to the sample enclosure for readout from four to one, as shown in Fig. B.2 (a). This change also coincides with the (unrelated) removal of IR filters mentioned above. The second modification is an extension to the qubit drive generation scheme for the purpose of driving IST sideband transitions in the

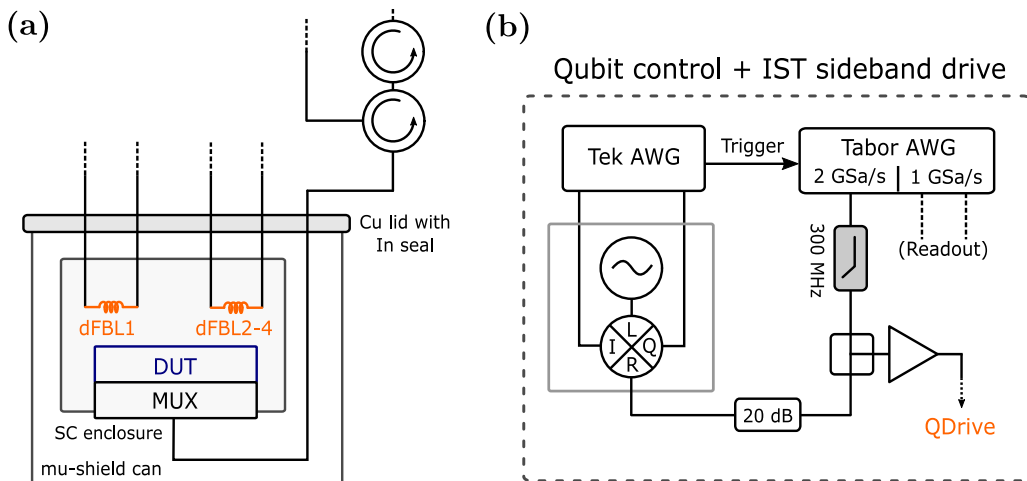


Figure B.2: Experimental setup – Modifications. (a) Modified readout circuitry employing a device-level 3D-integrated multiplexer (MUX) presented in Sec. 4.5. A single microwave line is used to read out all four resonators simultaneously. Cryogenic switch and IR filters have been removed. The rest of the setup is identical as in Fig. B.1. (b) Extended qubit control scheme used for IST sideband experiments in Chapter 6. A single AWG channel with sampling rate of 2 GSa/s implements red sideband drives up to around 300 MHz. These are combined with the standard microwave qubit drive, amplified and sent to the cryostat.

experiments of [Chapter 6](#). In order to apply sideband drives (usually in the range of a few 100 MHz for red sidebands) without losing the ability to apply conventional single qubit gates to the IST, we generate pulses from separate AWGs and combine them at room temperature before sending them to the dFBL inside the cryostat. This is shown schematically in [Fig. B.2](#) (b). Qubit Rabi drives are generated using the Tektronix AWG via up-conversion as before, while the sideband drives are generated directly from a single port of the Tabor AWG clocked at a sampling rate of 2 GSa/s to ensure sufficient sampling of the waveform. Synchronization of these pulses is achieved by using a trigger signal from the Tektronix to the Tabor. The use of a resistive power splitter/combiner reduces the signal power after combining the two sources, requiring an additional amplification stage before the signals are sent to the cryostat.

C.

Additional datasets




Figure C.1: Gradiometric transmon design variation. (a) Optical microscope image of gradiometric transmon variation with horizontal gradiometer orientation, i.e. perpendicular to the direction of Josephson junction deposition. Compare with reference design in Fig. 4.3. b) Detail of Josephson junction, showing the modified lead geometry.




Figure C.2: Flux crosstalk simulation. Simulated average flux crosstalk value $\bar{\chi}$ of the full crosstalk matrix $(\chi)_{ij}$ versus dFBL-qubit separation h_{dFBL} for two values of qubit metal permeability $\mu_{r,q}$. See Sec. 4.4.3 for details of the simulation method.

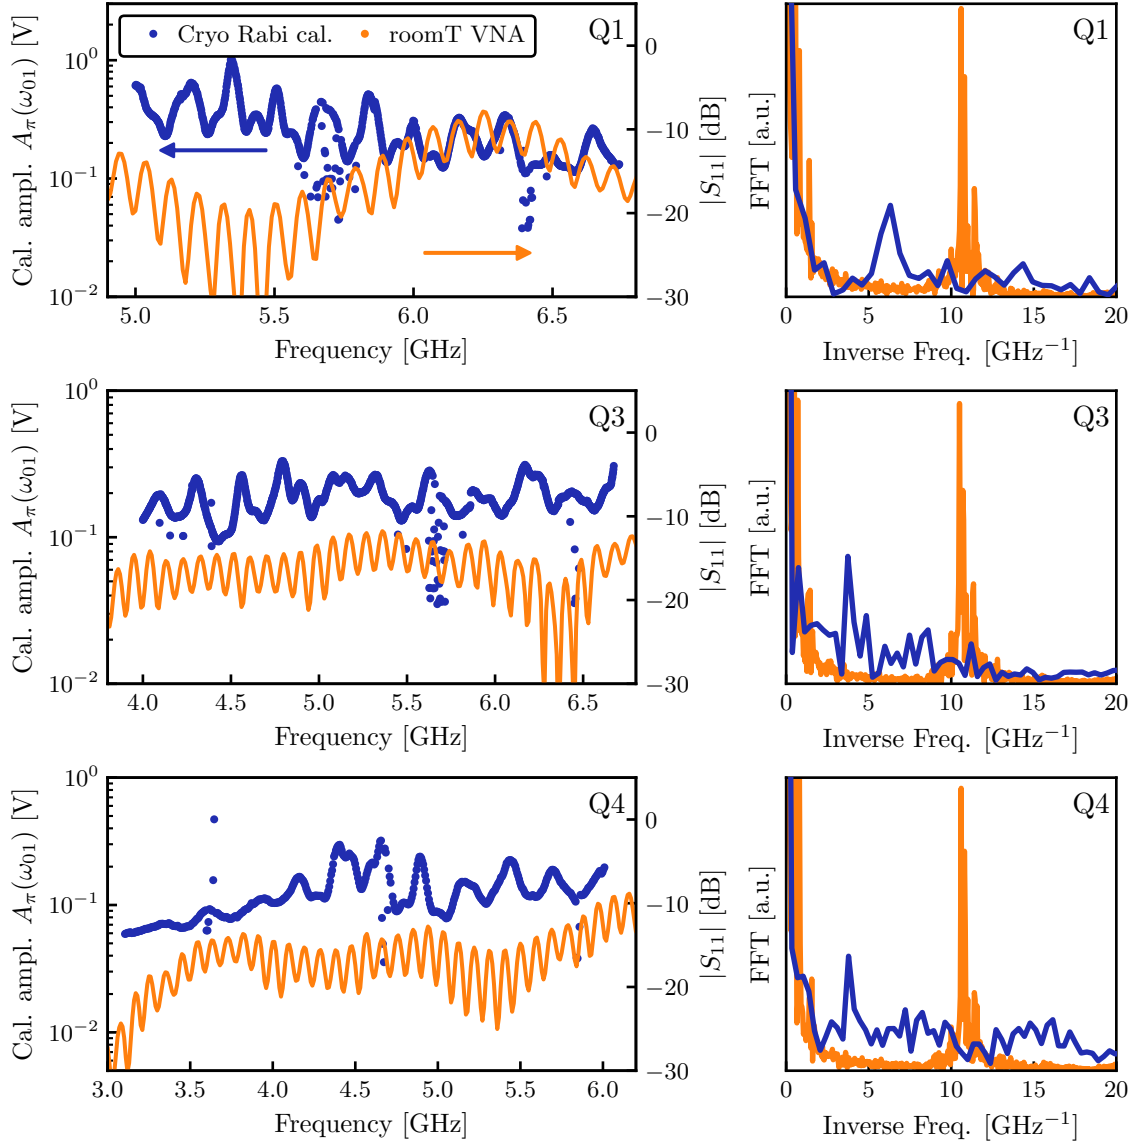


Figure C.3: Calibration of Rabi drive through dFBL. (a)–(c) AWG drive amplitude A_π required to perform a X_π pulse on qubits Q1, Q3, Q4 as a function of qubit frequency (blue points). A larger value corresponds to weaker coupling between AWG output and qubit. The reference VNA data from Fig. 4.7 in the relevant frequency range is added for comparison (orange). (d)–(f) FFT of the respective traces in (a)–(c).




Figure C.4: Direct comparison of T_1 and RO Rabi calibration. T_1 versus frequency (black) and Rabi pulse calibration through readout port (blue) for qubits Q2, Q3, and Q4 (top to bottom). The T_1 trace is obtained by applying a square window filter of width $n = 20$ to the raw data in Fig. 4.15. No readout port calibration data was obtained for Q1, which is therefore omitted here.

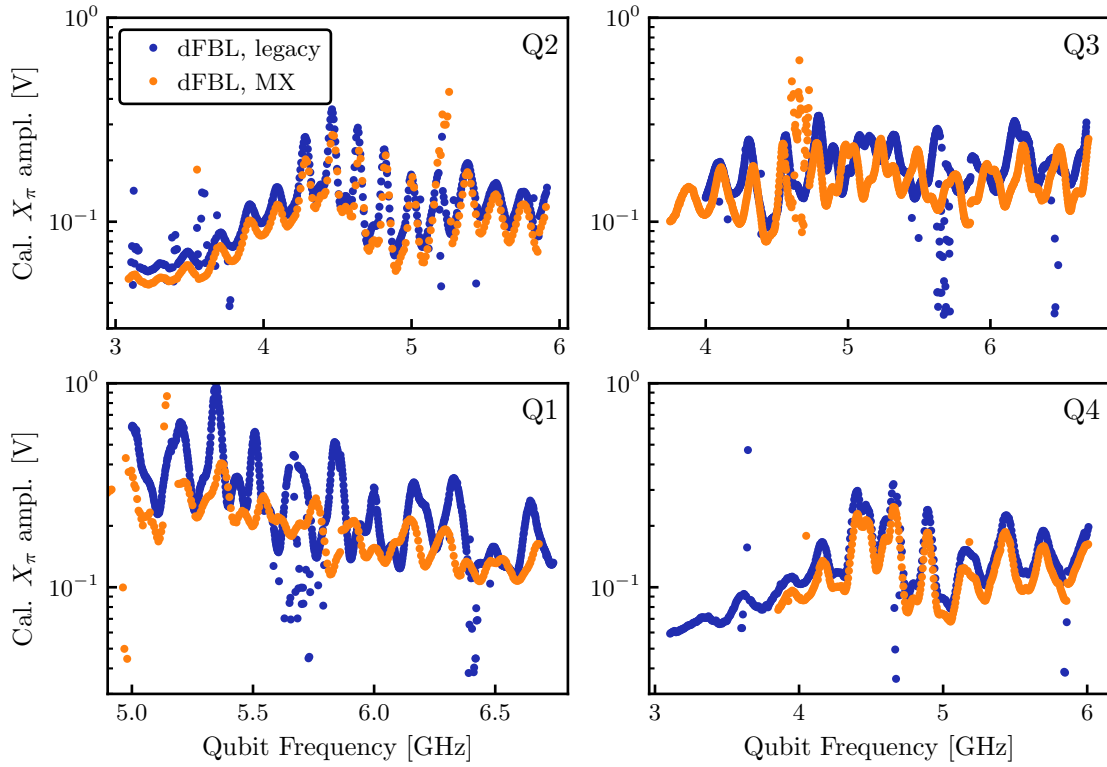


Figure C.5: Rabi drive calibration through dFBL in different cooldowns. AWG drive amplitude required to perform a X_π pulse on qubits Q1–Q4 of the gradiometric transmon device as described in Chapter 4. The drive is applied through the dFBL of each qubit. Data in blue refers to the first device cooldown using a sample enclosure with individual readout ports (legacy), while data in orange is for a second cooldown with the modified sample enclosure with readout multiplexer (MX).

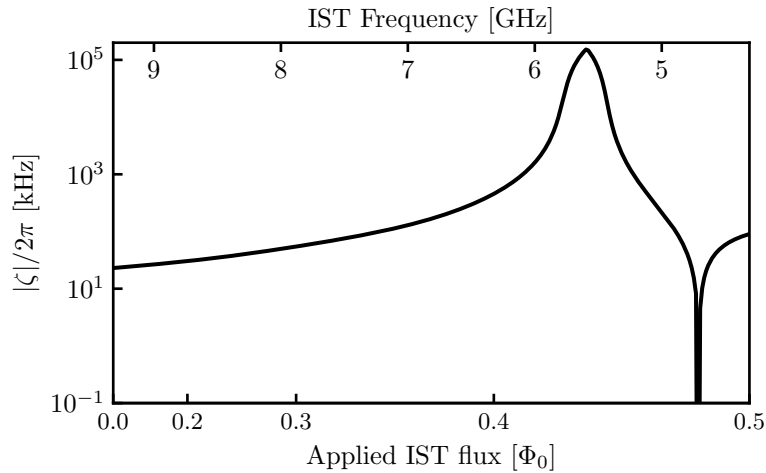


Figure C.6: Simulated static ZZ interaction between IST and transmon without parasitic modes. Same numerical simulation of the ZZ interaction as a function of IST flux as in Fig. 6.3, but excluding the parasitic modes of the IST inductor from the Hamiltonian.

Bibliography

- [1] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information: 10th Anniversary Edition* (Cambridge University Press, Cambridge, 2010).
- [2] R. P. Feynman, *Simulating Physics with Computers*, *International Journal of Theoretical Physics* **21**, 467 (1982).
- [3] J. Preskill, *Quantum Computing 40 Years Later*, [arxiv:2106.10522](https://arxiv.org/abs/2106.10522) (2023).
- [4] 'Plenty of Room' Revisited, *Nature Nanotechnology* **4**, 781 (2009).
- [5] P. W. Shor, *Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer*, *SIAM Journal on Computing* **26**, 1484 (1997).
- [6] D. Deutsch and R. Jozsa, *Rapid Solution of Problems by Quantum Computation*, *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* **439**, 553 (1997).
- [7] R. Cleve, A. Ekert, C. Macchiavello, and M. Mosca, *Quantum Algorithms Revisited*, *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences* **454**, 339 (1998).
- [8] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien, *Quantum Computers*, *Nature* **464**, 45 (2010).
- [9] S. S. Gill, A. Kumar, H. Singh, M. Singh, K. Kaur, M. Usman, and R. Buyya, *Quantum Computing: A Taxonomy, Systematic Review and Future Directions*, *Journal of Software: Practice and Experience* **52**, 66 (2022).
- [10] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, *Quantum Supremacy Using a Programmable Superconducting Processor*, *Nature* **574**, 505 (2019).
- [11] A. K. Fedorov, N. Gisin, S. M. Belousov, and A. I. Lvovsky, *Quantum Computing at the Quantum Advantage Threshold: A down-to-Business Review*, [arxiv:2203.17181](https://arxiv.org/abs/2203.17181) (2022).
- [12] D. P. DiVincenzo, *The Physical Implementation of Quantum Computation*, *Fortschritte der Physik* **48**, 771 (2000).

-
- [13] P. Hauke, H. G. Katzgraber, W. Lechner, H. Nishimori, and W. D. Oliver, *Perspectives of Quantum Annealing: Methods and Implementations*, [Reports on Progress in Physics](#) **83**, 054401 (2020).
- [14] B. K. Chakrabarti, H. Leschke, P. Ray, T. Shirai, and S. Tanaka, *Quantum Annealing and Computation: Challenges and Perspectives*, [Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences](#) **381**, 20210419 (2022).
- [15] J. I. Cirac and P. Zoller, *Goals and Opportunities in Quantum Simulation*, [Nature Physics](#) **8**, 264 (2012).
- [16] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, *Practical Quantum Advantage in Quantum Simulation*, [Nature](#) **607**, 667 (2022).
- [17] J. Preskill, *Quantum Computing in the NISQ Era and Beyond*, [Quantum](#) **2**, 79 (2018).
- [18] J. Preskill, *Fault-Tolerant Quantum Computation*, in *Introduction to Quantum Computation and Information* (WORLD SCIENTIFIC, 1998) pp. 213–269.
- [19] E. T. Campbell, B. M. Terhal, and C. Vuillot, *Roads towards Fault-Tolerant Universal Quantum Computation*, [Nature](#) **549**, 172 (2017).
- [20] L. Egan, D. M. Debroy, C. Noel, A. Risinger, D. Zhu, D. Biswas, M. Newman, M. Li, K. R. Brown, M. Cetina, and C. Monroe, *Fault-Tolerant Control of an Error-Corrected Qubit*, [Nature](#) **598**, 281 (2021).
- [21] G. Wendin, *Quantum Information Processing with Superconducting Circuits: A Review*, [Reports on Progress in Physics](#) **80**, 106001 (2017).
- [22] M. Kjaergaard, M. E. Schwartz, J. Braumüller, P. Krantz, J. I.-J. Wang, S. Gustavsson, and W. D. Oliver, *Superconducting Qubits: Current State of Play*, [Annual Review of Condensed Matter Physics](#) **11**, 369 (2020).
- [23] M. Gong, S. Wang, C. Zha, M.-C. Chen, H.-L. Huang, Y. Wu, Q. Zhu, Y. Zhao, S. Li, S. Guo, H. Qian, Y. Ye, F. Chen, C. Ying, J. Yu, D. Fan, D. Wu, H. Su, H. Deng, H. Rong, K. Zhang, S. Cao, J. Lin, Y. Xu, L. Sun, C. Guo, N. Li, F. Liang, V. M. Bastidas, K. Nemoto, W. J. Munro, Y.-H. Huo, C.-Y. Lu, C.-Z. Peng, X. Zhu, and J.-W. Pan, *Quantum Walks on a Programmable Two-Dimensional 62-Qubit Superconducting Processor*, [Science](#) **372**, 948 (2021).
- [24] Y. Wu, W.-S. Bao, S. Cao, F. Chen, M.-C. Chen, X. Chen, T.-H. Chung, H. Deng, Y. Du, D. Fan, M. Gong, C. Guo, C. Guo, S. Guo, L. Han, L. Hong, H.-L. Huang, Y.-H. Huo, L. Li, N. Li, S. Li, Y. Li, F. Liang, C. Lin, J. Lin, H. Qian, D. Qiao, H. Rong, H. Su, L. Sun, L. Wang, S. Wang, D. Wu, Y. Xu, K. Yan, W. Yang, Y. Yang, Y. Ye, J. Yin, C. Ying, J. Yu, C. Zha, C. Zhang, H. Zhang, K. Zhang, Y. Zhang, H. Zhao, Y. Zhao, L. Zhou, Q. Zhu, C.-Y. Lu, C.-Z. Peng, X. Zhu, and J.-W. Pan, *Strong Quantum Computational Advantage Using a Superconducting Quantum Processor*, [Physical Review Letters](#) **127**, 180501 (2021).

- [25] P. Ball, *First Quantum Computer to Pack 100 Qubits Enters Crowded Race*, [Nature](#) **599**, 542 (2021).
- [26] S. Krinner, N. Lacroix, A. Remm, A. Di Paolo, E. Genois, C. Leroux, C. Hellings, S. Lazar, F. Swiadek, J. Herrmann, G. J. Norris, C. K. Andersen, M. Müller, A. Blais, C. Eichler, and A. Wallraff, *Realizing Repeated Quantum Error Correction in a Distance-Three Surface Code*, [Nature](#) **605**, 669 (2022).
- [27] R. Acharya, I. Aleiner, R. Allen, T. I. Andersen, M. Ansmann, F. Arute, K. Arya, A. Asfaw, J. Atalaya, R. Babbush, D. Bacon, J. C. Bardin, J. Basso, A. Bengtsson, S. Boixo, G. Bortoli, A. Bourassa, J. Bovaird, L. Brill, M. Broughton, B. B. Buckley, D. A. Buell, T. Burger, B. Burkett, N. Bushnell, Y. Chen, Z. Chen, B. Chiaro, J. Cogan, R. Collins, P. Conner, W. Courtney, A. L. Crook, B. Curtin, D. M. Debroy, A. Del Toro Barba, S. Demura, A. Dunsworth, D. Eppens, C. Erickson, L. Faoro, E. Farhi, R. Fatemi, L. Flores Burgos, E. Forati, A. G. Fowler, B. Foxen, W. Giang, C. Gidney, D. Gilboa, M. Giustina, A. Grajales Dau, J. A. Gross, S. Habegger, M. C. Hamilton, M. P. Harrigan, S. D. Harrington, O. Higgott, J. Hilton, M. Hoffmann, S. Hong, T. Huang, A. Huff, W. J. Huggins, L. B. Ioffe, S. V. Isakov, J. Iveland, E. Jeffrey, Z. Jiang, C. Jones, P. Juhas, D. Kafri, K. Kechedzhi, J. Kelly, T. Khattar, M. Khezri, M. Kieferová, S. Kim, A. Kitaev, P. V. Klimov, A. R. Klots, A. N. Korotkov, F. Kostritsa, J. M. Kreikebaum, D. Landhuis, P. Laptev, K.-M. Lau, L. Laws, J. Lee, K. Lee, B. J. Lester, A. Lill, W. Liu, A. Locharla, E. Lucero, F. D. Malone, J. Marshall, O. Martin, J. R. McClean, T. McCourt, M. McEwen, A. Megrant, B. Meurer Costa, X. Mi, K. C. Miao, M. Mohseni, S. Montazeri, A. Morvan, E. Mount, W. Mruczkiewicz, O. Naaman, M. Neeley, C. Neill, A. Nersisyan, H. Neven, M. Newman, J. H. Ng, A. Nguyen, M. Nguyen, M. Y. Niu, T. E. O'Brien, A. Opremcak, J. Platt, A. Petukhov, R. Potter, L. P. Pryadko, C. Quintana, P. Roushan, N. C. Rubin, N. Saei, D. Sank, K. Sankaragomathi, K. J. Satzinger, H. F. Schurkus, C. Schuster, M. J. Shearn, A. Shorter, V. Shvarts, J. Skrzynny, V. Smelyanskiy, W. C. Smith, G. Sterling, D. Strain, M. Szalay, A. Torres, G. Vidal, B. Villalonga, C. Vollgraf Heidweiller, T. White, C. Xing, Z. J. Yao, P. Yeh, J. Yoo, G. Young, A. Zalcman, Y. Zhang, N. Zhu, and Google Quantum AI, *Suppressing Quantum Errors by Scaling a Surface Code Logical Qubit*, [Nature](#) **614**, 676 (2023).
- [28] S. Bravyi, O. Dial, J. M. Gambetta, D. Gil, and Z. Nazario, *The Future of Quantum Computing with Superconducting Qubits*, [Journal of Applied Physics](#) **132**, 160902 (2022).
- [29] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai, *Coherent Control of Macroscopic Quantum States in a Single-Cooper-pair Box*, [Nature](#) **398**, 786 (1999).
- [30] S. M. Girvin, M. H. Devoret, and R. J. Schoelkopf, *Circuit QED and Engineering Charge-Based Superconducting Qubits*, [Physica Scripta](#) **2009**, 014012 (2009).
- [31] W. D. Oliver and P. B. Welander, *Materials in Superconducting Quantum Bits*, [MRS Bulletin](#) **38**, 816 (2013).
- [32] P. Krantz, M. Kjaergaard, F. Yan, T. P. Orlando, S. Gustavsson, and W. D. Oliver, *A Quantum Engineer's Guide to Superconducting Qubits*, [Applied Physics Reviews](#) **6**, 021318 (2019).

- [33] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, *Circuit Quantum Electrodynamics*, [Reviews of Modern Physics](#) **93**, 025005 (2021).
- [34] B. D. Josephson, *Possible New Effects in Superconductive Tunnelling*, [Physics Letters](#) **1**, 251 (1962).
- [35] B. D. Josephson, *The Discovery of Tunnelling Supercurrents*, [Reviews of Modern Physics](#) **46**, 251 (1974).
- [36] M. Tinkham, *Introduction to Superconductivity* (Courier Corporation, 2004).
- [37] V. Bouchiat, D. Vion, P. Joyez, D. Esteve, and M. H. Devoret, *Quantum Coherence with a Single Cooper Pair*, [Physica Scripta](#) **1998**, 165 (1998).
- [38] T. Yamamoto, Y. A. Pashkin, O. Astafiev, Y. Nakamura, and J. S. Tsai, *Demonstration of Conditional Gate Operation Using Superconducting Charge Qubits*, [Nature](#) **425**, 941 (2003).
- [39] 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 (2004).
- [40] J. Koch, T. M. Yu, J. Gambetta, A. A. Houck, D. I. Schuster, J. Majer, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Charge-Insensitive Qubit Design Derived from the Cooper Pair Box*, [Physical Review A](#) **76**, 042319 (2007).
- [41] A. Blais, R.-S. Huang, A. Wallraff, S. M. Girvin, and R. J. Schoelkopf, *Cavity Quantum Electrodynamics for Superconducting Electrical Circuits: An Architecture for Quantum Computation*, [Physical Review A](#) **69**, 062320 (2004).
- [42] J. Majer, J. M. Chow, J. M. Gambetta, J. Koch, B. R. Johnson, J. A. Schreier, L. Frunzio, D. I. Schuster, A. A. Houck, A. Wallraff, A. Blais, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Coupling Superconducting Qubits via a Cavity Bus*, [Nature](#) **449**, 443 (2007).
- [43] N. Didier, E. A. Sete, M. P. da Silva, and C. Rigetti, *Analytical Modeling of Parametrically Modulated Transmon Qubits*, [Physical Review A](#) **97**, 022330 (2018).
- [44] 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*, [Physical Review Letters](#) **107**, 240501 (2011).
- [45] A. Nersisyan, S. Poletto, N. Alidoust, R. Manenti, R. Renzas, C.-V. Bui, K. Vu, T. Whyland, Y. Mohan, E. A. Sete, S. Stanwyck, A. Bestwick, and M. Reagor, *Manufacturing Low Dissipation Superconducting Quantum Processors*, [arxiv:1901.08042](#) (2019).

- [46] P. A. Spring, S. Cao, T. Tsunoda, G. Campanaro, S. Fasciati, J. Wills, M. Bakr, V. Chidambaram, B. Shteynas, L. Carpenter, P. Gow, J. Gates, B. Vlastakis, and P. J. Leek, *High Coherence and Low Cross-Talk in a Tileable 3D Integrated Superconducting Circuit Architecture*, [Science Advances](#) **8**, eabl6698 (2022).
- [47] A. P. M. Place, L. V. H. Rodgers, P. Mundada, B. M. Smitham, M. Fitzpatrick, Z. Leng, A. Premkumar, J. Bryon, A. Vrajitoarea, S. Sussman, G. Cheng, T. Madhavan, H. K. Babla, X. H. Le, Y. Gang, B. Jäck, A. Gyenis, N. Yao, R. J. Cava, N. P. de Leon, and A. A. Houck, *New Material Platform for Superconducting Transmon Qubits with Coherence Times Exceeding 0.3 Milliseconds*, [Nature Communications](#) **12**, 1779 (2021).
- [48] C. Wang, X. Li, H. Xu, Z. Li, J. Wang, Z. Yang, Z. Mi, X. Liang, T. Su, C. Yang, G. Wang, W. Wang, Y. Li, M. Chen, C. Li, K. Linghu, J. Han, Y. Zhang, Y. Feng, Y. Song, T. Ma, J. Zhang, R. Wang, P. Zhao, W. Liu, G. Xue, Y. Jin, and H. Yu, *Towards Practical Quantum Computers: Transmon Qubit with a Lifetime Approaching 0.5 Milliseconds*, [npj Quantum Information](#) **8**, 1 (2022).
- [49] L. DiCarlo, J. M. Chow, J. M. Gambetta, L. S. Bishop, B. R. Johnson, D. I. Schuster, J. Majer, A. Blais, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, *Demonstration of Two-Qubit Algorithms with a Superconducting Quantum Processor*, [Nature](#) **460**, 240 (2009).
- [50] V. Negîrneac, H. Ali, N. Muthusubramanian, F. Battistel, R. Sagastizabal, M. S. Moreira, J. F. Marques, W. J. Vlothuizen, M. Beekman, C. Zachariadis, N. Haider, A. Bruno, and L. DiCarlo, *High-Fidelity Controlled-Z Gate with Maximal Intermediate Leakage Operating at the Speed Limit in a Superconducting Quantum Processor*, [Physical Review Letters](#) **126**, 220502 (2021).
- [51] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, *Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits*, [Physical Review Letters](#) **103**, 110501 (2009).
- [52] M. Werninghaus, D. J. Egger, F. Roy, S. Machnes, F. K. Wilhelm, and S. Filipp, *Leakage Reduction in Fast Superconducting Qubit Gates via Optimal Control*, [npj Quantum Information](#) **7**, 1 (2021).
- [53] J. A. Schreier, A. A. Houck, J. Koch, D. I. Schuster, B. R. Johnson, J. M. Chow, J. M. Gambetta, J. Majer, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Suppressing Charge Noise Decoherence in Superconducting Charge Qubits*, [Physical Review B](#) **77**, 180502 (2008).
- [54] A. A. Houck, J. A. Schreier, B. R. Johnson, J. M. Chow, J. 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*, [Physical Review Letters](#) **101**, 080502 (2008).
- [55] M. D. Hutchings, J. B. Hertzberg, Y. Liu, N. T. Bronn, G. A. Keefe, M. Brink, J. M. Chow, and B. L. T. Plourde, *Tunable Superconducting Qubits with Flux-Independent Coherence*, [Physical Review Applied](#) **8**, 044003 (2017).

-
- [56] S. A. Caldwell, N. Didier, C. A. Ryan, E. A. Sete, A. Hudson, P. Karalekas, R. Marenti, M. P. da Silva, R. Sinclair, E. Acala, N. Alidoust, J. Angeles, A. Bestwick, M. Block, B. Bloom, A. Bradley, C. Bui, L. Capelluto, R. Chilcott, J. Cordova, G. Crossman, M. Curtis, S. Deshpande, T. E. Bouayadi, D. Girshovich, S. Hong, K. Kuang, M. Lenihan, T. Manning, A. Marchenkov, J. Marshall, R. Maydra, Y. Mohan, W. O'Brien, C. Osborn, J. Otterbach, A. Papageorge, J.-P. Paquette, M. Pelstring, A. Polloreno, G. Prawiroatmodjo, V. Rawat, M. Reagor, R. Renzas, N. Rubin, D. Russell, M. Rust, D. Scarabelli, M. Scheer, M. Selvanayagam, R. Smith, A. Staley, M. Suska, N. Tezak, D. C. Thompson, T.-W. To, M. Vahidpour, N. Vodrahalli, T. Whyland, K. Yadav, W. Zeng, and C. Rigetti, *Parametrically Activated Entangling Gates Using Transmon Qubits*, [Physical Review Applied](#) **10**, 034050 (2018).
- [57] A. Gyenis, A. Di Paolo, J. Koch, A. Blais, A. A. Houck, and D. I. Schuster, *Moving beyond the Transmon: Noise-Protected Superconducting Quantum Circuits*, [PRX Quantum](#) **2**, 030101 (2021).
- [58] F. Yan, Y. Sung, P. Krantz, A. Kamal, D. K. Kim, J. L. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver, *Engineering Framework for Optimizing Superconducting Qubit Designs*, [arxiv:2006.04130](#) (2020).
- [59] J. Koch, V. Manucharyan, M. H. Devoret, and L. I. Glazman, *Charging Effects in the Inductively Shunted Josephson Junction*, [Physical Review Letters](#) **103**, 217004 (2009).
- [60] J. E. Mooij, T. P. Orlando, L. Levitov, L. Tian, C. H. van der Wal, and S. Lloyd, *Josephson Persistent-Current Qubit*, [Science](#) **285**, 1036 (1999).
- [61] T. P. Orlando, J. E. Mooij, L. Tian, C. H. van der Wal, L. S. Levitov, S. Lloyd, and J. J. Mazo, *Superconducting Persistent-Current Qubit*, [Physical Review B](#) **60**, 15398 (1999).
- [62] I. Chiorescu, Y. Nakamura, C. J. P. M. Harmans, and J. E. Mooij, *Coherent Quantum Dynamics of a Superconducting Flux Qubit*, [Science](#) **299**, 1869 (2003).
- [63] F. Yan, S. Gustavsson, A. Kamal, J. Birenbaum, A. P. Sears, D. Hover, T. J. Gudmundsen, D. Rosenberg, G. Samach, S. Weber, J. L. Yoder, T. P. Orlando, J. Clarke, A. J. Kerman, and W. D. Oliver, *The Flux Qubit Revisited to Enhance Coherence and Reproducibility*, [Nature Communications](#) **7**, 12964 (2016).
- [64] V. E. Manucharyan, J. Koch, L. I. Glazman, and M. H. Devoret, *Fluxonium: Single Cooper-Pair Circuit Free of Charge Offsets*, [Science](#) **326**, 113 (2009).
- [65] L. B. Nguyen, Y.-H. Lin, A. Somoroff, R. Mencia, N. Grabon, and V. E. Manucharyan, *High-Coherence Fluxonium Qubit*, [Physical Review X](#) **9**, 041041 (2019).
- [66] D. G. Ferguson, A. A. Houck, and J. Koch, *Symmetries and Collective Excitations in Large Superconducting Circuits*, [Physical Review X](#) **3**, 011003 (2013).

- [67] 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*, [Physical Review B](#) **85**, 024521 (2012).
- [68] G. Viola and G. Catelani, *Collective Modes in the Fluxonium Qubit*, [Physical Review B](#) **92**, 224511 (2015).
- [69] C. T. Earnest, J. H. Béjanin, T. G. McConkey, E. A. Peters, A. Korinek, H. Yuan, and M. Mariani, *Substrate Surface Engineering for High-Quality Silicon/Aluminum Superconducting Resonators*, [Superconductor Science and Technology](#) **31**, 125013 (2018).
- [70] Y.-H. Lin, L. B. Nguyen, N. Grabon, J. San Miguel, N. Pankratova, and V. E. Manucharyan, *Demonstration of Protection of a Superconducting Qubit from Energy Decay*, [Physical Review Letters](#) **120**, 150503 (2018).
- [71] A. Somoroff, Q. Ficheux, R. A. Mencia, H. Xiong, R. V. Kuzmin, and V. E. Manucharyan, *Millisecond Coherence in a Superconducting Qubit*, [arxiv:2103.08578](#) (2021).
- [72] I. V. Pechenezhskiy, R. A. Mencia, L. B. Nguyen, Y.-H. Lin, and V. E. Manucharyan, *The Superconducting Quasicharge Qubit*, [Nature](#) **585**, 368 (2020).
- [73] M. Peruzzo, A. Trioni, F. Hassani, M. Zemlicka, and J. M. Fink, *Surpassing the Resistance Quantum with a Geometric Superinductor*, [Physical Review Applied](#) **14**, 044055 (2020).
- [74] M. Peruzzo, F. Hassani, G. Szep, A. Trioni, E. Redchenko, M. Žemlička, and J. M. Fink, *Geometric Superinductance Qubits: Controlling Phase Delocalization across a Single Josephson Junction*, [PRX Quantum](#) **2**, 040341 (2021).
- [75] C. Deng, J.-L. Orgiazzi, F. Shen, S. Ashhab, and A. Lupascu, *Observation of Floquet States in a Strongly Driven Artificial Atom*, [Physical Review Letters](#) **115**, 133601 (2015).
- [76] L. Verney, R. Lescanne, M. H. Devoret, Z. Leghtas, and M. Mirrahimi, *Structural Instability of Driven Josephson Circuits Prevented by an Inductive Shunt*, [Physical Review Applied](#) **11**, 024003 (2019).
- [77] R. Lescanne, L. Verney, Q. Ficheux, M. H. Devoret, B. Huard, M. Mirrahimi, and Z. Leghtas, *Escape of a Driven Quantum Josephson Circuit into Unconfined States*, [Physical Review Applied](#) **11**, 014030 (2019).
- [78] R. Shillito, A. Petrescu, J. Cohen, J. Beall, M. Hauru, M. Ganahl, A. G. Lewis, G. Vidal, and A. Blais, *Dynamics of Transmon Ionization*, [Physical Review Applied](#) **18**, 034031 (2022).
- [79] K. Zuo, Y. Urade, Z. Yan, S. Tamate, Y. Tabuchi, H. Terai, and Y. Nakamura, *Inductively shunted transmon qubit for ZZ interaction cancellation*, in [APS March Meeting Abstracts – X31.00010](#), APS March Meeting (2021).

-
- [80] F. Hassani, M. Peruzzo, L. N. Kapoor, A. Trioni, M. Zemlicka, and J. M. Fink, *A Superconducting Qubit with Noise-Insensitive Plasmon Levels and Decay-Protected Fluxon States*, [arxiv:2202.13917](#) (2022).
- [81] M. Steffen, S. Kumar, D. P. DiVincenzo, J. R. Rozen, G. A. Keefe, M. B. Rothwell, and M. B. Ketchen, *High-Coherence Hybrid Superconducting Qubit*, [Physical Review Letters](#) **105**, 100502 (2010).
- [82] J. R. Johansson, P. D. Nation, and F. Nori, *QuTiP: An Open-Source Python Framework for the Dynamics of Open Quantum Systems*, [Computer Physics Communications](#) **183**, 1760 (2012).
- [83] A. Noguchi, A. Osada, S. Masuda, S. Kono, K. Heya, S. P. Wolski, H. Takahashi, T. Sugiyama, D. Lachance-Quirion, and Y. Nakamura, *Fast Parametric Two-Qubit Gates with Suppressed Residual Interaction Using the Second-Order Nonlinearity of a Cubic Transmon*, [Physical Review A](#) **102**, 062408 (2020).
- [84] K. Banerjee, S. P. Bhatnagar, V. Choudhry, and S. S. Kanwal, *The Anharmonic Oscillator*, [Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences](#) **360**, 575 (1978).
- [85] J. B. Hertzberg, E. J. Zhang, S. Rosenblatt, E. Magesan, J. A. Smolin, J.-B. Yau, V. P. Adiga, M. Sandberg, M. Brink, J. M. Chow, and J. S. Orcutt, *Laser-Annealing Josephson Junctions for Yielding Scaled-up Superconducting Quantum Processors*, [npj Quantum Information](#) **7**, 1 (2021).
- [86] E. Hyppä, S. Kundu, C. F. Chan, A. Gunyhó, J. Hotari, D. Janzso, K. Juliusson, O. Kiuru, J. Kotilahti, A. Landra, W. Liu, F. Marxer, A. Mäkinen, J.-L. Orgiazzi, M. Palma, M. Savytskyi, F. Tosto, J. Tuorila, V. Vadimov, T. Li, C. Ockeloen-Korppi, J. Heinsoo, K. Y. Tan, J. Hassel, and M. Möttönen, *Unimon Qubit*, [Nature Communications](#) **13**, 6895 (2022).
- [87] S. Poletto, J. M. Gambetta, S. T. Merkel, J. A. Smolin, J. M. Chow, A. D. Córcoles, G. A. Keefe, M. B. Rothwell, J. R. Rozen, D. W. Abraham, C. Rigetti, and M. Steffen, *Entanglement of Two Superconducting Qubits in a Waveguide Cavity via Monochromatic Two-Photon Excitation*, [Physical Review Letters](#) **109**, 240505 (2012).
- [88] J. M. Chow, J. M. Gambetta, A. W. Cross, S. T. Merkel, C. Rigetti, and M. Steffen, *Microwave-Activated Conditional-Phase Gate for Superconducting Qubits*, [New Journal of Physics](#) **15**, 115012 (2013).
- [89] S. Sheldon, E. Magesan, J. M. Chow, and J. M. Gambetta, *Procedure for Systematically Tuning up Cross-Talk in the Cross-Resonance Gate*, [Physical Review A](#) **93**, 060302 (2016).
- [90] R. Barends, C. M. Quintana, A. G. Petukhov, Y. Chen, D. Kafri, K. Kechedzhi, R. Collins, O. Naaman, S. Boixo, F. Arute, K. Arya, D. Buell, B. Burkett, Z. Chen, B. Chiaro, A. Dunsworth, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, T. Huang, E. Jeffrey, J. Kelly, P. V. Klimov, F. Kostritsa, D. Landhuis, E. Lucero, M. McEwen, A. Megrant, X. Mi, J. Mutus, M. Neeley, C. Neill, E. Ostby, P. Roushan,

- D. Sank, K. J. Satzinger, A. Vainsencher, T. White, J. Yao, P. Yeh, A. Zalcman, H. Neven, V. N. Smelyanskiy, and J. M. Martinis, *Diabatic Gates for Frequency-Tunable Superconducting Qubits*, [Physical Review Letters](#) **123**, 210501 (2019).
- [91] A. Kandala, K. X. Wei, S. Srinivasan, E. Magesan, S. Carnevale, G. A. Keefe, D. Klaus, O. Dial, and D. C. McKay, *Demonstration of a High-Fidelity Cnot Gate for Fixed-Frequency Transmons with Engineered ZZ Suppression*, [Physical Review Letters](#) **127**, 130501 (2021).
- [92] E. A. Sete, N. Didier, A. Q. Chen, S. Kulshreshtha, R. Manenti, and S. Poletto, *Parametric-Resonance Entangling Gates with a Tunable Coupler*, [Physical Review Applied](#) **16**, 024050 (2021).
- [93] Q. Ficheux, L. B. Nguyen, A. Somoroff, H. Xiong, K. N. Nesterov, M. G. Vavilov, and V. E. Manucharyan, *Fast Logic with Slow Qubits: Microwave-Activated Controlled-Z Gate on Low-Frequency Fluxoniums*, [Physical Review X](#) **11**, 021026 (2021).
- [94] Y. Chen, K. N. Nesterov, V. E. Manucharyan, and M. G. Vavilov, *Fast Flux Entangling Gate for Fluxonium Circuits*, [Physical Review Applied](#) **18**, 034027 (2022).
- [95] F. Bao, H. Deng, D. Ding, R. Gao, X. Gao, C. Huang, X. Jiang, H.-S. Ku, Z. Li, X. Ma, X. Ni, J. Qin, Z. Song, H. Sun, C. Tang, T. Wang, F. Wu, T. Xia, W. Yu, F. Zhang, G. Zhang, X. Zhang, J. Zhou, X. Zhu, Y. Shi, J. Chen, H.-H. Zhao, and C. Deng, *Fluxonium: An Alternative Qubit Platform for High-Fidelity Operations*, [Physical Review Letters](#) **129**, 010502 (2022).
- [96] I. N. Moskalenko, I. A. Simakov, N. N. Abramov, A. A. Grigorev, D. O. Moskalev, A. A. Pishchimova, N. S. Smirnov, E. V. Zikiy, I. A. Rodionov, and I. S. Besedin, *High Fidelity Two-Qubit Gates on Fluxoniums Using a Tunable Coupler*, [npj Quantum Information](#) **8**, 1 (2022).
- [97] D. Weiss, H. Zhang, C. Ding, Y. Ma, D. I. Schuster, and J. Koch, *Fast High-Fidelity Gates for Galvanically-Coupled Fluxonium Qubits Using Strong Flux Modulation*, [PRX Quantum](#) **3**, 040336 (2022).
- [98] G. Burkard, M. J. Gullans, X. Mi, and J. R. Petta, *Superconductor–Semiconductor Hybrid-Circuit Quantum Electrodynamics*, [Nature Reviews Physics](#) **2**, 129 (2020).
- [99] R. Aguado, *A Perspective on Semiconductor-Based Superconducting Qubits*, [Applied Physics Letters](#) **117**, 240501 (2020).
- [100] J. Ku, X. Xu, M. Brink, D. C. McKay, J. B. Hertzberg, M. H. Ansari, and B. L. T. Plourde, *Suppression of Unwanted ZZ Interactions in a Hybrid Two-Qubit System*, [Physical Review Letters](#) **125**, 200504 (2020), [arxiv:2003.02775 \[cond-mat, physics:quant-ph\]](#) .
- [101] P. Zhao, P. Xu, D. Lan, J. Chu, X. Tan, H. Yu, and Y. Yu, *High-Contrast ZZ Interaction Using Superconducting Qubits with Opposite-Sign Anharmonicity*, [Physical Review Letters](#) **125**, 200503 (2020).

-
- [102] A. Ciani, B. M. Varbanov, N. Jolly, C. K. Andersen, and B. M. Terhal, *Microwave-Activated Gates between a Fluxonium and a Transmon Qubit*, [arxiv:2206.06203](#) (2022).
- [103] A. Maiani, M. Kjaergaard, and C. Schrade, *Entangling Transmons with Low-Frequency Protected Superconducting Qubits*, [PRX Quantum](#) **3**, 030329 (2022).
- [104] A. Finck, S. Carnevale, D. Klaus, C. Scerbo, J. Blair, T. McConkey, C. Kurter, A. Carniol, G. Keefe, M. Kumph, and O. Dial, *Suppressed Crosstalk between Two-Junction Superconducting Qubits with Mode-Selective Exchange Coupling*, [Physical Review Applied](#) **16**, 054041 (2021).
- [105] J. Wills, G. Campanaro, S. Cao, S. Fasciati, P. Leek, and B. Vlastakis, *Spatial Charge Sensitivity in a Multimode Superconducting Qubit*, [Physical Review Applied](#) **17**, 024058 (2022).
- [106] A. Gyenis, P. S. Mundada, A. Di Paolo, T. M. Hazard, X. You, D. I. Schuster, J. Koch, A. Blais, and A. A. Houck, *Experimental Realization of a Protected Superconducting Circuit Derived from the 0- π Qubit*, [PRX Quantum](#) **2**, 010339 (2021).
- [107] N. Schuch and J. Siewert, *Natural Two-Qubit Gate for Quantum Computation Using the XY Interaction*, [Physical Review A](#) **67**, 032301 (2003).
- [108] N. Didier, J. Bourassa, and A. Blais, *Fast Quantum Nondemolition Readout by Parametric Modulation of Longitudinal Qubit-Oscillator Interaction*, [Physical Review Letters](#) **115**, 203601 (2016).
- [109] S. Richer and D. DiVincenzo, *Circuit design implementing longitudinal coupling: A scalable scheme for superconducting qubits*, [Physical Review B](#) **93**, 134501 (2016).
- [110] B. K. Mitchell, R. K. Naik, A. Morvan, A. Hashim, J. M. Kreikebaum, B. Marinelli, W. Lavrijsen, K. Nowrouzi, D. I. Santiago, and I. Siddiqi, *Hardware-Efficient Microwave-Activated Tunable Coupling between Superconducting Qubits*, [Physical Review Letters](#) **127**, 200502 (2021).
- [111] S. Krinner, S. Lazar, A. Remm, C. Andersen, N. Lacroix, G. Norris, C. Hellings, M. Gabureac, C. Eichler, and A. Wallraff, *Benchmarking Coherent Errors in Controlled-Phase Gates Due to Spectator Qubits*, [Physical Review Applied](#) **14**, 024042 (2020).
- [112] P. Zhao, K. Linghu, Z. Li, P. Xu, R. Wang, G. Xue, Y. Jin, and H. Yu, *Quantum Crosstalk Analysis for Simultaneous Gate Operations on Superconducting Qubits*, [PRX Quantum](#) **3**, 020301 (2022).
- [113] M. A. Rol, F. Battistel, F. K. Malinowski, C. C. Bultink, B. M. Tarasinski, R. Vollmer, N. Haider, N. Muthusubramanian, A. Bruno, B. M. Terhal, and L. DiCarlo, *Fast, High-Fidelity Conditional-Phase Gate Exploiting Leakage Interference in Weakly Anharmonic Superconducting Qubits*, [Physical Review Letters](#) **123**, 120502 (2019).

- [114] P. Bertet, C. J. P. M. Harmans, and J. E. Mooij, *Parametric Coupling for Superconducting Qubits*, [Physical Review B **73**, 064512 \(2006\)](#).
- [115] Y. Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O'Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, M. R. Geller, A. N. Cleland, and J. M. Martinis, *Qubit Architecture with High Coherence and Fast Tunable Coupling*, [Physical Review Letters **113**, 220502 \(2014\)](#).
- [116] D. C. McKay, S. Filipp, A. Mezzacapo, E. Magesan, J. M. Chow, and J. M. Gambetta, *Universal Gate for Fixed-Frequency Qubits via a Tunable Bus*, [Physical Review Applied **6**, 064007 \(2016\)](#).
- [117] M. Roth, M. Ganzhorn, N. Moll, S. Filipp, G. Salis, and S. Schmidt, *Analysis of a Parametrically Driven Exchange-Type Gate and a Two-Photon Excitation Gate between Superconducting Qubits*, [Physical Review A **96**, 062323 \(2017\)](#).
- [118] E. A. Sete, A. Q. Chen, R. Manenti, S. Kulshreshtha, and S. Poletto, *Floating Tunable Coupler for Scalable Quantum Computing Architectures*, [Physical Review Applied **15**, 064063 \(2021\)](#).
- [119] F. Yan, P. Krantz, Y. Sung, M. Kjaergaard, D. L. Campbell, T. P. Orlando, S. Gustavsson, and W. D. Oliver, *Tunable Coupling Scheme for Implementing High-Fidelity Two-Qubit Gates*, [Physical Review Applied **10**, 054062 \(2018\)](#).
- [120] P. Mundada, G. Zhang, T. Hazard, and A. Houck, *Suppression of Qubit Crosstalk in a Tunable Coupling Superconducting Circuit*, [Physical Review Applied **12**, 054023 \(2019\)](#).
- [121] M. C. Collodo, J. Herrmann, N. Lacroix, C. K. Andersen, A. Remm, S. Lazar, J.-C. Besse, T. Walter, A. Wallraff, and C. Eichler, *Implementation of Conditional Phase Gates Based on Tunable ZZ Interactions*, [Physical Review Letters **125**, 240502 \(2020\)](#).
- [122] Y. Sung, L. Ding, J. Braumüller, A. Vepsäläinen, B. Kannan, M. Kjaergaard, A. Greene, G. O. Samach, C. McNally, D. Kim, A. Melville, B. M. Niedzielski, M. E. Schwartz, J. L. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver, *Realization of High-Fidelity CZ and ZZ-Free *i*SWAP Gates with a Tunable Coupler*, [Physical Review X **11**, 021058 \(2021\)](#).
- [123] J. Stehlik, D. M. Zajac, D. L. Underwood, T. Phung, J. Blair, S. Carnevale, D. Klaus, G. A. Keefe, A. Carniol, M. Kumph, M. Steffen, and O. E. Dial, *Tunable Coupling Architecture for Fixed-Frequency Transmon Superconducting Qubits*, [Physical Review Letters **127**, 080505 \(2021\)](#).
- [124] L. Viola and S. Lloyd, *Dynamical Suppression of Decoherence in Two-State Quantum Systems*, [Physical Review A **58**, 2733 \(1998\)](#).
- [125] L. Viola, E. Knill, and S. Lloyd, *Dynamical Decoupling of Open Quantum Systems*, [Physical Review Letters **82**, 2417 \(1999\)](#).

- [126] B. Pokharel, N. Anand, B. Fortman, and D. A. Lidar, *Demonstration of Fidelity Improvement Using Dynamical Decoupling with Superconducting Qubits*, [Physical Review Letters](#) **121**, 220502 (2018).
- [127] N. Sundaresan, I. Lauer, E. Pritchett, E. Magesan, P. Jurcevic, and J. M. Gambetta, *Reducing Unitary and Spectator Errors in Cross Resonance with Optimized Rotary Echoes*, [PRX Quantum](#) **1**, 020318 (2020).
- [128] H. Xiong, Q. Ficheux, A. Somoroff, L. B. Nguyen, E. Dogan, D. Rosenstock, C. Wang, K. N. Nesterov, M. G. Vavilov, and V. E. Manucharyan, *Arbitrary Controlled-Phase Gate on Fluxonium Qubits Using Differential Ac Stark Shifts*, [Physical Review Research](#) **4**, 023040 (2022).
- [129] K. X. Wei, E. Magesan, I. Lauer, S. Srinivasan, D. F. Bogorin, S. Carnevale, G. A. Keefe, Y. Kim, D. Klaus, W. Landers, N. Sundaresan, C. Wang, E. J. Zhang, M. Steffen, O. E. Dial, D. C. McKay, and A. Kandala, *Hamiltonian Engineering with Multicolor Drives for Fast Entangling Gates and Quantum Crosstalk Cancellation*, [Physical Review Letters](#) **129**, 060501 (2022).
- [130] X. Xu and M. Ansari, *ZZ Freedom in Two-Qubit Gates*, [Physical Review Applied](#) **15**, 064074 (2021).
- [131] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and J. M. Martinis, *Superconducting Quantum Circuits at the Surface Code Threshold for Fault Tolerance*, [Nature](#) **508**, 500 (2014).
- [132] D. M. Abrams, N. Didier, S. A. Caldwell, B. R. Johnson, and C. A. Ryan, *Methods for Measuring Magnetic Flux Crosstalk between Tunable Transmons*, [Physical Review Applied](#) **12**, 064022 (2019).
- [133] C. Neill, P. Roushan, K. Kechedzhi, S. Boixo, S. V. Isakov, V. Smelyanskiy, A. Megrant, B. Chiaro, A. Dunsworth, K. Arya, R. Barends, B. Burkett, Y. Chen, Z. Chen, A. Fowler, B. Foxen, M. Giustina, R. Graff, E. Jeffrey, T. Huang, J. Kelly, P. Klimov, E. Lucero, J. Mutus, M. Neeley, C. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, H. Neven, and J. M. Martinis, *A Blueprint for Demonstrating Quantum Supremacy with Superconducting Qubits*, [Science](#) **360**, 195 (2018).
- [134] R. Ma, B. Saxberg, C. Owens, N. Leung, Y. Lu, J. Simon, and D. I. Schuster, *A Dissipatively Stabilized Mott Insulator of Photons*, [Nature](#) **566**, 51 (2019).
- [135] Z. Yan, Y.-R. Zhang, M. Gong, Y. Wu, Y. Zheng, S. Li, C. Wang, F. Liang, J. Lin, Y. Xu, C. Guo, L. Sun, C.-Z. Peng, K. Xia, H. Deng, H. Rong, J. Q. You, F. Nori, H. Fan, X. Zhu, and J.-W. Pan, *Strongly Correlated Quantum Walks with a 12-Qubit Superconducting Processor*, [Science](#) **364**, 753 (2019).
- [136] J. Braumüller, A. H. Karamlou, Y. Yanay, B. Kannan, D. Kim, M. Kjaergaard, A. Melville, B. M. Niedzielski, Y. Sung, A. Vepsäläinen, R. Winik, J. L. Yoder, T. P. Orlando, S. Gustavsson, C. Tahan, and W. D. Oliver, *Probing Quantum Information Propagation with Out-of-Time-Ordered Correlators*, [Nature Physics](#) **18**, 172 (2022).

- [137] S. Krinner, S. Storz, P. Kurpiers, P. Magnard, J. Heinsoo, R. Keller, J. Lütolf, C. Eichler, and A. Wallraff, *Engineering Cryogenic Setups for 100-Qubit Scale Superconducting Circuit Systems*, *EPJ Quantum Technology* **6**, 1 (2019).
- [138] Z. Chen, A. Megrant, J. Kelly, R. Barends, J. Bochmann, Y. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, J. Y. Mutus, P. J. J. O'Malley, C. Neill, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, A. N. Cleland, and J. M. Martinis, *Fabrication and Characterization of Aluminum Airbridges for Superconducting Microwave Circuits*, *Applied Physics Letters* **104**, 052602 (2014).
- [139] J. A. Alfaro-Barrantes, M. Mastrangeli, D. J. Thoen, S. Visser, J. Bueno, J. J. A. Baselmans, and P. M. Sarro, *Superconducting High-Aspect Ratio Through-Silicon Vias With DC-Sputtered Al for Quantum 3D Integration*, *IEEE Electron Device Letters* **41**, 1114 (2020).
- [140] C. Navau, J. Prat-Camps, O. Romero-Isart, J. I. Cirac, and A. Sanchez, *Long-Distance Transfer and Routing of Static Magnetic Fields*, *Physical Review Letters* **112**, 253901 (2014).
- [141] O. Gargiulo, S. Oleschko, J. Prat-Camps, M. Zanner, and G. Kirchmair, *Fast Flux Control of 3D Transmon Qubits Using a Magnetic Hose*, *Applied Physics Letters* **118**, 012601 (2021).
- [142] M. Reed, *Entanglement and Quantum Error Correction with Superconducting Qubits*, Ph.D. thesis, Yale University (2013). Available at https://rsl.yale.edu/sites/default/files/files/RSL_Theses/reed.pdf.
- [143] R. Manenti, E. A. Sete, A. Q. Chen, S. Kulshreshtha, J.-H. Yeh, F. Oruc, A. Bestwick, M. Field, K. Jackson, and S. Poletto, *Full Control of Superconducting Qubits with Combined On-Chip Microwave and Flux Lines*, *Applied Physics Letters* **119**, 144001 (2021).
- [144] B. Foxen, J. Y. Mutus, E. Lucero, R. Graff, A. Megrant, Y. Chen, C. Quintana, B. Burkett, J. Kelly, E. Jeffrey, Y. Yang, A. Yu, K. Arya, R. Barends, Z. Chen, B. Chiaro, A. Dunsworth, A. Fowler, C. Gidney, M. Giustina, T. Huang, P. Klimov, M. Neeley, C. Neill, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, and J. M. Martinis, *Qubit Compatible Superconducting Interconnects*, *Quantum Science and Technology* **3**, 014005 (2017).
- [145] B. Johnson, *Controlling Photons in Superconducting Electrical Circuits*, Ph.D. thesis, Yale University (2011). Available at https://rsl.yale.edu/sites/default/files/files/RSL_Theses/johnson-thesis.pdf.
- [146] J. B. Majer, J. R. Butcher, and J. E. Mooij, *Simple Phase Bias for Superconducting Circuits*, *Applied Physics Letters* **80**, 3638 (2002).
- [147] M. J. Schwarz, J. Goetz, Z. Jiang, T. Niemczyk, F. Deppe, A. Marx, and R. Gross, *Gradiometric Flux Qubits with a Tunable Gap*, *New Journal of Physics* **15**, 045001 (2013).

- [148] J. Braumüller, M. Sandberg, M. R. Vissers, A. Schneider, S. Schlör, L. Grünhaupt, H. Rotzinger, M. Marthaler, A. Lukashenko, A. Dieter, A. V. Ustinov, M. Weides, and D. P. Pappas, *Concentric Transmon Qubit Featuring Fast Tunability and an Anisotropic Magnetic Dipole Moment*, [Applied Physics Letters](#) **108**, 032601 (2016).
- [149] F. C. Wellstood, C. Urbina, and J. Clarke, *Low-frequency Noise in Dc Superconducting Quantum Interference Devices below 1 K*, [Applied Physics Letters](#) **50**, 772 (1987).
- [150] F. Yoshihara, K. Harrabi, A. O. Niskanen, Y. Nakamura, and J. S. Tsai, *Decoherence of Flux Qubits Due to $1/f$ Flux Noise*, [Physical Review Letters](#) **97**, 167001 (2006).
- [151] R. C. Bialczak, R. McDermott, M. Ansmann, M. Hofheinz, N. Katz, E. Lucero, M. Neeley, A. D. O'Connell, H. Wang, A. N. Cleland, and J. M. Martinis, *$1/f$ Flux Noise in Josephson Phase Qubits*, [Physical Review Letters](#) **99**, 187006 (2007).
- [152] J. Braumüller, L. Ding, A. P. Vepsäläinen, Y. Sung, M. Kjaergaard, T. Menke, R. Winik, D. Kim, B. M. Niedzielski, A. Melville, J. L. Yoder, C. F. Hirjibehedin, T. P. Orlando, S. Gustavsson, and W. D. Oliver, *Characterizing and Optimizing Qubit Coherence Based on SQUID Geometry*, [Physical Review Applied](#) **13**, 054079 (2020).
- [153] R. H. Koch, D. P. DiVincenzo, and J. Clarke, *Model for $1/f$ Flux Noise in SQUIDs and Qubits*, [Physical Review Letters](#) **98**, 267003 (2007).
- [154] H. Wang, C. Shi, J. Hu, S. Han, C. C. Yu, and R. Q. Wu, *Candidate Source of Flux Noise in SQUIDs: Adsorbed Oxygen Molecules*, [Physical Review Letters](#) **115**, 077002 (2015).
- [155] P. Kumar, S. Sendelbach, M. A. Beck, J. W. Freeland, Z. Wang, H. Wang, C. C. Yu, R. Q. Wu, D. P. Pappas, and R. McDermott, *Origin and Reduction of $1/f$ Magnetic Flux Noise in Superconducting Devices*, [Physical Review Applied](#) **6**, 041001 (2016).
- [156] S. E. de Graaf, A. A. Adamyan, T. Lindström, D. Erts, S. E. Kubatkin, A. Ya. Tzalenchuk, and A. V. Danilov, *Direct Identification of Dilute Surface Spins on Al_2O_3 : Origin of Flux Noise in Quantum Circuits*, [Physical Review Letters](#) **118**, 057703 (2017).
- [157] S. Un, S. de Graaf, P. Bertet, S. Kubatkin, and A. Danilov, *On the Nature of Decoherence in Quantum Circuits: Revealing the Structural Motif of the Surface Radicals in α - Al_2O_3* , [Science Advances](#) **8**, eabm6169 (2022).
- [158] J. Bylander, S. Gustavsson, F. Yan, F. Yoshihara, K. Harrabi, G. Fitch, D. G. Cory, Y. Nakamura, J.-S. Tsai, and W. D. Oliver, *Noise Spectroscopy through Dynamical Decoupling with a Superconducting Flux Qubit*, [Nature Physics](#) **7**, 565 (2011).
- [159] T. Lanting, A. J. Berkley, B. Bumble, P. Bunyk, A. Fung, J. Johansson, A. Kaul, A. Kleinsasser, E. Ladizinsky, F. Maibaum, R. Harris, M. W. Johnson, E. Tolkacheva, and M. H. S. Amin, *Geometrical Dependence of the Low-Frequency Noise in Superconducting Flux Qubits*, [Physical Review B](#) **79**, 060509 (2009).

- [160] S. Sendelbach, D. Hover, A. Kittel, M. Mück, J. M. Martinis, and R. McDermott, *Magnetism in SQUIDs at Millikelvin Temperatures*, *Physical Review Letters* **100**, 227006 (2008).
- [161] J. Rahamim, T. Behrle, M. J. Peterer, A. Patterson, P. A. Spring, T. Tsunoda, R. Manenti, G. Tancredi, and P. J. Leek, *Double-Sided Coaxial Circuit QED with out-of-Plane Wiring*, *Applied Physics Letters* **110**, 222602 (2017).
- [162] A. Patterson, J. Rahamim, T. Tsunoda, P. Spring, S. Jebari, K. Ratter, M. Mergenthaler, G. Tancredi, B. Vlastakis, M. Esposito, and P. Leek, *Calibration of a Cross-Resonance Two-Qubit Gate Between Directly Coupled Transmons*, *Physical Review Applied* **12**, 064013 (2019).
- [163] V. Chidambaram, P. Spring, G. Campanaro, S. Cao, S. Fasciati, J. Wills, M. Bakr, B. Shteynas, and P. Leek, *Characterisation of a 3D-integrated 16-Qubit Superconducting Circuit*, in *APS March Meeting Abstracts – Z75.00011*, APS March Meeting (2023).
- [164] P. Spring, T. Tsunoda, B. Vlastakis, and P. Leek, *Modeling Enclosures for Large-Scale Superconducting Quantum Circuits*, *Physical Review Applied* **14**, 024061 (2020).
- [165] J. Rahamim, *Development of a Coaxial Circuit QED Architecture for Quantum Computing*, Ph.D. thesis, University of Oxford (2019). Available at <https://ora.ox.ac.uk/objects/uuid:c3311eef-7f31-4cfe-afd7-f3605b73ab36>.
- [166] C. Kurter, C. E. Murray, R. T. Gordon, B. B. Wymore, M. Sandberg, R. M. Shelby, A. Eddins, V. P. Adiga, A. D. K. Finck, E. Rivera, A. A. Stabile, B. Trimm, B. Wacaser, K. Balakrishnan, A. Pyzyna, J. Sleight, M. Steffen, and K. Rodbell, *Quasiparticle Tunneling as a Probe of Josephson Junction Barrier and Capacitor Material in Superconducting Qubits*, *npj Quantum Information* **8**, 1 (2022).
- [167] A. Dunsworth, A. Megrant, C. Quintana, Z. Chen, R. Barends, B. Burkett, B. Foxen, Y. Chen, B. Chiaro, A. Fowler, R. Graff, E. Jeffrey, J. Kelly, E. Lucero, J. Y. Mutus, M. Neeley, C. Neill, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, and J. M. Martinis, *Characterization and Reduction of Capacitive Loss Induced by Sub-Micron Josephson Junction Fabrication in Superconducting Qubits*, *Applied Physics Letters* **111**, 022601 (2017).
- [168] P. A. Spring, *Developing a Tileable Superconducting Circuit for Quantum Computation*, Ph.D. thesis, University of Oxford (2021). Available at <https://ora.ox.ac.uk/objects/uuid:0eb946e1-e9e3-45a8-ab2b-01035c71a127>.
- [169] M. F. Gely and G. A. Steele, *QuCAT: Quantum Circuit Analyzer Tool in Python*, *New Journal of Physics* **22**, 013025 (2020).
- [170] C. P. Poole, H. A. Farach, and R. J. Creswick, *Superconductivity* (Academic Press, 1995).
- [171] J. Zmuidzinas, *Superconducting Microresonators: Physics and Applications*, *Annual Review of Condensed Matter Physics* **3**, 169 (2012).

- [172] M. A. Rol, L. Ciorciaro, F. K. Malinowski, B. M. Tarasinski, R. E. Sagastizabal, C. C. Bultink, Y. Salathe, N. Haandbaek, J. Sedivy, and L. DiCarlo, *Time-Domain Characterization and Correction of on-Chip Distortion of Control Pulses in a Quantum Processor*, [Applied Physics Letters](#) **116**, 054001 (2020).
- [173] M. Jerger, A. Kulikov, Z. Vasselin, and A. Fedorov, *In Situ Characterization of Qubit Control Lines: A Qubit as a Vector Network Analyzer*, [Physical Review Letters](#) **123**, 150501 (2019).
- [174] E. M. Purcell, *Spontaneous emission probabilities at radio frequencies*, [Phys. Rev.](#) **69**, 681 (1946).
- [175] 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*, [Applied Physics Letters](#) **96**, 203110 (2010).
- [176] N. T. Bronn, E. Magesan, N. A. Masluk, J. M. Chow, J. M. Gambetta, and M. Steffen, *Reducing Spontaneous Emission in Circuit Quantum Electrodynamics by a Combined Readout/Filter Technique*, [IEEE Transactions on Applied Superconductivity](#) **25**, 1 (2015).
- [177] K. Koshino, S. Kono, and Y. Nakamura, *Protection of a Qubit via Subradiance: A Josephson Quantum Filter*, [Physical Review Applied](#) **13**, 014051 (2020).
- [178] S. Kono, K. Koshino, D. Lachance-Quirion, A. F. van Loo, Y. Tabuchi, A. Noguchi, and Y. Nakamura, *Breaking the Trade-off between Fast Control and Long Lifetime of a Superconducting Qubit*, [Nature Communications](#) **11**, 3683 (2020).
- [179] S. Filipp, P. Maurer, P. J. Leek, M. Baur, R. Bianchetti, J. M. Fink, M. Göppl, L. Steffen, J. M. Gambetta, A. Blais, and A. Wallraff, *Two-Qubit State Tomography Using a Joint Dispersive Readout*, [Physical Review Letters](#) **102**, 200402 (2009).
- [180] M. Jerger, S. Poletto, P. Macha, U. Hübner, E. Il'ichev, and A. V. Ustinov, *Frequency Division Multiplexing Readout and Simultaneous Manipulation of an Array of Flux Qubits*, [Applied Physics Letters](#) **101**, 042604 (2012).
- [181] J. Heinsoo, C. K. Andersen, A. Remm, S. Krinner, T. Walter, Y. Salathé, S. Gasparinetti, J.-C. Besse, A. Potočnik, A. Wallraff, and C. Eichler, *Rapid High-fidelity Multiplexed Readout of Superconducting Qubits*, [Physical Review Applied](#) **10**, 034040 (2018).
- [182] A. J. Annunziata, D. F. Santavicca, L. Frunzio, G. Catelani, M. J. Rooks, A. Frydman, and D. E. Prober, *Tunable Superconducting Nanoinductors*, [Nanotechnology](#) **21**, 445202 (2010).
- [183] H. Rotzinger, S. T. Skacel, M. Pfirrmann, J. N. Voss, J. Münzberg, S. Probst, P. Bushev, M. P. Weides, A. V. Ustinov, and J. E. Mooij, *Aluminium-Oxide Wires for Superconducting High Kinetic Inductance Circuits*, [Superconductor Science and Technology](#) **30**, 025002 (2016).
- [184] B. C. Wadell, *Transmission Line Design Handbook*, illustrated ed. (Artech Print on Demand, Boston, 1991).

-
- [185] P. Zhao, P. Xu, D. Lan, J. Chu, X. Tan, H. Yu, and Y. Yu, *High-Contrast ZZ Interaction Using Superconducting Qubits with Opposite-Sign Anharmonicity*, *Phys. Rev. Lett.* **125**, 200503 (2020).
- [186] A. Parra-Rodriguez, E. Rico, E. Solano, and I. L. Egusquiza, *Quantum Networks in Divergence-Free Circuit QED*, *Quantum Science and Technology* **3**, 024012 (2018).
- [187] M. Malekakhlagh, A. Petrescu, and H. E. Türeci, *Cutoff-Free Circuit Quantum Electrodynamics*, *Physical Review Letters* **119**, 073601 (2017).
- [188] 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*, *Physical Review Letters* **108**, 240502 (2012).
- [189] F. Solgun, D. W. Abraham, and D. P. DiVincenzo, *Blackbox Quantization of Superconducting Circuits Using Exact Impedance Synthesis*, *Physical Review B* **90**, 134504 (2014).
- [190] F. Solgun and D. P. DiVincenzo, *Multiport Impedance Quantization*, *Annals of Physics* **361**, 605 (2015).
- [191] Z. K. Mineev, Z. Leghtas, S. O. Mundhada, L. Christakis, I. M. Pop, and M. H. Devoret, *Energy-Participation Quantization of Josephson Circuits*, *npj Quantum Information* **7**, 1 (2021).
- [192] 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 (2004).
- [193] A. Wallraff, D. I. Schuster, A. Blais, J. M. Gambetta, J. Schreier, L. Frunzio, M. H. Devoret, S. M. Girvin, and R. J. Schoelkopf, *Sideband Transitions and Two-Tone Spectroscopy of a Superconducting Qubit Strongly Coupled to an On-Chip Cavity*, *Physical Review Letters* **99**, 050501 (2007).
- [194] K. Serniak, M. Hays, G. de Lange, S. Diamond, S. Shankar, L. D. Burkhardt, L. Frunzio, M. Houzet, and M. H. Devoret, *Hot Nonequilibrium Quasiparticles in Transmon Qubits*, *Physical Review Letters* **121**, 157701 (2018).
- [195] P. V. Klimov, J. Kelly, Z. Chen, M. Neeley, A. Megrant, B. Burkett, R. Barends, K. Arya, B. Chiaro, Y. Chen, A. Dunsworth, A. Fowler, B. Foxen, C. Gidney, M. Giustina, R. Graff, T. Huang, E. Jeffrey, E. Lucero, J. Y. Mutus, O. Naaman, C. Neill, C. Quintana, P. Roushan, D. Sank, A. Vainsencher, J. Wenner, T. C. White, S. Boixo, R. Babbush, V. N. Smelyanskiy, H. Neven, and J. M. Martinis, *Fluctuations of Energy-Relaxation Times in Superconducting Qubits*, *Physical Review Letters* **121**, 090502 (2018).
- [196] J. J. Burnett, A. Bengtsson, M. Scigliuzzo, D. Niepce, M. Kudra, P. Delsing, and J. Bylander, *Decoherence Benchmarking of Superconducting Qubits*, *npj Quantum Information* **5**, 1 (2019).

- [197] C. Müller, J. H. Cole, and J. Lisenfeld, *Towards Understanding Two-Level-Systems in Amorphous Solids: Insights from Quantum Circuits*, [Reports on Progress in Physics](#) **82**, 124501 (2019).
- [198] D. Gusenkova, F. Valenti, M. Spiecker, S. Günzler, P. Paluch, D. Rieger, L.-M. Pioraș-Tîmbolmaș, L. P. Zârbo, N. Casali, I. Colantoni, A. Cruciani, S. Pirro, L. Cardani, A. Petrescu, W. Wernsdorfer, P. Winkel, and I. M. Pop, *Operating in a Deep Underground Facility Improves the Locking of Gradiometric Fluxonium Qubits at the Sweet Spots*, [Applied Physics Letters](#) **120**, 054001 (2022).
- [199] S. Filipp, M. Göppl, J. M. Fink, M. Baur, R. Bianchetti, L. Steffen, and A. Wallraff, *Multimode Mediated Qubit-Qubit Coupling and Dark-State Symmetries in Circuit Quantum Electrodynamics*, [Physical Review A](#) **83**, 063827 (2011).
- [200] S. S. Hong, A. T. Papageorge, P. Sivarajah, G. Crossman, N. Didier, A. M. Polloreno, E. A. Sete, S. W. Turkowski, M. P. da Silva, and B. R. Johnson, *Demonstration of a Parametrically Activated Entangling Gate Protected from Flux Noise*, [Physical Review A](#) **101**, 012302 (2020).
- [201] M. Malekakhlagh, E. Magesan, and D. C. McKay, *First-Principles Analysis of Cross-Resonance Gate Operation*, [Physical Review A](#) **102**, 042605 (2020).
- [202] A. Morvan, L. Chen, J. M. Larson, D. I. Santiago, and I. Siddiqi, *Optimizing Frequency Allocation for Fixed-Frequency Superconducting Quantum Processors*, [Physical Review Research](#) **4**, 023079 (2022).
- [203] E. Magesan and J. M. Gambetta, *Effective Hamiltonian Models of the Cross-Resonance Gate*, [Physical Review A](#) **101**, 052308 (2020).
- [204] E. J. Zhang, S. Srinivasan, N. Sundaresan, D. F. Bogorin, Y. Martin, J. B. Hertzberg, J. Timmerwilke, E. J. Pritchett, J.-B. Yau, C. Wang, W. Landers, E. P. Lewandowski, A. Narasgond, S. Rosenblatt, G. A. Keefe, I. Lauer, M. B. Rothwell, D. T. McClure, O. E. Dial, J. S. Orcutt, M. Brink, and J. M. Chow, *High-Performance Superconducting Quantum Processors via Laser Annealing of Transmon Qubits*, [Science Advances](#) **8**, eabi6690 (2022).
- [205] N. E. Frattini, U. Vool, S. Shankar, A. Narla, K. M. Sliwa, and M. H. Devoret, *3-Wave Mixing Josephson Dipole Element*, [Applied Physics Letters](#) **110**, 222603 (2017).
- [206] N. E. Frattini, V. V. Sivak, A. Lingenfelter, S. Shankar, and M. H. Devoret, *Optimizing the Nonlinearity and Dissipation of a SNAIL Parametric Amplifier for Dynamic Range*, [Physical Review Applied](#) **10**, 054020 (2018).
- [207] P. J. Leek, S. Filipp, P. Maurer, M. Baur, R. Bianchetti, J. M. Fink, M. Göppl, L. Steffen, and A. Wallraff, *Using Sideband Transitions for Two-Qubit Operations in Superconducting Circuits*, [Physical Review B](#) **79**, 180511 (2009).
- [208] H. Xiong, Q. Ficheux, A. Somoroff, L. B. Nguyen, E. Dogan, D. Rosenstock, C. Wang, K. N. Nesterov, M. G. Vavilov, and V. E. Manucharyan, *Arbitrary Controlled-Phase Gate on Fluxonium Qubits Using Differential Ac-Stark Shifts*, [arxiv:2103.04491](#) (2021).

-
- [209] B. M. Vlastakis, *Controlling Coherent State Superpositions with Superconducting Circuits*, Ph.D. thesis, Yale University (2015). Available at https://rsl.yale.edu/sites/default/files/files/RSL_Theses/dissertation_vlastakis_0829.pdf.
- [210] E. Knill, D. Leibfried, R. Reichle, J. Britton, R. B. Blakestad, J. D. Jost, C. Langer, R. Ozeri, S. Seidelin, and D. J. Wineland, *Randomized Benchmarking of Quantum Gates*, *Physical Review A* **77**, 012307 (2008).
- [211] E. Magesan, J. M. Gambetta, and J. Emerson, *Scalable and Robust Randomized Benchmarking of Quantum Processes*, *Physical Review Letters* **106**, 180504 (2011).
- [212] D. C. McKay, C. J. Wood, S. Sheldon, J. M. Chow, and J. M. Gambetta, *Efficient Z Gates for Quantum Computing*, *Physical Review A* **96**, 022330 (2017).
- [213] J. M. Epstein, A. W. Cross, E. Magesan, and J. M. Gambetta, *Investigating the Limits of Randomized Benchmarking Protocols*, *Physical Review A* **89**, 062321 (2014).
- [214] J. M. Gambetta, A. D. Córcoles, S. T. Merkel, B. R. Johnson, J. A. Smolin, J. M. Chow, C. A. Ryan, C. Rigetti, S. Poletto, T. A. Ohki, M. B. Ketchen, and M. Steffen, *Characterization of Addressability by Simultaneous Randomized Benchmarking*, *Physical Review Letters* **109**, 240504 (2012).
- [215] M. Khezri, J. Dressel, and A. N. Korotkov, *Qubit Measurement Error from Coupling with a Detuned Neighbor in Circuit QED*, *Physical Review A* **92**, 052306 (2015).
- [216] D. M. Abrams, N. Didier, B. R. Johnson, M. P. da Silva, and C. A. Ryan, *Implementation of XY Entangling Gates with a Single Calibrated Pulse*, *Nature Electronics* **3**, 744 (2020).
- [217] A. D. Córcoles, J. M. Gambetta, J. M. Chow, J. A. Smolin, M. Ware, J. Strand, B. L. T. Plourde, and M. Steffen, *Process Verification of Two-Qubit Quantum Gates by Randomized Benchmarking*, *Physical Review A* **87**, 030301 (2013).
- [218] J. P. Gaebler, A. M. Meier, T. R. Tan, R. Bowler, Y. Lin, D. Hanneke, J. D. Jost, J. P. Home, E. Knill, D. Leibfried, and D. J. Wineland, *Randomized Benchmarking of Multiqubit Gates*, *Physical Review Letters* **108**, 260503 (2012).
- [219] E. Magesan, J. M. Gambetta, B. R. Johnson, C. A. Ryan, J. M. Chow, S. T. Merkel, M. P. da Silva, G. A. Keefe, M. B. Rothwell, T. A. Ohki, M. B. Ketchen, and M. Steffen, *Efficient Measurement of Quantum Gate Error by Interleaved Randomized Benchmarking*, *Physical Review Letters* **109**, 080505 (2012).
- [220] C. Leroux, A. Di Paolo, and A. Blais, *Superconducting Coupler with Exponentially Large On:Off Ratio*, *Physical Review Applied* **16**, 064062 (2021).
- [221] J. Cohen, A. Petrescu, R. Shillito, and A. Blais, *Reminiscence of Classical Chaos in Driven Transmons*, [arxiv:2207.09361](https://arxiv.org/abs/2207.09361) (2022).
- [222] Y. Yanay, J. Braumüller, S. Gustavsson, W. D. Oliver, and C. Tahan, *Two-Dimensional Hard-Core Bose-Hubbard Model with Superconducting Qubits*, *npj Quantum Information* **6**, 1 (2020).

- [223] G. J. Dolan and J. H. Dunsmuir, *Very Small (≈ 20 Nm) Lithographic Wires, Dots, Rings, and Tunnel Junctions*, [Physica B: Condensed Matter](#) **152**, 7 (1988).
- [224] K. Serniak, S. Diamond, M. Hays, V. Fatemi, S. Shankar, L. Frunzio, R. Schoelkopf, and M. Devoret, *Direct Dispersive Monitoring of Charge Parity in Offset-Charge-Sensitive Transmons*, [Physical Review Applied](#) **12**, 014052 (2019).
- [225] T. Connolly, P. D. Kurilovich, S. Diamond, H. Nho, C. G. L. Böttcher, L. I. Glazman, V. Fatemi, and M. H. Devoret, *Coexistence of Nonequilibrium Density and Equilibrium Energy Distribution of Quasiparticles in a Superconducting Qubit*, [arxiv:2302.12330](#) (2023).