

DISS. ETH NO. 27032

**SIMULATING MANY-BODY QUANTUM SYSTEMS:
QUANTUM ALGORITHMS AND EXPERIMENTAL
REALISATION**

A thesis submitted to attain the degree of
DOCTOR OF SCIENCES of ETH ZURICH

(Dr. sc. ETH Zurich)

presented by

Natalie Julie Pearson

MPhys. in Physics University of Oxford

born on 13.12.1992

citizen of the
United Kingdom of Great Britain and Northern Ireland (UK)

accepted on the recommendation of
Prof. Dr. Matthias Troyer, examiner
Prof. Dr. Charles M. Marcus, co-examiner
Prof. Dr. Renato Renner, co-examiner

2020

Natalie Julie Pearson
*Simulating Many-Body Quantum Systems:
Quantum Algorithms and Experimental Realisation*

Diss. ETH No. 27032

E-mail npearson@phys.ethz.ch

Abstract

The simulation of many-body quantum systems is an extensive field of research which has a large number of applications and has motivated numerous experimental, theoretical and computational studies. It hopes to answer fundamental questions about physical phenomena and aid in the development of new and exciting materials. At the frontiers of research, classical simulations grapple with increasingly large and complex systems which quantum simulators hope to alleviate. Even within quantum simulation there is great diversity as to the hardware used to perform calculations and whether simulations are implemented in an analogue or digital manner. In this thesis we explore elements from across the field.

We investigate the limits of classical simulation and the quantum resources that are required to outperform it. We include analysis of the behaviour of errors when the simulation is Trotterised which motivates an alternate perspective for the implications of this scheme. The execution of digital quantum simulation will be implemented on quantum hardware. Here, we present results of experiments on superconducting qubits of both the ‘traditional’ transmon type and a novel superconductor-semiconductor implementation of the transmon called the gatemon. We explore experiments aimed at increasing the connectivity of these qubits by use of a superconducting bus. A new realisation of these qubits is reported which enables them to be fabricated scalably and Bayesian inference is employed to efficiently characterise superconducting qubits.

Zusammenfassung

Die Simulation von Vielkörper-Quantensystemen ist ein umfangreiches Forschungsgebiet, mit einer Vielzahl von Anwendungen, das zahlreiche experimentelle, theoretische und rechnergestützte Studien motiviert hat. Diese zielen auf die Beantwortung grundlegender Fragen physikalische Phänomene ab und man verspricht sich Impulse in der Entwicklung neuer Materialien. Doch die klassischen Simulationen kommen in der aktuellen Forschung bei mit den stets größer und komplexer werdenden Systemen immer öfter an ihre Grenzen. Hier versprechen Quantensimulatoren diese Simulationshürden zu überwinden. Doch auch innerhalb der Quantensimulation gibt es große Unterschiede hinsichtlich der verwendeten Hardware, und ob die Simulationen auf einer analogen oder digitalen Plattform durchgeführt werden.

Wir erforschen die Grenzen der klassischen Simulation und die Quantenressourcen, die erforderlich sind, um diese zu übertreffen. Wir schließen hier explizit die Fehleranalyse bei der Trotterisierung der Simulation mit ein, was eine alternative Perspektive für die Auswirkungen dieses Schemas nahelegt. Die Ausführung der digitalen Quantensimulation wird auf Quanten-Hardware implementiert. Hier präsentieren wir die Ergebnisse von Experimenten mit supraleitenden Qubits, sowohl des traditionellen Transmon-Typs als auch einer neuartigen Supraleiter-Halbleiter-Implementierung des Transmons, dem so genannten Gatemon. Wir untersuchen Experimente, die darauf abzielen, die Konnektivität dieser Qubits mit Hilfe eines supraleitenden Busses zu erhöhen. Wir präsentieren eine neue Realisierung dieser Qubits, welche eine skalierbare Herzustellung ermöglicht, und verwenden die Bayes'sche Inferenz Methode um die supraleitenden Qubits effizient zu Charakterisieren.



Contents

Abstract	iii
Zusammenfassung	v
Contents	vii
1 Introduction	1
2 Fundamentals	7
2.1 The Bloch Sphere	7
2.2 Single-Qubit Gates	8
2.3 Two-qubit Gates	11
2.4 Trotter-Suzuki decomposition	12
3 Resource Estimation	17
3.1 Matrix Product States	17
3.2 Error Comparison	18
3.3 Fermi-Hubbard Model	20
3.4 Long-Range Ising Model	42
3.5 Conclusion and Outlook	48
4 Trotter Error Scaling	51
4.1 Introduction	52
4.2 Background	53
4.3 Numerical Methods: Thermal State Properties	55
4.4 Criticality of the Transverse Field Ising-Trotter-Suzuki Model	60
4.5 Local Observables	63
4.6 Non Local Observables	66
4.7 Conclusion and Outlook	67
5 Superconducting Qubits	71

5.1	Defining the qubit	72
5.2	Readout	79
5.3	Control and Characterisation	87
5.4	Two-Qubit Interaction	95
5.5	Summary	97
6	Connectivity: Voltage-controlled superconducting quantum bus	99
6.1	Introduction	99
6.2	Background	101
6.3	Experimental Realisation and Measurement Techniques	103
6.4	Results and Discussions	105
6.5	Conclusions and Outlook	108
6.6	Acknowledgements	108
7	Scalability: Superconducting gatemon qubit based on a proximitized two-dimensional electron gas	111
7.1	Introduction	112
7.2	Reliable fabrication	113
7.3	Single-Qubit Control and Characterisation	114
7.4	Two-Qubit Interaction	116
8	Calibration and Characterisation: Hamiltonian Learning	121
8.1	Introduction	122
8.2	Single Parameter Inference	124
8.3	n Parameter Inference	127
8.4	Model Comparison	130
8.5	Conclusion and Outlook	131
9	Conclusion and Outlook	133
A	Observable and error model choice for resource estimation	137
B	Q# code for gate count estimates	143
B.1	Long range Ising model	143
B.2	Fermi-Hubbard model	147
C	Fabrication details and setup for the voltage-controlled superconducting quantum bus	159
C.1	Device Fabrication	159
C.2	Setup wiring	159
D	Fabrication details and setup for the 2DEG gatemon	163

D.1 The sample	163
D.2 Qubit manipulation and readout.	164
List of Tables	165
List of Figures	167
List of publications	171
Bibliography	173
Acknowledgements	193

Chapter 1

Introduction

Since the birth of quantum mechanics in the early 1900s, the prospect of understanding the behaviour of the world around us at its most microscopic, and arguably most fundamental level, has formed a substantial part of Physics research, as well as that of many other disciplines. The behaviour of individual particles can be well understood but that of many-body quantum systems is exponentially more complex. The physics that we observe around us is inevitably the result of many particles, and as such the behaviour of many-body quantum systems and the phenomena that arise from their interactions is of great interest. Complex phenomena such as quantum magnetism [8], superconductivity [9], and topological states of matter [10, 11] are just a few of the manifestations of quantum behaviour which affect the macroscopic world and motivate research into understanding the underlying mechanisms which cause them.

The interplay between the many degrees of freedom and the competition between different energy scales gives rise to intricate emergent behaviour [12] which makes correlated many-body quantum systems both more interesting and correspondingly more difficult to treat both theoretically and computationally. We look to research in this field for explanations to poorly understood phenomena and to better understand our observations of the world. We also hope that it can provide new solutions to challenges we face by means of new materials, quantum computation, memories and sensing capabilities and that the research alone will provide us with new ways of thinking and problem-solving. Abundant fruit has already been borne of these endeavours across the field; we have gained understanding of the quantum phases of electrons in solids [13, 14], magnetic nanotechnology [15, 16], 2 dimensional graphene [17] and the realisation of a Bose-Einstein condensate among many other discoveries [18–21]

The challenges involved in simulating such systems lie in the vast amount of information that is required to fully characterise them. The resources required for such simulations

in the worst case grow exponentially with the system size and the assumptions and approximations that we can make to mitigate this problem will naturally have limiting effects on the obtained results. Methods such as density matrix renormalisation group (DMRG) [22] and matrix product states (MPS) [23–26] aim to simulate such systems as exactly as possible by truncating the size of the space simulated to contain only the ‘important’ parts while still faithfully representing the physics. Sampling methods such as Monte Carlo (MC) are frequently used across classical and quantum simulation to representatively sample distributions. Together with developments in update schemes these sampling methods can be used to find the properties of certain quantum systems which would otherwise be beyond the reach of classical simulation [27–34]. Dynamical mean field theory (DMFT) is another useful method for classically simulating strongly correlated many-body quantum systems by mapping the many-body problem to an effective quantum impurity problem with a self consistency condition [35, 36]. Research in these areas continues to be vigorous and rich, providing answers to the questions posed and also generating new ideas and approaches [37, 38].

Another contender for the simulation of quantum systems is to use a quantum system itself as a simulator. This was famously proposed by Richard Feynmann in 1982 [39], when he postulated that beyond a certain point we would need the simulators themselves to be quantum mechanical in order to match the behaviour and complexity of the systems we wish to simulate. This has been one of the most exciting possibilities expected of the field of quantum computation which also promises to revolutionise cryptography [40, 41], optimisation solutions [42] and even prediction of the weather [43]. The fundamental idea is to use a quantum system which is well understood, and over which we have control, to mimic the behaviour of the system of interest in such a way that we can study it and find its properties. One example of such a phenomena, that we do not have a full understanding of and struggle to simulate classically, is high-temperature superconductivity [44, 45], which aside from being interesting for its own sake offers tantalising real-world applications by bringing the benefits of superconductivity at a more feasible temperature for widespread use.

The candidates of hardware architectures from which we might be able to build a quantum computer are as wide and varied as the classical simulation techniques we have mentioned. Ion traps [46, 47], superconducting circuits [48–51] and electron spins [52, 53] are only a few of the contenders which have realised qubits with the potential to form quantum computers. Each have their own unique benefits and challenges, but it is not yet clear which will be the most generally applicable or if indeed we must satisfy ourselves with using different systems for different purposes. In the same way that classical computing was originally an analogue machine, the simplest quantum simulator is also analogue, matching the Hamiltonian of the simulator to the Hamiltonian of interest and letting the system evolve under it. This naturally means that unless the quantum

simulator is endlessly flexible we are limited in which systems we are able to implement and we are also susceptible to the accumulation of errors, without an effective way to correct them. Digital quantum computing instead uses a universal gate set to approximate evolution under a Hamiltonian. This approximation introduces its own errors, but the digitisation of the computation allows for error correction and the universality of the gate set means that the simulator is flexible to simulate any Hamiltonian. The subject of such error correction [54, 55] and the implementation of such a universal gate set [56] are themselves active fields of research.

This thesis will cover the breadth of the field, focussing in detail on some of the many interesting problems within it. Although far from a complete analysis we will give an overview of the challenges and their solutions at each level, from calibrating an individual qubit all the way to analysing the growth of error in simulation with system size. We will look at work done in the superconducting qubit community to build a scalable, connected quantum computer and estimate the resources required for a digital or analogue quantum computer to outperform classical simulations.

Thesis outline

Chapter 2: First we will introduce some fundamental aspects of quantum computing. This includes the representation of a state on the Bloch sphere and the single qubit operations we can use to manipulate it. Two-qubit gates will be introduced and for each operation the matrix and circuit representations are provided so that the more complex circuit diagrams we will later encounter are clear to the reader. Finally we will outline the method of Trotterisation for breaking up the evolution of a system under a Hamiltonian into steps that are feasible to execute both in classical simulation and on quantum hardware. We will assume some basic knowledge of quantum mechanics but will endeavour to be as explicit as possible throughout. In this chapter we aim to arm the reader with the notation and concepts that we will encounter throughout the thesis.

Chapter 3: In this chapter we investigate the necessary requirements for a quantum simulator to outperform classical computers at a meaningful problem. This subject is of course open to much debate and as such we will not offer any kind of rigorous proof but rather present the current most promising contenders in terms of using quantum analogue and digital simulators to outperform classical computers at the same task. The metrics we compare are errors on observables and the models we look at are the Fermi-Hubbard model and the long range transverse field Ising model. This analysis will focus on obtaining the gate counts for a digital simulation and details the optimisation schemes employed to minimise these. A Q# code accompanies this thesis which we hope will enable the reader to both verify the results and build upon them.

Chapter 4: Next, we move onto investigating the behaviour of errors on local and non-local observables around the critical point of the 1 dimensional transverse field Ising model as a function of the system size. The error in question is introduced by the Trotter break up, mentioned in Chapter 2, and is relevant both for classical and digital quantum simulation. We show that the upper bounds on this error are loose and propose to consider the error as a perturbation of the simulated Hamiltonian rather than an increasing error with evolution time. This is supported by our findings that the errors on both local and non-local observables saturate with system size. In the process of this investigation we find that the critical point of the nearest neighbour transverse field Ising model is unexpectedly invariant with Trotterisation. The combination of these observations provides motivation to consider the observables of interest and the model behaviour under Trotterisation when estimating the resources required. This makes us optimistic for the prospect of near term simulation.

Chapter 5: In this chapter we turn to the question of implementing quantum simulation on quantum hardware and introduce one of the most promising contenders for an architecture to realise this: the superconducting qubit. We detail how the qubit is composed and outline the physics of how it is prepared, operated, read out and coupled to other qubits. This is itself a very rich field and there is much literature on the subject. This chapter is intended to enable the reader to follow the subsequent chapters, where we will be detailing experiments exploring and improving the feasibility of this type of quantum hardware. We outline methods for calibrating the system as well as some of the strengths and weaknesses of this architecture. We introduce a specific type of superconducting qubit called the gatemon [57], made from superconductor-semiconductor nanowires rather than traditional superconductor-insulator junctions. This avoids the use of currents at low temperatures which can cause heating.

Chapter 6: The final three chapters examine superconducting qubits and detail experiments carried out in each of three of the main areas relevant to the prospect of using them for large scale computation. The first of these chapters describes the results of an experiment intended to increase the connectivity of a superconductor-semiconductor qubit device to facilitate all-to-all coupling of the qubits with on-off control. We measure two qubits coupled by a tunable bus which mediates the interactions between them. We are able to effectively turn on and off this coupling as desired, but operation on a time-scale necessary for implementing an algorithm on the qubits showed unexpected behaviour which would need to be resolved for this method to be used as intended.

Chapter 7: The second of these chapters describes an experiment in which we realise fabrication of a gatemon qubit using a 2 dimensional electron gas (2DEG) which increases the possibility to fabricate the gatemon qubit scalably as is necessary for building a useful quantum computer. We perform single and two-qubit calibration and report the results. We find that we can successfully fabricate qubits of this kind with

competitive coherence and relaxation times and capacitive nearest neighbour coupling. The coherence times are limited primarily by dielectric losses, as is the case for most superconducting qubit schemes, and as such we hope to benefit from further research in this area.

Chapter 8: In the final experimental chapter we implement Bayesian inference to efficiently learn the Hamiltonian of a superconducting qubit device. This method has many existing applications [58] and has recently been applied to the field of quantum Hamiltonian learning with the motivation to limit the amount of data we need to take and the number of experiments needed to accurately characterise parameters of the system [59, 60]. As quantum computers grow, so too will the challenge of calibration and characterisation of the system and so it is imperative that we explore optimal ways to learn about these systems so that we can operate them at their best as quickly as possible. Here we employ this method to learn about parameters of our superconducting qubit system and successfully learn its Hamiltonian using limited data.

We conclude the work of this thesis with a review of the results made and comment on the future outlook for the field. This ‘full stack’ work can hopefully give an overview of the field and additionally provide detailed information on specific topics at many levels.



Chapter 2

Fundamentals

Over the course of this thesis we will make extensive use of a number of building blocks which we will introduce in a concise fashion in this chapter. We will introduce the representation of a single quantum bit (qubit) on the Bloch sphere and the fundamental gates which we can use to traverse its surface, introducing notation and circuit diagrams as we go. We will then move onto two-qubit gates and their representations. Following this we will explain the decomposition used to break up evolution of a many-body system into small steps which we can then map to gates that can be executed on our qubits or in classical simulations, thus bridging the gap between the simulation we wish to implement and the physical execution of the algorithm on quantum or classical hardware.

2.1 The Bloch Sphere

The wavefunction of a single qubit with $d = 2$ requires two complex numbers α and β to define the state $|\psi\rangle$,

$$|\psi\rangle = \alpha |0\rangle + \beta |1\rangle, \quad (2.1)$$

where $|\alpha|^2 + |\beta|^2 = 1$ so that the wavefunction is normalised. We can further neglect the global $U(1)$ phase to reduce the number of real parameters characterising the wavefunction to 2. The wavefunction of a pure state is therefore commonly represented as a vector on the Bloch sphere as shown in Fig. 2.1, where these two real parameters are the angles θ and ϕ , correspond to the polar and azimuthal angles on the sphere:

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right) |0\rangle + e^{i\phi} \sin\left(\frac{\theta}{2}\right) |1\rangle. \quad (2.2)$$




Figure 2.1: The Bloch sphere which can be used to visualise the wavefunction of a single qubit. Pure states are found on the surface of the sphere and unitary evolutions implement rotations which can be decomposed into components acting around the x , y and z axes. The $|0\rangle$ state of the qubit is at the north pole of the sphere and $|1\rangle$ at the south pole.

When using superconducting qubits we refer to $|0\rangle$ and $|1\rangle$ as the ground and excited state respectively as this describes the two states of the system which we identify them with. This is discussed in more detail in Chapter 5.

2.2 Single-Qubit Gates

Unitary single-qubit operations allow us to traverse the surface of the Bloch sphere. Any of these operations can be decomposed into the Pauli matrices:

$$\begin{aligned}
 \sigma_I = I &= \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \\
 \sigma_x = X &= \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \\
 \sigma_y = Y &= \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \\
 \sigma_z = Z &= \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.
 \end{aligned} \tag{2.3}$$

On the Bloch sphere each of these corresponds to a rotation of π around the axis in question. Rotation by an arbitrary angle around each axis can also be achieved by executing a rotation gate. These are defined

$$R_\alpha(\theta) = e^{-i\frac{\theta}{2}\sigma_\alpha} = \cos\left(\frac{\theta}{2}\right)\sigma_I - i\sin\left(\frac{\theta}{2}\right)\sigma_\alpha, \tag{2.4}$$

such that

$$R_\alpha(\pi) = -i\sigma_\alpha \rightarrow X, Y, Z, \tag{2.5}$$

neglecting the global phase $-i$. The notation for a rotation by $\pm\pi/2$ is

$$R_\alpha(\pm\pi/2) \rightarrow \pm X/2, \pm Y/2, \pm Z/2, \tag{2.6}$$

such that two of such gates would be a π rotation about that axis.

To depict the execution of a rotation gate on a qubit we use the circuit element

$$\text{---} \boxed{R_\alpha(\theta)} \text{---}$$

where the system is in the initial state at the left and ‘moves through’ the R_α element, acquiring the rotation as it does so, to end in the final state on the right.

We can also combine the Pauli matrices to describe the creation and annihilation operators for an excitation on a qubit. These are the σ_+ and σ_- operators, where the former implements $|0\rangle \rightarrow |1\rangle$ and the latter does the reverse. They are neither Hermitian nor

unitary and so the Bloch sphere is no longer a useful tool for visualising their operation but they will be relevant when considering interaction of a qubit with another system and when using a qubit to simulate a system with such behaviour. They are defined as

$$\sigma_{\pm} = \frac{1}{2}(\sigma_x \pm i\sigma_y). \quad (2.7)$$

The Clifford gate set is useful for quantum computation. These gates affect permutations of the Pauli gates. They take a state from any one of the six points where an octahedron touches the surface of the Bloch sphere to any other of these points. In practice, in the single-qubit case, they are often implemented using the calibrated gates $\pm X(/2)$ and $\pm Y(/2)$ to apply the 24 possible operations by applying the following gate sequences

Table 2.1: Physical gate sequences corresponding to the 24 single-qubit Clifford rotations

I	X	Y	$X/2$	$-X/2$
$Y/2$	$-Y/2$	YX	$X/2 Y/2$	$X/2 -Y/2$
$-X/2 Y/2$	$-X/2 -Y/2$	$Y/2 X/2$	$Y/2 -X/2$	$-Y/2 X/2$
$-Y/2 -X/2$	$X Y/2$	$X -Y/2$	$Y X/2$	$Y -X/2$
$-X/2 Y/2 X/2$	$-X/2 -Y/2 X/2$	$X/2 Y/2 X/2$	$-X/2 Y/2 -X/2$	

Together with the so called T gate which implements a $\pi/4$ rotation around the z axis the Clifford plus T group is sufficient to reach all points on the Bloch sphere [61]. The T gate is


$$T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix} \sim R_x(\pi/4) \quad (2.8)$$

up to a global phase.

Specific Clifford gates of note include the Hadamard gate (H) which is used to change between the z and x basis by rotating π about an axis that lies between the z and x bases and is its own inverse. The Hadamard gate is

$$H = H^{-1} = X Y/2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \quad (2.9)$$

and has the circuit representation




Another Clifford gate used here is the Y^* gate which is used to change between the z and the y basis by rotating by $\pi/2$ around the x axis, the notation is used to signal this

basis change rather than the axis of rotation as is otherwise used.

$$\begin{aligned}
 Y^* &= X/2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ -i & 1 \end{pmatrix} \\
 Y^{*-1} &= -X/2 = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & i \\ i & 1 \end{pmatrix}.
 \end{aligned}
 \tag{2.10}$$


The circuit representation we use for this gate is



2.3 Two-qubit Gates


The two-qubit gates we use are the controlled- X gate (CNOT) and the controlled- Z gate (CZ). These apply the X or Z gate on the target qubit if the source is in $|1\rangle$ and do nothing otherwise. The controlled X gate is

$$\text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}
 \tag{2.11}$$



and the controlled Z is

$$\text{CZ} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}
 \tag{2.12}$$



In each of the circuit representations the upper line represents the source or control qubit, the operation being conditional on the state of this qubit. The lower line represents the target qubit on which the operation will be implemented in as much as the implementation can be said to be enacted on one of the two qubits given that they may be entangled.

2.4 Trotter-Suzuki decomposition

A quantum system with wavefunction $|\psi\rangle$ evolves in time, dependent on the Hamiltonian which governs it. For a time independent Hamiltonian, which for the remainder of the chapter we will use H to denote, the evolution proceeds according to the Schrödinger equation,

$$i \frac{\partial}{\partial t} |\psi\rangle = H |\psi\rangle \quad (2.13)$$

which is solved by

$$|\psi(t)\rangle = U(t) |\psi(0)\rangle \quad (2.14)$$

where $U(t) = e^{-iHt}$ and we have used $\hbar = 1$ as we will continue to do in this chapter.

It may be computationally or physically difficult to implement this unitary evolution matrix if calculating it requires the diagonalisation of a large matrix or if physically we are limited to single and two-qubit gates of the kind described above which do not correspond to what is potentially a large many-body operation. The Trotter-Suzuki decomposition enables us to implement many-body evolution by breaking it up into few-body terms which can then either be implemented directly in the native gates of the digital quantum computer or easily rewritten in terms of these native gates [62–64]. The ‘native’ gates of the system here means whatever gate set we have access to which may include the physical gates that we can apply or the logical gates which we are able to implement. These few-body terms are also easily diagonalised by a classical computer, thus overcoming the challenge of diagonalising large matrices which we otherwise face in exact classical simulation.

If the Hamiltonian itself contains only few-body terms then it is possible to split it into sections within which all the terms commute and the Trotter-Suzuki product decomposition can be used to simplify the process of evolving under this Hamiltonian. It is first necessary to split the evolution over time t into a series of small time-steps, each of size τ , using the property

$$U(t) = e^{-iHt} = \prod^{N_\tau} e^{-iH\tau}, \quad (2.15)$$

where $t = N_\tau\tau$ and N_τ is the number of time-steps.




Figure 2.2: Sketch of Trotterisation which breaks up an evolution under unitary $U(t)$ into N_τ steps which each implement $U_{\text{TS}}(\tau)$ which is approximately equal to $U(\tau)$. Hence the full evolution implemented $[U_{\text{TS}}(\tau)]_\tau^N \sim U(t)$. The individual time-steps can be implemented using native single and two-qubit gates (in orange) where the exact decomposition is flexible.

We then implement the first order Trotter-Suzuki decomposition which splits the time evolution operator into products of the non-commuting terms. We write $H = \sum_{k=1}^K H_k$, where the terms within H_k commute with each other but $[H_j, H_k]$ is nonzero unless $j = k$. The time evolution can now be written

$$e^{-iH\tau} = \prod_{k=1}^K e^{-iH_k\tau} + \mathcal{O}(N\tau^2), \quad (2.16)$$

where N is the number of sites in the system and we can now in practice use sequential

application of the time-step operator

$$U_{\text{TS}}^{(1)}(\tau) := \prod_{k=1}^K e^{-iH_k\tau} \quad (2.17)$$

to evolve the system, accumulating errors which can be limited by decreasing τ . The nature and effect of the error will be discussed in more detail in the next chapters. We can then further decompose without introducing additional errors, since within H_k all terms commute. Thus writing it as a sum over the sites, $H_k = \sum_{i=1}^N s_i$, where s_i is a local (usually single or two-qubit) operation, we see that

$$U_{\text{TS}}^{(1)}(\tau) = \prod_{k=1}^K \prod_{i=1}^N e^{-is_i\tau}, \quad (2.18)$$

where we have finally ended up with the task of simply exponentiating a small and local operator s_i .

The Trotter-Suzuki decomposition facilitates the evolution of a large system under a many-body Hamiltonian by the application of operations which only act on small parts of the system and are hence easy to exponentiate. A sketch of this process is shown in Fig. 2.2.

It is possible to further limit the errors by going to higher orders of the decomposition. This involves propagating back and forth for shorter time-periods to execute a full time-step of evolution at the expense of introducing a larger number of partial steps. The second order version uses a time-step evolution decomposition introducing errors $\mathcal{O}(N\tau^3)$ and is described by

$$U_{\text{TS}}^{(2)}(\tau) := \left(\prod_{k=1}^K e^{-iH_k\frac{\tau}{2}} \right) \left(\prod_{k=K}^1 e^{-iH_k\frac{\tau}{2}} \right), \quad (2.19)$$

which for $K = 2$ terms simplifies to

$$U_{\text{TS}}^{(2)}(\tau) = e^{-iH_1\frac{\tau}{2}} e^{-iH_2\tau} e^{-iH_1\frac{\tau}{2}}. \quad (2.20)$$

This simplifies further when subsequent time-steps are applied, as adjacent $e^{-iH_1\frac{\tau}{2}}$ terms can be combined into $e^{-iH_1\tau}$, so that ultimately to complete the full evolution only one additional term is required compared to the first order decomposition. This is the case when H can be represented as the sum of its x and z terms such that $H = H_x + H_z$. It is also applicable if only nearest neighbour terms are present, in which case H can be split into the part acting on odd qubits and even ones $H = H_{\text{odd}} + H_{\text{even}}$. The Trotter decomposition will affect the nature of the errors which arise but they can always be limited by decreasing τ .

To decrease errors further the fourth order decomposition can be used, the time-step decomposition for this is

$$U_{\text{TS}}^{(5)}(\tau) := \prod_{j=1}^5 \left(\prod_{k=1}^K e^{-iH_k \tau} \right) \left(\prod_{k=K}^1 e^{-iH_k \tau} \right) \left(p_j \frac{\tau}{2} \right), \quad (2.21)$$

where $p_{1,2,4,5} = p = \frac{1}{4-4^{\frac{1}{3}}}$ and $p_3 = 1 - 4p$. Errors introduced are limited to $\mathcal{O}(N\tau^5)$ at the expense of being five times more expensive than the second order scheme.

These are the fundamentals necessary for implementing simulation of a many-body quantum system on a digital quantum computer which has a specific gate set. The decompositions described are equally necessary for classical simulation when the matrix sizes involved exceed diagonalisation capabilities. The particular mapping of the Trotterised Hamiltonian to single and two-qubit gates will necessarily depend on the Hamiltonian itself, as we will explore in more detail in the next chapter. We are now prepared to address the challenges of implementing simulations of many-body quantum systems both theoretically and experimentally.

Chapter 3

Resource Estimation

The wide ranging potential of quantum computers to perform tasks beyond the scope of classical computation is exciting. However, it is not immediately obvious at which point a quantum computer will outperform classical computers to solve a useful problem. The simulation of many-body quantum systems is a promising contender for this task because of the challenges involved in simulating quantum systems using classical computers and the wide range of fields which rely heavily on such simulations [25, 26, 34, 65–68]. Quantum hardware, such as ultracold atoms, superconducting qubits and trapped ions, can be used to execute analogue quantum simulation. Digital quantum computation schemes on these hardwares also promise to implement such simulations but require sufficiently accurate operations with which to execute the digitised circuit [69–71]. In this chapter we establish the point beyond which exact classical simulation of a quantum system becomes infeasible for two hallmark systems. We introduce these two models and outline how to map simulation of these systems onto an algorithm for execution on a digital quantum computer. Finally, we carry out a detailed gate count analysis to match the accuracy achieved by state of the art ultracold quantum gas based analogue simulators. This chapter follows and expands upon Ref. [4].

3.1 Matrix Product States

The limits involved in running a quantum simulation on a classical computer include, but are not limited to, the space required to store the wavefunction. The Hilbert space grows exponentially with system size as d^N , where N is the number of sites in the system and d is the dimension of each site. As we have seen in Chapter 2, 2 real numbers are sufficient to represent the pure state of a single qubit. By the same logic for N qubits the number of real parameters needed to fully characterise the state is given by $2(d^N - 1)$.

For two qubits this equates to 6 parameters, for three qubits this already requires 14 parameters. Storing the exact wavefunction becomes quickly unmanageable, where an $N = 300$ system would require more bits than there are atoms in the universe, which is unlikely to be realisable. There are certain physical systems where the system's state can be characterised by only a few relevant degrees of freedom, making it possible to store an approximate wavefunction with much less space using tensor network algorithms such as matrix product states (MPS) [23–26, 72].

However, methods of this kind rely on being able to retain the majority of the information and represent the important physics of the system through compression, by reducing the complexity to only the most important degrees of freedom. This will inevitably introduce errors in the case that the entanglement exceeds that which can be faithfully represented under such compression schemes. For simulating dynamics typically the entanglement of a system grows, which limits such simulations to short or moderate times only.

In order to evolve the system under the Hamiltonian of interest we must find the solution to the Schrödinger equation Eq. (2.13). The Hamiltonian under which the system evolves, H , has the dimensions $d^N \times d^N$ and it is this that we need to exponentiate. Thus, in the case of a many-body system with large d , we have the challenge of diagonalising a large matrix in order to compute the evolution operator, $U(t)$, which is computationally expensive. As covered in Chapter 2, this is achievable by using the Trotter-Suzuki decomposition, which nevertheless introduces errors to the evolution.

3.2 Error Comparison

The errors introduced by using MPS are singular to the classical simulation, but the Trotter errors introduced by decomposing the evolution into ‘manageable’ operations are present both in classical simulation and in digital quantum simulation. Only in the case of analogue quantum simulation are Trotter errors not present but in this case there are likely to be limits to the calibration accuracy in the quantum hardware which will also introduce errors in the observables. The time-scale of the classical simulation is set by the growth of entanglement which limits how well the wavefunction can be represented under compression. In analogue quantum simulation the time-scale limitations are specific to the system in use. For the case of ultracold gases, which we will consider here, the limit is set by heating in the system. It should be noted that in an analogue simulator error correction is not possible. In digital quantum simulation the time-scale is set by the gate fidelities unless there is some other limiting condition in the physical system. In fault tolerant schemes the errors due to imperfect physical gates can be mitigated, making them theoretically the most attractive candidate for executing simulations to

long times. We will henceforth compare the errors on local observables introduced by Trotterising the Hamiltonian with the errors introduced by the calibration error on an analogue quantum simulator. The effect of implementing the simulation using an analogue simulator with calibration errors or using a Trotterised method on digital or classical hardware, is to change the Hamiltonian from the desired H to some effective H_{eff} . We evolve under H_{eff} and accumulate errors relative to evolution under H . We can then compare the errors on observables after evolution under H_{eff} compared to the exact values obtained by evolving under H . Implementing this comparison and calculating the resources required to obtain the same level of accuracy in the observables beyond the limits of classical simulation capabilities can tell us the resources required to execute an intractable simulation using quantum hardware.

The observables of interest generally include local or few-body correlations so we will focus on the accuracy with which we can find these values. A more detailed analysis of the errors in local and non-local observables introduced by Trotterisation is explored in depth in the next chapter. For the models that we will be using ultracold atoms in optical lattices provide a good platform for analogue simulation [69]. We will use this as the platform with which to establish the limits of analogue quantum simulation for comparison. We here establish the error introduced due to the Trotter decomposition compared to that introduced by calibration errors in the analogue simulation. It is possible to compute this using tensor network methods for ‘small’ systems of size $M \sim 20$. This allows us to establish the decomposition and step size needed for digital quantum computers to produce comparable errors to their analogue counterparts. From this we can then work out the total number of gates needed to perform evolution up to a time t and the circuit depth in each type of gate.

For the calibration error in analogue simulation we sample the model parameters from a normal distribution with a standard deviation given by the value of the calibration error. Repeating this process and comparing the measured results of the observable to that obtained when the mean model parameter values are used we can find the errors introduced due to imperfect calibration.

It is generally established that a simulation of $M > 100$ spins or particles entangled by dynamics will not be feasible on a classical computer so it is at this point that we are able to estimate the resources required for digital and analogue quantum simulators to outperform classical computers at this task. The work involved in computing the errors introduced on observables by Trotterising and those due to the imperfect calibration of the analogue quantum computer was done by Stuart Flannigan and Andrew Daley as part of Ref. [4]. The digitised circuits and gate count estimates for the digital simulations are my own.

3.3 Fermi-Hubbard Model

The Fermi-Hubbard model is a paradigmatic model in condensed matter, providing insight into the electronic and magnetic properties of materials [73]. It is of particular interest due to the questions of whether or not the 2 dimensional (2D) Fermi-Hubbard model can support superconductivity and might help us understand the mechanism underlying high-Tc superconductivity [74, 75]. It is the simplest model one can write down, that holds all ingredients one needs for a generic conductor. Despite decades of intense investigation the Hubbard model is still not fully understood and its full solution continues to elude the physics community [73, 76]. The Hamiltonian defining this model is

$$H = - \sum_{i,j} \sum_{\sigma} t_{i,j} (c_{i,\sigma}^{\dagger} c_{j,\sigma} + h.c) + U \sum_i n_{i,\uparrow} n_{i,\downarrow} - \mu \sum_i \sum_{\sigma} n_{i,\sigma}, \quad (3.1)$$

where $\sigma = \uparrow, \downarrow$ is the spin configuration, $c_{i,\sigma}$ and $c_{i,\sigma}^{\dagger}$ are the fermionic annihilation and creation operators, and $n_{i,\sigma}$ counts the excitations on the i th site with spin σ . i and j run from 1 to M , where M is the total number of sites. We assume open boundaries and that only nearest-neighbour hopping is allowed, with all hopping equal. This implies that $t_{i,j} = J$ for neighbouring i and j and 0 otherwise.




Figure 3.1: Quantitative comparison of the errors introduced to a local observable using analogue and digital approaches to implement quantum simulation of the Hubbard model. The observable is the off-diagonal correlation. The time dependence of the calibration errors (solid blue 1%, solid red 0.1%) in the off-diagonal correlation functions and trotter decomposition errors. For the digital we have plotted the 4th order decomposition, for the time steps $J\tau = 3$ (circles), $J\tau = 3/2$ (stars), $J\tau = 1$ (plus sign) and $J\tau = 1/2$ (crosses) for system size, $M = 20$. Calculated by Stuart Flannigan

By comparing the errors on the off-diagonal correlations introduced due to imperfect calibration to those introduced due to Trotter decomposition we can establish the Trotter

step size needed to execute a comparable digital quantum simulation to the analogue simulation. This is shown in Figure 3.1 for a 4th order Trotter decomposition. Further details about the observable choice and error model for the analogue simulation can be found in Appendix A.

Geometry	Trotter Order	Calibration Error (%)	τ	Total Sweeps (per t)
1D Chain	1 (1)	1	3.5×10^{-4}	2.9×10^3
		0.1	5.0×10^{-7}	2.0×10^6
		0.01	2.0×10^{-7}	5.0×10^6
	2 (2)	1	0.12	17
		0.1	0.014	150
		0.01	9.5×10^{-3}	210
	4 (10)	1	3.2	3.2
		0.1	1.0	10
		0.01	0.86	12
Ladder	1 (1)	1	0.018	55
		0.1	0.012	81
		0.01	2.2×10^{-3}	460
	2 (2)	1	0.046	44
		0.1	0.031	64
		0.01	5.9×10^{-3}	340
	4 (10)	1	0.96	10
		0.1	0.80	12
		0.01	0.38	26

Table 3.1: Values of the total number of sweeps for digital simulation to result in the same errors in values of a local observable as would result from calibration errors of 1, 0.1 and 0.01%. The number of sweeps per time step for the Trotter decompositions of order 1 2 and 4 are shown in brackets.

For different order decompositions a different number of sweeps back and forth is required to implement the operation equivalent to evolving for one time-step, τ . As a result the total number of time-steps, N_τ , the number number of sweeps per time-step, and the gate counts which implement a single ‘time-step-sweep’ should be multiplied together to find the total number of gates needed to implement the algorithm. The number of sweeps per step for the 1st, 2nd and 4th order Trotter decompositions are 1, 2 and 10 respectively. As can be seen from Table 3.1, the 4th order decomposition is favourable as, despite the increased number of sweeps per time-step, the errors introduced to the observable are sufficiently reduced as to result in a lower number of total sweeps to complete the time-step to the same level of accuracy. This is calculated by finding the value of τ for different decompositions required to produce levels of error in the observable competitive with that introduced by 1%, 0.1% and 0.01% calibration errors. The resultant number of sweeps per unit time is calculated (where time, t , is

measured in units of J). This is done for the case of the Hubbard model for a 1 dimensional (1D) chain and a ladder. In both cases the system size used is $M = 20$ and the simulations and resulting values of τ were provided by Stuart Flannigan.

This gives us the total number sweeps needed in order to simulate evolution under the Fermi-Hubbard model up to time, t , for different systems at a comparable accuracy to analogue quantum simulation. We can use this to calculate the total number of gates needed to execute the digital quantum simulation, together with the numbers and depths of each gate type which we will calculate.

Mapping to the digital circuit

In order to run this simulation on a digital quantum computer we must rewrite the operations which constitute a single Trotter time-step-sweep as those native to the digital quantum computer. The most common method to map from the fermionic creation/annihilation picture to the qubit picture is to use the Jordan Wigner mapping [77, 78]. The Hilbert space of a site of the Hubbard model is spanned by $|0\rangle$, $|\uparrow\rangle$, $|\downarrow\rangle$ and $|\uparrow\downarrow\rangle$ where the creation and annihilation operators, $c_{i,\sigma}^\dagger$ and $c_{i,\sigma}$ create or remove a spin of type σ on site i . Each site therefore has $d = 4$ and so must be represented by two qubits, each with $d = 2$ and which can be in $|0\rangle$ or $|1\rangle$. The mapping between these two systems is simply that state of one qubit is used to represent the presence or absence of one type of spin on the corresponding site, so the mapping for a single site is:

$$\begin{aligned} |0\rangle &= |00\rangle \\ |\uparrow\rangle &= |01\rangle \\ |\downarrow\rangle &= |10\rangle \\ |\uparrow\downarrow\rangle &= |11\rangle. \end{aligned} \tag{3.2}$$

In order to maintain the fermionic anticommutation relation, it is necessary to map the creation and annihilation operators to slightly more than just the σ_+ and σ_- operators which create or annihilate an excitation on the relevant qubit as described in Chapter 2. The mapping which achieves this is described by

$$\begin{aligned} c_k &= I^{\otimes k-1} \otimes \sigma_+ \otimes \sigma_z^{\otimes N-k} \\ c_k^\dagger &= I^{\otimes k-1} \otimes \sigma_- \otimes \sigma_z^{\otimes N-k} \end{aligned} \tag{3.3}$$

which by using $n_k = c_k^\dagger c_k$ results in

$$n_k = I^{\otimes k-1} \otimes \frac{1}{2}(I - \sigma_z) \otimes I^{\otimes N-k}. \tag{3.4}$$

It turns out to be computationally more efficient to rewrite the interaction term of the Hamiltonian as

$$U \sum_i (n_{i,\uparrow} - \frac{1}{2})(n_{i,\downarrow} - \frac{1}{2}) \quad (3.5)$$

and group the extra $U \sum_i \sum_\sigma \frac{1}{2} n_i$ into the chemical potential term so that it becomes

$$\sum_i \sum_\sigma (-\mu + \frac{1}{2}U) n_{i,\sigma}, \quad (3.6)$$

where we neglect constant energy terms, as they don't contribute to the dynamics of the system. This will result in the same Hamiltonian but will later save us some gates as now the single qubit gates implementing the chemical potential term and the Coulomb repulsion term can be combined.

This leads to the mappings:

Chemical potential term

$$n_k = \frac{1}{2}(I - \sigma_z) \rightarrow -\frac{1}{2}\sigma_z \quad (3.7)$$

Coulomb repulsion term

$$(n_k - \frac{1}{2})(n_l - \frac{1}{2}) = \frac{1}{4}\sigma_z^k \sigma_z^l \quad (3.8)$$

Hopping term

$$c_k^\dagger c_l - c_k c_l^\dagger = \frac{1}{2}(\sigma_x^k \sigma_x^l + \sigma_y^k \sigma_y^l) Z_{JW}, \quad (3.9)$$

where Z_{JW} represents the tensor product of σ_z on the sites between k and l , the so called Jordan Wigner strings


$$Z_{JW} = \sigma_z^{k+1} \sigma_z^{k+2} \dots \sigma_z^{l-1}. \quad (3.10)$$

We assume a native gate set of Clifford gates, CNOTs and an arbitrary rotation of angle θ clockwise around the z axis defined by $R_z(\theta) = e^{-i\theta\sigma_z/2}$ where these definitions are more explicitly given in Chapter 2. It should be noted that the Hadamard gate, H , is used to change from the z basis to the x basis and back and the Y^* gate is used to change from the z to y basis and its inverse to change back. This notation follows that of Ref. [79] as detailed in Chapter 2. A different gate set will of course change the numbers slightly but, given that two-qubit gates are usually the most challenging to implement in practice and that arbitrary rotations are the hardest to synthesise from error correctable gates, it is a fair assumption that any similar gate set will have comparable limitations. For example, if the native two-qubit gates were CZ then we could add two Hadamard gates to convert all CNOTs to CZs, which adds negligible cost when compared to the

two-qubit gates and the rotation gates. If our rotation gates are around another axis similarly a pair of Clifford gates can change this axis at negligible cost.


Armed with the knowledge of how to convert the constituent elements of the Hamiltonian into gates native to the digital quantum computer (or at least trivially mappable to those which are), we can write down the circuits which implement evolution for time τ under each of these terms of the Hamiltonian. In the following, k is the index of the qubit for the single qubit terms, while k and l are the indices of the source and target qubits for the multi-qubit terms.

Chemical potential term (from Eq. (3.7))



where $\theta_p = -(-\mu + \frac{1}{2}U)\tau$. This implements $e^{i(-\mu + \frac{1}{2}U)\tau\sigma_z^k/2}$ as required.

Coulomb repulsion term (from Eq. (3.8))



where k and l are indices for the qubits which represent the same physical site with opposite spins and $\theta_r = \frac{1}{2}U\tau$. This implements $e^{-iU\tau\sigma_z^k\sigma_z^l/4}$ which can be seen either by simply multiplying the matrices or, more intuitively, by observing that sandwiching the z rotation on qubit l by two CNOTs means that, at the moment when the rotation is applied, qubit l will be in the $|0\rangle$ state if the starting parity of the qubits is even (ie $|00\rangle$ or $|11\rangle$), and hence accumulate a phase θ_r through multiplication by the term $e^{-iU\tau/4}$. If the parity is odd (ie $|01\rangle$ or $|10\rangle$) then l will be in the $|1\rangle$ state at the point of the rotation and will accumulate a phase $-\theta_r$ through multiplication by $e^{iU\tau/4}$. As such the direction of the rotation is dictated by the parity of the qubits.

Hopping term (from Eq. (3.9))

With an extension of the logic above, the most naive implementation of this term is described by the circuit shown in Fig. 3.2, where here k and l will index two qubits which represent the occupation of neighbouring sites by species of the same type, between which the hopping may occur and where $\theta_h = -J$.




Figure 3.2: Circuit implementing the hopping term between qubits k and l .

This can be readily understood by splitting the circuit into two halves. The first half starts by applying Hadamard gates to qubits k and l to change their basis from x to z . Following this, operations on these two qubits flipping or rotating around the x axis can be viewed to flip or rotate around z and vice versa. Next, a string of CNOT gates implements the Jordan Wigner string, Z_{JW} , so that the direction of the rotation on qubit l depends on the parity of the number of excitations on qubits k , l and all intervening qubits. Due to the basis change the ‘parity’ of qubits k and l is taken in the x basis rather than z . This operation can be understood to be similar to that implementing the Coulomb repulsion terms except that here it is the parity of all of the qubits between the source and target, rather than just k and l , that determines the direction of rotation and that the basis for the source and target is different from that of the other qubits. After applying the rotation it is necessary to ‘uncompute’ the Z_{JW} string to return the qubits to their original state, having used the computed parity value to determine the direction of rotation on l . The source and target qubits are then returned to their original basis by a second pair of Hadamard gates and the first part of the circuit is concluded. At this point the $\sigma_x^k \sigma_x^l Z_{JW}$ part of the hopping has been completed. The second half of the circuit repeats this process but where instead the basis change is to switch from the z to the y basis for qubits k and l in order to implement the $\sigma_y^k \sigma_y^l Z_{JW}$ part of the hopping. A pair of Y^* gates is used to execute this basis change and a pair of Y^{*-1} gates to undo it.

As described in Ref. [79] it is useful to rewrite this circuit in order to minimise the number of CNOT gates needed to implement it. The motivation for this can be understood by observing that calculating the parity of the intervening qubits was done by implementing two sets of Z_{JW} CNOT chains and their corresponding ‘uncompute’ chains. Effectively, we calculated the same parity value twice which seems redundant. The second observation that helps us is to recognise that since we return all the inter-

vening qubits to their original state by this ‘uncompute’ chain it doesn’t matter how we calculate the parity of the system so long as l is rotated the right way around the right axis and that the other qubits are all returned to their original states. This allows us to calculate the parity of the qubits between k and l onto l by a chain of CNOT gates, starting at $k + 1$ and ending at l , *before* adding the parity contribution of qubit k . This alternate circuit diagram is shown in Fig. 3.3.




Figure 3.3: Circuit implementing same computation as in Fig. 3.2 but rearranged to enable CNOT cancelling.

Rewriting it in this way enables us to cancel the CNOT and CZ gates between the two halves of the circuit and to use the computed parity value for both rotations before uncomputing it as shown in Fig. 3.4.




Figure 3.4: Circuit implementing same computation as in Figs. 3.2 and 3.3 but now optimised to minimise number of CNOTs gates.

A CZ gate is introduced simply because the basis change now happens after the chain of CNOTs and not before and so this control gate must act to flip (or not flip) around the z axis in order to have the same effect as it previously did acting around the x axis after a Hadamard. If CZ gates are not native to the hardware then one can easily formed from a CNOT and two Hadamard gates. Since Clifford gates are likely to be ‘cheap’ in this scenario we will henceforth count a CZ gate as a CNOT as converting between them has negligible cost compared to rotation gates and two-qubit gates which are likely to be the limiting factor.

In the case that we need to implement the hopping between two qubits on either side of k and l , we can extend this observation even further to minimise the gate count. We do this by noting that the parity calculation from the Jordan Wigner string between qubits k and l is most of the calculation needed to find the parity of the excitations between qubits j and m for $j < k < l < m$. It is possible to compute the result of the Jordan

Wigner string between k and l onto an ancilla qubit. We then use this to implement the rotations on l in the correct direction, and then, through further computation, complete the computation of the j to m string onto the ancilla. This is then used to determine the direction of the rotations on m before uncomputing everything. The first step of this is to rewrite Fig. 3.4 using an ancilla which starts in $|0\rangle$ and is labelled with $|\alpha\rangle$ as shown in Fig. 3.5. This alone does not simplify anything and only adds another qubit but paves the way for us to reuse the parity calculation.




Figure 3.5: Circuit implementing hopping between qubits k and l as in Figs 3.2, 3.3 and 3.4 but with the addition of an ancilla qubit to store the parity calculation so that it can be reused for other hopping terms.

In order to keep the circuits simple, from now on the traditional swap circuit element is used represent the two-qubit circuit implementing the rotation, which we will call the ‘rotation sequence’.




Figure 3.6: Definition of circuit element to implement rotation sequence between two qubits.

Using this it is now simple to write the circuit for implementing the hopping between k and l , followed by that between j and m . This is shown in Fig. 3.7.

It is then trivial to cancel the CNOT gates, producing the simplified circuit in Fig. 3.8.

Here we only saved six CNOT gates (i.e. $2(l - k - 2)$), but since we can reuse this trick to implement the hopping term for the next outer pair of qubits and so on, the number of CNOT gates used to execute the hopping terms across the whole system can be drastically reduced depending on the structure and connectivity of the system and how the qubits representing it are arranged. We can go one step further by nesting terms, which will not reduce the gate count but the number of sequential gates of each type we need to apply. This is known as the circuit depth. Since the rotation sequence implementing hopping between qubits k and l , as shown in Fig. 3.6, doesn’t change the parity of the qubits between j and m we can commute the additional CNOTs and




Figure 3.7: Hopping between qubits k and l followed by that between j and m where we use the same ancilla qubit to store the parity of the Jordan Wigner string for both operations.




Figure 3.8: Executes the same evolution as Fig. 3.7 but with CNOT gates cancelled, thus reusing the calculation of the parity between k and l to calculate that between j and m .

CZ necessary to finish computing this parity through to before the rotation sequence. Doing this means that the two rotation sequences can be executed in parallel, which is useful if rotation gates are generated and split between many qubits as in Ref. [80]. It also means that this part of the circuit will only have a depth of 4 in CNOTs and 4 in Cliffords, whereas usually sequential execution would add a scaling $\propto N$ for the depth in each of these gates. The circuit reorganised in this way is shown in Fig. 3.9.

Coulomb repulsion

This term consists of two CNOT gates and one rotation gate. It includes only two qubits: those representing the up and down spins for a single site, and hence can be implemented across all qubits in parallel, as there are no shared qubits between sites.

Gate	Gate Count	Depth
$R_z(\theta_r)$	M	1
CNOT	$2M$	2

Table 3.3: Resources required to implement the Coulomb repulsion term on qubits simulating N sites of the Hubbard model.

The gate count and depth differs, however, for the hopping term as this encapsulates the physics of the different dimensions and couplings of the systems. The gate counts for the different structures are as follows.

Chain




Figure 3.10: Qubit ordering (i) and bond grouping (group 1 - red, group 2 - green) for the qubits corresponding to the occupation of a chain of $M = L = 6$ by one species of fermions. An identical chain represents the other species so the total number of qubits is $N_Q = 12$.

The simplest case is a 1D chain of sites. We arrange the qubits by site and then spin. This has no effect on the previous numbers calculated as we are assuming all-to-all coupling and so the CNOT gates of the Coulomb repulsion term, which includes the two qubits representing one site, commute with those acting on any other sites, whether or not the qubits themselves are neighbouring. Figure 3.10 illustrates the ordering of the qubits representing the sites of one species of the Hubbard model. Ordering the qubits in this way means that, for the 1D chain, qubits representing the same spin on neighbouring sites will neighbour each other and, since we are limited to only nearest-neighbour hopping, this eliminates the Jordan Wigner strings as there are no sites in between for us to ‘count the excitations on’. As such, we will implement the circuit described in Fig. 3.4 as there is no possibility to minimise gate counts by reusing calculations stored in an ancilla. We can execute terms which do not include the same sites simultaneously to minimise circuit depth. To achieve this we split bonds between sites into two groups, where the hopping terms within one group can be carried out simultaneously. This is shown in Figure 3.10 where the qubit ordering proceeds from

left to right, following the arrows, and the alternating coloured bonds represent the two groups that the hopping terms form. Within one group these terms commute, as the source and target qubits and the Jordan Wigner string joining them have no shared qubits. The two groups however do not commute with each other because they do share qubits. This therefore instructs us on how to parallelise the operation, by first applying the terms of group 1 simultaneously and then those of group 2. No ancilla qubits are needed, so the number of qubits, N_Q , is $2M$.

To execute the full hopping term, we then need simply to sum the operations required to complete the circuit shown in Fig. 3.4 with $l = k + 1$. This comes to 8 Cliffords, 2 R_z rotations and 2 CNOTs. The total number of gates will be this multiplied by the number of spin species and the total number of bonds, which comes to $2(M - 1)$. The depth will be the number of operations to execute this circuit multiplied by the number of groups, 2.

Hopping - chain

Gate	Gate Count	Depth
Cliffords (H, Y)	$16(M - 1)$	8
$R_z(\theta_h)$	$4(M - 1)$	4
CNOT	$8(M - 1)$	8

Table 3.4: Resources required to implement the hopping term on qubits simulating a chain of M sites under the Hubbard model.

Single sweep total - chain

Now incorporating the gates required to execute the Coulomb repulsion and chemical potential terms of Tables 3.3 and 3.2 this results in the total gate counts and depths per sweep

Gate	Gate Count	Depth
Cliffords (H, Y)	$16(M - 1)$	8
$R_z(\theta)$ (all)	$7M - 4$	6
CNOT	$2(5M - 4)$	10

Table 3.5: Total gate counts and depth to execute one time-step-sweep for a chain of sites.

Full evolution example - chain

To get a rough estimate of how many gates this corresponds to we look at the depth estimates for $M = 100$ and assume that we want to evolve to time $t = 10/J$. From Table 3.1 we can see that will need as few as 32 sweeps if we use the 4th order decomposition to result in an error comparable to a 1% calibration error. In this lower bound we can estimate the number and depth of gates needed to simulate a chain We will need $N_Q = 200$ qubits.

Gate	Count per Sweep	Total Count	Depth per Sweep	Total Depth
Cliffords	1584	50.6×10^3	8	256
$R_z(\theta_p)$	200	6.4×10^3	1	32
$R_z(\theta_r)$	100	3.2×10^3	1	32
$R_z(\theta_h)$	396	12.7×10^3	4	128
$R_z(\theta)$ (all)	696	22.3×10^3	6	192
CNOT	992	31.7×10^3	10	320

Table 3.6: Resource estimation for the evolution of a Hubbard chain of length $M = 100$ up to time $t = 10/J$ using a 4th order Trotter decomposition with Trotter errors comparable to a 1% calibration error.

Ladder

For a ladder there are two obvious ways to order the qubits. A snake pattern of ordering proceeds from the top left corner of a lattice and starts moving right, moving down a row reversing direction when it reaches the end of a row. As such, we can either look at a vertical ladder ordering, which minimises the maximum length of any Jordan Wigner strings, or a horizontal ladder ordering, which maximises the parallelisation of implementing the hopping terms. This will become more clear as we proceed with the example. First we look at the case of the ‘vertical ladder’, the ordering of qubits and grouping of commuting hopping terms as depicted in Fig. 3.11.

The terms are grouped into four groups, which do not commute with one another but within which terms do commute, as they act only on separate groups of qubits. As in the chain case it is clear that hopping terms within the groups can be executed in parallel for the groups 1 and 2, as no Jordan Wigner strings are required and the qubits affected do not overlap between the groups. In fact, these groups form the same chain as we have just looked at, so we can recycle the counting of Table 3.4 and need only add the counts for groups 3 and 4.

The terms within groups 3 and 4 each have a Jordan Wigner string spanning the two qubits between the source and target. This means that the gate count for each of the terms in groups three and four have an additional 2 CNOT gates and 2 CZ gates. It should be clear that the terms within group 3 do not share qubits and that the same is true within group 4. If the chain has length $L = M/2$ then we can count the number of each kind of term by using

$$A_{\pm}(L) = \frac{L - 1 \pm (L + 1)\%2}{2}, \quad (3.11)$$

which tells us the number of each type of bond in a chain of length L , where the bond types alternate along the chain. Naturally $A_-(L) + A_+(L) = L - 1$. As an example,




Figure 3.11: Vertical ladder qubit ordering and hopping term groups for qubits representing one species of fermions on a ladder of length $L = M/2 = 6$ with $M = 12$ sites. The black circles represent the qubits which are indexed following the snake of arrows starting at the top left and ending at the bottom left. Hopping terms along bonds which share the same colour can be executed simultaneously as each acts on a separate set of qubits. Solid lines indicate hopping terms between neighbouring qubits, requiring no Z_{JW} string, whereas dashed lines will require a Z_{JW} string following the solid arrows to join source and target qubits to execute the hopping between them.

for the ladder of Fig. 3.11 with $L = 6$, we have $A_+ = 3$ bonds of type 3 and $A_- = 2$ bonds of type 4. This is not crucial for gate count as we will always have $M/2 - 1$ terms of combined type 3 and 4 and they have the same gate counts as each other, but is a useful quantity for implementing each of the groups and will later be useful for the case of the 2D lattice. Each bond of type 3 or 4 will contribute 8 Cliffords, 2 R_z rotations, 6 CNOTs and 2 CZs (which we will count as CNOTs). These values summed with the depth counts of groups 1 and 2 give us the depth. Multiplied by the number of bonds of type 3 and 4 they can be added in to the gate counts.

Hopping - vertical ladder

Gate	Gate Count	Depth
Cliffords (H, Y)	$8(3M - 4)$	16
$R_z(\theta_h)$	$2(3M - 4)$	8
CNOT	$8(2M - 3)$	24

Table 3.7: Resources required to implement the hopping term on qubits simulating a ladder of length $L = M/2$ under the Hubbard model given site ordering which minimises the length of the Jordan Wigner strings, depicted in Fig. 3.11.

Now we look at the case of the ‘horizontal ladder’ where the qubits are indexed differently and require only three groups of hopping terms as shown in Fig. 3.12. Although this increases the length of the Jordan Wigner strings it also increases the option to reuse the calculations of parity and improves parallelisation, which will be a useful tool to implement when it is not possible to use only short Jordan Wigner strings, as in the 2D lattice.




Figure 3.12: Horizontal ladder with the according qubit ordering. The different colours show the different groups of hopping terms which can be executed in parallel. Solid (dashed) lines represent hopping terms between neighbouring (distant) qubits.

This way of indexing makes it less obvious as to why terms within group 3 can be executed in parallel. The reason that this is possible is because of the fact that executing a hopping term between qubits k and l does not change the parity of these qubits or any between them and so a Jordan Wigner string which crosses both of these qubits will be unchanged by whether or not the hopping has been executed and commutes with it. This can be conceptually understood as the Jordan Wigner string counting the number of excitations between j and m , where $j < k$ and $l < m$, and recognising that this number will remain unchanged by an excitation hopping between k and l . This recycling of the calculated parity was described in the previous section (3.3) and requires the addition of one ancilla qubit to optimise the gate count, where one per species will also optimise circuit depth. This can be visualised using Fig. 3.12, where the process is implemented by calculating the number of excitations between sites $i = L - 1$ and $i = L + 2$ and reusing this for the $i = L - 2$ to $i = L + 3$, which will in turn be reused to implement the $i = L - 3$ to $i = L + 4$ hopping and so on.

If, for group 3, we execute the circuit as described in Fig. 3.8 we use only one extra

qubit to store the parity and then reuse it for each bond in the group. This is executed in parallel for each spin type, requiring two ancillas in total. The chain of group 1 and 2 bonds, as before, are implemented using the circuit depicted in Fig. 3.4. For each bond of type 3, 8 Cliffords, 2 R_z rotations, 8 CNOTs, and 2 CZ gates are needed. There are $L - 1$ of these types of bond for each spin type. The parity calculation string of CNOTs must be implemented sequentially but the rotation sequence part of the hopping term can be executed simultaneously on all bonds of type 3. This includes 4 of the CNOTs, so the gate depth is reduced. This results in the gate count and depth in Table 3.8. Since we need 1 ancilla qubit per spin species to store the parity on this means we need $N_Q = 2(M + 1)$ for the simulation.

Hopping - horizontal ladder

Gate	Gate Count	Depth
Cliffords (H, Y)	$8(3M - 4)$	12
$R_z(\theta_h)$	$2(3M - 4)$	6
CNOT	$2(9M - 14)$	$3(M + 2)$

Table 3.8: Resources required to implement the hopping term on qubits simulating a ladder of length $L = M/2$ under the Hubbard model given the site ordering depicted in Figure 3.12 and using two additional ancilla qubits.

We can already see that implementing the hopping terms for the ladder in this way has a different gate count and depth to implementation using the vertical ladder, as calculated in Table 3.7. The gate counts for Clifford gates and R_z gates are the same for the two cases but the horizontal ladder has a smaller depth in these gates as it only has three groups of bonds which it needs to implement, compared to the four groups on the vertical ladder configuration. However, the number of CNOT gates needed is higher for the horizontal ladder as the longer Jordan Wigner strings are not completely compensated for by the trick of reusing the computed parity, especially as this is computed onto an additional ancilla qubit. This difference is more drastically seen in the depth of CNOTs, which in the case of the horizontal ladder has a dependence on M , such that for $M > 6$ the depth will also be greater for the horizontal ordering case and, crucially, that the depth in CNOT gates in the vertical ladder case is fixed at 24. As a result we conclude that minimising the length of the Jordan Wigner strings by using the vertical ladder ordering is the most advantageous. We move on to calculating the totals and making some estimates for the numbers needed for real simulations for the vertical organisation only.

Single sweep total - vertical ladder

Implementing the ladder as in Fig. 3.11 and including the contributions from the Coulomb repulsion and chemical potential (Tables 3.3 and 3.2) leads to the total gate counts and depths of Table 3.9. As no ancilla qubits are needed the total number of qubits needed is $N_Q = 2M$

Gate	Gate Count	Depth
Cliffords (H, Y)	$8(3M - 4)$	16
$R_z(\theta)$ (all)	$9M - 8$	10
CNOT	$3(6M - 8)$	26

Table 3.9: Total gate counts and depth to execute one time-step-sweep for a ladder, indexed as in Figure 3.11

Full evolution example - vertical ladder

For a ladder structure we can see we will need between 10 and 500 sweeps per t if we use a 4th order decomposition. We here estimate the full gate counts for the lower bound of 10 sweeps per t , giving us a total of 100 sweeps in order to reach $t = 10/J$. We use the gate counts based on the vertical ladder from Table 3.7 to minimise gate counts and keep the depth in CNOTs system size independent.

Gate	Count per Sweep	Total Count	Depth per Sweep	Total Depth
Clifford	2368	24.63×10^4	16	1664
$R_z(\theta)$ (all)	892	9.28×10^4	10	1040
CNOT	1776	18.47×10^4	26	2704

Table 3.10: Resource estimation for the evolution of a ladder of length $L = 50$ with $M = 100$ sites up to time $t = 10$ using a 4th order Trotter decomposition with Trotter errors comparable to 1% calibration error. $N_Q = 200$ qubits are used

2D Lattice

For an $L \times L$ 2D lattice with M sites we need $2M$ qubits to represent the system and to arrange them by spin and then by site. This is in order to limit the distance required for hopping. Here, $L = \sqrt{M}$ and the gates are applied using a snake pattern to join sites between which hopping is implemented as in Ref. [81]. An example of this qubit indexing and site grouping for a system with $L = 5$ is shown in Fig. 3.13.

The terms are grouped into four groups which do not commute with one another. Again executing groups 1 and 2 is the same process as for the chain and so we reuse the calculations of Section 3.3. In total there are $2(M - 1)$ of these terms. Groups 3 and 4 we




Figure 3.13: Lattice qubit ordering and hopping term groups for a qubits representing one spin species on a 2D lattice with $N = 25$ sites. Black circles represent the qubits, hopping terms which can be executed in parallel are grouped by colour and solid (dashed) lines represent hopping terms between neighbouring (distant) qubits which need to be implemented.

treat as we did the horizontal ladder, using one ancilla qubit per row (per spin species) onto which we compute the cumulative Jordan Wigner strings and implementing the rotation on the corresponding qubit pair one bond at a time, working from right to left for group 3 and from left to right for group 4. The different rows within group 3 [4] can be executed simultaneously as there is no overlap between the qubits involved. There are $L - 1$ bonds of type 3 [4] per row and $A_+(L)$ [$A_-(L)$] of such rows in total.

Hopping - lattice

The gate counts and depths to implement the hopping for the 2D lattice are shown in Table 3.11. We should note that, consistent with the explanation about minimising the number of CNOT gates required, there will be an extra $2A_+(\sqrt{N})$ ancilla qubits, as these will be needed to simultaneously implement the hopping of the terms in group 3, with one per row for each spin species. They can be reused to implement group 4, which has the same number of rows, or one fewer, participating. This means that in total we need $N_Q = 2(M + A_+(L))$ qubits.

3.3 Fermi-Hubbard Model

Gate	Gate Count	Depth
Cliffords (H, Y)	$32(M - \sqrt{M})$	16
$R_z(\theta_h)$	$8(M - \sqrt{M})$	8
CNOT	$4(7M - 10\sqrt{M} + 3)$	$4(3\sqrt{M} + 1)$

Table 3.11: Resources required to implement the hopping term on qubits simulating a lattice of size $L \times L$ with total number of sites $M = L^2$ under the Hubbard model given the site ordering depicted in Fig. 3.12.

Single sweep total - lattice

Incorporating the Coulomb repulsion gates of Table 3.3 and the chemical potential contribution of Table 3.2 we come to the full gate counts and depths required to execute one time-step-sweep in Table 3.13.

Gate	Gate Count	Depth
Cliffords (H, Y)	$32(M - \sqrt{M})$	16
$R_z(\theta_p)$	$2M$	1
$R_z(\theta_r)$	M	1
$R_z(\theta_h)$	$8(M - \sqrt{M})$	8
$R_z(\theta)$ (all)	$11M - 8\sqrt{M}$	10
CNOT	$2(15M - 20\sqrt{M} + 6)$	$6(2\sqrt{M} + 1)$

Table 3.12: Total gate counts and depth to execute one time-step-sweep for an $L \times L$ lattice with M sites.

Full evolution example - lattice

If we assume the same error estimation, and hence number of sweeps per time-step, as the ladder and estimate the number of gates needed to implement the 2D lattice with $L = 10$ and $M = 100$ sites up to $t = 10J$ using a 4th order Trotter decomposition, this results in the total gate counts and depths in Table 3.13. We will need $N_Q = 210$ qubits to run this simulation.

Gate	Count per Sweep	Total Count	Depth per Sweep	Total Depth
Clifford	2880	29.95×10^4	16	1664
$R_z(\theta_p)$	200	2.08×10^4	1	104
$R_z(\theta_r)$	100	1.04×10^4	1	104
$R_z(\theta_h)$	720	7.49×10^4	8	832
$R_z(\theta)$ (all)	1020	10.61×10^4	10	1040
CNOT	2612	27.15×10^4	126	13104

Table 3.13: Resource estimation for the evolution of a 2D $L \times L$ lattice of $M = 100$ sites with $L = 10$ up to time $t = 10J$ using 4th order Trotter decomposition with Trotter errors comparable to 1% calibration error.

Code in the Q# quantum programming language [82] implementing evolution under the Fermi-Hubbard model for each of the different structures mentioned here can be found at <https://github.com/nataliejpg/QuantumResourceEstimation> and in Appendix B. The evolution for small systems can be checked against exact diagonalisation and the resources required for larger systems can be found for a single time-step-sweep or the full evolution under 1st order decomposition. This resource means that it is easy to check the results and optimise further or expand on the structures or implementations available.

Fault Tolerance

Although these numbers are a not many orders of magnitude beyond the reach of current gate fidelities it clear that with $\sim 10^5$ CNOT gates and R_z gates the precision needed to execute this number of gates accurately requires errors is not within the current scope of noisy intermediate-scale quantum (NISQ) technology. It would require two-qubit gates and rotations with fidelities $> 99.999\%$. Currently it is still a challenge to realise two-qubit gates with fidelities above 99% with ion trap experiments achieving $\leq 99.92\%$ [46, 47] and superconducting qubits yet to attain that third 9 and remaining below 99.9% [48–50].

Another possibility is to use a fault tolerant scheme that has some redundancy so that physical qubits with lower fidelities can be combined to make up logical qubits which are not as susceptible to error. Stabilizer codes use an overhead in qubits to realise almost noise-free quantum computation, despite relying on noisy components [54, 55]. However, stabilized operations usually include the Clifford gate set which alone is not sufficient to realise universal quantum computation. The addition of the T gate to the gate set realises universal computation but its addition is not trivial as it requires large overhead [56]. However, a low-error T gate can be achieved by a method called

magic state distillation which requires the use of many ancillas on which noisy T gates are applied and a distilling process which results in a single low-noise logical ancilla in $T|+\rangle$, thus introducing the T gate to our logical computational space [83]

In order to construct the R_z gates in this scheme many T gates are used. If we want to achieve an error ϵ in the spectral norm after applying N_R of these gates then, as in Ref. [84], the number of T gates required is:

$$N_T \sim 1.15 \log_2\left(\frac{N_R}{\epsilon}\right) + 9.2. \quad (3.12)$$

For $\epsilon = 0.01$ this implies that using the total number of rotation gates calculated above, $N_R = 1020$, we need 40 T gates per rotation, giving a total of 4.24×10^6 .

Some recent results take advantage of the fact that many of the rotations are identical and executed in parallel [80]. They show that if we want to apply N_R identical $R_z(\theta)$ rotations in parallel, the cost is

$$4N_R + \log_2(N_R)(1.15 \log_2\left(\frac{1}{E} \log_2(N_R)\right) + 9.2), \quad (3.13)$$

where E is the error per rotation required. Assuming that we can tolerate an error of 1% we have $E = \frac{0.01}{1020} = 9.8 \times 10^{-6}$. This then gives us the number of T gates we need to execute the parallelised circuit.

Gate	Parallel Count (N_R)	T Gates per $R_z(\theta)$	Total T Gates
$R_z(\theta_p)$	200	6	12.48×10^4
$R_z(\theta_r)$	100	7	7.28×10^4
$R_z(\theta_h)$	90	7	52.42×10^4
			72.18×10^4

Table 3.14: T gates needed to execute the rotations necessary for the 2D lattice model simulation for $M = 100$ to $t = 10J$ as in Table 3.13

It should be noted that this parallel rotation implementation also requires $O(\log_2(N_R))$ extra qubits, where N_R is the number of rotations to be applied in parallel.

Summary

In order to implement the Hubbard model on a 2D $L \times L$ lattice to produce errors comparable to calibration errors of 0.1-1% we need a fault tolerant quantum computer with gate count estimates in Table 3.15 or 2 qubit gate errors $< 10^{-5}$. This gives an overview of the CNOT and R_z gates needed as these are likely to be the limiting gates to

Trotter Order	Error (%)		$L = 8$	$L = 10$
		M	64	100
		qubits	132	210
		tJ	8	10
2	1	τ	0.0458	
		sweeps	350	437
		CNOT	5.64×10^5	1.14×10^6
	$R_z(\theta)$	2.24×10^5	4.46×10^5	
	0.1	τ	0.0312	
		sweeps	513	642
CNOT		8.27×10^5	1.68×10^6	
4	1	τ	0.962	
		sweeps	84	104
		CNOT	1.35×10^5	2.72×10^5
	$R_z(\theta)$	5.37×10^4	1.06×10^5	
	0.1	τ	0.805	
		sweeps	100	125
CNOT		1.61×10^5	3.27×10^5	
$R_z(\theta)$	6.40×10^4	1.28×10^5		

Table 3.15: Summary table for the 2D lattice digital Hubbard simulation for an $L \times L$ lattice to achieve comparable errors on local observables to a 1% or 0.1% calibration error in an analogue simulation. The required Trotter time-step, τ , is given along with the number of sweeps required to reach $tJ = L = 10$ or $tJ = L = 8$ for Trotter decomposition of order 2 and 4.

implement. This Table is intended to give slightly more broad estimates of the numbers of these gates required dependent on the system size, the total evolution time, and the level of error the simulation must achieve.

In order to implement the R_z rotations of Table 3.15 in a fault tolerant scheme we need to use T gates which are expensive. If many of the rotations are the same (as is the case when $t_{i,j} = J$ for all nearest-neighbour hopping terms and when U and μ are the same for all sites) we can manufacture these identical rotations more efficiently where they occur in parallel at the expense of $\mathcal{O}(\log_2(N_R))$ extra qubits. I have labelled these two cases as ‘Different’ and ‘Same’ to describe whether this reflects an estimate for the case where we expect terms to be different or the same per site. This also shows the improvement that comes from using the method in Ref. [80]. The results are shown in Table 3.16.

These numbers are encouragingly only an order of magnitude or so beyond the reach of current fidelities. However, it is clear that in the near term analogue computers are a

Trotter Order	Error (%)	Different		Same	
		$L = 8$	$L = 10$	$L = 8$	$L = 10$
2	1	8.35×10^6	1.71×10^7	1.68×10^6	3.03×10^6
	0.1	1.37×10^7	2.81×10^7	2.50×10^6	4.46×10^6
4	1	1.87×10^6	3.82×10^6	4.03×10^5	7.22×10^5
	0.1	2.50×10^6	5.12×10^6	4.86×10^5	8.68×10^5

Table 3.16: Number of T gates necessary to implement the R_z rotations for evolution of a 2D $L \times L$ lattice under the Hubbard model to time $tJ = L$ for $L = 8$ and 10 using different order Trotter decompositions. Gate estimates are to execute simulation where the accuracy of a local observable must be comparable that due to calibration error 1% or 0.1% on an analogue simulator and where we have required 1% and 0.1% error introduced to the spectral by the rotation synthesis for these cases respectively.

more immediate solution to the problems of simulating these systems, despite their lack of flexibility.

3.4 Long-Range Ising Model

The Ising model is another extremely well studied model, which can be solved exactly in the case of only nearest-neighbour coupling [14, 85]. However, the long-range Ising model is not and so it is an interesting hallmark model to use in our comparison. The general Ising model is defined is

$$H = - \sum_{i,j>i} J_{ij} \sigma_z^i \sigma_z^j - \sum_i h_i \sigma_z^i - \sum_i g_i \sigma_x^i, \quad (3.14)$$

where i and j are integers up to the number of sites, M . J_{ij} is the Ising coupling strength, h_i the longitudinal field on site i , and g_i the transverse field on that site. σ_α^i are the Pauli matrices acting on site i . The model is short-range if J_{ij} is only non zero for nearby sites i and j but here we look at the long-range model as it is significantly more challenging to simulate classically and hence a more competitive field for analogue and digital quantum computing.

Again we compare the errors on the off-diagonal correlations, introduced due to imperfect calibration, to those introduced due to Trotter decomposition. We then establish the Trotter step size needed to execute a comparable digital quantum simulation to the analogue simulation. This is shown in Fig. 3.14 for a 4th order Trotter decomposition. We begin in a product state with all spins aligned along the x axis and evolve under the Hamiltonian of Eq. (3.14) with $J_{ij} = \frac{J}{4|i-j|^2} = \frac{1}{4|i-j|^2}$, $h_i = 0$ and $g_i = \frac{B}{2} = \frac{1}{2}$ and consider analogue calibration errors in J and B .




Figure 3.14: Quantitative comparison of the Analogue and Digital approaches for quantum simulation of the Ising model. The time dependence of the calibration errors (solid blue 1%, solid red 0.1%) in the off-diagonal correlation functions and trotter decomposition errors. For the digital we have plotted the 4th order decomposition, for the time-steps $J\tau = 4$ (crosses), $J\tau = 2$ (stars), $J\tau = 1$ (plus sign) and $J\tau = 1/2$ (crosses) for system size, $M = 20$. Calculated by Stuart Flannigan

From Figure 3.14 we can see that, for $J\tau = 2$, we are below the errors equivalent to those introduced by 1% calibration errors in the analogue simulator and for $J\tau = 1$ we are below those equivalent to 0.1% calibration errors. Consequently gate estimates based on these time-steps for the 4th order decomposition will perform competitively with an analogue quantum simulator at this task.

Mapping to the digital circuit

The long-range Ising model is more simply mapped to qubits and gates than the Fermi-Hubbard model as the native system is much closer to that we wish to simulate. Each site maps to one qubit so we will need $N_Q = M$ to simulate the system. The single site terms implementing the transverse and longitudinal fields are simply rotations on the corresponding qubit around that axis. The coupling terms are also simple as they are two-qubit interactions only, but the extent to which they can be executed in parallel is slightly more challenging. Evolution for time-step τ for each of the terms is as follows.


Longitudinal term

This is a rotation of qubit i around the z axis by angle $\theta_i = -2h_i\tau$.


$$i \text{ --- } \boxed{R_z(\theta_i)} \text{ ---}$$

Transverse term

This is a rotation of qubit i around the x axis by angle $\theta_t = -2g_i\tau$.




Alternatively, if we can only rotate around the z axis, sandwiching a rotation around z between two basis changing Hadamard gates achieves the same effect:



Coupling term

The coupling term between qubits i and j can be implemented as shown below, where $\theta_c = -2J_{ij}\tau$



Gate counts

The gate counts for the long-range Ising model are easy to calculate as the terms are simple to apply and the single qubit rotations within the longitudinal and transverse terms can be applied simultaneously. The only challenging part is in calculating the depth of the part of the circuit implementing the coupling terms as many of these can be executed simultaneously

Longitudinal term

Gate	Gate Count	Depth
$R_z(\theta_l)$	M	1

Table 3.17: Resources for implementation of longitudinal term of the long-range Ising model

Transverse term

If we assume that R_z rotations are possible then this terms contributes estimates from Table 3.18.

Gate	Gate Count	Depth
$R_x(\theta_t)$	M	1

Table 3.18: Resources for implementation of transverse term of the long-range Ising model

If, by contrast, we can only execute R_z rotations, then the gate count will be as in Table 3.19.

Gate	Gate Count	Depth
Cliffords (H)	$2M$	2
$R_x(\theta_t)$	M	1

Table 3.19: Resources for implementation of transverse term of the long-range Ising model with limited rotational gate axes

Coupling term

The coupling term is slightly more interesting because, although the number of gates is fixed, the depth of the circuit can be minimised by applying as many terms which do not share qubits as possible simultaneously. For all-to-all coupling this can be achieved with gate depths shown in Table 3.20. This gate depth can of course be improved upon if we have fewer than the full $M(M//2)$ bonds between qubits, where $//$ is integer division.

Gate	Gate Count	Depth
$R_z(\theta_c)$	$M(M//2)$	M
CNOT	$2M(M//2)$	$2M$

Table 3.20: Resources for implementation of coupling term of the long-range Ising model

Single sweep total

The sum of these tells us the number and depth of gates required to execute one time-step-sweep where we have here assumed that we do not have access to R_x rotations. If this is not the case then no Cliffords are needed and one should replace the $R_z(\theta_c)$ in the table with an $R_x(\theta_t)$ with the same count and depth.

3.4 Long-Range Ising Model

Gate	Gate Count	Depth
Cliffords (H)	2N	2
$R_z(\theta_l)$	N	1
$R_z(\theta_t)$	N	1
$R_z(\theta_c)$	$N(N//2)$	N
$R_z(\theta)$ (all)	$2N + N(N//2)$	$N + 2$
CNOT	$2N(N//2)$	2N

Table 3.21: Total resources required to implement single time-step-sweep of the long-range Ising model

Full evolution example

To get a rough estimate of how many gates this corresponds to we look at the depth estimates for $M = 100$ and assume that we want to evolve to time $t = 10/J$. If we use the 4th order decomposition, with $J\tau = 2$, then we will need 10 sweeps per time-step and a total of 100 sweeps. No ancilla qubits are necessary so we require only $N_Q = 100$ qubits.

Gate	Count per Sweep	Total Count	Depth per Sweep	Total Depth
Clifford	200	10×10^3	2	100
$R_z(\theta_l)$	100	5×10^3	1	50
$R_z(\theta_t)$	100	5×10^3	1	50
$R_z(\theta_c)$	5000	250×10^3	100	5000
$R_z(\theta)$ (all)	5200	260×10^3	102	5100
CNOT	10000	500×10^3	200	10000

Table 3.22: Resource estimation for the evolution of the long-range Ising model with $M = 100$ sites up to time $t = 10/J$ using a 4th order Trotter decomposition with Trotter errors comparable to a 1% calibration error.

Fault Tolerance

Again, the numbers of CNOT gates required is high enough to exceed current limits of the hardware, which motivates a fault tolerant approach. We use the same calculation as Eq. (3.12) to calculate the number of T gates needed to implement the R_z rotation gates in a fault tolerant system for $\epsilon = 0.01$. Here we have $N_R = 5200$ R_z gates per sweep, which corresponds to 38 T gates per rotation and a total of 9.88×10^6 . In this case, although the longitudinal and transverse term evolutions will result in many identical rotations being executed in parallel, the same cannot be said for the coupling term due to the order of execution and the fact that the couplings are not all equal. Consequently we will only be able to use the savings possible by using Eq. (3.13) from Ref. [80] for the transverse and longitudinal rotations. This means that the number of

T gates per $R_z(\theta_c)$ is 38 and so a total of 9.37×10^6 of these are needed for the full evolution. The error per rotation for the remaining gates is then $E = \frac{0.01}{200} = 5 \times 10^{-5}$ which gives us the number of T gates need to execute the parallelised circuit in Table 3.23.

Gate	Parallel Count (N_R)	T Gates per $R_z(\theta)$	Total T Gates
$R_z(\theta_l)$	100	6	3×10^4
$R_z(\theta_t)$	100	6	3×10^4
$R_z(\theta_c)$	1	38	950×10^4
			9.56×10^6

Table 3.23: T gates needed to execute the rotations necessary for the long-range Ising model simulation for $M = 100$ to $t = 10J$ as in Table 3.22

It is clear that this is only a moderate improvement as the biggest contribution to the rotation gates is the coupling terms, which we cannot optimise in this way and so it is unlikely to be worth the trade-off for the extra ancilla qubits necessary to implement this.

Again, the code implementing evolution under the long-range Ising model is available at <https://github.com/nataliejpg/QuantumResourceEstimation> and in Appendix B. For small systems a comparison to an exact diagonalisation solution is possible but more importantly resource estimation is included. The provision of this code makes it simple to experiment with implement different methods for applying the terms to minimise error or optimise gate counts further.

Summary

In order to implement the long-range Ising model of size M to produce errors comparable to calibration errors of 0.1-1% we need a fault tolerant quantum computer with gate count estimates in Table 3.24. This gives an overview of the CNOT and R_z gates needed as these are likely to be the limiting gates to implement.

The number of T gates required to implement the R_z gates of Table 3.24 is shown in Table 3.25, where we have not made the assumption that and gates can be done in parallel.

3.5 Conclusion and Outlook

Error (%)		$M = 64$	$M = 100$
	qubits	64	100
	tJ	8	10
1	τ	2	
	sweeps	40	50
	CNOT	1.64×10^5	5×10^5
	$R_z(\theta)$	8.7×10^4	2.6×10^5
0.1	τ	1	
	sweeps	80	100
	CNOT	3.28×10^5	1×10^6
	$R_z(\theta)$	1.74×10^5	5.2×10^5

Table 3.24: Summary table for the long-range Ising simulation for an $L \times L$ lattice to achieve comparable errors on local observables to a 1% or 0.1% calibration error in an analogue simulation. The required Trotter time-step, τ , is given along with the number of sweeps required to reach $tJ = \sqrt{M} = 10$ and $tJ = \sqrt{M} = 8$ for Trotter decomposition of order 4.

Error (%)	$M = 64$	$M = 100$
1	3.13×10^6	9.88×10^6
0.1	7.14×10^6	2.24×10^7

Table 3.25: Number of T gates necessary to implement the R_z rotations for evolution of the long-range Ising model to time $tJ = \sqrt{M}$ for $tJ = 8$ and 10 using different order Trotter decompositions.

3.5 Conclusion and Outlook

We have analysed the minimum gate counts required for digital quantum simulation to compete with current analogue simulators beyond the range of classical computation. We have investigated the results using two well known problems of interest in physics and materials science. We chose to investigate these at the quantum advantage point, where we go beyond dynamics that we can calculate classically. For analogue quantum simulators, the calibration level of current experiments is already sufficient to be in a regime of accurate simulations, obtaining observable values with errors bounded at the 1% percent level, for 2D systems where the same accuracy is not possible for any known classical algorithm. Note also that this comparison was carried out at the quantum advantage point, and as analogue systems are scaled up, even more digital logic gates would be required to match the accuracy that can be achieved in the analogue experiments. It is also worthwhile to observe that the maximum fidelity gates measured in digital quantum computers are often of individual qubits or at least in systems with N_Q much smaller than 100 as we discuss here so these high fidelities must also be achieved scalably. This demonstrates that analogue simulators are the best means to accurately compute observables in those classes of systems that can be realised (at this calibration

level) on analogue experiments, and confirms that existing experiments are already past the requirements for the quantum advantage point.

On the side of digital quantum simulation, we compare our findings to Ref. [86] which also estimated the number of gates needed to implement a simulation with quantum speed up. They focussed on the Heisenberg model in 1D and found that order 10^7 CNOT gates and 10^9 T gates would be needed to simulate a system of size 100. Our calculations find an improvement of 2 orders of magnitude on the number of CNOT gates needed for ‘quantum speed up’ and 3 orders of magnitude in the number of T gates. Since the models used are different there cannot be a direct comparison but in general this improvement arises from a combination of factors. One of the factors is that we use a 2D lattice instead of a 1D chain which means that shorter evolution times are needed to exceed the Lieb Robinson bounds as information can propagate in 2 dimensions, thus reaching all sites on the system faster than in 1D. Another factor is that in Ref. [86] they require 0.1% error on the full wavefunction rather than error on the values of the observables of interest as we do here. The final contributor to the lower gate count estimates can be attributed to gate optimisation schemes from Refs. [80, 81].

Although our digital estimates have assumed that all error comes from the Trotterisation further work could take into account the errors on single and two-qubit gates in digital quantum computers to further explore the feasibility of implementing these algorithms on NISQ machines. There may be ways to trade off errors in gates on NISQ machines with the Trotter errors - lowering the Trotter error by using more gates, and allowing some of the error budget to come from the gates themselves. It is additionally interesting to explore the optimisation of rotation gate generation in fault tolerant schemes as these will be useful beyond NISQ and bring us closer to being able to implement meaningful and useful computations on a quantum computer. It would also be meaningful to investigate the overhead introduced by removing the assumption of all-to-all coupling which we have imposed in our gate counts as this is likely to strongly impact the number of gates required if distant coupling of physical qubits is expensive. We look forward to further research in these directions as we move towards the realisation of universal quantum simulators.

Chapter 4

Trotter Error Scaling

The resources required to simulate a quantum many-body system scale with the number of time-steps used, motivating the use of as few as possible to achieve the required level of accuracy. Unfortunately, simulations of quantum many-body systems using Trotter-Suzuki product formulae require a time-step with size that decreases with the lattice size to produce constant accuracy results. However, the behaviour of the errors introduced on observables as we increase system size are not necessarily as pessimistic as the upper bounds suggest. In this chapter we will systematically study the impact of increasing the system size on the errors introduced to local and non-local observables in order to more rigorously explore the behaviour of these errors and ascertain the necessity of decreasing the time-step size.

Generally, the Trotter error is seen as an error incurred during the numerical evolution of the system, stemming from an imperfect approximation of the time-evolution propagator. In contrast, here we suggest to interpret the systematic errors caused by product formulae as a constant local perturbation to the system's Hamiltonian, instead of an error on the time-evolution operator. This shift of focus has a number of important consequences: the Trotter step size is fixed by the model calibration accuracy and is independent of the system size, the error on expectation values of time-dependent local observables is independent of the system size, while errors on imaginary time expectation values can depend on the system size only indirectly via the spectral gap. In particular the calibration error introduced can have concrete implications for the critical point of the system. This chapter largely follows Refs. [2] and [3] which is work done by Matthias Troyer, David Poulin and myself where I executed all simulations and David was responsible for the majority of the theoretical insight.

4.1 Introduction

The Trotter-Suzuki product decomposition [62, 63] is conceptually the simplest method to simulate a quantum many-body system using either widely used quantum Monte Carlo methods (QMC) [27–34] or quantum computer algorithms [86–92], and in the latter case this simplicity has enabled experimental demonstrations [93–95]. However, it results in a systematic error, ϵ , that can only be suppressed at the expense of considerably increasing the duration of the simulation by a factor $1/\epsilon^k$ for some constant k that depends on the details of the scheme [96, 97]. This conclusion is inevitable when we demand that the simulation produces a transformation which is close in operator norm to the (real or imaginary) time-evolution operator of the model quantum many-body system. This requirement is well justified mathematically, since the time-evolution operator contains all the dynamical and spectral information about the system. However, there is a sense in which the simulation should reproduce the right physics with a constant Trotter step even in the thermodynamic limit.

This line of inquiry is primarily motivated by the cost analysis of quantum simulation algorithms. To realize a constant-accuracy approximation to the time-evolution operator, the number of time-steps must increase with the system size [96, 97], c.f. Eq. (4.2). While we do not dispute this conclusion here, we argue that it sets the wrong target and that accurate physical information about the system can be obtained without incurring this additional simulation overhead.

We use two arguments to justify this change of target. The first one follows rather trivially from previous arguments [98, 99] but deserves reiterating because it applies to many physical settings, and simply shows that for a local time-evolved observable $O(t)$ and a fixed time t , the Trotter error is independent of the system size. This is a direct consequence of the Lieb-Robinson bound which sets a finite sound velocity in lattice models [72, 100].

The second argument is much more general and is based on the idea of attributing the Trotter error to the Hamiltonian rather than the time-evolution operator: product formulae produce an *exact* time-evolution operators corresponding to a slightly *perturbed* Hamiltonians. As we will see, this perturbation is physically well behaved in the sense that it is system size independent and exponentially decaying. From this point of view, the Trotter time-step should be chosen as a function of the model’s accuracy rather than the system size, as we now argue.

In situations where the simulation is used to obtain quantitative predictions of a material, the model is never exact, and comes with intrinsic calibration errors. These calibration errors stem from many sources. Often, they are made to simplify the model – limiting the Coulomb interaction to a few nearest-neighbours in the Hubbard model

is a widespread example in condensed-matter physics. Calibration errors are also made to reduce the dimension of the model's Hilbert space. A single-band Hubbard model illustrates this point [73]. More broadly, some quantum degrees of freedom are often integrated out and replaced by effective classical interactions to reduce the effective Hilbert space dimension – for instance the Coulomb interaction is an effective classical description of light-matter quantum interactions, and the nuclei location in a molecule are described classically in the Born-Oppenheimer approximation [101]. The computational intractability of *ab initio* calculations imposes the use of approximate methods like density functional theory that also contribute to calibration errors [102]. In Chapter 3 we saw the effect of calibration errors in the model which we could physically implement in simulation using an analogue quantum simulator. Lastly, the accuracy of a model's parameters are subject to the fundamental limits of quantum metrology [103], so in all cases the model only approximately describes a real system.

4.2 Background

The key message is that error resulting from the Trotter-Suzuki decomposition should be regarded as an additional source of calibration error, so the Trotter step should not depend on the system size. This simple observation has far-reaching consequences on quantum simulations, where it lowers the complexity by some power of the system size (the power depends on the order of the product formula). Under the eigenstate thermalisation hypothesis [104], it also shows that long-time local expectation values incur an error that is independent of the system size, which subsumes the conclusions of Ref. [105].

We consider a lattice system of M particles with some local model Hamiltonian

$$H = \sum_{\lambda \subset \Lambda} h_{\lambda}, \quad (4.1)$$

where $\Lambda \subset \mathbb{Z}^D$ denotes the set of vertices of some finite D -dimensional lattice and λ subsets thereof, interactions have bounded norm $\|h_{\lambda}\| < J$, and locality means that $h_{\lambda} = 0$ when the radius of the smallest ball on the lattice containing λ is greater than some constant r . The time-evolution operator is $U(t) = \exp\{-iHt\}$, where in general t can be complex.

The Trotter-Suzuki approximation requires grouping the different λ 's into sets k such that all $[h_{\lambda}, h_{\lambda'}] = 0$ for λ, λ' in the same k . Due to locality of the Hamiltonian, it is always possible to find a constant number K of such sets k , independent of the lattice size (but growing at most as r^D).

As discussed in Chapter 2 the second-order Trotter-Suzuki decomposition approximates $U(t)$ by a sequence of N_τ time-steps $[U_{\text{TS}}^{(2)}(t/N_\tau)]^{N_\tau}$ with $U_{\text{TS}}^{(2)}(\tau)$ given in Eq. (2.19). It can be shown that $\|U(\tau) - U_{\text{TS}}^{(2)}(\tau)\| \in \mathcal{O}(M\tau^3)$ for a lattice of size M , so the error for an evolution of time $t = N_\tau\tau$ is

$$\|U(t) - [U_{\text{TS}}^{(2)}(t/N_\tau)]^{N_\tau}\| \in \mathcal{O}(Mt\tau^2). \quad (4.2)$$

This bound suggests that we must set $N_\tau = \mathcal{O}(\sqrt{Mt^3/\epsilon})$ to achieve constant accuracy ϵ . Note that, as we increase the lattice size M , the bound Eq. (4.2) is only meaningful at short times, because the operator norm of the difference of two unitary operators such as $\|U - U_{\text{TS}}\|$ is bounded by 2. Thus, to see the saturation of Eq. (4.2), we must study the error at times t that scales with $1/\sqrt[3]{M}$.

While the above reasoning is correct, it is reached by demanding that the entire time-evolution operator is accurately reproduced. When the operator norm difference between U and U_{TS} is ϵ , it means that there exists an observable, O , whose time-dependent expectation value under the Trotter-Suzuki approximation will result in an error ϵ . In this sense the bound Eq. (4.2) is tight and physically meaningful because in principle there exists a physical measurement that will be affected by this error on a specific initial state. But there are two reasons why demanding an accurate time-evolution operator may be physically irrelevant. On the one hand, quantum many-body physics is typically concerned with *local* observables, or at best observables that involve only a few bodies (which could be far apart). On the other hand, time-evolution $U(t) = \exp\{-iHt\}$ is generated by a Hamiltonian, H , which encodes all the physics of the system, so it may be physically more relevant to think of the Trotter-Suzuki approximation in terms of the error on H instead of $U(t)$.

The bound Eq. (4.2) is tight at short times in general, so there exists an observable whose time-dependent expectation value will have an error proportional to M . But for a fixed observable – the same for all system sizes – the error on the expectation value cannot depend on the system size. This is a consequence of the Lieb-Robinson bound, which states that for an observable O_λ , supported on sites λ , the Heisenberg-picture observable $O_\lambda(t) := U^\dagger(t)O_\lambda U(t)$ acts (almost) trivially on all sites that are at a distance vt away from λ , where v is some velocity fixed by the microscopic details of the model but independent of the system size. Intuitively, this bound simply describes a finite sound velocity at which effects propagates in the system.

To apply this reasoning to the case of simulation errors, we express the Trotter-Suzuki time-evolution operator Eq. (2.19) as the *exact* time-evolution operator of some effective Hamiltonian

$$U_{\text{TS}}^{(2)}(\tau) =: e^{-iH_{\text{eff}}\tau}. \quad (4.3)$$

To leading order in τ , the effective Hamiltonian H_{eff} is equal to the model Hamiltonian H , but additional corrections contribute to an effective perturbation $V_{\text{eff}} := H_{\text{eff}} - H$.

The Baker–Campbell–Hausdorff (BCH) formula

$$e^A e^B = e^{A+B+\frac{1}{2}[A,B]+\frac{1}{12}([A,[A,B]]+[B,[B,A]])+\dots} \quad (4.4)$$

shows that V_{eff} can be expressed as a series of embedded commutators of the terms H_k entering the Trotter-Suzuki approximation Eq. (2.19). Using this formula recursively and keeping only the first contributions of order up to τ^3 results in

$$V_{\text{eff}} = \frac{\tau^2}{24} \sum_{k=1}^{K-1} [H_k, [H_k, H_{>k}]] - 2[H_{>k}, [H_{>k}, H_k]] + \mathcal{O}(\tau^3), \quad (4.5)$$

where $H_{>k} := \sum_{l>k} H_l$. Because each term H_k is the sum of terms h_λ acting on a ball of radius r , up to order τ^2 , V_{eff} is a sum of terms acting on a ball of radius $3r$.

This shows that the effective Hamiltonian H_{eff} is less local than the model Hamiltonian H . Not only does it include longer-range couplings, but the terms act on a growing number of sites. However these long-range many-body interactions are weak, since they are proportional to τ^3 .

4.3 Numerical Methods: Thermal State Properties

To find the thermal state properties of a system we often need to evaluate high dimensional integrals. In a quantum system the space grows exponentially with the number of particles, which makes sampling the system to the same accuracy increasingly challenging. Monte Carlo methods rely on randomly choosing points to approximate an integral:

$$\frac{1}{\Omega} \int f(x) dx \approx \frac{1}{N} \sum_i^N f(x_i), \quad (4.6)$$

where Ω is the integration volume and we have approximated the integral by choosing N points x_i . The error on this estimate is reduced by importance sampling where the points x_i are chosen from a probability distribution p as close to f as possible. This means that we sample more in the regions of higher ‘importance’.

The (pseudo) random numbers involved in choosing points can be found by following a Markov Chain. We start from an initial point x_0 and then move to the next point based on a transition probability W . We continue to generate new points in this way, such that we use a probability matrix W_{ij} to move from x_i to generate the next point x_j until N points have been generated. W_{ij} must be normalised and satisfy conditions of ergodicity and detailed balance in order to be able to traverse the full configuration space and asymptotically reach the desired probability distribution. This is described in more

detail in Ref. [34]. The simplest Monte Carlo algorithm is the Metropolis algorithm [27] where, starting from point x_i , a new point x' is proposed by choosing randomly from a fixed number of changes Δx , so that $x' = x_i + \Delta x$. This new point is accepted as the next point on the chain with probability $P = \min(1, \frac{p_{x'}}{p_{x_i}})$ which determines whether $x_{i+1} = x'$ or x_i . Ergodicity is achieved by making sure that the possible changes, Δx , allow all points in configuration space to be reached in a finite number of steps. Detailed balance is ensured by including, for each change, also its inverse, $-\Delta x$.

The Monte Carlo sampling techniques applied to many-body quantum systems allow us to find the thermal properties of the system. These properties are obtained from the density matrix of a thermal state

$$\rho(\beta) = e^{-\beta H}, \quad (4.7)$$

where to normalise we would need to divide by partition function

$$Z = \text{Tr} (e^{-\beta H}), \quad (4.8)$$

with β representing the inverse temperature. At this point it is useful to notice that $\rho(\beta)$ can be seen as the result of propagating some initial state at $\beta = 0$ in imaginary time under the imaginary time-evolution operator $U(-i\beta)$.

If we rewrite the density matrix in coordinate representation with matrix elements

$$\rho(R, R', \beta) = \langle R | e^{-\beta H} | R' \rangle, \quad (4.9)$$

then the partition function can be expressed as the integral of the diagonal matrix elements over all configurations

$$Z = \int \rho(R, R, \beta) dR. \quad (4.10)$$

We can apply the property that

$$e^{-\beta H} = (e^{-\tau H})^{N_\tau}, \quad (4.11)$$

where $\beta = N_\tau \tau$, to rewrite

$$\rho(R_1, R_{M+1}, \beta) = \int \cdots \int dR_2 dR_3 \cdots dR_M \rho(R_1, R_2, \tau) \rho(R_2, R_3, \tau) \cdots \rho(R_M, R_{M+1}, \tau). \quad (4.12)$$

If we have a small enough τ then we can often make a short time' approximation of $\rho(R_j, R_{j+1}, \tau)$, thus simplifying the problem to calculating a high dimensional integral

of a known function, which is exactly what our Monte Carlo algorithm is good at. The expression for an observable, O is

$$\langle O \rangle = \frac{1}{Z} \text{Tr} (O e^{-\beta H}) \quad (4.13)$$

and so it will also be possible to find the values of observables of the quantum system using the same techniques.

Application of the ‘short time’ approximation of the density matrix is the same as the process in Eq. (4.4), where we can most simplistically neglect all terms in H_{eff} beyond those linear in τ , which is exact in the limit that $N_\tau \rightarrow \infty$. However, the effective Hamiltonian H_{eff} itself depends on the the time-step size τ , and in particular on whether it is real or imaginary. The BCH formula enables us to express H_{eff} as a power series in τ , and the imaginary-time evolution operator $e^{-\tau H}$ is obtained from the real-time one $e^{-iH\tau}$ by a Wick rotation to imaginary time, substituting $\tau \rightarrow -i\tau$, so the κ 'th term in the BCH expansion is altered by phase $(-i)^\kappa$. We can then see that the process we have described is equivalent to evolving from infinite temperature to a finite β in steps of imaginary time. The Monte Carlo sampling methods applied are at this point called Quantum Monte Carlo (QMC) and are used to sample the path integrals efficiently in order to find the values of observables which can be obtained from the partition function.

It is worth noting that the simple method outlined above would not be sufficient to obtain results for our system without the aid of further techniques used to avoid critical slowing down of the algorithm as M is increased, by using cluster updates to move from x_i to x_{i+1} and retain ergodicity [31, 32]. Furthermore it is possible to go to the limit of $N_\tau \rightarrow \infty$ where, instead of storing configurations at all time-steps, we only store the times when the spins flip and where loop clusters and loop updates allow us to work with infinitesimal τ [28–30]. It is further worth noting that all terms in our integral are positive definite, where the introduction of the Fermi sign to simulate Fermi particles would result in an exponentially small signal to noise ratio, the so called ‘sign problem’, and hence require further methods to circumvent.

Care must be taken in a simulation based on a product formula that combines both imaginary-time-evolution to prepare an initial state and real-time-evolution to measure dynamical quantities, because the effective Hamiltonian interpretation of Trotter errors may demand a different effective Hamiltonian for real and imaginary time-evolution. It should be noted however that for quantum computer applications, the state preparation is often realized using real-time-evolution. This is because time-evolution is used jointly with quantum phase estimation [106] as a means of performing a projective energy measurement. This is the case for instance in the quantum Metropolis algorithm [107] which can be used to prepared a thermal state $\rho = \frac{1}{Z} e^{-\beta H}$, or more generally of adiabatic ground-state preparation [108, 109] or its discrete formulation [110] and in approaches

[111, 112] based on Grover’s search algorithm [113]. In any of these cases, the error caused by the Trotter-Suzuki approximation in both state preparation and dynamical evolution can be interpreted with the same effective Hamiltonian H_{eff} obtained from *real* time-evolution.

We now move to describing the model that we will use to give numerical examples which illustrate the effect of Trotter errors on quantum lattice systems. We also detail how the mapping to a classical system is achieved, which enables the use of QMC for different values of τ and system size. We use the transverse field Ising model (TFIM) as described in Eq. (3.14) in the case that there is no longitudinal field so $h_i = 0$ and that all longitudinal terms are the same so $g_i = g$. This is written:

$$H(g, L) = gH_x + H_z = -g \sum_{j=1}^L \sigma_x^j - \sum_{i<j} J_{ij} \sigma_z^i \sigma_z^j, \quad (4.14)$$

where L is the length of the ring of sites, with periodic boundary conditions such that $\sigma_\alpha^{L+i} = \sigma_\alpha^i$, and σ_α^i are Pauli matrices of type $\alpha \in x, y, z$ acting on the i -th spin. g is the strength of the transverse field and J_{ij} are coupling constants of the Ising interaction. The model can be solved exactly by Jordan-Wigner transformation when all $J_{ij} = 0$ except for the case of $j = i + 1$ [85, 114].

The partition function at inverse temperature β is given by

$$Z(\beta, g, L) = \text{Tr} (e^{-\beta H(g, L)}). \quad (4.15)$$

The problem is truly quantum mechanical because H_x and H_z don’t commute.

Decomposing this Hamiltonian using the Trotter-Suzuki decomposition we can define the partition function of the *transverse field Ising-Trotter-Suzuki model*

$$Z(\beta, g, L, \tau) = \text{Tr} ([e^{-\tau g H_x} e^{-\tau H_z}]^{N_\tau}). \quad (4.16)$$

It is well known that this is the same as the partition function of a classical Ising model with one extra “imaginary time” dimension, consisting of N_τ replicas of the system, Ising variables $s^{i,\tau}$ with $i = 1, \dots, L$ and $\tau = 1, \dots, N_\tau$, and the classical Hamiltonian

$$H^c = - \sum_{i<j} \sum_{\tau} J_{ij}^{\text{space}} s^{i,\tau} s^{j,\tau} - J^{\text{time}} \sum_i \sum_{\tau} s^{i,\tau} s^{i,\tau+1}, \quad (4.17)$$

where we assume periodic boundary conditions in imaginary time: $s^{i,\tau} = s^{i,\tau+N_\tau}$. The space-like and time-like couplings of this classical model are related to the couplings of the original quantum model by

$$\begin{aligned} \beta^c J_{ij}^{\text{space}} &= \beta J_{ij} / N_\tau = \tau J_{ij} \\ \beta^c J^{\text{time}} &= -\frac{1}{2} \tanh(\log(\tau g)) \end{aligned} \quad (4.18)$$




Figure 4.1: Coupling strengths of the 2D classical Ising model Eq. (4.17) required to simulate the 1D quantum Ising model Eq. (4.14). The solid lines correspond to J^{time} for different values of the quantum transverse field strength g and the dashed and dotted lines correspond to J_1^{space} and J_2^{space} respectively. If in the quantum model we choose to have no next-nearest-neighbour coupling so that $J_2 = 0$ we would of course have no corresponding J_2^{space} in the classical model.

where β^c is the inverse temperature of the classical system.

We will consider in particular the 1D case with

$$H_z = -J_1 \sum_{i=1}^L \sigma_z^i \sigma_z^{i+1} - J_2 \sum_{i=1}^L \sigma_z^i \sigma_z^{i+2}, \quad (4.19)$$

with nearest-neighbour and next-nearest-neighbour couplings J_1 and J_2 respectively. The values of the ‘space’ and ‘time’ classical Ising model couplings are shown in Fig. 4.1.

We can then use powerful Monte Carlo methods to sample this system and find properties of the ground state where the error analysis explored above is equally applicable to that of evolution under real time. Here too, for computational purposes, the Trotter-Suzuki formula is often used with a finite τ . In this case we map the quantum TFIM to a classical Ising model. In general this can produce systematic errors in physical predictions which we will explore further. We use the ALPS library for our simulations [66, 67] to find the ground state properties of the system and calculate the Binder ratio for a range of parameter values. To find the properties of the system with $\tau = 0$ we use the ‘looper’ code within ALPS to simulate a 1D quantum Ising model as described in Eq. (4.14) and Eq. (4.19), the results of which we can then compare to the finite τ parameter values. We use the ‘spinmc’ code within ALPS to simulate a 2D classical Ising

model as described by Eq. (4.17) and with couplings defined in Eq. (4.18) at $\beta^c = 1$, shown in Fig. 4.1. For small τ , not only do we need to increase the lattice sized used for simulation, but we also have a larger difference between the temporal and spatial couplings, which makes it harder for the system to thermally equilibrate as is necessary for our sampling techniques. As such we expect higher uncertainty for small but finite τ estimates. This is also true for large τ where again the difference between the couplings becomes large and the spatial coupling is much greater than the temporal one.

4.4 Criticality of the Transverse Field Ising-Trotter-Suzuki Model

The Trotter-Suzuki time-step τ defines a new parameter of the model, in addition to the coupling strength g of the transverse field Ising model. The transverse field Ising model is recovered when $\tau \rightarrow 0$. While in general the critical point depends on τ , we show that for the 1D nearest-neighbour model, the critical point is independent of τ . As a consequence, the critical point obtained from finite τ is exact. This observation is a consequence of the fact that the usual duality transformation of the classical two-dimensional Ising model is equivalent to the Kramers-Wannier duality transformation of the quantum model [115].

To accurately find the critical point of our model, we compute the Binder ratio, U , given by

$$U = 1 - \frac{\langle m^2 \rangle}{3\langle |m| \rangle^2}, \quad (4.20)$$

where

$$m = \sum_{i=1}^L \sum_{\tau=1}^m s^{i,\tau} \quad (4.21)$$

is the total magnetization of the classical Ising model derived from the TFIM. At the critical point U is independent of system size L . The crossing point of the curves of $U(g)$ for different L thus indicates the critical point as shown for both nearest-neighbour and next-nearest-neighbour models in Fig. 4.2. While in general, the critical value g_{cr} is a function of τ , we observe that for the nearest-neighbour model with $J_1 = 1$ and $J_2 = 0$, $g_{\text{cr}} = 1$ independent of τ .

While it is well known that the nearest-neighbour TFIM with $J_1 = 1$ is critical precisely at $g = 1$, our main result is that the corresponding free energy in the thermodynamic limit $L \rightarrow \infty$ and at zero temperature $\beta \rightarrow \infty$ is critical at $g = 1$ for *all values* of τ . In

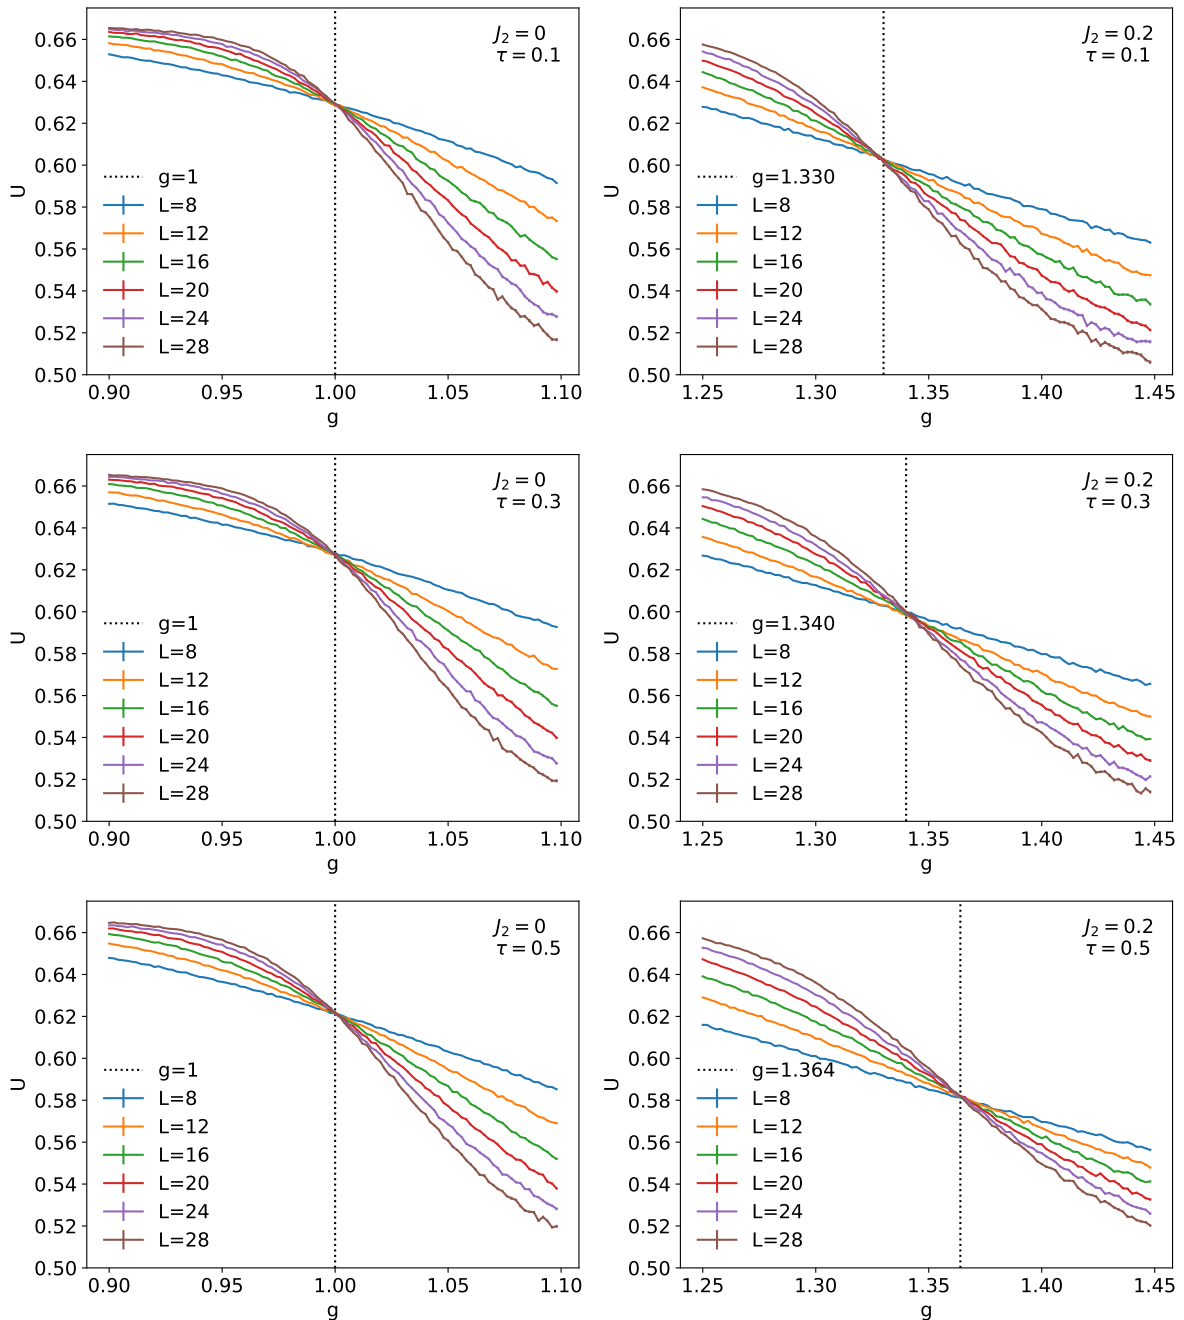


Figure 4.2: Binder cumulant values as a function of the transverse field strength g for Ising models derived from two different one-dimensional transverse field Ising models at different values of τ . The left column is the nearest-neighbour model with only a nearest-neighbour coupling $J_1 = 1$. The right column is for a model with an additional next-nearest-neighbour coupling $J_2 = 0.2$. Binder ratios are independent of the system size L at the critical point. The crossing point of the curves for various L thus indicates the location of the critical point. We observe that, while in general the critical point depends on τ , as can be seen in the bottom row, the nearest-neighbour model is unusual in that the location of the critical point does not depend on τ .

other words, τ parametrizes a critical line for the free energy

$$F_0(g, \tau) = - \lim_{\beta \rightarrow \infty} \lim_{L \rightarrow \infty} \log(Z(\beta, g, L, \tau)/\beta L) \quad (4.22)$$

in the (g, τ) plane. So, while a finite τ generally produces inaccurate physical predictions, it does not affect the location of the critical point nor the critical exponents in the case of the nearest-neighbour 1D TFIM.




Figure 4.3: The critical point for nearest-neighbour (a) and next-nearest-neighbour (b) coupling as a function of τ , showing how the location of the critical point is independent of Trotter step size for $J_2 = 0$ and increases with step size for $J_2 = 0.2$

The difference to the case of the model with $J_2 > 0$ can be seen clearly in Fig. 4.3 which shows the dependence of the critical point of the longer range non-integrable system on τ and the independence of the critical point of the nearest-neighbour system.

The reason for the independence of the nearest-neighbour system to τ is that, in the bulk, the Kramer-Wannier transformation exchanges H_x and H_z , so in the thermodynamic limit $L \rightarrow \infty$ it maps $Z(\beta, g, L, \tau)$ to $Z(\beta g, 1/g, L, \tau)$. This duality holds not only in the limit $\tau \rightarrow 0$, but for any finite value of τ . Therefore, in the limit $\beta \rightarrow \infty$, the model is critical at $g = 1$ for all τ .

This seems quite a trivial observation, but could perhaps be used to derive the critical conditions for the anisotropic classical 2D Ising model. More importantly, it indicates that caution is necessary when using the 1D TFIM as a prototypical quantum spin model when investigating the effects of Trotter-Suzuki approximations at finite time-steps, either in imaginary-time as discussed here or equivalently in real-time dynamics.

4.5 Local Observables

To investigate the effect of scaling the system size on the errors on a local observable we look at the local observable of the nearest-neighbour correlation $\langle \sigma_z^i \sigma_z^{i+1} \rangle$. The difference between the values of the observables obtained in the $\tau = 0$ case by loopex and in the $\tau > 0$ case by spinmc is due to the Trotter error introduced. Figure 4.4 shows the value of this observable for both types of system as a function of the transverse field strength g .




Figure 4.4: The nearest-neighbour correlation, $\langle \sigma_z^i \sigma_z^{i+1} \rangle$, as a function of the transverse magnetic field, g , for Trotter step sizes 0, 0.3 and 0.5 for varying system sizes with (a) $J_2 = 0$ and (b) $J_2 = 0.2$.

We choose a sufficiently high value of inverse temperature to obtain the ground state properties. As we increase the system size it is necessary to scale β with L in order to remain the same effective distance from zero temperature, as increasing the system size decreases the size of the spectral gap. We define the error on an observable O for an imaginary time-step $\tau = 0.5$, to be given by $|O(\tau = 0) - O(\tau = 0.5)|$. In Figure 4.5 we see that for $O = \sigma_z^i \sigma_z^{i+1}$ there is a Trotter error introduced which varies with the transverse field, g , but which is independent of system size, even at the critical point, for both $J_2 = 0$ and $J_2 = 0.2$. We attribute this to the fact that the observable itself is bounded, which will cause it's error to saturate and not increase with system size. If at $\tau = 0.5$ it has already reached this value then we will see no increase with system size.

The behaviour can be more rigorously investigated by looking at the error for different values of τ and looking at its scaling with τ as a function of L . This will reveal the behaviour of the error scaling before it saturates and tell us about the gradient of the error as a function of system size. We expect the error to increase proportionally to τ^2 as explained in Eq. (4.2) and so by simulating the system for different values of τ we

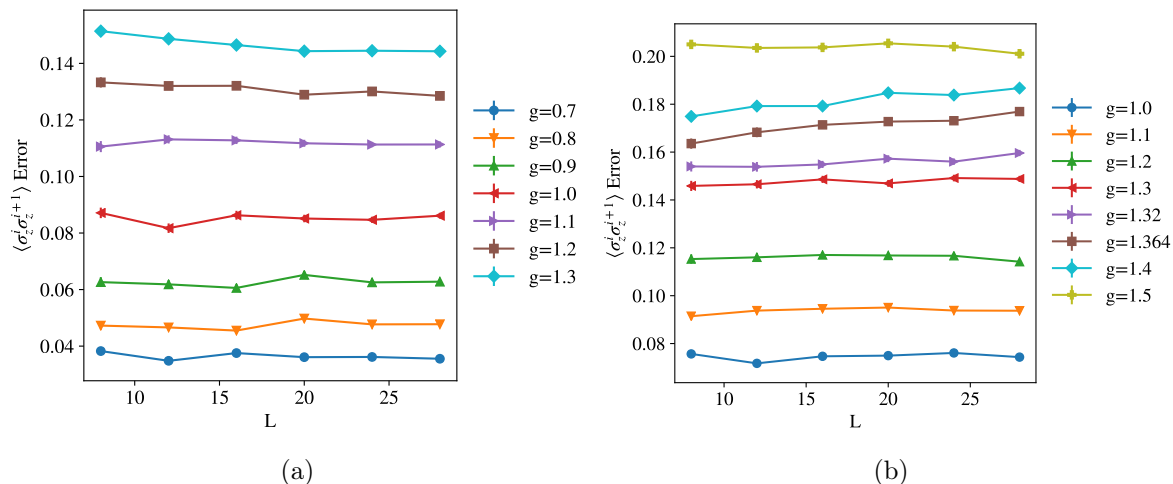


Figure 4.5: The error in observed correlation at $\tau = 0.5$ as a function of the system size for the case of (a) $J_2 = 0$ and (b) $J_2 = 0.2$.

extract the coefficient of this error for further study. For an observable, O , we fit

$$\langle O \rangle = a + b\tau^2, \quad (4.23)$$

where here $A = \sigma_z^i \sigma_z^{i+1}$ and we fix a as the $\tau = 0$ value of $\langle \sigma_z^i \sigma_z^{i+1} \rangle$. This process, shown in Fig. 4.6 enables us to obtain the coefficient of Trotter error, b . We can see that the error of $\langle \sigma_z^i \sigma_z^{i+1} \rangle$ follows the expected $\sim \tau^2$ scaling very well. In the shown range of values of τ we do not observe a saturation in the error of the observable, since we are still far below the value of $\tau = 0.5$ used in Fig. 4.5. Neither do we see any significant deviation from the linear regime due to the contribution of errors which scale with higher orders of τ .

Figure 4.7 shows the value of this coefficient, b , as a function of system size. We can see that, for both models considered here, this too is independent of system size so that not only the value at which the error saturates is independent of system size but so is the rate at which it reaches that saturated value. This implies that proximity to the critical point or moving of the critical point itself does not affect the behaviour of the error on this observable and we conclude that it is not a quantity that couples to critical fluctuations. Figure 4.4 shows that the behaviour of this quantity as we go through the critical point is in both cases a smooth function which may be a contributing factor. This is already an indication that in the case of a local observable the quantity of interest may behave in such a way that it is not necessary to scale N_τ with system size even for simulations near the critical point. This is supported by Appendix A, Fig. A.3.




Figure 4.6: A least squares fit of the error as a function of τ to Eq. (4.23) for $J_2 = 0.2$ at $g = 1.2$ with $L = 12$ where this fit is representative of those carried out to extract b of Eq. (4.23).

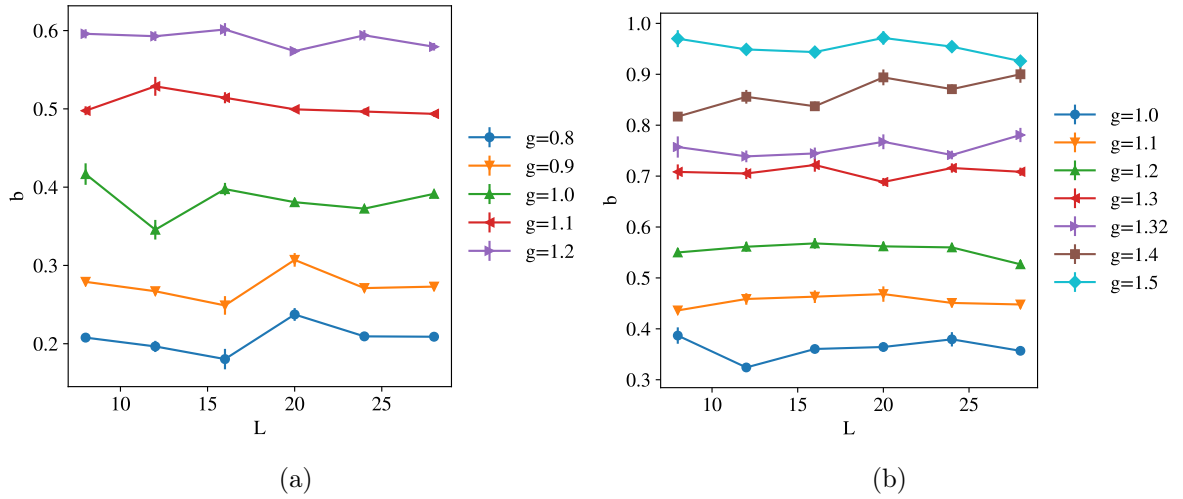


Figure 4.7: Trotter error coefficient b for $\langle \sigma_z^i \sigma_z^{i+1} \rangle$ as a function of system size for (a) $J_2 = 0$ and (b) $J_2 = 0.2$.

4.6 Non Local Observables

To amplify the effect of Trotter errors, we consider the behaviour of a non-local observable $\langle m^2 \rangle$ which is related to the magnetic susceptibility of the system.

Contrary to local observables such as $\langle \sigma_z^i \sigma_z^{i+1} \rangle$ which are bounded by constants, $\langle m^2 \rangle$ is unbounded so should enable a more direct observation of the error scaling with the system size. We also naturally expect non-local observables to be more susceptible to critical behaviour as they are sensitive to behaviour across the system which will be affected by the diverging correlation length. We look at the value of this observable as a function of g for both coupling scenarios for different values of τ in Figure 4.8. This is instructive as we see that the value of this observable changes dramatically as we cross the critical point. This is expected as it is related to the susceptibility of the model from which we expect two distinct regimes. We also see that the behaviour for the $J_2 = 0.2$ case is different from that of $J_2 = 0$ which can be understood to be a result of the critical point moving as a function of τ as discussed in Section 4.4.




Figure 4.8: The static susceptibility, $\langle m^2 \rangle$, as a function of the transverse magnetic field, g , for Trotter step sizes 0, 0.3 and 0.5 for varying system sizes with (a) $J_2 = 0$ and (b) $J_2 = 0.2$.

In this case the value of the observable is not bounded and changes dramatically with system size so we look at the relative error:

$$\frac{|O(\tau = 0) - O(\tau = 0.5)|}{|O(\tau = 0)|}, \quad (4.24)$$

and see in Fig. 4.9 that in both cases it initially increases and then, in the case of $J_2 = 0$ saturates for all values of g . This implies that the error does have some dependence on

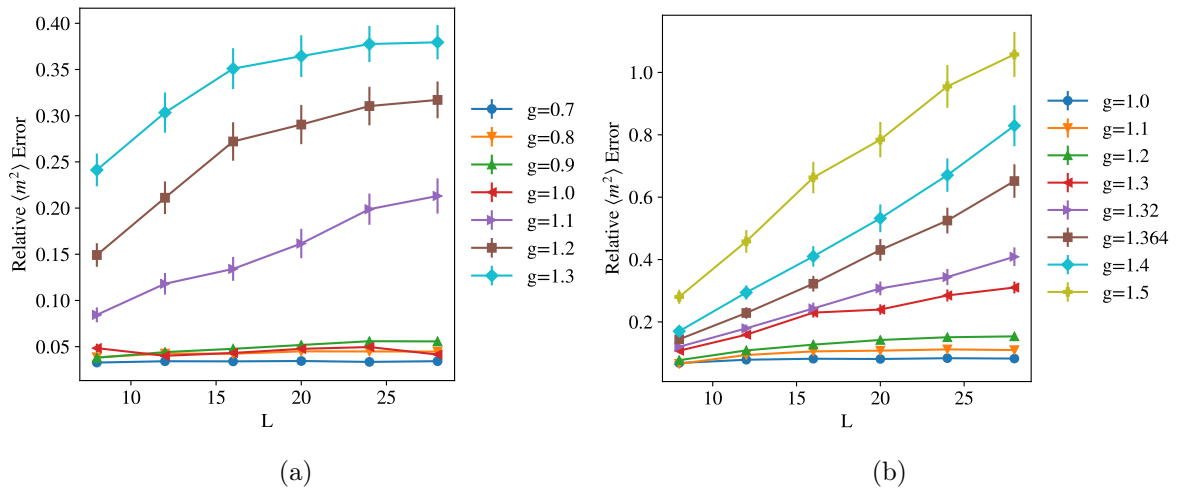


Figure 4.9: The relative error of the static susceptibility found for Trotter step size $\tau = 0.5$ for (a) $J_2 = 0$ and (b) $J_2 = 0.2$. Note that (b) includes data for both the $\tau = 0$ critical point at $g = g_{\text{cr}} = 1.32$ and the $\tau = 0.5$ critical point at $g = 1.364$.

system size but that the relative error saturates even at the critical point. The same appears to be true for $J_2 > 0$ but the behaviour is less well defined for $g > g_{\text{cr}}$.

We look at the τ^2 coefficient of the error on $\langle m^2 \rangle$ in Fig. 4.10 in the same way as we did for the local observable in order to investigate this further. We see that for $J_2 = 0$ the relative gradient b/a depends on system size for $g > g_c$ but saturates. In the case of $J_2 = 0.2$ it is less clear if there is saturation at and slightly above the critical point. This is because increasing τ increases the effective critical point, which has a large impact on the simulated values of $\langle m^2 \rangle$ as they are moved into the region of the effective critical point or through the phase transition. This demonstrates the importance of choosing τ such that it does not shift the critical point significantly relative to the region of interest in the phase diagram of the model.

The number of components in the sum which makes up $\langle m^2 \rangle$ is $(\beta L)^2$ so it is also interesting to look at how $b/(\beta L)^2$ varies with system size as we do in Fig. 4.11 where we see that in both scenarios it appears to saturate implying that the error per term does not increase with system size. This is consistent with the observation about local observables made above.

4.7 Conclusion and Outlook

We have seen that the effect of increasing the system size on the errors introduced by Trotterisation are not as dramatic as the upper bound might suggest. This supports the




Figure 4.10: The relative value of the coefficient of the Trotter error scaling, b , for the observable $\langle m^2 \rangle$ with system size for (a) $J_2 = 0$ and (b) $J_2 = 0.2$.

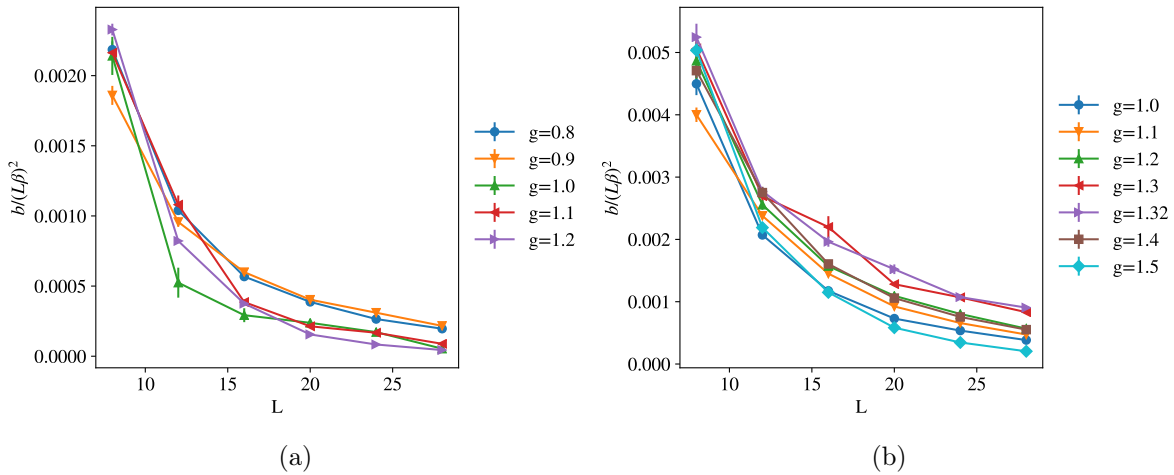


Figure 4.11: The coefficient of the Trotter error scaling, b , for the observable $\langle m^2 \rangle$ scaled by the number of terms in $\langle m^2 \rangle$ with system size for (a) $J_2 = 0$ and (b) $J_2 = 0.2$.

proposal that we treat the error introduced due to Trotterisation as constant calibration error on the Hamiltonian rather than an increasing error on the time-evolution operator. The implications of this are that we may be able to use a fixed Trotter time-step even as we increase the system size and expect to still see much of the same physics. This is supported by results showing that the error on local observables is not affected by system size and that even the error on non-local observables saturates with system size. Another observation of interest we make is that the position of the critical point may or may not be moved as a result of using a finite τ and, as such, τ is a parameter of the Hamiltonian. Specifically of note is that in the case of the quantum 1D TFIM with

nearest-neighbour coupling the critical point is insensitive to changing the time-step.

Further research is in progress to quantify the effective perturbation caused by the Trotter error and to investigate methods of Trotterisation which might allow a more varied set of models to benefit from critical point insensitivity. The exploration of the effect of increasing system size on other observables could be beneficial to verify the implications of these results. Furthermore, it could be beneficial to more rigorously quantify the point at which errors saturate in order to produce a useful guide for simulations of this kind based on the model of interest, the parameter values and the size of the system.

Chapter 5

Superconducting Qubits

In order to implement simulations on a digital quantum computer we of course need the hardware on which to run them. The necessary requirements were specified in the Divincenzo criteria in 2000 [116], but 20 years later the debate over what the best architecture is to build such a system is ongoing. A scalable system of well characterised qubits is required which can be initialised, manipulated by a universal gate set, and individually read out. We additionally require that their coherence times are long relative to the operation times and that the interaction between qubits is controllable. This represents an interesting and unique set of requirements. A simplistic way to look at one of the most problematic aspects of the challenge is that on the one hand we want a system that interacts very little with the environment, so that the computation is completed accurately and without interference, while on the other hand we need to have very precise control over the system including enabling many individual qubits to interact with each other on demand, and to be able to measure fast and accurately. The computation must also be fast enough to actually deliver any of the promised quantum advantage in a reasonable time.

The original possible contenders for fulfilling these requirements include optical photons where the two states are defined by the polarisation or the location of a photon. These are attractive because they are chargeless, can be transmitted over long distances with low loss and can be manipulated with phase shifters and beam splitters, leveraging the robust and well established techniques of the optics community. However they have the drawbacks that interaction between them is weak, loss of a photon is difficult to mitigate and measurement results in destruction of the qubit [117, 118]. They are nonetheless worth mentioning, not only to highlight an example of one of the early promising options, but also because they will serve as a crucial part of the coming chapter, albeit not themselves composing the qubit. Isolated atoms present another promising architecture, again using the wealth of knowledge from the lasers and optics

field to enable cooling and manipulation as either ultracold atoms or trapped by optical lattices or ion traps. Ion trap qubits boast the highest ratios of coherence time to gate time ($\sim 10^6$) and gate fidelities consistently above 99% [71]. Their longer two-qubit gate times and the challenges involved with scaling up the system remain areas of active research. The tantalising prospect of the topological protection which might be realised by majorana based qubits is another hotly researched area [119] although without yet the realisation of a qubit to benchmark. There are many other avenues through which useful digital quantum computers may be achieved and many fields of research have blossomed from this field.

One of the most popular and promising architectures however, which we will use here is the well known ‘transmon’ [51]. The transmon is a qubit comprised by the lowest two levels of an anharmonic oscillator coupled to a microwave cavity and with properties chosen to optimise performance. A Google Scholar search for ‘transmon’ returns over 9,000 results, testifying their popularity and if we consider the cost alone of the dilution refrigerators that are typically used to operate them it is fair to assume that a phenomenal amount of money and effort has gone into developing them. Transmon qubits have the benefit of being solid state and can be fabricated on chip which enables great flexibility in their parameters and possibly a more scalable architecture. They also have competitive gate fidelities and lifetimes [120, 121] and excitingly are the first architecture to execute a process not possible on a classical computer [122].

This chapter will outline the principles of the transmon qubit including how it is defined, controlled, characterised, coupled and measured. We will also attempt to give an overview of the pros and cons of this type of qubit. This should provide sufficient background and motivation to serve as an introduction to the following experiments which are motivated by the requirements to efficiently tune and characterise the system for operation, to have a scalable architecture and to enable precise control of two-qubit operations on an arbitrary pair of qubits.

Since this is such a widely studied system the contents of this chapter is already well documented. We will attempt to concisely outline the concepts necessary to explain the building blocks of the quantum systems we work with, drawing from the combined efforts and insights of many years of excellent and varied research [51, 70, 121, 123–126].

5.1 Defining the qubit

The basic building block of the superconducting qubit is an LC circuit. The classical Hamiltonian describing an LC oscillator with circuit diagram as in Fig 5.1a is

$$H = \frac{CV^2}{2} + \frac{LI^2}{2} \quad (5.1)$$

where C is the capacitance of the capacitor, L the inductance of the inductor and V and I are the voltage across and current around the loop.




Figure 5.1: (a) Simple LC circuit. (b) LC circuit with inductor replaced by Josephson Junction.

If we rewrite this in terms of charge $Q = CV$, and flux $\Phi = LI$, and raise these to the quantum operators (where for notational simplicity we are writing \hat{Q} as Q and $\hat{\Phi}$ as ϕ) we have the quantum Hamiltonian

$$H = \frac{Q^2}{2C} + \frac{\Phi^2}{2L}, \quad (5.2)$$

where $Q = -i\hbar \frac{\partial}{\partial \Phi}$ gives us the canonical commutation relation $[\Phi, Q] = i\hbar$. This forms a harmonic LC oscillator. The eigenenergies of such a system are

$$E = \hbar\omega(n + \frac{1}{2}) \quad (5.3)$$

where $n = a^\dagger a$ counts the number of excitations with a^\dagger and a being the excitation and annihilation operators respectively. $\omega = 1/\sqrt{LC}$ is the resonant frequency of the system.

Note that the circuit in Figure 5.1a has no resistor as this would serve to dissipate the energy stored in it, which is problematic if we are trying to store information in the number of excitations. A superconductor is a good platform for this as there is no resistance and through design we can ensure that the energy levels of our system are within the superconducting gap and so minimise the decay channels. We can ensure that this is fulfilled by making choices to set L and C to give an ω in this range. From now on we will assume that we are dealing with a superconducting circuit which means that electrons will be paired into Cooper pairs and we will have a single non-degenerate, macroscopic ground state [9].

We now turn to the question of how to isolate the lower two energy levels of this system to use as a qubit. This is achieved by replacing the inductor by a weak link coupling two superconductors called a Josephson Junction. This is commonly comprised of a superconductor-insulator-superconductor (S-I-S) connection [127] but a

superconducting-semiconducting-superconducting (S-Sm-S) configuration has also been demonstrated [7, 57, 128]. These systems obey the DC and AC Josephson relations [9, 129]

$$\begin{aligned} I &= I_C \sin(\delta) \\ V &= \frac{\Phi_0}{2\pi} \dot{\delta}, \end{aligned} \tag{5.4}$$

where $\delta = 2\pi\Phi(t)/\Phi_0$ is the phase difference across the junction and I_C is the critical current of the junction. The inductance of this component is defined by $V = -L_J \dot{I}$ which gives us

$$L_J = \frac{\Phi_0}{2\pi I_C \cos(\delta)}. \tag{5.5}$$

This circuit can be seen in Figure 5.1b where L_J and C_J are the inductance and capacitance of the junction. The energy stored in magnetic fields is

$$\begin{aligned} H_{\text{mag}} &= \int V I dt = \frac{I_C \Phi_0}{2\pi} \cos(\delta) \\ &= E_J \cos(\delta), \end{aligned} \tag{5.6}$$

where $E_J = I_C \Phi_0 / 2\pi$ is the Josephson energy. The energy stored in this part of the system is clearly nonlinear although for small δ we can recover linearity as $H_{\text{mag}} \approx E_J(1 - \delta^2/2)$. The effect of this nonlinear element is that we no longer have a parabolic potential with equally spaced energy levels but rather a cosine potential in $\hat{\Phi}$ which means that higher energy levels are closer in energy. The two potentials and their energy levels are shown in Figure 5.2. The important effect is that we can now address the lowest two energy levels individually as their separation is unique.

There are multiple ways to arrange this new element to form different ‘flavours’ of qubit. We may couple to the system by flux [130, 131], phase [132] or charge [133] and each has its own corresponding natural basis [126]. Coupling by charge relies on capacitance and is the most popular and includes the Cooper Pair Box which is the basis for the transmon [134]. The circuit for the Cooper Pair Box (CPB) can be seen in Figure 5.3.

The CPB is controlled by the variable gate voltage V_g . Increasing the gate voltage from zero will induce a polarisation which the charges on the island will attempt to cancel. It may become energetically favourable for Cooper pairs to tunnel across the Josephson junction onto the island in order to neutralise this charge. In the charge basis the states are numbered by the number of extra Cooper pairs on the island, N . This number will be affected by the capacitance to the island $C_\Sigma = C_g + C_J$ and the gate voltage V_g . The difference between the charge on and off the island is $2e(N - N_g)$ where N_g is the number of pairs off the island and the $2e$ is the charge of one Cooper pair. Therefore




Figure 5.2: Sketch of the energy potential of the harmonic (a) and anharmonic (b) LC oscillator as a function of Φ , where the qubit energy levels in each case are shown as well as the transition frequencies between levels. In the harmonic system all the energy levels are separated by E_r , whereas in the anharmonic system this is not the case and $E_{01} > E_{12} > E_{23}$ and so on. These values can be found by numerical analysis.




Figure 5.3: Cooper Pair Box: C_g is the gate charge which together with C_J describes the capacitance to the island (yellow). V_g is the gate voltage and L_J and C_J , as before, are the junction capacitance and inductance. N is the number of Cooper pairs on the island.

the energy stored in the electrical fields of this system is

$$\begin{aligned} H_{\text{el}} &= \frac{4e^2(N - N_g)^2}{2C_\Sigma} \\ &= 4E_C(N - N_g)^2, \end{aligned} \quad (5.7)$$

where $E_C = e^2/2C_\Sigma$ is the charging energy. This brings us to the Hamiltonian for the system which is a sum of the energy stored in the electric and magnetic fields:

$$\begin{aligned} H &= H_{\text{el}} + H_{\text{mag}} \\ &= 4E_C(N - N_g)^2 + E_J \cos(\delta). \end{aligned} \quad (5.8)$$

If $E_J/E_C \gg 1$ then δ will be localised around 0 and if phase is localised then charge and the number of excess Cooper pairs, N , must be correspondingly delocalised.

This can be numerically solved in the charge basis to give the energy spectrum as a function of N_g for different ratios of E_J/E_C where isolating the lower two levels gives us a qubit with transition energy $E_{01}(N_g) = E_1(N_g) - E_0(N_g)$. The main problem that arises is that for $E_J/E_C \lesssim 1$ the energy difference between the two lowest levels is highly sensitive to the value of N_g and so the qubit is correspondingly sensitive to charge noise. This means that while the number of excitations in the system is preserved the phase information, as stored in the ϕ angle of equation Eq. (2.2), is more easily lost. Increasing the value of E_J/E_C reduces the energy dispersion with respect to N_g because charge on the island is no longer well defined and so the effects of charge noise are reduced. The regime where $E_J/E_C \gg 1$ is called the transmon regime.




Figure 5.4: A measurement of a transmon qubit using a high and low drive power where the dips in the cavity magnitude response imply that there is a resonance at that frequency. The resonance at 5.52GHz is the E_{01} qubit resonance and is the only one visible at low driving powers as only single photon excitations are possible. At high power however the two photon transition E_{02} is possible so we see a resonance at $E_{02}/2$. The difference between these is necessarily $\alpha/2$.

In this regime we can expand the cosine so that

$$H_{\text{mag}} \approx E_J \left(1 - \frac{\delta^2}{2} + \frac{\delta^4}{24} \right) + \mathcal{O}(\delta^6), \quad (5.9)$$

where if we had also ignored the δ^4 term we would recover a harmonic oscillator. Having a large E_J/E_C reduces the difference between energy level spacing which we introduced by adding the nonlinear element. The deviation from equally spaced energy levels will set how easily we can isolate the lowest two levels which represent the $|0\rangle$ and $|1\rangle$ states of the qubit without leaving the computational space. The quantity that contains this




Figure 5.5: (a) A transmon qubit where the false colour green T shaped islands are the capacitor pads for neighbouring qubits which set E_C by their capacitance to the surrounding ground plane. The flux line and microwave drive lines to control the qubit are coloured blue and red respectively. Each qubit island is capacitively coupled to a readout resonator above it and also to the islands of the neighbouring qubit(s). At the bottom of each island a pair of JJs connects the island to ground, completing the anharmonic LC circuit necessary for a transmon qubit. A magnified in of one of these SQUID loops connecting the island to the ground plane is shown in (b) and a further magnified image of a single Josephson junction is shown in (c).

information is called the anharmonicity and is defined as the difference between the lowest two transition energies

$$\alpha = E_{12} - E_{01}. \quad (5.10)$$

In experiment we find the value of α by driving the system with a high power (many photon) microwave drive and sweeping the frequency. At high powers we can observe two-photon effects and so the $0 \rightarrow 2$ transition is possible and we see resonance at $E_{02}/2$. The drive and readout procedure will be explained in the next sections in more detail but Fig 5.4 shows this measurement for a qubit with $\alpha \sim 240\text{MHz}$.

For $E_J/E_C \gg 1$ the anharmonicity will be reduced but, while the charge dispersion decreases exponentially with increasing E_J/E_C , the anharmonicity decreases only algebraically [51]. In this regime, working in the charge basis and taking terms up to δ^4 in Eq. (5.9) the qubit frequency is

$$E_{01} = \sqrt{8E_J E_C} - E_C \quad (5.11)$$

and the anharmonicity

$$\alpha = -E_C, \quad (5.12)$$

as detailed in Refs. [51, 123]. Experimentally it is favourable to use a system with a smaller but finite anharmonicity in order to gain insensitivity to charge noise. To manufacture a CPB with $E_J/E_C \gg 1$ a shunting capacitor is used to increase C_Σ and hence decrease the value of E_C . The original proposal for these qubits included a transmission line resonator coupled to each qubit to form part of the shunting capacitance and facilitate readout, hence the name ‘transmon’. In our experiments $E_J/E_C \sim 50$ and $E_C \sim 200$ MHz. This large capacitance is predominantly due to the shunting capacitor island which can be seen in Fig. 5.5.




Figure 5.6: Data from a transmon qubit with an SQUID loop of two Aluminium Oxide Josephson junctions. There is a dip in the cavity magnitude response when the qubit drive frequency is on resonance with the E_{01} transition of the qubit and so as in Fig. 5.4 we can find the value of E_{01} . We do this while varying the voltage which we apply to a ‘flux line’ as seen in Fig 5.5a which induces a current near the SQUID loop, the magnitude of which dictates the flux through the loop and hence the qubit frequency according to Eq. (5.13) and Eq. (5.11).

The transmon frequency can be tuned by changing the effective E_J . One option is to tune δ by changing the flux through the junction. This is effectively possible by replacing the single Josephson Junction with a parallel pair of junctions where we can control the

flux going through this loop as shown in Fig. 5.5b. In this scenario we will have an effective

$$E_J^{\text{eff}} = (E_{J1} + E_{J2}) \cos\left(\pi \frac{\Phi}{\Phi_0}\right) \quad (5.13)$$

The spectrum of this qubit has the cosine behaviour seen in Figure 5.6

Another option which we will be looking at more closely is to use a superconductor-semiconductor-superconductor nanowire junction formed by a vapour-liquid-solid (VLS) nanowire [57, 128] or a two dimensional electron gas [7] where it is possible to use voltage to gate the junction by changing the value of I_C . This type of transmon qubit is called a ‘gatemon’ and the spectrum as a function of voltage is non-monotonic in the case of a VLS nanowire junction and approximately monotonic in the 2DEG nanowire junction case. The properties of this spectrum are discussed in more detail in Ref. [135].

5.2 Readout

Having established the basics of the qubit itself, it is now necessary to describe the readout process that allows us to measure its state. This works by entangling this state with another degree of freedom, which is then measured to infer the qubit state and project the qubit onto that state [123, 125]. This provides us with a method for quantum non-demolition (QND) readout as the qubit is left intact in the state it was measured.

We have already mentioned the harmonic LC oscillator and it is this which we will be using to discern the state of the qubit. The field of circuit quantum electrodynamics (cQED) predates the transmon but has developed alongside it as a way to isolate and readout qubits [136]. It is similar to the concept of cavity quantum electrodynamics where an atom is isolated from the environment by placing it in a cavity. The Jaynes Cummings Hamiltonian describes a two level system coupled to a cavity

$$H = \hbar\omega_c(a^\dagger a + \frac{1}{2}) + \hbar\omega_a\sigma_z + \hbar g(a^\dagger\sigma_- + a\sigma_+) + H_\kappa + H_\gamma, \quad (5.14)$$

where a^\dagger and a are the raising and lowering operators for the number of photons in the cavity and ω_c is the cavity frequency. ω_a is the atomic frequency and σ_z is the Pauli matrix as introduced in Chapter 2 and which has the ground and excited states of the atom as its eigenvectors. The interaction between atom and cavity is described by g , the coupling strength, which sets the likelihood of events where the atom goes from excited to ground state and produces a photon in the cavity: $a^\dagger\sigma_-$. The reverse event is $a\sigma_+$ with σ_\pm adding or removing an excitation to/from the qubit. H_κ describes events of photons escaping the cavity at rate κ and H_γ describes events of the atom decaying to

the ground state at rate γ and no photon being generated in the cavity [125]. In order to see coherent swapping of excitations between the cavity and the atom we must be in the strong coupling limit where $g \gg \gamma, \kappa$. Reaching this limit however can be challenging for real atoms which are very small although Rydberg atoms have seen results in this field [137].

If we move to the idea of putting superconducting circuits coupled to resonators instead of atoms, the same concepts apply but we are able to achieve strong coupling by device fabrication choices [136]. The resonator can be of a number of forms including the form of a transmission line resonator as originally envisaged [125], a three dimensional cavity [138], or a ‘hanging’ coplanar waveguide resonator [139], which is the kind used in our experiments. For a transmon qubit coupled capacitively to a harmonic LC resonator the Hamiltonian will be a sum of the resonator, qubit and the interaction between them so the full Hamiltonian is

$$\begin{aligned} H &= H_r + H_q + H_{\text{int}} \\ &= \hbar\omega_r a^\dagger a + 4E_C(N - N_g)^2 + E_J \cos \delta + 2\beta e V_{\text{rms}}^0 N(a^\dagger + a), \end{aligned} \tag{5.15}$$

where $\hbar\omega_r = \hbar/\sqrt{L_r C_r}$ is the resonator frequency, a^\dagger and a are the creation and annihilation operators for photons in the resonator. The transmon qubit part of the Hamiltonian is unchanged and the interaction is given by the root mean square voltage of the local oscillator $V_{\text{rms}}^0 = \sqrt{\hbar\omega_r/2C_r}$ multiplied by the charge on a Cooper pair $2e$ and the ratio of the gate capacitance to the total capacitance $\beta = C_g/C_\Sigma$, coupled via the number of excess Cooper pairs on the island N [51].




Figure 5.7: Transmon qubit coupled to a readout resonator where the resonator has capacitance C_r and inductance L_r and the qubit is similar to the CPB circuit of Figure 5.3 with a large shunting capacitance C_s in parallel to a SQUID loop of two Josephson Junctions. The qubit and resonator are capacitively coupled through C_g . The yellow island is the island of the CPB and corresponds physically to the large capacitor pad whose size sets C_s and consequently E_C .

Although this is not immediately comparable to the Jaynes Cummings Hamiltonian

Eq. (5.14), it can be rewritten in the basis of the uncoupled transmon states $|i\rangle$ as

$$H = \hbar\omega_r a^\dagger a + \hbar \sum_i \omega_i |i\rangle \langle i| + \hbar \sum_{i,j} g_{ij} |i\rangle \langle j| (a^\dagger + a), \quad (5.16)$$

where $g_{ij} = 2\beta e V_{\text{rms}}^0 \langle i| N |j\rangle$ are the dipole coupling energies and $\hbar\omega_i = E_i$ the qubit energy levels. In the transmon limit $\langle i| N |j\rangle \rightarrow 0$ for all $j \neq i \pm 1$. Making a rotating wave approximation (RWA) to neglect terms which do not conserve the number of excitations and ignoring higher levels of the transmon we arrive at

$$H = \hbar\omega_r a^\dagger a + \frac{\hbar\omega_q}{2} \sigma_z + \hbar \sum_{i,j} g(\sigma_- a^\dagger + \sigma_+ a) \quad (5.17)$$

where $g = g_{01}$ and $\omega_q = \omega_{01}$ and this now matches the Jaynes Cummings Hamiltonian for $\kappa = \gamma = 0$. Figure 5.7 shows the circuit diagram for this system.

Describing the difference between the qubit and resonator frequencies as $\Delta = |\omega_r - \omega_q|$ we can describe the regime where $g \ll \Delta$ as the dispersive limit. In this regime we can expand the Hamiltonian in powers of g/Δ to give the *dispersive Jaynes Cummings Hamiltonian*

$$H = \hbar(\omega'_r + \chi\sigma_z) a^\dagger a + \frac{\hbar\omega'_q}{2} \sigma_z, \quad (5.18)$$

where the resonator frequency and qubit frequency are renormalised by the so called Lamb shift to give $\omega'_r = \omega_r - \chi_{12}/2$ and $\omega'_q = \omega_q + \chi_{01}$, where $\chi_{ij} = g_{ij}^2/(\omega_{ij} - \omega_r)$. The resonator now has an effective frequency $\omega'_r \pm \chi$ which depends on the state of the qubit, with the shift given by $\chi = \chi_{01} - \chi_{12}/2$. The dispersive shift can be expressed in terms of the anharmonicity as

$$\chi = \frac{g^2}{\Delta} - \frac{\alpha/\hbar}{\Delta + \alpha/\hbar}, \quad (5.19)$$

which for a true two level system would reduce to g^2/Δ .

At this point it is necessary to introduce a finite κ as we need photons to be able to escape from the resonator in order to be able to readout the qubit state. The quality factor of a resonator is defined as

$$Q = \frac{\omega_r}{\kappa} = \frac{\text{input power}}{\text{dissipated power}} \quad (5.20)$$

In the dispersive regime the mismatch between the resonator and qubit frequency will serve to shield the qubit from decaying to the environment but even in the absence of




Figure 5.8: The full device image of a 4 qubit transmon device (a) where the qubits are each coupled to a readout resonator which is in turn coupled to a feedline. The flux and microwave lines for qubit control are coloured blue and red as in Fig. 5.5a. We send a microwave signal to the feedline and measure the response. If the signal is on resonance with one of the readout resonators then we see reduced magnitude response at that frequency as in (b). We refer to this as the cavity response where here in (b) we show the magnitude of this response. This response will depend on the state and frequency of the qubit which is how we readout the qubit state.

other decay channels there will still be some decay via the cavity, known as the Purcell effect [125, 140, 141]. The rate of decay in this way is $\gamma_\kappa = \kappa g^2 / \Delta^2$. As such we have some trade off between how easy it is for us to readout as making Δ smaller means a larger dispersive shift for us to detect but also increases the qubit decay rate γ which is undesirable. Typical parameters of $Q = 10000$, $2\pi\omega_r = 7\text{GHz}$, $2\pi\Delta = 2\text{GHz}$ and $2\pi g = 100\text{MHz}$ yield a lower limit on γ of $\sim 10000\text{s}^{-1}$. This can be avoided by adding a Purcell filter filtering to allow a higher κ whilst limiting γ and combined with techniques such as readout pulse shaping to ring up the cavity have achieved readout fidelities of $> 99\%$ in less than 100ns [142]. The quality factor of the resonator can be determined by the width of the resonance as seen in Fig. 5.8b which shows the S21 coefficient of the returning signal where the response is reduced when the signal sent to the sample is on resonance with one of the readout resonators.

In reality decaying via the cavity is not the only mechanism for the qubit decay. Dielectric losses to the substrate on which the circuit is fabricated have been shown to be a dominant source of decoherence [143] and the development of 3D resonators where the electric field lives mostly in vacuum rather than in the dielectric on chip [138]. Coupling to spurious modes can be mitigated by increasing the number of on and off chip bonds to prevent local fluctuations in potential. These on-chip bonds can be seen in Figs. 5.5a and 5.8a. The number of quasiparticles in the system is minimised at low temperature

so the decay and dephasing mechanisms opportuned by an odd number of electrons or breaking of Copper pairs are avoided by lowering the operating temperature. We must also couple our transmon to an external magnetic flux in the case of the S-I-S-junction transmon or to a voltage line in the case of the gatemon, in order to tune the qubit frequency, which provides another means for decay.

The readout resonator is coupled capacitively to a feedline as shown in Figure 5.8a, setting the minimum value for κ and enabling us to probe the readout resonator response. If we send a microwave pulse near the resonator frequency at $\omega_r' + \epsilon$ then the amount which it is detuned from the effective resonator frequency $\omega_r' \pm \chi$ due to the state dependent shift will result in an entanglement between the cavity and the qubit. The reflection coefficient will depend on the state of the qubit and the returning signal will have magnitude and phase dependent on this. This technique is most simply employed to find the frequency of the ω_{01} transition of the qubit as in Fig. 5.6. We transmit a signal close to the frequency of the resonator and observe the response at this frequency. A second signal with variable frequency is employed. When the second signal matches the qubit transition frequency there will be some probability that the qubit becomes excited. This change in the qubit state thus changes the effective resonator frequency and the detuning of our first signal from it. Measuring the response at the frequency of our first signal and averaging over many repetitions we therefore see a different resonator response when the second signal is resonant with the qubit, enabling us to identify the qubit frequency. In practice we first send the ‘second’ signal to excite the qubit and then the ‘readout’ signal as this minimises the chance of observing the effect of other short lived resonances on chip which might also affect the resonator response. When the qubit frequency is modulated by means of the flux, the background of the $|0\rangle$ response of the resonator is varied, as can be seen in Fig. 5.6. This is due to the fact that ω_r' depends on the qubit frequency due to the Lamb shift, and so we see this change in resonator response due to the qubit frequency moving as well as the change due to exciting the qubit. This is discussed in more detail in Refs. [121, 144, 145]. In more sophisticated measurements we can first apply a signal to the qubit which executes some desired operation and then measure the state of the qubit by applying a readout pulse at the resonator frequency and analysing the magnitude and phase of the response as we will discuss now.

We can determine the state of the qubit from a single measurement of the readout resonator, however the will uncertainty in this result will be increased if there is overlap in the gaussian distributions describing the $|0\rangle$ and $|1\rangle$ responses in magnitude-phase (or real-imaginary) space. A low signal-to-noise (SNR) ratio can also challenge readout and events such as the qubit decaying during readout may limit our ability to reliably find the initial qubit state [145]. We can quantify how well we are able to readout the qubit by finding the readout fidelity which we find by preparing the qubit in $|0\rangle$ or $|1\rangle$




Figure 5.9: Readout calibration process and readout fidelity calculation. The response of the readout resonator when a signal at the resonator frequency is incident on the resonator. The qubit is first prepared in $|0\rangle$ and the real (I) and imaginary (Q) parts of the response measured. This measurement is repeated 3000 times and the same is done for the qubit prepared in $|1\rangle$. The single-shot measurements of each type are shown in the left panel with the mean coordinates of the $|0\rangle$ prepared results indicated by a down triangle and those of the $|1\rangle$ prepared results indicated by an up triangle. The results of each type of preparation are projected onto the line joining the mean coordinates of the preparation results and then plotted in a histogram in the upper right panel. The cumulative sum of this is shown in the lower right panel where the maximum separation between the cumulative sum in each is the readout fidelity. The point at which this is found tells us how to decide whether future measurements should be assigned to $|0\rangle$ or $|1\rangle$ and is shown as a dashed line.

and then asking how likely we were to correctly ascertain this state from measurement. A totally uninformed guess would give us 50% success but even with all the information we are unlikely to reach 100% as if there is any thermal population or decay before measurement then even if we correctly find the true state every time this will not always match our expected state. This can be seen in Fig. 5.9 where there are always some ‘ $|0\rangle$ ’ results in the ‘ $|1\rangle$ ’ location and vice versa. This quantity tells us both about our ability to readout the state and also about our ability to initialise. Practically we find this by measuring the real (I) and imaginary (Q) response of the resonator averaged over the duration of the pulse for repeated measurements with the qubit in $|0\rangle$ and then $|1\rangle$. The average IQ response is found for each and the single-shot measurements projected onto the axis joining the two centres. A histogram of these points tells us how they are distributed along the line and the separation between the cumulative sums describes how well we can distinguish the state of the qubit from a result at that point. This

process is shown in Fig. 5.9, where this measurement finds a maximum separation of the cumulative sums to be $\mathcal{F} = 0.74$. This maximum separation is the readout fidelity and we also learn the axis onto which to project subsequent measurements for maximum contrast when averaging or the value on this axis which separates $|0\rangle$ and $|1\rangle$ in order to decide the qubit state when taking single-shot measurements.




Figure 5.10: The same data as used in Fig 5.9 without using any former knowledge about the prepared state of the qubit. 6000 single shots are taken with the qubit prepared equally in $|0\rangle$ and $|1\rangle$ and then the Kmeans algorithm is used to identify each single-shot measurement as corresponding to one of the two states according to which of the two learned ‘centres’ it is closer to. This is shown in the left panel where the centres are shown by up and down triangles. The line between the two centres is used to project the data onto and this projected data is shown histogrammed in the upper right panel. The line bisecting the one joining the two centres is depicted by a dashed line in all plots and is used to assign qubit state. The cumulative sum in the right panel illustrates how well this enables us to distinguish between the two states.

If we want to know how well we can distinguish the two states without being influenced by bad initialisation then the quantity known as ‘separation fidelity’ can be found by fitting the sum of two gaussians to the IQ data regardless of assigned state and then setting a threshold at the point where they intersect [145]. This quantity gives the mean probability that a measurement is identified correctly. In practice however, this relies on either fitting a two dimensional gaussian or projecting the data onto an axis defined using centres found as part of the other method and then fitting the sum of two gaussians. Here instead we use the KMeans algorithm from the scipy library [146] which makes clusters of data points and learns the centres of these clusters. We can then use this going forward to find the qubit state based on IQ readout coordinates but in practice

we use it to find the centres of the gaussians and use the midpoint as the decision value to determine the qubit state going forward. This alternate method can be seen in Fig. 5.10. The former method of course has the benefit of being a more rigorous test of how well we can initialise and readout our qubit but the latter method is often better for calibrating readout to distinguish the $|0\rangle$ and $|1\rangle$ state of the qubit. The second method also has the benefit of not requiring an X pulse to have been calibrated in advance. This is because it works by separating the points into two groups and requires no labelling so any experiment which has reasonable $|0\rangle$ and $|1\rangle$ population over the full experiment can be used to calibrate the readout in this way. The second method is not sensitive to initialisation problems as long as there is a significant portion of single shots in each state which is attractive for readout calibration, but it produces a less comprehensive quantity in terms of reporting the readout fidelity. An example is the Rabi experiment shown in Fig. 5.11, where the single-shot data for the experiment is be used to calibrate the readout and thus save a separate experiment to do so.

A low readout fidelity can be caused by bad initialisation, which itself could be due to thermal population, non-ideal gates, short qubit lifetime or paradoxically by long qubit lifetimes in the case where initialisation to $|0\rangle$ is done by waiting for the qubit to decay. Increased readout power decreases the width of the gaussian peaks which makes it easier to distinguish between states but increasing the power too much may drive the resonator into its bright state where there is no longer a state dependent shift. Interestingly however, the point at which this happens has spawned other readout methods [147]. Increasing the readout duration increases the distance between the centres of the gaussians, which helps with distinguishing the state, but taking too long to read out increases the chance that the qubit will decay during the process, so there has also been much research into optimising this procedure [142, 148].

Since there will be multiple qubits, each with its own unique frequency readout resonator, we need to generate signals at these frequencies to send into the feedline and then separately determine the IQ response at each of these frequencies in the output [148–150]. In practice this is achieved by generating a single frequency signal from a microwave source and then modulating it using signals from a waveform generator in order to generate several sidebands at the desired frequencies. This multi-frequency signal is then transmitted to the feedline of the sample. The result is mixed down with a tone from a local oscillator and the high frequency component filtered away. The resultant multi-low-frequency signal is then sampled by an analogue to digital converter (ADC) and demodulated at its constituent frequencies and the result averaged over the duration of the pulse. This results in a single IQ point for each frequency, enabling us to do simultaneous multi-qubit single-shot measurement or to repeat many times and instead find the average resonator response to the stimulus. One option for obtaining the excited qubit population $P_{|1\rangle}$ despite the fact that we only ever measure $|0\rangle$ or $|1\rangle$ is

to repeat single-shot measurements and for each measurement find the qubit state and then average these [103]. The other option is to average the IQ resonator response and then map this directly to the qubit state based on some initial calibration. The latter method is less rigorous but usually both experimentally and computationally faster.

5.3 Control and Characterisation

Another key component necessary to implement an algorithm is the ability to execute single-qubit gates. It is also necessary to characterise the properties of the qubits, both to tune up the gates for maximum gate fidelity but also to learn about sources of noise, in order to iteratively improve the quantum hardware and find the capabilities of a particular sample or architecture. We would like to be able to implement all the gates necessary to execute the algorithms of the previous two chapters, using the gates introduced in Chapter 2.

Single-Qubit Rotations

The first, and most simple, operation we wish to achieve is to rotate the qubit around its x axis in order to implement an $R_x(\theta)$ rotation gate, most simply a σ_x gate. This is achieved through the capacitive coupling of the qubit to an external voltage source which can be done via a direct voltage gate or through the feedline. Both options can be seen in Figure 5.8. This will modify the number of Cooper pairs off the island to be $N_g = N_{g0} + N_g(t)$, where we set $N_{g0} = 1/2$ and $N_g(t) = A \cos(\omega_d t + \phi)$. This modifies our qubit Hamiltonian to be

$$H_q = 4E_C A \cos(\omega_d t + \phi) \sigma_z - \frac{E_J}{2} \sigma_x, \quad (5.21)$$

which, by performing a static rotation to convert $\sigma_x \rightarrow \sigma_z$ and $\sigma_z \rightarrow -\sigma_x$, going into a frame rotating around our new z axis at ω_d , and making the RWA to discard fast rotating terms, results in

$$H'_q = \frac{\omega_q - \omega_d}{2} \sigma_z + \frac{\bar{A}}{2} [\cos(\phi) \sigma_x + \sin(\phi) \sigma_y], \quad (5.22)$$

where $\bar{A} = 4E_C A$. For $\omega_d = \omega_q$ this allows us rotation by any angle in the xy plane and by detuning the drive from the qubit frequency can introduce some z component. A more detailed derivation of this is available in Refs. [123, 124, 151]. For $\phi = 0$ the unitary operator expressing evolution under this Hamiltonian for time t is

$$U(t) = \begin{pmatrix} \cos(\Omega t/2) - i(\delta_{dq}/\Omega) \sin(\Omega t/2) & -i(\bar{A}/\Omega) \sin(\Omega t/2) \\ -i(\bar{A}/\Omega) \sin(\Omega t/2) & \cos(\Omega t/2) + i(\delta_{dq}/\Omega) \sin(\Omega t/2) \end{pmatrix} \quad (5.23)$$

where $\delta_{dq} = \omega_d - \omega_q$ and $\Omega = \sqrt{\bar{A}^2 + \delta_{dq}^2}$.

We are free to choose ϕ , \bar{A} , δ_{dq} and t and are thus able to realise X , $\pm X/2$, Y , $\pm Y/2$ and even the Hadamard gate, although with only rotations around the x and y axis we already have access to the whole Clifford gate set. Since we can also implement arbitrary angle rotations around these axes the whole Bloch sphere is theoretically accessible.

We start by tuning up the gates that implement X and $X/2$ operations by applying a pulse of duration τ_{pulse} to the qubit during which it will evolve under H'_q . This evolution will correspond to a rotation around the x axis of the Bloch sphere. We then apply a readout pulse to measure the state of the qubit and observe Rabi oscillations where the qubit state oscillates between $|0\rangle$ and $|1\rangle$ as a function of the duration of the drive pulse. Averaging over many measurements we observe the averaged magnitude response of the cavity to oscillate, mirroring the behaviour of $P_{|1\rangle}$. Sweeping the detuning δ_{dq} , we see the effect of changing the axis of rotation to include a z component. The result is the characteristic chevron pattern in Fig. 5.11a.

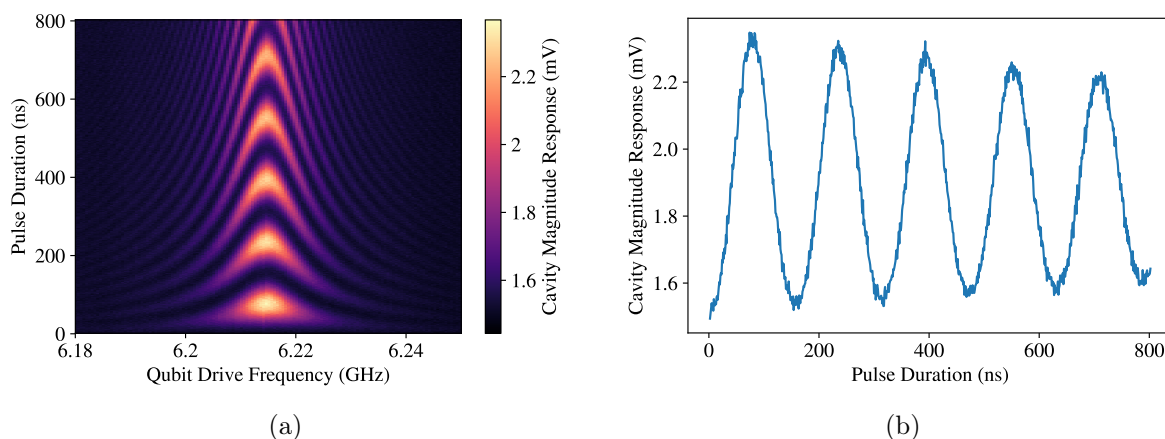


Figure 5.11: The averaged magnitude at the resonator frequency after a pulse with frequency $\omega_d/2\pi$ for duration τ_{pulse} . The characteristic chevron pattern described by Eq. (5.22) and Eq. (5.23) is shown in (a) and a cut at $\delta_{dq} = 0$ with $\omega_d/2\pi = \omega_q/2\pi = 6.215\text{GHz}$ is shown in (b) where we see coherent oscillations as the qubit state is transferred from the ground state to the excited state and back again. 10000 averages have been performed for each magnitude measurement.

If we instead take single-shot data and allocated each measurement to $|0\rangle$ or $|1\rangle$ based on the readout calibration we obtain the averaged $P_{|1\rangle}$ probability as shown in Fig. 5.12. Here we see that we only reach $P_{|1\rangle} = 0.8$ which could be because we are not driving

at exactly the right frequency but is also impacted by our imperfect readout. These important details are necessary to consider when designing an algorithm and are often neglected in the initial stages of benchmarking an architecture.




Figure 5.12: The averaged $P_{|1\rangle}$ of the qubit as a function of the duration of the signal applied to the it. The assignment was carried out based on readout calibration using the Kmeans algorithm on the data set. Each point plotted here is the result of 150 averages.

This process allows us to find a power, frequency and duration of signal that will complete state transfer and thus apply any $R_x(\theta)$ rotation. Changing the phase of the signal applying the gate from $\phi = 0$ to $\phi = \pi/2$ will change the axis of rotation to the y axis. Rotation around these two orthogonal axes enables us to apply any Clifford gate and hence any single-qubit rotation.

T_1 , T_2^* and $T_{2,\text{echo}}$

We now have all the tools necessary to execute the experiments necessary to characterise the behaviour of our qubit. We have already mentioned the limit on the lifetime of the qubit due to coupling to the environment through the resonator but this is not the only possible means of decay and it is of course important to verify this value in experiment. The number which quantifies the longitudinal decay of the qubit is T_1 , which is equal to the inverse of our decay rate quantity γ . In the most complete terms T_1 describes dissipation of the qubit to an environment with finite temperature but we will here assume that the infinite time state of the qubit will always be $|0\rangle$. This behaviour is described by the amplitude damping channel which serves to shrink the Bloch sphere, moving every point on it towards a fixed point at the north pole, representing $|0\rangle$ [117].

If we start in $|0\rangle$, apply an X gate to put the qubit in $|1\rangle$, and then wait for time τ_{wait} , then the probability for the qubit to be in the $|1\rangle$ state, $P_{|1\rangle}$, will decay as $\exp\{-\tau_{\text{wait}}/T_1\}$, where $T_1 = 1/\gamma$. This measurement and the least squares fit used to extract T_1 is shown




Figure 5.13: The averaged cavity response for varied wait time τ_{wait} between a pulse applying an X gate and readout. This fit is used to find the longitudinal decay time which here is $T_1 = 11\mu\text{s}$. Each point is the result of 5000 averages. The method of measuring the average cavity response and inferring from it the behaviour of $P_{|1\rangle}$ is detailed in the text.

in Figure 5.13 where we find a T_1 of $11\mu\text{s}$. Initialisation of the qubit can be done by simply leaving the system long enough that the qubit has with high probability decayed to the ground state. One of the few perks of having a low lifetime qubit is that it makes this process faster.

We must also quantify the transverse decay, which measures how long the phase information of the qubit can be retained. This information is lost if there is noise in a parameter which affects the qubit frequency as this introduces noise in the phase accumulated by the $|1\rangle$ state compared to that of the $|0\rangle$ state. The transverse decay time T_2 quantifies the phase damping of the qubit, the effect of which is to squeeze the Bloch sphere so that all points on its surface become closer to the axis joining the north and south pole. In practice it is this number which is often the limiting one, but in the case of low frequency noise it is possible to mitigate the effects of dephasing by applying dynamical-decoupling pulse sequences [152, 153].

The combined effect of T_1 and T_2 on the density matrix is

$$\rho = |\psi\rangle\langle\psi| = \begin{pmatrix} a & b \\ b^* & 1-a \end{pmatrix} \rightarrow \begin{pmatrix} (a-1)e^{-t/T_1} + 1 & be^{-t/T_2} \\ b^*e^{-t/T_2} & (1-a)e^{-t/T_1} \end{pmatrix}. \quad (5.24)$$

In practice T_2 is often split into the T_2^* and $T_{2,\text{echo}}$, where the former is the bare transverse coherence time and the latter is the refocussed transverse coherence time. This refocusing is achieved by implementing the Hahn echo pulse scheme [154], which mitigates slow noise in this decoherence mechanism. In both cases the qubit starts in $|0\rangle$, then an $X/2$ gate is applied to rotate it by $\pi/2$ around the x axis. Then, after some time τ_{wait} , a second $X/2$ gate is applied and then the qubit is measured. In both cases there

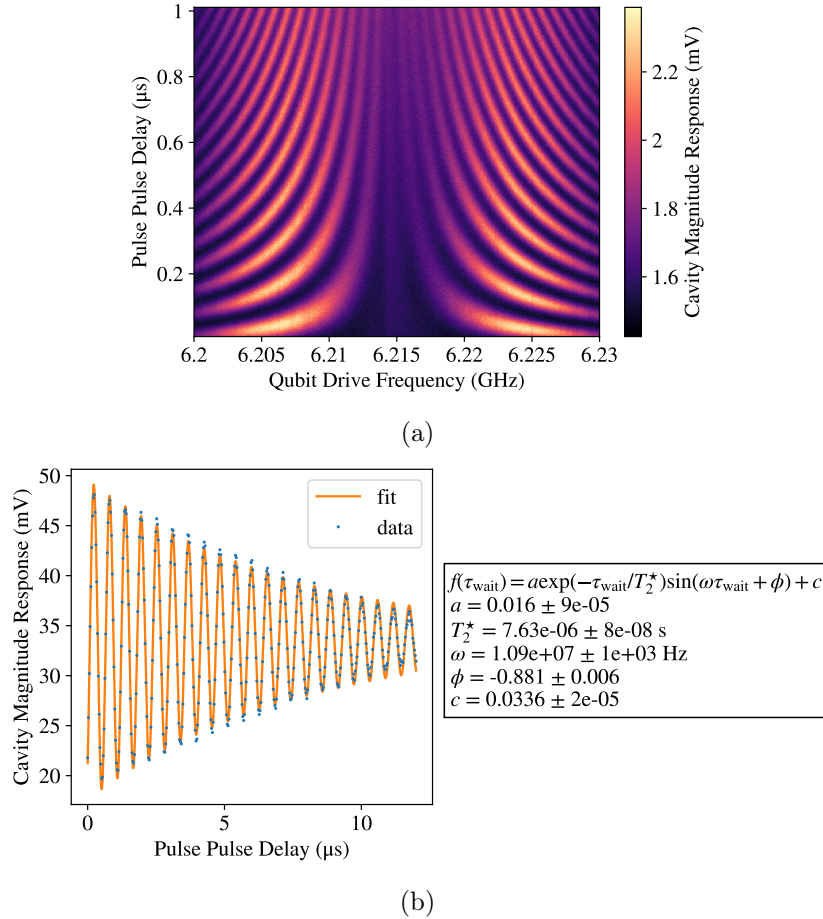


Figure 5.14: The averaged magnitude of the resonator response after a pulse implementing a rotation on the qubit, waiting for a time τ_{wait} , and then another identical pulse before reading out. When the pulses applied are at the qubit frequency the pulses will implement a $X/2$ rotation. In (a) we sweep δ_{dq} which is 0 at qubit drive frequency, $\omega_d/2\pi = 6.215\text{GHz}$. When $\delta_{dq} \neq 0$ a perfect $X/2$ is not implemented but rather a rotation by $\pi/2$ around an axis in the xz plane which means that not only do we not reach the $|1\rangle$ state after two such pulses but also that due to the rotating frame the qubit state effectively undergoes a z rotation during τ_{wait} and therefore its final state depends on the duration of the wait. These are the oscillations which can be seen away from $\delta_{dq} = 0$ in (a) and such a measurement from a different data set is seen in (b) where we have also executed a fit to the data to obtain the T_2^* time of $7.6\mu\text{s}$ as described in the text. 10000 averages have been performed for each magnitude measurement.

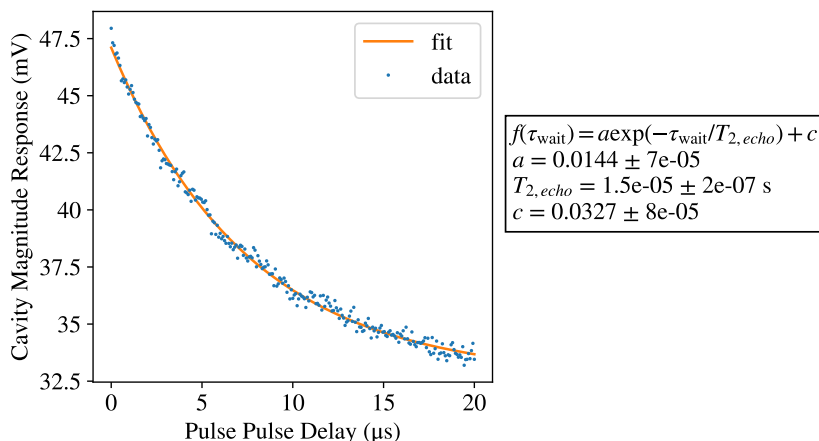


Figure 5.15: The averaged cavity response for varied wait time τ_{wait} between a pulse applying an $X/2$ gate and another $X/2$ before readout. An X gate is executed midway between the $X/2$ gates. The delay shown on the x axis of the plot is the delay between the first $X/2$ pulse and the central X pulse and so τ_{wait} is twice this value (where we take the duration of the central X pulse as negligible compared to the wait time). Each point is the result of 5000 averages. This Hahn echo sequence allows us to find the $T_{2,\text{echo}}$ time which here is $15\mu\text{s}$.

will be an exponential decay which characterises the T_2 time. For the T_2^* measurement the ‘ideal’ final state after two $X/2$ pulses is $|1\rangle$ but after a long time $P_{|1\rangle}$ will decay to 0.5. To measure this decay in experiment the $X/2$ gate is often detuned with $\delta_{dq} > 0$ so that there is an additional rotation around the z axis for the purpose that this is easier to fit this exponentially decaying sinusoid to find the T_2^* value. This experiment can be seen in Fig. 5.14, where the fit used to extract T_2^* is shown in b. For the $T_{2,\text{echo}}$ measurement an additional X gate is applied half way through τ_{wait} which serves to refocus any slow spread of the wavefunction and the experiment is executed at $\delta_{dq} = 0$. Thus the ideal final state is $|0\rangle$ and the long time final state has $P_{|1\rangle} = 0.5$. In practice we often fit an exponential decay to the cavity response rather than the qubit state as discussed previously. The data and fit can be seen in Fig. 5.15.

If decoherence is caused by low frequency noise, the effects of which can be mitigated by refocussing pulse schemes of this type, then the upper limit for $T_{2,\text{echo}}$ is $2T_1$ so finding $T_{2,\text{echo}}$ also enables us to characterise the type of noise experience by the qubit.

Randomised Benchmarking

A more general metric which is nevertheless useful across architectures and when estimating whether quantum hardware can meet the limits required to execute an algorithm

is the gate fidelity. The average fidelity of a quantum channel is

$$\mathcal{F}(\mathcal{E}) \equiv \int d\psi \langle \psi | \mathcal{E}(\psi) | \psi \rangle, \quad (5.25)$$

where \mathcal{E} is the trace preserving quantum operation action on the qubit ψ [155]. How similar the action of \mathcal{E} is to a quantum gate U is given by


$$\mathcal{F}(\mathcal{E}, U) \equiv \int d\psi \langle \psi | U^\dagger \mathcal{E}(\psi) U | \psi \rangle, \quad (5.26)$$

where $\mathcal{F}(\mathcal{E}, U) = 1$ if and only if \mathcal{E} exactly implements U . To find the average fidelity of a single Clifford gate we can implement a process called randomised benchmarking (RB) [48, 156–159], where sequence of gates of length m is chosen randomly from the gate set and then applied to a qubit in the $|0\rangle$ state. After this sequence has been implemented the single gate which returns the qubit to $|0\rangle$ is applied and the qubit is measured. The length of the sequence m is varied with n repetitions at each sequence length. The likelihood of returning to $|0\rangle$ will decay exponentially with the fidelity of the gates with


$$\mathcal{F}_{\text{seq}}(m) = Ap^m + B, \quad (5.27)$$

where $\mathcal{F}_{\text{seq}}(m)$ is the fidelity of the sequence of m gates compared their ideal operation and A and B are constants which absorb any preparation and measurement errors. This is the simplest formulation of the result which assumes that errors are gate independent. The ideal final state of the qubit is $|0\rangle$ but the long time $P_{|1\rangle} = 0.5$. We find the average fidelity per clifford gate p by fitting Eq. (5.27) to the data as in Fig 5.16b. Running this experiment while varying an experimental parameter can help us to tune the gates in order to maximise fidelity as function of the parameter [160, 161], as shown in Fig. 5.16a.

The average gate error r is given by $r = (1 - p)/2$ where in the case that the only errors introduced are due to the amplitude damping channel $r = t/3T_1$, where t is the time to execute a gate. For $T_1 = 10\mu\text{s}$ and a gate time of 50ns this would result in $r = 0.0017$ or a single-qubit Clifford gate fidelity of $p = 0.997$ so this can be used to set the limit on the minimum T_1 required to allow for high fidelity gates above some threshold. We note that a single Clifford gate may contain more than one ‘physical gate’ where physical gates are microwave pulses calibrated to induce a specific rotation on the qubit, commonly X , $\pm X/2$, Y and $\pm Y/2$. For us the average number of physical gates per Clifford gate is 1.83, so depending on the algorithm it is possible to have a higher fidelity than the Clifford gate fidelity if we can execute the algorithm using individually tuned native gates. It is also worth noting that the theory behind the development of RB is based on the assumption that the n repetitions will each consist of a unique random gate sequence but that in practice it is often experimentally faster to reuse gate



(a)



(b)

Figure 5.16: (a) Randomised benchmarking results, varying the frequency and power of the microwave source implementing single-qubit gates and extracting the Clifford fidelity at each point. Fits of the kind of those shown in (b) are performed at each point in (a) where the Clifford Fidelity is denoted by p . This is done by varying the number of Clifford gates, m which make up the sequence applied to a qubit before it is returned to $|0\rangle$ and then fitting Eq. (5.27) to the result. Each data point represents 5000 [100] averages in (a) [(b)].

sequences and repeat a specific sequence, only using a small number of unique sequences for different m as we do here. This is theoretically an inefficient way to find the average gate fidelity as there is a lot of redundancy but experimentally the data taking is often the fastest part of the experiment and generating different sequences the limiting factor.

The motivation behind RB was to find a single metric to quantify the performance of the quantum hardware. It is possible to report the minimum average fidelity required to reach a certain level of accuracy in a quantum computation for which such a metric is useful. However its simplicity means that it is also limited in its descriptive properties and usefulness. Methods such as state tomography and process tomography [162–164] provide more information about the action of operations on the qubit. They are therefore better suited to in depth analysis and error debugging but they are expensive in the amount of data and the number of experiments required. In between these two extremes exist methods of gate tune up and metrics such as restless tune up [165] and quantum volume [166] and processes such as AllXY [123] but the challenge of efficiently tuning and calibrating your quantum hardware remains a challenging field of research.

5.4 Two-Qubit Interaction

Finally, we need to implement two-qubit gates which requires some interaction between qubits. The most simple way to couple two transmon qubits to each other is to couple them capacitively by placing the large capacitor pads which contribute to the shunting capacitance near to each other. This is the same mechanism by which the qubit is coupled to the resonator and the Hamiltonian describing this coupling is similar. If we use the assumptions already made about the transmon being a two level system then the interaction part of the Hamiltonian between two neighbouring qubits will be

$$H_c = \hbar J(\sigma_- \sigma_+ + \sigma_+ \sigma_-), \quad (5.28)$$

with the coupling J defined as

$$J = \frac{C_{qq}\sqrt{f_{q1}f_{q2}}}{2C_{\Sigma 1}C_{\Sigma 2}}, \quad (5.29)$$

where C_{qq} is the mutual capacitance between the qubit islands, $f_{q1(2)}$ is the frequency of qubit 1(2) and $C_{\Sigma 1(2)}$ is the total capacitance of qubit 1(2). This means that when $f_{q1} = f_{q2}$ an excitation will swap between qubits 1 and 2 at rate J [167]. When the qubits are far detuned the effect will be suppressed. Figure 5.17 shows this effect.

In order to implement a two-qubit gate it is therefore possible to start with the qubits relatively detuned and then bring them into resonance for a time τ_{swap} where an excitation will swap between them and if τ_{swap} is chosen accordingly a $\sqrt{i\text{SWAP}}$ gate can be

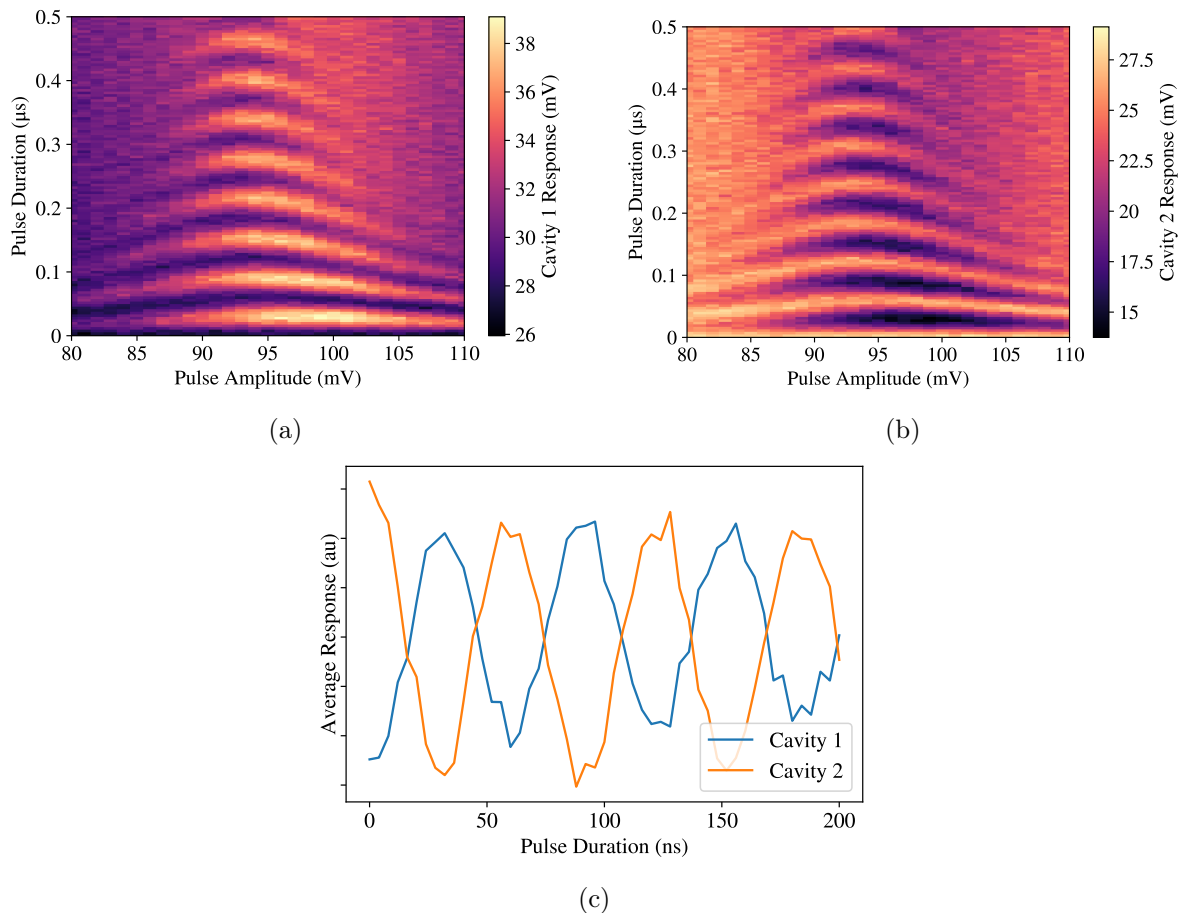


Figure 5.17: (a) [(b)] The response of the resonator coupled to qubit 1 [2] after the qubits have been tuned into resonance by a flux pulse for some duration. The qubits have been initialised in $|10\rangle$ and the excitation can be seen to swap between them when they are on resonance. This is shown more clearly in (c) where a high averaged cavity response corresponds to low $P_{|1\rangle}$ and vice versa for each qubit. All data was taken as part of the same experiment with (a) and (b) taken simultaneously and (c) showing a cut at amplitude 95mV. The amplitude is given in Volts as it is a voltage applied to a ‘flux line’ which induces a flux through the loop of qubit 2, changing its frequency for the duration of the pulse.

implemented which together with single-qubit gates is sufficient to complete a universal gate set. It is also possible to facilitate this interaction via a virtual photon in a cavity bus which has the same coupling Hamiltonian H_c but with

$$J = \frac{J_1 J_2}{2} \left(\frac{1}{\Delta_1} + \frac{1}{\Delta_2} \right), \quad (5.30)$$

where $J_{1(2)}$ is the coupling between qubit 1(2) and the bus as given by Eq. (5.29) and $\Delta_{1(2)}$ is the detuning between qubit 1(2) and the bus.

In practice the 02 transition of the qubits is more often used to implement a CZ gate

[168] by using the fact that the avoided crossing between $|11\rangle$ and $|02\rangle$ is reached before that of $|01\rangle$ and $|10\rangle$ so pulsing to near the $11 \leftrightarrow 02$ transition means that the $|11\rangle$ state will pick up a different phase than $|01\rangle$, $|10\rangle$ and $|00\rangle$, thus implementing a state dependent z rotation which for the duration of pulse will implement a CZ.

It is then necessary to find the fidelity of two-qubit gates in the same way as we found the fidelity of single-qubit gates. As the system size grows it necessarily becomes more challenging to characterise and tune up so there is a persistent need for development in algorithms and metrics to calibrate and report on the performance of many qubit systems.

5.5 Summary

This has been a simple introduction to the transmon qubit and is by no means a complete review but it is sufficient to describe the background we will be using in the following chapters. We have established how to use the nonlinearity of a Josephson Junction to isolate the lowest to levels of an anharmonic LC oscillator. Coupling it to a harmonic resonator facilitates dispersive QND readout and isolates the qubit from the environment. We can manipulate the qubit using microwave pulses applied through capacitive coupling which in a rotating frame execute rotations around the x , y and z axes of the Bloch sphere. Two-qubit interactions are enacted by direct capacitive coupling or via virtual photons in a cavity bus, and can be used to implement two-qubit gates. Outstanding challenges include scaling this system up both in terms of circuit depth capacity and the number of qubits which limits the width of the circuits executable on the hardware. We also face the challenge of having to work with increasingly complex systems which need to be tuned up and characterised efficiently.

The following chapters will address these challenges in more detail and our attempts to overcome them.

Chapter 6

Connectivity: Voltage-controlled superconducting quantum bus

Having established the basics of how a superconducting qubit works, we now turn to the practicalities of combining them to build a quantum computer. One of the main considerations to take into account is how qubits are coupled to one another. In the algorithms of Chapter 3 we assumed all-to-all coupling of the qubits when calculating resources to implement an algorithm. This also has implications for Chapter 4 where we showed that in many cases the Trotter error does not grow with system size but this may no longer be the bottleneck if the low connectivity means that the number of gates required to implement interactions between physically distant qubits limits our algorithms to shallow depth. In this chapter we will explore the implementation of a tuneable coupler intended to make coupling between qubits controllable and flexible for the gatemon qubits described in Chapter 5. The text and figures are taken in almost their entirety from Ref [6]. The fabrication and simulations of this device were done by Lucas Casparis and the device was measured by the two of us.

6.1 Introduction

A significant challenge to scaling any quantum processor architecture is controlling interactions between qubits for multi-qubit operations. Couplings between superconducting qubits are commonly controlled by detuning their transition frequencies [161, 169]. In this way, instead of changing the qubit-qubit coupling constant, the effective coupling can be suppressed by making the qubit energies nondegenerate [170]. As superconducting qubits scale to larger networks, however, the increasingly crowded spectrum of qubit transition frequencies will make this approach prohibitively difficult. Increased

frequency crowding makes residual couplings harder to suppress. Moreover, rearranging qubit frequencies, as is likely required during multi-qubit operations, can lead to state leakage, as described by Landau-Zener physics [171–174]. For reasonable device parameters this results in leakage of several percent [175]. On-chip switchable coupling is desirable, since there is a tradeoff between fast two-qubit gates and avoiding state leakage.

Tunable coupling schemes have been realised for nearest-neighbour-coupled flux-tunable qubits [176, 177], as well as fixed frequency qubits [178]. These schemes allow qubits to be isolated for certain operations, for instance frequency retuning or single-qubit rotations, while still enabling fast two-qubit gates. A tunable superconducting microwave resonator has also been proposed for selective qubit coupling [179]. Such an approach has the advantage that a superconducting resonator, acting as a quantum bus, can mediate long-range interactions between superconducting qubits, and also allows increased connectivity between qubits [136, 169, 170, 180]. Experimentally, flux control of resonators has been demonstrated [181, 182] and used to couple superconducting qubits to spin ensembles [183].

While superconducting qubit circuits often use on-chip current lines to generate fluxes for control, the recently introduced gatemon superconducting qubit [57, 128] is based on a voltage tunable semiconductor Josephson junction (JJ). Gatemons therefore allow for operation using voltages, which can be readily screened to minimize crosstalk and are compatible with semiconductor-based cryogenic control logic [184–186]. The advantage of voltage-controlled operation of semiconductor JJs suggests wider applications in a variety of superconducting circuits, such as superconducting field effect transistors (SFETs) [187].

Here, we implement a voltage controllable superconducting resonator - a tunable quantum bus - which is strongly coupled to two gatemon qubits. The bus is terminated by an SFET acting as a switch that allows *in situ* control of the resonator frequency and qubit-resonator coupling. We demonstrate that the coupling between the two gatemons can be switched between ‘on’ and ‘off’ states by controlling the SFET with on/off coupling ratios up to ~ 8 . We also show that when the coupling is turned off, the frequency of one qubit can be tuned through the other with a strong suppression of state leakage. Finally, we investigate switching the tunable bus on nanosecond time scales. Pulsing the coupler has a similar signature to exciting the qubits, albeit with suppressed phase coherence. The underlying mechanism behind this observation remains unclear.




Figure 6.1: Schematic and simulation of the voltage-controlled superconducting quantum bus. (a) Two gatemon qubits are capacitively coupled to a $\lambda/2$ resonator. One end of the resonator can be grounded through a voltage-controlled superconducting switch. Depending on the switch position being open (blue) or closed (orange), the rms voltage along the resonator length is changed, modifying the coupling between qubits by effectively turning the $\lambda/2$ resonator (blue) into a $\lambda/4$ resonator (orange). The dashed and solid orange lines represent the first and second modes of the $\lambda/4$ resonator respectively. (b) Simulated transmission through the feedline coupled to the tunable bus with the superconducting switch either open (blue) or closed (orange), supporting a large critical current (~ 250 nA). The dashed shaded region indicates the range for the operating frequency of the qubits, f_Q . ©2019 American Physical Society, published in [6].

6.2 Background

A schematic of the device is shown in Fig. 6.1a. Both qubits Q_1 and Q_2 are capacitively coupled to a $\lambda/2$ bus resonator with coupling strengths $g_{1(2)} \propto e\beta V_{\text{rms},1(2)}/\hbar$, where β is the ratio of coupling capacitance to total qubit capacitance, and $V_{\text{rms},1(2)}$ is the root-mean-square of the zero-point voltage fluctuations of the resonator at the location

of $Q1(2)$ [136]. With the qubits at the same frequency $f_{Q1}, f_{Q2} = f_Q$, detuned by $\Delta = 2\pi(f_{\text{res}} - f_Q)$ from the resonator frequency f_{res} , the bus-mediated qubit-qubit coupling $g_{12} = g_1 g_2 / \Delta$ [125, 188] can be controlled by changing either Δ or $g_1 g_2$.

An open switch gives a voltage antinode at the qubit end of the resonator (blue in Fig. 6.1(a)), which results in a large $V_{\text{rms},1(2)}$ with the resonator frequency given by the $\lambda/2$ mode, $f_{\text{res}} = f_{\lambda/2}$. With the SFET in this open state, and f_{Q1}, f_{Q2} close to $f_{\lambda/2}$ we expect that the cavity-mediated coupling is turned on. On the other hand, when the switch is closed, a voltage node is enforced at the qubit end of the resonator, with its fundamental mode changing from $\lambda/2$ to $\lambda/4$. This turns off the interqubit coupling by reducing $V_{\text{rms},1(2)}$ and moving the lowest bus modes to $f_{\lambda/2}/2$ and $3f_{\lambda/2}/2$, which are far detuned from the qubit frequencies [179].

We model the tunable bus as a transmission line terminated with an inductive load, given by the Josephson inductance $L_{J0} = \frac{\hbar}{2eI_c}$. The bus is capacitively coupled to a feedline and we calculate the feedline transmission as a function of frequency, shown in Fig. 6.1b. With the switch in the open state (blue in Fig. 6.1b), no current flows in the SFET and we find $f_{\text{res}} \sim 6\text{GHz}$ for the bare resonator, close to the typical qubit frequencies (dashed shaded region in Fig. 6.1b). From transport measurements using similar semiconductor JJs, we estimate that the SFET in the closed state has a critical current $I_{c,\text{closed}} \sim 250\text{nA}$, corresponding to $L_{J0} \sim 1\text{nH}$ (the SFET has five JJs connected in parallel, as discussed in Section 6.3). As expected, the simulation shows two resonances (orange in Fig. 6.1b) at frequencies approaching $f_{\text{res}}/2$ and $3f_{\text{res}}/2$ (dashed black lines in Fig. 6.1b). Using capacitance simulations we designed the qubit-bus coupling to be $g/2\pi \sim 80\text{MHz}$ for both qubits. With the switch in the open state and $\Delta/2\pi \sim 500\text{MHz}$, this results in $g_{12}/2\pi \sim 13\text{MHz}$. In the case where the switch is closed, the suppression of the coupling is determined by both the larger frequency detuning and reduction in $V_{\text{rms},1(2)}$ of the bus modes. We estimate this residual coupling by applying the ‘black box’ quantization formalism [189] and find that typical values for f_Q and $I_{c,\text{closed}} \sim 250\text{nA}$ result in $g_{12}/2\pi \sim 1\text{MHz}$. As the higher and lower bus modes contribute to $g_{12}/2\pi$ with opposite sign [190], we anticipate this residual coupling could be further suppressed by tuning the circuit, such that the two contributions cancel. The direct capacitive coupling between the qubits is estimated to be $< 1\text{MHz}$. A circuit diagram can be found in Appendix C, Fig. C.1.




Figure 6.2: Two-qubit device with switchable quantum bus. (a) Optical micrograph of the two-gatemon device with the $\lambda/2$ bus resonator terminated by a superconducting switch. Each qubit consists of a bar-shaped island and a gated Al-InAs-Al Josephson junction. (b) The superconducting switch consists of five parallel gated semiconducting weak link Josephson junctions controlled by a single gate voltage. (c) Scanning electron micrograph of the five top gated Al-InAs-Al Josephson junctions. ©2019 American Physical Society, published in [6].

6.3 Experimental Realisation and Measurement Techniques

Figure 6.2a shows an optical image of the tunable bus device. The JJs for both the cavity and the qubits are superconductor-semiconductor-superconductor (S-Sm-S) junctions with a few-channel Sm region [135], allowing the Josephson coupling energy E_J to be tuned using a gate voltage that controls the carrier density in the Sm region. The two transmon-type gatemon qubits each consist of a bar-shaped island with a single JJ to ground. The SFET at the end of the tunable bus is made from several gate tunable JJs in parallel as shown in Fig. 6.2b.

The device was fabricated following the recipe described in Ref. [191] and Appendix C, Section C.1. E_C/h of $Q1(2)$ was designed to be $\sim 200\text{MHz}$ with E_J/E_C tuned to 75-90 using the side gate voltage $V_{1(2)}$. To reduce the effective inductance of the bus switch when closed, five parallel JJs were used to form the SFET. As shown in Figure 6.2c, the five junctions were etched into a single wire (blue) and then covered with 15nm of ZrO₂ dielectric (yellow) deposited by atomic layer deposition. The SFET was controlled with a common top gate voltage V_{sw} (red).

The qubits were manipulated using phase-controlled microwave pulses for rotations around axes in the xy plane of the Bloch sphere and voltage pulses on $V_{1,2}$ for rota-




Figure 6.3: Switchable bus and qubit spectroscopy. (a) Normalized transmission as a function of bus drive frequency and V_{SW} . (b) $Q1$ resonance frequency as a function of V_{SW} . The $Q1$ readout resonator response was measured while a qubit microwave drive tone probed the $Q1$ transition frequency. ©2019 American Physical Society, published in [6].

tions around the z axis and fast frequency displacement. Rotations $R_\alpha(\theta) = e^{\pm i\sigma_\alpha\theta/2}$ ($\alpha = x, y, z$). Both xy and z control pulses were applied through each qubit's gate line. Measured lifetimes and inhomogeneous dephasing times of the two qubits were $\sim 4\mu s$ and $\sim 1 - 2\mu s$ respectively, for the bus in both the on and off states. The two qubits were coupled to individual $\lambda/4$ superconducting cavities (with resonant frequencies $f_{C1} \sim 6.87\text{GHz}$ and $f_{C2} \sim 6.80\text{GHz}$). These were then coupled to a common feedline for dispersive readout [139] with a superconducting travelling-wave parametric amplifier used to improve the signal-to-noise ratio [192]. The tunable bus was also coupled to the common feedline allowing an independent measurement of its resonance. The sample was placed inside an Al box, surrounded by a cryoperm shield and mounted at the mixing chamber of a cryogenfree dilution refrigerator with base temperature $\sim 20\text{mK}$ (see Appendix C, Section C.2).

6.4 Results and Discussions

Figure 6.3a shows vector network analyser measurements of the tunable bus resonance as a function of V_{SW} . At large negative V_{SW} , a resonance is observed at $f_{\lambda/2} \sim 5.6\text{GHz}$, which shows a quality factor $Q \sim 2000$, likely limited by internal losses and coupling to a dissipative environment via the SFET gate line. We anticipate that on-chip filtering of the gate line could increase the quality factor of the tunable bus [193]. While the Purcell effect could impose an upper bound on qubit lifetimes [140], for the qubits here with $T_1 \sim 4\mu\text{s}$ and detunings of several hundred MHz this is not a constraint. We attribute the asymmetry in the resonance line shape to impedance mismatch of the feedline input and output [194]. Going to more positive V_{SW} , the bus resonance disappears with some reentrant features indicating a nonmonotonic turn on of the SFET. We speculate that the disappearance of the resonance is due to the measurement excitation populating the bus with photons and thus driving the SFET normal, leading to a highly reduced Q factor. Although affecting our ability to directly track the bus frequency, it should not impact its role as a quantum bus for $Q1$ and $Q2$ as the coupling is mediated through virtual photons [125]. Interaction between the bus and the qubits renormalizes the qubit frequencies, allowing changes in the bus to be indirectly probed by measuring one of the qubits as shown in Fig. 6.3b. The push on f_{Q1} by the bus is given by $\chi_1 = g_1^2/(\Delta_1)$ (white arrow), where $\Delta_1 = 2\pi(f_{\text{res}} - f_{Q1})$. When the SFET is depleted, the qubit frequency is pushed by the resonator with $f_{\lambda/2} \sim 5.6\text{GHz}$. While closing the switch f_{Q1} increased, indicating that either the bus mode is moving up in frequency or g_1 is decreased, or both. We observed a crossing of the readout resonator with the bus resonator at around $V_{\text{SW}} = -0.5\text{V}$, characterized by a stripe in the spectroscopy data where the readout visibility is reduced. Both the continuous change of the qubit frequency and the crossing of a resonance with the readout resonator indicate that the first mode of the $\lambda/2$ resonator (switch open) turns continuously into the second mode of the $\lambda/4$ resonator (switch closed). For $V_{\text{SW}} > -0.5\text{V}$, the qubit frequency is roughly constant, indicating that either f_{res} no longer changes, or that g_1 is suppressed, although we cannot distinguish between these two effects.

Next, we turn to qubit coupling at fixed values of V_{SW} where the coupler is either on or off. We measured the spectrum while tuning $Q2$ into resonance with $Q1$ [Figs. 6.4a and 6.4b]. On resonance, the two-qubit states hybridize due to the bus-mediated coupling. As Fig. 6.4a illustrates, the splitting was small, although clearly nonzero, when the switch is closed. For an open switch the qubit coupling significantly increased, resulting in a larger splitting between hybridized states shown in Fig. 6.4b.

To further investigate the interqubit coupling, we performed experiments in the time domain. The two qubits were detuned by $\sim 400\text{MHz}$ and $Q1$ ($Q2$) was prepared in $|1\rangle$ ($|0\rangle$). A gate pulse was applied for a time τ to bring $Q2$ into resonance with $Q1$, which

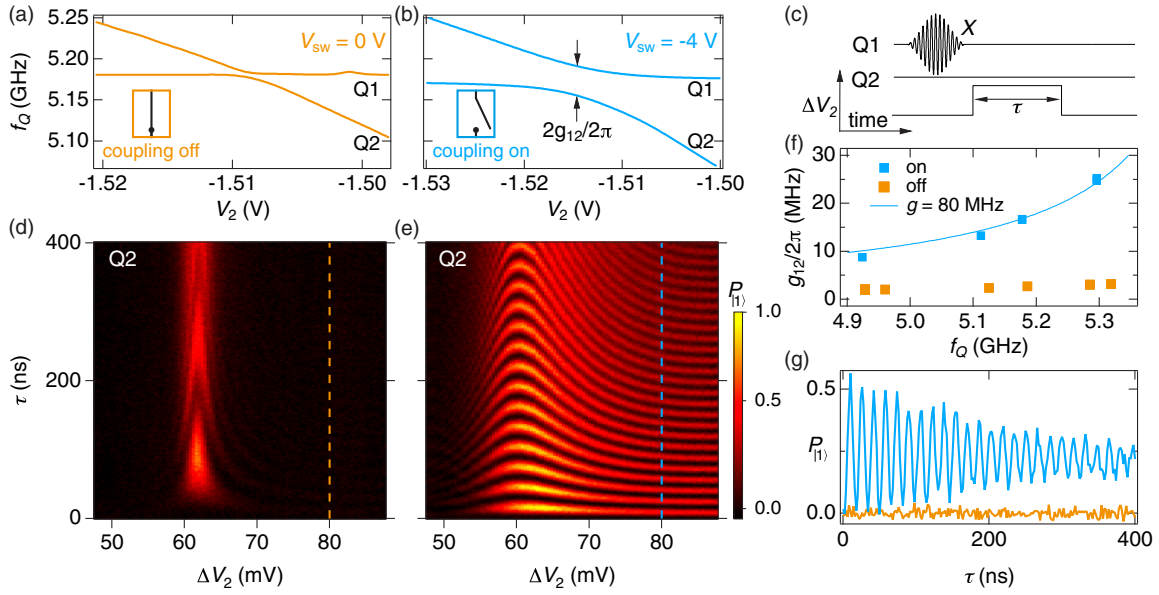


Figure 6.4: Tunable coherent gatemon coupling. (a) [(b)] Measurement of the avoided level crossing between $Q1$ and $Q2$ for the switch closed [open], corresponding to gatemon coupling off [on] as a function of the qubit drive and V_2 . (c) Pulse sequence to probe the coherent coupling between the qubits. With $Q1$ and $Q2$ detuned, $Q1$ is prepared in $|1\rangle$ and $Q2$ in $|0\rangle$. A square gate pulse with amplitude ΔV_2 is turned on for a time τ and brings $Q2$ close to or in resonance with $Q1$. (d) [(e)] The $|1\rangle$ state probability, $P_{|1\rangle}$, for $Q2$ as a function of ΔV_2 and τ for qubit coupling off [on]. (f) Extracted gatemon coupling strengths for on and off case as a function of qubit resonance frequency. The solid line is a fit to the function $g_{12}^{on} = g^2/\Delta$. (g) Cuts along the dashed lines in (d) and (e) at $\Delta V_2 = 80\text{mV}$. ©2019 American Physical Society, published in [6].

can be seen in Fig. 6.4c. Depending on τ and the pulse amplitude ΔV_2 elementary excitations swap between the two qubits. Figure 6.4d shows the swap oscillations with the coupler off and from sine fits to the oscillations, an interaction rate $g_{12}^{off}/2\pi \sim 3.2\text{MHz}$ is extracted, consistent with the avoided crossing measured in spectroscopy. With the coupler on, we observed significantly faster swap oscillations (Fig. 6.4e) and extract $g_{12}^{on}/2\pi \sim 18\text{MHz}$.

Figure 6.4f plots the gatemon coupling strength extracted from swap oscillations as a function of qubit frequency. As expected, g_{12}^{on} (blue) depended strongly on the detuning from the bus. Assuming $g_1 = g_2 = g$ and fitting the data to $g_{12}^{on} = g^2/\Delta$ yields $g/2\pi \sim 80\text{MHz}$. We measured a residual off state coupling $g_{12}^{off}/2\pi \sim 2 - 4\text{MHz}$, limiting the maximum on/off coupling ratio observed in this experiment to ~ 8 . While a larger than anticipated L_{J0} (due to a smaller $I_{c,\text{closed}}$) might explain the residual coupling, our model gives an upper bound of $\sim 1\text{MHz}$ for this coupling after accounting for the observation that the higher bus mode crosses the $Q1$ readout resonator. The dominant contributor to this residual coupling might then be spurious chip modes. Such modes

could be suppressed through more careful microwave engineering, for example, by using air bridges [195].

Figure 6.4g shows cuts from Figs. 6.4d and 6.4e where the $Q1$ frequency crossed through the $Q2$ frequency and then back with the coupler either on or off. These data illustrate that even a modest switching ratio $g_{12}^{\text{on}}/g_{12}^{\text{off}} \sim 6$ allows both strong suppression of state leakage when the coupler is off and fast swaps when on. For a double passage Landau-Zener-Stückelberg process [175], a maximum state leakage of $\sim 50\%$ in the on state (blue) indicates a level velocity of ~ 80 MHz/ns. Since the level velocity is the same for both coupler states, one can estimate a maximum state leakage of $\sim 2\%$ in the coupler off state, comparable to the measurement noise here.




Figure 6.5: Fast switch response. (a) Resonance frequency of $Q1$ as a function of V_{SW} , extracted from Fig. 6.3(b). (b) Ramsey pulse sequence to probe the fast response of the switch inserting a fast Gaussian switch pulse ($\sigma = 64$ ns) with amplitude V_{SW} between two Ramsey pulses. (c) $P_{|1\rangle}$ as a function of V_{SW} and delay τ . ©2019 American Physical Society, published in [6].

Finally, we investigated dynamic operation of the switch by pulsing V_{SW} . Figure 6.5a shows the change of the qubit frequency f_{Q1} while controlling the bus. Again, f_{Q1} is pushed down at large negative V_{SW} . We probed the effect that a fast voltage pulse on the switch has on $Q1$ through a Ramsey experiment. Two $X/2$ pulses were interleaved with a voltage pulse of the SFET gate in Fig. 6.5b. The Ramsey experiment is sensitive to the Lamb shift induced qubit frequency change. Sitting at a dc offset $V_{SW}^0 = -0.4$ V, for $V_{SW} > 0$ V the Ramsey fringes remained roughly constant, as f_{Q1} does not change, shown in Fig. 6.5c. At high pulse amplitudes the visibility of the fringes was reduced,

indicating reduced qubit coherence. We speculate that above certain amplitudes charge traps in the gate dielectric are excited and only relax on time scales comparable to the Ramsey experiment, causing decoherence, though further experiments would be needed to verify this. While applying negative pulses ($\Delta V_{\text{SW}} < 0\text{V}$) to change the qubit coupling on a fast time scale, f_{Q1} was lowered, reducing the period of the Ramsey fringes. For the negative pulses above a certain critical amplitude, $\Delta V_{\text{SW}}^c \sim -1.1\text{V}$, the readout response suggests that $Q1$ is excited into the $|1\rangle$ state and phase coherence is lost. The origin of this effect is presently unclear. We find that the value of ΔV_{SW}^c depends on both V_{SW}^0 and the shape of the switch pulse. It was observed that the time scale on which the qubit can be coherently manipulated after a switch pulse is somewhat shorter than the decay time of the qubit, possibly indicating a different mechanism than qubit excitation like impairment of the readout resonator. Similar effects have been observed in two other samples: one device identical to that presented here, and the other using a $\lambda/4$ switchable resonator as the quantum bus. We speculate that pulsing the SFET close to depletion nonadiabatically excites the qubit circuit [196]. Another possibility is that pulsing the SFET JJs towards depletion generates quasiparticles that induce decoherence [197].

6.5 Conclusions and Outlook

In summary, we have demonstrated a voltage-tunable superconducting quantum bus that can control the coherent coupling between two gatemons. The number of qubit pairs coupled through the tunable resonator could readily be increased, allowing for larger connectivity. This could be of interest for qubit architectures beyond the surface code geometry [198]. The continuously tunable coupling might also prove attractive for quantum simulation [199]. While dynamic operation of this voltage-controlled bus remains an outstanding problem, the potential advantages of this approach for coupling qubits motivates further investigations. Moreover, recent work integrating low loss microwave circuits with proximitized two-dimensional electron gases that support a wide range of critical currents provides an ideal platform to explore such voltage-controlled coupling schemes [6].

6.6 Acknowledgements

We acknowledge helpful discussions with S. Nigg and V. Shumeiko. This work was supported by Microsoft Project Q, the U.S. Army Research Office, and the Danish National Research Foundation. N.J.P. acknowledges support from the Swiss National

Science Foundation and NCCR QSIT. F.K. acknowledges support from the Danish Innovation Fund and C.M.M. acknowledges support from the Villum Foundation. The travelling-wave parametric amplifier used in this experiment was provided by MIT Lincoln Laboratory and Irfan Siddiqi Quantum Consulting (ISQC), LLC, via sponsorship from the U.S. Government.

Chapter 7

Scalability: Superconducting gatemon qubit based on a proximitized two-dimensional electron gas

Having explored an option for increasing the connectivity of the gatemon qubit we now turn to the challenge of a scalable implementation. The relevant inductance is either set by tailoring the metal oxide dimensions of single JJs, or magnetically tuned by parallelizing multiple JJs in superconducting quantum interference devices (SQUIDs) with local current-biased flux lines. JJs based on superconductor-semiconductor hybrids represent a tantalizing all-electric alternative in the gatemon [57, 128]. However the traditional fabrication of the gatemon uses VLS nanowires for the JJ which is a labour intensive process as placing nanowires is typically done by hand. In this chapter we demonstrate that semiconducting channels etched from a wafer-scale two dimensional electron gas (2DEG) are a suitable platform for building a scalable gatemon based quantum computer. We show that 2DEG gatemons meet the requirements [116] by performing voltage-controlled single-qubit rotations and two-qubit swap operations. We measure qubit coherence times up to $\sim 2\mu\text{s}$, limited by dielectric loss in the 2DEG substrate. The contents of this chapter are adapted from Ref. [7]. The fabrication was done by Lucas Casparis and Malcolm Connolly while I was jointly responsible for measuring the devices.

Figure 7.1 from Ref. [200] shows a schematic depiction of the different kinds of JJs commonly used to make transmons and the type of control necessary for each. This effectively visualises the differences between them and helps to contextualise this research within the field.




Figure 7.1: Transmon and gatemon qubits. (a) Schematic depiction of a 1D array of superconducting qubits. The qubits (blue T-shaped structures) couple capacitively to nearest neighbours, and control lines approach the qubits from below. This layout is compatible with both transmons and gatemons. (b) For transmons, the qubit frequency is tuned via a magnetic flux Φ through a superconducting control loop. This flux is applied through an on-chip current bias. (c) A semiconductor nanowire placed between two superconducting leads can be used as a voltage-tunable circuit element, which forms the basis of nanowire gatemons. (d) Another method to realize voltage-tunable circuit elements is to etch away a narrow strip of superconducting metal that sits on top of a wafer-scale 2DEG. Reprinted by permission from Springer Nature: Nature Nanotechnology, Gatemons get serious by Steven J Weber [200], ©2020.

7.1 Introduction

Fixed-frequency transmons that employ single metal-oxide JJs benefit from longer coherence times, but at the cost of slow (~ 150 ns) two-qubit gate operation times [201] and frequency crowding [202]. Frequency-tunable qubits allow faster two-qubit gates, but the susceptibility to flux noise in SQUIDs results in typical phase coherence times of $T_2^* \sim 5\mu\text{s}$ [202, 203]. Moreover, the milliamperic currents used to control the flux in the SQUIDs place additional demands on cooling power, introduce crosstalk and may complicate the integration with 3D architectures [204, 205]. In superconductor-

semiconductor JJs, E_J can be controlled by local capacitively coupled gates [57, 128, 206], which opens up the possibility to tune and modulate f_Q without the need for current-biased flux lines. Although this eliminates sensitivity to flux noise, charge fluctuations can couple to the qubit frequency via the gate electrode and dielectric substrate near the JJ. Recent experiments on nanowire (NW) gatemons demonstrated relaxation and coherence times that exceed $20\mu\text{s}$ and $4\mu\text{s}$, respectively [191, 207], which mitigates this concern in the near term. More immediate challenges to building NW-based architectures, however, are viable ways to scale and the upper limit on E_J imposed by their $\sim 100\text{nm}$ dimensions.

We leverage recent breakthroughs in using *in situ* epitaxy of Al (blue) on III-V semiconductors [208] to obtain a pristine high-transparency superconductor-semiconductor interface between a 50nm thick layer of superconducting aluminium and the 2DEG. Superconducting correlations are induced in the 2DEG by electron-hole (e-h) conversion at this interface, a process known as Andreev reflection. A JJ is formed by selectively etching away a narrow strip of the Al top layer to leave proximitised banks on either side of a normal region (see Appendix D for more fabrication details). Cooper pairs traverse this normal region via e-h pairs that repeatedly reflect from each side and form bound states between the proximitised 2DEG. E_J is modulated by the voltage applied to the top gate, which changes the carrier density in the normal region of the junction.

Figure 7.2a shows an optical micrograph of a typical device that hosts six 2DEG gatemon qubits. Each gatemon comprises an Al island shunted to the ground plane via a 2DEG JJ and capacitively coupled to a serpentine-shaped coplanar waveguide cavity. The self-capacitance C of the island together with the nonlinear inductance of the JJ creates an anharmonic potential for plasmon oscillations across the JJ. The ground $|0\rangle$ and excited $|1\rangle$ states of the qubit correspond to the lowest two harmonic oscillator states, which in the transmon limit ($E_J \gg E_C$) are separated in energy by a transition frequency $f_Q \approx \sqrt{8E_C E_J}/h$, where $E_C = e^2/2C$ is the charging energy and E_J is the Josephson energy [51, 126]. The 2DEG heterostructure is shown schematically in Fig. 7.2b. The 2DEG is formed in an InAs quantum well (green) encapsulated between InGaAs barriers (light grey). Figure 7.2c shows a false-coloured scanning electron micrograph of the JJ.

7.2 Reliable fabrication

First, we demonstrate that, unlike NW gatemons, 2DEG gatemons can be fabricated deterministically with a superior reproducibility. We fabricated three devices (S1, S2 and S3) each of which hosted six qubits with the junction width w increasing from 0.3 to $2.6\mu\text{m}$ (labelled Q1-Q6). To extract the E_J of the as-fabricated qubits, the corresponding cavity frequency (f_c) is measured before any voltage is applied to the gate. Due to

vacuum fluctuations in the electric field between the cavity and qubit, the cavity is Lamb shifted from its bare resonance frequency by $\chi = g_{\text{cav}}^2/\Delta_{\text{Q}}$, where $\Delta_{\text{Q}} = f_{\text{c}} - f_{\text{Q}}$ is the qubit-cavity detuning. The cavity coupling strength, $g_{\text{cav}}/2\pi \approx 100\text{MHz}$, is extracted from qubit spectroscopy in the dispersive limit (details are given below) and, together with numerical simulations for C ($E_{\text{C}}/h \approx 230\text{MHz}$), allows us to estimate E_{J} . Figure 7.3a plots the extracted E_{J} as a function of w for all the measured qubits. The data show that E_{J} increases for wider junctions, as expected with an increasing number of modes that participate in Cooper pair transport [135]. We omit the data from the widest JJs of $2.6\mu\text{m}$ ($Q6$), as the Lamb shift is too small to extract E_{J} . Such a precise control of E_{J} on a design parameter w represents an important step towards engineering scalable superconductor-semiconductor quantum information processors, improving on previous realizations, where w was limited by the 1D character of NWs [209].

7.3 Single-Qubit Control and Characterisation

Next, we show all-electric control by tuning the qubit transition frequency in Fig. 7.3b. We operate in the transmon regime, $E_{\text{J}}/E_{\text{C}} \approx 70 - 130$, and read out the qubit dispersively ($g_{\text{cav}} \ll |\Delta_{\text{Q}}|$) [136]. Using two-tone spectroscopy, we drive a single qubit ($Q3$) and identify its frequency as a function of gate voltage from the state-dependent push on the cavity. The frequency $f_{\text{Q}3}(V_3)$ is monotonic over a wider voltage range than for NWs [57, 191] and can be tuned by $\Delta f \approx 1\text{GHz}$ for 1V applied to the gate (V_j corresponds to the voltage applied to j th qubit Q_j). The dependence of qubit frequency on gate voltage can be optimized by changing the thickness of the dielectric layer and using 2DEGs with different field-effect mobilities. Higher-mobility 2DEGs might also allow further improvements to the reproducibility of the junction characteristics, and so simplify the control of larger-scale devices.

We next demonstrate the basic operations of individual qubits using time-domain manipulation and readout. Phase-controlled microwave pulses with drive frequency f_{d} are applied either via the cavity readout feedline or separately through the JJ top gate. The rotation about the x axis of the Bloch sphere is performed by applying the pulse for a time τ_{Rabi} and reading out the state via the cavity (pulse sequence, Fig. 7.3c). Plotting the probability to be in $|1\rangle$, $P_{|1\rangle}$, as a function of τ_{Rabi} and f_{d} , reveals Rabi oscillations (Fig. 7.3c), characteristic of the qubit rotation. These data are used to calibrate the pulse times and amplitudes to rotate by π and $\pi/2$ around the x axis (X and $X/2$ pulses, respectively). We next show the coherent accumulation of the dynamical phase by a controlled rotation of the qubit around the z axis. Figure 7.3d shows the pulse sequence that comprises a resonant ($f_{\text{d}} = f_{\text{Q}}$) $X/2$ pulse, a gate pulse with amplitude ΔV_3 and duration τ_{Ramsey} and a second $X/2$ pulse. When $\Delta V_3 = 0$ the qubit and drive are phase

locked, so the two $X/2$ pulses rotate the qubit to the $|1\rangle$ state. With increasing ΔV_3 , the qubit rotates around the z axis relative to the drive. Although further experiments, such as randomized benchmarking, are required to establish fidelities, these data establish the high degree of control afforded by electrostatically coupled gates.

To measure the relaxation time, T_1 , an X pulse excites the qubit (Fig. 7.4a, blue pulse sequence) and $P_{|1\rangle}$ is plotted as a function of τ , the time delay before readout. The probability decreases exponentially due to relaxation. Fitting the decay (blue) yields $T_1 = 1.1\mu\text{s}$. Assuming the qubit lifetime limits the precision of single-qubit gates, for a microwave pulse time of 20ns we estimate an upper bound of 99.4% for the fidelity [210]. To extract the dephasing time T_2^* , two slightly detuned $X/2$ pulses are applied (Fig. 7.4a, black pulse sequence), separated by the delay time τ . A fit to the decay of the resulting Ramsey fringes (Fig. 7.4a inset) gives a dephasing time of $T_2^* = 400\text{ns}$. To reduce the inhomogeneous dephasing due to low-frequency noise, we perform a Hahn echo sequence that comprises a refocusing X pulse between two $X/2$ pulses (Fig. 7.4a, red). The extracted $T_{2,\text{echo}} = 2.2\mu\text{s} \approx 2T_1$ indicates that 2DEG gatemon dephasing is dominated by low-frequency noise [211]. Figure 7.4b shows T_1 as a function of qubit frequency. Relaxation times vary between 0.2 and $2\mu\text{s}$ and fluctuate strongly with f_Q (the spectrum is plotted in the inset Fig. 7.4b). Owing to their periodicity, we attribute these fluctuations to on-chip modes, which is consistent with previous results from devices that lack crossover wire bonds.

An estimate for the dielectric loss of the qubit capacitor can be made using a test resonator coupled to the same feedline ($f_{\text{res}} = 5.35\text{GHz}$), which shows an internal quality factor $Q \approx 6.4 \times 10^4$ at a low photon number. Using the expression $T_1 = Q/(2\pi f_Q)$, we expect the relaxation time due to dielectric loss to follow the black dashed line in Fig. 7.4b [139]. The agreement between the measured T_1 times and this upper bound suggests the qubit lifetime is, indeed, limited by dielectric loss. Similar Q values are obtained on pure semi-insulating InP substrates, which suggests that the presence of the 2DEG does not introduce additional loss. Test resonators on GaAs and GaSb substrates showed significantly lower Q values, consistent with previous reports [212, 213]. Suitable solutions to reduce microwave loss compatible with InP-based 2DEGs include deep etching trenches [214] and flip-chip techniques used to host the qubit island on a low-loss substrate such as Si [205]. From the measured slope df/dV of 1GHzV^{-1} and $T_2^* = 400\text{ns}$, we estimate an equivalent root-mean-squared (r.m.s.) gate voltage noise of $\sim 0.5\text{mV}$ [215]. Previous r.m.s. gate-voltage fluctuation measurements of $\sim 8\mu\text{V}$ in III-V semiconductors [215] suggest that $T_2^* \sim 25\mu\text{s}$ could be achieved. This implies that the prospects are good for obtaining coherence times comparable with state-of-the-art flux-tunable transmons for which $T_1 \approx 30\mu\text{s}$ and $T_2^* \approx 5\mu\text{s}$ [165, 202, 203].

7.4 Two-Qubit Interaction

Finally, we coherently swap excitations using the capacitive coupling between neighbouring qubits. The monotonicity over a range of 0.5GHz of the qubit spectra established in Fig. 7.3 is useful for tuning adjacent qubits into resonance with each other. The signature of qubit-qubit coupling is a mutual push on the bare qubit frequencies due to hybridization. To detect this push, the qubits $Q2$ and $Q3$ are driven and read out through the feedline and their respective cavities. For clarity, the signals detected from both cavities are added to yield the sum V_Σ . Figure 7.5a shows V_Σ as a function of the qubit drive and V_2 . As expected, due to the absence of crosstalk, there are two peaks in V_Σ as a function of f_d , only one of which ($Q2$) is tuned by V_2 . When tuned onto resonance, the qubits anticross and a splitting of $2g/2\pi \approx 12\text{MHz}$ between the two hybridized states is observed, where g is the qubit-qubit coupling strength. Figure 7.5b shows the pulse sequence that exploits the anticrossing to transfer coherently an excitation between $Q2$ and $Q3$, the starting point for preparing arbitrary two-qubit states. With the two qubits detuned by $\sim 140\text{MHz}$ and $Q3$ idling, $Q2$ is prepared in $|1\rangle$. A gate pulse is then applied for time τ_{swap} and brings $Q2$ into resonance with $Q3$ [170]. Note that here the microwave pulses are applied through the gate line, which demonstrates qubit manipulation using a single control line per qubit. We emphasize that such single-gate control of rotations around the x , y and z axes is an important advantage of voltage-controlled qubits. The probability that an excitation swaps between the qubits depends on τ_{swap} and the pulse amplitude ΔV_2 . Figure 7.5c shows the typical chevron pattern of swap oscillations [216]. The lower panel in Fig. 7.5c shows $P_{|1\rangle}$ for each qubit separately. The anticorrelation confirms that the excitation transfers between $Q2$ and $Q3$ and demonstrates the possibility of generating entangled states using 2DEG gatemon qubits. From sinusoidal fits (Fig. 7.5b, solid lines) an interaction rate of $2g/2\pi = 14\text{MHz}$ is extracted, in good agreement with electrostatic simulations that yield $2g/2\pi = 15\text{MHz}$ for $f_Q = 5\text{GHz}$.

In summary, we have demonstrated that planar semiconductor materials and superconducting microwave circuits are compatible technologies that can be readily integrated while they maintain quantum coherence. This opens new possibilities for highly integrated quantum processors with on-chip components. Through a combination of geometry and applied voltages, E_J can be tailored to simultaneously suit qubits and peripheral control circuits that require a higher E_J , such as tunable couplers [177] [6] and on-chip microwave sources [217], and develop naturally into the 3D architectures required to implement fault-tolerant processing [205, 218]. Moreover, as 2DEG gatemons represent a perfect quantum counterpart to semiconductor-based cryogenic classical control logic [184–186], they take the first step towards realizing a scalable all-electric hybrid superconductor-semiconductor quantum processor.

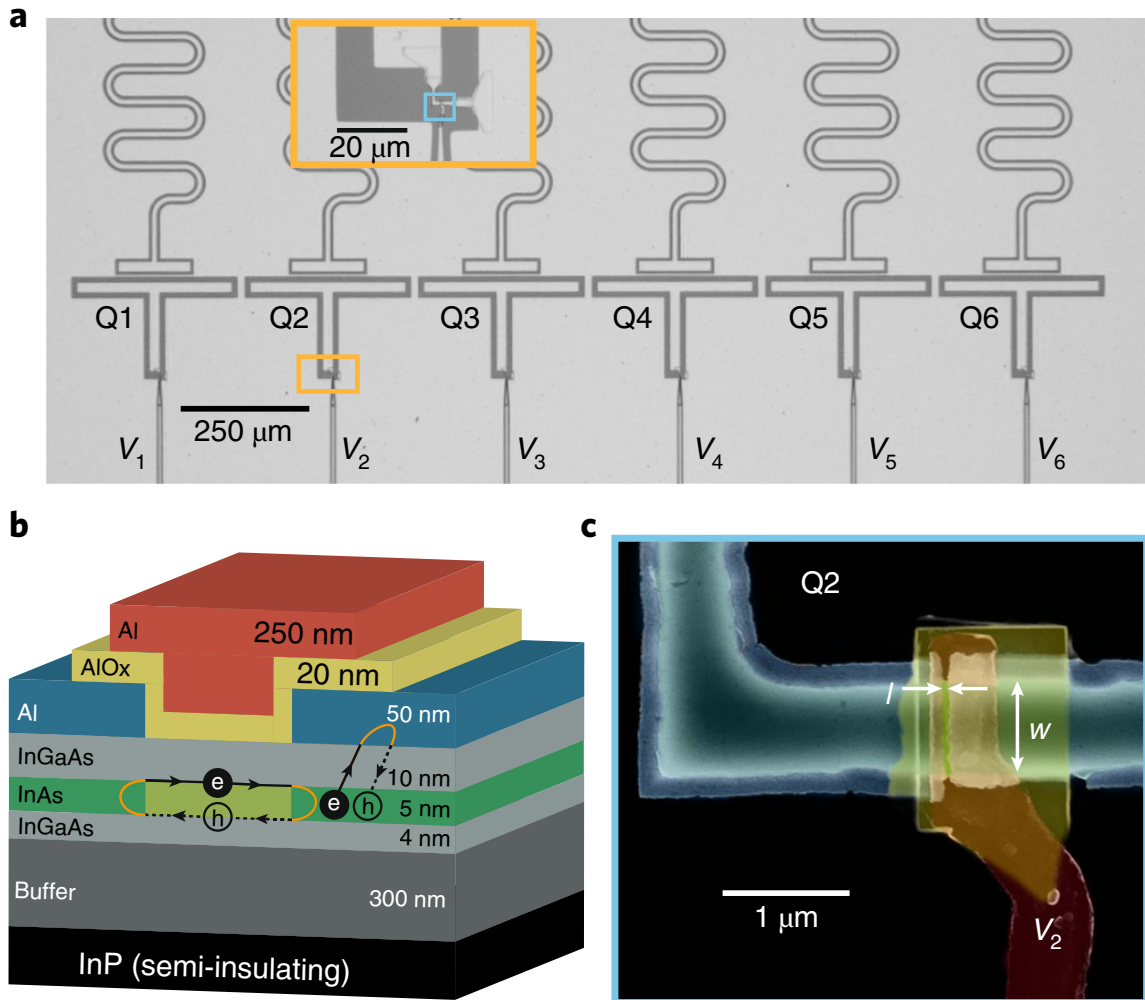


Figure 7.2: (a) Optical micrograph of a six qubit device. The 2DEG JJ is shunted by the T-shaped island to the surrounding ground plane and coupled to individual readout cavities. The gate voltage V_j changes the qubit frequency of Q_j . (b) Schematic of the wafer stack. The InAs quantum well (green) is proximitized by Andreev reflection (orange) at the highly transparent interface with the Al (blue). Cooper pairs traverse the JJ (light green) by Andreev bound states confined between the proximitized 2DEG regions under the Al (dark green). (c) False-coloured scanning electron micrograph of the gate-controlled 2DEG JJ of width w and length l . At the edge of the mesa the Al overhangs due to an undercut when etching the semiconductor buffer. ©2018 Springer Nature: Nature Nanotechnology, published in [7].

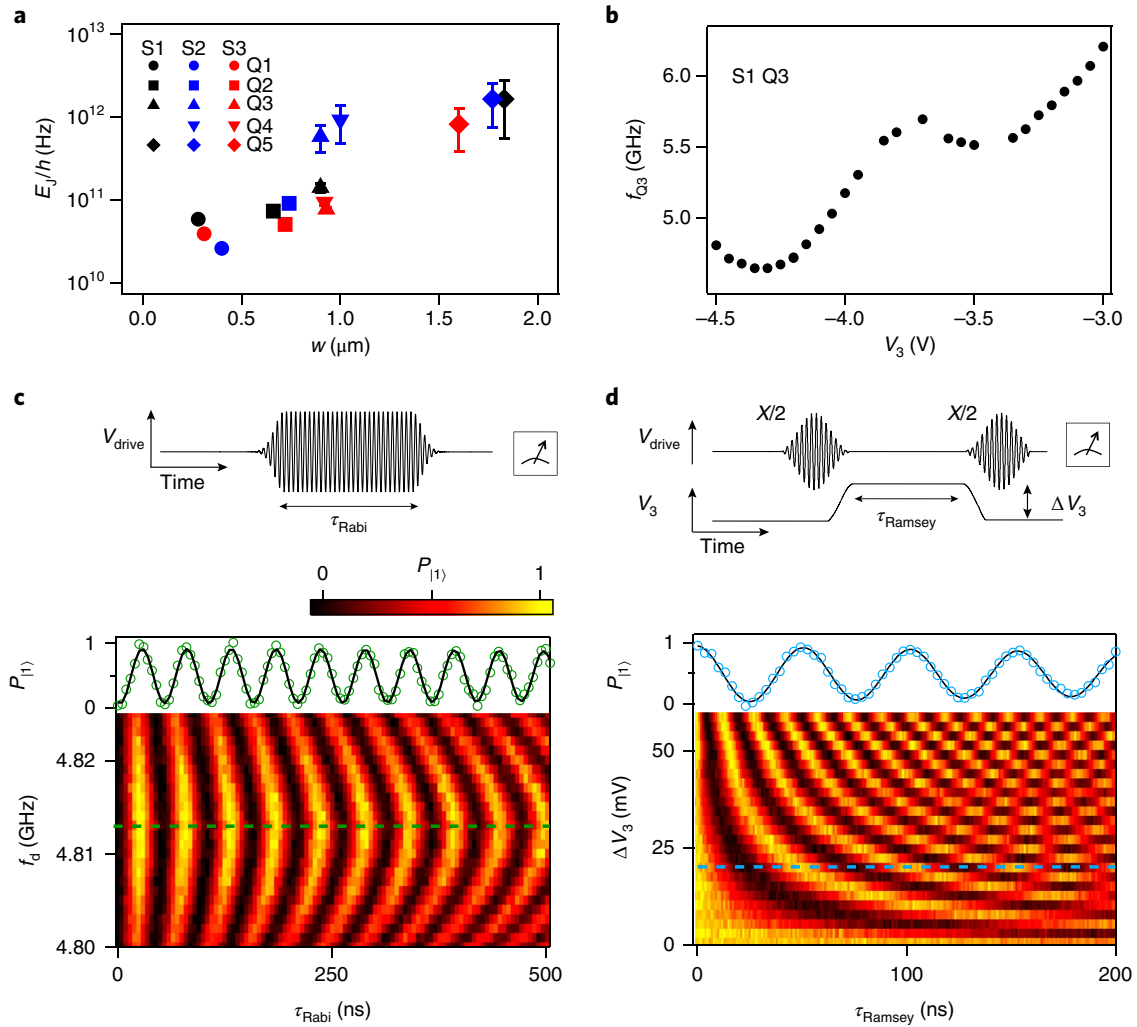


Figure 7.3: 2DEG gatemon. (a) Overview of E_J as a function of w for three devices, S1, S2 and S3, at zero gate voltage. (b) Frequency of S1 Q3 as a function of gate voltage. (c) Coherent Rabi oscillations performed at $V_3 = -4.5\text{V}$ by applying the microwave pulse sequence shown in the upper panel. The main panel shows qubit oscillations as a function of the qubit drive frequency f_d and τ_{Rabi} , and the inset shows a cut at the resonance frequency. The solid line is a fit to the data using a Gaussian-damped sinusoid. (d) Coherent qubit rotations around the z axis. The qubit is positioned on the equator with an $X/2$ pulse followed by a gate pulse with amplitude ΔV_3 and duration τ_{Ramsey} and finally rotated back by an $X/2$ pulse (upper panel). The main panel shows the coherent Z oscillation as a function ΔV_3 and τ_{Ramsey} , and the inset shows a cut at $\Delta V_3 = 20\text{mV}$. ©2018 Springer Nature: Nature Nanotechnology, published in [7].

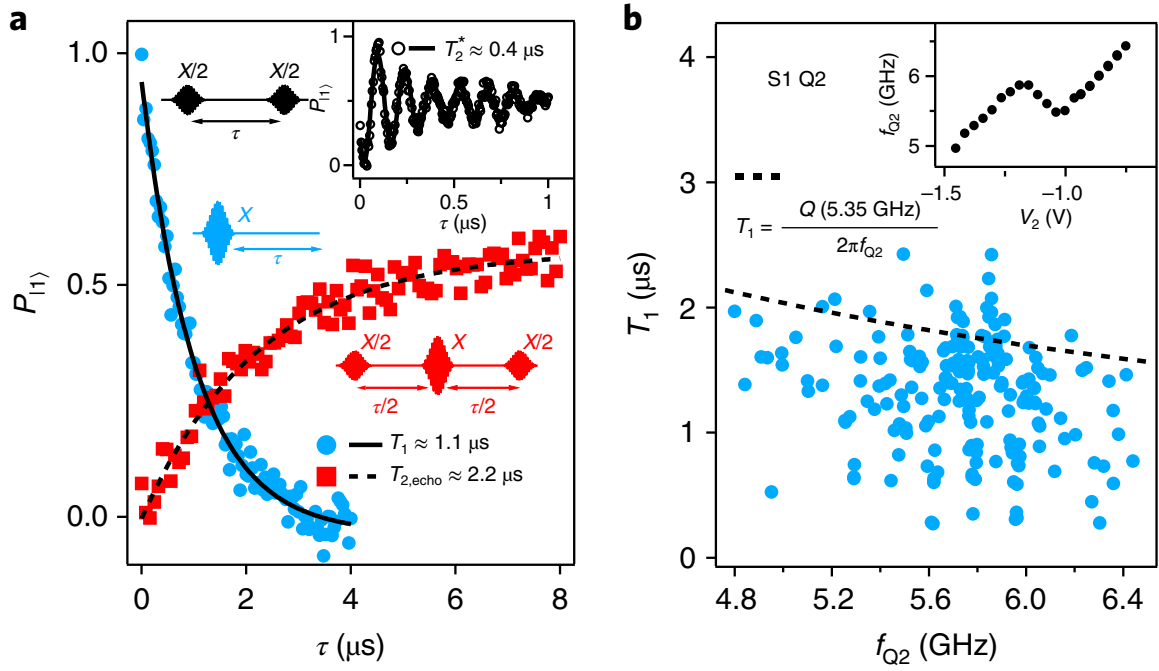


Figure 7.4: Coherence times. (a) Lifetime measurement for S1 Q2 (blue) with qubit resonance frequency $f_{Q2} \approx 5\text{GHz}$. We performed a Hahn echo experiment to determine $T_{2,\text{echo}}$ (red). The black lines (solid and dashed) are exponential fits. The pulse sequences for dephasing ($T_{2,\text{echo}}$) (red) and relaxation (T_1) (blue) measurements are shown next to their respective data. The inset shows a Ramsey experiment which was performed to determine T_2^* for Q2 with the pulse sequence shown next to the inset (black). The solid line in the inset is a fit to the data using an exponentially damped sinusoid. (b) Relaxation time measurements as a function of qubit frequency. The dashed line indicates the limit on the qubit lifetime for a quality factor $Q \approx 6.4 \times 10^4$. The inset shows the spectrum for Q2. ©2018 Springer Nature: Nature Nanotechnology, published in [7].




Figure 7.5: Coherent two-qubit interaction. (a) Measurement of the avoided level crossing between $Q2$ and $Q3$. The sum of the normalized heterodyne readout amplitude V_Σ for both qubits is shown as a function of qubit drive and V_2 . (b) Pulse sequence to probe the coherent coupling between the qubits. With $Q2$ and $Q3$ detuned, $Q3$ is prepared in the ground state and an X pulse prepares $Q2$ in $|1\rangle$. A gate pulse of amplitude ΔV_2 brings $Q2$ close to or in resonance with $Q3$ for time τ_{swap} . (c) The $|1\rangle$ state probability, $P_{|1\rangle}$, for $Q3$ as a function of ΔV_2 and τ_{swap} . Lower panel: $P_{|1\rangle}$ for both $Q2$ and $Q3$ at $\Delta V_2 \approx 71$ mV, which brings the two qubits into resonance. ©2018 Springer Nature: Nature Nanotechnology, published in [7].

Chapter 8

Calibration and Characterisation: Hamiltonian Learning

Chapters 6 and 7 show a possible pathway to large quantum processors. An increased number of qubits and interconnections require a more sophisticated approach in terms of data taking, data analysis and system modelling in order to implement both analogue and digital simulations. We have seen that a large number of connected qubits will be needed to execute interesting simulations and, in addition to the challenges involved in coupling these and scaling up the system, we must also characterise and tune up each qubit and interaction. This is necessary for our final quantum hardware requirements in order to implement accurate analogue or digital simulations. It is also important to learn about the properties of a device in order to systematically make improvements in both control and fabrication.

In Chapter 5 we discussed the methods commonly used to calibrate an X gate and the optimisation of the rest of the Clifford gates using randomised benchmarking (RB) [48, 156–159]. We also touched upon more rigorous but also more expensive methods using tomography to ascertain the exact operation of a gate on a quantum system [162–164]. However, these methods become increasingly expensive as the size of the system grows. We need to calibrate our system by means of Rabi experiments, Ramsey experiments, readout fidelity optimisation and other tune up procedures such as AllXY [123]. We also need find the values of the metrics readout fidelity, T_1 , T_2^* and $T_{2,\text{echo}}$ in order to characterise our system and report on it. Each of these processes typically requires its own experiment and fitting procedure to learn the parameter of interest and then, in the case of tune-up, optimise it. In some cases we have multiple degrees of freedom, requiring many repetitions of this process and motivating adaptive tune-up [165]. This is the minimal amount necessary to characterise a single qubit where a more thorough analysis would include tomographic analysis as well as including the option for

gate specific errors, time dependence and that of other parameters such as flux/voltage. As such, there is strong motivation to optimise the learning process and be as efficient as possible in data taking.

8.1 Introduction

In our system the initialisation, control and readout of a qubit takes $\sim 100\mu\text{s}$ and we usually need to average on the order of 10^2 to 10^4 times to overcome the metrology limitations and the uncertainty introduced by noise and decay to find the probability that the qubit is excited, $P_{|1\rangle}$, with sufficient confidence. We could reduce the experiment time by a factor of ~ 100 by employing active reset or ‘restless’ tune-up procedure but this itself adds complexity and is nontrivial [165]. Once we have an averaged data point for $P_{|1\rangle}$ which corresponds to one set of experiment parameters we then need to repeat the measurements, varying the relevant parameter(s) and conduct a fit to the result to ascertain the quantity of interest. A simple Rabi experiment thus might take 10s excluding setup procedures such as pulse sequence upload and data analysis. In theory this experiment can be carried out simultaneously on different qubits but for the corresponding two-qubit tune-up experiments the parallel options will be reduced. Furthermore in the scenario where crosstalk is present, and must be measured and accounted for, many configurations must be tested sequentially. Each measurement of the qubit state takes $\sim 1\mu\text{s}$ and requires ~ 500 measurements of the resonator as a function of time in order to execute the final mixing down step and ascertain the resonator response and from it the qubit state. This means that in total we must take and analyse $\sim 10^7$ data points, each comprising 12 bits. This comes to 75 Megabytes per Rabi experiment to be analysed and saved, although frequently only the averaged result is saved. This is computationally expensive and is the motivation for hardware solutions which mitigate the need for the software mixing down and developments which will move much of the calculation onto FPGA processors [219]. However it will always be beneficial to find ways to minimise the number of measurements required to learn the parameters of interest.

As explained, system parameters are extracted via numerical fits. The most frequently used method is to find either the average cavity response or $P_{|1\rangle}$ by taking many repetitions for a given set of values of the experiment parameters (examples include τ_{pulse} , τ_{wait} , gate voltage e.t.c.). This is then repeated for a different set of values. The results are then fit using the non-linear least squares fitting method which, given a model, finds the model parameters which minimise the sum of the squared residuals. Specifically the iterative Levenberg–Marquardt algorithm [220, 221] is used to find the local minima of the sum of the squares of deviations. In practice we find this very sensitive to the

initial parameter values provided from which to start iteration. This means that we must, in fact, use functions to generate good initial guesses which is increasingly difficult for more complex models. We can quantify the goodness of fit by calculating the chi-squared value but this is strictly meant to either confirm or deny the validity of the fit beyond a certain probability and in practice we usually visually check that the fit matches the data which is quite clearly not a scalable approach. All of these processes must be automated and executed as efficiently as possible.

The frequentist approach makes no assumptions about prior knowledge which is obviously a very defensible stance to take but, given that the models describing Rabi oscillations, longitudinal and transverse decay and RB decay are well established, we can also justify starting from these models with a relatively well known prior and using Bayesian techniques to update this prior and more quickly find the values of the parameters from fewer data points. Here we implement a Bayesian approach to learn about the properties of a superconducting qubit.

In probability theory Bayes' theorem gives the probability for event A to happen given that event B happens:

$$\begin{aligned} P(A|B) &= \frac{P(B|A)P(A)}{P(B)} \\ &= \frac{P(B|A)P(A)}{\int P(B|A)P(A)dA}, \end{aligned} \tag{8.1}$$

where $P(A)$ is the probability of observing A . Bayesian inference derives the posterior probability as a consequence of the prior probability and a likelihood function derived from the model. In our experiments A is the hypothesis for the model parameter values. $P(A)$ is the prior probability; an estimate of the probability of A without knowledge of B . B is then the data which we use to then find the posterior probability, $P(A|B)$, the probability of our hypothesis of the model parameters given the data. $P(B|A)$ is the likelihood function which tells us the likelihood of observing the data B given a set of model parameters A . The process we implement effectively performs Bayesian updates on a prior in response to new data to give us a posterior distribution from which we find the values of our model parameters and quantify our belief in these values. With each data point we analyse we can report our current model parameter estimate and the variance on that estimate. The probability distributions must be discretised in order to make the Bayesian updates tractable. This is achieved by the sequential Monte Carlo (SMC) approximation which draws samples from the initial prior distribution and then updates them based on Bayes' rule as data is collected [222, 223]

This method is motivated by the observations of Refs. [60, 222, 224] which show that this method is resilient against noise and that it is possible to reach optimal precision of parameters after a small number of trials. Experimental success in learning the

Hamiltonians of a quantum system has been achieved in nitrogen vacancy centres [59] where improvements of a factor 10 to 100 were achieved in the variance of five model parameters by using Bayesian methods to dictate which experimental to take compared to using pre-defined data. Another experiment used Bayesian inference to execute phase estimation using a silicon photonic chip [225] based on proposals from Ref. [226] and another used Bayesian inference to learn the Hamiltonian of a nitrogen vacancy centre via an interface with a silicon photonics quantum simulator [227]. Use of methods for even the most statistics-averse quantum physicist is facilitated by the Qinfer package [228, 229]. We here show that even for learning a only a few model parameters and without implementing the ‘online’ experiment methods of Ref. [59] we already see a large improvement in the amount of data needed to learn about our superconducting qubit system. This exemplifies the potential benefit of implementing non-traditional methods, even at the lowest level of calibration and tune-up and motivates further work to implement it at all levels and more rigorously find the trade off between increased computational complexity of the Bayesian methods and increased data taking in the experiment.

8.2 Single Parameter Inference

To demonstrate this method we apply it to a standard experiment and use it to find the Rabi frequency ω_{Rabi} . This is the frequency at which the qubit will rotate around the x axis if it undergoes a microwave drive at the qubit frequency as detailed in Chapter 5. If the qubit starts in $|0\rangle$ then the probability for it to be in the excited state after time τ_{pulse} is

$$P_{|1\rangle}(t) = 1 - \cos(\omega_{\text{Rabi}}\tau_{\text{pulse}})^2. \quad (8.2)$$

A least squares fit to data averaged 150 times can be seen in Fig. 8.1.

The data in Figure 8.1 is the result of 210×10^3 single shots. Although this number does not represent the optimised value as we could likely complete the fit using fewer averages or taking data with less resolution in pulse duration, it is representative of the scale of data points we would usually take to execute a calibration experiment of this kind, where the others include the T_1 experiment, Ramsey experiment and the $T_{2,\text{echo}}$ experiment explained in Chapter 5. Indeed the number is often considerably higher as we frequently need more averages or to sweep another parameter such as drive power or frequency to complete the full experiment. Frequently the number of single shots taken is of order 10^6 and crucially, although these numbers are not optimised, there is also no rigorous way to do so before taking the data.

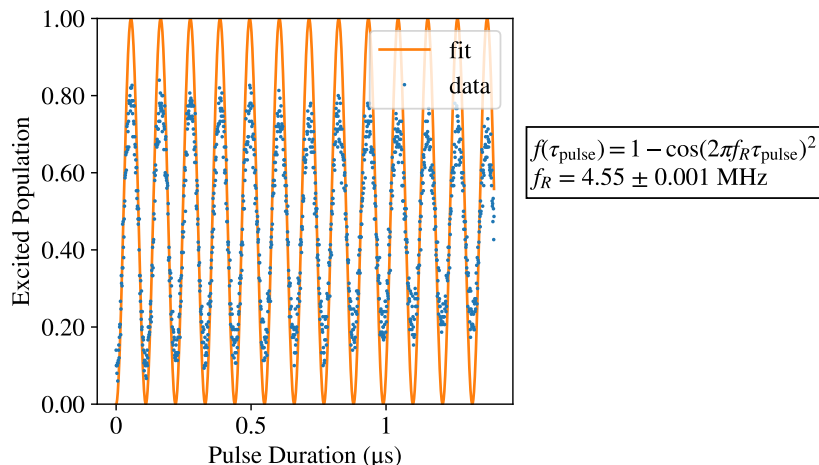


Figure 8.1: Averaged $P_{|1\rangle}$ of a qubit as a function of the duration of a pulse applied to it at the qubit frequency before measurement. These are Rabi oscillations where the frequency $f_R = \omega_{\text{Rabi}}/2\pi$ depends on the amplitude of the drive and the characteristics of the qubit. In order to apply an X pulse we must apply this drive for $\tau_{\text{pulse}} = 1/2f_R$ and so this experiment is used for rudimentary tune-up of an X gate.

In contrast, if we employ our Bayesian inference method we would first start by establishing a prior distribution for the model parameter, here the Rabi frequency f_R . Based on experiment parameters we expect this to be between 1MHz and 10MHz so we accordingly choose a uniform distribution between these values. The likelihood function we use to find $P(B|A)$ is Eq. (8.2) where in this case B is the single shot qubit measurements of 1 or 0. A is the Rabi frequency defined in Eq. (8.2) which will have a value given by the probability distribution of the prior $P(A)$. For each data point we find $P(A|B)$ and update our prior to reflect this new information. This posterior distribution of $P(A|B)$ becomes our new prior and we continue. After each update we have a distribution for A which reflects our best estimate for the value of the model parameter and also our certainty in this value. When we have reached the desired level of certainty, or when the level of confidence converges, we can stop taking data and report the value. This process is illustrated in Fig. 8.2.

The Bayesian inference method finds f_R to be $4.54 \pm 0.02\text{MHz}$ in comparison to the 4.52MHz found by the least squares fitting so the results are in good agreement. We can see from Fig 8.2 that, in comparison to the 210×10^3 single shots needed to find the least squares fit, we only needed 400 single shots to find f_R with Bayesian methods. We simultaneously obtain information about our certainty of this value and the point at which we do not improve our estimation of the model parameters by taking more data. The data used is shown in Figure 8.3 which highlights the relatively small amount of data needed to learn this parameter. In this case the single experiment is fast and effective to execute using either method but even at this proof of concept level the

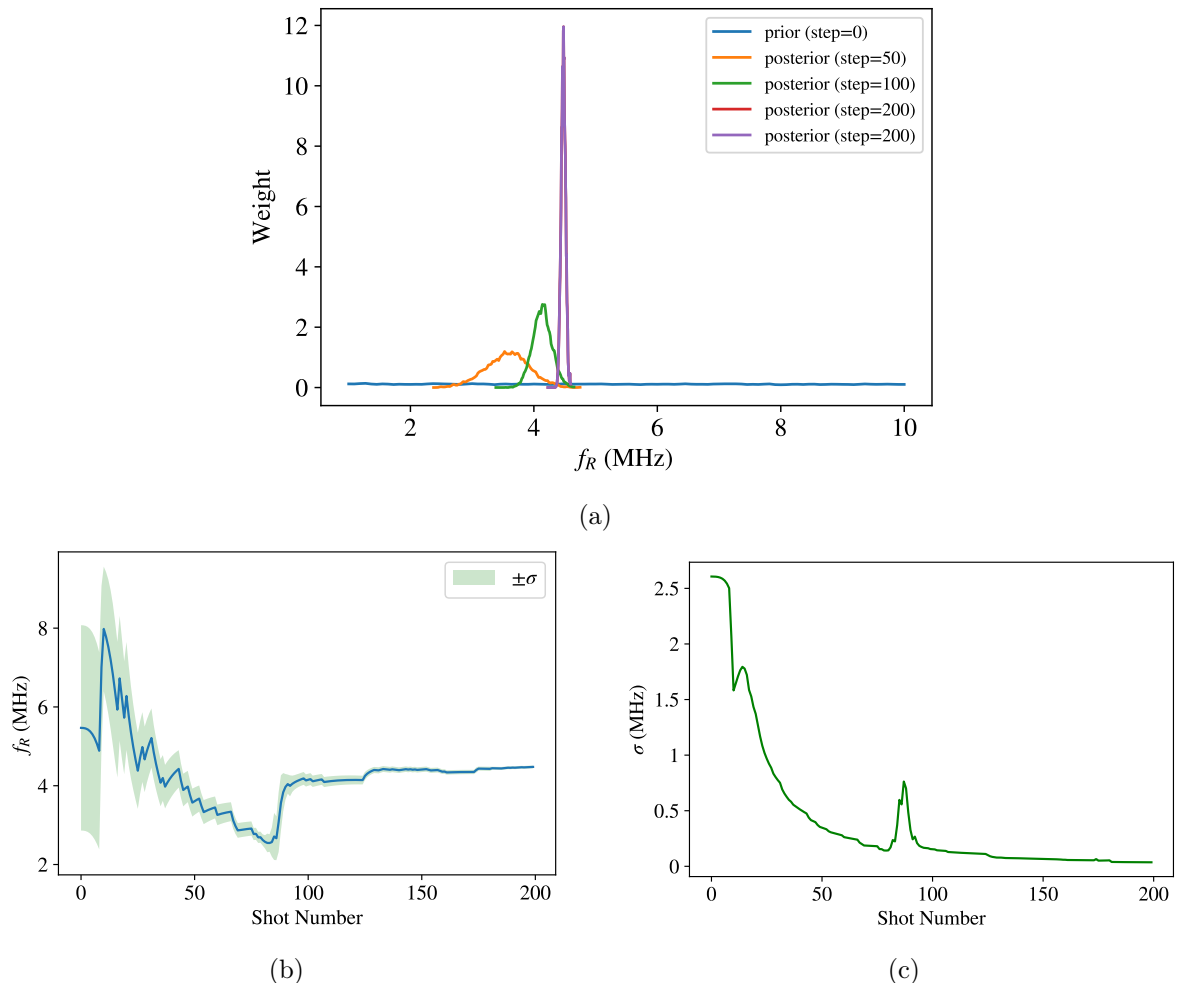


Figure 8.2: The process for learning the Rabi frequency f_R of Eq. (8.2) using Bayesian inference. (a) The prior distribution where we ascribe equal weight to all values of f_R between 1MHz and 10MHz. At each step a we add single shot 0 or 1 result of a qubit measurement and update the distribution to reflect this information and show the updated posterior probability distribution. Here we show these distributions after 50, 100, 200 and 400 steps. (b) The mean of the distribution of f_R as a function of the number of single shots analysed with one standard deviation σ shown. (c) The width of the peak of the distribution as a function of the number of shots analysed. This corresponds to the standard deviation shown in b) but can more clearly be seen to decrease and saturate at which point more data will not increase our confidence in the value of f_R .




Figure 8.3: The data used to implement the Bayesian inference of Fig. 8.2 where for each data point the qubit state is determined based on the readout calibration values calibrated previously. Then the qubit state and the pulse duration are provided to the Bayesian updater which updates the posterior distribution accordingly. This data should be compared to that of Fig. 8.1 where we averaged 150 times to find a single point and also went to longer pulse durations.

benefits of more rigorous analysis of data taking procedures can be seen.

8.3 n Parameter Inference

In the above example we only learned one parameter but we can easily extend the model to learn many parameters from this single experiment, for example to learn ω_R and T_1 at the same time we can use a phenomenological model

$$P_{|1\rangle}(t) = 1 - (1 + e^{-\tau_{\text{pulse}}\gamma}) \cos(2\omega_{\text{Rabi}}\tau_{\text{pulse}})/2, \quad (8.3)$$

where $T_1 = 1/\gamma$.

The only difference in procedure is that we now use Eq. (8.3) as the likelihood function instead of Eq. (8.2). The result of using this method to learn two parameters can be seen in Fig. 8.4 and the data used can be seen alongside the result of many averages in Fig. 8.5. Using 2400 single shots we find $f_R = 4.14 \pm 0.01\text{MHz}$ and $T_1 = 1.2 \pm 0.1\mu\text{s}$, which are in good agreement with traditional measurements of T_1 for this qubit which find that it varies between 1 and $2\mu\text{s}$.

This simple example shows that for the task of finding a small number of model parameters we already benefit by several orders of magnitude in terms of the amount of data needed. To the best of our knowledge this is the first time this method has been used for the characterisation of superconducting qubits.

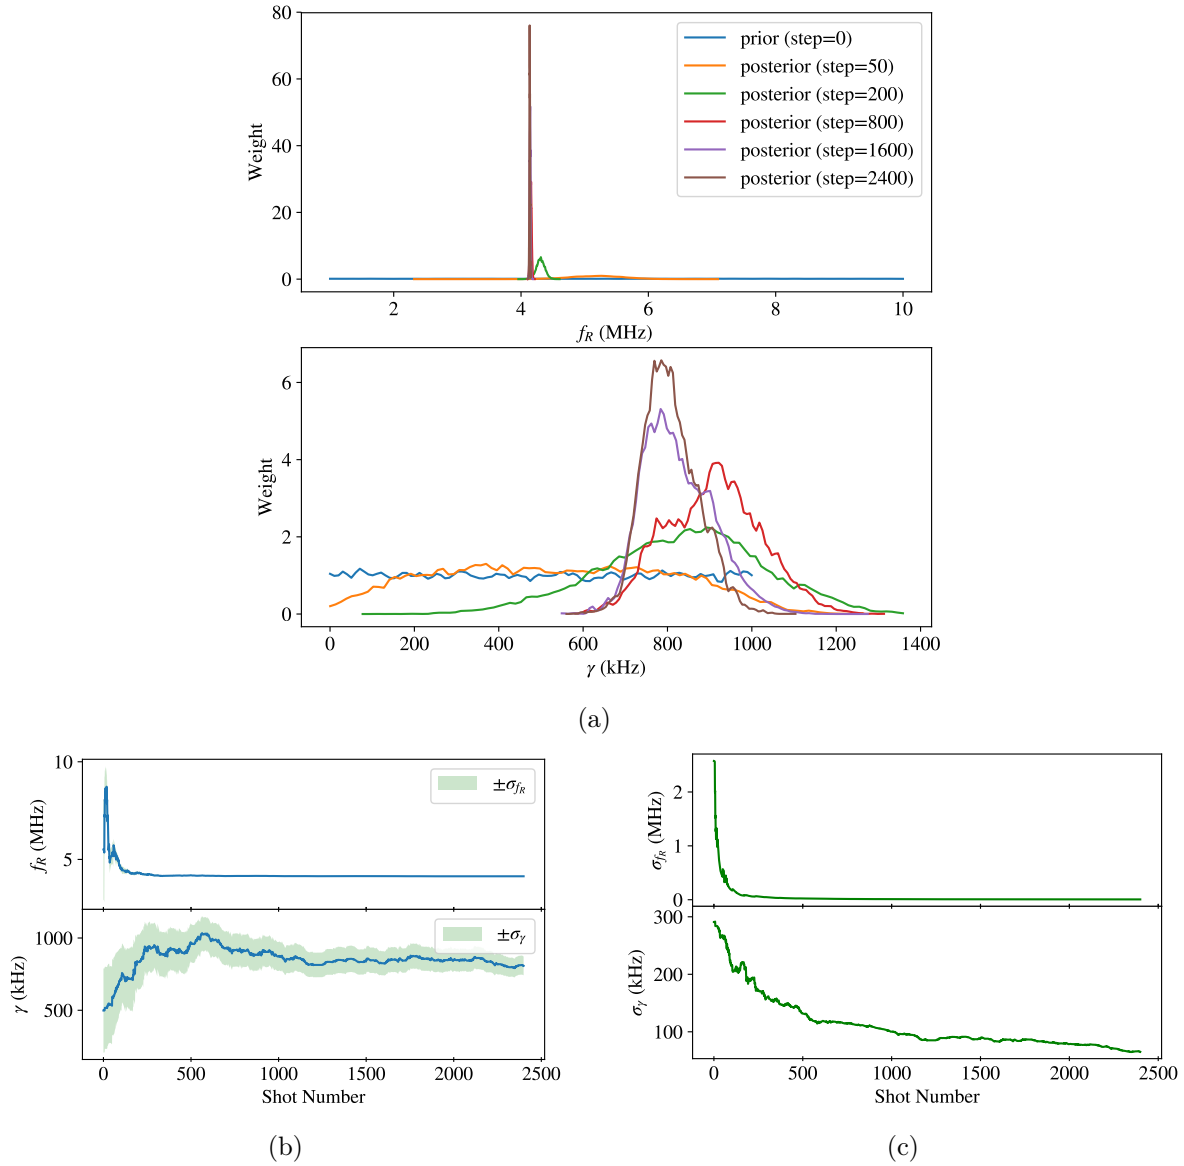


Figure 8.4: The process for learning the Rabi frequency f_R and T_1 of Eq. (8.3) using Bayesian inference. (a) The prior distribution where we ascribe equal weight to all values of f_R between 1MHz and 10MHz and all values of γ between 0 and 1MHz (i.e. $T_1 > 1\mu s$). The posterior distributions are shown for steps 50, 200, 800, 1600 and 2400. (b) The mean of the distributions of f_R and γ as a function of the number of single shots analysed with one standard deviation shown. (c) The width of the peak of the distributions as a function of the number of shots analysed where we can see that we quickly learn f_R but that we continue to gain certainty about γ right up to the 2400th shot.

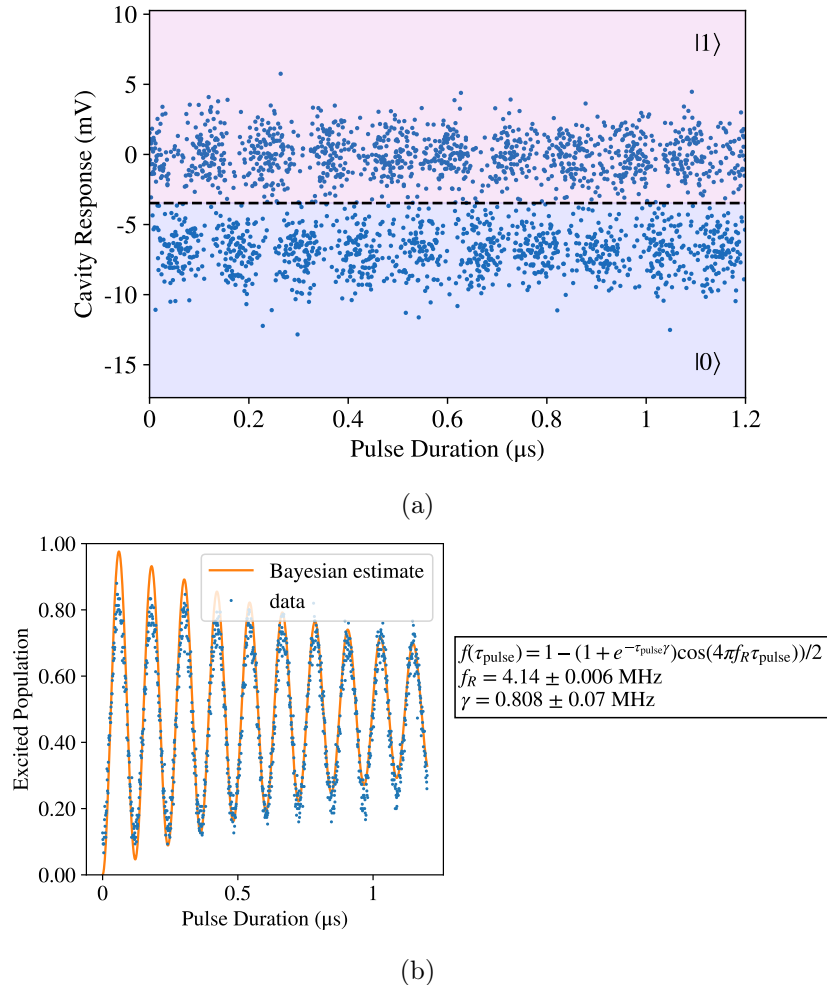


Figure 8.5: Data from a Rabi experiment where for each pulse duration, τ_{pulse} , the qubit is measured via the cavity and the state can be determined. There are 1200 values of τ_{pulse} , the cavity measurements for each of these values are shown in (a) and the qubit states determined from these comprise the the data used in Fig. 8.4 to learn f_R and T_1 of the phenomenological model described in Eq. (8.3). The decision process about the qubit state is indicated based on readout calibration values decided using the Kmeans algorithm described in Chapter 5. In (b) we show the result of this experiment after 150 averages together with the result of Eq. (8.3), given the parameters learned using the data in (a). This is meant to show that the values learned are consistent with the data without the need for many the many repetitions necessary for averaging.

8.4 Model Comparison

Another benefit of using the Bayesian approach is that it allows us to compare the likelihood of two models describing the same data. Bayes factor is a likelihood ratio between the marginal likelihood of the two models we would like to compare. The posterior probability of a model with the model parameters described by A , given the data B , is given by Bayes theorem of Eq. (8.1). During the computation we also calculated $P(B|A)$ which is the quantity describing how likely it is to observe the data, B , given the model and parameters of A . In the scenario that we must choose between the two models to decide which is more likely to result in the data we use Bayes factor which is given by

$$K = \frac{P(B|A_1)}{P(B|A_2)}, \quad (8.4)$$

where A_1 and A_2 describe the two competing models. If $K > 1$ then A_1 is more strongly supported by the data and vice versa. Happily, the Bayes factor naturally includes a penalty for using too many parameters in the model and so it is sensitive to overfitting, and thus reflects the simpler of the two models which best matches the data [222].

When using Qinfer to implement the Bayesian inference methods the values of $P(B|A)$ are naturally calculated in order to execute an update at each step [228, 229]. However it is numerically favourable to consider the property of the log total likelihood which is described by

$$\mathcal{L}(B|A) = \sum_i \log(P(b_i|A)), \quad (8.5)$$

where b_i are the individual data points that make up the data set. This is then related to the Bayes factor simply by

$$K = \exp(\mathcal{L}(B|A_1) - \mathcal{L}(B|A_2)). \quad (8.6)$$

We can apply this to the single and two parameter learning examples that we used above to demonstrate.

If we use Eq. (8.3) as model A_1 and Eq. (8.2) as model A_2 and analyse the data shown in Fig. 8.5, we obtain a Rabi frequency $f_R = 4.14\text{MHz}$ for both cases and a Bayes factor $K = 2.2 \times 10^{166} \gg 0$, which indicates that the decaying Rabi model of Eq. (8.3) is better supported by the data.

8.5 Conclusion and Outlook

We have demonstrated that Bayesian inference is a useful tool to efficiently characterise superconducting qubits. We successfully learned the values of parameters of our model for both a single and multi-parameter cases. This method allows us to update the confidence interval with every new measurement such that the measurement can be interrupted as soon as the confidence interval reaches a certain threshold. This enables us to report our confidence in the model parameter values as a function of the number of data points and to be much more conservative in the amount of data we take. In the examples shown it is not too onerous to perform the experiments with many averages, which we do here to provide a ‘sanity check’ that the Bayesian inference has correctly learned the parameter values. However, in scenarios where we are bottlenecked by the amount of time it takes to change between experiments or collect data, methods of this kind will become crucial for delivering calibration and tune-up of large quantum systems. As such, there remains much work to be done to integrate these methods throughout the process where they are found to be useful as well as considering how they interface with and compliment other methods for learning about the system and tuning it up. We hope that this minimal example provides motivation for such work.

Chapter 9

Conclusion and Outlook

In this thesis we have covered a very broad range of topics, spanning the field of the simulation of many-body quantum systems. Simulating these complex systems has long been of interest and in the last 30 years there have been significant advances in classical simulation [230] and, more recently, the promise of quantum simulators able to tackle tasks beyond the capabilities of classical computers. This exciting prospect may have dramatic implications for applications from high temperature superconductivity [44, 45] to weather forecasting [43]. One of the most promising of these applications is the simulation of correlated many-body quantum systems. We have highlighted particular areas of interest in the field to give an overall picture of the challenges remaining, while also investigating these areas in detail.

For two hallmark models we calculated the resources required for digital and analogue quantum simulation at the point where classical computers can no longer complete the simulation. We found that the realm of quantum advantage is already attainable by analogue quantum simulators and that the gate count estimates required to match this with digital quantum simulation is two orders of magnitude smaller than previously estimated. We then investigated the behaviour of the errors introduced in local and non-local observables by Trotterisation as a function of system size. We found that, despite diverging bounds on the fidelity of the state, these errors saturate with system size. This implies that larger system sizes than previously estimated can be simulated with finite Trotter step sizes. We used this result to justify our proposal that the Trotter error be formulated as a constant perturbation to the Hamiltonian rather than an increasing error on the evolution operator. This work can be expanded on by investigating the resources required for other systems and exploring the effects of relaxing some of the assumptions made in these chapters on the results. Further research or hardware developments may yet bring the estimates for resources required to implement useful simulations even closer to realisation, at which point it will be possible to have analogue and digital simulators

compete at the same task.

Having explored topics within analogue and digital quantum simulation, we then turned to the quantum hardware we might use to realise a digital quantum simulator. We focused on the superconducting qubit, a promising contender in the field. We introduced the physics required to understand the subsequent experimental chapters and to appreciate some of the benefits and drawbacks of this system, as well as the impressive developments in the field. In particular, the development of a superconductor-semiconductor qubit, the gatemon [57], which enables all-electric control of the qubit was discussed. Although each architecture will have its own unique set of challenges, they share the same requirements for use as a quantum computer capable of large-scale simulation. We have reported results from multiple aspects of meeting these requirements for our architecture of choice. First, we investigated a method for increasing the connectivity of the gatemon qubit by means of coupling qubits to a tuneable bus. This was successfully achieved although operation on a fast time scale was not possible. As such, connectivity of superconducting qubits is an area which deserves further research. It would be beneficial to understand better the failure mechanism for gatemon coupling by a tuneable bus and also to see other all-to-all proposals for superconducting qubits. This is highly beneficial in order to minimise the overhead of implementing operations on distant qubits and overcoming challenges such as frequency crowding.

Next, we reported results from research aimed to make this type of qubit more scalable by fabricating it using a 2 dimensional electron gas (2DEG), which enables reliable and reproducible fabrication *en masse*. The resulting qubit had competitive coherence times and we look forward to developments in selective area grown gatemon qubits [231] and their potential to use materials with less loss, while retaining the deterministic fabrication aspects.

Finally, we implemented Bayesian inference in order to characterise a superconducting qubit. We executed a simple Rabi experiment and successfully learned the parameters of the model describing it. This was a large improvement in efficiency over our previous methods for characterisation. However, more rigorous research is needed as our previous methods had not been optimised so it would first be important to find the optimal limits of traditional least squares fitting methods to do a quantitative comparison. Nevertheless, this research motivates the integration of Bayesian methods in the characterisation and calibration of superconducting qubit systems to increase the efficiency of these procedures and facilitate faster tune-up. This is likely to become more important as the size of the system being characterised increases or if we want to learn more complex models.

The simultaneous and coordinated development of algorithms and hardware is necessary for us to realise many-body quantum simulations with practical application. Bridging

the gap between experimental and theoretical aspirations is paramount and promises exciting and fruitful results. We have made significant steps in this direction but much work remains to be done.



Appendix A

Observable and error model choice for resource estimation

In Chapter 3 we calculate gate counts for digital quantum simulation based on the time-step τ which would result in errors on the off-diagonal correlation function comparable to those introduced by calibration errors in the analogue simulation. This appendix expands upon some of the decisions made and provides further simulation, all carried out by Stuart Flannigan and which we include for completeness.

For the Fermi-Hubbard model we can justify this choice by observing that the errors introduced due to calibration errors of the model parameters are greater than those introduced by the noise on the trapping lasers used in the experiment. This is achieved by performing simulations with time dependent parameter values for J and U which correspond to noise on the laser. The results are shown in Figure A.1

The errors due to imperfect model parameter calibration and laser fluctuations have the same qualitative time dependence and saturate quickly but the former are significantly larger and so it is these that we use to determine the maximum value for τ . We also compare different local errors which are shown to have similar behaviour, thus justifying the use of the off-diagonal correlation function as our observable. This is seen in Figures A.1 and A.2, where Figure A.2 shows the error on different observables in the analogue system as a function of time for different system sizes. Thus we can confidently use the off-diagonal correlation as a representative observable.

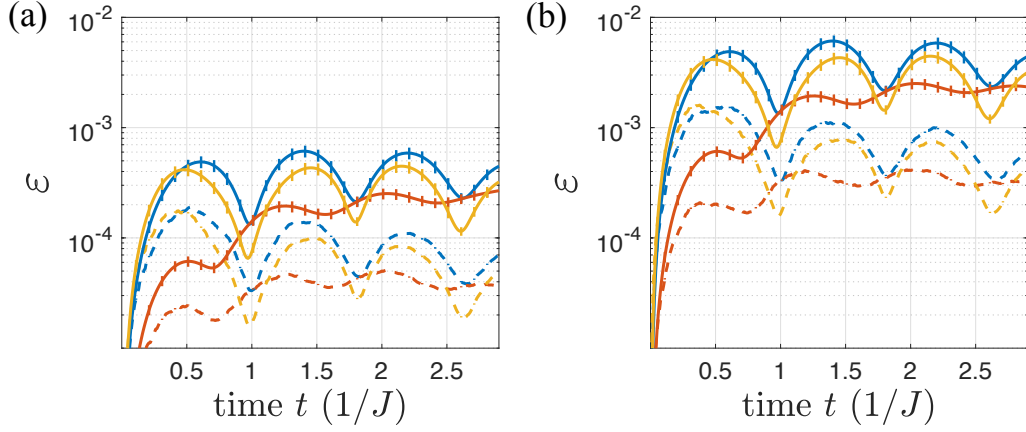


Figure A.1: Comparison of the calibration errors (dashed) in an analogue quantum simulation to the errors from laser fluctuations (solid) in the Hubbard model for different observables and magnitudes of error. The solid lines correspond to the errors in the observable in an analogue simulation comparing a calibration error of 1% (blue) and 0.1% (red). We compare the errors in the local occupation number (blue), the off-diagonal correlation functions (red) and the density-density correlations (yellow). For system size, $M = 20$.

The observables in question are defined:

$$\begin{aligned}
 O &= \frac{1}{\sqrt{M}} \sqrt{\sum_i |\langle \psi_{\text{sim}}(t) | c_{i,\uparrow}^\dagger c_{i,\uparrow} | \psi_{\text{sim}}(t) \rangle - \langle \psi_{\text{ex}}(t) | c_{i,\uparrow}^\dagger c_{i,\uparrow} | \psi_{\text{ex}}(t) \rangle|^2}, \\
 \text{CorrDens} &= \frac{2}{\sqrt{M}} \sqrt{\sum_i |\langle \psi_{\text{sim}}(t) | c_{\frac{N}{2},\uparrow}^\dagger c_{\frac{M}{2},\uparrow} c_{i,\uparrow}^\dagger c_{i,\uparrow} | \psi_{\text{sim}}(t) \rangle - \langle \psi_{\text{ex}}(t) | c_{\frac{M}{2},\uparrow}^\dagger c_{\frac{M}{2},\uparrow} c_{i,\uparrow}^\dagger c_{i,\uparrow} | \psi_{\text{ex}}(t) \rangle|^2}, \\
 \text{CorrOD} &= \frac{2}{\sqrt{M}} \sqrt{\sum_i |\langle \psi_{\text{sim}}(t) | c_{\frac{N}{2},\uparrow}^\dagger c_{i,\uparrow} | \psi_{\text{sim}}(t) \rangle - \langle \psi_{\text{ex}}(t) | c_{\frac{M}{2},\uparrow}^\dagger c_{i,\uparrow} | \psi_{\text{ex}}(t) \rangle|^2},
 \end{aligned} \tag{A.1}$$

where $|\psi_{\text{sim}}(t)\rangle$ is the simulated state at time t with errors due to miscalibration or Trotterisation and $|\psi_{\text{ex}}(t)\rangle$ is the exact solution. At $t = 0$ the state is initialised in an antiferromagnetic Néel state

Figure A.2 also contains information about the error introduced to the analogue simulation as a function of system size. The system size dependence comes from boundary effects, and if we only compare the density observables in the bulk of the system then these do not depend (as much) on system size. Figure A.3 contains this information for the digital simulation albeit in slightly different format. Here we see the critical time-step, τ , required for digital simulation to match analogue simulation when the latter has 1% calibration error. Errors in the density observables for a digital simulation depend

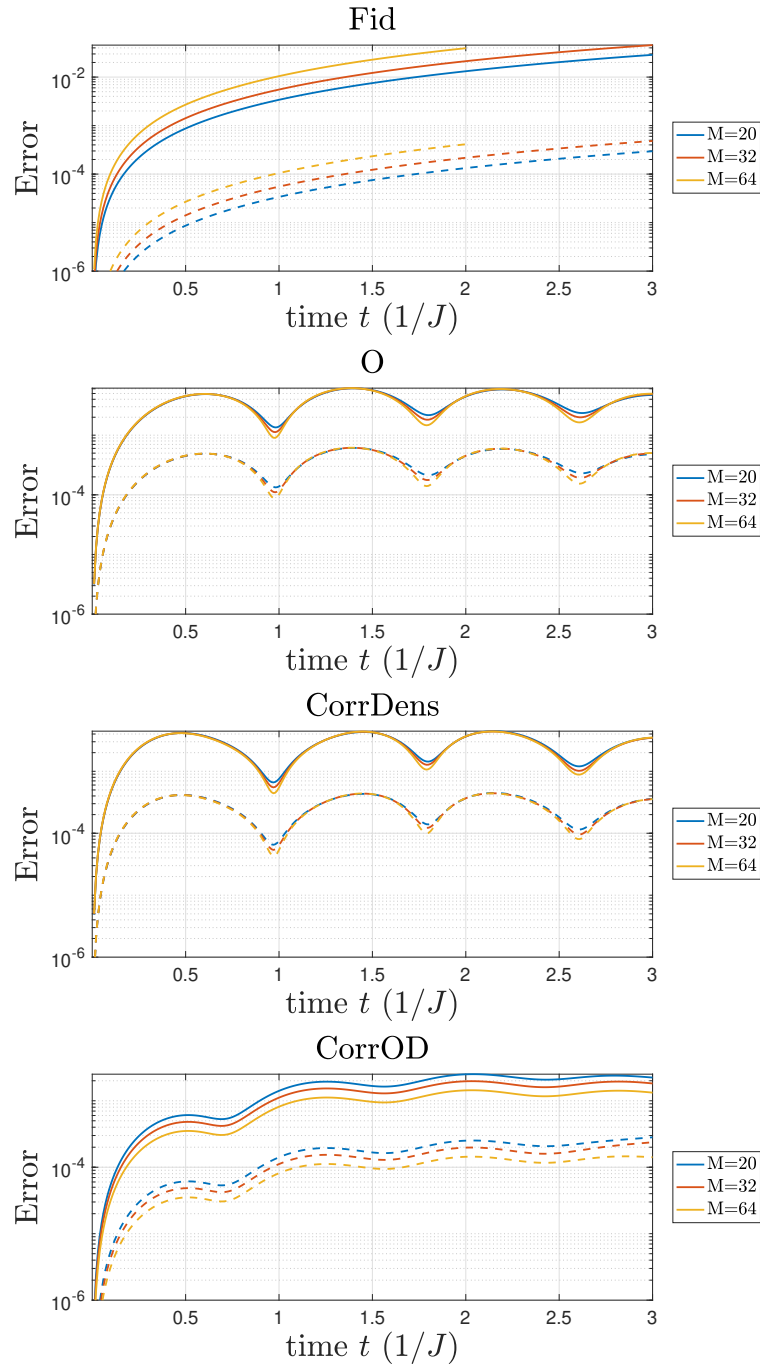


Figure A.2: Comparison of the calibration errors in an analogue quantum simulation for different system sizes, M . Solid lines are for a calibration error of 1% and dashed lines are for 0.1%.

on system size - leading to a critical time-step that grows with system size. The errors in the correlations for the digital simulation decrease with system size in the same way as the analogue leading to a critical time-step that is independent of system size.

The effects of increasing system size on the Trotter errors introduced are explored in more detail in Chapter 4 but these results are included for consistency to explain the choices made in Chapter 3 for τ in order to make a comparison to a realistic analogue simulation beyond the realm of classical simulation.

For the Ising model the off-diagonal correlation corresponds to the measurement

$$\text{CorrOD} = \frac{2}{\sqrt{M}} \sqrt{\sum_i |\langle \psi_{\text{sim}}(t) | S_+^{\frac{M}{2}} S_-^i | \psi_{\text{sim}}(t) \rangle - \langle \psi_{\text{ex}}(t) | S_+^{\frac{M}{2}} S_-^i | \psi_{\text{ex}}(t) \rangle|^2} \quad (\text{A.2})$$

where $S_{\pm}^i = \frac{1}{2}\sigma_{\pm}^i$ is the raising or lowering operator on site i .

These calculations were completed by Stuart Flannigan and Andrew Daley as part of [4].

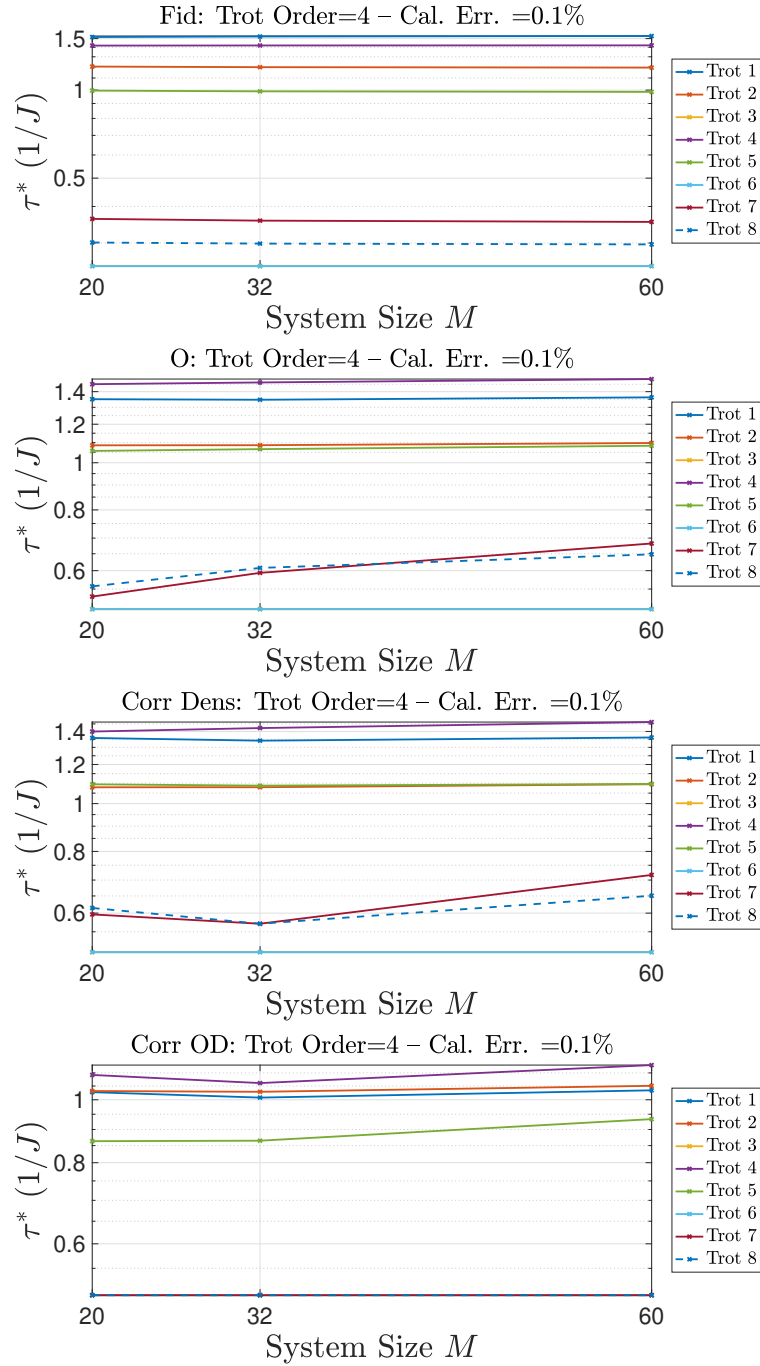


Figure A.3: The value of the critical time-step (at $tJ = 3$ for the observables and at $tJ = 2$ for the fidelities) for the digital simulation so that the errors are equal to a calibration error of 0.1% for the 4th order decompositions.



Appendix B

Q# code for gate count estimates

In Chapter 3 we outlined the necessary steps to execute the Fermi-Hubbard model and long-range Ising model on a digital quantum computer complete with gate count estimates. Here we include the Q# code which complements these calculations and shows how the execution is implemented in the Q# language. This allows us to simulate the computation for small system sizes to check that the implementation is correct. We can also directly find the resources required to implement the circuit in question which enables us to check the calculations of Chapter 3 and should enable the reader to make any improvements or changes and see the result both in terms of simulated observable values and resource estimates. This code and example notebooks for simulation are available at <https://github.com/nataliejpg/QuantumResourceEstimation> where we compare to exact evolution. In these examples only 1st order Trotter decomposition is implemented as the crucial numbers to calculate are the gate counts for a single time-step-sweep, from which we can easily find the total gate counts based on decomposition details. However for ease of use it would be worthwhile to implement higher order decompositions.

B.1 Long range Ising model

```
1
2 namespace Quantum.Isinglongrange {
3     open Microsoft.Quantum.Arrays;
4     open Microsoft.Quantum.Measurement;
5     open Microsoft.Quantum.Intrinsic;
6     open Microsoft.Quantum.Canon;
7     open Microsoft.Quantum.Math;
8     open Microsoft.Quantum.Convert;
```

```

9
10 // Implement evolution under Trotterised Ising model
11
12
13 /// # Summary
14 /// Applies coupling term for all qubits one site at a time
15 /// (ie sweep j for each i)
16 /// # Input
17 /// ## nSites
18 /// Number of sites in the Hamiltonian.
19 /// ## dt
20 /// Trotter time step size
21 /// ## J
22 /// 2d array of coupling coefficients Jij
23 /// ## qubits
24 /// Qubits that the encoded Ising Hamiltonian acts on.
25 operation EvolveCouplings(nSites: Int, dt: Double, J: Double[[[]],
26 qubits: Qubit[]): Unit {
27     for (i in 0 .. nSites - 2) {
28         for (j in i + 1 .. nSites - 1) {
29             ApplyWithCA(CNOT(qubits[i], _), Rz(-2.0 * J[i][j] * dt, _),
30 qubits[j]);
31         }
32     }
33 }
34
35 /// # Summary
36 /// Applies coupling term for all qubits minimising circuit depth
37 ///
38 /// # Input
39 /// ## nSites
40 /// Number of sites in the Hamiltonian.
41 /// ## dt
42 /// Trotter time step size
43 /// ## J
44 /// 2d array of coupling coefficients Jij
45 /// ## qubits
46 /// Qubits that the encoded Ising Hamiltonian acts on.
47 operation EvolveCouplingsNested(nSites: Int, dt: Double, J: Double[[[]],
48 qubits: Qubit[]): Unit {
49     let num = nSites/2 + nSites%2;
50     for (step in 0 .. num) {
51         for (ind in 0 .. nSites/2 - 1) {
52             let i = ind - step;
53             let j = nSites - (ind + step + 1);
54             ApplyWithCA(CNOT(qubits[i], _), Rz(-2.0 * J[i][j] * dt, _),
55 qubits[j]);
56         }
57     }
58     if ((nSites%2 == 0) or (step < nSites/2)){

```

```

58         for (ind in 0 .. nSites/2 - 1) {
59             let i = ind - step;
60             let j = nSites - (step + ind + 2);
61             if (ind == (num - 1)) {
62                 ApplyWithCA(CNOT(qubits[i], _),
63                     Rz(-2.0 * J[i][j + num] * dt, _),
64                     qubits[j + num]);
65             } else {
66                 ApplyWithCA(CNOT(qubits[i], _),
67                     Rz(-2.0 * J[i][j] * dt, _),
68                     qubits[j]);
69             }
70         }
71     }
72 }
73 }
74 }
75
76 /// # Summary
77 /// Applies evolution for a single timestep.
78 ///
79 /// # Input
80 /// ## nSites
81 /// Number of sites in the Hamiltonian.
82 /// ## dt
83 /// Trotter time step size
84 /// ## g
85 /// 1d array of transverse field coefficients, gj
86 /// ## h
87 /// 1d array of longitudinal field coefficients hj
88 /// ## J
89 /// 2d array of coupling coefficients Jij
90 /// ### nested
91 /// bool value of whether or not to reorder coupling terms to minimise
92 /// circuit depth
93 /// ## qubits
94 /// Qubits that the encoded Ising Hamiltonian acts on.
95 operation EvolveSingleTimestep(nSites: Int, dt: Double, g: Double[],
96 h: Double[], J: Double[][] , nested: Bool, qubits : Qubit[]): Unit {
97     for (idxSite in 0 .. nSites - 1) {
98         Rx((-2.0 * g[idxSite]) * dt, qubits[idxSite]);
99         Rz((-2.0 * h[idxSite]) * dt, qubits[idxSite]);
100     }
101     if (nested) {
102         EvolveCouplingsNested(nSites, dt, J, qubits);
103     } else {
104         EvolveCouplings(nSites, dt, J, qubits);
105     }
106 }

```

```

107
108     /// # Summary
109     /// Applies evolution for a single timestep starting with freshly
110     /// initialised qubits
111     /// (ie useful for gate count but not evolution)
112     ///
113     /// # Input
114     /// ## nSites
115     /// Number of sites in the Hamiltonian.
116     /// ## dt
117     /// Trotter time step size
118     /// ## g
119     /// 1d array of transverse field coefficients , gj
120     /// ## h
121     /// 1d array of longitudinal field coefficients hj
122     /// ## J
123     /// 2d array of coupling coefficients Jij
124     /// ### nested
125     /// bool value of whether or not to reorder coupling terms to minimise
126     /// circuit depth
127     operation EvolveSingleTimestepDummy(nSites: Int, dt: Double, g: Double[],
128     h: Double[], J: Double[][] , nested: Bool): Unit {
129         using (qubits = Qubit[nSites]) {
130             EvolveSingleTimestep(nSites, dt, g, h, J, nested, qubits);
131         }
132     }
133     /// # Summary
134     /// Applies full evolution in steps of dt.
135     ///
136     /// # Input
137     /// ## initialState
138     /// 1d array of initial states of each qubit in z basis (0 or 1)
139     /// ## time
140     /// Total time for evolution
141     /// ## dt
142     /// Trotter time step size
143     /// ## g
144     /// 1d array of transverse field coefficients , gj
145     /// ## h
146     /// 1d array of longitudinal field coefficients hj
147     /// ## J
148     /// 2d array of coupling coefficients Jij
149     /// ## nested
150     /// bool value of whether or not to reorder coupling terms to minimise
151     /// circuit depth
152     /// ## xinit
153     /// bool whether to apply an initial Hadamard so as to initialise in the
154     /// x basis
155     /// ## xmeas

```

```

156     /// bool whether to measure in the x basis (default is Z)
157     ///
158     /// # Output
159     /// ## finalState
160     /// 1d array of final states of each qubit in z basis (0 or 1)
161     operation Evolve(initialState: Int[], time: Double, dt: Double, g: Double[],
162 h: Double[], J: Double[][] , nested: Bool, xinit: Bool, xmeas: Bool): Result[] {
163         let nSites = Length(initialState);
164         using (qubits = Qubit[nSites]) {
165             for (idxSite in 0 .. nSites - 1) {
166                 if (initialState[idxSite] == 1) {
167                     X(qubits[idxSite]);
168                 }
169                 if (xinit) {
170                     H(qubits[idxSite]);
171                 }
172             }
173             let nSteps = Floor(time / dt);
174             for (idxIter in 0 .. nSteps - 1) {
175                 EvolveSingleTimestep(nSites, dt, g, h, J, nested, qubits);
176             }
177             if (xmeas) {
178                 for (q in qubits) {
179                     H(q);
180                 }
181             }
182             return ForEach(MResetZ, qubits);
183         }
184     }
185 }

```

B.2 Fermi-Hubbard model

```

1
2 namespace Quantum.Hubbard {
3     open Microsoft.Quantum.Arrays;
4     open Microsoft.Quantum.Measurement;
5     open Microsoft.Quantum.Intrinsic;
6     open Microsoft.Quantum.Canon;
7     open Microsoft.Quantum.Math;
8     open Microsoft.Quantum.Convert;
9
10    /// Implement evolution under Trotterised Hubbard model with nearest
11    /// neighbour hopping
12
13    /// # Summary
14    /// Applies Clifford X/2 rotation gate so 0 -> Y and Y -> 1
15    operation ClifY(qubit: Qubit): Unit is Adj + Ctl {

```

```

16     Rx(PI() / 2., qubit);
17 }
18
19 /// # Summary
20 /// Applies Clifford -X/2 rotation gate so 1 -> Y and Y -> 0
21 operation ClifYInv(qubit: Qubit): Unit is Adj + Ctl {
22     Rx(PI() / -2., qubit);
23 }
24
25 /// # Summary
26 /// Applies Controlled-Z gate between source q1 and target q2
27 operation CZ(q1: Qubit, q2: Qubit): Unit is Adj + Ctl {
28     H(q2);
29     CNOT(q1, q2);
30     H(q2);
31 }
32
33 /// # Summary
34 /// Applies chemical potential terms
35 ///
36 /// # Input
37 /// ## nSites
38 /// Number of sites in the Hubbard Hamiltonian.
39 /// ## dt
40 /// Trotter time step size
41 /// ## mu
42 /// Coefficient of the chemical potential, mu
43 /// ## U
44 /// Coefficient of the Coulomb repulsion, U
45 /// ## qubits
46 /// Qubits that the encoded Hubbard Hamiltonian acts on.
47 operation ApplyChemicalPotentialTerms(nSites: Int, dt: Double, mu: Double,
48 U: Double, qubits: Qubit[]): Unit {
49     for (idxSite in 0 .. nSites - 1) {
50         Rz((mu - U / 2.) * dt, qubits[idxSite]);
51         Rz((mu - U / 2.) * dt, qubits[idxSite + nSites]);
52     }
53 }
54
55 /// # Summary
56 /// Applies Coulomb repulsion terms
57 ///
58 /// # Input
59 /// ## nSites
60 /// Number of sites in the Hubbard Hamiltonian.
61 /// ## dt
62 /// Trotter time step size
63 /// ## U
64 /// Coefficient of the Coulomb repulsion, U

```

```

65  /// ## qubits
66  /// Qubits that the encoded Hubbard Hamiltonian acts on.
67  operation ApplyCoulumbRepulsionTerms(nSites: Int, dt: Double, U: Double,
68  qubits : Qubit[]): Unit {
69      for (idxSite in 0 .. nSites - 1) {
70          let q1 = qubits[idxSite];
71          let q2 = qubits[idxSite + nSites];
72          CNOT(q1, q2);
73          Rz((0.5 * U) * dt, q2);
74          CNOT(q1, q2);
75      }
76  }
77
78  /// # Summary
79  /// Applies controlled rotation part of a hopping term between two sites
80  /// of one species (ie w/o JW string), this is the full hopping term for
81  /// sites where the qubits representing the two sites are also adjacent.
82  ///
83  /// # Input
84  /// ## q1
85  /// Source qubit
86  /// ## q2
87  /// Target qubit
88  /// ## dt
89  /// Trotter time step size
90  /// ## J
91  /// Coefficient of the hopping, J
92  operation ApplyRotationSequence(q1: Qubit, q2: Qubit, dt: Double,
93  J: Double): Unit {
94      H(q1);
95      H(q2);
96      CNOT(q1, q2);
97      Rz((-1. * J) * dt, q2);
98      CNOT(q1, q2);
99      H(q1);
100     H(q2);
101     ClifY(q1);
102     ClifY(q2);
103     CNOT(q1, q2);
104     Rz((-1. * J) * dt, q2);
105     CNOT(q1, q2);
106     ClifYInv(q1);
107     ClifYInv(q2);
108 }
109
110 /// # Summary
111 /// Applies hopping term to a chain of qubits with nearest neighbour
112 /// hopping
113 ///

```

```

114  /// # Input
115  /// ## nSites
116  /// Number of sites in the Hubbard Hamiltonian.
117  /// ## dt
118  /// Trotter time step size
119  /// ## J
120  /// Coefficient of the hopping term, J
121  /// ## qubits
122  /// Qubits that the encoded Hubbard Hamiltonian acts on.
123  operation ApplyHoppingTermsChain(nSites: Int, dt: Double, J: Double,
124  qubits: Qubit[]): Unit {
125      for (idxSite in 0..2..nSites - 2) {
126          ApplyRotationSequence(qubits[idxSite],
127          qubits[idxSite + 1], dt, J);
128          ApplyRotationSequence(qubits[idxSite + nSites],
129          qubits[idxSite + nSites + 1],
130          dt, J);
131      }
132      for (idxSite in 1..2..nSites - 2) {
133          ApplyRotationSequence(qubits[idxSite],
134          qubits[idxSite + 1], dt, J);
135          ApplyRotationSequence(qubits[idxSite + nSites],
136          qubits[idxSite + nSites + 1],
137          dt, J);
138      }
139  }
140
141  /// # Summary
142  /// Applies the hopping term between two nearest neighbour sites on the
143  /// 'vertical ladder' which are represented by non adjacent qubits so a short
144  /// JW string must be computed
145  ///
146  /// # Input
147  /// ## q1Idx
148  /// Index of the source qubit
149  /// ## q2Idx
150  /// Index of the target qubit
151  /// ## dt
152  /// Trotter time step size
153  /// ## J
154  /// Coefficient of the hopping term, J
155  /// ## qubits
156  /// Qubits that the encoded Hubbard Hamiltonian acts on.
157  operation ApplyVerticalHoppingTerm(q1Idx: Int, q2Idx: Int, dt: Double,
158  J: Double, qubits: Qubit[]): Unit {
159      CNOT(qubits[q1Idx + 1], qubits[q1Idx + 2]);
160      CZ(qubits[q2Idx - 1], qubits[q2Idx]);
161      ApplyRotationSequence(qubits[q1Idx], qubits[q2Idx], dt, J);
162      CZ(qubits[q2Idx - 1], qubits[q2Idx]);

```



```

163     CNOT(qubits[q1Idx + 1], qubits[q1Idx + 2]);
164 }
165
166 /// # Summary
167 /// Applies the hopping term between two rows of qubits representing two
168 /// rows of sites of one species. The hopping is nearest neighbour and the
169 /// rotations are executed in parallel. The cumulative parity is stored on an
170 /// ancilla qubit.
171 ///
172 /// # Input
173 /// ## q1Idxs
174 /// Indices of the source qubits
175 /// ## q2Idxs
176 /// Indices of the target qubit
177 /// ## ancillaIdx
178 /// Index of the ancilla qubit
179 /// ## dt
180 /// Trotter time step size
181 /// ## J
182 /// Coefficient of the hopping term, J
183 /// ## qubits
184 /// Qubits that the encoded Hubbard Hamiltonian acts on plus ancilla(s)
185 operation ApplyRowHoppingTerms(q1Idxs: Int [], q2Idxs: Int [],
186 ancillaIdx: Int, dt: Double, J: Double, qubits: Qubit[]): Unit {
187     let rungs = Length(q1Idxs);
188     for (rIdx in 0..rungs - 1) {
189         CNOT(qubits[q1Idxs[rIdx] + 1], qubits[ancillaIdx]);
190         CNOT(qubits[q2Idxs[rIdx] - 1], qubits[ancillaIdx]);
191         CZ(qubits[ancillaIdx], qubits[q2Idxs[rIdx]]);
192     }
193     for (rIdx in 0..(rungs - 1)) {
194         let q1 = qubits[q1Idxs[rIdx]];
195         let q2 = qubits[q2Idxs[rIdx]];
196         ApplyRotationSequence(q1, q2, dt, J);
197     }
198     for (rIdx in (rungs - 1)..-1..0) {
199         CZ(qubits[ancillaIdx], qubits[q2Idxs[rIdx]]);
200         CNOT(qubits[q2Idxs[rIdx] - 1], qubits[ancillaIdx]);
201         CNOT(qubits[q1Idxs[rIdx] + 1], qubits[ancillaIdx]);
202     }
203 }
204
205 /// # Summary
206 /// Applies the hopping terms for a ladder architecture of the Hubbard
207 /// model assuming that the qubits are ordered in a snake on the
208 /// 'vertical' ladder.
209 ///
210 /// # Input
211 /// ## nSites

```

```

212  /// Number of sites in the Hubbard Hamiltonian.
213  /// ## dt
214  /// Trotter time step size
215  /// ## J
216  /// Coefficient of the hopping term, J
217  /// ## qubits
218  /// Qubits that the encoded Hubbard Hamiltonian acts on.
219  operation ApplyHoppingTermsLadderVertical(nSites: Int, dt: Double,
220  J: Double, qubits: Qubit[]): Unit {
221      ApplyHoppingTermsChain(nSites, dt, J, qubits);
222      for (idxSite in 0..4..nSites - 4) {
223          ApplyVerticalHoppingTerm(idxSite, idxSite + 3, dt, J, qubits);
224          ApplyVerticalHoppingTerm(idxSite + nSites, idxSite + nSites + 3,
225          dt, J, qubits);
226      }
227      for (idxSite in 2..4..nSites - 4) {
228          ApplyVerticalHoppingTerm(idxSite, idxSite + 3, dt, J, qubits);
229          ApplyVerticalHoppingTerm(idxSite + nSites, idxSite + nSites + 3,
230          dt, J, qubits);
231      }
232  }
233
234  /// # Summary
235  /// Applies the hopping terms for a ladder architecture of the Hubbard
236  /// model assuming that the qubits are ordered in a snake on the
237  /// 'horizontal' ladder.
238  ///
239  /// # Input
240  /// ## nSites
241  /// Number of sites in the Hubbard Hamiltonian.
242  /// ## dt
243  /// Trotter time step size
244  /// ## J
245  /// Coefficient of the hopping term, J
246  /// ## qubits
247  /// Qubits that the encoded Hubbard Hamiltonian acts on plus ancilla
248  operation ApplyHoppingTermsLadderHorizontal(nSites: Int, dt: Double,
249  J: Double, qubits: Qubit[]): Unit {
250      ApplyHoppingTermsChain(nSites, dt, J, qubits);
251      let length = nSites / 2;
252      let upperRowIndsup = Reversed(SequenceI(
253      0, length - 2));
254      let upperRowIndsdown = Reversed(SequenceI(
255      nSites, nSites + length - 2));
256      let lowerRowIndsup = SequenceI(
257      length + 1, nSites - 1);
258      let lowerRowIndsdown = SequenceI(
259      nSites + length + 1, 2 * nSites - 1);
260      ApplyRowHoppingTerms(upperRowIndsup, lowerRowIndsup,

```

```

261     2 * nSites, dt, J, qubits);
262     ApplyRowHoppingTerms(upperRowIndsdown, lowerRowIndsdown,
263     2 * nSites + 1, dt, J, qubits);
264 }
265
266 /// # Summary
267 /// Applies the hopping terms for a 2D lattice with qubits ordered in
268 /// snake pattern.
269 ///
270 /// # Input
271 /// ## nSites
272 /// Number of sites in the Hubbard Hamiltonian.
273 /// ## dt
274 /// Trotter time step size
275 /// ## J
276 /// Coefficient of the hopping term, J
277 /// ## qubits
278 /// Qubits that the encoded Hubbard Hamiltonian acts on plus ancillas
279 operation ApplyHoppingTermsLattice(nSites: Int, dt: Double, J: Double,
280 qubits: Qubit[]): Unit {
281     ApplyHoppingTermsChain(nSites, dt, J, qubits);
282     let length = Floor(Sqrt(IntAsDouble(nSites)));
283     for (row in 0..2..(length - 2)){
284         let upperRowIndsup = Reversed(SequenceI(
285         row * length,
286         row * length + length - 2));
287         let upperRowIndsdown = Reversed(SequenceI(
288         nSites + row * length,
289         nSites + row * length + length - 2));
290         let lowerRowIndsup = SequenceI(
291         row * length + length + 1,
292         row * length + 2 * length - 1);
293         let lowerRowIndsdown = SequenceI(
294         nSites + row * length + length + 1,
295         nSites + row * length + 2 * length - 1);
296         let ancillaIdxup = 2 * nSites + row;
297         let ancillaIdxdown = 2 * nSites + row + 1;
298         ApplyRowHoppingTerms(upperRowIndsup, lowerRowIndsup,
299         ancillaIdxup, dt, J, qubits);
300         ApplyRowHoppingTerms(upperRowIndsdown, lowerRowIndsdown,
301         ancillaIdxdown, dt, J, qubits);
302     }
303     for (row in 1..2..(length - 2)){
304         let upperRowIndsup = Reversed(SequenceI(
305         row * length,
306         row * length + length - 2));
307         let upperRowIndsdown = Reversed(SequenceI(
308         nSites + row * length,
309         nSites + row * length + length - 2));

```

```

310         let lowerRowIndsup = SequenceI(
311             row * length + length + 1,
312             row * length + 2 * length - 1);
313         let lowerRowIndsdown = SequenceI(
314             nSites + row * length + length + 1,
315             nSites + row * length + 2 * length - 1);
316         let ancillaIdxup = 2 * nSites + row - 1;
317         let ancillaIdxdown = 2 * nSites + row;
318         ApplyRowHoppingTerms(upperRowIndsup, lowerRowIndsup,
319             ancillaIdxup, dt, J, qubits);
320         ApplyRowHoppingTerms(upperRowIndsdown, lowerRowIndsdown,
321             ancillaIdxdown, dt, J, qubits);
322     }
323 }
324
325 /// # Summary
326 /// Applies evolution for a single timestep.
327 ///
328 /// # Input
329 /// ## nSites
330 /// Number of sites in the Hubbard Hamiltonian.
331 /// ## dt
332 /// Trotter time step size
333 /// ## mu
334 /// Coefficient of the chemical potential, mu
335 /// ## U
336 /// Coefficient of the Coulomb repulsion, U
337 /// ## J
338 /// Coefficient of the hopping term, J
339 /// ## structure
340 /// integer labeling structure, 1: chain, 2: vertical ladder,
341 /// 3: horizontal ladder, 4: 2D lattice
342 /// ## qubits
343 /// Qubits that the encoded Hubbard Hamiltonian acts on plus any
344 /// ancilla qubits
345 operation EvolveSingleTimestep(nSites: Int, dt: Double, mu: Double,
346 U: Double, J: Double, structure: Int, qubits: Qubit[]): Unit {
347     ApplyChemicalPotentialTerms(nSites, dt, mu, U, qubits);
348     ApplyCoulumbRepulsionTerms(nSites, dt, U, qubits);
349     if (structure == 1) {
350         ApplyHoppingTermsChain(nSites, dt, J, qubits);
351     } elif (structure == 2) {
352         ApplyHoppingTermsLadderVertical(nSites, dt, J, qubits);
353     } elif (structure == 3) {
354         ApplyHoppingTermsLadderHorizontal(nSites, dt, J, qubits);
355     } else {
356         ApplyHoppingTermsLattice(nSites, dt, J, qubits);
357     }
358 }

```

```

359
360 /// # Summary
361 /// Applies evolution for a single timestep starting with freshly
362 /// initialised qubits
363 /// (ie useful for gate count but not evolution)
364 ///
365 /// # Input
366 /// ## initialState
367 /// List of initial states of each site in the z basis
368 /// (0: down down, 1: down up, 2: up down, 3: up up)
369 /// ## dt
370 /// Trotter time step size
371 /// ## mu
372 /// Coefficient of the chemical potential, mu
373 /// ## U
374 /// Coefficient of the Coulomb repulsion, U
375 /// ## J
376 /// Coefficient of the hopping term, J
377 /// ## structure
378 /// integer labeling structure, 1: chain, 2: vertical ladder,
379 ///      3: horizontal ladder, 4: 2D lattice
380 operation EvolveSingleTimestepDummy(initialState: Int[], dt: Double,
381 mu: Double, U: Double, J: Double, structure: Int): Unit {
382     let nSites = Length(initialState);
383     mutable qubitNum = 2 * nSites;
384     if (structure == 3) {
385         set qubitNum += 2;
386     } elif (structure == 4) {
387         let length = Floor(Sqrt(IntAsDouble(nSites)));
388         set qubitNum += 2 * (length / 2);
389     }
390     using (qubits = Qubit[qubitNum]) {
391         EvolveSingleTimestep(nSites, dt, mu, U, J, structure, qubits);
392     }
393 }
394
395 /// # Summary
396 /// Applies full evolution in steps of dt on structure of choice,
397 /// allocating ancillas where necessary.
398 ///
399 /// # Input
400 /// ## initialState
401 /// List of initial states of each site in the z basis
402 /// (0: down down, 1: down up, 2: up down, 3: up up)
403 /// ## time
404 /// Total time for evolution
405 /// ## dt
406 /// Trotter time step size
407 /// ## mu

```

```

408  /// Coefficient of the chemical potential, mu
409  /// ## U
410  /// Coefficient of the Coulomb repulsion, U
411  /// ## J
412  /// Coefficient of the hopping term, J
413  /// ## structure
414  /// integer labeling structure, 1: chain, 2: vertical ladder,
415  /// 3: horizontal ladder, 4: 2D lattice
416  ///
417  /// # Output
418  /// ## finalState
419  /// list comparable to initialState of states of each site measured in
420  /// z basis
421  operation Evolve(initialState: Int[], time: Double, dt: Double,
422  mu: Double, U: Double, J: Double, structure: Int): Int[] {
423      let nSites = Length(initialState);
424      mutable qubitNum = 2 * nSites;
425      if (structure == 3) {
426          set qubitNum += 2;
427      } elif (structure == 4) {
428          let length = Floor(Sqrt(IntAsDouble(nSites)));
429          set qubitNum += 2 * (length / 2);
430      }
431      using (qubits = Qubit[qubitNum]) {
432          for (idxSite in 0 .. nSites - 1) {
433              if (initialState[idxSite] % 2 == 1) {
434                  X(qubits[idxSite + nSites]);
435              }
436              if (initialState[idxSite] > 1) {
437                  X(qubits[idxSite]);
438              }
439          }
440          let nSteps = Floor(time / dt);
441          for (idxIter in 0 .. nSteps - 1) {
442              EvolveSingleTimestep(nSites, dt, mu, U, J, structure, qubits);
443          }
444          let result = ForEach(MResetZ, qubits);
445          mutable finalState = new Int[nSites];
446          for (idxSite in 0 .. nSites - 1) {
447              if (result[idxSite] == One) {
448                  set finalState w/= idxSite <- 2;
449              }
450              if (result[idxSite + nSites] == One) {
451                  let newState = finalState[idxSite] + 1;
452                  set finalState w/= idxSite <- newState;
453              }
454          }
455          return finalState;
456      }

```

```
457     }  
458  
459 }
```


Appendix C

Fabrication details and setup for the voltage-controlled superconducting quantum bus

C.1 Device Fabrication

The sample is fabricated from a $\sim 100\text{nm}$ thick Al film on a high resistivity Si substrate. First the feed and control lines, gatemon islands, and readout resonators as well as windows for placing the nanowires ($20\mu\text{m}$ by $40\mu\text{m}$) are wet etched. Both the qubits and the tunable bus JJs were formed by selectively wet etching a segment of a $\sim 30\text{nm}$ thick Al shell that was epitaxially grown around a $\sim 75\text{ nm}$ diameter single-crystal InAs nanowire [232]. Subsequently, molecular-beam epitaxy-grown nanowires are transferred from the growth chip to the etched windows using a dry deposition technique. A $\sim 200\text{nm}$ segment of the Al shell is removed using a wet etch. The nanowire contacts are patterned from Al using a lift-off process with an ion mill step to remove the native Al_2O_3 prior to deposition. The tunable bus ZrO_2 gate dielectric is deposited using an atomic layer deposition lift-off process. Finally the qubit gates and the bus top gate are patterned from Al again using a lift-off process.

C.2 Setup wiring

Figure C.1 shows an electrical circuit diagram of the sample. All measurements presented in the paper are performed in a cryogen-free dilution refrigerator with a base temperature below 20mK . The details of the fridge shielding as well as the line filtering

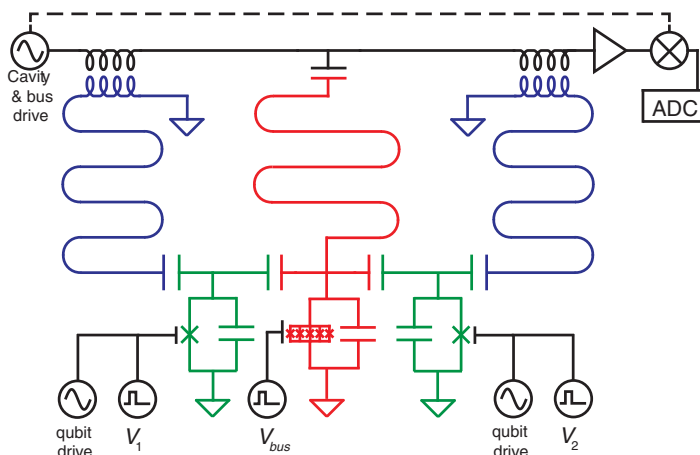


Figure C.1: Circuit diagram of the device, including readout resonators (blue), qubits (green), and the tunable bus (red). ©2019 American Physical Society, published in [6].

are shown in Fig. C.2. The sample is mounted inside an Al box to suppress magnetic fluctuations. This box is placed inside a Cu box used to mount the sample at the MC plate of the refrigerator. Both boxes are closed but not light tight; they are further surrounded by a cylindrical cryoperm shield, which is also thermally anchored to the mixing chamber. To manipulate an individual qubit, a coaxial line and a dc line are used (green in Fig. C.2). In contrast to earlier experiments, we use one single coax for both xy microwave control and fast gate voltage z control as illustrated in Fig. C.1. The coax line is filtered at high frequencies ($>300\text{MHz}$) by a Minicircuits VLF-320 low-pass filter and an ECCOSORB filter. A key feature of the Minicircuits filter is the increased transmission at typical qubit resonance frequencies. The dc line is added with a bias tee at low temperature. The tunable bus is controlled with a coaxial and dc line as well (red). Since there is no need for microwave control, the low-pass filter used (Minicircuits VLFX-300) filters high frequencies much more efficiently. For readout (blue), a signal line is used, which is heavily attenuated (60dB) to reduce both the thermal occupation of the resonator and noise to the sample. The readout line with magnetically shielded isolators allows signal out while suppressing any noise from the traveling-wave parametric amplifier (TWPA) and the cryogenic HEMT amplifier at the 4-K stage. The TWPA is driven with a microwave pump tone at $\sim 8\text{GHz}$ [192] and shielded by a separate cryoperm shield. We note that the scheme in Fig. C.2 only displays the setup for one qubit. For two qubit operation the green components have to be doubled. The data in all figures in the main text were acquired using parallel heterodyne detection in the dispersive regime. On the signal line we combine two drives with frequencies close to the resonance frequencies of cavity 1 and cavity 2 (blue). After passing through the TWPA and HEMT amplifiers and another amplification step at room temperature the

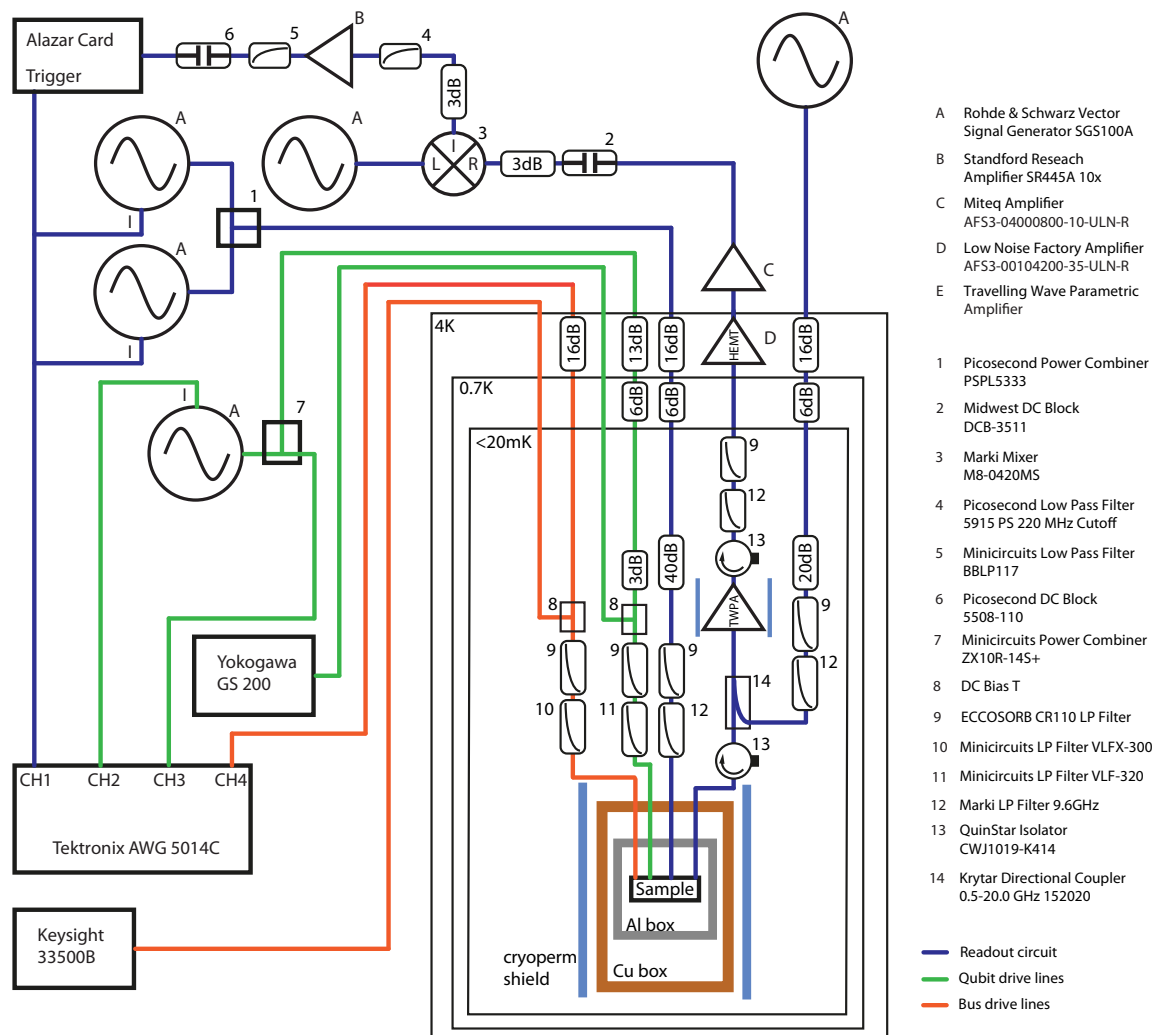


Figure C.2: Schematic of the experimental setup for readout of two qubits but manipulation of only one qubit. ©2019 American Physical Society, published in [6].

combined signal is mixed down to two intermediate frequencies with a local oscillator, before sampling and performing digital homodyne detection to extract the cavity magnitude response. Due to the low fidelity readout, qubit state measurements are obtained by averaging over many experimental runs. Qubit state assignments are calibrated using Rabi oscillations between the $|0\rangle$ and $|1\rangle$ states.

Appendix D

Fabrication details and setup for the 2DEG gatemon

D.1 The sample

Separate transport characterization shows that the 2DEGs exhibit a Hall mobility of approximately $2000\text{cm}^2\text{V}^{-1}\text{s}^{-1}$ and an induced gap of $200\mu\text{eV}$. Semi-insulating (Fe-doped) InP is used as a host substrate for the 2DEG buffer layer, which is etched away before patterning the qubit island and microwave control circuitry. The qubits were fabricated by first wet etching a mesa for the qubit JJ. The width w of the JJ was defined by the mesa etch. The JJ was then formed by selectively wet etching an $l \sim 100\text{nm}$ long segment of the $\sim 50\text{nm}$ thick Al. A 20nm thick AlOx layer (yellow in Fig. 7.2b,c) was deposited as a gate dielectric, followed by the evaporation of an Al top gate (red in Fig. 7.2b). The heterostructure and buffer were removed almost everywhere on the chip, to leave a mesa region a few micrometres large to form the active region of the qubit. The qubit islands, gate lines and readout cavities were defined in a lift-off process with a 100nm Al layer. Finally, the epitaxial Al layer on top of the mesa and the microwave circuit were connected in a contact step. For each qubit, E_C/h is determined by the capacitance of the T-shaped Al island to the surrounding ground plane and designed to be $\sim 230\text{MHz}$. All qubits were coupled to individual $\lambda/4$ superconducting cavities with resonant frequencies separated by 50MHz and centred around 7.25GHz . All six cavities were coupled to a common feed line [139]

D.2 Qubit manipulation and readout.

All measurements presented in the paper were performed in a cryogen-free dilution refrigerator with a base temperature below 50mK. The sample was mounted inside an Al box to suppress magnetic fluctuations. This box was placed inside a Cu box used to mount the sample at the mixing chamber plate of the refrigerator. Both boxes were closed but not light tight and further surrounded by a cylindrical cryoperm shield, which was also thermally anchored to the mixing chamber. The qubit was initialized in the $|0\rangle$ state by waiting for much longer than the relaxation time T_1 . To manipulate a single qubit, one coaxial line and a DC line were used: the coaxial line was filtered by a Minicircuits VLF-320 low-pass filter and an ECCOSORB filter to reduce the noise while allowing for gate pulses. At high frequencies ($>2\text{GHz}$), the filter was attenuated by roughly 20dB, which allowed direct driving of the qubit. The d.c. line was filtered with a resistor-capacitor filter and added with a bias tee at low temperature. For the X microwave control as well as readout, the pulses were shaped through IQ modulation of the microwave source using an arbitrary waveform generator channel for I (in-phase component of the waveform) and Q (quadrature component). For readout, the signal line was heavily attenuated (60dB) to reduce both the thermal occupation of the resonator and noise to the sample. After passing through a magnetically shielded isolator, a travelling wave parametric amplifier [192], another magnetically shielded isolator, a cryogenic Low Noise Factory high-electron-mobility transistor amplifier and another amplification stage at room temperature, the qubit readout signals were mixed down to intermediate frequencies with a local oscillator, before sampling and performing digital homodyne detection to extract the cavity magnitude response. Qubit state measurements were obtained by averaging over ~ 1000 experimental runs. We used the raw Rabi oscillation data for the qubit ensemble state assignments [57]. The data in Fig. 7.3a were acquired with a vector network analyser. The data in all the other figures were acquired using heterodyne detection in the dispersive regime. For Fig. 7.5 we combined two drives with frequencies close to the resonance frequencies of the cavities of Q_2 and Q_3 on the signal line

List of Tables

2.1	Physical gate sequences corresponding to the 24 single-qubit Clifford rotations	10
3.1	Trotter time steps required to implement simulation of the Fermi-Hubbard model with errors comparable to those introduced by calibration error.	21
3.2	Resources for implementation of chemical potential term of the Fermi-Hubbard model	29
3.3	Resources for implementation of Coulomb repulsion term of the Fermi-Hubbard model	30
3.4	Resources for implementation of hopping term of the Fermi-Hubbard model on a chain of sites	31
3.5	Total resources required to implement single time-step-sweep of the Fermi-Hubbard model on a chain	31
3.6	Total resources required to implement full evolution of the Fermi-Hubbard model on a chain	32
3.7	Resources for implementation of hopping term of the Fermi-Hubbard model on a vertical ladder of sites	34
3.8	Resources for implementation of hopping term of the Fermi-Hubbard model on a horizontal ladder of sites	35
3.9	Total resources required to implement single time-step-sweep of the Fermi-Hubbard model on a ladder	36
3.10	Total resources required to implement full evolution of the Fermi-Hubbard model on a ladder	36
3.11	Resources for implementation of hopping term of the Fermi-Hubbard model on a lattice of sites	38
3.12	Total resources required to implement single time-step-sweep of the Fermi-Hubbard model on a lattice	38
3.13	Total resources required to implement full evolution of the Fermi-Hubbard model on a ladder	39

3.14	Calculated T gate counts required to implement full evolution of the Fermi-Hubbard model on a lattice competitive with analogue simulation.	40
3.15	Summary table for resource estimation for simulation of the Fermi-Hubbard model on a 2D lattice	41
3.16	Summary of T gate resource estimation for simulation of the Fermi-Hubbard model on a 2D lattice	42
3.17	Resources for implementation of longitudinal term of the long-range Ising model	44
3.18	Resources for implementation of transverse term of the long-range Ising model	45
3.19	Resources for implementation of transverse term of the long-range Ising model with limited rotational gate axes	45
3.20	Resources for implementation of coupling term of the long-range Ising model	45
3.21	Total resources required to implement single time-step-sweep of the long-range Ising model	46
3.22	Total resources required to implement full evolution of the long-range Ising model	46
3.23	Calculated T gate counts required to implement full evolution of the long-range Ising model competitive with analogue simulation.	47
3.24	Summary table for resource estimation for simulation of the long-range Ising model	48
3.25	Summary of T gate resource estimation for simulation of the long-range Ising model.	48

List of Figures

2.1	The Bloch sphere	8
2.2	Sketch of Trotter-Suzuki implementation	13
3.1	Comparison of analogue and digital quantum simulation errors introduced to local observable in the Fermi-Hubbard model	20
3.2	Circuit diagram implementing the hopping term of the Fermi-Hubbard model	25
3.4	Circuit diagram for simplified implementation of hopping term of the Fermi-Hubbard model	26
3.5	Circuit diagram implementing the hopping term of the Fermi-Hubbard model with ancilla qubit for JW parity storage.	27
3.7	Circuit diagram implementing the hopping term of the Fermi-Hubbard model between two pairs of sites with ancilla qubit for JW parity storage.	28
3.9	Circuit diagram for simplified implementation of the hopping term of the Fermi-Hubbard model between two pairs of sites with ancilla qubit for JW parity storage	29
3.10	Sketch describing qubit ordering and bond grouping for chain of sites in the Fermi-Hubbard model	30
3.11	Sketch describing qubit ordering and bond grouping for ladder of sites ordered vertically in the Fermi-Hubbard model	33
3.12	Sketch describing qubit ordering and bond grouping for ladder of sites ordered horizontally in the Fermi-Hubbard model	34
3.13	Sketch describing qubit ordering and bond grouping for lattice of sites ordered horizontally in the Fermi-Hubbard model	37
3.14	Comparison of analogue and digital quantum simulation errors introduced to local observable in the long-range Ising model.	43
4.1	Coupling strengths of the 2D classical Ising model required to simulate the 1D quantum Ising model	59

4.2	Binder cumulant values of the 1D TFIM as a function of the transverse field strength g at different values of τ for nearest and next-nearest-neighbour coupling.	61
4.3	Critical point of the 1D TFIM for nearest and next nearest neighbour coupling	62
4.4	nearest-neighbour correlation as a function of the transverse magnetic field of the 1D TFIM	63
4.5	Error in the nearest-neighbour correlation as a function of system size	64
4.6	A linear fit of error as a function of τ^2 for nearest-neighbour correlation	65
4.7	τ^2 error coefficient for nearest-neighbour correlation as a function of system size	65
4.8	Static susceptibility as a function of the transverse magnetic field of the 1D TFIM	66
4.9	Relative error in the static susceptibility as a function of system size	67
4.10	Scaled τ^2 error coefficient for static susceptibility as a function of system size	68
4.11	τ^2 error coefficient for static susceptibility as a function of system size scaled by the number of terms	68
5.1	Linear and nonlinear LC circuit	73
5.2	Sketch of harmonic and anharmonic energy potentials	75
5.3	The Cooper Pair Box	75
5.4	Two tone spectroscopy measurement showing the 01 and 02/2 transitions of a transmon qubit	76
5.5	Device image of transmon qubit islands, zoom of SQUID loop and Josephson junction	77
5.6	Two tone spectroscopy showing transmon qubit frequency as a function of flux	78
5.7	Circuit diagram of transmon qubit coupled to a readout resonator	80
5.8	Full device image of 4 qubit transmon chip and measurement of a readout resonator	82
5.9	Readout calibration data and analysis for labelled data	84
5.10	Readout calibration data and analysis for unlabelled data	85
5.11	Cavity magnitude response data for Rabi experiment	88
5.12	Excited state probability for Rabi experiment	89
5.13	T_1 data and fit	90
5.14	Ramsey measurement and extracted T_2^* data and fit	91
5.15	$T_{2,\text{echo}}$ data and fit	92
5.16	Randomised benchmarking experiment and tune up procedure	94
5.17	Swaps experiment data for two capacitively coupled transmon qubits	96

6.1	Schematic and simulation of the voltage-controlled superconducting quantum bus	101
6.2	Two-qubit device with switchable quantum bus	103
6.3	Spectroscopy of switchable bus and qubit.	104
6.4	Tunable coherent gatemon coupling facilitated by the tunable bus	106
6.5	Fast switch response of tunable bus	107
7.1	Sketch of transmon and gatemon qubits overview	112
7.2	Device image, schematic of material stack and SEM of Josephson junction of 2DEG gatemon	117
7.3	Characteristic experiments of the 2DEG gatemon	118
7.4	Coherence times of the 2DEG gatemon	119
7.5	Two-qubit interaction of capacitively coupled 2DEG gatemons	120
8.1	Least squares analysis of Rabi experiment	125
8.2	Bayesian inference analysis of Rabi experiment	126
8.3	Data used for Bayesian inference analysis of Rabi experiment	127
8.4	n parameter learning by Bayesian inference on Rabi experiment	128
8.5	Data used in of n parameter learning by Bayesian inference on Rabi experiment	129
A.1	Calibration errors for different observables of the Fermi-Hubbard model	138
A.2	System size dependence of the calibration errors of the Fermi-Hubbard model	139
A.3	System size dependence of the critical time-step of the Fermi-Hubbard model	141
C.1	Circuit diagram of tunable coupler device	160
C.2	Schematic of the experimental setup	161

List of publications

- [1] **N. J. Pearson**, L. M. Austin, C. Grenade, A. Hertel, L. O. Anderson, N. Weibe, C. M. Marcus, and K. D. Petersson. Hamiltonian learning on a superconducting qubit using bayesian inference. *in preparation*.
- [2] **N. J. Pearson**, D. Poulin, and M. Troyer. Trotter errors in simulations of quantum lattice models. *in preparation*.
- [3] **N. J. Pearson**, D. Poulin, and M. Troyer. Criticality of the transverse field ising-trotter-suzuki model. *in preparation*.
- [4] S. Flannigan, **N. J. Pearson**, M. Troyer, and A. Daley. The practical quantum advantage point for analogue and digital quantum simulators. *in preparation*.
- [5] T. W. Larsen, M. E. Gershenson, L. Casparis, A. Kringhøj, **N. J. Pearson**, R. P. G. McNeil, F. Kuemmeth, P. Krogstrup, K. D. Petersson, and C. M. Marcus. Parity-protected superconductor-semiconductor qubit. *Phys. Rev. Lett.*, 125:056801, Jul 2020.
- [6] L. Casparis, **N. J. Pearson**, A. Kringhøj, T. W. Larsen, F. Kuemmeth, J. Nygård, P. Krogstrup, K. D. Petersson, and C. M. Marcus. Voltage-controlled superconducting quantum bus. *Phys. Rev. B*, 99:085434, Feb 2019.
- [7] L. Casparis, M. R. Connolly, M. Kjaergaard, **N. J. Pearson**, A. Kringhøj, T. W. Larsen, F. Kuemmeth, T. Wang, C. Thomas, S. Gronin, G. C. Gardner, M. J. Manfra, C. M. Marcus, and K. D. Petersson. Superconducting gatemon qubit based on a proximitized two-dimensional electron gas. *Nature Nanotechnology*, 13(10):915–919, 2018.

Bibliography

- [8] Patrik Fazekas. *Lecture notes on electron correlation and magnetism*, volume 5. World scientific, 1999.
- [9] Michael Tinkham. *Introduction to superconductivity*. Courier Corporation, 2004.
- [10] Shou-cheng Zhang. Topological states of quantum matter. *Physics*, 1:6, 2008.
- [11] M. Aidelsburger, M. Atala, M. Lohse, J. T. Barreiro, B. Paredes, and I. Bloch. Realization of the hofstadter hamiltonian with ultracold atoms in optical lattices. *Phys. Rev. Lett.*, 111:185301, Oct 2013.
- [12] Philip W Anderson. More is different. *Science*, 177(4047):393–396, 1972.
- [13] Subir Sachdev and Bernhard Keimer. Quantum criticality. *arXiv preprint arXiv:1102.4628*, 2011.
- [14] Subir Sachdev. *Quantum Phase Transitions*. 2 edition, 2011.
- [15] Mario Norberto Baibich, Jean Marc Broto, Albert Fert, F Nguyen Van Dau, Frédéric Petroff, P Etienne, G Creuzet, A Friederich, and J Chazelas. Giant magnetoresistance of (001) fe/(001) cr magnetic superlattices. *Physical review letters*, 61(21):2472, 1988.
- [16] Grünberg Binasch, Peter Grünberg, F Saurenbach, and W Zinn. Enhanced magnetoresistance in layered magnetic structures with antiferromagnetic interlayer exchange. *Physical review B*, 39(7):4828, 1989.
- [17] Kostya S Novoselov, Andre K Geim, Sergei V Morozov, D Jiang, Y_ Zhang, Sergey V Dubonos, Irina V Grigorieva, and Alexandr A Firsov. Electric field effect in atomically thin carbon films. *science*, 306(5696):666–669, 2004.
- [18] Herbert Kroemer. Band offsets and chemical bonding: the basis for heterostructure applications. *Physica Scripta*, 1996(T68):10, 1996.

- [19] Zh I Alferov. The history and future of semiconductor heterostructures. *Semiconductors*, 32(1):1–14, 1998.
- [20] Willard S Boyle and George E Smith. Charge coupled semiconductor devices. *Bell System Technical Journal*, 49(4):587–593, 1970.
- [21] Vinod Kumar Khanna. *Fundamentals of solid-state lighting: LEDs, OLEDs, and their applications in illumination and displays*. CRC press, 2014.
- [22] Steven R. White. Density matrix formulation for quantum renormalization groups. *Phys. Rev. Lett.*, 69(19):2863–2866, Nov 1992.
- [23] Ulrich Schollwöck. The density-matrix renormalization group in the age of matrix product states. *Annals of Physics*, 326(1):96–192, 2011.
- [24] Guifré Vidal. Efficient simulation of one-dimensional quantum many-body systems. *Physical review letters*, 93(4):040502, 2004.
- [25] Andrew John Daley, Corinna Kollath, Ulrich Schollwöck, and Guifré Vidal. Time-dependent density-matrix renormalization-group using adaptive effective hilbert spaces. *Journal of Statistical Mechanics: Theory and Experiment*, 2004(04):P04005, 2004.
- [26] Steven R White and Adrian E Feiguin. Real-time evolution using the density matrix renormalization group. *Physical review letters*, 93(7):076401, 2004.
- [27] Nicholas Metropolis, Arianna W Rosenbluth, Marshall N Rosenbluth, Augusta H Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092, 1953.
- [28] Hans Gerd Evertz, Gideon Lana, and Mihai Marcu. Cluster algorithm for vertex models. *Physical review letters*, 70(7):875, 1993.
- [29] Bernard Borodol Beard and U-J Wiese. Simulations of discrete quantum systems in continuous euclidean time. *Physical review letters*, 77(25):5130, 1996.
- [30] Beat Ammon, Hans Gerd Evertz, Naoki Kawashima, Matthias Troyer, and Beat Frischmuth. Quantum monte carlo loop algorithm for the t- j model. *Physical Review B*, 58(8):4304, 1998.
- [31] Ulli Wolff. Collective monte carlo updating for spin systems. *Physical Review Letters*, 62(4):361, 1989.
- [32] Robert H Swendsen and Jian-Sheng Wang. Nonuniversal critical dynamics in monte carlo simulations. *Physical review letters*, 58(2):86, 1987.

- [33] Lode Pollet. Recent developments in quantum monte carlo simulations with applications for cold gases. *Reports on progress in physics*, 75(9):094501, 2012.
- [34] Matthias Troyer and Uwe-Jens Wiese. Computational complexity and fundamental limitations to fermionic quantum monte carlo simulations. *Physical review letters*, 94(17):170201, 2005.
- [35] Antoine Georges, Gabriel Kotliar, Werner Krauth, and Marcelo J Rozenberg. Dynamical mean-field theory of strongly correlated fermion systems and the limit of infinite dimensions. *Reviews of Modern Physics*, 68(1):13, 1996.
- [36] Thomas Maier, Mark Jarrell, Thomas Pruschke, and Matthias H Hettler. Quantum cluster theories. *Reviews of Modern Physics*, 77(3):1027, 2005.
- [37] Philippe Corboz, Román Orús, Bela Bauer, and Guifré Vidal. Simulation of strongly correlated fermions in two spatial dimensions with fermionic projected entangled-pair states. *Physical Review B*, 81(16):165104, 2010.
- [38] Glen Evenbly and Guifré Vidal. Tensor network states and geometry. *Journal of Statistical Physics*, 145(4):891–918, 2011.
- [39] Richard P Feynman. Simulating physics with computers. *Int. J. Theor. Phys*, 21(6/7), 1982.
- [40] Marc Kaplan, Gaëtan Leurent, Anthony Leverrier, and María Naya-Plasencia. Breaking symmetric cryptosystems using quantum period finding. In *Annual International Cryptology Conference*, pages 207–237. Springer, 2016.
- [41] Mark Adcock and Richard Cleve. A quantum goldreich-levin theorem with cryptographic applications. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 323–334. Springer, 2002.
- [42] Ashley Montanaro. Quantum algorithms: an overview. *npj Quantum Information*, 2(1):1–8, 2016.
- [43] Frank Gaitan. Finding flows of a navier–stokes fluid through quantum computing. *npj Quantum Information*, 6(1):1–6, 2020.
- [44] J Georg Bednorz and K Alex Müller. Perovskite-type oxides—the new approach to high- T_c superconductivity. *Reviews of Modern Physics*, 60(3):585, 1988.
- [45] J George Bednorz and K Alex Müller. Possible high T_c superconductivity in the Ba-La-Cu-O system. *Zeitschrift für Physik B Condensed Matter*, 64(2):189–193, 1986.

- [46] John P Gaebler, Ting Rei Tan, Y Lin, Y Wan, R Bowler, Adam C Keith, S Glancy, K Coakley, E Knill, D Leibfried, et al. High-fidelity universal gate set for be 9+ ion qubits. *Physical review letters*, 117(6):060505, 2016.
- [47] C. J. Ballance, T. P. Harty, N. M. Linke, M. A. Sepiol, and D. M. Lucas. High-fidelity quantum logic gates using trapped-ion hyperfine qubits. *Phys. Rev. Lett.*, 117:060504, Aug 2016.
- [48] Rami Barends, Julian Kelly, Anthony Megrant, Andrzej Veitia, Daniel Sank, Evan Jeffrey, Ted C White, Josh Mutus, Austin G Fowler, Brooks Campbell, et al. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508(7497):500–503, 2014.
- [49] MA Rol, F Battistel, FK Malinowski, CC Bultink, BM Tarasinski, R Vollmer, N Haider, N Muthusubramanian, A Bruno, BM Terhal, et al. Fast, high-fidelity conditional-phase gate exploiting leakage interference in weakly anharmonic superconducting qubits. *Physical review letters*, 123(12):120502, 2019.
- [50] Niels Jakob S e Loft, Morten Kjaergaard, Lasse Bj orn Kristensen, Christian Kraglund Andersen, Thorvald W Larsen, Simon Gustavsson, William D Oliver, and Nikolaj T Zinner. Quantum interference device for controlled two-qubit operations. *npj Quantum Information*, 6(1):1–9, 2020.
- [51] Jens Koch, M Yu Terri, Jay Gambetta, Andrew A Houck, DI Schuster, J Majer, Alexandre Blais, Michel H Devoret, Steven M Girvin, and Robert J Schoelkopf. Charge-insensitive qubit design derived from the cooper pair box. *Physical Review A*, 76(4):042319, 2007.
- [52] Jason R Petta, Alexander Comstock Johnson, Jacob M Taylor, Edward A Laird, Amir Yacoby, Mikhail D Lukin, Charles M Marcus, Micah P Hanson, and Arthur C Gossard. Coherent manipulation of coupled electron spins in semiconductor quantum dots. *Science*, 309(5744):2180–2184, 2005.
- [53] Karl D Petersson, Louis W McFaul, Michael D Schroer, Minkyung Jung, Jacob M Taylor, Andrew A Houck, and Jason R Petta. Circuit quantum electrodynamics with a spin qubit. *Nature*, 490(7420):380–383, 2012.
- [54] Peter W Shor. Fault-tolerant quantum computation. In *Proceedings of 37th conference on foundations of computer science*, pages 56–65. IEEE, 1996.
- [55] Andrew M Steane. Active stabilization, quantum computation, and quantum state synthesis. *Physical Review Letters*, 78(11):2252, 1997.

- [56] Sergey Bravyi and Robert König. Classification of topologically protected gates for local stabilizer codes. *Physical review letters*, 110(17):170503, 2013.
- [57] Thorvald Wadum Larsen, Karl David Petersson, Ferdinand Kuemmeth, Thomas Sand Jespersen, Peter Krogstrup, Jesper Nygård, and Charles M Marcus. Semiconductor-nanowire-based superconducting qubit. *Physical review letters*, 115(12):127001, 2015.
- [58] George EP Box and George C Tiao. *Bayesian inference in statistical analysis*, volume 40. John Wiley & Sons, 2011.
- [59] Ian Hincks, Thomas Alexander, Michal Kononenko, Benjamin Soloway, and David G Cory. Hamiltonian learning with online bayesian experiment design in practice. *arXiv preprint arXiv:1806.02427*, 2018.
- [60] Alessandro Lumino, Emanuele Polino, Adil S Rab, Giorgio Milani, Nicolò Spagnolo, Nathan Wiebe, and Fabio Sciarrino. Experimental phase estimation enhanced by machine learning. *Physical Review Applied*, 10(4):044033, 2018.
- [61] A Yu Kitaev. Quantum computations: algorithms and error correction. *Russian Mathematical Surveys*, 52(6):1191, 1997.
- [62] H. F. Trotter. On the product of semi-groups of operators. *Proc. Amer. Math. Soc.*, 10:545, 1959.
- [63] Masuo Suzuki. Generalized trotter’s formula and systematic approximants of exponential operators and inner derivations with applications to many-body problems. *Commun. Math. Phys.*, 51(2):183–190, 1976.
- [64] Matthias Troyer. *Computational Quantum Physics*. 2015.
- [65] Jacob Jordan, Roman Orús, Guifre Vidal, Frank Verstraete, and J Ignacio Cirac. Classical simulation of infinite-size quantum lattice systems in two spatial dimensions. *Physical review letters*, 101(25):250602, 2008.
- [66] A. Albuquerque, F. Alet, P. Corboz, P. Dayal, A. Feiguin, S. Fuchs, L. Gamper, E. Gull, S. Gurtler, and A. Honecker. The ALPS project release 1.3: Open-source software for strongly correlated systems. *J. Magn. Magn. Mater.*, 310(2):1187–1193, March 2007.
- [67] B. Bauer and et al. (ALPS Collaboration). The ALPS project release 2.0: open source software for strongly correlated systems. *J. Stat. Mech.*, 2011(05):P05001, 2011.

- [68] David Poulin, Matthew B Hastings, Dave Wecker, Nathan Wiebe, Andrew C Doherty, and Matthias Troyer. The trotter step size required for accurate quantum simulation of quantum chemistry. *arXiv preprint arXiv:1406.4920*, 2014.
- [69] Immanuel Bloch, Jean Dalibard, and Wilhelm Zwerger. Many-body physics with ultracold gases. *Reviews of modern physics*, 80(3):885, 2008.
- [70] Morten Kjaergaard, Mollie E Schwartz, Jochen Braumüller, Philip Krantz, Joel I-J Wang, Simon Gustavsson, and William D Oliver. Superconducting qubits: Current state of play. *Annual Review of Condensed Matter Physics*, 11, 2019.
- [71] Colin D Bruzewicz, John Chiaverini, Robert McConnell, and Jeremy M Sage. Trapped-ion quantum computing: Progress and challenges. *Applied Physics Reviews*, 6(2):021314, 2019.
- [72] Matthew B Hastings. Locality in quantum and markov dynamics on lattices and networks. *Physical review letters*, 93(14):140402, 2004.
- [73] Tilman Esslinger. Fermi-hubbard physics with atoms in an optical lattice. 2010.
- [74] Ian Affleck and J Brad Marston. Large-n limit of the heisenberg-hubbard model: Implications for high-t c superconductors. *Physical Review B*, 37(7):3774, 1988.
- [75] Christoph J Halboth and Walter Metzner. d-wave superconductivity and pomeranchuk instability in the two-dimensional hubbard model. *Physical review letters*, 85(24):5162, 2000.
- [76] Steven R White, Douglas J Scalapino, Robert L Sugar, EY Loh, James E Gubernatis, and Richard T Scalettar. Numerical study of the two-dimensional hubbard model. *Physical Review B*, 40(1):506, 1989.
- [77] P. Jordan and E. Wigner. Über das Paulische Äquivalenzverbot. *Z. Phys.*, 47(9-10):631–651, 1928.
- [78] James D. Whitfield, Jacob Biamonte, and Alán Aspuru-Guzik. Simulation of electronic structure hamiltonians using quantum computers. *Molecular Physics*, 109(5):735–750, Mar 2011.
- [79] M. B. Hastings, D. Wecker, B. Bauer, and M. Troyer. Improving quantum algorithms for quantum chemistry, 2014.
- [80] Craig Gidney. Halving the cost of quantum addition. *Quantum*, 2:74, Jun 2018.
- [81] Dave Wecker, Matthew B. Hastings, Nathan Wiebe, Bryan K. Clark, Chetan Nayak, and Matthias Troyer. Solving strongly correlated electron models on a quantum computer. *Physical Review A*, 92(6), Dec 2015.

- [82] Krysta Svore, Alan Geller, Matthias Troyer, John Azariah, Christopher Granade, Bettina Heim, Vadym Kliuchnikov, Mariia Mykhailova, Andres Paz, and Martin Roetteler. Q# enabling scalable quantum computing and development with a high-level dsl. In *Proceedings of the Real World Domain Specific Languages Workshop 2018*, pages 1–10, 2018.
- [83] Sergey Bravyi and Alexei Kitaev. Universal quantum computation with ideal clifford gates and noisy ancillas. *Physical Review A*, 71(2):022316, 2005.
- [84] Alex Bocharov, Martin Roetteler, and Krysta M. Svore. Efficient synthesis of universal repeat-until-success quantum circuits. *Physical Review Letters*, 114(8), Feb 2015.
- [85] Pierre Pfeuty. The one-dimensional ising model with a transverse field. *ANNALS of Physics*, 57(1):79–90, 1970.
- [86] Andrew M Childs, Dmitri Maslov, Yunseong Nam, Neil J Ross, and Yuan Su. Toward the first quantum simulation with quantum speedup. *Proceedings of the National Academy of Sciences*, 115(38):9456–9461, 2018.
- [87] Richard P Feynman. Simulating physics with computers. *Int. J. Theor. Phys*, 21(6/7), 1999.
- [88] Seth Lloyd. Universal quantum simulators. *Science*, pages 1073–1078, 1996.
- [89] S Boixo, E Knill, and RD Somma. Fast quantum algorithms for traversing paths of eigenstates. *arXiv preprint arXiv:1005.3034*, 2010.
- [90] David Poulin, Angie Qarry, Rolando Somma, and Frank Verstraete. Quantum simulation of time-dependent hamiltonians and the convenient illusion of hilbert space. *Physical review letters*, 106(17):170501, 2011.
- [91] Sadegh Raeesi, Nathan Wiebe, and Barry C Sanders. Quantum-circuit design for efficient simulations of many-body quantum dynamics. *New Journal of Physics*, 14(10):103017, 2012.
- [92] Andrew M Childs and Yuan Su. Nearly optimal lattice simulation by product formulas. *Physical review letters*, 123(5):050503, 2019.
- [93] Ben P Lanyon, Cornelius Hempel, Daniel Nigg, Markus Müller, Rene Gerritsma, F Zähringer, Philipp Schindler, Julio T Barreiro, Markus Rambach, Gerhard Kirchmair, et al. Universal digital quantum simulation with trapped ions. *Science*, 334(6052):57–61, 2011.

- [94] Kenneth R Brown, Robert J Clark, and Isaac L Chuang. Limitations of quantum simulation examined by simulating a pairing hamiltonian using nuclear magnetic resonance. *Physical review letters*, 97(5):050504, 2006.
- [95] R Barends, L Lamata, J Kelly, L García-Álvarez, AG Fowler, A Megrant, E Jeffrey, TC White, D Sank, JY Mutus, et al. Digital quantum simulation of fermionic models with a superconducting circuit. *Nature communications*, 6(1):1–7, 2015.
- [96] Ish Dhand and Barry C Sanders. Stability of the trotter–suzuki decomposition. *Journal of Physics A: Mathematical and Theoretical*, 47(26):265206, 2014.
- [97] Suguru Endo, Qi Zhao, Ying Li, Simon Benjamin, and Xiao Yuan. Mitigating algorithmic errors in a hamiltonian simulation. *Physical Review A*, 99(1):012334, 2019.
- [98] David Poulin. Lieb-robinson bound and locality for general markovian quantum dynamics. *Physical review letters*, 104(19):190401, 2010.
- [99] Thomas Barthel and Martin Kliesch. Quasilocality and efficient simulation of markovian quantum dynamics. *Physical review letters*, 108(23):230504, 2012.
- [100] Elliott H Lieb and Derek W Robinson. The finite group velocity of quantum spin systems. In *Statistical mechanics*, pages 425–431. Springer, 1972.
- [101] Max Born and Robert Oppenheimer. Zur quantentheorie der molekeln. *Annalen der physik*, 389(20):457–484, 1927.
- [102] Walter Kohn, Axel D Becke, and Robert G Parr. Density functional theory of electronic structure. *The Journal of Physical Chemistry*, 100(31):12974–12980, 1996.
- [103] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. Advances in quantum metrology. *Nature photonics*, 5(4):222, 2011.
- [104] Mark Srednicki. Chaos and quantum thermalization. *Physical Review E*, 50(2):888, 1994.
- [105] Markus Heyl, Philipp Hauke, and Peter Zoller. Quantum localization bounds trotter errors in digital quantum simulation. *Science advances*, 5(4):eaau8342, 2019.
- [106] A Yu Kitaev. Quantum measurements and the abelian stabilizer problem. *arXiv preprint quant-ph/9511026*, 1995.

- [107] K. Temme, T.J. Osborne, K. Vollbrecht, David Poulin, and F. Verstraete. Quantum metropolis sampling. *Nature*, 471:87, March 2011.
- [108] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. Quantum computation by adiabatic evolution. *arXiv preprint quant-ph/0001106*, 2000.
- [109] Dorit Aharonov and Amnon Ta-Shma. Adiabatic quantum state generation and statistical zero knowledge. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 20–29, 2003.
- [110] Rolando D Somma, Sergio Boixo, Howard Barnum, and Emanuel Knill. Quantum simulations of classical annealing processes. *Physical review letters*, 101(13):130504, 2008.
- [111] David Poulin and Pawel Wocjan. Preparing ground states of quantum many-body systems on a quantum computer. *Physical review letters*, 102(13):130503, 2009.
- [112] David Poulin and Pawel Wocjan. Sampling from the thermal quantum gibbs state and evaluating partition functions with a quantum computer. *Physical review letters*, 103(22):220502, 2009.
- [113] Lov K Grover. Quantum mechanics helps in searching for a needle in a haystack. *Physical review letters*, 79(2):325, 1997.
- [114] Elliott Lieb, Theodore Schultz, and Daniel Mattis. Two soluble models of an antiferromagnetic chain. *Annals of Physics*, 16(3):407–466, 1961.
- [115] Hendrik A Kramers and Gregory H Wannier. Statistics of the two-dimensional ferromagnet. part i. *Physical Review*, 60(3):252, 1941.
- [116] David P DiVincenzo. The physical implementation of quantum computation. *Fortschritte der Physik: Progress of Physics*, 48(9-11):771–783, 2000.
- [117] M. A. Nielsen and I. L. Chuang. Quantum computation and quantum information. *Cambridge University Press*, 2010.
- [118] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O’Brien. Quantum computers. *Nature*, 464(7285):45–53, 2010.
- [119] David Aasen, Michael Hell, Ryan V. Mishmash, Andrew Higginbotham, Jeroen Danon, Martin Leijnse, Thomas S. Jespersen, Joshua A. Folk, Charles M. Marcus, Karsten Flensberg, and Jason Alicea. Milestones toward majorana-based quantum computing. *Phys. Rev. X*, 6:031016, Aug 2016.

- [120] Niels Jakob Loft, Morten Kjærgaard, Lasse Kristensen, Christian Kraglund Andersen, Thorvald Larsen, Simon Gustavsson, William Oliver, and Nikolaj Zinner. High-fidelity conditional two-qubit swapping gate using tunable ancillas. *APS*, 2019:L29–003, 2019.
- [121] Alexandre Blais, Steven M Girvin, and William D Oliver. Quantum information processing and quantum optics with circuit quantum electrodynamics. *Nature Physics*, pages 1–10, 2020.
- [122] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, 2019.
- [123] Matthew Reed. *Entanglement and quantum error correction with superconducting qubits*. PhD thesis, 2013.
- [124] Steven M Girvin. Circuit qed: superconducting qubits coupled to microwave photons. *Quantum Machines: Measurement and Control of Engineered Quantum Systems*, 113:2, 2011.
- [125] Alexandre Blais, Ren-Shou Huang, Andreas Wallraff, Steven M Girvin, and R Jun Schoelkopf. Cavity quantum electrodynamics for superconducting electrical circuits: An architecture for quantum computation. *Physical Review A*, 69(6):062320, 2004.
- [126] John Clarke and Frank K Wilhelm. Superconducting quantum bits. *Nature*, 453(7198):1031–1042, 2008.
- [127] Michel H Devoret and Robert J Schoelkopf. Superconducting circuits for quantum information: an outlook. *Science*, 339(6124):1169–1174, 2013.
- [128] G De Lange, B Van Heck, A Bruno, DJ Van Woerkom, A Geresdi, SR Plissard, EPAM Bakkers, AR Akhmerov, and L DiCarlo. Realization of microwave quantum circuits using hybrid superconducting-semiconducting nanowire josephson elements. *Physical review letters*, 115(12):127002, 2015.
- [129] Alexandre Avraamovitch Golubov, M Yu Kupriyanov, and E Il’Ichev. The current-phase relation in josephson junctions. *Reviews of modern physics*, 76(2):411, 2004.
- [130] Jonathan R Friedman, Vijay Patel, Wei Chen, SK Tolpygo, and James E Lukens. Quantum superposition of distinct macroscopic states. *nature*, 406(6791):43–46, 2000.

- [131] Caspar H Van Der Wal, ACJ Ter Haar, FK Wilhelm, RN Schouten, CJPM Harman, TP Orlando, Seth Lloyd, and JE Mooij. Quantum superposition of macroscopic persistent-current states. *Science*, 290(5492):773–777, 2000.
- [132] John M Martinis, S Nam, J Aumentado, and C Urbina. Rabi oscillations in a large josephson-junction qubit. *Physical review letters*, 89(11):117901, 2002.
- [133] Yu Nakamura, Yu A Pashkin, and Jaw Shen Tsai. Coherent control of macroscopic quantum states in a single-cooper-pair box. *nature*, 398(6730):786–788, 1999.
- [134] Vincent Bouchiat, D Vion, Ph Joyez, D Esteve, and MH Devoret. Quantum coherence with a single cooper pair. *Physica Scripta*, 1998(T76):165, 1998.
- [135] Anders Kringhøj, Lucas Casparis, Michael Hell, Thorvald Wadum Larsen, Ferdinand Kuemmeth, M Leijnse, Karsten Flensberg, P Krogstrup, Jesper Nygård, Karl David Petersson, et al. Anharmonicity of a superconducting qubit with a few-mode josephson junction. *Physical Review B*, 97(6):060508, 2018.
- [136] Andreas Wallraff, David I Schuster, Alexandre Blais, Luigi Frunzio, R-S Huang, Johannes Majer, Sameer Kumar, Steven M Girvin, and Robert J Schoelkopf. Strong coupling of a single photon to a superconducting qubit using circuit quantum electrodynamics. *Nature*, 431(7005):162–167, 2004.
- [137] Henning Labuhn, Daniel Barredo, Sylvain Ravets, Sylvain De Léséleuc, Tommaso Macrì, Thierry Lahaye, and Antoine Browaeys. Tunable two-dimensional arrays of single rydberg atoms for realizing quantum ising models. *Nature*, 534(7609):667–670, 2016.
- [138] Hanhee Paik, DI Schuster, Lev S Bishop, G Kirchmair, G Catelani, AP Sears, BR Johnson, MJ Reagor, L Frunzio, LI Glazman, et al. Observation of high coherence in josephson junction qubits measured in a three-dimensional circuit qed architecture. *Physical Review Letters*, 107(24):240501, 2011.
- [139] Rami Barends, Julian Kelly, Anthony Megrant, Daniel Sank, Evan Jeffrey, Yu Chen, Yi Yin, Ben Chiaro, Josh Mutus, Charles Neill, et al. Coherent josephson qubit suitable for scalable quantum integrated circuits. *Physical review letters*, 111(8):080502, 2013.
- [140] AA Houck, JA Schreier, BR Johnson, JM Chow, Jens Koch, JM Gambetta, DI Schuster, L Frunzio, MH Devoret, SM Girvin, et al. Controlling the spontaneous emission of a superconducting transmon qubit. *Physical review letters*, 101(8):080502, 2008.

- [141] Matthew D Reed, Blake R Johnson, Andrew A Houck, Leonardo DiCarlo, Jerry M Chow, David I Schuster, Luigi Frunzio, and Robert J Schoelkopf. Fast reset and suppressing spontaneous emission of a superconducting qubit. *Applied Physics Letters*, 96(20):203110, 2010.
- [142] Theodore Walter, Philipp Kurpiers, Simone Gasparinetti, Paul Magnard, Anton Potočnik, Yves Salathé, Marek Pechal, Mintu Mondal, Markus Oppliger, Christopher Eichler, et al. Rapid high-fidelity single-shot dispersive readout of superconducting qubits. *Physical Review Applied*, 7(5):054020, 2017.
- [143] John M Martinis, Ken B Cooper, Robert McDermott, Matthias Steffen, Markus Ansmann, KD Osborn, Katarina Cicak, Seongshik Oh, David P Pappas, Raymond W Simmonds, et al. Decoherence in josephson qubits from dielectric loss. *Physical review letters*, 95(21):210503, 2005.
- [144] A Lupaşcu, S Saito, T Picot, PC De Groot, CJPM Harmans, and JE Mooij. Quantum non-demolition measurement of a superconducting two-level system. *nature physics*, 3(2):119–123, 2007.
- [145] Evan Jeffrey, Daniel Sank, JY Mutus, TC White, J Kelly, R Barends, Y Chen, Z Chen, B Chiaro, A Dunsworth, et al. Fast accurate state measurement with superconducting qubits. *Physical review letters*, 112(19):190504, 2014.
- [146] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [147] MD Reed, L DiCarlo, BR Johnson, L Sun, DI Schuster, L Frunzio, and RJ Schoelkopf. High-fidelity readout in circuit quantum electrodynamics using the jaynes-cummings nonlinearity. *Physical review letters*, 105(17):173601, 2010.
- [148] Johannes Heinsoo, Christian Kraglund Andersen, Ants Remm, Sebastian Krinner, Theodore Walter, Yves Salathé, Simone Gasparinetti, Jean-Claude Besse, Anton Potočnik, Andreas Wallraff, et al. Rapid high-fidelity multiplexed readout of superconducting qubits. *Physical Review Applied*, 10(3):034040, 2018.
- [149] Yu Chen, D Sank, P O’Malley, T White, R Barends, B Chiaro, J Kelly, E Lucero, M Mariantoni, A Megrant, et al. Multiplexed dispersive readout of superconducting phase qubits. *Applied Physics Letters*, 101(18):182601, 2012.
- [150] JM Hornibrook, JI Colless, AC Mahoney, XG Croot, S Blanvillain, H Lu, AC Gosard, and DJ Reilly. Frequency multiplexing for readout of spin qubits. *Applied Physics Letters*, 104(10):103108, 2014.

- [151] Andreas Wallraff. *Lecture notes for Quantum Information Processing II: Implementations. Introduction to Quantum Information Science*. 2016.
- [152] Lorenza Viola and Emanuel Knill. Robust dynamical decoupling of quantum systems with bounded controls. *Physical Review Letters*, 90(3):037901, 2003.
- [153] Alexandre M Souza, Gonzalo A Alvarez, and Dieter Suter. Robust dynamical decoupling for quantum computing and quantum memory. *Physical review letters*, 106(24):240501, 2011.
- [154] Erwin L Hahn. Spin echoes. *Physical review*, 80(4):580, 1950.
- [155] Michael A Nielsen. A simple formula for the average gate fidelity of a quantum dynamical operation. *Physics Letters A*, 303(4):249–252, 2002.
- [156] Antonio D Córcoles, Jay M Gambetta, Jerry M Chow, John A Smolin, Matthew Ware, Joel Strand, Britton LT Plourde, and Matthias Steffen. Process verification of two-qubit quantum gates by randomized benchmarking. *Physical Review A*, 87(3):030301, 2013.
- [157] Kenton R Brown, Andrew C Wilson, Yves Colombe, C Ospelkaus, Adam M Meier, E Knill, D Leibfried, and David J Wineland. Single-qubit-gate error below 10^{-4} in a trapped ion. *Physical Review A*, 84(3):030303, 2011.
- [158] Easwar Magesan, Jay M Gambetta, and Joseph Emerson. Scalable and robust randomized benchmarking of quantum processes. *Physical review letters*, 106(18):180504, 2011.
- [159] Joel J Wallman and Steven T Flammia. Randomized benchmarking with confidence. *New Journal of Physics*, 16(10):103032, 2014.
- [160] Sarah Sheldon, Lev S Bishop, Easwar Magesan, Stefan Filipp, Jerry M Chow, and Jay M Gambetta. Characterizing errors on qubit operations via iterative randomized benchmarking. *Physical Review A*, 93(1):012301, 2016.
- [161] Julian Kelly, R Barends, B Campbell, Y Chen, Z Chen, B Chiaro, A Dunsworth, Austin G Fowler, I-C Hoi, E Jeffrey, et al. Optimal quantum control using randomized benchmarking. *Physical review letters*, 112(24):240504, 2014.
- [162] David Gross, Yi-Kai Liu, Steven T Flammia, Stephen Becker, and Jens Eisert. Quantum state tomography via compressed sensing. *Physical review letters*, 105(15):150401, 2010.

- [163] Seth T Merkel, Jay M Gambetta, John A Smolin, Stefano Poletto, Antonio D Córcoles, Blake R Johnson, Colm A Ryan, and Matthias Steffen. Self-consistent quantum process tomography. *Physical Review A*, 87(6):062119, 2013.
- [164] Robin Blume-Kohout, John King Gamble, Erik Nielsen, Kenneth Rudinger, Jonathan Mizrahi, Kevin Fortier, and Peter Maunz. Demonstration of qubit operations below a rigorous fault tolerance threshold with gate set tomography. *Nature communications*, 8(1):1–13, 2017.
- [165] MA Rol, CC Bultink, TE O’Brien, SR De Jong, LS Theis, Xiang Fu, F Luthi, RFL Vermeulen, JC de Sterke, Alessandro Bruno, et al. Restless tuneup of high-fidelity qubit gates. *Physical Review Applied*, 7(4):041001, 2017.
- [166] Lev S Bishop, Sergey Bravyi, Andrew Cross, Jay M Gambetta, and John Smolin. Quantum volume. *Quantum Volume. Technical Report*, 2017.
- [167] A Dewes, FR Ong, V Schmitt, R Lauro, N Boulant, P Bertet, D Vion, and D Esteve. Characterization of a two-transmon processor with individual single-shot qubit readout. *Physical review letters*, 108(5):057002, 2012.
- [168] T Yamamoto, M Neeley, E Lucero, RC Bialczak, J Kelly, M Lenander, Matteo Mariantoni, AD O’Connell, D Sank, H Wang, et al. Quantum process tomography of two-qubit controlled-z and controlled-not gates using superconducting phase qubits. *Physical Review B*, 82(18):184515, 2010.
- [169] Leonardo DiCarlo, Jerry M Chow, Jay M Gambetta, Lev S Bishop, Blake R Johnson, DI Schuster, J Majer, Alexandre Blais, Luigi Frunzio, SM Girvin, et al. Demonstration of two-qubit algorithms with a superconducting quantum processor. *Nature*, 460(7252):240–244, 2009.
- [170] J Majer, JM Chow, JM Gambetta, Jens Koch, BR Johnson, JA Schreier, L Frunzio, DI Schuster, Andrew Addison Houck, Andreas Wallraff, et al. Coupling superconducting qubits via a cavity bus. *Nature*, 449(7161):443–447, 2007.
- [171] L. D. Landau. Zur theorie der energieubertragung. *Phys. Z. Sowjetunion*, 2:46–51, 1932.
- [172] C. Zener. Non-adiabatic crossing of energy levels. *Proc. R. Soc. Lond.*, 137, Sep 1932.
- [173] E. C. G Stückelberg. *Helv. Phys. Acta* 5, 369, 1932.
- [174] Ettore Majorana. Atomi orientati in campo magnetico variabile. *Il Nuovo Cimento (1924-1942)*, 9(2):43–50, 1932.

- [175] CM Quintana, KD Petersson, LW McFaul, SJ Srinivasan, Andrew A Houck, and Jason R Petta. Cavity-mediated entanglement generation via landau-zener interferometry. *Physical review letters*, 110(17):173603, 2013.
- [176] RC Bialczak, M Ansmann, M Hofheinz, M Lenander, E Lucero, M Neeley, AD O’Connell, D Sank, H Wang, M Weides, et al. Fast tunable coupler for superconducting qubits. *Physical review letters*, 106(6):060501, 2011.
- [177] Yu Chen, C Neill, P Roushan, N Leung, M Fang, R Barends, J Kelly, B Campbell, Z Chen, B Chiaro, et al. Qubit architecture with high coherence and fast tunable coupling. *Physical review letters*, 113(22):220502, 2014.
- [178] David C McKay, Stefan Filipp, Antonio Mezzacapo, Easwar Magesan, Jerry M Chow, and Jay M Gambetta. Universal gate for fixed-frequency qubits via a tunable bus. *Physical Review Applied*, 6(6):064007, 2016.
- [179] Margareta Wallquist, VS Shumeiko, and Göran Wendin. Selective coupling of superconducting charge qubits mediated by a tunable stripline cavity. *Physical Review B*, 74(22):224506, 2006.
- [180] Mika A Sillanpää, Jae I Park, and Raymond W Simmonds. Coherent quantum state storage and transfer between two phase qubits via a resonant cavity. *Nature*, 449(7161):438–442, 2007.
- [181] Martin Sandberg, CM Wilson, Fredrik Persson, Thilo Bauch, Göran Johansson, Vitaly Shumeiko, Tim Duty, and Per Delsing. Tuning the field in a microwave resonator faster than the photon lifetime. *Applied Physics Letters*, 92(20):203501, 2008.
- [182] Agustín Palacios-Laloy, Francois Nguyen, Francois Mallet, Patrice Bertet, Denis Vion, and Daniel Esteve. Tunable resonators for quantum circuits. *Journal of Low Temperature Physics*, 151(3-4):1034–1042, 2008.
- [183] Yuimaru Kubo, Cecile Grezes, Andreas Dewes, T Umeda, Junichi Isoya, H Sumiya, N Morishita, H Abe, S Onoda, T Ohshima, et al. Hybrid quantum circuit with a superconducting qubit coupled to a spin ensemble. *Physical review letters*, 107(22):220501, 2011.
- [184] Daniel R Ward, DE Savage, MG Lagally, SN Coppersmith, and MA Eriksson. Integration of on-chip field-effect transistor switches with dopantless si/sige quantum dots for high-throughput testing. *Applied Physics Letters*, 102(21):213107, 2013.

- [185] H Al-Taie, LW Smith, B Xu, P See, JP Griffiths, HE Beere, GAC Jones, DA Ritchie, MJ Kelly, and CG Smith. Cryogenic on-chip multiplexer for the study of quantum transport in 256 split-gate devices. *Applied Physics Letters*, 102(24):243102, 2013.
- [186] JM Hornibrook, JI Colless, ID Conway Lamb, SJ Pauka, H Lu, AC Gossard, JD Watson, GC Gardner, S Fallahi, MJ Manfra, et al. Cryogenic control architecture for large-scale quantum computing. *Physical Review Applied*, 3(2):024010, 2015.
- [187] Zhenyi Qi, Hong-Yi Xie, Javad Shabani, Vladimir E Manucharyan, Alex Levchenko, and Maxim G Vavilov. Controlled-z gate for transmon qubits coupled by semiconductor junctions. *Physical Review B*, 97(13):134518, 2018.
- [188] Anders Sørensen and Klaus Mølmer. Quantum computation with ions in thermal motion. *Physical review letters*, 82(9):1971, 1999.
- [189] Simon E Nigg, Hanhee Paik, Brian Vlastakis, Gerhard Kirchmair, Shyam Shankar, Luigi Frunzio, MH Devoret, RJ Schoelkopf, and SM Girvin. Black-box superconducting circuit quantization. *Physical Review Letters*, 108(24):240502, 2012.
- [190] S Filipp, M Göppl, JM Fink, M Baur, R Bianchetti, L Steffen, and A Wallraff. Multimode mediated qubit-qubit coupling and dark-state symmetries in circuit quantum electrodynamics. *Physical Review A*, 83(6):063827, 2011.
- [191] Lucas Casparis, TW Larsen, MS Olsen, F Kuemmeth, P Krogstrup, Jesper Nygård, KD Petersson, and CM Marcus. Gatemon benchmarking and two-qubit operations. *Physical review letters*, 116(15):150505, 2016.
- [192] Chris Macklin, K O’Brien, D Hover, ME Schwartz, V Bolkhovsky, X Zhang, WD Oliver, and I Siddiqi. A near-quantum-limited josephson traveling-wave parametric amplifier. *Science*, 350(6258):307–310, 2015.
- [193] X Mi, JV Cady, DM Zajac, J Stehlik, LF Edge, and Jason R Petta. Circuit quantum electrodynamics architecture for gate-defined quantum dots in silicon. *Applied Physics Letters*, 110(4):043502, 2017.
- [194] Anthony Megrant, Charles Neill, Rami Barends, Ben Chiaro, Yu Chen, Ludwig Feigl, Julian Kelly, Erik Lucero, Matteo Mariantoni, Peter JJ O’Malley, et al. Planar superconducting resonators with internal quality factors above one million. *Applied Physics Letters*, 100(11):113510, 2012.

- [195] Zijun Chen, Anthony Megrant, Julian Kelly, Rami Barends, Joerg Bochmann, Yu Chen, Ben Chiaro, Andrew Dunsworth, Evan Jeffrey, JY Mutus, et al. Fabrication and characterization of aluminum airbridges for superconducting microwave circuits. *Applied Physics Letters*, 104(5):052602, 2014.
- [196] Yun-Pil Shim and Charles Tahan. Semiconductor-inspired design principles for superconducting quantum computing. *Nature communications*, 7(1):1–8, 2016.
- [197] U Patel, Ivan V Pechenezhskiy, BLT Plourde, MG Vavilov, and R McDermott. Phonon-mediated quasiparticle poisoning of superconducting microwave resonators. *Physical Review B*, 96(22):220501, 2017.
- [198] Tomas Jochym-O’Connor and Stephen D Bartlett. Stacked codes: Universal fault-tolerant quantum computation in a two-dimensional layout. *Physical Review A*, 93(2):022323, 2016.
- [199] Iulia Buluta and Franco Nori. Quantum simulators. *Science*, 326(5949):108–111, 2009.
- [200] Steven J Weber. Gatemons get serious. *Nature Nanotechnology*, 13(10):877–878, 2018.
- [201] Sarah Sheldon, Easwar Magesan, Jerry M Chow, and Jay M Gambetta. Procedure for systematically tuning up cross-talk in the cross-resonance gate. *Physical Review A*, 93(6):060302, 2016.
- [202] MD Hutchings, Jared B Hertzberg, Yebin Liu, Nicholas T Bronn, George A Keefe, Markus Brink, Jerry M Chow, and BLT Plourde. Tunable superconducting qubits with flux-independent coherence. *Physical Review Applied*, 8(4):044003, 2017.
- [203] Julian Kelly, Rami Barends, Austin G Fowler, Anthony Megrant, Evan Jeffrey, Theodore C White, Daniel Sank, Josh Y Mutus, Brooks Campbell, Yu Chen, et al. State preservation by repetitive error detection in a superconducting quantum circuit. *Nature*, 519(7541):66–69, 2015.
- [204] Charles Neill, Pedran Roushan, K Kechedzhi, Sergio Boixo, Sergei V Isakov, V Smelyanskiy, A Megrant, B Chiaro, A Dunsworth, K Arya, et al. A blueprint for demonstrating quantum supremacy with superconducting qubits. *Science*, 360(6385):195–199, 2018.
- [205] D Rosenberg, D Kim, R Das, D Yost, S Gustavsson, D Hover, P Krantz, A Melville, L Racz, GO Samach, et al. 3d integrated superconducting qubits. *npj quantum information*, 3(1):1–5, 2017.

- [206] Morten Kjaergaard, Henri Juhani Suominen, MP Nowak, AR Akhmerov, J Shabani, CJ Palmstrøm, Fabrizio Nichele, and Charles M Marcus. Transparent semiconductor-superconductor interface and induced gap in an epitaxial heterostructure josephson junction. *Physical Review Applied*, 7(3):034029, 2017.
- [207] F Luthi, T Stavenga, OW Enzing, A Bruno, C Dickel, NK Langford, MA Rol, Thomas Sand Jespersen, Jesper Nygård, P Krogstrup, et al. Evolution of nanowire transmon qubits and their coherence in a magnetic field. *Physical review letters*, 120(10):100502, 2018.
- [208] J Shabani, M Kjaergaard, HJ Suominen, Younghyun Kim, F Nichele, K Pakrouski, T Stankevic, Roman M Lutchny, P Krogstrup, R Feidenhans, et al. Two-dimensional epitaxial superconductor-semiconductor heterostructures: A platform for topological superconducting networks. *Physical Review B*, 93(15):155402, 2016.
- [209] Simon Abay, Henrik Nilsson, Fan Wu, HQ Xu, CM Wilson, and Per Delsing. High critical-current superconductor-inas nanowire-superconductor junctions. *Nano letters*, 12(11):5622–5625, 2012.
- [210] PJJ O’Malley, J Kelly, R Barends, B Campbell, Y Chen, Z Chen, B Chiaro, A Dunsworth, AG Fowler, I-C Hoi, et al. Qubit metrology of ultralow phase noise using randomized benchmarking. *Physical Review Applied*, 3(4):044009, 2015.
- [211] Jonas Bylander, Simon Gustavsson, Fei Yan, Fumiki Yoshihara, Khalil Harrabi, George Fitch, David G Cory, Yasunobu Nakamura, Jaw-Shen Tsai, and William D Oliver. Noise spectroscopy through dynamical decoupling with a superconducting flux qubit. *Nature Physics*, 7(7):565–570, 2011.
- [212] Jerzy Krupka, John G Hartnett, and Mirosław Piersa. Permittivity and microwave absorption of semi-insulating inp at microwave frequencies. *Applied Physics Letters*, 98(11):112112, 2011.
- [213] T. S. Frey. *Interaction between Quantum Dots and Superconducting Microwave Resonators*. PhD thesis, ETH Zurich, 2013. PhD Thesis.
- [214] A Bruno, G De Lange, S Asaad, KL Van Der Enden, NK Langford, and L DiCarlo. Reducing intrinsic loss in superconducting resonators by surface treatment and deep etching of silicon substrates. *Applied Physics Letters*, 106(18):182601, 2015.
- [215] OE Dial, Michael Dean Shulman, Shannon Pasca Harvey, H Bluhm, V Umansky, and Amir Yacoby. Charge noise spectroscopy using coherent exchange oscillations in a singlet-triplet qubit. *Physical review letters*, 110(14):146804, 2013.

- [216] Max Hofheinz, H Wang, Markus Ansmann, Radoslaw C Bialczak, Erik Lucero, Matthew Neeley, AD O'connell, Daniel Sank, J Wenner, John M Martinis, et al. Synthesizing arbitrary quantum states in a superconducting resonator. *Nature*, 459(7246):546–549, 2009.
- [217] MC Cassidy, A Bruno, S Rubbert, M Irfan, J Kammhuber, RN Schouten, AR Akhmerov, and LP Kouwenhoven. Demonstration of an ac josephson junction laser. *Science*, 355(6328):939–942, 2017.
- [218] Richard Versluis, Stefano Poletto, Nader Khammassi, Brian Tarasinski, Nadia Haider, David J Michalak, Alessandro Bruno, Koen Bertels, and Leonardo Di-Carlo. Scalable quantum circuit and control for a superconducting surface code. *Physical Review Applied*, 8(3):034021, 2017.
- [219] First commercial quantum computing control system for europe's quantum computer. *Zurich Instruments News*, 2018.
- [220] Kenneth Levenberg. A method for the solution of certain non-linear problems in least squares. *Quarterly of applied mathematics*, 2(2):164–168, 1944.
- [221] Donald W Marquardt. An algorithm for least-squares estimation of nonlinear parameters. *Journal of the society for Industrial and Applied Mathematics*, 11(2):431–441, 1963.
- [222] Nathan Wiebe, Christopher Granade, Christopher Ferrie, and David Cory. Quantum hamiltonian learning using imperfect quantum resources. *Physical Review A*, 89(4):042314, 2014.
- [223] Christopher E Granade, Christopher Ferrie, Nathan Wiebe, and David G Cory. Robust online hamiltonian learning. *New Journal of Physics*, 14(10):103013, 2012.
- [224] Nathan Wiebe, Christopher Granade, Christopher Ferrie, and David G Cory. Hamiltonian learning and certification using quantum resources. *Physical review letters*, 112(19):190501, 2014.
- [225] Stefano Paesani, Andreas A Gentile, Raffaele Santagati, Jianwei Wang, Nathan Wiebe, David P Tew, Jeremy L O'Brien, and Mark G Thompson. Experimental bayesian quantum phase estimation on a silicon photonic chip. *Physical review letters*, 118(10):100503, 2017.
- [226] Nathan Wiebe and Chris Granade. Efficient bayesian phase estimation. *Physical review letters*, 117(1):010503, 2016.

- [227] Jianwei Wang, Stefano Paesani, Raffaele Santagati, Sebastian Knauer, Antonio A Gentile, Nathan Wiebe, Maurangelo Petruzzella, Jeremy L O'Brien, John G Rarity, Anthony Laing, et al. Experimental quantum hamiltonian learning. *Nature Physics*, 13(6):551–555, 2017.
- [228] Christopher Granade, Christopher Ferrie, Steven Casagrande, Ian Hincks, Michal Kononenko, Thomas Alexander, and Yuval Sanders. QInfer: Library for statistical inference in quantum information, 2016.
- [229] Christopher Granade, Christopher Ferrie, Ian Hincks, Steven Casagrande, Thomas Alexander, Jonathan Gross, Michal Kononenko, and Yuval Sanders. Qinfer: Statistical inference software for quantum applications. *Quantum*, 1:5, 2017.
- [230] Immanuel Bloch, Jean Dalibard, and Sylvain Nascimbene. Quantum simulations with ultracold quantum gases. *Nature Physics*, 8(4):267–276, 2012.
- [231] Albert Hertel, Laurits Andersen, Natalie Pearson, Malcolm Connolly, Valentina Zannier, Lucia Sorba, Liu Yu, Peter Krogstrup, Geoffrey Gardner, Michael Manfra, et al. Superconducting gatemon qubits based on selective-area-grown semiconductor materials. *APS*, 2019:C29–012, 2019.
- [232] P Krogstrup, NLB Ziino, W Chang, SM Albrecht, MH Madsen, Erik Johnson, Jesper Nygård, Charles M Marcus, and TS Jespersen. Epitaxy of semiconductor–superconductor nanowires. *Nature materials*, 14(4):400–406, 2015.

Acknowledgements

The acknowledgements is of course the most difficult part of a thesis to write. As a result I must apologise in advance for them being so long but still far too short.

First, I must thank Matthias Troyer for offering me the incredible opportunity to be part of your group and having confidence in me throughout the PhD. I would like to thank you especially for your patience and generosity, no matter how many times the same question is asked, for always taking the time to explain things with the same level of care and enthusiasm. I have been uniquely lucky to have been equally supported by Charles Marcus who welcomed me into his extraordinary lab and whose great energy and passion for both science and people has been a constant source of motivation and conversation.

Although I have been at best a long distance group member in Zurich I would like to thank the people I had the pleasure of overlapping with and who made it such a welcoming place to be. I am especially grateful to Jana, Dominic, Mario, Bettina, Damian, Thomas, and Donjan, who were part of the final cohort of PhD students and whom I could always call on to help me with anything from a place to stay to literally anything computer related. I have no doubt that you will all be frighteningly successful in your endeavours.

Much of my time was spent in Copenhagen at QDev and I cannot adequately express how grateful I am to be part of this particular family. It is most accurate to say that this group is about as exceptional as they come, in its people, its facilities, and its fun quirks. In particular I would like to thank the original 'Transmon team', who welcomed me from the very beginning and who made me an experimentalist (quotes deliberately left out). Lucas, for your kindness and patience not least in teaching me to measure (often by feigning your own ignorance), and for your humour which can transform a bad day into an even worse one. Rob, for understanding my need for tea, for commiserating with me about the EU and for carrying so much more than your fair share of dragging me out of my house be it for BLM or Baywatch. Karl, for your persistence in tracking down those ground loops, for epitomising the deadpan and for teaching me to wait out a silence. Thorvald, for listening to me complain about code so much, and for being the

best person to sound out an idea to. Anders, the fam I will go to if ever I need to know how I look in profile, thank you for being there from noticing when I needed a hug to having the perfect put down when I needed to get over myself. Lillian, for your patience and perseverance in the face of any challenge, not least my indecipherable diagram/todo lists. Albert, the better half of my team and quite possibly a genuine angel, meme master, fridge whisperer, outstanding human. Outside of the Transmon team there are many people who make Qdev the amazing place it is and the list really is too long to fit but I would especially like to thank Alisa, Judith, Maren, Anne Helene, Katrine and Esteban who are not only my colleagues but also my friends and with whom I danced, ran, dined, exercised and partied and who brightened my days wherever they entered them (except the exercising, sometimes that was grim let's be honest).

The lab would not run without the capable work of of the admin and technical staff both in Copenhagen and Zurich; Benjamin, Jess, Katrin and many others, who are coincidentally all people with really great smiles. Thank you.

I was lucky enough to work with many incredibly talented physicists who were always extremely gracious and patient and whose valuable time I am very grateful for, among them Gorjan Alagic, Synge Todo and Bela Bauer. In particular I am grateful to have had the opportunity to work with David Poulin who was not only a dizzyingly impressive scientist but also simply a very nice person, who I feel extremely lucky to have known even a little.

I have loved living in Copenhagen which is in no small part due to the excellent company of my flatmates Dillen and Stefan, without whom I cannot imagine the last 4 years. I am so glad to have met you and find myself struggling to find words for my gratitude, which I hope is in itself a remarkable enough event for you to appreciate its magnitude. From tears to beers to avocados.

I would never have gotten to the start line let alone anywhere near the end if not for my long suffering and completely outstanding friends. Lucy, for keeping me going both metaphorically and literally when dragging me round half marathons. Jo, for always getting me up that hill and for having a card for every occasion. Katie, for watching Buffy with me and telepathetic communication via gif. Charlotte, for the fox socks and the fox Christmas. Mel, for being the sweet spot between serious and silly. Beanie, for the heart shaped sunglasses. Lexy, for the cat crockery and the podcast recommendations. Anna, for the happiest and most tiring day ever. Emma, for the notecards. Ju, for inventing the voice message monologue therapy session. All these people and many more, from the small treasures by post to the many joyful weekend visits, you have kept me myself. Whatever is done by only me is your doing.

To my mum, thank you for the infinite love. I'm pretty sure I'm not an actual genius but it is very nice to know that at least one person believes me one. I love you. To my

dad, thank you for the determination, humour, and for your love which I have never doubted. To Johannes, I probably could have done this without you but it would not have been worth it. That's not true, I couldn't have done this without you.

In as much as one can dedicate a thesis to anyone I dedicate this thesis to Emma. I doubt I will have another dedication opportunity so this is probably the best you will get but you are all my dedications forever. The wonder that keeps the stars apart.