

THESIS FOR THE DEGREE OF DOCTOR OF PHILOSOPHY

Control techniques for superconducting quantum computer hardware

*New radio-frequency quantum gate implementations,
and electrical annealing of Josephson junctions*

CHRISTIAN KRIŽAN



Department of Microtechnology and Nanoscience
Chalmers University of Technology
Gothenburg, Sweden, 2026

C. Križan (**Author**) warrants that he has authored this thesis (**Work**), and that the Work contains no material in violation of copyright law. The Author grants Chalmers University of Technology the non-exclusive right to publish, and store, the Work electronically and in a non-commercial purpose make the Work publicly accessible on the Internet. Third parties gaining publishing rights to the Work must be (1) notified of, and (2) yield to, the publishing and storage rights hereby granted to Chalmers University of Technology.

© CHRISTIAN KRIŽAN, 2026

Under the supervision of:

Jonas Bylander, Professor ¹

Giovanna Tancredi, Ph.D. ¹

Miroslav Dobšiček, Ph.D. ²

As examined by:

Per Delsing, Full Professor ¹

¹ Quantum Technology division, department of Microtechnology and Nanoscience

² RISE Research Institutes of Sweden

Doktorsavhandlingar vid Chalmers Tekniska Högskola. Ny serie, № 5949

DOI 10.63959/chalmers.dt/5949

ISBN 978-91-81034-92-9

ISSN 0346-718X

Department of Microtechnology and Nanoscience

Chalmers University of Technology

SE-412 96 Gothenburg, Sweden

Phone: +46 31 772 1000

www.chalmers.se

Cover art by Rickard Laurenus, MSE.

SEM micrographs by Irshad Ahmad, Ph.D.

Printed by Chalmers Digitaltryck

Gothenburg, Sweden, September 2026

Abstract

The thesis investigates three challenges in superconducting quantum computing: efficient quantum gate compilation, scalable microwave interfacing of quantum processors (QPUs), and electrical annealing of Josephson junctions.

The first part addresses an inefficiency in contemporary two-qubit gate compilation. Universal quantum computation requires only a single entangling two-qubit gate alongside single-qubit rotations, but this minimal gate set requires up to three additional two-qubit gates for every unsupported gate. I show how access to multiple inequivalent two-qubit gate classes overcomes this inefficiency. I experimentally realised an unsupported SWAP operation using one parametric CZ and one parametric iSWAP gate on a circuit-QED architecture. A Kraus-Cirac analysis shows that this combination belongs to the SWAP equivalency class, the fundamental reason for the reduction in two-qubit gates. The experiment demonstrates a rare case of simultaneous CZ and iSWAP support in the same quantum circuit, illustrating how access to multiple inequivalent gate classes can reduce compilation overhead on real hardware.

The second part explores direct digital synthesis of microwave signals for interfacing with quantum processors. Conventional control uses heterodyne upconversion with IQ mixers, requiring compensation of LO leakage and IQ skew. I investigated a mixerless architecture based on high-speed DAC/ADC systems to generate and analyse the required microwave signals. CZ and iSWAP gate experiments subsequently verify that the platform is equivalent to conventional QPU setups.

The third part investigates post-fabrication electrical tuning of Josephson junctions, in which voltage pulses modify junction resistance to correct fabrication-induced transmon-frequency variations. The speed of resistance change was found to depend exponentially on pulse amplitude. A novel stepwise protocol achieved a 270% resistance increase in an 11-month-old 11.5 k Ω junction, corresponding to a $\sim$ 2 GHz reduction in transmon frequency, the largest reported selective resistance modification to date. Simulations show that resistance, charging-energy, and superconducting-gap uncertainties contribute equally to the uncertainty in the final qubit frequency. The results support a picture where electrical manipulation modifies the energy landscape, accelerating natural junction ageing.

The thesis demonstrates that investigating gate-level compilation strategies, fully digital QPU control, and post-fabrication tuning, helps advance the understanding of how to approach scalable superconducting quantum processors.

Keywords: amorphous materials, direct digital synthesis, electrical tuning, Josephson junctions, superconducting microwave circuits, quantum information science, qubit.

Acronyms

AC:	Alternating current
ADC:	Analogue-to-digital converter
API:	Application programming interface
AWG:	Arbitrary waveform generator
BPF:	Band-pass filter
CNOT:	Controlled-NOT gate
CZ:	Controlled-Z gate
DAC:	Digital-to-analogue converter
dB:	decibel
DC:	Direct current
DDS:	Direct digital synthesis
DRAQ:	Derivative removal by adiabatic gate
EBL:	Electron-beam lithography
ESD:	Electrostatic discharge
FLOPS:	Floating point operations per second
HEMT:	High-electron-mobility transistor
HPF:	High-pass filter
IF:	Intermediate frequency
IQ:	In-phase and quadrature
LO:	Local oscillator
LPF:	Low-pass filter
MSE:	Mean squared error
NISQ:	Noisy intermediate-scale quantum computing

NP:	Nondeterministic polynomial time
P:	Probability
QED:	Quantum electrodynamics; also, circuit-QED
QPU:	Quantum processing unit
RF:	Radio-frequency
RO:	Readout
Sa/s:	Samples per second
SEM:	Scanning electron microscopy
SPAM:	State preparation and measurement
SQUID:	Superconducting quantum interference device
SWAP:	SWAP gate; distinct from bSWAP, fSWAP, iSWAP
XX, YY, ZZ:	XX, YY or ZZ Pauli interaction

List of publications

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

[A] **Christian Križan**, Janka Biznárová, Liangyu Chen, Emil Hogedal, Amr Osman, Christopher W. Warren, Sandoko Kosen, Hang-Xi Li, Tahereh Abad, Anuj Aggarwal, Marco Caputo, Jorge Fernández-Pendás, Akshay Gaikwad, Leif Grönberg, Andreas Nylander, Robert Rehammar, Marcus Rommel, Olga I. Yuzepovich, Anton Frisk Kockum, Joonas Govenius, Giovanna Tancredi, and Jonas Bylander. “Quantum SWAP gate realized with CZ and iSWAP gates in a superconducting architecture,” Published in *New Journal of Physics*, Jul. 2025.

[B] Mats O. Tholén, Riccardo Borgani, Giuseppe Ruggero Di Carlo, Andreas Bengtsson, **Christian Križan**, Marina Kudra, Giovanna Tancredi, Jonas Bylander, Per Delsing, Simone Gasparinetti, David B. Haviland. “Measurement and control of a superconducting quantum processor with a fully integrated radio-frequency system on a chip,” Published in *Review of Scientific Instruments*, Oct. 2022.

[C] Mats O. Tholén, Riccardo Borgani, **Christian Križan**, Jonas Bylander, David B. Haviland. “Characterization and benchmarking of a phase-sensitive two-qubit gate using direct digital synthesis,” arXiv preprint, Aug. 2023.

[D] **Christian Križan**, Maurizio Toselli, Irshad Ahmad, Hadi Khaksaran, Marcus Rommel, Nermin Trnjanin, Janka Biznárová, Mamta Dahiya, Emil Hogedal, Halldór Jakobsson, Andreas Nylander, Jonas Bylander, Per Delsing, and Giovanna Tancredi. “Electrical post-fabrication tuning of aluminum Josephson junctions at room temperature,” Published in *Materials for Quantum Technology*, Jul. 2026.

Other publications by the author outside the scope of this thesis:

[E] Andreas Bengtsson, Pontus Vikstål, Christopher Warren, Marika Svensson, Xiu Gu, Anton Frisk Kockum, Philip Krantz, **Christian Križan**, Daryoush Shiri, Ida-Maria Svensson, Giovanna Tancredi, Göran Johansson, Per Delsing, Giulia Ferrini, and Jonas Bylander. *Improved Success Probability with Greater Circuit Depth for the Quantum Approximate Optimization Algorithm*, Physical Review Applied **14**, 034010, Sep. 2020.

[F] Sandoko Kosen, Hang-Xi Li, Marcus Rommel, Tahereh Abad, Anuj Aggarwal, Janka Biznárová, Lert Chayanun, Liangyu Chen, Göran Johansson, Anton Kockum, **Christian Križan**, Eleftherios Moschandreu, Andreas Nylander, Amr Osman, Jorge Fernández Pendás, Robert Rehammar, Anita Roudsari, Daryoush Shiri, Giovanna Tancredi, Tom Vethaak, Christopher Warren, Per Delsing, Jonas Bylander. *Building blocks of a flip-chip integrated superconducting quantum processor*, Quantum Science and Technology **7**, 035018, Jun. 2022.

[G] Liangyu Chen, Hang-Xi Li, Yong Lu, Christopher W Warren, **Christian Križan**, Sandoko Kosen, Marcus Rommel, Shahnawaz Ahmed, Amr Osman, Janka Biznárová, Anita Fadavi Roudsari, Benjamin Lienhard, Marco Caputo, Kestutis Grigoras, Leif Grönberg, Joonas Govenius, Anton Frisk Kockum, Per Delsing, Jonas Bylander, Giovanna Tancredi. *Transmon qubit readout fidelity at the threshold for quantum error correction without a quantum-limited amplifier*, npj Quantum Information **9**, 26, Mar. 2023.

[H] Christopher W Warren, Jorge Fernández-Pendás, Shahnawaz Ahmed, Tahereh Abad, Andreas Bengtsson, Janka Biznárová, Kamanasish Debnath, Xiu Gu, **Christian Križan**, Amr Osman, Anita Fadavi Roudsari, Per Delsing, Göran Johansson, Anton Frisk Kockum, Giovanna Tancredi, Jonas Bylander. *Extensive characterization of a family of efficient three-qubit gates at the coherence limit*, npj Quantum Information **9**, 44, May 2023.

[I] Janka Biznárová, Amr Osman, Emil Rehnman, Lert Chayanun, **Christian Križan**, Per Malmberg, Marcus Rommel, Christopher Warren, Per Delsing, August Yurgens, Jonas Bylander, Anita Fadavi Roudsari. *Mitigation of interfacial dielectric loss in aluminum-on-silicon superconducting qubits*, npj Quantum Information **10**, 78, Aug. 2024.

[J] Sandoko Kosen, Hang-Xi Li, Marcus Rommel, Robert Rehammar, Marco Caputo, Leif Grönberg, Jorge Fernández-Pendás, Anton Frisk Kockum, Janka Biznárová, Liangyu Chen, **Christian Križan**, Andreas Nylander, Amr Osman, Anita Fadavi Roudsari, Daryoush Shiri, Giovanna Tancredi, Joonas Govenius, Jonas Bylander. *Signal Crosstalk in a Flip-Chip Quantum Processor*, PRX Quantum **5**, 030350, Sep. 2024.

Contribution to appended papers

The appended papers summarise my key scientific contributions during my PhD. I made additional contributions to work beyond these papers, but they represent my main body of research on which this thesis is based.

My contributions to Papers A and D span experimental design, execution, data collection, data analysis, post-processing, writing of the papers, and in both cases installing another dilution refrigerator.

Papers B and C were made in a scientific collaboration, and therefore merit a more detailed account of my specific contributions. These contributions are described below.

Paper A

The experimental concept stems from regular scientific dialogue with the neighbouring quantum theory division at our institution. I designed the experimental implementation and developed several parameter-estimation techniques for characterising two-qubit gates. I performed all experiments, analysed their data, and performed additional post-processing in collaboration with co-authors.

I assembled a substantial portion of the measurement infrastructure, ranging from cryogenic hardware to the room-temperature microwave interface. I also developed further the experimental control software, creating interface code for the AWG and digitiser, and extended the laboratory framework code to support the novel two-qubit gates.

The manuscript was written by me, with input from co-authors. Analysis relating to experimental error, readout misclassification, and qubit frame tracking was carried out in collaboration with co-authors.

Papers B and C

Apart from continuous scientific collaboration including regular feedback, visits, discussions, planning and suggestions, my specific work pertaining to these papers is the following.

- Investigation (supporting)
 - Planned and structured the measurement campaign for initial Presto-

-
- based experiments for qubit characterisation.
- Performed experiments using Presto prototypes at Chalmers.
 - Developed and shared measurement procedures and reference implementations for phase-coherent pulse sequences.
 - Contributed methodology for maintaining phase coherence in driven two-qubit operations. For instance, iSWAP-type interactions.
- Resources (supporting)
 - Made experimental hardware, including sample box construction, assembly, wiring, and wirebonding of chips. And, replaced damaged samples.
 - Installed and integrated the Presto prototypes at Chalmers.
 - Implemented DC biasing interface for protecting instrument ports.
 - Validation (supporting)
 - Cross-validated measurement implementations by exchanging and benchmarking experimental control code across sites.
 - Verified consistency of measurement results obtained using Presto across independent implementations.
 - Writing - review & editing (supporting)
 - Critically reviewed the manuscript focusing on evidential support for its claims.
 - Identified insufficient support for two-qubit gate demonstrations and contributed revisions to align claims with presented data.
 - Provided feedback on figures, analysis, and technical descriptions.
 - Other instrument-related development
 - Developed example control routines for phase-coherent waveform generation and timing-sensitive experiments.
 - Contributed to experimental strategies for implementing and evaluating two-qubit gate protocols using DDS-based control.
 - Broader contribution
 - In addition to the work directly underlying Paper B and Paper C, I con-

tributed to the development of the Presto platform through specification of its streaming mode prototype. I co-designed the communication protocol intended for use between host PC and instrument, and I supervised implementation work carried out in the context of two master's industry courses with improving the Presto being the end-deliverable. I also contributed to hardware-level system design decisions, including evaluation of biasing solutions for SQUID control and port layout requirements based on target multi-qubit device architectures, which I specified.

Paper D

The electrical tuning project originated as a master's thesis project by Maurizio Toselli [1], based on the junction annealing technique introduced in [2].

I extensively upgraded the experimental platform to make it ESD-safe. Other upgrades included a major revision of its measurement software. I also assembled the cryogenic experimental setup, with input from colleagues. With a small number of exceptions, the experiments reported in the paper were carried out by me; these exceptions are the ageing study, breakdown-voltage analysis, resistance-change measurements for the *Low-dose 1* and *Medium-dose 1* samples, and the liquid-nitrogen relaxation experiment.

Data analysis and all post-processing was performed by me, with two exceptions; the ageing study and liquid-nitrogen relaxation analysis were conducted with input from MT, while the breakdown-voltage post-processing was performed by Andreas Nylander.

I developed the theoretical work used for qubit-frequency prediction, including all related code and its formulaic implementation; it must be noted that the underlying derivation connecting junction resistance to qubit frequency is long-established within the field. Reference values for the superconducting energy gap and the ratio $R_{\text{cryogenic}}/R_{\text{room temperature}}$ were experimentally acquired by co-authors.

Several scientific directions emerged through discussions with collaborators: the liquid-nitrogen relaxation study originated in discussions with Jürgen Lisenfeld at KIT, while discussions with Prof. Per Delsing motivated investigation of resistance-manipulation speed as a function of voltage. The hypothesis that electrical tuning is fundamentally related to modifications of the barrier energy landscape emerged through discussions among multiple co-authors.

The manuscript was written by me, with input from co-authors.

Contents

Abstract	I
Acronyms	II
List of publications	IV
Contribution to appended papers	VII
 I Thesis	 XV
1 Introduction and methodological overview	1
1.1 Experimental workflow and thesis outline	3
1.1.1 Quantum gate implementation and an extension of the gate set	3
1.1.2 QPU input/output signal interface	5
1.1.3 Adjusting superconducting circuit parameters post-fabrication	7
1.1.4 Acquiring data	9
1.2 Introduction summary	10
 2 SWAP via the inequivalent CZ and iSWAP gate classes	 11
2.1 The taxonomy of two-qubit gates	12
2.2 Energy band dynamics of the two-qubit QPU	13
2.3 SWAP gate compilation	15
2.4 Kraus-Cirac decomposition of two-qubit gates	18

2.5	Local equivalency classes	20
2.6	Experimental verification via conditional- and cross-Ramsey interferometry	22
2.6.1	Readout state assignment setup	25
2.6.2	Error mitigation	27
2.7	Sources of experimental error	29
2.7.1	Energy relaxation and low-frequency noise	30
2.7.2	Readout error	32
2.8	Concluding remarks on experimental quantum informatics	33
3	Quantum processor control and readout using DDS	35
3.1	Direct digital synthesis approach to quantum computer interfaces . .	38
3.2	Quantum experiments for DDS verification	40
3.3	Concluding remarks on DDS control	44
4	Resistance tuning of Josephson junctions	45
4.1	Frequency crowding in superconducting quantum processors	46
4.2	Junction resistance and qubit frequency	48
4.3	Electrical resistance tuning	50
4.3.1	Sample information and manipulator setup	51
4.3.2	The voltage pulse technique	53
4.4	Active resistance manipulation	55
4.5	Residual resistance increase post-manipulation	60
4.5.1	Tuning frequency precision from the relaxation offset error . .	64
4.5.2	Impact from error in other parameters	65
4.6	Practical application	65
4.7	Concluding remarks on resistance manipulation	68
5	Conclusion and outlook	71
5.1	Future work	72
	References	75
II	Appended papers	87
A	Quantum SWAP gate realized with CZ and iSWAP gates in a superconducting architecture	A1
B	Measurement and control of a superconducting quantum processor with a fully integrated radio-frequency system on a chip	B1

C	Characterization and benchmarking of a phase-sensitive two-qubit gate using direct digital synthesis	C1
D	Electrical post-fabrication tuning of aluminum Josephson junctions at room temperature	D1

Part I

Thesis

CHAPTER 1

Introduction and methodological overview

Quantum computers promise to unlock solutions to certain computational problems that are unsolvable by conventional computers within the lifetime of the universe [3], [4]. Whether a problem is solvable depends on its computational complexity, that is, how the resources required to solve it, such as execution time, memory consumption, or the time needed to verify a solution, scale with the size of the input. Some problems grow only modestly in complexity. For instance, sorting a list of 20 names remains fairly straight-forward even if the list doubles in length. Other problems grow dramatically in complexity. Finding the shortest path through 40 cities is 10^{29} times harder than with 20 cities. Quantum computers are believed to provide advantages over conventional, classical computers for problems of scientific and industrial interest that fall within certain complexity classes. Such problems can be found [5]–[7], for example, in chemistry, materials science, logistics optimisation, and cryptanalysis.

This promise does not come about because of faster clock speed, but rather because quantum computers have access to exotic information processing techniques enabled by quantum principles, such as the no-cloning theorem and entanglement [6]. In the gate model of quantum computation, a quantum computer's memory consists of a register of n quantum bits of information, *qubits*. Each qubit is associated with a complex amplitude which describes its state within the space spanned by

the logical-zero $|0\rangle$ and logical-one $|1\rangle$. Crucially, the state of a qubit can be put in a quantum superposition, $\alpha|0\rangle + \beta|1\rangle$, with amplitudes α and β being complex coefficients, where $|\alpha|^2 + |\beta|^2 = 1$. The n -qubit entangled state, which is a quantum superposition states of several qubits, is represented by 2^n amplitudes: $|\psi\rangle = c_0|00\dots00\rangle + c_1|00\dots01\rangle + c_2|00\dots10\rangle + \dots + c_{2^n-1}|11\dots11\rangle$. A quantum algorithm orchestrates the quantum interference of entangled states, such that the qubit “waves” interfere to impart a high probability $|c_i|^2$ on the qubits that represent the answer, and essentially zero probability to the remaining qubits. This represents the essential advantage of quantum computers compared to a conventional algorithm, which can only be carried out on one state at a time.

Even as conventional, non-quantum computers continue to develop, quantum computer scientists believe that quantum computers will outperform conventional computers when applied to certain types of problems [3]. This gives us reason to believe that you and I can start asking, and answering, questions to which we previously would not have lived long enough to see the answer.

In order to implement a quantum computer in a physical device, we must represent the qubits in a quantum system with two distinct controllable states [3], or rather, a controllable two-dimensional subspace [8]. To realise practical quantum computation, several physical platforms have been pursued, which encode and manipulate quantum information in different ways [9] and therefore face different engineering trade-offs. There is not yet any universal consensus regarding which architecture will ultimately prove the most suitable for large-scale quantum computation.

Some of the most mature approaches to quantum information processing include trapped ions, neutral atoms, and spins in a semiconductor. Trapped ion computers encode information in long-lived, internal electronic states of individual ions confined in electromagnetic traps and are manipulated using laser interactions [8]–[10]. These systems are known for quantum states with minute-long coherence times and high quality of operations, that is, *fidelity*, albeit with physical scaling challenges [10], [11] and comparatively low speed of operations. Neutral atom quantum computing encodes qubits in selected atomic electronic states [12], while interactions between highly excited Rydberg states can be used to implement multi-qubit operations [13]–[16]. Semiconductor spin qubits instead encode information in electron or nuclear spin states, typically in quantum dots or donor atoms embedded in silicon, motivated in part by compatibility with established CMOS fabrication technologies [17]. Multiple additional approaches exist, such as photonic quantum computers [9] and solid-state defect systems [18]. A review of alternative platforms for quantum computation is outside the scope of this thesis.

My work pertains to several hardware aspects of superconducting quantum computation, the perhaps most successful platform to date [19]. This platform is pursued

by industrial actors such as IBM and Google, startup companies, and many research groups. Its maturity has been engineered to a state where cloud access platforms provide outside users an opportunity to configure and run their algorithms on physical quantum processors.

This thesis is based on four journal papers, summarised in the next section. Common for all chapters is the overview shown in figure 1.1, where I will highlight the subject matter of the different chapters. The papers contribute to the experimental areas of quantum information science, quantum processor control hardware, and post-fabrication electrical annealing of *Josephson junctions*; the Josephson junction is the fundamental non-linear, inductive circuit element that enables superconducting qubits and therefore underpins all superconducting quantum processors. An introduction for this component is given later.

I will frequently refer to the machine, which uses quantum phenomena to manipulate properties of information according to a list of instructions, as a *quantum processing unit* (QPU).

1.1 Experimental workflow and thesis outline

This section provides an overview of the experimental workflow and introduces the common infrastructure. All experiments follow a similar procedure, shown in figure 1.1: at the highest level, a computational task is translated into control signals, propagated into a cryogenic environment, executed on a superconducting quantum processor, and finally measured by studying the microwave transmission.

1.1.1 Quantum gate implementation and an extension of the gate set

Generally in computation, a computer carries out an algorithm, a sequence of instructions to be executed in order to solve a particular task, that is, a program. The QPU's program is known as a quantum circuit: a sequence of quantum operations that controllably manipulate a quantum state toward some desired state. Operations in these circuits are often represented graphically, as can be seen in figure 1.1(a).




Figure 1.1: Overview of an example setup for the quantum experiments shown throughout this thesis. (a) A quantum circuit contains the algorithmic description of the experiment the experimentalist wishes to run, see Chapter 2. This description of gates is compiled into instrument-specific instructions and uploaded onto (b) the QPU interface. The interface can be simplified as is done in (c), where much of the signal conditioning equipment is removed, see Chapter 3. The microwave signals are sent into (d, e) the cryogenic RF environment. (f) An example QPU chip is shown in figure 1.2, with figure 1.3 showing its fundamental elements: the Josephson junction, see Chapter 4.

The operations described in the quantum circuit must be translated into physical control signals that act on the qubits. In our superconducting qubit platform, each operation, that is, *quantum gate*, is implemented using microwave pulses that drive the transitions between qubit states. These pulses are characterised by waveform parameters such as pulse duration, amplitude, carrier frequency, and phase, all of which must be set correctly for that pulse to correspond to a particular gate.

Determining a gate’s pulse parameters is achieved through parameter estimation experiments, commonly referred to as *tune-up*, extensively covered in the literature [20], [21]. The hardware architecture of the QPU determines which operations are supported natively, collectively known as the QPU’s *gate set*. When compiling quantum circuits into some sequence of gates to be executed, a larger gate set can allow for a more efficient compilation: this is the topic of Paper A in Chapter 2.

The control layer, which determines which gates to implement in which order, is implemented using quantum circuit compiler software, such as Qiskit [22]. The resulting list of instructions is processed by a regular desktop computer, which acts as a translator between this list and hardware-specific control instructions.

Chapter 2 / Paper A: *Quantum SWAP gate realized with CZ and iSWAP gates in a superconducting architecture.*

Here I used a qubit architecture comprising fixed-frequency qubits and tunable qubit couplers to demonstrate CZ and iSWAP two-qubit gates in sequence, thereby realising a SWAP operation by using an extended gate set. The experiment demonstrates that an extended gate set enabled a more efficient decompilation of the SWAP operation than our platform would have otherwise required. This possibility emerges because the CZ and iSWAP gates belong to different so-called locally inequivalent two-qubit gate classes, further explained in Chapter 2. These gates are individually well established in the literature, but simultaneous utilisation of both in the same quantum circuit is rare, with previous experimental works focusing on supporting the gates in separate circuits [21], [23]. Simultaneous support for the iSWAP and CZ families has been shown previously in tunable-qubit architectures, using fSim gates whose settings made these gates CZ- and iSWAP-equivalent [24]. It has recently been proposed for a neutral-atom architecture [16]. The work is motivated by a desire to reduce quantum circuit depth, that is, the number of gates in sequence.

1.1.2 QPU input/output signal interface

Signalling interfaces to superconducting QPUs usually consist of microwave instruments placed outside of the refrigerator, which launch pulses into cables leading to the QPU placed at low temperature. This interface is illustrated in figure 1.1(b).

The list of instructions that make up the gate sequence is translated into instrument-specific control signals by instrument control software. Such software is often provided by the hardware vendors, where they expose instrument functions through an API. It is the task of an automation framework to steer the instruments into sending and receiving the microwave pulses as needed [25].

The relevant control frequencies for superconducting qubit architectures typically lie in the 0.2 GHz to 8 GHz range, depending on the state transitions being addressed; an example is shown in figure 2.1(a) on page 13. A conventional approach to generating signals in this band is to use arbitrary waveform generators (AWG) to synthesise waveforms at an intermediate frequency (IF) in the few-hundred MHz range. This IF content is subsequently upconverted to the target microwave band using analogue mixers. Such mixers may be implemented as stand-alone components or integrated into commercial instruments. Signal detection is done using digitisers that generate their own pilot tone, conceptually similar to a lock-in amplifier [26].

An alternative approach is highlighted in figure 1.1(c), where direct digital synthesis (DDS) is used to perform frequency conversion digitally, removing a lot of the hardware used in (b). The QPU control signals are synthesised directly in the target frequency band. This approach forms the basis of the QPU interface developed in Papers B and C. We demonstrated that an all-DDS control system can reproduce the QPU parameter estimations that are standard in superconducting quantum computing. We also interfaced to a circuit scheduler [22] for calculating the gate-level control signals needed.

Chapter 3 / Papers B and C: *Measurement and control of a superconducting quantum processor with a fully integrated radio-frequency system on a chip and Characterization and benchmarking of a phase-sensitive two-qubit gate using direct digital synthesis.*

These papers showcase an all-DDS QPU interface. Complexity and scalability aspects in control electronics remain one of the major technological challenges for large-scale quantum processors. The work therefore investigates DDS as an enabling technology for future quantum systems. The research was done in collaboration with the group of David Haviland at the KTH Royal Institute of Technology; the work contributed to technology transfer beyond academia, and the resulting instrument is now a product of the spin-off company Intermodulation Products from KTH.

The generated microwave control signals are launched into a microwave signal path setup inside a dilution refrigerator [29], illustrated in figure 1.1(d) and photographed in figure 1.1(e). The millikelvin temperatures of the setup (far below what is required for keeping the aluminium inside our QPUs superconducting) serves to suppress erroneous thermal excitations of higher quantum states, that is, the energy




Figure 1.2: (a) Wirebonded device for the experiments performed in Papers B and C, and the corresponding verification experiments done at Chalmers. (b) Micrograph of the chip, consisting of two fixed-frequency transmon qubits coupled together by a tunable coupler, with two resonator elements sharing a feedline for reading out the qubits' state. The numbers provided are in units of GHz, with the coupler stated in its flux-biased spectral position 4.740 GHz, and its unbiased frequency 6.254 GHz as experimentally acquired by fitting a flux-swept two-tone spectroscopy to equation (2) of [27], when biased to $0.25 \Phi/\Phi_0$. An equivalent schematic of this QPU will be shown in figure 2.1(b). The chip was originally fabricated by my master's thesis supervisor for [28], while the box design is by Marco Scigliuzzo.

separation between adjacent eigenstates $\hbar\omega_r$ must be greater than the thermal energy $k_B T$ [30]. The signal path itself consists of multiple stages of attenuation and filtering [31] serving to thermalise the incoming microwave lines to the temperature of each cryogenic plate inside the cryostat, while suppressing thermal noise originating from warmer plates. The signal path is therefore heavily filtered to suppress spurious spectral content and environmental noise.

The signals are delivered to the QPU, seen in figure 1.1(f). The processor used in Papers B and C is shown in greater detail in figure 1.2; this chip architecture was later redesigned into a flip-chip implementation [32], [33], subsequently used for the experiments in Paper A.

1.1.3 Adjusting superconducting circuit parameters post-fabrication

The cross-shaped features visible in 1.2(b) are specific to a particular superconducting qubit modality known as a *transmon* [34], and to be even more specific, the *Xmon* [35]; the cross constitutes the capacitive element of an inductance-capacitance (LC) cir-

cuit. To render the qubit energy spectrum anharmonic, and hence the state transitions distinct, a non-linear element is required. This requirement is solved by connecting a Josephson junction [36] at the base of each structure, acting as a non-linear inductance. A junction is highlighted in the micrographs of figure 1.3.




Figure 1.3: Micrographs showing progressively higher magnification of transmon qubit structures from a wafer¹ presented in Paper D. (a) Optical micrograph of an Xmon qubit's main capacitor "island". The highlighted region indicates (b), a scanning electron micrograph (SEM) of the Josephson junction. The junction is further magnified in the false-coloured SEM (c), revealing the two overlapping aluminium electrodes. The false-colouring corresponds to (d) a simplified, not-to-scale illustration of the junction geometry with approximate dimensions. The oxide formed between the electrodes is Al_2O_3 alumina. The illustration omits other oxides for clarity; an illustration including such native oxide formation can be seen in [37].

Josephson junctions are nanoscale structures with lateral dimensions of a few hundred nanometres. They are therefore difficult to manufacture with high precision, particularly in a university laboratory. During fabrication, small deviations appear in the junctions' cross-sectional area. These variations cause the electrical junction parameters to deviate from their intended design values [38], [39], particularly the transition frequency of the qubit. This variability introduces the risk of frequency

¹This wafer will be known as the *high-dose* wafer in Chapter 4.

crowding once QPU designs grow in size. Control signals intended for one qubit would inadvertently drive neighbouring qubits [40]. An example of this challenge will be shown graphically in figure 4.1 on page 48.

Therefore, it is interesting to explore methods enabling the manipulation of the qubit frequencies once the QPU has been manufactured. Post-fabrication correction of junction parameters has previously been performed mostly by heating the junction [41]–[49], where secondary effects are not yet fully understood [50]. A more recent alternative is *electrical tuning*, where voltage pulses are applied across the Josephson junction to modify its electrical properties [2], [51]–[53]. This approach remains comparatively new, and the physical mechanisms are incompletely understood. Increasing this understanding was the motivation for the research reported in Paper D, which took an empirical modelling approach to devise an electrical tuning method to adjust the frequency of a qubit post-fabrication.

Chapter 4 / Paper D: *Electrical post-fabrication tuning of aluminum Josephson junctions at room temperature.*

The chapter deals with the fabrication variability of Josephson junctions, and a method for compensating for this variability post-fabrication. A recent approach is electrical tuning, where the resistance of a Josephson junction is altered by exposing it to voltage pulses. The underlying physical mechanism remains insufficiently understood, and therefore Paper D presents a study of five different junction variants, aiming to provide an empirical characterisation. The goal was to assess the suitability of electrical tuning for mitigating qubit frequency crowding in a multi-qubit system. In my work, I demonstrated a selective resistance tunability range significantly beyond what had previously been reported in the literature. Furthermore, my work establishes that the speed of resistance change is exponentially dependent on the manipulation voltage, providing clues of the physical mechanisms. The study also shows that the spontaneous resistance drift that occurs after manipulation depends on the preceding intentional resistance change, which was previously assumed not to be the case. Post-fabrication adjustment of Josephson junctions may prove useful as quantum processors grow in size, where frequency crowding may be a challenge in certain QPU architectures.

1.1.4 Acquiring data

The final stage of an experiment is to extract information about the quantum state. In the predominant circuit-QED architecture, the qubit is coupled to a microwave resonator, whose resonant frequency depends on the state of the qubit. The qubit state can therefore be read out by probing the resonator: the outgoing signal acquires a state-dependent shift in amplitude and phase. These resonators are the

meandering structures in figures 1.1(e) and 1.2(b).

The readout approach resembles the principle of a lock-in amplifier [26], where a known reference signal is injected and its modification after interacting with a system is analysed. In our case, the readout signal is the reference signal, launched into the cryostat through the microwave readout chain.

Depending on the hardware architecture, qubit control signals may either share the same signal path, as was the case in Paper D; or, the control signals can be delivered through dedicated control lines, as was the case in Papers A, B, and C. Dedicated control lines allow for stronger coupling between the control signals and the qubits, as well as more selective control as each signal mainly reaches its target qubit [33]. Dedicated control lines increase wiring complexity and thermal load inside the cryogenic environment, but is generally of little concern for chips used in academic research.

Following interaction with the QPU, the weak microwave signal propagates back to room temperature through a chain of isolators, filters, and cryogenic amplifiers. Once back at room temperature, the signal is digitised and processed, with the resulting readout being a coordinate in magnitude-and-phase space [54], [55].

In this manner, a quantum experiment has completed its full cycle: a computational task expressed at room temperature was translated into microwave control signals, executed on a QPU chip, and finally reconstructed into experimentally accessible information.

1.2 Introduction summary

This introduction established the common experimental framework of the quantum experiments presented throughout this thesis. Starting from a quantum circuit, we followed how instructions are compiled into control signals, transmitted through a cryogenic measurement infrastructure, executed on a superconducting quantum processor, and finally converted back into experimentally accessible data through microwave readout.

The following chapters examine a different layer of this workflow in greater detail. Chapter 2 focuses on quantum gate compilation and a more efficient utilisation of the QPU's native gate set. Chapter 3 addresses the hardware interface between classical electronics and the quantum processor, with particular emphasis on using DDS for microwave control and readout. Chapter 4 moves further down the hardware stack to the Josephson junction itself, investigating electrical tuning as a means of mitigating frequency crowding.

CHAPTER 2

SWAP via the inequivalent CZ and iSWAP gate classes

This chapter expands on the theoretical modelling and experimental results presented in Paper A. There I demonstrated that a two-qubit SWAP operation, which is not native to our device platform, can be implemented more efficiently than what has been previously shown on this circuit-QED [30] architecture. This chapter elaborates on the theoretical background for this capability and discusses its relevance to more-efficient quantum circuit compilation.

The introduction in the previous chapter presented that the gate model of quantum computation denotes the set of supported operations as its *gate set*. Quantum computing platforms support different sets of two-qubit gates depending on their architecture. Different two-qubit gates in turn have advantages when faced with peculiarities such as the available qubit connectivity [56]–[59], or the presence of unmodelled interactions leading to coherent errors and incoherent errors from dominant noise channels [60]–[62]. Typically, a platform has access to some, but not all, types of two-qubit gates.

Access to single-qubit rotations together with an entangling two-qubit gate is theoretically sufficient for universal quantum computation [63], [64]. But, a larger native gate set can be advantageous for compiling quantum circuits into the shortest possible implementation on the targeted architecture. I will later detail how

two-qubit gates can be grouped into *local equivalency classes*, whose members can be transformed into one another using only single-qubit gates. Gates belonging to different local equivalency classes require the insertion of additional two-qubit operations for their implementation. Such a compilation strategy could be acceptable if all gates are short and execute with exceptional fidelity. However, decoherent (and coherent) errors represent a persistent problem for the implementation of quantum gates in superconducting platforms, which is why access to shorter circuits should be welcomed.

2.1 The taxonomy of two-qubit gates

Most superconducting quantum computing architectures currently support two-qubit gates from the Clifford group. This group contains four *classes* of locally inequivalent two-qubit gates [65]: CNOT-like gates, such as the CZ gate; iSWAP-like gates; SWAP-like gates; and Identity-like gates, which comprise all combinations of single-qubit gates [66]–[69].

The superconducting circuit-QED architecture used in Papers A, B and C, interconnects qubits with a parametrically modulated coupler. Such an architecture gives a highly tunable interaction landscape where distinct exchange interactions between transmon [34] states can be selectively activated. This architecture enables CZ-like, iSWAP-like, and other two-qubit gate classes [21], [27], [70]–[74]. But, the SWAP-like class is notably absent from most superconducting platforms [56], [75] despite playing a central role in architectures with limited qubit connectivity.

This missing SWAP gate class is useful: it enables interactions between distant qubits by sending states, along a chain of qubits, to some destination qubit where it is needed [12], [57], [59]. SWAP-like gates are also the only gates that do not introduce additional entanglement for any separable input state [76].

The common solution addressing this missing SWAP gate class is found in quantum circuit compilation, where the gate compiler can supplement a SWAP by inserting three CZ, iSWAP, or CNOT gates [8], [77]–[79]. This compilation strategy imposes a significant resource cost in NISQ-era devices, because two-qubit gates are slower and execute with lower fidelity.

In Paper A, we therefore demonstrated experimentally that this three-gate substitute strategy can be improved in our platform by leveraging the simultaneous availability of its CZ and iSWAP gates. Such reductions can directly reduce the circuit depth, and may also prove beneficial in quantum error correction settings, where SWAP-heavy circuits arise during encoding and decoding [80].




Figure 2.1: (a) Scale drawing of the energy-level diagram for our two-qubit system $|Q_1 Q_2\rangle$, where the coupler remains in the ground state for all experiments. The native two-qubit gate interactions used in Paper A are indicated. (b) Equivalent circuit-QED schematic of the hardware architecture, with resonators in green, transmons in red, the coupler highlighted in turquoise, and signal wiring in purple. The labelled ports correspond to those in figure 1.1(d) on page 4, with XY being signals for single-qubit excitation, Z for coupler signals, and RO the transmitted readout signal.

2.2 Energy band dynamics of the two-qubit QPU

The reason why different QPU chips support different two-qubit gates is because of the effective Hamiltonian of the system. The available components on-chip and how they are interconnected, define the interaction terms, and thus the native gate set. The equivalent circuit-QED schematic of the QPU chips used in this chapter is shown in figure 2.1 alongside an energy band diagram in the same figure. This schematic is also valid for the QPU in Chapter 3, albeit with different energy levels since the qubit frequencies are different, shown earlier in figure 1.2(b) on page 7.

The schematic in figure 2.1(b) equates to a three-body system of multi-level anharmonic oscillators [34], with the two lowest energy levels of the oscillators defining our qubits’ computational basis states. The turquoise coupler structure in the middle of figure 2.1(b) is constructed using a SQUID [81], [82], in turn also an anharmonic oscillator. The circuit’s behaviour is described using the following Hamiltonian:

$$\begin{aligned}
\frac{\hat{H}}{\hbar} = & \sum_{i=1}^2 \omega_i \hat{a}_i^\dagger \hat{a}_i + \frac{\eta_i}{2} \hat{a}_i^\dagger \hat{a}_i (\hat{a}_i^\dagger \hat{a}_i - 1) \\
& + \omega_c(\Phi) \hat{b}^\dagger \hat{b} + \frac{\eta_c}{2} \hat{b}^\dagger \hat{b} (\hat{b}^\dagger \hat{b} - 1) \\
& + \sum_{i=1}^2 g_{ic} (\hat{a}_i^\dagger + \hat{a}_i) (\hat{b}^\dagger + \hat{b})
\end{aligned} \tag{2.1}$$

where $\hat{a}$, $\hat{b}$, $\hat{a}^\dagger$, $\hat{b}^\dagger$ are annihilation and creation operators [83], with $\hat{a}_i$ acting on transmon i and $\hat{b}$ acting on the coupler; g_{ic} are the coupling strengths between qubits i and the coupler c ; ω_i are the angular frequencies of all three transmons, where the coupler transmon’s frequency is adjustable by a threading magnetostatic flux Φ ; η_i and η_c are the transmons’ anharmonicities [34], with suffix c denoting the coupler in particular; and $\hbar = \frac{h}{2\pi}$ is the reduced Planck constant. It is assumed that the direct qubit-qubit coupling $g_{12} \ll (g_{1,c}, g_{2,c})$. Throughout this thesis, I follow the notation $\omega_i = 2\pi f_i$, where ω denotes an angular frequency and f the ordinary frequency.

Microwave modulation of the coupler frequency mediates interactions between the two fixed-frequency qubits, that is, two-qubit operations are driven by sending microwave pulses onto the coupler to address transitions in the two-qubit manifold. Such parametric exchange coupling [27], [84], [85] enables the coupling of widely detuned, fixed-frequency qubits [27], although with a known drawback of being slow compared to alternatives such as cross-resonance gates [86].

The resulting two-qubit gates arise from distinct resonant processes. An iSWAP interaction relies on $XX + YY$ exchange interaction [87], activated by driving the $|01\rangle \rightarrow |10\rangle$ transition at frequency $\omega_{10} - \omega_{01}$, indicated in figure 2.1(a). The CZ gate is implemented by driving the $|11\rangle \rightarrow |20\rangle$ or $|11\rangle \rightarrow |02\rangle$ transition, producing a controlled interaction mediated by a brief excursion of the qubit state outside of the computational subspace and back again, for a gate duration τ_{CZ} ; the completed round-trip adds a geometric phase of π .

It should be mentioned that the bSWAP gate is also natively available via the $XX - YY$ exchange interaction, by driving a two-photon $|00\rangle \rightarrow |11\rangle$ transition [72], [87], [88], although this gate has not been explored in my work. Following the

procedure in section 2.5, the bSWAP gate can be shown to be an iSWAP-like gate, as is also claimed in [89]; this native gate would therefore not provide any benefit to the experiment shown in Paper A, since only the single-qubit gates used would change.

Our research group had previously shown parametric CZ support in our platform [28]; parametric CZ and iSWAP support in this architecture was known to the field before my experiments [21], [70], [71], thus we believed I would be able to add iSWAP support alongside the CZ gate, even if such simultaneous support in the same quantum circuit had not been shown previously to paper A in fixed-frequency-qubit architectures.

2.3 SWAP gate compilation

Assume right-handed rotations about Bloch sphere axis $N \in [X, Y, Z]$, where $|0\rangle$ is located at the north pole. *Increasing* a phase denotes a counter-clockwise Z-axis rotation of the Bloch vector when viewed from $|0\rangle$. The basis elements of the two-qubit computational states is given as $[|00\rangle, |01\rangle, |10\rangle, |11\rangle]^T$, and accordingly, the two-qubit gates in this thesis are defined as

$$\begin{aligned}
 U_{CZ} &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, & U_{\text{iSWAP}} &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \\
 U_{\text{SWAP}} &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, & U_{\text{CNOT}} &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.
 \end{aligned} \tag{2.2}$$

For this section, you should be aware of conjugate-transpose phase gates ($S^\dagger$) and Hadamard (H) gates,

$$S^\dagger = \begin{pmatrix} 1 & 0 \\ 0 & -i \end{pmatrix}, \tag{2.3}$$

$$H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \tag{2.4}$$

where the $S^\dagger$ gates are implemented as virtual-Z gates. And, in this section, I will follow the regular CNOT suffix convention of CNOT_{01} having a control qubit on the

top, and the target qubit on the bottom, and a CNOT_{10} thus interchanging these roles of the control and target qubits.

Section 2.1 explained that U_{SWAP} is not available. Similarly, U_{CNOT} is also not available, but will be used theoretically in this section. Moreover, the single-qubit Hadamard gate is also not natively available in our platform, and is supplemented at gate compilation as $H(\Psi) = [R_x(\frac{\pi}{2}) \circ R_z(\frac{\pi}{2}) \circ R_x(\frac{\pi}{2})](\Psi)$:

$$\begin{aligned}
 \Psi = |0\rangle &= \begin{pmatrix} 1 \\ 0 \end{pmatrix} \\
 &\xrightarrow{R_x(\pi/2)} \frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -i \end{pmatrix} \\
 &\xrightarrow{R_z(\pi/2)} \frac{1}{\sqrt{2}} \begin{pmatrix} e^{-i\pi/4} \\ -ie^{i\pi/4} \end{pmatrix} \\
 &\xrightarrow{R_x(\pi/2)} \frac{-i}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix} = -i|+\rangle,
 \end{aligned} \tag{2.5}$$

$$\begin{aligned}
 \Psi = |1\rangle &= \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\
 &\xrightarrow{R_x(\pi/2)} \frac{1}{\sqrt{2}} \begin{pmatrix} -i \\ 1 \end{pmatrix} \\
 &\xrightarrow{R_z(\pi/2)} \frac{1}{\sqrt{2}} \begin{pmatrix} -ie^{-i\pi/4} \\ e^{i\pi/4} \end{pmatrix} \\
 &\xrightarrow{R_x(\pi/2)} \frac{-i}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix} = -i|-\rangle,
 \end{aligned} \tag{2.6}$$

thus this decompiled sequence is equal to H to a global phase $-i$. Just like with the $S^\dagger$ gate, the R_z rotation is implemented as a virtual-Z gate, meaning practically that the final H operation produces two microwave pulses sent into the target qubit's XY line.

U_{SWAP} is compiled using the quantum circuits (2.7) to (2.10). The basis of this compilation is that any sequence of three consecutive conditional inversions, even in classical logic, equates to swapping bits in the input register. A quantum analogy

is seen with three CNOTs equating to a SWAP [8],

$$\begin{array}{c} \text{---} \times \text{---} \\ | \\ \text{---} \times \text{---} \end{array} = \begin{array}{c} \text{---} \bullet \oplus \text{---} \\ | \quad | \quad | \\ \oplus \quad \bullet \quad \oplus \end{array} \quad (2.7)$$

The two initial CNOT gates of (2.7) equals an iSWAP-like operation [65], which explains the three-to-two gate improvement from the compilation strategy overall. $\text{CNOT}_{10}(\text{CNOT}_{01}(\Psi))$ is compiled into its iSWAP equivalent,

$$\begin{array}{c} \text{---} \bullet \oplus \text{---} \\ | \quad | \\ \oplus \quad \bullet \end{array} = \begin{array}{c} \text{---} [H] [S^\dagger] \times \text{---} \\ | \\ \text{---} [S^\dagger] \times [H] \text{---} \end{array} \quad (2.8)$$

The last, remaining CNOT of (2.7) may be expressed using a CZ gate [90],

$$\begin{array}{c} \text{---} \bullet \text{---} \\ | \\ \text{---} \oplus \text{---} \end{array} = \begin{array}{c} \text{---} \bullet \text{---} \\ | \\ \text{---} [H] \bullet [H] \text{---} \end{array} \quad (2.9)$$

The last Hadamard gate of (2.8) cancels with the first Hadamard of (2.9), since $H(H(\Psi)) = I(\Psi)$, resulting in the equivalent quantum circuit for the SWAP sequence,

$$\begin{array}{c} \text{---} \times \text{---} \\ | \\ \text{---} \times \text{---} \end{array} = \begin{array}{c} \text{---} [H] [S^\dagger] \times \bullet \text{---} \\ | \quad | \\ \text{---} [S^\dagger] \times \bullet [H] \text{---} \end{array} \quad (2.10)$$

Note that the initial and last Hadamard gates are also unnecessary. An intuitive way to recognise their redundancy, would be to consider that the central portion of the circuit (2.10) already realises a SWAP operation, $(S^\dagger)^{\otimes 2} U_{\text{iSWAP-CZ}} = U_{\text{SWAP}}$. Any input state on the top qubit that sees the first Hadamard gate, is then transferred to the bottom qubit, where the final Hadamard gate simply undoes the first one. Nevertheless, the experimental implementation in Paper A used the sequence

as-derived in (2.10), further decompling the Hadamard gates as shown in (2.5) and (2.6).

An intuitive approach explaining this quantum circuit is that the iSWAP gate exchanges state information, albeit with unwanted added phases for some input states, which are conditionally corrected by the CZ gate. However, the order of the CZ and iSWAP gates is irrelevant as these gates commute,

$$\begin{aligned}
 U_{\text{CZ}}U_{\text{iSWAP}} &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \\
 &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \\
 &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \\
 &= U_{\text{iSWAP}}U_{\text{CZ}}.
 \end{aligned} \tag{2.11}$$

2.4 Kraus-Cirac decomposition of two-qubit gates

A theoretical validation of the SWAP sequence from section 2.3 can be done using Kraus-Cirac decomposition [65], [91]. Any two-qubit unitary $U \in \text{SU}(4)$ can be written as

$$U = (K_1 \otimes K_2) \text{Can}(t_x, t_y, t_z) (K_3 \otimes K_4), \quad K_i \in \text{SU}(2) \tag{2.12}$$

where

$$\text{Can}(t_x, t_y, t_z) = \exp \left[-i \frac{\pi}{2} (t_x X \otimes X + t_y Y \otimes Y + t_z Z \otimes Z) \right] \tag{2.13}$$

is known as a canonical gate. This method of two-qubit gate taxonomy reduces these gates to three parameters, namely the entangling coordinates t_x, t_y, t_z . The decomposition method assumes that global phases are factored out [65]. In equation (2.13), $\pi/2$ is a factored-out scalar to t_x, t_y, t_z , which is a matter of convention. As is seen in equation (2.13), these coordinates specify how much XX , YY , and ZZ interaction a two-qubit gate makes. The matrix representations of these gates

are,

$$XX = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}^{\otimes 2} = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \quad (2.14)$$

$$YY = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}^{\otimes 2} = \begin{pmatrix} 0 & 0 & 0 & -1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ -1 & 0 & 0 & 0 \end{pmatrix} \quad (2.15)$$

$$ZZ = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}^{\otimes 2} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 \\ 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (2.16)$$

Single-qubit gates cannot change how two-qubit gates are equivalent in terms of what entanglement they introduce, because single-qubit gates themselves do not change the entanglement between qubits: if U in (2.12) equals some combination of single-qubit gates applied onto the two qubits, then the corresponding canonical gate is

$$\text{Can}(0, 0, 0) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (2.17)$$

That is, any selection of single qubit gates making up U can be decomposed into only $K_{1..4}$ in equation (2.12).

The Weyl chamber taxonomy

Imagine a map of Earth. It is a two-dimensional sheet of paper with coordinates that corresponds to some spatial position on a three-dimensional sphere. There is a map where the axes instead correspond to algebra of two-qubit gates, and Earth is instead a Weyl chamber, illustrated in figure 2.2. The Weyl chamber itself is a geometric concept related to Lie algebra, not in scope for this thesis; for our purposes, the chamber forms a tool where we may geometrically draw two-qubit gate operations on a map, to see what entangling content the gates introduce between two qubits. The chamber allows us to quickly identify what selection of two-qubit gates would have to be supported in order to compile a maximally-short quantum circuit.




Figure 2.2: Weyl chamber of equivalent two-qubit operations, with geometric coordinates being the entanglement coordinates t_x, t_y, t_z . The locations of Identity, CNOT, iSWAP and SWAP are drawn. A trirectangular tetrahedron is drawn; by symmetry, any entanglement coordinates outside it correspond to a Can gate found inside the tetrahedron. Only the locally inequivalent gate classes relevant to this chapter are drawn.

For example, the CNOT gate to CZ gate conversion in (2.10) required no additional two-qubit gates, because the entanglement coordinates for both CNOT and CZ gates are $t_x, t_y, t_z = 0.5, 0, 0$, that is, they are both CNOT-like gates. The iSWAP gate in circuit (2.8) could similarly not be expressed as a single CNOT gate, since a single CNOT gate does not introduce both an XX and a YY component in equation (2.13), but an iSWAP gate does. This concept of local equivalency classes is the topic for the next section.

2.5 Local equivalency classes

Previously, we stated that two-qubit gates are locally equivalent if one can be converted into another using only single-qubit gates. This statement corresponds to using the same canonical gate for the two equivalent gates, with only the selection of $K_{1..4}$ being different [65], [91]. This means that all gates at the same entanglement coordinates in the Weyl chamber are interchangeable by the compiler without adding two-qubit gates.


As an example, the iSWAP gate belongs to the iSWAP class, defined by $\text{Can}(\frac{1}{2}, \frac{1}{2}, 0)$. We see this by selecting $K_{A1..A4} \in \text{SU}(2)$ to equal the single-qubit rotations $R_X(\pi)$,

$R_Z(\pi)$, $R_Z(\pi)$ and $R_X(\pi)$, respectively, resulting in

$$K_{A1} = \begin{pmatrix} 0 & -i \\ -i & 0 \end{pmatrix}, \quad K_{A3} = \begin{pmatrix} -i & 0 \\ 0 & i \end{pmatrix},$$

$$K_{A2} = \begin{pmatrix} -i & 0 \\ 0 & i \end{pmatrix}, \quad K_{A4} = \begin{pmatrix} 0 & -i \\ -i & 0 \end{pmatrix},$$

where we see that $U_{\text{iSWAP}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$ is locally equivalent to the iSWAP class,



$$= U_{\text{iSWAP}} \quad (2.18)$$

where the defining canonical gate of this class yields $\text{Can}(\frac{1}{2}, \frac{1}{2}, 0) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 0 & i \\ 0 & 1 & -i & 0 \\ 0 & -i & 1 & 0 \\ i & 0 & 0 & 1 \end{pmatrix}$.

Consider now a CZ gate executing after an iSWAP gate,

$$U_{\text{iSWAP-CZ}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \quad (2.19)$$

By selecting $K_{B1..B4} \in \text{SU}(2)$ to equal the single-qubit rotations $R_Z(\pi)$, $R_Z(\pi)$, $R_Z(3\pi/2)$ and $R_Z(3\pi/2)$, respectively, we get

$$K_{B1} = \begin{pmatrix} -i & 0 \\ 0 & i \end{pmatrix}, \quad K_{B3} = \begin{pmatrix} \frac{-1-i}{\sqrt{2}} & 0 \\ 0 & \frac{-1+i}{\sqrt{2}} \end{pmatrix},$$

$$K_{B2} = \begin{pmatrix} -i & 0 \\ 0 & i \end{pmatrix}, \quad K_{B4} = \begin{pmatrix} \frac{-1-i}{\sqrt{2}} & 0 \\ 0 & \frac{-1+i}{\sqrt{2}} \end{pmatrix},$$

and see that (2.19) is locally equivalent to the SWAP class,

$$\begin{array}{c}
 \text{---} \boxed{K_{B1}} \text{---} \boxed{\text{Can} \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2} \right)} \text{---} \boxed{K_{B3}} \text{---} \\
 \text{---} \boxed{K_{B2}} \text{---} \boxed{\phantom{\text{Can} \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2} \right)}} \text{---} \boxed{K_{B4}} \text{---}
 \end{array} = U_{\text{iSWAP-CZ}} \quad (2.20)$$

where this class is defined by $\text{Can} \left(\frac{1}{2}, \frac{1}{2}, \frac{1}{2} \right) = \begin{pmatrix} \frac{1-i}{\sqrt{2}} & 0 & 0 & 0 \\ 0 & 0 & \frac{1-i}{\sqrt{2}} & 0 \\ 0 & \frac{1-i}{\sqrt{2}} & 0 & 0 \\ 0 & 0 & 0 & \frac{1-i}{\sqrt{2}} \end{pmatrix}$, with the global phase of the right-hand side of equation (2.12) being $\frac{1+i}{\sqrt{2}}$.

Equation (2.20) thus tells us that a CZ gate executing after an iSWAP gate is a SWAP-class operation. Indeed, prepending one $S^\dagger$ gate onto each qubit results in

$$(S^\dagger)^{\otimes 2} = \begin{pmatrix} 1 & 0 \\ 0 & -i \end{pmatrix}^{\otimes 2} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -i & 0 & 0 \\ 0 & 0 & -i & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix} \quad (2.21)$$

$$(S^\dagger)^{\otimes 2} U_{\text{iSWAP-CZ}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} = U_{\text{SWAP}} \quad (2.22)$$

Recall the take-home message: using the two two-qubit gates in circuit (2.19) enabled a SWAP-like interaction in circuit (2.20) instead of requiring three two-qubit gates.

2.6 Experimental verification via conditional- and cross-Ramsey interferometry

The two-qubit interactions were experimentally studied in terms of their conditional phase accumulation and population transfer using conditional- and cross-Ramsey interferometry. The experimental results for the CZ, iSWAP, and SWAP operations are shown in figure 2.3.

In the quantum circuits at the top of this figure, the upper rail represents the control qubit q_c and the lower rail the target qubit q_t . Throughout this section,




Figure 2.3: Ramsey sequences demonstrating the populations and phases resulting from the two-qubit CZ (a, d), iSWAP (b, e), and SWAP (c, f) operations. In the upper row of the measured data, the two-qubit gates were performed on the states $|0\bar{i}\rangle$ and $|1\bar{i}\rangle$, whereas in the lower row, the roles of qubits 1 and 2 were swapped. The data points are post-processed to mitigate state preparation and measurement (SPAM) errors. The overlaid sinusoidal curves show the ideal outcome, and the dotted curves take T_1 relaxation into account. Blue-purple data points relate to conditional Ramsey experiments, while the red-green points relate to cross-Ramsey experiments. The mean squared error between the post-processed data and the ideal curves ranges between 0.001 and 0.020 at an average of 0.007, where the number of samples is 32 for the CZ data and 16 for the iSWAP and SWAP data.

all two-qubit states are written as $|q_c q_t\rangle$. Figures 2.3(a-c) and (d-f) differ in which physical qubit serves as q_c and q_t , that is, both qubits are verified in both roles.

For all experiments here, the target qubit q_t is initialized in $|\bar{i}\rangle$ by a $\sqrt{X}$ pulse. After executing the gate under test, a virtual-Z phase φ is applied, followed by a second $\sqrt{X}$ pulse for measuring in the Z basis. Sweeping φ produces Ramsey oscillations, from which the phase accumulated during the gate operation can be determined. Similarly, sweeping φ verifies that the output state takes the form $\cos \frac{\pi}{2} |0\rangle + e^{i\varphi} \sin \frac{\pi}{2} |1\rangle$, rather than its counterpart with phase $\pi - \varphi$.

The CZ gate is verified using conditional Ramsey interferometry, which confirms

that a π phase is applied to q_t only when $q_c = |1\rangle$, while no phase is acquired for $q_c = |0\rangle$. The iSWAP gate is verified with cross-Ramsey interferometry, where the Ramsey fringes oscillate in accordance with the expected outcome: from $|+0\rangle$ and $|-1\rangle$ at $\varphi = 0$, to $|0\rangle$ and $|1\rangle$ at $\varphi = \pi/2$, depending on the initial state of q_c . The SWAP gate is also verified with cross-Ramsey interferometry, where the operation transfers the state $|\bar{i}\rangle$ to q_c , producing the expected π -periodic Ramsey oscillations as the virtual-Z phase is swept.

A simplified intuitive way to distinguish the two interferometry experiments is that conditional Ramsey interferometry probes whether the target qubit acquires a phase conditioned on the state of the control qubit, whereas cross-Ramsey interferometry prepares the state on the control qubit and probes the resulting phase on the target qubit after the gate has successfully transferred the population. Importantly, all operations studied here are phase-sensitive, as further demonstrated by the measurements presented in Chapter 3.

These interferometry experiments can be understood in a simplified single-qubit picture: the conditional Ramsey experiment will modify the target if the conditional qubit is set; the cross-Ramsey experiment will prepare the intended final state on the wrong qubit initially to then transfer it onto the other qubit. Importantly, to be successful, these interferometry experiments require the transfer of phase information, which will be further emphasised in Chapter 3.

The approach to verifying the function of our CZ, iSWAP, and SWAP operations using Ramsey interferometry can be improved by applying a tomographically complete input set of states [92]. Such a verification uses the approach in Paper A as a necessary condition for certifying the CZ, iSWAP, and SWAP operations, while also determining these gate operations to a unique solution of the unitary gate matrices in equation (2.2).

Conventional Ramsey interferometry measures the accrued phase over time t of a single-qubit superposition state, $|0\rangle + e^{i\varphi(t)}|1\rangle$. A drawback of idling is that other time-dependent effects appear, such as energy relaxation. A good substitute is to replace the idling with a virtual-Z gate V_z , to induce a phase by shifting the phase of all subsequent control pulses. This approach reproduces the phase that would have been acquired during free evolution without exposing the qubit to decoherence during an idle period. The main notable effect being that the datapoints for all the respective subfigures in figure 2.3, are acquired at the same t counted from the start of each respective quantum circuit. t would not have been consistent between datapoints if using regular Ramsey interferometry, since a delay would have been required to accumulate phase after the two-qubit gate.

We will return to energy relaxation error later, in section 2.7.

2.6.1 Readout state assignment setup

The population axes of figure 2.3 are acquired using readout state discretisation. Each analogue readout is digitised and integrated, resulting in one complex coordinate in the magnitude-phase plane, and then assigned to the population centre that corresponds to the discrete cases where the corresponding qubit, for some probed resonator of interest, is in $|0\rangle$, $|1\rangle$ or $|2\rangle$. This approach is commonly referred to as state-assigned three-state readout, or *discretised* three-state readout.

Figure 2.4 illustrates the readout state-assignment procedure. For each prepared state Ψ_n into number state $|n\rangle$, the demodulated readout pulse is integrated to obtain a single complex-valued measurement outcome, represented as a point in the complex plane. Repeating the experiment many times yields three population centres corresponding to the prepared states $|0\rangle$, $|1\rangle$, and $|2\rangle$, seen in (a)-(c), respectively. State assignment is then performed by partitioning the complex plane into decision regions centred on the measured state populations, and classifying each single-shot measurement to its closest population centre as shown in (d). Once the complex coordinates are known, any acquired readout response during an experiment can be assigned, forming a statistical average from single-shot measurements, revealing the probability of reading out some n -qubit state. Here, the single-shot aspect is important, because averaging multiple single-shot measurements yields a complex coordinate that corresponds to skewing the population centres.

The resulting single-qubit classification performance is quantified by the confusion matrices shown in (e), whose diagonal elements correspond to successful state assignment quantified by the mean assignment fidelity,

$$F_{\text{Assign}\{|0\rangle,|1\rangle,|2\rangle\}} = \frac{|\langle 0|\Psi_0\rangle|^2 + |\langle 1|\Psi_1\rangle|^2 + |\langle 2|\Psi_2\rangle|^2}{3} , \quad (2.23)$$

and off-diagonal elements represent misclassification. This condition is usually framed as selecting the settings that separate the three population centres the most. The settings used for the readout signal¹ itself is then optimised in terms of frequency and amplitude, with parameters selected to maximise the mean assignment fidelity for each resonator.

The confusion matrices acquired for Paper A yielded assignment fidelities of 87.93% $\pm$ 4.95% and 88.73% $\pm$ 5.28%, which is comparable to the two-qubit planar reference design [28] on which that device was based, and 5-6% less than the uncorrected readout fidelity of our more recent devices [93]. We observed 7-15% classification error into $|0\rangle$: considering $T_1 \approx 77 \mu\text{s}$, the 2.3 μs readout pulse was expected to yield $\approx$ 3% energy relaxation error into $|0\rangle$ during the readout; the thermal

¹Pilot tone.




Figure 2.4: Readout state assignment. The top row shows single-shot population diagrams from resonator 1, while the bottom row shows resonator 2. Columns (a)-(c) show the demodulated single-shot readout responses obtained after preparing states $|00\rangle$, $|11\rangle$, and $|22\rangle$, respectively. Each point corresponds to one single experimental shot with no averaging. The horizontal and vertical axes denote the real and imaginary components of the integrated demodulated readout signal. The integrated signal is obtained by digitizing the readout pulse and integrating the acquired time trace to produce a single complex-valued measurement outcome. Column (d) shows all state-assignment regions, found using a nearest-population-centre classifier. Column (e) shows the resulting single-qubit confusion matrices. The data in columns (a)-(d) were acquired using 8,000 single-shot measurements per prepared state.

population error was estimated to $|\langle 1|\Psi_0\rangle|^2 + |\langle 2|\Psi_0\rangle|^2 \approx 2\%$; the remaining 2-10% error in $|0\rangle$ is believed to stem from gate error, overlap error, and readout-induced mixing as DRAG [94] was used in the experiments.

The state-assignment procedure described in this section is equally valid for figures 3.3(c, d) to be shown later on page 43, where states $|11\rangle$ and $|20\rangle$ are analysed. As would be expected, a readout result would be assigned into $|20\rangle$ if qubits 1 and 2 were classified into $|2\rangle$ and $|0\rangle$, respectively.

Following state-assignment, there are deviations between the experimental data and the ideal curves visible in figure 2.3. These discrepancies arise from a combination of various experimental error sources and imperfections present in the device, which will be discussed in section 2.7. Furthermore, the data in figure 2.3 has been post-processed to mitigate SPAM error, which is the topic of the next section.

2.6.2 Error mitigation

The key idea is to gather a large sample size of single-shot data, and under the assumption that the data collected during an experiment will on-average show a similar level of error, then correct for this error using the aforementioned statistics collection. This technique is widely used in the field of experimental quantum information.

Method

Let us consider any one of the twelve traces of data from figure 2.3. For each step along the x-axis, we measure some data point $\vec{y}$ that contains the observed probability of all nine considered two-qubit states, $(3 \text{ states})^{(2 \text{ qubits})}$:

$$\vec{y} = [|\alpha_{00}|^2, |\alpha_{01}|^2, |\alpha_{02}|^2, |\alpha_{10}|^2, |\alpha_{11}|^2, \dots, |\alpha_{22}|^2]^T \quad (2.24)$$

where α_i is the complex amplitude of the two-qubit state $|i\rangle \in \{0, 1, 2\}^2$ of the system's wave function Ψ at measurement.

Before the conditional and cross-Ramsey experiments, we collected the on-average distribution of error by intentionally preparing Ψ in the nine two-qubit states considered, followed by measuring the outcome from thousands of single-shot iterations. These collected statistics are used to construct a right stochastic [95] matrix

$$T = \begin{bmatrix} |\langle 00|\Psi_{00}\rangle|^2 & |\langle 01|\Psi_{00}\rangle|^2 & \cdots & |\langle 22|\Psi_{00}\rangle|^2 \\ |\langle 00|\Psi_{01}\rangle|^2 & |\langle 01|\Psi_{01}\rangle|^2 & & \\ \vdots & & \ddots & \\ |\langle 00|\Psi_{22}\rangle|^2 & & & |\langle 22|\Psi_{22}\rangle|^2 \end{bmatrix} \quad (2.25)$$

where each row shows the most likely distribution of measurement outcomes from observing the prepared two-qubit state Ψ_i , given that some two-qubit state $|i\rangle \in \{0, 1, 2\}^2$ was prepared for that row. Each element of the matrix shows the conditional probability

$$P(|j\rangle \text{ measured} \mid |i\rangle \text{ prepared}), \quad j, i \in \{0, 1, 2\}^2.$$

Figure 2.5 shows the experimentally acquired T . In Paper A, the CZ and iSWAP data reproduced in figure 2.3(a, b, d, e) were collected in a different experiment session than the SWAP data (c, f); two T were thus collected, figures 2.5(a) and (b), respectively.

Experimentally, T must be acquired by preparing and measuring all possible qubit state combinations. Assembling T from the single-qubit confusion matrices yields




Figure 2.5: Two separate multi-qubit-state confusion matrices acquired for post-processing the conditional and cross-Ramsey experiments in figure 2.3, corresponding to two different measurement sessions. These sessions experimented on either the CZ gate and iSWAP gate [row (f), left matrix], or the SWAP gate [row (f), right matrix]. These two-qubit confusion matrices were acquired using 25,000 (CZ, iSWAP) and 120,000 (SWAP) shots.

different results since effects from two-qubit interaction are effectively discarded; this concept will be returned to later in figure 2.6.

We now assume that our acquired data point $\vec{y}$ is the result of manipulating some “true value” data point $\vec{x}$,

$$\vec{x} = [|\beta_{00}|^2, |\beta_{01}|^2, |\beta_{02}|^2, |\beta_{10}|^2, |\beta_{11}|^2, \dots, |\beta_{22}|^2]^T \quad (2.26)$$

which is assumed to have been the true outcome of the experiment before $\vec{x}$ was manipulated by T , resulting in the measured output $\vec{y}$.

$$\vec{y} = T\vec{x} \quad (2.27)$$

Remember that for every data point in any given trace from figure 2.3, we acquired a vector $\vec{y} = T\vec{x}$ during the experiment itself. To extract $\vec{x}$, which is the number that the datapoints in figure 2.3 show, there are multiple tools available [96], [97] for minimizing the difference

$$\min \left(\frac{1}{2} \cdot |\vec{y} - T\vec{x}|^2 \right). \quad (2.28)$$

Specifically, since $\vec{y}$ and T were found empirically and are thus known, we find the vector $\vec{x}$ that minimizes equation (2.28).

Finally, the calculated $|\beta_{01}|^2$ and $|\beta_{11}|^2$ from equation (2.26) are put into figure 2.3 as datapoints, for the CZ conditional Ramsey traces. Similarly for the iSWAP and SWAP cross-Ramsey traces, $|\beta_{10}|^2$ and $|\beta_{11}|^2$ correspond to the datapoints shown.

Figure 2.3 shows several clusters of data points that deviate from the ideal traces. The next section will therefore discuss various sources of error that contribute to the non-ideal implementation observed our experiment.

2.7 Sources of experimental error

The deviating cluster of data shown in figure 2.3 were acquired close in time: the qubit is intermittently “struck out” of its predicted behaviour for durations between 2 min and 23 min, to later snap back at seemingly random intervals. At the time, we attributed this erroneous behaviour to two-level-system defects being strongly coupled to the qubits, which is to this day the dominant speculation.

To gauge the mismatch between the post-processed data and the ideal expectation, the mean squared error (MSE) was quantified in the appendix of Paper A. For ease of reference, I present this data in table 2.1. The on-average error is 0.007 where 0 corresponds to no error between datapoints and the ideal curve; where, the CZ datasets contained 32 samples per trace and the iSWAP and SWAP datasets contained 16 samples per trace.

A notable outlier is seen with the $|\beta_{11}|^2$ reconstruction for the CZ trace where $|q_c q_t\rangle = |Q_2 Q_1\rangle$ at $\text{MSE} = 0.02033$, which is twice as large as the second largest mean squared error. I attribute this outlier to the two notably deviating clusters of data in figure 2.3(d), $P(|11\rangle)$.

In principle, state discretisation cannot yield datapoints that fall out of the interval $\{0, 1\}$, because no readout can result in some state more than 100% of the time, or less than 0% of the time. The reason for datapoints being outside of this interval in figure 2.3 is due to the SPAM-mitigating post-processing technique used. Recall that in the previous section, I explained that this technique acts as a guess to what the true $\vec{x}$ would have been, under the assumption that experimental errors will on average show the same behaviour. The aforementioned TLS defect hypothesis is a typical example of a case where this assumption may be broken, and the deviating clusters of data could stray from the expected outcome.

The remainder of this section will refer to multiple device parameters. In ta-

Table 2.1: Conditional Ramsey and cross-Ramsey data of figure 2.3. P , *swing* is the fitted population swing where the ideal swing is from 0.0 to 1.0, that is, the expected ideal number is 1.0; P , $\Delta offset$ is the difference $m_{ideal} - m$ between the expected ideal offset $m_{ideal} = 0.5$ and the fitted offset m , where $m_{ideal} = 0.5$ denotes that the expected population swing is oscillating about 0.5, and P , $\Delta offset = 0.0$ means no difference between the expected ideal and the experimental data; $\Delta phase$ is similarly the deviation in the oscillating population’s phase, where $\Delta phase = 0.0$ means no difference between the expected phase and the experimental data; and MSE denotes the mean squared error: $MSE = \frac{\sum (y-p)^2}{N}$, where y is the datapoint, p is its predicted ideal value, and N is the number of datapoints. An MSE of 0 means no error between the datapoints and the ideal prediction.

Operation q_c, q_t	Output state	P, swing [-]	P, $\Delta offset$ [-]	$\Delta phase$ [mrad]	MSE [-]
CZ	$ 01\rangle$	0.987(1)	0.027(11)	37(8)	0.00117
q_1, q_2	$ 11\rangle$	0.967(3)	0.077(15)	57(15)	0.00632
CZ	$ 10\rangle$	1.008(13)	0.006(8)	71(12)	0.00114
q_2, q_1	$ 11\rangle$	0.868(41)	0.150(17)	16(52)	0.02033
iSWAP	$ 10\rangle$	1.079(18)	0.023(6)	33(16)	0.00548
q_1, q_2	$ 11\rangle$	1.064(29)	0.060(10)	85(26)	0.00915
iSWAP	$ 01\rangle$	1.041(5)	0.064(2)	85(4)	0.00401
q_2, q_1	$ 11\rangle$	1.057(9)	0.038(3)	75(8)	0.00570
SWAP	$ 10\rangle$	0.849(33)	0.049(16)	96(42)	0.00997
q_1, q_2	$ 11\rangle$	0.859(44)	0.023(27)	68(47)	0.00584
SWAP	$ 01\rangle$	0.954(39)	0.055(20)	223(44)	0.00953
q_2, q_1	$ 11\rangle$	0.993(23)	0.050(12)	103(26)	0.00327

ble 2.2, I reproduce the parameters I acquired experimentally in Paper A for ease of reference.

2.7.1 Energy relaxation and low-frequency noise

The device was analysed for regular decoherence parameters, T_1 , T_2^* and T_2^e , shown in table 2.2 with full distributions shown in the appendix of Paper A.

There is little reason to believe T_1 errors to be concerning: the dashed lines in figure 2.3 represent the T_1 compensated relaxation, and are significantly overlapping with the truly ideal expected outcomes drawn as solid traces. By comparison, the longest experiment in figure 2.3(c, f) executes in $\sim 4.2 \mu s$, that is, 5.3% of T_1 .

There is reason to suspect low-frequency noise being present, due to T_2^e being triple that of T_2^* [54].

Table 2.2: Physical parameters of the two-qubit system. The DRAG parameter α is defined using notation from [98, (3.12)]. The error bars reported for the frequency offsets Δ were calculated using simplified error propagation [99],

$$\text{Err}(\Delta f_{\text{CZ}}) = \sqrt{(-1)^2 \cdot \text{Err}(f_{12,Q1})^2 + (+1)^2 \cdot \text{Err}(f_{01,Q2})^2 + (-1)^2 \cdot \text{Err}(f_{\text{CZ}})^2},$$

$$\text{Err}(\Delta f_{\text{iSWAP}}) = \sqrt{(+1)^2 \cdot \text{Err}(f_{01,Q1})^2 + (+1)^2 \cdot \text{Err}(f_{01,Q2})^2 + (-1)^2 \cdot \text{Err}(f_{\text{iSWAP}})^2}.$$

The error reported for the three-state mean assignment fidelity, $F_{\text{Assign}\{|0\rangle,|1\rangle,|2\rangle\}}$, is the standard error from the mean (S.E.M. = $\sigma/\sqrt{3}$), and the variance is $\sigma^2 = (\sum_{i=1}^3 (F_{|i\rangle} - F_{\text{Assign}\{|0\rangle,|1\rangle,|2\rangle\}})^2) / (3 - 1)$.

Parameter	Q_1, R_1	Q_2, R_2	C_1
$f_{\text{RO}, 0\rangle}$	6.48084 GHz	6.25968 GHz	n/a
$f_{\text{RO}, 1\rangle}$	6.48068 GHz	6.25946 GHz	n/a
$f_{\text{RO}, 2\rangle}$	6.48056 GHz	6.25920 GHz	n/a
$f_{\text{RO,optimal}}$	6.48048 GHz	6.25960 GHz	n/a
$\chi_{ 0\rangle \rightarrow 1\rangle}$	80 kHz	110 kHz	
$\chi_{ 1\rangle \rightarrow 2\rangle}$	60 kHz	130 kHz	
τ_{RO}	2.3 μs		n/a
$F_{\text{Assign}\{ 0\rangle, 1\rangle, 2\rangle\}}$	0.8793	0.8873	n/a
	± 0.0494	± 0.0528	
$f_{ 0\rangle \rightarrow 1\rangle}$	3.86011 GHz	3.39741 GHz	4.863 GHz
	± 156 kHz	± 156 kHz	
$f_{ 1\rangle \rightarrow 2\rangle}$	3.60404 GHz	3.18992 GHz	
	± 190 kHz	± 190 kHz	
$\eta / 2\pi$	-256.07 MHz	-207.49 MHz	
	± 346 kHz	± 346 kHz	
$\tau_{1\text{QB}}$	20 ns		
α_{Motzoi}	$-770.5 \cdot 10^{-3}$	$-221.4 \cdot 10^{-3}$	
	$\pm 13 \cdot 10^{-3}$	$\pm 13 \cdot 10^{-3}$	
T_1	77 $\mu\text{s} \pm 27$ μs	79 $\mu\text{s} \pm 25$ μs	
T_2^*	37 $\mu\text{s} \pm 10$ μs	33 $\mu\text{s} \pm 10$ μs	
T_2^e	93 $\mu\text{s} \pm 27$ μs	105 $\mu\text{s} \pm 28$ μs	
$f_{\text{LO},XY}$	3.600000000 GHz		n/a
$f_{\text{LO},RO}$	6.600000000 GHz		n/a
Φ_{bias}	n/a		-0.336 Φ_0
			$\pm 0.003 \Phi_0$
$f_{ 0\rangle \rightarrow 1\rangle}, \Phi=0$	3.8610 GHz	3.3977 GHz	6.9373 GHz
$g_{q,c} / 2\pi$	37 MHz	35 MHz	
CZ			
τ_{CZ}	890 ns		
f_{CZ}	207.24 MHz ± 10 kHz		
Δf_{CZ}	-0.61 MHz ± 246 kHz		
$\varphi_{\text{comp.}}$	+0.61 rad	+1.03 rad	0 rad
iSWAP			
τ_{iSWAP}	640 ns		
f_{iSWAP}	461.51 MHz ± 82 kHz		
Δf_{iSWAP}	+1.19 MHz ± 235 kHz		
$\varphi_{\text{comp.}}$	+2.24 rad	+3.36 rad	+4.86 rad
SWAP			
τ_{SWAP}	1812 ns		
$\varphi_{\text{comp.}}$	-1.07 rad	+1.25 rad	= φ_{iSWAP}

2.7.2 Readout error

In terms of readout error, Paper A establishes that readout pertaining to any state where a qubit can decay from $|2\rangle \rightarrow |1\rangle$ dominates the misclassification error. In figure 2.5 we estimate an energy relaxation error of about 3% during the readout pulse, and a thermal population error of about 2%. From the remainder of the classification error in figure 2.5, we expect between 2-10% error in $|0\rangle$ to stem from gate error, overlap error, and readout-induced mixing.

Figure 2.6 further highlight effects pertaining to the two-qubit readout, such as measurement-induced state decay and readout crosstalk. This figure compares two-qubit confusion matrix data to the single-qubit confusion matrices, showing that these two-qubit confusion matrix data are overall very similar, differing mainly in $|2\rangle \rightarrow |1\rangle$ readout errors.




Figure 2.6: Matrix showing $C_{\text{Measured}} - C_{\text{Constructed}}$, where C_{Measured} is the confusion matrix measured prior to the SWAP experiments from figure 2.5(b), and $C_{\text{Constructed}}$ is the two-qubit confusion matrix acquired from combining the measured single-qubit confusion matrices from figure 2.4(e). This matrix highlights effects that stem from two-qubit readout, such as measurement-induced state decay and readout crosstalk, as seen below the diagonal. From preparing $|12\rangle$, we see that C_{Measured} yielded 0.19 more population above the diagonal than $C_{\text{Constructed}}$, which we believe hints at measurement-induced transition errors. Apart from $|\langle 12 | \Psi_{22} \rangle|^2$, the remaining six combinations where a qubit could decay $|2\rangle \rightarrow |1\rangle$, all accounted for the largest observed population increases: $\{0.36, 0.35, 0.26, 0.25, 0.21, 0.17\}$.

2.8 Concluding remarks on experimental quantum informatics

This chapter has shown that access to two locally inequivalent native two-qubit gate classes reduces the gate overhead arising during compilation when using minimum gate sets. Enriching one's native gate set with additional locally inequivalent interactions, provides practical advantages for superconducting quantum processors, valid beyond the NISQ era of quantum computation.

The SWAP gate is not natively available in the superconducting circuit-QED architecture pertaining to this thesis. It therefore had to be compiled from other two-qubit gates. The simultaneous availability of CZ and iSWAP interactions enabled its implementation using a substitution of two gates, improving upon the conventional three-gate decomposition based on gate sets only supporting CZ-like or iSWAP-like gates.

The theoretical framework presented in this chapter provides the basis for understanding why this reduction is possible. The Kraus-Cirac decomposition, and to some extent its associated Weyl chamber representation, was used to show how the iSWAP and CZ gates can be combined to form a SWAP-like interaction. These tools see usage in quantum circuit compilation when realising maximally-short circuits.

The experimental results of Paper A support the theoretical picture. Conditional- and cross-Ramsey interferometry verified the core concept by demonstrating phase accumulation and population transfer closely matching the values predicted. An error analysis was presented for consideration as to why the measured outcomes at places deviate from the ideal expectations.

The next chapter takes one step up in the hardware stack. Its focus will be on the microwave control electronics used for generating and analysing the signals interacting with the QPU, and the successful attempts at simplifying this previous chapter's conventional interface, into one based on direct digital synthesis.

CHAPTER 3

Quantum processor control and readout using DDS

For superconducting quantum computers to date, the control and readout methodology generally relies on the generation and analysis of precisely tuned and phase-synchronised radio-frequency (RF) signals in the 0.1 GHz to 8 GHz band. A generator located at room temperature launches pulsed signals through coaxial cables into a cryogenic refrigerator, where the pulses interact with the quantum hardware. When the carrier frequency is higher than the bandwidth of the available arbitrary waveform generators (AWG), such that direct digital synthesis (DDS) of the desired waveform is not possible, the pulses are shaped by quadrature mixing of local oscillators (LO) and AWG signals at the intermediate frequency (IF) [100]–[102]. A comparison between DDS and such mixing is shown in figure 3.1.




Figure 3.1: Comparison between a conventional superheterodyne approach (a) and the DDS approach (b) to a quantum computer interface used in Papers B, C and D. Conventional qubit-control systems require microwave mixers, local oscillators, and some degree of shaping the signal such as filtering and amplification [54]. DDS allows for integrating such signal conditioning, the mixers, and LOs, by moving their operations in the signal path into the digital domain.

Scaling to larger systems requires increasingly complex parameter estimation protocols and a control infrastructure which includes electronics, interconnects, thermalisation, and pulse-generation software, whose scaling with processor size is emerging as a major challenge for large-scale quantum computing [9], [103]–[105]. A solution may be developed out of recent industrial advances in high-speed digital electronics, which have produced integrated platforms capable of generating and analysing phase-coherent RF signals across hundreds of channels with high bandwidth and low distortion.

One such development is the radio-frequency system-on-a-chip (RFSoc) platform by Xilinx, now owned by AMD, which integrates high-speed analogue-to-digital converters (ADC) and digital-to-analogue converters (DAC) together with programmable logic on a single device [106]. Early demonstrations showed quantum computer control and readout using first-generation RFSoc evaluation platforms in conjunction with external mixers to cover the bandwidth required by their applications [100], [101], [107]. Later developments with the RFSoc enabled mixer-less architectures where the integrated ADCs and DACs directly interface the quantum processor without external frequency conversion [108]. One good, more recent example is the digital microwave control and readout platform known as Presto, developed by my collaborators in the group of Prof. David Haviland at KTH and commercialized by their spin-out company Intermodulation Products. The technology is demonstrated in Papers B and C, in which we demonstrated a quantum computer interface with microwave waveforms composed by only DDS as opposed to the conventional, more complex approach of analogue heterodyning. Figure 3.1 shows that the DDS setup removes the need for several of the components present in the heterodyning setup, which is achieved by replacing these components with mathematical operations in the digital domain.

In Paper B, I contributed to the experimental realisation and evaluation of the Presto platform by setting up and performing initial measurements on superconducting qubits, verifying against conventional measurement routines, and performing data analysis. For Paper C, my contribution was focused on co-developing a methodology for determining the iSWAP gate parameters using the Presto and comparing it to conventional tune-up methods based on heterodyne mixing. In section 3.2, I will provide complementary data to Papers B and C, showing additional experiments to more fully tune up parametric CZ and iSWAP gates. My work with this platform has also yielded a publicly available instrument control stack for quantum processor operation [109], used mainly in research [110], [111] and by the quantum device manufacturer ConScience.

Beyond simplifying conventional heterodyne signal chains, DDS has more fundamental implications for superconducting quantum computing. Multi-qubit opera-

tion places stringent demands on phase stability, timing determinism, and synchronisation across many control channels. The following section discusses how DDS can change the control and readout paradigm for quantum computer interfaces and some practical reasoning behind its adoption.

3.1 Direct digital synthesis approach to quantum computer interfaces

DDS is attractive in multi-qubit control systems because it removes LO phase constraints. In principle, DDS allows for absolute phase control, where the momentary phase of the AWG signals can be changed at will. As is seen in figure 3.1, DDS achieves this control by replacing the analogue heterodyne conversion in (a) with digitally generated RF signals derived from a common, global reference clock. Here I discuss some practical advantages and trade-offs of DDS vs. conventional LO-based interfaces.

Experimentally, it has often been a challenge to obtain sufficient phase synchronisation between multiple RF signals driving different transitions in a coupled-qubit system, see for example [112] and Paper A. In the two-qubit example of Paper A, shown in figure 2.1(b) on page 13, I used two independent mixers to upconvert AWG-generated IF signals to drive single-qubit gates for two qubits at 3.860 GHz and 3.397 GHz, respectively. Each mixer modulated its IF signal using its own LO source, each set to 3.6 GHz. The IF tones therefore remained within the AWG’s 300 MHz bandwidth when combined with quadrature mixing. However, the two-qubit gates used in Paper A required coupler signals at 207 MHz and 462 MHz¹ which were generated at baseband by the AWG. Therefore, the coupler does not share the same reference frame as the single-qubit gates in such a setup. Certain two-qubit gates, such as the iSWAP gate, rely on phase coherence between single-qubit and two-qubit operations. For such gates, care must therefore be taken to ensure that the period of the LO is commensurate with the duration of the experiment, that is, the execution time of the quantum circuit. Doing so resets the starting phase to the same value between consecutive quantum circuits, thereby ensuring proper phases throughout all iterations of an experiment. Without this step, the phase information would be lost when averaging the readout signal over several runs to obtain the expectation value of the qubit state at the end of the experiment. Parenthetically, on the other hand, the two-qubit CZ gate is immune to such phase scrambling with respect to the single-qubit gates [21].

¹The 462 MHz signal was generated by a 750 MHz DAC; the 300 MHz bandwidth specification stems from the operating mode of the output amplifier. This signal was therefore severely attenuated due to out-of-band roll-off.

This difficulty is removed by the DDS method, as all frequencies are derived from a common reference clock. DDS also does not require mixer calibration, effectively eliminating the need for LO leakage compensation and I/Q skew adjustment from the tune-up process [102].

Phase noise on an LO contributes to dephasing that is indistinguishable from environmental phase noise, because the qubit reference frame itself becomes unstable [113]. In DDS, since all signals are derived from a common master clock, the reference frame stability requirement enforces a master clock phase noise stability on par with that of LOs in heterodyne signal paths. One practical consequence can be seen in the Presto, which has to use a custom-designed clock module internally to achieve this phase noise requirement.

As is seen in figures 3.1(a) and (b), the number of physical components is reduced when mixing and filtering become mathematical operations in DDS. In principle, this reduction improves power efficiency as fewer components require electrical power [114]. However, the power consumption is shifted to high-speed digital electronics, where dynamic power losses dominate due to the high switching rates required. This trade-off can be limiting in practice: for example, the Presto gateway used in Papers B, C, and D is a so-called power-limited synthesis design, which means that the integrated electronics cannot deliver more electric power to implement additional features.

Both DDS and conventional heterodyning share scaling challenges in maintaining deterministic timing and phase coherence across large systems, specifically clock distribution and synchronisation between units. Contemporary offers from the likes of Qblox [115] and Zurich Instruments [116] approach the synchronisation problem by engineering rack backplanes for clock distribution and synchronisation between interconnected modules. The current state-of-the-art offer in QPU interfaces is from Zurich Instruments in the form of their ZQCS. This interface promises 1092 synchronised channels per rack, with DDS signal generation up to 10.5 GHz [116]; such an interface has an intended QPU size range in the 1000-qubit category. Achieving such signal synthesis is done with 24 GSa/s ADCs at 1.6 GHz instantaneous bandwidth; comparing to a conventional heterodyning approach, such a setup introduces hardware design challenges in thermal management due to the dynamic power losses of state-of-the-art data converters, solvable for instance by using water-cooled enclosures.

Whether future quantum computer interfaces ultimately employ fully digital synthesis, conventional heterodyning methods, or other alternatives such as cryogenic electronics, the transition of microwave signal generation from analogue to digitally defined architectures remains an important direction to consider for scaling superconducting quantum computing.

3.2 Quantum experiments for DDS verification

Papers B and C represent early demonstrations of quantum computer control using the Presto instrument, performed on the device [98] shown in figure 1.2 on page 7.

While the iSWAP gate tune-up demonstrated in Paper B was done by determining its driving frequency and pulse duration, here in figure 3.2 I show, additionally, how I determined the local phase accumulation of each of the two qubits based on the iSWAP tune-up methodology in [21]. Figures 3.2(a) and (b) show the quantum circuits used for determining these parameters. By [21], the experiment (a) as done in Paper B executes the unitary operator

$$U_{\text{iSWAP}}(\theta, \varphi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta/2 & ie^{i\varphi} \sin \theta/2 & 0 \\ 0 & ie^{-i\varphi} \sin \theta/2 & \cos \theta/2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} . \quad (3.1)$$

The iSWAP operation shown in equation (2.2) on page 15 requires that the diagonal elements of the $\{|01\rangle, |10\rangle\}$ subspace must be zero, determining a θ by

$$\cos \theta/2 = 0 \Leftrightarrow \theta = \Omega_{\text{eff.}} \cdot \tau_{\text{iSWAP}} = \pi \quad (3.2)$$

where $\Omega_{\text{eff.}}$ is an effective exchange rate between $|01\rangle \leftrightarrow |10\rangle$ when parametrically modulating the coupler for a gate duration τ_{iSWAP} [21]. Figures 3.2(c) and (d) highlight the pulse carrier frequency and duration where the population exchange is maximised between qubits 1 and 2, for an initially assumed pulse amplitude. The unitary executed using these parameters is

$$U_{\text{iSWAP}}(\theta = \pi, \varphi) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & ie^{i\varphi} & 0 \\ 0 & ie^{-i\varphi} & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} , \quad (3.3)$$

which equals an iSWAP only if $\varphi = 0$. Establishing a correct φ imposes additional experimental issues explained further in Paper A, such as compensating for the coupler drive dispersively shifting the coupler frequency out of the two-qubit frame.

Importantly, this shift also changes the qubit frequencies, causing phase accumulation on the single-qubit frames which must be compensated for with phase offsets φ_{Q1} and φ_{Q2} . In figure 3.2(b), I replicate one method [21] to determine these phase offsets: playing the imperfect U_{iSWAP} followed by its conjugate transpose $U_{\text{iSWAP}}^\dagger$ causes $U_{\text{iSWAP}}^\dagger(U_{\text{iSWAP}}) = I^{\otimes 2}$, but still leaves the single-qubit frame phase-shifted due to the dispersive shifts caused by these two coupler pulses. Figures 3.2(e) and

(f) establish the local phase offset required, found by sweeping φ of the virtual-Z gate in (b).

In conventional heterodyne setups, this virtual-Z gate is tuned by phase shifting the IF carrier, as was done in Paper A and explained on page 24 in section 2.6. In my complementary DDS experiments here, I instead achieve this phase shift by offsetting the phase accumulator before it is multiplied onto the digital equivalent of the LO.




Figure 3.2: Tune-up of the iSWAP gate on a DDS platform: complementary local phase measurements. (a-b) Quantum circuits for the determination of (a) gate duration and frequency for some initial estimate of the gate amplitude, and (b) for local phase adjustments. (c-d) Coherent population exchange measured on qubits 1 (c) and qubit 2 (d), resulting from running the circuit in (a). Qubits 1 and 2 are the top and bottom qubits in the circuits, respectively. (e-f) Determination of local phase corrections φ_{Q1} and φ_{Q2} for qubits 1 and 2, respectively, where the qubits take turns acting as the top qubit in (b). These phase corrections indicate that coherent population exchange alone is insufficient to tune up the gate operation.

Moreover, figure 3.3 shows the tune-up of a CZ gate using the same DDS interface. Here, I report on a sequence of gate parameter measurements necessary for ensuring a parametric CZ gate that imparts a π phase shift on the target qubit q_t , conditional on the control qubit q_c . Figure 3.3(a) and (b) highlight a parameter space scan for identifying where a population exchange between $|11\rangle \leftrightarrow |20\rangle$ exists (a), followed by the conditional Ramsey experiment (b) identifying the precise carrier frequency for a CZ gate. The two-qubit state probabilities in (c-d) are acquired through nearest-population-centre readout assignment, which was detailed in section 2.6.1 on page 25.

Figure 3.3(c) and (d) show a readout of states $|11\rangle$ and $|20\rangle$. An initial CZ gate time τ_{CZ} is selected, and the coupler pulse amplitude is swept as a function of its carrier frequency. The sweep identifies a cut through parameter space where the population transfer $|11\rangle \rightarrow |20\rangle$ is maximised, corresponding to a unitary [21],

$$U_{CPhase} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{-i\varphi_{01}} & 0 & 0 \\ 0 & 0 & e^{-i\varphi_{10}} & 0 \\ 0 & 0 & 0 & e^{-i(\varphi_g + \varphi_{ZZ} + \varphi_{10} + \varphi_{01})} \end{pmatrix}, \quad (3.4)$$

where φ_{10} and φ_{01} correspond to the local phase accumulation introduced onto the single-qubit frames akin to the dispersive shifting of the coupler shown previously in the iSWAP case; φ_{ZZ} is a phase shift introduced through the ZZ shift following the qubit state being $|11\rangle$, and φ_g being the geometric phase offset intentionally introduced by the $|11\rangle \rightarrow |20\rangle \rightarrow |11\rangle$ population transfer sequence. Executing this gate for q_c being $|0\rangle$ and $|1\rangle$ results in figure 3.3(e), demonstrating that U_{CPhase} does not equal the CZ unitary in (2.2); as with the iSWAP, the coupler drive itself dispersively shifts the single-qubit frames, introducing local phase offsets. These phase offsets are identified in a very similar method to 3.2(b), with further details in [21]. The local-phase compensated unitary follows [21] as

$$U_{CPhase'} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{-i(\varphi_g + \varphi_{ZZ})} \end{pmatrix}, \quad (3.5)$$

in which we now seek parameters for setting $e^{-i(\varphi_g + \varphi_{ZZ})} = \pi$. This conditional phase is found by sweeping the CZ gate frequency in (f), and selecting for a more precise frequency where the conditional phase equals π , that is, a CZ gate. For all datapoints in (f), the qubits are adjusted in local phase to ensure that the diagonal elements in the $\{|01\rangle, |10\rangle\}$ subspace of $U_{CPhase'}$ equal 1. Additional data pertaining to these iSWAP and CZ gate DDS measurements are shown in table 3.1.




Figure 3.3: Tune-up of the CZ gate: complementary parameter estimations. (a-b) Quantum circuits for the initial determination of gate parameters A and f (a), for some initial estimate of the gate time τ , and for the precise determination of f (b), such that the gate conditionally adds a phase π to the target qubit q_t . (c-d) Determination of A and f in circuit (a) by maximisation of the population exchange between $|11\rangle$ and $|20\rangle$ at half of the gate time, $\tau/2$. (e) Readout of qubit 2 determining the phase offset $\Delta\phi$ added to q_t in circuit (b), for the two preparations of the conditional qubit q_c being $|0\rangle$ or $|1\rangle$, corresponding to $U = I$ and $U = X$, respectively. (f) Determination of the carrier frequency f which imparts a $\Delta\phi = \pi$ phase offset on q_t .

Parameter	Qubit 1	Qubit 2	Coupler
f_R	6.166 GHz	6.028 GHz	
f_{01}	3.530 GHz	4.063 GHz	4.740 GHz
τ_{01}	30 ns	30 ns	
$\eta/2\pi$	-229 MHz	-221 MHz	
$2\chi/2\pi$	-199 kHz	-628 kHz	
$g/2\pi$	52 MHz	59 MHz	
T_1	49 μ s	40 μ s	
T_2^*	18 μ s	54 μ s	
$P_{\text{Therm.}}$	1.14%	2.06%	
T_Q	37.9 mK	43.7 mK	
Φ_C			0.254 Φ_0
f_{iSWAP}			535.47 MHz
τ_{iSWAP}			858 ns
f_{CZ}			760.93 MHz
τ_{CZ}			424 ns

Table 3.1: My measured qubit, resonator, and coupler parameters of the device in figure 1.2. f_R : dressed resonator frequency; f_{01} : qubit frequency; τ_{01} : single-qubit gate time; $\eta/2\pi$: anharmonicity; $2\chi/2\pi$: dispersive shift; $g/2\pi$: qubit-coupler coupling; T_1 : relaxation time; T_2^* : dephasing time, the Ramsey free-induction decay; $P_{\text{Therm.}}$: residual thermal population; T_Q : qubit temperature, measured using the Rabi $|1\rangle \leftrightarrow |2\rangle$ technique introduced in [117]; Φ_C : magnetostatic coupler flux; f_{iSWAP} : iSWAP frequency; τ_{iSWAP} : iSWAP gate time; f_{CZ} : CZ frequency; τ_{CZ} : CZ gate time, twice the $|11\rangle \leftrightarrow |20\rangle$ transfer time.

3.3 Concluding remarks on DDS control

This chapter demonstrated DDS as a simplified control and readout architecture for superconducting quantum processor interfaces, in comparison to the conventional approach of interfaces using heterodyne mixing. Using the Presto platform, I showed complementary parameter estimation experiments for both iSWAP and CZ gates, verifying that a DDS interface supports multi-qubit operations and their tune-up experiments.

The scaling of superconducting quantum computer interfaces is constrained not only by the choice of control infrastructure but also by the physical variability of the QPUs themselves. In particular, fabrication-induced variations in Josephson junction parameters lead to deviations in qubit frequencies and increased spectral crowding as a consequence, complicating the implementation of two-qubit gates in particular. The following chapter therefore turns to the problem of post-fabrication electrical annealing of Josephson junctions as a route towards mitigating frequency crowding in superconducting QPU architectures.

CHAPTER 4

Resistance tuning of Josephson junctions

In Chapter 2, I explained that there can be an advantage for the compilation and execution of quantum circuits on a quantum processor, when the device implementation supports multiple classes of two-qubit gates rather than only one particular two-qubit gate. Specifically, we saw how a non-native SWAP operation was realised using the iSWAP and CZ gates, the procedure of which was described in detail.

However, supporting multiple two-qubit gates worsens the problem of *frequency crowding*. Recall from the introduction that transitions between different states in the system are addressed by supplying the energy difference between them. Where, a transition frequency between states i and j would be given as $f_{ij} = \frac{E_i - E_j}{h}$ via the Planck-Einstein relation $E = hf$. Figure 2.1(a) on page 13 showed the resulting energy band diagram for two qubits interconnected by one coupler, using our quantum processor architecture.

This chapter begins by explaining the implications of frequency crowding for superconducting quantum processors, to then introduce post-fabrication resistance manipulation of Josephson junctions as a means of mitigating this limitation.

4.1 Frequency crowding in superconducting quantum processors

Frequency crowding occurs when distinct transitions become insufficiently separated in frequency: the spectral width of the applied microwave tone intended to address one transition will also off-resonantly drive nearby transitions, leading to leakage and loss of fidelity.

The qubit frequencies as-manufactured are not fixed precisely at their design values. Device fabrication introduces variations in the Josephson junctions, translating into a spread in junction parameters such as their resistance, and a concomitant spread in the qubit frequencies and the spectrum of the CZ and iSWAP gate interactions.

Section 2.2 in Chapter 2 explained that driving a CZ interaction in our architecture is expected to work at two main frequencies, denoted here as $f_{\text{CZ},|02\rangle}$ and $f_{\text{CZ},|20\rangle}$, respectively. These frequencies correspond to the energy difference between the $|11\rangle$ and $|02\rangle$ states, or the $|11\rangle$ and $|20\rangle$ states. With η being the transmon anharmonicity [34], $f_{\text{CZ},|02\rangle}$ and $f_{\text{CZ},|20\rangle}$ are expected at

$$\begin{aligned} f_{\text{CZ},|20\rangle} &= |(2f_{Q_1} + \eta_{Q_1}) - (f_{Q_1} + f_{Q_2})| = |f_{Q_1} + \eta_{Q_1} - f_{Q_2}| \\ f_{\text{CZ},|02\rangle} &= |(2f_{Q_2} + \eta_{Q_2}) - (f_{Q_1} + f_{Q_2})| = |f_{Q_2} + \eta_{Q_2} - f_{Q_1}|, \end{aligned} \quad (4.1)$$

the former of which was also demonstrated experimentally in section 3.2 in Chapter 3. Similarly, the iSWAP interaction is expected at the difference frequency between $|01\rangle$ and $|10\rangle$

$$f_{\text{iSWAP}} = |f_{Q_1} - f_{Q_2}|, \quad (4.2)$$

also shown experimentally in that section.

Let's assume a Josephson junction resistance fabrication error of $\pm 2\%$ based on what is currently reported on a state-of-the-art level [118], and $\frac{\eta_{Q_1}}{2\pi} = \frac{\eta_{Q_2}}{2\pi} = -280.0 \text{ MHz}$ anharmonicity for all qubits. Given these assumptions, figure 4.1(a) shows an example architecture with target qubit frequencies 5827.6 MHz and 5494.9 MHz, and the resulting spectral crowding, when considering the CZ and iSWAP gate frequencies of equations (4.1) and (4.2). I calculate qubit frequencies using equation (4.6), which will be detailed later in this chapter.

We now extend this system with one qubit at 5255.3 MHz as shown in figure 4.1(b), along with the resulting spectral crowding. We see that small systems can still be considered fairly straight-forward to design in terms of frequency crowding, provided that the AWG for generating drive pulses is fundamentally limited by its IF baseband bandwidth, that is, the range of frequencies it can produce in the absence of mixing,

as was discussed in Chapter 3. This consideration applies generally, regardless of whether upconversion is implemented using external local oscillators or through the digital synthesis discussed in that chapter.

I also outlined why the different drive frequencies should be generated using as few LO signals as possible, to maintain phase synchronicity between qubits. As a result, transitions of interest in the spectrum should be reachable using only the bandwidth of the AWGs’ IF. For the coupler drive, we assume direct baseband synthesis without the use of mixers, consistent with the control approach used in Paper A. Under these assumptions, figure 4.1(b) shows that parametric CZ and iSWAP gates can be achieved provided that the IF baseband bandwidth of the coupler drive port is at least ~ 340 MHz.

Figures 4.1(c) further adds two qubits at 5050.3 MHz and 4850.2 MHz, with the spectrum getting further crowded by the resulting two-qubit gate transitions. Repeating the experiment from Paper A on every qubit would for this scenario require 2.3 times the IF bandwidth of the instrument I used in that paper, along with a 637 MHz mixer, set by the Q_1Q_3 CZ₂₀ and the Q_1Q_5 iSWAP gates. For this simple design of 5 qubits, the resulting band spacing is 31.4 MHz. The minimum qubit resistance fabrication precision required before a spectral collision occurs is 8.75%, corresponding to frequency precision better than 240.1 MHz in this scenario.

In figure 4.1(d), a much more complicated system of 7 qubits using three-qubit couplers is considered, by appending the two qubits $f_{Q_{6..7}} = 6077.6$ MHz, and 6522.4 MHz to the design. The resulting spectral crowding shows that the minimum distance between any two 2-qubit-gate bands is 8.4 MHz between qubits 2 and 4 executing a CZ₀₂ gate, and between qubits 2 and 7 also executing a CZ₀₂ gate. This spacing is small considering that a state-of-the-art fabrication precision was assumed for the qubit frequencies, and still required an “above average” IF bandwidth for the control AWGs for the design to be realisable. Such experiments are in principle possible using the instrument of Chapter 3, which can be configured to output an IF baseband bandwidth of 2 GHz.

To this, it should be mentioned that a 100 qubit QPU that still uses the simpler pairwise coupler architecture, approaches 0% probability of a collision-free spectrum under comparable assumptions made in this section [40].

Since the qubit frequency is ultimately set by the Josephson energy, and therefore by the junction resistance, *post-fabrication resistance manipulation* provides a potential route for frequency correction and mitigating frequency crowding, the core motivation behind the research presented in this chapter. The next section outlines the theoretical connection between the junction resistance we measure, and its corresponding qubit frequency.




Figure 4.1: Frequency crowding in different QPU architectures considering iSWAP, CZ_{20} , and CZ_{02} gate operations, using equations (4.1) and (4.2). Each spectrum has the architecture drawn as an insert, with transmons in black and the couplers in purple. (a) Two-qubit design employing a single pairwise coupler. (b) Three-qubit design using pairwise couplers. (c) Five-qubit design using pairwise couplers. (d) Seven-qubit design using three-qubit couplers. The coloured bands indicate the spread arising from a 2% fabrication precision of the Josephson junction resistance.

4.2 Junction resistance and qubit frequency

To evaluate how resistance manipulation translates into frequency changes, I first review the relation between the junction resistance and the transmon transition frequency. I base this derivation on the same one I made for Paper D, but supply additional details here.

For a transmon qubit with charging energy $E_C = e^2/(2 \cdot C)$ where $4E_C = \frac{(2e)^2}{2C}$ corresponds to the energy of a Cooper pair stored in the transmon's capacitor [34], [119], and Josephson energy E_J , the transition frequency f_{01} between the ground and first excited state $|0\rangle \rightarrow |1\rangle$ can be approximated [34], [120] to

$$h f_{01} = \sqrt{8 E_J E_C} - E_C (1 + \xi \frac{1}{4} + \xi^2 \frac{21}{128}), \quad (4.3)$$

where $\xi = \sqrt{2 E_C/E_J}$, and h is the Planck constant. The Josephson energy E_J is related to the critical current I_C via

$$E_J = \frac{\Phi_0}{2\pi} I_C = \frac{\hbar}{2e} I_C, \quad (4.4)$$

where Φ_0 is the magnetic flux quantum and e the electron charge. The critical current I_C relates to the normal state resistance R_N of the junction through the Ambegaokar-Baratoff [121], [122] relation

$$I_C = \frac{\pi \Delta_g(T)}{2e R_N} \tanh \left(\frac{\Delta_g(T)}{2k_B T} \right), \quad (4.5)$$

where $\Delta_g(T)$ is the superconducting energy gap, k_B is the Boltzmann constant, and T is the superconductor temperature.

Later in this chapter, the best-case frequency precision of the resistance tuning process will be estimated. It is worth noting at this stage that using the simplified expression [34] $h f_{01} = \sqrt{8 E_J E_C} - E_C$ instead of equation (4.3), introduces an error of 6.7 MHz. I will show that the best-case precision of the tuning process is estimated to be 11 MHz, that is, the simplified model introduces a significant error.

For Paper D, we measured the superconducting energy gap of a junction cooled to $T \approx 10$ mK in a dilution refrigerator. This experiment was a voltage-biased I-V measurement. Voltage-biasing was chosen to avoid having to collect hysteresis statistics for the switching current. The two highest peaks for $\frac{dV}{dI}$ spaced by $4\Delta_g/e$ yielded $\Delta_g(T \approx 10 \text{ mK}) = 174.3 \text{ } \mu\text{eV}$, inferring that $\tanh \left(\frac{\Delta_g(T)}{2k_B T} \right) \approx 1$ in equation (4.5). We note that this superconducting energy gap is close to the textbook number for bulk aluminium $\Delta_g(T) \approx 178.2 \text{ } \mu\text{eV}$ [123], resulting in $\frac{2\Delta_g(0.01)}{\alpha_{\text{BCS}} \cdot k_B} = T_c \approx 1.16 \text{ K}$, where $\alpha_{\text{BCS}} = 3.5$ is the ideal gap ratio [124]. For reference, the film thicknesses of this junction are 39 nm for the bottom electrode, and 78 nm for the top electrode. These thicknesses stem from optimisation of our fabrication process, with these thicknesses allowing for a high yield during the *lift-off* fabrication step, with the 1:2 thickness ratio being commonly used in the field for engineering the gap energy. These thicknesses have taken into account that the wafer undergoes angle evaporation at a 45° angle.

Substituting equation (4.5) into equation (4.4), and subsequently into equation (4.3), completes a relation between R_N and f_{01} ,

$$hf_{01} = \sqrt{8E_C \frac{\hbar\pi\Delta_g(T)}{4e^2R_N}} - E_C \left(1 + \frac{1}{4} \sqrt{\frac{2E_C}{\frac{\hbar\pi\Delta_g(T)}{4e^2R_N}}} + \frac{21}{128} \sqrt{\frac{2E_C}{\frac{\hbar\pi\Delta_g(T)}{4e^2R_N}}} \right). \quad (4.6)$$

Note that the normal state resistance R_N is a characteristic parameter that represents the Josephson junction's resistance in the cryogenic environment, for a normal electron current after gap-onset [125]. I use an empirical conversion rate $R_N = R \cdot 1.1375$, where R represents the resistance measured at room temperature. This correction is derived from measurements reported in Paper D, and is generally not to be considered universal.

We now understand how the qubit frequency is predominantly set by Δ_g , E_C , and R . The next section introduces the voltage pulse technique used for its post-fabrication adjustment.

4.3 Electrical resistance tuning

Almost all post-fabrication approaches to tuning superconducting qubit frequencies rely on heating the Josephson junction. Uniform wafer heating can significantly change junction resistance, with reported increases of up to 400% in Nb-AlO_x-Nb junctions [42] and 290% in Al-AlO_x-Al junctions [126]. However, global heating does not allow selective tuning of individual junctions, and is thus impractical for controlling frequency crowding in QPUs. Localised heating can instead be achieved through laser irradiation [43]–[46], [50] or electron-beam exposure [48], [49].

Electrical tuning [51]–[53], [127] is an alternative approach, where probes are connected to selected Josephson junctions, and an electric pulse train is applied to increase their resistance. Heating is often used in combination with electrical tuning, since it is known that heating accelerates the process [2], [128]. Recent work shows that heating can be omitted [1], [52], [127], reducing setup complexity at the expense of longer manipulation times.

The resistance increase appears to be primarily driven by two processes: the intentional manipulation of resistance, referred to here as *active resistance manipulation*, and an ensuing spontaneous¹ resistance increase referred to as *relaxation*. Paper D was a collection of empirical studies done to better understand both of these segments. A third segment exists which is the *initial resistance drop*, where

¹An automatic resistance increase that occurs without user input.

the junction resistance drops when electrical tuning starts. The physical mechanism behind this resistance drop is currently not understood, and was out of scope for Paper D.

4.3.1 Sample information and manipulator setup

The experiments presented in this chapter were performed on five Josephson junction variants fabricated on three wafers using distinct fabrication processes, summarised in table 4.1. The junction variants are classified as *low-dose*, *medium-dose*, and *high-dose*, according to the increasing oxidation dose employed during junction fabrication. Here, the oxidation dose D_{ox} denotes the empirical parameter introduced by [129], defined as a weighted combination of oxidation pressure and oxidation time used during junction evaporation. All presented data presented in the remainder of this chapter refer to these sample variants, unless stated otherwise.

Table 4.1: Summary of the five Josephson junction variants fabricated on three wafers of different oxidation doses, denoted as *low-dose*, *medium-dose*, and *high-dose*. R_W denotes the mean initial room-temperature resistances of the unmanipulated junctions. The electrode width $\sqrt{A}$ is defined assuming junctions of square cross-sections, and $R_W \cdot A$ is the corresponding resistance-area product. The oxidation parameters are the pressure p_{ox} , time t_{ox} , and dose D_{ox} , with D_{ox} being calculated using the empirical formula in [129]. Age specifies the time elapsed between fabrication and the first measurement, where t_0 is defined as the time of the deposition of the second junction electrode, corresponding to the junction *birth time* [130].

Sample	Low-dose 1	Low-dose 2	Medium-dose 1	High-dose 1	High-dose 2
R_W [Ω]	11662	5513	11057	7840	6281
$\sqrt{A}$ [nm]	200	300	350	318	354
$R_W \cdot A$ [$n\Omega m^2$]	0.47	0.50	1.35	0.79	0.79
Age [days]	> 182	> 359	> 364	> 102	> 102
p_{ox} [mbar]	2	2	10	10	10
t_{ox} [min]	20	20	60	150 ²	150 ²
D_{ox} [$mbar^{0.43} \cdot min^{0.65}$]	9.4	9.4	38.5	224.0	224.0

Measuring the room-temperature resistance of Josephson junctions can be done by monitoring an applied voltage and a consequently sourced current, if the applied voltage is small enough for the junctions to behave like a linear resistor [1]. The junction variants in table 4.1 have such a linear resistance R_W on the order of $k\Omega$, when biased with a few mV. The corresponding currents are therefore on the order of μA at the extremes of the measurement range, and reach nA near zero bias. Such sensitive voltage and current requirements are conveniently fulfilled by a source-

²The oxidation time for the *high-dose* wafer is an estimate, due to a cleanroom power failure with no machine logs being saved for the actual time.

measurement unit (SMU), which combines a programmable voltage source with a current measurement circuit in a single instrument. Notably, AC mains interference is therefore a source of measurement error, but can be mitigated by choosing the measurement duration to be an integer multiple of $\frac{1}{50\text{ Hz}} = 20\text{ ms}$.

The junction manipulation and resistance measurements were performed with a Keysight B2902B SMU connected to a Keyfactor SimplePS manual probe station, equipped with MPI MP40 probes using PT-W-S-2MIC tungsten needles. This setup is shown in figure 4.2(a), while (b) shows a rendering of a junction being probed on its microscopic test patches. The SMU can conveniently also apply the substantially larger DC biases required for junction manipulation. The same electrical contacts could therefore be retained on the junction test patches throughout an experiment, also while monitoring the resistance changing during relaxation, with the SMU alternating between manipulation and measurement. All measurements were performed in an electrostatic discharge (ESD) protected environment, with an ioniser fan continuously directed towards the sample to mitigate charge build-up.

The junction resistances were measured using I-V sweeps from -13 mV to $+13\text{ mV}$, in steps of $500\text{ }\mu\text{V}$. I determined the value of 13 mV experimentally by selecting the largest voltage range for which all junction variants fit to a linear regression model $V = R \cdot I$ with a high coefficient of determination $r^2 \geq 0.999$.

Manipulating the resistance of the junction while performing the resistance measurement is not possible with the setup as-described, since the manipulation itself applies voltages far beyond the resistor-like voltage interval. However, if the resulting I-V measurement could instead be fitted to the non-resistor-like behaviour of a junction at room temperature, then larger measurement ranges might be possible. This direction was not explored further in my work.

Two chips were also made for testing qubit performance following resistance manipulation of the qubits' junctions. The layout of these samples is shown in figure 4.2(c), comprised of eight fixed-frequency transmon qubits with readout resonators coupled to a common feedline. For cryogenic characterisation, these samples were thermalised at the 10 mK stage of a Bluefors LD250 dilution refrigerator with heavily attenuated RF lines following the setup approach described in [31]. The samples used for qubit tuning experiments shared a common input/output readout line in all cryogenic measurements, with coaxial microwave switches used to select between samples. Waveform generation and readout were implemented using the QPU interface instrument developed in Papers B and C, previously reported on in Chapter 3.




Figure 4.2: Manipulator setup. (a) SMU and probe station setup. (b) Rendered illustration showing the four tungsten probes in contact with a junction test structure during resistance manipulation and measurement. The upper two leads are the positive force and sense terminals, while the lower being the negative force and sense terminals, respectively. The negative force terminal is grounded to the probe station and mains earth, with the SMU held floating otherwise. Typical test structures can be seen in the upper right and lower left corners of (c) the quantum devices tested in Paper D. These chips comprise eight fixed-frequency transmon qubits coupled to readout resonators, sharing one feedline.

4.3.2 The voltage pulse technique

The experimental procedures and measurement setup described in this section apply to all experiments presented in this chapter, unless otherwise stated.

I manipulated the Josephson junction resistances using a bipolar voltage pulse train consisting of a positive-bias square pulse, a blanking interval, a negative-bias square pulse, and a second blanking interval of equal duration. This approach closely follows the alternating-bias assisted annealing technique developed at Rigetti [2], [52], [127], with modifications to the probing and measurement procedure. The main reason for using a bipolar pulse train follows the work by [1], [2], which has shown that bipolar pulses increase the speed and available range of resistance manipulation, compared to single-sided manipulation schemes.

The pulse sequence is illustrated in figure 4.3(a). Each manipulation experiment

consists of N iterations. In every iteration, a number M of square-wave voltage pulses are applied, separated by a 100 ms blanking interval, after which the junction resistance is measured. The pulse amplitude I used was $\pm V_a = 0.85$ V and $M = 6$ pulses per iteration. The square-wave pulses have a period of 200 ms with symmetric positive and negative bias amplitudes. I selected these values experimentally: since Paper D studied multiple junction variants, I experimented with various voltages and pulse counts until I found a protocol that could be applied to all variants studied without breaking the junctions. The idea was to use the same protocol for all studies, motivated by consistency across comparisons between junctions.

Figure 4.3(b) shows the typical junction response during active manipulation. As was hinted to earlier, we observe three distinct regimes: (i) an initial decrease in resistance, denoted R_{Drop} ; (ii) the active manipulation phase, during which the resistance is intentionally increased to R_{Active} ; and (iii) a post-manipulation relaxation characterised by a rapid rise followed by a long-tailed saturation after the manipulation is halted at time t_{Stop} . The indicated value R_{30} will be used when discussing the relaxation behaviour later, and represents the junction resistance measured 30 min after the manipulation procedure has been stopped.




Figure 4.3: Active resistance manipulation of a Josephson junction using two different sequences of voltage pulses. (a) Schematic illustration of the *active* resistance manipulation protocol. (b) An example resistance trace of a junction manipulated up to the time t_{Stop} with the scheme in (a), showing the initial resistance drop, active manipulation, and subsequent relaxation. The inset shows a magnified view of the initial resistance drop.

4.4 Active resistance manipulation

To characterise and quantify resistance changes we define the resistance increase as

$$\Delta R(t) = \frac{R(t)}{R(0)} - 1 \quad , \quad (4.7)$$

where $R(t)$ is the junction resistance measured at time t , with $R(0)$ being the initial resistance. We express $\Delta R(t)$ in units of % to enable comparison across junctions with different initial resistances and wafer origins. This method of expressing the resistance increase of a junction has emerged as a standard convention in the field. However, where to define the initial $t = 0$ differs depending on how to interpret the initial resistance drop. An example is to define $t = 0$ at the lowest point of this resistance drop [127], while I set $t = 0$ at the first resistance measurement as can be seen in figure 4.3(b).

When experimenting for Paper D, I noticed that the speed of the resistance increase was difficult to predict. This speed relation had to be quantified in order to be able to use electrical tuning for achieving some calculated qubit frequency. All junction variants were thus tuned at different voltage amplitudes, resulting in figures 4.4(a-e). Across all devices, increasing V_a always yielded a higher rate of resistance change over time, with a maximum increase of 96.2% achieved in 5 min.

Since the underlying physical mechanism is not known, we approached this problem empirically. The active resistance manipulation can be fitted to a second-order polynomial equation

$$\Delta R(t) = \alpha(V) \cdot t + \beta(V) \cdot t^2 \quad . \quad (4.8)$$

A comparison to higher order models was done in Paper D, and resulted in near-identical fit errors at the cost of extending the model to higher orders. I therefore proceeded with using a second-order polynomial empirical model motivated by it being simpler, that is, the empirical model relies less on a large number of unexplained higher orders for fitting any arbitrary data with a good match. Similarly, power-law and exponential models were also explored, but were worse predictors when compared to experimental data.

This active resistance manipulation model was intended for the active resistance manipulation regime. The initial resistance drop was therefore discarded prior to fitting, illustrated in figures 4.4(f, g), and also discussed further in Paper D.




Figure 4.4: Resistance increase during active manipulation at multiple pulse amplitudes for (a) *low-dose 1*, (b) *low-dose 2*, (c) *high-dose 1*, (d) *high-dose 2*, and (e) *medium-dose 1* junctions. Manipulation at voltages exceeding the highest value shown in each panel resulted in junction failure. The time-dependent resistance change is dominated by a linear term $\alpha(V)$ with a quadratic contribution $\beta(V)$, as defined in equation (4.8). (f) and (g) show the duration of the manually excluded initial resistance drop for each device set in (a-e), as a function of the manipulation voltage (f) and the maximum resistance change (g).

The two factors $\alpha(V)$ and $\beta(V)$, being linear and quadratic time dependencies, are therefore analogous to the speed and acceleration of the resistance increase. Figure 4.5 shows the parameters obtained from fitting the data in figures 4.4(a-e) with equation (4.8).

The predominantly negative values of $\beta(V)$ indicate a decreasing rate of resistance change over time, with occasional positive values attributed to fitting uncertainty near zero. The crossover between $\alpha(V)$ and $\beta(V)$ occurs at a median time of 26 min, mean 50 min, across the junctions in figure 4.5. The implication is that the linear term is dominant over the experimental timescales investigated in Paper D. A second implication is the existence of a time where the resistance stops changing. Figures 4.5(c-e) initially suggest a weak anticorrelation between $\alpha(V)$ and $\beta(V)$, suggesting that junctions that increase in resistance faster also stop changing resistance sooner; however, the limited data set does not permit making such a conclusion reliably, but may suggest an opening for a future study.

Since $\alpha(V)$ was dominant, this parameter was further studied in Paper D to establish its dependency on voltage. $\alpha(V)$ in figure 4.5(a) was found to have an exponential dependence on voltage; I reached this conclusion after successfully modelling $\alpha(V)$ as

$$\alpha(V) = \alpha_0 \cdot e^{V/V_0} \quad , \quad (4.9)$$

where V_0 is the characteristic voltage corresponding to an e -fold increase in $\alpha(V)$, with α_0 being a prefactor modifying the voltage sensitivity. The extracted parameters are listed in table 4.2; increasing the pulse amplitude by approximately 200 mV results in at least an order-of-magnitude increase in the resistance-change rate for the *low-dose* and *high-dose* junctions, whereas the *medium-dose* junctions exhibit a more modest increase of about a factor of two over the measured range. Figure 4.5(a) also shows that samples from the same wafer cluster according to their resistance-area product $R_W \cdot A$, rather than junction area. This observation is consistent with previous reports indicating that electrical tuning exhibits only weak dependence on junction area [2].

Furthermore, the *medium-dose* manipulation results show that the resistance-area product alone is not sufficient for predicting the exponential voltage dependence, since the resistance-area product deviates from the trend held otherwise, where a higher resistance-area product equates to a sample showing a weaker response to increased manipulation amplitude. An additional factor not addressed in Paper D is that $R_W \cdot A$ is dependent on the natural ageing of the junction since R_W is not normalised, and will spontaneously change over time. This issue does not change the aforementioned conclusion that the $R_W \cdot A$ trend in figure 4.5(a) does not hold. However, future attempts at identifying a dependency between $R_W \cdot A$ and $\alpha(V)$ should take this ageing constraint into consideration.

As seen in table 4.2, the prefactor α_0 spans over three orders of magnitude across the junctions while the extracted characteristic voltage V_0 lies in the range 55–91 mV. Interestingly, these values are higher but comparable to the thermal voltage at room temperature, $V_T \equiv k_B T / e \approx 25$ mV, consistent with slow resistance changes due to ageing at room temperature.




Figure 4.5: Extracted parameters (a) $\alpha(V)$ and (b) $\beta(V)$ obtained by fitting equation (4.8) to the resistance traces in figure 4.4. Panels (c-e) compare the voltage dependence of $\alpha(V)$ and $\beta(V)$. The linear coefficient $\alpha(V)$ exhibits an exponential dependence on voltage. The resistance-area product $R_W A$ [nΩ m²] for the five junction variants is indicated in the legends of (a) and (b).

Active manipulation always carries the risk of breaking the junction, most commonly through the formation of a short-circuit across the junction. This breaking mechanism has been linked to the glass-like behaviour in Al-AlO_x-Al Josephson junctions during annealing and relaxation [130]. The suggestion was therefore that allowing the junction to relax between successive manipulation periods may reduce the accumulation of structural strain. To investigate this possibility, I introduced a *stepped* active manipulation protocol, in which the active resistance manipulation is interleaved by defined relaxation intervals. This stepped protocol is the topic of the following section.

Table 4.2: Parameters α_0 and V_0 as extracted by fitting the data in figure 4.5(a), and α'_0 being the value for α_0 converted into units of Ωs^{-1} by multiplying with R_W . For α_0 , the % denotes the resistance increase percentage. The characteristic voltage V_0 represents the additional voltage required for the rate of resistance change to increase by a factor e .

	Low-dose 1	Low-dose 2	Medium-dose 1	High-dose 1	High-dose 2
V_0 [mV]	72.5(24)	70.4(34)	91.4(43)	57.5(50)	55.3(31)
α_0 [10^{-9} %/s]	824(351)	420(271)	1002(531)	0.79(119)	0.37(37)
α'_0 [$\mu\Omega s^{-1}$]	98.2(23)	21.5(13)	108.4(33)	0.0650(5)	0.0253(5)

Stepped active manipulation

For these experiments in Paper D, I performed active resistance manipulation until a fixed resistance increment $\Delta R(t)$ had been reached. The junction was then allowed to relax for a waiting time t_{Relax} before the manipulation was resumed. The protocol is illustrated in figure 4.6.




Figure 4.6: (a) Schematic illustration of the *stepped* resistance manipulation protocol. Active resistance manipulation is followed by a relaxation interval of duration t_{Relax} , during which the junction resistance is repeatedly measured with a period t_{Probe} . The manipulation-relaxation sequence is then repeated. (b) Measured resistance evolution during three successive active manipulation periods and the intervening relaxation intervals. The resistance increase during each active manipulation period follows the same second-order polynomial behaviour shown in figure 4.3(b), but appear vertical here due to the timescale. The inset in (b) shows the beginning of the third relaxation event.

The stepped active manipulation protocol was used to increase the junction resistance beyond the 96% increase achieved in figure 4.4(a). Figure 4.7 shows the resistance evolution of a *low-dose 1* junction, aged 348 days at the time of measurement, during successive manipulation steps. The resistance increased by 269.5%, corresponding to a transmon frequency shift of $\Delta f = -1964$ MHz according to equation 4.6. To my knowledge, this is the largest resistance change reported for

an individually targeted Josephson junction. Upon further stepping, the junction ultimately failed during the initial resistance drop of the final manipulation step. Two systematic changes are observed with successive manipulation steps: the rate of resistance increase decreases, and the depth of the initial resistance drop increases.




Figure 4.7: Stepped active manipulation of a 200 nm *low-dose* junction. Each step consists of an active resistance manipulation sequence that increases the junction resistance by $+0.10 \cdot R(0)$, followed by a 3 h relaxation period. The steps are repeated until junction failure, achieving a resistance change of $\Delta R = +270\%$. Using equation (4.6), this corresponds to a transmon frequency shift of $\Delta f = -1964$ MHz.

A substantial portion of the total resistance increase appeared during the relaxation intervals. This relaxation phenomenon can not be discarded for the electrical tuning technique, and must be understood for the technique to be accurate. In Paper D I therefore attempted to model the relaxation behaviour, which is the topic of the next section.

4.5 Residual resistance increase post-manipulation

Seemingly after any amount of electrical tuning, the Josephson junctions undergo a relaxation phase where they spontaneously increase resistance. This phase is shown as the long-tailed logarithmic decay in figure 4.3(b), starting just after the junction resistance has been intentionally increased by R_{Active} . Paper D showed that this decay is logarithmic and seemingly without any saturation point within the first 24 hours after manipulation; being able to model the relaxation is therefore crucial for actually using it to predict where the final junction resistance will end up. Recent work in the field [127] certifies that the variance of this final resistance is better than the junction resistance variance acquired as-fabricated.

Consistently through my work in Paper D it appeared as if the relaxation mechanism was a separate function that appeared regardless of the active resistance

manipulation. This trend held even during the stepped active manipulation, whose relaxation intervals are shown in figure 4.8. Here, (a) shows the relaxation response normalised to the initial resistance at the beginning of the experiment $R(t=0)$, where each trace corresponds to the relaxation phase of a single manipulation step in figure 4.7. It is seen that the relaxation behaviour appears inversely logarithmic; other models also exist in the field that instead further divide it into piecewise regions [2], [52].

Specifically with the Paper D result in figure 4.8(b), I wanted to answer the question “does the relaxation stack between manipulations”. Indeed, by studying the resistance increase at the markers in (a), I do claim that the relaxation speeds up when repeated. It so would seem that previous relaxation intervals do contribute to subsequent relaxations, but an exact number was not established in Paper D.




Figure 4.8: Analysis of relaxation during the stepped active manipulation in figure 4.7. (a) Lin-log₁₀ plot of the relaxation after each step, normalised to $R(0)$. The relaxation traces mainly follow the same logarithmic time dependence as the first step, indicating repeated activation of the same relaxation process rather than a new one from a different resistance baseline. (b) Resistance increase evaluated at the 30 min and 180 min markers indicated in (a) as dashed vertical lines, shown as stars and diamonds, respectively. The increasing slope across steps indicates a growing relaxation rate with successive manipulations, suggesting ongoing relaxation from previous steps accumulates.

Further studying the relaxation, specifically the degree it appeared to be independent of the active resistance manipulation, we noticed that prior work [51] had reported that the relaxation would always add a fixed amount of resistance increase regardless of how much active resistance manipulation had been performed. Such an observation would have proved very useful for being able to model and use electrical tuning, and I therefore wanted to reproduce this result.

For this hypothesis to be correct, one could create a plot of the total resistance increase as a function of the active resistance increase. A large set of junctions would be tuned to various resistance increases, and their resistance at some determined time t would be recorded. A linear slope of 1 in such a plot would confirm the hypothesis, with the trace offset corresponding to the resistance increase gathered from relaxation. Recall that the relaxation initially appeared to never saturate; I therefore defined some time t for which I would study the junctions' relaxation. Since the probe station was manual, I settled for $t = 30$ min, which is shown as R_{30} in figure 4.3(b).

The results from Paper D are reproduced in figure 4.9, showing R_{30} as a function R_{Active} for (a) *low-dose 1*, (b) *medium-dose 1*, (c) *high-dose 1* and (d) *high-dose 2* junctions. The *high-dose* samples reproduced the previously reported behaviour, whose slopes both fitted to 1.00 within the error bar. The *low-dose 1* and *medium-dose 1* exhibit larger slopes of 1.13 ± 0.01 and 1.08 ± 0.02 , respectively. With this result, Paper D showed that the relaxation process can indeed depend on the amount of active resistance manipulation performed. And, that such behaviour was device-dependent, and also linear; the linear property meant that the resistance added from relaxation could in principle still be predicted, linearly, based on the offset and slope of some sacrificial set of example junctions.

The *low-dose 2* junctions are notably absent from this dataset, the main reason being that *low-dose 2* was the last junction variant to be studied for Paper D, and a larger focus was being put on studying the *low-dose 1* junctions over time, which will be presented later. The data for *low-dose 1* had already disproved the *relaxation is independent of manipulation* hypothesis at this point in time, since the slope was greater than 1.

To supplement the slope and offset study, I wanted to know to which extent such a relaxation predictor would change over time. I therefore wanted a plot showing the linear fit slope k and offset m over time. I made this analysis on the *low-dose 1* junctions, for which the largest data set was available. The result was figure 4.10, showing that neither the slope nor the offset parameter stabilises within 24 h.

To be able to use the time-dependent k and m , I modelled their time dependence empirically to a logarithmic growth function,

$$y(t') = a + b \cdot \ln(1 + t'/\tau) \quad , \quad (4.10)$$

which was also compared to a power-law fit reported in Paper D. The variable t' signals that relaxation has commenced, that is, $t = t_{\text{Active}} + t'$, where t is the time since the tuning was started. The extracted fit parameters are shown in the legends. The key purpose of attaching such a model to the time-dependent linear




Figure 4.9: Linear fit of the *total* resistance manipulation at R_{30} as a function of the *active* resistance manipulation R_{Active} for (a) *low-dose 1*, (b) *medium-dose 1*, (c) *high-dose 1*, and (d) *high-dose 2* junctions. The vertical axes show $R(t_{\text{Stop}} + 30 \text{ min})/R(0) - 1$, while the horizontal axes show $R(t_{\text{Stop}})/R(0) - 1$. A slope of unity indicates that the relaxation contributes a fixed resistance increase offset independent of the actively induced resistance increase R_{Active} .

fit slope $k(t')$ and offset $m(t')$ parameters, was to be able to use these parameters when predicting the final resistance at some time, rather than trying to identify the underlying physical mechanism behind relaxation. Figure 4.10(a) shows that the resistance increase gathered only from the amount of active manipulation made is substantial: in addition to the manipulation-independent relaxation resistance increase, approximately a quarter of R_{Active} is added to ΔR_{Total} .

The implications are that both the slope and offset of the relaxation process must be characterised to achieve accurate electrical tuning, requiring multi-day monitoring of statistically significant ensembles of sacrificial junctions.




Figure 4.10: Time dependence of the slope (a) and offset (b) extracted from linear fits of total versus active resistance manipulation. An example of the $t' = 30$ min fit is shown in figure 4.9(a).

4.5.1 Tuning frequency precision from the relaxation offset error

Figure 4.10(b) showed the portion of the total resistance increase that is independent of the preceding active manipulation. Here, the resistance increase was $3.29 \pm 0.19\%$. The uncertainty of $\pm 0.19\%$ is important for frequency targeting: it represents an intrinsic uncertainty in the final resistance after a junction has been tuned. For comparison, the error bar of $\pm 0.19\%$ should be compared to recent results reporting similar post-tuning values of $\pm 0.4\%$ [127].

Since the qubit frequency depends on the junction resistance by equation (4.6), this uncertainty establishes a lower bound on the frequency precision achievable using the current protocol. This section will show how this uncertainty sets a lower bound of 11 MHz; I feel obliged to stress that Paper D presents this number as the best-case lower bound of the precision: without further understanding of the relaxation mechanism, there is no way the tuning process of Paper D can reliably promise a qubit frequency tuning better than 11 MHz, due to the error bar of the relaxation stemming from unknown processes.

I repeat this calculation here using the updated value $R_N/R = 1.1375$ adopted throughout this chapter, rather than the value $R_N/R = 1.1385$ used in Paper D. The latter value was based on earlier measurements from our group [118].

Using the experimental mean values from Paper D, namely $f_{01} = 5.4910$ GHz and $\eta/2\pi = -203.0$ MHz for qubits $Q_{5..8}$, together with the measured superconducting

gap $\Delta_g = 174.3 \mu\text{eV}$, E_C can be obtained from equation (4.11):

$$\frac{\eta}{2\pi} = -E_C - \frac{9}{16}(E_C \xi) - \frac{81}{128}(E_C \xi^2) - \frac{3645}{4096}(E_C \xi^3) - \frac{46899}{32768}(E_C \xi^4) \quad . \quad (4.11)$$

Here, ξ is calculated using the same definition as in equation (4.3). Now enforcing $\eta/2\pi = -203.0 \text{ MHz}$ in equation (4.11) while simultaneously solving (4.6) to $f_{01} = 5.4910 \text{ GHz}$, creates a bound solution for legal values of R . I find R numerically to be 5526.3Ω , which is 4.9Ω different from the value obtained in Paper D using $R_N/R = 1.1385$. Note that the new R still yields $E_J = 21.63 \text{ GHz}$ from equations (4.4) and (4.5), with $E_C = 186.7 \text{ MHz}$ from equation (4.11), since the boundary conditions based on measured numbers for η , f_{01} , and Δ_g , restrict the solution such that it is R that changes when R_N/R changes. I repeat the calculation for $R \pm 0.19\%$, which still yields a frequency error of $5496.6 - 5485.3 = 11.3 \text{ MHz}$, despite the different $R_N/R = 1.1375$ used here.

4.5.2 Impact from error in other parameters

While Paper D focused on the resistance aspect of Josephson junctions, I show here that other parameters contribute on a very similar order, when attempting to target specific qubit frequencies. In figure 4.11, I simulate and show the result of a Monte Carlo analysis where 1000000 junctions were virtually manufactured with randomly sampled parameters as input into equation (4.6). To gauge the impact that randomising the parameters had, a representative qubit was given parameters $R = 5275 \Omega$, $E_C = 183 \text{ MHz}$, $\Delta_g = 174.3 \mu\text{eV}$, $T = 10 \text{ mK}$, and $R_N/R = 1.1375$. Each parameter was then given a variance corresponding to 2% of its initial value. The simulated results are shown in table 4.3: the precision when measuring the resistance of the junction, the precision in the conversion factor between room temperature resistance and the cryogenic normal state resistance, and the precision in stating the superconducting energy gap, are all equally important for determining the qubit frequency from theory. Moreover, the precision in stating the charging energy, carries a significance on a similar order but slightly lower impact, at $\sim 7\%$ less variance given our representative values. Error in the temperature of the cryogenic environment is of no significance, with simulations yielding a $16.2 \mu\text{Hz}$ variance.

4.6 Practical application

A practical implementation of the resistance manipulation technique requires a wafer-level calibration strategy. Given the wafer-to-wafer variability observed in Paper D, each fabricated wafer should include a statistically meaningful set of dedicated test junctions that are geometrically identical to the junctions incorporated




Figure 4.11: Monte Carlo simulations where 1000000 junctions are fabricated virtually, and the resulted qubit frequency is calculated using equation (4.6). To gauge the significance of each parameter, a variance of 2% is applied onto the (a) resistance of the Josephson junction as measured in room temperature, (b) the difference between this resistance and the true normal state resistance of the junction in its cryogenic environment, (c) the charging energy E_C , and (d) the superconducting energy gap Δ_g . A 2% variation in T yields a qubit frequency variance of $16.2 \mu\text{Hz}$ as-simulated, and is not shown.

Table 4.3: Qubit frequency variance when varying parameters by $\pm 2\%$ for 1000000 simulated junctions.

Varied parameter	$\pm$ Variation	f_Q variance [MHz]
Resistance [Ω]	5275×0.02	57.71
E_C [MHz]	183×0.02	53.69
Δ_g [μeV]	174.3×0.02	57.66
Temperature [K]	0.010×0.02	1.62×10^{-11}
R_N/R ratio [-]	1.1375×0.02	57.75

into qubits. These test junctions serve as per-wafer calibration standards. Following fabrication, the wafer should be left undisturbed to undergo natural ageing, which

we estimate to stabilise sufficiently after approximately 20 days at room temperature. Existing work suggests that ageing effects can be sped up by heating the wafer [126]. Once aged, a subset of its test junctions can be subjected to voltage pulses of incrementally increasing amplitude. By extracting $\alpha(V)$ and $\beta(V)$, and subsequently fitting an exponential voltage dependence for parameters α_0 and V_0 , the wafer-specific parameters governing active manipulation can be determined. This step establishes the effective tuning voltage window and quantifies the amplitude sensitivity, that is, the voltage increment required to change the manipulation rate by a factor of e . Subsequently, a larger ensemble on the order of 100 junctions should be manipulated to varying target resistance increases. Then, their post-manipulation relaxation must be monitored over timescales exceeding 24 hours. Because the relaxation follows a logarithmic time dependence which we've shown can be dependent on the preceding manipulation done, reliable parameter extraction requires long-duration monitoring to disentangle the actively induced change from the passive relaxation contribution. Only once both the active manipulation dynamics and the relaxation behaviour are characterised for the wafer can qubit junctions be targeted with predictable frequency corrections. Additionally, the superconducting energy gap for the wafer must be characterised, thus a subset of the test junctions must undergo I-V measurements, preferably so in a dilution refrigerator to remove lingering uncertainty over the temperature dependence of $\Delta_g(T)$.

A remaining parameter to characterise is E_C . As was shown in figure 4.11(c), characterisation of E_C carries a near-equal importance for reliably predicting the qubit frequency as does the resistance. I claim future work is needed to map E_C as a function of junction tuning. Experimentally, E_C can be measured following the tune-up done in Paper A, that is, determining the $|1\rangle \rightarrow |2\rangle$ transition frequency in Ramsey interferometry with $|1\rangle$ as the input state. Then, E_C can be experimentally determined using equation (4.11).

An actual machine where electrical tuning is codified into a finished product carries additional caveats. From experimentation we know that the devices are light sensitive, that is, ambient light changes the as-measured resistance in every measurement on the order of a few ohms; such change is immediately impactful for the intended purposes, with $\pm 5\ \Omega$ corresponding to 5.7 MHz for the junction parameters simulated in figure 4.11. Similarly, as was experienced first-hand, Josephson junctions perish wildly when subjected to ESD, and any machine for tuning requires a by-the-book rigid control of ESD, due to the breakdown voltage of a single volt or less.

4.7 Concluding remarks on resistance manipulation

The fabrication of Josephson junctions introduces a variation in their resistance, which translates into a variation of qubit frequencies, which can become problematic in the crowded spectrum of a multi-qubit system as inadvertent driving of a non-target qubit transition can easily occur. The required frequency precision depends on multiple factors such as chip scale and architecture. For example, the seven-qubit layout shown in figure 4.1(d), with 1 GHz AWG bandwidth, requires 8 MHz precision to enable both iSWAP and CZ operations.

In this chapter, I have presented my work on electrical post-fabrication adjustment of junction resistances. While several post-fabrication techniques rely on thermal annealing, electrical annealing appears to operate through a different mechanism that is not primarily driven by heating. I characterised the electrical junction manipulation through a series of experiments, where three distinct regimes were seen during electrical annealing: an initial *resistance drop*, an *active resistance manipulation* phase in which most of the resistance change occurs, and a subsequent *relaxation* process where the junction continues to increase in resistance after manipulation has ended.

I found that the active resistance change over time can be modelled as a second-order polynomial, and that the *speed* of resistance change depends exponentially on the applied voltage. This model is purely empirical, and further work would be needed to pinpoint the physical mechanisms resulting in the observed behaviour. The achievable tuning range is substantial, with examples of nearly 100% resistance increase in 5 minutes, and 270% if done sequentially over the course of 70 hours, corresponding to a transmon qubit changing its frequency by nearly 2 GHz. The relaxation that follows the active manipulation has a decaying logarithmic behaviour and does not settle for over 24 hours, which means that junctions must be monitored for a long time to determine how much resistance a particular batch of junctions will gain post-manipulation. Once the long-term behaviour is known, figure 4.10 establishes that the relaxation is reproducible to a precision of $\pm 0.2\%$ ΔR for the *low-dose 1* junction variant. I showed that the magnitude of the relaxation can depend strongly on the active manipulation, with observed resistance increases of up to 1.2 times the intentionally induced resistance change.

From my observations, I estimated that our tuning process today has a best-case frequency precision of 11 MHz, currently limited by variance in the relaxation process. This precision represents a five-fold improvement over our current fabrication precision for junctions of size 300 nm $\times$ 300 nm. However, as I showed in the beginning of this chapter, small 5-qubit grid designs do not require this tuning technology to avoid frequency crowding, provided a 9% resistance precision as-fabricated, a 1 GHz IF bandwidth of the control AWG, and thoughtful selection of the qubit

frequencies. In contrast, Osman et al. [40] have argued that 100-qubit designs attain 80 percent probability of collisions in the spectrum at 40 MHz qubit frequency imprecision as-manufactured.

As I have shown in this chapter, resistance alone is insufficient for accurate frequency targeting. Other parameters, including the superconducting energy gap, E_C , and the conversion between room-temperature and cryogenic resistance, contribute comparably to the total frequency uncertainty to a nearly identical degree. This means that, for the electrical tuning process to be reliable, being able to measure and control these other factors is paramount for the frequency prediction model itself to be reliable.

CHAPTER 5

Conclusion and outlook

The common conclusion of this thesis is that scalability aspects of the hardware stack are important and will ultimately triumph in practice over ad-hoc systems employed today. Building on my findings, I present the following conclusions.

I conclude that quantum gate compilation stands to benefit from larger gate sets. The work on quantum gate compilation shows a reduction of the quantum circuit depth, an improvement that surpasses the NISQ era, since settling for the bare-minimum number of quantum gates needed for universality always imposes the risk of three extra two-qubit gates for unsupported operations. Experimentally, I demonstrated a realisation of this enlarged gate set of iSWAP and CZ gates on a fixed-frequency qubit tunable-frequency coupler architecture. Several different quantum processing platforms stand to benefit from the tools developed in the niche field of gate inequivalency, such as the Weyl chamber map of locally equivalent gates, useful for predicting which two-qubit operations can efficiently be replaced with each other at compilation.

I conclude that DDS should be investigated further as a reasonable contender for larger QPU interfaces. DDS interfaces can replicate those in conventional setups, including single- and two-qubit gate support. The key advantage presented by the DDS approach is found in the reduction of hardware complexity, where multiple

components are removed from the signal path.

I conclude that electrical tuning of Josephson junctions is clearly worthy of further investigation as a post-manufacture remedy of frequency crowding in QPUs. The technique offers selective tuning of junctions in the fixed-qubit-frequency QPU architectures studied empirically in my work. My results present a pathway forward for how to use the technique in a practical setting, but the effects of electrical tuning on material properties other than resistance remain to be established. The relationship between tuning, ageing, and thermal annealing remains to be established, as well as mechanisms underlying relaxation. My work motivates further investigation of the physical mechanism at play during electrically-induced resistance change in Josephson junctions, particularly the emerging hypothesis that the resistance change phenomenon is Arrhenius-like in nature, supported in preliminary hints of a thermally activated mechanism.

5.1 Future work

I end this thesis with a brief list of ideas for future ambitions enabled by my work.

SWAP gate executed by a bichromatic application of an iSWAP gate and a CZ gate. The follow-up experiment from Paper A is to combine the iSWAP and CZ gate pulses into a bichromatic SWAP pulse. Such a gate pulse would effectively create a SWAP gate for the fixed-qubit tunable-coupler platform. Adding the bichromatic SWAP to the gate set would enable quantum circuits for this platform where any two-qubit gate from the Clifford group can be executed with at-most a single two-qubit gate substitution. The theoretical framework underlying such bichromatic driving has already been introduced [131], with previous work highlighting that simultaneous support for the iSWAP-like and CPHASE-like classes risks skewing what entanglement coordinates get executed on certain architectures [24].

Compiler efficiency impact from supporting iSWAP-like and CZ-like gates. Paper A highlights that compiling for a non-supported two-qubit gate *risks* introducing three additional gates. Multiple special cases exist where such compilation does not happen; for instance, $U_{CZ}(|00\rangle) = U_{iSWAP}(|00\rangle)$. A study is warranted where the prevalence of such special cases is studied for quantum circuits of interest; is it common that an unsupported SWAP gate is substituted for three two-qubit gates?

Experiments in quantum error correction enabled by the CZ and iSWAP classes.

Since the CXSWAP gate is locally equivalent to the iSWAP gate, and the CX gate being locally equivalent to the CZ gate, simultaneous support for iSWAP and CZ gates allows for experiments with Vine codes on the common square-grid architecture without compiler overhead from two-qubit gates [132].

DDS-specific interface disadvantages. Larger DDS interfaces would expose where the technology breaks; more specific examples of understudied topics in DDS interfaces include bitness versus qubit fidelity. Previous work establishes that data converter bitness impacts the qubit reference frame [113], alas a systematic study is warranted where such bitness is artificially lowered and studied as a function of fidelity. The motivation for such work is scalability. Another understudied example is a systematic power consumption study comparing DDS to heterodyne interfaces.

Map electrical tuning to the Arrhenius model Building on the results from my work on electrical tuning, and particularly the hypothesis of electrical tuning showing an Arrhenius-like behaviour, future work should attempt to fully map electrical tuning to this model; for instance, by the application of heat, and by modelling some “equivalent activation energy” given various pulse settings. A complement to these studies should be the identification of the chemical orderliness of the barrier pre- and post-tuning, in turn building on previous results hinting that this orderliness increases from tuning [2], [52]. A side-effect of tuning is still the observed resistance drop phenomenon [51], which remains to be scientifically studied. Contribution from additional effects also remain to be studied, for instance whether there is electromigration.

Establish working procedures for characterising and tracking the cryogenic resistance conversion ratio Figure 4.11 highlights that the resistance conversion ratio 1.1375, and its errors, are of equal importance as the resistance and superconducting energy gap parameters when predicting the final qubit frequency. This ratio equals the increase in normal-state resistance of the junction at cryogenic temperatures. Still, our working procedures regarding this resistance conversion ratio assumes a mere estimate of 14%, with little care taken to characterise this number across fabricated wafers, and, its variance. A promising procedure was developed in Paper D, where junctions are cooled in liquid nitrogen alongside a thermistor. The procedure then uses a Simmons model for acquiring the conversion ratio experimentally. This procedure should be explored further to establish the variance of the ratio for predicting the final qubit frequency after electrical tuning.

References

- [1] M. Toselli, “Post-fabrication Frequency Tuning in Superconducting Transmon Qubits,” Master’s thesis, Chalmers University of Technology, Gothenburg, Sweden, 2024.
- [2] D. P. Pappas, M. Field, C. J. Kopas, *et al.*, “Alternating-bias assisted annealing of amorphous oxide tunnel junctions,” *Communications Materials*, vol. 5, no. 1, Aug. 2024. DOI: 10.1038/s43246-024-00596-z.
- [3] J. Preskill, *Quantum computing and the entanglement frontier*, Rapporteur talk at the 25th Solvay Conference on Physics, “The Theory of the Quantum World”, Brussels, 19-22 October 2011, 2012. DOI: 10.48550/ARXIV.1203.5813.
- [4] F. Arute, K. Arya, R. Babbush, *et al.*, “Quantum supremacy using a programmable superconducting processor,” *Nature*, vol. 574, no. 7779, pp. 505–510, Oct. 2019, ISSN: 1476-4687. DOI: 10.1038/s41586-019-1666-5.
- [5] B. Bauer, S. Bravyi, M. Motta, and G. K.-L. Chan, “Quantum Algorithms for Quantum Chemistry and Quantum Materials Science,” *Chemical Reviews*, vol. 120, no. 22, pp. 12 685–12 717, Oct. 2020, ISSN: 1520-6890. DOI: 10.1021/acs.chemrev.9b00829.
- [6] A. M. Dalzell, S. McArdle, M. Berta, *et al.*, *Quantum Algorithms: A Survey of Applications and End-to-end Complexities*. Cambridge University Press, Apr. 2025, ISBN: 9781009639668. DOI: 10.1017/9781009639651.
- [7] S. P. Jordan, *Quantum Algorithm Zoo*, <https://quantumalgorithmzoo.org/>, Accessed: 2026-07-17, 2025.
- [8] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, 10th Anniversary. Cambridge University Press, 2010, ISBN: 978-11-07002-17-3.

-
- [9] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien, "Quantum computers," *Nature*, vol. 464, no. 7285, pp. 45–53, Mar. 2010, ISSN: 1476-4687. DOI: 10.1038/nature08812.
 - [10] B. Lekitsch, S. Weidt, A. G. Fowler, *et al.*, "Blueprint for a microwave trapped ion quantum computer," *Science Advances*, vol. 3, no. 2, Feb. 2017, ISSN: 2375-2548. DOI: 10.1126/sciadv.1601540.
 - [11] Z. D. Romaszko, S. Hong, M. Siegle, *et al.*, "Engineering of microfabricated ion traps and integration of advanced on-chip features," *Nature Reviews Physics*, vol. 2, no. 6, pp. 285–299, May 2020, ISSN: 2522-5820. DOI: 10.1038/s42254-020-0182-8.
 - [12] L. Henriët, L. Beguin, A. Signoles, *et al.*, "Quantum computing with neutral atoms," *Quantum*, vol. 4, p. 327, Sep. 2020, ISSN: 2521-327X. DOI: 10.22331/q-2020-09-21-327.
 - [13] H. Levine, A. Keesling, G. Semeghini, *et al.*, "Parallel Implementation of High-Fidelity Multiqubit Gates with Neutral Atoms," *Physical Review Letters*, vol. 123, no. 17, Oct. 2019, ISSN: 1079-7114. DOI: 10.1103/physrevlett.123.170503.
 - [14] M. Saffman, T. G. Walker, and K. Mølmer, "Quantum information with Rydberg atoms," *Reviews of Modern Physics*, vol. 82, no. 3, pp. 2313–2363, Aug. 2010, ISSN: 1539-0756. DOI: 10.1103/revmodphys.82.2313.
 - [15] M. Saffman, "Quantum computing with atomic qubits and Rydberg interactions: progress and challenges," *Journal of Physics B: Atomic, Molecular and Optical Physics*, vol. 49, no. 20, p. 202001, Oct. 2016, ISSN: 1361-6455. DOI: 10.1088/0953-4075/49/20/202001.
 - [16] P. Ildefonso, A. Byun, A. Konovalov, J. Kazemi, M. Schuler, and W. Lechner, "Expanding the neutral-atom gate set: Native iSWAP and exchange gates from dipolar Rydberg interactions," *Physical Review Research*, vol. 8, no. 3, Sep. 2026, ISSN: 2643-1564. DOI: 10.1103/z46h-mgh4.
 - [17] N. Dumoulin Stuyck, A. Saraiva, W. Gilbert, *et al.*, "CMOS compatibility of semiconductor spin qubits," *Nature Reviews Electrical Engineering*, vol. 3, no. 5, pp. 300–315, Apr. 2026, ISSN: 2948-1201. DOI: 10.1038/s44287-026-00283-w.
 - [18] J. H. Wesenberg, K. Mølmer, L. Rippe, and S. Kröll, "Scalable designs for quantum computing with rare-earth-ion-doped crystals," *Physical Review A*, vol. 75, no. 1, Jan. 2007, ISSN: 1094-1622. DOI: 10.1103/physreva.75.012304.
 - [19] M. Kjaergaard, M. E. Schwartz, J. Braumüller, *et al.*, "Superconducting Qubits: Current State of Play," *Annual Review of Condensed Matter Physics*, vol. 11, no. 1, pp. 369–395, Mar. 2020, ISSN: 1947-5462. DOI: 10.1146/annurev-conmatphys-031119-050605.

-
- [20] Z. Chen, “Metrology of Quantum Control and Measurement in Superconducting Qubits,” PhD Dissertation, University of California, Santa Barbara, Jan. 2018.
- [21] M. Ganzhorn, G. Salis, D. J. Egger, *et al.*, “Benchmarking the noise sensitivity of different parametric two-qubit gates in a single superconducting quantum computing platform,” *Physical Review Research*, vol. 2, no. 3, Sep. 2020. DOI: 10.1103/physrevresearch.2.033447.
- [22] IBM Quantum, *Qiskit: An open-source SDK for quantum computing*, Accessed: 2026-03-20, 2026.
- [23] A. Noguchi, A. Osada, S. Masuda, *et al.*, “Fast parametric two-qubit gates with suppressed residual interaction using the second-order nonlinearity of a cubic transmon,” *Physical Review A*, vol. 102, no. 6, Dec. 2020, ISSN: 2469-9934. DOI: 10.1103/physreva.102.062408.
- [24] B. Foxen, C. Neill, A. Dunsworth, *et al.*, “Demonstrating a Continuous Set of Two-qubit Gates for Near-term Quantum Algorithms,” *Physical Review Letters*, vol. 125, no. 12, Sep. 2020, ISSN: 1079-7114. DOI: 10.1103/physrevlett.125.120504.
- [25] C. Križan, “Instrument and measurement automation for classical control of a multi-qubit quantum processor,” Master’s thesis, Chalmers University of Technology, SE-412 96 Gothenburg, Sweden, 2019, ch. 4, pp. 51–70.
- [26] J. Heinsoo, C. K. Andersen, A. Remm, *et al.*, “Rapid High-fidelity Multiplexed Readout of Superconducting Qubits,” *Physical Review Applied*, vol. 10, no. 3, Sep. 2018, ISSN: 2331-7019. DOI: 10.1103/physrevapplied.10.034040.
- [27] 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*, vol. 6, no. 6, Dec. 2016. DOI: 10.1103/physrevapplied.6.064007.
- [28] A. Bengtsson, P. Vikstål, C. Warren, *et al.*, “Improved Success Probability with Greater Circuit Depth for the Quantum Approximate Optimization Algorithm,” *Physical Review Applied*, vol. 14, no. 3, Sep. 2020, ISSN: 2331-7019. DOI: 10.1103/physrevapplied.14.034010.
- [29] P. Das, R. B. de Ouboter, and K. W. Taconis, “A Realization of a London-Clarke-Mendoza Type Refrigerator,” in *Low Temperature Physics LT9*. Springer US, 1965, pp. 1253–1255, ISBN: 9781489964434. DOI: 10.1007/978-1-4899-6443-4_133.
- [30] A. Blais, A. L. Grimsmo, S. M. Girvin, and A. Wallraff, “Circuit quantum electrodynamics,” *Reviews of Modern Physics*, vol. 93, no. 2, May 2021. DOI: 10.1103/revmodphys.93.025005.

-
- [31] S. Krinner, S. Storz, P. Kurpiers, *et al.*, “Engineering cryogenic setups for 100-qubit scale superconducting circuit systems,” *EPJ Quantum Technology*, vol. 6, no. 1, May 2019. DOI: 10.1140/epjqt/s40507-019-0072-0.
- [32] S. Kosen, H. Li, M. Rommel, *et al.*, “Building blocks of a flip-chip integrated superconducting quantum processor,” *Quantum Science and Technology*, vol. 7, no. 3, p. 035 018, Jun. 2022. DOI: 10.1088/2058-9565/ac734b.
- [33] S. Kosen, H. Li, M. Rommel, *et al.*, “Signal Crosstalk in a Flip-Chip Quantum Processor,” *PRX Quantum*, vol. 5, no. 3, Sep. 2024. DOI: 10.1103/prxquantum.5.030350.
- [34] J. Koch, T. M. Yu, J. Gambetta, *et al.*, “Charge-insensitive qubit design derived from the Cooper pair box,” *Physical Review A*, vol. 76, no. 4, Oct. 2007. DOI: 10.1103/physreva.76.042319.
- [35] R. Barends, J. Kelly, A. Megrant, *et al.*, “Coherent Josephson Qubit Suitable for Scalable Quantum Integrated Circuits,” *Physical Review Letters*, vol. 111, no. 8, Aug. 2013, ISSN: 1079-7114. DOI: 10.1103/physrevlett.111.080502.
- [36] B. D. Josephson, “Possible new effects in superconductive tunnelling,” *Physics Letters*, vol. 1, no. 7, pp. 251–253, Jul. 1962. DOI: 10.1016/0031-9163(62)91369-0.
- [37] A. Mamgain, S. S. Khaire, U. Singhal, *et al.*, “A Review of Developments in Superconducting Quantum Processors,” *Journal of the Indian Institute of Science*, vol. 103, no. 2, p. 653, Oct. 2022, Figure 18(d), ISSN: 0019-4964. DOI: 10.1007/s41745-022-00330-z.
- [38] A. Osman, J. Simon, A. Bengtsson, *et al.*, “Simplified Josephson-junction fabrication process for reproducibly high-performance superconducting qubits,” *Applied Physics Letters*, vol. 118, no. 6, Feb. 2021. DOI: 10.1063/5.0037093.
- [39] A. A. Pishchimova, N. S. Smirnov, D. A. Ezenkova, *et al.*, “Improving Josephson junction reproducibility for superconducting quantum circuits: Junction area fluctuation,” *Scientific Reports*, vol. 13, no. 1, Apr. 2023. DOI: 10.1038/s41598-023-34051-9.
- [40] A. Osman, J. Fernández-Pendás, C. Warren, *et al.*, “Mitigation of frequency collisions in superconducting quantum processors,” *Physical Review Research*, vol. 5, no. 4, Oct. 2023. DOI: 10.1103/physrevresearch.5.043001.
- [41] A. B. Pavolotsky, D. Dochev, and V. Belitsky, “Aging- and annealing-induced variations in Nb/Al-AlO_x/Nb tunnel junction properties,” *Journal of Applied Physics*, vol. 109, no. 2, Jan. 2011. DOI: 10.1063/1.3532040.
- [42] T. Lehnert, D. Billon, C. Grassl, and K. H. Gundlach, “Thermal annealing properties of Nb-Al/AlO_x-Nb tunnel junctions,” *Journal of Applied Physics*, vol. 72, no. 7, pp. 3165–3168, Oct. 1992. DOI: 10.1063/1.351479.
- [43] C. Granata, A. Vettoliere, L. Petti, *et al.*, “Localized laser trimming of critical current in niobium based Josephson devices,” *Applied Physics Letters*, vol. 90, no. 23, Jun. 2007. DOI: 10.1063/1.2746060.

-
- [44] S. van der Meer, “Improving Frequency Targeting of Transmon Qubits Through Automated Laser Annealing,” Accessed: 2025-07-11, Master’s thesis, Delft University of Technology, Oct. 2021.
- [45] J. B. Hertzberg, J. S. Orcutt, H. Paik, S. Rosenblatt, and M. O. Sandberg, *Frequency tuning of multi-qubit systems*, US Patent US12414482B2, Sep. 2025.
- [46] J. B. Hertzberg, E. J. Zhang, S. Rosenblatt, *et al.*, “Laser-annealing Josephson junctions for yielding scaled-up superconducting quantum processors,” *npj Quantum Information*, vol. 7, no. 1, Aug. 2021. DOI: 10.1038/s41534-021-00464-5.
- [47] E. J. Zhang, S. Srinivasan, N. Sundaresan, *et al.*, “High-performance superconducting quantum processors via laser annealing of transmon qubits,” *Science Advances*, vol. 8, no. 19, May 2022. DOI: 10.1126/sciadv.abi6690.
- [48] T. R. Groves, “Theory of beam-induced substrate heating,” *Journal of Vacuum Science & Technology B: Microelectronics and Nanometer Structures Processing, Measurement, and Phenomena*, vol. 14, no. 6, pp. 3839–3844, Nov. 1996. DOI: 10.1116/1.588678.
- [49] Y. Balaji, N. Acharya, R. Armstrong, *et al.*, *Electron-beam annealing of Josephson junctions for frequency tuning of quantum processors*, 2024. DOI: 10.48550/ARXIV.2402.17395.
- [50] H. Kim, C. Jünger, A. Morvan, *et al.*, “Effects of laser-annealing on fixed-frequency superconducting qubits,” *Applied Physics Letters*, vol. 121, no. 14, Oct. 2022, ISSN: 1077-3118. DOI: 10.1063/5.0102092.
- [51] D. Pappas, M. Field, E. A. Sete, J. Y. Mutus, and X. Wang, *Voltage-Assisted Annealing to Alter Tunnel Junction Properties*, US Patent US20250204275A1, Figures 5b, 8 and 9., Jun. 2025.
- [52] X. Wang, J. Howard, E. A. Sete, *et al.*, “Precision Frequency Tuning of Tunable Transmon Qubits Using Alternating-Bias Assisted Annealing,” in *2024 IEEE International Conference on Quantum Computing and Engineering (QCE)*, vol. 01, IEEE, Sep. 2024, pp. 1315–1323. DOI: 10.1109/qce60285.2024.00157.
- [53] University of Chicago, *Technologies for tuning superconducting Josephson junctions*, US Patent US20220140221A1, May 2022.
- [54] 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*, vol. 6, no. 2, Jun. 2019. DOI: 10.1063/1.5089550.
- [55] M. O. Tholén, R. Borgani, G. R. Di Carlo, *et al.*, “Measurement and control of a superconducting quantum processor with a fully integrated radio-frequency system on a chip,” *Review of Scientific Instruments*, vol. 93, no. 10, Oct. 2022. DOI: 10.1063/5.0101398.

-
- [56] F. Leymann and J. Barzen, “The bitter truth about gate-based quantum algorithms in the NISQ era,” *Quantum Science and Technology*, vol. 5, no. 4, p. 044 007, Sep. 2020. DOI: 10.1088/2058-9565/abae7d.
- [57] D. M. Abrams, N. Didier, B. R. Johnson, M. P. da Silva, and C. A. Ryan, “Implementation of XY entangling gates with a single calibrated pulse,” *Nature Electronics*, vol. 3, no. 12, pp. 744–750, Dec. 2020.
- [58] M. McEwen, D. Bacon, and C. Gidney, “Relaxing Hardware Requirements for Surface Code Circuits using Time-dynamics,” *Quantum*, vol. 7, p. 1172, Nov. 2023. DOI: 10.22331/q-2023-11-07-1172.
- [59] G. Yan, W. Wu, Y. Chen, *et al.*, *Quantum Circuit Synthesis and Compilation Optimization: Overview and Prospects*, 2024. DOI: 10.48550/ARXIV.2407.00736.
- [60] J. K. Iverson and J. Preskill, “Coherence in logical quantum channels,” *New Journal of Physics*, vol. 22, no. 7, p. 073 066, Jul. 2020, ISSN: 1367-2630. DOI: 10.1088/1367-2630/ab8e5c.
- [61] T. Abad, J. Fernández-Pendás, A. Frisk Kockum, and G. Johansson, “Universal Fidelity Reduction of Quantum Operations from Weak Dissipation,” *Physical Review Letters*, vol. 129, no. 15, Oct. 2022. DOI: 10.1103/physrevlett.129.150504.
- [62] T. Abad, Y. Schattner, A. F. Kockum, and G. Johansson, “Impact of decoherence on the fidelity of quantum gates leaving the computational subspace,” *Quantum*, vol. 9, p. 1684, Apr. 2025. DOI: 10.22331/q-2025-04-03-1684.
- [63] A. Y. Kitaev, A. H. Shen, and M. N. Vyalyi, “*Classical and quantum computation*”. *MTsNMO-CheRo*, ch. 8, 1999, (Transl.: in L. J. Senechal, *Classical and Quantum Computation*. Providence, RI, USA: American Mathematical Society, 2002, ch. 8, pp. 65-82).
- [64] A. Barenco, C. H. Bennett, R. Cleve, *et al.*, “Elementary gates for quantum computation,” *Physical Review A*, vol. 52, no. 5, pp. 3457–3467, Nov. 1995. DOI: 10.1103/physreva.52.3457.
- [65] G. E. Crooks, “*Gates, States, and Circuits*,” Normal Computing Corporation, 418 Broadway, Albany, NY 12207, USA, White paper, Tech. Note 014 v0.11.0 Beta, Mar. 2024.
- [66] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, “Geometric theory of nonlocal two-qubit operations,” *Physical Review A*, vol. 67, no. 4, Apr. 2003. DOI: 10.1103/physreva.67.042313.
- [67] Y. Sung, L. Ding, J. Braumüller, *et al.*, “Realization of High-Fidelity CZ and ZZ-Free iSWAP Gates with a Tunable Coupler,” *Physical Review X*, vol. 11, no. 2, Jun. 2021. DOI: 10.1103/physrevx.11.021058.
- [68] L. Lao, A. Korotkov, Z. Jiang, W. Mruczkiewicz, T. E. O’Brien, and D. E. Browne, “Software mitigation of coherent two-qubit gate errors,” *Quantum Science and Technology*, vol. 7, no. 2, p. 025 021, Mar. 2022.

-
- [69] C. Huang, T. Wang, F. Wu, *et al.*, “Quantum Instruction Set Design for Performance,” *Physical Review Letters*, vol. 130, no. 7, Feb. 2023. DOI: 10.1103/physrevlett.130.070601.
- [70] X. Jin, Z. Parrott, K. Cicak, *et al.*, “Superconducting architecture demonstrating fast, tunable high-fidelity CZ gates with parametric control of ZZ coupling,” *Physical Review Applied*, vol. 24, no. 6, Dec. 2025. DOI: 10.1103/kmls-1gp5.
- [71] T. Roy, Z. Li, E. Kapit, and D. I. Schuster, “Two-Qutrit Quantum Algorithms on a Programmable Superconducting Processor,” *Physical Review Applied*, vol. 19, no. 6, Jun. 2023. DOI: 10.1103/physrevapplied.19.064024.
- [72] 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*, vol. 96, no. 6, Dec. 2017. DOI: 10.1103/physreva.96.062323.
- [73] Y. Xu, J. Chu, J. Yuan, *et al.*, “High-Fidelity, High-Scalability Two-Qubit Gate Scheme for Superconducting Qubits,” *Physical Review Letters*, vol. 125, no. 24, Dec. 2020. DOI: 10.1103/physrevlett.125.240503.
- [74] J. Stehlik, D. M. Zajac, D. L. Underwood, *et al.*, “Tunable Coupling Architecture for Fixed-Frequency Transmon Superconducting Qubits,” *Physical Review Letters*, vol. 127, no. 8, Aug. 2021. DOI: 10.1103/physrevlett.127.080505.
- [75] J.-X. Han, J. Zhang, G.-M. Xue, H. Yu, and G. Long, “Protecting logical qubits with dynamical decoupling,” *Physical Review Applied*, vol. 24, no. 2, Aug. 2025. DOI: 10.1103/lm8x-r48q.
- [76] A. T. Rezakhani, “Characterization of two-qubit perfect entanglers,” *Physical Review A*, vol. 70, no. 5, Nov. 2004. DOI: 10.1103/physreva.70.052313.
- [77] M. Lukac, M. Perkowski, and M. Kameyama, “Evolutionary quantum logic synthesis of Boolean reversible logic circuits embedded in ternary quantum space using structural restrictions,” in *IEEE Congress on Evolutionary Computation*, IEEE, Jul. 2010, p. 6. DOI: 10.1109/cec.2010.5585969.
- [78] N. Sangouard, X. Lacour, S. Guérin, and H. R. Jauslin, “Fast SWAP gate by adiabatic passage,” *Physical Review A*, vol. 72, no. 6, Dec. 2005. DOI: 10.1103/physreva.72.062309.
- [79] D. Venturelli, M. Do, E. Rieffel, and J. Frank, “Compiling quantum circuits to realistic hardware architectures using temporal planners,” *Quantum Science and Technology*, vol. 3, no. 2, p. 025004, Feb. 2018. DOI: 10.1088/2058-9565/aaa331.
- [80] A. Mondal and K. K. Parhi, “Quantum Circuits for Stabilizer Error Correcting Codes: a Tutorial,” *IEEE Circuits Syst. Mag.*, vol. 24, no. 1, pp. 33–51, 2024. DOI: 10.1109/mcas.2024.3349668.

-
- [81] R. C. Jaklevic, J. Lambe, A. H. Silver, and J. E. Mercereau, “Quantum interference effects in Josephson tunneling,” *Physical Review Letters*, vol. 12, no. 7, pp. 159–160, Feb. 1964. DOI: 10.1103/physrevlett.12.159.
- [82] J. Clarke and A. I. Braginski, “SQUID Theory,” in *The SQUID Handbook*. Wiley-VCH Verlag GmbH & Co. KGaA, 2004, ch. 2.1, pp. 29–87, ISBN: 978-35-27402-29-8.
- [83] C. Gerry and P. Knight, “Field quantization,” in *Introductory Quantum Optics*. Cambridge University Press, 2004, ch. 2, pp. 10–42.
- [84] A. O. Niskanen, Y. Nakamura, and J.-S. Tsai, “Tunable coupling scheme for flux qubits at the optimal point,” *Physical Review B*, vol. 73, no. 9, Mar. 2006. DOI: 10.1103/physrevb.73.094506.
- [85] M. Reagor, C. B. Osborn, N. Tezak, *et al.*, “Demonstration of universal parametric entangling gates on a multi-qubit lattice,” *Science Advances*, vol. 4, no. 2, Feb. 2018. DOI: 10.1126/sciadv.aao3603.
- [86] K. Kubo and H. Goto, “Fast parametric two-qubit gate for highly detuned fixed-frequency superconducting qubits using a double-transmon coupler,” *Applied Physics Letters*, vol. 122, no. 6, Feb. 2023. DOI: 10.1063/5.0138699.
- [87] M. Roth, N. Moll, G. Salis, *et al.*, “Adiabatic quantum simulations with driven superconducting qubits,” *Physical Review A*, vol. 99, no. 2, Feb. 2019. DOI: 10.1103/physreva.99.022323.
- [88] S. Poletto, J. M. Gambetta, S. T. Merkel, *et al.*, “Entanglement of Two Superconducting Qubits in a Waveguide Cavity via Monochromatic Two-Photon Excitation,” *Physical Review Letters*, vol. 109, no. 24, Dec. 2012. DOI: 10.1103/physrevlett.109.240505.
- [89] K. X. Wei, I. Lauer, E. Pritchett, W. Shanks, D. C. McKay, and A. Javadi-Abhari, “Native Two-Qubit Gates in Fixed-Coupling, Fixed-Frequency Transmons Beyond Cross-Resonance Interaction,” *PRX Quantum*, vol. 5, no. 2, May 2024. DOI: 10.1103/prxquantum.5.020338.
- [90] J. Zhang, J. Vala, S. Sastry, and K. B. Whaley, “Minimum Construction of Two-Qubit Quantum Operations,” *Physical Review Letters*, vol. 93, no. 2, Jul. 2004. DOI: 10.1103/physrevlett.93.020502.
- [91] G. Fazio, J. He, and M. G. Paris, “Cartan quantum metrology,” *Physics Open*, vol. 23, p. 100260, May 2025, ISSN: 2666-0326. DOI: 10.1016/j.physo.2025.100260.
- [92] C. H. Baldwin, A. Kalev, and I. H. Deutsch, “Quantum process tomography of unitary and near-unitary maps,” *Physical Review A*, vol. 90, no. 1, Jul. 2014. DOI: 10.1103/physreva.90.012110.
- [93] L. Chen, H. Li, Y. Lu, *et al.*, “Transmon qubit readout fidelity at the threshold for quantum error correction without a quantum-limited amplifier,” *npj Quantum Information*, vol. 9, no. 1, Mar. 2023. DOI: 10.1038/s41534-023-00689-6.

-
- [94] F. Motzoi, J. M. Gambetta, P. Rebentrost, and F. K. Wilhelm, “Simple Pulses for Elimination of Leakage in Weakly Nonlinear Qubits,” *Physical Review Letters*, vol. 103, no. 11, Sep. 2009. DOI: 10.1103/physrevlett.103.110501.
- [95] F. E. Szabo, “Stochastic matrix,” in *The Linear Algebra Survival Guide*, F. E. Szabo, Ed., Boston: Academic Press, 2015, p. 363, ISBN: 978-01-24095-20-5. DOI: <https://doi.org/10.1016/B978-0-12-409520-5.50026-6>.
- [96] C. Baumgarten, M. Brett, E. Burovski, *et al.*, `scipy.optimize.lsq_linear`, SciPy 1.10.1, The SciPy community, The SciPy community, Feb. 2023.
- [97] M. A. Branch, T. F. Coleman, and Y. Li, “A Subspace, Interior, and Conjugate Gradient Method for Large-Scale Bound-Constrained Minimization Problems,” *SIAM Journal on Scientific Computing*, vol. 21, no. 1, pp. 1–23, Jan. 1999. DOI: 10.1137/S1064827595289108.
- [98] A. Bengtsson, “Quantum information processing with tunable and low-loss superconducting circuits,” ISBN 978-91-79052-53-9, Ph.D. dissertation, Chalmers University of Technology, Department of Microtechnology and Nanoscience, Quantum Technology laboratory, Apr. 2020.
- [99] H. H. Ku, “Notes on the use of propagation of error formulas,” *Journal of Research of the National Bureau of Standards, Section C: Engineering and Instrumentation*, vol. 70C, no. 4, p. 264, Oct. 1966. DOI: 10.6028/jres.070c.025.
- [100] U. Singhal, S. Kalipatnapu, P. K. Gautam, *et al.*, “SQ-CARS: A Scalable Quantum Control and Readout System,” *IEEE Transactions on Instrumentation and Measurement*, vol. 72, pp. 1–15, 2023, ISSN: 1557-9662. DOI: 10.1109/tim.2023.3305656.
- [101] L. Stefanazzi, K. Treptow, N. Wilcer, *et al.*, “The QUICK (Quantum Instrumentation Control Kit): Readout and control for qubits and detectors,” *Review of Scientific Instruments*, vol. 93, no. 4, Apr. 2022, ISSN: 1089-7623. DOI: 10.1063/5.0076249.
- [102] S. W. Jolin, R. Borgani, M. O. Tholén, D. Forchheimer, and D. B. Haviland, “Calibration of mixer amplitude and phase imbalance in superconducting circuits,” *Review of Scientific Instruments*, vol. 91, no. 12, Dec. 2020, ISSN: 1089-7623. DOI: 10.1063/5.0025836.
- [103] D. Riste, S. Fallek, B. Donovan, and T. A. Ohki, “Microwave Techniques for Quantum Computers: State-of-the-Art Control Systems for Quantum Processors,” *IEEE Microwave Magazine*, vol. 21, no. 8, pp. 60–71, Aug. 2020, ISSN: 1557-9581. DOI: 10.1109/mm.2020.2993477.
- [104] S. Asaad, C. Dickel, N. K. Langford, *et al.*, “Independent, extensible control of same-frequency superconducting qubits by selective broadcasting,” *npj Quantum Information*, vol. 2, no. 1, Aug. 2016, ISSN: 2056-6387. DOI: 10.1038/npjqi.2016.29.

-
- [105] R. McDermott, M. G. Vavilov, B. L. T. Plourde, *et al.*, “Quantum–classical interface based on single flux quantum digital logic,” *Quantum Science and Technology*, vol. 3, no. 2, p. 024004, Jan. 2018, ISSN: 2058-9565. DOI: 10.1088/2058-9565/aaa3a0.
 - [106] *Zynq UltraScale+ RFSoc ZCU208 and ZCU216 RF Data Converter Evaluation Tool User Guide (UG1433)*, version 1.4, Released 2024-09-26, Advanced Micro Devices, Sep. 2024.
 - [107] R. Gebauer, N. Karcher, D. Gusenkova, *et al.*, “State preparation of a fluxonium qubit with feedback from a custom FPGA-based platform,” in *Fifth International Conference on Quantum Technologies (ICQT-2019)*, AIP Publishing, 2020. DOI: 10.1063/5.0011721.
 - [108] K. H. Park, Y. S. Yap, Y. P. Tan, *et al.*, “ICARUS-Q: Integrated control and readout unit for scalable quantum processors,” *Review of Scientific Instruments*, vol. 93, no. 10, Oct. 2022, ISSN: 1089-7623. DOI: 10.1063/5.0081232.
 - [109] C. Križan, *Indecorum*, version Indev 2026-04-15, 2026.
 - [110] J. Biznárová, A. Osman, E. Rehnman, *et al.*, “Mitigation of interfacial dielectric loss in aluminum-on-silicon superconducting qubits,” *npj Quantum Information*, vol. 10, no. 1, Aug. 2024, ISSN: 2056-6387. DOI: 10.1038/s41534-024-00868-z.
 - [111] J. Biznárová, J. Rivera Hernández, D. Forchheimer, J. Bylander, D. B. Haviland, and G. Andersson, “Intermodulation spectroscopy and the nonlinear response of two-level systems in superconducting coplanar-waveguide resonators,” *Physical Review Applied*, vol. 22, no. 1, Jul. 2024, ISSN: 2331-7019. DOI: 10.1103/physrevapplied.22.014063.
 - [112] C. W. Warren, J. Fernández-Pendás, S. Ahmed, *et al.*, “Extensive characterization and implementation of a family of three-qubit gates at the coherence limit,” *npj Quantum Information*, vol. 9, no. 1, May 2023. DOI: 10.1038/s41534-023-00711-x.
 - [113] H. Ball, W. D. Oliver, and M. J. Biercuk, “The role of master clock stability in quantum information processing,” *npj Quantum Information*, vol. 2, no. 1, Nov. 2016, ISSN: 2056-6387. DOI: 10.1038/npjqi.2016.33.
 - [114] J. P. G. van Dijk, B. Patra, S. Pellerano, E. Charbon, F. Sebastiano, and M. Babaie, “Designing a DDS-Based SoC for High-Fidelity Multi-Qubit Control,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 67, no. 12, pp. 5380–5393, Dec. 2020, ISSN: 1558-0806. DOI: 10.1109/tcsi.2020.3019413.
 - [115] *Cluster*, Online: Accessed 2026-06-01, Qblox B.V., Delftechpark 22, 2628 XH, Delft, NL, Jun. 2026.
 - [116] *ZQCS Quantum Control System*, Online: Accessed 2026-06-01, Zurich Instruments AG, Technoparkstrasse 1, 8005 Zürich, CH, Jun. 2026.

-
- [117] X. Jin, A. Kamal, A. P. Sears, *et al.*, “Thermal and Residual Excited-State Population in a 3D Transmon Qubit,” *Physical Review Letters*, vol. 114, no. 24, Jun. 2015, ISSN: 1079-7114. DOI: 10.1103/physrevlett.114.240501.
 - [118] A. Osman, “Scaling Superconducting Quantum Processors: Coherence, Frequency Targeting and Crosstalk,” Ch. 2.5, Ph.D. dissertation, Department of Microtechnology and Nanoscience (MC2), Gothenburg, Sweden, Nov. 2024, ISBN: 978-91-8103-139-3.
 - [119] J. Rahamim, “Development of a Coaxial Circuit QED Architecture for Quantum Computing,” Trinity term, Ph.D. dissertation, University of Oxford, Pembroke College, 2019.
 - [120] N. Didier, E. A. Sete, M. P. da Silva, and C. Rigetti, “Analytical modeling of parametrically modulated transmon qubits,” *Physical Review A*, vol. 97, no. 2, Feb. 2018. DOI: 10.1103/physreva.97.022330.
 - [121] V. Ambegaokar and A. Baratoff, “Tunneling Between Superconductors,” *Physical Review Letters*, vol. 10, no. 11, pp. 486–489, Jun. 1963, ISSN: 0031-9007. DOI: 10.1103/physrevlett.10.486.
 - [122] V. Ambegaokar and A. Baratoff, “Errata: Tunneling Between Superconductors,” *Physical Review Letters*, vol. 11, no. 2, pp. 104–104, Jul. 1963, Erratum. DOI: 10.1103/physrevlett.11.104.
 - [123] P. Mangin and R. Kahn, “Cooper pairs - principal results of BCS theory,” in *Superconductivity: An introduction*. Cham, Switzerland: Springer International Publishing, 2017, p. 214, ISBN: 978-33-19505-27-5. DOI: 10.1007/978-3-319-50527-5_9.
 - [124] J. Bardeen, L. N. Cooper, and J. R. Schrieffer, “Theory of Superconductivity,” *Physical Review*, vol. 108, no. 5, pp. 1175–1204, Dec. 1957. DOI: 10.1103/physrev.108.1175.
 - [125] H. Rashid, V. Desmaris, A. Pavolotsky, and V. Belitsky, “Harmonic and reactive behavior of the quasiparticle tunnel current in SIS junctions,” *AIP Advances*, vol. 6, no. 4, Apr. 2016, ISSN: 2158-3226. DOI: 10.1063/1.4947133.
 - [126] P. J. Koppinen, L. M. Väistö, and I. J. Maasilta, “Complete stabilization and improvement of the characteristics of tunnel junctions by thermal annealing,” *Applied Physics Letters*, vol. 90, no. 5, Jan. 2007. DOI: 10.1063/1.2437662.
 - [127] X. Wang, M. Field, T. Zhang, *et al.*, *Scaling Alternating-Bias-Assisted Annealing for Precision Transmon Frequency Targeting on Superconducting Quantum Processors*, 2026. DOI: 10.48550/ARXIV.2608.27789.
 - [128] K. Chaves, V. Iaia, D. Rash, Y. Rosen, L. Alegria, and A. Abelson, “Oral: Expanded Parameter Search of Alternating-Bias Assisted Annealing (ABAA),” Presented by Loren Alegria at the APS Global Physics Summit, session MAR-C18: Advanced Fabrication for Superconducting Qubits, Anaheim Convention Center, March 17, 2025, 2025.

-
- [129] L. J. Zeng, S. Nik, T. Greibe, *et al.*, “Direct observation of the thickness distribution of ultra thin AlOx barriers in Al/AlOx/Al Josephson junctions,” *Journal of Physics D: Applied Physics*, vol. 48, no. 39, p. 395 308, Sep. 2015. DOI: 10.1088/0022-3727/48/39/395308.
- [130] J. R. Nesbitt and A. F. Hebard, “Time-dependent glassy behavior of interface states in Al-AlOx-Al tunnel junctions,” *Physical Review B*, vol. 75, no. 19, May 2007. DOI: 10.1103/physrevb.75.195441.
- [131] Z. Jiang and M. H. Ansari, “Concurrent fermionic simulation gate,” *Physical Review Applied*, vol. 24, no. 2, Aug. 2025, ISSN: 2331-7019. DOI: 10.1103/m5z3-jgk8.
- [132] G. M. Nixon, C. K. McLauchlan, and C. C. L. van Rest, *Vine Codes: Low-Overhead Quantum LDPC Codes on a Planar Square Grid*, 2026. DOI: 10.48550/ARXIV.2606.20263.