THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Quantum information processing with tunable and low-loss superconducting circuits

ANDREAS BENGTSSON

Quantum Technology Laboratory
Department of Microtechnology and Nanoscience (MC2)
CHALMERS UNIVERSITY OF TECHNOLOGY

Göteborg, Sweden 2020

Quantum information processing with tunable and low-loss superconducting circuits
ANDREAS BENGTSSON
ISBN 978-91-7905-253-9

Doktorsavhandlingar vid Chalmers tekniska högskola
Ny serie nr. 4720
ISSN 0346-718X

Quantum Technology Laboratory
Department of Microtechnology and Nanoscience (MC2)
Chalmers University of Technology
SE-412 96 Göteborg
Sweden
Telephone: +46 (0)31-772 1000

Chalmers Digitaltryck
Göteborg, Sweden 2020

Quantum information processing with tunable and low-loss superconducting circuits
ANDREAS BENGTSSON
Quantum Technology Laboratory
Department of Microtechnology and Nanoscience (MC2)
Chalmers University of Technology

Abstract

The perhaps most promising platform for quantum information processing is the circuit-QED architecture based on superconducting circuits representing quantum bits. These circuits must be made with low losses so that the quantum information is retained for as long as possible. We developed fabrication processes achieving state-of-the-art coherence times of over 100 µs. We identified the primary source of loss to be parasitic two-level systems by studying fluctuations of qubit relaxation times.

Using our high-coherence circuits, we implemented a quantum processor built on fixed-frequency qubits and frequency-tunable couplers. The tunable couplers were lumped-element LC resonators, where the inductance came from a superconducting quantum interference device (SQUID). We achieved a controlled-phase gate with a fidelity of 99 % by parametric modulation of the coupler frequency. Using this device, and another similar to it, we demonstrated two different quantum algorithms, the quantum approximate optimization algorithm, and density matrix exponentiation. We achieved high algorithmic fidelities, aided by our carefully calibrated gates.

Additionally, we researched parametric oscillations using frequency-tunable resonators. Previously, degenerate parametric oscillations have been demonstrated by modulation of the resonant frequency at twice that frequency. We use this phenomenon to implement a readout method for a superconducting qubit with a fidelity of 98.7 %.

We demonstrated correlated radiation in nondegenerate parametric oscillations by modulating at the sum of two resonant frequencies of a multimode resonator. We showed an excellent quantitative agreement between the classical properties of the oscillations with a theoretical model. Moreover, we studied higher-order modulation at up to five times their resonant frequencies. These types of parametric oscillation states might be used as a quantum resource for continuous-variable quantum computing.

Keywords: superconducting circuits, quantum information, circuit quantum electrodynamics, high coherence, parametric modulation

List of publications

This thesis is based on the work contained in the following papers:

Paper A J. Burnett, A. Bengtsson, D. Niepce, and J. Bylander, "Noise and loss of superconducting aluminium resonators at single photon energies", Journal of Physics: Conference Series **969**, 012131 (2018)

Paper B 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**, 9 (2019)

Paper C Y. Lu, A. Bengtsson, J. J. Burnett, E. Wiegand, B. Suri, P. Krantz, A. F. Roudsari, A. F. Kockum, S. Gasparinetti, G. Johansson, and P. Delsing, "Characterizing decoherence rates of a superconducting qubit by direct microwave scattering", arXiv:1912.02124 (2019)

Paper D M. Scigliuzzo, L. E. Bruhat, A. Bengtsson, J. Burnett, A. F. Roudsari, and P. Delsing, "Phononic loss in superconducting resonators on piezoelectric substrates", Accepted for publication in New Journal of Physics (2020)

Paper E P. Krantz, A. Bengtsson, M. Simoen, S. Gustavsson, V. Shumeiko, W. Oliver, C. Wilson, P. Delsing, and J. Bylander, "Single-shot read-out of a superconducting qubit using a Josephson parametric oscillator", Nature Communications **7**, 1–8 (2016)

Paper F M. Kjaergaard, M. E. Schwartz, A. Greene, G. O. Samach, A. Bengtsson, M. O'Keeffe, C. M. McNally, J. Braumüller, D. K. Kim, P. Krantz, M. Marvian, A. Melville, B. M. Niedzielski, Y. Sung, R. Winik, J. Yoder, D. Rosenberg, K. Obenland, S. Lloyd, T. P. Orlando, I. Marvian, S. Gustavsson, and W. D. Oliver, "A quantum instruction set implemented on a superconducting quantum processor", arXiv:2001.08838 (2020)

Paper G A. Bengtsson, P. Vikstål, C. Warren, M. Svensson, X. Gu, A. F. Kockum, P. Krantz, C. Križan, D. Shiri, I.-M. Svensson, G. Tancredi, G. Johansson, P. Delsing, G. Ferrini, and J. Bylander, "Quantum approximate optimization of the exact-cover problem on a superconducting quantum processor", arXiv:1912.10495 (2019)

Paper H I.-M. Svensson, A. Bengtsson, P. Krantz, J. Bylander, V. Shumeiko, and P. Delsing, "Period-tripling subharmonic oscillations in a driven superconducting resonator", Physical Review B **96**, 174503 (2017)

Paper I I.-M. Svensson, M. Pierre, M. Simoen, W. Wustmann, P. Krantz, A. Bengtsson, G. Johansson, J. Bylander, V. Shumeiko, and P. Delsing, "Microwave photon generation in a doubly tunable superconducting resonator", Journal of Physics: Conference Series **969**, 012146 (2018)

Paper J I.-M. Svensson, A. Bengtsson, J. Bylander, V. Shumeiko, and P. Delsing, "Period multiplication in a parametrically driven superconducting resonator", Applied Physics Letters **113**, 022602 (2018)

Paper K A. Bengtsson, P. Krantz, M. Simoen, I.-M. Svensson, B. Schneider, V. Shumeiko, P. Delsing, and J. Bylander, "Nondegenerate parametric oscillations in a tunable superconducting resonator", Physical Review B **97**, 144502 (2018)

Paper L B. Schneider, A. Bengtsson, I. Svensson, T. Aref, G. Johansson, J. Bylander, and P. Delsing, "Observation of broadband entanglement in microwave radiation from a single time-varying boundary condition", arXiv:1802.05529 (2018)

Paper M S. R. Sathyamoorthy, A. Bengtsson, S. Bens, M. Simoen, P. Delsing, and G. Johansson, "Simple, robust, and on-demand generation of single and correlated photons", Physical Review A **93**, 063823 (2016)

Other papers that are outside the scope of this thesis:

Paper I Y. Sung, A. Vepsäläinen, J. Braumüller, F. Yan, J. I.-J. Wang, M. Kjaergaard, R. Winik, P. Krantz, A. Bengtsson, A. J. Melville, B. M. Niedzielski, M. E. Schwartz, D. K. Kim, J. L. Yoder, T. P. Orlando, S. Gustavsson, and W. D. Oliver, "Multi-level quantum noise spectroscopy", arXiv:2003.02782 (2020)

Contents

Acknowledgements

First of all, a big thank you to my supervisor Jonas Bylander for your great support and leadership. This thesis and our results would not exist without your guidance, knowledge, and attention to detail. Secondly, I wish to express my gratitude to my examiner and co-supervisor, Per Delsing. Thank you for leading our team, supporting me, and answering all my questions about superconductivity and low-temperature physics.

To Vitaly Shumeiko, Göran Johansson, Giulia Ferrini, and the rest of the theory team, thank you for all your help, without you the results wouldn't have been nearly as good or well understood.

A special mention goes to Philip Krantz and Michaël Simoen for helping me so much at the beginning of my Ph.D. Thank you for all your encouragement, and for letting me help and learn from you. Also to Christian Križan for proofreading this thesis.

Thanks to all members of the EQUS group at MIT for your hospitality during my stay in Cambridge. You are both great colleges and friends, and you are all doing fantastic work. look forward to seeing where your careers will end up. We will definitely see each other somewhere around the globe. An extra thank you goes to Simon Gustavsson, Morten Kjaergaard, and Bharath Kannan for making sure that I felt at home from day one.

None of the devices in this thesis would exist without the excellent support of the staff at the Nanofabrication Laboratory. My gratitude, especially to Niclas Lindvall, Henrik Fredriksen, Mats Hagberg, and Johan Andersson for all the help with lithography, thin film deposition, and etching.

It doesn't help to have excellent devices unless you have some way to mount them in your dilution refrigerator. Lars Jönsson helped us machine all the parts for sample boxes, mounting plates, coils, and other things. Sorry for all the urgent jobs with late notice! Your work is very much appreciated.

The administrative support from Susannah Carlsson and Linda Brånell was invaluable! Thank you for all the help to make sure that being a Ph.D. student is as easy as possible. You really allow us to focus on the research, knowing that the rest is taken care of. Same goes for Henric Fjellstedt, our IT-manager Thanks for always being on our side in the battle for good IT solutions.

To my parents, Ingrid and Anders, and my siblings, Sandra, Kaij, Amelie and Tobias, thank you for all the support during my 23 years of school! My friends, you are too many to mention by name. I love you all, and thank you for being there and filling my life with so much joy.

Finally, a big thank you to all current and former employees at QTL, AQP, and QDP for making it a great working place.

Andreas Bengtsson, Göteborg, March 2020

Thesis

CHAPTER I

Introduction

Computers are used everywhere in today's society. We use them to predict the weather, find the shortest route between two points, keep track of our money, and watch videos of cats on the Internet. On a fundamental level, a computer is a combination of logical gates operating on a set of bits. However, not many computer programmers express their programs in terms of 'microcode,' which describes how the computer should modify its bits. Instead, high-level programming languages and compilers are used to allow the programmer to express their program and its intended function more easily.

A computer program needs to implement some algorithm that takes input data and produces the correct solution. There can be many algorithms for solving the same problem, but the time complexity of each algorithm can be different. Time complexity describes how the time it takes to solve a problem scales with the size of the input data. For example, a brute-force algorithm for searching for a number in a list would be to go through and check all numbers in the list, resulting in linear time complexity, $\mathcal{O}(N)$, where N is the number of numbers in the list. However, if we know that the list is sorted, we can search through it in $\mathcal{O}(\log N)$ time by checking whether our number is lower or higher than the number in the middle of the list. We can then discard half of the list, and continue doing so until we find our number. This algorithm is known as a binary search and is exponentially faster than the brute-force search. On the other hand, if we need to search the list many times, it would be better to store the numbers in a different data structure than a list. For instance, a hash table would allow for searching in $\mathcal{O}(1)$ time.

As we just saw from this simple example, the choice of algorithm and the resulting time complexity depend on how the data is structured and what we want to do with the data. The same is true for quantum algorithms executed on quantum computers. Quantum computers will not give a universal speed-up to all problems and algorithms. However, it can give speed-ups for specific problems in certain applications [1].

1.1 Quantum information

A quantum computer is predicted to be able to solve problems that today are intractable for classical computers, such as prime factorization of large integers [2] and quantum simulation of molecules [3]. The holy grail for a quantum computer is exponential speed-up over a classical computer. For instance, Shor's algorithm for factorization has a polynomial-time complexity [2], which can be compared to the best *known* classical algorithm which runs in sub-exponential time [4] (however, still much slower than polynomial). Therefore, Shor's algorithm provides an almost exponential speed-up.

The fundamental building block of a quantum computer is the quantum bit (qubit). It takes inspiration from the classical bit, which can represent either a zero or a one. In the context of qubits, Dirac's ket notation is commonly used to represent the two states, $|0\rangle$ and $|1\rangle$. The laws of quantum mechanics allow for superposition states, meaning that each state is associated with a probability amplitude. Superposition states are sometimes (wrongly) referred to as a qubit being in more than one state at the same time and the source of parallel-processing power in quantum computers. However, superposition states do enable parallel access to all 2^n states available to n qubits, which is a part of why quantum algorithms can outperform classical algorithms.

For instance, take two qubits with four possible states: $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$. It is possible to create an equal superposition of the four states, where each measurement of the system would randomly yield only one of the states with a uniform probability (essentially a random number generator). A computer that gives random answers is not a good computer. We might have been in a superposition of all possible states, but the outcome is not useful. The trick is not only to "try" all states at the same time; we also need to be in the state corresponding to the solution to the problem by the time of measurement. Therefore, a quantum algorithm must manipulate the qubits so that they first utilize the large state space available to them, and then converge into to final state at the end of the algorithm. The process to converge into one particular state relies on the fact that the probability amplitudes can actually be negative (or even complex), which allows for destructive interference between different state amplitudes (cf. the famous double-slit experiment). This interference between different states is what gives quantum computers their power [5]. Without the interference effect, the computer would not be a quantum computer, but instead a probabilistic Turing machine [6], which can be used to factorize integers [7], but not in polynomial time.

Another example of a quantum algorithm is the HHL algorithm for solving linear systems of equations [8], which it does in $\mathcal{O}(\log N)$ time, where N is the number of variables. Classically, this problem would require $\mathcal{O}(N)$ time. So, HHL provides an exponential speed-up, even if the classical algorithm is polynomial in time. HHL does come with some caveats, it does, for instance, not provide the solution vector directly, and it requires a quantum random access memory. In general, quantum algorithms can provide exponential speed-ups for linear algebra; however, one should read the fine print of the practical limitations [9].

There are other quantum algorithms that "only" provide polynomial speed-ups. The most famous one being Grover's search algorithm, which can search through an unsorted list in $\mathcal{O}(\sqrt{N})$ time. Sometimes, the phrase 'Grover-type speed-up' is used to describe

quadratic speed-up. One practical application of Grover's algorithm would be to break symmetric encryption (e.g., AES) by brute force (i.e., by guessing each possible combination). However, these types of attacks are rather weak, since a doubling in the encryption key size would erase the quadratic speed-up gained by Grover's algorithm.

To run useful versions of Shor's and HHL would require large-scale fault-tolerant quantum computers. Large-scale meaning thousands of quantum bits, and fault-tolerant meaning that gates can be executed on those qubits without any errors [10, 11]. It is unlikely that physical qubits will ever have low enough error rates on their own; instead, quantum error correction (QEC) will be needed [5, 12]. Briefly, QEC encodes the state of a logical qubit into many physical qubits, which leads to reduced error rates but also a significant overhead in the number of qubits needed. As an example, factorizing a 2048 bit number with Shor's algorithm would require 6144 logical qubits. However, even when taking into account state-of-the-art physical error rates of 10^{-3}, the overhead of QEC brings the number of physical qubits up to 20 million [13, 14].

Current quantum processors are not fault-tolerant and are usually referred to as noisy intermediate-scale quantum (NISQ) devices [15]. Nevertheless, such processors hold great promise; for instance, they might allow for the execution of heuristic quantum algorithms solving combinatorial optimization problems. On small-scale quantum processors, these algorithms can serve as technology demonstrators. While there is no proven speed-up (hence, heuristic), one argument for why quantum computers could be good at combinatorial optimization problems is the enormous state-space available to them. For instance, the number of different combinations that 60 people can sit on 60 chairs exceeds the number of atoms in the known universe, and it would be impossible for a classical computer to try all the different combinations. On the other hand, a quantum computer with 272 qubits can form a superposition of that many states. The difficulty lies in designing an algorithm that gives a high probability of measuring the optimal state at the end of the algorithm.

Finding useful algorithms for quantum computers is no easy task. Proving their performance and time complexity can be even harder. For each specific problem, a quantum algorithm needs to be found that not only utilizes the quantum properties of the qubits (i.e., superposition and state-amplitude interference) but also yields the correct result upon a measurement of the qubits' states.

In 2019, Arute et al. [16], demonstrated that a 53 qubit quantum computer could perform a computation that is intractable on even the most powerful supercomputer in the world. It is important to note that it is not proven that it cannot exist a classical algorithm that does the same computation efficiently. Moreover, the computation itself does not have any practical application. Nevertheless, it most likely demonstrates that quantum computers can do better than classical computers on some problems (although not every problem).

To build a successful quantum computer, both small- and large-scale, several criteria need to be fulfilled [17]. These criteria are related to, among other things, coherence, control, and readout of qubits, which all are topics of the research included in this thesis.

Quantum bits can be realized with many different physical systems. Some common ones include nitrogen-vacancies in diamonds [18, 19], quantum dots in semiconductors [20], trapped ions [21, 22], semiconducting nanowires [23] and 2D electron gases [24],

superconducting tunnel junctions, and perhaps topological protected Majorana fermions [25]. In this work, we solely use superconducting tunnel junctions and circuits to build our devices.

1.2 Superconducting devices and circuit QED

As the name suggests, superconductivity is a phenomenon where the conductivity is very large. In fact, it is infinite for direct currents. However, there is a critical current, I_c, which is the maximum current that a superconductor can sustain without returning to its normal (non-superconducting) state. The microscopics behind superconductivity is modeled as an attractive interaction between the electrons in the material. Two electrons pair up due to this interaction and create a so-called Cooper pair. Cooper pairs, in contrast to electrons, are bosons. Meaning that below the critical temperature T_c, they condense into a common ground state. Due to the binding between two electrons in a Cooper pair, a gap in the density of states is opened where electrons cannot exist. Due to this gap, the Cooper pairs can move without scattering inside the superconductor, yielding zero resistivity. More information about superconductivity and its microscopic theory can be found in Ref. [26].

A Josephson junction is formed by separating two superconductors via a thin insulator. The insulator acts as a tunneling barrier, which Cooper pairs can tunnel through without dissipating energy, making it possible to draw current through the barrier without applying any voltage across it. Moreover, the Josephson junction also acts as an inductor, with an inductance given by

$$L = \frac{\hbar}{2e}\frac{1}{I_c \cos\varphi}, \tag{1.1}$$

where φ is the superconducting phase difference between the two superconductors, $\hbar$ is the reduced Planck constant, and e is the electron charge. φ is related to the current I through the junction,

$$I = I_c \sin(\varphi). \tag{1.2}$$

Superconducting circuits and devices [26, 27] are today found in a variety of systems, such as MRI and MEG machines at hospitals, particle accelerators, and power grids. Three main features make superconducting circuits so versatile and useful. Firstly, the absence of electrical resistance allows for low-loss electrical circuits, as well as the generation of high magnetic fields. Secondly, tunable nonlinear elements can be implemented and used as, for example, signal amplifiers, mixers, and detectors. Finally, superconducting circuits have a high susceptibility to changes in their environment, enabling a broad range of different detectors and sensors for energies close to the quantum level [28].

In this thesis, we use superconducting circuits for quantum information processing. Superconducting qubits are one of the most promising architectures for realizing a useful quantum computer [29–36]. During the last two decades, superconducting circuits have played a significant role in the development of what is now circuit quantum electrodynamics (cQED) [37–39]. In other words, the study of light-matter interactions at the quantum level using circuits (instead of actual atoms). For instance, coherence times exceeding

100 µs [Paper B], gate fidelities above 99% [40], readout in less than 50 ns [41], and scalable 3D-integration [42] have all been demonstrated.

In cQED, there are three main building blocks. The first one is the resonator. A typical example of a resonator is the LC circuit, which can be used as a storage for energy and is characterized by its resonant frequency and quality factor. A quality factor is a measure of the number of cycles that energy can be stored in the resonator. In this work, resonators are built out of planar superconducting circuits, which lets us achieve quality factors of several millions [43].

The second block is the superconducting quantum interference device (SQUID) [44]. A SQUID consists of two superconducting Josephson tunnel junctions connected in parallel, see Fig. 1.1 (c-d). A magnetic field threading the loop formed by the two junctions will induce a circulating current, which effectively decreases the critical current of the SQUID. Assuming identical junctions, Eq. (1.1) modifies to

$$L = \frac{\hbar}{2e} \frac{1}{2I_{\mathrm{c}}|\cos(\pi\Phi/\Phi_0)|\cos\varphi}, \tag{1.3}$$

where Φ is the magnetic flux threading the SQUID loop and $\Phi_0 = h/(2e)$ is the magnetic flux quantum. Therefore, the SQUID can be used as a magnetic flux tunable inductor. For a comprehensive review of SQUIDs, see Ref. [44].

The tunability allows for precise and rapid modulation of different circuit parameters, such as resonant frequencies [45] and couplings [46]. The nonlinearity of the Josephson junction with respect to the current I, and the fact that it can be made strong, is what allows for the generation of nonclassical states. A strong nonlinearity is one of the reasons for the success of cQED, as it allows for experiments not possible with optical photons where nonlinearities are much weaker.

The third and final block is the quantum bit (qubit). A qubit is a two-level system typically used to store quantum information. Moreover, a qubit can be used as an artificial atom to study how light and atoms interact. The benefits of using circuits and not real atoms are that it is easier to isolate and manipulate a circuit, and that is easier to engineer its properties.

1.2.1 Planar resonators

Microwave resonators are fundamental parts of many electrical circuits. Resonators can be used to store electromagnetic energy in the microwave regime, to filter noisy signals, to stabilize oscillators, and to achieve high electric fields [47]. Superconducting resonators are often used in applications requiring very low losses (high quality factors), such as in particle accelerators [48], radiation detectors [49], and quantum devices [32]. In quantum physics, the electromagnetic field inside a resonator can be quantized and instead described using particles known as photons [50]. The quantization of the electromagnetic field leads to the presence of vacuum fluctuations, which means that even at zero temperature, the resonator is still oscillating.

A piece of a transmission line can realize a microwave resonator. By fixing the length of the transmission line, and the two boundary conditions (usually open and grounded), see Fig. 1.1 (a-b), a certain mode structure is achieved with some resulting resonant

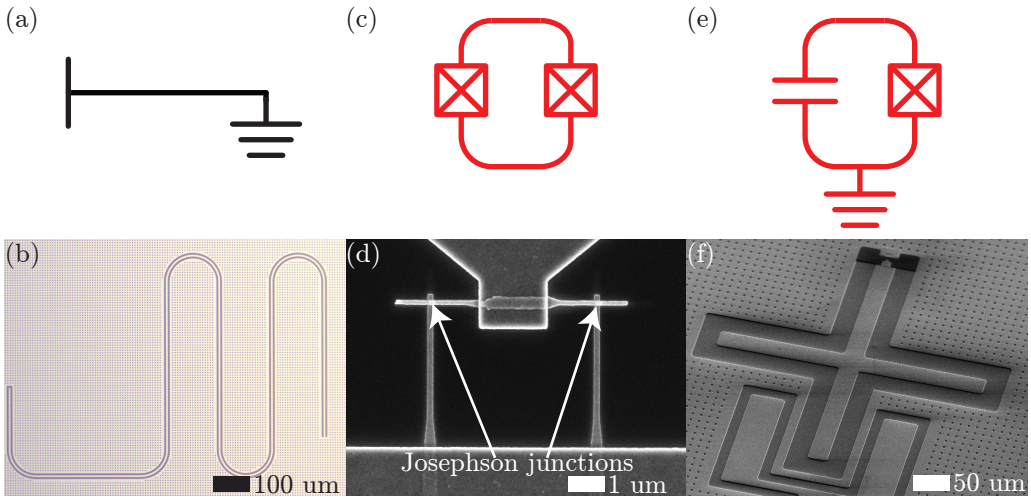

Figure 1.1: *Three building blocks of circuit quantum electrodynamics. The top row shows the equivalent electrical circuits, and the bottom row shows images of fabricated devices. (a-b) A distributed, quarter wavelength, microwave resonator. In practice, the resonator is meandered to reduce its footprint. (c-d) A superconducting quantum interference device (SQUID) consisting of two Josephson junctions in parallel. (e-f) A transmon qubit. The x-shape forms the capacitor while the Josephson junction is located (but not visible) in the darker area near the top.*

frequencies. Usually, one of the resonators' ends is used to couple the resonator to another transmission line with a strength κ, which is used to drive and measure the resonator. Refs. [51–53] cover the derivation of many resonator properties (e.g., resonant frequencies and coupling strengths) as a function of their geometry.

In addition to being a useful element for cQED, resonators themselves can be used to study physics. For instance, they are valuable tools for probing surface defects (e.g., magnetic spins [54] and quasiparticle recombination time [55]). Resonators are also used as high-sensitivity detectors for radiation ranging from infrared to x-rays [56].

1.2.2 The transmon qubit

Superconducting quantum bits (qubits) have seen tremendous development since their first demonstration in 1999 [57]. Their coherence times, the gate and readout fidelities, and the number of coupled qubits have all improved by orders of magnitude. However, there is still a long way to go before a large-scale universal quantum computer can be demonstrated. The largest improvement has been in the coherence time, which is now in the 100 µs range, compared to a few nanoseconds in the first qubit. With that problem under control, for now, the focus has shifted to scaling up the systems from just a few qubits to tens or hundreds of qubits. The currently largest, and published, superconducting quantum processor consists of 53 qubits [16]. A significant challenge is to maintain the long coherence times while scaling up to more qubits due to the complex 3D geometries needed to provide control access to all qubits [42].

There are many different types of superconducting qubits [31, 35]. Almost all superconducting qubits have in common that their transition frequencies are in the low GHz range and that they utilize the Josephson nonlinearity in some way. One of the most popular qubits is the transmon [58], which is a capacitively shunted Josephson junction, see Fig. 1.1 (e-f). The transmon can be modeled as a highly nonlinear LC resonator. Since the inductance of the Josephson junction is nonlinear, it introduces anharmonicity in the spectrum of the LC circuit, see Fig. 1.2 (a). If the anharmonicity is large enough, the two lowest energy levels of the oscillator, denoted $|0\rangle$ and $|1\rangle$, can be isolated and used as a qubit. For the transmon, the anharmonicity is negative (the energy difference between $|1\rangle$ and $|2\rangle$ is lower than the one between $|0\rangle$ and $|1\rangle$) and its magnitude is typically 200–300 MHz.

The transmon is characterized by two energies, the Josephson energy E_J and charging energy E_C. The Josephson energy is related to the critical current (and inductance) of the Josephson junction,

$$E_\mathrm{J} = \frac{\Phi_0 I_\mathrm{c}}{2\pi}. \tag{1.4}$$

The charging energy is related to the sum of all capacitances in the system, C_Σ,

$$E_\mathrm{C} = \frac{e^2}{2C_\Sigma}. \tag{1.5}$$

$E_\mathrm{J}/E_\mathrm{C} \gg 1$ is required for being in the so-called transmon regime, where the qubit frequency is insensitive to electric charge [58]. The sensitivity to electric charge is known as the charge dispersion and is exponentially suppressed in $E_\mathrm{J}/E_\mathrm{C}$. A low charge dispersion

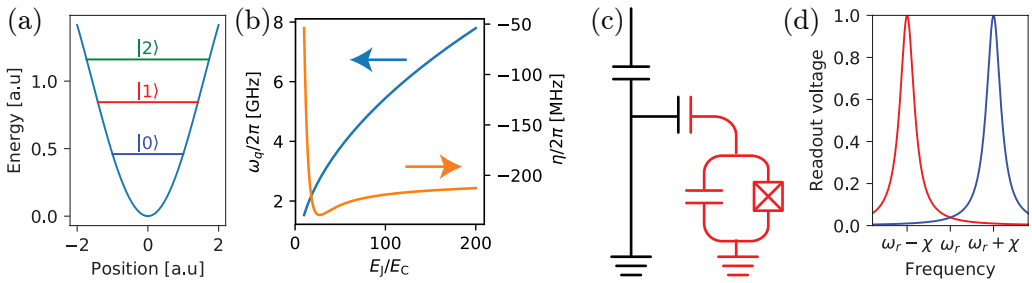

Figure 1.2: *(a) Illustration of the transmon spectrum including the three lowest energy levels. Note that the distance between* $|1\rangle$ *and* $|2\rangle$ *is smaller than the distance between* $|0\rangle$ *and* $|1\rangle$*. (b) Transmon frequency and anharmonicity as a function of* E_J/E_C *with* $E_C/2\pi = 200\,\text{MHz}$ *(c) A transmon qubit (red) coupled to a microwave resonator (black). (d) Readout voltage as a function of frequency. The dispersive shift* χ *moves the resonant frequency depending on the qubit's state.*

is needed to achieve a long coherence time as inevitable charge noise would otherwise cause dephasing. Noise, relaxation, and dephasing will be discussed in more detail in Chapter 2.

Far into the transmon regime, the transition frequency between the ground and the excited state is

$$\omega_q = \frac{\sqrt{8E_J E_C} - E_C}{\hbar}, \tag{1.6}$$

and the anharmonicity is

$$\eta = -\frac{E_C}{\hbar}. \tag{1.7}$$

However, for finite and practical E_J/E_C ratios, numerical solutions of the full transmon Hamiltonian in a phase basis using Mathieu functions are typically needed for accurate estimations of ω_q, η, and the charge dispersion [59]. In Fig. 1.2 (b), we plot ω_q and η as a function of E_J/E_C with $E_C/2\pi = 200\,\text{MHz}$. For the qubit frequency, Eq. (1.6) has an error smaller than 1 % and is not visible on the scale of this plot. However, for η, Eq. (1.7) predicts a constant anharmonicty of $-2\pi \times 200\,\text{MHz}$. which is an error of 10–20 % for $30 < E_J/E_C < 200$.

1.2.3 A coupled system of resonators and qubits

In the circuit quantum electrodynamics architecture, a qubit is strongly coupled to one or several resonators [60], as shown in Fig. 1.2 (c). The resonators allow for fast and accurate state readout, qubit–qubit coupling, and gives protection from the environment. A two-level system (qubit) coupled to a harmonic oscillator (microwave resonator) is typically described by the Jaynes-Cummings Hamiltonian, which after a rotating-wave approximation equals to [37]

$$\hat{H}/\hbar = \omega_r \hat{a}^\dagger \hat{a} + \frac{\omega_q}{2}\hat{\sigma}^z + g\left(\hat{\sigma}^+ \hat{a} + \hat{\sigma}^- \hat{a}^\dagger\right), \tag{1.8}$$

where $\hat{a}$ ($\hat{a}^\dagger$) is the annihilation (creation) operator for the harmonic oscillator, $\hat{\sigma}^\pm$ are the raising and lowering operators for the qubit, and g is the coupling strength.

To enable measurements of the qubit's state, the dispersive coupling regime is typically used. This regime is valid when $|\Delta| \gg g$, where $\Delta = \omega_a - \omega_r$ is the qubit–resonator detuning. In this limit, the Jaynes-Cummings Hamiltonian can be rewritten as

$$\hat{H}_{\text{disp}}/\hbar = (\omega_r + \chi\hat{\sigma}^z)\,\hat{a}^\dagger\hat{a} + \frac{1}{2}\left(\omega_q + \frac{g^2}{\Delta}\right)\hat{\sigma}^z. \tag{1.9}$$

For a true two-level system, $\chi = g^2/\Delta$. However, for the transmon with its rather weak anharmonicity, the higher levels also contribute [58],

$$\chi = \frac{g^2}{\Delta}\frac{\eta}{\Delta+\eta}. \tag{1.10}$$

χ is known as the dispersive shift, and its effect is that the resonator frequency depends on the qubit state, as illustrated in Fig. 1.2 (d). Therefore, the dispersive shift can be used to readout the state of the qubit.

In general, Eq. (1.8) is not enough to describe the full behavior of an implementation of a qubit. The environment interacts with the qubit and causes, for example, relaxation back into the ground state or thermal excitation to higher levels. It is challenging to model the microscopics of the environment due to the many degrees of freedom involved. Instead, statistical properties of the qubit–environment interaction are used. These include the qubit's effective temperature T and relaxation time T_1. If the qubit is initialized to its excited state, it will relax back into its ground state with the characteristic time scale T_1, which is defined as the time where the probability of finding the qubit in its excited state is equal to $1/e$, in the case of exponential decay.

1.2.4 Frequency tunability and nonlinearities

The resonator and qubit described so far are fixed in frequency, which in some applications is sufficient. However, in this work, resonators with a tunable resonant frequency are needed. To achieve fast frequency tunability of resonators, a SQUID is placed at the grounding point of the resonator, where the current is at its maximum [45]. The inductance of the SQUID gives an additional phase shift, which is equivalent to an increase in the resonator's electrical length. The resonant frequency ω_r of the SQUID-terminated resonator can be approximated by [61]

$$\omega_r(\Phi) \approx \frac{\omega_{\lambda/4}}{1 + \gamma_{\text{l}}/|\cos(\pi\Phi/\Phi_0)|}, \tag{1.11}$$

where $\omega_{\lambda/4}$ is the resonant frequency without the SQUID, and γ_{l} is the inductive participation ratio defined as the ratio between the SQUID and bare resonator inductances.

Incorporation of a SQUID into the resonator does not only enable frequency modulation, but it also adds nonlinearity to the system due to the dependence between the current through and phase across the junction. A Taylor expansion of Eq. (1.2) yields

$$I \approx I_{\text{c}}(\varphi - \varphi^3/6). \tag{1.12}$$

The cubic term is known as the Duffing term, or Kerr term in optics. This nonlinearity causes a power-dependent frequency shift of the resonator, which can be quantified by α_n, the amount of negative frequency shift per photon in the resonator [62]. Nonlinear dynamical systems are popular in textbooks, where effects such as harmonic generation, bifurcation, and chaos are studied [63, 64].

The transmon qubit can also be made frequency tunable by replacing its single Josephson junction with a SQUID.

1.3 Continuous-variable quantum computing

Another architecture for quantum computing is based on continuous variables (CV) where harmonic oscillators (resonators) replace the qubits [65, 66]. Harmonic oscillators have an infinite Hilbert (state) space, in contrast to qubits, which have two energy levels. However, a harmonic oscillator is usually thought of as a classical system. So to do quantum information processing, nonclassical states of the harmonic oscillator are created, such as superpositions of Fock states [67], entangled two-mode squeezed states [68–70], and multi-photon cat states [71]. In the creation of these states, parametric phenomena play a prominent role.

While this thesis includes results related to parametric pumping of quantum states that could be useful for CV quantum computing, we do not utilize them in that context.

1.3.1 Parametric effects

Parametric phenomena have a long history in a vast number of areas. A common, everyday example is that of a child standing on a swing and increasing the swing amplitude by moving up and down with twice the frequency of the swing. Parametric pumping is the modulation (pumping) of one system parameter, for example, the frequency or nonlinearity of a resonator. If done at certain frequencies, the parametric pump can provide amplification, damping, or frequency conversion. In the quantum regime, it can also be used to create entanglement between optical, microwave, or mechanical modes. Parametric amplifiers also add the lowest amount of noise possible and are vital to measurements of superconducting qubits and gravitational waves.

There are a few concepts related to parametric pumping that are important to know. First, energy must, of course, be conserved. In the case of 'degenerate' amplification of a signal via three-wave-mixing, the pump has an energy $\hbar\omega_p$, and the signal $\hbar\omega_s$. To conserve energy, $\omega_p = 2\omega_s$. In a quantum picture, one pump photon gets split into two signal photons with half the energy. In the nondegenerate case, the signal is slightly detuned from $\omega_p/2$, which implies that an idler is created at ω_i to conserve energy, $\omega_p = \omega_s + \omega_i$.

There is also four-wave-mixing, where the pump frequency is close to the signal frequency, and $2\omega_p = \omega_s + \omega_i$. Three- and four-wave refers to how many photons are involved in the parametric process.

Parametric pumping can also perform frequency conversion between two modes of different frequencies. If done via three-wave-mixing, $\omega_p = \omega_1 - \omega_2$, where $\omega_{1/2}$ are the frequencies of the respective modes.

Superconducting circuits can implement all of the above parametric phenomena. Typically, the parametric modulation is achieved by modulating the flux through the SQUID loop of a frequency tunable microwave resonator. A microwave transmission line is placed close to the SQUID so that the current through the transmission line induces a magnetic field around itself and into the SQUID loop.

Parametric pumping refers to the pumping (modulation) of one of the parameters of a system. In the case of a resonator, there are two possibilities, either to modulate the damping or the resonant frequency. When placing a SQUID at the current antinode of the resonator, its resonant frequency will depend on the magnetic flux through the SQUID. By modulating the magnetic field with a sinusoidal signal, and therefore the resonant frequency, different parametric effects such as amplification [72] and frequency conversion [73] can be realized.

In this work, a phenomenon knows as parametric oscillation is studied and utilized. By frequency modulating a resonator at twice its resonant frequency, and with a large amplitude, the ground state of the resonator becomes unstable. Therefore, the resonator starts to oscillate at its resonant frequency. These oscillations are said to be self-sustained since there is no external drive of the system at the resonant frequency. Parametric oscillations have been observed in a variety of systems, such as optical and microwave cavities [74, 75]. We demonstrate a different regime by using a multimode resonator and a frequency modulation at the sum of two resonant frequencies. This regime is referred to as nondegenerate parametric oscillations [76, 77].

1.4 Measurement techniques

Superconducting qubits and devices need to operate at cryogenic temperatures. The reason is two-fold: the critical temperature of the most commonly used superconductors is on the order of 1 to 10 Kelvin, and for resonators and qubits to occupy their ground states, their energy $\hbar\omega$ must be much less than the thermal energy $k_{\mathrm{B}}T$. Most superconducting qubits are partly made out of aluminum, which has a critical temperature of 1.2 K, while the temperature corresponding to the qubit energy is around 200 mK. Therefore, it is desirable to operate the device below 20 mK.

Commercially available dilution refrigerators reach temperatures below 10 mK. Cooling to a temperature of 3 K is done using the second stage of a pulse-tube cryocooler [78, 79]. Then, a dilution unit in the refrigerator achieves temperatures below 10 mK via a transfer of helium-3 through a phase boundary in a mixture of helium-3 and helium-4 [80].

The electrical measurement setup can be divided into two parts: inside and outside of the cryostat. The cryostat is equipped with a set of coaxial cables for transmitting and receiving microwave signals to and from devices. The input signals are attenuated on their way down from room temperature to 10 mK. This serves three purposes: first, it thermally anchors the lines to the different stages of the cryostat. Secondly, it yields the appropriate signal powers at the resonator input. Finally, it also attenuates the thermal noise generated at the different temperature stages inside the cryostat.

The setup outside of the cryostat consists of various microwave instruments. For instance, signal generators provide monochromatic microwave signals in the GHz range,

and vector network analyzers are used to measure scattering matrices. Due to the complexity of the measurements, such as the amount of data and instruments involved, computers are needed for synchronous hardware control and data logging. This work uses the software Labber for all instrument control, measurement automation, and data logging[1]. Each physical instrument has its own software instrument driver responsible for translating from a general command to an instrument-specific command. For example, the user might request that the frequency of a signal generator be set to 5 GHz. The driver then translates this to a command that the instrument understands. Most instruments are controlled via the standard commands for programmable instruments (SCPI); however, some instruments have their own APIs. Most instrument drivers are found on the Labber driver repository[2].

1.5 Thesis outline

This thesis serves both as an introduction to, and a summary of, 13 journal articles in the subject area of superconducting circuits and quantum information. The articles can be grouped into three main categories with extensive overlap. Papers A to D concern the fabrication and characterization of superconducting devices. Paper E is about a readout technique for quantum bits. Paper F and Paper G show experimental implementations of two quantum information algorithms. Finally, we look at how superconducting circuits can be used to create different microwave photon states in Papers H to M.

This thesis will not include all the results of all the articles. Instead, all papers are appended at the end and cited as needed in the thesis. The thesis focuses on some of the background and experimental details omitted from the papers, as well as on putting the appended papers in a broader context.

[1] See http://labber.org for more information.
[2] See https://github.com/Labber-software/Drivers

CHAPTER II

High-coherence superconducting circuits

This chapter describes some of the different loss and dephasing mechanisms in superconducting circuits. This knowledge is required for understanding how to best design circuits and their measurement environment. It is also important to fabricate the devices with as few defects and lossy interfaces as possible. One section of this chapter is dedicated to explaining our fabrication procedure. Finally, this chapter provides a summary of the results in Papers A to D, which deal with the characterization of losses and noise in a few different systems.

High-coherence circuits are needed to be able to store quantum information in the circuits, as decoherence is equivalent to the loss of information. In quantum computing, information is stored in the complex state (amplitude and phase) of a qubit. Therefore, both amplitude and phase need to be preserved for as long as possible. Loss of amplitude is usually referred to as relaxation and can be caused by electrical dissipation in the case of superconducting qubits. The time with which the qubit relaxes is denoted by T_1. In the context of quantum information, relaxation can be seen as a probability of having the qubit state flip from $|1\rangle$ to $|0\rangle$.

On the other hand, a qubit's phase is lost by uncertainty in the qubit frequency. The time on which this happens is known as the dephasing time T_ϕ. Analogously to bit flips, dephasing can be seen as a probability of having a phase flip between $|1\rangle$ and $-|1\rangle$. The combination of relaxation and dephasing is what causes decoherence and loss of quantum information, which happens with the decoherence time T_2, related to T_1 and T_ϕ by

$$\frac{1}{T_2} = \frac{1}{2T_1} + \frac{1}{T_\phi}. \tag{2.1}$$

There are two ways of extending T_2. Either by reducing the mechanisms responsible for relaxation and dephasing or by the use of error-correcting codes [81]. Quantum error correction (QEC) is a research topic by itself and is not featured in this work. Basically, QEC relies on encoding the information of one logical qubit into many physical qubits. By some measurement, errors can be detected and subsequently corrected for, resulting in

a longer coherence time of the logical qubit compared to each physical qubit. In this work, we are focused on identifying and reducing the relaxation and dephasing mechanisms, which allow us to realize the high-coherence qubits needed to enable QEC and large-scale quantum computers in the future.

2.1 Loss in superconducting devices

Our circuits are made out of superconductors for minimal loss due to electrical resistance. While superconductors have zero resistance for DC currents, they are not necessarily lossless at higher frequencies. For instance, the presence of quasiparticles (unpaired electrons) or magnetic vortices (non-superconducting regions) in the superconductor can induce loss for alternating currents. Additionally, there is loss due to electromagnetic interaction with parasitic two-level systems (TLS) that can absorb and dissipate energy, and due to radiation into other electromagnetic modes present in the device. The dominating loss mechanism for a given circuit mainly depends on its geometry and environment; however, it is not trivial to experimentally determine the dominating mechanism. This can make it challenging to improve circuits, as changes to a non-dominant mechanism cannot be observed without collecting a significant amount of statistics.

Over the last two decades, superconducting resonators and qubits have increased their relaxation times into the 100 µs range, corresponding to quality factors well above one million [82]. Achieving these long lifetimes was enabled via improved circuit design [58, 83, 84], reduced loss to quasiparticles and magnetic vortices by improved shielding [85, 86], and improved materials with fewer TLS [87, 88]. In Paper A and Paper B, we demonstrated that the loss in our resonators and qubits, are mainly due to TLS.

For a qubit, we define the relaxation rate to its steady-state,

$$\Gamma_1 = \frac{1}{T_1} = \Gamma_{1\to0} + \Gamma_{0\to1}, \tag{2.2}$$

where $\Gamma_{1\to0}$ is the transition rate from excited to ground state and $\Gamma_{0\to1}$ from ground to excited state. At low temperatures ($T \ll \hbar\omega/k_B$), $\Gamma_{0\to1}$ is negligible. We typically observe that our qubits have a ground state population above 99%, consistent with $\Gamma_{0\to1} \ll \Gamma_{1\to0}$

Unwanted qubit-state transitions are due to noise or signal in some parameter λ, which can be calculated via Fermi's golden rule [89],

$$\Gamma_{1\to0} = \frac{1}{\hbar^2}\sum_{\lambda}\left|\langle 1|\frac{\mathrm{d}\hat{H}_\lambda}{\mathrm{d}\lambda}0\rangle\right|^2 S_\lambda(\omega_q), \tag{2.3}$$

where $S_\lambda(\omega_q)$ is the spectral density at the qubit frequency, and the sum goes over all possible parameters. To decrease the decay rate, we can either decrease the sensitivity to the noise (the matrix element) or the noise itself. Some matrix elements cannot decrease too much since we need them to control the qubit itself. One example of reducing a matrix element can be seen with the evolution from the single Cooper-pair box into the transmon. The transmon exponentially suppresses the qubit's sensitivity to electric charge, implying that decoherence due to charge noise is also suppressed [58].

Every possible loss mechanism contributes linearly to the total relaxation rate. For instance, we could have a combination of TLS, quasiparticle, magnetic and radiative (Purcell) loss,

$$\Gamma_1 = \Gamma_{\mathrm{TLS}}(P,T) + \Gamma_{\mathrm{qp}}(T) + \Gamma_{\mathrm{mag}}(B) + \Gamma_{\mathrm{p}}(T), \tag{2.4}$$

where each mechanism has some dependence on the external parameters P (circulating power in the resonant circuit), T (temperature), and B (magnetic field). As long as the external parameter dependence is not identical for two or more mechanisms, we can change the external parameters to extract the rate for each mechanism. For instance, we can measure the loss of a resonator as a function of P to extract Γ_{TLS} as it is the only mechanism with a power dependence. Typically, we perform experiments where we measure loss as a function of power and temperature.

2.1.1 Radiative loss

A qubit coupled directly to a wideband transmission line would spontaneously emit into said transmission line with a rate Γ_{r}, where Γ_{r} is the coupling strength between qubit and transmission line. To prevent this rapid decay, we use a frequency-detuned readout resonator as a filter between the qubit and the transmission line, which results in a reduced effective decay rate Γ_{p}. This rate is known as the Purcell rate [90]. Another way to look at the Purcell decay is that noise in the transmission-line induces qubit relaxation via Eq. (2.3). The resonator acts as a filter and reduces the noise at the qubit frequency, therefore reducing the decay rate.

The Purcell rate depends on the inter-system couplings (g and κ) and frequency detuning ($\Delta = \omega_q - \omega_r$). For a transmon qubit, a common formula for the Purcell rate (at zero temperature) is [58],

$$\Gamma_{\mathrm{p}} = \kappa \frac{g^2}{\Delta^2}. \tag{2.5}$$

To account for finite temperature, Eq. (2.5) should be multiplied by $\coth(\hbar\omega_q/2k_{\mathrm{B}}T)$ [89], meaning that the Purcell rate increases with temperature. This increase is understood as stimulated emission due to increased thermal noise in the transmission line.

However, Eq. (2.5) does not tell the full story. There is a frequency correction factor ω_q/ω_r [91], and g is not detuning independent as it originates from a capacitive coupling capacitance C_g with a frequency-dependent impedance $1/j\omega C_g$. A modified Purcell rate expressed with circuit parameters is

$$\Gamma_{\mathrm{p}} = \kappa \frac{C_g^2}{4C_qC_r} \frac{\omega_q\omega_r}{(\omega_q - \omega_r)^2}. \tag{2.6}$$

Moreover, both Eq. (2.5) and Eq. (2.6) assume that the qubit is coupled to a single-mode resonator. However, our readout resonators are typically distributed coplanar waveguide resonators containing modes at harmonics of its fundamental resonant frequency. These higher frequency modes introduce an asymmetry in Δ for the Purcell rate, increasing the decay rate for positive detunings [83]. In addition, g and κ increase with the square-root and square of the mode number, respectively. These scalings with the mode number result in a significant contribution to the total Purcell rate from the higher modes [83].

A numerical investigation of a full circuit model taking the higher modes into account was performed in Ref. [83] and showed good agreement with experimentally observed relaxation times. With that said, we use $\Delta < 0$ to minimize the Purcell rate, and Eq. (2.6) is accurate enough in that regime.

If we cannot reach a low enough Purcell rate due to other constraints on g, κ, and Δ, we may add additional circuit elements to change the environment seen by the qubit to reduce the rate even further. For instance, a second resonator [91, 92] or an impedance transformer [93], can be added between the readout resonator and the transmission line to act as a so-called Purcell filter. Purcell filters are typically used when fast qubit readout is needed (i.e., large κ). In the case of using a second resonator as a Purcell filter, the filter is resonant with the readout resonator, which reduces the coupling between the qubit and transmission line (analogously to a second-order filter where its transfer function is squared, and a sharper cut-off is achieved). A Purcell filter with a quality factor Q_F and resonant with the readout resonator decreases the Purcell rate with a factor of $(\omega_r/2\Delta Q_F)^2$.

Finally, it is possible to couple a Purcell filter to more than one readout resonator, which is beneficial for frequency-multiplexed readout in a multi-qubit architecture. In that case, the bandwidth of the Purcell filter (ω_r/Q_F) should be wide enough so that the readout resonators can be spaced far enough apart in frequency to avoid any significant overlap between two or more resonators. In our work, only Paper F uses a Purcell filter, where it is coupled to three readout resonators.

2.1.2 Quasiparticle loss

Here, quasiparticles refer to unpaired electrons in a superconductor. They can cause a loss in both resonators and qubits by absorbing energy and then relaxing via a phonon, so it is important to minimize the number of quasiparticles as much as possible.

Just below the critical temperature of the superconductor, only a small amount of all electrons form Cooper-pairs, but as the temperature decreases, more and more electrons pair up. At zero temperature, all electrons are paired up (except one if the total number of electrons is odd). If the number of quasiparticles follows what is expected for a certain temperature, they are said to be equilibrium or thermal quasiparticles. However, the temperature of the quasiparticles might be elevated above the temperature of the cryostat (10 mK), sometimes referred to as cold nonequilibrium quasiparticles. For equilibrium and cold nonequilibrium quasiparticles, the loss rate for resonators and transmons is given by [94, 95]

$$\Gamma_{\mathrm{qp}} = \frac{L_k \omega_r}{L\pi} \sqrt{\frac{2\Delta_0}{\hbar\omega_r}} x_{\mathrm{qp}}, \tag{2.7}$$

where L_k is the kinetic inductance from the superconductor (which depends on the superconductor used and the geometry of the resonator), L is the total inductance, Δ_0 is the superconducting gap, and x_{qp} is the ratio between the density of quasiparticles and the density of Cooper-pairs. For transmons, $L_k = L$ and $\omega_r = \omega_q$. At low temperatures, we have that [95],

$$x_{\mathrm{qp}} = \frac{\sqrt{2\pi\Delta_0 k_{\mathrm{B}} T}}{\Delta_0} e^{-\Delta_0/k_{\mathrm{B}}T}. \tag{2.8}$$

Combining Eq. (2.7) and Eq. (2.8) shows that the loss to quasiparticles is linear in the number of quasiparticles, which decreases exponentially with temperature.

There are also hot nonequilibrium quasiparticles. These are believed to originate from high-energy photons. A photon with an energy larger than the superconducting gap can impinge the superconductor and break a Cooper pair into two quasiparticles. If the energy of the photon is much larger than the gap, it will also provide a large momentum to the quasiparticles, enough for the quasiparticles to stimulate emission and absorption of qubits. These quasiparticles will not follow a distribution given by a system in equilibrium with some temperature, and we cannot use Eq. (2.8). In general, we would need to integrate over the exact distribution of quasiparticles, which we do not know, to find x_{qp}.

The quasiparticle population does not grow with time since two quasiparticles can again pair-up and form a Cooper pair. This mechanism is known as recombination and happens with some timescale depending on the superconductor [96] and substrate [55] used. Therefore, x_{qp} depends on the balance between the rate that incoming photons and thermal fluctuations generate quasiparticles, and the rate with they are removed [94]. It is also possible to create traps for quasiparticles to reduce their population [97, 98].

2.1.3 Two-level systems

The term 'two-level systems' is a very general concept that refers to any microscopic system with two predominant states in which the system can switch between the two states. A TLS can, therefore, be characterized by the energy difference and switching rate between its two states. The TLS model was originally developed for amorphous solids and their low-temperature properties. The simplest TLS model, the standard tunneling model (STM), assumes an ensemble of TLS with a uniform energy distribution, and importantly, no TLS–TLS interaction.

The microscopic origin of TLS is not fully known and has been a long-lasting problem in condensed matter physics. Recently, atomic hydrogen [99] and molecular oxygen [100] have been shown to be two sources of TLS. For a TLS to absorb energy from a qubit, the TLS must couple to electrical fields (e.g., via an electric dipole moment). The coupling strength between a TLS and a qubit strongly depends on the physical location of the TLS. For instance, TLS on the surface of the capacitor couples less than a TLS inside the insulating barrier for a Josephson junction. For qubits with large or many Josephson junctions, the probability of having strongly coupled TLS in the junction barrier is higher than for a transmon with only one junction.

The electrical dipole moment of a TLS allows for tuning of the TLS with external electrical fields, which in turn can be used to infer the physical location of TLS [101, 102]. In Ref. [102], they found that 46 % of the TLS reside on the surface of the device.

For resonant circuits (e.g., qubits and resonators), TLS with similar resonant frequencies can exchange energy with the circuit. An excited TLS then relaxes by emitting a phonon, essentially removing energy from the circuit. Since a single TLS can only absorb a single energy quantum at a time, there is a maximum loss rate that the TLS ensemble can induce. By increasing the circulating power in the circuit, the TLS gets saturated, and the loss rate decreases. In superconducting resonators, we can probe their losses as a function of the circulating power in the resonator P. The STM predicts that the TLS

loss rate as a function of P is

$$\Gamma_{\mathrm{TLS}}(P,T) = \omega_r F \delta_{\mathrm{TLS}} \frac{\tanh \frac{\hbar\omega_r}{2k_B T}}{\left(1+\frac{P}{P_c}\right)^{\beta}}, \tag{2.9}$$

where F is a filling factor equal to the ratio of electric field in the TLS region to the total electric field, δ_{TLS} is the TLS loss tangent, and P_c is a critical power needed to saturate one TLS. The exponent β should be $1/2$ according to the STM; however, experimentally β is usually closer to 0.2 for superconducting resonators. There are various theories for explaining this discrepancy, such as spectral diffusion of strongly interacting TLS or the presence of different types of TLS ensembles with different P_c [103].

In the case of qubits, $P \ll P_c$ for typical P_c, thus Eq. (2.9) can be approximated as

$$\Gamma_{\mathrm{TLS}}(T) = \omega_q F \delta_{\mathrm{TLS}} \tanh \frac{\hbar\omega_r}{2k_B T}. \tag{2.10}$$

Therefore, we can use measurements of loss in superconducting resonators at low powers to estimate the expected loss for qubits with similar filling factors [104]. This is advantageous as measuring the loss of resonators is simpler and faster than measuring the loss of qubits.

Equation (2.9) predicts that the loss to TLS should decrease with increasing temperatures, as a higher temperature can saturate the TLS in a similar way to high circulating power. This effect has been observed in resonators [105]. We might, therefore, conclude that qubits should perform better at an elevated temperature; however, the temperature needed to observe a substantial decrease in TLS loss is above 100 mK. At that temperature, other qubit relaxation mechanisms would be increased, effectively nullifying the benefits of TLS saturation, see Fig. 2.1.

Finally, we compare the temperature dependence of TLS [Eq. (2.10)] to quasiparticle [Eq. (2.7)] and radiative loss [Eq. (2.5)], see Fig. 2.1. It is clear that at low temperatures, the loss due to quasiparticles is many orders of magnitude smaller than for both TLS and Purcell. However, the quasiparticles loss is calculated assuming only equilibrium quasiparticles, which is generally not the case. At low temperatures where qubits typically operate, the loss is dominated by TLS. For the calculations in Fig. 2.1, we assumed aluminum as the superconductor ($\Delta_0/h = 44\,\mathrm{GHz}$), $F\delta_{\mathrm{TLS}} = 5 \times 10^{-}7$, $\omega_q/2\pi = 4\,\mathrm{GHz}$, and a Purcell rate at zero temperature of $1/\Gamma_{\mathrm{p}} = 300\,\mu\mathrm{s}$.

2.2 Qubit dephasing

Energy loss is not the only decoherence mechanism for superconducting qubits, see Eq. (2.1). If the frequency of a qubit is not stable (hint, it is not), the qubit state will evolve in time with a non-deterministic phase. This phase gives a probability of having a phase flip of the qubit state. The larger the frequency fluctuations are, and the longer the qubit state evolves under these fluctuations, the larger the phase error becomes. In the long evolution time limit, all phase information is lost, and the qubit is in a fully mixed state.

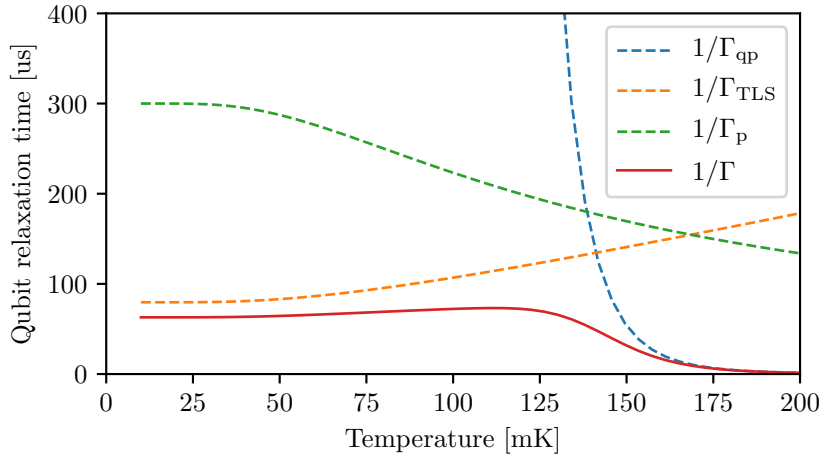

Figure 2.1: *The temperature dependencies of energy loss due to quasiparticles, two-level systems, and radiation (Purcell). The solid line corresponds to the combined loss. At low temperatures, TLS dominate, while at high temperatures, the loss is dominated by quasiparticles.*

As for loss, dephasing can be related to a noise spectral density and a matrix element. In contrast to Eq. (2.3), the dephasing rate depends on the noise within a window of frequencies, resulting in a strong dependence on the shape of the noise spectrum. The exact window depends on the pulse sequence used to measure dephasing, and a qubit may, therefore, be used to measure the noise spectrum itself [106]. For instance, flux and charge noise usually have $1/f$ (pink) spectra, whereas thermal resonator noise gives a constant spectrum (white). In the case of A/f noise, where A is the frequency noise amplitude at 1 Hz, the resulting dephasing rate is equal to A [58].

Historically, three primary sources of dephasing were considered for qubits based on Josephson junctions: charge, flux, and critical-current noise. To a great extent, charge noise was mitigated by the invention of the transmon qubit, which made the qubit frequency much less sensitive to charge [58]. Critical-current noise is interesting since it was estimated that transmons should be limited to dephasing times of 35 µs [58], which is not the case as dephasing times far exceeding that has been demonstrated, as we will see later. This must imply that the critical-current noise in modern circuits is much less than what has been previously reported in the literature [107].

Frequency-tunable circuits are common in circuit-QED. This tunability is typically achieved by tuning the magnetic flux through a SQUID acting as a tunable inductance. Any frequency tunability necessarily induces frequency fluctuations due to fluctuations in the control parameter. For circuits with SQUIDs, dephasing is caused by flux noise, a well-researched topic due to the application of SQUIDs as detectors [44]. Typically, we distinguish between flux noise in the macroscopic environment surrounding the device and the microscopic environment on the surface of the SQUID itself. By proper filtering and shielding, the macroscopic flux noise can be reduced to a level well below the microscopic.

It is noteworthy that the level of microscopic flux noise is very similar for both detectors and qubits, implying that the two communities can benefit from each other and collaborate on reducing flux noise.

One simple solution to mitigate the contribution of flux noise to qubit dephasing is to remove the SQUID and use a single Josephson junction, which removes the qubit-frequency tunability. In Papers B and G, we use these fixed-frequency qubits to achieve high coherence. As we will see in later chapters, we introduce tunable couplers to overcome some of the limitations of having fixed-frequency qubits. These couplers do not have the same requirement of high coherence and can have SQUIDs even in the presence of flux noise. If the flux tunability and its associated noise is an issue, an asymmetric SQUID can be used to reduce the sensitivity to flux noise [108].

A sometimes overlooked dephasing mechanism is capacitance noise, which can be caused by interacting TLS. Such noise has been observed in planar superconducting resonators, as in Paper A and Ref. [109], and there is no fundamental reason why transmons should be protected from this type of noise. Further investigations are needed into how capacitance noise could limit transmon qubits.

An additional dephasing mechanism is varying photon numbers in, for instance, readout resonators. A residual photon population could originate from previous measurements (coherent states) or nonideal thermalization of the resonator (thermal states). Both thermal and coherent states contain superpositions of different numbers of photons, implying that there is an uncertainty in the number of photons in the resonator. This uncertainty yields frequency fluctuations of qubits coupled to the resonator through the AC-stark shift. The effective dephasing rate from a thermal state with a small population $\bar{n} \ll 1$ can be expressed as [110],

$$\Gamma_\phi^{\mathrm{th}} = \frac{\kappa^2}{\kappa^2 + 4\chi^2} \frac{4\chi^2}{\kappa} \bar{n}. \tag{2.11}$$

Finally, fluctuations in qubit frequencies could be confused with frequency fluctuations in the signals used to drive the qubits. However, all microwave equipment can be phase-locked to a common atomic clock, yielding frequency stabilities far greater than the stability of any current superconducting qubit frequency [111].

2.3 Qubit and resonator design

From our understanding of loss and dephasing mechanisms, we set out to design practical circuits optimized for low loss and dephasing. Here, practical means that we can measure and control the qubits with high fidelity in a reasonable time. For instance, we need $1/\kappa \approx 500\,\mathrm{ns}$ to achieve a sufficiently fast qubit readout and an anharmonicity larger than 200 MHz to be able to drive sufficiently fast qubit gates without leaking to higher qubit states. Another practical consideration is the physical size of the qubit. A physically large capacitor can reduce the filling factor F in Eq. (2.10), effectively decreasing loss due to TLS [87]. However, since the ultimate goal is to create a device with multiple qubits on the same chip, there is a trade-off between the filling factor and the physical size of the chip. The first transmons used compact interdigitated capacitors with an approximate

size of 0.005 mm^2 [112]. While they were physically small, their filling factors were high, yielding relaxation times of $T_1 \approx 2\,\mu s$ [113]. By increasing the finger width of the capacitor (and reducing the number of fingers to keep the capacitance constant), relaxation times of up to 18 µs was observed in Ref. [114] using similar materials as Ref. [113]. In the limit of only two fingers, the capacitor consists of two large rectangular plates. In this limit, using a 0.25 mm^2 capacitor, Paik et al. [84] demonstrated $T_1 \approx 60\,\mu s$. Later on, Barends et al. [115], developed the 'xmon', a cross-shaped transmon suitable for scaling up to a large 2D-grid of qubits while having a reduced filling factor, demonstrating $T_1 \approx 44\,\mu s$ with an area of 0.05 mm^2.

To operate in the transmon regime, we need a Josephson energy much greater than the charging energy, $E_J/E_C \gg 1$; however, it is not practical to go too deep into the transmon regime. There are three main trade-offs for E_J and E_C: lower qubit frequency yields higher coherence; greater anharmonicity allows for faster qubit operations; and larger E_J/E_C reduces the qubit frequency sensitivity to charge. These three criteria cannot be fulfilled at the same time. A good trade-off is $E_J = 12\,GHz$ and $E_C = 180\,MHz$, which gives $E_J/Ec = 67$, a qubit frequency around 4 GHz, an anharmonicity of $-200\,MHz$, and a charge dispersion less than 400 Hz. With these parameters, we typically achieve relaxation times $T_1 \approx 60\,\mu s$, decoherence times $T_2 \approx 100\,\mu s$, and operations can safely be run in 20 ns.

With the qubit frequencies fixed, we proceed to the design of the resonator, which includes its frequency ω_r (and the detuning $\Delta = \omega_q - \omega_r$) and the two coupling strengths g and κ. Here, we trade-off the speed of the readout $(1/\kappa)$ to the Purcell rate (Eq. (2.6)), while also maintaining a dispersive shift χ similar to κ. For $\kappa \approx \chi$, the dephasing due to thermal photons in the resonator, Eq. (2.11), simplifies to $\Gamma_\phi^{th} \approx \kappa\bar{n}$. Even though a cryogenic refrigerator cools the devices to below 10 mK, $\bar{n}$ will not thermalize to the value expected from a Boltzmann distribution at 10 mK, $\bar{n} < 10^{-12}$. Instead, an elevated effective temperature is usually observed, yielding $\bar{n}$ on the order of 10^{-3} [116]. There is therefore a trade-off between the speed of the qubit readout $(1/\kappa)$ and dephasing time $(1000/\kappa)$.

We summarize our design constraints in Table 2.1 and give equations for converting between system and circuit parameters in Table 2.2. A typical device with one resonator and one qubit is seen in Fig. 2.2 (a), with the resonator–transmission line coupling in Fig. 2.2 (b), and the resonator–qubit coupling along with the qubit in Fig. 2.2 (c). The qubit is of the xmon style [115].

One of the significant advantages of transmon qubits is their close resemblance to lumped-element LC resonators. By replacing the Josephson junction with a linear inductance, we may simulate the transmon's parameters and its response to external signals using linear electromagnetic field solvers, such as COMSOL and HFSS. Importantly, all involved capacitances can be simulated using electrostatics, vastly reducing the simulation time.

Superconductors in the presence of magnetic fields can exhibit loss due to the movement of so-called Abrikosov vortices [26, 117]. In such a vortex, superconductivity is suppressed, and the vortex can interact via a Lorentz force with currents flowing through the superconductor. An oscillating current, as in our microwave circuits, can, therefore, dissipate through the motion of these vortices. Therefore, we introduce flux-trapping

Table 2.1: Design constraints for a high-coherence and practical circuit-QED system.

Constraint	Motivation
$\omega_q/2\pi < 5\,\mathrm{GHz}$	Lower qubit frequency implies lower Γ_{TLS}
$\eta/2\pi < -200\,\mathrm{MHz}$	Minimize leakage to $\|2\rangle$
$E_{\mathrm{J}}/E_{\mathrm{C}} > 60$	Low enough charge dispersion
$\omega_q < \omega_r$	The Purcell rate is lower for negative detunings
$1/\kappa < 1\,\mu\mathrm{s}$	Minimize relaxation during readout
$\chi > \kappa$	Maximize readout contrast
$1/\Gamma_p > 150\,\mu\mathrm{s}$	Not to be limited by Purcell decay
$1/\Gamma_\phi^{\mathrm{th}} > 200\,\mu\mathrm{s}$	Not to be limited by thermal photons in the resonator

Table 2.2: Equations for converting between circuit and system parameters. C_q and C_r are the qubit and resonator capacitances, C_g is the coupling capacitance between resonator and qubit, and C_κ is the coupling capacitance between resonator and transmission line. Z_r is the characteristic impedance of the resonator and Z_e is the environmental impedance seen by the resonator.

Parameter	Comment
$E_{\mathrm{C}} = e^2/2C_q$	Charging energy
$E_{\mathrm{J}} = \Phi_0 I_{\mathrm{c}}/2\pi$	Josephson energy
$g = C_g\sqrt{\omega_q\omega_r}/(2\sqrt{C_qC_r})$	Capacitive coupling between two resonant circuits
$\omega_q = \sqrt{8E_{\mathrm{J}}E_{\mathrm{C}}} - E_{\mathrm{C}}$	Qubit frequency
$\omega_r = \pi c/(2l\sqrt{\epsilon_{\mathrm{eff}}})$	Quarter-wavelength resonator frequency
$\epsilon_{\mathrm{eff}} = (1+\epsilon_{\mathrm{r}})/2$	Effective dielectric constant. For silicon, $\epsilon_r = 11.7$
$C_r = \pi/4\omega_r Z_r$	Effective capacitance for a quarter-wavelength resonator
$\kappa = \omega_r^2 C_\kappa^2 Z_e/C_r$	Coupling between a resonator and a transmission line

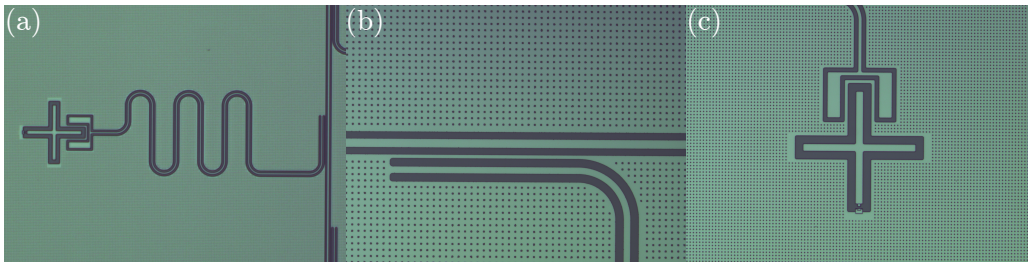

Figure 2.2: *(a) A coplanar waveguide resonator coupled to a transmon qubit. (b) Inductive resonator to transmission line coupling. Visible is also the flux-trapping holes. (c) Resonator to qubit coupling. The ground plane extends all around the qubit, ensuring a good microwave environment.*

holes in our superconducting ground planes, which allow vortices to form without creating dissipative normal metal cores [118]. It is important to have these holes close to the qubits. However, there is a risk of introducing additional TLS loss due to the exposed substrate in the flux-trapping hole. Chiaro et al. [119] investigated the loss dependence as a function of the distance between the holes and resonators. The flux-trapping holes used in our work are 2 µm in diameter with a minimum distance of 8 µm to any microwave structures, see Fig. 2.2 (b).

2.4 Fabrication techniques

Concerning qubit coherence, there are a few important things to consider when fabricating devices: use substrates with low dielectric loss, minimize the amount of surface oxide [120] and resist residues [121], and avoid jagged transmission lines with sharp corners that can enhance local electric fields, hence inducing additional loss [122].

As a dielectric substrate, we use intrinsic high-resistivity (100) silicon. It is essential that the silicon is truly intrinsic (not counter-doped), so that, in practice, no charge carriers exist in the substrate that could induce loss. Another commonly used substrate is sapphire. Even though bulk sapphire has a lower dielectric loss than bulk silicon [123], the dominating surface losses bring the two substrates to comparable effective loss [87]. From a fabrication point of view, sapphire has one major disadvantage for scalable quantum circuits. Its hardness and chemical inertness make it almost impossible to etch; therefore, it is not feasible to trench, or create vias through, the sapphire substrate. Vias are needed to suppress spurious microwave modes when scaling to more qubits and larger chips [42].

The fabrication starts with an empty 2-inch silicon wafer. First, any organic residues are cleaned from the silicon by solvents and a hot mixture of hydrogen peroxide and ammonium hydroxide. We then submerge the wafer into a bath of hydrofluoric acid (HF). HF removes the native silicon surface oxide (SiO_2) that is known to host a large density of two-level systems. However, the oxide is regrown quickly if the bare silicon is exposed to air [105]. To prevent the regrowth of silicon dioxide, we load the wafer into the vacuum of a deposition system within minutes of the HF dip.

As the deposition system, we use an electron-beam evaporator equipped only with aluminum targets to minimize the probability of cross-contamination. To remove any residual water from the wafer, we immediately ramp up the temperature of an infrared heater until the silicon has reached a temperature of 300 °C. The wafer is then allowed to cool down to room temperature while the evaporation chamber is evacuated to its base pressure 4×10^{-8} mbar. We evaporate a 150 nm thin film at a rate of $1.0\,\mathrm{nm\,s^{-1}}$ and at room temperature. A scanning electron microscope (SEM) image of the aluminum film is seen in Fig. 2.3(a). The grains are clearly visible, as expected from our quite high deposition rate [124]. However, it is not clear if the grain size has any effect on the qubit coherence. The grain size and surface roughness of evaporated aluminum depend on the deposition rate and temperature [124, 125], and it would be interesting to study the relationship between these parameters and coherence in the future.

We proceed by patterning all microwave structures using a direct-write laser lithography system. The laser writer has a resolution below 1 µm that requires careful calibration of the laser intensity before every exposure. When combined with an almost omnidirectional etch process that widens any pattern, it becomes challenging to achieve sub-micron features routinely. Not to push the limit of the system, we keep all features above 2 µm. The aluminum is etched using a hot mixture of phosphoric, nitric, and acetic acids. Other groups have had success with reactive ion etching (RIE) of aluminum using chlorine-based chemistry [43]. Using RIE to etch aluminum also comes with the issue of chlorine residues that eventually form HCl and corrodes the aluminum. The chlorine residues may be removed by immediately rinsing the wafer in water after etching. We tried using RIE as well, but the measured quality factors of resonators were never as high as with the acid-based etch.

We have an optional step to reduce losses to TLS. After the aluminum etching, we may subject the wafer to RIE with a fluorine-based gas (e.g., NF_3 or SF_6). Fluorine will etch and remove silicon in the gaps of the coplanar waveguides and the qubit capacitor, so-called 'trenching.' By trenching the substrate, we move the lossy silicon surface further away from the electrical field of the device, effectively lowering the filling factor and the associated loss rate [105, 126]. The reason why we have this step as an optional one is that it adds a lithography step since we do not want to trench underneath Josephson junctions, and we have not yet conclusively proved that trenching improves the coherence of our circuits. The resonator in Paper A, qubit B in Paper B, the qubit in Paper C, and all resonators in Paper D use trenched substrates. Trenching also reduces the effective dielectric constant of the coplanar waveguides, implying higher impedances and lower capacitances that need to be accounted for when designing the circuits.

We realize a Josephson junction by patterning two narrow, orthogonal, and crossing lines by electron-beam lithography. The lines have widths between 100 and 200 nm, which is less than the resist thickness. If we evaporate aluminum at a 45° angle, any aluminum not evaporated along one of the lines will get deposited on the resist wall and therefore removed when the resist is later dissolved. This effect allows us to evaporate the first junction electrode along with one of the lines, do controlled oxidation to create the insulating part of the junction, and then rotate the wafer 90° to evaporate the second junction electrode along the other line. We use static oxidation using 2 mbar of pure oxygen for 20 minutes, and control the tunnel junction resistance by the size of junction.

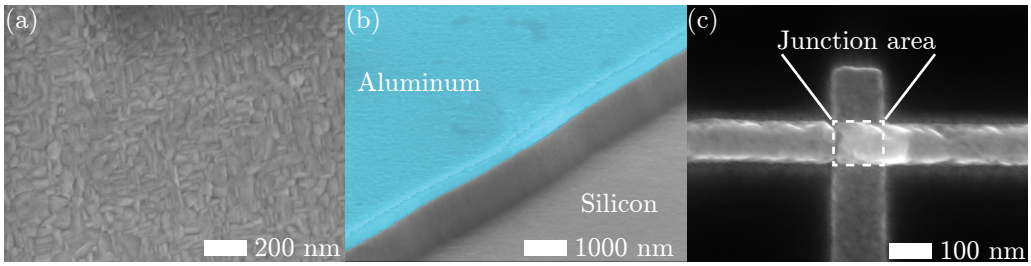

Figure 2.3: *Scanning electron microscopy images. (a) The* 150 nm *thin aluminum film that constitutes all microwave structures. The grains are clearly visible. (b) An instance of a* 1 µm *trench into the silicon. (c) A Josephson junction. The insulator is sandwiched between the two overlapping aluminum leads, which forms a cross.*

Our oxidation gives a normal-state resistance of $3.05 \times 10^8\,\Omega\,\mathrm{nm}^{-2}$. This fabrication process for Josephson junctions is 'bridge-less' as it does not rely on the formation of a resist bridge to shadow the evaporation [127]. The bridge-less process is less sensitive to variation in the resist development and its associated undercut, which makes our process more robust against variations in development rates and resist thicknesses.

The junction geometry and fabrication used here avoids the creation of any extra Josephson junctions. While such junctions do not participate in the electrical circuit, they have been shown to contain a large density of TLS [101, 102].

A final lithography step is used to connect the Josephson junctions' electrodes to the rest of the circuit. We pattern rectangles covering parts of the electrodes and the connecting circuit. We ensure good electrical contact by using an in-situ ion mill to remove native aluminum oxide before evaporating 200 nm of aluminum. Even though these rectangles are several micrometers large, we use electron-beam lithography for the pattering. The low glass-transition temperature of conventional photoresists makes ion-milling difficult as the resist will reflow and cross-link due to heating from a neutralizing filament and the implantation of argon ions. E-beam resists are not as sensitive, thus yielding better results with cleaner surfaces.

As the last step, the wafer is diced into individual chips and thoroughly cleaned in solvents and an ozone atmosphere.

2.5 Measurement setup

As discussed previously, the environment around the device may induce loss and dephasing of qubits and resonators. An ideal environment would be to isolate the device as much as possible by enclosing it in a vacuum-sealed superconducting box to screen any electromagnetic fields (e.g., static, microwave, infrared, and optical). However, to measure the device, we need to send and receive microwaves from the device. Additionally, for some experiments, we also need static magnetic fields and electrical currents, implying that we need to trade-off isolation and ease of measurement. The exact setup used for each experiment in this work is found in the corresponding appended paper.

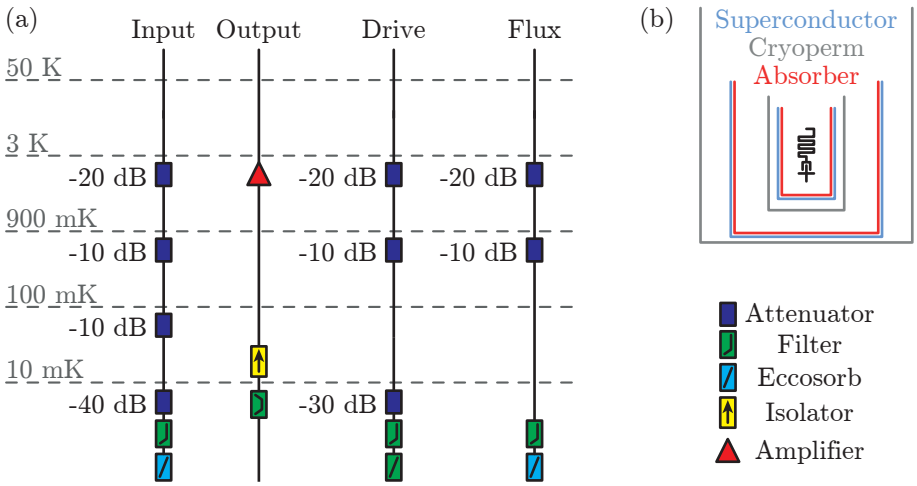

Figure 2.4: *Measurement setup for high-coherence superconducting circuits. (a) Attenuation and filtering scheme for the four types of coaxial lines described in the main text. (b) Shielding setup. Two cryoperm shields and two superconducting shields protect the sensitive circuits from any external magnetic fields. The superconducting shields are painted on the inside with infrared-absorbing paint to reduce the amount of quasiparticle generating radiation.*

As the devices operate below 10 mK and at GHz frequencies, it is crucial to attenuate the Johnson-Nyquist noise from room temperature and warmer stages of the cryostat. Otherwise, the effective device temperature may be much higher than the cryostat temperature. Typically, we have four different types of coaxial lines to our devices: inputs and outputs for readout, drive lines for single-qubit gates, and flux lines for frequency-tunable elements. The required attenuation is different for each type, see Fig. 2.4 (a), as the noise in each line will couple with different matrix elements in Fermi's golden rule, Eq. (2.3). Ideally, all attenuation is placed at the mixing chamber of the cryostat to achieve the lowest possible amount of noise. However, the cooling power of the mixing chamber is limited and not enough to handle the dissipated power if all attenuators where to be placed there.

For all lines except the readout output lines, we place a 20 dB attenuator at the 3 K stage of the cryostat, where the cooling power is more than 1000 times greater than at the mixing chamber. The readout input line is the most attenuated, since qubit dephasing is very sensitive to residual thermal photons in the readout resonators, Eq. (2.11). The qubit drive lines are also heavily attenuated since the qubits will thermalize to the same effective temperature as the drive line. However, the drive line is slightly less attenuated than the readout input to allow for fast qubit gates. In contrast to a drive line, a flux line is grounded close to the tunable element and does not couple to electric fields of the device. Therefore, the attenuation constraint is relaxed for the flux line, and it is enough with 30 dB in total and none at the mixing chamber. We do not put any attenuation at the mixing chamber as we want to be able to use the flux line for DC-biasing of SQUIDs,

therefore requiring currents of the order of 1 mA to pass through the line. Any attenuation at the mixing chamber would dissipate too much power at such currents. Finally, the readout output line does not have any attenuation at all, as that would decrease the signal-to-noise ratio of the readout. To shield the device from high-temperature noise, we use microwave isolators to provide an effective attenuation of at least 60 dB in the reverse direction, while providing almost none in the forward direction. Isolators only work in a specific frequency band, so we use low- and high-pass filters to shield the device outside of the isolator band.

In our multi-qubit setups, we use dedicated drive lines for each qubit. A drive line is weakly coupled to a qubit so that the qubit does not relax into the drive line faster than the qubit's intrinsic relaxation rate. We design the coupling so that $1/\Gamma_1$ into the drive line is above 250 µs. This weak coupling implies that the drive signal needs a high amplitude to drive a fast (10–30 ns) single-qubit gate. However, we still need to heavily attenuate any thermal noise in the drive line not to cause a significant population of the $|1\rangle$ state. To achieve a low thermal population and fast gates puts a tight window of viable attenuation values that can drive a fast gate with typical output powers of microwave signal generators. We use −60 dB.

To minimize any detrimental effects from stray magnetic fields, we employ flux-trapping holes. However, for flux tunable circuits, we also need to minimize any flux noise due to fluctuations in any external magnetic fields. We shield our devices with several layers of magnetic shielding, Fig. 2.4(b). We have experimented with several configurations of shielding and converged at two high magnetic permeability (mumetal and cryoperm) and two superconducting shields. The outer mumetal and superconducting shield mainly reduce magnetic fields from sources outside the cryostat (e.g., the earth magnetic field). The two inner shields are there to reduce magnetic fields from sources inside the cryostat, for instance, microwave isolators containing permanent magnets. Inside the inner shields, we exclusively use non-magnetic hardware (e.g., screws, connectors, and cables).

Photons with an energy higher than twice the superconducting gap can break Cooper-pairs into quasiparticles, which cause relaxation and dephasing. It is believed that there are two main paths for high-energy photons to enter the cryostat. The first is through the coaxial cables themselves. Although the lines are low-pass filtered to cut-off all frequencies above 8 GHz, such filters are never perfect, especially at many times the cut-off frequency. At hundreds of gigahertz, both filters and the dielectric of the cables are transparent, and infrared radiation from room temperature can travel unhindered through the control lines to the sample where it creates quasiparticles. To solve this issue, we use home-made eccosorb filters [86], which have exponentially increasing attenuation with respect to frequency.

The second path is by direct line-of-sight between hot surfaces and the sample. The cryostat has several temperature stages, where the hottest one is at room temperature. All stages will emit black-body radiation, and it is important to shield the sample from this radiation. Therefore, the cryostat itself has shields on each temperature stage (except at the 100 mK stage) to shield from radiation from the previous stage. As a final precaution, we use infrared-absorbing paint on the shields surrounding the device [94, 128].

2.6 Characterization of loss and dephasing

2.6.1 Coherence and frequency noise of resonators

In Paper A and Paper D, we specifically study the energy loss of resonators, but it is characterized in almost all the appended papers. For resonators, we typically express loss as quality factors $Q_i = \omega_r/\Gamma_1$, where ω_r is the resonant frequency. Similarly, the coupling between resonator and transmission line is expressed as $Q_c = \omega_r/\kappa$. To experimentally find the resonant frequencies and quality factors, we measure either the reflection coefficient S_{11} or transmission coefficient S_{21} as a function of frequency. The coefficient we use depends on the setup and varies across the different papers. By fitting the measured coefficients to either

$$S_{11} = \frac{2\delta\omega/\omega_r - i\,(1/Q_i - 1/Q_c)}{2\delta\omega/\omega_r - i\,(1/Q_i + 1/Q_c)}, \tag{2.12}$$

or

$$S_{21} = 1 - \frac{1/Q_c}{1/Q_i + 1/Q_c + 2i\delta\omega/\omega_r}, \tag{2.13}$$

we can determine the resonant frequency and internal and coupling quality factors.

In Paper A, we measure a single microwave resonator made of aluminum on a silicon substrate. We change the power in the resonator and extract the internal quality factor for each power. We see a monotonically increasing quality factor with circulating power. Since TLS is the only loss with a power dependence, we compare the internal quality factors with

$$Q_i(P) = \frac{1}{\Gamma_{\mathrm{TLS}}(P)/\omega_r + \delta_{\mathrm{other}}}, \tag{2.14}$$

where δ_{other} contains all other loss mechanisms. From this measurement and Eq. (2.9), we extract $F\delta_{\mathrm{TLS}} = 1/(0.87 \times 10^6)$ and $\delta_{\mathrm{other}} = 1/(3.5 \times 10^6)$. The agreement between data and the TLS model is excellent and demonstrates that the resonator's dominant loss mechanism at low powers is TLS. Which loss mechanism that dominates at high power cannot be inferred from this measurement. If we were to measure the resonator's quality factor against temperature and magnetic field, we might be able to understand more about the high-power loss.

In Paper A, we also study fluctuations of the resonant frequency. The resonant frequency is determined by the resonator geometry and dielectric constant, which should not change. However, the movement of TLS in frequency space due to TLS interaction with low-frequency TLS causes the effective dielectric constant of the resonator to fluctuate [109, 129, 130].

We perform the study at a circulation power close to that of a single photon, making this study relevant to qubits, as it will set a limit on the dephasing for qubits with similar design and fabrication. We measure the frequency fluctuations using a Pound loop, which is a frequency locked loop that can monitor the resonant frequency in real-time [131]. We find that the noise in our aluminum resonator is well described by a fractional frequency fluctuation spectrum

$$S_y(f) = \frac{3.5 \times 10^{-15}}{f^{1.05}} + 2.5 \times 10^{-16}. \tag{2.15}$$

We can estimate the dephasing time for a qubit experiencing a $1/f$ noise with that amplitude, $T_\phi \approx 520\,\mu s$.

Our noise amplitudes are higher than those observed in niobium resonators [109], although those were measured at higher powers and temperatures, which decreases the noise substantially.

In Paper D, we study similar resonators but on a piezoelectrical substrate, gallium arsenide. Piezoelectricity is the creation of mechanical strain from electric fields, and vice versa. A varying electric field will inevitably create propagating phonons, both in the bulk and at the surface of the substrate. This effect can be desired, such as in the field of quantum acoustics [132, 133]. However, it is also detrimental for high-coherence circuits, as the conversion from photons to phonons can severely increase the loss in such circuits.

In Paper D, the power dependencies of the resonators' quality factors are much less pronounced than on silicon, with $F\delta_{\mathrm{TLS}} = 1/(0.35 \times 10^6)$ and $\delta_{\mathrm{other}} = 1/(0.046 \times 10^6)$, indicating that TLS is not the dominant loss mechanism. In the paper, we present a numerical simulation of a model where the varying electrical field creates bulk and surface acoustic waves that propagate in and on the substrate. This photon to phonon conversion predicts a linear relationship between quality factor and frequency, which we also observe by measuring a large set of resonators with different frequencies. The absolute values of δ_{other} agree well with the simulated values, lending confidence in that the piezoelectric substrate is the dominant loss mechanism for these resonators.

A related topic is that any surface is piezoelectric, even if the bulk of the material itself is not, which means that the varying electric fields generated by qubits and resonators on silicon will generate phonons [58, 134]. Currently, this effect is believed to be small, and there has not been any observation of it yet. Nevertheless, it is worth taking it into consideration as other sources of loss are reduced.

2.6.2 Coherence fluctuations in circuit QED

In the circuit-QED architecture, a qubit is usually coupled to a resonator and not to any transmission lines, Fig. 2.2(a), which prohibits us from probing the qubit properties directly. The lack of coupling is by design, as it increases the coherence of the qubit. Thankfully, we can use the resonator to measure the qubit state, and therefore the qubit properties as well. Depending on which property we want to investigate, we apply different evolutions, or gate sequences, to the qubit. First, to measure the time constant of the qubit relaxation, T_1, we prepare the qubit in its excited state, idle for some time, and measure the qubit state. We repeat the experiment multiple times and extract the excited state probability. The probability as a function of idling time can be fitted to some decay model (usually a single decaying exponential e^{-t/T_1}), from where the relaxation time T_1 is extracted.

In Fig. 2.5 (a), we show one instance of a measured qubit with $T_1 = 63.5\,\mu s$. This qubit had a frequency of $\omega_q/2\pi = 4.4\,\mathrm{GHz}$, corresponding to a quality factor of 1.8×10^6. The qubit was fabricated very similarly to the resonator in Paper A, so one could expect them to have similar quality factors. The discrepancy of a factor of 2 can partly be explained by that the resonator has smaller dimensions than the qubit, implying a higher filling factor for the resonator. The resonator also had a fabrication issue with residual photoresist

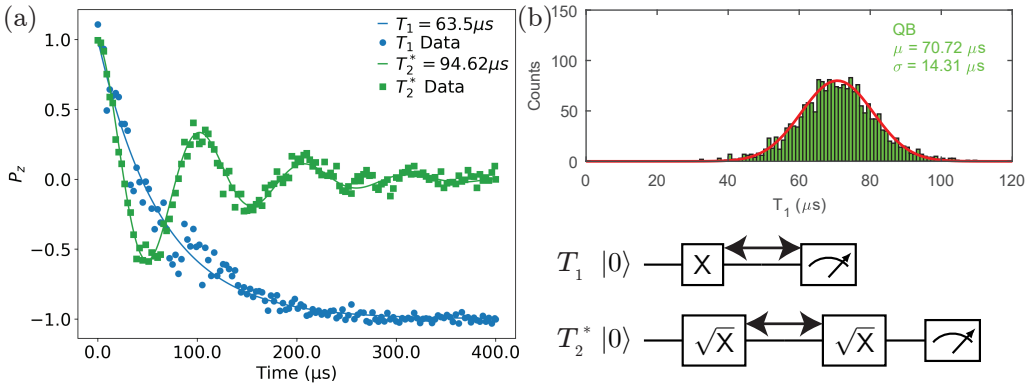

Figure 2.5: *(a) Relaxation and coherence measurement (Ramsey free-induction decay) of a fixed-frequency transmon qubit operating at* 4.4 GHz. *The corresponding gate sequences are shown to the right. The arrows indicate where the time delay is inserted. (b) Histogram over 2000 measurements of the qubit relaxation time, and a fit to a Gaussian distribution with a mean of* 71 µs *and a standard deviation of* 14 µs.

on the resonator, leading to a higher δ_{TLS}. The lower filling factor and the lower loss tangent for the qubit explain why it has a higher quality factor than the resonator has at low powers. In principle, we could make new resonators with the same filling factor as the qubit, and their quality factors should agree [104].

The sequence for measuring the relaxation time has the qubit confined to the measurement basis; thus, the qubit phase evolution does not affect the measurement outcome. To be sensitive to phase errors and measure dephasing rates, we instead prepare the qubit in a superposition state (e.g., $|+\rangle$ by applying a $\sqrt{\text{Y}}$ gate), idle for some time t, rotate into the measurement basis (by another $\sqrt{\text{Y}}$ gate) and perform a measurement. This sequence is referred to as Ramsey, named after Norman Ramsey. If the qubit phase changes during the idling time, the subsequent gate will no longer bring the qubit to $|1\rangle$. For instance, if the qubit is in $|-\rangle$ after the idling period, a $\sqrt{\text{Y}}$ gate will bring the qubit to $|0\rangle$. If the qubit drive is detuned in frequency from the qubit by δ, the qubit acquires a deterministic phase shift equal to δt, causing the Ramsey measurement to oscillate between $|1\rangle$ and $|0\rangle$ as t is increased. However, if δ fluctuates, the qubit phase will no longer be deterministic, so when averaged over many experimental trials, the oscillations will have a smaller amplitude. As t increases, the qubit phase will be more random, leading to a decaying amplitude (envelope).

The decay envelope is used to extract the decoherence time T_2^* (the star is used to denote that it is measured using a Ramsey sequence), from which the dephasing time T_ϕ is calculated via Eq. (2.1). The exact shape of the envelope depends on the noise spectrum experienced by the qubit. A $1/f$ spectrum would cause a Gaussian envelope $e^{-(t/T_2^*)^2}$, whereas a white spectrum gives a single exponential e^{-t/T_2^*}. Generally, a combination of the two models can be used if neither $1/f$ or white noise is dominating. For the edge case where T_ϕ is infinite and $T_2^* = 2T_1$, the decay is exponential.

Figure 2.5 (a) shows one instance of a Ramsey measurement with $T_2^* = 94.6$ µs and

$T_\phi = 370\,\mu s$. The decay shape used here only contains an exponential and agrees well with the data. In our fixed-frequency qubits, we typically only observe exponential decaying Ramsey fringes, indicating that the noise seen by the qubit is mostly white during the duration of the experiment. The source of the white noise could be thermal photons in the resonator.

Only measuring and reporting single instances of coherence measurements is not enough. Coherence times are known to fluctuate [135–137]. The main result of Paper B is the extensive statistics collected and analyzed. For example, we measure T_1 uninterrupted for three days, giving us a time series of 2001 T_1 values. The histogram of the extracted T_1 values are shown in Fig. 2.5 (b). We observe a difference between the minimum and maximum values of almost a factor of 3.

We analyze the relaxation time fluctuations in detail by transforming the time series to the frequency domain. The spectrum can give insights into which loss mechanism is responsible for the fluctuations. Additionally, in the time domain, we use the overlapping Allan deviation [138], a tool commonly used in frequency metrology. First, we consider three common noise processes, expressed in Allan deviation as [138],

$$\sigma_{T_1}(\tau) = \left(\frac{2\pi^2}{3} h_{-2}\right)^{\frac{1}{2}} \tau^{\frac{1}{2}} + \left(2\ln(2) h_{-1}\right)^{\frac{1}{2}} + \left(\frac{h_0}{2}\right)^{\frac{1}{2}} \tau^{-\frac{1}{2}}, \tag{2.16}$$

where h_{-2}, h_{-1}, and h_0 are the amplitudes of random walk, $1/f$, and white noise processes, respectively. In the frequency domain, these processes translate into power laws

$$S_{T_1}(f) = \frac{h_{-2}}{f^2} + \frac{h_{-1}}{f} + h_0, \tag{2.17}$$

from which the amplitude notation becomes obvious. Importantly, no power-law process will generate a peak in the Allan-deviation (i.e. a stationary point with negative second-derivative). To achieve such a feature in the Allan-deviation, we must include a Lorentzian noise process [139],

$$\sigma_L(\tau) = \frac{A\tau_0}{\tau} \left(f e^{-\tau/\tau_0} - e^{-2\tau/\tau_0} - 3 + 2\tau/\tau_0\right)^{\frac{1}{2}} \tag{2.18}$$

with the corresponding frequency spectrum

$$S_L(f) = \frac{4A^2\tau_0}{1 + (2\pi f \tau_0)^2}, \tag{2.19}$$

where A is the Lorentzian amplitude and $1/\tau_0$ its switching rate.

In Paper B, we found that most T_1 fluctuations could be explained by a combination of two Lorentzians and white noise. In Fig. 2.6 (a) and (b), we see an example of such a data set. Both the Allan deviation and the frequency spectrum agrees well with the measured data. In Fig. 2.6 (a), we see a clear peak in the Allan deviation, corresponding to Lorentzian noise processes.

The Lorentzians' switching rates were in the range 70–1850 µHz, well in agreement with TLS switching rates found in amorphous solids [140, 141] and other qubit experiments

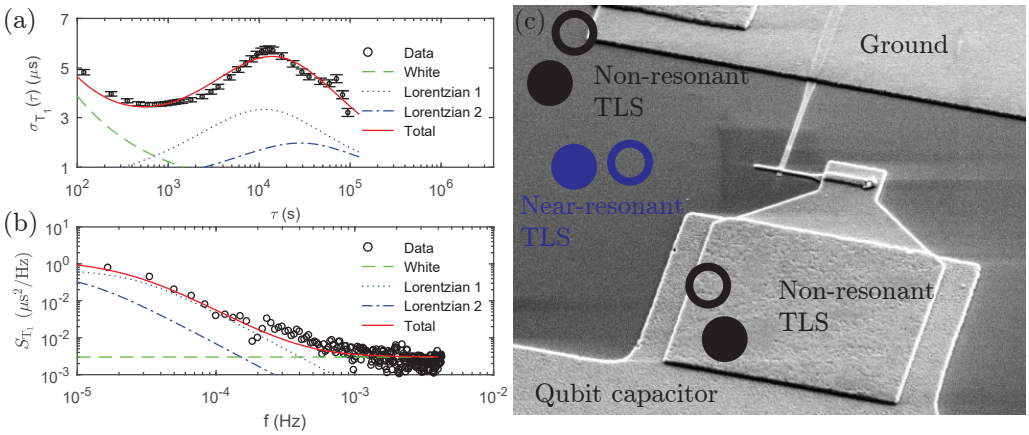

Figure 2.6: *Relaxation-time fluctuations due to interacting two-level systems (TLS). (a) Overlapping Allan deviation containing two Lorentzians and white noise. (b) Spectral density of the same fluctuations and noise models as in (a). (c) A possible explanation of the two Lorentzians. A single near-resonant TLS interacts with two neighboring off-resonant TLS. The states of the off-resonant TLS affect the frequency of the near-resonant TLS, causing the qubit to see a different amount of noise depending on the states of the off-resonant TLS. The switching rates in (a) and (b) correspond to the switching rates of the off-resonant TLS.*

.

[135, 136]. We interpret the two Lorentzians as the qubit being coupled to either one near-resonant TLS coupled to two off-resonant TLS [see Fig. 2.6 (c)], or two near-resonant TLS coupled to one off-resonant TLS each. In the interacting-TLS-model [129], TLS can interact with off-resonant TLS, which can have slow switching rates. When a coupled TLS switches its state, it induces a frequency shift of the near-resonant TLS, changing the noise seen by the qubit, which in turn changes its relaxation time. The observed switching rates are the switching rates of the off-resonant TLS.

These slow switching rates highlight the importance of measuring coherence across long periods as the TLS can be in the same configurations for hours before switching. We argue that any coherence study should be performed during at least 5 hours and with more than 1000 individual samples to acquire a sufficient amount of statistics. Moreover, fluctuations and their spectra, and not only mean values, should be reported and compared.

2.6.3 Measuring coherence in waveguide QED

Experiments not aiming to store quantum information in qubits can use qubits directly coupled to transmission lines. The coupling to the transmission line, Γ_r, is typically made large such it dominates the qubit lifetime, $\Gamma_1 \approx \Gamma_\mathrm{r}$. However, there are applications where Γ_r has to be kept reasonably low, and the intrinsic loss rate (or non-radiative decay rate) Γ_nr can come into play, together with the dephasing rate. In general, for a

waveguide-QED system, we have

$$\Gamma_1 = \Gamma_\mathrm{r} + \Gamma_\mathrm{nr}, \tag{2.20}$$

$$\Gamma_2 = \frac{\Gamma_1}{2} + \Gamma_\phi = \frac{\Gamma_\mathrm{r} + \Gamma_\mathrm{nr}}{2} + \Gamma_\phi. \tag{2.21}$$

Typically, Γ_r and Γ_2 are extracted from the coherent scattering at low signal powers. However, this measurement cannot distinguish Γ_nr from Γ_ϕ. In Paper C, we demonstrate several methods to overcome this deficit and extract all involved rates. We find $1/\Gamma_\mathrm{nr} = 3\,\mu\mathrm{s}$ and $1/\Gamma_\phi = 53\,\mu\mathrm{s}$, both which are an order of magnitude smaller than the T_1 and T_ϕ from a circuit-QED system. We currently do not have an answer to why, but we are speculating that it could be to a different measurement setup for the waveguide-QED system or just an outlier. Additional measurements and statistics are needed to confirm if the discrepancy actually exists.

2.7 Further improvements to qubit coherence

There are three problems regarding the coherence of current superconducting qubits. They are "only" on the order of one-hundred microseconds, they fluctuate in time, and there are frequencies where the coherence is much worse. The fluctuations can lead to downtime and the need for recalibration of quantum processors, while the existence of harmful frequencies is challenging for achieving high fidelity gates across a large grid of qubits [16]. As the number of qubits grows, there will always be some qubits close to such frequencies, severely limiting the performance of gates involving those qubits.

Luckily, all three problems are currently related to the presence of two-level systems. So if we can eliminate TLS, we should solve all three problems. Of course, then the next loss mechanism will limit the coherence instead.

How do we go about identifying the source and location of TLS? Similarly to Refs. [99–101], we could perform different types of spectroscopy techniques to identify the atom or molecule responsible. The location of TLS can be inferred by the method in Ref. [102]. Identifying the source is not enough; we also need to remove the source. Depending on where on the device they are located, that can be challenging. Typical hosts of TLS are oxides and other surface defects. It is challenging to prevent oxide from growing on the surface of aluminum and silicon. It is possible to cover the aluminum surface with a noble metal (e.g., gold, palladium, or platinum). Noble metals corrode less, but can still attract impurities on their surfaces. It is crucial that the cover is thin enough so that it becomes superconducting due to the proximity effect [26].

An alternative method is to use a superconductor that oxidizes less, such as titanium nitride [114, 122]. Recently, tantalum has also shown good promise, reaching relaxation times of 300 µs [88].

Additionally, other types of superconducting qubits show great promise in terms of coherence times. For example, the fluxonimum has shown $T_1 > 8\,\mathrm{ms}$ [142] and $T_2 > 400\,\mu\mathrm{s}$ [143], and the capacitively shunted flux qubit has shown coherence times similar to the transmons in this work [116]. These types of qubit have other disadvantages, such as that

they have to operate under specific bias conditions, are difficult to fabricate reproducibly, and make it difficult to couple qubits together.

One can ask, how long coherence do we need? One answer might be, as long as possible. Typical quantum algorithms like Shor's require error rates below 10^{-15} [13]. If we can perform gates in 10 ns, we would need coherence times on the order of 10 days to achieve such low error rates, which is not feasible for superconducting circuits. Therefore, no matter how much we improve our coherence, error correction will be needed for such algorithms. That does not mean that there is no point in improving coherence times. The overhead for the error correction scales with the error rates of the physical gates. For example, physical error rates on the order of 0.01 yield 4761 physical qubits per logical qubit if we want logical error rates below 10^{-15}. If we can improve the errors with a factor of 10, then the overhead is approximately four times lower [144].

CHAPTER III

High-fidelity control and readout of quantum bits

A vital aspect of quantum bits is to be able to control and measure their states with high accuracy. Control errors degrade the overall performance of a quantum computer and limit how deep quantum circuits we are able to run. This chapter summarizes the ideas behind qubit control and readout (including Paper E about a new type of readout). We go through the single and two-qubit gates used in Paper F and Paper G and how to tune them up to state-of-the-art fidelities, 0.999 for single-qubit gates and 0.99 for a two-qubit gate.

For quantum computers to have any chance of having an advantage over their classical counterpart, the qubits need to utilize the large Hilbert space given by superposition. To initialize a qubit into an arbitrary superposition state, $|\psi\rangle = c_0|0\rangle + c_1|1\rangle$, we use a sinusoidal drive at the qubit frequency. The amplitude, phase, and length of the drive determine the complex state amplitudes c_i. The probability of measuring a certain state $|i\rangle$ is $P_{|i\rangle} = |c_i|^2$, thus, $|c_0|^2 + |c_1|^2 = 1$ is required for proper normalization.

A single-qubit state can be represented by the position on a Bloch sphere where $|0\rangle$ is on one pole and $|1\rangle$ on the other, with the z-axis intersecting the two poles. For two (or more qubits) with any entanglement between them, Bloch-spheres are no longer good representations, and we instead use a state vector with a complex amplitude for each state,

$$|\psi\rangle = c_{00}|00\rangle + c_{01}|01\rangle + c_{10}|10\rangle + c_{11}|11\rangle = \begin{pmatrix} c_{00} \\ c_{01} \\ c_{10} \\ c_{11} \end{pmatrix}, \tag{3.1}$$

where c_{ij} is the complex amplitude for the $|ij\rangle$ state. For n qubits, the state vector contains 2^n amplitudes.

A controlled change of a state vector is known as a quantum gate. All quantum gates are reversible and can be represented by their unitary matrices U. The new state vector after the application of U is equal to $U|\psi\rangle$. For example, the single-qubit X gate has the

unitary

$$\mathsf{X} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \tag{3.2}$$

which swaps the amplitudes of the $|0\rangle$ and $|1\rangle$ states. Two examples of two-qubit gates are the controlled-phase (CZ) gate with the unitary

$$\mathsf{CZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \tag{3.3}$$

and the iSWAP gate, which has the following unitary

$$\mathsf{iSWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}. \tag{3.4}$$

A CZ gate applies a phase shift of π to one qubit if the other is in the excited state, while the iSWAP gate swaps the state amplitudes between the two qubits but with a $pi/2$ phase shift on both qubits. Both the CZ and iSWAP gates conserve the number of excitations in the system, and together with arbitrary single-qubit gates, constitute universal gate sets, which means that any multi-qubit unitary can be decomposed into the gates in those sets [145].

In Chapter 2, we discussed two different types of errors, bit and phase flips, due to relaxation and dephasing of the qubit state. These types of errors are probabilistic, known as incoherent errors, and their associated error rates are related to the gate, relaxation, and dephasing times. In the case of exponentially decaying decoherence, the incoherent error rate for a single-qubit gate is [146]

$$\epsilon = \frac{t_{\text{gate}}}{3} \left(\Gamma_1 + \Gamma_\phi \right), \tag{3.5}$$

where t_{gate} is the time duration of the gate. Due to their probabilistic nature, incoherent errors cannot be described by a single unitary matrix. Instead, they are represented as sets of Kraus operators [5].

It is worth noting that in the presence of decoherence and incoherent errors, a state vector can no longer be used to describe the resulting mixed quantum state. Instead, we need to use a density matrix ρ [5]. For a single qubit, a mixed state can be seen in the Bloch sphere by a vector with a length less than unity. The state is then said to be inside the Bloch sphere.

There are also coherent errors. These errors are deterministic and due to some nonideality in the applied unitary. For instance, it could be that the drive amplitude of a single-qubit gate is too high, leading to an over-rotation. The name, coherent errors, refers to the fact that the qubit is still well described by a state vector since no coherence was lost due to the error. The coherent-error rates are strongly dependent on how well

we know and calibrate the system. For example, An over-rotation of an X gate by 10% would have the nonideal unitary matrix

$$\mathsf{X}' = \begin{pmatrix} 0.02 - 0.15j & 0.98 + 0.15j \\ 0.98 + 0.15j & 0.02 - 0.15j \end{pmatrix}, \tag{3.6}$$

which should be compared with Eq. (3.2). For very short gate times, $t_{\text{gate}} \lesssim 2\pi/\eta$, coherent errors are not as simple as an error in the drive amplitude. For these short gate times, higher states of the qubit come into play, and the coherent error rate is more related to how well we can compensate for effects due to these higher states.

To quantify how good a gate is, we define a fidelity number, F, which is

$$F = \left(\text{tr}\sqrt{\sqrt{U'}U\sqrt{U'}} \right)^2 = \left(\text{tr}\sqrt{\sqrt{\rho'}\rho\sqrt{\rho'}} \right)^2, \tag{3.7}$$

where U is the ideal unitary and U' is the performed unitary, or in the case of density matrices, ρ is the ideal state after the gate and ρ' is the actual state [147]. Several techniques can be used to measure or infer qubit states and performed unitaries, such as quantum state and process tomography [148, 149], and cross-entropy benchmarking [150]. In Paper F, we use process tomography to investigate what unitary was implemented using a sequence of gates.

In general, longer gate times imply fewer coherent, but more incoherent, errors. The lowest total error rate occurs when the coherent and incoherent error rates are equal. However, there can be applications where one type of error is more detrimental than others. For instance, in the context of quantum error correction, coherent errors are more damaging than incoherent ones [151]. In our work, we run with a bit slower gates where the incoherent error rates are higher than the coherent ones since slower gates make calibration easier.

3.1 Single-qubit gates

To realize single-qubit gates, we need some way to couple transversely and exchange energy with the qubit. The transmon qubit couples well to electric fields and a time-varying voltage, $V(t)$, yields the following Hamiltonian [58]

$$\frac{\hat{H}}{\hbar} = \frac{\omega_q}{2}\hat{\sigma}^z + CV(t)\left(\hat{\sigma}^+ + \hat{\sigma}^-\right), \tag{3.8}$$

where C is some constant to account for the coupling strength between the voltage source and the qubit. The signal used to drive the qubit can be described as a motion in a complex plane characterized by the in-phase $I(t)$ and quadrature $Q(t)$ voltages, together with the drive frequency ω_d,

$$V(t) = I(t)\cos(\omega_d t + \phi) + Q(t)(t)\sin(\omega_d t + \phi). \tag{3.9}$$

We enter a rotating frame at ω_d, take $\phi = 0$, and make the rotating wave approximation, which transforms Eq. (3.8) into

$$\frac{\hat{H}}{\hbar} = \delta\hat{\sigma}^z + \frac{C}{2}I(t)\hat{\sigma}^x + \frac{C}{2}Q(t)\hat{\sigma}^y, \tag{3.10}$$

where $\delta = \omega_q - \omega_d$. In the Bloch-sphere picture, $I(t)$ determines rotations of the qubit state around the X-axis. To achieve rotations around the Y-axis, we can either set $\phi = \pi/2$ or use $Q(t)$. In general, the drive phase ϕ determines the rotation axis for the in-phase voltage.

3.1.1 X and Y gates

To implement controlled rotations with some angle θ around the X-axis, we apply a pulse to $I(t)$. We choose to use a pulse envelope with a cosine shape, see Fig. 3.1 (a), which is parameterized by

$$I(t) = \frac{A}{2}\left(1 - \cos\frac{2\pi t}{\tau}\right), \tag{3.11}$$

where A is the pulse amplitude, and τ is the pulse length. Other pulse shapes, such as a Gaussian, are also commonly used. One good feature of the cosine shape is that it starts and ends at zero, whereas a Gaussian extends infinitely and needs to be truncated, leading to a broader spectral content.

The area under the pulse envelope is proportional to the rotation angle θ. By changing A or τ, we control θ. In Fig. 3.1, we plot the measured qubit state as a function of A and observe that the qubit is oscillating between $|0\rangle$ and $|1\rangle$, so-called Rabi oscillations. This type of plot provides a map between the rotation angle θ and A, allowing us to realize arbitrary rotations around the X-axis. For instance, $A \approx 0.5$ and $A \approx 1$ correspond to $\sqrt{\mathsf{X}}$ ($\theta = \pi/2$) and X ($\theta = \pi$) gates, respectively. The reason that $P_{|1\rangle}$ does not reach 1.0 is not due to an error in the gate but to a limited readout fidelity. More details about the limitations of readout and how it can be calibrated are discussed in Section 3.3.

Compensation of nonidealities due to higher qubit levels

The quadrature $Q(t)$, can be used to compensate for nonidealities due to the presence of higher qubit levels. We apply the Derivative Removal by Adiabatic Gate (DRAG) scheme [152] by using a signal on the quadrature voltage set to

$$Q(t) = -\frac{\alpha}{\eta}\frac{\mathrm{d}X(t)}{\mathrm{d}t} = -\frac{\pi A\alpha}{\tau\eta}\sin\frac{2\pi t}{\tau}, \tag{3.12}$$

where η is the qubit anharmonicity, and α is a real-valued scaling parameter. $Q(t)$ is anti-symmetric around the center of the pulse, implying that its area is zero. In an ideal system, the DRAG component would not affect the qubit state at all since it just reverses the $\hat{\sigma}^y$ trajectory after half the pulse. However, for a qubit with more states than just $|0\rangle$ and $|1\rangle$, such as the transmon, errors can occur due to those higher states.

During a resonant drive of the transition between $|0\rangle$ and $|1\rangle$, there is also an off-resonant drive of the $|1\rangle$ to $|2\rangle$ transition. Even if the probability of populating $|2\rangle$ is decreased due to the drive being off-resonant, it is still finite and causes leakage out of the qubit subspace. Moreover, the drive induces an AC-Stark shift, causing the $|1\rangle$ and $|2\rangle$ states to repel and therefore reduces the transition frequency ω_q during the drive. The effect of a detuning between the drive and qubit frequencies results in a $\hat{\sigma}^z$ component in the interaction Hamiltonian, which gives a phase error on the final qubit state.

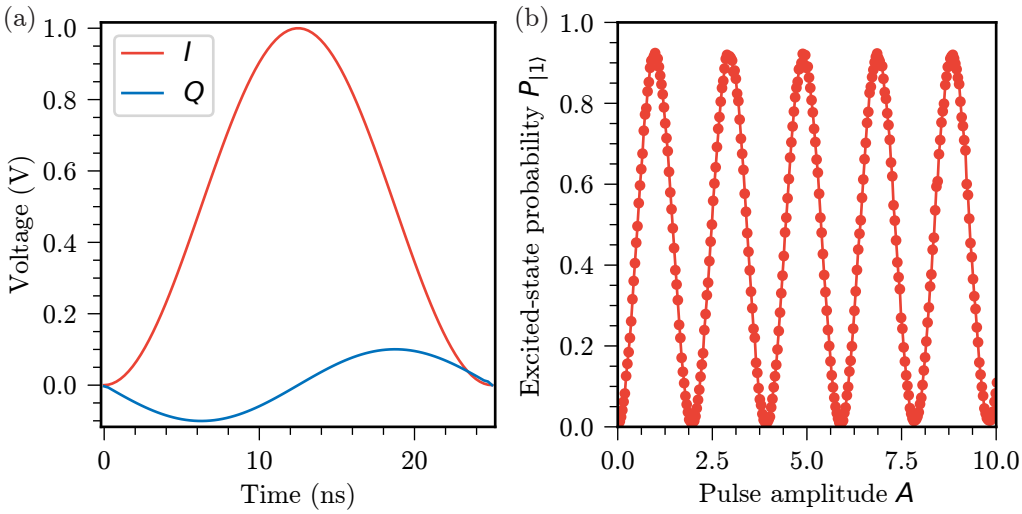

Figure 3.1: *Single-qubit X control. (a) The pulse envelope with its two components $I(t)$ (red) and $Q(t)$ (blue). Here, the pulse length is $\tau = 25\,\mathrm{ns}$ and the amplitude is $A = 1$. The Q component is the DRAG compensation, Eq.* (3.12)*. (b) Experimentally observed Rabi oscillations. The probability of measuring $|1\rangle$ as a function of the pulse amplitude A.*

DRAG can compensate for each of the two types of errors. DRAG with $\alpha = 1/2$ should minimize phase-errors [153, 154], and $\alpha = 1$ should minimize leakage [155]. It is possible to reduce both errors by using $\alpha = 1$ and adding a frequency shift to ω_d to compensate for the phase error. In our work, we use $\alpha = 1/2$ as leakage is sufficiently reduced by using a long enough pulse, $\tau \gg 2\pi/\eta \approx 5\,\mathrm{ns}$. Experimentally, the optimal α can differ from $1/2$ and has to be found empirically, see Fig. 3.3 (a).

Pulse generation using IQ mixers

To physically realize Eq. (3.9), we use two channels of an arbitrary waveform generator (AWG), a microwave signal generator at ω_{LO}, and an IQ-mixer, see Fig. 3.2 (a). The AWG is programmed to play $I(t)$ and $Q(t)$, whereas the mixer and signal generator are used to upconvert the signals to ω_d. However, physical IQ-mixers are never ideal. There are three considerations: the output voltage is not linear with input voltages; zero input voltage does not yield zero output voltage; and an amplitude and phase imbalance between the I and Q ports.

If the output voltage is not zero when $I = Q = 0$, a part of the local oscillator is transmitted to the qubit, so-called bleedthrough. This bleedthrough drives unwanted qubit transitions if the local oscillator is resonant with the qubit. To reduce the impact, we shift the local-oscillator frequency by an intermediate frequency, ω_{IF}, away from the qubit frequency, so that any bleedthrough is off-resonant with the qubit. However, only shifting the local oscillator frequency would lead to the drive not being on resonance with the qubit. To keep the drive resonant, we transform the I and Q voltages played by the

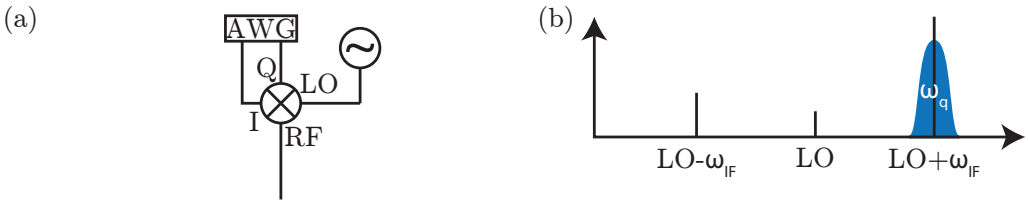

Figure 3.2: *Single-sideband upconversion for driving single-qubit gates. (a) Microwave setup for generating single-qubit pulses. The AWG plays two voltages $I(t)$ and $Q(t)$, which get upconverted to a frequency $\omega_{\mathrm{LO}} + \omega_{\mathrm{IF}}$. (b) The resulting spectrum of a nonideal upconversion. There are unwanted signals at ω_{LO} and $\omega_{\mathrm{LO}} - \omega_{\mathrm{IF}}$.*

AWG according to

$$I^*(t) = I(t)\cos(\omega_{\mathrm{IF}}t + \phi) - Q(t)\sin(\omega_{\mathrm{IF}}t + \phi), \tag{3.13}$$
$$Q^*(t) = Q(t)\cos(\omega_{\mathrm{IF}}t + \phi) + I(t)\sin(\omega_{\mathrm{IF}}t + \phi). \tag{3.14}$$

As a mixer acts as a multiplier between its I/Q and LO ports, the output signal will be at the frequencies $\omega_{\mathrm{LO}} \pm \omega_{\mathrm{IF}}$. The above transformation to I and Q creates signal at the positive sideband, $\omega_{\mathrm{LO}} + \omega_{\mathrm{IF}}$, only. Any signal at the image sideband, $\omega_{\mathrm{LO}} - \omega_{\mathrm{IF}}$, is due to imbalance between the I and Q ports of the mixer.

Even if there should not be any qubit transitions at ω_{LO} or $\omega_{\mathrm{LO}} - \omega_{\mathrm{IF}}$, see Fig. 3.2 (b), we still want to minimize the signals at those frequencies as it makes the mixer behavior more ideal. We calibrate a mixer by applying two sinusoids to its I and Q ports. There are four parameters to calibrate: two DC-offsets, the ratio between the sinusoidal amplitudes, and the phase between the two sinusoids. The offsets determine the local-oscillator bleedthrough, so by varying the offsets and measuring the amplitude at the local-oscillator frequency, the bleedthrough can be minimized. We can typically suppress the local oscillator by more than $-80\,\mathrm{dB}$.

Similarly, the amplitude ratio and phase difference can be optimized by minimizing the amplitude at the image sideband. For instance, we typically use $\omega_{\mathrm{IF}} = -100\,\mathrm{MHz}$, which means that we should minimize the signal at $\omega_{\mathrm{LO}} + 100\,\mathrm{MHz}$. If executed properly, the image sideband is at least $-50\,\mathrm{dB}$ compared to the signal sideband.

Proper mixer calibration is crucial for achieving the high-fidelity gates in Paper F and Paper G, along with the demonstration of the feasibility of the single-photon source proposed in Paper M. To speed-up the mixer calibration, we feed the measured amplitudes to an optimization algorithm, which finds the optimal mixer parameters in a few seconds.

3.1.2 The Z gate

For a spin in a magnetic field, the spin precesses along the magnetic field axis with the Larmor frequency. A qubit is mapped onto the same spin Hamiltonian, including the term $\omega_q\hat{\sigma}^z$, implying that the qubit state precesses around the Z-axis of the Bloch sphere with the frequency ω_q. However, we typically enter a rotating frame at ω_q, which removes any precession, Eq. (3.10). If we were to change ω_q by δ, the rotating frame is no longer at

the qubit frequency, and a precession around the Z-axis with a rate equal to δ is obtained. Therefore, by detuning the qubit frequency for some time in a controlled manner, a Z gate can be realized.

The ability to adjust the frequency of a transmon qubit requires a magnetic-flux sensitive SQUID. However, fixed-frequency qubits generally have longer coherence than frequency-tunable ones. So, how do we implement Z gates for fixed-frequency qubits? If we cannot change the qubit frequency, we instead have to change the rotating-frame frequency. The rotating-frame frequency is set by the frequency of the microwave signal used to drive X and Y rotations, ω_d. Technically, we could change that frequency by δ for some time t. However, the net effect would just be a phase shift equal to δt, and that is equivalent (up to a global phase) to changing the phase ϕ of subsequent gates by δt. This implementation of a Z gate is entirely made in software to the signals uploaded to the AWG and is sometimes referred to as a 'virtual Z gate.' Consequently, virtual Z gates have zero time duration and fidelities equal to unity [156]. So even if we could control the qubit frequency, virtual Z gates would be preferred.

As an example, if we have the gate sequence $\sqrt{\mathsf{X}}$, Z, X. The first gate is implemented as described in the previous subsection. The X gate, however, needs to have its phase ($\phi = \pi$) adjusted so that the Z gate is effectively performed.

3.1.3 Error amplification and tune-up

For a single-qubit gate, we use three control parameters: the drive frequency ω_d, the amplitude A, and the DRAG coefficient α. We measure the detuning between ω_d and ω_q using a Ramsey sequence, whose oscillation frequency is equal to the detuning. By measuring a Ramsey sequence extending up to three times the coherence time, we make sure that our knowledge of the qubit frequency is accurate enough not to cause a significant error during that time.

From the previous Rabi experiment, we have a good idea of the pulse amplitude corresponding to an X gate. However, to optimize the amplitude can be challenging since the sensitivity to an error in the amplitude is rather small. To amplify the error and increase the sensitivity, we may apply sequences of gates. For instance, a sequence of ten X gates with some over-rotation will over-rotate ten times more than a single gate would.

Since we already have a good idea of the pulse amplitude corresponding to an X gate, we first optimize the DRAG coefficient before returning to fine-tune the amplitude. DRAG can compensate for either phase errors, leakage, or both. We restrict the pulse length to $\tau > 25\,\text{ns}$, such that leakage is minimal even without DRAG [155]. To amplify any phase errors, we apply an alternating sequence of X and X^{-1} gates. After an even number of gates, the qubit should remain in its ground state, even in the presence of any over- or under-rotations since they will cancel out due to the alternation between positive and negative gates. Contrarily, phase errors accumulate since the AC-Stark shift does not depend on the sign of the gate, causing oscillations between the $|0\rangle$ and $|1\rangle$ states as the number of gates is increased. In Fig. 3.3 (a), we measure this alternating sequence as a function of the gate count and the DRAG-coefficient. The optimal $\alpha \approx 0.08$ is where no oscillations are observed as the number of gates is increased.

To optimize the pulse amplitude, we apply a similar sequence as for the DRAG-

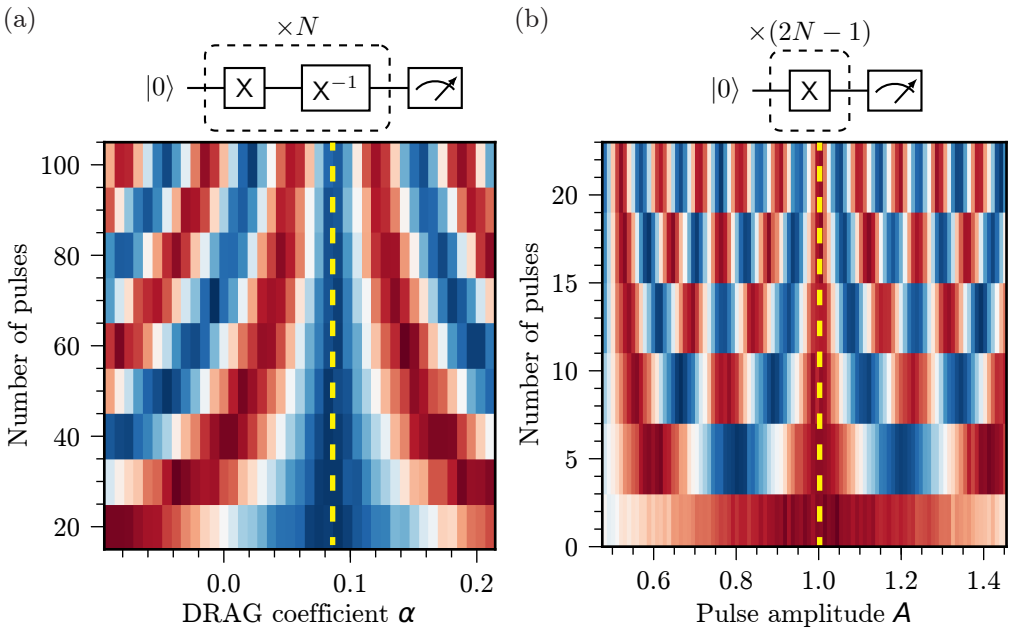

Figure 3.3: *Tune-up for single-qubit gates. Blue corresponds to a high population of the* $|0\rangle$ *state, and red to the* $|1\rangle$ *state. (a) An even number of an* X *gates and its inverse is applied such that phase errors accumulate. The dashed yellow line marks the optimal DRAG coefficient. (b) An odd number of* X *gates is applied such that over- and under-rotations accumulate. The dashed yellow line indicates the optimal pulse amplitude.*

coefficient, the differences being that the sequence does not alternate the signs of the gates and that we vary the pulse amplitude instead of α. An error in the amplitude leads to an accumulation of over- or under-rotations, which causes oscillations between $|1\rangle$ and $|0\rangle$. In Fig. 3.3 (b), we measure this sequence as a function of A and the number of X gates to find the optimal amplitude. Here, the optimal amplitude is $A \approx 1.01$.

3.1.4 Benchmarking of single-qubit gate fidelities

To benchmark gates, we could directly use the fidelity metric defined in Eq. (3.7). However, accurately measuring the unitary of an applied gate is no easy task. Quantum state and process tomography [5] are only as good as the gates used to perform the tomography, which is an issue when those gates are the same as the ones we want to measure the fidelity of. To overcome this issue, the randomized benchmarking protocol was developed [157]. Randomized benchmarking can deconvolve gate errors from state-preparation and measurement (SPAM) errors and estimate an accurate average gate fidelity equal to the definition in Eq. (3.7).

Randomized benchmarking is performed by applying sequences of random unitaries, but with a net effect equal to the identity matrix, to one or more qubits. It is enough to perform measurements in the computational basis instead of a full tomographic reconstruction of the quantum state. This more straightforward measurement protocol is sufficient since, due to the randomness, any coherent gate errors turn into incoherent ones when averaged over many different randomizations. Since the net effect should be the identity, a system initialized in its ground state would always come back to the ground state by the end of the sequence if no errors occurred. By measuring how the probability of occupying the ground state decreases with the number of random unitaries, the average unitary fidelity can be estimated separately from any SPAM errors.

To simplify and make the protocol more efficient, we restrict the unitaries to be drawn from the Clifford group [147]. The single-qubit Clifford group contains all unitaries which transforms between the six Pauli operator eigenstates $(|0\rangle, |1\rangle, |+\rangle, |-\rangle, |+i\rangle, |-i\rangle)$. There are 24 unitaries that have that property. We decompose those unitaries into a combination of maximally three $\sqrt{\mathsf{X}}$, X, $\sqrt{\mathsf{Y}}$, and Ygates. On average, a single-qubit Clifford is composed of 1.875 gates. For the exact decompositions used, see the Supplementary material of Ref. [40].

In practice, we perform Clifford-based randomized benchmarking by the following steps:

1. Draw m random unitaries from the Clifford group.
2. Calculate the combined unitary of those unitaries and the inverse of that combined unitary.
3. Initialize the qubit in $|0\rangle$ and apply the randomly selected unitaries followed by the inverse and measure the probability of the $|0\rangle$ state.
4. Repeat steps 1–3 k times for different randomizations and take the average of $P_{|0\rangle}$.
5. Repeat steps 1–4 for different values of m and record $P_{|0\rangle}$.

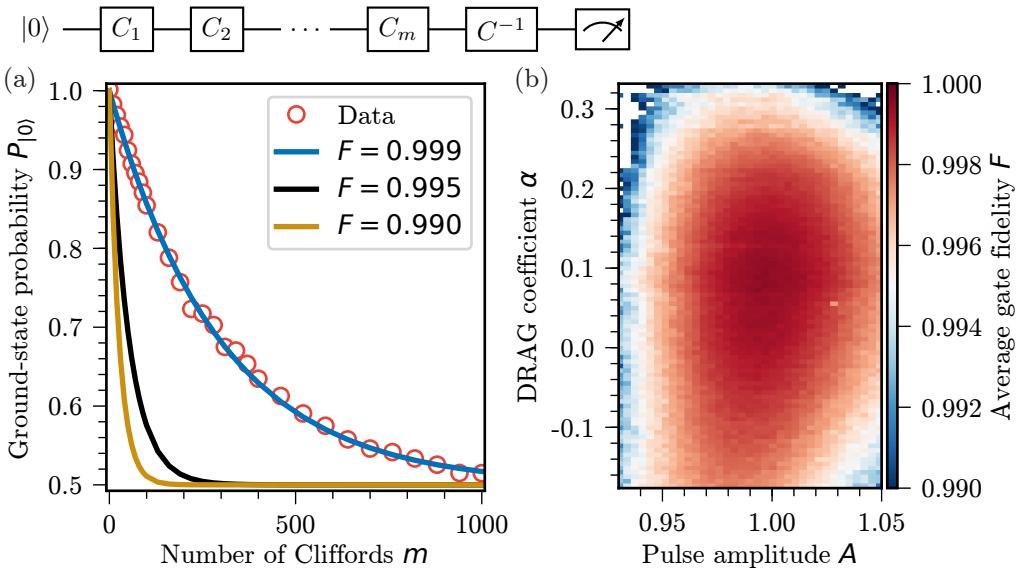

Figure 3.4: *Single-qubit randomized benchmarking. m Cliffords followed by an inverse gate are applied to a qubit. (a) Probability of measuring the ground state as a function of the number of random Cliffords. Circles are experimental data, and lines are theoretical curves for three different average gate fidelities. (b) The average gate fidelity for $m = 100$ Cliffords as a function of two pulse parameters. The highest recorded fidelity is 0.9995.*

6. Fit the decaying probability to $ap^m + b$, where a and b are constants related to the SPAM errors, and p is a decay constant.

The decay constant p is related to the error probability per Clifford as $(p+1)/2$. Since there are 1.875 gates per Clifford, the average gate fidelity from randomized benchmarking is

$$F = 1 - \frac{1-p}{2 \cdot 1.875}. \tag{3.15}$$

In Fig. 3.4 (a), we show randomized benchmarking performed with up to 1000 random Cliffords. From the data, we extract an average gate fidelity of 0.999. We also show the sensitivity of the protocol by plotting how the decay would look for lower gate fidelities. At a given value of m, the state probability increases monotonically with F. In Fig. 3.4 (b), we estimate the gate fidelity as a function of the pulse amplitude and DRAG coefficient. Since a and b should not change with the pulse parameters, we measure $P_{|0\rangle}$ at $m = 100$ and estimate the decay constant via $p = [(P_{|0\rangle} - b)/a]^{1/m}$. The fact that $P_{|0\rangle}$ increases monotonically with F allows for a fast estimation between changes in pulse parameters and their effect on the gate fidelity. This estimation can be used in a feedback loop to optimize the pulse parameters numerically [158, 159]. The landscape of the optimization cost function would look like Fig. 3.4 (b).

The benchmarking protocol just described yields the average gate fidelity for $\sqrt{\mathsf{X}}$, X, $\sqrt{\mathsf{Y}}$, and Y. To extract the fidelity of each gate individually (or any single-qubit Clifford

Table 3.1: Interleaved randomized benchmarking fidelities for nine different single-qubit gates.

X	X^{-1}	$\sqrt{\mathsf{X}}$	$\sqrt{\mathsf{X}}^{-1}$	Y	Y^{-1}	$\sqrt{\mathsf{Y}}$	$\sqrt{\mathsf{Y}}^{-1}$	I
0.9988	0.9988	0.9993	0.9994	0.9990	0.9981	0.9978	0.9976	0.9986

for that matter), we use interleaved randomized benchmarking [160]. The process is very similar, with the addition that after each random single-qubit Clifford, the gate of interest is inserted, and the decay constant $p_{\text{interleaved}}$ is extracted. The fidelity of the interleaved gate is

$$F_{\text{interleaved}} = \frac{1 + p_{\text{interleaved}}/p}{2}. \tag{3.16}$$

A typical set of interleaved randomized benchmarking results is found in Table 3.1. We see that there is some variation between the fidelities, even if there is no reason why any particular gate should have a higher or lower fidelity than the others. The variation in gate fidelities is most likely due to fluctuations of the qubit coherence during the benchmarking [Paper B].

We could expect the identity gate to have a better fidelities than the other gates. The identity gate is performed by idling the qubit for a time equal to the pulse length τ. Since the I gate requires no active control, it should not have any coherent errors. Therefore, interleaved randomized benchmarking of the I gate can serve as a metric for the incoherent error rate. However, we see that the I-gate fidelity is lower than for the Y gate, most likely due to coherence fluctuations of the device, making the incoherent errors temporally larger during the benchmarking of the I gate. Another way to benchmark the incoherent-error rate would be with a protocol known as purity benchmarking [161], where the trace of the qubit's density matrix is measured (using quantum state tomography) as a function of the number of Cliffords.

3.2 Two-qubit gates

The implementation of single-qubit gates in transmon qubits is quite straightforward, and all research teams do it more or less identically. On the other hand, there is a large variation in how two-qubit gates are achieved. Both in the physical implementation and the actual gate (e.g., CZ, CNOT, or iSWAP). The reason for this variation is because there is no clear winner. All implementations have their respective advantages and disadvantages. For a comparison between the current performance of different two-qubit gates with superconducting qubits, see the review of Kjaergaard et al. [36].

One of the difficulties in realizing a high-fidelity two-qubit gate is the on-off-ratio problem, that is, to achieve a high ratio between the coupling strength when the coupling should be on and when it should be off. A controllable interaction between the qubits is needed to accomplish a high on-off ratio. For instance, the interaction could be achieved

by capacitive coupling between two qubits, leading to the following Hamiltonian,

$$\frac{\hat{H}}{\hbar} = \frac{\omega_1 \hat{\sigma}_1^z}{2} + \frac{\omega_2 \hat{\sigma}_2^z}{2} + J(\hat{\sigma}_1^+ \hat{\sigma}_2^- + \hat{\sigma}_1^- \hat{\sigma}_2^+), \tag{3.17}$$

where $\hat{\sigma}_i^z$, $\hat{\sigma}_i^+$, and $\hat{\sigma}_i^-$ are the Pauli-Z, and the raising and lowering operators for the i^{th} qubit, respectively, and J is the qubit–qubit coupling strength.

If $\omega_1 = \omega_2$, the evolution of Eq. (3.17) leads to the iSWAP unitary in Eq. (3.4) [162, 163]; however, the issue is that the interaction is always on. To turn off the interaction, the qubit's frequency could be made tunable so that the qubit frequencies can be detuned from each other. To achieve a high ratio between the on and off couplings, the detuning usually needs to be larger than 1 GHz. A qubit that far detuned from its flux sweet spot will have a significant sensitivity to magnetic fields and flux noise. Moreover, while tuning the qubit frequency by several hundreds of megahertz along some trajectory, the qubit will pick up a considerable phase shift during the gate. While this phase can be calibrated and compensated for, any small deviation from the trajectory will quickly degrade the gate fidelity. The trajectory in frequency space might also contain regions where the coherence is severely reduced [136], leading to more incoherent errors during the gate.

A CZ gate can be implemented similarly. Instead of having the two qubits on resonance with each other, one qubit's $|0\rangle$–$|1\rangle$ transition is nearly resonant with the other's $|1\rangle$–$|2\rangle$. If they are precisely on resonance, swapping between the $|11\rangle$ and $|02\rangle$ states occurs. If timed correctly so that all population returns to $|11\rangle$, the $|11\rangle$ state will have accumulated a minus sign, corresponding to a CZ gate [164]. The minus sign on the $|11\rangle$ state is easily understood as two iSWAP gates between $|11\rangle$ and $|02\rangle$. By squaring the iSWAP unitary, Eq. (3.4), both the involved states acquire a minus sign; however, $|02\rangle$ is not a computational state and is therefore not an issue. The drawback of this diabatic implementation is that we temporally use a non-computational state and risk leaving some population there (i.e., leakage) [165]. Moreover, this type of gate suffers from the same off problem as before, and the relaxation of the $|2\rangle$ state is a factor of 2 faster than for the $|1\rangle$ state [166].

An adiabatic implementation of this gate can be made [40, 167] and is used in Paper F. There, an optimized trajectory brings $|11\rangle$ and $|02\rangle$ nearly into resonance, but in a way that no swapping into $|02\rangle$ occurs. Due to the off-resonant interaction between the two states, $|11\rangle$ still acquires a phase-shift that can be made to be equal to π, which is a CZ gate.

Two-qubit gates based on direct coupling are typically fast but have issues with turning off the interaction. To reduce the off coupling, we can insert a detuned resonant circuit (a 'coupler') between the two qubits, as illustrated and shown in Fig. 3.5 (a) and (b). The effective coupling strength between the two qubits is

$$J = \frac{g_1 g_2}{2} \left(\frac{1}{\omega_1 - \omega_c} + \frac{1}{\omega_2 - \omega_c} \right), \tag{3.18}$$

where ω_c is the frequency of the coupler. While the off coupling is vastly reduced, so is also the on coupling.

To turn on the interaction, either the coupler is tuned to a point where J is large so that the qubits are strongly coupled [16, 168, 169], or the coupler is first tuned into the

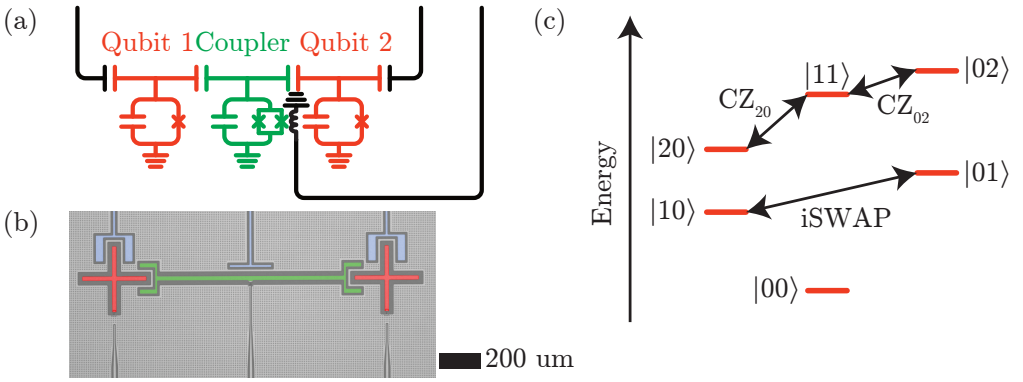

Figure 3.5: *Two fixed-frequency transmon qubits coupled via a frequency-tunable transmon coupler. (a) Electrical circuit of two qubits (red) and one coupler (green). In black are two drive lines and one flux line. (b) Micrograph of an actual device. In addition to the qubits, coupler, and control lines, there are also readout resonators in blue. (c) The energy diagram of the two transmons truncated to their three lowest states. The three possible interactions are indicated by black arrows.*

resonance of one qubit, and then the other [61], or the coupler is parametrically modulated to induce a strong coupling [170, 171]. Alternatively, a microwave drive to either the coupler itself [172] or one of the qubits [173] can be used to induce interactions.

3.2.1 A parametric and diabatic CZ gate

In our work, we have chosen to focus on a parametric two-qubit gate. We use a parametrically-driven tunable coupler to mediate the interaction between two fixed-frequency transmon qubits. We have the ability to implement either an iSWAP, a CZ, or both gates. The corresponding modulation frequencies are shown in Fig. 3.5 (c). Here, and in Paper G, we focus on the implementation of a CZ gate.

The following Hamiltonian describes the two qubits and the coupler in the dispersive regime [174]

$$\frac{\hat{H}}{\hbar} = \frac{\omega_1 \hat{\sigma}_1^z}{2} + \frac{\omega_2 \hat{\sigma}_2^z}{2} + \frac{\omega_c \hat{\sigma}_c^z}{2} + g_1(\hat{\sigma}_1^+ \hat{\sigma}_c^- + \hat{\sigma}_1^- \hat{\sigma}_c^+) + g_2(\hat{\sigma}_2^+ \hat{\sigma}_c^- + \hat{\sigma}_2^- \hat{\sigma}_c^+) \tag{3.19}$$

where $\omega_c(t) = \omega_{c0}\sqrt{|\cos(\pi\Phi(t)/\Phi_0)|}$. By eliminating the coupler, we end up with the same Hamiltonian as for two directly coupled qubits Eq. (3.17), where the coupling strength J given by Eq. (3.18).

To induce an interaction, we should modulate the frequency of the coupler at the correct frequency. In the case of a CZ, the modulation frequency should be the difference between the $|11\rangle$ and $|02\rangle$ states, or $|11\rangle$ and $|20\rangle$ states. Without loss of generality, we choose the former, which is $\omega_2 - \omega_1 + \eta_2$. However, due to the nonlinear relation between coupler frequency and magnetic flux, the coupler and qubit frequencies will shift

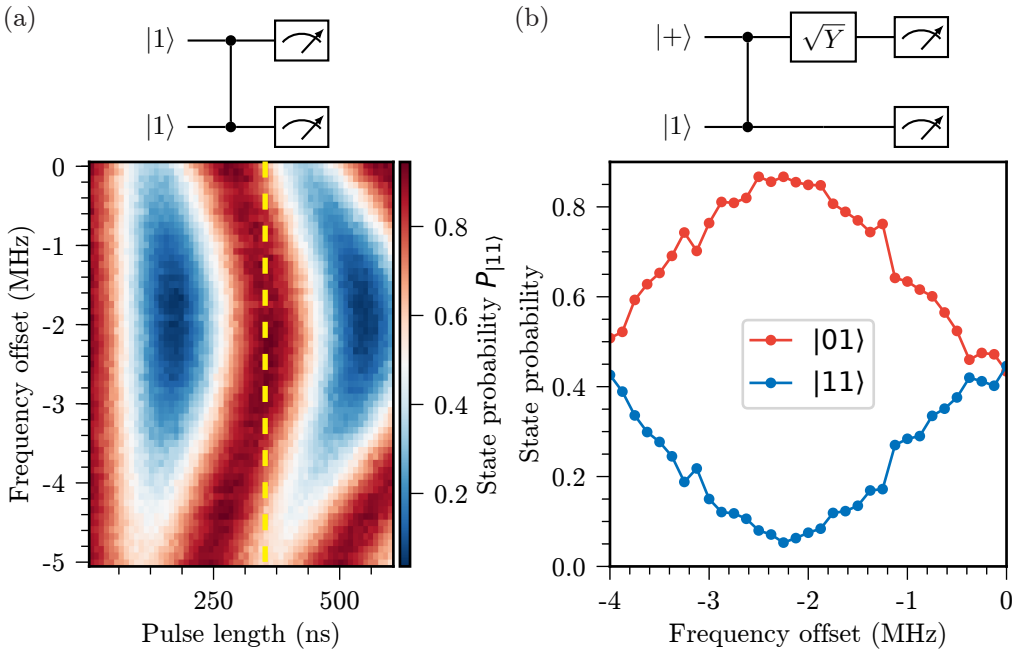

Figure 3.6: *(a) Parametrically induced swapping between* $|11\rangle$ *and* $|02\rangle$*, as measured by the gate squence above the data plot. (b) Gate sequence and data for measuring the acquired phase of the* $|11\rangle$ *state by changing the modulation frequency along the dashed line in (a). In both panels, the state probabilities do not reach 1 since they are limited by the readout fidelities.*

during modulation [170]. We need to compensate for this shift, modifying the modulation frequency to $\omega_2 - \omega_1 + \eta_2 + \xi$, where ξ is the frequency shift due to the modulation.

To characterize the interaction, we initialize the system in $|11\rangle$ and apply a parametric modulation with some frequency and duration. In Fig. 3.6 (a), we plot the probability of measuring the $|11\rangle$ state as a function of modulation duration and frequency offset ξ. We observe swapping from $|11\rangle$ to $|02\rangle$, which exhibits the typical Chevron pattern. Around $\xi = -2\,\text{MHz}$, the oscillation with pulse length has the largest contrast and lowest frequency, which means that the drive is resonant with the transition and that all the population of $|11\rangle$ is transferred into $|02\rangle$ and then back again. This is where we want to operate the CZ gate.

To measure the phase acquired by the $|11\rangle$ state during the conditional-phase gate, we prepare the system in $|01\rangle + |11\rangle = |+1\rangle$, apply the parametric modulation followed by a $\sqrt{\text{Y}}$ gate on the first qubit, as shown in Fig. 3.6 (b). If the $|11\rangle$ state did not acquire any phase, the final system state should be $|11\rangle$. However, if a CZ gate was performed, the final state should instead be $|01\rangle$. In Fig. 3.6 (b), we plot the probabilities of these two states as a function of the modulation frequency offset. Around $-2.2\,\text{MHz}$, we observe that most of the system is in $|01\rangle$, corresponding to a CZ gate.

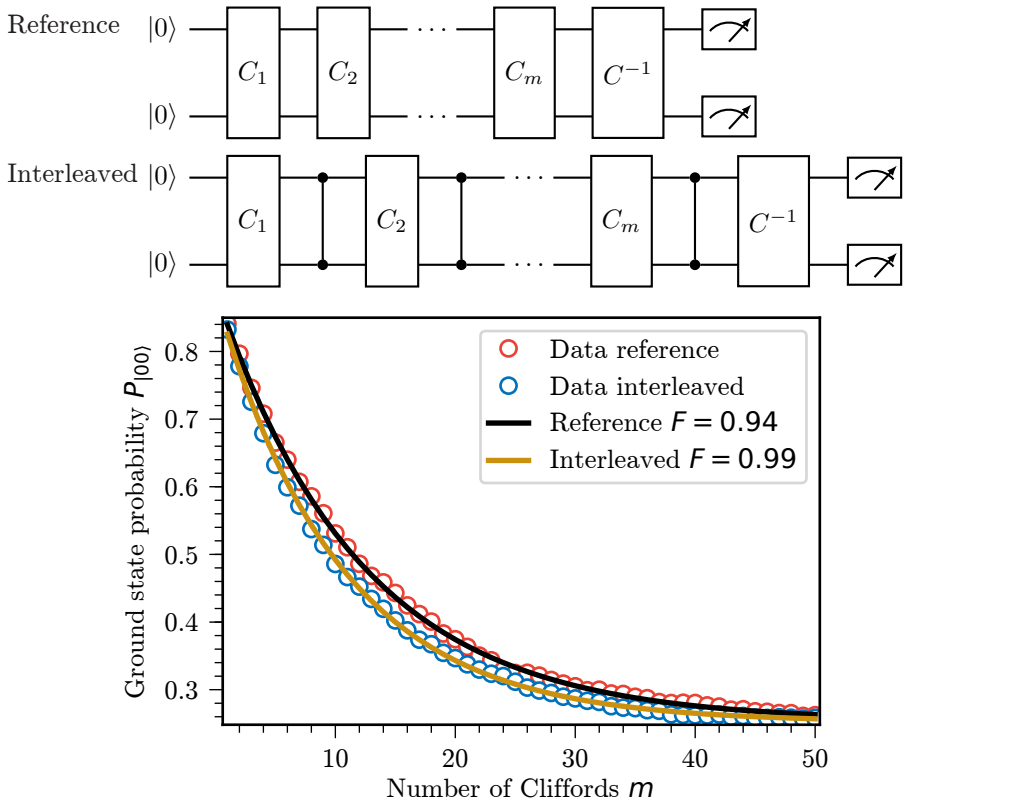

Figure 3.7: *Two-qubit randomized benchmarking of a parametric CZ gate. Red circles with a solid black line show the reference data (only two-qubit Cliffords) with a fidelity of* 94 %*. Blue circles and a solid orange line have CZ gates interleaved between each Clifford, showing a CZ fidelity of* 99 %*.*

During the parametric modulation, the coupler frequency will negatively detune due to the nonlinear relation between flux and frequency. If we Taylor expand the relation, we get a quadratic term. The square of a cosine contains a zero-frequency component, which is the reason for the lower coupler frequency during the gate. Due to the coupling between the coupler and the qubits, the qubits frequencies also become detuned during the gate, making the $|01\rangle$ and $|10\rangle$ states acquire phase shifts. These are deterministic phase shifts, meaning that they are independent of the actual qubit states. Since the phases are deterministic, we can compensate by applying virtual Z gates on each qubit after performing the CZ gate.

3.2.2 Benchmarking of two-qubit gate fidelities

Similarly to the benchmarking of single-qubit gates, we apply the protocol of randomized benchmarking to two-qubit gates. The two-qubit Clifford group consists of four different

subgroups [175]: the single-qubit group (576 elements), the CNOT-like group (5184 elements), the iSWAP-like group (5184 elements), and the SWAP-like group (576 elements). Each group starts with a single-qubit Clifford on both qubits, then follows either a CNOT, an iSWAP, or a SWAP. For the CNOT and iSWAP groups, an additional single-qubit gate is added to both qubits. An iSWAP gate can be decomposed into two CNOTs, and a SWAP into three CNOTs. This gives an average of 1.5 CNOTs per two-qubit Clifford. Note that a CNOT can be decomposed into one CZ gate and two single-qubit gates (i.e., 1.5 CZ gates per two-qubit Clifford [40]). If one had access to both CZ and iSWAP gates, as we do in the parametric scheme, only one two-qubit gate per two-qubit Clifford would be required, on average.

To perform the benchmarking, one randomly draws m two-qubit Cliffords and append the inverse gate at the end that returns the system to $|00\rangle$. Then, we measure the probability of $|00\rangle$ as a function of m and for different randomizations. We fit the resulting decay to the same model as for a single-qubit gate. From p we calculate the average two-qubit Clifford fidelity $(3p + 1)/4$. From just the two-qubit Clifford fidelity, we cannot extract the CZ fidelity since a two-qubit Clifford contains both single- and two-qubit gates.

To extract only the CZ fidelity, we run interleaved randomized benchmarking. After each random Clifford, we insert a CZ gate and then repeat the benchmarking protocol. From the resulting decay constant p_{CZ} we can calculate the CZ fidelity as

$$F = \frac{3}{4}(1 - p_{\mathsf{CZ}}/p), \tag{3.20}$$

In Fig. 3.7, we plot both the reference (no interleaved CZ gates) and the interleaved benchmarking, with a two-qubit Clifford fidelity of 94.0 % and a CZ fidelity of 99.0 %.

A two-qubit fidelity of 99.0 % is close to the state-of-the-art and above the threshold for the error-correcting surface code [13, 40], which means that we could use our gate to implement error correction. However, error correction needs more than two qubits, which our current processor does not have. Nevertheless, two qubits are enough to demonstrate quantum algorithms, and our gate fidelity is more than enough as well. In Chapter 4, we demonstrate two quantum algorithms.

To increase the gate fidelity further, we would need to operate the gate faster. The current gate time of 260 ns is rather slow, which means that we are mainly limited by incoherent errors. We are confident that increasing g_i by a factor of $\sqrt{2}$ will allow us to achieve gate times of 100 ns and fidelities of 99.5 %. There is a trade-off between short gate times and coherent errors due to off-resonant driving of other transitions [174]. Therefore, achieving gate times approaching 10 ns might be challenging.

We are working on increasing the number of qubits and couplers on our processors. It is possible to couple more than two qubits to the same coupler [170], which can be advantageous as it reduces the amount of hardware needed. One issue when scaling to more qubits might be crosstalk between different couplers. Currently, the crosstalk between two SQUIDs can be as high as 10 %; however, this number is generally vastly reduced when moving to a 3D-integrated architecture [16, 42].

3.3 Readout of qubit states

Finally, we turn to the last quantum operation, the measurement of a qubit's state. A measurement is performed using the readout resonator that is coupled to the qubit. As a result of the dispersive coupling between the two, the effective resonator frequency is $\omega_r + \chi\hat{\sigma}^z$, where χ is the dispersive shift, Eq. (1.10). The dispersive shift makes the resonant frequency dependent on the qubit's state, allowing the state of the qubit to be inferred by probing the resonator with a weak microwave signal. Experimentally, a single microwave tone is applied to the feedline at a frequency close to ω_r, and the transmitted signal is measured. Depending on where the resonant frequency is, the transmission coefficient, Eq. (2.13), is different, and the qubit's state can be inferred.

A readout pulse is generated by the same upconversion scheme as a single-qubit pulse. However, here we also need to downconvert the transmitted signal to be able to sample it with an analog-to-digital converter (ADC). The up- and downconversions are done using the same local oscillator [Fig. 3.8 (a)], which provides a stable phase relationship between generation and detection. As for the single-qubit drive, we use an intermediate frequency to avoid having the local-oscillator frequency at the readout-resonator frequency. Otherwise, in the presence of bleedthrough, that would lead to an excess population of the resonator and qubit dephasing, Eq. (2.11). The use of IQ mixers allows us to distinguish between positive and negative intermediate frequencies, essentially doubling the available bandwidth.

Accurate triggering of when the readout should start is crucial. The readout should start immediately after the last quantum gate to minimize any decoherence of the system state. At the end of the last gate, a trigger signal is sent from the qubit AWG to the readout AWG and ADC. The ADC then acquires both channels for a specific time and integrates the signal over this time window, resulting in two voltages I and Q. To compensate for the intermediate frequencies of the readout signal, we use integration weights proportional to $\cos(\omega_{\mathrm{IF}}t)$ and $\sin(\omega_{\mathrm{IF}}t)$. By a proper choice of integration weights, only one of the sidebands of the downconverted signal is measured. The gate sequence and measurement can be repeated multiple times to build up statistics of the qubit state.

It is important to point out that even with no noise on the readout signal, statistics are needed to infer the qubit populations. According to the Copenhagen interpretation of quantum mechanics, a measurement projects a quantum state to one of its eigenstates. So if a qubit is in a superposition state (e.g., $|0\rangle + |1\rangle$), a single measurement would only yield $|0\rangle$ or $|1\rangle$. It is only by statistics (and state tomography) that we could tell that the qubit is, in fact, in a superposition state. In Fig. 3.8 (b), we plot a two-dimensional histogram over the readout I and Q voltages for a qubit prepared in $|0\rangle + |1\rangle$. We see two clouds of almost equal size, each corresponding to one of the two-qubit states.

In Fig. 3.8 (c), we plot the integrated I voltages for a qubit prepared in the ground and excited states. Ideally, each state should have a voltage distribution described by a narrow Gaussian. The width of each Gaussian is proportional to the total system noise, and the separation between is proportional to the signal amplitude and readout duration. Additionally, we see that each distribution is the sum of two Gaussians, with the second one having a smaller magnitude. The origin of the second Gaussian is due to some qubit-state error since the second Gaussian has the same mean value as the primary

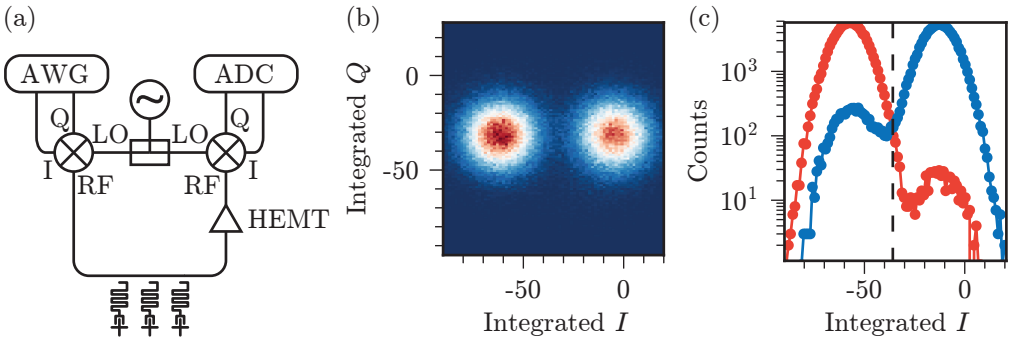

Figure 3.8: *(a) Setup used to generate and detect readout pulses. An arbitrary waveform generator (AWG) is used to generate the readout signal, and an analog-to-digital converter (ADC) is used to detect it. The local oscillator is shared between the up- and downconversion mixers. Here, three readout resonators are coupled to a common feedline for multiplexed qubit readout. (b) 2D histogram of acquired complex readout voltages of a qubit in the* $|0\rangle + |1\rangle$ *state. The left and right clouds correspond to* $|0\rangle$ *and* $|1\rangle$*, respectively. (c) Histogram of measured I voltages with the qubit in* $|0\rangle$ *(red) and* $|1\rangle$ *(blue). The dashed line corresponds to a threshold value used for deciding the qubit state.*

Gaussian of the other state. The second Gaussian of the red trace (qubit ideally prepared in $|0\rangle$) is most likely due to a thermal population of the qubit. The second Gaussian of the blue trace originates from a combination of thermal population, single-qubit gate error, and qubit relaxation before and during the readout, explaining why the second Gaussian is more significant for the blue trace than for the red. By identifying the mean I and Q values for the qubit in its ground and excited state, it is possible to define a threshold in-between those means and label a measurement as a $|0\rangle$ or $|1\rangle$.

A readout fidelity F is defined via the overlap between the measurement distributions with the qubit in the ground respective excited state. This overlap equals to $1 - F = P(|0\rangle||1\rangle) - P(|1\rangle||0\rangle)$, where P is the conditional probability of measuring the wrong state. Ideally, the only contribution to the system noise would be from the vacuum. However, in most practical applications, the semiconducting HEMT amplifier used adds noise that dominates over the vacuum. This causes more overlap between the readout voltage distributions corresponding to the qubit in the ground and the excited state.

To reduce the noise and the width of the Gaussians, we can either use a longer integration time or amplifiers with less noise. We are already using the best HEMT amplifiers on the market with a noise temperature (2 K) approximately ten times that of the Heisenberg limit ($\hbar\omega/2k_\mathrm{B} \approx 150\,\mathrm{mK}$). To decrease the system noise temperature further and approach the quantum limit [176], we could use custom-made superconducting amplifiers based on parametric interactions. For example, a Josephson parametric amplifier (JPA) has been used as the first-stage amplifier to demonstrate single-shot readout [177, 178]. JPAs provide high gain but low bandwidth and saturation power. To combat these drawbacks, the traveling wave parametric amplifier (TWPA) was developed [179, 180]. In Paper C and Paper F, we use TWPAs to boost the signal-to-noise ratio of the readout

signal.

In linear systems, the SNR can be increased by just increasing the signal. Here, that corresponds to increasing the amplitude of the microwave tone used for readout. However, if the amplitude and the number of photons inside the resonator are too large, the approximations performed to the dispersive Hamiltonian Eq. (1.9) breaks down [37], and state mixing and transitions are induced [181]. The critical photon number where state mixing starts to occur is

$$n_{\mathrm{c}} = \frac{\Delta^2}{4g^2} \approx 200, \tag{3.21}$$

for our qubits. Note that Δ and g are strongly dependent on the readout speed needed. For faster readouts than ours, n_{c} quickly approaches 1-10 photons, matching the noise added by the best HEMT amplifiers currently available.

To build a quantum computer, we need more than one qubit. If we were to multiply the readout setup by the number of qubits, that would be very impractical and expensive. One advantage of the transmission configuration [Fig. 3.8 (a)] is that multiple readout resonators can be coupled to the same feedline, allowing for a simultaneous readout of several qubits using frequency multiplexing.

There is a practical limit on the number of resonators coupled to the same feedline. In essence, the frequency spacing between each resonator must be sufficiently large compared to the bandwidth of the resonators (κ) so that there is minimal crosstalk and all resonant frequencies must be within the bandwidth of the detection chain. In our case, that bandwidth is set by the analog-to-digital converter and is 1200 MHz. For example, we could couple ten resonators with a spacing of 120 MHz, which is sufficient for typical κ values.

3.3.1 Tuneup of dispersive readout

In practice, we have to tuneup and calibrate our readout. Three main parameters are associated with a readout pulse: amplitude, frequency, and duration. The amplitude should be high enough to overcome the system noise, but not higher than the critical photon number. The optimal frequency is typically close to ω_r and can be found by maximizing the separation of the Gaussian distributions [Fig. 3.8 (c)] as a function of frequency. The optimal duration is found by analyzing the readout fidelity as a function of the integration window. Too long duration and the qubit relaxes during the readout. Too short duration and the separation error is significant. In Paper G, we do not use any parametric amplifiers. However, we can still reach readout fidelities above 0.9, since the readout can be made relatively long (2.3 µs), well below our relaxation times (several tens of microseconds).

To refine our accuracy of measuring state probabilities in the presence of limited readout fidelity, we collect statistics of the measured qubit population as a function of the rotation angle θ around the X-axis on the Bloch sphere, see Fig. 3.9. Since the measured population increases monotonically with the angle, we can renormalize the populations to account for the limited readout fidelity (95%). We fit the measured probabilities to $\beta_{\langle\hat{\sigma}^z\rangle}\cos(\theta) + \beta_I$. The β parameters allow us to convert between measured state probabilities and the actual expectation values of a qubit. The mathematical relation

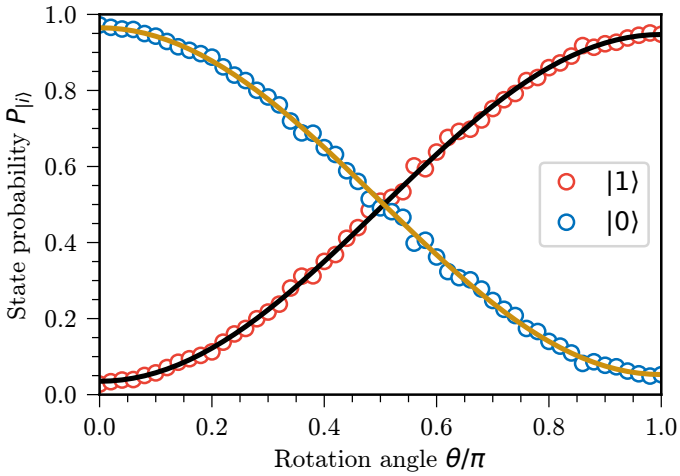

Figure 3.9: *Measured state probabilities as a function of the rotation angle θ. The $|1\rangle$ probability increases monotonically with the rotation angle, allowing us to renormalize the state probabilities.*

between state probabilities and expectation values is

$$\begin{pmatrix} P_{|0\rangle} \\ P_{|1\rangle} \end{pmatrix} = \begin{pmatrix} \beta_I & \beta_{\hat{\sigma}^z} \\ \beta_I & -\beta_{\hat{\sigma}^z} \end{pmatrix} \begin{pmatrix} \langle I \rangle \\ \langle \hat{\sigma}^z \rangle \end{pmatrix}. \tag{3.22}$$

This conversion is crucial, as it allows for quantum state tomography in the presence of limited readout fidelities [153]. Moreover, for more than one qubit, it is still enough to calibrate the β parameters separately for each qubit. The system β matrix is then given by the tensor product between all individual β matrices.

3.3.2 Readout using a degenerate parametric oscillator

In Paper E, we demonstrate a slightly different readout protocol, which was first proposed in Ref. [62]. It still relies on the dispersive coupling between a qubit and a resonator, but instead of probing the resonator using a weak signal, we parametrically modulate the readout resonator. Under the right conditions, the parametric pumping yields strong parametric oscillations (see Chapter 5).

The excitation of oscillations in a degenerate parametric oscillator is sensitive to the pump parameters: the pump amplitude must exceed an instability threshold, and the frequency must be close to $2\omega_r$. By coupling a superconducting qubit, such as the transmon, to a parametric oscillator (Fig. 3.10), the dispersive shift can move the resonant frequency so that the second criterion is no longer fulfilled. The oscillator will then remain in its ground state. In this way, the ground state of the qubit can be mapped to an oscillating state and the excited state to a silent state, or vice versa. Since the oscillating state has a large number of photons in the resonator, there will be a significant

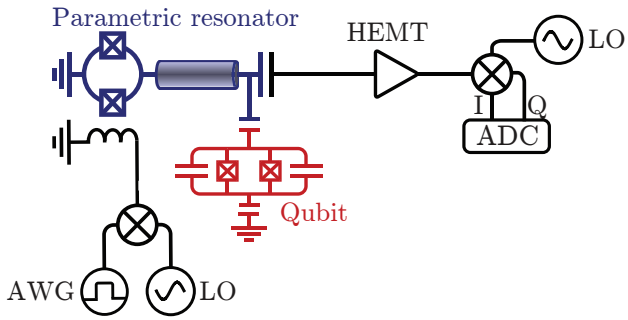

Figure 3.10: *Experimental setup for the parametric oscillation readout. A transmon qubit (red) is coupled to a parametric oscillator (blue). The parametric readout pump is controlled by a microwave mixer and an arbitrary waveform generator. The outgoing field is amplified, downconverted, and sampled by an ADC.*

measurement contrast between the outputs corresponding to the qubit states, and therefore a high signal-to-noise ratio.

In Fig. 3.11 (a) and (b), the output power of the oscillator is presented as a function of pump detuning and amplitude when the qubit is prepared in the ground and excited states, respectively. For the ground state, we see a typical parametric region around zero detuning between half the pump frequency and the resonator frequency. For the excited state, the region has moved to lower frequencies due to the dispersive shift χ. However, a faint region is still seen in the original position. This is due to qubit relaxation back to its ground state before the readout, or that the excitation pulse was not successful.

It is clear that if the parametric oscillator is operated close to zero detunings, the finite-amplitude state is mapped to $|0\rangle$ and the zero-amplitude state to $|1\rangle$. If performed at negative detunings instead, the situation is reversed. The optimal operating point is found by extracting the readout fidelity at each point in the parametric regions.

For the operating point with the highest readout fidelity, 81.5%, the readout voltage distributions are analyzed in detail in Fig. 3.11 (c). With the qubit in the ground state, the oscillator is mostly in its zero amplitude state (the Gaussian at the origin). In the excited state, the oscillator has a finite amplitude and shows a combination of three Gaussians. The two outer ones are the oscillating and the most probable states. The reason for two Gaussians at finite voltages is that the oscillator has two degenerate states, which are shifted by π radians. The presence of the central peak with the qubit in $|1\rangle$ is due to qubit errors. For example, when the qubit relaxes before the start of the readout, the excited state measurement shows three Gaussians instead of two. If the qubit instead relaxes during the readout, it introduces an asymmetry in the Gaussians.

Performing a careful analysis of the acquired histograms can give insights into the origins of the fidelity loss. Due to the probabilistic nature of the effects limiting the readout fidelity, we use a Monte-Carlo based simulation of the experiment. The simulation starts with the qubit initialized to its ground state. Several random numbers between 0 and 1 are generated and compared to the probability of specific processes, deciding if that process will happen during one cycle of the experiment. The final qubit state is recorded,

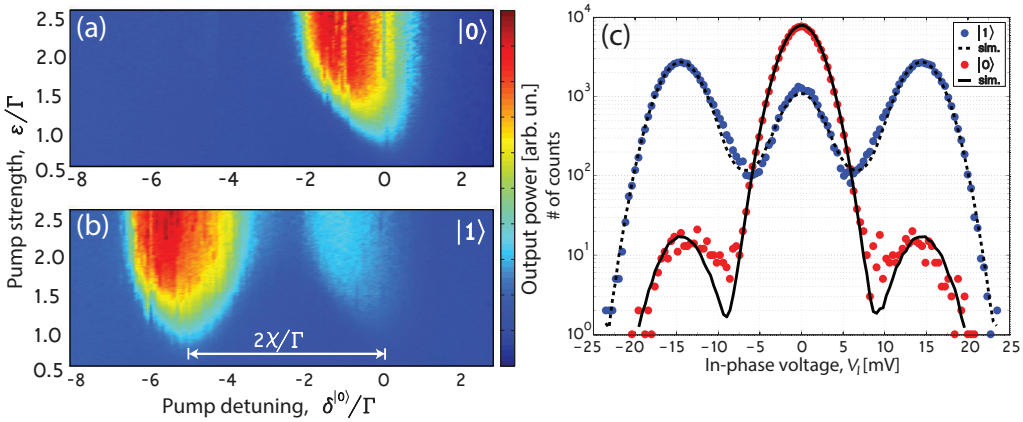

Figure 3.11: *Measurements of the output power of the parametric oscillations as a function of normalized pump parameters, when the qubit is prepared in the ground state (a) and excited state (b). Note that in the context of parametric oscillators, ϵ is the pump amplitude, and Γ is the resonator bandwidth. (c) Histogram of measured I quadrature voltages, together with simulation results.*

and a readout voltage is drawn from the associated Gaussian distribution. The simulation is then repeated, but with different random numbers. Over time, a distribution of qubit states and readout voltages are built up.

There are five sources of error taken into account in the simulation:

1. The qubit is not in its ground state initially. The Boltzmann distribution gives this probability for a certain qubit frequency and temperature.

2. The qubit is not excited by the π pulse, which is modeled as a gate fidelity.

3. The qubit relaxes before readout. The probability of decay after a given delay time t_d between the excitation pulse and the start of the readout is given by $1 - e^{-t_d/T_1}$.

4. The qubit relaxes during readout. The probability of decay in a given time interval Δt is given by $e^{-\Delta t/T_1}\Delta t/T_1$.

5. The parametric oscillator switches during readout. The probability is given by the switching rate.

The switching rate and relaxation time T_1 are measured separately, while the gate fidelity and the temperature of the qubit are treated as fitting parameters.

In Fig. 3.11 (c), the result of the simulation and the experimental data are compared. The side peaks in the red data set correspond to the qubit initially being in the excited state due to thermal excitation, and the temperature T in the simulation is adjusted to match the data. The central peak in the blue data set corresponds to either relaxation before the readout, an unsuccessful π pulse, or the qubit being in the excited state initially and then de-excited by the π pulse. Since T_1 is measured separately, and T is already

fixed, only the gate fidelity is adjusted in the simulation. These two adjustments give an overall good agreement with experimental data and simulation. The loss of fidelity associated with each error channel is:

- 16.1% to relaxation and preparation errors.
- 1.1% to thermal population.
- 1.2% to phase switching of the parametric oscillator.
- 0.002% due to overlap between Gaussians.
- In total, 18.4% of error.

From this analysis, we conclude that the fidelity due to the readout method itself is 98.7% and limited by the phase switching. What determines the switching rate is still unknown to a large extent, but two things could be done to improve it. Either, the readout duration is decreased so that the probability of switching during the readout is lower, or the oscillations can be injection-locked to one of the states by applying a small resonant signal, as done in Ref. [182] and Paper K.

The qubit induced readout errors are mainly due to the low coherence time compared to the readout duration. This experiment was carried out with an older qubit with coherence times on the order of single microseconds. Our more modern qubits have coherence times at least ten times longer than the one used here, and should, therefore, reduce errors 3 and 4 substantially. A more in-depth description and analysis of the parametric readout is found in the Ph.D. thesis of Philip Krantz, Ref. [183]. Also, this readout method is not used in any of the other appended papers.

CHAPTER IV

Two quantum algorithms and their implementation

So far, we have discussed how to design and fabricate qubits with high coherence and how to control them with high fidelity. We now turn to the actual quantum information *processing* part of this thesis. This chapter summarizes the results of Paper F and Paper G, which are about two different quantum algorithms. We cover the basics of each algorithm and how they are implemented on actual hardware.

Current quantum processors are far from large or fault-tolerant, which means that factoring large integers via Shor's algorithm is not possible at the moment. The textbook algorithms, which have provable speed-up over the best known classical algorithms, all require thousands of error-free qubits. In parallel with working towards building such systems, it is important to investigate heuristic algorithms that can run on current or near-term hardware. Such algorithms are referred to as noisy intermediate-scale quantum (NISQ) algorithms [15]. While we cannot prove their speed-up, or expect an exponential speed-up over classical algorithms, they might provide polynomial or even just constant speed-up. In many industrial applications, 10, 20, or 50 % speed-up would still be of great value. In this chapter, we will see one NISQ algorithm, the quantum approximate optimization algorithm [184], and one non-NISQ, density matrix exponentiation.

The quantum processors used here contain two qubits. While it is trivial to simulate such a processor, and even up to tens of qubits, on a classical computer, it is still important to run algorithms on actual hardware. First of all, it demonstrates that there is nothing fundamentally wrong with quantum mechanics and the algorithm. Second, it sparks interest in the research of new quantum algorithms and how they could be used to solve real-world problems, for example, in the airline industry [185, 186]. Third, small-scale algorithms can act as benchmarks for how a specific quantum processor is performing. Finally, implementing algorithms on small and noisy processors forces us to come up with smart and creative ways to utilize the limited resources available (e.g., efficient gate compilation and reuse of qubits), which we will greatly benefit from even when larger processors become available.

4.1 Density matrix exponentiation

Density matrix exponentiation (DME) is a unitary transformation, $e^{-i\rho\theta}$, of a density matrix ρ, and an angle θ. To classically implement DME would require a full tomographic reconstruction of ρ. This, in turn, requires $\mathcal{O}(2^{2n}/\epsilon^2)$ copies of ρ, where n is the number of qubits in the density matrix and ϵ is the desired precision [187].

If ρ represent some data matrix, then DME combined with quantum phase estimation [5] can be used to extract the dominant eigenvalues and eigenvectors of ρ [188], which provides exponential speed-up for principle component analysis [189] and quantum machine learning [190]. Furthermore, if ρ is a large entangled state, DME efficiently reveals its entanglement spectrum without having to perform full quantum state tomography [191].

It has been shown that a quantum version of DME asymptotically outperforms any tomographic strategy to implement $e^{-i\rho\theta}$ [188, 189]. DME, as implemented directly on quantum hardware, requires only $\mathcal{O}(\theta^2/\epsilon)$ copies of ρ, resulting in an exponential reduction in the resource requirements [189].

Another way of seeing density matrix exponentiation is as a quantum instruction set. By changing the state of an instruction qubit, we achieve different operations on a target qubit. The opposite would be a classical instruction set, where the operation of the target qubit is determined by a classical computer and the actual waveforms sent to the qubit. A quantum instruction set has a variety of applications, including quantum simulation [188] and quantum emulation [192].

A practical protocol for implementing the quantum version of DME utilizes partial SWAP gates between a target state σ and N copies of an instruction state ρ. We will denote this protocol DME_N. The SWAP unitary exchanges the states of two qubits according to $\mathsf{SWAP}(\sigma \otimes \rho) = \rho \otimes \sigma$. The DME_N algorithm relies on the observation that:

$$\mathrm{Tr}_\rho \left[e^{-i\mathsf{SWAP}\delta} \sigma \otimes \rho e^{i\mathsf{SWAP}\delta} \right] = \sigma - i\delta[\rho, \sigma] + \mathcal{O}(\delta^2) = e^{-i\rho\delta} \sigma e^{i\rho\delta} + \mathcal{O}(\delta^2).$$

That is, after a partial SWAP, defined as $\delta\mathsf{SWAP} \equiv e^{-i\mathsf{SWAP}\delta}$ [193], σ undergoes unitary evolution of the form $e^{-i\rho\delta}$ (up to first order in δ). Applying this technique serially by introducing a new copy of ρ for each of N $\delta\mathsf{SWAP}$ operations with $\delta = \theta/N$ implements the operator

$$\mathsf{DME}_N(\rho, N, \theta) \to e^{-i\rho\theta} + \mathcal{O}\left(\frac{\theta^2}{N}\right). \tag{4.1}$$

For intuition, if ρ is a single-qubit pure state, DME_N rotates σ about the axis defined by the Bloch sphere vector of ρ, with an angle θ. The more copies of ρ we have, the more we can reduce the algorithmic error.

For small δ, the state of the target and instruction qubits are nearly unaffected by a $\delta\mathsf{SWAP}$. In this case, a measurement of ρ in its eigenbasis approximately 'resets' the qubit by projecting it onto one of the eigenstates. As the effect of $\delta\mathsf{SWAP}$ is small, the probability of projecting into the same eigenstate every time is high, and the qubit will, therefore, stay in that eigenstate, an effect that is known as the quantum Zeno effect [194]. So, by measuring the instruction qubit after each $\delta\mathsf{SWAP}$, we can use the same instruction qubit to the next round of $\delta\mathsf{SWAP}$; thus, removing the need for N physical copies of ρ.

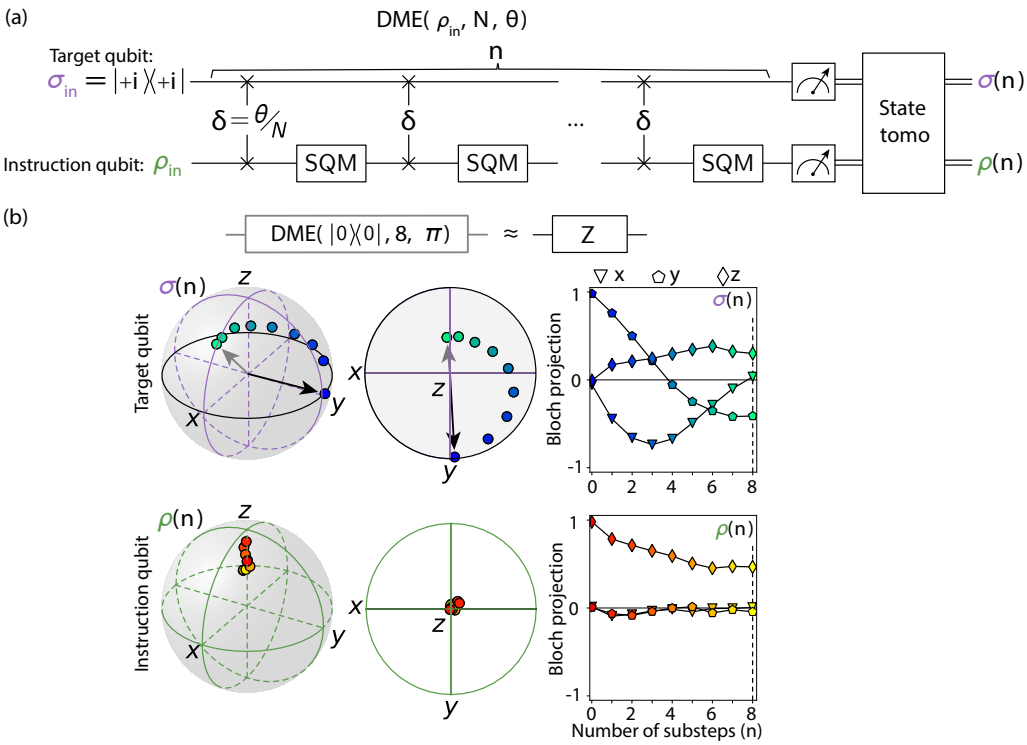

Figure 4.1: *(a) Two-qubit DME implementation using simulated quantum measurements (SQM) to approximately reinitialize the instruction qubit to* ρ_{in}*. The substep parameter* n *is stepped from 0 to* N*. We perform* n *rounds of* $\delta SWAP + SQM$*, measure the two-qubit density matrix, and trace over each subsystem to extract the individual density matrices* $\sigma(n)$ *and* $\rho(n)$*. (b) Substeps of* $DME(|0\rangle\langle 0|, 8, \pi)$*, corresponding to a* Z *gate on the target qubit at* $n = N$*. Black lines are guides to the eye.*

This resource-efficient protocol, which we denote DME, trades a moderate increase in algorithmic error for a significant reduction in the required number of qubits. Because the reset of ρ is approximate, it introduces additional error on the same order of magnitude as the algorithmic error in DME_N:

$$\mathsf{DME}(\rho, N, \theta) \quad \rightarrow \quad e^{-i\rho\theta} + \underbrace{\mathcal{O}(\tfrac{\theta^2}{N})}_{\text{Finite } N} + \underbrace{\mathcal{O}(\tfrac{\theta^2}{N})}_{\text{Reset}}. \tag{4.2}$$

The reset protocol does not depend on us knowing the outcome of the measurement. It is the effect of the measurement itself that performs the approximate reset. Instead of performing the measurement, which can take several microseconds, we use a novel probabilistic operation, which we call a simulated quantum measurement (SQM). SQM mimics a measurement by constructing a dephasing channel via a random application of either an identity (I) or a Z gate:

$$-\boxed{\mathsf{SQM}}- = \begin{cases} \mathsf{I} & \text{with } p = 0.5 \\ \mathsf{Z} & \text{with } p = 0.5 \end{cases} \tag{4.3}$$

Here, we focus on instruction states in the z-basis $|0\rangle, |1\rangle$; however, SQM can be extended also to the x- and y-bases. When averaged over many randomizations, SQM is identical to a measurement whose outcomes are ignored, while only requiring the time span of a single-qubit gate (tens of nanoseconds).

The circuit for our implementation of DME is shown in Fig. 4.1 (a). We interleave δSWAP operations with SQM on the instruction qubit, executing many instantiations of the circuit with random choices of I and Z for each SQM, and averaging together the outcomes. We implement a δSWAP using single-qubit and CZ gates.

In Paper F, we extensively characterize and benchmark our DME implementation. Here, in Fig. 4.1 (b), we demonstrate the working principle. We choose a target state $\sigma = |+i\rangle\langle+i|$, an instruction state $\rho = |0\rangle\langle 0|$, number of steps $N = 8$, and total phase $\theta = \pi$, yielding $\mathsf{DME}(|0\rangle\langle 0|, 8, \pi)$. Since the instruction is z-polarized, this operation should perform a Z gate on the target qubit. We interrupt the algorithm after n substeps of $\delta\mathsf{SWAP} + \mathsf{SQM}$ and perform quantum state tomography, averaging together all SQM randomizations to produce a single density matrix. We see that $\sigma(n)$ undergoes partial rotation at each step about an axis defined by ρ; $\rho(n)$ maintains its z-polarization but undergoes depolarization due to gate errors. This demonstrates and validates the fundamental underlying principle of DME.

DME is not a NISQ algorithm, as we cannot allow for any errors to happen during the execution, and the number of qubits needs to be large to be able to solve interesting problems. Nevertheless, its demonstration is important. Just as factoring 15 using a quantum computer is not useful by itself [195], it is an important demonstration of the underlying theory and technology. As quantum computers grow in size and become fault-tolerant, we hope that DME can prove its usefulness in quantum simulation and machine learning.

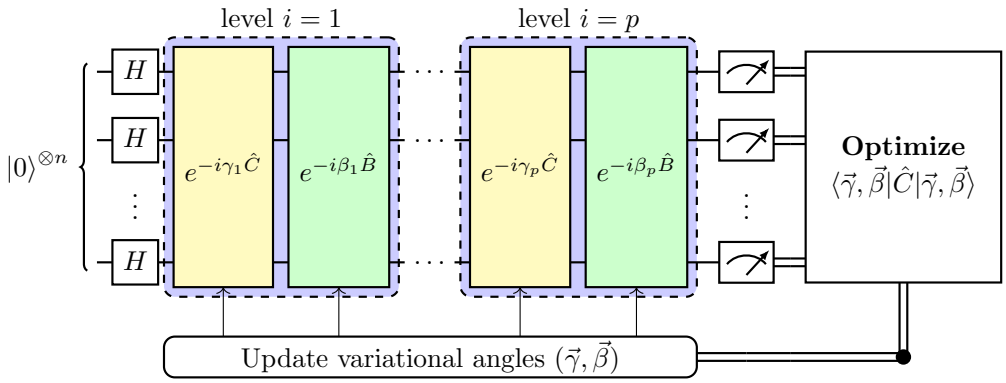

Figure 4.2: *The quantum approximate optimization algorithm (QAOA) for a problem specified by an Ising Hamiltonian $\hat{C}$. An alternating sequence of two Hamiltonians ($\hat{C}$ and $\hat{B}$) is applied to an equal superposition of n qubits. After measurement of the qubit states, a cost is calculated, which a classical optimization algorithm minimizes by varying the angles $\vec{\gamma}, \vec{\beta}$.*

4.2 Quantum approximate optimization

The quantum approximate optimization algorithm (QAOA) is a heuristic algorithm for solving NP-complete optimization problems [184, 196, 197]. It is not believed that quantum computers will be able to provide an exponential speed-up for NP-complete problems, as that would imply that all NP problems would be efficiently solvable on a quantum computer. However, quantum computers might still bring polynomial speed-up (cf. Grover's search algorithm [198]). It has been shown that QAOA should provide a polynomial speed-up to the solution of specific problems encoded in a quantum Hamiltonian [199, 200].

QAOA is executed on both a classical and a quantum computer, making it a hybrid algorithm. The quantum part consists of a circuit with p levels and some variational angles, where better approximations to the solution of the encoded problem are generally achieved with higher p. A classical optimizer is then used to minimize some cost function measured on the quantum processor.

All NP-complete problems can be formulated in terms of finding the ground state of an Ising Hamiltonian [201],

$$\hat{C} = \sum_{i=1}^{n} h_i \hat{\sigma}_i^z + \sum_{i<j} J_{ij} \hat{\sigma}_i^z \hat{\sigma}_j^z, \tag{4.4}$$

where h_i and J_{ij} are real coefficients. QAOA aims at finding this state by applying two non-commuting Hamiltonians, $\hat{B}$ and $\hat{C}$, in an alternating sequence (with length p) to an

equal superposition state of n qubits (visualized in Fig. 4.2),

$$\left|\vec{\gamma},\vec{\beta}\right\rangle = \prod_{i=1}^{p}\left[e^{-i\beta_i\hat{B}}e^{-i\gamma_i\hat{C}}\right]\left(\frac{|0\rangle+|1\rangle}{2}\right)^{\otimes n}, \tag{4.5}$$

where γ_i and β_i are variational angles. The second Hamiltonian is a transverse field (mixing) Hamiltonian defined by

$$\hat{B} = \sum_{i=1}^{n}\hat{\sigma}_i^x, \tag{4.6}$$

The ground state of Eq. (4.4) corresponds to the lowest-energy state. Therefore, we aim to find the minimum value of the energy expectation value of Eq. (4.5),

$$F(\vec{\gamma},\vec{\beta}) = \left\langle\vec{\gamma},\vec{\beta}\right|\hat{C}\left|\vec{\gamma},\vec{\beta}\right\rangle = \sum_{i=1}^{n}h_i\langle\hat{\sigma}_i^z\rangle + \sum_{i<j}J_{ij}\langle\hat{\sigma}_i^z\hat{\sigma}_j^z\rangle. \tag{4.7}$$

We use F as a cost function for a classical optimization algorithm aiming to find the optimal varitional angles $\vec{\gamma}^*,\vec{\beta}^*$. F is evaluated by repeatedly preparing and measuring $|\vec{\gamma},\vec{\beta}\rangle$ on a quantum processor. For a high enough p, $|\vec{\gamma}^*,\vec{\beta}^*\rangle$ is equal to the ground state of $\hat{C}$ and the answer to the optimization problem [184]. However, for algorithms executed on real hardware without error correction, noise will inevitably limit the circuit depth, implying that there is a trade-off between algorithmic errors (too low p) and gate errors (too high p). The angles $\vec{\gamma}^*,\vec{\beta}^*$ themselves are not interesting, as long as they yield the lowest-energy state. This gives some robustness against coherent gate errors since any over- or under-rotations can be compensated for by a change in the variational angles [202].

In Paper G, we report on using our superconducting quantum processor to demonstrate QAOA with up to $p = 2$. We solve four instances of the NP-complete exact-cover problem. Let us consider one of the four instances where $n = 2$, $J_{12} = 1/2$, $h_1 = -1/2$ and $h_2 = 0$. It is easy to check that the corresponding ground state is $|10\rangle$. For $p = 1$, we apply a simple grid (61×61) search of $\beta_1, \gamma_1 \in [0, \pi[$ while recording 5000 measurements of each qubit. From these, we calculate $\langle\sigma_i^z\rangle$, $\langle\sigma_1^z\sigma_2^z\rangle$, the cost function F, and the occupation probability for each of the four possible states. The grid search allows us to explore the shape of the optimization landscape, which may bring an important understanding of the difficulty of finding global minima for black-box optimizers.

In Fig. 4.3 (a), we show the measured cost function. Due to the normalization of h_i and J_{12}, the ground state corresponds to $F = -1$. We observe that the cost function never reaches below -0.5. To achieve costs approaching -1, additional levels ($p > 1$) are needed. Moreover, the existence of a local minimum around $\gamma_1 \approx \beta_1 \approx 3\pi/4$ could cause difficulties for optimizers trying to find the global minimum.

In Fig. 4.3 (b), we take linecuts along the dashed lines in Fig. 4.3 (a) and benchmark our measured cost functions and state probabilities against those of an ideal quantum computer without any noise. We see excellent agreement between measurement and theory: the measured positions of each minimum and maximum are aligned with those of the theory, consistent with low coherent-error rates. In addition, we observe excellent agreement

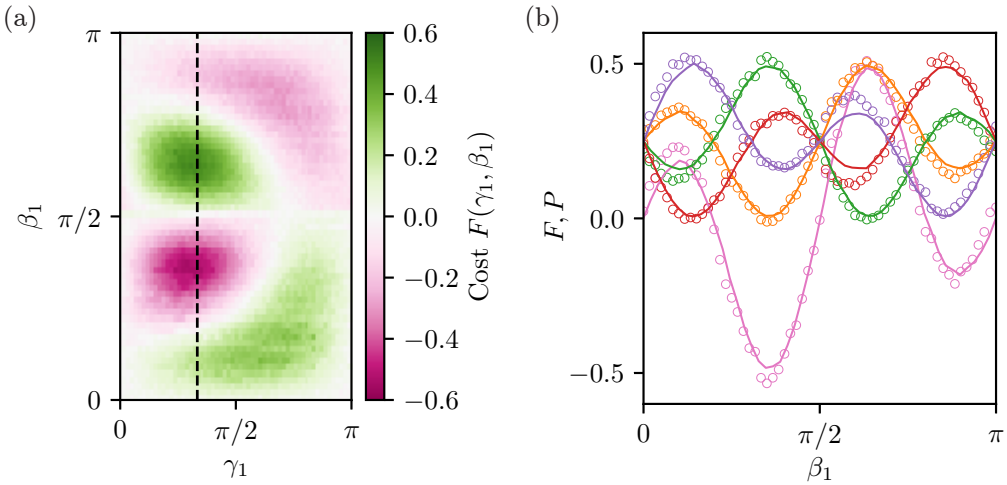

Figure 4.3: *(a) Cost function $F(\vec{\gamma}, \vec{\beta})$ for QAOA with $p = 1$ and $n = 2$. Each experimental data point is evaluated from the average of 5000 measurements on our quantum processor. (b) A comparison between experiment (open circles) and theory (solid lines). Each color corresponds to either a state probability or the value of the cost function F. The linecuts are taken at the vertical dashed line in (a). The theory curves are calculated assuming an ideal quantum processor.*

between the absolute values at the minima and maxima, indicating low incoherent-error rates as well.

Even with high gate fidelities, high algorithmic fidelity is not guaranteed. Randomized benchmarking gives the average fidelity over a large number of random gates, which transforms any coherent errors into incoherent ones. For real quantum algorithm circuits, the gates are generally not random. Therefore, any coherent errors can quickly add up and yield algorithmic performance far lower than expected from randomized benchmarking fidelities alone [203].

To increase the success probability, we add an additional level ($p = 2$). For $p > 1$, a grid search to map out the full landscape becomes unfeasible due to the many parameters (equal to $2p$). Therefore, we instead use a black-box optimizer to find the optimal variational angles. Here, we use Bayesian optimization with Gaussian processes (BGP). We choose BGP due to its ability to find global minima.

We evaluate the optimizer performances by running 200 independent optimizations with random starting values $(\vec{\gamma}, \vec{\beta} \in [0, \pi[)$. We set a threshold for convergence at $F < -0.95$ and count the number of converged optimization runs as well as the number of calls to the quantum processor (function calls) required to converge. We also record the success probability of measuring the problem solution ($P_{|10\rangle}$). For the optimizer presented in Fig. 4.4, we converged to the ground state 61.5% of the time, with an average number of function calls of 44.

We study the individual optimization trajectories in Fig. 4.4. For each run, we plot the

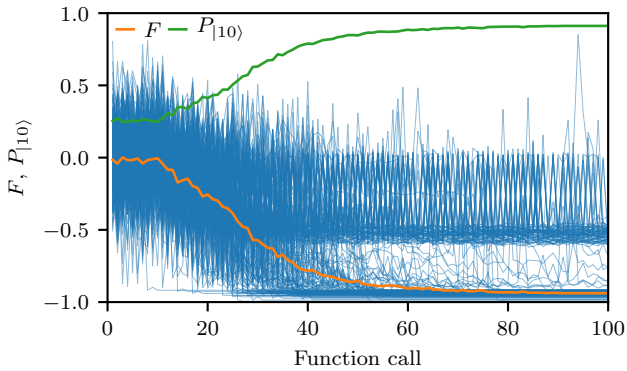

Figure 4.4: *QAOA using $p = 2$ and Bayesian optimization with Gaussian processes. We run the optimization 200 times with random starting parameters. Plotted as blue lines are the individual optimization trajectories for the cost F. In orange and green are F and the success probability ($P_{|10\rangle}$) averaged over the converged runs.*

cost F and the success probability $P_{|10\rangle}$, along with their average over all converged runs. The first function calls are random points, which explains why there is no improvement in the beginning. After that, the cost function decreases rapidly, accompanied by an expected increase in $P_{|10\rangle}$. We see an indication of a local minimum at $F \approx -0.5$, and that the optimizer sometimes manages to escape from this minimum, showcasing the strength of Bayesian optimization.

At the end of the optimization, the highest recorded probability of generating the correct state is 96.6%. The success probability is limited by imperfect gates (we have verified that an ideal quantum computer and $p = 2$ can achieve $P_{|10\rangle} = 1$). We compare our measured success probability to what we would expect from the randomized-benchmarking fidelities. The quantum circuit, together with the fidelities for each gate, predicts a total fidelity of 96.3%, in good agreement with the measured fidelity, considering experimental uncertainties (e.g., fluctuations in qubit coherence and gate fidelities). Note that $p = 3$ would not yield a higher success probability, since adding more gates would lower the total fidelity further (predicted to be 94.2%).

Finally, let us discuss success probabilities. The important question is, how high does it need to be? If we were to guess solutions randomly, we would have a success probability of $1/2^n$. Of course, there are better classical algorithms than random guessing; however, since these are heuristic algorithms, it is difficult to quantify how high success probability QAOA needs to outperform a classical computer. As an example, even if the success probability of measuring the correct state is only 5%, we could measure 100 instances and still attain a probability higher than 99% of finding the correct state at least once in those 100 measurements. We can therefore use QAOA to generate a shortlist of possible solutions, which then can be checked efficiently on a classical computer. The fact that we can tolerate errors by re-running the algorithm, and that we potentially only need 50 to 100 qubits to solve something that a classical computer cannot, is why QAOA is a NISQ

algorithm.

4.3 Gate compilation

There are two points of gate compilation: to translate an arbitrary sequence to gates that are available on the hardware, and to minimize the number of gates used. Since the number of gates that we can run while maintaining a high algorithmic fidelity is limited, it is essential to compile any quantum algorithm into as few gates as possible. However, the gate compilation itself should not be exponentially costly, as that would remove any possible speed-up of the quantum algorithm. The gates in the available gate set might not have the same fidelity, and a compiler should also minimize the number of low-fidelity gates.

In this section, we will see how the gate compilation of DME was implemented. DME utilizes partial SWAP gates. A full SWAP can be implemented by [5]

$$\mathsf{SWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = \quad . \tag{4.8}$$

From this construction, a partial δSWAP interaction is generated as

$$\delta\mathsf{SWAP} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & (1+e^{i2\delta})/2 & (1-e^{i2\delta})/2 & 0 \\ 0 & (1-e^{i2\delta})/2 & (1+e^{i2\delta})/2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = \quad , \tag{4.9}$$

where

$$\delta = \mathsf{CZ}_\delta = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{-i\delta} \end{pmatrix} = \quad \tag{4.10}$$

is a partial CZ gate. This decomposition together with Eq. (4.9) would allows us to rely on just CNOT gates (together with single-qubit gates), instead of having to tune up high fidelity versions of CZ_δ for each value of δ [204]. However, this decomposition introduces two CNOTgates for each CZ_δ gate, making it a total of four CNOT gates per δSWAP, which is an overhead in the circuit depth.

To optimize the δSWAP implementation further, we utilize the fact that any two-qubit unitary can be decomposed into a circuit with maximally three CNOT and four single-qubit gates [5, 205]:

$$U = \quad , \tag{4.11}$$

where $\mathsf{R}_{i,j}$ is a single-qubit gate acting on qubit i at moment j in the circuit. So, for each value for δ in δSWAP, we can numerically calculate the rotation axes and angles for $\mathsf{R}_{i,j}$. This is also referred to as Cartan's KAK decomposition [206].

Finally, we turn the CNOT gates into our physically available two-qubit gate, the CZ gate, by using the identity

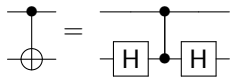

(4.12)

In the context of Eq. (4.11), we can absorb the Hadamard gates into the neighboring single-qubit gates $\mathsf{R}_{i,j}$.

To construct the $\mathsf{DME}(\rho, N, \theta)$ circuit, we append N copies of the compiled δSWAP gate using $\delta = \theta/N$, interleaving SQM on qubit 2 (the instruction qubit, ρ) to emulate the effect of measurements. Rows 1 and 2 in Fig. 4.5 shows the generic structure and gate decomposition of our implementation of DME. The final layer of single-qubit gates in the δSWAP at step n can be recompiled together with the SQM and the first layer of single-qubit gates in the δSWAP at step $n + 1$. We slice out these three layers (Row 2 in Fig. 4.5) of single-qubit gates, recompile them into a single layer (Row 3 in Fig. 4.5), and reinsert them (Row 4 in Fig. 4.5).

Our compilation enables us to achieve high algorithmic fidelity at significant circuit depth since it relies upon a restricted set of gates that are readily characterized and numerically optimized. In particular, the final compiled circuit has a regular structure (each CZ is followed by exactly one layer of single-qubit gates), amenable to generic tuneup protocols for reducing coherent error buildup.

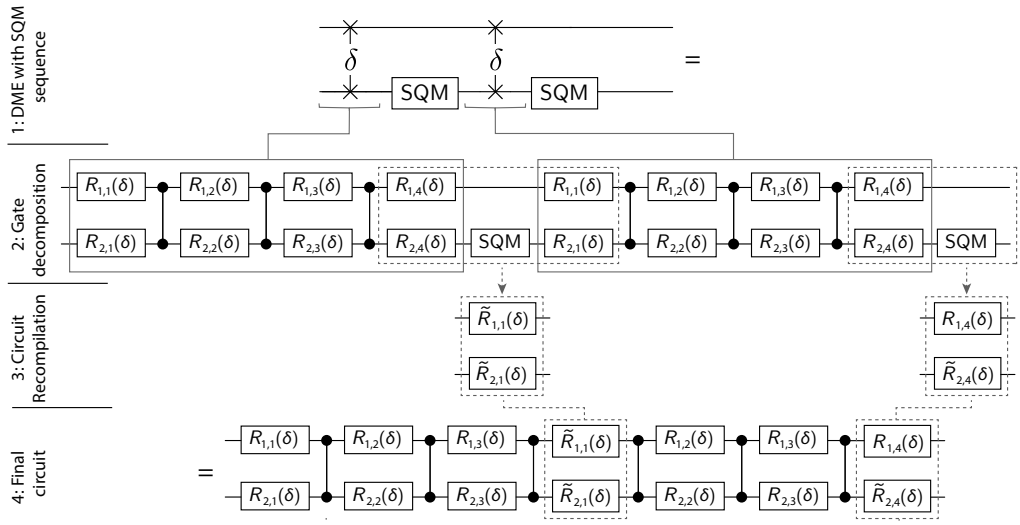

Figure 4.5: ***Row 1***. *Two steps of the density matrix exponentiation algorithm implemented using partial SWAP operations and the simulated quantum measurement (SQM) gate.* ***Row 2***. *Decomposing each δSWAP according to Eq.* (4.11). *Each substep at this step requires 8 layers of gates (7 for δSWAP decomposition and 1 for SQM).* ***Row 3***. *The three layers of single-qubit gates stemming from the the end of the δSWAP of step n, followed by SQM, and the first layer of single-qubit gates in δSWAP of step $n+1$ can be recompiled into a single layer.* ***Row 4***. *The recompiled gates are reinserted into the algorithm result in the optimal structure of exactly one CZ gate, followed by a single layer of single-qubit gates.*

CHAPTER V

Generation of microwave photon states

This chapter summarizes our work on the creation of different microwave states. We use two different methods to achieve this, parametric pumping of microwave resonator (Paper H–Paper L), and controlled exciation an decay of quantum bits (Paper M).

Quantum information can not only be processed by qubits (a system with only two levels). It is also possible to use harmonic oscillators with an infinite number of levels, so-called continuous-variable quantum computing [65, 66]. However, the typical excitation of a harmonic oscillator is the coherent state, which resembles the kind of excitation seen in pendulums, guitar strings, and other common examples of harmonic oscillators. To perform universal quantum computing with continuous variables, something more than just coherent states are needed. Typically, non-Gaussian states are used; such as Fock states [50], photon subtracted coherent states [207, 208], or cat states [71, 209].

If the states are Gaussian (e.g., coherent), it can still be possible to perform universal quantum computing. In that case, either non-Gaussian evolution (i.e., non-quadratic Hamiltonians) or detection (e.g., single-photon detectors) is needed. If the states, evolution, and detection are all Gaussian, the evolution can be efficiently simulated on a classical computer [210]. In this chapter, we will see how non-quadratic Hamiltonians can be achieved using parametric modulations of nonlinear microwave resonators, and a proposal for a single-photon generator in the microwave domain.

5.1 Multimode resonators

The resonators used in this work are quarter-wavelength resonators made of pieces of coplanar waveguides that are open in one end and shorted to ground in the other. Coplanar waveguide resonators are of the distributed type, which means that higher harmonics also exist. As long as the voltage profile in the resonator fulfills the given boundary condition, the voltage can have an arbitrary number of nodes inside the piece of a transmission line, as illustrated in Fig. 5.1 (a). For resonators with nonlinearities from embedded

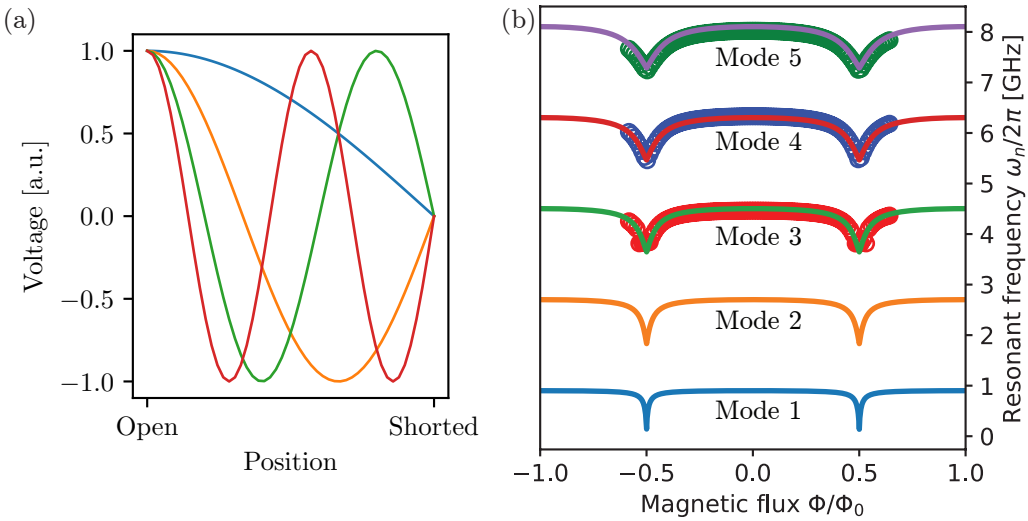

Figure 5.1: *A multimode resonator. (a) The voltage profiles for the four lowest modes of a quarter-wavelength resonator. (b) Resonant frequencies of the five lowest modes, as a function of magnetic flux through the SQUID loop. These frequencies are calculated for a resonator with* $\omega_{\lambda/4}/2\pi = 919\,\mathrm{MHz}$, $\gamma_l = 0.02$, *and* $\gamma_c = 0.009$. *Open circles are experimental data, showing good agreement with the theoretically calculated frequencies.*

Josephson junctions, the harmonics exist at $\omega_n \approx (2n-1)\omega_1$, where $n = 1, 2, 3, ...$ is the mode number. This relation is only approximate since the phase drop across the junctions modifies the mode frequencies differently for each mode, creating a slightly anharmonic spectrum. The spectrum of a tunable multimode resonator can be found by solving a dispersion equation [76, 211],

$$k_n \tan k_n = \frac{|\cos(\pi\Phi/\Phi_0)|}{\gamma_l} - k_n\gamma_c, \tag{5.1}$$

where $k_n = \pi\omega_n/2\omega_{\lambda/4}$, and γ_c (γ_l) are the ratios between the capacitance (inductance) of the SQUID and the resonator. In Fig. 5.1 (b), the five lowest modes calculated using Eq. (5.1), together with experimental data for three of the modes, are seen as a function of magnetic flux through the SQUID loop, Φ. Each mode of the resonator is characterized by its resonant frequency ω_n, external loss rate Γ_{n0}, total loss rate Γ_n, and Duffing nonlinearity α_n. In relation to the previously discussed coupling rate between a resonator and a transmission line, κ, we have that $\kappa = 2\Gamma_{n0}$.

The resonators used for the parametrically pumped circuits are fabricated slightly differently from the high-coherence quantum bits seen previously in this thesis. Here, we used an older fabrication process where the superconductor for the resonators is niobium instead of aluminum. A thin film of niobium (80-90 nm) is sputtered onto a substrate using a near-UHV, DC magnetron sputtering system. The deposited niobium covers the entire surface of the wafer. Subsequently, the niobium is patterned using either an electron-beam or laser lithography system. The pattern is transferred to the niobium

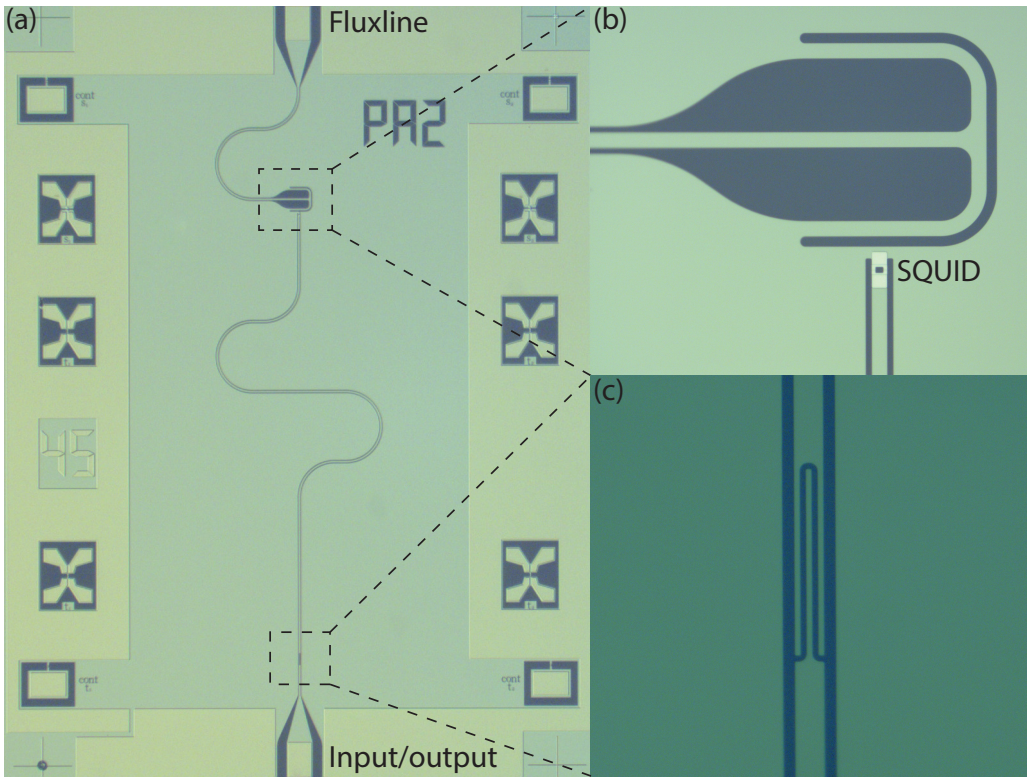

Figure 5.2: *Micrographs of a typical tunable resonator. (a) The full sample with an on-chip flux line at the top, and input/output port with a coupling capacitor at the bottom. The resonator is meandered to fit in a smaller chip. (b) The SQUID together with the inductively coupled flux line. (c) The coupling capacitor between the resonator and the transmission line.*

using an inductively-coupled reactive ion plasma of NF_3. The Josephson junctions are still made out of aluminum.

5.2 Parametric pumping of nonlinear resonators

Parametric oscillations are well-studied phenomena, with applications in amplification, quantum optics, and quantum information processing. They can occur as a parameter of a system, such as the resonant frequency, is modulated or "pumped" by an external field. A nonlinearity of the system can then transfer power from the pumping frequency to two frequencies known as signal and idler. When the signal falls within a resonance, and when the pump amplitude exceeds an instability threshold, parametric oscillations occur.

The first observation of self-sustained parametric oscillations was done in a second-order nonlinear ($\chi^{(2)}$) optical cavity, by pumping the cavity at twice its resonant frequency (de-

generately) [74]. When the pump amplitude exceeded an instability threshold, oscillations were observed at half the pump frequency. Pumping at the sum of two resonant frequencies (nondegenerately) leads to correlated oscillations in both modes [212]. Entanglement between the two modes was demonstrated [213], and also extended to include the pump mode, yielding a three-mode entanglement [214].

In the microwave domain, frequency-tunable superconducting resonators based on Josephson junctions are versatile tools. Two parallel Josephson junctions form a superconducting quantum interference device (SQUID), which acts as a magnetic-flux tunable inductance, therefore enabling frequency modulation via inductively coupled microwaves [45]. When a frequency-tunable resonator is driven parametrically, below the threshold amplitude, it provides amplification at the quantum limit [72, 215, 216] and entangled two-mode squeezed states [68, 70]. Above the threshold amplitude, parametric oscillations occur [75].

Another but related phenomenon is subharmonic oscillations. These oscillations occur when a nonlinear resonator is driven strongly. The strong drive is also referred to as a current pump since a drive result in current flowing in the resonator with the same frequency as the drive (in contrast to a flux pump, which does not induce any current at the pump frequency).

In a quantum picture, the process that underlies parametric and subharmonic oscillations is a decay of a single photon into two, three, or more photons. For instance, the Hamiltonian for a degenerately pumped resonator is [77]

$$\frac{\hat{H}}{\hbar} = -\frac{\alpha_n}{2}(\hat{a}_n^\dagger \hat{a}_n)^2 - \frac{\epsilon}{2}(\hat{a}_n^2 + \hat{a}_n^{\dagger 2}) = -\frac{\alpha_n}{2}\left(\hat{a}_n^{\dagger 2} + \frac{\epsilon}{\alpha_n}\right)\left(\hat{a}_n^2 + \frac{\epsilon}{\alpha_n}\right) + \frac{\epsilon^2}{2\alpha_n}, \tag{5.2}$$

where $\hat{a}_n$ is the annihilation operator for the n^{th} mode. The term $\hat{a}_n^{\dagger 2}$ creates a pair of photons in the mode.

By remembering that a coherent state is the eigenstate of $\hat{a}_n$, it is easy to see that the eigenstates of Eq. (5.2) are two coherent states with equal amplitude but opposite phases, $|\pm i\sqrt{\epsilon/\alpha_n}\rangle$. Therefore, the steady-state of a parametric oscillator is a coherent superposition of the two coherent states.

These kinds of states are also known as cat states, as they are a superposition of two "macroscopic" states [71]. This analysis generalizes to higher-order terms as well, where superpositions of N coherent states can be created by Hamiltonians on the following form [217]

$$\frac{\hat{H}}{\hbar} = -\frac{\alpha_n}{2}(\hat{a}_n^\dagger \hat{a}_n)^N - \frac{\epsilon}{2}(\hat{a}_n^N + \hat{a}_n^{\dagger N}). \tag{5.3}$$

There is one caveat. The analysis above relies on the resonator having no loss of photons. However, to measure and utilize the states, we need to connect the resonator to some kind of measurement apparatus, which will introduce single-photon loss. In the presence of this loss, the coherence of the superposition will be lost, and the resonator steady-state will be a mixed state of the two coherent states. Puri et al. [217] showed that by introducing a two-photon loss process, the coherence of the cat state could be maintained even in the presence of single-photon loss.

In the presence of loss and couplings to external drives B_n, the full system can be described by Langevin equations. Wustmann and Shumekio derived such equations for

a variety of different parametric pumping schemes in Refs. [62, 76, 77]. In general, it is possible to derive quantum Langevin equations for $\hat{a}_n$; however, in the case of large oscillation amplitudes, the states are almost classical, and the equation can be written in terms of a classical amplitude A_n instead. We will see examples of such Langevin equations in later sections.

As the internal field couples into the measurement line, it creates an output amplitude C_n, which is related to the external drive and resonator mode amplitudes as

$$C_n = B_n - i\sqrt{2\Gamma_{n0}}A_n. \tag{5.4}$$

The normalization of the fields is such that $|A_n|^2$ is equivalent to the average number of photons in the n^{th} mode; hence $|B_n|^2$ and $|C_n|^2$ are in- and outgoing photon rates.

5.3 Measurement techniques

To measure the outgoing field C_n from a resonator, a vector digitizer is used to downconvert the signal using heterodyne mixing to an IF frequency of 187.5 MHz, which is then sampled with an analog-to-digital converter (ADC) at 250 MS/s. After digitization, an onboard FPGA filters and processes the signal by digital down conversion to zero frequency, as well as decimating it to an effective sampling frequency [218]. It is important to choose the effective sampling rate high enough, so all signals of interest are captured, but not much higher than that to minimize the amount of noise in the measurement. From the digitizer, the quadratures of the complex voltage, $I_n(t)$ and $Q_n(t)$, are transferred to a computer for further processing. The total power is calculated as $P_n = \langle I_n(t)^2\rangle + \langle Q_n(t)^2\rangle$, where the voltages are assumed to be over a $50\,\Omega$ resistor.

One complication when measuring power instead of amplitude is that the power of noise, $X(t)$, doesn't average to zero (i.e., $\langle X(t)\rangle = 0 \not\Longrightarrow \langle X(t)^2\rangle = 0$). This implies that a measurement of the power of parametric oscillations will also include the power of the noise. However, by assuming that the noise is constant in time, the noise can be measured separately and subtracted from subsequent measurements. The measured output power can then be translated to output field intensity via $|C_n|^2 = (P_n - P_{\text{noise}})/G\hbar\omega_n$, where P_{noise} is the system noise power and G is the gain of the system.

For simultaneous detection of two modes far separated in frequency (as in Paper K and Paper L), two vector digitizers are used. The digitizers use two different, but phase-locked, local oscillators. This ensures phase stability between the digitizers, which is essential when correlating signals between the modes. One might ask, how can there even exist a phase relationship between two signals with different frequencies? The answer is that the phase relation only exists when the frequencies are connected by a third frequency, which acts as a clock (e.g., a parametric pump at the sum of two mode frequencies, $\omega_p = \omega_n + \omega_m$). It is also crucial to start the data acquisition for the digitizers simultaneously. This is achieved by a clock module, which distributes a common trigger signal to both digitizers in cables of equal length.

5.4 Gain and loss calibrations

Calibration of the system loss and gain is needed for proper measurements of the outgoing fields C_n, as well as for proper normalization of the incoming fields B_n. It is enough to know either the loss or the gain since the other can be inferred by using that a resonator is fully reflecting far-off resonance. Subsequently, the number of photons inside the mode can be calculated using Eq. (5.4).

The gain is calibrated using a shot noise tunnel junction (SNTJ) [219, 220] connected to a microwave switch at the mixing chamber of the refrigerator. The SNTJ is a tunnel junction made of a normal metal–insulator–normal metal. When a current is applied through the junction, shot noise is produced. The amount of shot noise depends on the resistance of the junction, and the current through the junction. Moreover, there is also thermal, vacuum, and amplifier noise. By measuring the total noise as a function of the current through the junction, the contribution of the different noise processes can be extracted, together with the gain between the SNTJ and the detector. For practical reasons, the designed value of the SNTJ resistance is $50\,\Omega$. For more information about the SNTJ and how it is used, see Ref. [79].

In Paper L, we do not use an SNTJ, but instead, we use the Josephson junctions of the device as the setup allows us to apply currents through the junctions. This makes the calibration very precise since we can calibrate the gain referenced to the junctions instead of the microwave switch. For the project in Paper L, which required small error bars, a calibration using an SNTJ would yield too large uncertainty.

5.5 Degenerate parametric oscillations

Parametric pumping at twice a resonant frequency is referred to as degenerate pumping. This is effectively three-wave mixing, where a signal and an idler are at the same frequency. In a quantum picture, it can be seen as a split of one pump photon with frequency ω_p into two equal signal photons in mode n of frequency $\omega_n = \omega_p/2$, implying that there is always an even number of photons produced at ω_n. When the pump strength overcomes an instability threshold set by the mode loss rate Γ_n, the resonator starts oscillating at ω_n, even in the absence of an external drive at that frequency.

In Ref. [62], the equation of motion for the mode amplitude A_n for degenerate parametric pumping in a rotating frame at $\omega_p/2 = \omega_n + \delta$ is derived as

$$i\dot{A}_n + \left(\delta + \alpha_n |A_n|^2\right) A_n + \epsilon A_n^* + i\Gamma_n A_n = 0, \tag{5.5}$$

where $\delta = \omega_p/2 - \omega_n$ is the pump detuning, ϵ is the parametric pump strength, α_n is the Duffing parameter, and A_n^* denotes the complex conjugate of A_n. The term $\alpha_n |A_n|^2$ is effectively increasing δ, equivalent to a lower resonant frequency. This is the Duffing, or Kerr, effect due to the cubic nonlinearity of the current-phase relation of the SQUID.

In the steady state, $\dot{A}_n = 0$, Eq. (5.5) has three solutions

I. The trivial solution $A_n = 0$. However, it is only stable for $\epsilon < \Gamma_n$ or $|\delta| > \sqrt{\epsilon^2 - \Gamma_n^2}$.

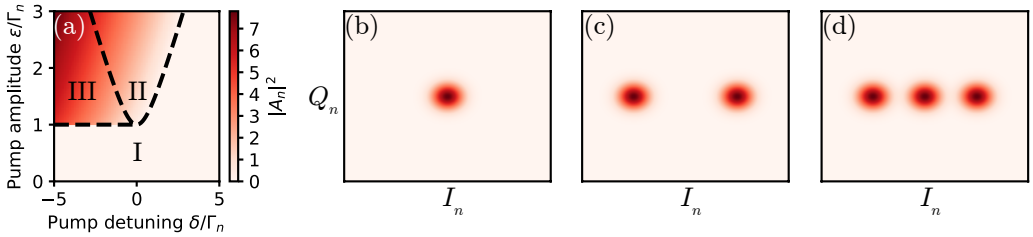

Figure 5.3: *(a) Theoretical steady-state intensities $|A_n|^2$ of a degenerate Josephson parametric oscillator as a function of pump detuning δ and amplitude ϵ. Inside region I, only the ground state is stable (b). In region II, only the excited state is stable (c). Finally, in region III, boths states are stable (d). (b-d) Phase-space distributions for the three different stability regions.*

II. The excited state solution $|A_n|^2 = \left(-\delta + \sqrt{\epsilon^2 - \Gamma_n^2}\right)/\alpha_n$. This solution is stable for $\epsilon \geq \Gamma_n$ and $\delta \leq \sqrt{\epsilon^2 - \Gamma_n^2}$.

III. The third solution has the same amplitude and stability as 2, although its phase is shifted by π radians.

The regions of stability can be visualized by plotting the excited state amplitudes, see Fig. 5.3 (a). In the absence of nonlinearities, the oscillation amplitude would diverge and increase indefinitely. However, due to the Kerr effect, δ increases as the amplitude $|A_n|$ increases, which stabilizes the oscillations.

The excited state of the degenerate parametric oscillator has two π shifted solutions, see Fig. 5.3(c). Switching between the two states is possible [221] and has been observed [75]. It has been proposed that one could create a coherent superposition of the two states [217], forming a so-called Schrödinger's cat state, which is a resource in quantum information processing [209]. The frequency of the parametric oscillations is always centered around $\omega_p/2$, and the switching rate between the states gives the frequency width.

In Paper I, we demonstrate degenerate parametric oscillations in a doubly pumped half-wavelength resonator with one SQUID in each end. This configuration allows for both 'vibrating' and 'breathing' pumping schemes, depending on the relative phase between the pumps for the two SQUIDs. We observe that the threshold for exciting parametric oscillations depends on the relative phase, with a minimum for a breathing mode. For the vibrating mode, the threshold grows to infinity (or at least larger than what we can experimentally measure).

5.6 Nondegenerate parametric oscillations

Nondegenerate pumping refers to when the pump frequency equals the sum of two resonant frequencies. Then, the signal and idler are separated into two different modes, and the pump frequency is $\omega_p = \omega_n + \omega_m + 2\delta$, where $m > n$, and δ is the pump detuning. In the

quantum picture, the photons are again created in pairs, so that if a photon is detected in mode n, it is certain that there is a photon also in mode m. It is, therefore, easy to understand that there could be entanglement in such a system.

Nondegenerate parametric oscillations have shown both two- and three-mode entanglement in the optical regime. The microwave nondegenerate parametric oscillations demonstrated here have an important qualitative similarity to the optical ones, namely a continuous degeneracy of the oscillator state. This degeneracy should lead to large fluctuations under the effect of vacuum noise, and as a result of this, the state might be non-Gaussian [222]. A non-Gaussian state is necessary for universal quantum computing with continuous variables [223].

Similarly to the degenerate case, a quantum Hamiltonian (for the quantum operators $\hat{a}_n$ and $\hat{a}_m$) and a Langevin equation (for the classical mode fields A_n and A_m) can be derived in the corresponding rotating frames $\omega_{n,m} + \delta$ [76]

$$\frac{\hat{H}}{\hbar} = -\frac{\alpha_n}{2}(\hat{a}_n^\dagger \hat{a}_n)^2 - \frac{\alpha_m}{2}(\hat{a}_m^\dagger \hat{a}_m)^2 - 2\alpha_\mathrm{g}(\hat{a}_n^\dagger \hat{a}_n \hat{a}_m \dagger \hat{a}_m) - \epsilon(\hat{a}_n \hat{a}_m + \hat{a}_n^\dagger \hat{a}_m^\dagger), \tag{5.6}$$

and

$$\begin{aligned} i\dot{A}_n + (\zeta_n + i\Gamma_n)A_n + \epsilon A_m^* = 0, \\ i\dot{A}_m + (\zeta_m + i\Gamma_m)A_m + \epsilon A_n^* = 0, \end{aligned} \tag{5.7}$$

where n and m correspond to two different modes, and ζ_n and ζ_m are effective detunings equal to

$$\begin{aligned} \zeta_n = \delta + \alpha_n |A_n|^2 + 2\alpha_\mathrm{g} |A_m|^2, \\ \zeta_m = \delta + \alpha_m |A_m|^2 + 2\alpha_\mathrm{g} |A_n|^2. \end{aligned} \tag{5.8}$$

For simplicity, the geometric means of the dampings $\Gamma_\mathrm{g} = \sqrt{\Gamma_n \Gamma_m}$, and Duffing terms $\alpha_\mathrm{g} = \sqrt{\alpha_n \alpha_m}$, are introduced. The second terms in Eq. (5.8) are again the Duffing or Kerr effects. Similarly, the third terms are equivalent to cross-Kerr effects. Both the Kerr and cross-Kerr lead to larger effective detunings ζ_n, resulting in lower resonant frequencies, as A_n increases.

In the steady state, $\dot{A}_n = \dot{A}_m = 0$, Eq. (5.7) can be solved to produce the following expressions for the excited state intensities $|A_n|^2$ and $|A_m|^2$:

$$|A_n|^2 = \frac{2\Gamma_m(\delta_\mathrm{th}(\epsilon) - \delta)}{R}, \quad |A_m|^2 = \frac{2\Gamma_n(\delta_\mathrm{th}(\epsilon) - \delta)}{R}, \tag{5.9}$$

where R is a constant $R = \alpha_n \Gamma_m + \alpha_m \Gamma_n + 2\alpha_\mathrm{g}(\Gamma_n + \Gamma_m)$, and the pump-dependent threshold detuning δ_th is given by

$$\delta_\mathrm{th}(\epsilon) = \frac{\Gamma_n + \Gamma_m}{2}\sqrt{\frac{\epsilon^2}{\Gamma_\mathrm{g}^2} - 1}. \tag{5.10}$$

The finite steady-state amplitudes is only stable for $\epsilon \geq \Gamma_\mathrm{g}$ and $\delta \leq \delta_\mathrm{th}$. The ground state $A_n = 0$ is stable for $\epsilon < \Gamma_\mathrm{g}$ and $|\delta| > \delta_\mathrm{th}$, thus three different regions exist in δ-ϵ space:

I. For $\epsilon < \Gamma_g$ or $\delta > \delta_{th}$ the ground states $A_n = A_m = 0$ is stable.

II. For $\epsilon \geq \Gamma_g$ and $|\delta| \leq \delta_{th}$ the excited states $|A_n|$ and $|A_m|$ are stable.

III. For $\epsilon \geq \Gamma_g$ and $\delta < -\delta_{th}$ both states are stable.

These are the same regions as for the degenerate case, see Fig. 5.3.

Where the phase of the degenerate parametric oscillations had a bi-stability, the nondegenerate parametric oscillations have a continuous degeneracy. The steady-state phases of the mode fields are denoted by θ_n and θ_m. Their sum is given by

$$\theta_n + \theta_m = \Theta \in \{\pi/2, \pi\}, \quad \tan\Theta = -1/\sqrt{\epsilon^2/\Gamma_g^2 - 1}, \tag{5.11}$$

while their difference $\theta_n - \theta_m$ is undefined. Therefore, the intermode phases have continuous degeneracies between $-\pi$ and π.

In Paper K, we demonstrated nondegenerate parametric oscillations in a superconducting microwave resonator. The output intensities $|C_n|^2$ are measured as a function of the pump detuning δ and amplitude ϵ, see Fig. 5.4(a-b). The oscillations start at a certain pump power, $\epsilon \approx \Gamma$, defining the instability threshold, and grow rapidly above that. The observed intensities can be compared with the theoretical steady-state amplitudes Eq. (5.9), see Fig. 5.4(c-d). We find good agreement inside regions I and II. However, in region III, where the ground and excited states are both stable, the discrepancy is larger. The model does not predict the probability of occupying the excited state; consequently, Eq. (5.9) is the maximally achievable intensities. In the experiment, the resonator can switch between the two states, yielding a lower intensity on average. For large negative detunings, the observed probability for oscillations tends to zero.

We study the phase-space distributions for the two modes at the point in δ-ϵ space indicated by the white circles in Fig. 5.4(a-b). We acquire quadrature voltages $I_n(t)$ and $Q_n(t)$, and calculate 2-D histograms, see Fig. 5.5 (a-b). The oscillations have a finite average amplitude, while the phase is random, hence the large frequency noise.

To quantify if the observed state is quantum or classical, we would typically study the variances of the different quadratures. However, the large fluctuations in phase impose a problem as they introduce large variances of the individual quadratures. To theoretically study the effect of vacuum noise on the system, and more specifically, its quantum properties such as entanglement, we would linearize around the classical solution and consider small deviations due to the noise. Here, however, the large fluctuations in phase impose a problem as they are present already in the classical part, implying that it is not possible to linearize the model. Therefore, we are not able to perform any quantitative analysis of the noise in the quantum regime.

However, we can study the cross-quadrature histograms I_3, I_4 and Q_3, Q_4 in Fig. 5.5 (c-d). From those histograms, it is clear that I_3 and I_4 have equal signs, while Q_3 and Q_4 have opposite signs, implying that the phase sum is fixed as predicted from Eq. (5.11).

The question is, what drives the phase evolutions? In Fig. 5.5(e), the phase evolution in time of the two modes is plotted, and the clear anti-correlation is observed. For illustration, the phases of the local oscillators are set so that $\langle\theta_3 + \theta_4\rangle = 0$. From the phase evolution, the frequency noise spectrum $S_y(f)$ is extracted, see Fig. 5.5(f). The

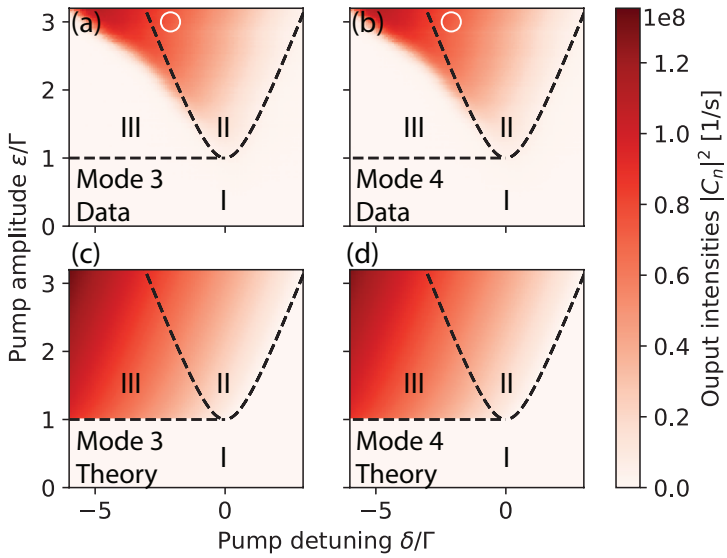

Figure 5.4: *Nondegenerate pumping of modes $n = 3$ and 4, at frequency $\omega_p = \omega_3 + \omega_4 + 2\delta$. (a-d) Experimentally observed and theoretical output intensities $|C_n|^2$ vs. detuning δ and amplitude ϵ of the applied pump tone. I-III indicate the three different stability regions described in the main text.*

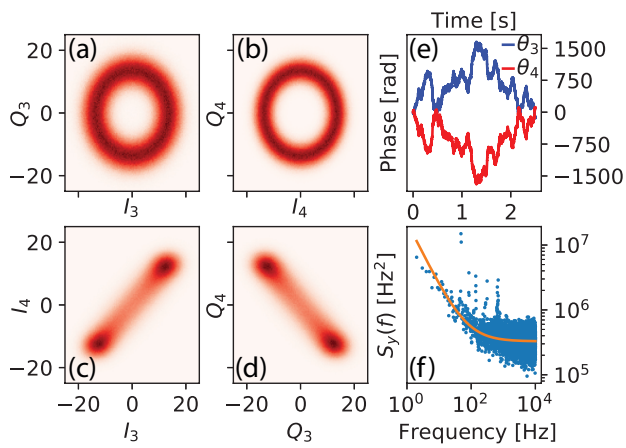

Figure 5.5: *Quadrature histograms of the nondegenerate parametric oscillations. Panels (a-b) show the phase-space distributions for the two modes measured at the point indicated by the white circles in Fig. 5.4. The color scale is proportional to the number of counts in each bin of the digitized output. (c-d) Two out of four cross-quadrature histograms. All histograms consist of 1 million samples each. (e) Evolution of the phases θ_n in time. (f) The frequency spectrum of the nondegenerate parametric oscillations in mode 3. The solid line is a combination of white and flicker noise.*

spectrum is a combination of $1/f$ and white noise. It is difficult to pinpoint the origin of the noise, but it is well known that SQUIDs have a $1/f$ flux noise spectrum [224]. Flux noise translates directly into frequency noise of the resonator, and therefore also into frequency noise of the parametric oscillations.

5.7 Subharmonic oscillations

In Paper H, we study period-tripling subharmonic oscillations in a driven superconducting resonator. We detect the output field at frequencies near the fundamental mode when driving the resonator close to three times the fundamental mode. With an amplitude exceeding an instability threshold, we observe three stable radiative states with equal amplitudes, phase-shifted by $2\pi/3$ rad.

Subharmonic oscillations are described by nonperturbative solutions to dynamical equations, which appear abruptly and always coexist with a stable vacuum state. In this respect, subharmonic oscillations distinctly differ from conventional parametric oscillations, which gradually emerge as a result of vacuum instabilities. Although period-multiplication in nonlinear equations is theoretically explained in textbooks, experimental demonstrations of the phenomenon are not common. A few early observations of subharmonic generation were performed in classical electrical circuits based on saturable inductors [225] and varactors [226].

The Kerr nonlinearity in our circuits generates the downconversion from 3ω to ω. However, that is not enough to describe the full picture of our device. The presence of a mode close to 3ω leads to an enhancement of the effect, and we have to describe the system by two coupled equations, one for each mode,

$$\begin{aligned} i\dot{A}_n + (\delta_n + i\Gamma_n + \alpha_n A_n^\dagger A_n + 2\alpha A_m^\dagger A_m)A_n + \tilde{\alpha} A_n^{*^2} &= 0 \\ i\dot{A}_m + (\delta_m + i\Gamma_m + \alpha_m A_m^\dagger A_m + 2\alpha A_n^\dagger A_n)A_m + \frac{\tilde{\alpha}}{3} A_n^3 &= \sqrt{2\Gamma_{m0}} B_m. \end{aligned} \tag{5.12}$$

The quantum Hamiltonian for a closed cavity (no photon loss) would be

$$\begin{aligned} \frac{\hat{H}}{\hbar} = &- \frac{\alpha_n}{2}(\hat{a}_n^\dagger \hat{a}_n)^2 - \frac{\alpha_m}{2}(\hat{a}_m^\dagger \hat{a}_m)^2 - 2\alpha_g(\hat{a}_n^\dagger \hat{a}_n \hat{a}_m \dagger \hat{a}_m) \\ &- \frac{\tilde{\alpha}}{3}(\hat{a}_n^{\dagger^3} \hat{a}_m + \hat{a}_n^3 \hat{a}_m^\dagger) + \sqrt{2\Gamma_{m0}}(B_2 \hat{a}_2^\dagger + B_2^* \hat{a}_2), \end{aligned} \tag{5.13}$$

where the $\hat{a}_n^{\dagger^3} \hat{a}_m$ term is responsible for the creation of three signal photons from one pump photon.

Solutions to Eq. (5.12) include the trivial solution $A_n = 0$, which is always stable. Three nontrivial solutions exist, all with the same amplitude but with their relative phases of $2\pi/3$ with respect to each other. The nontrivial solutions exist within an interval of effective pump strengths, where the upper limit is given by

$$|B_m|^2 \approx \frac{2}{7}(3\omega_n - \omega_m)^2 \frac{|\delta_1|}{\alpha_1 \Gamma_2}, \tag{5.14}$$

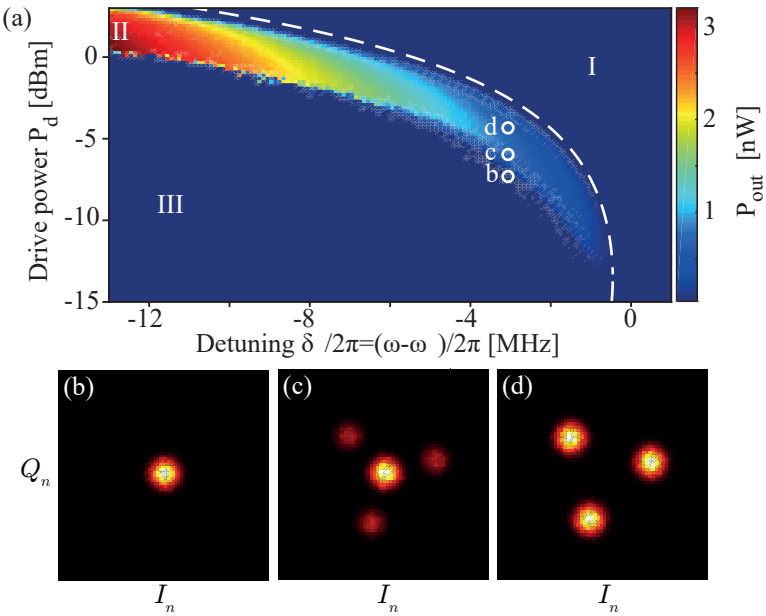

Figure 5.6: *(a) The intensity of the subharmonic oscillation output signal as a function of drive power and detuning. The oscillations are detected in region II. The dashed white line that separates regions I and II corresponds to the theoretical boundary of existence for the subharmonic oscillations. In region III, the oscillations, although they exist as a solution to Eq.* (5.12)*, are not visible because of the oscillator switching to the ground state. (b)–(d) Histograms of the detected radiation quadrature voltages from the operating points indicated by the white circles in (a). These histograms reveal three dynamic states: (b) the ground state; (c) the ground state (in the middle) and the three excited states with equal amplitudes and with phases differing by* $2\pi/3$*; and (d), only the excited states.*

which is plotted as the dashed white line in Fig. 5.6 (a). The lower limit is far below the visible area. The reason why we do not see any signal in area III is that the oscillator is predominantly in its ground state, even if the nontrivial states exist, similar to region III for the nondegenerate parametric oscillations.

In Fig. 5.6 (a), we observe a distinct region in pump power and frequency detuning where the subharmonic oscillations occur. Phase-space distributions for three different pump conditions are shown in Fig. 5.6 (b-d). In (b), which is just inside region III, the oscillator is solely in its ground state (even though the finite-amplitude solutions are stable). In (c), the finite solutions are observed together with the ground state. Finally, in (d), only the finite-amplitude solutions are observed. We see the clear $2\pi/3$ phase-shift between the three states.

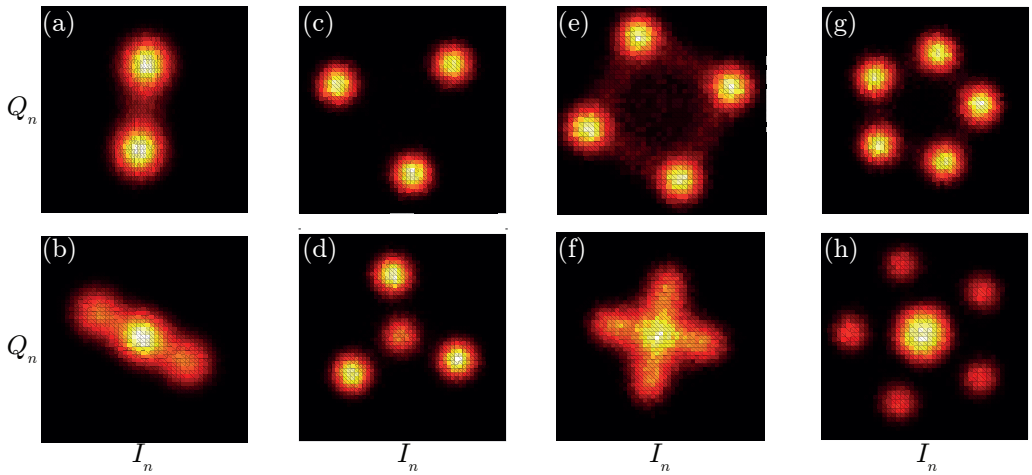

Figure 5.7: *Phase-space histograms of period multiplication oscillations generated by applying a microwave signal to the flux pump line at multiples N of the fundamental mode frequency. (a) and (b) $N = 2$; (c) and (d) $N = 3$; (e) and (f) $N = 4$; (g) and (h) $N = 5$. The output signals correspond to around 60 photons.*

5.8 Period multiplication

So far, we have seen the creation of two and three photons from one pump photon. In Paper J, we investigate parametric pumping at higher multiples, N, of a resonant frequency. Again, a quantum Hamiltonian for a closed cavity without loss can be written as

$$\frac{H}{\hbar} = -\frac{\alpha_1}{2}(a_1^\dagger a_1)^2 - \frac{\epsilon_n}{N}(a_1^{\dagger N} + a_1^N). \tag{5.15}$$

The $a_1^{\dagger N}$ term will create N photons in the resonator. The resonator state would be a superposition state of N coherent states with relative phases of $2\pi/N$.

We apply a pump tone for $N = 2, 3, 4, 5$, and record phase-space distributions, see Fig. 5.7. We observe a qualitative agreement between the observed data and the model of N photon creation. Depending on the exact parameters of the pump, the ground state of the oscillator can also be stable (bottom row of Fig. 5.7). Since our measurements are performed on an open cavity with photon loss, we do not expect to see any coherent superpositions of the different coherent states (an N-headed cat state).

5.9 Entanglement in the dynamical Casimir effect

All parametric effects seen so far in this work have been implemented using one or more resonances. An extreme limit is when the resonator–transmission line coupling is very strong, which approaches the case of not having a resonator at all (i.e., the SQUID is

coupled directly to the transmission line). Let us consider the case of nondegenerate parametric oscillations. If we were to take $\Gamma_{n0} = \Gamma_{m0} \to \infty$, we would never be able to reach the instability threshold and observe parametric oscillations. However, below the threshold, there is parametric amplification [76, 211]. A SQUID coupled at the end of a transmission line can, therefore, be used to amplify both signals and vacuum fluctuations.

Another view of a flux-pumped SQUID at the end of a transmission line is as a moving mirror. A change of inductance is mapped to a change in electrical length. Using a SQUID, effective length changes of more than 100 µm can be achieved. Together with a fast modulation at GHz frequencies, we can achieve effective velocities of more than 10 % of the speed of light.

At such high velocities, relativistic effects appear. One such effect is the dynamical Casimir effect (DCE), which predicts that a moving mirror will create photons out of vacuum [227]. A qualitative explanation is that in a quantum vacuum, virtual photons get created and subsequently annihilated in pairs. If a mirror moves fast enough, it can insert itself between such a photon pair and prevent the annihilation, creating two real photons in the process. The DCE was demonstrated using SQUIDs in Refs. [228, 229]

To first order and assuming an ideal transmission line and zero temperature, the output photon flux density at a frequency ω from the DCE radiation is [228, 230, 231]

$$n(\omega) = \frac{l^2}{c^2}\omega(\omega_p - \omega) = \frac{v^2}{c^2}\left(\frac{1}{4} - \frac{\delta^2}{\omega_p^2}\right), \tag{5.16}$$

where l is the displacement length of the mirror, c is the effective speed of light in the transmission line, $v = l\omega$ is the speed of the mirror, and $\delta = \omega - \omega_p/2$. The relativistic effect is obvious as n goes to zero when the speed of light goes to infinity.

In Fig. 5.8 (a), we plot the measured n as a function of the flux modulation amplitude, with $\Phi_{\mathrm{AC}} = 20\ \mathrm{m}\Phi_0$ corresponding to $v/c_0 = 0.31$. Simply applying Eq. (5.16) with the data starts showing discrepancy for higher modulation amplitudes. We explain this discrepancy by the nonlinear relationship between flux and electrical length. At high modulation amplitude, there is a significant generation of harmonics at $2\omega_p, 3\omega_p, ...$ This leads to additional photon creation, and if we sum up the contributions of the first three orders, we see good agreement up to the last data point.

In Paper L, we demonstrate an additional property of DCE that was not shown in Ref. [228] due to insufficent calibration of the system gain and noise. Since the photons get created in pairs symmetrically around $\omega_p/2$, we expect to see entanglement between negative and positive detunings δ. Therefore, we use two digitizers to simultaneously measure the quadratures at two different frequencies symmetric around $\omega_p/2$, I_+, Q_+, I_-, Q_-. By careful calibration of the system gain and temperature, we can extract the correlations between two frequencies in units of photons.

Entanglement of two electromagnetic modes manifests itself as two-mode squeezing below the Heisenberg limit (0.5 photons per mode), according to the Duan criterion [232]. We define

$$\delta_{IQ+} = \langle (I_+ + I_-)^2 \rangle + \langle (Q_+ - Q_-)^2 \rangle \tag{5.17}$$

$$\delta_{IQ-} = \langle (I_+ - I_-)^2 \rangle + \langle (Q_+ + Q_-)^2 \rangle. \tag{5.18}$$

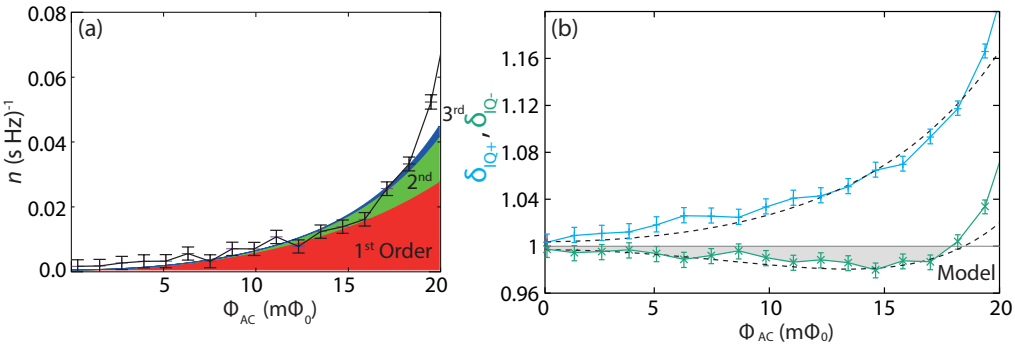

Figure 5.8: *Demonstration of entanglement in the dynamical Casimir effect. (a) Average photon spectral density generated as a function of flux-pump amplitude. A flux modulation of* 20 $\mathrm{m}\Phi_0$ *corresponds to an equivalent mirror speed of* 31 % *of the speed of light. A model is used to fit the resulting photon rate taking the first three orders of the SQUID nonlinearity into account. (b) The variance of two combinations of quadratures of two modes. A variance of 1 corresponds to the vacuum fluctuations set by the Heisenberg limit. A variance below that implies that the two modes are entangled.*

We plot these two quantities as a function of the modulation amplitude in Fig. 5.8 (b). Since we are combining the variances of two quadratures, the Heisenberg limit (the vacuum fluctuations) is equal to 1. We observe that δ_{IQ-} goes below 1, well outside the error bars, proving the entanglement between the two modes. At high modulation amplitudes, the higher-order terms dilute the entanglement until it is no longer measurable. The dilution is understood as two modes that are symmetric around $\omega_p/2$, are not symmetric around $2\omega_p/2$ or $3\omega_p/2$. The photons created by the harmonics are therefore not entangled at the detection frequencies, but still get mixed in with the entangled ones from the fundamental pump.

5.10 A single-photon generator

Finally, in Paper M, we propose a method of generating propagating single photons using a superconducting qubit coupled to a transmission line. Propagating refers to that the photon will be traveling in a wave packet along a microwave transmission line. This is different from having a qubit or resonator in the $|1\rangle$ state. However, a single propagating photon can be created by the decay of a qubit into a transmission line.

In our proposal, a transmon qubit (but it could be any superconducting qubit) is strongly coupled at the end of an open transmission line. The open end acts as a mirror to microwaves and ensures that the decay of the qubit is unidirectional. A qubit coupled to a two-sided transmission line would emit to either side half of the time (however, it would be unknown prior to measurement where the photon went). The goal is to first excite the qubit by sending a coherent state, α_{in}, which drives half a Rabi-oscillation between $|0\rangle$ and $|1\rangle$, and then let the qubit decay into the transmission line. The issue

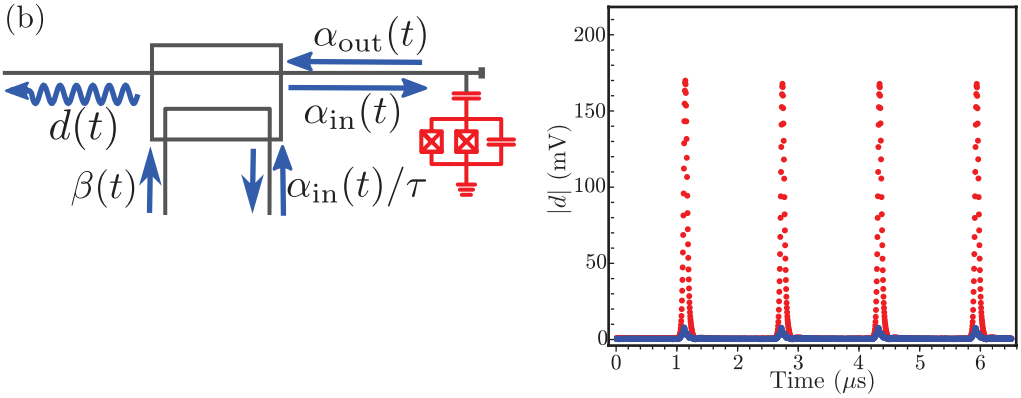

Figure 5.9: *(a) A proposed single-photon generate using a transmon qubit (red) and a directional coupler. A coherent state* $\alpha_{\rm in}(t)$ *excites a qubit at the end of a transmission line to the state* $|1\rangle$*. The reflected field* $\alpha_{\rm out}(t)$ *contains a coherent part and the emission from the qubit. (b) Demonstration of the pulse cancellation with a qubit present. Measurement of the output* $d(t)$ *with (red) and without (blue) the cancellation drive* $\beta(t)$*. The pulses are* 100 ns *long and show a cancellation of* -34 dB.

with this approach is that the coherent state will reflect in the mirror and get mixed with the single photon. Our scheme, shown in Fig. 5.9 (a), uses a directional coupler, which acts as an unbalanced beam splitter. The idea is that after the first coherent state $\alpha_{\rm in}$ interacts with and excites the qubit, a second coherent state with opposite sign, β, can be sent through the second arm of the beam splitter to cancel the first coherent state, leaving only the single-photon in the output.

The outgoing field, $d(t)$, is related to the inputs as

$$d = i\sqrt{1-\tau^2}\left(\alpha_{\rm in}(t) + \sqrt{\Gamma_{\rm r}}\sigma^-(t)\right) + \tau\beta(t), \tag{5.19}$$

where τ is the transmission coefficient of the beam splitter, $\Gamma_{\rm r}$ is the radiative coupling rate between the qubit and the transmission line, and σ^- is the lowering operator acting on the qubit. By choosing $\beta(t) = -i\sqrt{1-\tau^2}\alpha_{\rm in}(t)$, the coherent part is completely cancelled and the output contains only the qubit contribution $i\sqrt{1-\tau^2}\sqrt{\Gamma_{\rm r}}\sigma^-(t)$. In the limit of a perfect qubit excitation and $\tau \to 0$, the probability of having a single photon in the output d is unity.

In Fig. 5.9 (b), we demonstrate the pulse cancellation scheme (in the absence of a qubit), by sending 100 ns pulses through both ports of a directional coupler with $\tau = 0.1$. By adjusting the relative phase and amplitude between the ports, we achieve more than -34 dB of cancellation.

The advantage of our proposed single-photon generator is its simplicity. It does not rely on a fixed-frequency resonator, as in Refs. [233, 234]. This also implies that our generator can change the frequency of its photon in-situ, by tuning the magnetic flux threading the SQUID loop of the qubit, something that is less trivial in a resonator-based setup where the frequency detuning between resonator and qubit is important. In Paper M, we show

that the efficiency of generating single photons should approach 99 %, which would be state-of-the-art for microwave single-photon generators.

Paper C utilizes a setup up like the one proposed here and is one step on the way to realizing the suggested single-photon generator.

CHAPTER VI

Summary and paper contributions

In conclusion, we have successfully developed fabrication processes and measurement setups yielding low-loss and high-coherence superconducting circuits. By incorporating magnetic flux tunability, we have used these circuits to build a variety of different tunable devices, such as resonators, qubits, and mirrors. Then, we used these devices to create different microwave states and to demonstrate important concepts for gate-based quantum computers. Achieving high-fidelity control and readout of our qubits has allowed us to implement two quantum algorithms. Our low-loss circuits and high-fidelity control have shown that it will be feasible to build an intermediate-scale quantum-information processor using tunable superconducting circuits in the microwave domain, something that is sought after by industry and academia alike all over the world.

In **Paper A**, we primarily studied the frequency noise of superconducting aluminum resonators. While this type of study was not new, we extended the results by measuring the noise at an excitation energy in the resonator equivalent to that of a single photon. The energy level makes our study relevant to superconducting transmon qubits, where frequency noise sets an upper bound on the qubit coherence time. I designed and fabricated the sample, and contributed to the experimental setup, measurements, and writing process.

In **Paper B**, we performed extensive benchmarking of qubit-coherence fluctuations across days. By careful analysis of the fluctuations, we found that the spectrum of the qubit relaxation time followed a combination of two Lorentzian noise processes. We attribute these two Lorenzians to interacting two-level systems. I designed and fabricated the sample. I also performed parts of the measurements, data analysis, and manuscript writing. Note shared the first authorship.

In **Paper C**, we developed techniques for measuring the intrinsic relaxation and dephasing rates of superconducting qubits directly coupled to transmission lines (waveguide-QED). I designed and fabricated the device. I provided input on the experimental setup, the measurements, the data analysis, and the manuscript.

In **Paper D**, we studied the quality factors of superconducting resonators fabricated

on a piezoelectric substrate. We found that the quality factors are mainly limited by a photon to phonon conversion. I contributed ideas to the sample design, fabrication, measurements, analysis, and manuscript.

In **Paper E**, we demonstrated a single-shot readout of a superconducting qubit using a degenerate parametric oscillator. The extracted readout fidelity was 98.7 %. I contributed to device fabrication, experimental measurements, and data analysis. The Monte-Carlo simulation of the readout fidelity was done solely by me.

In **Paper F**, we implemented density matrix exponentiation on a superconducting quantum processor. I developed much of the software needed to perform the experiments, specifically the virtual-Z gates and the gate compilation. I also contributed to the tuneup protocols and the calibration of the quantum processor.

In **Paper G**, we showed that the NP-complete exact-cover problem could be solved using the quantum approximate optimization algorithm. Using our low-loss two-qubit processor with high gate fidelities, we iterated the algorithm up to level 2 with a success probability of 96.6 %. I designed and fabricated the device, built the measurement setup, performed and analyzed the experiments, and wrote the manuscript.

In **Paper H**, we observed subharmonic oscillation of a resonator when driven at three times its resonant frequency. I developed much of the measurement software and provided input on the data analysis and the manuscript.

In **Paper I**, we used a double tunable resonator to perform degenerate parametric oscillations by pumping it using two pumps. We observed the predicted threshold dependence on the pump phase difference. I developed much of the measurement software and provided input on the sample design and fabrication.

In **Paper J**, we extended degenerate parametric oscillations to higher-order pumping. We modulated a frequency-tunable resonator at up to five times its resonant frequency. I developed much of the measurement software and performed preliminary measurements.

In **Paper K**, we observed nondegenerate parametric oscillations, which manifest themselves as correlated signals in two modes. The nondegenerate parametric oscillations were achieved by a parametric modulation at the sum of the two mode frequencies. We studied the classical properties of the oscillations and showed an excellent agreement with a theoretical model. I designed and fabricated the device, built the measurement setup, and developed the necessary software. I also analyzed the data and wrote the manuscript.

In **Paper L**, we demonstrated that the radiation coming from a mirror moving in vacuum at 10 % of the speed of light creates entangled photons. This is called the dynamical Casimir effect. I fabricated the device. I also contributed to the device design, experimental setup, data analysis, and the manuscript.

In **Paper M**, we proposed a new single-photon generator in the microwave regime. Our proposal is built on a superconducting qubit and a directional coupler acting as a beam splitter. The beam splitter provides cancellation of the signal used to excite the qubit that then can decay and generate a single photon. I came up with the idea on how to implement the idea in superconducting circuits and provided experimental data that demonstrated the feasibility of the proposal.

References

[1] C. H. Bennett, E. Bernstein, G. Brassard, and U. Vazirani, "Strengths and weaknesses of quantum computing", SIAM Journal on Computing **26**, 1510–1523 (1997).

[2] P. W. Shor, "Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer", SIAM Review **41**, 303–332 (1999).

[3] D. S. Abrams and S. Lloyd, "Simulation of many-body fermi systems on a universal quantum computer", Physical Review Letters **79**, 2586 (1997).

[4] C. Pomerance, "A tale of two sieves", Biscuits of Number Theory **85**, 175 (2008).

[5] M. A. Nielsen and I. Chuang, *Quantum computation and quantum information* (American Association of Physics Teachers, 2002).

[6] J. Gill, "Computational complexity of probabilistic turing machines", SIAM Journal on Computing **6**, 675–695 (1977).

[7] W. A. Borders, A. Z. Pervaiz, S. Fukami, K. Y. Camsari, H. Ohno, and S. Datta, "Integer factorization using stochastic magnetic tunnel junctions", Nature **573**, 390–393 (2019).

[8] A. W. Harrow, A. Hassidim, and S. Lloyd, "Quantum algorithm for linear systems of equations", Physical Review Letters **103**, 150502 (2009).

[9] S. Aaronson, "Read the fine print", Nature Physics **11**, 291–293 (2015).

[10] J. Preskill, "Fault-tolerant quantum computation", in *Introduction to quantum computation and information* (World Scientific, 1998), pp. 213–269.

[11] D. Gottesman, "Theory of fault-tolerant quantum computation", Physical Review A **57**, 127 (1998).

[12] E. Knill, R. Laflamme, and L. Viola, "Theory of quantum error correction for general noise", Physical Review Letters **84**, 2525 (2000).

[13] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland, "Surface codes: towards practical large-scale quantum computation", Physical Review A **86**, 032324 (2012).

[14] C. Gidney and M. Ekerå, "How to factor 2048 bit rsa integers in 8 hours using 20 million noisy qubits", arXiv:1905.09749 (2019).

[15] J. Preskill, "Quantum computing in the nisq era and beyond", Quantum **2**, 79 (2018).

[16] F. Arute et al., "Quantum supremacy using a programmable superconducting processor", Nature **574**, 505–510 (2019).

[17] D. P. DiVincenzo, "Quantum computation", Science **270**, 255–261 (1995).

[18] J. Weber, W. Koehl, J. Varley, A. Janotti, B. Buckley, C. Van de Walle, and D. D. Awschalom, "Quantum computing with defects", Proceedings of the National Academy of Sciences **107**, 8513–8518 (2010).
[19] L. Childress and R. Hanson, "Diamond nv centers for quantum computing and quantum networks", MRS bulletin **38**, 134–138 (2013).
[20] R. Hanson, L. P. Kouwenhoven, J. R. Petta, S. Tarucha, and L. M. Vandersypen, "Spins in few-electron quantum dots", Reviews of Modern Physics **79**, 1217 (2007).
[21] D. Kielpinski, C. Monroe, and D. J. Wineland, "Architecture for a large-scale ion-trap quantum computer", Nature **417**, 709–711 (2002).
[22] H. Häffner, C. F. Roos, and R. Blatt, "Quantum computing with trapped ions", Physics reports **469**, 155–203 (2008).
[23] S. M. Frolov, S. R. Plissard, S. Nadj-Perge, L. P. Kouwenhoven, and E. P. Bakkers, "Quantum computing based on semiconductor nanowires", MRS bulletin **38**, 809–815 (2013).
[24] L. Casparis et al., "Superconducting gatemon qubit based on a proximitized two-dimensional electron gas", Nature Nanotechnology **13**, 915–919 (2018).
[25] D. Aasen et al., "Milestones toward majorana-based quantum computing", Physical Review X **6**, 031016 (2016).
[26] M. Tinkham, *Introduction to superconductivity* (Courier Corporation, 2012).
[27] T. Van Duzer and C. W. Turner, *Principles of superconductive devices and circuits* (Edward Arnold, 1981).
[28] C. L. Degen, F. Reinhard, and P. Cappellaro, "Quantum sensing", Reviews of Modern Physics **89**, 035002 (2017).
[29] M. H. Devoret and J. M. Martinis, "Implementing qubits with superconducting integrated circuits", in *Experimental aspects of quantum computing* (Springer, 2005), pp. 163–203.
[30] R. J. Schoelkopf and S. M. Girvin, "Wiring up quantum systems", Nature **451**, 664–669 (2008).
[31] J. Clarke and F. K. Wilhelm, "Superconducting quantum bits", Nature **453**, 1031 (2008).
[32] M. H. Devoret and R. J. Schoelkopf, "Superconducting circuits for quantum information: an outlook", Science **339**, 1169–1174 (2013).
[33] X. Gu, A. F. Kockum, A. Miranowicz, Y. X. Liu, and F. Nori, "Microwave photonics with superconducting quantum circuits", Physics Reports **718-719**, 1–102 (2017).
[34] G. Wendin, "Quantum information processing with superconducting circuits: a review", Reports on Progress in Physics **80**, 106001 (2017).
[35] 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).
[36] 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**, 369–395 (2020).
[37] 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).

[38] R. H. Hadfield and G. Johansson, *Superconducting devices in quantum optics* (Springer, 2016).
[39] A. Blais, S. M. Girvin, and W. D. Oliver, "Quantum information processing and quantum optics with circuit quantum electrodynamics", Nature Physics **16**, 247–256 (2020).
[40] R. Barends et al., "Superconducting quantum circuits at the surface code threshold for fault tolerance", Nature **508**, 500–503 (2014).
[41] T. Walter et al., "Rapid high-fidelity single-shot dispersive readout of superconducting qubits", Physical Review Applied **7**, 054020 (2017).
[42] D. Rosenberg et al., "3D integrated superconducting qubits", npj Quantum Information **3**, 42 (2017).
[43] A. Megrant et al., "Planar superconducting resonators with internal quality factors above one million", Applied Physics Letters **100**, 113510 (2012).
[44] J. Clarke and A. I. Braginski, *The squid handbook: applications of squids and squid systems* (John Wiley & Sons, 2006).
[45] M. Sandberg, C. M. Wilson, F. Persson, T. Bauch, G. Johansson, V. Shumeiko, T. Duty, and P. Delsing, "Tuning the field in a microwave resonator faster than the photon lifetime", Applied Physics Letters **92**, 203501 (2008).
[46] Y. Yin et al., "Catch and release of microwave photon states", Physical Review Letters **110**, 107001 (2013).
[47] D. M. Pozar, *Microwave engineering* (John Wiley & Sons, 2009).
[48] H. Padamsee, "The science and technology of superconducting cavities for accelerators", Superconductor science and technology **14**, R28–R51 (2001).
[49] J. Baselmans, "Kinetic inductance detectors", Journal of Low Temperature Physics **167**, 292–304 (2012).
[50] R. Loudon, *The quantum theory of light* (OUP Oxford, 2000).
[51] B. A. Mazin, "Microwave kinetic inductance detectors", PhD thesis (California Institute of Technology, 2005).
[52] R. Barends, "Photon-detecting superconducting resonators", PhD thesis (Delft University of Technology, 2009).
[53] D. T. Sank, "Fast, accurate state measurement in superconducting qubits", PhD thesis (University of California, Santa Barbara, 2014).
[54] S. d. Graaf, A. Danilov, A. Adamyan, T. Bauch, and S. Kubatkin, "Magnetic field resilient superconducting fractal resonators for coupling to free spins", Journal of Applied Physics **112**, 123905 (2012).
[55] N. Vercruyssen, R. Barends, T. M. Klapwijk, J. Muhonen, M. Meschke, and J. P. Pekola, "Substrate-dependent quasiparticle recombination time in superconducting resonators", Applied Physics Letters **99**, 062509 (2011).
[56] J. Zmuidzinas, "Superconducting microresonators: physics and applications", Annual Review of Condensed Matter Physics **3**, 169–214 (2012).
[57] Y. Nakamura, Y. A. Pashkin, and J. S. Tsai, "Coherent control of macroscopic quantum states in a single-cooper-pair box", Nature **398**, 786–788 (1999).
[58] J. Koch et al., "Charge-insensitive qubit design derived from the cooper pair box", Physical Review A **76**, 042319 (2007).

[59] M. H. Devoret, A. Wallraff, and J. M. Martinis, "Superconducting qubits: a short review", arXiv:cond-mat/0411174 (2004).
[60] A. Wallraff et al., "Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics", Nature **431**, 162–167 (2004).
[61] M. Wallquist, V. Shumeiko, and G. Wendin, "Selective coupling of superconducting charge qubits mediated by a tunable stripline cavity", Physical Review B **74**, 224506 (2006).
[62] W. Wustmann and V. Shumeiko, "Parametric resonance in tunable superconducting cavities", Physical Review B **87**, 184501 (2013).
[63] R. Z. Sagdeev, D. Usikov, and G. M. Zaslavskii, *Nonlinear physics: from the pendulum to turbulence and chaos* (Harwood Academic, 1990).
[64] S. H. Strogatz, *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering* (Hachette UK, 2014).
[65] S. L. Braunstein and P. Van Loock, "Quantum information with continuous variables", Reviews of Modern Physics **77**, 513 (2005).
[66] S. L. Braunstein and A. K. Pati, *Quantum information with continuous variables* (Springer Science & Business Media, 2012).
[67] M. Hofheinz et al., "Synthesizing arbitrary quantum states in a superconducting resonator", Nature **459**, 546–549 (2009).
[68] C. Eichler, D. Bozyigit, C. Lang, M. Baur, L. Steffen, J. M. Fink, S. Filipp, and A. Wallraff, "Observation of two-mode squeezing in the microwave frequency domain", Physical Review Letters **107**, 113601 (2011).
[69] E. Flurin, N. Roch, F. Mallet, M. H. Devoret, and B. Huard, "Generating entangled microwave radiation over two transmission lines", Physical Review Letters **109**, 183901 (2012).
[70] C. S. Chang et al., "Generating multimode entangled microwaves with a superconducting parametric cavity", Physical Review Applied **10**, 044019 (2018).
[71] B. Vlastakis et al., "Deterministically encoding quantum information using 100-photon Schrödinger cat states", Science **342**, 607–610 (2013).
[72] T. Yamamoto, K. Inomata, M. Watanabe, K. Matsuba, T. Miyazaki, W. D. Oliver, Y. Nakamura, and J. S. Tsai, "Flux-driven Josephson parametric amplifier", Applied Physics Letters **93**, 042510 (2008).
[73] E. Zakka-Bajjani, F. Nguyen, M. Lee, L. R. Vale, R. W. Simmonds, and J. Aumentado, "Quantum superposition of a single microwave photon in two different colour states", Nature Physics **7**, 599–603 (2011).
[74] J. A. Giordmaine and R. C. Miller, "Tunable coherent parametric oscillation in $LiNbO_3$ at optical frequencies", Physical Review Letters **14**, 973 (1965).
[75] C. Wilson, T. Duty, M. Sandberg, F. Persson, V. Shumeiko, and P. Delsing, "Photon generation in an electromagnetic cavity with a time-dependent boundary", Physical Review Letters **105**, 233907 (2010).
[76] W. Wustmann and V. Shumeiko, "Nondegenerate parametric resonance in a tunable superconducting cavity", Physical Review Applied **8**, 024018 (2017).
[77] W. Wustmann and V. Shumeiko, "Parametric effects in circuit quantum electrodynamics", Low Temperature Physics **45**, 848–869 (2019).

[78] A. T. A. M. de Waele, "Basic Operation of Cryocoolers and Related Thermal Machines", Journal of Low Temperature Physics **164**, 179–236 (2011).
[79] M. Simoen, "Parametric interactions with signals and the vacuum", PhD thesis (Chalmers University of Technology, 2015).
[80] G. K. White and P. Meeson, *Experimental Techniques in Low-Temperature Physics*, 4th ed. (OUP Oxford, 2002), pp. 153–164.
[81] J. Roffe, "Quantum error correction: an introductory guide", Contemporary Physics **60**, 226–245 (2019).
[82] W. D. Oliver and P. B. Welander, "Materials in superconducting quantum bits", MRS bulletin **38**, 816–825 (2013).
[83] A. A. Houck et al., "Controlling the spontaneous emission of a superconducting transmon qubit", Physical Review Letters **101**, 080502 (2008).
[84] H. Paik et al., "Observation of high coherence in Josephson junction qubits measured in a three-dimensional circuit qed architecture", Physical Review Letters **107**, 240501 (2011).
[85] H. Wang et al., "Improving the coherence time of superconducting coplanar resonators", Applied Physics Letters **95**, 233508 (2009).
[86] J. Baselmans and S. Yates, "Long quasiparticle lifetime in aluminum microwave kinetic inductance detectors using coaxial stray light filters", in AIP Conference Proceedings, Vol. 1185, 1 (American Institute of Physics, 2009), pp. 160–163.
[87] C. Wang, C. Axline, Y. Y. Gao, T. Brecht, Y. Chu, L. Frunzio, M. H. Devoret, and R. J. Schoelkopf, "Surface participation and dielectric loss in superconducting qubits", Applied Physics Letters **107**, 162601 (2015).
[88] A. P. Place et al., "New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds", arXiv:2003.00024 (2020).
[89] R. Schoelkopf, A. Clerk, S. Girvin, K. Lehnert, and M. Devoret, "Qubits as spectrometers of quantum noise", in *Quantum noise in mesoscopic physics* (Springer, 2003), pp. 175–203.
[90] E. Purcell, "Spontaneous emission probabilities at radio frequencies", Physical Review **69**, 681 (1946).
[91] E. Jeffrey et al., "Fast accurate state measurement with superconducting qubits", Physical Review Letters **112**, 190504 (2014).
[92] E. A. Sete, J. M. Martinis, and A. N. Korotkov, "Quantum theory of a bandpass purcell filter for qubit readout", Physical Review A **92**, 012325 (2015).
[93] 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).
[94] R. Barends et al., "Minimizing quasiparticle generation from stray infrared light in superconducting quantum circuits", Applied Physics Letters **99**, 113507 (2011).
[95] G. Catelani, R. J. Schoelkopf, M. H. Devoret, and L. I. Glazman, "Relaxation and frequency shifts induced by quasiparticles in superconducting qubits", Physical Review B **84**, 064517 (2011).
[96] W. Eisenmenger, K. Lassmann, H.-J. Trumpp, and R. Krauß, "Intrinsic and experimental quasiparticle recombination times in superconducting films", Applied Physics **12**, 163–171 (1977).

[97] L. Sun et al., "Measurements of quasiparticle tunneling dynamics in a band-gap-engineered transmon qubit", Physical Review Letters **108**, 230509 (2012).
[98] R.-P. Riwar and G. Catelani, "Efficient quasiparticle traps with low dissipation through gap engineering", Physical Review B **100**, 144514 (2019).
[99] S. De Graaf, A. Adamyan, T. Lindström, D. Erts, S. Kubatkin, A. Y. Tzalenchuk, and A. Danilov, "Direct identification of dilute surface spins on Al2O3: origin of flux noise in quantum circuits", Physical Review Letters **118**, 057703 (2017).
[100] P. Kumar et al., "Origin and reduction of 1/f magnetic flux noise in superconducting devices", Physical Review Applied **6**, 041001 (2016).
[101] J. Lisenfeld et al., "Electric field spectroscopy of material defects in transmon qubits", npj Quantum Information **5**, 1–6 (2019).
[102] A. Bilmes, A. Megrant, P. Klimov, G. Weiss, J. M. Martinis, A. V. Ustinov, and J. Lisenfeld, "Resolving the positions of defects in superconducting quantum bits", Scientific Reports **10**, 1–6 (2020).
[103] 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).
[104] A. Dunsworth et al., "Characterization and reduction of capacitive loss induced by sub-micron Josephson junction fabrication in superconducting qubits", Applied Physics Letters **111**, 022601 (2017).
[105] A. Bruno, G. De Lange, S. Asaad, K. Van der Enden, N. Langford, and L. DiCarlo, "Reducing intrinsic loss in superconducting resonators by surface treatment and deep etching of silicon substrates", Applied Physics Letters **106**, 182601 (2015).
[106] J. Bylander et al., "Noise spectroscopy through dynamical decoupling with a superconducting flux qubit", Nature Physics **7**, 565–570 (2011).
[107] D. J. Van Harlingen, T. Robertson, B. Plourde, P. Reichardt, T. Crane, and J. Clarke, "Decoherence in josephson-junction qubits due to critical-current fluctuations", Physical Review B **70**, 064517 (2004).
[108] M. Hutchings, J. B. Hertzberg, Y. Liu, N. T. Bronn, G. A. Keefe, M. Brink, J. M. Chow, and B. Plourde, "Tunable superconducting qubits with flux-independent coherence", Physical Review Applied **8**, 044003 (2017).
[109] J. Burnett et al., "Evidence for interacting two-level systems from the 1/f noise of a superconducting resonator", Nature Communications **5**, 1–6 (2014).
[110] A. Clerk and D. W. Utami, "Using a qubit to measure photon-number statistics of a driven thermal oscillator", Physical Review A **75**, 042302 (2007).
[111] H. Ball, W. D. Oliver, and M. J. Biercuk, "The role of master clock stability in quantum information processing", npj Quantum Information **2**, 1–8 (2016).
[112] D. Schuster et al., "Resolving photon number states in a superconducting circuit", Nature **445**, 515–518 (2007).
[113] J. Schreier et al., "Suppressing charge noise decoherence in superconducting charge qubits", Physical Review B **77**, 180502 (2008).
[114] J. B. Chang et al., "Improved superconducting qubit coherence using titanium nitride", Applied Physics Letters **103**, 012602 (2013).
[115] R. Barends et al., "Coherent Josephson qubit suitable for scalable quantum integrated circuits", Physical Review Letters **111**, 080502 (2013).

[116] F. Yan et al., "The flux qubit revisited to enhance coherence and reproducibility", Nature Communications **7**, 1–9 (2016).
[117] C. Song, T. W. Heitmann, M. P. DeFeo, K. Yu, R. McDermott, M. Neeley, J. M. Martinis, and B. L. Plourde, "Microwave response of vortices in superconducting thin films of Re and Al", Physical Review B **79**, 174512 (2009).
[118] D. Bothner, T. Gaber, M. Kemmler, D. Koelle, and R. Kleiner, "Improving the performance of superconducting microwave resonators in magnetic fields", Applied Physics Letters **98**, 102504 (2011).
[119] B. Chiaro et al., "Dielectric surface loss in superconducting resonators with flux-trapping holes", Superconductor Science and Technology **29**, 104006 (2016).
[120] R. Barends, H. Hortensius, T. Zijlstra, J. J. Baselmans, S. Yates, J. Gao, and T. M. Klapwijk, "Noise in NbTiN, Al, and Ta superconducting resonators on silicon and sapphire substrates", IEEE transactions on applied superconductivity **19**, 936–939 (2009).
[121] C. Quintana et al., "Characterization and reduction of microfabrication-induced decoherence in superconducting quantum circuits", Applied Physics Letters **105**, 062601 (2014).
[122] M. Sandberg, M. R. Vissers, J. S. Kline, M. Weides, J. Gao, D. S. Wisbey, and D. P. Pappas, "Etch induced microwave losses in titanium nitride superconducting resonators", Applied Physics Letters **100**, 262605 (2012).
[123] J. M. Sage, V. Bolkhovsky, W. D. Oliver, B. Turek, and P. B. Welander, "Study of loss in superconducting coplanar waveguide resonators", Journal of Applied Physics **109**, 063915 (2011).
[124] K. Bordo and H.-G. Rubahn, "Effect of deposition rate on structure and surface morphology of thin evaporated Al films on dielectrics and semiconductors", Materials Science **18**, 313–317 (2012).
[125] M. Higo, X. Lu, U. Mazur, and K. Hipps, "Preparation of atomically smooth aluminum films: characterization by transmission electron microscopy and atomic force microscopy", Langmuir **13**, 6176–6182 (1997).
[126] W. Woods et al., "Determining interface dielectric losses in superconducting coplanar-waveguide resonators", Physical Review Applied **12**, 014012 (2019).
[127] G. J. Dolan, "Offset masks for lift-off photoprocessing", Applied Physics Letters **31**, 337–339 (1977).
[128] A. D. Córcoles et al., "Protecting superconducting qubits from radiation", Applied Physics Letters **99**, 181906 (2011).
[129] L. Faoro and L. B. Ioffe, "Internal loss of superconducting resonators induced by interacting two-level systems", Physical review letters **109**, 157005 (2012).
[130] C. T. Earnest, J. H. Béjanin, T. G. McConkey, E. A. Peters, A. Korinek, H. Yuan, and M. Mariantoni, "Substrate surface engineering for high-quality silicon/aluminum superconducting resonators", Superconductor Science and Technology **31**, 125013 (2018).
[131] T. Lindström, J. Burnett, M. Oxborrow, and A. Y. Tzalenchuk, "Pound-locking for characterization of superconducting microresonators", Review of Scientific Instruments **82**, 104706 (2011).

[132] T. Aref et al., "Quantum acoustics with surface acoustic waves", in *Superconducting devices in quantum optics* (Springer, 2016), pp. 217–244.
[133] P. Delsing et al., "The 2019 surface acoustic waves roadmap", Journal of Physics D: Applied Physics **52**, 353001 (2019).
[134] L. Ioffe, V. Geshkenbein, C. Helm, and G. Blatter, "Decoherence in superconducting quantum bits by phonon radiation", Physical review letters **93**, 057001 (2004).
[135] C. Müller, J. Lisenfeld, A. Shnirman, and S. Poletto, "Interacting two-level defects as sources of fluctuating high-frequency noise in superconducting circuits", Physical Review B **92**, 035442 (2015).
[136] P. Klimov et al., "Fluctuations of energy-relaxation times in superconducting qubits", Physical Review Letters **121**, 090502 (2018).
[137] S. Schlör, J. Lisenfeld, C. Müller, A. Bilmes, A. Schneider, D. P. Pappas, A. V. Ustinov, and M. Weides, "Correlating decoherence in transmon qubits: low frequency noise by single fluctuators", Physical Review Letters **123**, 190502 (2019).
[138] E. Rubiola, *Phase noise and frequency stability in oscillators* (Cambridge University Press, 2009).
[139] C. M. Van Vliet and P. H. Handel, "A new transform theorem for stochastic processes with special application to counting statistics", Physica A: Statistical Mechanics and its Applications **113**, 261–276 (1982).
[140] D. Salvino, S. Rogge, B. Tigner, and D. Osheroff, "Low temperature ac dielectric response of glasses to high dc electric fields", Physical Review Letters **73**, 268 (1994).
[141] S. Ludwig, P. Nalbach, D. Rosenberg, and D. Osheroff, "Dynamics of the destruction and rebuilding of a dipole gap in glasses", Physical Review Letters **90**, 105501 (2003).
[142] I. M. Pop, K. Geerlings, G. Catelani, R. J. Schoelkopf, L. I. Glazman, and M. H. Devoret, "Coherent suppression of electromagnetic dissipation due to superconducting quasiparticles", Nature **508**, 369–372 (2014).
[143] 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).
[144] J. M. Martinis, "Qubit metrology for building a fault-tolerant quantum computer", npj Quantum Information **1**, 1–3 (2015).
[145] C. M. Dawson and M. A. Nielsen, "The Solovay-Kitaev algorithm", arXiv:quant-ph/0505030 (2005).
[146] P. O'Malley et al., "Qubit metrology of ultralow phase noise using randomized benchmarking", Physical Review Applied **3**, 044009 (2015).
[147] E. Magesan, J. M. Gambetta, and J. Emerson, "Characterizing quantum gates via randomized benchmarking", Physical Review A **85**, 042311 (2012).
[148] D. F. V. James, P. G. Kwiat, W. J. Munro, and A. G. White, "Measurement of qubits", Physical Review A **64**, 052312 (2001).
[149] J. Chow et al., "Randomized benchmarking and process tomography for gate errors in a solid-state qubit", Physical Review Letters **102**, 090502 (2009).
[150] B. Foxen et al., "Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms", arXiv:2001.08343 (2020).

[151] D. Greenbaum and Z. Dutton, "Modeling coherent errors in quantum error correction", Quantum Science and Technology **3**, 015007 (2017).
[152] 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).
[153] J. M. Chow, L. DiCarlo, J. M. Gambetta, F. Motzoi, L. Frunzio, S. M. Girvin, and R. J. Schoelkopf, "Optimized driving of superconducting artificial atoms for improved single-qubit gates", Physical Review A **82**, 040305 (2010).
[154] E. Lucero et al., "Reduced phase error through optimized control of a superconducting qubit", Physical Review A **82**, 042339 (2010).
[155] Z. Chen et al., "Measuring and suppressing quantum state leakage in a superconducting qubit", Physical Review Letters **116**, 020501 (2016).
[156] 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).
[157] J. Emerson, R. Alicki, and K. Życzkowski, "Scalable noise estimation with random unitary operators", Journal of Optics B: Quantum and Semiclassical Optics **7**, S347 (2005).
[158] J. Kelly et al., "Optimal quantum control using randomized benchmarking", Physical Review Letters **112**, 240504 (2014).
[159] M. Rol et al., "Restless tuneup of high-fidelity qubit gates", Physical Review Applied **7**, 041001 (2017).
[160] E. Magesan et al., "Efficient measurement of quantum gate error by interleaved randomized benchmarking", Physical Review Letters **109**, 080505 (2012).
[161] J. Wallman, C. Granade, R. Harper, and S. T. Flammia, "Estimating the coherence of noise", New Journal of Physics **17**, 113020 (2015).
[162] M. Neeley et al., "Generation of three-qubit entangled states using superconducting phase qubits", Nature **467**, 570–573 (2010).
[163] A. Dewes, F. Ong, V. Schmitt, R. Lauro, N. Boulant, P. Bertet, D. Vion, and D. Esteve, "Characterization of a two-transmon processor with individual single-shot qubit readout", Physical Review Letters **108**, 057002 (2012).
[164] L. DiCarlo et al., "Demonstration of two-qubit algorithms with a superconducting quantum processor", Nature **460**, 240–244 (2009).
[165] R. Barends et al., "Diabatic gates for frequency-tunable superconducting qubits", Physical Review Letters **123**, 210501 (2019).
[166] M. J. Peterer et al., "Coherence and decay of higher energy levels of a superconducting transmon qubit", Physical Review Letters **114**, 010501 (2015).
[167] J. M. Martinis and M. R. Geller, "Fast adiabatic qubit gates using only σ_z control", Physical Review A **90**, 022307 (2014).
[168] Y. Chen et al., "Qubit architecture with high coherence and fast tunable coupling", Physical Review Letters **113**, 220502 (2014).
[169] 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).

[170] 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).
[171] S. Caldwell et al., "Parametrically activated entangling gates using transmon qubits", Physical Review Applied **10**, 034050 (2018).
[172] H. Paik et al., "Experimental demonstration of a resonator-induced phase gate in a multiqubit circuit-qed system", Physical Review Letters **117**, 250502 (2016).
[173] J. M. Chow et al., "Simple all-microwave entangling gate for fixed-frequency superconducting qubits", Physical Review Letters **107**, 080502 (2011).
[174] 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).
[175] A. D. Córcoles, J. M. Gambetta, J. M. Chow, J. A. Smolin, M. Ware, J. Strand, B. L. Plourde, and M. Steffen, "Process verification of two-qubit quantum gates by randomized benchmarking", Physical Review A **87**, 030301 (2013).
[176] A. A. Clerk, M. H. Devoret, S. M. Girvin, F. Marquardt, and R. J. Schoelkopf, "Introduction to quantum noise, measurement, and amplification", Reviews of Modern Physics **82**, 1155 (2010).
[177] R. Vijay, D. H. Slichter, and I. Siddiqi, "Observation of quantum jumps in a superconducting artificial atom", Physical Review Letters **106**, 110502 (2011).
[178] Z. R. Lin, K. Inomata, W. D. Oliver, K. Koshino, Y. Nakamura, J. S. Tsai, and T. Yamamoto, "Single-shot readout of a superconducting flux qubit with a flux-driven Josephson parametric amplifier", Applied Physics Letters **103**, 132602 (2013).
[179] C. Macklin, K. O'Brien, D. Hover, M. Schwartz, V. Bolkhovsky, X. Zhang, W. Oliver, and I. Siddiqi, "A near–quantum-limited Josephson traveling-wave parametric amplifier", Science **350**, 307–310 (2015).
[180] T. White et al., "Traveling wave parametric amplifier with Josephson junctions using minimal resonator phase matching", Applied Physics Letters **106**, 242601 (2015).
[181] D. Sank et al., "Measurement-induced state transitions in a superconducting qubit: beyond the rotating wave approximation", Physical Review Letters **117**, 190503 (2016).
[182] Z. R. Lin, K. Inomata, K. Koshino, W. D. Oliver, Y. Nakamura, J.-S. Tsai, and T. Yamamoto, "Josephson parametric phase-locked oscillator and its application to dispersive readout of superconducting qubits", Nature Communications **5** (2014).
[183] P. Krantz, "The Josephson parametric oscillator-from microscopic studies to single-shot qubit readout", PhD thesis (Chalmers University of Technology, 2016).
[184] E. Farhi, J. Goldstone, and S. Gutmann, "A quantum approximate optimization algorithm", arXiv:1411.4028 (2014).
[185] P. Vikstål, M. Grönkvist, M. Svensson, M. Andersson, G. Johansson, and G. Ferrini, "Applying the quantum approximate optimization algorithm to the tail assignment problem", arXiv:1912.10499 (2019).
[186] T. Stollenwerk, E. Lobe, and M. Jung, "Flight gate assignment with a quantum annealer", in International workshop on quantum technology and optimization problems (Springer, 2019), pp. 99–110.

[187] J. Haah, A. W. Harrow, Z. Ji, X. Wu, and N. Yu, "Sample-optimal tomography of quantum states", IEEE Transactions on Information Theory **63**, 5628–5641 (2017).
[188] S. Kimmel, C. Y.-Y. Lin, G. H. Low, M. Ozols, and T. J. Yoder, "Hamiltonian simulation with optimal sample complexity", npj Quantum Information **3**, 1–7 (2017).
[189] S. Lloyd, M. Mohseni, and P. Rebentrost, "Quantum principal component analysis", Nature Physics **10**, 631–633 (2014).
[190] J. Biamonte, P. Wittek, N. Pancotti, P. Rebentrost, N. Wiebe, and S. Lloyd, "Quantum machine learning", Nature **549**, 195–202 (2017).
[191] H. Pichler, G. Zhu, A. Seif, P. Zoller, and M. Hafezi, "Measurement protocol for the entanglement spectrum of cold atoms", Physical Review X **6**, 041033 (2016).
[192] I. Marvian and S. Lloyd, "Universal quantum emulator", arXiv:1606.02734 (2016).
[193] Y. Salathé et al., "Digital quantum simulation of spin models with circuit quantum electrodynamics", Physical Review X **5**, 021027 (2015).
[194] W. M. Itano, D. J. Heinzen, J. Bollinger, and D. Wineland, "Quantum zeno effect", Physical Review A **41**, 2295 (1990).
[195] E. Lucero et al., "Computing prime factors with a josephson phase qubit quantum processor", Nature Physics **8**, 719–723 (2012).
[196] J. Otterbach et al., "Unsupervised machine learning on a hybrid quantum computer", arXiv:1712.05771 (2017).
[197] G. Pagano et al., "Quantum approximate optimization with a trapped-ion quantum simulator", arXiv:1906.02700 (2019).
[198] L. K. Grover, "From Schrödinger's equation to the quantum search algorithm", Pramana **56**, 333–348 (2001).
[199] Z. Jiang, E. G. Rieffel, and Z. Wang, "Near-optimal quantum circuit for grover's unstructured search using a transverse field", Physical Review A **95**, 062317 (2017).
[200] M. Y. Niu, S. Lu, and I. L. Chuang, "Optimizing qaoa: success probability and runtime dependence on circuit depth", arXiv:1905.12134 (2019).
[201] A. Lucas, "Ising formulations of many np problems", Frontiers in Physics **2**, 5 (2014).
[202] P. J. O'Malley et al., "Scalable quantum simulation of molecular energies", Physical Review X **6**, 031007 (2016).
[203] K. Michielsen, M. Nocon, D. Willsch, F. Jin, T. Lippert, and H. De Raedt, "Benchmarking gate-based quantum computers", Computer Physics Communications **220**, 44–55 (2017).
[204] R. Barends et al., "Digital quantum simulation of fermionic models with a superconducting circuit", Nature Communications **6**, 1–7 (2015).
[205] F. Vatan and C. Williams, "Optimal quantum circuits for general two-qubit gates", Physical Review A **69**, 032315 (2004).
[206] R. R. Tucci, "An introduction to cartan's kak decomposition for qc programmers", arXiv:quant-ph/0507171 (2005).
[207] M. Kim, E. Park, P. Knight, and H. Jeong, "Nonclassicality of a photon-subtracted Gaussian field", Physical Review A **71**, 043805 (2005).

[208] A. Ourjoumtsev, A. Dantan, R. Tualle-Brouri, and P. Grangier, "Increasing entanglement between Gaussian states by coherent photon subtraction", Physical Review Letters **98**, 030502 (2007).

[209] M. Mirrahimi, Z. Leghtas, V. V. Albert, S. Touzard, R. J. Schoelkopf, L. Jiang, and M. H. Devoret, "Dynamically protected cat-qubits: a new paradigm for universal quantum computation", New Journal of Physics **16**, 045014 (2014).

[210] S. D. Bartlett, B. C. Sanders, S. L. Braunstein, and K. Nemoto, "Efficient classical simulation of continuous variable quantum information processes", Physical Review Letters **88**, 097904 (2002).

[211] M. Simoen, C. Chang, P. Krantz, J. Bylander, W. Wustmann, V. Shumeiko, P. Delsing, and C. Wilson, "Characterization of a multimode coplanar waveguide parametric amplifier", Journal of Applied Physics **118**, 154501 (2015).

[212] A. Heidmann, R. J. Horowicz, S. Reynaud, E. Giacobino, C. Fabre, and G. Camy, "Observation of quantum noise reduction on twin laser beams", Physical Review Letters **59**, 2555 (1987).

[213] A. S. Villar, L. S. Cruz, K. N. Cassemiro, M. Martinelli, and P. Nussenzveig, "Generation of bright two-color continuous variable entanglement", Physical Review Letters **95**, 243603 (2005).

[214] A. S. Coelho, F. A. S. Barbosa, K. N. Cassemiro, A. S. Villar, M. Martinelli, and P. Nussenzveig, "Three-color entanglement", Science **326**, 823–826 (2009).

[215] B. Yurke, L. R. Corruccini, P. G. Kaminsky, L. W. Rupp, A. D. Smith, A. H. Silver, R. W. Simon, and E. Whittaker, "Observation of parametric amplification and deamplification in a Josephson parametric amplifier", Physical Review A **39**, 2519 (1989).

[216] N. Roch, E. Flurin, F. Nguyen, P. Morfin, P. Campagne-Ibarcq, M. H. Devoret, and B. Huard, "Widely tunable, nondegenerate three-wave mixing microwave device operating near the quantum limit", Physical Review Letters **108**, 147701 (2012).

[217] S. Puri, S. Boutin, and A. Blais, "Engineering the quantum states of light in a Kerr-nonlinear resonator by two-photon driving", npj Quantum Information **3**, 1–7 (2017).

[218] R. G. Lyons, *Understanding digital signal processing* (Pearson Education India, 2004).

[219] L. Spietz, K. W. Lehnert, I. Siddiqi, and R. J. Schoelkopf, "Primary electronic thermometry using the shot noise of a tunnel junction", Science **300**, 1929–1932 (2003).

[220] L. Spietz, R. J. Schoelkopf, and P. Pari, "Shot noise thermometry down to 10 mk", Applied Physics Letters **89**, 183123 (2006).

[221] M. I. Dykman, C. M. Maloney, V. Smelyanskiy, and M. Silverstein, "Fluctuational phase-flip transitions in parametrically driven oscillators", Physical Review E **57**, 5202 (1998).

[222] K. V. Kheruntsyan and K. G. Petrosyan, "Exact steady-state wigner function for a nondegenerate parametric oscillator", Physical Review A **62**, 015801 (2000).

[223] N. C. Menicucci, P. van Loock, M. Gu, C. Weedbrook, T. C. Ralph, and M. A. Nielsen, "Universal quantum computation with continuous-variable cluster states", Physical Review Letters **97**, 110501 (2006).

[224] R. H. Koch, J. Clarke, W. Goubau, J. M. Martinis, C. M. Pegrum, and D. J. Van Harlingen, "Flicker (1/f) noise in tunnel junction dc squids", Journal of Low Temperature Physics **51**, 207–224 (1983).

[225] C. F. Spitzer, "Sustained subharmonic response in non-linear series circuits", Journal of Applied Physics **16**, 105–111 (1945).

[226] P. S. Linsay, "Period doubling and chaotic behavior in a driven anharmonic oscillator", Physical Review Letters **47**, 1349 (1981).

[227] G. T. Moore, "Quantum theory of the electromagnetic field in a variable-length one-dimensional cavity", Journal of Mathematical Physics **11**, 2679–2691 (1970).

[228] C. M. Wilson, G. Johansson, A. Pourkabirian, M. Simoen, J. R. Johansson, T. Duty, F. Nori, and P. Delsing, "Observation of the dynamical Casimir effect in a superconducting circuit", Nature **479**, 376–379 (2011).

[229] P. Lähteenmäki, G. Paraoanu, J. Hassel, and P. J. Hakonen, "Dynamical Casimir effect in a Josephson metamaterial", Proceedings of the National Academy of Sciences **110**, 4234–4238 (2013).

[230] J. R. Johansson, G. Johansson, C. Wilson, and F. Nori, "Dynamical Casimir effect in a superconducting coplanar waveguide", Physical Review Letters **103**, 147003 (2009).

[231] J. Johansson, G. Johansson, C. Wilson, and F. Nori, "Dynamical Casimir effect in superconducting microwave circuits", Physical Review A **82**, 052509 (2010).

[232] L.-M. Duan, G. Giedke, J. I. Cirac, and P. Zoller, "Inseparability criterion for continuous variable systems", Physical Review Letters **84**, 2722 (2000).

[233] A. A. Houck et al., "Generating single microwave photons in a circuit", Nature **449**, 328–331 (2007).

[234] M. Pechal, L. Huthmacher, C. Eichler, S. Zeytinoğlu, A. Abdumalikov Jr, S. Berger, A. Wallraff, and S. Filipp, "Microwave-controlled generation of shaped single photons in circuit quantum electrodynamics", Physical Review X **4**, 041010 (2014).

Appended Papers

Paper A

Noise and loss of superconducting aluminium resonators at single photon energies

Paper B

Decoherence benchmarking of superconducting qubits

Paper C

Characterizing decoherence rates of a superconducting qubit by direct microwave scattering

Paper D

Phononic loss in superconducting resonators on piezoelectric substrates

Paper E

Single-shot read-out of a superconducting qubit using a Josephson parametric oscillator

Paper F

A quantum instruction set implemented on a superconducting quantum processor

Paper G

Quantum approximate optimization of the exact-cover problem on a superconducting quantum processor

Paper H

Period-tripling subharmonic oscillations in a driven superconducting resonator

Paper I

Microwave photon generation in a doubly tunable superconducting resonator

Paper J

Period multiplication in a parametrically driven superconducting resonator

Paper K

Nondegenerate parametric oscillations in a tunable superconducting resonator

Paper L

Observation of broadband entanglement in microwave radiation from a single time-varying boundary condition

Paper M

Simple, robust, and on-demand generation of single and correlated photons