

Experimental Quantum Simulation with noisy intermediate-scale superconducting processors

Sagastizabal, R.E.

DOI

[10.4233/uuid:f6e98794-be32-4b47-9866-7fd659d128f4](https://doi.org/10.4233/uuid:f6e98794-be32-4b47-9866-7fd659d128f4)

Publication date

2023

Document Version

Final published version

Citation (APA)

Sagastizabal, R. E. (2023). *Experimental Quantum Simulation with noisy intermediate-scale superconducting processors*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:f6e98794-be32-4b47-9866-7fd659d128f4>

Important note

To cite this publication, please use the final published version (if applicable). Please check the document version above.

Copyright

Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights. We will remove access to the work immediately and investigate your claim.



Experimental Quantum Simulation with Noisy Intermediate-Scale Superconducting Processors

RAMIRO SAGASTIZABAL



Experimental Quantum Simulation with Noisy Intermediate-Scale Superconducting Processors

RAMIRO SAGASTIZABAL

Casimir series 2014-2023
ISBN: 978-94-6419-969-7

**EXPERIMENTAL QUANTUM SIMULATION WITH NOISY
INTERMEDIATE-SCALE SUPERCONDUCTING PROCESSORS**

EXPERIMENTAL QUANTUM SIMULATION WITH NOISY INTERMEDIATE-SCALE SUPERCONDUCTING PROCESSORS

Dissertation

for the purpose of obtaining the degree of doctor
at Delft University of Technology
by the authority of the Rector Magnificus prof.dr.ir. T.H.J.J. van der Hagen,
chair of the Board for Doctorates,
to be defended publicly on
Day 6th of November 2023 at 17:30 o'clock

by

Ramiro Ernesto SAGASTIZABAL

Master of Science in Physics,
Université d'Aix-Marseille, Marseille (France),
born in Esperanza, Santa Fe, Argentina.

This dissertation has been approved by the promotor.

Composition of the doctoral committee:

Rector Magnificus	chairperson
Prof.dr. L. DiCarlo	Delft University of Technology, promotor
Prof.dr.ir. L.P. Kouwenhoven	Delft University of Technology, promotor

Independent members:

Prof.dr. G. A. Steele	Delft University of Technology, Netherlands
Dr.Ir. M. Veldhorst	Delft University of Technology, Netherlands
Dr. C.K. Andersen	Delft University of Technology, Netherlands
Dr. J.J. Garcia-Ripoll	Centro Superior de Investigaciones Cientificas, Spain
Prof.dr. Y.M. Blanter	Delft University of Technology, Netherlands, Reserve member



Printed by: Gildeprint, Enschede – www.gildeprint.nl

Cover: Artistic impression of a house of cards, representing the challenging balance pursued in the operation and calibration of quantum processors.
Design by Yamila Santoro.

Copyright © 2023 by R. Sagastizabal

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without prior permission from the copyright owner.

Casimir PhD Series, Delft-Leiden 2020-24

ISBN 978-94-6419-969-7

An electronic version of this dissertation is available at <http://repository.tudelft.nl/>.

To my family,

CONTENTS

Summary / Samenvatting	xi
1 Introduction	1
1.1 Quantum computers: the idea	1
1.2 Transmon qubits	3
1.3 Thesis overview	4
2 Control and signals for algorithm execution	7
2.1 The Transmon qubit	8
2.2 The coupling in the dispersive regime	10
2.3 The dispersive readout	12
2.4 Qubit states and gates.	13
2.5 Single-qubit gates	16
2.6 Two-qubit gates	17
2.7 Tomography	19
2.8 Typical noise mechanisms	20
2.8.1 Relaxation and dephasing times	20
2.8.2 Thermal population	21
2.8.3 Leakage	21
2.8.4 Residual ZZ coupling.	22
2.9 Quantum Algorithms	22
2.9.1 Trotterized Evolution	22
2.9.2 Variational Quantum Eigensolvers	23
2.9.3 Quantum Approximate Optimization Algorithm	24
3 Experimentally simulating the dynamics of quantum light and matter at deep-strong coupling	27
3.1 Introduction	28
3.2 Results	29
3.2.1 Digital quantum Rabi simulator with phase-controlled tuning.	29
3.2.2 Comparing qubit and resonator parity dynamics.	31
3.2.3 Resonator photon-number dynamics	33
3.2.4 Resonator phase-space dynamics	34
3.2.5 Demonstrating qubit-resonator entanglement	34
3.2.6 Quantum Rabi dynamics in the nondegenerate-qubit case	34
3.3 Discussion.	35

3.4	Methods	36
3.4.1	Phase-controlled Trotterisation of the quantum Rabi model	36
3.4.2	Trotter step	38
3.4.3	Qubit control	39
3.4.4	Wigner tomography reconstructions	39
3.5	Supplementary Notes	39
3.5.1	Supplementary Note 1: Experimental setup	39
3.5.2	Supplementary Note 2: Device fabrication	41
3.5.3	Supplementary Note 3: Device operating parameters and qubit performance	41
3.5.4	Supplementary Note 4: Calibration of the flux distortions	43
3.5.5	Supplementary Note 5: Calibration of the photon meters	45
3.5.6	Supplementary Note 6: Calibration of Wigner tomography	47
3.5.7	Supplementary Note 7: Analog vs digital Jaynes-Cummings dynamics	47
3.5.8	Supplementary Note 8: Trotter simulation with excited and ground initial states	49
3.5.9	Supplementary Note 9: Trotterisation performance vs Trotter order	49
3.5.10	Supplementary Note 10: Trotterisation performance vs Trotter step size	50
3.5.11	Supplementary Note 11: Qubit entropy dynamics	50
4	Experimental Error Mitigation via Symmetry Verification in a Variational Quantum Eigensolver	67
4.1	Introduction	68
4.2	The Variational Eigensolver	68
4.3	Mitigation by Symmetry Verification	69
4.4	Results	72
4.5	Physicality of the state reconstructed	74
4.6	Summary and conclusions	77
4.7	Supplementary Notes	77
4.7.1	Supplementary Note 1: Device fabrication	77
4.7.2	Supplementary Note 2: Experimental setup	77
4.7.3	Supplementary Note 3: Measured device parameters	78
4.7.4	Supplementary Note 4: Pulses for gates	80
4.7.5	Supplementary Note 5: Tomographic reconstruction and limitations	82
4.7.6	Supplementary Note 6: Constraining the positivity of reduced density matrices	84
4.7.7	Supplementary Note 7: Theoretical modeling of the experiment	85
5	Variational preparation of finite-temperature states on a quantum computer	87
5.1	Introduction	88
5.2	Results	88
5.2.1	Theory	88
5.2.2	Experiment	90

5.3	Discussion	92
5.4	Methods	93
5.4.1	Quantum Circuit	93
5.4.2	Modeling and Simulations	93
5.4.3	Leakage in Transmons	94
5.5	Supplementary Notes	94
5.5.1	Supplementary Note 1: Optimization of the cost function	94
5.5.2	Supplementary Note 2: Circuit compilation	95
5.5.3	Supplementary Note 3: Device and transmon parameters at bias point	98
5.5.4	Supplementary Note 4: Experimental setup	98
5.5.5	Supplementary Note 5: Gate performance	100
5.5.6	Supplementary Note 6: Residual ZZ coupling at bias point	100
5.5.7	Supplementary Note 7: Measurement models, cost function evaluation, and two-qubit state tomography	101
5.5.8	Measurement models	103
5.5.9	Cost function evaluation	104
5.5.10	Two-qubit state tomography	105
5.5.11	Supplementary Note 8: Impact of leakage on two-qubit tomography	106
5.5.12	Supplementary Note 9: Error model for numerical simulations	107
5.5.13	Supplementary Note 10: Classical optimizer under ideal conditions	109
6	Conclusions and Outlook	111
6.1	Summary	112
6.2	The promise of the variational approach	112
6.3	A personal forecast	114
	Acknowledgements	115
	Curriculum Vitæ	117
	List of Publications	119
	References	125

SUMMARY

Superconducting qubits have seen a tremendous progress in the last two decades, and yet they remain unable to extract quantum advantage for application scenarios. While algorithms like Shor's demonstrate quantum advantage in theory, they do not seem to fit modern noisy quantum processors. Then, in order to achieve quantum advantage either modern quantum processors catch up to their expected behaviour, or modern algorithms are tailored to fit the expected behaviour of modern chips.

This thesis focuses on the second approach, which is to explore the implementation of algorithms on modern quantum processors. Along all these implementations, we study the errors that cause these algorithms to derail from their ideal results. We attempt to understand, quantify and control these errors, in the hope that this provides useful insights into how to design algorithms for the modern hardware.

This thesis starts by introducing the topic of superconducting quantum processors and modern algorithms in the first two chapters. Then we move onto the three experiments, one chapter each, detailing our findings.

The first experiment cover an digital-analog implementation of a quantum simulation of light-matter interaction. We present the implementation that makes use of both digital (gates) and analog (evolution) blocks. The accuracy of the Trotterization technique is studied in detail, as well as the capability to study the photon population in the resonator. We manage to implement up to 90 Trotter steps and reproduce the behaviour in the ultra-strong coupling regime.

The second experiment presents an error mitigation technique, on an application of great interest to the field (molecular simulations). This application is a fully digital one, within the hot topic of variational algorithms for ground-state preparation. The mitigation technique, which is an invention of our own team (see referenced theoretical works) manages to reduce the algorithm error over an order of magnitude. In order to demonstrate this level of control, we quantify the error through accurate simulations of the quantum process and independent quantification of the parameters involved.

The third experiment presents another variational algorithm, this time to produce thermal states rather than ground states. Again, we pursue a detailed study of the many error mechanisms involved, in order to quantify and match the results obtained. We go beyond incoherent errors and add a coherent error mechanism common to our hardware architecture, the residual ZZ coupling.

Finally, we reflect on the final chapters about how to continue towards implementations that make the most out of modern, noisy, hardware.

SAMENVATTING

Supergeleidende qubits hebben de afgelopen twintig jaar een enorme vooruitgang geboekt, en toch zijn ze nog steeds niet in staat om kwantumvoordeel te behalen voor toepassingsscenario's. Hoewel algoritmen zoals die van Shor in theorie een kwantumvoordeel aantonen, lijken ze niet te passen in moderne luidruchtige kwantumprocessors. Om vervolgens kwantumvoordeel te behalen, halen moderne kwantumprocessors hun verwachte gedrag in, of worden moderne algoritmen op maat gemaakt om te passen bij het verwachte gedrag van moderne chips.

Dit proefschrift richt zich op de tweede benadering, namelijk het onderzoeken van de implementatie van algoritmen op moderne kwantumprocessors. Bij al deze implementaties bestuderen we de fouten die ervoor zorgen dat deze algoritmen hun ideale resultaten niet behalen. We proberen deze fouten te begrijpen, kwantificeren en beheersen, in de hoop dat dit nuttige inzichten oplevert in hoe algoritmen voor de moderne hardware kunnen worden ontworpen.

Dit proefschrift begint met de introductie van het onderwerp supergeleidende kwantumprocessors en moderne algoritmen in de eerste twee hoofdstukken. Daarna gaan we verder met de drie experimenten, elk één hoofdstuk, waarin onze bevindingen gedetailleerd worden beschreven.

Het eerste experiment omvat een digitaal-analoge implementatie van een kwantumsimulatie van licht-materie-interactie. We presenteren de implementatie die gebruik maakt van zowel digitale (poorten) als analoge (evolutie) blokken. De nauwkeurigheid van de Trotterization-techniek wordt in detail bestudeerd, evenals de mogelijkheid om de fotonenpopulatie in de resonator te bestuderen. We slagen erin om tot 90 Trotter-stappen te implementeren en het gedrag te reproduceren in het ultrasterke koppelingsregime.

Het tweede experiment presenteert een techniek voor het beperken van fouten, op een toepassing die van groot belang is voor het veld (moleculaire simulaties). Deze applicatie is volledig digitaal en valt binnen het populaire onderwerp van variatie-algoritmen voor de voorbereiding van de grondtoestand. De mitigatietechniek, een uitvinding van ons eigen team (zie theoretische werken waarnaar wordt verwezen), slaagt erin de algoritmefout over een orde van grootte te verminderen. Om dit niveau van controle aan te tonen, kwantificeren we de fout door middel van nauwkeurige simulaties van het kwantumproces en onafhankelijke kwantificering van de betrokken parameters.

Het derde experiment presenteert een ander variatiealgoritme, dit keer om thermische toestanden te produceren in plaats van grondtoestanden. Ook hier voeren we een gedetailleerde studie uit van de vele betrokken foutmechanismen, om de verkregen resultaten te kwantificeren en te matchen. We gaan verder dan incoherente fouten en voegen een coherent foutmechanisme toe dat gebruikelijk is in onze hardwarearchitectuur: de resterende ZZ-koppeling.

Ten slotte reflecteren we op de laatste hoofdstukken over hoe verder te gaan in de richting van implementaties die het meeste uit moderne, luidruchtige hardware halen.

INTRODUCTION

Life is not about waiting for the storm to pass. It is about learning to dance in the rain

Author disputed, commonly attributed to either Mahatma Gandhi or Vivian Greene.

1.1 Quantum computers: the idea

The first mention of quantum computing was made by R. Feynman in a very conceptual fashion [1]. His core idea was to propose that ultimately, to board the complexities of simulating quantum systems, it would be best to have another quantum system which we control. That gave rise to the idea of expanding classical computation with devices that are governed by quantum mechanics.

But what is quantum mechanics and what can it offer to computation devices? The first part of this question is something that appears much clearer than the second. Quantum mechanics is the physical theory modern scientists use to model and understand phenomena like the atomic orbitals, which could not be properly explained with previous theories (we call this *classical physics*). It was introduced in the early XX century by Bohr, Einstein, Schroedinger, Heisenberg and many more physicists [2], nowadays celebrities of the field. The core concept was that of representing particles by a wave function $|\psi(t)\rangle$, instead by vectorial coordinates for space and velocity $\vec{x}(t)$, $\vec{v}(t)$ as in classical mechanics. The Lagrangian and Hamiltonians are extended into operators acting on the wave functions representing the states. Schroedinger's biggest contribution was the equation (which carries his name) under which the Hamiltonian operator $\hat{H}(t)$ dictates the dynamics of a system via differential equations:

$$i\hbar |\dot{\psi}(t)\rangle = \hat{H}(t) |\psi(t)\rangle \quad (1.1)$$

Typically, systems are modelled by writing a Hamiltonian in classical mechanics terms, and then quantizing by what is called first (or second) quantization. However, quantum mechanics brought about the possibility to model degrees of freedom without a classical counterpart: spin systems. Along with such systems, quantum mechanics brings several concepts that do not appear in classical mechanics: superposition, entanglement and collapse of the wave function by measurement. Wave functions can be the sum of several different states,

with complex values as pre-factors, and we call this a superposition state. The system is, until other event changes it, simultaneously in every state of the superposition. This leads to entanglement, which is a quality of states where subparts of the system are correlated in a non-trivial way. For example, a superposition that contains all states where spins A and B point in the same direction is an entangled one. This correlation can give rise to unexpected effects, which puzzled Einstein and other scientists [3], motivating them to question the validity of quantum mechanics. Finally, and maybe the weirdest feature of all: the collapse of the wave function. Measuring involves a target quantity like energy or position, represented by an operator M . The quantum state of the measured system needs to be thought of in the measured operator's eigenbasis. Upon measurement, this state *collapses*, randomly selecting only one of the eigenstates in the superposition and returning as a measured value the eigenvalue associated. The probability of the state collapsing onto each eigenvalue is determined by the pre-factors absolute value squared. Naturally, this imposes a normalization constraint to ensure probabilities add up to 100 %. This is the mechanism that allowed Bohr to explain the energy levels in the atom. The strangeness of this effect is the stochasticity. The outcome of measurement is not deterministic anymore, but stochastic.

The three features above mentioned are the main ones expected to give quantum computers an advantage over classical ones. The first step is using them to extend the concept of bits to a quantum level: qubits. While bits are typically electrical circuits designed to have two different voltage states (0 or 1), qubits are quantum systems which can be in two different quantum states called $|0\rangle$ and $|1\rangle$. As quantum mechanics allows, qubits can be not only in each one of these states, but also in a superposition of both. This introduces a first idea of advantage: the possibility to compute different bit strings simultaneously. Another advantage is that of entanglement, which can produce high correlation between the state of two qubits.

All this train of thought is interesting but lacks concreteness so far. Perhaps the best example of a concrete quantum algorithm is one of the first to come: Shor's algorithm. Published in 1995 by Peter Shor[4], this procedure manages to find the period of periodic functions, allowing to factorize numbers into their prime factors more efficiently than classical algorithms. If quantum computers were to meet its requirements, this algorithm would break most publickey cryptography used nowadays, compromising even the security of online banking.

Although nowadays quantum computers are not big enough to run Shor's algorithm at a useful level, its existence shows that quantum computers can be programmed to produce a quantum advantage, at least in theory. Ever since, many more algorithms have appeared, and this is an active area of research. A fairly complete list of proposed algorithms can be found in [5].

More precise things can be said about the challenges real quantum computers have to implement algorithms. It starts with the fact that concepts and operations that we take for granted in classical bits need considerably more effort to be delivered with qubits. Basic things like initializing a qubit in the state $|0\rangle$, changing a state with an operation and measuring qubits to assign them a state, are operations that have to be benchmarked, demonstrating control over the quantum system implementing the processor. Such was the idea of

Di Vincenzo, who formalized the requirements in 2000, in a set now known as *Di Vincenzo criteria* [6].

This set of requirements sets a blueprint or checklist that any technology attempting to implement quantum processors should demonstrate. Amongst the most promising technologies fulfilling said criteria, the most prominent nowadays are those implementing qubits with superconducting circuits, trapped ions and spin degrees of freedom.

1.2 Transmon qubits

In this thesis we focus on a specific case of superconducting qubits: transmons. Their initial design, named Cooper-Pair Box (CPB) [7], introduced a resonant element, defined by a capacitor and a non-linear superconducting inductor called Josephson Junction [8, 9]. Such system can be designed into several regimes [10], making different uses of the charge and magnetic degrees of freedom to encode, control and readout the state of the qubit with electromagnetic fields. While a big first step, these designs suffered from many loss mechanisms that deteriorate the state of the qubits namely dephasing and energy relaxation. Dephasing occurs due to noise along the several parameters that couple these qubits to the environment and is quantified by the characteristic time at which the loss of phase occurs. Energy relaxation occurs when the qubit spontaneously emits an excitation to the environment. This was further improved by Cavity QED [11, 12], where a resonating cavity regulates the relaxation events via the Purcell effect [13, 14]. Quantum coherence of the system's states is affected by several mechanisms that enforce fluctuations on the system frequency: quasiparticle fluctuations [15], photons in the resonator [16] and $1/f$ flux noise [17].

The transmon qubit [18] is a design that reduces the dephasing impact of charge noise by making the frequency charge insensitive. This is accomplished by making the qubit's Josephson energy, E_J , 30–50 times larger than its charging energy, E_C . Transmons have shown a rapid growth since their introduction in 2007, and progressed fulfilling all the Di Vincenzo criteria [19–22].

Ever since all the DiVincenzo criteria were met by demonstrating a two-qubit gate, the natural continuation to implement computations has been attempted [22]. In spite of a continuous improvement in coherence and relaxation times [23], errors occurring (mainly due to dephasing and relaxation) remain the biggest impediment to scaling these efforts in terms of qubit count. To that extent, there exist several efforts ever since pushing for the implementation of error correction codes to overcome these problems [24–27]. However, these efforts seem several years away from delivering logical protected qubits, capable of running algorithms in a fault-tolerant way. In spite of this immediate impediment, there is no obvious reason why errors could not be circumvented, and several teams focus their efforts into exploring this avenue by implementing algorithms without logical qubits [28–31]. The present thesis details our efforts in that direction.

Within the field of quantum algorithms, the first interesting implementing were quantum simulations. By quantum simulations, one means the reproduction of a quantum system characteristics, via the implementation of certain operations in a controlled quantum processor.

This splits into two big families of quantum simulations: firstly those that *emulate* the target system, by enforcing upon the qubits (or other processor components) a Hamiltonian that mimics the one of the target system; secondly, those that encode the problem in a digital domain using the qubits, in analogy to binary representations in classical computation, and solve the problem by the action of several gates involving one or several qubits. The first group is called *analog quantum simulations* and the second is called *digital quantum simulations*, in clear analogy to analog and digital logic in classical computation.

This thesis has focused in digital quantum simulations, with a special case being that of Chapter 3, where one of our gates is motivated in an analog implementation of the target system. Quantum systems can be solved both in time or frequency domains, as the interesting quantities might be the dynamics or rather the spectrum of the target system. Following that idea, digital quantum simulations can be subdivided into two big groups: those that aim at solving the dynamics of a target quantum system, and those that aim at solving the spectrum. One of the first applications found for quantum computers was that of Trotterization, which allows to digitally implement the evolution under a target Hamiltonian. This remains the most prominent approach to the subgroup of time-domain solutions. Also related, the algorithm called Quantum Phase Estimation uses Trotterized evolution to obtain the target system spectrum from the phase acquired by the evolution of eigenstates

In parallel to the start of this thesis, a new family of algorithms was introduced: *variational algorithms*. Their main feature was to bring complexity by parameterizing the quantum circuits, which were optimized to a target solution afterwards, via a classical optimization routine. Their first applications focused on obtaining ground states, as convergence was easily guaranteed via the Variational Principle of quantum mechanics. This family of algorithms is called *Variational Quantum Eigensolvers* (VQE). However, such paradigm was understood to be more general, as it simply parameterizes a target state and interprets it via some mapping linking measurements to the quantity of interest. Another quite general family of variational algorithms, also used to simulate quantum states, is that of *Quantum Alternating Operator Ansatz* (QAOA). For QAOA algorithms, the circuit is built as repeated layers, alternating the implementation of target operators with certain parameterization. While there are many more algorithms, the ones mentioned here are amongst the most studied in the literature and correspond to those included in the results of this thesis.

1.3 Thesis overview

The epigraph of this chapter is a motivational quote, reflecting that circumstances are non-optimal a lot times, but that one should live and progress disregarding them. Quantum computers were a theorist dream for several decades until the developments of the last twenty years gave us working prototypes. However, there is still, as the proverb says, rain: Modern processors are incapable, in terms of size and errors, of running the dream algorithms of Shor or Grover. My motivation for the works of this thesis has been, as the proverb recommends, to dance under the rain: to attempt the execution of algorithms and face the many challenges

encountered to make the most out of them. The following chapters of this thesis present the various shots at it conducted in our lab at QuTech, TU Delft.

In Chapter 2 (Chapter 2) we present the basic concepts of transmon qubits, their physical description and the several approximations involved. We progress along the various control concepts (readout, single- and two-qubit gates and tomography), and board the noise mechanisms that affect and determine the performance of transmon devices. We finally introduce quantum algorithms and the three specific cases run in the following chapters of the thesis.

In Chapter 3 (Chapter 3) we use a device with two transmons and three resonators to simulate the dynamics of light-matter coupling in a high coupling regime, known as Ultra-Strong Coupling (USC). We present a scheme of Trotterized evolution capable of simulating both the atomic and cavity degrees of freedom. Implementing up to 90 trotter steps, we validate the expected excitation number violations of USC as well as the entangled nature of its eigenstates.

In Chapter 4 (Chapter 4) we conceive and present a technique to mitigate the effect of several noise mechanisms. We implement it on the simulation of the ground state of a H₂ molecule and improve its performance by an order of magnitude.

In Chapter 5 (Chapter 5) we use four transmons in a seven-qubit processor, to produce Gibbs states with a variational approach. In order to do this, we implement a trick that allows the production, with purely unitary control, of mixed states in a controllable fashion, therefore representing finite-temperature physics.

In Chapter 6 (Chapter 6) we briefly discuss the key lessons gathered from the experiments reported and I offer my view on the outlook of digital quantum simulations.

In order to understand the quantum algorithms presented in the coming chapters, one must introduce first the necessary concepts. In this chapter, we introduce the transmon qubit and the regime in which it operates. Then we move on to the control operations (single- and multi-qubit gates as well as readout and state reconstruction). Next we cover briefly the dominating noise sources introducing errors in Transmon operations. Finally we introduce the algorithms that will be studied in the coming chapters.

2.1 The Transmon qubit

The works in this thesis make use of several transmon qubits. The transmon is one of the simplest and most used qubits in quantum computing. It consists of a superconducting capacitor in parallel with a Josephson Junction (JJ), designed within a specific parameter regime that makes it charge insensitive. In the following paragraphs, we introduce the many aspects necessary to understand the next chapters.

The physical implementation of the Transmon qubit consists of superconducting leads defining capacitor islands, and a JJ connecting the superconducting leads [Figure 2.1(a)]. In the devices used in this thesis, the superconducting materials are NbTiN for every structure, except the JJs which are made of Al and Aluminum Oxide. Since the JJ is a non-linear inductor, we can model the Transmon qubit as a non-linear LC resonant circuit, consisting of a capacitance C_Σ in parallel with the non-linear inductor. The capacitance here is $C_\Sigma = C_s + C_{JJ} + C_g$, where C_s is the shunting capacitance, C_{JJ} is the capacitance of the JJ, and C_g is the coupling capacitance of the circuit with the ground plane. We can write a Hamiltonian for the charges (Cooper pairs) in this system as [18]

$$H = 4E_C (\hat{n} - n_g)^2 - E_J \cos \hat{\phi}, \quad (2.1)$$

where $E_C = e^2/2C_\Sigma$ is the charging energy, E_J is the Josephson energy, $\hat{n}$ is the Cooper-pair imbalance between the islands, $\hat{\phi}$ the superconducting phase across the islands and conjugated variable to $\hat{n}$, and finally n_g is an offset *gate* charge that fluctuates according to the environment.

The Transmon is a variation of the Cooper-pair Box (CPB) [7], where parameters E_J and E_C are chosen to minimize the system dependence on n_g . The CPB resonant frequency is strongly dependent on n_g , whose fluctuations contribute to diminish the coherence times ($\sim 0.5 \mu s$ [32]). By choosing a ratio $\frac{E_J}{E_C} \sim 50$, the frequency dependence with n_g diminishes [Figure 2.1(c-d)]. This choice of regime also has an impact on the qubit anharmonicity, which is a key requirement to independently address the energy levels of the structure. Specifically, defining anharmonicity α as the difference between the two first transitions $\alpha = E_{12} - E_{01}$, it scales with $\frac{E_J}{E_C}$ as $\alpha \approx -\left(\frac{E_J}{E_C}\right)^{-1/2}$. Commonly, and in this thesis as well, Transmon qubits are designed to have an anharmonicity $\frac{\alpha}{h}$ of ~ -300 MHz and a ratio $\frac{E_J}{E_C}$ in the range 40 – 50. For this, we target $\frac{E_C}{h}$ in the range 250 – 350 MHz and $\frac{E_J}{h}$ in the range of 8 – 20 GHz.

Finally, transmons can be made tunable in frequency by replacing the JJ for two of them in parallel. Such structure is called a Superconducting QUantum Interference Device (SQUID). The name points to the fact that the superconducting phases in the SQUID loop interfere with the flux threading through it. Placing a current nearby the SQUID loop, we can control with it the magnetic flux threading the SQUID and therefore the effective inductance. With this mechanism we can tune the qubit frequency accordingly. The effective Josephson Energy E_J of a SQUID loop with two junctions is given by:

$$E_J = (E_{J_A} + E_{J_B}) \left\| \cos \left(\frac{\pi \Phi}{\Phi_0} \right) \right\| \sqrt{1 + d^2 \tan^2 \left(\frac{\pi \Phi}{\Phi_0} \right)}, \quad (2.2)$$




Figure 2.1: **The Transmon qubit.** (a) Optical micrograph of a Transmon qubit, defined by a shunting interdigitated capacitor and a SQUID loop (top of the capacitor). The qubit couples capacitively to a CPW resonator on its bottom-right end. A flux-bias line couples inductively to the SQUID loop on the top. All the surfaces are made of NbTiN on top of a Si substrate, except for the JJs and airbridges (stitch-like structures on top of CPW lines), which are made of Al. (b) SEM picture of a similar device centered on the SQUID loop. The JJs are evaporated on top of the square pattern defining the SQUID. The flux-bias line is placed assymmetrically, in order for its current to affect the magnetic flux threading inside the SQUID loop. (c-d) The first three energy levels of a transmon qubit for $\frac{E_C}{\hbar} = 1$ GHz and (c) $\frac{E_J}{\hbar} = 1$ GHz, (d) $\frac{E_J}{\hbar} = 10$ GHz and (e) $\frac{E_J}{\hbar} = 50$ GHz. As E_J/E_C increases, the spectrum dependence on n_g decreases exponentially. The qubit transitions frequencies are set by the difference between these energy levels.

where Φ is the magnetic flux threading the SQUID loop, E_{JA} , E_{JB} are the respective values for each JJ, and the asymmetry $d = \frac{E_{JA} - E_{JB}}{E_{JA} + E_{JB}}$. This dependance takes the effective E_J of our circuit from a maximum of $E_{JA} + E_{JB}$ at $\Phi = 0$ to a minimum of $|E_{JA} - E_{JB}|$ at $\Phi = \Phi_0/2$, where Φ_0 is the flux quantum. This in turns determines the qubit frequency ω_q via the equation:

$$\hbar\omega_q \approx \sqrt{8E_J(\Phi)E_C} - E_C. \quad (2.3)$$

2.2 The coupling in the dispersive regime

In order to readout the state of the qubit, we make use of a harmonic oscillator, modeled with the Hamiltonian

$$H_{reso} = \hbar\omega_r \left(\frac{1}{2} + a^\dagger a \right). \quad (2.4)$$

The operators $a^\dagger$ (a) are the standard creation (annihilation) operators of quantum mechanics and ω_r is the resonant frequency of the circuit. For a harmonic oscillator, its energy levels are equally spaced by $\hbar\omega_r$. The excitation subspace spanned by the creation and annihilation operators is denoted as $|n_r\rangle$, and represents the photons stored in the resonant system. The *number of photons* in the system is modeled with the operator $n_r = a^\dagger a$.

The harmonic oscillator is physically implemented via Coplanar Waveguides (CPW). A CPW consists of a central conductor, surrounded by a ground plane on each side. The dimensions and distance between these planes defines the capacitance per unit-length c of such a waveguide, as well as its inductance per unit length l . These two parameters define the impedance of the circuit, $Z = \sqrt{\frac{l}{c}}$. We make a CPW resonant by enforcing boundary conditions on its edges: A short imposes zero electric field, while an open imposes zero current [33]. While total capacitance C_r of the CPW is defined by its geometry (length), its inductance L_r is defined by both a geometric component L_g and a kinetic component L_k (known as kinetic inductance [34]). Finally, the resonant frequency is given by $\omega_r = \frac{1}{\sqrt{L_r C_r}}$.

These two components (Transmon and CPW resonator) are coupled capacitively. The term in the Hamiltonian can be modeled with the Quantum Rabi Model [35] as

$$H_{\text{coupl.}}^R = \hbar g \left(a + a^\dagger \right) \left(\sigma^+ + \sigma^- \right), \quad (2.5)$$

where $a^\dagger$ and σ^+ are the creation operators for resonator and qubit excitation, and a and σ^- are their respective annihilation counterparts. Following some approximations [18], for the transmon qubit we have $g = 2 \frac{C_g}{C_\Sigma} \sqrt{\frac{\hbar\omega_r}{2C_r}}$. We simplify this model further by using the rotating-wave approximation (RWA) and obtaining

$$H_{\text{coupl.}} = \hbar g \left(\sigma^+ a + \sigma^- a^\dagger \right). \quad (2.6)$$

The approximated Hamiltonian $H_{\text{coupl.}}$ has a symmetry that $H_{\text{coupl.}}^R$ does not: It conserves the number of excitations in the system. Specifically, it only allows creation of a photon via $a^\dagger \sigma^-$ which in turn annihilates a qubit excitation. The opposite happens with the second term of the coupling Hamiltonian. When the coupling is not small compared to the detuning

$\Delta = (\omega_r - \omega_q)$, this approximation is no longer valid, and we call this the *Ultra-strong coupling regime* (USC). Typically, we work with couplings on the order of 50 – 100MHz and detunings $|\Delta|/2\pi > 1\text{GHz}$, which results well within this approximation. However, we explore the USC in further details, via a quantum simulation, in Chapter 3 of this thesis.

Using the the diagonal base $|j\rangle$ for the transmon qubit, we have the Hamiltonian for the full system as

$$H = \hbar \sum_j \omega_j |j\rangle \langle j| + \hbar \omega_r a^\dagger a + \hbar \left[\sum_i g_{i,j} a^\dagger |i\rangle \langle i+1| + \text{h.c.} \right], \quad (2.7)$$

where the coupling scales as $g_{i,j} = g \langle i | \hat{n} | i+1 \rangle$.

By limiting ourselves to the subspace of the two lowest-energy states of the Transmon, we finally obtain a subspace of two levels that we can refer to as qubit. In that subspace, the JC Hamiltonian is written as:

$$H = -\hbar \frac{\omega_q}{2} \sigma_z + \hbar \omega_r a^\dagger a + \hbar g \left(\sigma^- a^\dagger + \sigma^+ a \right). \quad (2.8)$$

For detunings much larger than the coupling, $g \ll |\Delta|$, the dispersive approximation is valid for this Hamiltonian, allowing yet another simplification, obtaining the so-called dispersive model:

$$H = -\hbar \frac{\omega_q}{2} \sigma_z + \hbar \omega_r a^\dagger a + \hbar \chi \sigma^z a^\dagger a, \quad (2.9)$$

where χ is called the *dispersive shift* and it is $\chi \approx \frac{g}{\Delta^2}$ up to first order of $\frac{g}{|\Delta|}$. As the coupling from Equation (2.7) scales with the photon number n , these approximations will break when a high-enough number of photons is present. A critical photon number n_{crit} can be defined [36] as

$$n_{\text{crit}} = \frac{\Delta^2}{4g^2}. \quad (2.10)$$

While this sets an upper limit in the number of photons necessary to break the first-order approximation for the dispersive limit, there still remains another possibility where the resonator inherits enough non-linearity from the qubit. When the resonator frequency is moved by more than a linewidth κ , a second critical photon number can be defined as n_κ following [36]

$$n_\kappa = \frac{\kappa \Delta}{\chi^2}. \quad (2.11)$$

The coupling term can be grouped into either the qubit or the resonator terms, and interpreted as a *dispersive shift in the frequency*, quantified by χ . While grouping χ in to the resonator, we obtain a frequency shift dependent on the qubit state. We make use of this as a readout mechanism, where we associate a change in the resonant frequency of the harmonic oscillator to a change in the qubit state. This is called dispersive readout.

This term can be also interpreted as a qubit frequency shift, dependent on the number of photons in the resonator $\hat{n}_r = a^\dagger a$. This effect is called the *Stark shift* [37]. Conversely, this term can be used to measure the number of photons in the resonator indirectly by measuring the qubit frequency. Such a scheme is proposed and implemented in Chapter 3 of this thesis.

2.3 The dispersive readout

In order to read the qubit state, we take advantage of the dispersive shift on the resonator. A microwave signal is sent and collected at either another or the same input port. The resonance, whose frequency depends on the qubit state affects both the transmission and phase of the collected output signal [Figure 2.2]. The input tone needs to be below both n_{crit} and n_K so that the dispersive Hamiltonian model is valid. The output tone is collected with acquisition hardware, which samples it with a fixed rate, resulting in a read-out trace $S(t)$ of voltage as a function of time.

Each trace contains the information of a single shot at the qubit readout. The readout procedure continues with a recipe to assign a corresponding qubit state $|0\rangle$ or $|1\rangle$ to each traced obtained. We start by integrating the trace as a function of time, with a weight function $w(t)$.

$$V_{\text{int}} = \frac{1}{\tau_{\text{int}}} \int_{t=0}^{\tau_{\text{int}}} S(t) w(t) dt, \quad (2.12)$$

where τ_{int} is the integration duration. The weight function acts as a filter, perhaps the most simple case being $w(t) = \cos(\omega t)$, which only computes signals at frequency ω . Both S and w are phasors, with In-phase and Quadrature components, to be treated as complex numbers, and consequently V_{int} is too. The readout procedure uses V_{int} to assign the state, relying on our capability of discerning between the probability distribution of the integrated signal for each qubit state $V_{\text{int},|0\rangle}$ and $V_{\text{int},|1\rangle}$. There exists an optimal choice $w(t) = S_{|1\rangle}(t) - S_{|0\rangle}(t)$ that makes V_{int} real, and maximizes its ability to discern between $S_{|1\rangle}$ and $S_{|0\rangle}$. Under this choice, the probability distributions for $V_{\text{int},|0\rangle}$ and $V_{\text{int},|1\rangle}$ can be seen in a one-dimensional plot, and have a gaussian profile [Figure 2.2]. The state is assigned by implementing a threshold in voltage, separating the voltage in sides above or below the threshold, corresponding to $|0\rangle$ and $|1\rangle$.

We evaluate the performance of the readout procedure by preparing both $|0\rangle$ and $|1\rangle$, and collecting the statistics. The average assignment fidelity of the procedure is defined as

$$\mathcal{F}_a = 1 - \frac{\epsilon_0 + \epsilon_1}{2}, \quad (2.13)$$

where ϵ_0 and ϵ_1 are the probabilities of mistakenly assigning $|0\rangle$ and $|1\rangle$, respectively. As gaussians have an infinite tail, there is always some overlap between both distributions, causing a minimal missassignment probability. By looking at the cumulative distribution functions (CDF) [Figure 2.2], one can identify the fidelity $F = 2\mathcal{F}_a - 1$ and these probabilities from comparing moving the distributions further apart. One can do this by integrating the signal for a longer period τ_{int} . However, this cannot be done infinitely as other errors occur.

There are three main types of physical effects that complicate the simple picture explained above: thermal excitations, energy relaxation and induced transitions. The first one, thermal excitations, is a direct consequence of the system not being at 0 K. From thermodynamics, one can model the state of such two-level system as a thermal state, with a population on the first eigenstate given by

$$|\langle 1 | \Psi \rangle|^2 = \frac{1}{1 + e^{\hbar\omega_q / T k_B}}, \quad (2.14)$$

where $|\Psi\rangle$ is the qubit state, $|1\rangle$ its first eigenstate, k_B is the Boltzmann constant and T the qubit temperature. The effect of this consideration is directly observable [Figure 2.2] while measuring the gaussian distributions for V_{int} : there is a second small gaussian distribution for $|0\rangle$ centered around $V_{\text{int},|1}\rangle$.

The second physical effect, energy relaxation, occurs spontaneously. The qubit, after being prepared in $|1\rangle$ randomly losses its energy via the emission of a photon to the environment. This is suppressed by the coupling g to the resonator and its linewidth κ , and this is known as the Purcell effect. The modelling of this error mechanism is explained further in section 2.8.1. This effect causes the analog and opposite symptom of thermal excitations: the readout signal for $|1\rangle$ exhibits a small gaussian distribution right below the main $|0\rangle$ gaussian distribution.

Finally, while using photons to probe the qubit states, one might mediate transitions between [38] the states in its multi-level spectrum. We minimize this effect by keeping the readout power as low as possible.

With the above three mechanisms in mind, a trade-off appears: the faster we readout, the less separated the gaussian distributions will be. At the same time, the longer we read out, the more the state $|1\rangle$ will relax into $|0\rangle$. To quantify this, one can model the signal to noise ratio following [39]. Experimentally, we sweep the measurement power P and the integration length τ_{int} to obtain the parameters for best fidelity $\mathcal{F}_a$. As the resonator frequency is slightly dependent on the photon number, a readout frequency sweep also results useful for maximizing $\mathcal{F}_a$.

2.4 Qubit states and gates

Qubits states belong in two-dimensional Hilbert spaces, and after taking into consideration the normalization and global-phase constraints one can write any general qubit state as

$$|\Psi(\theta, \phi)\rangle = \cos \theta |0\rangle + e^{i\phi} \sin \theta |1\rangle. \quad (2.15)$$

This parametrization comes in very handy when trying to visualize the qubit state. The reader might find an immediate analogy with the surface of a unit radius sphere, θ, ϕ being the polar and azimuthal angles, respectively. This sphere analogy is a key concept allowing intuitive understanding and representing qubit states, and is called *Bloch sphere* [Figure 2.3(a)].

Since we have already covered readout of a qubit state, the next necessary tool is operations that change the state of one or several qubits. Those operations are called *gates*, and implemented by turning on specific Hamiltonian terms $\hat{H}_{\text{gate}}$, and allowing for the unitary evolution $\hat{U} = e^{-it\hat{H}_{\text{gate}}/\hbar}$ to implement the desired change on the qubit state.

Another core concept with respect to quantum computing gates is that of a *universal set of gates*. This is: a set of gates out of which we can compile any desired unitary action on the state. There is not only one, but several such sets, for example that made of H , S , T and CNOT gates, known as the *Clifford-T group*. The first three gates change the state of a single qubit, while the latter is a two-qubit gate.

Single-qubit gates can be understood very intuitively from the Bloch sphere as a map between its states. However, not all operations are allowed, for example reflections of states




Figure 2.2: **Readout control** (a) Average assignment fidelity $\mathcal{F}_a$ as a function of readout pulse duration τ for different readout pulse amplitudes A . $\mathcal{F}_a$ increases monotonically as amplitude or pulse improve, until a point of maximal fidelity where $\mathcal{F}_a$ starts worsening with increasing τ (see yellow curve) (b) Average assignment fidelity $\mathcal{F}_a$ as a function of readout pulse frequency for different readout pulse amplitudes A . $\mathcal{F}_a$ is maximal around a specific frequency. (c) Histograms for the signal V_{int} integrated after preparing the qubit in the states $|0\rangle$ and $|1\rangle$. Solid lines correspond to gaussians fit centered at voltage V_0 and V_1 for each state.

against a certain plane are not valid operations. Rigorously, the set of valid operations is the mathematical group $SU(2)$, and has the geometrical interpretation of being all the rotations over the Bloch sphere. This is why, when speaking about single-qubit gates, we will refer to them in terms of the axis and angle of the rotation they describe [Figure 2.3(b-c)]. Mathematically, This can be parameterized as

$$\hat{U}(\theta, \phi, \lambda) = \begin{pmatrix} \cos \theta/2 & -ie^{i\lambda} \sin \theta/2 \\ -ie^{i\phi} \sin \theta/2 & e^{i(\lambda+\phi)} \cos \theta/2 \end{pmatrix} = \hat{Z}_\phi \cdot \hat{X}_\theta \cdot \hat{Z}_\lambda, \quad (2.16)$$

where θ, ϕ, λ are the Euler angles defining the rotation along the Z-X-Z basis. This consideration simplifies further the problem of single-qubit operations: any list of rotations, however long, amounts to the action of just one effective rotation. This is mathematically supported in the fact that $SU(2)$ is a group, and the product of its elements is just another element within the group.

The next step in qubit control are two-qubit gates: operators that act on the state of two qubits. While simply combining simultaneous single-qubit operations is a valid two-qubit operation, it still does not give us full control over two-qubit states. One can show that single-qubit rotations do not change the entanglement of a two qubit state [2]. Therefore, if we want to change the entanglement of the state, we need to make use of other operators. The most commonly found two-qubit gate in textbooks is the *Controlled-NOT* gate or CNOT. This




Figure 2.3: **Bloch sphere.** (a) State representation on the Bloch sphere. The eigenstates for each basis are shown: $|0\rangle, |1\rangle$ (green), $|+\rangle, |-\rangle$ (red) and $|+i\rangle, |-i\rangle$ (blue). (b) Visual example of the states spanned while rotating $R_X(\theta)$ (red) and $R_Z(\phi)$ (blue). A rotation around the X axis (red vector) of an angle θ , starting from $|0\rangle$ spans the red points. A rotation around the Z axis (blue vector) of an angle ϕ , starting from $|+\rangle$ spans the blue points

gate takes one of the qubits as a *control*, and only acts with a NOT gate ($\hat{X}_\pi$) on the second qubit (*target*) for the states where the control qubit is in the state $|1\rangle$. One can represent this operator as

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

where the matrix is defined in the computational basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$. The control qubit here is the most-significant one (left-most), and the target the least-significant one (right-most). However, in superconducting qubits it is more common to perform a similar type of gate called *Controlled-Z* or CZ. It is analogous to the above, but instead of a NOT = $\hat{X}_\pi$ gate, a $\hat{Z}_\pi$ is performed on the target qubit. This gate has an immediate connection to the CNOT, as it is the same controlled rotation, but along the Z axis instead of the X axis. Therefore, one can easily link both gates by a basis change (using an H gate) from X to Z on the control qubit

$$\text{CNOT} = H \text{ CZ } H^\dagger. \quad (2.18)$$

This formula holds if we use $\hat{Y}_{\pi/2}$ instead of H , or any other rotation that takes the Z axis onto the X one. This formula is relevant, as it links the most used two-qubit gate in textbooks and theory with gates that are native to transmon architectures: CZ and $\hat{Y}_{\pi/2}$.

2.5 Single-qubit gates

Transmons support two main types of single-qubit gates: microwave XY gates and flux Z-gates. The first ones are accomplished by applying an electromagnetic pulse in the vicinity of the transmon capacitor. When the electromagnetic field has the right frequency $\hbar\omega_q$, it provides the necessary energy to trigger the transition from $|0\rangle$ to $|1\rangle$. We implement this in pulses in a scheme called DRAG [40]. Specifically, these pulses are shaped in time as a gaussian function, whose area directly controls how much of the $|0\rangle$ state is transferred to the $|1\rangle$ state. The DRAG scheme adds an extra component: The derivative of the gaussian pulse, on the other quadrature of the electromagnetic phasor, scaled by a *DRAG coefficient*. This coefficient can be empirically adjusted to minimize the undesired frequency components of the pulse that introduce phase errors, and additionally trigger transitions on the higher transmons levels. Specifically, it is of key interest to minimize the undesired transition from $|1\rangle$ to $|2\rangle$ known as *leakage*.

DRAG pulses implement a family of unitaries describing rotations of angle θ around an axis at the equator of the Bloch sphere ϕ parameterized by Equation (2.16) as $\hat{U}(\theta, \phi, \lambda = -\phi)$, where θ defines the angular aperture and ϕ the equatorial axis $\hat{a} = \cos\phi\hat{X} + \sin\phi\hat{Y}$ of the rotation. The phase ϕ is physically implemented by controlling the phase of the phasor that describes the microwave pulse.

We tune these gates by choosing a determined gate duration. Specifically, we implement the gaussian up to a range of $\pm 2\sigma$, setting a duration of 4σ . We then calibrate the pulse amplitude necessary for a complete transfer from $|0\rangle$ to $|1\rangle$, which is called a π -pulse amplitude, as it reflects a rotation of $\theta = \pi$ around any equatorial axis. We then scale this amplitude proportionally to obtain the desired θ . As said above, ϕ is controlled with the pulse phase. We also make a sweep of DRAG coefficients to minimize phase error and leakage. For this purpose we use a calibration routine called AIIXY [20].

Transmon qubits also have a mechanism for implementing rotations around the Z axis. By changing the flux threading through the qubit SQUID loop, we can move its frequency for a certain period of time t . The population of the $|1\rangle$ state evolves now in a different dynamic frame, detuned from the reference frame by $\Delta(t)$. The accumulated effect of this along time results in a dynamic phase $\phi = \int_0^t \Delta(t) dt$ for the state $|1\rangle$. This corresponds to a rotation of angular aperture ϕ around the Z axis and is represented by the unitary

$$\hat{U}_Z(\phi) = \begin{pmatrix} e^{i\phi/2} & 0 \\ 0 & e^{-i\phi/2} \end{pmatrix}. \quad (2.19)$$

2.6 Two-qubit gates

Two-qubit gates are implemented by coupling pairs of qubits, and they provide us access to entanglement. Generating entanglement requires the action of a two-body Hamiltonian term, coupling both qubits. There exist several ways of doing this: direct capacitance between qubits [41], mediated via a third component like a resonator [22] or another transmon [42, 43] for instance. In the following chapters we make use of CPW resonators, capacitively coupled




Figure 2.4: **Two Qubit gates.** (a) Spectrum for two transmons at typical frequencies (5 GHz and 6.6 GHz), with $E_C = 300$ MHz, connected to a common bus with qubit-qubit coupling $J_1 = 30$ MHz. (b) Zoom in on the avoided crossing of levels $|20\rangle$ and $|11\rangle$, working area of the CZ gate. (c) Zoom in on the avoided crossing of levels $|01\rangle$ and $|10\rangle$, working area of the $\sqrt{i}$ SWAP gate.

at each end to a transmon qubit.

A system of two multi-level transmons coupled to a common resonator (commonly referred to as a bus) is described by the Tavis-Cummings Hamiltonian [44]:

$$H = \omega_c a^\dagger a + \sum_{q \in [1,2]} \left(\sum_{j=0}^N \omega_{0-j}^q |j\rangle_q \langle j|_q + (a^\dagger + a) \sum_{j,k=0}^N g_{jk} |j\rangle_q \langle k|_q \right), \quad (2.20)$$

where ω_c is the bare bus frequency, ω_{0-j}^q is the transition energy between states 0 and j for transmon q and g_{jk} is the qubit-cavity couplings. Such a Hamiltonian couples the energy levels of the qubits, giving rise to avoided crossings between them. Figure 2.4 depicts the energy levels of two transmons as a function of the flux applied to one of the qubits. As levels are coupled they do not cross each other, mixing into eigenvectors with a superposition of both levels at the would-be crossing point. This provides a direct mechanism to implement two-qubit gates, by pulsing the qubits to an interaction zone. The system there evolves under the action of the coupling Hamiltonian, which defines a different basis. There are two main examples of this: the crossing of $|01\rangle$ and $|10\rangle$, and the crossing of $|11\rangle$ and $|20\rangle$, giving rise to two different two-qubit gates.

The crossing of $|01\rangle$ and $|10\rangle$ gives rise to a coupling Hamiltonian term of the form [45]

$$H = J_1 \left(\sigma_+^{(1)} \sigma_-^{(2)} + \sigma_-^{(1)} \sigma_+^{(2)} \right), \quad (2.21)$$

where $J_1 = \frac{g_1 g_2 (\Delta_1 + \Delta_2)}{2\Delta_1 \Delta_2}$ is the effective qubit-qubit coupling, g_i the coupling between qubit i and the bus resonator and $\Delta_i = \omega_i - \omega_b$ the detuning between the frequencies of qubit i and the bus resonator. The evolution operator for this coupling Hamiltonian is

$$U_{J_1}(t) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos(J_1 t) & i \sin(J_1 t) & 0 \\ 0 & i \sin(J_1 t) & \cos(J_1 t) & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (2.22)$$

and at the time $t = \pi/(4J_1)$, $U_{J_1}(t = \pi/(4J_1))$ implements the operator for an $\sqrt{i\text{SWAP}}$ gate. This is fully entangling gate that takes $|01\rangle$ and $|10\rangle$ respectively into the odd Bell states $|\Psi^+\rangle$ and $|\Psi^-\rangle$. In Chapter 4 we implement the operator from Equation (2.22) to generate moderate to full entanglement starting from $|01\rangle$, in order to solve the ground state of an H_2 molecule.

The second two-qubit gate we cover is the controlled-phase gate or CZ. This is a far more adopted gate in the context of transmon architectures [46], as it has a very simple connections to the gate CNOT = HCZH, commonly used in theory approaches to formulate algorithms. This gate is implemented by bringing into resonance the levels $|11\rangle$ and $|20\rangle$. The operator acts exactly the same as in the $\sqrt{i\text{SWAP}}$ gate, but as $|20\rangle$ is outside of the computational subspace, we need to expand our basis to include it $[|00\rangle, |01\rangle, |10\rangle, |11\rangle, |20\rangle]$:

$$U_{J_2}(t) = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & \cos(J_2 t) & i \sin(J_2 t) \\ 0 & 0 & 0 & i \sin(J_2 t) & \cos(J_2 t) \end{pmatrix}. \quad (2.23)$$

Here, the coupling that sets the gate speed limit is labeled $J_2 \approx \frac{(\sqrt{2}g_1)g_2(\Delta_{20}+\Delta_{11})}{2\Delta_{20}\Delta_{11}}$, and detunings Δ now refer to the frequencies of the levels involved ($|11\rangle, |20\rangle$). Note the appearance of a factor multiplying g_1 , as the coupling between the higher level ($|2\rangle$) of the transmon with the bus is increased by $\sim \sqrt{2}$. At $t = \pi/J_2$ we obtain a CZ gate. This was first implemented in reference [22] and has since been pushed forward by several efforts [47–49].

A topic of major concern for the performance of this gate is that of leakage out of the computational subspace: population remaining on the $|20\rangle$ as a result of the process. A first approach to minimize the impact of this introduced the parameterization of pulses in an adiabatic fashion, minimizing diabatic transitions [47]. Further works [48] extended the adiabatic pulses introducing a net-zero approach to mitigate distortion effects on the flux control line and include a built-in echo mechanism. The net-zero scheme destructively interferes the leakage occurring at each half of the gate, minimizing its impact even more.

In Chapter 5 we make use of CZ gates in the net-zero style [48] in order to simulate finite-temperature states on an Ising chain.

2.7 Tomography

Characterizing a quantum state is not a trivial task. Consider the quantum state of a single qubit, following Equation (2.15) we need two parameters θ, ϕ to describe it. But a measurement only yields us one value, stochastically, according to the probability amplitudes of the state. This is solved by repeating the experiment several times and computing a mean from the collected statistics.

However, a problem remains: a measurement m_Z , along the Z basis, will only gain us knowledge of θ according to $\bar{m}_Z = (\cos \theta)^2$. In order to gain knowledge about the degree of freedom ϕ , we need to perform a different measurement, one along the X basis (or any other axis in the equatorial plane of the Bloch sphere). Instead of rotating the axis of our readout (which couples to σ_z), we counter-rotate the axis of our qubit. For the example at hand, a measurement along the X axis, consider a rotation on the Bloch Sphere that takes the X axis onto the Z axis. Then whatever we measure as $\langle Z \rangle$ will correspond to $\langle X \rangle$ of the state before this rotation was introduced. We can then easily get ϕ from $\hat{m}_X = (\cos \phi)^2$.

From the above simple example, one can derive a further intuition: as quantum states scale, they span more and more degrees of freedom. To properly derive them from measurements, one needs not a single one, but rather a *set of measurements* along different bases that provide the data required to derive the state compatible with the experimental observations. Such a procedure is called *Quantum State Tomography* (QST) and the rotations used to change basis prior to measurement are called *tomographic pre-rotations*. As Quantum State Tomography forces the state to collapse via its measurement, it is typically the last part of a quantum circuit.

Since Quantum State Tomography aims at assigning a quantum state based on a set of measurements, it turns important that we address states in a more realistic way: by using the density matrix formalism and allowing for states to not be pure. As a hermitian unit-trace operator, $\hat{\rho}$ can be expressed in that vector space, along the basis of Pauli operators $P = [\hat{I}, \hat{X}, \hat{Y}, \hat{Z}]$ as

$$\hat{\rho} = \sum_{\hat{\sigma} \in P} \text{Tr}(\hat{\sigma} \hat{\rho}) \hat{\sigma}. \quad (2.24)$$

Under this more general description for mixed states, measurements along a specific basis get described by

$$\bar{m}_b = \text{Tr} \left(\hat{M} R_b^\dagger \rho R_b \right), \quad (2.25)$$

where $\hat{M}$ is the measurement operator, and R_b the tomographic pre-rotation. For transmons, the most basic modelling for $\hat{M}$ is the $\hat{Z}$ operator of the qubit being measured. However, this does not consider crosstalk (sensing other qubit states via this readout). We detail a more comprehensive model for $\hat{M}$ in Chapters 4 and 5.

Solving QST consists in assigning a state ρ , consistent with the datapoints measured for several bases and Equation (2.25). The simplest way is to solve the linear system of equations. We cover the details of this procedure are covered in Chapters 4 and 5.

2.8 Typical noise mechanisms

While Quantum Computing models can look simple, reality deviates significantly when considering noise mechanisms. These are undesired physical effects that alter the qubit state configuring a state error. Noise mechanisms are separated into coherent and incoherent ones.

Coherent noise mechanisms are those which can be modelled with a unitary operator, mapping pure states into other (*undesired*) pure states. An example of such error is a residual ZZ coupling, where a parasitic undesired coupling between two qubits continuously adds a coherent phase to their $|11\rangle$ state, much like a slow and always-on CZ gate.

Incoherent noise mechanisms are those which produce a detriment in the state purity, taking an initial pure state into mixed ones. To model them with a unitary evolution, one needs to address the environmental degrees of freedom that *take* the coherence away from the qubit. Instead, it is a common practice to model them with a Lindbladian formalism of superoperators. An example of such error mechanisms is the qubit relaxation, also referred to as T_1 .

2.8.1 Relaxation and dephasing times

For Transmon qubits the dominant coherence loss mechanisms are relaxation and dephasing. To model and understand these effects, we make use of the Pauli Transfer Matrix (PTM) formalism [2]. As in the tomography section, we model the state by its coefficients P_i in the bases of the Pauli set. Then, to model an operation over a state a PTM $\Lambda(P_i)$ is defined, acting on the state as

$$(R_\Lambda)_{i,j} = \frac{1}{d} \text{Tr} [P_i \Lambda(P_j)], \quad (2.26)$$

where R_Λ is a linear map acting on the P_i vectors within the Pauli-set subspace.

Following the above framework, we define the PTM for the relaxation of a single qubit as

$$R_{T_1} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \sqrt{1-p} & 0 & 0 \\ 0 & 0 & \sqrt{1-p} & 0 \\ p & 0 & 0 & 1-p \end{pmatrix}, \quad (2.27)$$

where p is the probability of a relaxation event. We experimentally quantify this via a T_1 experiment, where the qubit is prepared in $|1\rangle$ and measured after a certain wait time τ which is moved along the experiment. Stochastically, the qubit spontaneously decays from $|1\rangle$ to $|0\rangle$ emitting a photon with the respective energy $\hbar\omega_q$, in what we define as an relaxation event. As the experiment is repeated for different wait times, the increased probability of decay events produces an exponential decay, with characteristic time T_1 related to p by $p = 1 - e^{-t/T_1}$.

Similarly, the qubit dephasing is modelled as

$$R_{T_\phi} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1-p & 0 & 0 \\ 0 & 0 & 1-p & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (2.28)$$

where in turn, we define the characteristic time T_ϕ $p = 1 - e^{-t/T_\phi}$. This is characterized via a Ramsey experiment, where the excitation is put along the X axis, initially preparing the qubit in the state $|+\rangle$, and (also after the wait time τ) returning it to the Z axis for measurement. To be more precise, the Ramsey experiment measures a characteristic time T_2^* , combining the effects of both T_1 and T_ϕ as $\frac{1}{T_2^*} = \frac{1}{T_\phi} + \frac{1}{2T_1}$.

2.8.2 Thermal population

Transmons do not reach absolute zero temperature but rather a small finite one. While typical dilution refrigerators can reach 10–20 mK, this temperature is measured at the fridge plates, and not exactly at the qubits position. Because of this finite temperature, the qubit will not be at its ground state $|0\rangle$ but at a thermal state, corresponding to the Boltzman distribution over all its possible states, as quantified by Equation (2.14). The undesired thermal population at $|1\rangle$ is referred to as *thermal excitation*, and acts as an error mechanism for any algorithm that relies on the qubit being initialized in the $|0\rangle$ state. We can characterize this paying special attention to the readout characterization experiment of single shots as shown in Figure 2.2(a).

This problem can be actively avoided by *initialization by measurement*, where before running the algorithm/circuit of interest, the qubit is measured. Then, we can either post-select the data where the qubit started at $|0\rangle$ or actively perform a reset operation (π pulse) before the algorithm. One can also address this problem passively, by designing algorithms that are robust to thermal excitations. We show a procedure like this on Chapter 4.

2.8.3 Leakage

Transmons are not two-level systems, they are multi-level systems where we make use of only the two lowest energy levels $|0\rangle$, $|1\rangle$. This implies an important worry, which is to monitor that our operations do not transfer some of the populations to other energy levels, escaping the computational subspace.

The most relevant state is $|2\rangle$ because of its spectral proximity to the computational subspace ($|0\rangle$, $|1\rangle$). On one hand, its transition frequency ω_{12} sits very close to the qubit transition frequency ω_{01} , and the single-qubit microwave pulses occur at ω_{01} . On the other hand, when performing two-qubit gates, we bring the second-excited state levels closer in frequency to the computational level $|11\rangle$. This becomes extreme on CZ gates, where $|02\rangle$ is brought

into resonance with $|11\rangle$.

A way to address this is by actively measuring the $|2\rangle$ state. Then, Randomized Benchmarking protocols can tell us the average rate at which the operation in question (either single- or two-qubit gate) leaks to the second-excited state.

2.8.4 Residual ZZ coupling

The residual ZZ coupling is an undesired effect of coupling the qubits. When coupling qubits one has to consider the on/off ratio of this coupling with whatever turning knob is provided. In this thesis, we make use of transmons coupled via CPW resonators. The controlling knob to turn on and off this interaction is the qubit frequency, controlled by the flux threading through its flux-bias line. However, as the qubits are not infinitely detuned from each other at their flux sweetspot, there is still a residual effect of this coupling. While designed to be small, a mistargeting in qubit frequency or capacitive couplings can significantly alter this effect.

This effect is modeled by the Hamiltonian:

$$H_{ZZ} = \zeta_{i,j} |11\rangle \langle 11| = \zeta_{i,j} (II - IZ - ZI + ZZ), \quad (2.29)$$

where $\zeta_{i,j}$ is the residual coupling strength (expression up to second order in g^2/Δ).

2.9 Quantum Algorithms

Now that we can operate qubits, the question arises of what exactly do we want to do and how can we use this to compute something interesting. In this thesis we explore the usage of algorithms to simulate other systems with our quantum computer. However, we do this in three different ways: Trotterized Evolution, Variational Quantum Eigensolvers (VQE) and Quantum Alternating Operator Ansatz (QAOA).

2.9.1 Trotterized Evolution

A typical problem is that of simulating the action of a certain Hamiltonian H on the state $|\Psi\rangle$ of a system. If we define a mapping between the states in the system of interest and our qubits, we can as well do it for H and an equivalent unitary U in the qubit space. Finally, we can implement U in the qubit system, therefore simulating the system of interest.

A common challenge to implement the unitary H is that it often appears defined as a multi-term equation similar to $H = H_A + H_B$. While typically terms H_A and H_B have a simple correspondence to available qubit operations, the propagation operator $e^{-itH/\hbar}$ is not a commutative operation. This makes the implementation of H not straightforward as $e^{-itH/\hbar} \neq e^{-itH_A/\hbar} e^{-itH_B/\hbar}$ for any arbitrary A, B . To circumvent this one can make

use of the Trotter-Suzuki expansion formula. For that, we start from the Baker-Campbell-Hausdorff formula which states that

$$\log e^{\delta A} e^{\delta B} = \delta(A+B) + \frac{\delta^2}{2} [A, B] + \frac{\delta^3}{12} [A, [A, B]] - \frac{\delta^3}{12} [B, [A, B]] + \mathcal{O}(\delta^4), \quad (2.30)$$

where the brackets denote the quantum commutator operation. For small δ , this expression can be approximated to first order, yielding the Trotter-Suzuki approximation

$$e^{\delta(A+B)} = \lim_{\delta \rightarrow 0} \left[e^{\delta A} e^{\delta B} \right] + \mathcal{O}(\delta^2), \quad (2.31)$$

with an error given by the first commutator term $\frac{\delta^2}{2} [A, B]$. Furthermore, this approximation is exact for commuting operators A and B . Note that, if needed, one can push further the order of the approximation.

With the above expansion, and provided we can implement the individual terms H_A and H_B , one can in principle simulate any given system on our qubit platform. We explore this algorithm in an experimental setting on Chapter 3 for the specific case of the Rabi model in an ultra-strong coupling regime.

2.9.2 Variational Quantum Eigensolvers

A challenging problem in classical computing that may benefit from quantum computers is that of obtaining the eigenvalues of a system. Classical algorithms for diagonalization of an $M \times M$ matrix scale as $\mathcal{O}(M^3)$. While this might seem efficient as it is polynomial, when one considers Hilbert spaces, $M = 2^N$ with N the number of qubits, therefore being exponential on N . The same occurs while considering molecular orbitals instead of qubits: the size of the Hilbert space, spanned by the orbital bases, grows exponentially with N . The idea of Variational Quantum Eigensolvers (VQE) is to encode such systems (molecules) into the qubit Hilbert space, and use the quantum processor to solve the spectrum of the system with its various couplings (off-diagonal terms).

The key ingredient that guarantees the (ideal) success of VQEs is the variational principle of quantum mechanics. It states that given a parameterized family of states $|\Psi(\vec{\theta})\rangle$, the energy of any of these states is bounded by the ground state E_0 of the Hamiltonian:

$$\langle \Psi(\vec{\theta}) | H | \Psi(\vec{\theta}) \rangle \geq E_0 \quad \forall \theta. \quad (2.32)$$

This equation can be phrased in a simple way: the ground state of a Hamiltonian is the lowest energy of any state, and any combination of eigenvectors (arbitrary physical state) will not be able to violate that. For VQEs, this works as a guarantee, that disregarding the parameters $\vec{\theta}$ that control the production of states $|\Psi(\vec{\theta})\rangle$, the most one can do is get closer to the ground state energy *from above*.

The VQE algorithms are built around this principle as follows: One first encodes the Hamiltonian of interest (e.g. the molecule one) in the qubit Hilbert space, determining a map between target Hamiltonian terms, and those terms written in the Pauli basis of the qubit Hilbert space. Then one runs a circuit, where its gates implement the control parameters $\vec{\theta}$. Finally,

one measures the necessary elements (qubit correlators) in the Pauli basis, necessary to reconstruct the target Hamiltonian via the map. By doing this, one can obtain an experimental estimation for Equation (2.32). One can then close the loop by letting a classical optimization algorithm minimize $\langle \Psi(\vec{\theta}) | H | \Psi(\vec{\theta}) \rangle$ obtaining the best approximation to the ground state possible within the given parameterization.

As an example, let us consider the case of the H_2 molecule. Mapping the molecular orbitals starts from a basis definition, which then needs to be integrated over its spatial degrees of freedom and mapped onto the qubit space [50]. This needs to be done while enforcing anti-commutation rules for fermions, not enforced in our qubit subspace [51]. The procedure is nowadays automated and available from open-source software [52]. When mapping the molecular orbitals of H_2 onto two-qubits one obtains the Hamiltonian

$$H_{H_2} = h_{II}II + h_{ZI}ZI + h_{IZ}IZ + h_{XX}XX + h_{YY}YY + h_{ZZ}ZZ,$$

where the factors h_i are the result of the integrations explained above, and queried from *Open Fermion* [52]. This specific case, being the minimal example of VQE algorithms applied to quantum chemistry has been extensively studied with superconducting qubits [30, 31, 53] as well as other systems like photonics and trapped ions [54, 55]. In this case, the Bravyi-Kitaev transformation used encodes the presence of one electron into the bonding and anti-bonding orbitals into the states $|01\rangle$ and $|10\rangle$. This specific property gives rise to a symmetry: the solution for an electron in the H_2 molecule needs to be within the specific subspace in question. The knowledge of such a symmetry may assist in the experimental implementation of the algorithm: specific symmetry-respecting gates might be used [56], and it might be used to mitigate some errors [57]. In Chapter 4 we experimentally demonstrate the technique proposed in Ref. [57] improving by over one order of magnitude the accuracy of the results. Such symmetry successfully explains not only the improvements in our results [58] but also those obtained via Quantum Subspace Expansion [59].

2.9.3 Quantum Approximate Optimization Algorithm

The Quantum Approximate Optimization Algorithm (QAOA) is another variational approach to solving problems with a quantum computer [60]. Its formulation starts from the adiabatic theorem, which poses that the state of a quantum system under perturbation will remain in its instantaneous eigenstate superposition, provided the time scale of the evolution is slow enough.

This means that we start from a Hamiltonian H , varying in time from $H_0 = H(t=0)$ and eigenbasis $|\psi_j^0\rangle$ to $H_f = H(t=T)$ with eigenbasis $|\psi_j^f\rangle$. The intuitive idea behind the adiabatic principle is that if we manage to change the Hamiltonian very slowly, populations for each eigenstate should smoothly change from the initial to the final basis. As states might cross along the time trajectory, guaranteeing this is not straightforward and involves an additional requirement [61]: i) that the eigenstates i, j keep a minimal distance

$$\epsilon_{i,j} = \min_{0 \leq t \leq T} (E_i - E_j),$$

along the adiabatic trajectory and that ii) this is much bigger [62] than the derivative of the Hamiltonian element linking both

$$g_{i,j} = \max_{0 \leq t \leq T} \langle \Psi_i(t) | \frac{dH}{dt} | \Psi_j(t) \rangle,$$

obtaining $\epsilon_{i,j} \ll g_{i,j}^2$.

To link this with the original formulation of QAOA [60] we need to identify the initial Hamiltonian as H_M and the final Hamiltonian as H_P . The first one is referred to as the *mixing Hamiltonian*, typically chosen as the transverse field model $H_P = \sum_{j=1}^{N_q} X_j$. The latter is called *phase Hamiltonian*, and encodes the cost function f as $H_P |j\rangle = f(j) |j\rangle$. Finally, one needs to parameterize $H(t, \vec{\theta})$ slowly going from H_M to H_P , as parameterized by $\vec{\theta}$. This takes the ground state $|\Psi_0^0\rangle$ to a final state $|\Psi(t, \vec{\theta})\rangle = e^{-iTH(t, \vec{\theta})} |\Psi_0^0\rangle$. Following the adiabatic principle, and provided we have satisfied the $|\Psi(t, \vec{\theta})\rangle$ approximates the ground state of H_P , minimizing the eigenvalue f . The variational approach to this is to optimize with a classical algorithm the controls $\vec{\theta}$ for which f is minimized.

A first approach could be to trotterize the evolution of a Hamiltonian similar to $H = (1 - t/T) H_M + (t/T) H_P$. However, this is not necessarily the optimal trajectory. One can take another approach [63], by using Pontryagin's principle [64]: For the alternating application of terms like $e^{-i\alpha H_M}$, $e^{-i\gamma H_P}$, the optimal choice for the control parameters is either the minimum or maximum values possible for α, γ . This principle is commonly referred to as *the bang-bang principle* as it causes the evolution to maximally apply each Hamiltonian term after the other, obtaining:

$$|\Psi(t, \vec{\theta})\rangle = \left(\prod_{t=0}^T e^{-i\alpha_t H_M} e^{-i\gamma_t H_P} \right) |\Psi_0^0\rangle. \quad (2.33)$$

The reason why this formulation is more interesting than that of simply trotterizing an explicit parameterization is that it only relies on the bang-bang principle, and does not require the system to remain on the ground state for all t . Furthermore, parameters $\vec{\theta}$ have been unfolded into two sets $\vec{\alpha}, \vec{\gamma}$, not necessarily constraint by the time-parameterization of $H(t)$, and independent of each other at every time step t .

Recent results have shown applications of QAOA algorithms for several graph problems [43, 65] as well as magnetic systems [66]. In Chapter 5 of this thesis we use transmons to implement a QAOA algorithm for the simulation of finite-temperature Gibbs states on a spin chain.

EXPERIMENTALLY SIMULATING THE DYNAMICS OF QUANTUM LIGHT AND MATTER AT DEEP-STRONG COUPLING

The quantum Rabi model describing the fundamental interaction between light and matter is a cornerstone of quantum physics. It predicts exotic phenomena like quantum phase transitions and ground-state entanglement in ultrastrong and deep-strong coupling regimes, where coupling strengths are comparable to or larger than subsystem energies. Demonstrating dynamics remains an outstanding challenge, the few experiments reaching these regimes being limited to spectroscopy. Here, we employ a circuit quantum-electrodynamics chip with moderate coupling between a resonator and transmon qubit to realise accurate digital quantum simulation of deep-strong coupling dynamics. We advance the state of the art in solid-state digital quantum simulation by using up to 90 second-order Trotter steps and probing both subsystems in a combined Hilbert space dimension ~ 80 , demonstrating characteristic Schrödinger-cat like entanglement and large photon build-up. Our approach will enable exploration of extreme coupling regimes and quantum phase transitions, and demonstrates a clear first step towards larger complexities such as in the Dicke model.

3.1 Introduction

Targetted digital quantum simulations [68] are likely to provide the first demonstrations of quantum advantage for small-scale quantum computers, with applications in fields such as quantum chemistry [30, 69] and condensed-matter physics [28, 70–72]. In a digital quantum simulator, sequences of discrete interaction components synthesise the evolution of an artificial Hamiltonian, allowing access to more exotic dynamics than the simulator can realise naturally. Systems involving ultrastrong light-matter interactions raise significant challenges for both theoretical analysis [36, 73–77] and experimental study [78], making them ripe candidates for exploration via quantum simulation.

Ultrastrong coupling (USC) [73] of light and matter has been achieved in a range of physical systems, including circuit quantum electrodynamics (cQED) [79–81], semiconductor quantum-well systems [82], terahertz electron cyclotron transitions [83–85] and photochromic molecules [86]. Some experiments have demonstrated spectroscopic signatures deep into USC [80, 81, 84] where the coupling-to-frequency ratio $g/\omega \gtrsim 1$ (at so-called deep-strong coupling, or DSC [87]), but a dynamical signature has only been measured at $g/\omega \sim 0.09$ [85].

The standard quantum Rabi model (QRM) [36] describes the coupling of a two-level atom (energy $\hbar\omega_q^R$) to a quantum harmonic field mode (energy $\hbar\omega_r^R$) by a field-dipole interaction (energy $\hbar g^R$):

$$\frac{H_R}{\hbar} = -\frac{\omega_q^R}{2}\sigma_z + \omega_r^R a^\dagger a + g^R (a + a^\dagger) (\sigma^+ + \sigma^-), \quad (3.1)$$

where $a = \sqrt{n}|n-1\rangle\langle n|$ and $\sigma^- = |g\rangle\langle e|$ are annihilation operators for field mode and atom, respectively (with creation operators $a^\dagger$ and σ^+), and $\sigma_z = |g\rangle\langle g| - |e\rangle\langle e|$ is the Pauli z-basis operator. Under small coupling ($g^R \ll \omega_q^R, \omega_r^R$), this reduces to the Jaynes-Cummings (JC) model via the rotating-wave approximation:

$$\frac{H_{JC}}{\hbar} = -\frac{\omega_q}{2}\sigma_z + \omega_r a^\dagger a + g (a\sigma^+ + a^\dagger\sigma^-), \quad (3.2)$$

which contains only the excitation-number-conserving interaction terms, $a\sigma^+$ and $a^\dagger\sigma^-$, and has an exact solution. In the USC regime ($g^R \sim \omega_q^R, \omega_r^R$), however, the excitation-nonconserving terms $a\sigma^-$ and $a^\dagger\sigma^+$ cannot be neglected and only total parity $[\sigma_z \sum_n (-1)^n |n\rangle\langle n|]$ is conserved [87]. Without the strong symmetry of number conservation, the full QRM becomes difficult to solve [74], predicting phenomena such as ground-state entanglement and large ground-state photon populations which have not yet been observed experimentally. Theory suggests that simulations of the quantum Rabi model (QRM) could explore widely varied coupling regimes in architectures like cQED [88–90], cold atoms [91] and trapped ions [92]. Simulated QRM dynamics have been observed in restricted regimes in trapped ions, including the Dirac equation ($\omega_r^R=0, \omega_q^R \neq 0$) [93, 94] and coupling only ($\omega_r^R=0, \omega_q^R=0$) [95, 96] regimes. A classical analog simulation of evolution in a restricted subspace of the QRM has been performed in photonic waveguide systems [87, 97].

Here, we implement an accurate experimental simulation of quantum Rabi model dynamics well into the deep-strong coupling regime using a cQED quantum simulator with only

moderate atom-cavity coupling. To achieve this, we implement a digital protocol [88] with up to 90 second-order Trotter steps. In particular, we significantly extend the protocol by developing a phase-controlled method for tuning the target system parameters which allows us to explore a wide range of relative coupling strengths with high precision. Combining this control with versatile measurements of atom, cavity and joint system properties, we carry out a comprehensive study of quantum Rabi dynamics from ultrastrong to extreme deep-strong coupling. We first investigate the restricted case with zero atomic frequency ($\omega_q^R = 0$) to demonstrate key signatures verifying the simulation of deep-strong coupling. These include the characteristic collapses and revivals in both atom and cavity parities, coherent oscillations in cavity population reaching large photon numbers, and opposing cavity phase-space trajectories. We then show that the simulated deep-strong coupling leads to conditional non-classical Schrödinger cat states in the cavity, which verifies the presence of the atom-cavity entanglement arising from coherent deep-strong coupling dynamics. Finally, we study deep-strong coupling dynamics for several nonzero values of atomic frequency ($g^R/\omega_q^R \gtrsim 1$). This shows that our simulation is able to access the full complexity of the quantum Rabi model, and allows us to develop a heuristic understanding of the expected dynamics in terms of a competition between deep-strong coupling and Jaynes-Cummings dynamics.

3.2 Results

3.2.1 Digital quantum Rabi simulator with phase-controlled tuning

Deep-strong coupling dynamics can produce nontrivial quantum states and significant build-up of photon numbers [87]. Many characteristic dynamical features of DSC can already be seen in the degenerate-qubit limit $\omega_q^R = 0$. Here, the interaction-picture Hamiltonian

$$H_{R, \text{int}} = \sigma_x \left(\hbar g^R e^{-i\omega_r^R t} a + \hbar g^R e^{i\omega_r^R t} a^\dagger \right) \quad (3.3)$$

is a coherent drive on the oscillator mode, with an amplitude $\pm g^R e^{i\omega_r^R t}$ conditioned on the σ_x basis state of the atom ($\sigma_x = \sigma^+ + \sigma^-$ is the Pauli x-basis operator). The conditional coupling $\pm g^R$ coherently displaces the field, but in a continuously rotating direction given by $e^{i\omega_r^R t}$, creating two diametrically opposite circular trajectories in phase space [see later in Fig. 3.4(a), final panel]. Relating the diameter and circumference of these trajectories, $\pi\alpha_{\text{max}} = \dot{\alpha}T^R$, with field displacement rate $\dot{\alpha} = g^R$ and period $T^R = 2\pi/\omega_r^R$, gives a maximum amplitude $\alpha_{\text{max}} = 2r$ set by the relative coupling ratio $r \equiv g^R/\omega_r^R$. Figure 3.2(a) illustrates the atomic and photonic parity dynamics (σ_z and $\sum_n (-1)^n |n\rangle\langle n|$, respectively) for characteristic coupling regimes, starting in an eigenstate of the uncoupled system, $|e\rangle_q \otimes |0\rangle_r$. Because this is a superposition of the σ_x eigenstates $|\pm\rangle_q \sim |g\rangle_q \pm |e\rangle_q$, evolution gives rise to an atom-field entangled state (Bell-cat state) [98], $|+, +\alpha\rangle_{q,r} + |-, -\alpha\rangle_{q,r}$. For $r \ll 1$, the two trajectories remain virtually indistinguishable, giving evolution closely approximating simple JC dynamics with an atom-field detuning equal to ω_r^R [cf. Supplementary Note 7]. As r increases, the curves start distorting from the sinusoidal JC exchange oscillations (USC regime), until reaching DSC ($r \gtrsim 1$), where the parities exhibit a characteristic




Figure 3.1: Digital-analog quantum Rabi simulation using phase-controlled parameter tuning. (a) Parity dynamics of the ideal quantum Rabi model in the degenerate-qubit case ($\omega_q^R = 0$) for qubit (green) and resonator (red) in coupling regimes: $r = g^R/\omega_r^R = 0.1$ (dotted), 0.5 (dashed) and 1.0 (solid). In this example, g was chosen to match the experimentally observed value of $g/2\pi = 1.79$ MHz. (b) Two-transmon, three-resonator cQED chip (detailed description in Supplementary Notes 2 and 3). (c) Sequence schematic for second-order Trotterisation. The rotating frame defining the simulated resonator frequency (ω_r) is controlled via the Q_R bit-flip pulse phases. (d) Example simplified experimental pulse sequence for 5 Trotter steps followed by a photon parity measurement.

Gaussian-shaped “collapse”, followed by flat plateaus and periodic revivals at multiples of T^R . The cross-over between these dynamical regimes is related to the maximum distinguishability of the two coherent states of the field. When the paths separate completely, the qubit appears to be in a mixed state, with parity 0.5.

Our circuit QED Rabi simulator uses a hybrid digital-analog encoding of the atom and field mode, respectively, in a transmon qubit (Q_R) [18] and a coplanar waveguide resonator (R_R) (energies $\hbar\omega_q$ and $\hbar\omega_r$) [device shown in Fig. 3.2(b)]. Because the transmon is only weakly anharmonic ($\omega_q^{0-1} - \omega_q^{1-2} \ll \omega_q^{0-1}$), directly increasing the qubit-resonator coupling g leads to a breakdown in its qubit behaviour at small r , and full circuit quantization shows that DSC cannot be reached for any circuit parameters [99]. Instead, building on the proposal in Ref. [88], we perform a digital simulation of the QRM for arbitrarily large r using a coupling in the manifestly non-USC regime ($r < 10^{-3}$). The full Rabi Hamiltonian can be decomposed into two JC-like interactions [88]:

$$H_R(g^R, \omega_r^R, \omega_q^R) = H_{JC}(g, \Delta_r, \Delta_q^{JC}) + H_{AJC}(g, \omega_r, \Delta_q^{AJC}),$$

where $H_{AJC} = \sigma_x H_{JC} \sigma_x$ contains only counter-rotating interaction terms, and the effective Rabi parameters $g^R = g$, $\omega_r^R = 2\Delta_r$ and $\omega_q^R = \Delta_q \equiv \Delta_q^{JC} - \Delta_q^{AJC}$ are not related to the natural circuit frequencies, but defined relative to a nearby rotating frame

($\Delta = \omega - \omega_{\text{RF}}$), and can be arbitrarily small. Using the standard method of Trotterization [68], Rabi dynamics can therefore be simulated into the DSC regime by decreasing Δ_r and Δ_q . Figure 3.2(c) illustrates the second-order Trotter step used here [see Methods and Supplementary Note 9]. An asymmetric transmon with two flux-insensitive “sweet” spots [18] is driven and measured at its lower sweet spot (5.452 GHz) far below the resonator (6.381 GHz), with digital π pulses being interleaved with short analog JC interaction blocks applied by fast frequency-tuning flux pulses [22]. [See Supplementary Note 3 and Supplementary Table 1 for details of the experimental scheme and Supplementary Note 4 for details of how the flux-pulse distortion compensations are calibrated.] Experimentally, a rotating frame is usually defined by the frequency of a drive tone. Here, the choice of rotating frame specifies the required rotation axis of the π pulses which create the AJC interaction. By appropriately updating the pulse phases, which are controlled with high precision, we can therefore arbitrarily select the rotating frame detuning from the resonator, even though these pulses are applied far from both resonator and rotating frame [see Methods].

Numerical modelling of the digital Rabi protocol highlighted several challenges for device design and fabrication [Supplementary Note 2]. Most significantly, due to practical flux-pulsing bandwidths which limit the shortest achievable Trotter step, it is challenging to digitise fast compared with the dynamics. Reaching acceptably low Trotter error in interesting regimes of r therefore required small qubit-resonator coupling (here, $g/2\pi = 1.95$ MHz). This also placed constraints on other device parameters, including coherence (for long simulation times), flux-tuning precision and qubit-resonator frequency targetting (due to a very narrow resonance). An extra qubit Q_W was strongly and dispersively coupled to R_R to probe the intraresonator quantum state via its photon-dependent frequency shift (-1.26 MHz per photon) using pulse sequences based on Ramsey interferometry. We used Q_W to implement a range of photon measurements: average photon number with a controllable dynamic range (number meter), average photon parity [100, 101] (parity meter) and, combining parity measurements with coherent field displacements through an external input coupler, direct Wigner tomography of the resonator. [Full details of the operating principles and calibrations of these different photon measurements are provided in Supplementary Notes 5 and 6.] Qubits were driven and measured through dedicated read-out resonators. A full description of the experimental setup is provided in Supplementary Note 1.

3.2.2 Comparing qubit and resonator parity dynamics

We first experimentally simulate the QRM for the degenerate-qubit case over a wide range of r , covering the USC and DSC regimes from $r \sim 0.3$ to $r \rightarrow \infty$ (Fig. 3.2). We use 60 Trotter steps to simulate $1.2 \mu\text{s}$ of dynamics ($gt = 4.68\pi$) and measure either qubit or photon parity after each step. (Simulations start in the state $|1\rangle_q \otimes |0\rangle_r$ for all results in the main text, but Supplementary Note 8 shows that the features of DSC dynamics are observed also for $|0\rangle_q \otimes |0\rangle_r$.) A simplified pulse sequence is illustrated in Fig. 3.2(d). The qubit and photon parity dynamics [Fig. 3.2(e, g)] show very similar qualitative behaviour, consistent with parity conservation. At all large couplings, the measurements exhibit the Gaussian-shaped parity




Figure 3.2: **Qubit and resonator parity dynamics of the quantum Rabi model in the degenerate-qubit case.** Measured dynamical landscapes for (a) qubit and (b) photon parity dynamics for a broad range of parameters up to 60 Trotter steps, with the extreme DSC regime in the centre decreasing to weaker USC near the edges. The data show clear Gaussian-shaped collapses for all r , along with the characteristic plateaus of DSC. Qubit revivals are observed up to $r \sim 0.8$, while photon parity shows clear revivals up to $r \sim 1.8$. Vertical line cuts for (c) qubit and (d) photon parity are plotted for evenly spaced ω_r^R/g^R between the red and blue dashed lines in (a) and (b), respectively. For $r \gtrsim 1.5$, some deviation from the expected revival time in the photon parity results from a small residual Kerr-type nonlinearity in the resonator [see also Supplementary Fig. 9] and is correlated with significant photon populations. Arrows in (c) and (d) show expected revival times for each slice. In this and following figures, coupling ratios were calculated using the observed simulated coupling of $g/2\pi = 1.79$ MHz.

collapse (set by the simulated g^R) and flat plateau which is a key signature of DSC dynamics. Fitting the initial qubit data points, we calculate an average $g^R \approx 2\pi \times 1.79$ MHz, slightly lower than the expected $g^R = g \approx 2\pi \times 1.95$ MHz determined from independent spectroscopy and vacuum Rabi oscillations. This is consistent with a small residual flux-pulse distortion and provides the best estimate for the simulated g^R achieved in these experiments. The revival periods T^R are in excellent agreement with the predictions of USC Rabi dynamics (dashed curves), and strikingly different from those predicted for a pure Jaynes-Cummings interaction with the equivalent qubit-resonator detuning ($T^{\text{JC}} = 2\pi/\sqrt{4g^2 + \Delta_{\text{q-r}}^2}$) (dotted curves) [Supplementary Note 7].

From the observation of parity revivals, combined with the simulated g^R , we can estimate the range of r reached in these simulations. For $g^R/2\pi = 1.79$ MHz and $r = 1$ (archetypal DSC), the expected revival time is $0.56 \mu\text{s}$. Line cuts for the qubit parity dynamics [Fig. 3.2(f)] show revivals beyond $0.4 \mu\text{s}$ ($r \sim 0.7$). Photon parity revivals, however, persist beyond

$1.0 \mu\text{s}$ ($r \sim 1.8$). This difference again results from photon decay, as shown by excellent agreement with numerical modelling which includes cavity decay but no other decoherence (not shown). Photon decay becomes increasingly critical at larger couplings, because even a single decay destroys the qubit-resonator entanglement, and losing a photon becomes increasingly likely for larger photon numbers. The qubit parity revivals rely on entanglement being maintained. This is supported by measurements of reduced qubit entropy, which show that the qubit state collapses to the mixed state, before displaying a revival in purity [Supplementary Note 11]. The resonator parity dynamics, however, are more robust to decay and provide a more direct measure of DSC dynamics. Photon parity collapses and revivals prove the field undergoes large-amplitude excursions through phase space even during a single cycle of the resonator period. The difference between qubit and photon parity dynamics is a quantitative signature of break-down in parity conservation, caused by resonator decay.

3.2.3 Resonator photon-number dynamics

We next directly explore the build-up of large photon populations (Fig. 3.3), another feature of DSC dynamics that contrasts strikingly with the excitation-conserving dynamics expected under weak coupling. Using a Ramsey pulse sequence with small separation τ , the excitation probability in Q_W becomes a measure of average photon number in the resonator [Supplementary Note 5]. The dynamic range and sensitivity of this number meter are controlled via τ [Figs 3.3(a, b)]. Measured with a linear range of ~ 0 –8 photons [Fig. 3.3(c)], the resonator displays the complementary build-up of photons which causes the collapse of qubit and photon parity, clearly demonstrating the violation of number conservation expected for the QRM. As with photon parity, clear oscillations can be seen out to $r \sim 1.8$ [Fig. 3.3(e)]. The large central feature appears to deviate from the expected trend, but is in fact due to photon number exceeding the dynamic range of the number meter. To explore this region further, we extended the linear range to ~ 0 –20 photons using a number meter with a non-centred refocussing pulse [Fig. 3.3(d)] and simulated up to 90 Trotter steps ($gt = 7.0\pi$), allowing photon oscillations beyond $1.5 \mu\text{s}$ to be observed. This range operated at the limits of approximately uniform driving given the bandwidth of the 12 ns (4σ) QW pulses. At $r \gtrsim 2$, the photon dynamics in Figs 3.3(c) and (d) are clearly skewed, causing the observed oscillations to deviate from the expected revival period T^R (also observable in the photon parity [Fig. 3.2(g)]. This results from a residual Kerr nonlinearity in R_D inherited from the dispersively coupled ancilla qubit [102].

Exploring the resonator oscillations more quantitatively, the maximum photon number in each vertical (constant- r) slice [Fig. 3.3(f)] compares well with the expected ideal behaviour. The discrepancy between the two curves in the overlapping region results from bandwidth limitations in the high-dynamic-range (HDR) number meter and the limits in linearity of the number-to-probability mapping for Q_W . Because of the sinusoidal conversion, the calibrated value at either end of the range compressed slightly towards the centre from the real photon number. The measurement saturates at the highest r even for the HDR meter, suggesting that we observe more than 30 photons (average) building up in the resonator for the strongest

DSC regions. Given the Poissonian statistics expected for coherent states, this accesses a resonator subspace of dimension ~ 40 (i.e., a subspace larger than that of 5 qubits). This ability to access large Hilbert spaces with a simple system is an advantage of the analog resonator encoding.

3.2.4 Resonator phase-space dynamics

Combining the parity measurement with coherent displacements from an external drive allows observation of resonator phase-space dynamics using direct Wigner tomography [100, 101]. Figure 3.4(a) shows unconditional maximum-likelihood tomograms [ignoring the state of Q_R ; see Methods] measured after each Trotter step with $r \sim 0.9$ [full movie available in Supplementary Movie 1], with the full trajectory obtained from two-dimensional double-Gaussian fits of the raw data. The resonator state displays the clear signatures of DSC dynamics, first separating into two distinct Gaussian (coherent-state) peaks which follow opposite circular trajectories before re-coalescing at the origin. The peaks do not return perfectly to the origin because of photon decay, in agreement with a numerical simulation at $g^R/2\pi = 1.79$ which includes $T_{1,r} = 3.5 \mu\text{s}$ (green curves).

3.2.5 Demonstrating qubit-resonator entanglement

By capturing the complete resonator quantum state, the Wigner function also enables the demonstration of coherence in DSC dynamics, by contrast with photon parity and number measurements, which are largely insensitive to coherence. Observing this requires correlating the resonator and qubit states, because the coherence is stored in entanglement. We did this in two ways. First, we measured the Wigner function after 10 Trotter steps for $r \sim 0.9$ with Q_R initialised in states $|g\rangle$, $|e\rangle$, $|+\rangle$ and $|-\rangle$ [Figs 3.4(b–e)]. This showed that the resonator and qubit were correlated, consistent with the expected Bell-cat entanglement. Second, we ran the simulation for $r \sim 0.9$ and 2.1 (8 Trotter steps) with the qubit prepared in the excited state, conditioning the Q_W measurement on the state of Q_R in the σ_z basis (Fig. 3.5). For the expected Bell-cat state, an outcome of $|g\rangle$ ($|e\rangle$) for Q_R leaves the resonator in an odd (even) Schrödinger cat state ($|\alpha\rangle \mp |-\alpha\rangle$). Numerical modelling shows that only in the DSC regime is negativity in the Wigner function observed for both Q_R measurement outcomes. The negative regions observed in all the Wigner functions demonstrate nonclassicality for all resonator cat states, which arises from coherence in the underlying Bell-cat entanglement. Reduced visibility is again caused primarily by photon decay, but also by single-shot read-out infidelity (here, $\sim 85\text{--}90\%$) and experimental drift over the long measurements. These different measurements provide clear evidence of qubit-resonator entanglement arising from coherent DSC dynamics.

3.2.6 Quantum Rabi dynamics in the nondegenerate-qubit case

Finally, by detuning the qubit frequency during the AJC half of the Trotter steps [Fig. 3.2(c)], we also experimentally simulate dynamics for the nondegenerate-qubit case of the QRM for effective qubit frequencies $g^R/\omega_q^R \sim 4, 2$ and 1 (Fig. 3.6). Deviation from the degenerate-qubit

case occurs primarily when $\omega_r^R \lesssim \omega_q^R$ [103] and these regimes access the full complexity of QRM dynamics. To develop a rough intuition for the expected dynamics, we overlay the plotted landscapes with the expected revival times for both pure degenerate-qubit QRM dynamics and pure nondegenerate-qubit Jaynes-Cummings (exchange) dynamics (centred around the effective qubit frequency). This illustrates that the ideal dynamics (no decay) [Fig. 3.6(right)] can be thought of as a competition between the two cases. As qubit frequency increases, standard Jaynes-Cummings dynamics begin to emerge, with qubit population oscillations (and increasingly pronounced positive-parity regions) appearing in the collapse-revival dynamics characteristic of the DSC regime. This interpretation and trend become clearer for qubit frequencies ω_q^R larger than the coupling g^R , where the standard Jaynes-Cummings exchange dynamics start to dominate [numerical modelling shown for $g^R/\omega_q^R \sim 0.48$ in Supplementary Fig. 12]. The measured dynamics [Fig. 3.6(left)] capture many features of the ideal case [Fig. 3.6(right)], even up to $r \gg 1$. Numerical modelling of the digital QRM simulation including the measured $T_{1,r}$ [Fig. 3.6(centre)] confirms that simulation fidelity is primarily limited by resonator decay.

3.3 Discussion

Demonstrating stabilisation by decreasing step sizes will be an important part of validating the behaviour of future complex digital simulators achieving quantum advantage [104]. In Supplementary Notes 9 and 10, we showed that using second-order Trotterisation and decreasing the Trotter step size both significantly improved performance. This indicates that the simulation is not limited by an error-per-gate noise floor as in previous cQED simulations [28], and enables us to linearly increase the number of Trotter steps for increasing simulated time, rather than keeping the number fixed [28, 30, 72]. This is an important step towards the quadratic scaling needed for universal quantum simulation [68]. In combination, these achievements advance solid-state quantum simulators based on cQED to a digital performance previously attained only in trapped-ion systems [71].

Interestingly, a QRM simulator even has some direct advantages over natural USC systems. Although large couplings can lead to ground-state entanglement and significant ground-state photon populations, these potentially interesting ground states are not readily accessible in natural USC systems [78, 99, 105] without the ability to rapidly (nonadiabatically) tune or switch off the ultrastrong coupling. In systems where the coupling reaches many gigahertz, tuning system parameters on this timescale represents a significant technical challenge [80, 81]. In our simulator, however, cavity photons are always real (not virtual), detectable and usable, and it is straightforward to nonadiabatically tune system parameters to implement quantum quenches [106]. This makes a cQED chip with natural JC interactions an ideal platform to explore the preparation of interesting ground states in future experiments. The challenge is that the simulator decay processes differ from those in a natural USC system and do not move the system towards the USC ground state [75]. This highlights the need to improve $T_{1,r}$ so that photon decay does not limit the dynamics. It should be possible to improve $T_{1,r}$ ten-fold using novel processing methods [107]. However, an interesting next

step will be to determine the effective USC decay resulting from simulator-frame resonator decay.

Finally, the phase technique we have developed to define a rotating frame via single-qubit pulses introduces a precise and flexible paradigm for engineering artificial Hamiltonians which can be applied across architectures such as trapped ions and cold atoms [71, 91, 92]. In combination with the number of Trotter steps demonstrated, the technique will allow accurate simulation of the time-dependent Hamiltonians [28, 71, 108] required to perform adiabatic preparation of USC ground states. It is therefore ideally suited for exploring novel quantum phase transitions relying on extreme coupling regimes recently identified for the QRM [91, 109, 110]. Further, by extending to small-scale Dicke-model systems [88, 90], it will avoid the problem of additional nonlinear evolution terms [90] which have been suggested to prevent the onset of a long-predicted superradiant phase transition in a range of physical systems [76, 77, 99, 111].

3.4 Methods

3.4.1 Phase-controlled Trotterisation of the quantum Rabi model

In the digital QRM simulation proposed in Ref. [88], the effective parameters of the simulated Rabi Hamiltonian are $g^R = g$, $\omega_r^R = 2\Delta_r$ and $\omega_q^R = \Delta_q^{JC} - \Delta_q^{AJC}$, where $\Delta_r = \omega_r - \omega_{RF}$ and $\Delta_q = \omega_q - \omega_{RF}$ are defined relative to a rotating frame. This rotating frame is essential to reaching DSC with weakly anharmonic transmon qubits, by allowing us to tune the simulated ω_r^R and ω_q^R . Typically, the frequency of a rotating frame is set by a physical generator or drive signal that defines a rotation or a measurement basis. In the digital simulation, the rotating frame is still abstract, since no drive is used to induce an interaction. Here, we describe a method we have developed for controlling the frequency of the rotating frame which is simple, high-resolution and flexible.

The basic intuition is that the bit flips in Ref. [88], which convert every second JC interaction into an effective AJC interaction, are the only concrete operations which take place in the otherwise abstract rotating frame. In any Trotter step, the frequency of the rotating frame is therefore defined by the rotation axes of the bit-flip pulses (i.e., the absolute pulse phase), but these flips are driven by microwave pulses at a frequency far (~ 1 GHz) below the resonator, at the qubit's bottom sweet spot. Nevertheless, while the drive generator's phase continuously and rapidly rotates relative to the resonator, the drive pulses can be effectively locked to the resonator frequency by discretely updating the pulse phase at each pulse. This is achieved by advancing the phase of each pulse by an amount proportional to the elapsed time between pulses. An arbitrary offset frequency from the resonator is then straightforwardly achieved by correcting this phase advance by an amount proportional to the Trotter step size. Interestingly, in the scheme of Ref. [88], because the simulated resonator frequency (but not the qubit frequency) is sensitive to the absolute detuning from the rotating frame, this effective qubit offset frequency tunes the frequency of the resonator (but not the qubit).

We now derive the analytical relation between the bit-flip pulse phases and the rotating-frame frequency in the simulation. We start by writing down the full Trotter step and then

derive the effective Hamiltonian implemented by this step given the lowest-order Trotter approximation. The symmetric, second-order Trotter step for the digital QRM simulation is:

$$U_R^{\text{Tr}}(\tau) = U_{\text{JC}}^{\frac{1}{2}}(\tau) U_{\text{AJC}}(\tau) U_{\text{JC}}^{\frac{1}{2}}(\tau), \quad (3.4)$$

where $U_{\text{JC}}(\tau) = \exp(-iH_{\text{JC}}\tau/\hbar)$ and an arbitrary AJC step

$$U_{\text{AJC}}(\tau) = R_{\phi_2}(\pi) \exp\left(\frac{-iH_{\text{JC}}\tau}{\hbar}\right) R_{\phi_1}(\pi), \quad (3.5)$$

is defined by the phases used to set the rotation axes $\phi_{1,2}$ of the bit flips $R_\phi(\pi)$. Writing the JC Hamiltonian in the rotating frame of the resonator, and using the identity $R_\phi(\pi) = R_z(\phi) R_x(\pi) R_z(-\phi) = R_z(2\phi) \sigma_x = \sigma_x R_z(-2\phi)$, gives:

$$U_{\text{AJC}}(\tau) = R_z(2\phi_2) \sigma_x \exp\left[-i\Delta_{q-r}\tau\sigma_z/2 - i\epsilon\left(a\sigma^+ + a^\dagger\sigma^-\right)\right] \sigma_x R_z(-2\phi_1), \quad (3.6)$$

$$= \exp(-i\Delta\phi\sigma_z/2) \exp(-i\phi_\Sigma\sigma_z/2) \exp\left[i\Delta_{q-r}\tau\sigma_z/2 - i\epsilon\left(a\sigma^- + a^\dagger\sigma^+\right)\right] \quad (3.7)$$

$$\exp(i\phi_\Sigma\sigma_z/2) \exp(-i\Delta\phi\sigma_z/2), \quad (3.8)$$

$$= \exp(-i\Delta\phi\sigma_z/2) \exp\left[i\Delta_{q-r}\tau\sigma_z/2 - i\epsilon\left(a\sigma^- e^{-i\phi_\Sigma} + a^\dagger\sigma^+ e^{i\phi_\Sigma}\right)\right] \quad (3.9)$$

$$\exp(-i\Delta\phi\sigma_z/2), \quad (3.10)$$

where $\epsilon = g\tau$, $\phi_\Sigma = \phi_1 + \phi_2$, $\Delta\phi = \phi_2 - \phi_1$, $\Delta_{q-r} = \Delta_{q-r}^{\text{JC}} - \Delta_{q-r}^{\text{AJC}}$, and we have set $\Delta_{q-r}^{\text{JC}} = 0$. Equation (3.10) is reached by noting that $e^{-i\phi_\Sigma\sigma_z/2}\sigma^\pm e^{i\phi_\Sigma\sigma_z/2} = \sigma^\pm e^{\pm i\phi_\Sigma}$.

Next, noting that $\Delta\phi = \pi\omega_R^R\tau \ll 1$ if $\tau \ll 1/\omega_R^R$, and providing the Trotter conditions $\epsilon = g\tau \ll 1$ and $\Delta_{q-r}\tau \ll 1$ are fulfilled, we can combine exponentials in Eq. (3.10) using a Trotter approximation to give:

$$U_{\text{AJC}}(\tau) \approx \exp\left[-i\Delta\phi\sigma_z + i\Delta_{q-r}\tau\sigma_z/2 - i\epsilon\left(a\sigma^- e^{-i\phi_\Sigma} + a^\dagger\sigma^+ e^{i\phi_\Sigma}\right)\right]. \quad (3.11)$$

Combining the JC and AJC steps with a further Trotter approximation then gives the full Trotter step

$$U_R^{\text{Tr}}(\tau) \approx \exp\left[i(-2\Delta\phi + \Delta_{q-r}\tau)\frac{\sigma_z}{2} - i\epsilon\left(a\sigma^+ + a^\dagger\sigma^- + a\sigma^- e^{-i\phi_\Sigma} + a^\dagger\sigma^+ e^{i\phi_\Sigma}\right)\right]. \quad (3.12)$$

So far, we have considered arbitrary ϕ_1 and ϕ_2 . In the experiment, however, we keep $\Delta\phi$ constant for all sequential pairs of bit flips. Specifically, for the n th Trotter step, the two phases are $\phi_1 = \phi_0 + (2n-2)\Delta\phi$ and $\phi_2 = \phi_0 + (2n-1)\Delta\phi$, where the choice of ϕ_0 has no effect on the dynamics. Setting $\phi_0 = 3\Delta\phi/2$ gives $\phi_\Sigma = 4n\Delta\phi$, and the n th Trotter step can be rewritten in terms of a frequency $\omega_0 = 2\Delta\phi/\tau$ and a simulated time $t_n = n\tau$:

$$U_R^{(n)}(\tau) = \exp\left[i(-\omega_0 + \Delta_{q-r})\tau\frac{\sigma_z}{2} - i\epsilon\left(a\sigma^+ + a^\dagger\sigma^- + a\sigma^- e^{-i2\omega_0 t_n} + a^\dagger\sigma^+ e^{i2\omega_0 t_n}\right)\right]. \quad (3.13)$$

which corresponds to an effective Hamiltonian:

$$\frac{\tilde{H}_{\text{eff}}}{\hbar} = (\omega_0 - \Delta_{\text{q-r}}) \frac{\sigma_z}{2} + g \left(a\sigma^+ + a^\dagger\sigma^- + a\sigma^- e^{-i2\omega_0 t} + a^\dagger\sigma^+ e^{i2\omega_0 t} \right). \quad (3.14)$$

Until this point, the calculation has been carried out with both qubit and resonator in a frame rotating with the resonator. We now transform $\tilde{H}_{\text{eff}}$ into a rotating frame where both qubit and resonator are rotating at frequency $(-\omega_0)$, i.e., with $H_0 = -\hbar\omega_0 \left(-\sigma_z/2 + a^\dagger a \right)$, giving a new effective Hamiltonian:

$$\frac{H_{\text{eff}}}{\hbar} = -\Delta_{\text{q-r}} \frac{\sigma_z}{2} + \omega_0 a^\dagger a + g \left(a + a^\dagger \right) \left(\sigma^+ + \sigma^- \right). \quad (3.15)$$

This completes the mapping of the phase-controlled Trotterisation into the form of a simulated Rabi Hamiltonian and we can now identify the effective simulated parameters $g^{\text{R}} = g$, $\omega_{\text{q}}^{\text{R}} = \Delta_{\text{q-r}}$ and $\omega_{\text{r}}^{\text{R}} = \omega_0 = 2\Delta\phi/\tau$. Note that the final frame transformation takes place in the *simulated Hilbert space*, i.e., with frequency ω_0 defined relative to simulated time. Consequently, the frequency of the abstract rotating frame in Ref. [88], defined in the laboratory reference frame of the cQED simulator, is less by a factor 2, i.e., $\omega_{\text{RF}} = \omega_0/2$.

Here, we have shown how to engineer a virtual rotating frame by applying virtual phase corrections via updating the rotation axis of subsequent drive pulses [112] in the stroboscopic context of Trotterised digital quantum simulations. This technique should be broadly applicable in the context of Trotterised quantum simulations, although some details or interpretation may vary depending on the specific simulation. For example, it could be applied virtually unmodified to implement the digital Ising model simulations with interacting spins from Ref. [72], where phase gates were instead implemented via physical detunings of the qubits (as also done in [71]). More generally, in Trotterised dynamics, a continuous frequency detuning is to lowest order identical to a discrete phase gate applied in each Trotter step. In any case where a gate is implemented using an exchange-type interaction, frequency detunings can therefore be effectively transferred between different circuit elements and mapped onto the most easily controllable element. This turns the theoretical aide of moving between interaction pictures into a concrete experimental tool. If the Trotter step also includes single element control pulses, then these can often be modified to also incorporate the phase gate. If this option is not available [e.g., see the digital Jaynes-Cummings simulation in Supplementary Note 7] then the phase correction can still be implemented directly. In our case, a simulated frequency detuning was applied to a resonator (which was not easily tunable) by *virtually* applying a discrete phase update to the qubit via the drive phase of the bit-flip pulses.

3.4.2 Trotter step

For a second-order Trotter step with simulated time τ , the Trotter step consists of 3 flux pulses ($\tau/2$, τ and $\tau/2$) and 2 single-qubit rotations with buffers separating the different gates. Adjacent $\tau/2$ flux pulses from neighbouring Trotter steps are implemented as a single flux pulse of length τ . Each flux pulse was followed by a 5 ns phase-compensation flux pulse [Supplementary Note 7]. For most of the data presented in this work, the simulated $\tau = 20$ ns. The qubit drive pulses on Q_{R} were 16 ns total duration (4σ) and the pulses

buffers were 10 ns. The total Trotter step for $\tau = 20$ ns was therefore $\tau_{\text{step}} = 122$ ns. In addition to the drive-pulse phase advance required to define ω_r^R , another linear phase advance $\Delta\phi = (\omega_q^{\text{drive}} - \omega_r)\tau_{\text{step}}/2$ is required to compensate the rapid rotation of the qubit drive with respect to the resonator frequency.

3.4.3 Qubit control

Qubit rotations were implemented using DRAG pulses [21, 41], with a Gaussian envelope in the X quadrature and a derivative-of-Gaussian envelope in the Y quadrature. The 4σ pulse durations were 16 ns for Q_R and 12 ns for Q_W . The performance of the Trotter sequences, which contained up to 180 bit-flip pulses, was very sensitive to details of the Q_R pulse calibrations. In particular, the drive amplitude was calibrated using a sequence of 50 π -pulse pairs preceding a single $\pi/2$ pulse. All parameters were typically calibrated just before launching a long measurement. The drive amplitude was intermittently recalibrated during the scans. Because only 2 or 3 pulses were applied to Q_W for the photon measurements, it was optimised using the AllXY sequence [113] of 21 combinations of two σ_x and σ_y rotations (either $\pi/2$ or π). The frequency of Q_W was regularly calibrated during photon measurements using Ramsey sequences.

3.4.4 Wigner tomography reconstructions

Tomograms shown in Figs 3.4 and 3.5 are maximum likelihood reconstructions [114, 115] of the resonator quantum state from direct Wigner tomography measurements [101]. The Wigner function at a phase-space position α is:

$$W(\alpha) = \frac{2}{\pi} \text{Tr} \left[\Pi D^\dagger(\alpha) \rho_r D(\alpha) \right] = \frac{2}{\pi} \text{Tr} [M_\alpha \rho_r], \quad (3.16)$$

where ρ_r is the resonator density matrix, $\Pi = \sum_n (-1)^n |n\rangle \langle n|$ is the photon parity operator and $D(\alpha)$ is the coherent displacement operator. For each measured α , we calculated $M_\alpha = D(\alpha) \Pi D^\dagger(\alpha)$ using an operator dimension much larger than the largest $|\alpha|^2$ in the measured phase space, to avoid edge effects when calculating $D(\alpha)$. The M_α were then truncated to a maximum photon number sufficient to capture all of the reconstructed state, but small enough to allow fast reconstructions and ensure an informationally complete set of operators ($n_{\text{max}} = 12$ and 8 for tomograms in Figs 3.4 and 3.5, respectively). The maximum likelihood reconstruction was carried out using convex optimisation [116, 117]. In Fig. 3.4, a systematic phase correction was applied to the density matrices to correct for a miscalibration of the resonator drive phase used in the coherent displacement. Finally, the reconstructed density matrix was then used to calculate the plotted Wigner functions.

3.5 Supplementary Notes

3.5.1 Supplementary Note 1: Experimental setup

The sample and low-temperature microwave components were mounted inside magnetic and infrared radiation shielding consisting of two layers of cryogenic mu metal around a layer of

aluminium, with an internal layer of copper foil coated in a mixture of silicon carbide and Stycast (2850 FT) [118]. Microwave coaxial cables are connected to the PCB-mounted chip via non-magnetic SMP connectors (Rosenberger).

The qubit drive and read-out tones are sent through two dedicated feedlines which are connected via a short coaxial cable off-chip. The input line for the qubit drives is filtered at the mixing chamber with 30 dB cold attenuation, a small home-built inline eccosorb filter and a 10 GHz low-pass filter (K& L 6L250-10000/T20000-0/0). [The resonator input line filter is 8 GHz low-pass (K& L 6L250-8000/T18000-0/0).] The output line passes through two 3 – 12 GHz isolators (Pamtech CWJ1019K) and a circulator (Quinstar CTH0408KCS) mounted above the mixing chamber on the way to a 4 – 8 GHz cryogenic HEMT amplifier (Low-Noise Factory LNF-LNC4_8A), two room-temperature amplifiers (Miteq AFS3-04000800-10-ULN, then AFS3-00101200-35-ULN-R), RF demodulation (Marki 0618LXP IQ mixer) and amplification, and finally digitised in a data acquisition card (AlazarTech ATS9870). The flux-bias lines are filtered at the mixing chamber with 1.35 GHz low-pass filters (Minicircuits VLFX-1350) followed by home-built eccosorb filters. All input lines are thermalised with 20 dB attenuators mounted at the 4 K plate. The microwave input lines and output line are connected to the fridge through a DC block.

Qubit and resonator drive pulses are created via single-sideband modulation with IQ mixers and generated by two arbitrary waveform generators (AWGs; Tektronix AWG5014). We use a 3–7 GHz IQ mixer (Marki 0307MXP) for the resonator and two custom-built 4–8.5 GHz IQ mixers (QuTech F1c: DC–3.5 GHz IF bandwidth) for the qubit drives. The qubit drive pulses were amplified by a high-power (35 dB) microwave amplifier (Minicircuits ZV-3W-183) before passing through a 5.5 GHz low-pass filter (Minicircuits LFCN 5500+) to minimise amplifier noise at the readout resonator frequencies.

Most microwave units receive a 10 MHz reference from a microwave generator (Agilent E8257D) via a home-built distribution unit. However, the generators used for driving Q_R and R_R (R&S SGS100A) synchronised directly via a 1 GHz reference. This was critical to achieving the phase stability required to measure R_R Wigner functions during measurement runs lasting up to 40 hours. The frequencies for these two generators were also always set to a multiple of the trigger repetition rate (5 kHz), to ensure a stable phase relationship. For phase-sensitive measurements, a 500 MHz scope (Rigol DS4034) monitored the relative trigger timing between the master and slave AWGs to select consistent delay configurations between the AWG outputs.

Home-built low-noise current sources mounted in a TU Delft IVVI-DAC2 rack provided precision DC bias currents for flux tuning of the qubit frequencies. The DC bias for Q_R was combined with the amplified output of one channel of the master AWG (the same as used for generating Q_R drive pulses) using a reactive bias tee (Minicircuits ZFBT-6GW+). The flux pulses from the AWG were amplified using a home-built 2 V/V flux-pulse amplifier.

3.5.2 Supplementary Note 2: Device fabrication

The device was fabricated using a method similar to that of Ref. [119], but with several specific improvements:

1. The transmon design includes a rounded spacing between the shunt capacitor plates [Fig. 3.8(a)] to avoid the regions of high electric field which can increase sensitivity to interface two-level fluctuators [120].
2. The flux-bias line was centred between the transmon capacitor plates to symmetrise the capacitive coupling with the goal of decoupling the qubits from possible decay-inducing effects of voltage noise fluctuations on the flux-bias lines.
3. As in our previous work [119], the transmon qubits were patterned with niobium titanium nitride (NbTiN) capacitor plates to further reduce susceptibility to noise from two-level fluctuators. Prior to evaporation of the aluminium (Al) junction layers, a short hydrogen-fluoride (HF) dip removed surface oxides to facilitate a good contact between the evaporated Al and NbTiN thin film. To avoid contact problems caused by unwanted etching into the silicon substrate during patterning of the NbTiN, we: 1) optimised the reactive-ion etch (RIE) recipe and duration to minimise the substrate etch and eliminate underetch (under the NbTiN); and 2) introduced a narrow bay in the NbTiN fingers at the contact point to create a softer etch for more reliable contact [Fig. 3.8(c)].
4. The junction development process and double-angle evaporation parameters were optimised to improve the reliability of the very small junction sizes needed for the asymmetric qubit [Fig. 3.8(b)].

3.5.3 Supplementary Note 3: Device operating parameters and qubit performance

Figure 3.9(a) shows the frequencies for the two qubits and three resonators on the device as a function of the applied qubit flux in units of the flux quantum $\Phi_0 = h/2e$, along with the operating points for both qubits during the quantum simulation experiments. Measured device parameters are summarised in Tab. 3.1. Qubit T_1 , T_2^{echo} and T_2^* decay times are shown as a function of qubit frequency in Fig. 3.9(b,c,d).

At the operating point, the Rabi qubit Q_R was designed to sit below the resonator R_R and be pulsed up into resonance with it to avoid continually crossing the resonator with the Q_R 's 1–2 transition during the long flux-pulse sequence. Because of significant protocol times and two operating points, an asymmetric qubit design with two flux-insensitive “sweet” spots was used for Q_R [18], with drive pulses applied at its bottom sweet spot. The first-order flux insensitivity at this point also mitigated some of the impact of rapid, long-range flux-pulsing on the qubit pulse tuning. The maximum and minimum frequencies for Q_R in the final cooldown were 6.670 GHz and 5.451 GHz, respectively.

The asymmetric design also minimised the stringent challenge of targetting the qubit frequency to resonator closely on the scale of the very small coupling frequency. Ideally, the resonator would have been closer to the qubit top sweet spot to maximise phase coherence also

Component	Frequency domain		Time domain	
Q_R	$f_{\max}$	6.670GHz	At operating point:	
	$f_{\min}$	5.451GHz	T_1	20 – 30 μs
	α (asymmetry)	0.68	T_2^{echo}	30 – 60 μs
	$E_C/2\pi$	–281 MHz	T_2^*	20 – 50 μs
	f_{readout}	7.026 GHz		
	$g_{\text{readout}}/2\pi$	43 MHz		
R_R	f	6.381GHz	$T_{1,r}$	3 – 4 μs
	$g_r/2\pi$ (to Q_R)	1.92MHz	$g_r/2\pi$	1.95MHz
	χ_w/π (to Q_W)	–1.26 MHz		
Q_W	$f_{\max}$	5.653GHz	At operating point:	
	f_{exp}	5.003 GHz	T_1	30 – 40 μs
	E_C	—	T_2^{echo}	5 – 7 μs
	f_{readout}	6.940GHz	T_2^*	1.5 – 1.8 μs
	$g_{\text{readout}}/2\pi$	42 MHz	At top sweet spot:	
			T_2^{echo}	30 – 60 μs
			T_2^*	20 – 50 μs

Table 3.1: Measured device parameters and qubit and resonator performance. The coupling strength between Q_R and R_R was measured both by spectroscopy of the avoided crossing, and time-domain measurement of the vacuum Rabi oscillation frequency. For both qubits, Ramsey sequences measured at the sweet spots exhibited beating consistent with quasiparticle tunnelling [15]. T_2^* s reported here were measured by fitting a decaying double sinusoid to a long, beating Ramsey signal and represents the underlying coherence of the qubits. At the operating point for Q_W far from the sweet spot, no beating was observed in the Ramsey measurements.

during the interaction pulses. However, with the asymmetric design, the reduced flux gradient relaxes this constraint. With an asymmetry parameter of $\alpha = (E_J^{\max} - E_{J, \min}) / (E_J^{\max} + E_{J, \min}) \sim 0.68$, the Ramsey time T_2^* for Q_R did not typically drop below a few microseconds, even at the positions with steepest flux gradient.

The asymmetry of Q_R was smaller than targetted, with the result that the bottom sweet spot was also lower in frequency than intended. The ancilla qubit Q_W (a standard symmetric transmon) was therefore operated around 650 MHz below its own maximum-frequency sweet spot of 5.653 GHz. At this operating point, its T_2^* was typically $\gtrsim 1.5 \mu s$. Because we were able to drive Q_W and achieve good photon-sensitive operation at this lower position, we chose not to rapidly tune its frequency up to the sweet spot to perform the photon meter measurements.

To identify the flux operating point that positioned Q_R precisely at the bottom sweet spot, we applied the following procedure. We first decoupled the applied DC qubit fluxes, applying

the appropriate linear correction to compensate for flux cross-talk. Then, after positioning Q_W roughly at its selected operating point, we applied a simple excitation swapping sequence for Q_R with R_R with fixed swap time (near a full swap) and varying amplitudes of positively and negatively directed pulses. Finally, we varied the applied flux on Q_R and identified the operating point as the symmetric flux point where the qubit hit the resonance for positive and negative pulses of equal amplitude. We were able to identify this point to 1 part in 5000. Because the precise choice of operating frequency for Q_W was not critical, any slight shift in frequency due to residual DC cross-talk remaining after the flux decoupling measurements was unimportant.

3.5.4 Supplementary Note 4: Calibration of the flux distortions

Implementing the digital Trotterisation of the Rabi model proposed in Ref. [88] required tuning the qubit frequency with a long series of square interaction pulses. To achieve this, it was necessary to compensate for the filtering effects of electronics and microwave components in the line (Fig. 3.7) [121]. One of the particular challenges of an experiment using a long train (up to 10 μ s) of very short pulses (10–20 ns) is that the system is sensitive to both short- and long-time pulse distortions. These effects included the intrinsic bandwidth of the AWG and the flux-pulse amplifier, the high-pass characteristics of the bias tee, a range of low-pass effects including the Minicircuits and eccosorb filters and filtering from the skin effect of the coaxial cabling, pulse bounces at impedance mismatches, as well as more intangible effects such as transient decays in step responses. Subject to the system operating in a linear regime (e.g., the AWG operating in a comfortable amplitude range), this could be achieved by applying predistortions to the target fluxing sequence.

Figure 3.10 illustrates the calibration process used in this experiment. Rather than building a single, comprehensive model for all flux distortions, we took a divide-and-conquer approach, applying a series of corrections to compensate individual effects. For processes outside the fridge, we calculated the required compensations by directly measuring the system step response using a fast oscilloscope (R&S RTO1024, 10 Gs/s sampling rate and 2 GHz bandwidth). We applied predistortion corrections sequentially, at each step correcting the longest-time behaviour and zooming in to shorter time scales once the longer-time response is successfully corrected. Once measuring through the fridge, we optimised on the shape of the two-dimensional flux-pulse resonance, the so-called “chevron”. Again, we typically focussed initially on correcting the coarse features before zooming in to finer details.

The procedure we used to calculate the external corrections was:

- Sample a measured step response at a period τ : $x[n] = x(n\tau)$.
- Construct the system impulse response function according to: $h[n] = x[n] - x[n-1]$.
- Construct the system transfer matrix H from $h[n]$ (H is a lower-triangular matrix with $h[j]$ in every position on the j^{th} lower diagonal).
- Invert H to find the transfer matrix of the so-called predistortion kernel and calculate the step response of the predistortion kernel as $Hu[n]$, where $u[n]$ is the discrete

Heaviside function. This numerical matrix inversion step limits the length of the step response that can be treated in this way. The sampling period τ is chosen to ensure the sampled step response covers the region of interest.

- Fit the numerically inverted kernel step response using a simple functional form which can then be used to construct a high-resolution predistortion kernel (the impulse response calculated as above from a high-resolution step response). The down-sampling of the step response reduces the fit function dependence on high-frequency effects. For each step, we varied the sampling period to check that the fit parameters were relatively robust to details of the sampling.

Figure 3.10(a) shows the step response from the AWG measured after the home-built flux-pulse amplifier (see Fig. 3.7), with a zoom into the top of the step in (b). In this case, the longest-time response was actually an effectively linear ramp over the long step response. Here, we used a slightly modified procedure to the one above, fitting a linear function directly to the measured step response. Using Laplace transformations, it is possible to show that a step response with a linear ramp, $(1 + \alpha t) u(t)$, can be corrected using a predistortion kernel with an exponentially decaying step response $\exp(-\alpha t) u(t)$. After this linear correction, we then implemented a series of three corrections with “exponential-approach” predistortion step responses of the form $(1 + \alpha \exp(-t/\tau)) u(t)$ with τ values between 5 μs and 500 ns (various amplitudes), determined using the above procedure. Figure 3.10(c) shows the corrected step function measured after applying the four initial corrections. The small but distinct sawtooth structure in the otherwise flat step response is due to the vertical resolution of the AWG.

After correcting for distortions from the AWG and flux-pulse amplifier, we measured the step response after the bias tee, at the fridge input. Figures 3.10(d, e) show the measured step response and sampled predistortional kernel step response calculated using the above procedure (with $\tau = 50$ ns). The high-pass characteristics of a reactive bias tee’s RF input naïvely predict a kernel step response with a full initial step followed by a continually increasing linear voltage ramp. From Fig. 3.10(e), however, it is clear that the kernel step response is not completely linear. We instead fit the step response to a quadratic form and proceed as above. The step response measured after compensating for the bias tee is shown in Fig. 3.10(f).

Inside the fridge, we calibrated the flux-pulse predistortions to optimize the shape of the flux chevron [Figs 3.10(g–i)], which probes the excitation-swapping exchange interaction between qubit Q_R and resonator R_R as a function of flux-pulse amplitude and interaction time. When the qubit is exactly on resonance, the swapping interactions are expected to be slowest and strongest. As it moves off resonance, the oscillations speed up and reduce in amplitude. Interestingly, despite the good performance of the bias-tee correction when measured outside the fridge, the chevron measured with the same corrections [Fig. 3.10(g)] showed a clear ramp in the start of the interaction signal (the lateral skew), consistent with an under-compensated bias tee. We do not understand the cause of this discrepancy, but corrected it

empirically by adjusting the linear coefficient of the bias-tee correction. The chevron measured after optimising this correction (final linear coefficient corresponded to a time constant $\tau = 9.7 \mu\text{s}$) showed the characteristic asymmetric signature of low-pass filtering from the skin effect [Fig. 3.10(h)]. This was corrected by applying a kernel numerically calculated from a step response of the form $(1 - \text{erf}(\alpha_1 \text{ GHz}/21\sqrt{t+1})) u(t)$ [122], using $\alpha_1 \text{ GHz} = 1.7 \text{ dB}$. Finally, we implemented another series of exponential-approach kernels with values of τ between 1500 ns and 30 ns, to achieve the result in Fig. 3.10(i).

3.5.5 Supplementary Note 5: Calibration of the photon meters

Using a photon meter based on a Ramsey sequence's sensitivity to qubit frequency and Q_W 's dispersive frequency dependence on resonator photon number allows detection of average photon number with controllable sensitivity and dynamic range. Suppose the resonator is in the state $\psi = \sum_j \alpha_j |j\rangle$. To implement the photon meter, we apply a Ramsey pair of $\pi/2$ pulses with pulse separation τ on Q_W at a frequency $\Omega_W^d = \Omega_W^0 - d2\chi$, corresponding to the d^{th} photon peak. Different photon-number frequency components accrue different phases during the variable delay between pulses, given by $\theta_j = (j-d)2\chi\tau$. By driving first around σ_x and then around σ_y , the d^{th} photon term ends up on the equator of the Bloch sphere. Measuring the excitation of Q_W then gives a measurement probability

$$P_W^e = \sum_j \frac{|\alpha_j|^2}{2} (1 + \sin \theta_j).$$

Provided τ is chosen such that θ_j is small for all photon components j present in the photon state,

$$\begin{aligned} P_W^e &= \frac{1}{2} \left(1 + \sum_j (j-d)2\chi\tau |\alpha_j|^2 \right), \\ &= \frac{1}{2} (1 + 2\chi\tau(\bar{n} - d)). \end{aligned}$$

Increasing τ therefore increases the sensitivity of measured probability to average photon number, but decreases the accessible range of photon numbers for which the linearity condition $\sin \theta_j \approx \theta_j$ holds. An accurate calibration of the photon meter also requires an accurate calibration of the single-photon dispersive frequency shift 2χ and Q_W 's zero-photon frequency (which determines Ω_W^d). Here, we describe a self-consistent calibration of our photon meters which does not rely on quantities derived from other measurements, such as spectroscopy, and relies primarily on knowing drive-pulse frequencies, probably the most accurate control parameter we have in the experiment. At each stage, we first calibrate Q_W 's zero-photon frequency using a standard Ramsey sequence. With the performance of Q_W at the operating point (dephasing time $T_2^* \sim 1.5 \mu\text{s}$), we routinely achieved frequency accuracy better than 10 kHz.

To calibrate the single-photon dispersive shift [sequence shown in Fig. 3.11(a)], a calibrated SWAP pulse on Q_R transfers an excitation into R_R , before the resonator photon

number is probed via Q_W . The single-photon excitation in R_R dispersively shifts the frequency of Q_W by 2χ . Driving Q_W at the calibrated zero-photon frequency around σ_x and then σ_y , the correct parity condition corresponds to the point where the curve crosses 0.5 excitation probability [Fig. 3.11(d): 383 ns wait time]. This measurement is robust to both the relatively short resonator photon decay time $T_{1,r} \sim 3.5\mu\text{s}$ and the short dephasing time of Q_W at its operating point ($T_2^* \sim 1.5\text{--}1.8\mu\text{s}$ at ~ -650 MHz detuned from its top sweet spot), because these processes both reduce the visibility of the curve, but not the oscillation period, and therefore do not affect the value of the crossing point. The zero-photon frequency calibration is the main limitation, because that calibration limits the accuracy with which the crossing point represents the correct delay time between $\pi/2$ pulses.

The wait time identified above specifies the time between the end of the first pulse and the beginning of the second required to realise a photon parity measurement, but this does not account for the finite pulse duration. To calibrate the effective value of τ , we fix the pulse separation and sweep the frequency of the Q_W drive generator this time without loading any photons into the resonator [Fig. 3.11(b)]. For a pulse separation of 383 ns, the effective τ is ~ 398 ns [Fig. 3.11(e)]. Note that the difference here is not quite the same as the drive pulse width used in the experiment ($4\sigma = 12$ ns). This value of τ is related to the dispersive shift of R_R on Q_W in the usual way: $\tau = \pi/2\chi$, giving $2\chi/2\pi = -1.26$ MHz. Note that, when used directly as a parity meter, the read-out of Q_W was calibrated using a parity pulse pair either with the usual phase on the second pulse, or a phase shifted by π radians. This accounted for the reduced parity visibility from the short T_2^* of Q_W at its operating point and helped to track any fluctuations in the correct parity extremes as a result of drift in qubit frequency and T_2^* .

Figures 2(a, b) show the pulse sequences for two different photon meters used in the experiment, one with the standard Ramsey sequence [calibrations in Figs 3.11(f, g)] and one an unbalanced “echo”-like sequence with an off-centre refocussing pulse (calibrations not shown). The mapping between average photon number and qubit excitation is approximately valid provided the phase advance/delay is less than 30 degrees, which corresponds to a qubit excitation of 0.25. We select the appropriate Ramsey pulse separation by driving the qubit at the frequency corresponding to the mid-point of the desired range (here, the 4-photon position), calculated from the dispersive shift and the calibrated zero-photon frequency, and choosing the separation which gives the target excitation probability of 0.25 [Fig. 3.11(f)], here 4 ns. The effective τ was calibrated, as above, to be ~ 19 ns. Moving to the smaller τ necessary for a higher photon number dynamic range requires frequency refocussing. Ultimately, the main limitation to the range achievable with such a photon meter is set by the bandwidth of the drive pulse.

We used a photon number meter calibrated using the above procedure to follow the excitation-swapping oscillations of a vacuum-Rabi exchange between Q_R and R_R , plotted as a function of the duration of the flux pulse on Q_R [Figure 3.11(h); sequence in Fig. 3.11(a)]. The drifting baseline results from pulsed flux cross-talk between Q_R and Q_W . To correct this, we repeated the same measurement without initially exciting Q_R in order to avoid exciting photons in R_R [Fig. 3.11(c)]. This curve was compensated by adjusting the drive phase of

the second Ramsey pulse in the photon meter (on Q_W), leading to the compensated measurement in Fig. 3.11(j). To maximise the sensitivity of the cross-talk calibration, during the calibration, Q_W can be driven at the zero-photon frequency, which then places the expected “null” measurement result on the equator of the Bloch sphere. A modified version of this procedure can be carried out for all flux-pulse sequences of interest. Note that cross-talk compensation was also necessary to ensure an accurate calibration of the parity condition in Fig. 3.11(d) above.

3.5.6 Supplementary Note 6: Calibration of Wigner tomography

We implement Wigner tomography using the direct method of Ref. [101]. After the algorithm part of the pulse sequence [represented in Fig. 3.12(a) by a swap], a 50 ns square pulse applies a coherent displacement to the resonator photon state before the usual parity readout pulses. The phase-sensitive resonator drive tone is created via single-sideband modulation in an IQ mixer. We calibrate the drive frequency and amplitude using the already calibrated photon meter (Figs 3.12(b, c), respectively). The drive amplitude is calibrated in the middle of the linear range, where we expect the best performance. Figure 3.12(c) illustrates the breakdown of the linear mapping between average photon number and Q_W excitation probability both towards the edge of the linear regime and above the range, as the higher photon components wrap around in phase. In the digital QRM simulation, for phase-sensitive Wigner tomograms (e.g., Figs 3 and 4), it was critical to maintain phase stability between the drives on Q_R and R_R during the measurement. To achieve this, the two microwave generators were synchronised using a 1 GHz reference, with frequencies set as a multiple of the 5 kHz experimental repetition rate.

Figure 3.12 shows one- and two-dimensional Wigner tomograms of a zero-photon (d, e) and one-photon (f, g) state (scaled in terms of photon parity). The maximum visibilities in Figs 3.12(f, g) do not reach the expected values, because these tomograms were measured without an accompanying full set of parity meter calibrations. However, the radial symmetry observed in these tomograms demonstrates the correct behaviour of the coherent resonator drive.

The curves in Figs 3.12(d, g) show fits to the data of a classical mixture of zero-photon and one-photon Wigner function cross-sections, with a free x -axis scaling parameter has been included in the fits. These fits demonstrate that the measured tomograms agree well with theoretical expectations, subject to an x -axis scaling error of $\sim 5\%$. That is, the fits indicate that the amplitude calibrations result in a small systematic overestimate in displacement by 5%. This also agrees with two-dimensional double Gaussian fits of individual frames of the unconditional Wigner movie in Fig. 3(a) of the main text, which give an average Gaussian width $\bar{\sigma} = 0.526 \pm 0.003$, compared with the expected value of 0.5.

3.5.7 Supplementary Note 7: Analog vs digital Jaynes-Cummings dynamics

Simple modelling of the Trotterised version of the full Rabi model shows that high-quality simulations require both slow dynamics and short Trotter steps (i.e., fast flux pulsing). Such

an experiment is sensitive to both short-time and long-time effects in the flux-pulse shaping. A simpler experiment which verifies the performance of this flux pulsing is to implement a digital simulation of the standard Jaynes-Cummings (JC) interaction underlying the standard excitation-swapping experiments demonstrated with single flux pulses (Fig. 3.10).

In the standard continuous-wave (single-pulse) version of a JC excitation-swapping interaction, resonance between the qubit and resonator frequencies gives rise to maximum visibility oscillations of the excitation moving between the two components. When detuned, the different phases accrued by the qubit and resonator during the interaction decrease the oscillation visibility, while increasing the oscillation frequency. This gives rise to the characteristic shape of the flux chevron. Significant care is required, however, to accurately reproduce the (analog) JC interaction with a digital pulse train.

Figures 3.13(a, b) show analog and digital versions of the JC interaction (viewed through the qubit excitation) under otherwise identical conditions. The digital chevron shows a series of resonances which do not appear in analog measurements (not shown), and there is also no chevron visible at the natural resonance condition around 2.45 Vpp.

The new features relate to the extra “interaction off” times in the digital version. The regular spacing between neighbouring satellite resonances is around 50MHz (after converting AWG amplitude to qubit frequency), which is the inverse pulse duration. During the interaction time, the qubit-resonator relative phase evolves as expected. However, in the “off” time between interaction pulses, the qubit accrues phase at a different rate, and will hence not have the required phase at the beginning of the next pulse for the interaction to pick up where it left off at the end of the previous pulse. Therefore, the necessary condition for observing a chevron feature at exactly the position of the natural resonance is that the qubit phase accrued (relative to the resonator) during the “off” time should be a multiple of 2π . The observation of multiple satellite resonances is a form of digital aliasing, where the interaction will build up constructively from pulse to pulse provided the relative phase accrued between qubit and resonator during the “on” time of the pulse again differs only by an integer multiple of 2π . However, this is an aliasing of the dynamics itself, not just an aliasing of the measurement, which could also occur in natural continuous-wave (CW) chevrons and would never lead to the observation of extra satellite peaks.

This pulsed interaction can also be viewed as a Trotterised simulation of the CW interaction. While successive interaction pulses obviously commute with each other, they do not necessarily commute with the “off” pulses. The condition on qubit-resonator phase during the “off” pulse can be understood as the condition where the Trotter error vanishes, because the Hamiltonian term resulting from the qubit detuning coincides with the identity. The satellites arise because the phase contribution from the qubit detuning in the “on” pulse is identical if the frequency change matches a multiple of 2π phase.

To compensate for the phase error accrued in the qubit during the “off” pulses, we apply a 5 ns compensation flux pulse between interaction pulses. Using the flux-pulse amplitude which corresponds to the centre of the CW chevron, the amplitude of the compensation pulse was swept to identify the correct compensation point. In this way, very good agreement was achieved between the digital JC dynamics and the traditional analog version [Figs 3.13(d–

i)). The main differences are a slightly reduced visibility because of the increased experiment time, and a slightly lower effective coupling frequency ($g/2\pi \sim 1.8\text{MHz}$, instead of $\sim 1.95\text{MHz}$). The latter most likely arises from residual short-time pulse imperfections which do not contribute significantly to the long interactions in the analog form.

3.5.8 Supplementary Note 8: Trotter simulation with excited and ground initial states

In the degenerate-qubit case, when understood in terms of the cavity trajectories in phase space, it is clear that the structure of the dynamics of the full Rabi model with USC should not depend on whether the qubit starts in the ground or excited state. This contrasts with the JC interaction, where the $|g, 0\rangle$ state is decoupled from the rest of the system and the system will only undergo nontrivial dynamics if an initial excitation is loaded in the system. Indeed, in a natural USC system, if it were possible to turn the coupling on and off rapidly, it would be extremely interesting to watch an uncoupled-system ground state evolve into a state with excitations in the qubit and cavity. In this digital simulation, however, this is less satisfying, since the protocol in any case involves regularly injecting excitation into the system in the form of qubit flipping pulses. Most of the results reported here therefore take the more conservative position of initialising the system with an excitation, with the motivation that observing a difference between the simulated dynamics and what would be expected in a weak-coupling scenario could then only result from the simulated counter-rotating terms. Although there were some stability issues during the measurement with ground-state initialisation, there is nevertheless extremely good agreement between the two cases, for example with the timing of the revivals in both cases agreeing with the theoretical predictions. For this particular measurement of ground-state initialisation, qubit revivals are observed even out to $r \equiv g^R/\omega_Q^R \sim 1$.

3.5.9 Supplementary Note 9: Trotterisation performance vs Trotter order

As discussed already, initial modelling of a Trotterised Rabi simulation showed that unusually low qubit-resonator coupling between Q_R and R_R was required to be able to achieve reasonable simulation fidelities given the hard bandwidth limitations of flux-based fast frequency tuning. This, however, required longer experimental times for the simulations, which in turn placed significant constraints on qubit and resonator coherence. Indeed, the shorter-than-anticipated resonator coherence time proved to be the biggest limitation. As a result, it was critical to use all available measures to minimize the Trotter error in our simulations, given the limits on the shortest achievable Trotter step sizes.

The accuracy of the Trotter approximation is set by the amount of non-commutativity between different components in the step [68]. While first-order Trotterisations [$\exp(A + B) \approx \exp(A)\exp(B)$] lead to Trotter errors that scale with single commutators (quadratically with simulation time), higher-order Trotterisations can be used to eliminate lower orders of Trotter error. For example, the symmetry of a second-order Trotterisation [$\exp(A + B) \approx \exp(A/2)\exp(B)\exp(A/2)$] ensures that first-order error terms (related to single commutators) cancel, pushing the largest Trotter error terms out to third order in simulation time. For

two-part Hamiltonians, however, second-order Trotterisation in practice only involves modifying the pulses in the first and last Trotter steps. All the results in the main text were obtained using a second-order Trotterisation. The plots in Fig. 3.15 illustrate that this was absolutely critical in order to extend the simulations deep into the ultrastrong coupling regime. The first-order and second-order Trotterisation agree reasonably well at $r < 0.5$, but behave fundamentally differently at the higher values. The first-order simulation starts to show qualitatively different behaviour for relative coupling strengths $r \gtrsim 0.5$. In particular, only in the second-order case are the characteristic plateaus and revivals of the USC regime observable.

3.5.10 Supplementary Note 10: Trotterisation performance vs Trotter step size

As illustrated in Fig. 3.15, the effects of Trotter error are most visible in the high r regimes, which is reasonable, considering that for low r , the Rabi model is well approximated by the JC model where the excitation-nonconserving terms (non-commuting with the excitation-conserving terms) do not play a significant role. This was also visible when studying the performance of the simulation as a function of the Trotter step size.

Measurements and numerical simulations show significant reduction in Trotter error as the number of Trotter steps over $1.2 \mu\text{s}$ increased from 24 to 60. The Trotter error shows up in two ways, namely the central features departing from the expected plateaus, and a tendency for the dynamical landscape to “break apart”, even out into the lower coupling regimes. In the measured results and the simulation with decay, the fine details do not appear as strongly, but the effect appears to wash out the oscillation dynamics more rapidly. Only at the smallest step size are these effects absent from the measured results, and in the ideal simulations (without decoherence) there are even then central features which only disappear at a still smaller 10 ns step size. The measured results agree very closely with the numerical Trotter dynamics which include only the effect of photon decay, again highlighting that the primary limiting factor in our experiments was T_1 , r . It is clear from these results that moving towards the smallest possible Trotter steps will be a key challenge for reaching quantum supremacy in complex quantum simulations.

3.5.11 Supplementary Note 11: Qubit entropy dynamics

In the Rabi model, as the resonator states separate, the qubit-resonator entanglement causes the reduced qubit state to collapse towards the maximally mixed state. A revival occurs in the qubit purity only if the underlying entanglement is still present when the resonator states re-coalesce at the origin in phase space. While many possible uninteresting effects may cause an initial collapse in qubit purity, a revival in purity is a signature of entanglement with another system, in this case the resonator. After each Trotter step, a tomographically complete set of measurements on Q_R was used to reconstruct its reduced state using maximum-likelihood tomography. We use the von Neumann entropy to characterise the purity of the reduced qubit state and observe revivals in qubit purity out to $r > 0.8$ [Fig. 3.17(a)], consistent with the observed revivals in qubit parity. While the observed revivals shown in the slices [Fig. 3.17(b)] appear smaller than the qubit parity revivals, in fact this is deceiving, resulting from the fact

that purity (as with other entropy measures) is a quadratic function of the qubit population difference. The inset shows that the background noise of this signal is small and that the revivals are quite distinct. Moreover, plotting an appropriate square root of the entropy (not shown) shows that the revivals are consistent with the qubit parity case.




Figure 3.3: Photon number dynamics of the quantum Rabi model in the degenerate-qubit case. (a, b) Average photon number is probed by applying Ramsey and echo-like pulses to Q_W . The effective Ramsey pulse separation τ determines the photon dynamic range. Because of finite pulse widths, reaching the small τ needed for high dynamic ranges (b) requires an unbalanced “echo”-like sequence. (c, e) Measured photon number dynamics up to 60 Trotter steps using a low-dynamic-range (LDR) photon number meter ($\tau \sim 18.7$ ns) with a linear range of ~ 0 –8 photons [indicated by grey regions in (e)]. Large photon populations in the resonator highlight the non-conservation of excitation number in the quantum Rabi model. The resonator displays clear oscillations up to $r > 1.8$ in good agreement with the expected qubit revival times (dashed curves). The red feature in the middle reflects the upper limit on the number meter’s dynamic range set by Q_W “population wrapping” at high photon numbers. (d) Measured photon dynamics up to 90 Trotter steps using a high-dynamic-range (HDR) number meter with $\tau \sim 6.5$ ns and a linear range of ~ 0 –20 photons, allowing observation of photon oscillations beyond 1.5μ s of simulated time (more than 75 Trotter steps). This data shows the effect of a residual Kerr nonlinearity at high values of r . (e) Line slices are plotted for evenly spaced resonator-coupling frequency ratios between the red and blue dashed lines shown in (c). Grey regions delineate the linear range of the number meter. (f) Maximum measured average photon number for each value of r for both LDR and HDR number meters.




Figure 3.4: Photon dynamics in phase space in the DSC regime (degenerate-qubit case) from maximum-likelihood Wigner tomography. (a) Selected frames from a “movie” (measured over ~ 40 hours) showing the phase-space evolution of the resonator reduced state for $r \sim 0.9$ (frames labelled by Trotter step n and simulated time), with the final panel showing the full trajectories determined from 2D double-Gaussian fits to the raw data [the full movie is provided in the Supplementary Movie 1 of the published work]. Plotted tomograms are maximum-likelihood reconstructions of direct Wigner tomography measured data with a systematic phase correction [see Methods]. When the effective drive on the intracavity field created by the Rabi interaction has a strength comparable to the resonator’s natural frequency (i.e., $g^R \sim \omega_f^R$), this drive is able to create a significant displacement of the cavity field before the phase-space rotation caused by ω_f^R brings the field back towards the origin. This effect is observed clearly here in the creation of two well-resolved, rotating peaks and subsequent re-coalescence which are characteristic signatures of DSC dynamics. Deviation from the ideal circular trajectories (orange curves) arises from photon decay. The measured trajectory shows excellent agreement with a numerical Trotter simulation at $g^R/2\pi = 1.79$ MHz which includes resonator $T_1, r = 3.5 \mu\text{s}$ (green curves). From the fits, we calculate an estimated Wigner function width $\sigma = 0.526 \pm 0.003$, instead of the predicted 0.5, indicating a displacement calibration error of $\sim 5\%$ [Supplementary Note 6]. Background noise arises from phase instability of microwave sources and frequency stability of the Wigner qubit over the long measurement. (b-e) Conditional phase-space evolution illustrated by the resonator Wigner function for different initial states of Q_R : (b) $|0\rangle$, (c) $|1\rangle$, (d) $|+\rangle$ and (e) $|-\rangle$. The phase-space trajectory of R_R depends on the qubit state in the σ_x basis, consistent with creation of Bell-




Figure 3.5: **Nonclassical resonator cat states from conditioned DSC-driven entanglement (degenerate-qubit case).** The plots show Wigner functions of nonclassical Schrödinger cat states in the Rabi resonator, reconstructed from maximum-likelihood state tomography for two different DSC coupling strengths with $g^R/\omega_r^R \sim 0.9$ (top, $n = 10$ Trotter steps) and $g^R/\omega_r^R \sim 2.1$ (bottom, $n = 8$ Trotter steps), conditioned on measuring Q_R in $|0\rangle$ (left) and $|1\rangle$ (right). The regions of negativity and visibility of several fringes between the well-resolved coherent state peaks are clear signatures of nonclassicality in the Rabi field mode and demonstrates the coherence and entanglement of the underlying qubit-resonator state. Combined with the qubit conditioning shown in Fig. 3.4, observing clear cat states for both outcomes of the Q_R measurement is a clear signature of coherent DSC dynamics.




Figure 3.6: **Measured and numerical quantum Rabi model qubit dynamics for nondegenerate qubit frequency.** The cases implemented are $g^R/\omega_q^R \sim 4$ (top), ~ 2 (middle) and ~ 1 (bottom), with the plots showing measured qubit dynamics (left), numerically simulated dynamics of a Trotterised QRM with the measured $T_{1,r} \sim 3.5 \mu\text{s}$ included (centre), and ideal Rabi dynamics (right). The results illustrate that the nondegenerate-qubit dynamics do not deviate significantly from the degenerate-qubit case in the regime where $\omega_r^R \gg \omega_q^R$. The measured dynamics exhibit many qualitative features in good agreement with the ideal QRM and show excellent agreement with the numerical Trotter simulation with decay, indicating that the fidelity of the measured results to the ideal case is limited primarily by resonator decay.




Figure 3.7: Experimental schematic showing the connectivity of microwave electronics and components in and outside the dilution refrigerator. The sample mounted below the mixing chamber typically remained at around 30 mK. Qubit and resonator drive lines and flux-bias lines were thermalised and attenuated at the 4 – K and 30 – mK stages and were low-pass filtered before arriving at the sample. The qubits and resonator drive pulses were generated by AWGs and IQ mixers. Home-built low-noise current sources provided DC bias currents for qubit frequency tuning, which were combined with fast frequency-tuning bias pulses using reactive bias tees. AWG markers provided the gating for pulse-modulated measurement pulses.




Figure 3.8: SEM images of a sister device with added false colour. (a) Rabi qubit (Q_R) with coupling to the Rabi resonator (R_R , above) and readout resonator (below), showing the centred flux-bias line and displaced SQUID loop. Q_R is coupled to R_R near its shorted end in order to achieve the required small coupling g . (b, c) Josephson junctions are contacted to the NbTiN SQUID loop fingers using small bays to achieve better contact. In (b), it is possible to see the large asymmetry in junction size, with a zoom on the small junction in (c).




Figure 3.9: Schematic showing measured spectral arrangement for the digital Rabi quantum simulator device and the qubit coherence times. (a) Measured data for the 0–1 transition of the Rabi qubit Q_R (green curve) and the Wigner qubit Q_W (blue curve) are plotted as a function of applied flux in units of Φ_0 . Also shown are the frequencies of the Rabi resonator R_R (red: $\omega_r = 6.381\text{GHz}$) and the readout resonators for Q_R (green dashed: $\sim 7.03\text{GHz}$) and Q_W (blue dashed: $\sim 6.94\text{GHz}$). The operating points of the qubits for the Trotter simulation are given by the green and blue dotted lines for Q_R and Q_W , respectively. (b, c, d) Time constants measured for Q_R (green) and Q_W (blue) for (b) T_1 , (c) T_2^{echo} and (d) T_2^* . Note that, at the sweet spots, measured qubit T_2^* times here are limited by slow frequency-switching processes in the qubits such as quasiparticle tunnelling [15].




Figure 3.10: Calibration of the flux distortions. (a,b) Step response of the amplified AWG flux channel output with a zoom in (b), measured using a fast oscilloscope. (c) Corrected step response achieved using one linear response correction and three exponential decay corrections with parameters (τ, α) : $(5.1\mu\text{s}, 0.0012)$, $(670\text{ns}, 0.015)$ and $(520\text{ns}, -0.00037)$ (see text for details). (d, e) Measured step response (d) and numerically calculated predistortion step response (e) after the bias-tee. (f) Corrected step response achieved using a quadratic bias tee correction (see text for details). (g) Distorted flux “chevron” measured with the corrections applied in (e). (h) Dramatically improved chevron obtained after sweeping one parameter in the bias-tee correction (that corresponding to the standard RC time constant). The asymmetric signature observed here is characteristic of the low-pass filtering effect produced by the skin effect in the coaxial cables. (i) A well-compensated chevron obtained after applying a correction for the skin effect and several more exponential decay corrections with (τ, α) : $(350\text{ns}, -0.0063)$, $(600\text{ns}, -0.0037)$, $(1500\text{ns}, -0.002)$, $(100\text{ns}, -0.0017)$ and $(30\text{ns}, 0.0036)$.




Figure 3.11: Calibration of the photon meter. (a–c) Measurement sequences used for calibrating the parity meter, specifically: (a) the dispersive shift of R_R on Q_W , (b) the effective delay time τ corresponding to a particular pulse separation, and (c) high-frequency flux cross-talk between flux pulses on Q_R and the flux offset of Q_W . (d) Calibrating the parity condition, identified as the first crossing point of a Ramsey experiment with one photon in the resonator, giving a pulse separation of 383 ns. (e) Calibrating the effective delay time τ for a particular pulse separation. Using parity pulses separated by 383 ns, we calibrated the effective separation τ to be 398 ns, corresponding to a dispersive shift $2\chi/2\pi = -1.26$ MHz. (f) Configuring an average photon number meter for a specific dynamic range of 0–8 photons. Driving at the midpoint of the 0–8 photon frequency range, the Ramsey pulse separation is chosen to lie on the edge of the linear region. For 0–8 photons, we chose to use a separation of 4 ns. (g) Calibrating the photon meter effective τ . Repeating the measurement described in (e), the effective pulse delay for a 4 ns separation was ~ 19 ns. Comparing the oscillation period of the curves in (e) and (g) highlights the different sensitivity of the two photon meters. (h–j) Calibrating high-frequency flux cross-talk. The flux cross-talk is calibrated by measuring the photon meter without loading excitations into the resonator and corrected by adjusting the phase of the second photon meter pulse.




Figure 3.12: Calibration of Wigner tomography. (a) Pulse sequence used to make the displaced photon parity measurement which provides a direct measurement of the Wigner function at a particular position in phase space. (b) This plot shows the response of R_R to the drive pulse as a function of drive frequency, as recorded by the Q_W photon meter, centred at 6.3814 GHz, with a FWHM of ~ 21 MHz, in reasonable agreement with the 18 MHz expected for a 50 ns square pulse. (c) The pulse displacement amplitude is also calibrated using a low-dynamic-range photon meter with a linear range of 0–8 photons. We fit the data in the centre of the linear range, where the photon meter mapping is most accurate, with a function of the form $\langle n \rangle = k_A A^2$, finding $k_A = 9.10$. (e, f) Measured direct Wigner tomograms of zero-photon (e) and one-photon (f) states (one-photon state prepared using a calibrated SWAP pulse between Q_R and R_R). (d, g) Direct Wigner tomogram slices of zero-photon (d) and one-photon (g) states measured using the full parity meter calibrations.




Figure 3.13: Comparison of analog and digital versions of a Jaynes-Cummings interaction. (a) Standard analog JC chevron showing the resonant excitation swapping between qubit and resonator after the qubit is initialised in the excited state as a function of flux-pulse amplitude (x axis) and duration (y axis) (qubit-resonator detuning and interaction time, respectively). The x -axis location of the chevron (~ 2.445 Vpp) therefore defines the qubit-resonator on-resonance condition. (b) Digital JC chevron (measured under otherwise identical conditions) using a pulse duration of 20 ns showing a series of equally spaced resonances with different apparent interaction strengths. (c) We scan amplitude of a 5 ns compensation flux pulse to identify the value which enforces that the digital chevron is centred around the natural resonance position. (d-f) Standard analog and (g-i) digital JC chevrons measured by probing: (d, g) the excited state probability for Q_R , (e, h) the average photon number in R_R (linear range 0–2 photons), and (f, i) the photon parity of R_R .




Figure 3.14: Comparison of the simulation between initialising in the excited (left) versus the ground state (right). (a, b) These plots directly verify the symmetrical behaviour of the simulated Rabi model. (c,d) Line slices are plotted at evenly spaced frequencies between the red and blue dashed lines in (a, b). Arrows in (c, d) show the expected time for the first revival.




Figure 3.15: Comparison of simulation performance for asymmetric, first-order (a–d) and symmetric, second-order (e–h) Trotterisation. (a, e) Pulse sequences for the first-order (a) and second-order (e) Trotterisation. (b, f) Numerical simulations of the Trotterised Rabi model for the ideal case with no decay. Note that the sharp features in the centre of the plots (deep in the ultrastrong coupling regime) are not artifacts of the numerics, but Trotter error related to the 20 ns step size (these features disappear for 10 ns pulses). (c, g) Experimental quantum simulations for first-order (c) and second-order (g) Trotterisation, showing very good agreement with the numerical results in (b, f). (d, h) Vertical line slices are plotted for evenly spaced resonator frequencies between the red and blue dashed lines in plots (c) and (g).




Figure 3.16: Comparison of simulation performance for various Trotter step sizes, showing measurements (left), numerical simulations with no decay (middle) and numerical simulations with the measured $T_{1,r} = 3.5\mu\text{s}$: (a) 20 ns steps (60 Trotter steps), (b) 30 ns steps (40 Trotter steps), (c) 40 ns steps (30 Trotter steps) and (d) 50 ns steps (24 Trotter steps).




Figure 3.17: Tomography of the reduced state of qubit Q_R as a function of simulation time and relative resonator-coupling frequency. (a) Image plot showing the dynamics of qubit quantum von Neumann entropy over different ultrastrong-coupling regimes. (b) Line slices are plotted at evenly spaced frequencies between the blue and red dashed lines. Inset: Zoom showing revivals.

EXPERIMENTAL ERROR MITIGATION VIA SYMMETRY VERIFICATION IN A VARIATIONAL QUANTUM EIGENSOLVER

Variational quantum eigensolvers offer a small-scale testbed to demonstrate the performance of error mitigation techniques with low experimental overhead. We present successful error mitigation by applying the recently proposed symmetry verification technique to the experimental estimation of the ground-state energy and ground state of the hydrogen molecule. A finely adjustable exchange interaction between two qubits in a circuit QED processor efficiently prepares variational ansatz states in the single-excitation subspace respecting the parity symmetry of the qubit-mapped Hamiltonian. Symmetry verification improves the energy and state estimates by mitigating the effects of qubit relaxation and residual qubit excitation, which violate the symmetry. A full-density-matrix simulation matching the experiment dissects the contribution of these mechanisms from other calibrated error sources. Enforcing positivity of the measured density matrix via scalable convex optimization correlates the energy and state estimate improvements when using symmetry verification, with interesting implications for determining system properties beyond the ground-state energy.

4.1 Introduction

Noisy intermediate-scale quantum (NISQ) devices [123], despite lacking layers of quantum error correction (QEC), may already be able to demonstrate quantum advantage over classical computers for select problems [124, 125]. In particular, the hybrid quantum-classical variational quantum eigensolver (VQE) [54, 126] may have sufficiently low experimental requirements to allow estimation of ground-state energies of quantum systems that are difficult to simulate purely classically [50, 127–129]. To date, VQEs have been used to study small examples of the electronic structure problem, such as H_2 [30, 31, 55, 56, 59, 130], HeH^+ [54, 131], LiH [31, 55, 130], and BeH_2 [31], as well as exciton systems [132], strongly correlated magnetic models [130], and the Schwinger model [133]. Although these experimental efforts have achieved impressive coherent control of up to 20 qubits, the error in the resulting estimations has remained relatively high due to performance limitations in the NISQ hardware. Consequently, much focus has recently been placed on developing error mitigation techniques that offer order-of-magnitude accuracy improvement without the costly overhead of full QEC. This may be achieved by using known properties of the target state, e.g., by checking known symmetries in a manner inspired by QEC stabilizer measurements [57, 134], or by expanding around the experimentally-obtained state via a linear (or higher-order) response framework [135]. The former, termed symmetry verification (SV), is of particular interest because it is comparatively low-cost in terms of required hardware and additional measurements. Other mitigation techniques require understanding the underlying error models of the quantum device, allowing for an extrapolation of the calculation to the zero-error limit [53, 136, 137], or the summing of multiple calculations to probabilistically cancel errors [136, 138, 139].

In this Rapid Communication, we experimentally demonstrate the use of SV to reduce the error of a VQE estimating the ground-state energy and the ground state of the H_2 molecule by one order of magnitude on average across the bond-dissociation curve. Using two qubits in a circuit QED processor, we prepare a variational ansatz state via an exchange gate that finely controls the transfer of population within the single-excitation subspace while respecting the underlying symmetry of the problem (odd two-qubit parity). We show that SV improves the energy and state estimates by mitigating the effect of processes changing total excitation number, specifically qubit relaxation and residual qubit excitation. We do this through a full density-matrix simulation that matches the experimental energy and state errors with and without SV, and then using this simulation to dissect the contribution of each error source. Finally, we explore the limitations of SV arising from statistical measurement noise, and find that enforcing the positivity of the fermionic 2-reduced density matrix ties the improvement in energy estimation from SV to the improvement in ground-state fidelity (which was previously not the case).

4.2 The Variational Eigensolver

A VQE algorithm [54, 126] approximates the ground state $\rho^{(0)}$ of a Hamiltonian $\hat{H}$ by a variational state $\rho^{(\text{raw})}(\vec{\theta})$, with $\vec{\theta}$ a set of parameters that control the operation of a quantum device. These parameters are tuned by a classical optimization routine to minimize the varia-

tional energy $E(\vec{\theta}) = \text{Tr}[\rho^{(\text{raw})}(\vec{\theta})\hat{H}]$. In practice, this is calculated by expanding $\rho^{(\text{raw})}(\vec{\theta})$ and $\hat{H}$ over the N -qubit Pauli basis $\mathbb{P}^N := \{I, X, Y, Z\}^{\otimes N}$,

$$\rho^{(\text{raw})}(\vec{\theta}) = \frac{1}{2^N} \sum_{\hat{P} \in \mathbb{P}^N} \rho_{\hat{P}}^{(\text{raw})}(\vec{\theta}) \hat{P}, \quad \hat{H} = \sum_{\hat{P} \in \mathbb{P}^N} h_{\hat{P}} \hat{P}, \quad (4.1)$$

where the Pauli coefficients are given by $\rho_{\hat{P}}^{(\text{raw})}(\vec{\theta}) = \text{Tr}[\hat{P}\rho^{(\text{raw})}]$. The variational energy may then be calculated as

$$E^{(\text{raw})}(\vec{\theta}) = \sum_{\hat{P} \in \mathbb{P}^N} \rho_{\hat{P}}^{(\text{raw})}(\vec{\theta}) h_{\hat{P}}. \quad (4.2)$$

For example, consider the H_2 molecule studied in this work. Mapping the Hamiltonian of this system (in the STO-3G basis) onto four qubits via the Bravyi-Kitaev transformation [140] and then further reducing dimensions by projecting out two non-interacting qubits [30] gives

$$\begin{aligned} \hat{H}_{H_2} = & h_{II}II + h_{ZI}ZI + h_{IZ}IZ \\ & + h_{XX}XX + h_{YY}YY + h_{ZZ}ZZ, \end{aligned} \quad (4.3)$$

where coefficients $h_{\hat{P}}$ depend on the interatomic distance R . These coefficients may be determined classically using the OpenFermion [52] and psi4 [141] packages. The Pauli coefficients $\rho_{\hat{P}}^{(\text{raw})}$ of the density matrix $\rho^{(\text{raw})}$ are extracted by repeated preparation and (partial) tomographic measurements of the ansatz state. As one only needs those Pauli coefficients $\rho_{\hat{P}}^{(\text{raw})}$ with non-zero corresponding Hamiltonian coefficients $h_{\hat{P}}$, one need not perform full tomography of $\rho^{(\text{raw})}$. However, in a small-scale experiment, full state tomography of $\rho^{(\text{raw})}$ may still be feasible, and may provide useful information for the purposes of benchmarking. In particular, the fidelity of $\rho^{(\text{raw})}$ to $\rho^{(0)}$,

$$F^{(\text{raw})} = \text{Tr}[\rho^{(\text{raw})}\rho^{(0)}], \quad (4.4)$$

is a more rigorous measure of the ability to prepare the ground state than the energy error,

$$\Delta E^{(\text{raw})} = \text{Tr} \left[\left(\rho^{(\text{raw})} - \rho^{(0)} \right) \hat{H} \right]. \quad (4.5)$$

Error mechanisms such as decoherence pull $\rho^{(\text{raw})}$ away from $\rho^{(0)}$, decreasing F and increasing ΔE .

4.3 Mitigation by Symmetry Verification

These errors may be mitigated by using internal symmetries $\hat{S} \in \mathbb{P}^{N-1}$ of the target problem, such as parity checks [57, 134]. These checks project $\rho^{(\text{raw})}$ to a symmetry verified matrix $\rho^{(\text{SV})}$ that lies in the $\langle \hat{S} \rangle = s$ subspace of the symmetry. This projection could be performed

¹As described in Refs. [57, 134], one does not require $\hat{S}$ to be a Pauli operator, however this makes the SV procedure significantly simpler.

via direct measurement of $\hat{S}$ on the quantum device, but one may instead extract the relevant terms of the density matrix $\rho^{(\text{SV})}$ in post-processing:

$$\rho_{\hat{P}}^{(\text{SV})} = \frac{\rho_{\hat{P}}^{(\text{raw})} + s\rho_{\hat{S}\hat{P}}^{(\text{raw})}}{1 + s\rho_{\hat{S}}^{(\text{raw})}}, \quad (4.6)$$

The right-hand side may be obtained by partial tomographic measurement of the ansatz state, with at most twice the number of Pauli coefficients that need to be measured. This upper bound is not always achieved. For example, the $\hat{H}_{\text{H}_2}$ Hamiltonian has a $\hat{S} = ZZ$ symmetry, which maps the non-zero Pauli terms in $\hat{H}_{\text{H}_2}$ to other non-zero Pauli terms in $\hat{H}_{\text{H}_2}$. Symmetry verification in this problem then does not require any additional measurements to estimate $E^{(\text{SV})}$ beyond those already required to estimate $E^{(\text{raw})}$. Even when it does require additional measurements, SV remains attractive because it does not require additional quantum hardware or knowledge of the underlying error model. One can show that the SV state $\rho^{(\text{SV})}$ may be equivalently obtained via a variant of the quantum subspace expansion (QSE) [135], suggesting an alternative name of S-QSE [57].

One may further minimize the error in a quantum algorithm by tailoring the quantum circuit or the gates within. In a VQE, one wishes to choose a variational ansatz motivated by the problem itself [30, 142] while minimizing the required quantum hardware [31]. To balance these considerations, we suggest constructing an ansatz from an initial gate-set that is relevant to the problem at hand. For example, in the electronic structure problem, the quantum state is generally an eigenstate of the fermion number. When mapped onto qubits, this often corresponds to a conservation of the total qubit excitation number. Gates such as single-qubit Z rotations, two-qubit C-Phase [22], and two-qubit iSWAP [143] gates preserve this number, making these gates a good universal gate set (within the target subspace [144]) for quantum simulation of electronic structure. In the example of H_2 , the total two-qubit parity (ZZ) is indeed conserved and the ground state at any R may be generated by applying to $|01\rangle$ or $|10\rangle$ an exchange gate

$$U_\theta = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \cos \theta & i \sin \theta & 0 \\ 0 & i \sin \theta & \cos \theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (4.7)$$

with R -dependent optimal exchange angle θ and a follow-up phase correction on one qubit.

We now experimentally investigate the benefits of SV in the VQE of H_2 using two of three transmon qubits in a circuit QED quantum processor (see details in Supplementary Notes). The two qubits (Q_0 and Q_1) are coupled by a common bus resonator, and have dedicated microwave lines for single-qubit gating, flux bias lines for local and ns-scale control of their frequency, and dedicated readout resonators coupling to a common feedline for independent readout by frequency multiplexing. We prepare the ansatz state with an efficient circuit [Fig. 4.1(a)] that first excites Q_1 with a π pulse to produce the state $|10\rangle$, and then applies a square flux pulse of fixed duration and amplitude to Q_0 , bringing it into or near resonance with Q_1 to coherently exchange the excitation population. A plot of population exchange as a function of flux-pulse amplitude and duration [Fig. 4.1(b)] reveals the expected




Figure 4.1: Quantum circuit and energy landscape of the variational eigensolver. (a) Quantum circuit for generating and measuring the variational ansatz state. (b) Coherent excitation exchange, produced as Q_0 is fluxed into resonance with Q_1 by a square flux pulse of fixed amplitude (x axis) and duration (y axis). The amplitude controls the frequency to which Q_0 is pulsed (~ 1.428 V bringing it on resonance with Q_1). (c) Zoom-in of (b) into the region used in the experiment to control the exchange of population between Q_0 and Q_1 . Colored lines illustrate the combinations of square-pulse amplitudes and duration used to achieve fine adjustment of $\tilde{\theta}$. (d) Excitation of Q_0 for the combinations of pulse amplitudes and duration marked by colored lines in (c), showing the matching of the experimentally-defined $\tilde{\theta}$ to the target θ defined in Eq. (4.7) (black dashed curve). Colors [matching (c)] correspond to pulse duration. (e) Landscape of energies $E^{(raw)}(\tilde{\theta}, R)$ as a function of the experimentally-defined $\tilde{\theta}$ angle and the interatomic distance R .

chevron pattern that is the hallmark of coherent population exchange between the two qubits, albeit with some asymmetry arising from the bandwidth limitation of the flux-control line. We make use of the square-pulse duration (1 ns resolution) and amplitude (0.5 mV resolution) as coarse and fine knobs, respectively, to control population exchange. We choose 1500 combinations of pulse duration and amplitude settings to parametrize an experimental knob $\tilde{\theta}$ [Fig. 4.1(d)] capable of finely controlling population exchange like θ in Eq. (4.7) over the range $[0, \pi/4]$ [Fig. 4.1(c)]. The circuit concludes with simultaneous pre-rotation gates on both qubits followed by simultaneous measurement of both qubits, in order to perform tomography of the prepared ansatz state. To fully reconstruct the state, we use an overcomplete set of 36 pre-rotation pairs and extract estimates of the average measurement for each qubit as well as their shot-to-shot correlation using N_{meas} measurements per pre-rotation. Note

while the flux pulse implements the exchange gate of Eq. (4.7) with additional single-qubit phase rotations, the correction for these phase rotations can be performed virtually from the fully-reconstructed state.

4.4 Results

We now optimize the VQE to approximate the ground-state energy and ground state of H_2 . At each chosen R , we employ the covariance matrix adaptation evolution strategy (CMA-ES) optimization algorithm [145], using $E^{(\text{raw})}$ as cost function and $\tilde{\theta}$ as single variational parameter. The evolutionary strategy optimizes $\tilde{\theta}$ over repeated generations of $N_{\text{pop}} = 10$ samples of $E^{(\text{raw})}(\tilde{\theta})$, each calculated from a raw density matrix $\rho^{(\text{raw})}$ using linear inversion of $N_{\text{meas}} = 10^3$ ². Optimizations have a hard-stop criterium of [Fig. 4.2(a) inset] 20 generations (~ 2 hours). The converged state is finally reconstructed with greater precision, using $N_{\text{meas}} = 10^5$. Figure 2 shows the resulting energy estimate for twelve values of R and the reconstructed optimized state at three such distances. These tomographs show that the optimal solutions are concentrated in the single-excitation subspace of the two qubits, with two-qubit entanglement increasing as a function of R .

Performing the described symmetry verification procedure on the converged states shows improvement across the entire bond-dissociation curve. To quantify the improvement, we focus on the energy error ΔE and the infidelity $1 - F$ to the true ground state, with and without SV (Fig. 4.3). SV reduces the energy error by an average factor ~ 10 and reduces the infidelity by an average factor ~ 9 . In order to quantitatively understand the limits of the VQE optimization, and to clearly pinpoint the origin of the SV improvement, we simulate the experiment via the density-matrix simulator *quantumsim* [146], using an error model built from independently measured experimental parameters [147]. We build the error model incrementally, progressively adding: optimization inaccuracy (the difference between the state ideally produced by the converged θ and the true ground state); dephasing on both qubits (quantified by the measured Ramsey dephasing times T_2^*); relaxation on both qubits (quantified by the measured relaxation times T_1); residual qubit excitations (measured from single-shot histograms with each qubit prepared in $|0\rangle$); and increased dephasing of Q_0 during the exchange gate (quantified by its reduced T_2^* when tuned into the exchange interaction zone). By plotting the errors from each increment of the model, we are able to dissect the observed experimental error into its separate components without [Fig. 4.3(c)] and with [Fig. 4.3(b)] SV. Measured temporal fluctuations of dephasing, relaxation and residual excitation are used to obtain simulation error bars.

The simulation using the full error model shows fairly good matching with experiment for both the ground-state energy error [Figs. 4.3(a,b)] and the state infidelity [Fig. 4.3(c)], without and with SV. The error model dissection shows that the energy error when not using SV is dominated by residual qubit excitations. This is remarkable as the calibrated residual excitations are only 0.25% for Q_0 and 1.34% for Q_1 (see Supplementary Notes). The improvement

²We have chosen to reconstruct $\rho^{(\text{raw})}$ via linear inversion for the purposes of scalability; methods such as maximum-likelihood estimation have a classical computation cost similar to that of solving the electronic structure problem itself, which would negate a quantum speedup in future experiments.




Figure 4.2: Convergence of the VQE algorithm. (a) Experimental VQE estimate of H_2 ground-state energy as a function of interatomic distance R . At each chosen R , we minimize the raw energy $E^{(raw)}$ (blue data points) over the variational parameter $\tilde{\theta}$ using the CMA-ES evolutionary algorithm [145]. Applying SV to the converged solution (orange data points) lowers the energy estimate towards the exact solution (dashed curve). Inset: A typical optimization trace for the convergence of the energy estimate. (b-d) The reconstructed density matrices of the converged states at (b) $R = 0.25$ Å, (c) $R = 0.80$ Å, and (d) $R = 2.00$ Å, showing that the converged states lie mostly in the single-excitation subspace, and that entanglement increases with the interatomic distance R .

from SV results from the mitigation of errors arising from these residual excitations and from qubit relaxation. This is precisely as expected: these error mechanisms change total qubit excitation number and violate the underlying ZZ symmetry. Using SV changes the dominant error mechanism from residual qubit excitation to optimization inaccuracy, which is bounded by the sampling noise during the optimization itself (where $N_{\text{meas}} = 10^3$), rather than the sampling noise from the final step (where $N_{\text{meas}} = 10^5$). This error could be reduced experimentally by increasing N_{meas} during the optimization, at the cost of increased convergence time. The improvement in state infidelity by SV can be explained along similar lines. We observe some increased deviations between the observed and simulated state infidelity at large R . We attribute these to limitations in our modeling of error during the exchange gate (whose duration increases with R).

4

4.5 Physicality of the state reconstructed

VQEs rely on variational bounding to ensure that the obtained approximation to the ground-state energy is accurate, but this is only guaranteed when the experimental results correspond to a physical state. Our method for calculating the ground-state energy [Eq. (4.1)] independently estimates each Pauli coefficient of the density matrix with error $\propto N_{\text{meas}}^{-1/2}$. Such estimation cannot guarantee a set of Pauli coefficients that could have come from a positive density matrix. This in turn breaks the variational lower bound on the energy estimate, and increases the error in estimates of other properties of the true ground state [148, 149]. As experimental error is reduced, $\rho^{(\text{raw})}$ tends towards a rank-1 density matrix, increasing its chance of being unphysical [149]. Moreover, $\rho^{(\text{SV})}$ is a lower-rank density matrix than $\rho^{(\text{raw})}$ (being projected onto a subspace of the Hilbert space), which implies that unphysicality may be enhanced by SV. The variance in a given term $\rho_{\hat{P}}$ post-SV can be calculated as

$$\text{Var}[\rho_{\hat{P}}^{(\text{SV})}] \approx \frac{3N_{\text{meas}}}{N_{\text{meas}}(1 + \text{Tr}[\rho^{(\text{raw})}\hat{S}])}. \quad (4.8)$$

SV has maximal impact on the quantum state precisely when this denominator is small, so this represents a natural bound for the power of SV as an error mitigation strategy.

The effect of sampling noise may be mitigated somewhat by restricting the fermionic 2-reduced density matrix to be positive (which may be completed in polynomial time) [148]. To investigate the effect of such mitigation, we bin the data used for final tomography of converged states to construct 100 density matrices with $N_{\text{meas}} = 10^3$ at each R , thus increasing the sampling noise by a factor of 10. We wish to study the relative improvement of SV in the two figures of merit, which we quantify as

$$\eta_E = \frac{|\Delta E^{(\text{raw})}|}{|\Delta E^{(\text{SV})}|} \quad \text{and} \quad \eta_F = \frac{|1 - F^{(\text{raw})}|}{|1 - F^{(\text{SV})}|} \quad (4.9)$$

when physicality of the raw density matrices is enforced and not. To enforce physicality, we employ a convex optimization routine to find the closest positive semidefinite matrix to the experimentally measured $\rho^{(\text{raw})}$ (closest in the L^2 norm sense on the space induced by the the Pauli basis). We then apply symmetry verification to the post-processed density matrix.




Figure 4.3: Impact of SV in ground-state energy and state fidelity, and dissected error budget. (a) Experimental (solid circles) energy error ΔE without and with SV compared to the result (empty circles and dashed line) of a full density-matrix simulation using the full error model. The contributions from optimizer inaccuracy, qubit dephasing, qubit relaxation, residual qubit excitations and increased Q_0 dephasing during the exchange gate are shown as shaded regions for the case of no SV applied. Without SV, ΔE is clearly dominated by residual qubit excitation. (b) Zoom-in on experimental and simulated ΔE with SV and corresponding error budget. With SV, the effects of residual excitation and qubit relaxation are successfully mitigated, as predicted in Ref. [57]. The remaining energy error is dominated by optimizer inaccuracy. Simulation error bars are obtained by modelling measured fluctuations of T_1 , T_2^* , and residual excitation. (c) Experimental (solid circles with error bars) infidelity to the true ground state without and with SV compared to simulation using the full error model (empty circles and dashed line). Error bars are propagated through the linear inversion procedure for experiment and calculated from sampling noise for simulations. For simulations, error bars are smaller than the markers.




Figure 4.4: Constraining positivity with symmetry verification to mitigate the effect of sampling noise. The experimental data from Fig. 4.3 is split into 100 sample simulations for each R , increasing the sampling noise by a factor of 10 and making it comparable to other sources of experimental error. For each sample, we plot (red) the relative energy error and infidelity [Eq. (4.9)]. Values below 1 (dashed lines) indicate that SV has not provided an improvement, as may be the case when the density matrix has negative eigenvalues. We restore the improvement from SV by constraining the positivity of the 2-reduced density matrix [148] (green). Histograms on the top and right axes show the marginal distribution of the two scatter plots. When the density matrices are constrained to be positive, we observe the points fall along the line $y = x$ (blue dashed line), indicating that SV improves both metrics by the same amount.

Figure 4.4 shows a scatter plot of η_E and η_F , and relative histograms of each. Without enforcing physicality, SV makes no significant improvement to the state fidelity, although it almost always improves the energy error. However, when positivity is enforced, SV greatly improves the overlap with the true ground state. We also find that the improvement in the energy from SV is equal to the improvement in fidelity when the starting state is physical, but is relatively uncorrelated when the starting state is not. This makes sense, as the energy gain from SV given a physical matrix comes directly from substituting higher energy states with density on the ground state. It is unclear whether such a strong trend will continue in larger systems without requiring too stringent a positivity constraint. As this is a four-orbital two-electron system, enforcing the positivity of the 2-reduced density matrix enforces positivity on the entire density matrix (which is exponentially difficult in the system size [150]). Testing this scalability is a clear direction for future research³.

³Note that, for this system, enforcing positivity of the 1-reduced density matrix corresponds to ensuring that all expectation values are bounded between -1 and 1 , and so this does not provide any additional data.

4.6 Summary and conclusions

In summary, we have experimentally demonstrated the use of SV to mitigate errors in the VQE of H_2 with two transmon qubits. We implemented an efficient variational ansatz based on an exchange gate producing finely adjustable population transfer in the single-excitation subspace, respecting the ZZ symmetry of the H_2 Hamiltonian. Verification of this symmetry reduced the error of the estimated ground-state energy and the ground state by one order of magnitude on average over the full dissociation curve. A full density-matrix simulation of our system allowed us to budget the contributions from known experimental error mechanisms. We observe that SV mitigates the effect of processes that affect total qubit excitation number, specifically qubit relaxation and residual excitation. Finally, we have investigated the effect of reconstructing density matrices via linear tomographic inversion in the presence of sampling, which voids the guarantee of positivity and in turn the guarantee that SV improves estimation of the ground state. Intriguingly, we observe that when physicality is enforced, the reduction in energy error from SV is directly linked to the increase in fidelity to the ground state. If this observation extends to larger systems, a user can be confident that symmetry-verified Pauli coefficients are accurate for calculations beyond the ground-state energy.

4.7 Supplementary Notes

4.7.1 Supplementary Note 1: Device fabrication

A high-resistivity intrinsic silicon wafer was cleaned with acetone and 2-isopropanol, and stripped of native oxides using buffered oxide etch solution (BOE 7 : 1). The wafer was subjected to HMDS vapor and sputtered with 200 nm of NbTiN followed by dicing into smaller dies. The device plane was spun with a high-contrast positive tone resist and patterned using e-beam lithography. The exposed base layer was subtractively patterned using reactive ion etching and the resist was stripped. This was followed by spinning of a bilayer resist for fabrication of Josephson junctions by double-angle shadow evaporation. For the fabrication of airbridges, a 6 μm thick e-beam resist was patterned and subjected to reflow. A 450 nm thick layer of aluminum was deposited using an e-beam evaporator. The chip was diced and wirebonded to a printed circuit board.

4.7.2 Supplementary Note 2: Experimental setup

The device was mounted on a copper sample holder attached to the mixing chamber of a Leiden Cryogenics CF-650 dilution refrigerator with ~ 22 mK base temperature. For radiation shielding, the cold finger was enclosed by a copper can coated with a mixture of Stycast 2850 and silicon carbide granules (15 to 1000 nm diameter) used for infrared absorption. To shield against external magnetic fields, the can was enclosed by an aluminum can and two Cryophy cans. Microwave lines were filtered using 60 dB of attenuation with both commercial cryogenic attenuators and home-made Eccosorb filters for infrared absorption. Flux-bias lines were also filtered using commercial low-pass filters and Eccosorb filters with a stronger

absorption. Fast flux-pulses were coupled to the flux-bias lines via room-temperature bias tees.

Amplification of the readout signal was done in three stages: first a TWPA (provided by MIT-LL [151]) located at the mixing chamber plate, then a Low Noise Factory HEMT at the 4 K plate, and finally a Miteq amplifier at room temperature. The TWPA was mounted on a separate sample holder with the same shielding layers as the device.

Room-temperature electronics used both commercial hardware and custom hardware developed in QuTech. Rhode & Schwarz SGS100 sources provided all microwave signals for single-qubit gates and readout. The DC bias was provided by home-built current sources (IVVI racks). QuTech arbitrary waveform generators (QWG) generated the modulation envelopes for single-qubit gates and the flux pulse for the exchange gate. A Zurich Instruments UHFQA was used to perform independent readout of both qubits as well as their correlation. QuTech mixers were used for frequency up- and down-conversion. The QuTech Central Controller Light (CCL) coordinated the triggering of QWGs and UHFQA.


All measurements were controlled at the software level with qCoDeS [152] and PycQED [153] packages. The QuTech OpenQL compiler translated high-level Python code into the eQASM code [154] forming the input to the CCL.

4.7.3 Supplementary Note 3: Measured device parameters

Qubit	Q_1	Q_0
Readout resonator frequency (GHz)	8.0005	7.7377
Qubit sweetspot frequency (GHz)	5.1468	5.9207
T_1 (μ s)	9.8 ± 1.0	11.7 ± 0.6
T_2^* (μ s)	9.0 ± 1.3	17.3 ± 1.0
Residual qubit excitation (%)	1.34 ± 0.20	0.25 ± 0.09
Single-qubit gate fidelity (%)	99.86	99.91
Coupling $\frac{J_1}{2\pi}$ (MHz)	20.9	

Table 4.1: Measured device parameters. Resonator and qubit frequencies were measured by spectroscopy, while relaxation and dephasing times, T_1 and T_2^* , respectively, were measured by standard time-domain experiments. Error bars on T_1 and T_2^* correspond to the standard deviation of 56 repeated measurements performed over a 24-hour period. See text for the procedure used to quantify residual qubit excitations. Single-qubit gate fidelity was measured by randomized benchmarking. The qubit-qubit coupling strength was measured both by spectroscopy and time-domain measurements.

We ran a series of characterization experiments to extract the device parameters needed as inputs to the error model used in our density-matrix simulation. These are summarized in Table 5.1. The qubit relaxation time T_1 and dephasing time T_2^* for each qubit were measured using standard time-domain sequences. The reduced dephasing time $T_2^{*,\text{red}}$ of Q_0 during the exchange gate was measured by DC biasing Q_0 to 5.1468 GHz while DC biasing Q_1 suf-



4

Figure 4.5: Device and wiring schematic. The $2\text{ mm} \times 7\text{ mm}$ chip contains three Starmon [46] qubits. Qubit pairs Q_0 - Q_1 and Q_1 - Q_2 are coupled by bus resonators. Qubits have dedicated microwave drive lines, flux bias lines, and dispersively-coupled readout resonators. Readout resonators are coupled to a common feedline allowing independent qubit readout by frequency multiplexing. See text for details of cryogenic system and wiring. In this experiment, we only make use of Q_0 and Q_1 . The unused leftmost qubit, Q_2 , is parked at its sweetspot frequency (4.128 GHz) throughout.

ficiently far away from its sweetspot. We extract $T_2^{*,\text{red}} = 0.995\text{ }\mu\text{s}$ from a standard Ramsey time-domain experiment (Fig. 4.6). Single-qubit gate fidelity was extracted from randomized benchmarking of each qubit separately.

We quantified residual qubit excitations from a subset of the measurement set used to calibrate the measurement operators in the post-convergence tomographic reconstruction at each value of R in Fig. 3 (see section below). The measurement set consists of 7×10^5 measurements with the two qubits nominally prepared in each of the four computational states $|00\rangle$, $|01\rangle$, $|10\rangle$, and $|11\rangle$. Since all homodyne voltage shots were stored (not just their av-




Figure 4.6: Ramsey experiment with qubit Q_0 DC-flux biased to the frequency of the resonant exchange interaction (5.1468 GHz, the Q_1 sweetspot frequency), but with qubit Q_1 biased away to a lower frequency. The best fit value $T_2^{*,\text{red}} = 0.995 \mu\text{s}$ is used for modelling the increased dephasing of Q_0 during the exchange gate.

erage), we could construct histograms of the measurements for $|00\rangle$ before 1-bit digitization. The residual excitation was then extracted from a double Gaussian fit [Fig. 4.7(a-b)]. We performed this procedure to extract residual qubit excitations for every data point in Fig. 4.7(c).

4.7.4 Supplementary Note 4: Pulses for gates

All single-qubit gates are performed with standard DRAG-type [41] microwave pulses of 12 ns duration, resonant on the qubit transition. Gaussian in-phase and derivative-of-gaussian in-quadrature envelopes correspond to 3 ns gaussian width. The gaussian envelope is offset to ensure continuity at the start and end of the pulse.

The exchange gate is performed with a square pulse of fixed amplitude and duration. The linear-dynamical distortions from the components in the flux-control line are compensated by first measuring the step response and then defining pre-distortion filters to approach the ideal step. This procedure is performed in two steps. First we address the components at room temperature, and then those inside the dilution refrigerator.

The distortions caused by room-temperature components (mainly a bias tee) are measured (Fig. 4.8[a]) with a Rhode-Schwarz RTO1024 oscilloscope, with 2 Gs/s and 4 GHz bandwidth. To improve the step response, we follow the procedure and models as explained in the Supplementary Material of Ref. [67]. We iteratively fit simple models with few parameters to the measured step response. We then use the best estimate of these models to compute the impulse response of the inverse filter. The pre-distortion filter convolves the input with this computed impulse response.

Using the pre-distortion correction for room-temperature components, we next measure the distortions from components inside the dilution refrigerator (Fig. 4.8[c]) using the method of Ref. [155]. For completeness, we briefly discuss the procedure here. We apply square pulses of variable duration between the $\pi/2$ pulses of a Ramsey sequence. This Ramsey sequence allows us to extract the quantum phase acquired by the qubit as a function of the pulse duration. Taking the derivative of this phase, we obtain the instantaneous detuning of the qubit as a function of the pulse duration. Using the known dependence of qubit frequency with flux amplitude, we can thus extract the voltage-to-flux step response. We fit simple models with few parameters to the measured step response, and compute the impulse response of the corresponding inverse filter. We convolve the computed impulse responses of the inverting filters for distortions outside and inside the refrigerator to produce the impulse response of the pre-distortion filter for the complete flux-control line.

The set of corrections summarized in Table 4.2 allow us to produce a net step response accurate to 1% after 5 ns (Fig. 4.8[d]). As in Ref. [67], we mitigate the effect of residual




Figure 4.7: Residual qubit excitations. Single-shot readout histograms for (a) Q_1 and (b) Q_0 nominally prepared in $|0\rangle$ and $|1\rangle$. We extract a residual qubit excitation from the best-fit double gaussian on the ground-state histogram (solid curves). (c) Residual excitations extracted from calibration data in the final, post-VQE-convergence in the dataset of Figs. 2 and 3 at each value of R . The average (standard deviation) of the residual excitation is 0.25% (0.09%) for Q_0 and 1.34% (0.20%) for Q_1 .

Model	Step-response distortion	Parameters
High-pass	$e^{-t/\tau} u(t)$	$\tau = 40.9 \mu s$
Exp.	$\left(1 + ae^{-t/\tau}\right) u(t)$	$a = 0.13129 ; \tau = 15.027 \mu s$
Exp.	$\left(1 + ae^{-t/\tau}\right) u(t)$	$a = 0.99088 ; \tau = 6.408 \mu s$
Exp.	$\left(1 + ae^{-t/\tau}\right) u(t)$	$a = -0.0747 ; \tau = 3.541 \mu s$
Exp.	$\left(1 + ae^{-t/\tau}\right) u(t)$	$a = -0.0176 ; \tau = 1.059 \mu s$
Exp.	$\left(1 + ae^{-t/\tau}\right) u(t)$	$a = 0.0659 ; \tau = 21.7 \text{ ns.}$
Exp.	$\left(1 + ae^{-t/\tau}\right) u(t)$	$a = -0.0045 ; \tau = 13.0 \text{ ns.}$
Exp.	$\left(1 + ae^{-t/\tau}\right) u(t)$	$a = -0.015 ; \tau = 3.0 \text{ ns.}$
Bounce	$(1 + a)u(t) - au(t - \tau)$	$a = 0.02 ; \tau = 15.0 \text{ ns}$
Bounce	$(1 + a)u(t) - au(t - \tau)$	$a = -0.045 ; \tau = 8.0 \text{ ns}$

Table 4.2: Models used to match the step response of components outside (top) and inside the refrigerator (model), and corresponding fit values (HPF = high-pass filter, Exp. = exponential approach, Bounce = transmission-line bounce). Here, $u(t)$ is the Heaviside step function.

uncompensated long timescale distortions by adding negative copies of the applied flux pulse at the end of each sequence (after measurement) to ensure zero net area.

4.7.5 Supplementary Note 5: Tomographic reconstruction and limitations

Tomographic reconstruction was performed with the same technique described in [156]. We provide a brief description here for completeness. For each measurement channel (measurement of Q_1 , measurement of Q_0 , and their correlation), the average measurement outcome is given by $\langle m_i \rangle = \text{Tr}(\hat{M}_i \rho)$, with operator

$$\hat{M}_i = \beta_{II}^i \hat{I} \hat{I} + \beta_{IZ}^i \hat{I} \hat{Z} + \beta_{ZI}^i \hat{Z} \hat{I} + \beta_{ZZ}^i \hat{Z} \hat{Z}, \quad (4.10)$$

and real-valued coefficients β_j^i . Single-shot measurements of Q_0 and Q_1 are 1-bit digitized before correlation and before averaging each of the three channels.

The simultaneously applied measurement pre-rotations R_0 and R_1 consist of the 36 pairs created by drawing each rotation separately from the set $\{I, X_\pi, X_{\pi/2}, Y_{\pi/2}, X_{-\pi/2}, Y_{-\pi/2}\}$. These measurement pre-rotations effectively change the measurement operator to

$$M_i^{k,l} = \text{Tr} \left(R^{k,l,\dagger} \hat{M}_i R^{k,l} \right).$$

There are thus 108 linear equations (36 per channel) linking the averaged measurement to the 15 nontrivial 2-qubit Pauli coefficients (we force $\langle \hat{I} \hat{I} \rangle = 1$). We then extract the Pauli coefficients by performing least-squares linear inversion. Prior to the linear inversion, the measurements are scaled to approximately match the noise in the three channels.

The coefficients β_j^i are obtained from standard calibration measurements. The two qubits are nominally prepared in the four computational states and measured. In total, we perform

$7 \times N_{\text{meas}}$ measurements per computational state. The matrix relating the four measurement averages of a channel to the coefficients has elements of the form $\langle \hat{I}\hat{I} \rangle$, $\pm \langle \hat{I}\hat{Z} \rangle$, $\pm \langle \hat{Z}\hat{I} \rangle$ and $\pm \langle \hat{Z}\hat{Z} \rangle$. By taking into account the calibrated residual qubit excitations, which reduce the magnitude of $\langle \hat{I}\hat{Z} \rangle$, $\langle \hat{Z}\hat{I} \rangle$, and $\langle \hat{Z}\hat{Z} \rangle$ from unity, we ensure that the coefficients β_j^i and thus also the operator M_i are not corrupted by residual excitation [157].




Figure 4.8: (a) Measured (blue) voltage step response of room-temperature components of the flux-control line, and fit model (orange) using the parameters on the top half of Table 4.2. (b) Measured voltage-step response after applying the pre-distortion filter for room-temperature components. (c) Measured (blue) flux-step response of the complete flux-control line (using only pre-distortion filter for room-temperature components), and fit model (orange) using the parameters on the lower half of Table 4.2. (d) Final voltage-to-flux step response of the complete flux-control line using pre-distortion for both components outside and inside the dilution refrigerator. The response is within 1% of the ideal step (green) after 5 ns.




Figure 4.9: Minimal eigenvalue of density matrices obtained from linear tomographic reconstruction with different N_{meas} . Here, the state preparation targets a Bell state using our variational ansatz (inset) with $\theta = \pi/4$ (producing a $\sqrt{\text{iSWAP}}$ gate). Physicality constraints on density matrices restrict their eigenvalues to be non-negative. We observe negative minimum eigenvalues over the entire range of N_{meas} . A *quantumsim* simulation produces a similar trend, asymptoting to a physical state by $N_{\text{meas}} \sim 50,000$.

Tomography by linear inversion does not ensure physicality of the reconstructed density matrix. We investigate this effect by performing tomography with variable N_{meas} on the state produced by our ansatz with $\theta = \pi/4$ and extracting the minimum eigenvalue of the reconstructed density matrix (Fig. 4.9). A negative minimum eigenvalue manifests unphysicality over the N_{meas} range covered. Our *quantumsim* simulation produces a similar trend, asymptoting to a physical state by $N_{\text{meas}} = 5 \times 10^4$. These observations led us to choose $N_{\text{meas}} = 10^5$ for the final state tomography post VQE convergence in Fig. 3, and to further investigate (in Fig. 4) how unphysicality can violate the variational principle, producing reductions in energy from imprecise state reconstruction rather than algorithmic precision.

4.7.6 Supplementary Note 6: Constraining the positivity of reduced density matrices

Testing whether a N -qubit density matrix ρ is positive is in general QMA-hard [150]. However, if we trace out all but a polynomial number of degrees of freedom of ρ , testing positivity of the resulting reduced density matrix $\rho^{(\text{red})}$ is tractable on classical hardware, and obtaining the closest nearby positive matrix is similarly so. This gives a set of necessary but insufficient physicality conditions for ρ , but enforcing k -local constraints (on a density matrix from a

VQE) tends to be sufficient to variationally bound the resulting energies [148]. Following the reduction, we write $\rho^{(\text{red})}$ as a vector over the Pauli basis,

$$\rho^{(\text{red})} = \sum_{\hat{P}} \rho_{\hat{P}}^{(\text{red})} \hat{P}. \quad (4.11)$$

Then, we attempt to find the density matrix $\tilde{\rho}^{(\text{red})}$ closest to $\rho^{(\text{red})}$ in the L^2 -norm

$$\sum_{\hat{P}} \left(\rho_{\hat{P}}^{(\text{red})} - \tilde{\rho}_{\hat{P}}^{(\text{red})} \right)^2, \quad (4.12)$$

subject to the conditions $\tilde{\rho}^{(\text{red})} \succ 0$, and $\tilde{\rho}_I^{(\text{red})} = 1$. This gives a quadratic minimization problem with cone inequality and linear equality constraints, which we solve using interior point methods.

4.7.7 Supplementary Note 7: Theoretical modeling of the experiment

We use our full-density-matrix simulator *quantumsim* to model the experiment. The error model takes as input parameters the measured values of T_1 , T_2^* and residual excitation for both qubits, and $T_2^{*,\text{red}}$ for Q_0 . We also include the effect of fluctuations on the device parameters by Monte Carlo sampling.

Numerical simulations

The simulations are performed by extracting the full-density-matrix $\rho^{(\text{sim})}$ at the end of the circuit. We use the converged value of $\tilde{\theta}$ at each R to generate the quantum state and extract the Pauli coefficients $\rho_{\hat{P}}^{(\text{sim})}(\tilde{\theta}) = \text{Tr}[\hat{P}\rho^{(\text{sim})}]$. We add sampling noise to each coefficient, drawn from a zero-mean Gaussian distribution with variance $(1 + \rho_{\hat{P}}^{(\text{sim})})(1 - \rho_{\hat{P}}^{(\text{sim})})/N_{\text{meas}}$, where $N_{\text{meas}} = 4 \times 10^5$. Note that this is greater than the number of measurements per tomographic prerotation in the experiment, as data from multiple tomographic prerotations is used to estimate each Pauli coefficient. To account for fluctuations on the device parameters T_1 , T_2^* , and residual excitations, we average over 10^4 simulations for every R . For each simulation, we draw parameters from independent normal distributions using values in Table 5.1. As the dephasing noise T_2^* depends on T_1 ,

$$\frac{1}{T_2^*} = \frac{1}{T_{\text{QE}}} + \frac{1}{2T_1}, \quad (4.13)$$

it is more appropriate to sample the pure dephasing rate T_{QE} independently. We calculate the T_{QE} mean and variance ($\bar{T}_{\phi}$, $\text{Var}[\bar{T}_{\phi}]$) from T_2^* ,

$$\begin{aligned} \bar{T}_{\phi} &= \frac{1}{\frac{1}{T_2^*} - \frac{1}{2T_1}}, \\ \text{Var}[\bar{T}_{\phi}] &= \bar{T}_{\phi}^2 (T_2^*)^{-2} \left[\text{Var}[T_2^*] - \frac{\text{Var}[T_1]}{2T_1^2 (T_2^*)^{-2}} \right]. \end{aligned} \quad (4.14)$$

From the 10^4 simulations we obtain 95%-confidence error bars for ΔE and F as twice the population standard deviation.

Exchange gate

Due to quasi-static flux noise, the angle of the unitary exchange gate (Eq. 7) differs between subsequent applications. Assuming that these fluctuations are fast on the scale of the 2 hour optimization, this may be simulated by integrating over the range of applied gates, resulting in an incoherent noise model. As the gate is not repeatedly applied during a single-shot experiment, this incoherent approximation does not lead to an error in the final result. To perform the integration, we convert our unitary U_θ into a Pauli Transfer Matrix representation (PTM) [158, 159]:

$$[R_\theta]_{i,j} = \frac{1}{2} \text{Tr}[\hat{P}_i U_\theta \hat{P}_j U_\theta^\dagger], \quad \hat{P}_i, \hat{P}_j \in \mathbb{P}^N, \quad (4.15)$$

which may then be integrated over a probability distribution in the deviation δ from the target angle θ :

$$[\tilde{R}]_{i,j} = \int d\delta p(\delta) [R_{\theta+\delta}]_{i,j}. \quad (4.16)$$

We choose for $p(\delta)$ a Gaussian distribution: $p(\delta) = e^{-\frac{\delta^2}{2\sigma^2}}$. In order to obtain the distribution width σ^2 , we note that the same effect causes single-qubit dephasing of Q_0 when fluxed to the exchange point when Q_1 is fluxed away. We may thus estimate σ as

$$\sigma^2 = 1 - e^{-\frac{t_{\text{int}}}{T_2^{*,\text{red}}}}, \quad (4.17)$$

where t_{int} is the exchange gate duration and $T_2^{*,\text{red}}$ the dephasing time of Q_0 at the exchange point (with Q_1 fluxed away). The final gate simulation also includes the effect of amplitude damping on both qubits, and the dephasing of Q_1 at the sweet spot as discrete error channels of duration $\frac{t_{\text{int}}}{2}$ on either side of the gate $\tilde{R}$.

VARIATIONAL PREPARATION OF FINITE-TEMPERATURE STATES ON A QUANTUM COMPUTER

The preparation of thermal equilibrium states is important for the simulation of condensed-matter and cosmology systems using a quantum computer. We present a method to prepare such mixed states with unitary operators, and demonstrate this technique experimentally using a gate-based quantum processor. Our method targets the generation of thermofield double states using a hybrid quantum-classical variational approach motivated by quantum-approximate optimization algorithms, without prior calculation of optimal variational parameters by numerical simulation. The fidelity of generated states to the thermal-equilibrium state smoothly varies from 99 to 75% between infinite and near-zero simulated temperature, in quantitative agreement with numerical simulations of the noisy quantum processor with error parameters drawn from experiment.¹

¹This chapter has been published in *npj Quantum Information* **7** 130 (2021).

5.1 Introduction

The potential for quantum computers to simulate other quantum mechanical systems is well known [1], and the ability to represent the dynamical evolution of quantum many-body systems has been demonstrated [68]. However, the accuracy of these simulations depends on efficient initial state preparation within the quantum computer. Much progress has been made on the efficient preparation of non-trivial quantum states, including spin-squeezed states [160] and entangled cat states [101]. Studying phenomena like high-temperature superconductivity [161] requires preparation of thermal equilibrium states, or Gibbs states. Producing mixed states with unitary quantum operations and measurements is not straightforward, and has only recently begun to be explored [162, 163]. In this work, we demonstrate the use of a variational quantum-classical algorithm to realize Gibbs states using (ideally unitary) gate control on a transmon quantum processor.

Our approach is mediated by the generation of thermofield double (TFD) states, which are pure states sharing entanglement between two identical quantum systems with the characteristic that when one of the systems is considered independently (by tracing over the other), the result is a mixed state representing equilibrium at a specific temperature. TFD states are of interest not only in condensed matter physics but also for the study of black holes [164, 165] and traversable wormholes [166, 167]. We use a variational protocol [168] motivated by quantum-approximate optimization algorithms (QAOA) that relies on alternation of unitary intra- and inter-system operations to control the effective temperature, eliminating the need for a large external heat bath. Other methods have been studied for generation of Gibbs states, such as quantum metropolis sampling [169] and imaginary time evolution using variational quantum simulation [170, 171]. However the advantage of QAOA compared to these proposals is that the form of the ansatz is relatively straightforward and low-depth, whereas the metropolis sampling involves phase estimation which leads to a high-depth circuit, and the imaginary time evolution proposal does not have a clear proposal for the form of the ansatz. Recently, verification of TFD state preparation was demonstrated on a trapped-ion quantum computer [162]. Our work experimentally demonstrates the generation of finite-temperature states in a superconducting quantum computer by variational preparation of TFD states in a hybrid quantum-classical manner.

5.2 Results

5.2.1 Theory

Consider a quantum system described by Hamiltonian H with eigenstates $|j\rangle$ and corresponding eigenenergies E_j :

$$H|j\rangle = E_j|j\rangle. \quad (5.1)$$

The Gibbs state ρ_{Gibbs} of the system is

$$\rho_{\text{Gibbs}}(\beta) = \frac{1}{Z} \sum_j \exp(-\beta E_j) |j\rangle \langle j|, \quad (5.2)$$

where $\beta = 1/k_B T$ is the inverse temperature, k_B is the Boltzmann constant, and

$$Z = \sum_j \exp(-\beta E_j) \quad (5.3)$$

is the partition function. Except in the limit $\beta \rightarrow \infty$, the Gibbs state is a mixed state and thus impossible to generate strictly through unitary evolution. To circumvent this, we define the TFD state [168] on two identical systems A and B as

$$|\text{TFD}(\beta)\rangle = \frac{1}{\sqrt{Z}} \sum_j \exp\left(\frac{-\beta E_j}{2}\right) |j\rangle_B |j\rangle_A. \quad (5.4)$$

Tracing out either system yields the desired Gibbs state in the other.

To prepare the TFD states, we follow the variational protocol proposed by [168] and consider two systems each of size n . In the first step of the procedure, the TFD state at $\beta = 0$ is generated by creating Bell pairs $|\Phi_i^+\rangle = (|0\rangle_{B_i} |0\rangle_{A_i} + |1\rangle_{B_i} |1\rangle_{A_i}) / \sqrt{2}$ between corresponding qubits i in the two systems. Tracing out either system yields a maximally mixed state on the other, and vice versa. The next steps to create the TFD state at finite temperature depend on the relevant Hamiltonian. Here, we choose the transverse field Ising model in a one-dimensional chain of n spins [172], with $n = 2$ [??(a)]. We map spin up (down) to the computational state $|0\rangle$ ($|1\rangle$) of the corresponding transmon. The Hamiltonian describing system A is

$$H_A = ZZ_A + gX_A, \quad (5.5)$$

where $ZZ_A = Z_{A2}Z_{A1}$, $X_A = X_{A2} + X_{A1}$, and g is proportional to the transverse magnetic field. The Hamiltonian for system B is the same. We focus on $g = 1$, where a phase transition is expected in the transverse field Ising model at large n [173]. We use a QAOA-motivated variational ansatz [168, 174], where intra-system evolution is interleaved with a Hamiltonian enforcing interaction between the systems:

$$H_{BA} = XX_{BA} + ZZ_{BA}, \quad (5.6)$$

where $XX_{BA} = X_{B2}X_{A2} + X_{B1}X_{A1}$, and analogously for ZZ_{BA} . For single-step state generation, the unitary operation describing the TFD protocol is

$$U(\vec{\alpha}, \vec{\gamma}) = U_{\text{inter}}(\vec{\alpha}) U_{\text{intra}}(\vec{\gamma}), \quad (5.7)$$

where

$$U_{\text{intra}}(\vec{\gamma}) = \exp[-i\gamma_2 (ZZ_B + ZZ_A)/2] \times \exp[-i\gamma_1 (X_B + X_A)/2], \quad (5.8)$$

$$U_{\text{inter}}(\vec{\alpha}) = \exp(-i\alpha_2 ZZ_{BA}/2) \exp(-i\alpha_1 XX_{BA}/2). \quad (5.9)$$

The variational parameters $\vec{\gamma} = (\gamma_1, \gamma_2)$, $\vec{\alpha} = (\alpha_1, \alpha_2)$ are optimized by the hybrid classical-quantum algorithm to generate states closest to the ideal TFD states. A single step of intra- and inter-system interaction ideally produces the state $|\psi(\vec{\alpha}, \vec{\gamma})\rangle = U(\vec{\alpha}, \vec{\gamma}) \left(|\Phi_2^+\rangle \otimes |\Phi_1^+\rangle \right)$ [175].

The variational algorithm extracts the cost function after each state preparation. We engineer a cost function $\mathcal{C}$ to be minimized when the generated state is closest to an ideal TFD state [175]. Following recent work on the concentration of control parameters for QAOA [176, 177], we expect engineering the cost function based on the target state for small-sized systems to lead to a general expression for the cost function of an arbitrary-sized system. The engineered cost function is given by:

$$\begin{aligned} C(\beta) = & \langle X_A \rangle + \langle X_B \rangle + 1.57 (\langle ZZ_A \rangle + \langle ZZ_B \rangle) \\ & - \beta^{-1.57} (\langle XX_{BA} \rangle + \langle ZZ_{BA} \rangle). \end{aligned} \quad (5.10)$$

We compare the performance of this engineered cost function $\mathcal{C}_{1.57}$ to that of the non-optimized cost function $\mathcal{C}_{1.00}$, using the reduction of infidelity to the Gibbs state as the ultimate metric of success [see Supplementary Note 1]. The engineered cost function achieves an average improvement of 54% across the β range covered ($[10^{-2}, 10^2]$ in units of $1/g$), as well as a maximum improvement of up to 98% for intermediate temperatures ($\beta \sim 1$). Our choice of the class of cost functions to optimize lets us trade off a slight decrease in low-temperature performance with a significant increase in performance at intermediate temperatures. See [175] for further details on the theory.

The quantum portion of the algorithm prepares the state according to a given set of angles $(\vec{\alpha}, \vec{\gamma})$, performs the measurements, and returns these values to the classical portion. The classical portion then evaluates the cost function according to the returned measurements, performs classical optimization, generates and returns the next set of variational angles to evaluate on the quantum portion.

5.2.2 Experiment

We implement the algorithm using four of seven transmons in a monolithic quantum processor [??(a)]. The four transmons (labelled A_1 , A_2 , B_1 , and B_2) have square connectivity provided by coupling bus resonators, and are thus ideally suited for implementing the circuit in ??(b). Each transmon has a microwave-drive line for single-qubit gating, a flux-bias line for two-qubit controlled- Z (CZ) gates, and a dispersively coupled resonator with dedicated Purcell filter [26, 178]. The four transmons can be simultaneously and independently read out by frequency multiplexing, using the common feedline connecting to all Purcell filters. All transmons are biased to their flux-symmetry point (i.e., sweetspot [179]) using static flux bias to counter residual offsets. Device details and a summary of measured transmon parameters are provided in Supplementary Note 3. Details on the experimental setup can be found in Supplementary Note 4.

In order to realize the theoretical circuit in ??(b), we first map it to the optimized depth-13 equivalent circuit shown in ??(b), which conforms to the native gate set in our control architecture. This gate set consists of arbitrary single-qubit rotations about any equatorial axis of the Bloch sphere, and CZ gates between nearest-neighbor transmons. Conveniently, all variational angles are mapped to either the axis or angle of single-qubit rotations. Further details on the compilation steps are reported in Methods and Supplementary Note 2. Bases

pre-rotations are added at the end of the circuit to first extract all the terms in the cost function $\mathcal{C}$ and finally to perform two-qubit state tomography of each system.

Prior to implementing any variational optimizer, it is helpful to build a basic understanding of the cost-function landscape. To this end, we investigate the cost function $\mathcal{C}$ at $\beta = 0$ using two-dimensional cuts: we sweep $\vec{\gamma}$ while keeping $\vec{\alpha} = 0$ to study the effect of U_{intra} and vice versa to study the effect of U_{inter} . Note that owing to the $\beta^{-1.57}$ divergence, the cost function reduces to $-\langle H_{\text{BA}} \rangle$ in the $\beta = 0$ limit. Consider first the landscape for an ideal quantum processor, which is possible to compute for our system size. The $\vec{\gamma}$ landscape at $\vec{\alpha} = 0$ is π -periodic in both directions due to the invariance of $|\text{TFD}(\beta = 0)\rangle$ under bit-flip (X) and phase-flip (Z) operations on all qubits. The cost function is minimized to -4 at even multiples of $\pi/2$ on γ_1 and γ_2 : $|\text{TFD}(\beta = 0)\rangle$ is a simultaneous eigenstate of XX_{BA} and ZZ_{BA} with eigenvalue $+2$ due to the symmetry of the constituting Bell states $|\Phi_i^+\rangle$. In turn, the cost function is maximized to $+4$ at odd multiples of $\pi/2$, at which the $|\Phi_i^+\rangle$ are transformed to singlets $|\Psi_i^-\rangle = (|0\rangle_{\text{Bi}}|1\rangle_{\text{Ai}} - |1\rangle_{\text{Bi}}|0\rangle_{\text{Ai}})/\sqrt{2}$. The $\vec{\alpha}$ landscape at $\vec{\gamma} = 0$ is constant, reflecting that $|\text{TFD}(\beta = 0)\rangle$ is a simultaneous eigenstate of XX_{BA} and ZZ_{BA} and thus also of any exponentiation of these operators. The corresponding experimental landscapes show qualitatively similar behavior. The $\vec{\gamma}$ landscape clearly shows the π periodicity with respect to both angles, albeit with reduced contrast. The $\vec{\alpha}$ landscape is not strictly constant, showing weak structure particularly with respect to α_2 . These experimental deviations reflect underlying errors in our noisy intermediate-scale quantum (NISQ) processor, which include transmon decoherence, residual ZZ coupling at the bias point, and leakage during CZ gates. We discuss these error sources in detail further below.

The task of the variational algorithm is to balance the mixture of the states at each β , in order to generate the corresponding Gibbs state. Although thermal states are well understood, it is challenging to accurately generate them in NISQ devices for studies of finite temperature systems. When working with small systems, it is possible and tempting to predetermine the variational parameters at each β by a prior classical simulation and optimization for an ideal or noisy quantum processor. We refer to this common practice [162, 180] as *cheating*, since this approach does not scale to larger problem sizes and skips the main quality of variational algorithms: to arrive at the parameters variationally. Here, we avoid cheating altogether by starting at $\beta = 0$, with initial guess the obvious optimal variational parameters for an ideal processor ($\vec{\gamma} = \vec{\alpha} = 0$), and using the experimentally optimized $(\vec{\alpha}, \vec{\gamma})$ at the last β as an initial guess when stepping β in the range $[0, 5]$ (in units of $1/g$). This approach only relies on the assumption that solutions (and their corresponding optimal variational angles) vary smoothly with β . At each β , we use the Gradient-Based Random-Tree optimizer of the *scikit-optimize* [181] Python package to minimize $\mathcal{C}$, using 4096 averages per tomographic pre-rotation necessary for the calculation of $\mathcal{C}$. After 200 iterations, the optimization is stopped. The best point is remeasured two times, each with 16384 averages per tomographic pre-rotation needed to perform two-qubit quantum state tomography of each system. A new optimization is then started for the next β , using the previous solution as the initial guess.

To begin comparing the optimized states ρ_{Exp} produced in experiment to the target Gibbs states ρ_{Gibbs} , we first visualize their density matrices (in the computational basis) for a sampling of the β range covered (??). Starting from the maximally-mixed state $\mathbb{I}/4$ at $\beta = 0$, the Gibbs state monotonically develops coherences (off-diagonal terms) between all states as β increases. Coherences between states of equal (opposite) parity have 0 (π) phase throughout. Populations (diagonal terms) monotonically decrease (increase) for even (odd) parity states. By $\beta = 5$, the Gibbs state is very close to the pure state $|\Upsilon\rangle\langle\Upsilon|$, where $|\Upsilon\rangle \approx \sqrt{0.36}(|01\rangle + |10\rangle) - \sqrt{0.14}(|00\rangle + |11\rangle)$. The noted trends are reproduced in ρ_{Exp} . However, the matching is evidently not perfect, and to address this we proceed to a quantitative analysis.

We employ two metrics to quantify experimental performance: the fidelity F of ρ_{Exp} to ρ_{Gibbs} and the purity P of ρ_{Exp} , given by

$$F = \text{Tr} \left(\sqrt{\rho_{\text{Gibbs}}^{1/2} \rho_{\text{Exp}} \rho_{\text{Gibbs}}^{1/2}} \right), \quad (5.11)$$

$$P = \text{Tr}(\rho_{\text{Exp}}^2). \quad (5.12)$$

At $\beta = 0$, $F = 99\%$ and $P = 0.262$, revealing a very close match to the ideal maximally-mixed state. However, F smoothly worsens with increasing β , decreasing to 92% at $\beta = 1$ and 75% by $\beta = 5$. Simultaneously, P does not closely track the increase of purity of the Gibbs state. By $\beta = 5$, the Gibbs state is nearly pure, but P peaks at 0.601.

In an effort to quantitatively explain these discrepancies, we perform a full density-matrix simulation of a four-qutrit system using *quantumsim* [182]. Our simulation incrementally adds calibrated errors for our NISQ processor, starting from an ideal processor (model 0): transmon relaxation and dephasing times at the bias point (model 1), increased dephasing from flux noise during CZ gates (model 2), crosstalk from residual ZZ coupling at the bias point (model 3), and transmon leakage to the second-excited state during CZ gates (model 4). The experimental input parameters for each increment and details of modeling are described in Methods and Supplementary Notes 5–9. The added curves in ?? clearly show that model 4 quantitatively matches the observed dependence of F and P over the full β range, and identifies leakage from CZ gates as the dominant error.

5.3 Discussion

The power of variational algorithms relies on their adaptability: the optimizer is meant to find its way through the variational parameter space, adapting to mitigate coherent errors as allowed by the chosen parametrization. For completeness, we compare in ?? the performance achieved with our variational strategy to that achieved by cheating, i.e., using the pre-calculated optimal $(\vec{\alpha}, \vec{\gamma})$ for an ideal processor. Our variational approach, whose sole input is the obvious initial guess at $\beta = 0$, achieves comparable performance at all β . This aspect is crucial when considering the scaling with problem size, as classical pre-simulations will require prohibitive resources beyond ~ 50 qubits, but variational optimizers would not. Given the dominant role of leakage as the error source, which cannot be compensated by

the chosen parametrization, it is unsurprising in hindsight that both approaches yield nearly identical performance.

In summary, we have presented the generation of finite-temperature Gibbs states in a quantum computer by variational targeting of TFD states in a hybrid quantum-classical manner. The algorithm successfully prepares mixed states for the transverse field Ising model with Gibbs-state fidelity ranging from 99% to 75% as β increases from 0 to $5/g$. The loss of fidelity with decreasing simulated temperature is quantitatively matched by a numerical simulation with incremental error models based on experimental input parameters, which identifies leakage in CZ gates as dominant. This work demonstrates the suitability of variational algorithms on NISQ processors for the study of finite-temperature problems of interest, ranging from condensed-matter physics to cosmology. Our results also highlight the critical importance of continuing to reduce leakage in two-qubit operations when employing weakly-anharmonic multi-level systems such as the transmon.

During the preparation of this manuscript, we became aware of related experimental work [180] on a trapped-ion system, applying a non-variationally prepared TFD state to the calculation of a critical point.

5.4 Methods

5.4.1 Quantum Circuit

We map the theoretical circuit in ??(b) to an equivalent circuit conforming to the native gate set in our control architecture and exploiting virtual Z -gate compilation [183] to minimize circuit depth. Single-qubit rotations $R_{XY}(\phi, \theta)$, by arbitrary angle θ around any equatorial axis $\cos(\phi)\mathbf{x} + \sin(\phi)\mathbf{y}$ on the Bloch sphere, are realized using 20 ns DRAG pulses [21, 41]. Two-qubit CZ gates are realized by baseband flux pulsing [22, 184] using the Net Zero scheme [48?], completing in 80 ns. In the optimized circuit [??(b)], CZ gates only appear in pairs. These pairs are simultaneously executed and tuned as one block. Single-qubit rotations R_1 - R_4 are used to change the measurement bases, as required to measure $\mathcal{C}$ during optimization and to perform two-qubit tomography [58] in each system to extract F and P . A summary of single- and two-qubit gate performance [see Supplementary Note 5] and a step-by-step derivation of the optimized circuit are provided [see Supplementary Note 2].

5.4.2 Modeling and Simulations

Noiseless simulations were performed prior to experiments, for verification of algorithm convergence. During simulations, experimental conditions were maintained exactly for the algorithm and the control software, while the outcome from readout hardware (Zurich Instruments UHFQC) was replaced with a simulated readout [see Supplementary Note 10 for details].

The models used to simulate the performance of the algorithm are incremental: model k contains all the noise mechanisms in model $k - 1$ plus one more, which we use for labeling in ??. Model 0 corresponds to an ideal quantum processor without any error. Model 1 adds the measured relaxation and dephasing times measured for the four transmons at their bias point.

These times are tabulated in Supplementary Table 1. Model 2 adds the increased dephasing that flux-pulsed transmons experience during CZ gates. For this we extrapolate the echo coherence time T_2^{echo} to the CZ flux-pulse amplitude using a $1/f$ noise model [185, 186] with amplitude $\sqrt{A} = 1\mu\Phi_0$. This noise model is implemented following [187]. Model 3 adds the idling crosstalk due to residual ZZ coupling between transmons. This model expands on the implementation of idling evolution used for coherence times: the circuit gates are simulated to be instantaneous, and the idling evolution of the system is trotterized. In this case, the residual ZZ coupling operator uses the measured residual ZZ coupling strengths at the bias point [see Supplementary Note 6 for details]. Finally, model 4 adds leakage to the CZ gates, based on randomized benchmarking with modifications to quantify leakage [48, 119], and implemented in simulation using the procedure described in [187].

5.4.3 Leakage in Transmons

Leakage to transmon second-excited states is found essential to quantitatively match the performance of the algorithm by simulation. To reach this conclusion it was necessary to thoroughly understand how leakage affects the two-qubit tomographic reconstruction procedure employed. The readout calibration only considers computational states of the two transmons involved. Moreover, basis pre-rotations only act on the qubit subspace, leaving the population in leaked states unchanged. Using an overcomplete set of basis pre-rotations for state tomography, comprising both positive (X, Y, Z) and negative $(-X, -Y, -Z)$ bases for each transmon, leads to the misdiagnosis of a leaked state as a maximally mixed state qubit state for that transmon. See Supplementary Note 8 for further details.

5.5 Supplementary Notes

5.5.1 Supplementary Note 1: Optimization of the cost function

We optimize the cost function to maximize fidelity of the variationally-optimized state $|\psi(\vec{\alpha}, \vec{\gamma})\rangle$ to the TFD state $|\text{TFD}(\beta)\rangle$ (assuming an ideal processor). We consider the class of cost functions defined by

$$\mathcal{C}_\varsigma(\beta) = X_A + X_B + \varsigma(ZZ_A + ZZ_B) - \beta^{-\varsigma}(XX_{BA} + ZZ_{BA}), \quad (5.13)$$

and perform nested optimization of parameter ς to minimize the infidelity of the variationally-optimized state to the TFD state over a range of inverse temperatures

$$\mathcal{B} = \left\{ 10^{x/2} : (x \in \mathbb{Z}) \wedge (-8 \leq x \leq 8) \right\}. \quad (5.14)$$

Here we use ideal TFD states to tailor the cost function for variational preparation of Gibbs states [168, 175]. We define the minimization quantity of interest as

$$\Xi(\varsigma) = \sum_{\beta \in \mathcal{B}} \sum_{o \in \mathcal{O}} |\langle \text{TFD}(\beta) | o | \text{TFD}(\beta) \rangle - \langle \Psi(\varsigma, \beta) | o | \Psi(\varsigma, \beta) \rangle|, \quad (5.15)$$

where $\mathcal{O}$ is the set of operators

$$\mathcal{O} = \{X_A, Y_A, Z_A, X_B, Y_B, Z_B, XX_A, YY_A, ZZ_A, XX_B, YY_B, ZZ_B, XX_{BA}, YY_{BA}, ZZ_{BA}\}, \quad (5.16)$$

and $|\Psi(\varsigma, \beta)\rangle$ is the state optimized using $\mathcal{C}_\varsigma(\beta)$. We find the minimum value of Ξ at $\varsigma = 1.57$ [see Figure 5.1(a)].

We compare the performance of the optimized cost function $\mathcal{C}_{1.57}$ to that used in prior work, $\mathcal{C}_{1.00}$, in two ways. First, we compare the simulated infidelity to the TFD state of states optimized with both cost functions in the range $\beta \in [0.1, 10]$. The optimized cost function $\mathcal{C}_{1.57}$ performs better over the entire range. Finally, we compare the simulated fidelity F of the reduced state of system A to the targeted Gibbs state. As shown in Figure 5.2, using $\mathcal{C}_{1.57}$ significantly reduces the infidelity $1 - F$ for $\beta \gtrsim 3$. We also observe that the purity of the reduced state tracks that of the Gibbs state more closely when using $\mathcal{C}_{1.57}$.




Figure 5.1: Optimization of the cost function. (a) Plot of Ξ versus ς . The minimum of Ξ is found near $\varsigma = 1.57$, indicating the closest match between optimized states and ideal TFD states across the range of β considered. (b) Simulated infidelity to the ideal TFD state of states optimized using the cost function $\mathcal{C}_{1.00}$ (blue) and the optimized cost function $\mathcal{C}_{1.57}$ (red).

5.5.2 Supplementary Note 2: Circuit compilation

In this section we present the step-by-step transformation of the circuit in Fig. 1(b) into the equivalent circuit in Fig. 2(b) realizable with the native gate set in our control architecture.

Exponentiation of ZZ and XX: We first substitute the standard decomposition of the operations $e^{-i\phi ZZ/2}$ and $e^{-i\phi XX/2}$ using controlled-NOT (CNOT) gates and single-qubit rotations, shown in Figure 5.3. The decomposition of $e^{-i\phi ZZ/2}$ uses an initial CNOT to transfer the two-qubit parity into the target qubit, followed by a rotation $R_Z(\phi)$ on this target qubit, and a final CNOT inverting the parity. The decomposition of $e^{-i\phi XX/2}$ simply dresses the transformations above by pre- and post-rotations transforming from the X basis to the Z basis and back, respectively. The result of these substitutions is shown in Figure 5.4.




Figure 5.2: Cost function performance comparison. (a) Simulated infidelity of the reduced state of system A to the Gibbs state for states optimized using the cost function $C_{1.00}$ (blue) and the optimized cost function $C_{1.57}$ (red). (b) Corresponding purity of the reduced state. The purity of the Gibbs state is also shown for comparison (black).

5




Figure 5.3: Exponentiation of ZZ and XX . (a) Standard compilation of $e^{-i\phi ZZ/2}$ using CNOT gates and single-qubit rotations. (b) Standard compilation of $e^{-i\phi XX/2}$ using additional basis pre- and post-rotations.




Figure 5.4: Compilation step 1. Depth-14 circuit obtained by replacing the ZZ and XX exponentiation steps in Fig. 1(b) with the circuits of Figure 5.3.

Compilation using native gate set: The native gate set consists of single-qubit rotations of the form $R_{XY}(\phi, \theta)$ and CZ gates. We compile every CNOT in Figure 5.4 as a circuit using native gates, shown in Figure 5.5. Note that $R_Y(\theta) = R_{XY}(90^\circ, \theta)$. Using this replacement together with the identities $R_Y(-90^\circ)R_Z(\phi)R_Y(90^\circ) = R_X(-\phi)$ and $R_Y(-90^\circ)R_X(\phi)R_Y(90^\circ) = R_Z(\phi)$ leads to the circuit in Figure 5.6.

Reduction of circuit depth: Exploiting the commutations in Figure 5.7 together with the identities $R_Y(-90^\circ) = R_Y(180^\circ)R_Y(90^\circ)$ and $R_Y(180^\circ)R_X(\phi)R_Y(180^\circ) = R_X(-\phi)$, we




Figure 5.5: Compilation of CNOT gate using native gates: one CZ sandwiched by single-qubit rotations on the target qubit.




Figure 5.6: Compilation step 2. Depth-16 circuit obtained by replacing every CNOT in Figure 5.4 with the circuit of Figure 5.5 and using simple identities.

can bring two identical pairs of CZ gates back-to-back and cancel them out (since CZ is its own inverse). This leads to the circuit in Figure 5.8.




Figure 5.7: Commutation of $R_Y(180^\circ)$ and CZ gates.

Elimination of Z rotations: All the R_Z gates in Figure 5.8 can be propagated to the beginning of the circuit using the commutation relation

$$R_Z(\alpha)R_{XY}(\phi, \theta) = R_{XY}(\phi + \alpha, \theta)R_Z(\alpha) \quad (5.17)$$

and noting that R_Z commutes with CZ. State $|0\rangle$ is an eigenstate of all R_Z rotations, so we can ignore all R_Z gates at the start because they only produce a global phase. This action leads to the final depth-11 circuit shown in Figure 5.9, which matches that of Fig. 2(b) upon adding measurement pre-rotations and final measurements on all qubits.




Figure 5.8: Compilation step 3. Depth-12 circuit obtained by applying the commutation rule in Figure 5.7 and simple identities to the circuit of Figure 5.6.




Figure 5.9: Compilation step 4. Depth-11 circuit obtained by propagating all R_Z gates in Figure 5.8 to the beginning of the circuit and then eliminating them.

5.5.3 Supplementary Note 3: Device and transmon parameters at bias point

Our experiment makes use of four transmons with square connectivity within a seven-transmon processor. Figure 5.10 provides an optical image zoomed in to this transmon patch. Each transmon has a flux-control line for two-qubit gating, a microwave-drive line for single-qubit gating, and dispersively-coupled resonator with Purcell filter for readout [26, 178]. The readout-resonator/Purcell-filter pair for B_1 is visible at the center of this image. A vertically running common feedline connects to all Purcell filters, enabling simultaneous readout of the four transmons by frequency multiplexing. Air-bridge crossovers enable the routing of all input and output lines to the edges of the chip, where they connect to a printed circuit board through aluminum wirebonds. The four transmons are biased to their sweetspot using static flux bias to counter any residual offset. Table 5.1 presents measured transmon parameters at this bias point.

5.5.4 Supplementary Note 4: Experimental setup

The device was mounted on a copper sample holder attached to the mixing chamber of a Bluefors XLD dilution refrigerator with 12 mK base temperature. For radiation shielding, the cold finger was enclosed by a copper can coated with a mixture of Stycast 2850 and silicon




Figure 5.10: Optical image of the device, zoomed in to the four transmons used in the experiment. Added false color highlights the transmon pair of system A (blue, A₁, A₂) the transmon pair of system B (red, B₁, B₂), and the dedicated bus resonators used to achieve intra-system (read and blue) and inter-system coupling (purple).

Transmon	B ₁	B ₂	A ₁	A ₂
Sweetspot frequency (GHz)	6.433	5.771	5.887	4.534
Relaxation time T_1 (μ s)	32.1	40.7	64.0	33.7
Echo dephasing time T_2^{echo} (μ s)	29.9	40.5	45.9	68.8
Readout frequency (GHz)	7.493	7.225	7.058	6.913
Average assignment fidelity (%)	96.5	96.5	97.0	93.8

Table 5.1: Summary of measured transmon parameters at bias point.

carbide granules (15–1000 nm diameter) used for infrared absorption. To shield against external magnetic fields, the can was enclosed by an aluminum can and two Cryophy cans. Microwave-drive lines were filtered using ~ 60 dB of attenuation with both commercial cryogenic attenuators and home-made Eccosorb filters for infrared absorption. Flux-control lines were also filtered using commercial low-pass filters and Eccosorb filters with stronger absorption. Flux pulses for CZ gates were coupled to the flux-bias lines via room-temperature bias tees. Amplification of the readout signal was done in three stages: a travelling-wave parametric amplifier (TWPA, provided by MIT-LL [151]) located at the mixing chamber plate, a Low Noise Factory HEMT at the 4 K plate, and finally a Miteq amplifier at room temperature.

Room-temperature electronics used both commercial hardware and custom hardware developed in QuTech. Rohde & Schwarz SGS100 sources provided all microwave signals for single-qubit gates and readout. Home-built current sources (IVVI racks) provided static flux biasing. QuTech arbitrary waveform generators (QWG) generated the modulation envelopes for single-qubit gates and a Zurich Instruments HDAWG-8 generated the flux pulses for CZ gates. A Zurich Instruments UHFQA was used to perform independent readout of the four qubits. QuTech mixers were used for all frequency up- and down-conversion. The QuTech Central Controller (QCC) coordinated the triggering of the QWG, HDAWG-8 and UHFQA. See Figure 5.11 for a wiring diagram of the experimental setup.

All measurements were controlled at the software level with QCoDeS [152] and PycQED [153] packages. The QuTech OpenQL compiler translated high-level Python code into the eQASM code [154] forming the input to the QCC.

5.5.5 Supplementary Note 5: Gate performance

The gate set in our quantum processor consists of single-qubit rotations $R_{XY}(\phi, \theta)$ and two-qubit CZ gates. Single-qubit rotations are implemented as DRAG-type microwave pulses with total duration $4\sigma = 20$ ns, where σ is the Gaussian width of the main-quadrature Gaussian pulse envelope. We characterize single-qubit gate performance by single-qubit Clifford randomized benchmarking (100 seeds per run) with modifications to detect leakage, keeping all other qubits in $|0\rangle$. Two-qubit CZ gates are implemented using the Net Zero flux-pulsing scheme, with strong pulses acquiring the conditional phase in 70 ns and weak pulses nulling single-qubit phases in 10 ns. Intra-system and inter-system CZ gates were simultaneously tuned in pairs (using conditional-oscillation experiments as in [48]) in order to reduce circuit depth. However, we characterize CZ gate performance individually using two-qubit interleaved randomized benchmarking (100 seeds per run) with modifications to detect leakage, keeping the other two qubits in $|0\rangle$. Figure 5.12 presents the extracted infidelity and leakage for single-qubit gates (circles) and CZ gates (squares).

5.5.6 Supplementary Note 6: Residual ZZ coupling at bias point

Coupling between nearest-neighbor transmons in our device is realized using dedicated coupling bus resonators. The non-tunability of these couplers leads to residual ZZ coupling between the transmons at the bias point. We quantify the residual ZZ coupling between every pair of transmons as the shift in frequency of one when the state of the other changes from $|0\rangle$ to $|1\rangle$ [188]. We extract this frequency shift using the simple time-domain measurement shown in Figure 5.13(a): we perform a standard echo experiment on one qubit (the echo qubit), but add a π pulse on the other qubit (control qubit) halfway through the free-evolution period simultaneous with the refocusing π pulse on the echo qubit. An example measurement with B_2 as the echo qubit and B_1 as the control is shown in Figure 5.13(b). The complete results for all Echo-qubit, control-qubit combinations are presented as a matrix in Figure 5.13(c). We observe that the residual ZZ coupling is highest between B_1 and the mid-frequency qubits B_2 and A_1 . This is consistent with the higher (lower) absolute de-




Figure 5.11: Fridge wiring and electronic setup.

tuning and the lower (higher) transverse coupling between A_2 (B_1) and the mid-frequency transmons.

5.5.7 Supplementary Note 7: Measurement models, cost function evaluation, and two-qubit state tomography

In this section we present detailed aspects of measurement as needed for evaluation of $\mathcal{C}$ and for performing two-qubit state tomography. We begin by characterizing the fidelity and




Figure 5.12: Single- and two-qubit gate performance. (a) Infidelity and (b) leakage of single-qubit gates (circles) and CZ gates (squares), extracted by randomized benchmarking.




Figure 5.13: Characterization of residual ZZ coupling. (a) Modified echo experiment to determine the shift in frequency of one qubit (echo qubit) when another (the control qubit) changes state from $|0\rangle$ to $|1\rangle$. (b) Example data for pair B_2 (echo) - B_1 (control). (c) Table of extracted frequency shifts for all pairs of echo and control qubits.

crosstalk of simultaneous single-qubit measurements using the cross-fidelity matrix as defined in [27]:

$$F_{ji} = 1 - \text{Prob}(e_j | I_i) - \text{Prob}(g_j | \pi_i), \quad (5.18)$$

where e_j (g_j) denotes the assignment of qubit j to the $|1\rangle$ ($|0\rangle$) state, and π_i (I_i) denotes the preparation of qubit i in $|1\rangle$ ($|0\rangle$). The measured cross-fidelity matrix for the four qubits is shown in Figure 5.14. From diagonal element F_{ii} we extract the average assignment fidelity for qubit i , the latter given by $1/2 + F_{ii}/2$ and quoted in Table 5.1. The magnitude of the

off-diagonal elements F_{ji} with $j \neq i$ quantifies readout crosstalk, and is below 2% for all pairs. This low level of crosstalk justifies using the simple measurement models that we now describe.




Figure 5.14: Measured readout cross-fidelity matrix.

5

5.5.8 Measurement models

The evaluation of the cost function and two-qubit state tomography require estimating the expected value of single-qubit and two-qubit Pauli operators. We do so by least-squares linear inversion of the experimental average of single-transmon measurements and two-transmon correlation measurements.

When measuring transmon i , we 1-bit discretize the integrated analog signal for its readout channel at every shot, outputting $m_i = +1$ when declaring the transmon in $|0\rangle$ and $m_i = -1$ when declared it in $|1\rangle$. The expected value of $\langle m_i \rangle$ is given by $\bar{m}_i = \text{Tr}(M_i \rho_{\text{Exp}})$, where the measurement operator (in view of the low crosstalk) is modelled as

$$M_i = \sum_{k=0} c_k^i |k_i\rangle \langle k_i|, \quad (5.19)$$

where k denotes transmon excitation level, and $c_k^i \in [-1, 1]$ are real-valued coefficients. Making use of the 1-qubit Pauli operators $I_i = |0_i\rangle \langle 0_i| + |1_i\rangle \langle 1_i|$ and $Z_i = |0_i\rangle \langle 0_i| - |1_i\rangle \langle 1_i|$ and truncating to three transmon levels, we can rewrite the measurement operator in the form

$$M_i = c_I^i I_i + c_Z^i Z_i + c_2^i |2_i\rangle \langle 2_i|, \quad (5.20)$$

also with real-valued c_I^i and $c_Z^i \in [-1, 1]$.

When correlating measurements on transmons i and j , we compute the product of the 1-bit discretized output for each transmon. The expected value of $m_{ji} = m_j \times m_i$ is given by $\langle m_{ji} \rangle = \text{Tr}(M_{ji} \rho_{\text{Exp}})$, where the measurement operator (also in view of the low crosstalk) is modelled as

$$M_{ji} = \sum_{k,l} c_{lk}^{ji} |l_j k_i\rangle \langle l_j k_i|, \quad (5.21)$$

with real-valued coefficients $c_{lk}^{ji} \in [-1, 1]$. Making use of the 2-qubit Pauli operators given by tensor products, and again truncating to three transmon levels, we can rewrite the measurement operator in the form

$$M_{ji} = c_{II}^{ji} I_j I_i + c_{IZ}^{ji} I_j Z_i + c_{ZI}^{ji} Z_j I_i + c_{ZZ}^{ji} Z_j Z_i + c_{2I}^{ji} |2_j\rangle\langle 2_j| I_i + c_{2Z}^{ji} |2_j\rangle\langle 2_j| Z_i + c_{2I}^{ji} I_j |2_i\rangle\langle 2_i| + c_{2Z}^{ji} I_j |2_i\rangle\langle 2_i| + c_{22}^{ji} |2_j 2_i\rangle\langle 2_j 2_i|. \quad (5.22)$$

In experiment, we calibrate the coefficients c_P^i and c_{PQ}^{ji} ($P, Q \in I, Z$) by linear inversion of the experimental average of single-transmon and correlation measurements with the four transmons prepared in each of the 16 computational states (for which $\langle P_i \rangle = \pm 1$ and $\langle Q_j P_i \rangle = \pm 1$). We do not calibrate the coefficients c_2^i , c_{2P}^{ji} , c_{Q2}^{ji} , or c_{22}^{ji} .

Measurement pre-rotations change the measurement operator as follows: $R_X(180^\circ)$ on transmon i transforms $Z_i \rightarrow -Z_i$; $R_X(\pm 90^\circ)$ transforms $Z_i \rightarrow \pm Y_i$; and $R_Y(\pm 90^\circ)$ transforms $Z_i \rightarrow \mp X_i$. Pre-rotations do not transform the projectors $|2_i\rangle\langle 2_i|$ as they only act on the qubit subspace.

5

5.5.9 Cost function evaluation

To evaluate the cost function $\mathcal{C}$, we must estimate the expected value of all single-qubit Pauli operators X_i , the two intra-system two-qubit Pauli operators $Z_j Z_i$, and the inter-system two-qubit Pauli operators $X_j X_i$ and $Z_j Z_i$ [the latter only between corresponding qubits in the two systems (e.g., B_1 and A_1)]. We estimate these by linear inversion of the experimental averages (based on 4096 measurements) of single-transmon and relevant correlation measurements with the transmons measured in the bases specified in Table 5.2. As an example, Figure 5.15 shows the raw data for the estimation of $\mathcal{C}$ with variational parameters $(\vec{\alpha}, \vec{\gamma}) = 0$. Note that every evaluation of the cost function includes readout calibration measurements to extract measurement-operator coefficients c_P^i and c_{PQ}^{ji} ($P, Q \in I, Z$).

#	Z basis measurements	X basis measurements
1	$+Z_{B2}, +Z_{B1}, +Z_{A2}, +Z_{A1}$	$+X_{B2}, +X_{B1}, +X_{A2}, +X_{A1}$
2	$+Z_{B2}, +Z_{B1}, +Z_{A2}, -Z_{A1}$	$+X_{B2}, +X_{B1}, +X_{A2}, -X_{A1}$
3	$+Z_{B2}, +Z_{B1}, -Z_{A2}, +Z_{A1}$	$+X_{B2}, +X_{B1}, -X_{A2}, +X_{A1}$
4	$+Z_{B2}, -Z_{B1}, +Z_{A2}, +Z_{A1}$	$+X_{B2}, -X_{B1}, +X_{A2}, +X_{A1}$
5	$-Z_{B2}, +Z_{B1}, +Z_{A2}, +Z_{A1}$	$-X_{B2}, +X_{B1}, +X_{A2}, +X_{A1}$
6	$+Z_{B2}, +Z_{B1}, -Z_{A2}, -Z_{A1}$	$+X_{B2}, +X_{B1}, -X_{A2}, -X_{A1}$
7	$-Z_{B2}, -Z_{B1}, +Z_{A2}, +Z_{A1}$	$-X_{B2}, -X_{B1}, +X_{A2}, +X_{A1}$
8	$+Z_{B2}, -Z_{B1}, +Z_{A2}, -Z_{A1}$	$+X_{B2}, -X_{B1}, +X_{A2}, -X_{A1}$
9	$-Z_{B2}, +Z_{B1}, -Z_{A2}, +Z_{A1}$	$-X_{B2}, +X_{B1}, -X_{A2}, +X_{A1}$

Table 5.2: List of the basis measurements used to evaluate all the terms in the cost function $\mathcal{C}$. The bases in the left column are used to estimate intra- and inter-system terms $\langle Z_j Z_i \rangle$, while those on the right are used to extract single-qubit terms $\langle X_i \rangle$ and inter-system terms $\langle X_j X_i \rangle$.




Figure 5.15: Example raw data for evaluating the cost function with variational angles $(\vec{\alpha}, \vec{\gamma}) = 0$. Experimental average (from 4096 shots) of single-qubit and correlation measurements for the 18 measurement bases in Table 5.1 and for calibration of the measurement operators (following preparation of the 16 computational states of the four transmons). Panels show the experimental averages of (a-d) single-transmon measurements, (e,h) intra-system correlations and (f,g) inter-system correlations. (i) Bar graph of the experimental estimate of the expected values of all terms in $\mathcal{C}$, obtained by linear inversion on the measurement averages. For comparison, we also show are the expected values of all the terms for an ideal processor.

5.5.10 Two-qubit state tomography

After optimization, we perform two-qubit state tomography of each system separately to assess performance. To do this, we obtain experimental averages (from 16394 shots) of single-transmon and correlation measurements using an over-complete set of measurement bases.

This set consists of the 36 bases obtained by using all combinations of bases for each transmon i and j , drawn from the set $\{+X, +Y, +Z, -X, -Y, -Z\}$. The expectation values of single-qubit and two-qubit Pauli operators are estimated by least-squares linear inversion. Finally, these are used to construct

$$\rho_{\text{Exp}} = \frac{1}{4} \left(I_j I_i + \sum_{P \in X, Y, Z} \langle P_i \rangle_{\text{est}} I_j P_i + \sum_{P \in X, Y, Z} \langle P_j \rangle_{\text{est}} P_j I_i + \sum_{P, Q \in X, Y, Z} \langle Q_j P_i \rangle_{\text{est}} Q_j P_i \right). \quad (5.23)$$

5.5.1.1 Supplementary Note 8: Impact of leakage on two-qubit tomography

Our linear inversion procedure for converting measurement averages into estimates of the expected value of one- and two-qubit Pauli operators is only valid for $\langle |l_j k_i\rangle \langle l_j k_i| \rangle = 0$ whenever either k or $l \geq 2$, i.e., there is no leakage on either transmon. It is therefore essential, particularly for simulation model 4, to understand precisely how leakage in either or both transmons infiltrates our extraction of the two-qubit density matrix ρ_{Exp} .

First we consider the estimation of expected values of single-qubit Pauli operators $\langle P_i \rangle$, taking $\langle Z_i \rangle$ as a concrete example. The expected value of all measurements on this transmon for the basis combinations (6 in total) where this specific transmon is measured in the $+Z$ basis is:

$$\langle m_i^{(+)} \rangle = c_I^i \langle I_i \rangle + c_Z^i \langle Z_i \rangle + c_2^i \langle |2_i\rangle \langle 2_i| \rangle. \quad (5.24)$$

In turn, the expected value of all measurements on this transmon for the basis combinations (6 in total) where this specific transmon is measured in the $-Z$ basis is:

$$\langle m_i^{(-)} \rangle = c_I^i \langle I_i \rangle - c_Z^i \langle Z_i \rangle + c_2^i \langle |2_i\rangle \langle 2_i| \rangle. \quad (5.25)$$

Note that the contribution from the leakage term is unchanged because the pre-rotation (I_i or $R_X(180^\circ)$) only acts on the qubit subspace. Our least-squares linear inversion of these 12 experimental averages to estimate $\langle Z_i \rangle$ is

$$\langle Z_i \rangle_{\text{est}} = \frac{1}{12c_Z^i} \left(\sum_{+} \bar{m}_i^{(+)} - \sum_{-} \bar{m}_i^{(-)} \right). \quad (5.26)$$

Clearly, owing to the balanced nature of this linear combination (all coefficients of equal magnitude, 6 positive and 6 negative), this estimator is not biased by c_2^i . In other words, the average of $\langle Z_i \rangle_{\text{est}}$ is independent of the value of c_2^i .

Consider now the estimation of the expected value of two-qubit Pauli operators $\langle Q_j P_i \rangle$, taking $\langle X_j Z_i \rangle$ as a concrete example. There are four correlation measurements that contain this term. For measurement bases $+X_j$ and $+Z_i$,

$$\begin{aligned} \langle m_{ji}^{(++)} \rangle = & c_{II}^{ji} \langle I_j I_i \rangle + c_{IZ}^{ji} \langle I_j Z_i \rangle + c_{ZI}^{ji} \langle X_j I_i \rangle + c_{ZZ}^{ji} \langle X_j Z_i \rangle + c_{II}^{ji} \langle |2_j\rangle \langle 2_j| I_i \rangle + \\ & c_{Z2}^{ji} \langle |2_j\rangle \langle 2_j| Z_i \rangle + c_{I2}^{ji} \langle I_j |2_i\rangle \langle 2_i| \rangle + c_{Z2}^{ji} \langle X_j |2_i\rangle \langle 2_i| \rangle + c_{22}^{ji} \langle |2_j 2_i\rangle \langle 2_j 2_i| \rangle. \end{aligned} \quad (5.27)$$

For measurement bases $+X_j$ and $-Z_i$,

$$\begin{aligned} \langle m_{ji}^{+-} \rangle = & c_{II}^{ji} \langle I_j I_i \rangle - c_{IZ}^{ji} \langle I_j Z_i \rangle + c_{ZI}^{ji} \langle X_j I_i \rangle - c_{ZZ}^{ji} \langle X_j Z_i \rangle + c_{II}^{ji} \langle |2_j\rangle \langle 2_j| I_i \rangle - \\ & c_{ZI}^{ji} \langle |2_j\rangle \langle 2_j| Z_i \rangle + c_{I2}^{ji} \langle I_j |2_i\rangle \langle 2_i| \rangle + c_{Z2}^{ji} \langle X_j |2_i\rangle \langle 2_i| \rangle + c_{22}^{ji} \langle |2_j 2_i\rangle \langle 2_j 2_i| \rangle. \end{aligned} \quad (5.28)$$

For measurement bases $-X_j$ and $+Z_i$,

$$\begin{aligned} \langle m_{ji}^{-+} \rangle = & c_{II}^{ji} \langle I_j I_i \rangle + c_{IZ}^{ji} \langle I_j Z_i \rangle - c_{ZI}^{ji} \langle X_j I_i \rangle - c_{ZZ}^{ji} \langle X_j Z_i \rangle + c_{II}^{ji} \langle |2_j\rangle \langle 2_j| I_i \rangle + \\ & c_{ZI}^{ji} \langle |2_j\rangle \langle 2_j| Z_i \rangle + c_{I2}^{ji} \langle I_j |2_i\rangle \langle 2_i| \rangle - c_{Z2}^{ji} \langle X_j |2_i\rangle \langle 2_i| \rangle + c_{22}^{ji} \langle |2_j 2_i\rangle \langle 2_j 2_i| \rangle. \end{aligned} \quad (5.29)$$

Finally, for measurement bases $-X_j$ and $-Z_i$,

$$\begin{aligned} \langle m_{ji}^{--} \rangle = & c_{II}^{ji} \langle I_j I_i \rangle - c_{IZ}^{ji} \langle I_j Z_i \rangle - c_{ZI}^{ji} \langle X_j I_i \rangle + c_{ZZ}^{ji} \langle X_j Z_i \rangle + c_{II}^{ji} \langle |2_j\rangle \langle 2_j| I_i \rangle - \\ & c_{ZI}^{ji} \langle |2_j\rangle \langle 2_j| Z_i \rangle + c_{I2}^{ji} \langle I_j |2_i\rangle \langle 2_i| \rangle - c_{Z2}^{ji} \langle X_j |2_i\rangle \langle 2_i| \rangle + c_{22}^{ji} \langle |2_j 2_i\rangle \langle 2_j 2_i| \rangle. \end{aligned} \quad (5.30)$$

Our least-squares linear inversion of these 4 experimental averages to estimate $\langle X_j Z_i \rangle$ is

$$\langle X_j Z_i \rangle_{\text{est}} = \frac{1}{4c_{ZZ}^{ji}} \left(\bar{m}_{ji}^{++} - \bar{m}_{ji}^{+-} - \bar{m}_{ji}^{-+} + \bar{m}_{ji}^{--} \right). \quad (5.31)$$

Clearly, owing to the balanced nature of this linear combination (all coefficients of equal magnitude, 2 positive and 2 negative), this estimator is not biased by c_{2I}^{ji} , c_{2Z}^{ji} , c_{I2}^{ji} , c_{Z2}^{ji} , and c_{22}^{ji} . In other words, the average of $\langle X_j Z_i \rangle_{\text{est}}$ is independent of the value of these coefficients.

We are finally in position to describe how leakage in the two-transmon system infiltrates into our two-qubit tomographic reconstruction procedure. Evidently, the complete description of the two-transmon system would be a two-qutrit density matrix $\rho_{2\text{Qutrit}}$, but our procedure returns a two-qubit density matrix ρ_{Exp} . It is therefore key to understand how elements of $\rho_{2\text{Qutrit}}$ are mapped onto ρ_{Exp} . Table 5.3 summarizes these mappings and Figure 5.16 illustrates them, including several examples. We have verified the mappings by exactly replicating the tomographic procedure in our numerical simulation using *quantumsim*. To incorporate this into the simulations, we have made use of the measurement coefficients c_i experimentally obtained. These leakage mappings have also been used when adding leakage in simulation model 4.

5.5.12 Supplementary Note 9: Error model for numerical simulations

Our numerical simulations use the *quantumsim* [182] density-matrix simulator with the error model described in the *quantumsim_dclab* subpackage.

Single-qubit gates are modeled as perfect rotations, sandwiched by two 10 ns idling blocks. The idling model takes into account amplitude damping (T_1), phase damping (T_2^{echo}) (noise model 1 of the main text), and residual ZZ crosstalk (noise model 3). To implement

Elements of $\rho_{2\text{Qutrit}}$	Mapping onto ρ_{Exp}
$ l'_j, k'_i\rangle \langle l_j, k_i $	$ l'_j, k'_i\rangle \langle l_j, k_i $
$ l'_j, 2_i\rangle \langle l_j, 2_i $	$\frac{1}{2} l'_j\rangle \langle l_j I_i$
$ 2_j, k'_i\rangle \langle 2_j, k_i $	$\frac{1}{2} I_j k'_i\rangle \langle k_i $
$ 2_j, 2_i\rangle \langle 2_j, 2_i $	$\frac{1}{4} I_j I_i$
All other elements	0

Table 5.3: Mapping of the elements of $\rho_{2\text{Qutrit}}$ to ρ_{Exp} by our two-qubit state tomography procedure. Here, $k, l, k', l' \in \{0, 1\}$ denote computational (unleaked) states.




Figure 5.16: Illustration of the impact of leakage on two-qubit state tomography. (a) Mapping of two-qutrit density matrix elements onto elements of the two-qubit density matrix ρ_{Exp} returned by our tomographic reconstruction procedure. Elements in the qubit subspace (blue) are mapped correctly. Elements corresponding to one transmon in the leaked state (orange and green) are mapped onto a two-qubit state with the corresponding qubit fully mixed. The state with leakage on both transmons (red) is mapped to the fully mixed two-qubit state. (b-e) Example mappings for the states (b) $|--\rangle$, (c) $|+, 2\rangle$, (d) $|2, -i\rangle$, and (e) $|22\rangle$, where $|--\rangle = \frac{1}{\sqrt{2}}(|0\rangle - |1\rangle)$ and $|\pm i\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm i|1\rangle)$.

it, we first split the idling intervals into slices of 10 ns or less. These slides include amplitude and phase damping. Between these slices, we add instantaneous two-qubit gates capturing the residual coupling described by the Hamiltonian [188]:

$$H = \zeta_{ij} |11\rangle \langle 11| = -\frac{\zeta_{ij}}{2} (1 - Z_i - Z_j + Z_i Z_j). \quad (5.32)$$

The measurements of T_1 , T_2^{echo} and ζ_{ij} are detailed in previous sections.

The error model for CZ gates is described in detail in [187]. The dominant error sources are identified to be leakage of the fluxed transmon to the second-excited state through the $|11\rangle \leftrightarrow |02\rangle$ channel and increased dephasing (reduced T_2^{echo}) due to the fact that the fluxed transmon is pulsed away from the sweetspot. These two effects implement noise models 4 and 2, respectively. The two-transmon process is modeled as instantaneous, and sandwiched by two idling blocks of 35 ns with decreased T_2^{echo} on the fluxed transmon. Quasistatic flux noise is suppressed to first order in the Net Zero scheme and is therefore neglected. Residual ZZ crosstalk is not inserted during idling for the transmon pair, because it is absorbed by the gate calibration. The error model described in [187] allows for higher-order leakage effects, e.g., so-called leakage conditional phases and leakage mobility. We do not include these effects.

The simulation finishes by including the effect of leakage on the tomographic procedure as discussed in Supplementary Note 8. The density matrix is obtained at the qutrit level, and the correct mapping for the density-matrix elements is applied. We take special care to use the experimental readout coefficients c_i to model the readout signal for the simulated density matrix, according to Equation (5.20) and Equation (5.21). The simulation produces data for the same basis set as shown in Figure 5.15. Afterwards, the same tomographic state reconstruction routine as in the experiment is applied to these data. In this way, noise model 4 properly accounts for the imperfect reconstruction of leaked states, providing a fair comparison to experiment.

5.5.13 Supplementary Note 10: Classical optimizer under ideal conditions

Our simulation efforts so far have addressed noise sources originating from physical effects (e.g., qubit relaxation, dephasing, ZZ crosstalk, and leakage). It is important to also verify that the classical optimizer does not limit the performance of the experiment. To this end, we make use of the same optimization code [153] as in experiment while replacing qubit readout data with simulated readout data for a noiseless quantum processor. In this way, we can establish the extent to which our classical optimization limits performance. We observe that the achieved infidelity under these conditions [Figure 5.17(a)] is three orders of magnitude lower than in the experiment. Furthermore, we observe that it closely follows the trend expected from cost-function design [Figure 5.2(a)]. We therefore confirm that the addressed physical noise sources dominate in experiment, while optimizer limitations would only become relevant upon reducing infidelity by approximately three orders of magnitude.




Figure 5.17: Test of the classical optimizer. By feeding simulated readout data for an ideal processor (instead of actual measurements) into the software stack used to control our experiment [153], we verify the limitations of the optimizer under the exact configuration and conditions used in the experiment. Both figures include the curves from Fig. S3 (red) as reference. (a) Simulated infidelity (green) of the reduced state of system A to the Gibbs state after optimization. (b) Corresponding purity of the reduced state (green). The purity of the Gibbs state is also shown for comparison (black).

6.1 Summary

During this thesis we have tried several runs at executing quantum algorithms on the available hardware. As expected from simulations, quantum mechanics (via several noise mechanisms) alters the performance of such algorithms and in many cases renders them useless. Circumventing these errors is an unavoidable necessity in order to use quantum computers without fault-tolerant codes. The experiments of Chapter 3, Chapter 4 and Chapter 5 allowed for the build-up of simulation methods, capable of explaining and anticipating how theoretical formulations for algorithms will behave (and possibly fail) because of the uncorrected errors quantum processors possess. Not only we learned to predict such errors, but also to mitigate them, introducing the Symmetry Verification technique. While using similar tools as stabilizers in surface codes, we avoided the prohibitive resources such codes would need, while still mitigating certain errors.

Looking backwards, it is useful to reflect upon what would be interesting to do, if repeating the experiment. For the case of Chapter 3, it would be of high interest to attempt the preparation of the system ground state, either by adiabatic preparation in a Trotterized way, or with a QAOA approach. The later would be especially interesting, as the Trotterized evolution was so carefully implemented, and would allow an interesting comparison to the performance of repeated QAOA iterations. In the experiment of Chapter 4, it would certainly be better to use Symmetry Verification directly on the optimization routine, instead of afterwards. Furthermore, the Hydrogen case allows for the scaling of the register used for the algorithm, from 1 to at least 4 qubits (and of course extendable to more). Such study was only very recently conducted by the Google Team [43] [191]. Finally, for the experiment of Chapter 5, it would also be very informative to understand how the application of several layers of the QAOA unitaries improve the performance of the algorithm and its matching with our understanding of the leading error mechanisms.

6.2 The promise of the variational approach

Another important topic to reflect upon is that this thesis contains two very different types of quantum simulations in terms of the algorithmic strategy. On Chapter 3 we produced a trotterized implementation of an atom-cavity system. While Trotterization is a very clear and straight-forward method for implementing arbitrary Hamiltonians, it remains experimentally challenging as its guarantees of performance are asymptotic, and push the coherence requirements of the algorithm. Specifically in Chapter 3, we managed to produce qubits with very competitive coherence at that point in time, and push the state of the art with 90 Trotter steps. However, as system size increases the cost of every Trotter step does as well, making it very difficult to scale up. On the other hand, around the time of this thesis, other family of algorithms started taking attention from the community: the variational algorithms. The idea was very well received because it was pragmatic: instead of focusing on circuit depths state of the art could not afford, it focused on what could be done at the time. It managed very well to minimize the set of concepts necessary to implement it: (i) a low-depth circuit optimized by a classical code, (ii) a mapping between qubit states and solutions spaces and

(iii) a target metric whose minimization converges to a solution of the given problem. These set of concepts divided the algorithm in items to develop and improve upon independently. Improving these items to the levels required for useful applications seem plausible in a much shorter time-scale than pushing coherences by one or two orders of magnitude as required for Trotterized algorithms.

While variational algorithms seem very promising and currently hold the lead towards useful applications, it remains clear that more improvements are necessary for the several items composing the algorithm. One of the most solid progresses so far has been the effort of packages like OpenFermion [52], that automates the mapping of quantum chemistry problems onto Hamiltonians for N qubit processors. The Hydrogen molecule is a good example of how by sacrificing simplicity of the cost function, one can reduce the qubits required from 4 to even 1. This is achieved by, instead of directly representing one orbital per qubit, resorting to fancier mapping techniques such as Bravyi-Kitaev one [51]. Furthermore, one can sacrifice a bit more complexity by adding extra qubit dimensions with symmetries [57] or subspace expansion terms [126]. Handling such a procedure is a bookkeeping task, which a layer of software is ideally suited to perform. The further development of such software packages is an ongoing effort, that will most certainly contribute to the first commercial attempts at using quantum computers.

Another important feature is that of the ansatzes, and the gates composing them. While transmon architectures already achieved universal gate-sets via CZ gates, further two-qubit gates might be very helpful in reducing the necessary depth for certain unitaries. Specifically in the context of quantum chemistry, swap-like gates provide (up to single-qubit phases) a native implementation of the population balance needed to solve between orbitals [56]. Recently, the Google Team presented a continuous family of gates, parametrized with two degrees of freedom and including both CZ and iSWAP gates [189]. As depth reduction directly affects the impact on coherence, using such tricks to further reduce the depth can provide good returns in terms of performance.

The methods by which we construct an ansatz were also pushed in the recent years. The IBM team [31] introduced the concept of *Hardware efficient ansatz*, where instead of implementing a theory-motivated circuit, one simply states the most general unitary possible within the architecture. In the case of transmon architectures, it is a layer of single qubit gates, where unitaries are totally arbitrary $SU(2)$ ones implemented via virtual-Z compilation, followed by a layer of the entangling gates available. In principle, as the gate-set is universal, this should suffice to implement any unitary, provided enough depth is accessible. However, such a pragmatic approach needs to be exercised with care: the excessive parametrization of such an ansatz can lead to plateaus for the classical optimization of the parameters, stagnating the variational process [142]. The matter of an optimal procedure to deciding which parameters to include in the optimization and which not, remains an open question. It is my personal view that, towards scaling up these algorithms, such a procedure will be necessary to avoid stagnation via Baren plateaus.

In hindsight, all experiment chapters share a common insight: performance would have been much worse, had it not been for inventive small tricks. The experiment of Chapter 3

uses a second-order Trotterization, Chapter 4 introduced a brand new mitigation technique, and Chapter 5 avoided the need of bath engineering for controlling mixture in states.

All of this tricks operate in the same essence as *hacks* in classical computation. A hack¹ occurs when a classical computing system is manipulated to perform a task B for which it was not designed, instead of the task A it was designed for. In the case of quantum computers it gets a complication: We design a quantum processor to do a task A , but it only manages to produce a noisy implementation A_0 . Our hacking task consists in making a task B , which is closer to A than A_0 , improving our performance. In Chapter 4 we realized a symmetry dictating how the Hilbert space is used, and in Chapter 5 we artificially mixed a state by erasing half of it. It is my personal view and conclusion that such *hacks* are going to be necessary if we aim at squeezing the most we can out of the computing power of pre-fault-tolerant quantum processors.

Furthermore, it is my personal view that such tricks, along with the continuous improvements on every aspect of variational approaches (mapping, ansatz, classical optimizers), should allow us to implement useful applications. All this occurring much earlier than the landmark of thousands of qubits that seems necessary to start having small quantum error correction codes. I finish this thesis with high hopes of reading about useful applications for *quantum customers*, running within the next 5 – 10 years.

6

6.3 A personal forecast

To conclude this thesis, I prefer to be more precise, and share in more detail what my vision of such a near-future application looks like.

For starters, current efforts focus on cloud-based platforms granting access to the quantum processors. Although I believe defence and military applications will rather keep their own hardware, I find it hard to imagine this paradigm breaking apart. Secondly, the hardware particularities: I believe features like specific gate-sets, compilation and transpilation tricks and small architecture features will be an active competition ground for the several platforms disputing the market. Offering gate-sets and tools that best implement the necessary connectivities and unitaries will be a topic of importance for any quantum programmer deciding which platform to use.

An integrating piece of software will be needed to orchestrate the several quantum programming layers. While mappings are already quite solved by software like OpenFermion, further layers of code will have to efficiently decide appropriate ansatzes solving the trade-off to place parameters without running into the Barren plateau. Finally, some sort of error mitigation will need to be implemented, symmetry-based ones being currently the most promising at the moment. Finally, the most likely applications seem to be in quantum chemistry and logistics, where the degrees of freedom to model can be modest, but their correlation and interdependence might profit from quantum processors. All this scheme, properly integrated with an adequate engineering effort, seems to be today not only achievable, but a low-hanging fruit ready to be pursued giving further impulse into the usage of quantum computers.

¹While defining the term *hack* may require a much more detailed and exhaustive approach, the quick definition provided here is enough for the purpose of the discussion at hand.

ACKNOWLEDGEMENTS

A thesis is not a project, is a life-path, and no man is an island. This thesis would not have been possible without the many people that walked that path with me.

First of all, **Leo**, a long time has passed since you answered my application for a masters project back in 2014. None of this would have been possible without that opportunity. The role of a lab leader is challenging, and you did it with passion. You mentored us not through a thesis, but through a life-path. You taught me a degree of precision and detail-attention I had never seen, and which improved me on my Achilles heel. But you passed onto us more than that, you passed onto us criterium: to prioritize, so critical in a lab, and for problem-solving, that I can assure plays its part in the more than 4 startups that rely on (many founded by) your alumni. If I had to choose, I'd do it again! Thanks forever.

I would also like to thank to all the committee, **Meno, Chris, Gary, Yaroslav and Juanjo**. This thesis would not be possible without your participation, and I wish you enjoy reading it and the ceremony as well.

My paranympths, **Chris**, we met back in 2014 when I was looking for my masters project, and you ended up being far more than my supervisor. You were the best friend and colleague I could have had. I cannot even recount how many things happened on our path, just that I'm glad you were there. And **Lucas**, I met you at the end of my Dutch times, and yet I had never met a friend this side of the ocean with whom I share so much of my projects. To both of you, I won't say goodbye-like wishes for good. I wish you the best, but I also wish to be there to see it.

There was so many lab-mates I have to thank: **Ale**, you are the ultimate Rick, a true crazy and awesome scientist, I am so glad to have seen your growth from postdoc to fabman of Europe and responsible father and husband. **Nathan**, you walked me on my first steps, and I am super grateful for the time and attention you dedicated me, not only at a technical level. I wish you the best in Australia. **Florian**, with you I learnt the true Swiss precision. Long live the Luthi-meister. **Adriaan**, thanks for so much time shared, you were always a good friend along this path. I wish you the best as a businessman. **Niels**, thanks for all the shared moments, lab and outside. I'm glad you succeeded in your Qblox dream, and I am glad to be so close to it. **Jules**, the same goes to you, I am very glad to remain close in your path. **Thijs**, you started in the lab at around the same time, and you also enjoyed the Argentinian culture for a while. You have been the best Dutch-Argentinian friend I could have asked. I wish you the best in your new quantum job!. **Nandini**, we met before your DiCarlo lab times, and you've been an awesome friend. I wish you the best in your new path as well. **Marios**, we started together but soon diverged our paths, I will always treasure your friendship. **Stefano**, good times! We spent much time together on characterization, I am glad you found your place in the US. **Hany, Jorge**, the new generation! I am super-happy that you decided to join the lab

and hope you the best on your next path. **Malay**, you were my first master student and whole new responsibility. I am glad you enjoyed our field and found your way in Munich. **Berend**, happy to see you to staying in our field, I will always remember those days with the "busy hat" and the "gym-bro lunches". **Viktor**, your fast pass through the lab did not go unnoticed. I wish you the best of lucks. **David**, the gangztah from Portland! I will forever treasure such good times in the cleanroom troops. Hopefully I'll see you around at the ceremony. **Sarwan, Jaap, Rene, Luc, Marc**, you were the youth of the lab, it's nice to see many of you stayed in this field.

But the lab expanded, soon we were more than just experimentalists. We had engineers:

Miguel, Jacob, Wouter and **Nadja**, thanks for taking the lab further than just academics figuring out the next paper. We also shared a lot of efforts with theorists: **Xavi, Tom, Brian, Slava** and **Boris**, you were all not only fun to work with, but also amazing friends outside of the lab.

QuTech also did their part as a whole, with so many friends that turned this path into a community: **Michiel, Damaz, Willemijn, Nodar, Fokko, Srijit** and **Jelmer**, my time in Delft would not have been as enjoyable without you guys. And my favourite theory gang, **Victoria, Glaucia, Jonas** and **Valentina**, always enjoyed hanging out with you guys.

Beyond the institute, I met some incredible people: **Deniz, Raj, Jochem** and all the Wednesdays Jam gang at Chris'. Work-life balance would not have been the same without you.

To my roommates, **James** I moved in with you quite early, I was doubtful of looking for a roommate and I found one of the best friends I could have had. I will always remember the amazing cooking and the Sundays. I wish you and Deniz the best. **Arian**, the quarantine would not have felt the same without you as roommate. We hit it off on many things, the finance topics, and those coffee & waffles Sundays. So glad to see you found your place with Vincy, wish you the best of fortune.

Also to those that supported me at a distance, my lifetime friends **Nacho, Franco, Fede, Edu, Marcos, Miguel, Luis** and **Sofi**, whom I saw barely once a year but remained in touch online. My first physics professor, who encouraged me much earlier than being my teacher, **Jose**, thanks for helping to steer in me in this direction. And of course my family my **family**, who supported me all along.

Ramiro Ernesto SAGASTIZABAL

26-03-1990 Born in Esperanza, Santa Fe, Argentina.

Education

2003–2008 Technical High School
Escuela de Educación Técnica, General don José de San Martín
Esperanza, Santa Fe, Argentina

2009–2014 Licenciatura in Physics (Undergraduate)
Universidad Nacional de Córdoba, Córdoba
Córdoba, Argentina
Lic. thesis: Leyes de escala en problemas de transmisión de esta-
dos cuánticos
(Scaling laws in quantum-state transfer problems)
Supervisor: Dr. Omar Osenda

2014–2015 MSc. in Physics
Université d'Aix-Marseille, Marseille
Provence-Alpes-Côte d'Azur, France
Thesis: First steps towards remote entanglement with transmon
qubits
Supervisors: Dr. Ir. Christian Dickel and Prof. dr. L. DiCarlo

2015–2023 PhD. Physics
Delft University of Technology, Delft, the Netherlands
Thesis: Experimental Quantum Simulation with Noisy
Intermediate-scale Superconducting Processors
Promotors: Prof. dr. L. DiCarlo and Prof. dr. ir. L. P. Kouwenhoven

LIST OF PUBLICATIONS

8. **R. Sagastizabal**, S.P. Premaratne, B.A. Klaver, M.A. Rol, V. Negîrneac, M.S. Moreira, X. Zou, S. Johri, N. Muthusubramanian, M. Beekman, C. Zachariadis, V.P. Ostroukh, N. Haider, A. Bruno, A.Y. Matsuura, L. DiCarlo, *Variational preparation of finite-temperature states on a quantum computer*, npj Quantum Information **7**, 1, 1-7 (2021)
7. V. Negirneac, H. Ali, N. Muthusubramanian, F. Battistel, **R. Sagastizabal**, M.S. Moreira, J.F. Marques, W.J. Vlothuizen, M. Beekman, C. Zachariadis, N. Haider, A. Bruno, L. DiCarlo, *High-Fidelity Controlled- Gate with Maximal Intermediate Leakage Operating at the Speed Limit in a Superconducting Quantum Processor*, Physical Review Letters **126**, 22, 220502 (2021)
6. M.A. Rol, L. Ciorciaro, F. K Malinowski, B.M. Tarasinski, **R. Sagastizabal**, C.C. Bultink, Y. Salathe, N. Haandbæk, J. Sedivy, L. DiCarlo, *Time-domain characterization and correction of on-chip distortion of control pulses in a quantum processor*, Applied Physics Letters **116**, 5, 054001 (2020)
5. T.E. O'Brien, B. Senjean, **R. Sagastizabal**, X. Bonet-Monroig, A. Dutkiewicz, F. Buda, L. DiCarlo, L. Visscher, *Calculating energy derivatives for quantum chemistry on a quantum computer*, npj Quantum Information **5**, 1, 1-12 (2019)
4. **R. Sagastizabal**, X. Bonet-Monroig, M. Singh, M.A. Rol, C.C. Bultink, X. Fu, C.H. Price, V.P. Ostroukh, N. Muthusubramanian, A. Bruno, M. Beekman, N. Haider, T.E. O'Brien, L. DiCarlo, *Experimental error mitigation via symmetry verification in a variational quantum eigensolver*, Physical Review A **100**, 1, 010302 (2019)
3. X. Bonet-Monroig, **R. Sagastizabal**, M. Singh, T.E. O'Brien, *Low-cost error mitigation by symmetry verification*, Physical Review A **98**, 6, 062339 (2018)
2. C. Dickel, J.J. Wesdorp, N.K. Langford, S. Peiter, **R. Sagastizabal**, A. Bruno, B. Criger, F. Motzoi, L. DiCarlo, *Chip-to-chip entanglement of transmon qubits using engineered measurement fields*, Physical Review B **97**, 6, 064508 (2018)
1. N.K. Langford, **R. Sagastizabal**, M. Kounalakis, C. Dickel, A. Bruno, F. Luthi, D.J. Thoen, A. Endo, L. DiCarlo, *Experimentally simulating the dynamics of quantum light and matter at deep-strong coupling*, Nature communications, **8**, 1 (2017)

.

.

REFERENCES

- [1] R. P. FEYNMAN. *Simulating physics with computers*. International Journal of Theoretical Physics, **21** (6-7), 467–488, 1982.
- [2] M. A. NIELSEN and I. L. CHUANG. Quantum computation and quantum information. Cambridge University Press, Cambridge, 2000.
- [3] A. EINSTEIN, B. PODOLSKY, and N. ROSEN. *Can quantum-mechanical description of physical reality be considered complete?* Phys. Rev., **47**, 777–780, 1935.
- [4] P. W. SHOR. Scheme for reducing decoherence in quantum computer memory. Phys. Rev. A, **52**, R2493, 1995.
- [5] S. JORDAN. Quantum algorithm zoo. <https://quantumalgorithmzoo.org>.
- [6] D. P. DIVINCENZO. The physical implementation of quantum computation. Fortschritte der Physik, **48**, 771, 2000.
- [7] Y. NAKAMURA, Y. PASHKIN, and J. TSAI. *Coherent control of macroscopic quantum states in a single-Cooper-pair box*. Nature, **398**, 786, 1999.
- [8] B. D. JOSEPHSON. *Coupled superconductors*. Rev. Mod. Phys., **36**, 216–220, 1964.
- [9] B. D. JOSEPHSON. *Possible new effects in superconductive tunnelling*. Phys. Lett., **1** (7), 251–253, 1962.
- [10] M. DEVORET, A. WALLRAFF, and J. MARTINIS. *Superconducting qubits: A short review*. arXiv:cond-mat/0411174.
- [11] A. BLAIS, J. GAMBETTA, A. WALLRAFF, D. I. SCHUSTER, S. M. GIRVIN, M. H. DEVORET, and R. J. SCHOELKOPF. Quantum-information processing with circuit quantum electrodynamics. Phys. Rev. A, **75**, 032329, 2007.
- [12] A. WALLRAFF, D. I. SCHUSTER, A. BLAIS, L. FRUNZIO, R.-S. HUANG, J. MAJER, S. KUMAR, S. M. GIRVIN, and R. J. SCHOELKOPF. *Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics*. Nature, **431**, 162–167, 2004.
- [13] E. M. PURCELL, H. C. TORREY, and R. V. POUND. *Resonance absorption by nuclear magnetic moments in a solid*. Phys. Rev., **69**, 37–38, 1946.
- [14] A. A. HOUCK, J. A. SCHREIER, B. R. JOHNSON, J. M. CHOW, J. KOCH, J. M. GAMBETTA, D. I. SCHUSTER, L. FRUNZIO, M. H. DEVORET, S. M. GIRVIN, and R. J. SCHOELKOPF. Controlling the spontaneous emission of a superconducting transmon qubit. Phys. Rev. Lett., **101**, 080502 (pages 4), 2008.
- [15] D. RISTÈ, C. C. BULTINK, M. J. TIGGELMAN, R. N. SCHOUTEN, K. W. LEHNERT, and L. DICARLO. *Millisecond charge-parity fluctuations and induced decoherence in a superconducting transmon qubit*. Nat. Commun., **4**, 1913, 2013.
- [16] A. P. SEARS, A. PETRENKO, G. CATELANI, L. SUN, H. PAIK, G. KIRCHMAIR, L. FRUNZIO, L. I. GLAZMAN, S. M. GIRVIN, and R. J. SCHOELKOPF. *Photon shot noise dephasing in the strong-dispersive limit of circuit QED*. Phys. Rev. B, **86**, 180504, 2012.
- [17] J. BYLANDER, S. GUSTAVSSON, F. YAN, F. YOSHIHARA, K. HARRABI, G. FITCH, D. G. CORY, Y. NAKAMURA, J.-S. TSAI, and W. D. OLIVER. *Noise spectroscopy through dynamical decoupling with a superconducting flux qubit*. Nat. Phys., **7**, 565, 2011.

- [18] J. KOCH, T. M. YU, J. GAMBETTA, A. A. HOUCK, D. I. SCHUSTER, J. MAJER, A. BLAIS, M. H. DEVORET, S. M. GIRVIN, and R. J. SCHOELKOPF. *Charge-insensitive qubit design derived from the Cooper pair box*. Phys. Rev. A, **76**, 042319, 2007.
- [19] D. RISTÈ, J. G. VAN LEEUWEN, H.-S. KU, K. W. LEHNERT, and L. DICARLO. *Initialization by measurement of a superconducting quantum bit circuit*. Phys. Rev. Lett., **109**, 050507, 2012.
- [20] M. D. REED, L. DICARLO, B. R. JOHNSON, L. SUN, D. I. SCHUSTER, L. FRUNZIO, and R. J. SCHOELKOPF. *High-fidelity readout in circuit quantum electrodynamics using the Jaynes-Cummings nonlinearity*. Phys. Rev. Lett., **105**, 173601, 2010.
- [21] J. M. CHOW, L. DICARLO, J. M. GAMBETTA, F. MOTZOI, L. FRUNZIO, S. M. GIRVIN, and R. J. SCHOELKOPF. *Optimized driving of superconducting artificial atoms for improved single-qubit gates*. Phys. Rev. A, **82**, 040305, 2010.
- [22] L. DICARLO, J. M. CHOW, J. M. GAMBETTA, L. S. BISHOP, B. R. JOHNSON, D. I. SCHUSTER, J. MAJER, A. BLAIS, L. FRUNZIO, S. M. GIRVIN, and R. J. SCHOELKOPF. *Demonstration of two-qubit algorithms with a superconducting quantum processor*. Nature, **460**, 240, 2009.
- [23] A. P. PLACE, L. V. RODGERS, P. MUNDADA, B. M. SMITHAM, M. FITZPATRICK, Z. LENG, A. PREMKUMAR, J. BRYON, S. SUSSMAN, G. CHENG, *et al.* *New material platform for superconducting transmon qubits with coherence times exceeding 0.3 milliseconds*. arXiv preprint arXiv:2003.00024, 2020.
- [24] D. RISTÈ, S. POLETTI, M. Z. HUANG, A. BRUNO, V. VESTERINEN, O. P. SAIRA, and L. DICARLO. *Detecting bit-flip errors in a logical qubit using stabilizer measurements*. Nat. Commun., **6**, 6983, 2015.
- [25] J. KELLY, R. BARENDT, A. G. FOWLER, A. MEGRANT, E. JEFFREY, T. WHITE, D. SANK, J. MUTUS, B. CAMPBELL, Y. CHEN, B. CHIARO, A. DUNSWORTH, I.-C. HOI, *et al.* *State preservation by repetitive error detection in a superconducting quantum circuit*. Nature, **519** (7541), 66–69, 2015.
- [26] C. C. BULTINK, T. E. O'BRIEN, R. VOLLMER, N. MUTHUSUBRAMANIAN, M. W. BEEKMAN, M. A. ROL, X. FU, B. TARASINSKI, V. OSTROUKH, B. VARBANOV, A. BRUNO, and L. DICARLO. *Protecting quantum entanglement from leakage and qubit errors via repetitive parity measurements*. Science Advances, **6** (12), 2020.
- [27] C. K. ANDERSEN, A. REMM, S. LAZAR, S. KRINNER, N. LACROIX, G. J. NORRIS, M. GABUREAC, C. EICHLER, and A. WALLRAFF. *Repeated quantum error detection in a surface code*. Nat. Phys., **16** (8), 875–880, 2020.
- [28] R. BARENDT, L. LAMATA, J. KELLY, L. GARCÍA-ÁLVAREZ, A. FOWLER, A. MEGRANT, E. JEFFREY, T. WHITE, D. SANK, J. MUTUS, B. CAMPBELL, Y. CHEN, Z. CHEN, *et al.* *Digital quantum simulation of fermionic models with a superconducting circuit*. Nat. Commun., **6**, 7654, 2015.
- [29] J. ZHANG, M.-H. YUNG, R. LAFLAMME, A. ASPURU-GUZI, and J. BAUGH. *Digital quantum simulation of the statistical mechanics of a frustrated magnet*. Nat. Commun., **3**, 880, 2012.
- [30] P. J. J. O'MALLEY, R. BABBUSH, I. D. KIVLICHAN, J. ROMERO, J. R. MCCLEAN, R. BARENDT, J. KELLY, P. ROUSHAN, A. TRANTER, N. DING, B. CAMPBELL, Y. CHEN, Z. CHEN, *et al.* *Scalable quantum simulation of molecular energies*. Phys. Rev. X, **6**, 031007, 2016.
- [31] A. KANDALA, A. MEZZACAPO, K. TEMME, M. TAKITA, M. BRINK, J. M. CHOW, and J. M. GAMBETTA. *Hardware-efficient variational quantum eigensolver for small molecules and quantum magnets*. Nature, **549**, 242–246, 2017.
- [32] R. BARENDT, J. KELLY, A. MEGRANT, D. SANK, E. JEFFREY, Y. CHEN, Y. YIN, B. CHIARO, J. MUTUS, C. NEILL, P. O'MALLEY, P. ROUSHAN, J. WENNER, *et al.* *Coherent Josephson qubit suitable for scalable quantum integrated circuits*. Phys. Rev. Lett., **111**, 080502, 2013.
- [33] A. WALLRAFF, D. I. SCHUSTER, A. BLAIS, L. FRUNZIO, J. MAJER, M. H. DEVORET, S. M. GIRVIN, and R. J. SCHOELKOPF. *Approaching unit visibility for control of a superconducting qubit with dispersive readout*. Phys. Rev. Lett., **95**, 060501, 2005.

- [34] D. M. POZAR. Microwave engineering. John Wiley & Sons, Hoboken, 3 edition, 2005.
- [35] R. BARENDs. Photon-detecting superconducting resonators. PhD Dissertation, Delft University of Technology, 2009.
- [36] I. RABI. *On the process of space quantization*. Phys. Rev., **49** (4), 324, 1936.
- [37] D. I. SCHUSTER. Circuit quantum electrodynamics. PhD Dissertation, Yale University, 2007.
- [38] D. I. SCHUSTER, A. WALLRAFF, A. BLAIS, L. FRUNZIO, R.-S. HUANG, J. MAJER, S. M. GIRVIN, and R. J. SCHOELKOPF. *ac Stark shift and dephasing of a superconducting qubit strongly coupled to a cavity field*. Phys. Rev. Lett., **94**, 123 602, 2005.
- [39] D. SANK, Z. CHEN, M. KHEZRI, J. KELLY, R. BARENDs, B. CAMPBELL, Y. CHEN, B. CHIARO, A. DUNSWORTH, A. FOWLER, E. JEFFREY, E. LUCERO, A. MEGRANT, *et al.* *Measurement-induced state transitions in a superconducting qubit: Beyond the rotating wave approximation*. Phys. Rev. Lett., **117**, 190 503, 2016.
- [40] J. GAMBETTA, A. BLAIS, M. BOISSONNEAULT, A. A. HOUCK, D. I. SCHUSTER, and S. M. GIRVIN. *Quantum trajectory approach to circuit QED: Quantum jumps and the Zeno effect*. Phys. Rev. A, **77** (1), 012 112, 2008.
- [41] F. MOTZOI, J. M. GAMBETTA, P. REBENTROST, and F. K. WILHELM. *Simple pulses for elimination of leakage in weakly nonlinear qubits*. Phys. Rev. Lett., **103**, 110 501, 2009.
- [42] M. KOUNALAKIS, C. DICKEL, A. BRUNO, N. K. LANGFORD, and G. A. STEELE. *Tuneable hopping and nonlinear cross-Kerr interactions in a high-coherence superconducting circuit*. npj Quantum Information, **4** (1), 1–7, 2018.
- [43] F. ARUTE, K. ARYA, R. BABBUSH, D. BACON, J. C. BARDIN, R. BARENDs, S. BOIXO, M. BROUGHTON, B. B. BUCKLEY, D. A. BUELL, B. BURKETT, N. BUSHNELL, Y. CHEN, *et al.* *Quantum approximate optimization of non-planar graph problems on a planar superconducting processor*. ArXiv:2004.04197, 2020.
- [44] M. TAVIS and F. W. CUMMINGS. *Exact solution for an n-molecule-radiation-field hamiltonian*. Phys. Rev., **170**, 379–384, 1968.
- [45] J. M. CHOW. Quantum information processing with superconducting qubits. PhD Dissertation, Yale University, 2010.
- [46] R. VERSLUIS, S. POLETTI, N. KHAMMASSI, B. TARASINSKI, N. HAIDER, D. J. MICHALAK, A. BRUNO, K. BERTELS, and L. DICARLO. *Scalable quantum circuit and control for a superconducting surface code*. Phys. Rev. Applied, **8**, 034 021, 2017.
- [47] J. M. MARTINIS and M. R. GELLER. *Fast adiabatic qubit gates using only σ_x control*. Phys. Rev. A, **90**, 022 307, 2014.
- [48] M. A. ROL, F. BATTISTEL, F. K. MALINOWSKI, C. C. BULTINK, B. M. TARASINSKI, R. VOLLMER, N. HAIDER, N. MUTHUSUBRAMANIAN, A. BRUNO, B. M. TERHAL, and L. DICARLO. *Fast, high-fidelity conditional-phase gate exploiting leakage interference in weakly anharmonic superconducting qubits*. Phys. Rev. Lett., **123**, 120 502, 2019.
- [49] V. NEGÎRNEAC, H. ALI, N. MUTHUSUBRAMANIAN, F. BATTISTEL, R. SAGASTIZABAL, M. S. MOREIRA, J. F. MARQUES, W. J. VLOTHUIZEN, M. BEEKMAN, C. ZACHARIADIS, N. HAIDER, A. BRUNO, and L. DICARLO. *High-fidelity controlled-z gate with maximal intermediate leakage operating at the speed limit in a superconducting quantum processor*. Phys. Rev. Lett., **126**, 220 502, 2021.
- [50] R. BABBUSH, N. WIEBE, J. MCCLEAN, J. MCCLAIN, H. NEVEN, and G. K.-L. CHAN. *Low-depth quantum simulation of materials*. Phys. Rev. X, **8**, 011 044, 2018.

- [51] A. KITAEV, A. SHEN, and M. VYALYI. Classical and quantum computation. American Mathematical Society, 2002.
- [52] J. R. MCCLEAN, I. D. KIVLICHAN, K. J. SUNG, D. S. STEIGER, Y. CAO, C. DAI, E. SCHUYLER FRIED, C. GIDNEY, B. GIMBY, P. GOKHALE, T. HÄNER, T. HARDIKAR, V. HAVLIČEK, *et al.* *OpenFermion: The Electronic Structure Package for Quantum Computers*. ArXiv:1710.07629.
- [53] Y. LI and S. C. BENJAMIN. *Efficient variational quantum simulator incorporating active error minimization*. Phys. Rev. X, **7**, 021050, 2017.
- [54] A. PERUZZO, J. R. MCCLEAN, P. SHADBOLT, M. YUNG, X. ZHOU, P. J. LOVE, A. ASPURU-GUZI, and J. L. O'BRIEN. *A variational eigenvalue solver on a photonic quantum processor*. Nat. Commun., **5**, 4213, 2014.
- [55] C. HEMPEL, C. MAIER, J. ROMERO, J. MCCLEAN, T. MONZ, H. SHEN, P. JURCEVIC, B. P. LANYON, P. LOVE, R. BABBUSH, A. ASPURU-GUZI, R. BLATT, and C. F. ROOS. *Quantum chemistry calculations on a trapped-ion quantum simulator*. Phys. Rev. X, **8**, 031022, 2018.
- [56] M. GANZHORN, D. EGGER, P. BARKOUTSOS, P. OLLITRAULT, G. SALIS, N. MOLL, M. ROTH, A. FUHRER, P. MUELLER, S. WOERNER, I. TAVERNELLI, and S. FILIPP. *Gate-efficient simulation of molecular eigenstates on a quantum computer*. Phys. Rev. App., **11**, 044092, 2019.
- [57] X. BONET-MONROIG, R. SAGASTIZABAL, M. SINGH, and T. E. O'BRIEN. *Low-cost error mitigation by symmetry verification*. Phys. Rev. A, **98**, 062339, 2018.
- [58] R. SAGASTIZABAL, X. BONET-MONROIG, M. SINGH, M. A. ROL, C. C. BULTINK, X. FU, C. H. PRICE, V. P. OSTROUKH, N. MUTHUSUBRAMANIAN, A. BRUNO, M. BEEKMAN, N. HAIDER, T. E. O'BRIEN, *et al.* *Error Mitigation by Symmetry Verification on a Variational Quantum Eigensolver*. Phys. Rev. A, **100**, 010302(R), 2019.
- [59] J. I. COLLESS, V. V. RAMASESH, D. DAHLEN, M. S. BLOK, M. E. KIMCHI-SCHWARTZ, J. R. MCCLEAN, J. CARTER, W. A. DE JONG, and I. SIDDIQI. *Computation of molecular spectra on a quantum processor with an error-resilient algorithm*. Phys. Rev. X, **8**, 011021, 2018.
- [60] E. FARHI, J. GOLDSTONE, and S. GUTMANN. *A quantum approximate optimization algorithm*, 2014.
- [61] E. FARHI, J. GOLDSTONE, S. GUTMANN, and M. SIPSER. *Quantum computation by adiabatic evolution*, 2000.
- [62] J. C. BUDICH and B. TRAUZETTEL. *From the adiabatic theorem of quantum mechanics to topological states of matter*. physica status solidi (RRL) – Rapid Research Letters, **7** (1-2), 109–129, 2013. <https://onlinelibrary.wiley.com/doi/pdf/10.1002/pssr.201206416>.
- [63] Z.-C. YANG, A. RAHMANI, A. SHABANI, H. NEVEN, and C. CHAMON. *Optimizing variational quantum algorithms using pontryagin's minimum principle*. Phys. Rev. X, **7**, 021027, 2017.
- [64] L. PONTRYAGIN. *Mathematical theory of optimal processes*. Routledge, 2017.
- [65] A. BENGTSSON, P. VIKSTÅL, C. WARREN, M. SVENSSON, X. GU, A. F. KOCKUM, P. KRANTZ, C. KRIŽAN, D. SHIRI, I.-M. SVENSSON, G. TANCREDI, G. JOHANSSON, P. DELSING, *et al.* *Improved success probability with greater circuit depth for the quantum approximate optimization algorithm*. Phys. Rev. Appl., **14**, 034010, 2020.
- [66] D. ZHU, S. JOHRI, N. M. LINKE, K. A. LANDSMAN, C. H. ALDERETE, N. H. NGUYEN, A. Y. MATSUURA, T. H. HSIEH, and C. MONROE. *Generation of thermofield double states and critical ground states with a quantum computer*. Proceedings of the National Academy of Sciences, **117** (41), 25402–25406, 2020. <https://www.pnas.org/doi/pdf/10.1073/pnas.2006337117>.
- [67] N. K. LANGFORD, R. SAGASTIZABAL, M. KOUNALAKIS, C. DICKEL, A. BRUNO, F. LUTHI, D. THOEN, A. ENDO, and L. DICARLO. *Experimentally simulating the dynamics of quantum light and matter at deep-strong coupling*. Nat. Commun., **8**, 1715, 2017.

- [68] S. LLOYD. *Universal quantum simulators*. Science, **273** (5278), 1073, 1996.
- [69] I. KASSAL, S. P. JORDAN, P. J. LOVE, M. MOHSENI, and A. ASPURU-GUZIK. *Polynomial-time quantum algorithm for the simulation of chemical dynamics*. P. Natl Acad. Sci. USA, **105** (48), 18 681–18 686, 2008.
- [70] D. S. ABRAMS and S. LLOYD. *Simulation of many-body fermi systems on a universal quantum computer*. Phys. Rev. Lett., **79** (13), 2586, 1997.
- [71] B. P. LANYON, C. HEMPEL, D. NIGG, M. MÜLLER, R. GERRITSMAN, F. ZÄHRINGER, P. SCHINDLER, J. T. BARREIRO, M. RAMBACH, G. KIRCHMAIR, M. HENNRICH, P. ZOLLER, R. BLATT, *et al.* *Universal digital quantum simulation with trapped ions*. Science, **334** (6052), 57–61, 2011.
- [72] Y. SALATHÉ, M. MONDAL, M. OPPLIGER, J. HEINSOO, P. KURPIERS, A. POTOČNIK, A. MEZZA-CAPO, U. LAS HERAS, L. LAMATA, E. SOLANO, S. FILIPP, and A. WALLRAFF. *Digital quantum simulation of spin models with circuit quantum electrodynamics*. Phys. Rev. X, **5** (2), 021 027, 2015.
- [73] C. CIUTI, G. BASTARD, and I. CARUSOTTO. *Quantum vacuum properties of the intersubband cavity polariton field*. Phys. Rev. B, **72** (11), 115 303, 2005.
- [74] D. BRAAK. *Integrability of the rabi model*. Phys. Rev. Lett., **107** (10), 100 401, 2011.
- [75] F. BEAUDOIN, J. M. GAMBETTA, and A. BLAIS. *Dissipation and ultrastrong coupling in circuit qed*. Phys. Rev. A, **84** (4), 043 832, 2011.
- [76] P. NATAF and C. CIUTI. *No-go theorem for superradiant quantum phase transitions in cavity qed and counter-example in circuit qed*. Nat. Commun., **1**, 72, 2010.
- [77] O. VIEHMANN, J. VON DELFT, and F. MARQUARDT. *Superradiant phase transitions and the standard description of circuit qed*. Phys. Rev. Lett., **107** (11), 113 602, 2011.
- [78] J. LOLLI, A. BAKSIC, D. NAGY, V. E. MANUCHARYAN, and C. CIUTI. *Ancillary qubit spectroscopy of vacua in cavity and circuit quantum electrodynamics*. Phys. Rev. Lett., **114** (18), 183 601, 2015.
- [79] T. NIEMCZYK, F. DEPPE, H. HUEBL, E. MENZEL, F. HOCKE, M. SCHWARZ, J. GARCIA-RIPOLL, D. ZUECO, T. HÜMMER, E. SOLANO, A. MARX, and R. GROSS. *Circuit quantum electrodynamics in the ultrastrong-coupling regime*. Nat. Phys., **6** (10), 772–776, 2010.
- [80] P. FORN-DÍAZ, J. GARCÍA-RIPOLL, B. PEROPADRE, J.-L. ORGIAZZI, M. YURTALAN, R. BELYANSKY, C. WILSON, and A. LUPASCU. *Ultrastrong coupling of a single artificial atom to an electromagnetic continuum in the nonperturbative regime*. Nat. Phys., **13**, 44–47, 2017.
- [81] F. YOSHIHARA, T. FUSE, S. ASHHAB, K. KAKUYANAGI, S. SAITO, and K. SEMBA. *Superconducting qubit–oscillator circuit beyond the ultrastrong-coupling regime*. Nat. Phys., **13**, 44–47, 2017.
- [82] G. GÜNTHER, A. A. ANAPPARA, J. HEES, A. SELL, G. BIASIOL, L. SORBA, S. DE LIBERATO, C. CIUTI, A. TREDICUCCI, A. LEITENSTORFER, *et al.* *Sub-cycle switch-on of ultrastrong light–matter interaction*. Nature, **458** (7235), 178–181, 2009.
- [83] G. SCALARI, C. MAISSEN, D. TURČINKOVÁ, D. HAGENMÜLLER, S. DE LIBERATO, C. CIUTI, C. REICHL, D. SCHUH, W. WEGSCHEIDER, M. BECK, *et al.* *Ultrastrong coupling of the cyclotron transition of a 2d electron gas to a thz metamaterial*. Science, **335** (6074), 1323–1326, 2012.
- [84] C. MAISSEN, G. SCALARI, F. VALMORRA, M. BECK, J. FAIST, S. CIBELLA, R. LEONI, C. REICHL, C. CHARPENTIER, and W. WEGSCHEIDER. *Ultrastrong coupling in the near field of complementary split-ring resonators*. Phys. Rev. B, **90** (20), 205 309, 2014.
- [85] J. ZHANG, P. W. HESS, A. KYPRIANIDIS, P. BECKER, A. LEE, J. SMITH, G. PAGANO, I.-D. POTIRNICHE, A. C. POTTER, A. VISHWANATH, N. Y. YAO, and C. MONROE. *Observation of a Discrete Time Crystal*. arXiv:1609.08684, 2016.
- [86] T. SCHWARTZ, J. A. HUTCHISON, C. GENET, and T. W. EBBESEN. *Reversible switching of ultrastrong light-molecule coupling*. Phys. Rev. Lett., **106** (19), 196 405, 2011.

- [87] J. CASANOVA, G. ROMERO, I. LIZUAIN, J. J. GARCÍA-RIPOLL, and E. SOLANO. *Deep strong coupling regime of the jaynes-cummings model*. Phys. Rev. Lett., **105** (26), 263 603, 2010.
- [88] A. MEZZACAPO, U. LAS HERAS, J. PEDERNALES, L. DICARLO, E. SOLANO, and L. LAMATA. *Digital quantum rabi and dicke models in superconducting circuits*. Sci. Rep., **4**, 7482, 2014.
- [89] D. BALLESTER, G. ROMERO, J. J. GARCÍA-RIPOLL, F. DEPPE, and E. SOLANO. *Quantum simulation of the ultrastrong-coupling dynamics in circuit quantum electrodynamics*. Phys. Rev. X, **2** (2), 021 007, 2012.
- [90] L. LAMATA. *Digital-analog quantum simulation of generalized dicke models with superconducting circuits*. Sci. Rep., **7**, 43 768, 2017.
- [91] S. FELICETTI, G. ROMERO, E. SOLANO, and C. SABÍN. *Quantum rabi model in a superfluid bose-einstein condensate*. Phys. Rev. A, **96**, 033 839, 2017.
- [92] J. PEDERNALES, I. LIZUAIN, S. FELICETTI, G. ROMERO, L. LAMATA, and E. SOLANO. *Quantum rabi model with trapped ions*. Sci. Rep., **5**, 15 472, 2015.
- [93] R. GERRITSMAN, G. KIRCHMAIR, F. ZÄHRINGER, E. SOLANO, R. BLATT, and C. ROOS. *Quantum simulation of the dirac equation*. Nature, **463** (7277), 68–71, 2010.
- [94] R. GERRITSMAN, B. LANYON, G. KIRCHMAIR, F. ZÄHRINGER, C. HEMPEL, J. CASANOVA, J. J. GARCÍA-RIPOLL, E. SOLANO, R. BLATT, and C. F. ROOS. *Quantum simulation of the klein paradox with trapped ions*. Phys. Rev. Lett., **106** (6), 060 503, 2011.
- [95] H.-Y. LO, D. KIENZLER, L. DE CLERCQ, M. MARINELLI, V. NEGNEVITSKY, B. C. KEITCH, and J. P. HOME. *Spin-motion entanglement and state diagnosis with squeezed oscillator wavepackets*. Nature, **521** (7552), 336–339, 2015.
- [96] D. KIENZLER, C. FLÜHMANN, V. NEGNEVITSKY, H.-Y. LO, M. MARINELLI, D. NADLINGER, and J. HOME. *Observation of quantum interference between separated mechanical oscillator wave packets*. Phys. Rev. Lett., **116** (14), 140 402, 2016.
- [97] A. CRESPI, S. LONGHI, and R. OSELLAME. *Photonic realization of the quantum rabi model*. Phys. Rev. Lett., **108** (16), 163 601, 2012.
- [98] B. VLASTAKIS, A. PETRENKO, N. OFEK, L. SUN, Z. LEGHTAS, K. SLIWA, Y. LIU, M. HATRIDGE, J. BLUMOFF, L. FRUNZIO, *et al.* *Characterizing entanglement of an artificial atom and a cavity cat state with bell's inequality*. Nat. Commun., **6**, 8970, 2015.
- [99] T. JAAKO, Z.-L. XIANG, J. J. GARCIA-RIPOLL, and P. RABL. *Ultrastrong-coupling phenomena beyond the dicke model*. Phys. Rev. A, **94** (3), 033 850, 2016.
- [100] P. BERTET, A. AUFEVES, P. MAIOLI, S. OSNAGHI, T. MEUNIER, M. BRUNE, J.-M. RAIMOND, and S. HAROCHE. *Direct measurement of the wigner function of a one-photon fock state in a cavity*. Phys. Rev. Lett., **89** (20), 200 402, 2002.
- [101] B. VLASTAKIS, G. KIRCHMAIR, Z. LEGHTAS, S. E. NIGG, L. FRUNZIO, S. M. GIRVIN, M. MIRRAHIMI, M. H. DEVORET, and R. J. SCHOELKOPF. *Deterministically encoding quantum information using 100-photon Schrödinger cat states*. Science, **342** (6158), 607–610, 2013.
- [102] J. BOURASSA, F. BEAUDOIN, J. M. GAMBETTA, and A. BLAIS. *Josephson-junction-embedded transmission-line resonators: From kerr medium to in-line transmon*. Phys. Rev. A, **86**, 013 814, 2012.
- [103] V. V. ALBERT, G. D. SCHOLES, and P. BRUMER. *Symmetric rotating-wave approximation for the generalized single-mode spin-boson system*. Phys. Rev. A, **84** (4), 042 110, 2011.
- [104] J. I. CIRAC and P. ZOLLER. *Goals and opportunities in quantum simulation*. Nature Physics, **8** (4), 264–266, 2012.
- [105] C. K. ANDERSEN and A. BLAIS. *Ultrastrong coupling dynamics with a transmon qubit*. New Journal of Physics, **19** (2), 023 022, 2017.

- [106] W. ZUREK. *Cosmological experiments in superfluid helium?* Nature, **317** (6037), 505–508, 1985.
- [107] A. BRUNO, G. DE LANGE, S. ASAAD, K. L. VAN DER ENDEN, N. K. LANGFORD, and L. DICARLO. *Reducing intrinsic loss in superconducting resonators by surface treatment and deep etching of silicon substrates.* App. Phys. Lett., **106**, 182601, 2015.
- [108] R. BARENDTS, A. SHABANI, L. LAMATA, J. KELLY, A. MEZZACAPO, U. LAS HERAS, R. BABBUSH, A. FOWLER, B. CAMPBELL, Y. CHEN, *et al.* *Digitized adiabatic quantum computing with a superconducting circuit.* Nature, **534** (7606), 222–226, 2016.
- [109] M.-J. HWANG, R. PUEBLA, and M. B. PLENIO. *Quantum phase transition and universal dynamics in the rabi model.* Phys. Rev. Lett., **115** (18), 180404, 2015.
- [110] R. PUEBLA, M.-J. HWANG, and M. B. PLENIO. *Excited-state quantum phase transition in the rabi model.* Phys. Rev. A, **94** (2), 023835, 2016.
- [111] S. DE LIBERATO. *Light-matter decoupling in the deep strong coupling regime: The breakdown of the purcell effect.* Phys. Rev. Lett., **112** (1), 016401, 2014.
- [112] M. STEFFEN, L. M. K. VANDERSYPEN, and I. L. CHUANG. *Simultaneous soft pulses applied at nearby frequencies.* J. Magn. Reson., **146**, 369, 2000.
- [113] M. REED. *Entanglement and quantum error correction with superconducting qubits.* PhD Dissertation, Yale University, 2013.
- [114] Z. HRADIL. *Quantum-state estimation.* Phys. Rev. A, **44** (11-12), R1561, 1997.
- [115] D. F. V. JAMES, P. G. KWIAT, W. J. MUNRO, and A. G. WHITE. *Measurement of qubits.* Phys. Rev. A, **64** (5), 052312, 2001.
- [116] S. BOYD and L. VANDENBERGHE. *Convex optimization.* Cambridge University Press, 2004.
- [117] N. K. LANGFORD. *Errors in quantum tomography: diagnosing systematic versus statistical errors.* New J. Phys., **15** (3), 035003, 2013.
- [118] R. BARENDTS, J. WENNER, M. LENANDER, Y. CHEN, R. C. BIALCZAK, J. KELLY, E. LUCERO, P. O'MALLEY, M. MARIANTONI, D. SANK, H. WANG, T. C. WHITE, Y. YIN, *et al.* *Minimizing quasi-particle generation from stray infrared light in superconducting quantum circuits.* App. Phys. Lett., **99**, 113507, 2011.
- [119] S. ASAAD, C. DICKEL, S. POLETO, A. BRUNO, N. K. LANGFORD, M. A. ROL, D. DEURLOO, and L. DICARLO. *Independent, extensible control of same-frequency superconducting qubits by selective broadcasting.* npj Quantum Inf., **2**, 16029, 2016.
- [120] J. GAMBETTA, C. MURRAY, Y.-K.-K. FUNG, D. MCCLURE, O. DIAL, W. SHANKS, J. SLEIGHT, and M. STEFFEN. *Investigating surface loss effects in superconducting transmon qubits.* arXiv preprint arXiv:1605.08009, 2016.
- [121] B. R. JOHNSON. *Controlling photons in superconducting electrical circuits.* PhD Dissertation, Yale University, 2011.
- [122] R. WIGINGTON and N. NAHMAN. *Transient analysis of coaxial cables considering skin effect.* Proc. IRE, **45** (2), 166–174, 1957.
- [123] J. PRESKILL. *Quantum Computing in the NISQ era and beyond.* Quantum, **2**, 79, 2018.
- [124] S. MCARDLE, A. MAYOROV, X. SHAN, S. BENJAMIN, and X. YUAN. *Quantum computation of molecular vibrations.* arXiv:1811.04069, 2018.
- [125] Y. CAO, J. ROMERO, J. OLSON, M. DEGROOTE, P. JOHNSON, M. KIEFEROVÁ, I. KIVLICHAN, T. MENKE, B. PEROPADRE, N. SAWAYA, S. SIM, L. VIES, and A. ASPURU-GUZI. *Quantum chemistry in the age of quantum computing.* ArXiv:1812.09976, 2018.

- [126] J. R. McCLEAN, J. ROMERO, R. BABBUSH, and A. ASPURU-GUZI. *The theory of variational hybrid quantum-classical algorithms*. New Journal of Physics, **18** (2), 023023, 2016.
- [127] D. POULIN, A. KITAEV, D. S. STEIGER, M. B. HASTINGS, and M. TROYER. *Quantum algorithm for spectral measurement with a lower gate count*. Phys. Rev. Lett., **121**, 010501, 2018.
- [128] D. W. BERRY, M. KIEFEROVÁ, A. SCHERER, Y. R. SANDERS, G. H. LOW, N. WIEBE, C. GIDNEY, and R. BABBUSH. *Improved techniques for preparing eigenstates of fermionic hamiltonians*. npj Quant. Inf., **4**, 22, 2018.
- [129] I. D. KIVLICHAN, J. McCLEAN, N. WIEBE, C. GIDNEY, A. ASPURU-GUZI, G. K.-L. CHAN, and R. BABBUSH. *Quantum simulation of electronic structure with linear depth and connectivity*. Phys. Rev. Lett., **120**, 110501, 2018.
- [130] A. KANDALA, K. TEMME, A. D. CORCOLES, A. MEZZACAPO, J. M. CHOW, and J. M. GAMBETTA. *Error mitigation extends the computational reach of a noisy quantum processor*. Nature, **567**, 491–495, 2019.
- [131] Y. SHEN, X. ZHANG, S. ZHANG, J.-N. ZHANG, M.-H. YUNG, and K. KIM. *Quantum implementation of the unitary coupled cluster for simulating molecular electronic structure*. Phys. Rev. A, **95** (020501), 2017.
- [132] R. SANTAGATI, J. WANG, A. GENTILE, S. PAESANI, N. WIEBE, J. McCLEAN, S. MORLEY-SHORT, P. SHADBOLT, D. BONNEAU, J. SILVERSTONE, D. TEW, X. ZHOU, J. O'BRIEN, *et al.* *Witnessing eigenstates for quantum simulation of hamiltonian spectra*. Sci. Adv., **4** (1), eaap9646, 2018.
- [133] C. KOKAIL, C. MAIER, R. VAN BIJNEN, T. BRYDGES, M. JOSHI, P. JURCEVIC, C. MUSCHIK, P. SILVI, R. BLATT, C. ROOS, and P. ZOLLER. *Self-verifying variational quantum simulation of lattice models*. Nature, **569**, 355–360, 2019.
- [134] S. McARDLE, X. YUAN, and S. BENJAMIN. *Error-mitigated digital quantum simulation*. Phys. Rev. Lett., **122**, 180501, 2019.
- [135] J. R. McCLEAN, M. E. KIMCHI-SCHWARTZ, J. CARTER, and W. A. DE JONG. *Hybrid quantum-classical hierarchy for mitigation of decoherence and determination of excited states*. Phys. Rev. A, **95**, 042308, 2017.
- [136] K. TEMME, S. BRAVYI, and J. M. GAMBETTA. *Error mitigation for short-depth quantum circuits*. Phys. Rev. Lett., **119**, 180509, 2017.
- [137] M. OTTEN and S. GRAY. *Recovering noise-free quantum observables*. Phys. Rev. A, **99** (012338), 2018.
- [138] S. ENDO, S. BENJAMIN, and Y. LI. *Practical quantum error mitigation for near-future applications*. Phys. Rev. X, **8** (031027), 031027, 2018.
- [139] M. HUO and Y. LI. *Temporally correlated error tomography and mitigation in quantum computer*. ArXiv:1811.02734, 2018.
- [140] S. BRAVYI and A. KITAEV. *Fermionic quantum computation*. ArXiv:quant-ph/0003137, 2002.
- [141] R. M. PARRISH, L. A. BURNS, D. G. A. SMITH, A. C. SIMMONETT, A. E. DEPRINCE, E. G. HOHENSTEIN, U. BOZKAYA, A. Y. SOKOLOV, R. DI REMIGIO, R. M. RICHARD, J. F. GONTHIER, A. M. JAMES, H. R. McALEXANDER, *et al.* Psi4 1.1: An open-source electronic structure program emphasizing automation, advanced libraries, and interoperability. Journal of Chemical Theory and Computation, **13**, 3185, 2017.
- [142] J. McCLEAN, S. BOIXO, V. SMELYANSKIY, R. BABBUSH, and H. NEVEN. *Barren plateaus in quantum neural network traininglandscapes*. Nat. Commun., **9**, 4812, 2018.
- [143] J. MAJER, J. M. CHOW, J. M. GAMBETTA, B. R. JOHNSON, J. A. SCHREIER, L. FRUNZIO, D. I. SCHUSTER, A. A. HOUCK, A. WALLRAFF, A. BLAIS, M. H. DEVORET, S. M. GIRVIN, and R. J. SCHOELKOPF. *Coupling superconducting qubits via a cavity bus*. Nature, **449**, 443, 2007.

- [144] D. J. BROD and A. M. CHILDS. The computational power of matchgates and the xy interaction on arbitrary graphs. *Quantum Info. Comput.*, **14** (11–12), 901–916, 2014.
- [145] N. HANSEN. *Benchmarking the nelder-mead downhill simplex algorithm with many local restarts*. In *Proceedings of the 11th Annual Conference Companion on Genetic and Evolutionary Computation Conference: Late Breaking Papers*, GECCO '09, pages 2403–2408. ACM, New York, NY, USA, 2009.
- [146] T. E. O'BRIEN, B. M. TARASINSKI, and L. DICARLO. *Density-matrix simulation of small surface codes under current and projected experimental noise*. *npj Quantum Information*, **3** (39), 2017.
- [147] See supplemental material at [insert URL] for additional data.
- [148] N. C. RUBIN, R. BABBUSH, and J. MCCLEAN. *Application of fermionic marginal constraints to hybrid quantum algorithms*. *New J. Phys.*, **20**, 053020, 2018.
- [149] R. BLUME-KOHOUT. *Optimal, reliable estimation of quantum states*. *New J. Phys.*, **12**, 043034, 2010.
- [150] Y.-K. LIU, M. CHRISTANDL, and F. VERSTRAETE. *Quantum computational complexity of then-representability problem: Qma complete*. *Phys. Rev. Lett.*, **98**, 110503, 2007.
- [151] C. MACKLIN, K. O'BRIEN, D. HOVER, M. E. SCHWARTZ, V. BOLKHOVSKY, X. ZHANG, W. D. OLIVER, and I. SIDDIQI. *A near-quantum-limited josephson traveling-wave parametric amplifier*. *Science*, **350** (6258), 307–310, 2015.
- [152] A. JOHNSON, G. UNGARETTI, *et al.* *QCoDeS*, 2016.
- [153] M. A. ROL, C. DICKEL, S. ASAAD, C. C. BULTINK, R. SAGASTIZABAL, N. K. L. LANGFORD, G. DE LANGE, B. C. S. DIKKEN, X. FU, S. R. DE JONG, and F. LUTHI. *PycQED*, 2016.
- [154] X. FU, L. RIESEBOS, M. A. ROL, J. VAN STRATEN, J. VAN SOMEREN, N. KHAMMASSI, I. ASHRAF, R. F. L. VERMEULEN, V. NEWSUM, K. K. L. LOH, J. C. DE STERKE, W. J. VLOTHUIZEN, R. N. SCHOUTEN, *et al.* *eQASM: An executable quantum instruction set architecture*. In *Proceedings of 25th IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pages 224–237. IEEE, 2019.
- [155] M. A. ROL, L. CIORCIARO, F. K. MALINOWSKI, B. M. TARASINSKI, R. E. SAGASTIZABAL, C. C. BULTINK, Y. SALATHE, N. HAANDBAEK, J. SEDIVY, and L. DICARLO. *Time-domain characterization and correction of on-chip distortion of control pulses in a quantum processor*. *App. Phys. Lett.*, **116** (5), 054001, 2020.
- [156] O.-P. SAIRA, J. P. GROEN, J. CRAMER, M. MERETSKA, G. DE LANGE, and L. DICARLO. *Entanglement genesis by ancilla-based parity measurement in 2D circuit QED*. *Phys. Rev. Lett.*, **112**, 070502, 2014.
- [157] C. DICKEL, J. J. WESDORP, N. K. LANGFORD, S. PEITER, R. SAGASTIZABAL, A. BRUNO, B. CRIGER, F. MOTZOI, and L. DICARLO. *Chip-to-chip entanglement of transmon qubits using engineered measurement fields*. *Phys. Rev. B*, **97**, 064508, 2018.
- [158] J. M. CHOW, J. M. GAMBETTA, A. D. CÔRGOLES, S. T. MERKEL, J. A. SMOLIN, C. RIGETTI, S. POLETO, G. A. KEEFE, M. B. ROTHWELL, J. R. ROZEN, M. B. KETCHEN, and M. STEFFEN. *Universal quantum gate set approaching fault-tolerant thresholds with superconducting qubits*. *Phys. Rev. Lett.*, **109**, 060501, 2012.
- [159] D. GREENBAUM. *Introduction to quantum gate set tomography*. arXiv:1509.02921, 2015.
- [160] O. HOSTEN, N. J. ENGELSEN, R. KRISHNAKUMAR, and M. A. KASEVICH. *Measurement noise 100 times lower than the quantum-projection limit using entangled atoms*. *Nature*, **529** (7587), 505–508, 2016.
- [161] P. A. LEE, N. NAGAOSA, and X.-G. WEN. *Doping a mott insulator: Physics of high-temperature superconductivity*. *Rev. Mod. Phys.*, **78**, 17–85, 2006.
- [162] D. ZHU, S. JOHRI, N. M. LINKE, K. A. LANDSMAN, N. H. NGUYEN, C. H. ALDERETE, A. Y. MATSUURA, T. H. HSIEH, and C. MONROE. *Generation of thermofield double states and critical ground states with a quantum computer*, 2016.

- [163] A. N. CHOWDHURY, G. H. LOW, and N. WIEBE. *A variational quantum algorithm for preparing quantum gibbs states*, 2020.
- [164] W. ISRAEL. *Thermo-field dynamics of black holes*. Physics Letters A, **57** (2), 107–110, 1976.
- [165] J. MALDACENA. *Eternal black holes in anti-de sitter*. Journal of High Energy Physics, **2003** (04), 021, 2003.
- [166] J. MALDACENA, D. STANFORD, and Z. YANG. *Diving into traversable wormholes*. Fortschritte der Physik, **65** (5), 1700 034, 2017. <https://onlinelibrary.wiley.com/doi/pdf/10.1002/prop.201700034>.
- [167] P. GAO, D. L. JAFFERIS, and A. C. WALL. *Traversable wormholes via a double trace deformation*. Journal of High Energy Physics, **2017** (12), 151, 2017.
- [168] J. WU and T. H. HSIEH. *Variational thermal quantum simulation via thermofield double states*. Phys. Rev. Lett., **123**, 220 502, 2019.
- [169] K. TEMME, T. J. OSBORNE, K. G. VOLLBRECHT, D. POULIN, and F. VERSTRAETE. *Quantum metropolis sampling*. Nature, **471** (7336), 87–90, 2011.
- [170] X. YUAN, S. ENDO, Q. ZHAO, Y. LI, and S. C. BENJAMIN. *Theory of variational quantum simulation*. Quantum, **3**, 191, 2019.
- [171] M. MOTTA, C. SUN, A. T. K. TAN, M. J. O’ROURKE, E. YE, A. J. MINNICH, F. G. S. L. BRANDÃO, and G. K.-L. CHAN. *Determining eigenstates and thermal states on a quantum computer using quantum imaginary time evolution*. Nature Physics, **16** (2), 205–210, 2020.
- [172] W. W. HO and T. H. HSIEH. *Efficient variational simulation of non-trivial quantum states*. SciPost Phys., **6**, 29, 2019.
- [173] O. F. DE ALCANTARA BONFIM, B. BOECHAT, and J. FLORENCIO. *Ground-state properties of the one-dimensional transverse Ising model in a longitudinal magnetic field*. Phys. Rev. E, **99**, 012 122, 2019.
- [174] S. HADFIELD, Z. WANG, B. O’GORMAN, E. G. RIEFFEL, D. VENTURELLI, and R. BISWAS. *From the Quantum Approximate Optimization Algorithm to a Quantum Alternating Operator Ansatz*. Algorithms, **12** (2), 2019.
- [175] S. P. PREMARATNE and A. Y. MATSUURA. *Engineering the cost function of a variational quantum algorithm for implementation on near-term devices*, 2020.
- [176] T. STAVENGA and L. DICARLO. *Lower-temperature fabrication of airbridges by grayscale lithography to increase yield of nanowire transmons in circuit qed quantum processors*, 2023.
- [177] M. STREIF and M. LEIB. *Training the quantum approximate optimization algorithm without access to a quantum processing unit*, 2019.
- [178] J. HEINSOO, C. K. ANDERSEN, A. REMM, S. KRINNER, T. WALTER, Y. SALATHÉ, S. GASPARINETTI, J.-C. BESSE, A. POTOČNIK, A. WALLRAFF, and C. EICHLER. *Rapid high-fidelity multiplexed readout of superconducting qubits*. Phys. Rev. App., **10**, 034 040, 2018.
- [179] J. A. SCHREIER, A. A. HOUCK, J. KOCH, D. I. SCHUSTER, B. R. JOHNSON, J. M. CHOW, J. M. GAMBETTA, J. MAJER, L. FRUNZIO, M. H. DEVORET, S. M. GIRVIN, and R. J. SCHOELKOPF. *Suppressing charge noise decoherence in superconducting charge qubits*. Phys. Rev. B, **77**, 180 502(R), 2008.
- [180] A. FRANCIS, D. ZHU, C. HUERTA ALDERETE, S. JOHRI, X. XIAO, J. FREERICKS, C. MONROE, N. M. LINKE, and A. KEMPER. *Many body thermodynamics on quantum computers via partition function zeros*. ArXiv:2009.04648, 2020.
- [181] T. HEAD, M. KUMAR, H. NAHRSTAEDT, G. LOUPPE, and I. SCHERBATYI. *Scikit-optimize*, 2020.
- [182] B. M. TARASINSKI, V. P. OSTROUKH, X. BONET-MONROIG, T. E. O’BRIEN, and B. VARBANOV. *quantumsim*. <https://gitlab.com/quantumsim>, 2016.

- [183] D. C. MCKAY, C. J. WOOD, S. SHELDON, J. M. CHOW, and J. M. GAMBETTA. *Efficient z gates for quantum computing*. Phys. Rev. A, **96**, 022330, 2017.
- [184] F. W. STRAUCH, P. R. JOHNSON, A. J. DRAGT, C. J. LOBB, J. R. ANDERSON, and F. C. WELLSTOOD. *Quantum logic gates for coupled superconducting phase qubits*. Phys. Rev. Lett., **91**, 167005, 2003.
- [185] F. WELLSTOOD, C. URBINA, and J. CLARKE. *Excess noise in dc squids from 4.2 k to 0.022 k*. IEEE transactions on magnetics, **23** (2), 1662–1665, 1987.
- [186] E. PALADINO, Y. M. GALPERIN, G. FALCI, and B. L. ALTSHULER. *$1/f$ noise: Implications for solid-state quantum information*. Rev. Mod. Phys., **86**, 361–418, 2014.
- [187] B. M. VARBANOV, F. BATTISTEL, B. M. TARASINSKI, V. P. OSTROUKH, T. E. O'BRIEN, L. DI-CARLO, and B. M. TERHAL. *Leakage detection for a transmon-based surface code*. npj Quantum Information, **6** (1), 102, 2020.
- [188] M. A. ROL. *Control for programmable superconducting quantum systems*. PhD Dissertation, TU Delft, 2020.
- [189] B. FOXEN, C. NEILL, A. DUNSWORTH, P. ROUSHAN, B. CHIARO, A. MEGRANT, J. KELLY, Z. CHEN, K. SATZINGER, R. BARENDs, F. ARUTE, K. ARYA, R. BABBUSH, *et al.* *Demonstrating a continuous set of two-qubit gates for near-term quantum algorithms*. ArXiv:2001.08343, 2020. 2001.08343.

