Scalable Quantum Computing with Two-Dimensional Arrays of Trapped Ions Enabled by Fast Gates

Zain Mehdi

A thesis submitted for the degree of
Bachelor of Science (Advanced) with Honours in Physics at
The Australian National University

October 2019

Declaration

The research presented in this thesis has been completed on the lands of the Ngunnawal and Ngambri people, the traditional owners of the land of the ANU. I acknowledge that this land was stolen and sovereignty was never ceded, and pay my respects to the elders past, present, and emerging.

This thesis is an account of research undertaken between February 2019 and October 2019 at the Department of Quantum Science, Research School of Physics, The Australian National University, Canberra, Australia.

Except where acknowledged in the customary manner, the material presented in this thesis is, to the best of my knowledge, original and has not been submitted in whole or part for a degree at any other university.

Zain Mehdi

October 2019

Acknowledgements

While nominally my name is the only one on this document, this thesis would not exist without the support of my friends, family, and colleagues (at least not in any comprehensible or useful form).

To start, I would like to thank my supervisor Joe Hope, who in addition to being an excellent physicist, cutting-edge educator, pop-culture connoisseur, and world-class pizza chef, has found time to debug many of my confusions and give deep insight into the world of trapped ions and beyond.

I would also like to thank Alex Ratcliffe, who has taken the time to answer the hundreds (potentially thousands) of questions I have thrown his way. Alex has an immense working knowledge of the subtleties of controlling trapped ions, which has proven invaluable over the course of this year. Importantly, Alex has been very consistent in his excitement every time I have shown him a pretty new plot (regardless of whether or not the data was correct), which has been very validating and thoroughly appreciated.

In addition to Joe and Alex, I am also very thankful to the other members of our group who have supported me throughout this year in many ways. In no particular order, thank you Simon, Stuart, Angela, Jay, Ruvi, Tim, and fellow honours students, Matt and Tiggy. It has been a pleasure to work, play table tennis, and eat burgers alongside you all.

I think if I were to try and list the friends I am thankful for, and all of the times I have felt supported and loved by them throughout this year, I would likely break some sort of record for the longest thesis ever written. Suffice it to say that I have an amazing group of friends without whom I could not imagine the past year. In particular I would like to mention my friends Nupur, Tanika, and Aisha, who are not only the best friends I could ask for, but also incredibly fierce activists. It has been humbling and heart-warming to watch you all make this world a better place.

As always, I am incredibly thankful for my family. Amma, Abba, and Gia; thank you for standing (sometimes sitting) with me every step of the way. Finally, thank you Lauren, for being endlessly supportive, and for inspiring me every day.

Abstract

Realising large-scale quantum computation in the near future will require increasing the number of low-error two-qubit gates that can be implemented on a quantum computer before decoherence. One of the biggest challenges facing current trapped ion quantum computers is implementing high-speed two-qubit operations, whilst increasing the number of qubits. One of the most promising proposals for overcoming current limitations is the use of ultra-fast pulses to implement fast two-qubit gates between nearest-neighbour pairs of ions. In this thesis, I investigate these 'fast gates' in two-dimensional arrays of microtraps, each containing a single ion. I argue that two-dimensional architectures allow for a significant reduction in the number of two-qubit gates required for a particular computation, as compared to one-dimensional ion chains. I demonstrate this reduction for a quantum simulation of a 40-mode Fermi-Hubbard Hamiltonian. I develop an efficient scheme that allows fast gates to be numerically optimised for two-dimensional geometries. I find that this optimisation scheme is capable of designing gates that are faster, higher fidelity, and require lower laser repetition rates. Using this scheme, I find that high-speed two-qubit gates can be optimised for two-dimensional architectures, with fidelities well above thresholds required for fault tolerant error correction, around 99.99%. Furthermore, I find that fast gates in these architectures are robust to the presence of large numbers of surrounding ions. Following previous studies [1, 2] which have identified pulse imperfections as a dominant source of error in fast gates, I perform a worst-case error analysis. I find the fast gates presented in this thesis to require very small errors in single-qubit rotations, and I present recommendations for achieving those requirements. I also investigate other experimental considerations, and make recommendations for overcoming other technical challenges in realising fast gates.

Contents

Chapter 1

Introduction

Quantum computing is one of the most rapidly progressing fields of physics today, and is receiving significant attention across academic, corporate, and government bodies. At its core, quantum computing is an attempt to exploit the true quantum-mechanical nature of the universe to process information in ways that classical computers (i.e. Turing machines) simply cannot. The hype behind the development of the field is understandable; quantum computers can perform certain tasks with significantly fewer resources than their classical analogues.

One of the most well-known examples of this is Shor's algorithm [3] for factoring large prime numbers with only polynomial requirements on computational resources. Equivalent classical algorithms require exponentially more resources, which is the basis for classical cryptography, and thus quantum computing has powerful implications for code-breaking and information security [4]. However, to break even a 2048-bit RSA encryption key with Shor's algorithm, one would require a quantum computer with thousands of quantum bits (*qubits*) [5]. Furthermore, a task of this size would almost certainly require use of error-correction routines [6], which would require upwards of tens of thousands of qubits. This is well out of reach from current generations of quantum computers that have, at best, tens of qubits.

However there are still incredibly appealing applications for near-future quantum computers with $50 - 100$ qubits, that may not require error-correction [7]. Perhaps the most appealing, at least for physicists, is the prospect of being able to efficiently simulate the dynamics of a many-body quantum system, often referred to as *quantum simulation* [8].

1.1 Simulating quantum systems

Simulating physical systems on a computer is an incredibly powerful theoretical tool, used across fields of physics, chemistry, engineering, and biology. However, due to

the exponential scaling of the complexity of quantum systems with their size, this tool can be extended only to the most simple of quantum systems. Richard Feynman was the first to recognise that this exponential barrier may be overcome by flipping the problem on its head: if one would like to simulate a quantum system, then the platform on which it is simulated must indeed be quantum-mechanical [8].

Today, Feynman's vision is being realised in two very different types of quantum simulators: *analog* and *digital*. The analog approach aims to simulate a many-body quantum system through studying a simpler 'analog' system that can be described by the same mathematical model. The ideal choice of the analog system is one that can be well controlled in an experimental setting, with some freedom in parameter choices that can be tuned to the particular problem. This has been realised with cold atoms in optical potentials, as well as on a variety of other platforms including trapped ions [9]. However, this approach is limited by the ability to find a suitable analog system, and practical barriers in its isolation and control in an experimental setting.

In the digital approach, the many-body quantum system is fully abstracted onto a set of interacting qubits in a programmable quantum computer [10]. The dynamics of the system are encoded in a circuit of logical operations that, when applied to a carefully prepared initial state of the qubits, returns a final state from which information about the evolved quantum system can be extracted. The main advantage of the digital approach is that any quantum system can be abstracted onto a sufficiently large set of qubits, and any Hamiltonian evolution can be simulated by programming the logical operations. Digital quantum simulations and their algorithms have been proposed and experimentally demonstrated on a wide range of quantum information platforms [4, 11–14], however a large-scale quantum simulation (i.e. one that is inaccessible to a classical computer) has yet to be demonstrated. In this thesis, I will investigate the feasibility of implementing a large-scale simulation of a fermionic Hamiltonian, on near-future trapped ion quantum computers.

1.2 The road to quantum utility

The milestone of implementing some computation on a quantum computer that is inaccessible to even the most powerful supercomputer is sometimes referred to as *quantum supremacy* [7]. A recent paper by the AI team at Google [15] claims to have achieved quantum supremacy by implementing a task in 200 seconds that would take up to 10000 years on a classical supercomputer. The task itself - sampling a randomly generated quantum circuit a million times - has no known application,

save for demonstrating supremacy. The next milestone, which is perhaps more interesting, is *quantum utility*. This is the goal of implementing a task on a quantum computer that is not only intractable to classical computers, but is also *useful* for some practical application. This is also sometimes phrased as *quantum advantage*.

In this thesis, I will consider digitally simulating a quantum system of 40 interacting fermions to be a task that is both useful and inaccessible to classical computing resources, i.e. a candidate for achieving quantum utility. To motivate this, it is sufficient to consider memory requirements. To store the state of 40 two-level atoms requires $2^{40} \approx 10^{12}$ classical bits, corresponding to about $\sim$1 terabyte of memory. To simulate its time evolution for a single time-step requires computing a matrix with up to $2^{40} \times 2^{40} \approx 10^{24}$ elements, in full generality. Even if a particular fermionic system of interest had a highly-sparce matrix and the non-zero elements were possible to calculate with a powerful supercomputer, even just storing the state of a slightly larger system (with say, $\sim$50 fermions) would be completely out of reach.

Scalability is one of the biggest barriers to quantum utility, and is a main limitation of all currently proposed quantum information platforms. In this context, scalability both refers to number of qubits as well as number of logical operations (formally called *gates*) that can be performed before the quantum system decoheres through interactions with the environment, or spontaneous internal processes. The characteristic timescale for this decoherence is called the coherence time. In fact, there currently exist quantum computers with upwards of 40 qubits [15–18][1], suggesting that the main barrier to quantum utility is the number of low-error operations that can be performed within the coherence time.

1.3 State of the art quantum computers

In realisations of quantum computers, information is stored in *qubits*, which are states of a two-level quantum system. In analogy to classical information processing, the computational basis states are canonically referred to as the $|0\rangle$ and the $|1\rangle$. Analogous to *logic gates* that are used in classical computers to manipulate information stored in bits, *quantum logic gates* are unitary operators applied to one or more qubits, to change their state for the purposes of computation. In order to implement any arbitrary quantum circuit, gates that entangle pairs of qubits together (entangling gates) are required, in conjunction with single qubit gates. In

[1] IBM and Intel have both announced quantum computing chips with upwards of 50 qubits, but currently do not have any publications where they are used. Citations [16, 17] are of the press releases from each company declaring the production of their respective processors.

Platform	Coherence time (s)	Two-qubit gate fidelity	Two-qubit gate time (s)
Superconducting circuits	10^{-4} [22]	99.4%[23]	3×10^{-8} [24]
Trapped ions	10^{0} [21]	99.9% [25]	5.4×10^{-7} [26]

Table 1.1: Comparison of superconducting (transmon) qubits and trapped ion qubits in terms of achievable gate times and fidelities, with respect to coherence time. All quantities here are the best that have been achieved - the achieved gate times listed here do not necessarily correspond to the achieved gate fidelities.

practice, these entangling gates are constructed from physical processes which can only approximate the ideal unitary operation. The *fidelity* is defined as the overlap between the realised and ideal qubit state, and is a useful measure of the 'accuracy' of a gate operation. I will use the fidelity measure extensively throughout this thesis. For computation, a readout of the final qubits state is required, performed by measurement on the quantum state [4].

The fundamental aspects of quantum computing (single and two-qubit gates, state initialisation, and read-out) have been experimentally demonstrated with a range of physical platforms, such as trapped ions, superconducting circuits, NMR, linear optics, and solid-state systems [4]. To date, the largest scale demonstrations of quantum computing have been with superconducting circuits and trapped ion qubits [19, 20]. There are disadvantages and advantages to both these technologies: for example, trapped ions have excellent coherence times but absolute gate times is generally slow [21], as compared to superconducting qubits which generally have much faster gate operations but are more prone to error [20]. For both these platforms, the main thing limiting large-scale computation is gate quality, whether it be speed or fidelity. For comparison, I have tabulated current state-of-the-art two-qubit gate fidelities and speeds in Table 1.1.

Of the other technologies, solid state systems have great potential in adapting the industrial scale manufacturing techniques of the semiconductor industry. Qubits realised in defects in diamonds and single spins in silicon are emerging as a promising candidate for robust quantum information processing. Notably, nitrogren-vacancy (NV) centers in diamond have the longest demonstrated coherence times of any physical qubit [27], and are the only solid-state technology to have demonstrated small-scale computation at room temperature [28]. However current demonstrations have only been on a small-scale with a few qubits, and the pathway to scaling to tens or hundreds of qubits remains unclear.

1.4 Trapped ion quantum computers

In this thesis, I will focus on the trapped ion platform for quantum computing. In a trapped ion quantum computer, the electronic states of ions trapped in electromagnetic potentials serve as the qubits. Collective motion of the ions are used as a 'bus' for the exchange and transfer of quantum information between qubits, controlled by interactions with laser light. The use of trapped ions as a platform for quantum information processing has been extensively studied and well-demonstrated experimentally [19, 21, 25, 29, 30]. Entangled states of ions have been demonstrated that remain coherent for tens of seconds, single qubits can be controlled with errors of the order 10^{-6}, and final states can be read-out with near unit efficiency [21].

Two-qubit entangling gates have also been well demonstrated on trapped ions, which generally entangle the qubits using spectroscopic methods to excite transitions of a single motional-sideband [29, 31]. These entangling gates have been experimentally demonstrated with fidelities over 99.9% [32]. However, resolving these motional sidebands requires gate times that are adiabatic with respect to the characteristic oscillation frequency of the ions. This imposes a fundamental speed limit on these gates, on the order of tens of microseconds. Furthermore, as more ions are added to the system, these transitions require increasingly longer timescales to resolve, resulting in longer gate operation times. This limits the number of two-qubit operations that can be done before decoherence, thus precluding this class of entangling gates from being useful in implementing large-scale computation.

1.4.1 Fast gates

An alternate scheme for realising two-qubit gates on trapped ion processors involve the excitation of *multiple* motional modes of the ions, which are disentangled from the qubit states only at the end of the gate operation. There is no adiabatic timescale required for this scheme, and thus these gates have no fundamental limitation on speed, for which they are known as *fast gates*. Fast gates entangle ions by orchestrating trajectories through phase space that result in state-dependent phase accumulation [33]. Fast gates have been experimentally implemented in recent years [26, 34], however these demonstrations have been unable to achieve high-fidelities for gates faster than the oscillation frequency of the ions. In this thesis I will focus on a particular set of proposals for realising fast gates, that use groups of resonant counter-propagating pulse pairs to entangle qubits together [1, 2, 33–36]. These gate schemes have been well-studied theoretically, but are yet to be experimentally realised. Unless otherwise stated, I will use the term 'fast gates' to exclusively refer

to this class of gate schemes for the remainder of this thesis.

A significant barrier to realising high-fidelity fast gates is the optimisation of the number of pulse pairs in each pulse group and their timings, which is a difficult computational challenge. This has been a key focus within these proposals, which make the numerical optimisation tractable by imposing constraints on the form of the gate [33, 35, 36]. Within these gate schemes higher fidelities can always be achieved, but only at the cost of requiring higher laser repetition rates to implement, which is a technological barrier. However, imposing constraints on the optimisation for computational tractability comes with the possibility of throwing away high-fidelity gate solutions with low repetition rate requirements that are more difficult to find. Recent unpublished work by Evan Gale [37] has described a method for optimising the number of pulse pairs in each pulse group that can find high-fidelity solutions for two qubits with modest laser repetition rate requirements. I will build on this work, and develop a variant that is computationally cheaper for optimising fast gates in the presence of more than two ions.

A recent paper by Ratcliffe *et al.* [38] has demonstrated that these fast gate schemes are capable of entangling nearest-neighbour ions in separate microtraps, even when large numbers of surrounding ions are present. This opens up the possibility of realising a scalable quantum computer with two-dimensional microtrap arrays. However, there is no existing literature studying fast gates in two-dimensional ion crystals. In this thesis I will address some major questions that are currently unanswered in this area: Can fast gates be optimised for high-fidelity implementation in two-dimensional architectures of ions? Does the notion of nearest-neighbour optimality extend to ions that are diagonally separated in a two-dimensional lattice? Does the performance of these gates decay indefinitely in increasingly larger arrays? What are practical barriers to the experimental realisation of these fast gates?

1.5 Thesis outline

This thesis is structured as follows. Chapters 2 and 3 cover background, and Chapters 4-6 contain the bulk of my original calculations. Chapters 7 and 8 contain discussions, conclusions, and suggestions for future work. In Chapter 2, I introduce the fundamentals of quantum information processing. The formalism of qubits and quantum logic gates are described in detail, with a particular focus on two-qubit entangling operations. The general fidelity measure is introduced, as well as the state-averaged fidelity, which is used extensively in this thesis as a metric for comparing gates.

In Chapter 3, I introduce trapped ions as a candidate platform for quantum information processing. I derive the Hamiltonian for the interaction of ions with light under the rotating-wave approximation. The motional-sideband transitions are explicitly shown in the context of adiabatic entangling gates. Conditions for implementing an ideal fast gate are presented, and existing optimisation schemes for finding solutions to these conditions are described. Finally, an algorithm for implementing the quantum simulation of an interacting fermionic system with nearest-neighbour fast gates is described.

In Chapter 4, I study the implementation of the simulation algorithm for the Fermi-Hubbard model with 40 fermionic modes. I calculate the number of entangling gates required to realise the algorithm with nearest-neighbour fast gates in both one-dimensional and two-dimensional architectures. Finally, I place thresholds on the two-qubit gate speeds and infidelitues required to efficiently implement the simulation within coherence times.

In Chapter 5, I present an adapted form of the fast gate optimisation scheme devised by Evan Gale that has an imposed anti-symmetry. This anti-symmetric scheme is shown to be better suited to designing fast gates in complex geometries, allowing for higher fidelity gate solutions to be found for a linear trap with increasing number of ions as compared to current best schemes.

In Chapter 6, I study fast gates in two-dimensional microtrap architectures. I show that the mode structure of a simple 2×2 array can be described by a single non-dimensional parameter. I present the results of gate optimisations for adjacent gates in this four-ion array, across a variety of gate operation times and values of the non-dimensional parameter. Optimisation results are also presented for diagonal gates, and compared to equivalent operations constructed of adjacent gates. Finally, gate performance is analysed as the number of ions in the two-dimensional array is increased.

In Chapter 7, I discuss considerations for experimentally realising fast gates. I address imperfections in pulse area as a dominant source of error, and perform a worst-case analysis to assess its impact on gate performance. I also consider the application of fast gates to a wide variety of candidate ions, and make a series of recommendations for experimentally implementing robust, high-fidelity fast gates.

Finally in Chapter 8, I discuss the possibility of realistic implementation of the simulation algorithm introduced in Chapter 4, with fast gates optimised for one-dimensional and two-dimensional architectures. I summarise the results of this thesis, and suggest avenues for future work.

Chapter 2

Background I: Quantum Computation

In this chapter, I will briefly introduce some of the fundamentals of quantum information theory, in the context of quantum computation. This will only cover the essentials for understanding and appreciating the results of this thesis, and will follow the canonical text by Nielson and Chuang [4], where the reader is directed for comprehensive treatment on the topics discussed. Furthermore, I will assume basic knowledge of quantum mechanics including density matrix formalism for describing mixed states, and I direct readers unfamiliar with these concepts to Chapter 2 of the text referenced above [4].

2.1 Quantum bits (qubits)

The fundamental unit of information in a quantum information processor is the quantum bit, or *qubit*. Here I will first discuss the formalism for describing a single qubit, and then generalise this to multi-qubit systems.

2.1.1 Single qubits and the Bloch sphere representation

A qubit is simply a normalised quantum state of a two-level quantum system. Mathematically, a qubit $|\psi\rangle$ can be understood as a unit vector in a two-dimensional Hilbert space,

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle\ , \tag{2.1}$$

where $|0\rangle\,,\ |1\rangle$ are the computational basis states, and $\alpha,\ \beta$ are their respective complex amplitudes. Normalisation of the qubit can be expressed as

$$\langle\psi|\psi\rangle = 1\,, \tag{2.2}$$

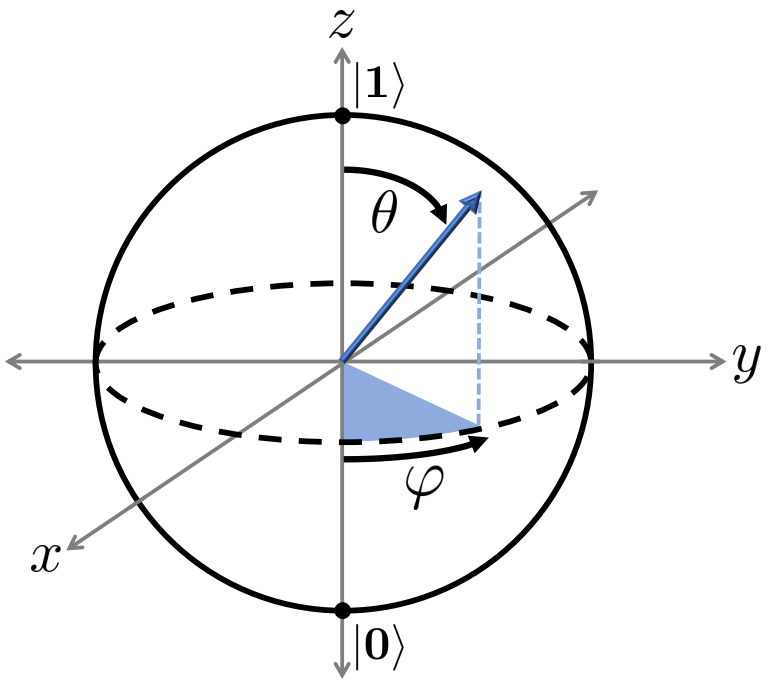

Figure 2.1: In the Bloch sphere representation, a qubit state is given by a unit vector to a point on the 2−sphere. The north and south pole of the sphere represent the computational $|1\rangle$ and $|0\rangle$ states, respectively. The polar (θ) and azimuthal (φ) angles are annotated.

which imposes the constraint on the coefficients

$$|\alpha|^2 + |\beta|^2 = 1\,. \tag{2.3}$$

Equivalently, the qubit can be represented in spherical co-ordinates as

$$|\psi\rangle = \cos\left(\frac{\theta}{2}\right)|0\rangle + e^{i\varphi}\sin\left(\frac{\theta}{2}\right)|1\rangle\ , \tag{2.4}$$

up to a non-physical global phase. Here (θ, φ) are the regular polar and azimuthal angles respectively, and define the *Bloch sphere representation*. In this representation the state of the qubit can be geometrically represented by a point on the 2-sphere, called the *Bloch sphere*. This is pictured in Figure 2.1, and is a very powerful tool for understanding and controlling individual qubits.

2.1.2 Multiple qubits

Describing multiple qubits is similar to describing any composite quantum system. Recall that the Hilbert space of a composite quantum system is given by the tensor

product of the Hilbert spaces of the individual systems, i.e.

$$\mathcal{H} = \bigotimes_j \mathcal{H}_j \,, \tag{2.5}$$

where $\mathcal{H}_j$ is the Hilbert space of the j-th component system. Applying this formalism to two-dimensional Hilbert spaces allows for the description of multi-qubit states. For example, the general form of a two-qubit state is

$$|\psi\rangle = \alpha_{00}\,|00\rangle + \alpha_{01}\,|01\rangle + \alpha_{10}\,|10\rangle + \alpha_{11}\,|11\rangle \ , \tag{2.6}$$

where normalisation requires

$$|\alpha_{00}|^2 + |\alpha_{01}|^2 + |\alpha_{10}|^2 + |\alpha_{11}|^2 = 1 \,. \tag{2.7}$$

I have used a short-hand notation here to describe the tensor product states, for example $|0\rangle_1 \otimes |1\rangle_2 \equiv |01\rangle$ describes the first qubit in state $|0\rangle$ and the second in state $|1\rangle$. Note that the two-qubit state is described by four complex amplitudes. A general n-qubit state will require 2^n complex amplitudes to characterise its state. This is an exponential growth in complexity as more qubits are added, which is in stark contrast to classical computers which require just n numbers to describe each bit.

Entanglement

A consequence of equation (2.6) is the existence of states that cannot be written as a tensor product of single-qubit states, i.e.

$$|\psi\rangle \neq |\phi\rangle \otimes |\gamma\rangle \tag{2.8}$$

for any two single-qubit states $|\phi\rangle$ and $|\gamma\rangle$. These are called *entangled* quantum states. A quantum system is called *separable* if it is not entangled. Entanglement is a uniquely quantum mechanical phenomena, and as I will describe in the next section and later chapters, is fundamental for quantum information processing.

Perhaps the most well-known entangled states are the two-qubit *Bell states*, which

are maximally entangling.[1]

$$\left|\Phi^+\right\rangle = \frac{1}{\sqrt{2}}\Big(\,|00\rangle + |11\rangle\,\Big) \tag{2.9}$$

$$\left|\Phi^-\right\rangle = \frac{1}{\sqrt{2}}\Big(\,|00\rangle - |11\rangle\,\Big) \tag{2.10}$$

$$\left|\Psi^+\right\rangle = \frac{1}{\sqrt{2}}\Big(\,|01\rangle + |10\rangle\,\Big) \tag{2.11}$$

$$\left|\Psi^-\right\rangle = \frac{1}{\sqrt{2}}\Big(\,|01\rangle + |10\rangle\,\Big) \tag{2.12}$$

Bell states are used widely in quantum information applications, from superdense coding to quantum cryptography [4].

2.2 Quantum gates

Just as classical computers manipulate bits using logic gates, quantum computers use unitary operations as the analogous *quantum gates* (often referred to simply as 'gates'). The requirement that these operations be unitary is important to note, as it allows for every gate to have an inverse, which is not always the case in classical computing[2]. Just as in the previous section, I will first consider gates on single qubits, and then generalise the discussion to multi-qubit gates.

2.2.1 Single-qubit gates

In order to describe gate operations, it is useful to move to a vector notation description of the qubits. In the basis of the computational $|0\rangle$ and $|1\rangle$ states, a general two-qubit state can be written as the vector $(\alpha, \beta)^T$, corresponding to equation (2.1). Unitary operations on a single qubit can then be represented by 2×2 unitary matrices.

It is useful to introduce the Pauli matrices, not only because they represent common single qubit operations (called the Pauli operations or Pauli gates), but because they also form a basis for all 2×2 unitary matrices, along with the identity operation $\mathbf{1}$.

[1] There are many measures of entanglement, and thus the exact meaning of maximal entanglement is not quite clear and is not particularly relevant to the contents of this thesis. In this case, I am referring to maximally entangling by the measure of entropy of entanglement [39].

[2] It is worth noting that reversible forms of classical computing do exist [40], however are generally considered *unconventional*. Almost all classical computers involve non-reversible logic operations.

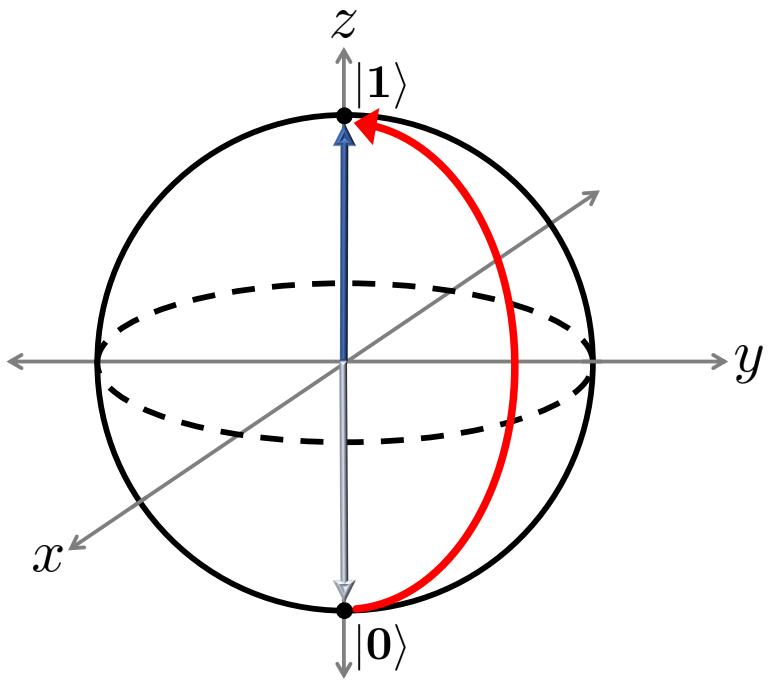

Figure 2.2: Representation of the Pauli-X gate on the Bloch sphere. Also known as the quantum NOT gate, the operation inverts the state of a qubit, corresponding to a rotation of π about the x axis. Here a qubit initially in the $|0\rangle$ state is rotated to the $|1\rangle$ state.

$$X \equiv \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \; Y \equiv \sigma_y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \; Z \equiv \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \tag{2.13}$$

The Pauli-X gate (σ_x) is equivalent to the classical NOT gate, as it inverts the qubit states $|0\rangle \leftrightarrow |1\rangle$, i.e.

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle \xrightarrow{X} |\psi'\rangle = \alpha\,|1\rangle + \beta\,|0\rangle \; . \tag{2.14}$$

Geometrically, the Pauli-X gate corresponds to a rotation of π around the x-axis in the Bloch sphere representation, as shown in Figure 2.2. Similarly, any single-qubit gate can be understood as a rotation on the Bloch sphere with respect to some axis. The matrices representing the Pauli X, Y, and Z gates are the generators of rotations about the respective x, y, and z axes of the Bloch sphere.

2.2.2 Multi-qubit gates

More generally, a gate on n-qubits can be described by a unitary $2^n \times 2^n$ matrix. It is only necessary to use this formalism when considering gates that are entangling, as any non-entangling (separable) multi-qubit operation can be decomposed into a

tensor product of single-qubit gates. Furthermore, any n-qubit operation can be constructed with only two-qubit and single-qubit gates. To be precise, any entangling two-qubit gate and the set of all single-qubit gates form a *universal set of gates* [41]. Therefore it is sufficient to generalise the discussion above only to two-qubit gates.

A two-qubit gate can be described by a unitary 4×4 matrix in the basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$. Perhaps the most commonly known two-qubit gate is the *controlled-NOT* (CNOT) operation where a Pauli-X gate is applied to a target qubit, conditional on the state of a control qubit. This has the matrix form

$$U_{\text{CNOT}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} . \tag{2.15}$$

This form of the matrix describes the NOT operation applied to the second qubit, but only in elements of the superposition where the first qubit is in the $|1\rangle$ state.

In this thesis I will consider a different two-qubit operation called the *controlled-phase* gate (CPH)[3], which has the matrix form

$$U_{\text{CPH}}(\phi) = \begin{pmatrix} e^{i\phi} & 0 & 0 & 0 \\ 0 & e^{-i\phi} & 0 & 0 \\ 0 & 0 & e^{-i\phi} & 0 \\ 0 & 0 & 0 & e^{i\phi} \end{pmatrix} , \tag{2.16}$$

which can also be written by exponentiation of the Pauli-X matrices on the two qubits (indexed as 1 and 2 here for simplicity):

$$U_{\text{CPH}}(\phi) = \exp\left(i\phi\sigma_x^1 \otimes \sigma_x^2\right) . \tag{2.17}$$

Specifically I will be considering the case of $\phi = \pi/4$, which implements a maximially entangling gate, $U_{\text{CPH}}\left(\frac{\pi}{4}\right)$, that is equivalent to the CNOT operation, up to single-qubit rotations. Both this controlled-phase gate and the CNOT gate can be the two-qubit element in a universal gate set.

A useful two-qubit gate to introduce here is the SWAP operation, which swaps the

[3]Sometimes called the *controlled-phase flip* (CPF) gate or the *Ising XX* gate.

state of one qubit with another. This has the matrix form

$$U_{\text{SWAP}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} . \tag{2.18}$$

This operation can be implemented with a sequence of three CNOT gates, or equivalently, three controlled-phase gates and single-qubit rotations. This is a property that will be exploited in later chapters.

2.2.3 The fidelity measure

In experimental realisations of quantum computers, quantum gates are realised by some sort of interaction that manipulate the physical system which represent the qubits. An ideal gate can only ever be approximated by the physical operation, and thus it is important to define some sort of measure of 'closeness' between the desired final state and the actual final state.

In this thesis I will use the *fidelity* measure, which captures the 'distance' between two quantum states by their overlap. For two mixed states represented by the density matrices ρ and σ, the fidelity is defined as [42]:

$$F(\rho, \sigma) = \left(\text{Tr}[\sqrt{\rho^{1/2} \sigma \rho^{1/2}}] \right)^2 . \tag{2.19}$$

Notably, this reduces to an inner product for a pure state $\rho = |\psi\rangle \langle\psi|$,

$$F(|\psi\rangle \langle\psi| , \sigma) = \langle\psi| \sigma |\psi\rangle . \tag{2.20}$$

The fidelity measure has the useful property as being bounded by $0 \leq F(\rho, \sigma) \leq 1$, and is only unit-valued when $\rho = \sigma$ [42]. Furthermore, the expression (2.19) is symmetric in its inputs, although that is certainly not clear purely from inspection.

The fidelity can be recast as a measure between real (U_{re}) and ideal (U_{id}) unitary operations on some initial state $|\psi_0\rangle$, by expressing the density operators as

$$\rho = U_{re} |\psi_0\rangle \langle\psi_0| U_{re}^\dagger \tag{2.21}$$

$$\sigma = U_{id} |\psi_0\rangle \langle\psi_0| U_{id}^\dagger . \tag{2.22}$$

The fidelity expression then becomes

$$F = |\langle\psi_0| U_{re}^{\dagger} U_{id} |\psi_0\rangle|^2 . \tag{2.23}$$

This form of the fidelity allows one to describe the probability of a physical implementation of a gate, described by the unitary U_{re}, being successful in implementing some desired unitary operation, U_{id}.

However, measure is dependent on the choice of initial state $|\psi_0\rangle$, and thus may give much higher fidelities for some initial states than others. In this thesis I will investigate implementations of entangling gates, and seek to compare their performance for arbitrary initial states. Therefore I will use a modified form of (2.23), called the *state-averaged infidelity* [43], where the fidelity is averaged over all initial states:

$$F_{av} = \frac{1}{N} \int_{\psi_0} |\langle\psi_0| U_{re}^{\dagger} U_{id} |\psi_0\rangle|^2 d|\psi_0\rangle \ . \tag{2.24}$$

In this expression, the integral is over the surface of the unit hypersphere, and the fidelity is renormalised by a factor $N \equiv \int_{\psi_0} d|\psi_0\rangle$. Unless otherwise stated, all fidelities reported in this thesis are state-averaged.

Chapter 3

Background II: Fast Gates in Trapped Ion Computing

In this chapter I will introduce trapped ions as a platform for quantum computation, and describe the representation of qubits and implementation of multi-qubit gates. In Section 3.1 I will describe the representation of qubits in trapped ions, and their trapping architectures. In 3.2, I describe the quantisation of the motional dynamics, and in 3.3 the Hamiltonian for ion-light interactions is introduced. These sections include the theory required for describing various methods for implementing two-qubit entangling gates on trapped ions. For a more comprehensive review of these topics, the reader is directed to Refs. [21, 44]. In Section 3.4, I describe the class of entangling gates that require weak coupling to the laser field. Section 3.5 describes the formalism for 'fast' gates which are the focus of this thesis, and discusses different schemes for their numerical optimisation. Finally, in Section 3.6, I introduce an algorithm for simulating Hamiltonians of interacting fermions on trapped ion quantum computers, which I will investigate in Chapter 4.

3.1 Trapped ions

In trapped ion experiments, ions are confined in potentials generated by an electromagnetic trap. As charged particles cannot be trapped in three-dimensions by a purely static potential [45], experimental realisations of ion traps typically involve the use of time-dependent electromagnetic fields which generate an effective harmonic potential [4]. The most common trapping architecture is the radio-frequency (RF) *Paul trap* [30], pictured in Figure 3.1. A Paul trap consists of four electrodes in a quadropole arrangement: two are grounded to a static potential of U_0, and the other two are driven by a rapidly oscillating voltage, generating a potential of the

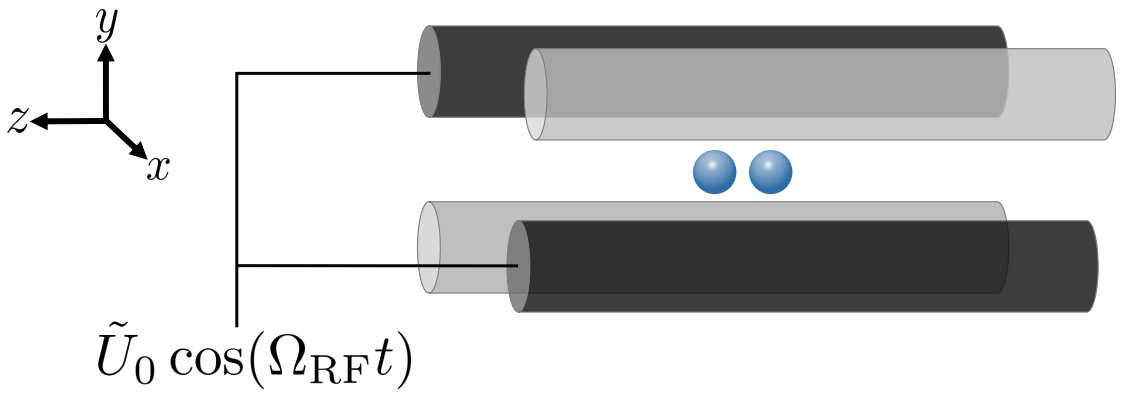

Figure 3.1: Radio-frequency (RF) Paul trap consisting of four cylindrical electrodes in a quadrapole arrangement. An oscillating voltage with frequency Ω_{RF} is applied to the dark electrodes, and the other two are grounded by a capacitor (unshown). This generates an effective trap for the ions (blue) to be confined in. Adapted from Ref. [44].

form [44]:

$$\Phi = \frac{1}{2}U_0\left(\alpha^2x^2 + \beta^2y^2 + \gamma^2z^2\right) + \tilde{U}_0\cos(\Omega_{RF}t)\left(\alpha'^2x^2 + \beta'^2y^2 + \gamma'^2z^2)\right) . \quad (3.1)$$

As shown in Ref. [30], averaging over the rapidly-oscillating terms leads to an effective harmonic potential of the form

$$V_T(x,y,z) = \frac{1}{2}M\left(\omega_x^2x^2 + \omega_y^2y^2 + \omega_z^2z^2\right) \quad (3.2)$$

where M is the mass of the trapped ion, and typically $\omega_x, \omega_y \gg \omega_z$. The oscillation of the potential also induces a rapid oscillation in the motion of the trapped ions called *micromotion* [44], which is often assumed to be negligible for low kinetic energies. I will work under this assumption for the scope of this thesis.

3.1.1 Qubit representation

In trapped ion quantum computers, qubits are represented in the internal (electronic) states of ions. Candidate ions have to be chosen such that they have an isolated pair of electronic states that are well-described by the two-level Hamiltonian

$$\hat{H}_{el} = \frac{1}{2}\hbar\omega_a\hat{\sigma}_z\,, \quad (3.3)$$

where ω_a is the frequency separation of the levels. These levels must also have long coherence times, to ensure the robustness of qubits and the quantum information they hold [30]. Broadly, there are two types of qubit representations in trapped ions,

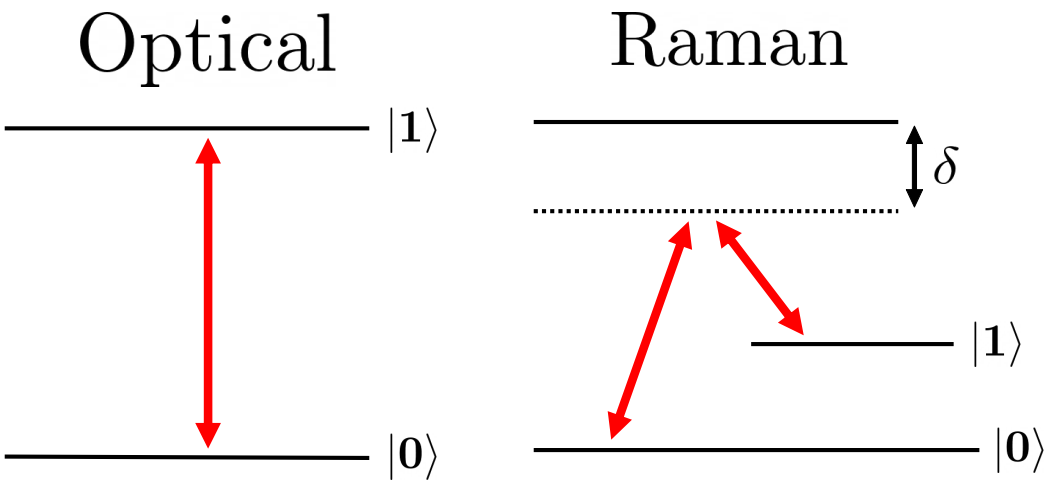

Figure 3.2: Optical qubits (left) and hyperfine qubits (right) represented in the internal states of trapped ions, where the computational $|0\rangle$ and $|1\rangle$ are annotated. Optical qubits use direct optical transitions between levels. In hyperfine qubits, a detuning δ to a third, auxiliary, level is used in Raman transitions between computational states.

as shown in Figure 3.2: *optical* qubits, and *hyperfine* qubits. Optical qubits consist of a ground and long-lifetime metastable excited state, and optical transitions can be directly excited [4]. In hyperfine qubits, the qubit is stored in a pair of hyperfine split low energy levels, and Raman transitions via a detuning on an auxiliary state is used to drive transitions [21]. $^{40}Ca^{+}$ is an example of the former [46, 47], and $^{9}Be^{+}$ and $^{171}Yb^{+}$ are examples of the latter [48, 49]. For the majority this thesis, I will consider only $^{40}Ca^{+}$, although I discuss this choice in Chapter 7. The level scheme for $^{40}Ca^{+}$ is shown in Figure 3.3. The ground-state $S_{1/2}$ and metastable $D_{5/2}$ will be used as the computational $|0\rangle$ and $|1\rangle$ states for the qubit. This choice is motivated by the long-lifetime of the $D_{5/2}$ state which has a coherence time of the order $\sim$1 second [21], as well as the availability of auxiliary short-lived states that can be used for state read-out [44].

3.1.2 Multi-ion architectures

Trapped ion chains are often realised by linearly arranging multiple ions in a single common Paul trap, as shown in Figure 3.4 for two ions. These architectures do not allow arbitrary numbers of ions without buckling of the ion chain. In order to prevent buckling, the radial and longitudinal trap frequencies, ω_r and ω_t, must

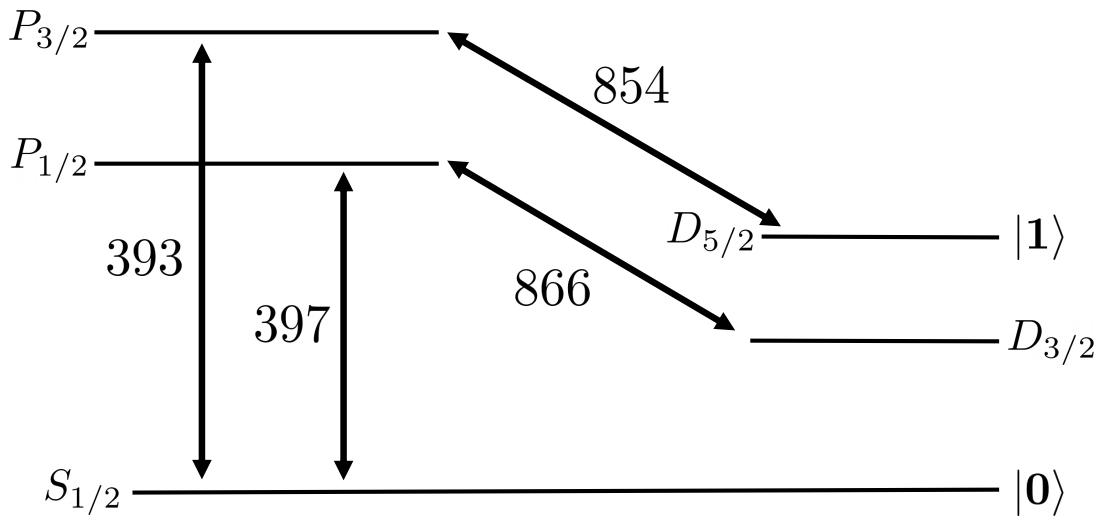

Figure 3.3: Level scheme for $^{40}Ca^{+}$ with the wavelengths of electric-dipole transitions annotated (in nano-meters), taken from Ref [50]. The ground state $S_{1/2}$ and the 1.2 second lifetime metastable $D_{5/2}$ state will be used as the computational basis states, as annotated.

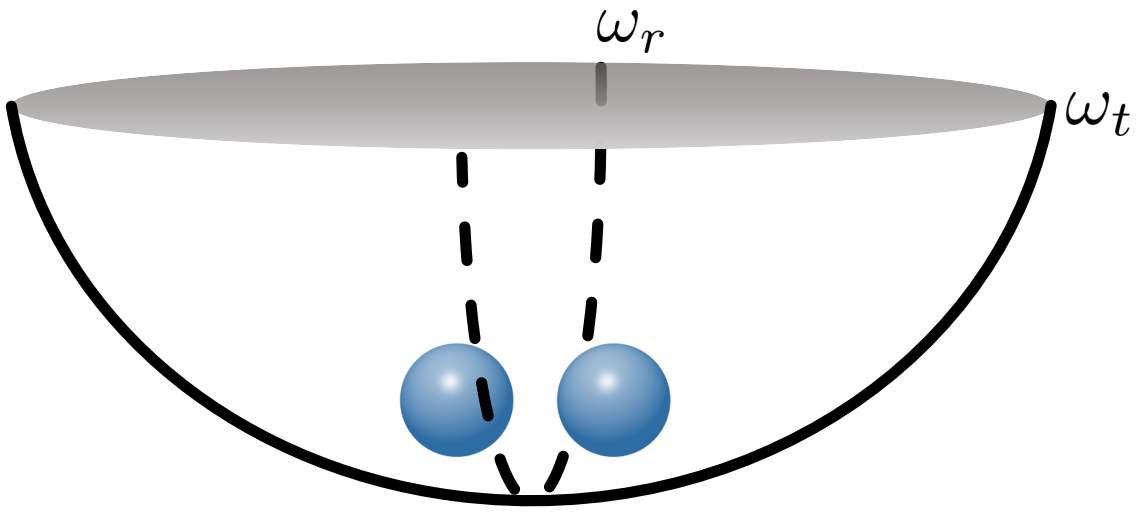

Figure 3.4: Two ions trapped in a harmonic potential generated by a linear Paul trap with longitudinal/axial trap frequency ω_t and radial trap frequency ω_r.

satisfy

$$\frac{\omega_r}{\omega_t} > 0.63L^{0.865} \tag{3.4}$$

where L is the number of ions in the trap [51, 52]. This means that as the number of ions increases, the trap frequency must be lowered, which results in longer timescales required for entangling gates operations [2], which in turn reduces the number of

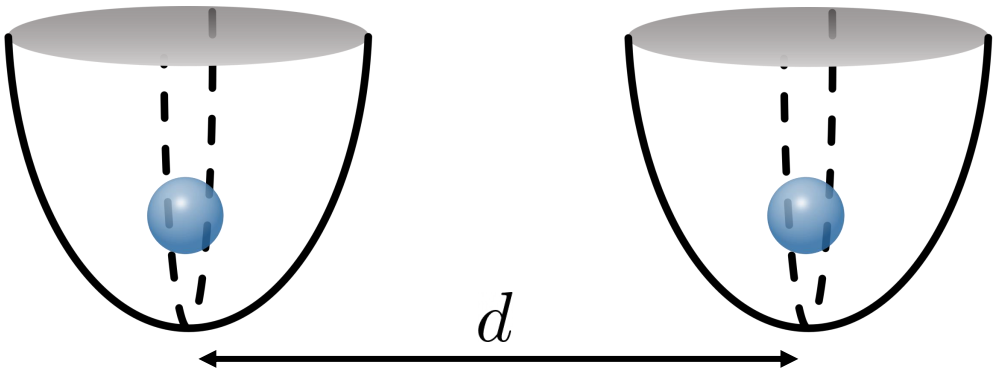

Figure 3.5: A microtrap chain of two ions, individually trapped. The center (or equivalently, minima) of the microtraps are separated by a distance, d, typically $\sim 100\mu$m [54].

gates that can be implemented within the coherence time. Even for fast entangling gates that are not limited to timescales adiabatic with respect to the trap period (which I will discuss later in this chapter), computation with more than ~ 40 ions in a single Paul trap becomes intractable with current laser technologies [2]. Further, ions in the middle of the chain are pushed closer together as more ions are added to the trap, which poses an experimental challenge in individually addressing them with laser light.

A promising alternative for a scalable trapped ion architecture is an array where ions are confined in individual Paul traps with radial symmetry (i.e. $\omega_r = \omega_t$) called *microtraps* [53]. This is shown in Figure 3.5. In recent years, Ratcliffe *et al.* have demonstrated that fast entangling gates can be implemented directly between nearest-neighbour microtraps, and scale favourably as the number of ions in a one-dimension chain of microtraps is increased [38]. Two-dimensional microtrap architectures have recently been realised by Kumph *et al.* [55]. Studying fast gates in two-dimensional microtrap architectures will be a key focus of this thesis.

3.2 Quantised motion of ions

In trapped ion quantum information processors, ions are coupled together by their collective motion due to the shared Coulomb interaction. These collective oscillations can be decomposed into their eigenstates, which are known as the *motional modes*. Excitation of the motional modes are fundamental to multi-qubit gates in trapped ion quantum computers, and so it is important to understand their structure. For simplicity I will begin by deriving the motional Hamiltonian for a single

ion in a harmonic trapping potential, and then extend the discussion to multi-ion systems.

3.2.1 Single ion in one-dimension

The Hamiltonian of the motion of a single ion in a one-dimensional external trapping potential $V_T(\hat{x})$ can be written as,

$$\hat{H}_{mot} = \frac{\hat{p}^2}{2m} + V_T(\hat{x})\,. \tag{3.5}$$

Here m is the mass of the ion. In the limit where the ion is cooled such that the amplitude of its oscillation in the trap is small, the trapping potential can be treated as quadratic around the equilibrium. In fact in trapped ion experiments, typically great effort is gone into ensuring the harmonicity of the trap as well, as in keeping the ions cool [56, 57], and so this approximation is typically quite robust [2, 58]. The motional Hamiltonian then becomes that of the quantum harmonic oscillator,

$$\hat{H}_{mot} = \frac{\hat{p}^2}{2m} + \frac{1}{2}m\omega_t^2\hat{x}^2\,, \tag{3.6}$$

where ω_t is the harmonic frequency of the trapping potential. It is well known that this Hamiltonian can be transformed into the basis of energy eigenstates by introducing the dimensionless ladder operators [59],

$$\hat{a} = \sqrt{\frac{m\omega_t}{2\hbar}}\left(\hat{x} + \frac{i}{m\omega_t}\hat{p}\right) \tag{3.7}$$

$$\hat{a}^\dagger = \sqrt{\frac{m\omega_t}{2\hbar}}\left(\hat{x} - \frac{i}{m\omega_t}\hat{p}\right) \tag{3.8}$$

which when acted on the n-th eigenstate, $|n\rangle$, give

$$\hat{a}\,|n\rangle = \sqrt{n}\,|n-1\rangle \tag{3.9}$$

$$\hat{a}^\dagger\,|n\rangle = \sqrt{n+1}\,|n+1\rangle\,. \tag{3.10}$$

These operators can be interpreted as annihilating and creating quanta of energy, respectively. Importantly, the motional Hamiltonian can be written in terms of these 'annihilation' and 'creation' operators as

$$\hat{H}_{mot} = \hbar\omega_t\left(\hat{a}^\dagger\hat{a} + \frac{1}{2}\right)\,. \tag{3.11}$$

In this form, it is clear that the Hamiltonian describes excitations of a single motional mode that has an oscillation frequency of ω_t. This is known as the *common-motional* mode.

3.2.2 Motional modes

For a crystal of multiple trapped ions, describing the ions' motion becomes more complicated. This is because the ions interact with one another through the Coulomb force, and thus see an effective potential

$$V_{eff}(\hat{\mathbf{x}}) = V_T(\hat{\mathbf{x}}) + V_C(\hat{\mathbf{x}})\,, \tag{3.12}$$

where V_C is the Coulomb potential, and $\hat{\mathbf{x}}$ is an operator-valued vector of the positions of each ion, typically defined as a displacement from the equilibrium of V_{eff}. This interaction is the mechanism behind the coupling between ions, which leads to collective motion that can be described in terms of the motional modes of the system. For two ions, there are two motional modes[1]: the common-motional mode in which both ions oscillate identically about equilibrium, and the breathing mode in which the ions have the same magnitudes of oscillation but with opposite direction. It is useful to note that the oscillation frequency of the common-motional mode is given by the trap frequency ω_t, and therefore the timescale for motional dynamics is characterised by the trap geometry. In Appendix B, I describe the calculation of the motional mode frequencies, ω_p, and coupling vectors $\mathbf{b_p}$.

Quantum mechanically, excitations of these motional modes are discrete and can be treated as bosonic quasi-particles (often referred to as phonons). Under the approximation of a harmonic potential, each of these modes has the structure of a quantum harmonic oscillator. Therefore, the motional Hamiltonian can be written in terms of operators $\hat{a}_p^\dagger$ and $\hat{a}_p$, which create and annihilate a single phonon in the p-th motional mode, respectively, taking an analogous form to (3.11)

$$\hat{H}_{mot} = \sum_p \hbar\omega_p \left(\hat{a}_p^\dagger \hat{a}_p + \frac{1}{2}\right)\,. \tag{3.13}$$

Here ω_p is the oscillation frequency of the p-th motional mode, in angular units. The bound on the sum is the number of motional modes, which is a function of the number of ions, and the dimension of the crystal.

[1]To be precise, there are two motional modes for each spatial dimension. However the modes of each spatial dimension entirely decouple from each other in a harmonic trap, so it is sufficient to only consider the one-dimensional case.

For sufficiently small displacements about equilibrium, the position of the m-th ion can be described by

$$\hat{x}^{(m)} = \sum_p b_p^{(m)} \hat{Q}_p \tag{3.14}$$

where b_p^m is the m-th component of the eigenvector of the p-th mode, and can be interpreted as the coupling of the m-th ion to that mode. $\hat{Q}_p$ is the position operator of the p-th motional mode [30], given by

$$\hat{Q}_p = \sqrt{\frac{\hbar}{2M\omega_p}} \left(\hat{a}_p^\dagger + \hat{a}_p\right) . \tag{3.15}$$

3.3 Ion-light interactions

To understand how logical operations are performed on trapped ion qubits, it is necessary to first describe how laser light interacts with two-level systems. This interaction can be described by the *Jaynes-Cumming Hamiltonian*,

$$\hat{H}_{int} = \hbar g \left(\hat{\sigma}_+ + \hat{\sigma}_-\right) \left(\hat{c}\, e^{i(k\hat{x}-\omega_L t)} + \hat{c}^\dagger e^{-i(k\hat{x}-\omega_L t)}\right) . \tag{3.16}$$

This Hamiltonian describes the interaction between an ion with quantised position $\hat{x}$, and a single quantised mode of a light field with creation/annihilation operators $\hat{c}^\dagger/\hat{c}$ [60]. In Appendix A, I derive this Hamiltonian from a more general multi-mode description of the light field.

The number of photons in a lasing mode is typically incredibly large, particularly for trapped ion experiments, and thus it is reasonable to approximate the lasing mode as a coherent state with complex amplitude $\beta e^{i\phi}$. Here β and ϕ correspond to the magnitude and phase of the laser, respectively. This is known as the *semi-classical approximation*, treating the light-field as an infinite reservoir that is unaffected by the atom. The semi-classical Hamiltonian can be calculated by taking the expectation value over the Hilbert space of the light field,

$$\hat{H}_{int} \approx \hbar g \left(\hat{\sigma}_+ + \hat{\sigma}_-\right) \left(\beta e^{i\phi} e^{i(k\hat{x}-\omega_L t)} + \beta e^{-i\phi} e^{-i(k\hat{x}-\omega_L t)}\right) \tag{3.17}$$

$$= \frac{\hbar\Omega}{2} \left(\hat{\sigma}_+ + \hat{\sigma}_-\right) \left(e^{i(k\hat{x}-\omega_L t+\phi)} + e^{-i(k\hat{x}-\omega_L t+\phi)}\right) , \tag{3.18}$$

where $\Omega = 2\beta g$ is the *Rabi frequency*.

To better understand the dynamics of this Hamiltonian in this regime, it is useful to transform into the interaction picture with respect to the electronic Hamiltonian:

$$\hat{H} \rightarrow \hat{H}' = e^{i\hat{H}_{el}t/\hbar}\,\hat{H}_{int}\,e^{-i\hat{H}_{el}t/\hbar}\,. \tag{3.19}$$

Using the identity $(\hat{\sigma}_z)^2 = \hat{1}$ and the commutation relation $[\hat{\sigma}_z, \hat{\sigma}_+] = \hat{\sigma}_+$, one can show this can be expressed as

$$\begin{aligned} \hat{H}' &= \frac{\hbar\Omega}{2}\Big(\hat{\sigma}_+ e^{-i(k\hat{x}-(\omega_L+\omega_a)t+\phi)} + \hat{\sigma}_+ e^{i(k\hat{x}-\Delta t+\phi)} && (3.20)\\ &+ \hat{\sigma}_- e^{i(k\hat{x}-(\omega_L+\omega_a)t+\phi)} + \hat{\sigma}_- e^{-i(k\hat{x}-\Delta t+\phi)}\Big)\,, && (3.21)\end{aligned}$$

where I have defined the detuning $\Delta \equiv \omega_L - \omega_a$. Under the typical *rotating-wave approximation*, fast-oscillating terms of the form $e^{\pm i(\omega_L+\omega_a)}t$ are neglected, i.e.

$$\hat{H}' \approx \hat{H}'_{RWA} = \frac{\hbar\Omega}{2}\Big(\hat{\sigma}_+ e^{i(k\hat{x}-\Delta t+\phi)} + \hat{\sigma}_- e^{-i(k\hat{x}-\Delta t+\phi)}\Big)\,. \tag{3.22}$$

This approximation is valid for QIP applications with trapped ions where typically $\Delta \ll \omega_a$, and has been shown to be very robust for the fast gate operations that I will focus on in this thesis [58].

3.3.1 The weak-coupling regime

A key requirement for many methods of implementing entangling two-qubit gates in trapped ion computers is that the laser light only weakly interacts with the ion [61–63]. This is known as the *weak-coupling regime* and is closely related to the satisfaction of the Lamb-Dicke criterion, which for a single ion is

$$\eta\sqrt{\langle(\hat{a}+\hat{a}^\dagger)^2\rangle} \ll 1\,. \tag{3.23}$$

Here, η is the Lamb-Dicke parameter

$$\eta = k\sqrt{\frac{\hbar}{2M\omega_t}}\,, \tag{3.24}$$

where k is the wavenumber of the laser light as introduced earlier. Satisfying this criterion physically corresponds to the amplitude of the ion's oscillation being much smaller than the wavelength of the incoming light [30], and implies the weaker condition $\eta \ll 1$. In the regime where this criterion is satisfied, equation (3.22) can be expanded to first order [21]:

$$\hat{H}' = \frac{\hbar\Omega}{2}\Big[\hat{\sigma}_+ e^{-i(\Delta t-\phi)} + \hat{\sigma}_- e^{i(\Delta t-\phi)} + i\eta(\hat{\sigma}_+ e^{-i(\Delta t-\phi)} - \hat{\sigma}_- e^{i(\Delta t-\phi)})(\hat{a}+\hat{a}^\dagger)\Big] + \mathcal{O}(\eta^2)\,. \tag{3.25}$$

Motional sideband spectrum

For certain values of the detuning Δ, resonant transitions in this Hamiltonian appear, which form the *sideband* spectrum. I will explicitly show how this spectrum arises for a single trapped ion. To study the motional dynamics, it is useful to transform (3.25) into the interaction picture with respect to the motional Hamiltonian (3.13),

$$\begin{aligned} \hat{H}'' &= e^{i\omega_t(\hat{a}^\dagger\hat{a}+\frac{1}{2})t}\hat{H}'e^{-i\omega_t(\hat{a}^\dagger\hat{a}+\frac{1}{2})t} \\ &= \frac{\hbar\Omega}{2}\Big[\hat{\sigma}_+e^{-i(\Delta t-\phi)}+\hat{\sigma}_-e^{i(\Delta t-\phi)}+i\eta(\hat{\sigma}_+e^{-i(\Delta t-\phi)}-\hat{\sigma}_-e^{i(\Delta t-\phi)})(\hat{a}e^{-i\omega_t t}+\hat{a}^\dagger e^{i\omega_t t})\Big]. \end{aligned} \tag{3.26}$$

Applying a second rotating wave approximation, rotating terms are discarded:

$$\hat{H}'' \approx \frac{\hbar\Omega}{2}e^{-i(\Delta t-\phi)}\hat{\sigma}_+\Big[1+i\eta(\hat{a}e^{-i\omega_t t}+\hat{a}^\dagger e^{i\omega_t t})\Big]+h.c.\,. \tag{3.27}$$

This approximation is only valid when $\Omega \ll \omega_t$, i.e. in the regime where only one transition is relevant at any given time [21]. For detunings at integer multiples of the trap frequency, resonant sideband transitions appear. The most interesting of these lie close to the trap frequency [64].

For $\Delta = 0$, the *carrier* transition is performed, corresponding to resonant transition of the internal electronic states:

$$\hat{H}_{carrier} = \frac{\hbar\Omega}{2}(\hat{\sigma}_+e^{i\phi}+\hat{\sigma}_-e^{-i\phi})\,. \tag{3.28}$$

For $\Delta = +\omega_t$, the blue sideband transition is performed,

$$\hat{H}_{blue} = i\eta\frac{\hbar\Omega}{2}(\hat{a}^\dagger\hat{\sigma}_+e^{i\phi}+\hat{a}\hat{\sigma}_-e^{-i\phi})\,. \tag{3.29}$$

This is commonly known as the heating transition as phonons are created as the electronic state is excited [21]. Analogous, the red (cooling) transition is given by $\Delta = -\omega_t$,

$$\hat{H}_{red} = i\eta\frac{\hbar\Omega}{2}(\hat{a}\hat{\sigma}_+e^{i\phi}+\hat{a}^\dagger\hat{\sigma}_-e^{-i\phi})\,. \tag{3.30}$$

The carrier, blue, and red sideband transitions have effective Rabi frequencies of Ω, $\Omega\eta\sqrt{n+1}$, and $\Omega\eta\sqrt{n}$, respectively, where n is the number of phonons excited [21, 30, 64].

Adiabatic timescales for population transfer

The carrier Hamiltonian in the Bloch sphere qubit representation describes rotations between the excited $|e\rangle$ and ground $|g\rangle$ states, with angle

$$\theta = \frac{1}{2}\int_0^\tau \Omega(t)dt \tag{3.31}$$

accumulated over some interaction time τ. As any unitary operation on a single qubit can be understood as a rotation on the Bloch sphere, any single-qubit gate can be implemented on a trapped ion with a single laser pulse of a well-chosen duration τ.

A complete population inversion between $|e\rangle$ and $|g\rangle$ corresponds to a $\theta = \pi/2$ rotation on the Bloch sphere. This is known as a π-*pulse*, as for a constant Rabi frequency

$$\Omega = \frac{\pi}{\tau} \tag{3.32}$$

characterises the time required for the transfer [30]. Recall that I stated $\Omega \ll \omega_t$ must be satisfied in order for a single sideband transition to be resolved. This means the timescale τ for a full population transfer using these transitions must be *adiabatic* with respect to the motional dynamics, i.e. $\tau \gg \frac{1}{\omega_t}$. This will become important shortly, when considering two-qubit gates.

3.4 Sideband-resolving entangling gates

Entangling gates between two-qubits are essential for universal computation on any quantum computer. As mentioned in the introduction, there are two distinct classes of entangling gates in trapped ion quantum computing: adiabatic sideband-resolving gates and fast gates. In this thesis I will study the latter, and so I will only briefly describe the sideband-resolving gates here. A curious reader can find more detailed explanations of these gate schemes in any of several review papers [21, 30, 64].

In the original proposal of trapped ion quantum computing, Cirac and Zoller outlined a method for implementing two-qubit entangling gates with sideband transitions [61]. These gates are implemented by mapping the electronic states of the ions to the common-motional (COM) vibrational mode, which is shared by all ions in the trap. The connectivity of the COM mode is then used to perform operations on a target qubit conditioned on the state of a control qubit. By mapping the quantum information back onto the electronic states at the end of the operation, a CNOT gate can be implemented. The method they proposed involves π and 2π pulses

on the red sideband transition, and requires that each motional mode is cooled to its ground state. Experimental realisations of this entangling gate scheme quickly followed this seminal paper [65, 66]. Notably, this method requires cooling of each motional mode to its ground state and is thus very sensitive to thermal effects.

Shortly after, Mølmer and Sørensen proposed an alternative gate scheme that also uses sideband transitions, but is more robust to experimental implementation [63, 67]. As opposed to the Cirac-Zoller gate which uses a laser to excite motional transitions on a single ion, the Mølmer-Sørensen gate uses bichromatic lasers tuned near the red and blue sidebands, addressing both ions to couple their internal states. By adjusting the phases of the light field, a controlled phase gate can be implemented between the ions, which is equivalent to a CNOT up to single-qubit operations. Importantly, this gate scheme does not require that the motional modes to be cooled to their ground states, only to the Lamb-Dicke regime [21].

The main limitation of both these gate schemes is that they require resolution of individual sideband transitions, and thus can only operate in the weak coupling regime $\Omega \ll \omega_t$. As described above, this means the timescales of these gates must be adiabatic to respect to the motional dynamics of the ions, and thus gate speed is limited by the trap geometry. Furthermore, the requirement that the Lamb-Dicke criterion (3.23) be satisfied requires increasingly longer gate times as more ions are present in a trap [35], and so these sideband-resolving gates do not scale to large ion crystals.

3.5 Fast gates

There exist a wealth of proposals for implementing entangling gates where multiple motional modes are excited during the gate operation [33, 35, 36, 68]. Because these gates operate in the strong-coupling regime, they are not subject to the same speed restrictions as the Cirac-Zoller and Mølmer-Sørensen gates. In very recent years, fast gates have been experimentally demonstrated [26, 34], however high-fidelity (i.e. above 99%) gates significantly faster than the trap frequency are yet to be realised.

In the strong-coupling regime, individual motional modes cannot be excited independently of one another. Fast gates involve the excitation of many motional modes simultaneously with state-dependent momentum kicks from ultra-fast laser pulses that excite state-dependent trajectories of the ions through phase space. These trajectories can be orchestrated to implement a controlled-phase gate (which is equiv-

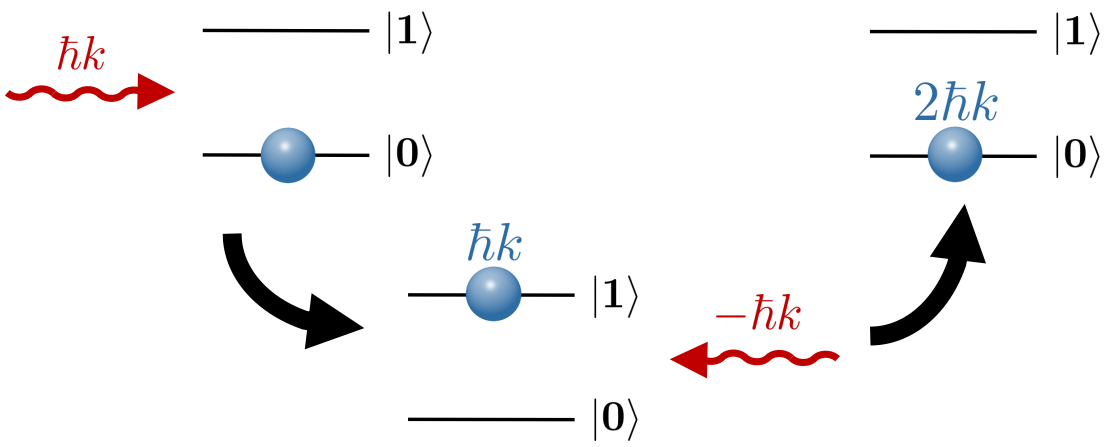

Figure 3.6: The effect of a counter-propagating pair of π pulses (red) with wavenumber k on an ion (blue) with initial internal state $|0\rangle$. The ion absorbs momentum $\hbar k$ from the π pulse which excites a transition $|0\rangle \rightarrow |1\rangle$. The second π pulse is counter-propagating, which is equivalent to a change in sign of the wavenumber, and causes the transition $|1\rangle \rightarrow |0\rangle$, i.e. restoring the ion to its initial internal state. This imparts a further momentum kick of $\hbar k$ on the ion, resulting in a total momentum kick of $2\hbar k$. If the initial state is $|1\rangle$, a similar process results in a total momentum kick of $-2\hbar k$.

alent to a CNOT, or any other maximally entangling two-qubit gate), provided that the motional and internal ion states are disentangled at the end of the gate operation.

3.5.1 Fast gate formalism

As presented in the original fast gate proposal [33], fast gates use counter-propagating π pulse pairs to excite state-dependent motion in an ion. The effect of each pulse pair is to impart a momentum transfer of $\pm 2\hbar k$ (where the sign depends on the initial state) leaving the internal state of the ion untouched. This is diagrammatically represented in Figure 3.6. In fast gates, state-dependent kicks on the ions are implemented by groups of counter-propagating pulse pairs (*pulse groups*), separated temporally by periods of free evolution. By carefully choosing the timings and magnitudes of these momentum kicks the controlled-phase operation,

$$\hat{U}_{\mathrm{CPH}} = \exp\left\{ i\frac{\pi}{4}\hat{\sigma}_z^A \otimes \hat{\sigma}_z^B \right\}, \tag{3.33}$$

can be implemented. Here $\hat{\sigma}_z^i$ is the Pauli Z operator acting on the internal state of the i-th ion. The ideal unitary for a fast gate of duration t_G is

$$\hat{U}_I = \hat{U}_{\mathrm{CPH}} \prod_p e^{-i\omega_p t_G \hat{a}_p^\dagger \hat{a}_p} \,. \tag{3.34}$$

This equation can be understood as the implementation of a controlled phase gate, explicitly including the restoration of the motional modes. By restoration here I mean that at the end of the gate, each motional mode is exactly in the state that it would be if there had only been free evolution over the period t_G, and no kicks at all [33]. Importantly, this means that the ideal implementation has no requirements on temperature, as the unitary is independent of initial motional state.

Fast gate fidelity

In Appendix C, I derive explicit conditions that need to be satisfied for ideal gate implementation. The first condition,

$$\frac{\pi}{4} = 8 \sum_{p=1}^{L} \eta^2 \frac{\omega_t}{\omega_p} b_1^{(p)} b_2^{(p)} \sum_{m=2}^{N} \sum_{k=1}^{m-1} z_m z_k \sin\left(\omega_p (t_m - t_k)\right) \tag{3.35}$$

describes the requirement of acquiring a relative phase difference of $\pi/4$. Here, $b_i^{(p)}$ is the coupling of the i-th ion to the p-th motional mode, and t_m and z_m are the timing and number of pulse pairs in the m-th pulse group. Geometrically, this phase difference can be understood as the difference in areas enclosed by different state-dependent trajectories in phase space. Additionally, the motional restoration conditions require satisfaction:

$$\left| 2\eta \sqrt{\frac{\omega_t}{\omega_p}} \sum_{k=1}^{N} z_k e^{i\omega_p t_k} \right| = 0 \,. \tag{3.36}$$

In general the unitary that is actually implemented in a fast gate, $\hat{U}_{re}$, is only an approximation to the ideal operation. As introduced in the previous chapter, the state-averaged fidelity is a useful measure to characterise how well the ideal unitary is approximated. Following the Supplementary Material in Ref. [36], the state-averaged fidelity for a fast gate can be expressed as

$$F_{av} = \frac{1}{\int_{\psi_0} d\,|\psi_0\rangle} \int_{\psi_0} \mathrm{Tr}_m \left[\langle\psi_0| U_{id}^\dagger U_{re} (|\psi_0\rangle \langle\psi_0| \otimes \rho_m) U_{re}^\dagger U_{id} |\psi_0\rangle \right] d\,|\psi_0\rangle \tag{3.37}$$

where the partial trace is taken over the motional states. The motional density

matrix is assumed to be given by a thermal product state for each mode,

$$\rho_m = \bigotimes_{p=1} \rho_m^{(p)}, \text{ where} \tag{3.38}$$

$$\rho_m^{(p)} = (1 - e^{-\hbar\omega_p/k_BT}) \sum_{n=0}^{\infty} e^{-n\hbar\omega_p/k_BT} |n\rangle \langle n| , \tag{3.39}$$

where T is the temperature of the ion system. In this thesis I will focus on gates with fidelities extremely close to unity, and so I will report the infidelity $1 - F_{av}$. Refs. [38, 68] present a truncated form of (3.37)

$$1 - F \approx \frac{2}{3}|\Delta\phi|^2 + \frac{4}{3}\sum_p (\frac{1}{2} + \bar{n}_p) \left((b_p^A)^2 + (b_p^B)^2 \right) |\Delta P_p|^2 , \tag{3.40}$$

where $\bar{n}$ is the average phonon occupation of each mode. Here, $\Delta\phi$ is the phase-mismatch and ΔP_p is the unrestored motion, which are given by

$$|\Delta\phi| = \left| 8\eta^2 \frac{\omega_t}{\omega_p} b_p^A b_p^B \sum_{i \neq j} z_i z_j \sin\left(\omega_p |t_i - t_j|\right) \right| - \frac{\pi}{4} \tag{3.41}$$

$$|\Delta P_p| = 2\eta \sqrt{\frac{\omega_t}{\omega_p}} \sqrt{ \left(\sum_i z_i \cos\left(\omega_p t_i\right) \right)^2 + \left(\sum_i z_i \sin\left(\omega_p t_i\right) \right)^2 } . \tag{3.42}$$

This truncated form is extremely robust for high-fidelity gates [2], and is computationally efficient for numerical optimisation of gates [38]. Note that for perfect phase acquisition ($\Delta\phi = 0$) and perfect motional restoration ($\Delta P_p = 0$ for all p), the ideal gate is implemented ($F = 1$).

3.5.2 Optimality of nearest-neighbour gates

For high-fidelity gates, each motional mode must be restored at the end of the gate operation. For multi-ion chains, this is somewhat complicated by the fact that distant ions are connected via the Coulomb interaction, and thus state-dependent momentum kicks on a pair of ions necessarily perturb the motional modes of all other ions in the chain.

This perturbation can be made small by performing the entire gate operation faster than the trap period, which is the characteristic timescale of the motional oscillations (on the order of 1 μs). In the limit of an infinitely fast gate, these motional perturbations vanish. However, this is only true for neighbouring pairs of ions - non-neighbouring pairs of ions are connected by the motion of the ions in between, and thus do not have this fast-speed limit. This has been studied by Bentley *et al.*

[36] for chains of ions in a linear Paul trap. The authors found that it is always optimal to connect distant ions by a series of *nearest-neighbour* fast gates.

It is worth noting that fast gates can in fact be directly performed on non-neighbouring pairs of ions, however in general it is better to use multiple nearest-neighbour gates. For the purposes of this thesis, I will assume that this holds for one-dimensional microtrap architectures, in which non-nearest neighbour fast gates have not been explicitly studied. The scaling arguments in [36] suggest that it is unclear whether nearest-neighbour gates are still optimal for diagonally separated ions in two-dimensional arrays. This is a key question of this research, and is addressed in Chapter 6.

3.5.3 Gate schemes for efficient optimisation

As fast gates are not constrained by trap geometry [1], the main limitation in their implementation is the available laser power [33]. For a given maximium laser repetition rate (the rate at which the laser emits pulses), the challenge is to design gate schemes that implement the ideal fast gate unitary (3.34), which implicitly requires satisfaction of the phase acquisition (3.35) and motional state restoration (3.36) conditions.

Unfortunately, non-trivial analytical solutions that satisfy these conditions are not known, and so finding gates that implement the ideal unitary of a controlled phase-gate with high fidelity and with short gate time becomes an optimisation problem. The aim of this optimisation problem is to find values of free variables, which are the timings and number of pulse pairs in each pulse group, that implement the ideal unitary for a controlled-phase gate (3.34), as close to unit fidelity as possible. As the desired fidelities are close to unity, this can be framed as the minimisation of a cost function, where the cost function is typically chosen to be gate infidelity. For tractable computation however, it is favourable to work with the truncated form of the state-averaged infidelity (3.40).

Optimising fast gates can be understood as designing a gate such that the trajectory of the ions in phase space results in: (a) accumulation of a relative phase as close as possible $\frac{\pi}{4}$, (b) the restoration of each motional mode close to where it would have been (in phase space) if there were no momentum kicks. These trajectories depend on two aspects of the gate: the state-dependent momentum kicks, and the timings at which they arrive at the ions. As described earlier, these state-dependent kicks can be implemented through groups of counter-propagating pulse pairs, with each pulse pair contributing a momentum kick of $\pm 2\hbar k$ (where the sign depends on the

internal state of the ion).

Naively, for a given number of pulse groups, N, this approach requires optimisation over the $2N$ elements of the vectors:

$$\mathbf{Z} = \{z_1, z_2, \ldots, z_N\}; \tag{3.43}$$

$$\mathbf{t} = \{\tau_1, \tau_2, \ldots, \tau_N\}, \tag{3.44}$$

where z_k represents the number of counter-propagating pulse pairs in the k-th pulse group (i.e. the effective magnitude of the k-th momentum kick), and τ_k is the timing of the k-th kick in trap period units, i.e. in units of $\frac{2\pi}{\omega_t}$. Global optimisation over this fully general solution space is completely intractable. This can be addressed by placing additional constraints on the search.

Optimisation over pulse group timings

The conventional approach thus far has been to place fixed ratios on the z_k elements and optimise over an anti-symmetrised $\mathbf{t}$-vector [33, 35, 36]. That is, optimisations are done over the elements $\{n_k, \tau_1, \ldots, \tau_{N/2}\}$ for a gate scheme of the form

$$\mathbf{Z} = n_k\{-a, -b, -c, \ldots, c, b, a\}, \tag{3.45}$$

$$\mathbf{t} = \{-t, -t_2, \ldots, -t_{N/2}, t_{N/2}, \ldots, t_2, t_1\}. \tag{3.46}$$

Anti-symmetric schemes have the benefit of guaranteeing restoration of the velocity of each motional state, which simplifies the expression for motional restoration to

$$2\eta\sqrt{\frac{\omega_t}{\omega_p}} z_k \sin(\omega_p t) = 0. \tag{3.47}$$

For a six-pulse group scheme, this reduces the unconstrained twelve dimensional space of solutions to a four dimensional optimisation over $\{n_k, \tau_1, \tau_2, \tau_3\}$. The original fast gate scheme proposed by Garciá, Zoller, and Cirac (GZC) falls into this category, specifying (a, b, c) values of $(1, -2, 2)$ [33]. Several years later, Bentley *et al.* [36] proposed a scheme that falls with (a, b, c) values of $(2, -3, 2)$, which they found to be more robust and able to achieve higher-fidelities for given gate times. This has been dubbed the 'Fast Robust Antisymmetric Gate' scheme (FRAG), and is the optimal choice for gate optimisation currently in the literature. Ratcliffe *et al.* [69] have demonstrated that optimisations in which the ordering of a, b, and c is not fixed allows for high-fidelity gates to be found for a wider range of gate times and experimental parameters. In unpublished work, Lachlan Oberg studied the re-

moval of strict ratios on pulse groups, extending the optimisation to search over integer values of (a, b, c). While this was successful in improving achievable gate fidelities, the larger solution space limits extensions to larger number of pulse groups for computationally tractable optimisation.

The main limitation to the general approach of optimising over pulse timings is that it is very inefficient. The cost function is sinusoidal in pulse timings, i.e. $1 - F \sim \sin^2(\omega\tau_k)$, which means the corresponding solution space is dense with local minima. An optimisation in this space is far less likely to approach a global minima, and thus requires a very large search density that scales unfavourably with the number of dimensions, effectively limiting the number of pulse groups that can be optimised with finite computational resources. Furthermore, even searches with small n_k can result in gate solutions that require restrictively large finite repetition rates to resolve, as by nature of searching over pulse timings, there is no cost of pulse groups being arbitrarily close together.

Generalised Pulse Group (GPG) scheme

In an unpublished honours dissertation [37], Evan Gale addressed these limitations by considering a gate scheme in which pulse groups are regularly timed with a fixed gate speed, and the optimisation is done over elements of the $\mathbf{Z}$ vector. In this thesis I will refer to it as the Generalised Pulse Group (GPG) scheme. For a total gate time τ_G and N pulse groups, the GPG scheme has the form

$$\mathbf{Z} = \{z_1, z_2, \ldots, z_N\}; \tag{3.48}$$

$$\mathbf{t} = \tau_G \{0, \frac{1}{N}, \frac{2}{N}, , \ldots, \frac{N-1}{N}, \}, \tag{3.49}$$

where the search is over the elements $\{z_1, z_2, \ldots, z_N\}$. There are several advantages of this approach. The first is that the solution space contains less minima as it is quartic in this space, i.e. $1 - F \sim \mathcal{O}(z_k^4)$, and is thus much less dense. The regular spacing of the timings of the pulse groups is also favourable as it minimises the repetition rate required to resolve a gate with prescribed z_k values, and therefore lends itself to less strict laser power requirements for fast, high-fidelity gates. However, even with the simplified solution space this scheme becomes computationally intractable for more than $N{\sim}16$ pulse groups. Whilst this dimensionality restriction still allows for extremely high-fidelity gates to be designed for two isolated ions which only require restoration of two motional modes, it becomes insufficient when a large number of modes require restoration. This will become apparent in the course of this thesis when I consider optimisations of fast gates in systems with more than

two ions.

Gale also presented a method for further local optimisation using a set of ODEs that characterise the classical motion of the ions. These ODEs can be locally optimised over pulse group timings, which can be used to further improve the fidelity of a gate solution. Furthermore, this description includes non-linearities in the Coulomb interaction, and can also be extended to include non-linearities in the trapping potential as well. The motivation for the classical treatment of the ions motion comes from the fact that the ion's motional wavefunction is typically well-approximated as a Gaussian wave-packet, the center of which follows a classical trajectory [44]. The fast gate momentum kicks can be understood as displacement operators that shift the Gaussian state without changing its shape [36], and so the system is well-described by a classical description of the motion of the centroid (i.e. an ODE description). I will not utilise ODE optimisation within the scope of this thesis, but will note that it is an essential step in any practical realisation of fast gates in Chapter 7.

Including finite repetition rate

In all the optimisation schemes described above, it is assumed that the momentum kick imparted by each pulse group is *instantaneous*. This is unrealistic, as a real laser will have some finite rate, f_{rep} at which pulses are emitted. This is known as the laser *repetition rate*, and is closely related to its power [70]. As the achievable gate speed is directly related to the number of pulse pairs available [33], finite repetition rates is a key limiting factor for fast gates. Current state-of-the-art lasers have repetition rates as high as $1 - 5$ GHz [50, 70].

Duan was the first to explicitly include finite repetition rate effects, breaking up pulse groups into individual pulse pairs separated by $t_{rep} = 1/f_{rep}$ [35]. In the Duan scheme, the delay time between pulses is the only free evolution. While this lends itself to minimising gate time, it leaves only a single parameter free to be optimised and thus the scheme is unable to satisfy the condition equations (3.35) and (3.36) [71]. Other fast gate schemes have since been well-studied with finite-repetition rate effects included [1, 2, 38]. It is also possible to optimise a gate for a particular repetition rate by enforcing pulses to be snapped to a grid of timings (with spacing given by $1/f_{rep}$), and locally optimising over the timings of the individual pulses. This is best done with an ODE description of the ions (as described above). For the results presented in this thesis, finite repetition rate will not be included in any gate optimisation. Fast gates tend to be robust to finite repetition rates as long as pulse groups do not overlap [37, 38], and thus I will find it sufficient to report the minimum

repetition requirement such that pulse groups can be individually resolved, within the scope of this thesis. I will refer to this as the *minimum resolving repetition rate*, f_{min}.

3.6 Quantum simulation of fermions with trapped ions

In order to demonstrate the suitability of fast gates to enable scalable quantum-computing in the near future, it is useful to consider particular examples of quantum computation. As mentioned in the introductory chapter, one such task that has been well-studied on trapped ion systems is *quantum simulation* [2, 11, 14, 72–74]. Of particular interest is the simulation of interacting fermionic systems. Classically simulating fermionic systems is known to be notoriously difficult due to the fermion sign-problem [75] that precludes the use of quantum Monte-Carlo methods, as well as the exponential scaling of the Hilbert space with the number of fermionic modes [8].

Algorithms for quantum simulation can in general be understood as a mapping of some quantum system of interest to a system of interacting qubits, and simulating its dynamics through some quantum circuit. The mapping itself is done through one of several transformations. I will focus on the Jordan-Wigner transformation for mapping fermionic operators to Pauli operators that act on qubits [76]. The time evolution operator for the system of interest is discretised such that each step can be implemented through a circuit of logical gate operations, typically via Trotterisation (described below) [77].

Here, I will briefly outline the details of a scheme originally proposed by Lamata *et al.* [74] for simulating fermionic systems, i.e. systems with Hamiltonians that can be written in terms of fermionic creation/annihilation operators $\hat{b}^\dagger/\hat{b}$. The protocol can be broken into three steps.

1. The fermionic Hamiltonian is mapped onto a spin representation by the Jordan Wigner transform (JWT)[2]

$$\hat{H} \xrightarrow{JWT} \sum_j \hat{H}_j = \sum_{i,j,k,\ldots} g_{ijk\ldots} \left(\hat{\sigma}_i^1 \otimes \hat{\sigma}_j^2 \otimes \hat{\sigma}_k^3 \ldots \right) \tag{3.50}$$

   where $\hat{\sigma}_j^m$ is the Pauli $j \in \{\mathbb{1}, x, y, z\}$ operation on the m-th qubit. In general,

[2]More detail on the Jordan Wigner transform can be found in Appendix D, with respect to the transformation of the Fermi-Hubbard Hamiltonian considered in the next chapter.

many of the terms $\hat{H}_j$ will be composed of non-local spin-operations.

2. The time evolution unitary can be decomposed into products of exponentials of each $\hat{H}_j$:

$$\hat{U}(t) = e^{-\frac{i}{\hbar}\sum_j \hat{H}_j t} \approx \left(\prod_j e^{-i\hat{H}_j \delta} \right)^n . \tag{3.51}$$

This decomposition neglects non-commutativity of the individual terms in the sum[3], and thus has an error on the order $O\left((\frac{t}{n})^2\right)$, which can be made arbitrarily small for sufficiently large n [10]. Here $\delta \equiv t/n$ is the *Trotter step*, and is essentially the discretised time-step for numerical simulation of the evolution.

3. Each term in the product can be implemented with a pair of non-local multi-qubit entangling gates and a single-qubit rotation on some m-th qubit:

$$e^{-\frac{i}{\hbar}\hat{H}_j \frac{t}{n}} = \hat{U}^\dagger_{UMQ}\, e^{-\frac{i}{\hbar}\phi\hat{\sigma}^m_z}\, \hat{U}_{UMQ}\,, \tag{3.52}$$

where $\hat{U}_{UMQ}$ is the *ultrafast multi-qubit* (UMQ) gate

$$\hat{U}_{UMQ} = \exp\left(-\frac{i}{\hbar}\frac{\pi}{4}\hat{\sigma}^m_z \sum_{j\neq m} \hat{\sigma}^j_z \right). \tag{3.53}$$

Note that the summation in (3.53) is only over the qubits specified in the non-local multi-qubit spin operators in (3.50).

The naming of the gate reflects the fact that UMQ operations can be constructed out of fast gates [2, 74]. This will be explicitly shown in the next chapter where I investigate a quantum simulation of the Fermi-Hubbard model with nearest-neighbour fast gates in microtrap ion architectures, using the protocol described above.

[3]The Trotter decomposition described here is only a first order method - higher-order methods are possible with Trotter-Suzuki decomposition methods [78], but will not be considered in this thesis.

Chapter 4

Quantum simulation of the Fermi-Hubbard model with fast gates

In order to investigate the suitability of fast gates for near-future quantum computing, it is useful to consider a particular task that requires large-scale computation. As alluded to in the introductory chapter, the digital simulation of a fermionic system with more than 40 modes is a good example of a task feasible for a near-future quantum computer that is largely intractable to classical computing resources. Here I will consider a fermionic system described by the *Fermi-Hubbard model*, which is perhaps the simplest model that contains strongly correlated dynamics that cannot be simulated classically. Originally developed to describe electrons in solids [79], the Fermi-Hubbard model can been applied to certain chemical systems [80, 81], and as a toy model for superconducting cuprates [82].

In this chapter, I will study the use of the quantum algorithm outlined in Ref. [2] for simulating the Fermi-Hubbard model. I will describe how the algorithm can be built up of fast gates between pairs of trapped ions in both one-dimensional and two-dimensional architectures, and calculate the total number of fast gates required to implement the simulation. As discussed in Section 3.5.2, I will assume the use of nearest-neighbour gates to always be the optimal over non-local operations. Following the approach of [2], I will build non-local operations from sequences of nearest-neighbour fast gates. In Section 4.4, I will present a discussion on the results of this chapter, and estimate bounds on gate requirements to implement a simulation that outperforms what can be achieved on a classical computer [74].

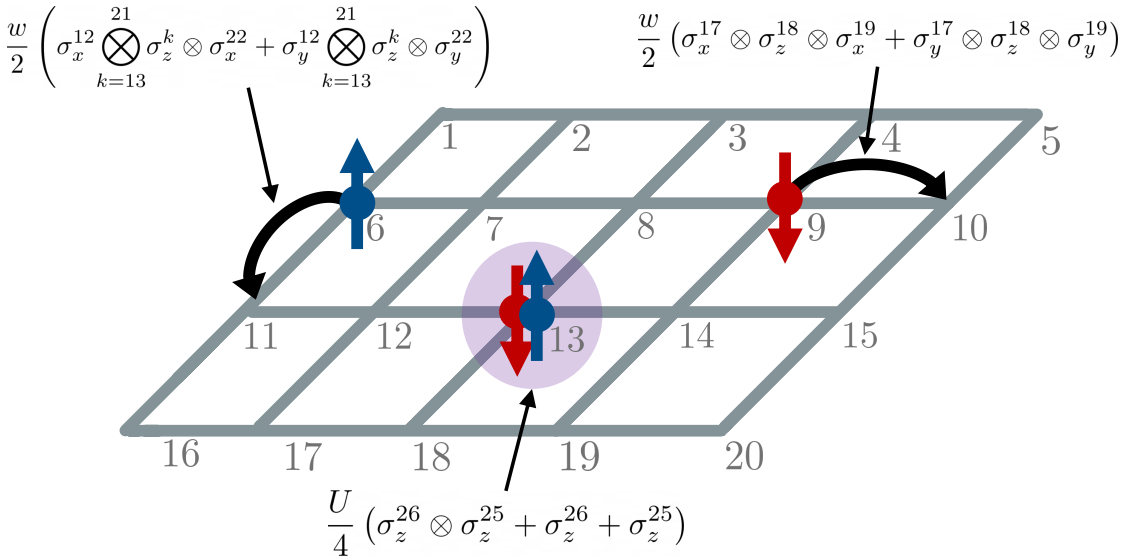

Figure 4.1: Dynamics of the Fermi-Hubbard model of spin-$\frac{1}{2}$ particles: hopping between nearest neighbour sites (black arrows), and onsite interaction (purple) between spin up and spin down fermions. Examples of the different types of terms in the Jordan-Wigner mapped Hamiltonian are shown. Conjugate terms are not explicitly visualised. For simulating the model for the 20-site lattice shown, 40 qubits would be required - one for each spin occupancy of each site.

4.1 The Fermi-Hubbard Hamiltonian

In Section 3.6, I described a general protocol for implementing a simulation of a fermionic system using ultra-fast multi-qubit (UMQ) gates, originally proposed in Ref. [74]. Here I will investigate the use of nearest-neighbour fast gates to implement this protocol for the simulation of a 40 mode Fermi-Hubbard Hamiltonian. Specifically, I will investigate the Fermi-Hubbard model for 40 spin-$\frac{1}{2}$ fermions interacting with nearest-neighbours in a two-dimensional 5×4 lattice. In second quantised notation[1] the Hamiltonian I will consider has the form

$$H = w \sum_{\langle i,j\rangle,\sigma}^{20} (\hat{b}_{i,\sigma}^\dagger \hat{b}_{j,\sigma} + \text{h.c.}) + U \sum_{j=1}^{20} \hat{b}_{j,\uparrow}^\dagger \hat{b}_{j,\uparrow} \hat{b}_{j,\downarrow}^\dagger \hat{b}_{j,\downarrow}\,, \tag{4.1}$$

where $\hat{b}^\dagger$ $\left(\hat{b}\right)$ is the fermionic creation (annihilation) operator, $\sigma = \uparrow, \downarrow$ is the spin index, and $\langle i, j\rangle$ denotes nearest-neighbour pairing. The first term represents tunnelling of fermions between neighbouring sites, and the second term describes the potential generated by on-site occupation.

This Hamiltonian can be mapped onto a system of 40 interacting qubits by the

[1] Also known as number notation. A comprehensive introduction to this notation can be found in any modern quantum mechanics textbook, such as Ref. [59].

Jordan-Wigner transformation. In Appendix D, I show that equation (4.1) has the following mapping,[2]

$$\hat{H} \rightarrow \sum_j H_j = w \sum_{\lambda=x,y} \left(\sum_{j=2}^{39} \sigma_\lambda^{j-1} \otimes \sigma_z^j \otimes \sigma_\lambda^{j+1} + \sum_{j=1}^{30} \sigma_\lambda^j \bigotimes_{k=j+1}^{j+9} \sigma_z^k \otimes \sigma_\lambda^{j+10} \right) + U \left(20 + \sum_j^{20} \sigma_z^{2j} \otimes \sigma_z^{2j-1} + \sum_k^{40} \sigma_z^k \right) . \quad (4.2)$$

In total there are 185 terms in this expression, including a single constant term, 40 single-body terms, 20 two-body terms, 64 three-body terms, and 60 ten-body terms. Figure 4.1 provides a visualisation of some of these terms. As per the Trotter expansion step of the formula, each of these terms correspond to a unitary $e^{-\frac{i}{\hbar}H_j\delta}$. The product of these unitaries forms the time-evolution operator for a single Trotter step, which can each be simulated with multi-qubit gates on the trapped ions.

Naive applications of the simulation protocol would suggest that each of these terms require UMQ gates of form (3.53) to simulate, however many of these terms can be implemented in a simpler manner. The constant term can be neglected as it simply amounts to a non-physical global phase of the Fermi-Hubbard system being simulated, and each of the single-body terms can be simulated with rotations on single-qubits. Further reduction can be achieved by noting that all the two-body terms can be simulated with a single fast gate between neighbouring qubits.[3] This leaves just the three-body and ten-body terms, which will require UMQ operations.

4.2 Implementation on a one-dimensional chain of ions

Taylor *et al.* [2] reported 268 UMQs required per Trotter step for simulating the Fermi-Hubbard model described above, with $4N - 7$ nearest-neighbour fast gates required to construct each UMQ, for $N = 40$ qubits. This analysis was built from the worst case assumption that each UMQ involves entangling qubits as far as opposite ends of the chain through nearest-neighbour fast gates and SWAP operations. However, such terms would only appear when simulating fermionic Hamiltonians

[2] In order to differentiate between operators represented by fermionic creation/annihilation operators, and Pauli-operators, I will omit hats on operators in the spin representation.

[3] Technically, these terms are equivalent to fast gates between qubits that have adjacent *indexing*, which in one-dimension trivially corresponds to nearest-neighbouring ions. This association is less trivial in two-dimensions which I will address later in the chapter.

with highly non-local interactions and even in such extreme cases there will always be terms that involve fewer than N qubits. In most practical applications, only smaller subsets of the qubits will be required for each UMQ.

I will first investigate the number of gates required for simulating the Hamiltonian (4.1) on a one-dimensional chain of ions in a microtrap array. This is in contrast to the approach of Ref. [2] where the authors considered fast gates between ions all in a single linear Paul trap. This is motivated by recent work in Ref. [38], where the scaling of microtrap chains was shown to be favourable to that of ions in a single trap.

Three-body terms

I have already shown that UMQs are only required to simulate terms involving three and ten qubits, which is already a significant reduction in complexity from the worst case. Furthermore, the UMQ operations for simulating the three-body terms involve only local pairings of qubits, and thus can be constructed out of just nearest-neighbour gates without any SWAP operations. For example, consider the three-body term $\sigma_x^4 \otimes \sigma_z^5 \otimes \sigma_x^6$. The unitary corresponding to this term can be simulated with a pair of forward/reverse UMQ operations $\{U_{\mathrm{UMQ}}^{\mathrm{three}}, \left(U_{\mathrm{UMQ}}^{\mathrm{three}}\right)^\dagger\}$ and a local rotation on ion 5,

$$e^{i\phi\sigma_x^4\otimes\sigma_z^5\otimes\sigma_x^6} = \left(U_{\mathrm{UMQ}}^{\mathrm{three}}\right)^\dagger e^{-i\phi\sigma_z^5} U_{\mathrm{UMQ}}^{\mathrm{three}} = e^{+i\frac{\pi}{4}(\sigma_x^4\otimes\sigma_x^5+\sigma_x^5\otimes\sigma_x^6)} e^{-i\phi\sigma_z^5} e^{-i\frac{\pi}{4}(\sigma_x^4\otimes\sigma_x^5+\sigma_x^5\otimes\sigma_x^6)}. \tag{4.3}$$

By observation of the RHS, this UMQ can be built up of nearest-neighbour fast gates between qubits 4 and 5, and qubits 5 and 6, up to single-qubit rotations to transform between σ_z and σ_x. Three-body terms involving σ_y, such as $\sigma_y^4 \otimes \sigma_z^5 \otimes \sigma_y^6$, can be similarly constructed, with a different set of single-qubit rotations. There are 64 three-body terms, each requiring a pair of UMQs which in turn can each be constructed with 2 fast gates. In total, implementing the three-body terms requires 256 fast gates per Trotter step.

Ten-body terms

The ten-body terms do, however, require non-local couplings between ions. In other words, building the ten-body from nearest-neighbour fast gates requires SWAP operations. These terms are of the form $\sigma_\lambda^i \otimes \cdots \otimes \sigma_\lambda^{i+10}$ for every consecutive set of

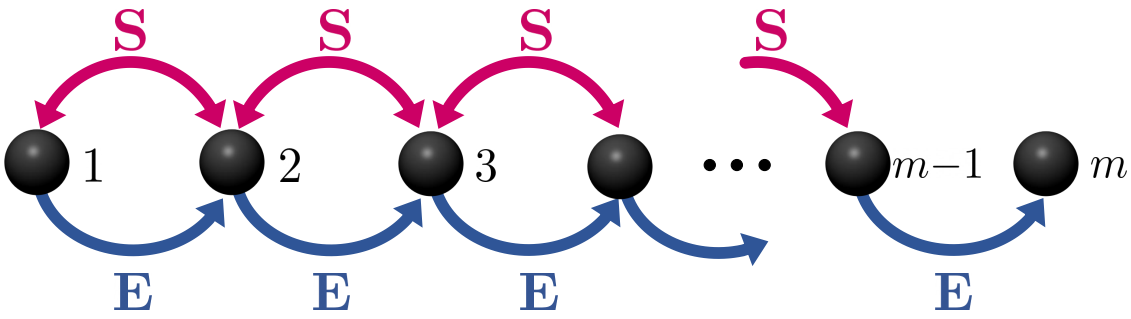

Figure 4.2: Construction of a m-body UMQ with nearest-neighbour fast gates and SWAP operations. SWAP operations (pink) and nearest-neighbour fast gate entangling operations (blue) are depicted. In total, $m-2$ SWAP operations and $m-1$ nearest-neighbour fast gates are required.

10 qubits in the chain, where $\lambda = x, y$. The corresponding UMQ has the form,

$$U_{\mathrm{UMQ}}^{\mathrm{ten}} = \exp\left\{-i\frac{\pi}{4}(\sigma_\lambda^i \otimes \sigma_\lambda^{i+1} + \cdots + \sigma_\lambda^i \otimes \sigma_\lambda^{i+10})\right\}. \tag{4.4}$$

This can be interpreted as a series of non-local controlled-phase (CPH) gates between pairs of qubits i and $i+1$, i and $i+2$, up to i and $i+10$. Figure 4.2 visualises the manner in which this operation can be implemented with nearest-neighbour fast gates: a fast gate entangles the first and second qubits, and then the two are swapped. The first qubit is now nearest-neighbours with the third qubit, and the two are entangled with a fast gate. This process is repeated, until the first qubit is nearest-neighbours with the eleventh qubit, which are then coupled together with a fast gate. This is sufficient to construct the ten-body UMQ 4.4, up to single-qubit rotations on each qubit.

I note that while this UMQ operation changes the positions of the qubits in the chain, the reverse UMQ operation will restore each qubit to their original ions. For a ten-body UMQ, this process requires 9 nearest-neighbour fast gates and 8 SWAP operations. There are 60 of these ten-body terms, each of which requires a forward and a reverse ten-body UMQ. Each SWAP gate can be constructed with 3 nearest-neighbour fast gates, up to local rotations. The total number of fast gates contributed by the ten-body terms is $2 \times (9 + 3 \times 8) \times 60 = 3960$. In total, $20 + 256 + 3960 = 4236$ nearest-neighbour fast gates are required for the simulating the Hamiltonian (4.1) for a single Trotter step, on a one-dimensional chain of ions.

4.3 Implementation on a two-dimensional array of ions

In the one-dimensional case studied above, the resource requirements for the simulation, in terms of number of fast gates needed, are dominated by the SWAP operations used to connect distant ions. In higher-dimensional arrays, ions are naturally closer together, which promises a significant reduction in the total number of nearest-neighbour gates required to connect distant ions. I will now consider a two-dimensional array of ions, and explicitly describe the construction of UMQs with nearest-neighbour gates.

Due to the results of Chapter 6, I will work under the assumption that fast gates between diagonally separated ions are not only possible in two-dimensional arrays, but superior to equivalent operations constructed from multiple 'adjacent' (i.e. horizontal/vertically nearest-neighbour) fast gates. Implementing the algorithm solely with adjacent fast gates is a straightforward extension of the following calculation.

For simplicity I will consider a two-dimensional array of micro-traps, with the distance between horizontally and vertically neighbouring micro-traps equal, i.e. $d_x = d_y = d$. In general I will refer to such arrays as *symmetric*. Connectivity of ions is key to reducing the number of SWAP operations required for non-local couplings, and so square arrays (i.e. arrays of size $M \times M$) are optimal. I will consider a 6×7 array of ions to realise the 40 qubits required for the simulation. This is chosen to maximise connectivity, and is pictured in Figure 4.3.

The indexing of the qubits in this array is non-trivial in two-dimensions, and corresponds to a choice in the mapping of the fermionic Hilbert space to the Hilbert space of interacting qubits. I have chosen a particular non-sequential indexing, as shown in Figure 4.3, to ensure that the i-th qubit is always a nearest-neighbour to both the $(i-1)$-th and $(i+1)$-th qubits. This is desirable as, just as in the one-dimensional case, that the two-body terms can each be implemented with a single fast gate between adjacent nearest-neighbours. Furthermore, the UMQs that simulate the three-body terms can similarly be constructed out of a pair of fast gates between adjacent nearest-neighbours. Therefore, just as in the one-dimension case, these terms require a total of $20 + 256 = 276$ fast gates to be simulated, per Trotter step.

Because this choice of indexing the qubits in the array is non-trivial, it is no longer possible to describe how the ten-body UMQs can be constructed using SWAP gates with a single general case. I have identified four unique types of ten-body UMQs that

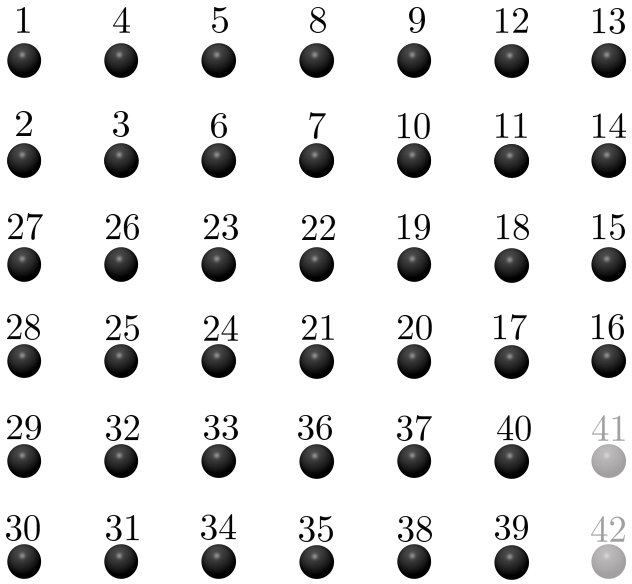

Figure 4.3: A visualisation of the 6×7 array of ions, chosen for implementation of the simulation. The indexing of the qubits is deliberately chosen as non-sequential to ensure that the two-body and three-body terms do not involve couplings between non-neighbouring qubits. The two ions marked in grey are not required for the algorithm, and are only included here for completeness of the array.

each require different numbers of adjacent and diagonal nearest-neighbour gates to construct. I describe the manner in which they are constructed from SWAP and fast gate operations diagrammatically in Figure 4.4. I have tabulated the total number of fast gates required for the simulation, per timestep, in Table 4.1, noting that there are two terms $\sigma_x^i \otimes \cdots \otimes \sigma_x^{i+10}$, $\sigma_y^i \otimes \cdots \otimes \sigma_y^{i+10}$ for each i, and two UMQs are required for each term.

It is useful to note here that the diagonal gates allow for a significantly reduced number of SWAP operations for the ten-body UMQs, and thus fewer fast gates required

Term	Adjacent fast gates	Diagonal fast gates
Pair terms	20	0
Three-body terms	256	0
Ten-body (A)	800	40
Ten-body (B)	416	160
Ten-body (C)	416	64
Ten-body (D)	160	96
Total	2068	360

Table 4.1: Number of nearest-neighbour fast gates required to implement the simulation algorithm, per term in the mapped Hamiltonian (4.2). The contribution of terms that require only single-qubit rotations have been neglected.

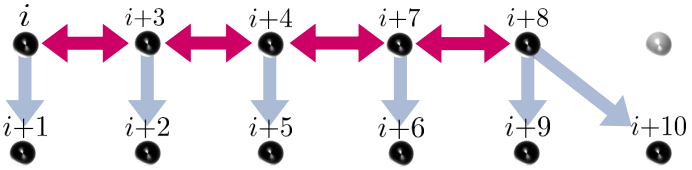

(a) Type A: $i \in \{1, 2, 3, 4, 15, 16, 17, 18, 29, 30\}$.

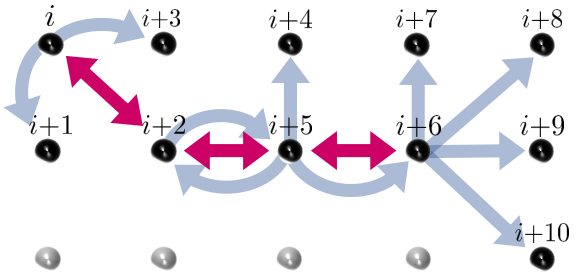

(b) Type B: $i \in \{5, 6, 13, 14, 19, 20, 27, 28\}$.

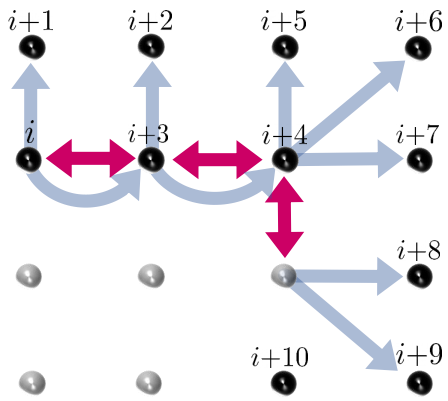

(c) Type C: $i \in \{7, 8, 11, 12, 21, 22, 25, 26\}$.

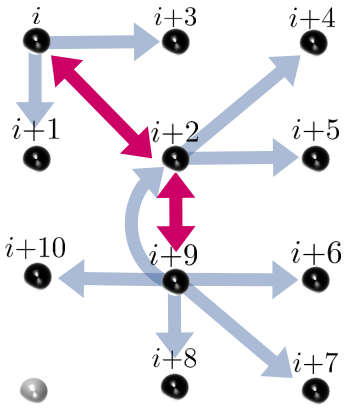

(d) Type D: $i \in \{9, 10, 23, 24\}$.

Figure 4.4: Four distinct types of ten-body UMQ gates are shown in terms of their construction out of adjacent and diagonal nearest-neighbour gates. SWAP operations (pink) are shown as double ended arrows, and fast-gate operations (blue) are visualised as a single-ended arrow beginning from the current location of the i-th qubit.

for the total simulation. It is possible to construct these UMQs purely from horizontal and vertical nearest-neighbours, and it is a straightforward extension to show that each ten-body UMQ would require five or fewer adjacent SWAP operations. In this case, 3154 fast gates are required per Trotter step, which is significantly larger than the totals of Table 4.1.

4.4 Discussion

A fast gate typically involves hundreds or even thousands of laser pulses. Therefore it is not necessary to track single-qubit operations which require only single laser pulses, when commenting on overall simulation infidelity and time costs. Therefore, despite the fact that many single-qubit operations are involved in the simulation, I will assume that resource requirements for the simulation are dominated by the two-qubit operations. I will also assume contributions from other single-pulse operations are negligible, such as the shelving of qubits into long-lifetime internal states (explored further in Section 7.2).

The total numbers of fast gates required for the simulation algorithm in the one-dimensional and two-dimensional cases studied in this chapter are both significantly reduced from those reported in Ref. [2], where the authors reported a requirement of 268 UMQs $\times$ 153 gates per UMQ $\sim 40,000$ nearest neighbour fast gates per Trotter step. As such it is useful to re-investigate the gate requirements to implementing the simulation such that it exceeds the capacity of classical computers. Lamata *et al.* [74] suggest that simulating the 40-mode Fermi-Hubbard model considered in this chapter would require a minimum of 10 Trotter steps. Taylor *et al.* [2] used a threshold of 70% total simulation fidelity as a minimum for determining a successful implementation. However this is largely arbitrary, as the requirements for individual gate fidelity will be almost identical for a 90% threshold, as the total fidelity is a product of such a large number of gate operations. I will estimate the fidelity of the total computation by multiplying the individual fast gate fidelities together, to place bounds on gate requirements. This estimate is valid for small gate errors, which is necessarily true for fidelities required for large-scale computation.

The other aspect that needs to be taken into account is the total time taken by the simulation algorithm. This must be considerably less than to the coherence times of trapped ion qubits, which can be on the order of seconds. Most of the studies of fast gates have been done considering $^{40}Ca^{+}$ ions, which has a meta-stable $D_{5/2}$ state that can be used for qubit storage with lifetimes on the order of $\sim$2s [21]. A stronger bound can be placed by considering heating rates, which for current traps are on

the order of $10s^{-1}$ [2], which limits available time for the entire implementation of the simulation to $\sim$100ms.

For implementation on a chain of 40 ions each in its own micro-trap, I found that 4236 nearest neighbour fast gates were required per Trotter step. To achieve a total simulation fidelity of 70% then, individual gate infidelities of the order $\sim 7 \times 10^{-6}$ are required. I found that in two-dimensional arrays, even fewer nearest neighbour fast gates are required to implement the algorithm, with only 2468 gates needed per Trotter step. For 10 steps, this corresponds to an individual gate infidelity requirement of $\sim 10^{-5}$ to achieve a total simulation fidelity of 70%. This is a slightly weaker upper bound than the one-dimensional case. Importantly the significant reduction of the number of gates means the entire simulation would be less susceptible to experimental errors and can be implemented in almost half the time.

In the final chapter of this thesis, I will discuss the feasibility of implementing this simulation algorithm on a trapped ion quantum computer in the near future. This discussion will involve calculated gate fidelities that will be presented in the next two chapters for one and two-dimensional microtrap architectures, modified by realistic experimental errors that will be explored in Chapter 7.

Chapter 5

Efficient fast gate optimisation for scaled ion crystals

In order to realise the simulation algorithm introduced in the previous chapter, or indeed any large-scale computation on a trapped ion processor, fast gates must perform well in large arrays of ions. This is where the current optimal gate schemes fall short. The FRAG and Generalised Pulse Group (GPG) schemes were invented to optimise gates for two-ion systems, and are not necessarily optimal for multi-ion systems. These schemes have been shown to be sufficient for one-dimensional chains of ions in microtraps [38], however not for architectures of multiple ions in a single Paul trap [36]. Furthermore, designing fast gates for even the simplest two-dimensional array requires optimisation over a 2×2 ion system, and so it becomes necessary to develop an optimisation scheme well suited for large ion crystals.

In Section 5.1 of this chapter, I will present an efficient gate scheme that builds on the General Pulse Group (GPG) scheme introduced in the previous chapter, which I will call the Antisymmetric Pulse Group (APG) scheme. In this scheme, optimisation is done over the number of pulse pairs in each pulse group on a grid of regularly-spaced timings. Additionally, an anti-symmetry is imposed in the sign of the momentum kicks from the pulse pairs, further restricting the solution space. I will present comparisons of optimisations of this scheme to the FRAG scheme for two ions in a) a linear microtrap array in Section 5.2.3, and b) a single linear Paul trap in Section 5.2.4. Further, in Section 5.3 I will show that APG gates between ions in a linear trap scales more favourably as more ions are added to the trap as compared with gates optimised under the FRAG and GPG schemes.

5.1 Antisymmetric Gate Scheme

I propose a modification to the GPG scheme, which imposes anti-symmetrisation of the **Z** vector. I will refer to this as the Antisymmetric Pulse Group (APG) scheme. The form of the scheme is

$$\mathbf{Z} = \{-z_{N/2}, \ldots, -z_2, -z_1, z_1, z_2 \ldots, z_{N/2}\}, \tag{5.1}$$

$$\mathbf{t} = \frac{\tau_G}{2}\{-1, -\frac{N-2}{N}, -\frac{N-4}{N}, \ldots, -\frac{2}{N}, \frac{2}{N}, \ldots, \frac{N-4}{N}, \frac{N-2}{N}, 1\}. \tag{5.2}$$

For some fixed gate time τ_G, gate optimisation under this scheme is then over the elements $\{z_1, \ldots, z_{N/2}\}$, i.e. over a $\frac{N}{2}$-dimensional search space. This reduction of dimensionality allows for gates with more pulse groups to be tractably optimised than the GPG allows for. However, the imposition of anti-symmetry is a restriction on the more general GPG scheme. Formally, this implies that given enough computational time, the GPG scheme is likely to find higher-fidelity solutions in general. However, as the number of ions is increased, there are an increased number of motional modes that need to be restored, i.e. more conditions (*c.f.* equation (3.36)) that need to be satisfied. As a result, the cost function gets increasingly more complicated, and the general N-dimensional space of the GPG scheme requires increasingly high density of sampling to find these higher-fidelity solutions. Furthermore, the anti-symmetry of the APG scheme guarantees velocity restoration of each motional mode, which results in a simpler cost function, and so is more likely to find optimal gate solutions in regimes where the cost function of the GPG scheme is difficult to efficiently search.

5.2 Comparison of gate schemes for two ions

For bench-marking the APG gate scheme, it is useful to compare it to the current best gate schemes, FRAG and GPG. For brevity, I will adopt the shorthand:

[Scheme Name](Number of Pulse Groups)

for describing the APG and GPG gate schemes. For example, a GPG scheme with $N = 12$ pulse groups will be denoted as GPG(10).

5.2.1 Metrics for comparing gates

Gates will be compared on the basis of three things: their infidelity, minimum resolving repetition rate, and time, all of which are preferred to be as small as possible. In practise, high-fidelity gates that require low repetition rates to resolve generally come at the cost of increasing the gate time, and the converse is also true [36]. Thus fidelity will be the main metric for comparing gates, and will be presented across ranges of gate time and minimum resolving repetition rates. For bench-marking, I will compare infidelities to the commonly cited 'fault tolerance threshold' of $1 - F = 2 \times 10^{-4}$ [83, 84], which can be understood as a indicative required error rate in order to implement fault-tolerant error correction using a Bacon-Shor code with a depth of 10 [6, 85].

5.2.2 Numerical optimisation considerations

For consistency with the literature, particularly with regards to Ref. [36], I will consider the 393 nm $S_{1/2} \rightarrow P_{3/2}$ transition in $^{+}\mathrm{Ca}^{40}$ as the one used for the state-dependent momentum kicks (this is discussed in Chapter 7). I will also assume a constant average occupation of 0.1 for each of the motional modes, which is reasonable for current ion trapping experiments [86]. The parameters corresponding to these choices are:

$$\omega_t = 2\pi \times 1.2 \text{ MHz (angular units)} \tag{5.3}$$

$$\eta = 0.16 \tag{5.4}$$

$$\bar{n}_1 = \bar{n}_2 = 0.1\,. \tag{5.5}$$

The results below will largely be reported in dimensionless units, i.e. the gate time will be reported in trap periods and minimum resolving repetition rates will be reported in trap frequency units of $\nu_t = \omega_t/2\pi$. Furthermore, I will assume that the time-delay between the counter-propagating pulse pairs is not limited by laser repetition rate, corresponding to a pulse splitting set-up (discussed in Chapter 7). The minimum resolving repetition rates of gates that I will report are calculated by:

$$f_{min} = \max_k \left(\frac{|z_k| + |z_{k+1}|}{2\,\delta t_k} \right) \tag{5.6}$$

where δt_k is the time between the k-th and $k+1$-th pulse groups.

Following Ref. [38], the effect of the crystal geometry on a two-ion system can be

characterised by the dimensionless parameter

$$\chi = \frac{\omega_b - \omega_t}{\omega_t}\,, \tag{5.7}$$

where ω_t and ω_b are the frequencies of the common-motional and breathing modes, respectively. The linear microtrap array I will consider in Section 5.2.3 is defined by $\chi = 1.8 \times 10^{-4}$ and the linear Paul trap I will consider in Section 5.2.4 is defined by $\chi = \sqrt{3} - 1$. In the case of the microtrap array, the value $\chi = 1.8 \times 10^{-4}$ corresponds to an intertrap distance of ∼90 microns for $\omega_t = 2\pi \times 1.2$ MHz.

I use the Limited-Memory Broyden Fletcher Goldfarb Shanno (LBFGS) algorithm for numerical optimisation of the cost-function (3.40) for each of the gate schemes in this section. This is a typical algorithm used for computationally intensive global minima searches, and combines techniques of gradient descent with stochastic sampling of the search space, to avoid being trapped in local minima [87]. The highest fidelity gate is selected from a collection of local-gradient searches for restricted boundaries, with respect to the parameters being searched over (pulse timings for FRAG, number of pulse pairs per group for GPG and APG). These boundaries are then expanded by a small amount, within which again minima are found. This process is iteratively repeated until the boundaries hit a specified maximum. It is important to note that this search is over a continuous solution-space, and so for the APG and GPG schemes, I implement an extra step to the algorithm where found minima are snapped to integer values of z_k. This assumes that the integer-snapped solutions are close to the minima found in the algorithm, which does pose the risk of missing fine detail in the search space. The C++ code used for these calculations is adapted (with permission) from code originally developed by Alexander Ratcliffe at the ANU.

In this chapter, the results of all APG and GPG optimisations presented were calculated with search densities chosen for comparable computational times, and the maximum number of pulse pairs in any one pulse group is capped at 200 in all optimisations. For consistency, FRAG searches are capped at $n_k = 100$, which is equivalent to placing a cap of 200 pulse pairs in any single pulse group. I will present results for particular gate times, which correspond to fixed parameters in the GPG and APG schemes, and maximum boundaries in FRAG optimisations.

5.2.3 Microtrap results

In Figure 5.1, I present the comparison between gate optimisations of two ions in a microtrap array chatacterised by $\chi = 1.8 \times 10^{-4}$ for FRAG, GPG(10), APG(10), and

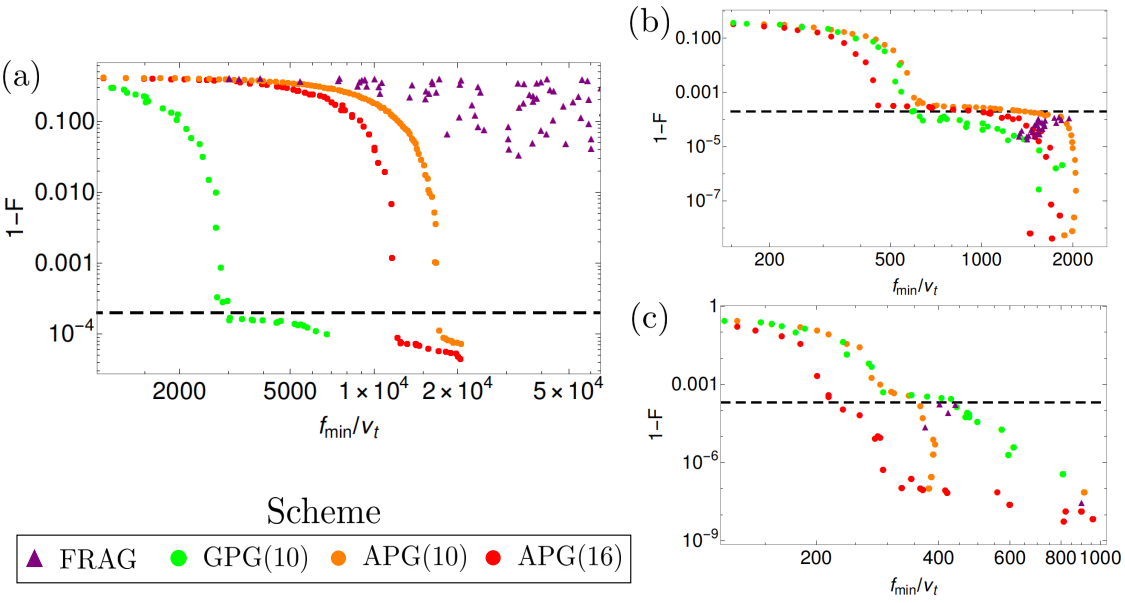

Figure 5.1: Infidelity is plotted as a function of minimum resolving repetition rate for optimisation of the FRAG (purple, triangle), APG with 16 pulse groups (red), APG with 10 pulse groups (orange), and GPG (green) with 10 pulse groups. Gates are optimised for a gate time of (a) 0.45 trap periods, (b) 1.0 trap periods, and (c) 1.75 trap periods. For the FRAG scheme, these caps are implemented by direct restrictions on the pulse group timings, and thus the FRAG gates may be faster than the compared APG and GPG gates. The black dashed line shows the fault tolerant threshold of 2×10^{-4}, corresponding to an fidelity of 99.98%. For each gate time, the APG and GPG schemes clearly find solutions with lower infidelity than the FRAG scheme, for given repetition rates.

APG(16) schemes. Gates were optimised for short, medium, and long gate times of 0.45, 1.0, and 1.75 trap periods, respectively.[1]

In all cases APG and GPG optimised gates clearly outperform FRAG gates, achieving lower infidelities for lower resolving repetition rates across all three gate times. For the sub trap-period gate time, the FRAG optimisation was not able to find any high-fidelity gates at all within the constraints placed on number of pulses. Further, the FRAG data is more scattered across each of the plots, indicating that the solution space was not adequately searched in the computational time allowed. In comparison, almost monotonic trends are present in the APG and GPG data for similar computational time, suggesting the solution space was more effectively sampled.

The comparison between the APG and GPG gate schemes is more interesting. Even

[1]Note that I use 'long' only with respect to the trap period, 1.75 trap periods corresponds to a gate time of 1.5μs. This is on the limit of what has been achieved in state-of-the-art experiments that use different gate schemes and significantly larger infidelities of about 10^{-3} [26].

with the larger number of pulse groups, gates designed under the APG(16) scheme fall short of the GPG(10) scheme for the short gate time. While the APG(16) optimisations found lower infidelity solutions overall, gates optimised under the GPG(10) scheme pass beyond the fault tolerant threshold for almost an order-of-magnitude lower resolving repetition rates, which is likely due to the gate time being close to a half-integer multiple of the trap period. Anti-symmetric gates are unfavourable at (half-)integer multiples of the trap frequency where motional restoration is difficult, as the gate is operating on very similar timescales to common-motional oscillations.[2] This effect can also be seen for $\tau_G = 1.0$ trap period, with the GPG scheme outperforming both APG schemes in a small region between $f_{min} = 600\nu_t$ and $f_{min} = 1500\nu_t$, and the APG(16) scheme being optimal in all other regions. For $\tau_G = 1.75$ trap periods, both APG schemes clearly outperform the GPG scheme, achieving lower infidelities for lower resolving repetition rates.

5.2.4 Linear Paul trap results

Similar optimisations were done for fast gates between two ions in a single linear Paul trap; for short, medium, and long gate times of 0.05, 0.15, 0.55 trap periods, respectively. This is shown in Figure 5.2. Note that these gate times are significantly shorter than the ones fixed for the microtrap study, as it has been shown that linear Paul traps lend themselves to faster gate times [36]. This can be understood in terms of the larger separation in the frequencies of the common-motional and breathing modes, which leads to more rapid accumulation of phase, and thus shorter time required to achieve the phase of $\pi/4$ required for the controlled-phase gate. This can be seen by observing the dependence of the mode frequency ratios in Equation (3.35).

For the ultra-short gate time of 0.05 trap periods, only the GPG and APG schemes were able to find high-fidelity solutions within the boundaries allowed, similar to the microtrap case. However for the other two gate times, FRAG gates are able to achieve significantly lower infidelities of $\sim 10^{-12}$ whereas the APG and GPG gates struggle to fall below 10^{-8}. This suggests the solution space of the linear Paul trap architecture has lots of 'structure' in terms of the pulse timing variables that is obscured by the enforcement of regular pulse timings in the APG and GPG schemes. This is unsurprising as the fast phase accumulation due to large mode separation, which results in 'sharper' minima in the solution space. This is an issue for the

[2] Motional restoration of the common-motional mode at a trap period is guaranteed by default if the momentum kicks are spaced out by (half-)integer multiples of the trap period, however this is generally not the case because of the competing needs for phase restoration.

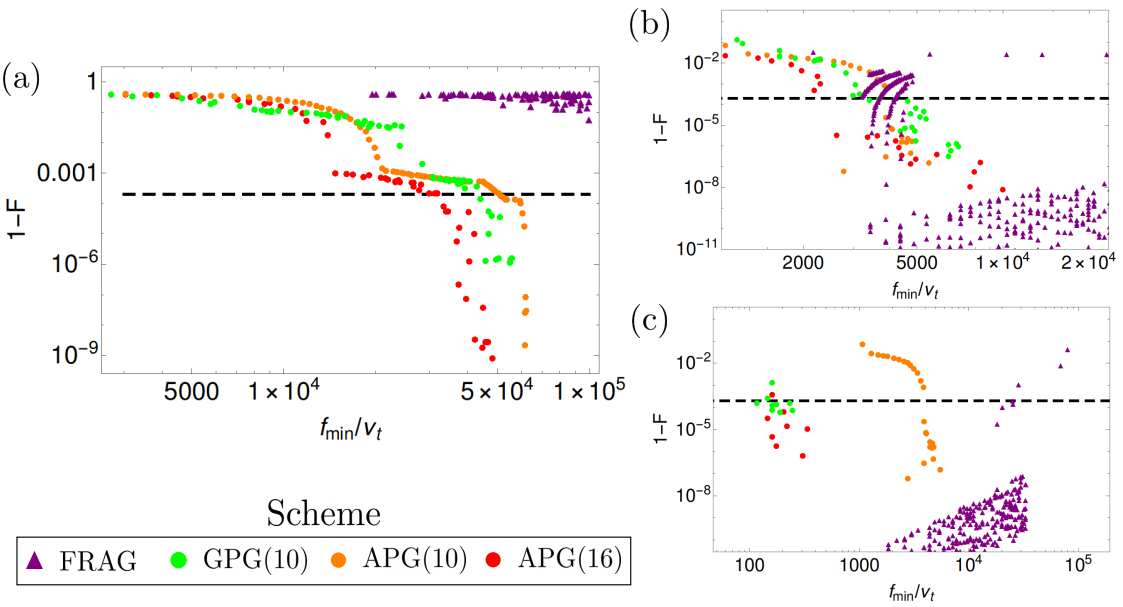

Figure 5.2: Infidelity is plotted as a function of minimum resolving repetition rate for optimisation of the FRAG (purple), APG(16) (red), APG(10) (orange), and GPG (green) with 10 pulse groups. Gates are optimised for a gate time of (a) 0.05 trap periods, (b) 0.15 trap periods, and (c) 0.55 trap periods. Across all gate times, the APG and GPG schemes are clearly optimal, as they find higher-fidelity solutions for lower repetition rates than FRAG gates.

integer-snapping step in the APG and GPG optimisations, which assume integerised solutions are very close to the non-integer numerical solutions found. Regardless, the APG and GPG gates are able to achieve infidelities below the 2×10^{-4} fault tolerant threshold for dramatically lower repetition rates.

The gates optimised under the APG(16) scheme clearly have lower repetition rate requirements for high fidelity gates for each of the three gate times, which arguably is the most important aspect of gate comparison for experimental realisation. Interestingly, for the gate time of 0.55 trap periods, there is a significant difference between the general and anti-symmetric searches for gates with $N = 10$ pulse groups. This suggests that the optimal solution for a ten pulse group gate is not anti-symmetric for gate times near half a trap period in linear Paul trap architectures. This is unsurprising and can be attributed to difficulties with motional restoration at half-integer and integer-multiples of the trap period, as described earlier. The results for $\tau_G = 0.15$ suggest that a 180 nano-second gate with fidelity above $1 - 10^{-6}$ can be implemented with a laser repetition rate of ∼3GHz. This is significantly faster and higher fidelity than any experimentally realised two-qubit gate on trapped ions, and requires an order of magnitude lower repetition rate to resolve than previously reported by studies of the FRAG scheme [36].

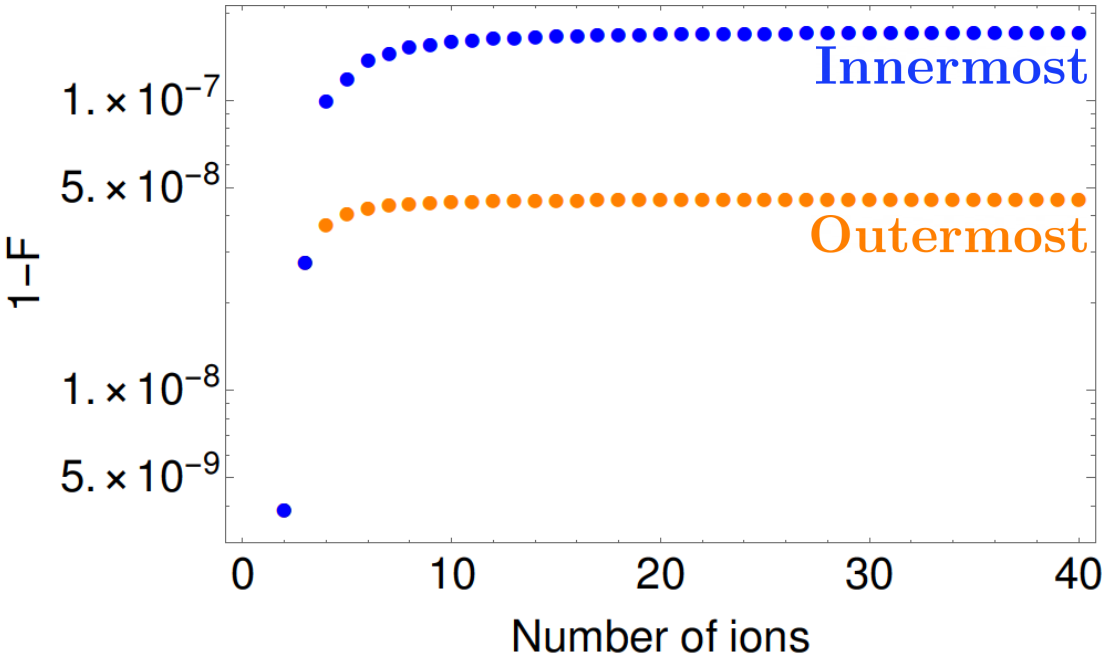

Figure 5.3: Scaling of a fast gate between two ions in a linear microtrap array as the number of ions is increased, between the innermost (blue) and outermost (orange) pairs of ions. The infidelity presented is for a GPG(10) gate optimised for a two-ion microtrap array with $\chi = 1.8 \times 10^{-4}$, that has a gate time of $\tau_G = 1.2$ trap periods and a minimum resolving repetition rate of $f_{min} = 1$GHz. For large numbers of ions, the infidelity does not increase indefinitely and plateaus to $\sim 10^{-7}$.

5.3 Ion scaling

The use of fast-gates in any large-scale computation with trapped ions requires an understanding of how the infidelity and speed of gates scale with the number of ions. In this section, I will report the scaling behaviour of fast gates as the number of ions in a one-dimensional chain is increased. This will be done for both microtrap and linear Paul trap architectures.

5.3.1 Microtrap scaling

While scaling analysis has been done recently for microtraps by Ratcliffe *et al.* [38], I will reproduce the results here as it will serve as a useful comparison to scaling in a Paul trap. The approach taken by Ratcliffe *et al.* was to study the performance of a gate optimised for two ions in a system with more ions. This is a worst-case option that is computationally inexpensive. The robustness of this approach can be motivated in part by the difference between the mode frequencies being very small: the biggest mode splitting is between the common-motional and breathing

modes, which is fixed to $\chi = 1.8 \times 10^{-4}$ for the system I have been studying in this chapter. This implies that motional restoration of a single motional mode results in the almost complete restoration of each of the other motional modes as well. While this approach is unlikely to find the highest-fidelity gates for a scaled ion system, it is certainly more computationally tractable to only optimise for two-ion systems. Furthermore it is experimentally favourable to not have to change pulse schemes as ions are added or removed.

The fidelity of the two-ion gate is calculated using the truncated infidelity (3.40), evaluated individually for each system with varying number of ions in a one-dimensional array of microtraps each separated by a distance $d = 100\,\mu$m. This includes the motional modes of all of the ions, which are calculated by linearisation of the potential of the system, as described in Appendix B. The results are plotted in Figure 5.3, from which it can be seen that the fidelity of the gate initially decays as more ions are added, but not indefinitely. For more than $\sim$10 ions in the chain the fidelity plateaus to a constant value. The plateauing behaviour can be understood as the additional motional modes added from added ions contributing less to the dynamics of the gate in increasingly large ion chains. This value depends on which two ions are involved in the gate, with an order of magnitude difference in fidelity between the outermost and innermost ions.

5.3.2 Linear Paul trap scaling

Unfortunately, the same approach fails for studying scaling of linear ion chains in a single Paul trap. This is because the splitting between mode frequencies is much larger than the microtrap case, and so the extra motional modes from increasing the number of ions are non-trivial to restore. Bentley *et al.* [36] found that this can be avoided by increasing the speed of the gates so they only excite local motion of ions, but the repetition rates required were prohibitively high. Thus it becomes necessary to optimise gates individually for each number of ions in the trap, which is the approach I will take here. I will, however, enforce a capped gate time in each of these optimisations and restrict this analysis to gate infidelity, as gate speed will not change as the ion crystal is scaled.

In Figure 5.4, I have plotted the infidelities gates optimised for different numbers of ions in a Paul trap with fixed longitudinal frequency under the FRAG, APG(16), and GPG(10) schemes. The reported infidelities are for gates that were optimised under each scheme for comparable computational time. Gate times were capped at half a trap period, which for the FRAG scheme this corresponds to the maximum

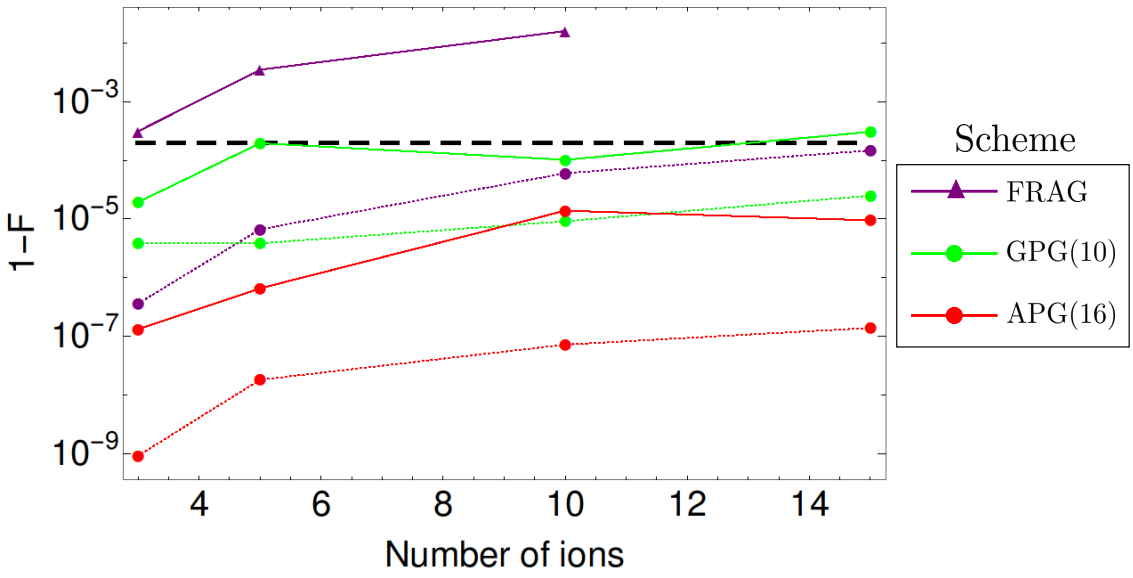

Figure 5.4: Infidelities of FRAG, APG, and GPG gates for increasing numbers of ions in a single Paul trap is increased, for fixed axial frequency of $\omega_t = 2\pi \times 1.2$ MHz. Gates were optimised for each multi-ion system with gate times capped at $\tau_G = 0.6$ trap periods. The infidelities reported here are for minimum resolving repetition rates of 1 GHz (solid) and 10 GHz (dashed), respectively. The FRAG optimisation for 15 ions was not able to find gates that are resolvable by 1 GHz repetition rate.

boundary placed on the search. In the APG and GPG schemes, optimisation was done over a range of fixed gate times lower than 0.6 of a trap period, and the best-performing gate was selected. Just as with the rest of the gate optimisations presented in this chapter, finite repetition rate was not enforced in the search but calculated afterwards in terms of minimum resolving repetition rate from Equation (5.6). For more than 15 ions, optimisation was found to be intractable within the scope of this thesis: each additional ion adds an extra condition for motional restoration and thus the cost function becomes increasingly complex.

There appear to be similar trends across all three gate schemes as more ions are added to the trap, with infidelity growing. Higher repetition rates are able to resolve higher-fidelity solutions across all three gate schemes, which is likely due to increase gate speeds allowable with 10 GHz as opposed to 1 GHz. As compared to the results shown in Figure 5.2, there seems to be a sudden jump in the infidelity: for example FRAG was able to achieve extremely low-infidelities of 10^{-12} for a two-ion system with a 1 GHz repetition rate laser - for even three ions in a trap this drops to 10^{-4}. This drop appears to be less dramatic for the APG and GPG schemes, which for given laser repetition rates are able to achieve orders of magnitude better fidelities. This can likely be attributed to the fact that the FRAG scheme has only three free parameters that are optimised over, which is well suited to two-ion systems, where

only three conditions need to be satisfied (one phase condition, and restoration of two motional modes). As more ions are added, the increased motional modes result in more restoration conditions that need to be satisfied, and thus the FRAG scheme is no longer an exact solution to the equations (3.35) and (3.36). In contrast, the APG and GPG schemes have an increased number of free variables (in this case, eight and ten, respectively) and thus have more freedom to find solutions that satisfy all motional conditions.

There is also a significant difference between the APG(16) and GPG(10) schemes in this scaling plot. While the previous sections saw the GPG achieving comparable (and occasionally better) fidelities for given repetition rates, Figure 5.4 tells a different story. The fidelities of gates optimised under the APG(16) scheme appear to be between one to two orders of magnitude better than those optimised under the GPG(10) scheme, for a given resolving repetition rate. This validates the initial motivation of developing an anti-symmetric scheme that optimises over z_k elements: it combines the anti-symmetry of FRAG which simplifies motional restoration, and the increased number of free variables in the GPG scheme. As mentioned previously, as the GPG scheme is a generalised form of the APG scheme, it should always be able to achieve better fidelities if sampled sufficiently densely. This suggests that the cost-function that includes the motional modes of three or more ions, is unable to be sampled densely in GPG optimisation for the computational time allowed in this analysis. Therefore the APG scheme is very well suited to finding high-fidelity solutions to multi-ion systems with finite computational resources. One may note that unlike the FRAG trends, the infidelity of the APG and GPG gates do not monotonically grow as a function of ion number, with some 'kinks' in the trends around 10 ions. This likely is not due to any physical phenomena, rather corresponding to the integerisation of the optimised solutions (described in Section 5.2.2). If each gate was locally optimised for its pulse timings, these kinks would likely disappear. It is also possible these kinks are sign of an under-sampled search space, and may disappear for higher density optimisations.

5.4 Conclusions

The analysis presented in this chapter lays the groundwork for optimising fast gates in scaled ion crystals for large scale computation. I have discussed a fast gate optimisation scheme that is a variant of the Generalised Pulse Group (GPG) scheme with the imposition of anti-symmetry, which I have dubbed the Antisymmetric Pulse Group (APG) scheme. The APG scheme has been compared against the GPG

scheme and the FRAG scheme, which is the current optimal gate scheme in fast-gate literature, across a range of gate times and repetition rates. For almost all gate optimisations in a simple two-ion system, I have found the APG scheme is optimal, due to the simplification of the cost function from the anti-symmetry constraint. Notable exceptions are in microtrap architectures, where the GPG scheme is optimal for gate times that are integer or half-integer multiples of the trap period. Further, for increasing numbers of ions in a single linear Paul trap the APG scheme has been shown to achieve orders of magnitude better fidelities for fixed resolving repetition rates than both the FRAG and GPG scheme. This suggests that the APG scheme is well suited for optimising gates in trapped ion architectures where motional restoration of multiple modes is non-trivial. As such, the APG scheme will be used in the following chapter to optimise gates for a two-dimensional microtrap architectures, which has a mode structure significantly more complicated than the one-dimensional cases studied here.

Chapter 6

Fast gates in two-dimensional symmetric ion crystals

In Chapter 4, I showed that the increased connectivity of ions in two-dimensions allowed for significant reduction in the number of nearest-neighbour fast gates required to implement non-local entanglement operations. While two-dimensional microtrap arrays have been experimentally realised [55], there has been little analysis of fast gates in these architectures. Even in the simplest cases, the mode structures of two-dimensional architectures are sufficiently different to their one-dimensional counterparts that it is not clear whether fast gates can be implemented with high-fidelity, if at all. Furthermore, it is unclear whether it is preferable to implement fast gates directly between diagonally separated ions in a two-dimensional lattice, or connect them via a series of nearest-neighbour fast gates.

In this chapter I will investigate fast gates in square-symmetric ion crystals. To be precise, square-symmetric refers to an $N \times N$ array of ions in a square arrangement with respect to a set of orthogonal axes. This symmetry allows for a non-dimensional treatment for the simplest case of a 2×2 arrangement of ions, which I will describe in Section 6.1. I will present the results of gate searches for fast gates between nearest-neighbour ions in Section 6.2.2 for variations of this parameter. I will compare fast gates between diagonally separated ions, and compare to equivalent operations constructed out of (adjacent) nearest-neighbour gates. Finally, in Section 6.4, I will consider the performance of these nearest-neighbour gates as the size of the two-dimensional crystal is scaled.

6.1 2D microtrap architecture

The physical architecture I will consider is the two-dimensional generalisation of the linear microtrap array studied in Chapter 5, where ions are individually trapped

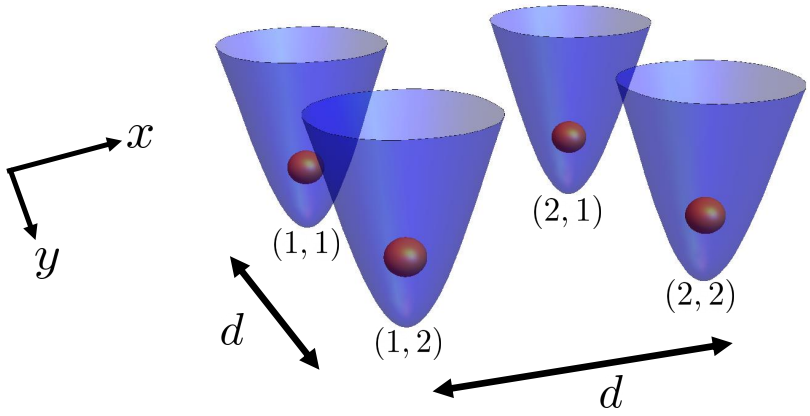

Figure 6.1: Diagram of the 2×2 symmetric ion crystal, where ions are represented as (red) spheres in microtrap potentials (blue). The indexing of the ions corresponds to column and row, respectively. The vertical (y) and horizontal (x) axes correspond to column and row, respectively. The crystal is symmetric, in that the distance d between neighbouring traps along each column and each row is the same, as indicated.

in symmetric microtraps, on a regular grid. In this geometry, the distance between neighbouring traps along the vertical and horizontal axes are equal, i.e. $d_x = d_y = d$, and the individual microtraps are symmetric in their radial and longitudinal trapping frequencies, i.e. $\omega_x = \omega_y = \omega_t$. This is represented for the simple 2×2 crystal in Figure 6.1

6.1.1 Non-dimensionalisation of mode structure

Just as Ratcliffe *et al.* [38] presented a non-dimensional analysis of fast gate schemes for two ions in a linear array, the symmetry of this two-dimensional geometry allows for an analogous non-dimensional treatment. By inspection of the condition equations (3.35) and (3.36), the ideal implementation of a fast gate is the function of the frequencies of the collective motional modes, which are defined completely by the trap geometry.

In analogy to the one-dimensional case, the effect of the trap geometry on the dynamics of a fast gate should be able to be captured by the dimensionless parameter $\xi_0 = \frac{\alpha}{d^3\omega_t^2}$, where $\alpha = \frac{1}{M}\frac{e^2}{4\pi\epsilon_0}$. Naively, one might expect the dynamics to be dominated by the common-motional and breathing modes. The smallest possible crystal in two-dimensions is a four-ion system in a 2×2 arrangement, and so I will begin by modelling this smallest sub-unit of a 2D ion crystal. In contrast with the analogous one-dimensional case, there are extra motional modes present that are orthogonal to both the common-motional and breathing modes. Furthermore, the common-

motional and breathing modes are degenerate, and thus not unique. As such, some care needs to be taken in finding an appropriate dimensionless parameter that is a good predictor of the system dynamics.

In order to circumvent these subtleties, I take a brute force approach by analytically computing the Hessian matrix of the system, pulling out common factors, and forcing them to be non-dimensional by factors of ω_t and d. In this manner, I identify that the suitable dimensionless parameter is the normalised difference between the *squared* oscillation frequencies of the common-motional and breathing modes,

$$\xi = \frac{\omega_{BR}^2}{\omega_t^2} - 1\,. \tag{6.1}$$

Importantly this can be expressed in terms of the more fundamental geometry parameter ξ_0,

$$\xi = \pi^2 \xi_0 \left(\frac{\delta^{1/3} + 18 - 8\sqrt{2}\delta^{-1/3} + 2\left(\sqrt{2}-4\right)}{3\sqrt{2}\left(2\sqrt{2}-1\right)} + 1 \right)^{-3}, \tag{6.2}$$

where

$$\begin{aligned} \delta \;=\; & 12\pi\sqrt{21}\sqrt{\left(401\sqrt{2}-596\right)\xi_0 - 189\left(72\sqrt{2}-113\right)\pi^2\xi_0^2} \\ & -\xi_0 756\left(4\sqrt{2}-9\right)\pi^2 - 50\sqrt{2} + 88\,. \end{aligned} \tag{6.3}$$

In terms of this dimensionless parameter, the frequencies of the motional modes of the four-ion system can be written as

$$\omega_p^2 = \omega_t^2 \left\{ 1, 1, \xi+1, \xi+1, -\frac{\xi}{2\sqrt{2}} - \xi + 1, -\frac{\xi}{2\sqrt{2}} + 2\xi + 1, \frac{\xi}{\sqrt{2}} - \xi + 1, \frac{\xi}{\sqrt{2}} + 2\xi + 1 \right\}, \tag{6.4}$$

where the first two elements correspond to the two common-motional modes, and the second two correspond to the breathing modes. Visualisation of the corresponding mode coupling vectors $\mathbf{b_p}$ are shown in Figure 6.2.

6.2 Nearest-neighbour gate searches

This mode structure is significantly different to the motional spectrum of one-dimensional chains of ions, and thus the generalisation of gate optimisation is entirely non-trivial. In this Section, I will present results of gate optimisation based on the 2×2 ion system. An additional complexity in two-dimensions is the possibility of

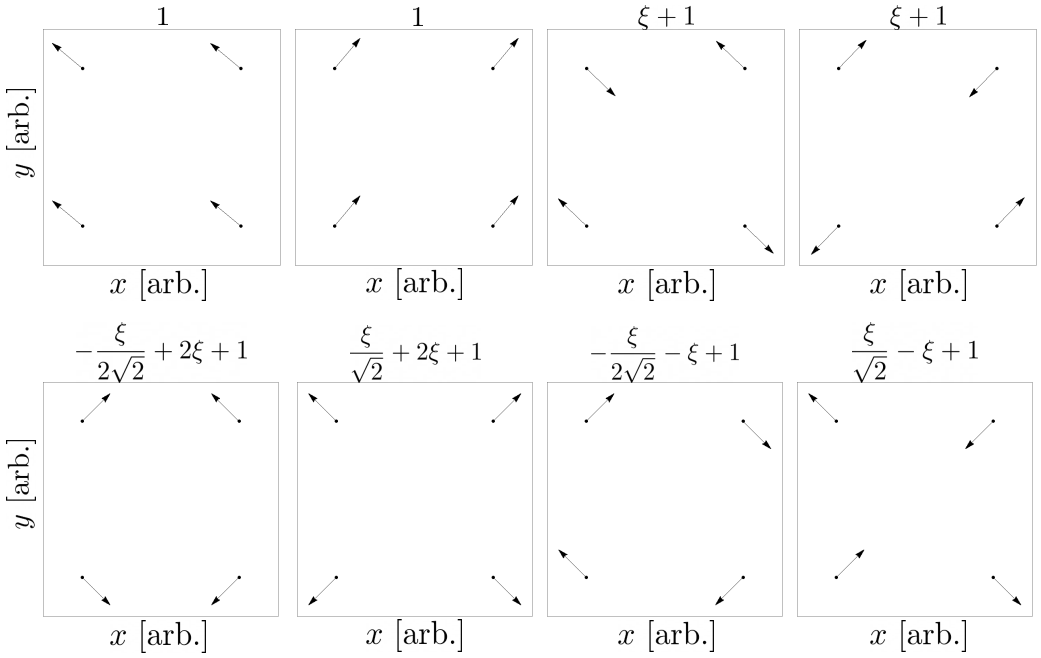

Figure 6.2: Visualisation of the couplings of the motional modes to the ions in a 2×2 symmetric crystal. The displacements in the x and y directions are represented in arbitrary units as the coupling vectors $\mathbf{b_p}$ are normalised. The corresponding squared frequency of each mode is annotated, ω_p^2/ω_t^2.

performing gates across the diagonal of the crystal, which I will study in the next section. To distinguish between these two possibilities, I will refer to gates along the horizontal and the vertical (i.e. non-diagonal gates) as *adjacent*. The symmetry of this considered architecture is implicit in this nomenclature, as gates along the horizontal axis and gates along the vertical axes are treated as equivalent.

6.2.1 Cost-function adjustments

Thus far in this thesis, I have only considered gates where the momentum kicks from the counter-propagating π-pulses are administered along a single trapping axis, which is implicit in the cost-function expression (3.40). This is too restrictive in two-dimensions, where it may be preferable to have these kicks in more general directions. This will be useful when I investigate diagonal gates. In order to make this generalisation, I will introduce a vector $\mathbf{K}$ to describe the *kick direction*. This is a normalised vector indexed by the x, y co-ordinates of the i-th ion in the j-th row of the crystal lattice, i.e.

$$\mathbf{x}_{i,j} = \{x_{1,1}, y_{1,1}, x_{1,2}, y_{1,2}, x_{2,1}, y_{2,1}, x_{2,2}, y_{2,2}\}. \tag{6.5}$$

These coordinates correspond to the (i, j) indexing of the ion positions, as shown in Figure 6.1.

For example, the vector

$$\mathbf{K} = (0, 1, 0, 0, 0, 0, 0, 0)^T \tag{6.6}$$

characterises momentum kicks of the top left ion along the y-axis. Using this definition, the truncated infidelity expression (3.40) for a fast gate between two ions, A and B, can be written in a more general form,

$$1 - F \approx \frac{2}{3}|\Delta\phi|^2 + \frac{4}{3}\sum_p (\frac{1}{2} + \bar{n}_p)\left((\mathbf{K_A}\cdot\mathbf{b_p})^2 + (\mathbf{K_B}\cdot\mathbf{b_p})^2\right)|\Delta P_p|^2 \,. \tag{6.7}$$

Furthermore, the expressions for the phase mismatch, $\Delta\phi$, and motional restoration, ΔP_p, can be written in this more general form as well:

$$|\Delta\phi| = \left|8\eta^2\frac{\omega_t}{\omega_p}(\mathbf{K_A}\cdot\mathbf{b_p})(\mathbf{K_B}\cdot\mathbf{b_p})\sum_{i\neq j} z_i z_j \sin\left(\omega_p|t_i - t_j|\right)\right| - \frac{\pi}{4} \tag{6.8}$$

$$|\Delta P_p| = 2\eta\sqrt{\frac{\omega_t}{\omega_p}}\sqrt{\left(\sum_i z_i \cos\left(\omega_p t_i\right)\right)^2 + \left(\sum_i z_i \sin\left(\omega_p t_i\right)\right)^2} \,. \tag{6.9}$$

By inspection of (6.8) and (6.9) it can be noted that the choice of sign of $\mathbf{K}$ is arbitrary.

This truncated expression will be used as the cost-function for optimising gates in this chapter, in analog to the one-dimensional gate searches in Chapter 5. Parameter values will be kept the same as in Chapter 5, unless otherwise stated. It is important to note that this cost-function is more complicated for two-dimensional gate searches due to the presence of extra modes that require restoration. For this reason, I choose to use the APG(16) gate scheme for optimisation as by the anti-symmetric nature of this gate scheme, momentum restoration of each motional mode is guaranteed.

In the APG scheme, both the pulse mismatch $|\Delta\phi|$ and the motional restoration terms $|\Delta P_p|$ depend only on values of $z_i z_j$ and the oscillation frequencies of the motional modes, as the pulse timings t_i are fixed. This suggests that the infidelity can be described by n_{max}^2 and ξ for fixed η and gate time τ_G, where $n_{max} = \max(z_i)$.

6.2.2 Adjacent gate results

I have plotted the infidelity of gates optimised under the APG scheme as a function of $n_{max}^2\xi$ in Figure 6.3. For a given gate time, the infidelity appears monotonic as a

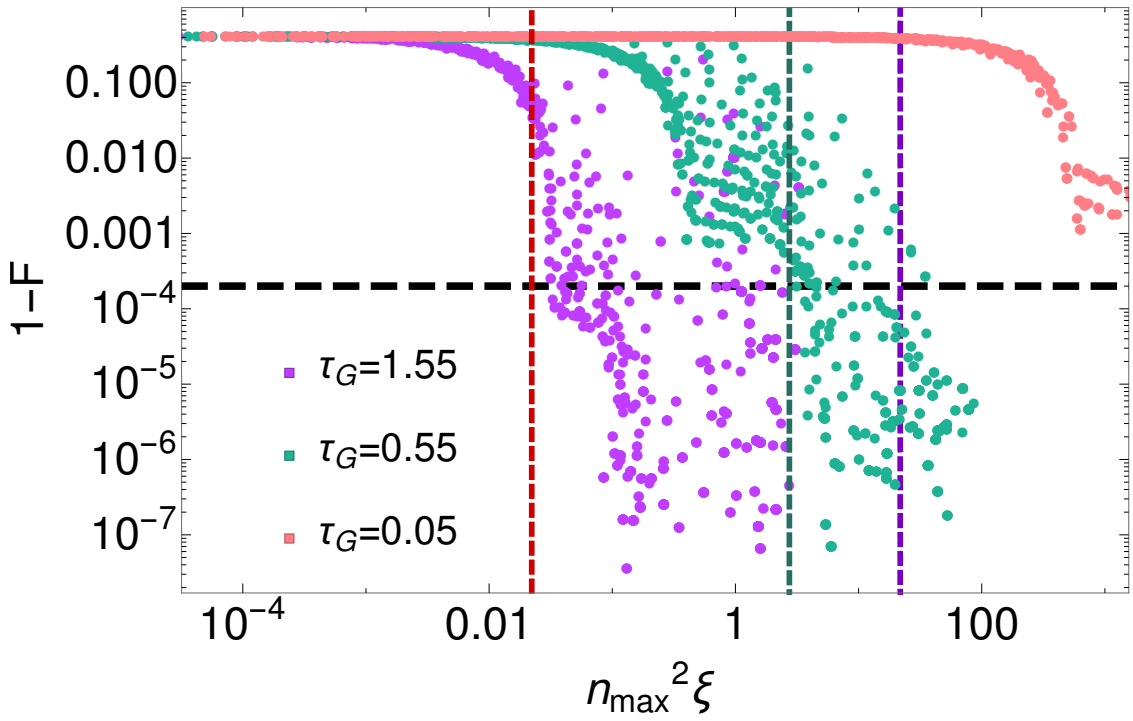

Figure 6.3: Infidelities of optimised gates are plotted as a function of $n_{max}^2\xi$, where n_{max} is the largest number of pulse pairs in a given group for a given gate, plotted for several different gate times (in trap periods). As in the previous chapter, the black dashed line marks the fault tolerant threshold of 2×10^{-4}. For each gate time, the infidelity of a gate is well described by the parameter $n_{max}^2\xi$ until it falls below approximately $1-F=10^{-2}$. The vertical lines correspond to the maximum value of $n_{max}^2\xi$ achievable by a 5 GHz repetition rate laser, and a trap geometry with $\omega_t = 2\pi\times1.2$ MHz and $d=100$ microns (corresponding to $\xi=1.7\times10^{-3}$), for each gate time. This data suggests that high-fidelity gates with operation times longer than ∼0.6 trap periods can be realised with a 1 GHz repetition rate laser in such a geometry.

function of $n_{max}^2\xi$ until it falls below $\sim10^{-2}$, after which $n_{max}^2\xi$ fails to characterise the gate infidelity. This is similar to characteristics of fast gates between a pair of microtraps in a linear array, analysed in Ref. [38], based on the analogous dimensionless parameter $\chi=\frac{\omega_{BR}}{\omega_t}-1$. This suggests that the solution-space of fast gates in two-dimensional microtrap arrays has a similar structure to its one-dimensional counterparts, at least for the simplest crystals. The data also demonstrates that a 5 GHz repetition rate is capable of resolving gates as fast as 500 nano-seconds[1] with high-fidelity in an experimentally feasible geometry with $\omega_t = 2\pi\times1.2$ MHz and trap separation of $d=100$ microns.

It is perhaps an unsurprising result that, qualitatively at least, optimisation of fast

[1]I have converted between trap periods and dimensional time for the given trapping frequency, by calculating $\frac{2\pi}{\omega_t}\tau_G$.

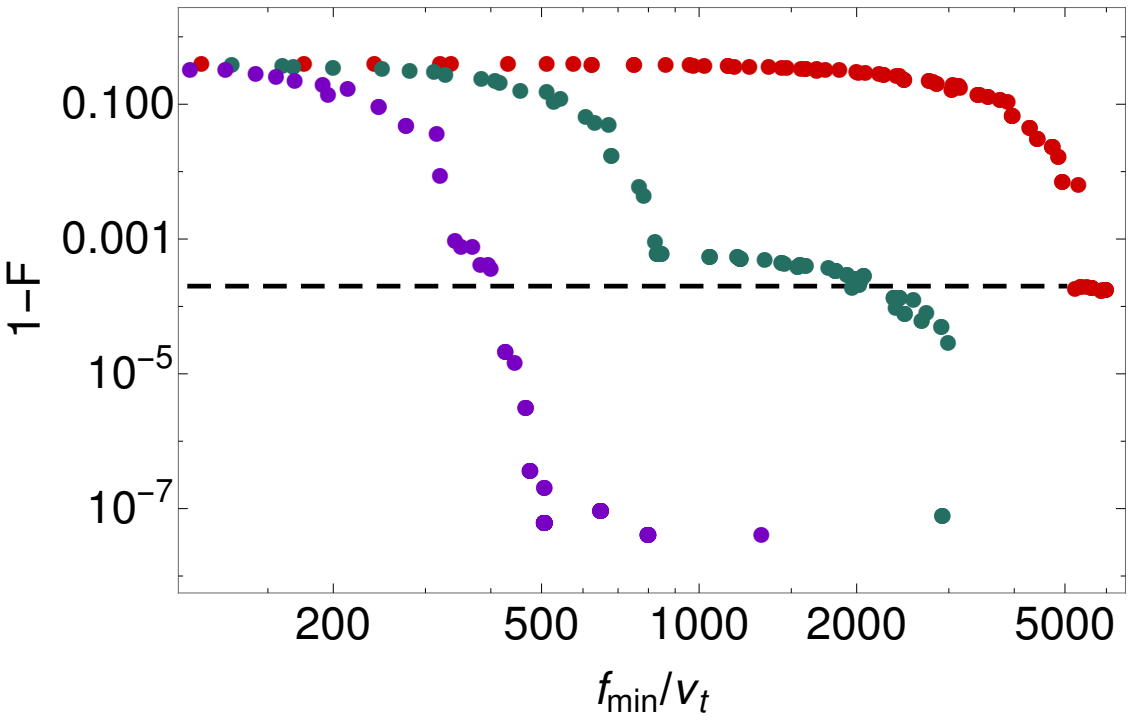

Figure 6.4: Infidelities of optimised gates are plotted against their minimum resolving repetition rate (in units of trap frequencies). These gates are optimised for $\xi = 1.7 \times 10^{-3}$, and for gate times of 0.5 (red), 1.0 (green), and 1.8 (purple) trap periods. The 2×10^{-4} fault tolerance benchmark is marked as a black dashed line. For each gate time, this threshold is achieved within the bounds placed on the search.

gates between micro-traps share similar structure in one-dimensional chains and two-dimensional arrays, as the motional dynamics can be reduced in both cases to a single dimensionless parameter. While the parameters are not exactly equivalent - χ is the (normalised) difference in the common-motional and breathing mode frequencies and ξ is the *squared* difference - the parameters play an analogous role to one another as they characterise the relative rate of phase acquisition between these two modes. This rate is, of course, also dependent on the magnitude of the momentum kicks which can be parameterised by n_{max}, and thus depends on the available laser power. Importantly, the fact that the gate infidelity scales with the square of this quantity (n_{max}^2) suggests that increasing pulse repetition rate is the most effective way to improve gate performance in a practical setting.

In Figure 6.4, I plot the results of gate optimisations in terms of the infidelity as a function of minimum resolving repetition rate, for several different gate times. This is done for a fixed value of $\xi = 1.7 \times 10^{-3}$, corresponding to an experimentally-feasible inter-trap distance of $d = 100\,\mu$m, and trapping frequency $\omega_t = 2\pi \times 1.2$ MHz. The gate infidelities are almost monotonic with respect to the repetition rate, and follow similar trends to the one-dimensional microtrap array (*c.f.* Figure 5.1).

By comparison to the one-dimensional case, the repetition rates required to achieve a given infidelity appear to be slightly higher. Furthermore, the overall achievable fidelities appear to be not as high as the one-dimensional system. This can be explained by the increased connectivity of two-dimensional crystals, which results in more motional mode excitation during the gate operation. The requirements of phase accumulation and motional restoration constitute nine conditions in the two-dimensional 2×2 crystal, as opposed to three in the one-dimensional analog, thus resulting in a more complicated solution space. This additional complexity also scales with the size of the crystal: in a general two-dimensional array there are twice the number of motional modes as ions, in contrast with one-dimensions where there are only as many motional modes as ions.

6.3 Diagonal gates

The additional dimensionality of two-dimensional architectures also allows the possibility of fast gates directly between diagonally separated ions being preferred to an equivalent sequence of adjacent gates. This is in contrast to previous notions of nearest-neighbour gates always being optimal in one-dimensional chains of ions (see Section 3.5.2). However, the poor distance scaling of fast gates (see Ref. [36]) may pose a difficultly for implementing fast gates directly between diagonally separated ions, which are a factor of $\sqrt{2}$ further from each other than adjacent neighbouring ions. Fast gates between diagonally-neighbouring ions thus require larger momentum kicks, and consequently higher repetition rate lasers, to achieve comparable fidelities to gates between adjacent nearest-neighbours. However, a more concerning aspect of diagonal gates is the motion they may excite in nearby ions, and whether diagonal gates can be designed to effectively restore all the motional modes involved is not clear.

6.3.1 Entangling diagonal qubits using adjacent gates

The question is then, is it better to construct an 'effective' gate between the diagonal ions composed of adjacent gates (gates between adjacent nearest-neighbour ions), or simply directly do a fast gate directly between them. The process for implementing an effective-diagonal gate between the $(1,1)$ and $(2,2)$ is visualised in Figure 6.5, and involves the following:

1. A SWAP operation exchanges the states of qubits $(1,1)$ and $(1,2)$. In the

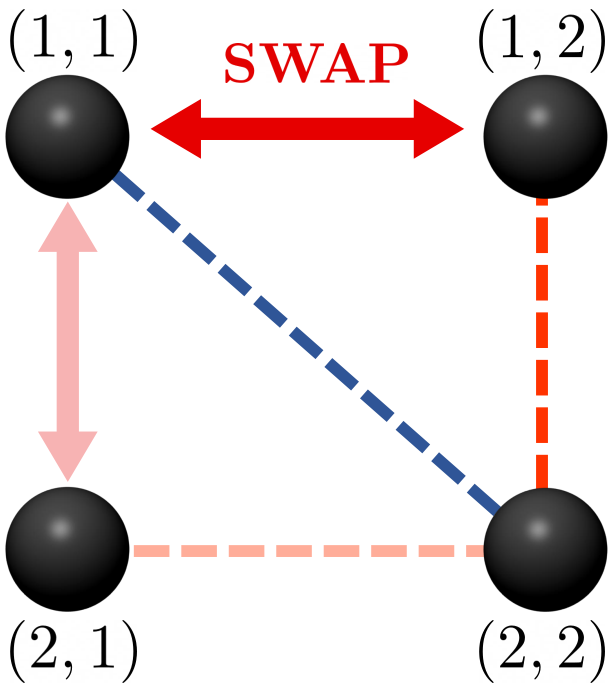

Figure 6.5: Visualisation of fast gate operations between diagonally separated ions $(1,1)$ and $(2,2)$. Fast gates are represented as dashed lines, and the swap gate is represented as a solid double-ended arrow. The direct-diagonal fast gate is shown in blue, and the equivalent non-position preserving operation constructed with adjacent gates through $(1,2)$ is shown in red. Equivalently the operation can be done through $(2,1)$, shown in light red.

computational basis this operation has the form

$$\hat{U}_{\text{SWAP}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{6.10}$$

and can be constructed with three (vertical) adjacent fast gates, and four single-qubit gates [2]. Note that the positions of the ions themselves have not changed, rather the qubits are now represented by different ions.

2. Qubits $(1,1)$ and $(2,2)$ are now (horizontally) adjacent to each other, and are entangled with a single (vertical) adjacent gate.

It is important to note that while this process implements an effective diagonal gate, it is *non-position preserving* as it does not restore the swapped qubits $(1,1)$ and $(1,2)$ to their original locations in the crystal. To do so would require an additional SWAP gate, and thus an extra three adjacent fast gates.

For comparison to direct fast gates, I will neglect the contribution of the single-

qubit gates to constructing these effective-diagonal operations. As discussed in Chapter 4, this is a very robust approximation as single-qubit gates can typically be implemented with a single laser pulse whereas fast gates require on the order of hundreds of pulses [2]. Furthermore, I will work under the approximation that the fidelities of each of the four constituent adjacent gates are uncorrelated, and can be calculated by multiplying the fidelities of each individual gate together, i.e. $F \approx (F_{adjacent})^4$. Similarly, I will calculate the total duration for the operation by summing the individual adjacent gate times, i.e. $\tau_G = 4\tau_{adjacent}$. These assumptions describe a best-case implementation of this effective-diagonal fast gate, and thus place an upper-bound on its performance.

6.3.2 Comparing direct-diagonal and effective-diagonal fast gates

The above protocol for implementing an effective-diagonal gate can be compared to a direct-diagonal gate by comparing repetition rate requirements to achieve high-fidelity. In order to present this comparison, direct-diagonal gates are optimised between ions $(1,1)$ and $(2,2)$. Note that a gate between $(1,2)$ and $(2,1)$ would be identical due to symmetry. I performed numerical optimisations for the 2×2 lattice characterised by $\xi = 1.7 \times 10^{-3}$, using the APG(16) scheme introduced in the previous chapter. The truncated infidelity expression (6.7) was used as the cost-function for this optimisation, with the kick vectors

$$\mathbf{K_{1,1}} = \frac{1}{\sqrt{2}}(1,1,0,0,0,0,0,0) \tag{6.11}$$

$$\mathbf{K_{2,2}} = \frac{1}{\sqrt{2}}(0,0,0,0,0,0,1,1)\,. \tag{6.12}$$

These kick vectors correspond to kicking the ions on opposite corners of the 2×2 array toward each other along the diagonal. This maximises the closeness of the ions during the gate and subsequently increases the strength of the inter-ion Coulomb interaction. This in turn leads to larger trajectories through phase space, and thus more rapid phase accumulation. This choice of kick direction is also motivated by the results for gates in a linear Paul trap (see 5.2.4), where I showed that the increased closeness of ions leads to faster speeds for lower repetition rate requirements.

For fair comparison of the direct-diagonal and effective-diagonal gates, I calculated the minimum repetition rates required to resolve gates with infidelities below the 2×10^{-4} fault tolerant threshold across a range of different gate times. This is shown in Figure 6.6. It is clear here that direct-diagonal gates are superior to equivalent

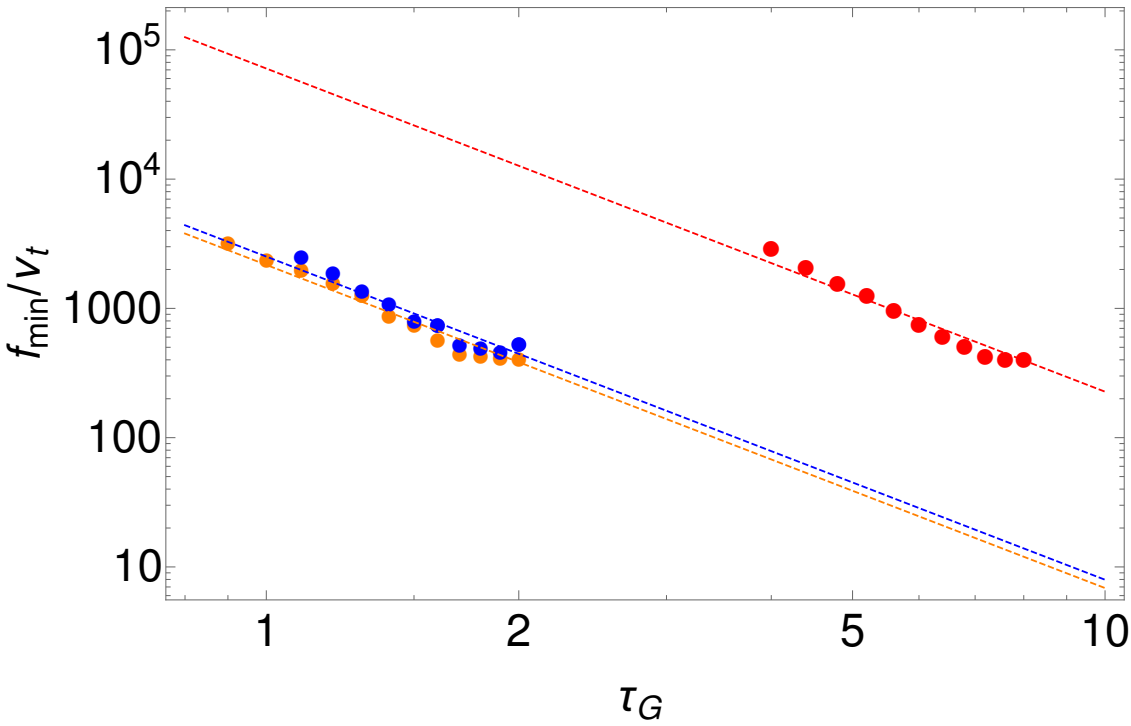

Figure 6.6: Minimum repetition rates required to resolve gates with infidelity below fault-tolerant error correction threshold of 2×10^{-4} as a function of gate time for adjacent gates (orange), diagonal gates (blue). Each data point corresponds to a gate optimised under the APG scheme with 16 pulse groups, in a 2D crystal with $\xi = 1.7 \times 10^{-3}$. Note, $\nu_t = \omega_t/2\pi$. Results for non-position preserving diagonal gates constructed with adjacent gates (red) are also presented, which were calculated as having the same required repetition rate of adjacent gates and the combined gate time for four adjacent gates (a SWAP gate, and a single entangling gate). All data points are fitted to the trend $f_{min}/\nu_t = a\ \tau_G^{-2.5}$ with $a = 2170.42$, 2516.75, and 71712.2 for the adjacent, diagonal, and effective-SWAP gate data, respectively.

operations constructed with adjacent gates. For a given repetition rate, high-fidelity diagonal gates can be resolved that are roughly five times faster than their effective-diagonal counterparts.

The data is fitted to the trend $\tau_G = a(f_{min}/\nu_t)^{-2/5}$, which has been shown to be the optimal scaling of gate times for FRAG gates [36]. Data points clearly fall on the trend line except for half-integer and integer multiples of the trap period, which is an artifact of the APG scheme being anti-symmetric. This suggests that the APG scheme generally also follows the scaling behaviour:

$$\tau_G \propto f_{min}^{-2/5}\,. \tag{6.13}$$

The values of the fitted parameter a suggest that for a given gate time, the resolv-

ing repetition rate requirements are approximately 3000% lower for direct-diagonal gates. The difference would be even more dramatic for position-preserving diagonal operation which would require an additional three adjacent gates (for the second SWAP operation). From here it is clear that, for symmetric two-dimensional crystals, it is far more preferable to do gates directly between diagonally nearest-neighbour ions, as opposed to composing equivalent operations with adjacent gates.

6.4 Gate performance in scaled crystals

Thus far, the analysis I have presented is limited to the smallest two-dimensional ion crystal with four ions in a 2×2 arrangement. However, for any practical application of quantum computing, more qubits are required. Therefore it is important to study the performance of fast gates in scaled crystals. In this section I will consider only $N \times N$ ion crystals, however the results presented place an upper bound on more general $N \times M$ arrays for $N \geq M$.

A brute-force approach to this analysis would be to optimise gates individually for crystals of different sizes, and compare fidelities for similar gates of similar operational time and minimum repetition rate resolution. However, the complexity of the truncated cost function (6.7) scales with the square of the number of motional modes, which in turn scales with the number of ions in the crystal. Thus in practise, gate optimisations are computationally infeasible for all but the smallest crystals.

An alternative approach is to optimise a gate for the 2×2 system, and then calculate its performance in a larger crystal. Motivated by the success of this approach for one dimensional geometries (*c.f.* Figure 5.3), this is the approach I will take here. This approach has the benefit of consistency in that the minimum resolving repetition rate and operational gate time do not change as the size of the crystal is changed. Thus the infidelities I will report in this section will be calculated taking into account the motional states all of the ions, for gates that are optimised only for the motional modes of a 2×2 crystal. This is equivalent to reporting a lower-bound on achievable fidelities as the gates optimised for the simple 2×2 array will in general not be the optimal gate for larger crystals. I will assume that all motional modes in each crystal size considered will have an average mode occupancy of $\bar{n}_p = 0.1$. This effectively assumes that larger crystals are hotter, which is a worst-case scenario.

In Figure 6.7 I have plotted the infidelity of an adjacent gate optimised for a 2×2 ion array as a function of crystal size (charactered by the number of rows N, for an

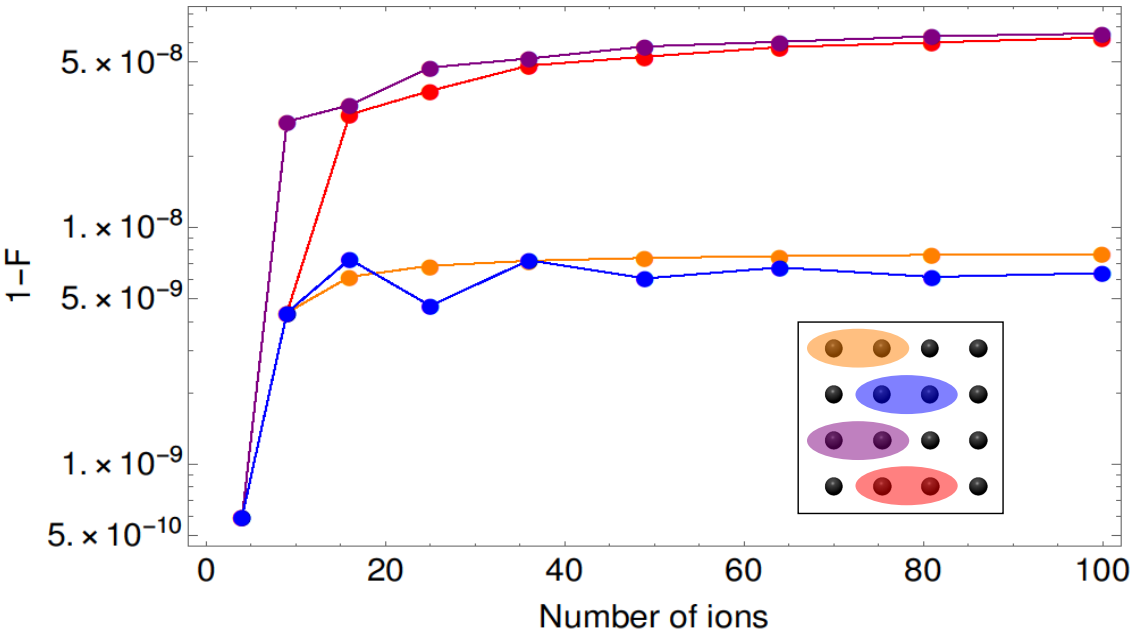

Figure 6.7: Scaling of nearest-neighbour adjacent gates in different locations in a square ion lattice of different sizes. The locations are pictorially represented for the 4×4 array. The chosen gate is optimised under the APG scheme with $\tau_G = 1.7$ and is resolvable by a 1 GHz repetition rate. Fast gates at each of the locations scale favourably, with infidelities plateauing to a constant value of $\sim 10^{-8}$ at high number of ions.

$N \times N$ crystal).[2] Notably, even for the worst-case placement the fidelity of the gate does not decay endlessly and plateaus at around $1 - 5 \times 10^{-8}$, which is still very close to unity. The plots also suggest this behaviour occurs regardless of which pair of neighbouring ions the gate is operated on. This is a very powerful result, as it shows that fast gates enable computation in scaled crystals without significant loss of gate fidelity or speed.

In the one-dimensional case, the scaling behaviour was perfectly captured by the performance of an optimised two-qubit gate at the edge of the chain and by one placed in the middle of the chain. In two-dimensional arrays, there are several distinct locations corresponding to whether the gate is between ions in the middle or edge rows/columns, which can be understood by considering the 4×4 ion crystal. This is visualised in each of the Figure, which shows that there are two clear trends. Surprisingly, the innermost and outermost ion pairs appear to follow the same trend, despite the distinct differences in their locations. I claim that this is due to two competing effects: the innermost ions experience increased Coulomb interaction from the surrounding ions, which results in a more rapid phase accumulation, however also involves motion of more surrounding ions, which is harder to restore. The out-

[2] This discussion extends to diagonal gates, which I found to have identical scaling behaviour.

ermost ions are surrounded by fewer ions, so experience less Coulomb interaction but also involve fewer motional modes. It appears that these two effects balance for the outermost and innermost ion pairs, resulting in comparable fidelities. A similar effect is visible for the ion pairs at the middle of the bottom of the lattice and the middle of the side of the lattice, albeit less extreme, corresponding to a slightly higher fidelity overall.

6.5 Conclusions

In this chapter, I have demonstrated that fast gates are well suited to symmetric two-dimensional crystals of ions in individual micro-traps. I showed that the performance of fast gates in this crystal can be characterised by the dimensionless parameter ξ, and the maximum number of pulses in a given pulse group. Notably, for experimentally reasonable values of the trap geometry parameters ($d = 100\,\mu$m and $\omega_t = 2\pi \times 1.2$ MHz), I found that extremely high-fidelity gates between neighbouring ions are resolvable by reasonable repetition rate lasers, even for gates with operation times smaller than the trap period. For example, a 5 GHz repetition rate laser is sufficient to resolve a 500 nano-second gate between adjacent ions with infidelity well below 2×10^{-4}. This demonstrates that fast gates are well suited to two-dimensional architectures of trapped ions in microtraps.

The notion of nearest-neighbour fast gates was extended to gates that are diagonally connected in this two-dimensional geometry. I found that diagonal gates only require slightly larger repetition rates to resolve than adjacent gates of a similar fidelity for the same gate time. Furthermore, I showed that diagonal gates are superior to equivalent operations constructed with adjacent gates which required over an order of magnitude larger repetition rates to resolve high-fidelity gates of a given duration. In general, diagonal gates will reduce the number of operations that need to be done to connect distant ions, thus making implementation of quantum algorithms more accessible.

The scaling behaviour of fast gates in these geometries I presented is also of great significance to practical applications of quantum computers. I found that gates optimised for a small 2×2 ion crystal perform very well as more ions are added to the system, with infidelities plateauing to the order of 10^{-8} for very large crystals in the worst-case. Further, this behaviour was identified even under the worst-case assumption that there is equal mean occupation in all motional modes of the scaled crystals. This places a strong upper bound on the achievable fidelities of nearest-neighbour fast gates in very large crystal geometries, well above error correction

thresholds, and enables large-scale computation in two-dimensional ion crystals.

Chapter 7

Experimental considerations

While the results of the previous chapters have indicated a promising outlook for the future of fast-gate enabled trapped ion quantum computing, it is important to note that the particular type of fast gates considered in this thesis (based on ultra-fast counter-propagating pulse pairs) have not yet been experimentally realised. Thus it is useful to consider practical challenges in realising fast gates, particularly as several experimental groups are working on its realisation [50, 70].

In this Chapter, I will address some of these challenges and make recommendations for experimentally realising robust, high-fidelity fast gates. In Section 7.1, I perform a worst-case analysis of the effect of pulse imperfections and make recommendations for improving the robustness of their population transfer. I then investigate issues specific to choice of candidate ions in section 7.2, focusing on the energy levels used for the momentum kick and their lifetimes. Finally, in section 7.3, I make a series of experimental recommendations based on the results of this thesis.

7.1 Effect of pulse imperfection

It has been identified in several analyses that errors in population transfer are the main limitations in experimentally implementing a high-fidelity fast gate [2, 38, 58, 88]. Specifically, this is a coherent source of error that is dominated by imperfections in the pulse areas of the ultra-fast pulses which lead to variability in the amount of rotation on the Bloch sphere, thereby resulting in unwanted internal state populations and incorrect motional states. Furthermore, Bentley *et al.* [58] identified that errors in a small number of pulses compound significantly with the number of pulses. This is fairly unfortunate, as the theoretical performance of fast gates, in both speed and fidelity, improves with the number of pulses.

7.1.1 Worst-case analysis

Here, I will assume that each the pulses in each counter-propagating pair are perfectly correlated, which is reasonable if they are generated by pulse splitting (this will be addressed later in this chapter in Section 7.3). Under this assumption, it is reasonable to set the phase of the laser to zero, $\phi = 0$, as phase will cancel out (see Appendix C).

Errors from a single pulse pair

The unitary for the effect of imperfect transition from a forward-propagating pulse is

$$\hat{U}_{\rightarrow} = \left(\hat{\sigma}_{+} e^{-ik\hat{x}} + \hat{\sigma}_{-} e^{ik\hat{x}}\right) \cos\theta + i\,\hat{\mathtt{i}} \sin\theta \tag{7.1}$$

where $\mathtt{i}$ is the identity, θ is some parameterisation of the error, and the factor of i is required for unitarity. Note that when $\theta = 0$, this reduces to the ideal unitary for a single π-pulse (*c.f.* equation (3.32)). The unitary for the counter-propagating pulse can be obtained by inverting the sign of the wave-number, $k \rightarrow -k$:

$$\hat{U}_{\leftarrow} = \left(\hat{\sigma}_{+} e^{ik\hat{x}} + \hat{\sigma}_{-} e^{-ik\hat{x}}\right) \cos\theta + i\,\hat{\mathtt{i}} \sin\theta\,. \tag{7.2}$$

The total unitary for the counter-propagating pair can then be obtained by simply multiplying these two unitaries together, which takes the following form in the $\{|g\rangle, |e\rangle\}$ basis:

$$\hat{U}_{\text{pair}} = \hat{U}_{\rightarrow}\hat{U}_{\leftarrow} = \begin{pmatrix} e^{2ik\hat{x}}\cos^2(\theta) - \sin^2(\theta) & i\cos(k\hat{x})\sin(2\theta) \\ i\cos(k\hat{x})\sin(2\theta) & e^{-2ik\hat{x}}\cos^2(\theta) - \sin^2(\theta)\,. \end{pmatrix} \tag{7.3}$$

By expanding to second order in θ, the unitary can be expressed as

$$\hat{U}_{\text{pair}} = (1-\theta^2)\hat{U}_0 + 2\theta\,\hat{U}_{error} - \theta^2\,\hat{\mathtt{i}} + \mathcal{O}(\theta^3) \tag{7.4}$$

where

$$\hat{U}_0 = \begin{pmatrix} e^{2ik\hat{x}} & 0 \\ 0 & e^{-2ik\hat{x}} \end{pmatrix}, \text{ and} \tag{7.5}$$

$$\hat{U}_{error} = \begin{pmatrix} 0 & i\cos(k\hat{x}) \\ i\cos(k\hat{x}) & 0 \end{pmatrix}, \tag{7.6}$$

and $\hat{\mathtt{i}}$ is the identity operation. Here, $\hat{U}_0$ is the ideal unitary for the counter-propagating pulse pair, as derived in Appendix C. It describes a momentum kick of

$\pm 2\hbar k$, with the sign dependent on the internal state of the ion. The 'error' unitary $\hat{U}_{error}$ describes a failure to restore the original internal state, and also results in an incorrect motional state. States generated by the application of the error unitary will always be orthogonal to states generated by the ideal unitary, by virtue of the internal electronic states being orthogonal to one another. The identity operation results in only an incorrect motional state, not an incorrect internal state.

Errors from N pulse pairs

The total unitary for a series of N pulse pairs is simply the product of each individual unitary, i.e.

$$\begin{aligned} \hat{U}_{total} &= \left(\hat{U}_{pair}\right)^N && (7.7)\\ &= \left((1-\theta^2)\hat{U}_0 + 2\theta\,\hat{U}_{error} - \theta^2\,\hat{\mathbb{1}} + \mathcal{O}(\theta^3)\right)^N && (7.8)\\ &\approx (1-2N\theta^2)\,\hat{U}_0^N + 2N\theta\hat{U}_\perp && (7.9) \end{aligned}$$

where in the last line I have grouped all non-ideal unitaries into $\hat{U}_\perp$. I will assume that $\hat{U}_\perp$ always produces states orthogonal to states produced by application of $\hat{U}_0^N$. This is a worst-case assumption, as it neglects any terms that may result in some overlap with the ideal final state.

This expression of the unitary can be used to calculate a fidelity for the total N pulse pair operation. For some initial state $|\psi_0\rangle$ the fidelity takes the form,

$$F_{\text{N pairs}} = \left|\,\langle\psi_0|\,\hat{U}^\dagger_{\text{target}}\hat{U}_{\text{actual}}\,|\psi_0\rangle\,\right|^2. \tag{7.10}$$

Substituting (7.9) for the actual unitary, and noting the target unitary is simply $\hat{U}_0^N$, this expression becomes

$$F_{\text{N pairs}} = \left|(1-2N\theta^2)\,\langle\psi_0|\,(\hat{U}_0^N)^\dagger\hat{U}_0^N\,|\psi_0\rangle + 2N\theta\,\langle\psi_0|\,(\hat{U}_0^N)^\dagger\hat{U}_\perp\,|\psi_0\rangle\,\right|^2. \tag{7.11}$$

The second term in this expression is zero by virtue of how I have defined $\hat{U}_\perp$, and thus the fidelity is simply given by (to second order in θ):

$$F_{\text{N pairs}} = \left|(1-2N\theta^2)\right|^2 \approx 1 - 4N\theta^2 \tag{7.12}$$

The parameter θ can be understood as an amplitude for the part of the state that is orthogonal to the desired state due to imperfect population transfer. This can be seen from inspection of equation (7.1), and noting $\sin\theta \approx \theta$ for $\theta \ll 1$. Therefore θ^2 can be interpreted as a probability that an error in the population transfer is

				$\epsilon = 10^{-5}$	$\epsilon = 10^{-6}$	$\epsilon = 10^{-7}$	$\epsilon = 10^{-8}$
	τ_G	$1 - F_0$	N	$1 - F_{real}$			
	0.45	1.0×10^{-4}	1552	6.2×10^{-2}	6.3×10^{-3}	7.2×10^{-4}	1.6×10^{-4}
(a)	1.0	6.3×10^{-9}	640	2.6×10^{-2}	2.6×10^{-3}	2.6×10^{-4}	2.6×10^{-5}
	1.75	2.4×10^{-7}	191	7.6×10^{-3}	7.6×10^{-4}	7.7×10^{-5}	7.9×10^{-6}
	0.65	3.2×10^{-5}	64	2.6×10^{-3}	2.9×10^{-4}	5.7×10^{-5}	3.4×10^{-5}
(b)	0.85	3.3×10^{-5}	72	2.9×10^{-3}	3.2×10^{-4}	6.5×10^{-5}	3.6×10^{-5}
	1.25	2.2×10^{-6}	46	1.8×10^{-3}	1.9×10^{-4}	2.1×10^{-5}	4.1×10^{-6}
	0.9	8.4×10^{-5}	2306	9.2×10^{-2}	9.3×10^{-3}	1.0×10^{-3}	1.8×10^{-4}
(c)	1.2	2.9×10^{-8}	1690	6.8×10^{-2}	6.8×10^{-3}	6.8×10^{-4}	6.8×10^{-5}
	2.0	1.1×10^{-9}	910	3.6×10^{-2}	3.6×10^{-3}	3.6×10^{-4}	3.6×10^{-5}
	0.9	2.3×10^{-4}	2322	9.3×10^{-2}	9.5×10^{-3}	1.2×10^{-3}	3.2×10^{-4}
(d)	1.2	1.7×10^{-7}	2060	8.2×10^{-2}	8.2×10^{-3}	8.2×10^{-4}	8.3×10^{-5}
	2.0	1.7×10^{-9}	866	3.5×10^{-2}	3.5×10^{-3}	3.5×10^{-4}	3.5×10^{-5}

Table 7.1: Infidelities including worst-case pulse errors for gates (a) in a two-ion linear microtrap chain, and (b) between two-ions in a single linear Paul trap. (c) and (d) tabulate infidelities for gates in a 2×2 microtrap array between adjacent and diagonal pairs of ions, respectively. This is reported for a range of gate times τ_G (in trap periods), and different values of the transition error ϵ. Here, F_0 is the infidelity of the optimised gate not including pulse errors, and N is the number of pulse pairs in the gate.

measured. I will define this to be the transition error, $\epsilon \equiv \theta^2$.

For a fast gate, realistic fidelities can thus be calculated for a gate with N pulse pairs by multiplying the theoretical fidelity F_0 (the fidelity assuming perfect population transfer, as reported in previous chapters) by (7.12):

$$F_{real} = (1 - 4N\epsilon)\, F_0\,. \tag{7.13}$$

This forms a lower bound on achievable gate fidelities. For a gate with $N{\sim}1000$ pulse pairs, the transition error needs to be on the order of 10^{-6} in order to achieve fidelities beyond 99%. This clearly shows that robust population transfer is vital for realising high-fidelity fast gates.

In Table 7.1, I present the infidelities of gates that have been presented in the previous chapters for both linear Paul trap and microtrap architectures, for a range of transition error values. Across a range of gate times, it is is clear that transition errors are required on the order of $\epsilon = 10^{-8}$ or smaller in order to maintain infidelities below the fault tolerant threshold of 2×10^{-4}. Single qubit operations in trapped

ions on optical transitions have been demonstrated with infidelities as low as 10^{-5} [19], which is a good order-of-magnitude approximation of the achievable values of ϵ. Notably these fidelities were achieved using shaped pulses. If single-qubit operations can be performed to this fidelity on the ultra-fast timescales required for fast gates, $1-2$ trap period fast gates with fidelities above 99% in microtrap architectures are achievable. Notably, $\epsilon = 10^{-5}$ enables sub-microsecond gates with fidelity on the order of 99.9% between two ions in linear Paul trap. However to go beyond this value and towards gates that enable fault-tolerant error-correction schemes in large crystals, there needs to be substantial improvement in minimising errors in population transfer.

7.1.2 Achieving high-fidelity population transfer

Fast gates are composed of ultra-fast pulses that each perform a population inversion on the qubit state. The simplest method of implementing a population inversion is with a single, square π-pulse. I will refer to this as the *primitive* approach. For a laser with intensity fluctuations of ΔI around some steady-state intensity I, the transition error ϵ for a square pulse is,

$$\epsilon = \frac{\pi^2}{8}\frac{\Delta I}{I}\,. \tag{7.14}$$

This means that square pulses are highly sensitive to fluctuations in intensity fluctuations in a laser. Current state-of-the-art trapped ion experiments are able to get intensity fluctuations as low as 10^{-3} [89], which can be further improved by calibration of the laser intensity to reduce the effect of long-term intensity drift. Given that there are typically $\mathcal{O}(10^{10})$ photons per pulse, the existence of shot noise would suggest that achieving transition errors of 10^{-6} or lower will require a scheme that is robust to first order fluctuations in intensity.

Methods for robust population transfer in two-level systems have been well studied, not only in the context of quantum information, but in many other fields as well, particularly in nuclear magnetic resonance (NMR). Here I will outline several suggestions for improving population inversions that can be applied to current trapped-ion experiments. The general concept of each of these proposals is to reduce the sensitivity of the population transfer to laser intensity fluctuations.

Shaping of single pulses

One way to improve population transfer is to use non-square pulse shapes that are less sensitive to fluctuations in intensity. This generally involves use of a diffraction grating to split the pulse into its frequency components. By applying different phase shifts to different components, and recombining them, the profile (waveform) of a pulse can be shaped. Pulse shaping techniques have been extensively developed in the context of NMR, and are known to improve robustness of population transfer to intensity fluctuations as well as other pulse imperfections such as timing or frequency errors [90, 91]. There has also been extensive work in using machine learning to optimise pulse shaping for robustness against noise sources common in trapped ion experiments [92], which have promising applications to the regimes relevant to fast gates. However, the world of NMR is typically concerned with the radio-frequency domain, i.e. pulse duration on the order of μs. Fast gates require pulses on the order of pico-seconds, which are technically far more challenging to shape, particularly for ultra-violet and optical frequencies [93].

Over the last decade, there has been great progress in developing pulse shaping methods suited to ultra-fast pulses as fast as femtoseconds [93, 94]. Simple shaping methods have recently been applied to pico-second single-qubit operations in trapped ions [34, 95], albeit only achieving fidelities slightly above 99%. While this still falls short of the transition errors required to implement the gates in Table 7.1, the use of more technically challenging techniques are promising for improving ultra-fast single qubit operations significantly in the near future.

Composite pulses

Another way in which population transfer errors can be reduced is through the use of *composite pulse sequences*, which have been adapted from the field of NMR for quantum computation (an excellent review of NMR techniques adapted for quantum computation can be found in Ref. [96]). These protocols describe sets of composite pulses that each perform different rotations, but implement the same net rotation as a single primitive pulse. Most relevant for use in fast gates, is the BB1 family of protocols, which describe simple pulse sequences that can cancel the effect of laser intensity fluctuations on population transfer, to first order. These protocols have been shown to be effective in correcting for both static and time-dependent noise channels [97, 98]. There are also protocols for cancelling higher order intensity fluctuations [99], but at the cost of requiring more complicated pulse sequences that may not be practical on an ultra-fast timescale.

In a recent paper, Edmunds *et al.* [100] demonstrated the use of the CORPSE (Compensation of Off-Resonance with a Pulse SEquence), WAMF (Walsh Amplitude Modulated Filter), and the BB1 protocols to implementing composite-pulse single-qubit gate operations on a multi-ion system. They found that these protocols reduced the magnitude of error of each single-gate, and also reduced 'cross-talk' with ions not involved in the gate. Notably, the use of composite pulse methods is compatible with pulse shaping techniques; each pulse in the composite-sequence can be shaped for further robustness to intensity fluctuations.

Rapid adiabatic passage (RAP) with chirped laser pulses

High-fidelity population transfer can also be achieved by use of rapid adiabatic passage (RAP) with chirped laser pulses [101]. In RAP, Rabi oscillations are excited by a pulse that has an initial detuning far above the resonant frequency of the two-level system, which is frequency-swept until the detuning is far below resonance. If the change in the detuning is adiabatic with respect to the Rabi frequency, population inversion can be done with 100% theoretical efficiency [102], and is far more robust to fluctuations in laser intensity than the primitive approach. Furthermore, using linearly chirped laser pulses, this process can be done as fast as tens of femtoseconds [103]. Notably this process can be modified for a Raman configuration [104].

RAP has been demonstrated in trapped ion control [105], but only in the context of the weak-coupling regime, which is necessarily much longer than the timescales required for fast gates. While ultra-fast population transfer on the order of picoseconds is certainly possible with the use of chirped laser pulses [103], it is unclear whether there are practical barriers to its experimental implementation.

7.2 Candidate ions

Thus far in this thesis I have taken $^{40}Ca^{+}$ to be the candidate ion for gate calculations, and in particular have focused on laser pulses resonant on the $S_{1/2} \rightarrow P_{3/2}$ transition. This choice is made largely for convenience of comparison to previous studies in fast gates, namely Refs. [2, 36, 38], and does not necessarily reflect the ideal choice for experimental realisations. Ultimately choice of candidate ions will come down to practical factors to do with the specifics of the experiment, available laser resources, and monetary costs. It is important to note that the implementation of fast gates is not specific to a particular species of ion. In this section, I will investigate the effect the choice of candidate of ion may have on fast gate perfor-

mance. In principle, the only difference relavant to fast gates is the wavelength of the transitions used to drive the momentum kicks. However there are other practical considerations such as state life-time, and availability of electronic states suitable for read-out of the qubit.

7.2.1 Big kicks versus long lifetimes

As alluded to previous chapters, bigger momentum kicks are favourable for faster and higher fidelity fast gates. This is because the speed of a fast gate is limited by how quickly the relative phase difference of $\pi/4$ can be accrued between the computational basis states. This corresponds to the area enclosed in the ions' trajectories through phase space, which in turn is directly related to the size of the momentum kicks. Therefore it is clear that it is favourable to use pulses resonant on a high-energy (strong) transition. Unfortunately, high-energy states tend to have short lifetimes, and so there is a trade-off to be made.

Resonant (single-photon) transitions

One approach is to use a high-energy transition for the momentum kicks, and use a separate long-lived state accessible by a lower energy transition to store the computational $|1\rangle$ state. For the calculations made in this thesis I have taken this approach for $^{40}Ca^{+}$ ions, using the $S_{1/2} \rightarrow P_{3/2}$ transition for the momentum kicks, and the meta-stable $D_{5/2}$ level for the computational $|1\rangle$ state. Here I will outline two methods for implementing this approach: I will use $^{40}Ca^{+}$ as an example, however these techniques will work for a general ion species with controllable transitions to a low-energy long-lifetime electronic level that can be used as the $|1\rangle$, and a high-energy state that can be used for momentum kicks.

One method of realising this approach is to excite a transition between the $|1\rangle$ state and the high-energy state before and after each pulse group in the fast gate. In doing so, the $|1\rangle$ state can be stored in the long-lived state at all times except during the times when the pulses are being applied to ensure that the full momentum transfer is achieved. Care must be taken, however, to ensure that this '(un-)shelving' process does not impart extra momentum onto the ion that is not accounted for in gate optimisation. I propose the following procedure in the context of particular transitions in $^{40}Ca^{+}$, but is applicable to other transitions as well as other ion species.

1. Prior to the arrival of each pulse group, a pair of counter-propagating $\pi/2$-

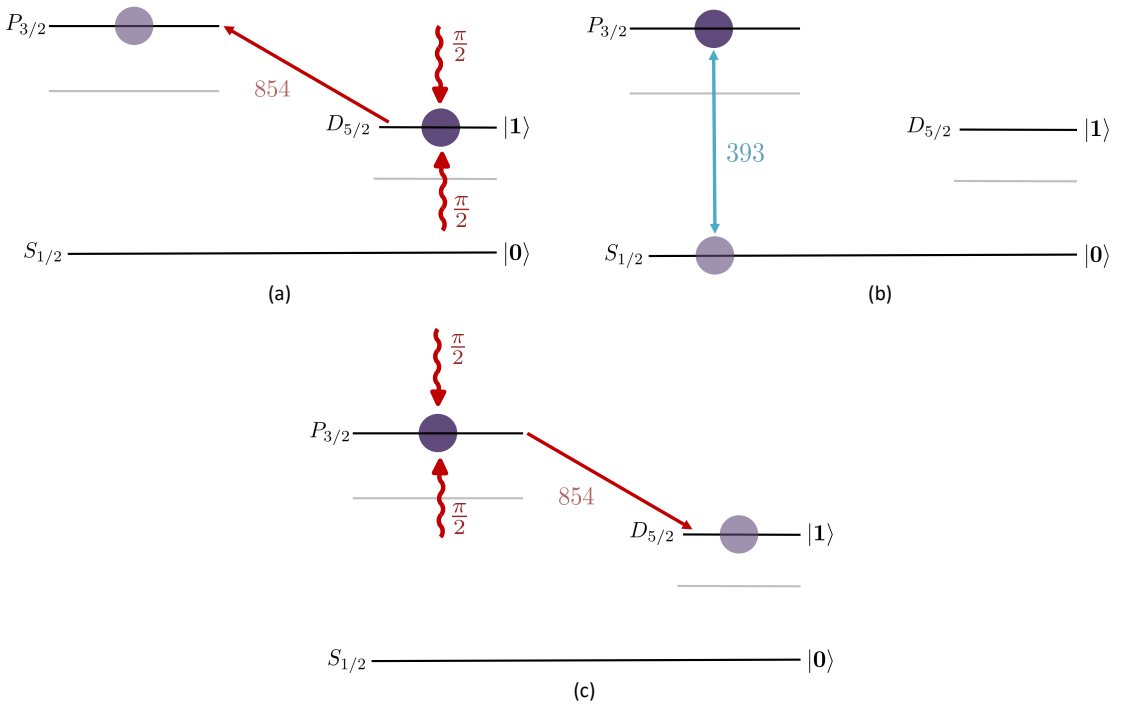

Figure 7.1: Protocol for un-shelving the $|1\rangle$ part of the qubit from the metastable $D_{5/2}$ level in $^{40}Ca^{+}$, to the high-energy $P_{3/2}$ that can be used for momentum kicks in the fast gate. (a) A pair of counter-propagating $\pi/2$ pulses with wavelength 854nm invert the state population between the $D_{5/2}$ to the $P_{3/2}$ levels, and imparts no net momentum. (b) The ion is now in a two-level system of the $S_{1/2}$ and $P_{3/2}$ levels. A given group of counter-propagating pulse pairs with ultra-short wavelengths of 393nm apply a momentum kick to the ion, as part of the fast gate. (c) A second pair of counter-propagating pulses with wavelength 854nm excites a transition between the $P_{3/2}$ to the $D_{5/2}$, restoring the excited state of the ion to the computational $|1\rangle$ state. Additional energy levels of $^{40}Ca^{+}$ not involved in this process are indicated in grey.

pulses are applied to the ion, resonant on the $D_{5/2} \rightarrow P_{3/2}$ transition. This gives a population inversion (effective π-pulse) from the computational $|1\rangle$ state to the high-energy $P_{3/2}$ with no net momentum transfer due to the counter-propagation.

2. The counter-propagating pulse pairs that compose the pulse group are applied resonant to the high-energy transition $S_{1/2} \leftrightarrow P_{3/2}$. This is the step corresponding to the momentum kick in the fast gate.

3. Directly after the end of the momentum kick, a second pair of counter-propagating $\pi/2$-pulses resonant on the $P_{3/2} \rightarrow D_{5/2}$ transition are applied, inverting the population from the high-energy $P_{3/2}$ state back to the long-lifetime computational $|1\rangle$ state (in this case, the $D_{5/2}$ level).

This process is diagramatically presented in Figure 7.1. The advantage of this method is that the full momentum kick is obtained, and the long-lifetime of the low-energy state can be exploited for storing the $|1\rangle$ basis for the qubit. Unfortunately this method is significantly limited by the repetition rate of the laser which dictates the time taken for application of each pulse group, and thus the amount of time the ion is in the high-energy state. Even for a pulse group with $\sim$10 pulse pairs enabled by a state-of-the-art 5 GHz laser, the time spent in the high-energy state would be on the order of $\sim$2 ns. For the example of $^{40}Ca^{+}$, the high-energy $P_{3/2}$ level has a lifetime of 6.9 ns [50], and so this protocol is highly susceptible to decoherence through spontaneous emission and decay to lower energy levels. This effect may be mitigated for a different choice of high-energy transition, but ultimately limits the effectiveness of this protocol.

Perhaps a simpler method is to just use the high-energy ($S_{1/2} \rightarrow P_{3/2}$) transition for the fast gate, leaving the computational $|1\rangle$ state untouched. This is the approach taken by Ref. [50], and results in an effective halving of the magnitude of each of the momentum kicks [36]. Thus, while this approach has the benefit of simplicity and longevity of the $|1\rangle$ state, it will result in slower achievable gates for a given laser repetition rate. This method is roughly equivalent to simply exciting a lower energy transitions between the $|0\rangle$ and $|1\rangle$ states, which would have a weaker momentum kick and thus would require greater repetition rate to maintain gate speed. This will be more explicitly considered later in this chapter (see section 7.2.2). Furthermore, in this method there is still some time where the qubit is in the short-lifetime (high-energy) state, which is the time delay between the arrival of the first and second pulses in each counter-propagating pair. The fast gates I have considered in this thesis typically have on the order of $\sim$1000 pulse pairs (see Table 7.1), and thus the delay between the counter-propagating pulse pair will have to be as fast as a picosecond to avoid decoherence from spontaneous emission of the excited state. This is a restrictively short timescale that will likely pose a challenge to experimentally realise.

In essence both these methods are limited by the timescales of the laser pulses, and are likely not scalable to implementing multiple fast gates in succession. To go beyond these limitations then, I suggest the use of two-photon methods.

Raman (two-photon) transitions

A more robust approach is to stimulate transitions between the computational $|0\rangle$ and $|1\rangle$ states via a third high-energy auxiliary $|e\rangle$ state using a Raman transition (introduced briefly in Section 3.1). A diagram of a Raman transition driven by two

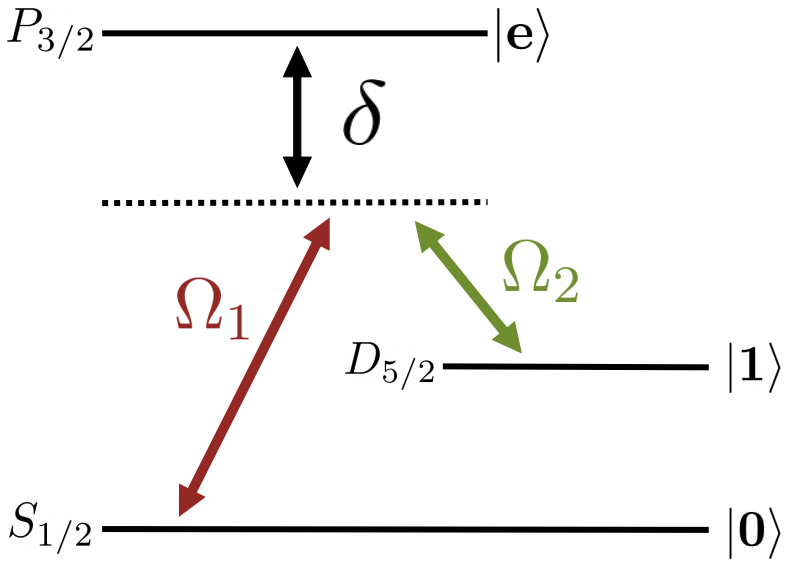

Figure 7.2: Raman transition between $S_{1/2}$ ($|0\rangle$) and $D_{5/2}$ ($|1\rangle$) states via a detuning δ on the high-energy auxiliary $P_{3/2}$ ($|e\rangle$) state in $^{40}Ca^{+}$. The transition is enabled by a pair of lasers with Rabi frequencies of Ω_1 and Ω_2.

lasers is shown in Figure 7.2. The benefit of using a Raman transition is that while the net transition is between states of small energy separation, the change in kinetic energy resulting from the transition can be made much higher. This can be seen by considering the effective wave-vector of the transition, which is given by

$$\mathbf{k}_{\text{eff}} = \mathbf{k}_1 - \mathbf{k}_2 \,, \tag{7.15}$$

where $\mathbf{k}_1$ and $\mathbf{k}_2$ are the wave-vectors of lasers with respective Rabi frequencies Ω_1 and Ω_2. Importantly, this means that if the two lasers are counter-propagating, then the magnitude of the wave-vector is simply the sum of the wave-vectors of each transition,

$$|\mathbf{k}_{\text{eff}}| = |\mathbf{k}_1| + |\mathbf{k}_2| \,. \tag{7.16}$$

In essence, the use of Raman transitions allows for the magnitude of the state-dependent momentum kicks to be doubled. This leads to a stronger coupling of the light field to the ion, resulting in a larger effective η.

The use of Raman transitions can also be used to improve individual addressing of ions in a multi-ion system. This is because the Raman transition requires both lasers to be on in order to drive the transition, and will not stimulate excitations if only one of the two lasers is applied. This allows for the possibility of two-dimensional addressing, where one of the lasers is aimed down one row of the ion crystal and the other is oriented along a column, thus targeting a single ion. Further, it may not

be the case that both lasers require high-repetition rates: continuously addressing an ion (or even multiple ions) with a continuous laser, and pulsing the second laser may be sufficient for experimental implementations of fast gates.

The use of counter-propagating Raman lasers in trapped ion experiments has been well demonstrated, and recently implemented in ultra-fast π-pulses in regimes relevant to the fast gates studied in this thesis [34, 95]. Raman transitions are also compatible with the methods outlined earlier for improving pulse imperfections (composite pulses, pulse shaping, and adiabatic rapid passage with chirped pulses). For these reasons, and those outlined above, the use of Raman transitions is likely necessary for the implementation of high-fidelity fast gates.

7.2.2 Gate dependence on η

I have alluded to the lower energy transitions corresponding to higher repetition rate requirements for implementing fast gates, and here I will explicitly show this. By inspection of the fast gate condition equations (3.35) and (3.36), the dependence of the gate fidelity on the light field and the choice of candidate ion is only in the Lamb-Dicke parameter,

$$\eta = k\sqrt{\frac{\hbar}{2M\omega_t}}\,, \tag{7.17}$$

which describes the coupling of the light-field to the motion of the ions. In other words, the Lamb-Dicke parameter dictates the magnitude of the excited motion in the ion by each laser pulse. Recall that in this expression, k is the wavenumber of the laser pulse and M is the mass of the candidate ion. In Table 7.2, I have listed candidate ions and their transition wavelengths commonly used in trapped ion quantum computing, and calculated the Lamb-Dicke parameter for each assuming a trapping frequency of $\omega_t = 2\pi \times 1.2$ MHz and a light field resonant on each transition. While these may not be the true Lamb-Dicke parameters in an experimental realisation, which may use a different trapping frequency or a Raman transition with a different wavenumber, this indicates the range of Lamb-Dicke parameters of relevance: from $\eta{\sim}0.02$ to $\eta{\sim}0.43$.

To understand the effect changing the value of η has on gate performance, optimisations were run for a range of values for a one-dimensional microtrap array with two ions. This was done for gate times of 667 ns, 1.0 μs, and 1.3 μs. The results of the optimisations are shown in Figure 7.3. These results qualitatively show that higher values of η correspond to lower required repetition rates to resolve high-fidelity (i.e. above the fault-tolerant threshold of 99.98%) for each given gate time. In fact, this is roughly linear: the repetition rates required to resolve gates with infidelities below

Candidate ion	Transition wavelength (nm)	Lamb-Dicke† (η)
$^{9}Be^{+}$	313	0.43
$^{25}Mg^{+}$	280	0.29
	397	0.16
$^{40}Ca^{+}$	393	0.16
	729	0.09
	422	0.10
$^{87}Sr^{+}$	408	0.11
	674	0.06
$^{111}Cd^{+}$	226	0.17
	214	0.18
	493	0.07
$^{133}Ba^{+}$	455	0.08
	1762	0.02
	369	0.08
$^{171}Yb^{+}$	329	0.09
	411	0.08

Table 7.2: Transitions in ions of interest for quantum computing, wavelength data taken from Ref. [25]. Each transition can be addressed with a Raman and/or optical transition. The decay rates of the excited levels have been omitted here, as this table is simply indicative of Lamb-Dicke parameters one may have in an experimental set-up. †Lamb-Dicke parameter calculated for a trapping frequency of $\omega_t = 2\pi \times 1.2$ MHz. Recall that I have defined the Lamb-Dicke parameter as $\eta = \sqrt{\frac{\hbar}{2M\omega_t}}$, which is independent of the motional mode structure.

$\sim 10^{-4}$ are approximately double for $\eta = 0.08$ than for $\eta = 0.16$.

The relationship between repetition rate and η is quantitatively presented in Figure 7.4, where I have plotted the minimum repetition rate to resolve gate infidelities below 2×10^{-4} as a function of η. This is presented for both two ions in individual microtraps (separated by $d = 100$ μm), as well as in a linear Paul trap. Clear monotonic trends are present in these figures, which show a significant correlation between the value of η and minimum resolving repetition rate. Reasonably higher values of η than the value studied in this thesis ($\eta = 0.16$), such as those corresponding to $^{9}Be^{+}$ and $^{25}Mg^{+}$ in Table 7.2, correspond to dramatically lower repetition rates required to resolve sub-microsecond gates with high fidelity. Notably, the data suggests a 3 GHz laser can resolve high-fidelity gates as fast as 83 ns in a linear Paul trap with an axial trapping frequency of 1.2 MHz, which is an order of magnitude faster than any entangling gate that has been realised to date [25] on a trapped ion processor and is comparable to gate times in superconducting qubits [106]. A similar result is shown for the microtrap case, with the data suggesting a 300 MHz repetition rate is able to resolve a high-fidelity gate as fast as 667 ns between traps separated by 100

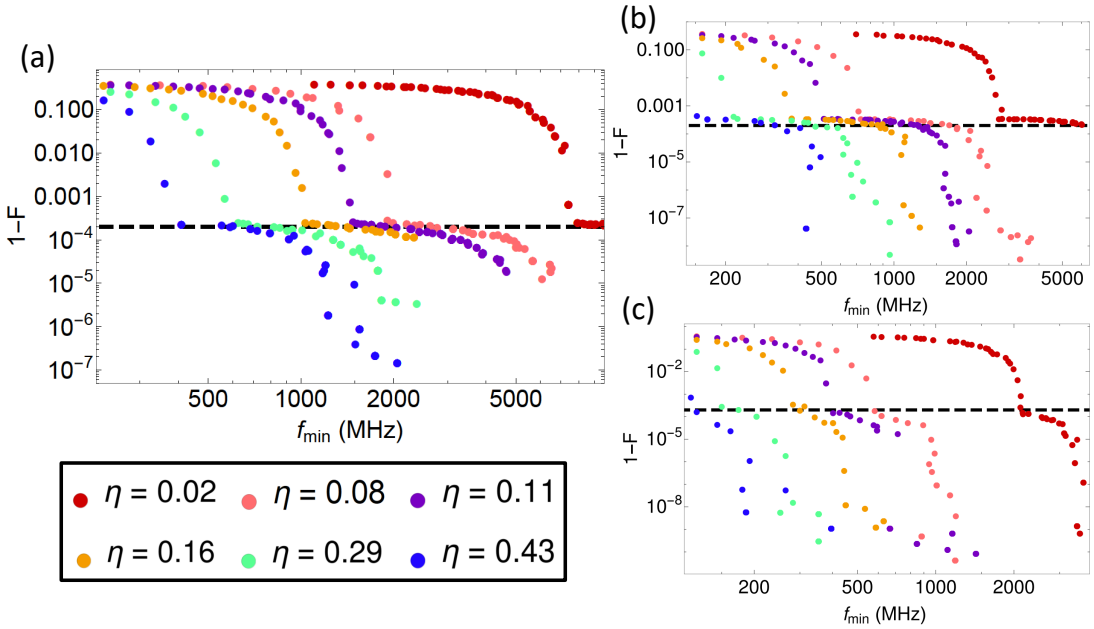

Figure 7.3: Fast gate optimisations for a two-ion system in a linear microtrap array ($\omega_t = 2\pi \times 1.2$ MHz, $d = 100$ μm), across a range of different values of η. Gate infidelity is shown as a function of minimum resolving repetition rate. Gate time is fixed to 667 ns, 1.0 μs, and 1.3 μs in (a), (b), and (c), respectively. The APG scheme was used for optimisation with 16 pulse groups. The fault tolerance threshold of 2×10^{-4} is indicated by the black dashed line.

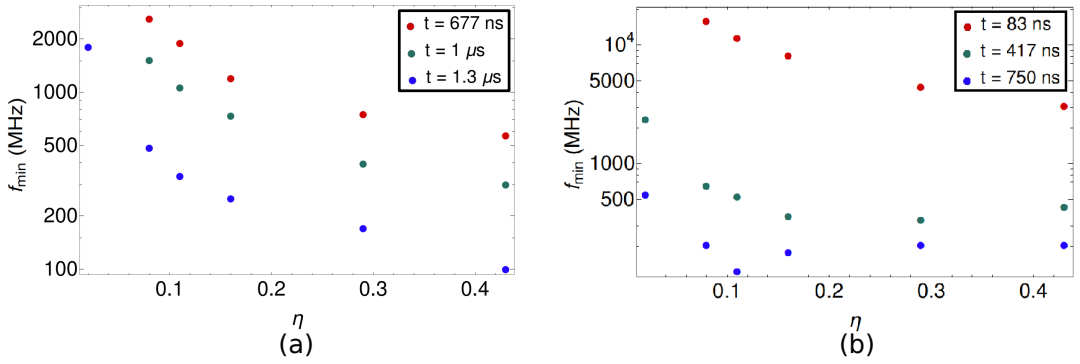

Figure 7.4: Minimum resolving repetition rate plotted as a function of η for a range of gate times for (a) two ions in a linear microtrap array with $d = 100$ μm and (b) two ions in a single linear Paul trap. Each data point corresponds to a gate solution with fidelity above 99.98% (or equivalently, an infidelity below 2×10^{-4}). Repetition rate requirements decrease significantly as η is made larger, allowing for faster gates to be resolved.

microns. Perhaps of more importance is the fact that fast gates can be performed effectively with η values as low as 0.02 as fast as 1 μs between two microtraps or 417 ns in a linear Paul trap, with a 3 GHz repetition rate laser.

Regardless, it seems favourable to, if possible, maximise the value of η in any ex-

perimental realisation of a fast gate. Again, this is not something that can be done arbitrarily. The choice of candidate ion goes beyond the effectiveness of two-qubit gates, and includes other considerations such as ease of state-preparation, efficiency of measurement, and availability of suitable lasers for the particular transitions. Decreasing the trap frequency will always increase η but will also change the timescale of the oscillation of the ions. This in turn will lead to slower fast gates for a particular repetition rate laser. This can be seen from inspection of the fast gate condition equations, (3.35) and (3.36), from which the trap frequency can be identified as the timescale for fast gate dynamics. Therefore once a candidate ion is chosen and a particular transition is identified for the momentum kicks, it is unlikely that further modifications to the value of η are likely to improve the fast gate performance, at least in terms of gate speed.

7.3 Recommendations for implementation

Thus far in this chapter I have considered several possible ways in which fast gates may be made more robust. Here I will present my recommendations for experimentally realising fast gates based on the results of this chapter, and this thesis as a whole. I will omit discussion on improving pulse imperfections as they have been well-described in Section 7.1.2 and are compatible with the recommendations I will outline below.

Choosing the ions

I will not comment extensively on the choice of candidate ions, as there are many more factors to be considered than have been studied here. However, I will comment on some of the desirable qualities one should consider. The first is and most obvious is the existence of a long-lifetime electronic level to use as the computational $|1\rangle$ state, which can be given by either a meta-stable excited state or a low-energy hyperfine state. There should be an accessible optical or Raman transition between these levels as well. A lighter ion is preferable in general as it corresponds to a larger η value, although trapping lighter species may prove more difficult than heavier ions like $^{171}\mathrm{Yb}^{+}$.

Trapping architecture

The choice between an architecture of ions trapped in a linear Paul trap or one in which ions are in individual microtraps largely depends on the purpose of the experiment. For a proof-of-principle experiment that only involves two ions, a linear Paul trap architecture is favourable, as it enables ultra-fast gates as fast as tens of nanoseconds for state-of-the-art lasers with GHz repetition rates. For an experiment looking to implement a non-trivial computation, such as a quantum simulation algorithm, microtrap architectures are much more favourable. As shown in Chapters 5 and 6, the fidelity of fast gates do not drop significantly as microtrap architectures are scaled, in contrast with linear Paul traps where even chains of $\sim$5 require restrictively high repetition rates to implement gates with high fidelity.

For large-scale computation with fast gates, two-dimensional microtrap architectures are favourable, as the increased connectivity allows for a significant reduction in the number of gate operations (exemplified in Chapter 4). This is a scalable architecture, allowing for high-speed and high-fidelity gates to be implemented between nearest-neighbour qubits in scaled architectures with experimentally feasible repetition rate lasers.

Gate optimisation

For the gate fidelities that I have presented in this thesis, I have made several assumptions to make the global optimisation process more computationally efficient. In practise, these assumptions can be addressed by a further set of optimisations. I propose three distinct steps in overall gate design, to optimise gates for implementation in a given experimental set-up.

First, global optimisations of the truncated infidelity function (3.40) should be done, following the method outlined in Chapter 5. For this to be done, the harmonic frequencies of the trap must be known. The GPG and APG schemes should both be used in this optimisation, with the optimal solution taken from both optimisations. For optimisation of systems with more than two-ions, the APG scheme will generally be preferable. Furthermore, this optimisation should be done over a range of fixed gate times to identify the fastest gate that can be implemented with high-fidelity by the laser repetition rate available. A starting point for selecting which gate times to optimise for can be obtained from the scaling behaviour $\tau_G \propto f_{min}^{-2/5}$ identified in Chapter 6.

Next, local optimisations of the best gate from the previous step should be carried

out using a set of ODEs that describe the classical motion of each ion (as described in section 3.5.3). Each pulse group in the solution should be expanded out into individual pulse pairs occurring on a finite grid of timings specified by the laser repetition rate. The local optimisations should be done on the pulse timings, snapped to a finite grid. As the ODE description of the ions' motion does not assume a harmonic potential, measured anharmonicities of the trap can, and should, be included in this detailed optimisation. The output of this step is a gate solution that is specific to the experimental trapping potential and the repetition rate of the laser.

In the final step, an *online* optimisation of the gate (from the previous step) should be done. In an online optimisation the gate is experimentally implemented, the resulting two-ion quantum state is tomographically reconstructed, and the infidelity of the gate is calculated. This is then repeated for small variations of the timings of the pulses in the gate. By identifying a gradient in the calculated fidelities, a slope of steepest descent in the parameter space can be identified. Further modifications of the gate timings along this slope, and repetition of this process implements a local optimisation of the gate scheme. This online optimisation allows for aspects of the experiment which are not well-known *a priori* to be accounted for. The resulting gate is therefore optimised for the exact nature of the experiment, which promises the highest fidelity results.

Pulse splitting

The final recommendation I will make is to use pulse splitting techniques. Specifically, I recommend that pulse splitting is used for the counter-propagating pulse pairs: have the laser emit 2π-pulses and then split each into a pair of counter-propagating π-pulses which follow paths of different lengths to the ion, and thus are delayed with respect to each other. This delay can be much shorter than the time between pulses emitted from the laser, and can be tuned to be effectively instantaneous. This is the approach taken by current leading groups working on implementing fast gates [34, 50, 70].

The advantage of this method is two-fold. First, the alternative of simply picking out every other pulse emitted from the laser to be the counter-propagating is inefficient and will lead to a doubling of repetition rate requirements to match the beam-splitting method. More importantly, however, is that this method makes the gate insensitive to phase noise: the phase contribution of each π pulse in a counter-propagating pair perfectly cancels, as they are identical by virtue of being split from the same 2π pulse. Thus phase drift between pulses emitted from the laser has no effect on the fast gate, and is removed as a potential source of error.

Chapter 8

Discussion and conclusions

The field of quantum computing is moving out of its infancy with experimental demonstrations of quantum algorithms rapidly approaching the capabilities of even the most powerful supercomputers. However, it remains unclear what the physical platform that quantum computers of the future will be based upon. Superconducting circuits and trapped ions are currently leading experimental demonstrations of small-scale computation, but a truly scalable platform remains to be implemented.

8.1 Summary of results

In this thesis I have presented a blueprint for a scalable architecture for a quantum information processor, where fast gates enable computation in a two-dimensional array of ion traps. I found that the increased connectivity of trapped ion qubits in two-dimensions can be used to significantly reduce the number of gates that need to be implemented for large-scale quantum computation. I have shown this explicitly for a quantum simulation algorithm of the Fermi-Hubbard model with 40 fermionic modes in Chapter 4, where I have calculated that a two-dimensional array of ions allows for the number of gates per Trotter step to be reduced from 4236 to 2468. For larger computations involving N qubits, the reduction in number of gates from moving to a two-dimensional array will be more dramatic. In the large N limit, the magnitude of improvement will scale as $\sqrt{N}$.

Fast gates have not been well-studied in two-dimensional arrays, and in order to study their performance in these architectures I have developed an efficient search scheme for numerical optimisation. The Antisymmetric Pulse Group (APG) scheme builds upon a more generalised gate scheme where optimisation is done over the number of pulse pairs in each pulse group of a fast gate, and imposes an anti-symmetry constraint. I find that the anti-symmetry constraint improves gate optimisations as it results in a simplification of the cost function, and is well suited to optimisations

of systems with more than two ions. The development of the APG scheme enables tractable optimisations of fast gates in two-dimensional trapped ion arrays.

I found that fast gates can be optimised for two-dimensional arrays of ions using the APG scheme, without requiring significantly higher repetition rates than in two-ion systems. In contrast to previous studies that have found gates between nearest-neighbour qubits to always be preferable in one-dimensional chains of ions [36], I showed that in two-dimensional arrays it is optimal to, where possible, connect ions via diagonal gates. Importantly, I demonstrated that the fidelity of fast gate operations does not decay as the size of the two-dimensional array is scaled: gates optimised for a simple 2×2 array are able to achieve ultra-high fidelities (upwards of $1 - 10^{-7}$) even in very large crystals with upwards of 100 ions. This enables this platform to be suitable for large-scale computation, without requiring restrictively intense optimisations of the gate scheme. This extends previous studies of fast gates, which were limited to one-dimensional ion chains and largely focused on two-ion systems [1, 2, 33, 35–38].

Based on previous analyses identifying errors in pulse area to be the dominant source of error for fast gates, I performed a worst case analysis to quantify the impact on gate fidelity. Unfortunately, I found that the transition error required to implement high-fidelity gates involving $\sim$1000 pulses to be restrictively smaller than has currently been demonstrated, on the order of $10^{-6} - 10^{-8}$. Achieving these low errors with primitive square pulses is not feasible, as they are very sensitive to fluctuations in laser intensity. I have suggested how these error thresholds can be experimentally achieved using well-studied techniques from other fields, notably the use of pulse shaping and rapid adiabatic passage using chirped laser pulses. These techniques allow for increased robustness against fluctuations in laser intensity and other pulse imperfections, and are likely reduce transition errors to a regime relevant to fast gates in microtrap architectures.

I considered also how fast gate performance changes for different choices of candidate ions and the transition used for the momentum kicks. While some ion species are more naturally suited to fast gates, in that they require lower resolving repetition rates to implement ultra-fast gates with high-fidelity, I showed that fast gates can be effectively performed with any ion species. Even for the heaviest ions, microsecond gates were shown to be possible that require only 100s of MHz in laser repetition rate, which are already available in current experiments [50, 70].

Finally, I presented a series of experimental recommendations on choosing the candidate ion and trapping architecture, optimising the gate scheme for the specific nature of the experiment, and implementing the counter-propagating pulse pairs required

for fast gates by pulse-splitting techniques. This is based on both investigations in this thesis and analyses that currently exist in the literature.

8.2 Discussion

The results presented in this thesis have profound implications for the prospect of realising a scalable quantum information processor. Here I will assess the feasibility for realising a large-scale computation on the studied trapped ion architectures, and discuss the place of fast gates for the future of quantum computers.

8.2.1 Realistic implementation of simulation algorithm

An algorithm for simulating the Fermi-Hubbard model with 40 fermionic modes was investigated in Chapter 4 to demonstrate two things: the increased connectivity of two-dimensional architectures allows for a significant reduction in the number of gate operations, and fast gates enable the realistic implementation of the simulation on a near-future quantum computer. While the former aspect has been discussed in depth (see Section 4.3), I have deliberately delayed the discussion of the latter till now so that the error analysis of the previous chapter can be included.

For implementing the simulation algorithm on one-dimensional chain of ions, I found that 42360 fast gates are required for 10 Trotter steps. Theoretically, this requires individual gate fidelities on the order of $1-10^{-6}$ to achieve a total simulation fidelity above 90%. The results of Chapter 5 suggest that $1-2\,\mu$s gates can be implemented with experimentally feasible laser repetition rates between 500 MHz - 2 GHz, with fidelities on the order of 10^{-8} even in scaled chains of > 40 ions. This corresponds to a total simulation time of $50-100\,$ms, which is roughly the same timescale of typical trap heating rates [2]. Fortunately, trap heating is not a fundamental limitation, but a technical one as heating rates can be significantly suppressed to the order of $1\,\mathrm{s}^{-1}$ in cryrogenic traps [18].

The reduced number of gate operations that are required to implement the simulation (24680) in a two-dimensional microtrap array allows for the total time taken to be significantly smaller; $1-2\,\mu$s gates can be implemented with $500\,\mathrm{MHz}-3\,\mathrm{GHz}$ lasers with fidelities that plateau to about $1-10^{-7}$ in large crystals (as shown in Chapter 6). This places the total simulation time around $25-50\,$ms, which is a factor of $2-4$ faster than typical trap heating timescales ($\sim$100 ms) and over an order of magnitude faster than typical coherence timescales of trapped ion qubits

($\sim$1 s). A state-of-the-art 5 GHz repetition rate laser [50] enables sub-microsecond gates, and a total simulation time closer to 15 ms. A gate infidelity of 10^{-7} corresponds to a total (theoretical) simulation fidelity above 99.7%; simulation fidelities of $\sim$80% require infidelities as low as $\sim 10^{-5}$.

Of course these theoretical gate infidelities do not include any errors in the experimental implementation. Including the effect of population transfer errors (ϵ) arising from laser intensity fluctuations allows for a more realistic consideration of this simulation algorithm. Based on the worst-case analysis in the previous chapter (*c.f.* equation (7.12)), achieving a total simulation fidelity of about 80% requires the population transfer error of each pulse to be below 10^{-9}, when implemented on a one-dimensional chain of ions. For implementation on a two-dimension array, the requirement is slightly more flexible, allowing for transition errors on the order of 10^{-8}. As mentioned in the previous chapter, currently demonstrated errors in the population transfer for single pulses are only on the order of $10^{-5} - 10^{-6}$ [25] which are not sufficient for a computation of this scale.

A previous study by Taylor *et al.* [2] found large-scale computation with fast gates to be limited in the near future by restrictively laser high repetition required to implement a similar simulation algorithm on a chain of ions in a linear Paul trap. However in the two-dimensional microtrap architecture I have studied here, this is not the case. The repetition rates required to implement the simulation of the Fermi-Hubbard model, as discussed above, are well in the regime of what is currently achievable, ranging between 300 MHz and 5 GHz for implementing the simulation in tens of milliseconds. The limitation that I have identified here is the compounding effect of errors in population transfer from the π pulses, which need to be orders of magnitude lower than what is currently achievable in order to implement a large-scale computation. Fortunately, there is strong outlook for improvement in this area, particularly if the recommendations made in 7.1.2 can be implemented to make population transfers first-order insensitive to laser intensity fluctuations.

8.2.2 Speeding up the quantum processor

In this thesis I have focused on gate speeds that are faster, or very close to the trapping frequency, which generally require laser repetition rates on the order of 1 GHz. While a 5 GHz repetition rate laser designed specifically for realising fast gates has indeed been experimentally demonstrated [50], it is important to note that this is the absolute state-of-the-art. For most ion trap experimental groups, the range of accessible repetition rates is closer to $50-100$ MHz [107]. While fast gates can indeed

be implemented with these lower repetition rate lasers, the achievable gate times are on the order of a few trap periods, around $3 - 5\,\mu$s. This is still significantly faster than the speed of current generation two-qubit gates based on sideband resolution. The reader may then be curious; if this regime is relevant to more experimental teams, then why have I chosen to focus on regimes accessible only to a select few?

My answer, simply put, is that the goal of this thesis is to investigate an architecture that will enable the *future* of quantum computing with trapped ions. I have focused on gates faster than the trap frequency specifically because it is a regime where current (sideband-resolving) entangling gates are fundamentally unable to access. More importantly, it is exactly in this regime that near-future quantum computation will need to take place. Until quantum computers reach the scales required to implement fault tolerant error correction routines, there will always be a desire for faster operations so larger computations can be done within the time limit imposed by decoherence. In this thesis I have investigated what can be achieved with the best of current laser technologies, which lays the groundwork for realising large-scale computation in the near future. All the results I have presented can be straightforwardly generalised to gate times achievable with lower repetition rate lasers via the scaling relation $f_{rep} \propto \tau_G^{-2.5}$.

8.3 Avenues for future work

This thesis can be considered a blueprint for scalable quantum computing with trapped ions, but is certainly not a complete description of all the detail that is required for its realisation. Certainly there are several avenues for future inquiry that need to be pursued, particularly as more experimental groups become interested in realising fast gates.

Perhaps the most natural extension of the analyses presented in this thesis is the inclusion of an ODE description of the ions motion in gate design. This follows from the suggestions for gate design in Section 7.3. In this thesis I have assumed that gate fidelity will not be significantly affected by finite repetition rate as long as pulse groups do not overlap, and any losses in fidelity can be restored by local optimisation using the ODE description. Furthermore it has been assumed that any losses in fidelity from linearising the Coulomb interaction can be similarly accounted for by ODE optimisation. It is important these assumptions are validated by explicitly including these effects in future work to ensure robustness of fast gates, particularly for in two-dimensional arrays where they have not been previously studied.

An ODE description of fast gate dynamics would also enable the inclusion of micromotion in gate design. Recent work by Ratcliffe *et al.* [68] has demonstrated that micromotion can significantly damage fast gate fidelity if unaccounted for, but if included in the gate design, it can lead to an *enhancement* of gates designed under the FRAG scheme. Extending this analysis with the APG scheme developed in this thesis promises a further improvement of gate speed and fidelity. This extension is particularly important for two-dimensional trapped ion architectures, where micromotion is necessarily present in at least one axis of the traps. If the microtrap design is similar to the cylindrical ring traps used in the two-dimensional trap arrays reported in Ref. [55], then micromotion is present symmetrically in all directions, and thus it is extremely important that it is accounted for in gate design.

Another avenue of future work is the study of fast gates in alternative two-dimensional architectures, such as one with rows of adjacent linear Paul traps each containing some number of ions. Fast gates between ions in adjacent traps will likely be similar to fast gates between ions in individual microtraps, however it is unclear how these architectures will scale and what the optimal number of ions per trap should be. This kind of architecture may be well suited to implementing some sort of error-correction code, with the possibility of ions in a particular linear Paul trap functioning as a single logical qubit. Speaking more generally, the scale at which microtrap architectures are preferable to linear Paul trap architectures is unclear. As most current ion trap experiments use linear Paul traps, it is of great experimental relevance to report this threshold.

From a practical standpoint, the most pressing avenue of future work is further investigation into challenges that are likely to occur in experimental implementations of fast gates, and the development of methods to mitigate these issues. As pulse stability has been identified as a likely source of difficulty, it would be of use to study the methods suggested for improving pulse imperfections in greater depth. In particular, a further study of composite pulse sequences for suppressing the effect of laser intensity fluctuations. It is possible that composite pulse sequences can be designed specifically for fast gate operations using the methods outlined in [99], which would promise increased robustness from correlated errors between subsequent fast gate operations in a large-scale computation.

8.4 Outlook

Fast gates are a promising candidate for large-scale computation on trapped ion quantum computers. Given arbitrary laser power and stability, gate speeds well

beyond any practical requirements can be achieved with extremely high-fidelity. Furthermore, they can be optimised for any trap geometry, and are well suited to large-scale computation in microtrap arrays.

However, fast gates are very sensitive to errors in population transfer in individual pulses. For square pulses, the level to which intensity fluctuations need to be suppressed is incredibly restrictive. This is a technological barrier that places a large cost on experimentally realising fast gate schemes, both in terms of time and money. Fortunately, there is great hope in improving robustness of fast gate operations to intensity fluctuations by exploiting techniques from the world of NMR and beyond. The challenge is now to implement these techniques for picosecond pulses in the optical and ultra-violet frequencies. Doing so is likely to propel fast gates to the forefront of trapped ion quantum computing, and pave the way for large-scale computation in the near future.

Appendix A

Deriving the Jaynes-Cumming Hamiltonian

In this Appendix, I will derive the Jaynes-Cumming Hamiltonian, which describes the interaction between a single quantised mode of light and a two-level atom.

Recall that in the absence of interaction, the Hamiltonian for an atom-light system can be decomposed as

$$\hat{H}_0 = \hat{H}_{el} + \hat{H}_{light} \tag{A.1}$$

where H_{el} is the internal Hamiltonian describing the electronic states of the atom and H_{light} is the Hamiltonian of a light-field in free space. For a two-level atom the electronic Hamiltonian is given by

$$\hat{H}_{el} = \frac{1}{2}\hbar\omega_a\hat{\sigma}_z\,. \tag{A.2}$$

In full generality, the Hamiltonian for a light field in free space can be written as

$$\hat{H}_{light} = \sum_{\lambda}\int d^3\mathbf{k}\,\hbar\omega_{\mathbf{k}}\left(\hat{c}^{\dagger}_{\mathbf{k},\lambda}\hat{c}_{\mathbf{k},\lambda} + \frac{1}{2}\right) \tag{A.3}$$

where $\mathbf{k}$ is the wave-vector of an optical mode with frequency $\omega_{\mathbf{k}}$, and $\hat{c}^{\dagger}_{\mathbf{k},\lambda}$ ($\hat{c}_{\mathbf{k},\lambda}$) creates (annihilates) a photon with wave-vector $\mathbf{k}$ and polarisation λ.

The dipole approximation

In the regime where the wavelength of the light field is significantly larger than the spatial extent of the atomic wave-function (determined by the Bohr radius), the interaction between a two-level atom and light can be described by the interaction Hamiltonian

$$\hat{H}_{int} = -q\,\hat{\mathbf{r}}\cdot\hat{\mathbf{E}}\,, \tag{A.4}$$

where $\hat{\mathbf{E}}$ is the electric field observable. This is known as the *dipole approximation* [108], which holds for optical frequencies typically used in trapped ion experiements. In the basis of energy eigenstates of the bare Hamiltonian $|n\rangle$, the term can be expressed as

$$\hat{H}_{int} = -q \left(\sum_{i,j} \langle i| \hat{\mathbf{r}} |j\rangle |i\rangle \langle j| \right) \cdot \hat{\mathbf{E}} \tag{A.5}$$

where the diagonal elements vanish due to the parity of the electronic eigenstates (this is a consequence of the symmetry of any central potential). For a two-level system of a ground state $|g\rangle$ and excited state $|e\rangle$ this then becomes

$$\hat{H}_{int} = -\left(\mathbf{d}_{eg}\hat{\sigma}_+ + \mathbf{d}^*_{eg}\hat{\sigma}_-\right) \cdot \hat{\mathbf{E}} \tag{A.6}$$

$$\tag{A.7}$$

where I have defined the coefficient $\mathbf{d}_{eg} = q \langle e| \hat{\mathbf{r}} |g\rangle$ and the raising and lowering operators on the internal states $\hat{\sigma}_+ \equiv |e\rangle \langle g|$ and $\hat{\sigma}_- \equiv (\hat{\sigma}_+)^\dagger$.

The electric field observable can be expressed in the basis of energy eigenstates of the light field as

$$\hat{\mathbf{E}}(\hat{\mathbf{r}}) = \sum_k \int d^3\mathbf{k} \sqrt{\frac{\hbar\omega_k}{2\epsilon_0}} \left(\hat{c}_{\mathbf{k},\lambda}\mathbf{u}_{\mathbf{k},\lambda}(\hat{\mathbf{r}}) + \hat{c}^\dagger_{\mathbf{k},\lambda}\mathbf{u}^*_{\mathbf{k},\lambda}(\hat{\mathbf{r}})\right), \tag{A.8}$$

where $\hat{c}^\dagger_{\mathbf{k},\lambda}$ and $\hat{c}_{\mathbf{k},\lambda}$ are the respective creation and annihilation operators for a plane wave with wave-vector $\mathbf{k}$ and polarisation λ (which are eigenstates of $\hat{H}_{light}$). Here $\mathbf{u}_{\mathbf{k},\lambda}$ are the corresponding eigenstate wave-functions

$$\mathbf{u}_{\mathbf{k},\lambda}(\hat{\mathbf{r}}) = \mathbf{E}_\lambda e^{i(\mathbf{k}\cdot\hat{\mathbf{r}})} \tag{A.9}$$

with vector-valued amplitude $\mathbf{E}_\lambda$ that is dependent on the polarisation.

The Jaynes-Cumming Hamiltonian

I will now consider a light-field dominated by a single optical mode with frequency ω_L. Furthermore, I will choose a coordinate basis in which the field is polarised along the y axis and travelling transversely along the x axis. The interaction term of the Hamiltonian is then

$$\hat{H}_{int} = -E_y \sqrt{\frac{\hbar\omega_L}{2\epsilon_0}} \left(d_{eg}\hat{\sigma}_+ + d^*_{eg}\hat{\sigma}_-\right)\left(\hat{c}\, e^{ik\hat{x}} + \hat{c}^\dagger e^{-ik\hat{x}}\right) \tag{A.10}$$

where the index on the creation and annihilation operators is suppressed as there is only one mode under consideration. Here $\hat{x}$ is the position of the ion within the light field, with respect to some steady-state equilibrium location. It is common to take d_{eg} to be real valued, and absorb all the constants into a single coupling constant g, such that

$$\hat{H}_{int} = \hbar g \left(\hat{\sigma}_+ + \hat{\sigma}_-\right) \left(\hat{c}\, e^{ik\hat{x}} + \hat{c}^\dagger e^{-ik\hat{x}}\right) . \tag{A.11}$$

Moving into the interaction picture with respect to the light field, i.e.

$$\hat{H}_{int} \to e^{i\omega_L t \hat{c}^\dagger \hat{c}} \hat{H}_{int} e^{-i\omega_L t \hat{c}^\dagger \hat{c}} \tag{A.12}$$

we arrive at the Jaynes-Cumming Hamiltonian:

$$\hat{H}_{int} = \hbar g \left(\hat{\sigma}_+ + \hat{\sigma}_-\right) \left(\hat{c}\, e^{i(k\hat{x} - \omega_L t)} + \hat{c}^\dagger e^{-i(k\hat{x} - \omega_L t)}\right) . \tag{A.13}$$

Appendix B

Linearisation and normal modes

In this appendix, I will describe the calculation of the mode structure of a one-dimensional chain of trapped ions via linearisation of the Coulomb potential. This calculation can be straightforwardly generalised to two or three-dimensions. I will follow use a classical mechanics approach, following Goldstein [109].

The potential for a one-dimensional chain of ions with charge $+e$ can be expressed as

$$V = \sum_{a \neq b} \frac{e^2}{4\pi\epsilon_0} \frac{1}{r_{ab}} + \frac{1}{2} M\omega^2 \sum_i (x^i)^2 \tag{B.1}$$

where $r_{ab} \equiv |x_a - x_b|$ is the distance between the a-th and b-th ions.[1] The first step towards linearisation is finding the equilibrium positions x_0^i i ions. This can be done by brute force by solving the set of simulatenous equations

$$\left\{ \frac{\partial V}{\partial x^i} = 0 \right\}. \tag{B.2}$$

This can be understood as a minimisation of the potential. It is worth noting that only real solutions of this set of equations correspond to ions being trapped at their equilibrium positions.

In studying oscillations of the ions around equilibrium, it is sufficient to consider only small deviations q to describe the ions' positions, i.e.

$$x^i = x_0^i + q^i \,. \tag{B.3}$$

Expanding the potential around equilibrium (i.e. $\eta^i = 0$) to second order gives

$$V = V(x_0^1, x_0^2, \dots) + \left.\frac{\partial V}{\partial x^i}\right|_{\mathbf{x}_0} q^i + \frac{1}{2} \left.\frac{\partial^2 V}{\partial x^i \partial x^j}\right|_{\mathbf{x}_0} q^i q^j + \mathcal{O}(q^3) \,, \tag{B.4}$$

[1] The trap frequency ν relates to the angular trap frequency that appears in the potential and almost everywhere else as $\omega = 2\pi\nu$.

where I have adopted Einstein summation notation over repeated indicies. Note that by virtue of equation (B.2), the second term in this expression vanishes. Further, $V(x_0^1, x_0^2, \dots)$ can be neglected by choosing the potential to be 0 at equilibrium.

The classical motion of the ions can be obtained by considering the Lagrangian for the potential truncated to second order,

$$L = \frac{\mathbf{p}^2}{2M} - V(x) \tag{B.5}$$

$$\approx \frac{1}{2}\left(\sum_i M(\dot{q}^i)^2 + V_{ij}q^i q^j\right), \tag{B.6}$$

where I have defined the coefficients $V_{ij} \equiv \frac{\partial^2 V}{\partial x^i x^j}\Big|_0$. Dynamics can be obtained directly from Euler-Lagrange equations

$$\frac{d}{dt}\left(\frac{\partial L}{\partial \dot{x}^i}\right) = \frac{\partial L}{\partial x^i}. \tag{B.7}$$

For the potential truncated to second order, this leads to the equations

$$\ddot{q}^i = \frac{V_{ij}}{M}q^j, \tag{B.8}$$

noting the sum over the j index on the right hand side. Assuming there exist p oscillatory solutions of the form $q_p^i = b_p^i e^{i(\omega t+\varphi)}$ for each of the p motional modes, this reduces to

$$\frac{V_{ij}}{M}b_p^j = \omega_p^2 b_p^i. \tag{B.9}$$

It is convenient here to introduce the Hessian matrix $\underline{\mathbf{H}}$ as the matrix defined by coefficients $\frac{1}{M}V_{ij}$. Equation (B.9) then can be expressed as an *eigenvalue problem*,

$$\underline{\mathbf{H}} \cdot \mathbf{b_p} = \omega_p^2 \mathbf{b_p}. \tag{B.10}$$

Therefore, calculating the mode mode structure reduces to simply solving for the eigensystem of the Hessian matrix: the eigenvectors give the coupling vectors $\mathbf{b_p}$ and the eigenvalues give the squared mode frequencies ω_p^2.

Appendix C

Condition equations for fast gate implementation

Authors note: This appendix derives the conditions for implementation of a controlled phase gate with a fast gate operation, following [33, 36]. In traditional style of derivations, I will use 'we' to refer to myself and the readers following along. I will also omit putting hats on operators to avoid convoluted expressions.

Here we derive the conditions for fast gates acting on two ions, labelled 1 and 2, of a chain of ions. We begin by considering the atom-light Hamiltonian in the interaction picture for a single ion (3.22), which for resonant transitions ($\Delta = 0$) is given by

$$H = \frac{\hbar\Omega(t)}{2}\left(\sigma_+ e^{i(kx+\phi)} + \sigma_- e^{-i(kx+\phi)}\right). \tag{C.1}$$

For a constant Rabi frequency $\Omega(t) = \Omega$ (corresponding to a square pulse), this generates the unitary

$$U(\Omega, k) = \begin{pmatrix} \cos(\Omega t/2) & -i\sin(\Omega t/2)e^{-i(kx+\phi)} \\ -i\sin(\Omega t/2)e^{i(kx+\phi)} & \cos(\Omega t/2)\,. \end{pmatrix} \tag{C.2}$$

For a π-pulse (i.e. $\Omega t = \pi$) this gives the unitary

$$U_\pi(k) = \begin{pmatrix} 0 & e^{-i(kx+\phi)} \\ e^{i(kx+\phi)} & 0\,. \end{pmatrix} \tag{C.3}$$

where I have discarded $-i$ as it contributes only a global phase. This can equivalently be expressed as

$$U_\pi(k) = \sigma_+ e^{-i(kx+\phi)} + \sigma_- e^{i(kx+\phi)}\,. \tag{C.4}$$

These gates are composed of momentum kicks from groups of counter propagating π-pulse pairs. The kicks are assumed to be fast relative to the motion of the ions,

and are thus treated as instantaneous. Each kick performs the following evolution on the state,

$$U_{kick} = \Big(U_\pi(-k)U_\pi(+k)\Big)^{2z}_{\text{ion }1}\Big(U_\pi(-k)U_\pi(+k)\Big)^{2z}_{\text{ion }2} = \mathrm{e}^{-i2zk(x_1\sigma_1^z + x_2\sigma_2^z)} \tag{C.5}$$

where z is the number of pulse pairs in the pulse group, k is the wavenumber of the laser, x_1 and x_2 are the positions operators of each ion, and σ^z is the Pauli-z operator. Note that the use of σ^z relates the inversion of the state by the π-pulse, and the factor of 2 comes from the fact that they come in pairs. Also, it is notable that the phase ϕ cancels out between the forward and counter-propagating pulses.

The position of the ion can be written in terms of the mode position operator Q_p, and the mode couplings $b_i^{(p)}$,

$$x_i = \sum_{p=1}^{L} b_i^{(p)} Q_p \tag{C.6}$$

where L is the total number of motional modes of the trapped ion system. Q_p can be written in terms of the mode annihilation and creation operators, and the Lamb-Dicke parameter η_p[1].

$$Q_p = \frac{\eta_p}{k}(a_p^\dagger + a_p) \tag{C.7}$$

The evolution of the state from the momentum kicks can thus be expressed in terms of the displacement operator for the L motional modes

$$\begin{aligned} U_{kick} &= \exp\left[-i2zk\sum_{p=1}^{L}(b_1^{(p)}\sigma_1^z + b_2^{(p)}\sigma_2^z)Q_p\right] &\text{(C.8)}\\ &= \prod_{p=1}^{L}\exp\left[-i2z\eta_p(b_1^{(p)}\sigma_1^z + b_2^{(p)}\sigma_2^z)(a_p + a_p^\dagger)\right]. &\text{(C.9)} \end{aligned}$$

At this point it becomes convenient to introduce the displacement operator for each motional mode:

$$D_p(\alpha) = \exp\left[\alpha a_p^\dagger - \alpha^* a_p\right]. \tag{C.10}$$

It is then clear that U_{kick} can be expressed by a product of displacement operators,

$$U_{kick} = \prod_{p=1}^{L} D_p\left(-i2z\eta_p(b_1^{(p)}\sigma_1^z + b_2^{(p)}\sigma_2^z)\right). \tag{C.11}$$

In between momentum kicks, the motional state is allowed to freely evolve, which

[1]Note that this Lamb-Dicke parameter is *mode-dependent*, defined as $\eta_p = \sqrt{\frac{\hbar}{2M\omega_p}}$. This is in contrast with the definition used in the main body of this thesis, which is independent of modes. The two are perfectly equivalent, and related by $\eta_p = \sqrt{\frac{\omega_t}{\omega_p}}\eta$

corresponds to a rotation in the phase space of each mode.

$$U_{mot,p} = \exp\left[-i\omega_p \delta t_k a_p^\dagger a_p\right] \tag{C.12}$$

Here ω_p is the frequency of the pth mode, and δt_k is the time between the kth and $(k+1)$th kicks.

The total gate is given by N momentum kicks, each composed of z_k pulse pairs, with free evolution between pulse groups:

$$U_{gate} = \prod_{k=1}^{N} \prod_{p=1}^{L} D_p\left(-ic_{pk}\right) \exp\left[-i\omega_p \delta t_k a_p^\dagger a_p\right] = \prod_{p=1}^{L} \mathcal{U}_p, \tag{C.13}$$

where $\mathcal{U}_p$ is the evolution of the pth mode, and the coefficients c_{pk} are defined as

$$c_{pk} \equiv 2 z_k \eta_p (b_1^{(p)} \sigma_1^z + b_2^{(p)} \sigma_2^z). \tag{C.14}$$

Let us consider the action of $\mathcal{U}_p$ on a single-mode coherent state $|\alpha\rangle$,

$$\begin{aligned} \mathcal{U}_p |\alpha\rangle &= \prod_{k=1}^{N} D_p(-ic_{pk}) \exp\left[-i\omega_p \delta t_k a_p^\dagger a_p\right] |\alpha\rangle & \text{(C.15)} \\ &= D_p(-ic_{pN}) e^{-i\omega_p \delta t_N a_p^\dagger a_p} \ldots D_p(-ic_{p1}) e^{-i\omega_p \delta t_1 a_p^\dagger a_p} |\alpha\rangle & \text{(C.16)} \end{aligned}$$

Lets consider at the first term in this series of products.

$$D_p(-ic_{p1}) e^{-i\omega_p \delta t_1 a_p^\dagger a_p} |\alpha\rangle \tag{C.17}$$

First we note that, $e^{-i\omega_p \delta t_1 a_p^\dagger a_p} |\alpha\rangle = \left|e^{-i\omega_p \delta t_1} \alpha\right\rangle$. This can be easily shown by expanding the coherent state as an infinite sum of number states and noting $a_p^\dagger a_p |n\rangle = n |n\rangle$.

$$D_p(-ic_{p1}) e^{-i\omega_p \delta t_1 a_p^\dagger a_p} |\alpha\rangle = D_p(-ic_{p1}) \left|e^{-i\omega_p \delta t_1} \alpha\right\rangle \tag{C.18}$$

Next we note that any coherent state can be described in terms of a displacement operator acting on the vacuum:

$$|\alpha\rangle = D(\alpha) |0\rangle \tag{C.19}$$

$$\Rightarrow D_p(-ic_{p1}) \left|e^{-i\omega_p \delta t_1} \alpha\right\rangle = D_p(-ic_{p1}) D_p\left(e^{-i\omega_p \delta t_1} \alpha\right) |0\rangle . \tag{C.20}$$

Finally, we consider fact that the product of two displacement operators is a dis-

placement operator with an added phase factor,

$$D(\alpha)D(\beta) = e^{\frac{1}{2}(\alpha\beta^* - \alpha^*\beta)} D(\alpha+\beta) = e^{2i\Im\mathfrak{m}(\alpha\beta^*)} D(\alpha+\beta) \tag{C.21}$$

Now lets return to equation (C.18), and use these notes to evaluate the impact of the kick on the coherent state.

$$\begin{aligned} D_p(-ic_{p1}) \left| e^{-i\omega_p \delta t_1}\alpha \right\rangle &= D_p(-ic_{p1}) D_p\left(e^{-i\omega_p\delta t_1}\alpha\right)|0\rangle & \text{(C.22)} \\ &= e^{i\Im\mathfrak{m}(-ic_{p1}e^{i\omega_p\delta t_1}\alpha^*)} D\left(-ic_{p1} + e^{-i\omega_p\delta t_1}\alpha\right)|0\rangle & \text{(C.23)} \\ &= e^{-i\Im\mathfrak{m}(ic_{p1}e^{i\omega_p\delta t_1}\alpha^*)} \left| -ic_{p1} + e^{-i\omega_p\delta t_1}\alpha \right\rangle & \text{(C.24)} \end{aligned}$$

This suggests the the entire sequence of displacements and rotations acting on $|\alpha\rangle$ will result in a new coherent state and a phase factor ξ_p.

$$\mathcal{U}_p |\alpha\rangle = e^{i\xi_p} |\tilde{\alpha}\rangle \tag{C.25}$$

For simplicity, we will first find an expression for $|\tilde{\alpha}\rangle$. Expanding the first few displacements and rotations, and ignoring the phase factor from (C.21) gives:

$$\begin{aligned} |\tilde{\alpha}\rangle &= \ldots D_p(-ic_{p3}) e^{-i\omega_p\delta t_3 a_p^\dagger a_p} D_p(-ic_{p2}) e^{-i\omega_p\delta t_2 a_p^\dagger a_p} D_p(-ic_{p1}) e^{-i\omega_p\delta t_1 a_p^\dagger a_p} |\alpha\rangle \\ &= \ldots D_p(-ic_{p3}) e^{-i\omega_p\delta t_3 a_p^\dagger a_p} D_p(-ic_{p2}) e^{-i\omega_p\delta t_2 a_p^\dagger a_p} \left| -ic_{p1} + e^{-i\omega_p\delta t_1}\alpha \right\rangle & \text{(C.26)} \\ &= \ldots D_p(-ic_{p3}) e^{-i\omega_p\delta t_3 a_p^\dagger a_p} \left| -ic_{p2} + e^{-i\omega_p\delta t_2}(-ic_{p1} + e^{-i\omega_p\delta t_1}\alpha) \right\rangle & \text{(C.27)} \\ &= \ldots \left| -ic_{p3} + e^{-i\omega_p\delta t_3}(-ic_{p2} + e^{-i\omega_p\delta t_2}(-ic_{p1} + e^{-i\omega_p\delta t_1}\alpha)) \right\rangle & \text{(C.28)} \end{aligned}$$

By expanding out the brackets in the last line, we can infer a pattern for the impact of the total sequence of displacement and rotations.

$$|\tilde{\alpha}\rangle = \left| -i\sum_{k=1}^{N} c_{pk} e^{-i\omega_p \sum_{j=k}^{N} \delta t_j} + \alpha e^{-i\omega_p \sum_{k=1}^{N} \delta t_k} \right\rangle \tag{C.29}$$

If we are considering the momentum kicks to be instantaneous, we have that $\sum_{j=1}^{m} \delta t_j = t_m$, which is to say the time up to the $(m+1)$th kick.

Therefore, we have

$$\tilde{\alpha} = \alpha e^{-i\omega_p T_G} - i\sum_{k=1}^{N} c_{pk} e^{i\omega_p(t_k - T_G)} \tag{C.30}$$

as the coherent state produced, $\mathcal{U}_p |\alpha\rangle \to |\tilde{\alpha}\rangle$, where T_G is total gate time.

$$\sum_{k=1}^{N} \delta t_k = T_G. \tag{C.31}$$

Now lets look at the associated phase, which recall comes from the product of two displacement operators (C.21). Again, let us expand the first few displacements and rotations and find a pattern.

$$\begin{aligned}
\mathcal{U}_p |\alpha\rangle &= D_p(-ic_{pN}) e^{-i\omega_p \delta t_N a_p^\dagger a_p} \dots D_p(-ic_{p1}) e^{-i\omega_p \delta t_1 a_p^\dagger a_p} |\alpha\rangle && \text{(C.32)} \\
&= e^{i\mathfrak{Im}(-ic_{p1} e^{i\omega_p \delta t_1 \alpha^*})} \dots D_p(-ic_{p2}) e^{-i\omega_p \delta t_2 a_p^\dagger a_p} \left| -ic_{p1} + e^{-i\omega_p \delta t_1} \alpha \right\rangle && \text{(C.33)} \\
&= \dots e^{i\mathfrak{Im}(-ic_{p3} + e^{i\omega_p \delta t_3}(ic_{p2} + e^{i\omega_p \delta t_2}(ic_{p1} + e^{i\omega_p \delta t_1}\alpha^*)))} && \text{(C.34)} \\
& \quad e^{i\mathfrak{Im}(-ic_{p2} + e^{i\omega_p \delta t_2}(ic_{p1} + e^{i\omega_p \delta t_1}\alpha^*))} e^{i\mathfrak{Im}(-ic_{p1} e^{i\omega_p \delta t_1}\alpha^*)} |\tilde{\alpha}\rangle && \text{(C.35)}
\end{aligned}$$

By expanding the products of exponentials in the last line, we can infer the following pattern for the phase,

$$\xi_p = \mathfrak{Im}\left(\sum_{k=1}^{N} (-ic_{pk}) e^{i\omega_p \sum_{m=1}^{k} \delta t_k} \alpha^* + \sum_{l=2}^{N} \sum_{m=1}^{l-1} (-ic_{pl})(ic_{pm}) e^{i\omega_p \sum_{q=m+1}^{l} \delta t_k} \right) \quad \text{(C.36)}$$

which, noting that $\mathfrak{Im}(e^{i\theta}) = \sin\theta$, reduces to:

$$\xi_p = \sum_{m=2}^{N} \sum_{k=1}^{m-1} c_{pk} c_{pm} \sin\left(\omega_p (t_m - t_k)\right) - \mathfrak{Re}\left(\alpha \sum_{k=1}^{N} c_{pk} e^{i\omega_p t_k} \right). \quad \text{(C.37)}$$

In the ideal case, we demand that the motional state after the gate's evolution of the state is the same as if there were no kicks, i.e. that net displacement of each motional mode is zero.

$$\begin{aligned}
\sum_{k=1}^{N} c_{pk} e^{i\omega_p t_k} = -i \sum_{k=1}^{N} 2 z_k \eta_p \left(b_1^{(p)} \sigma_1^z + b_2^{(p)} \sigma_2^z \right) e^{i\omega_p t_k} &= 0 && \text{(C.38)} \\
\Rightarrow \sum_{k=1}^{N} z_k e^{i\omega_p t_k} &= 0 && \text{(C.39)}
\end{aligned}$$

This condition removes the phase term $e^{i\mathfrak{Re}\left(\sum_{k=1}^{N} c_{pk} e^{i\omega_p t_k} \alpha\right)}$. Noting this condition, we expand (C.37) for the remaining phase term.

$$\begin{aligned}
\xi_p &= \sum_{m=2}^{N} \sum_{k=1}^{m-1} c_{pk} c_{pm} \sin\left(\omega_p (t_m - t_k)\right) && \text{(C.40)} \\
&= 4 z_m z_k \eta_p^2 \left((b_1^{(p)} \sigma_1^z)^2 + (b_2^{(p)} \sigma_2^z)^2 + 2 b_1^{(p)} b_2^{(p)} \sigma_1^z \sigma_2^z \right) && \text{(C.41)}
\end{aligned}$$

We are only concerned with the terms that are dependent on the internal state of both ions involved in the gate, so we can ignore the global phase terms, and thus,

$$\xi_p = 8 \eta_p^2 b_1^{(p)} b_2^{(p)} \sigma_1^z \sigma_2^z \sum_{m=2}^{N} \sum_{k=1}^{m-1} z_m z_k \sin\left(\omega_p (t_m - t_k)\right). \quad \text{(C.42)}$$

Ideally, the product of kicks and displacements for each mode will result in the evolution of the state according to the ideal unitary,

$$\prod_{p=1}^{L} e^{i\xi_p} = e^{i\Theta\sigma_1^z\otimes\sigma_2^z}. \tag{C.43}$$

We would like to implement the controlled-phase gate, which has $\Theta = \frac{\pi}{4}$:

$$\frac{\pi}{4} = 8\sum_{p=1}^{L}\eta_p^2 b_1^{(p)} b_2^{(p)} \sum_{m=2}^{N}\sum_{k=1}^{m-1} z_m z_k \sin\left(\omega_p(t_m - t_k)\right). \tag{C.44}$$

This equation, together with our demand for motional restoration (C.38), gives a set of $L+1$ conditions for the implementation of a maximally entangling controlled-phase gate between two qubits.

Appendix D

Jordan-Wigner transformation of Fermi-Hubbard Hamiltonian

D.1 Jordan Wigner transform

In order to 'digitize' a quantum system on a quantum computer for the purposes of simulation, its Hamiltonian needs to be mapped to Pauli operators that act on qubits. For a fermionic system form of the mapping is non-trivial, as the fermionic anticommutation relations require satisfaction even in mapped form. One such mapping that is commonly used is the Jordan-Wigner transformation, which has the following form:

$$b_j = -\left(\bigotimes_{n=1}^{j-1} \sigma_z^n\right) \otimes \sigma_-^j \tag{D.1}$$

$$b_j^\dagger = -\left(\bigotimes_{n=1}^{j-1} \sigma_z^n\right) \otimes \sigma_+^j \,. \tag{D.2}$$

Note that this form of transformation shows explicitly that j fermionic modes can be mapped to j interacting qubits. By substitution, one can check that this transformation maintains the anticommutation relations $\{b_j, b_k^\dagger\} = \delta_{jk}$. In doing so, it is useful to keep in mind some of the properties of the Pauli operators:

- $\sigma_m^2 = \mathbb{1}$ for $m = x, y, z$
- $[\sigma_a, \sigma_b] = 2i\varepsilon_{abc}\sigma_c$ where ε_{abc} is the Levi-Civita tensor
- $\sigma_a\sigma_b = \delta_{ab} + i\varepsilon_{abc}\sigma_c$

D.2 Mapping the Fermi-Hubbard hamiltonian

In this appendix, I will consider mapping of the Fermi-Hubbard hamiltonian introduced in Chapter 4. Recall that in second quantised form, the Hamiltonian can be expressed:

$$H = w \sum_{\langle i,j\rangle,\sigma}^{20} (b^\dagger_{i,\sigma} b_{j,\sigma} + \text{h.c.}) + U \sum_{j}^{20} b^\dagger_{j,\uparrow} b_{j,\uparrow} b^\dagger_{j,\downarrow} b_{j,\downarrow}\,, \tag{D.3}$$

where $\sigma = \uparrow, \downarrow$ and $\langle i, j\rangle$ denotes pairing between nearest neighbour sites. This represents a *5 by 4* lattice, with each site supporting up to two fermions: a spin up, and spin down.

There are three distinctly different types of terms in this Hamiltonian (D.3):

1. Row tunnelling terms, e.g. $b^\dagger_{3,\uparrow} b_{4,\uparrow} + \text{h.c.}$
2. Column tunnelling terms, e.g. $b^\dagger_{5,\uparrow} b_{10,\uparrow} + \text{h.c.}$
3. Interaction potential terms, e.g. $b^\dagger_{5,\uparrow} b_{5,\uparrow} b^\dagger_{5,\downarrow} b_{5,\downarrow}$

In order to do the Jordan-Wigner transformation of each term, it is useful to combine the fermion lattice site and spin index together. For this purpose, I will implement the following change in notation:

$$b_{j,\uparrow} \rightarrow b_{2j} \tag{D.4}$$

$$b_{j,\downarrow} \rightarrow b_{2j-1}\,. \tag{D.5}$$

Row tunnelling terms

I will explicitly work through the mapping of the row-tunnelling term $b^\dagger_{2,\uparrow} b_{1,\uparrow} + \text{h.c.}$, from which the form of the mapped term can be extracted.

$$\begin{aligned} b^\dagger_{2,\uparrow} b_{1,\uparrow} + \text{h.c.} \rightarrow b^\dagger_4 b_2 + \text{h.c.} &= \left(\sigma^1_z \otimes \sigma^2_z \otimes \sigma^3_z \otimes \sigma^4_+\right) \otimes \left(\sigma^1_z \otimes \sigma^2_-\right) && \text{(D.6)} \\ &= \left(\sigma^2_z \otimes \sigma^2_- \otimes \sigma^3_z \otimes \sigma^4_+\right) && \text{(D.7)} \\ &= \sigma^2_- \otimes \sigma^3_z \otimes \sigma^4_+ && \text{(D.8)} \\ &= \frac{1}{2}\left(\sigma^2_x \otimes \sigma^3_z \otimes \sigma^4_x + \sigma^2_y \otimes \sigma^3_z \otimes \sigma^4_y\right) && \text{(D.9)} \end{aligned}$$

There are terms like $i\sigma^2_x \sigma^3_z \sigma^4_y$ that have been omitted in the expansion between the last two lines because they exactly cancel with their hermitian conjugate pair. This pattern is present for all of the row-tunnelling terms, which appear in the mapped

Hamiltonian as:

$$\frac{w}{2}\sum_{j=2}^{39}\left(\sigma_x^{j-1}\otimes\sigma_z^{j}\otimes\sigma_x^{j+1}+\sigma_y^{j-1}\otimes\sigma_z^{j}\otimes\sigma_y^{j+1}\right). \tag{D.10}$$

Column tunnelling terms

The column-tunnelling terms in the Hamiltonian have a similar mapping, however are slightly more complicated by the fact that difference between the nearest-neighbour site indexes are larger. This will essentially result in more σ_z's wedged between the σ_x's and σ_y's, by comparison to the row-tunneling terms.

For example, for the term $b^\dagger_{1,\uparrow}b_{6,\uparrow}+\text{h.c.}$, the expansion looks remarkably similar

$$\begin{aligned} b^\dagger_{1,\uparrow}b_{6,\uparrow}+\text{h.c.}\rightarrow b_2^\dagger b_{12}+\text{h.c.} &= \sigma_-^2\otimes\sigma_z^3\otimes\cdots\otimes\sigma_z^{11}\otimes\sigma_+^{12}+\text{h.c.} \\ &= \frac{1}{2}\left(\sigma_x^2\otimes\sigma_z^3\otimes\cdots\otimes\sigma_z^{11}\otimes\sigma_x^{12}+\sigma_y^2\otimes\sigma_z^3\otimes\cdots\otimes\sigma_z^{11}\otimes\sigma_y^{12}\right). \end{aligned} \tag{D.11}$$

The column tunnelling terms then have the following general form in the mapped Hamiltonian:

$$\frac{w}{2}\sum_{j=1}^{30}\left(\sigma_x^j\bigotimes_{k=j+1}^{j+9}\sigma_z^k\otimes\sigma_x^{j+10}+\sigma_y^j\bigotimes_{k=j+1}^{j+9}\sigma_z^k\otimes\sigma_y^{j+10}\right). \tag{D.12}$$

Potential terms

The onsite-potential energy terms of the Hamiltonian are perhaps the simplest to demonstrate the mapping for. This is because these terms are quadratic in the fermionic operators, and thus all σ_z contributions cancel (noting $\sigma_z^2=\mathbb{1}$).

$$\begin{aligned} U\sum_j^{20}b^\dagger_{j,\uparrow}b_{j,\uparrow}b^\dagger_{j,\downarrow}b_{j,\downarrow} &\rightarrow U\sum_j^{20}(\sigma_+^{2j}\otimes\sigma_-^{2j})\otimes(\sigma_+^{2j-1}\otimes\sigma_-^{2j-1}) \\ &= \frac{U}{8}\sum_j^{20}\left(1+\sigma_z^{2j}\right)\otimes\left(1+\sigma_z^{2j-1}\right) \\ &= \frac{U}{4}\left(20+\sum_j^{20}\sigma_z^{2j}\otimes\sigma_z^{2j-1}+\sum_k^{40}\sigma_z^k\right) \end{aligned} \tag{D.13}$$

In the second last line I have used the property $\sigma_+\otimes\sigma_-=1+\sigma_z$.

Mapped hamiltonian

Combining each of these terms, the 20-site Fermi-Hubbard model (D.3) can be fully expressed in terms of Pauli operators on a Hilbert space of 40 qubits.

$$\begin{aligned} H &= w \sum_{\lambda=x,y} \left(\sum_{j=2}^{39} \sigma_\lambda^{j-1} \otimes \sigma_z^j \otimes \sigma_\lambda^{j+1} + \sum_{j=1}^{30} \sigma_\lambda^j \bigotimes_{k=j+1}^{j+9} \sigma_z^k \otimes \sigma_\lambda^{j+10} \right) \\ &\quad + U \left(20 + \sum_j^{20} \sigma_z^{2j} \otimes \sigma_z^{2j-1} + \sum_k^{40} \sigma_z^k \right) \end{aligned} \tag{D.14}$$

Note that I have re-scaled the parameters w and U for convenience.

Bibliography

[1] C. D. B. Bentley, R. L. Taylor, A. R. R. Carvalho, and J. J. Hope. Stability thresholds and calculation techniques for fast entangling gates on trapped ions. *Phys. Rev. A*, 93:042342, Apr 2016. doi: 10.1103/PhysRevA.93.042342. URL https://link.aps.org/doi/10.1103/PhysRevA.93.042342.

[2] Richard L. Taylor, Christopher D. B. Bentley, Julen S. Pedernales, Lucas Lamata, Enrique Solano, Andre R. R. Carvalho, and Joseph J. Hope. A study on fast gates for large-scale quantum simulation with trapped ions. *Scientific Reports*, 7(1), 2017. doi: 10.1038/srep46197.

[3] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, October 1997. ISSN 0097-5397. doi: 10.1137/S0097539795293172. URL http://dx.doi.org/10.1137/S0097539795293172.

[4] Michael A. Nielsen and Isaac L. Chuang. *Quantum Computation and Quantum Information: 10th Anniversary Edition*. Cambridge University Press, New York, NY, USA, 10th edition, 2011. ISBN 1107002176, 9781107002173.

[5] Stephane Beauregard. Circuit for Shor's Algorithm using 2n+3 qubits. *Quantum Info. Comput.*, 3(2):175–185, March 2003. ISSN 1533-7146. URL http://dl.acm.org/citation.cfm?id=2011517.2011525.

[6] P. W. Shor. Fault-tolerant quantum computation. In *Proceedings of 37th Conference on Foundations of Computer Science*, pages 56–65, Oct 1996. doi: 10.1109/SFCS.1996.548464.

[7] John Preskill. Quantum Computing in the NISQ era and beyond. *Quantum*, 2:79, August 2018. ISSN 2521-327X. doi: 10.22331/q-2018-08-06-79. URL https://doi.org/10.22331/q-2018-08-06-79.

[8] Richard P. Feynman. Simulating physics with computers. *International Journal of Theoretical Physics*, 21(6-7):467–488, 1982. doi: 10.1007/bf02650179.

[9] Immanuel Bloch. Quantum simulations come of age. *Nature Physics*, 14(12):

1159–1161, December 2018. ISSN 1745-2481. doi: 10.1038/s41567-018-0371-x. URL https://www.nature.com/articles/s41567-018-0371-x.

[10] Seth Lloyd. Universal Quantum Simulators. *Science*, 273(5278):1073–1078, August 1996. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.273.5278.1073. URL https://science.sciencemag.org/content/273/5278/1073.

[11] M Müller, K Hammerer, Y L Zhou, C F Roos, and P Zoller. Simulating open quantum systems: from many-body interactions to stabilizer pumping. *New Journal of Physics*, 13(8):085007, August 2011. ISSN 1367-2630. doi: 10.1088/1367-2630/13/8/085007. URL http://stacks.iop.org/1367-2630/13/i=8/a=085007?key=crossref.c603ed99dca23916cb3304b9c5071948.

[12] Daniel S. Abrams and Seth Lloyd. Simulation of Many-Body Fermi Systems on a Universal Quantum Computer. *Physical Review Letters*, 79(13):2586–2589, September 1997. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.79.2586. URL http://arxiv.org/abs/quant-ph/9703054. arXiv: quant-ph/9703054.

[13] Lucas Lamata, Antonio Mezzacapo, Jorge Casanova, and Enrique Solano. Efficient quantum simulation of fermionic and bosonic models in trapped ions. *EPJ Quantum Technology*, 1(1), 2014. doi: 10.1140/epjqt9.

[14] Cornelius Hempel, Christine Maier, Jonathan Romero, Jarrod McClean, Thomas Monz, Heng Shen, Petar Jurcevic, Ben P. Lanyon, Peter Love, Ryan Babbush, Alán Aspuru-Guzik, Rainer Blatt, and Christian F. Roos. Quantum chemistry calculations on a trapped-ion quantum simulator. *Phys. Rev. X*, 8:031022, Jul 2018. doi: 10.1103/PhysRevX.8.031022. URL https://link.aps.org/doi/10.1103/PhysRevX.8.031022.

[15] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C. Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando G. S. L. Brandao, David A. Buell, Brian Burkett, Yu Chen, Zijun Chen, Ben Chiaro, Roberto Collins, William Courtney, Andrew Dunsworth, Edward Farhi, Brooks Foxen, Austin Fowler, Craig Gidney, Marissa Giustina, Rob Graff, Keith Guerin, Steve Habegger, Matthew P. Harrigan, Michael J. Hartmann, Alan Ho, Markus Hoffmann, Trent Huang, Travis S. Humble, Sergei V. Isakov, Evan Jeffrey, Zhang Jiang, Dvir Kafri, Kostyantyn Kechedzhi, Julian Kelly, Paul V. Klimov, Sergey Knysh, Alexander Korotkov, Fedor Kostritsa, David Landhuis, Mike Lindmark, Erik Lucero, Dmitry Lyakh, Salvatore Mandrà, Jarrod R. McClean, Matthew McEwen, Anthony Megrant, Xiao Mi, Kristel Michielsen, Masoud Mohseni, Josh Mutus, Ofer Naaman, Matthew Neeley,

Charles Neill, Murphy Yuezhen Niu, Eric Ostby, Andre Petukhov, John C. Platt, Chris Quintana, Eleanor G. Rieffel, Pedram Roushan, Nicholas C. Rubin, Daniel Sank, Kevin J. Satzinger, Vadim Smelyanskiy, Kevin J. Sung, Matthew D. Trevithick, Amit Vainsencher, Benjamin Villalonga, Theodore White, Z. Jamie Yao, Ping Yeh, Adam Zalcman, Hartmut Neven, and John M. Martinis. Quantum supremacy using a programmable superconducting processor. *Nature*, 574(7779):505–510, oct 2019. doi: 10.1038/s41586-019-1666-5. URL `https://doi.org/10.1038%2Fs41586-019-1666-5`.

[16] IBM News room. *IBM News room*, Nov 2017. URL `https://www-03.ibm.com/press/us/en/pressrelease/53374.wss`.

[17] Intel News room. *Intel Newsroom*, Jan 2018. URL `https://newsroom.intel.com/news/intel-advances-quantum-neuromorphic-computing-research/#gs.9iyre3`.

[18] G Pagano, P W Hess, H B Kaplan, W L Tan, P Richerme, P Becker, A Kyprianidis, J Zhang, E Birckelbaw, M R Hernandez, Y Wu, and C Monroe. Cryogenic trapped-ion system for large scale quantum simulation. *Quantum Science and Technology*, 4(1):014004, oct 2018. doi: 10.1088/2058-9565/aae0fe. URL `https://doi.org/10.1088%2F2058-9565%2Faae0fe`.

[19] A. Bermudez, X. Xu, R. Nigmatullin, J. O'Gorman, V. Negnevitsky, P. Schindler, T. Monz, U. G. Poschinger, C. Hempel, J. Home, F. Schmidt-Kaler, M. Biercuk, R. Blatt, S. Benjamin, and M. Müller. Assessing the progress of trapped-ion processors towards fault-tolerant quantum computation. *Phys. Rev. X*, 7:041061, Dec 2017. doi: 10.1103/PhysRevX.7.041061. URL `https://link.aps.org/doi/10.1103/PhysRevX.7.041061`.

[20] Morten Kjaergaard, Mollie E. Schwartz, Jochen BraumÃijller, Philip Krantz, Joel I-Jan Wang, Simon Gustavsson, and William D. Oliver. Superconducting qubits: Current state of play, 2019.

[21] H Haffner, C Roos, and R Blatt. Quantum computing with trapped ions. *Physics Reports*, 469(4):155–203, December 2008. ISSN 03701573. doi: 10.1016/j.physrep.2008.09.003. URL `https://linkinghub.elsevier.com/retrieve/pii/S0370157308003463`.

[22] Chad Rigetti, Jay M. Gambetta, Stefano Poletto, B. L. T. Plourde, Jerry M. Chow, A. D. Córcoles, John A. Smolin, Seth T. Merkel, J. R. Rozen, George A. Keefe, Mary B. Rothwell, Mark B. Ketchen, and M. Steffen. Superconducting qubit in a waveguide cavity with a coherence time approaching 0.1 ms. *Phys.*

Rev. B, 86:100506, Sep 2012. doi: 10.1103/PhysRevB.86.100506. URL https://link.aps.org/doi/10.1103/PhysRevB.86.100506.

[23] R. Barends, J. Kelly, A. Megrant, A. Veitia, D. Sank, E. Jeffrey, T. C. White, J. Mutus, A. G. Fowler, B. Campbell, Y. Chen, Z. Chen, B. Chiaro, A. Dunsworth, C. Neill, P. O'Malley, P. Roushan, A. Vainsencher, J. Wenner, A. N. Korotkov, A. N. Cleland, and John M. Martinis. Superconducting quantum circuits at the surface code threshold for fault tolerance. *Nature*, 508:500, April 2014. URL https://doi.org/10.1038/nature13171.

[24] Yu Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O'Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, Michael R. Geller, A. N. Cleland, and John M. Martinis. Qubit architecture with high coherence and fast tunable coupling. *Phys. Rev. Lett.*, 113:220502, Nov 2014. doi: 10.1103/PhysRevLett.113.220502. URL https://link.aps.org/doi/10.1103/PhysRevLett.113.220502.

[25] Jeremy M. Sage Robert McConnell John Chiaverini Colin D. Bruzewicz. Trapped-Ion Quantum Computing: Progress and Challenges. (arXiv:1904.04178v1 [quant-ph]). *arXiv Physics General*. doi: arXiv:1904.04178v1. URL http://arxiv.org/abs/1904.04178.

[26] V. M. Schaefer, C. J. Ballance, K. Thirumalai, L. J. Stephenson, T. G. Ballance, A. M. Steane, and D. M. Lucas. Fast quantum logic gates with trapped-ion qubits. *arXiv Quantum Physics*. doi: arXiv:1709.06952v2. URL http://arxiv.org/abs/1709.06952.

[27] Lilian Childress and Ronald Hanson. Diamond NV centers for quantum computing and quantum networks. *MRS Bulletin*, 38(2):134–138, 2013. doi: 10.1557/mrs.2013.20.

[28] Yang Wu, Ya Wang, Xi Qin, Xing Rong, and Jiangfeng Du. A programmable two-qubit solid-state quantum processor under ambient conditions. *npj Quantum Information*, 5(1):1–5, January 2019. ISSN 2056-6387. doi: 10.1038/s41534-019-0129-z. URL https://www.nature.com/articles/s41534-019-0129-z.

[29] J. I. Cirac and P. Zoller. Quantum Computations with Cold Trapped Ions. *Physical Review Letters*, 74(20):4091–4094, May 1995. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.74.4091. URL https://link.aps.org/doi/10.1103/PhysRevLett.74.4091.

[30] D.J. Wineland, C. Monroe, W.M. Itano, D. Leibfried, B.E. King, and D.M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *Journal of Research of the National Institute of Standards and Technology*, 103(3):259, May 1998. ISSN 1044677X. doi: 10.6028/jres.103.019. URL https://nvlpubs.nist.gov/nistpubs/jres/103/3/j33win.pdf.

[31] Anders Sorensen and Klaus Molmer. Entanglement and quantum computation with ions in thermal motion. *Physical Review A*, 62(2), July 2000. ISSN 1050-2947, 1094-1622. doi: 10.1103/PhysRevA.62.022311. URL http://arxiv.org/abs/quant-ph/0002024. arXiv: quant-ph/0002024.

[32] T. P. Harty, D. T. C. Allcock, C. J. Ballance, L. Guidoni, H. A. Janacek, N. M. Linke, D. N. Stacey, and D. M. Lucas. High-fidelity preparation, gates, memory, and readout of a trapped-ion quantum bit. *Phys. Rev. Lett.*, 113:220501, Nov 2014. doi: 10.1103/PhysRevLett.113.220501. URL https://link.aps.org/doi/10.1103/PhysRevLett.113.220501.

[33] J. J. Garcia-Ripoll, P. Zoller, and J. I. Cirac. Fast and robust two-qubit gates for scalable ion trap quantum computing. *Physical Review Letters*, 91 (15), October 2003. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.91.157901. URL http://arxiv.org/abs/quant-ph/0306006. arXiv: quant-ph/0306006.

[34] J. D. Wong-Campos, S. A. Moses, K. G. Johnson, and C. Monroe. Demonstration of two-atom entanglement with ultrafast optical pulses. *Phys. Rev. Lett.*, 119:230501, Dec 2017. doi: 10.1103/PhysRevLett.119.230501. URL https://link.aps.org/doi/10.1103/PhysRevLett.119.230501.

[35] L.-M. Duan. Scaling ion trap quantum computation through fast quantum gates. *Phys. Rev. Lett.*, 93:100502, Sep 2004. doi: 10.1103/PhysRevLett.93.100502. URL https://link.aps.org/doi/10.1103/PhysRevLett.93.100502.

[36] C D B Bentley, A R R Carvalho, and J J Hope. Trapped ion scaling with pulsed fast gates. *New Journal of Physics*, 17(10):103025, October 2015. ISSN 1367-2630. doi: 10.1088/1367-2630/17/10/103025. URL http://stacks.iop.org/1367-2630/17/i=10/a=103025?key=crossref.2a63d802f18ceec31219bd3e5dcb491c.

[37] Evan Parnell Goldfarb Gale. *Limits of fast gate performance in trapped ion quantum computing.* Honours thesis, Australian National University, 2018.

[38] Alexander K. Ratcliffe, Richard L. Taylor, Joseph J. Hope, and Andre R. R.

Carvalho. Scaling trapped ion quantum computers using fast gates and microtraps. *Phys. Rev. Lett.*, 120:220501, May 2018. doi: 10.1103/PhysRevLett.120.220501. URL https://link.aps.org/doi/10.1103/PhysRevLett.120.220501.

[39] Dominik Janzing. *Entropy of Entanglement*, pages 205–209. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. ISBN 978-3-540-70626-7. doi: 10.1007/978-3-540-70626-7_66. URL https://doi.org/10.1007/978-3-540-70626-7_66.

[40] C. H. Bennett. Logical reversibility of computation. *IBM Journal of Research and Development*, 17(6):525–532, November 1973. doi: 10.1147/rd.176.0525. URL https://doi.org/10.1147/rd.176.0525.

[41] David P. DiVincenzo. Two-bit gates are universal for quantum computation. *Phys. Rev. A*, 51:1015–1022, Feb 1995. doi: 10.1103/PhysRevA.51.1015. URL https://link.aps.org/doi/10.1103/PhysRevA.51.1015.

[42] Alexei Gilchrist, Nathan K. Langford, and Michael A. Nielsen. Distance measures to compare real and ideal quantum processes. *Phys. Rev. A*, 71:062310, Jun 2005. doi: 10.1103/PhysRevA.71.062310. URL https://link.aps.org/doi/10.1103/PhysRevA.71.062310.

[43] Line Hjortshøj Pedersen, Niels Martin Møller, and Klaus Mølmer. The distribution of quantum fidelities. *Physics Letters A*, 372(47):7028 – 7032, 2008. ISSN 0375-9601. doi: https://doi.org/10.1016/j.physleta.2008.10.034. URL http://www.sciencedirect.com/science/article/pii/S0375960108015041.

[44] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. Quantum dynamics of single trapped ions. *Rev. Mod. Phys.*, 75:281–324, Mar 2003. doi: 10.1103/RevModPhys.75.281. URL https://link.aps.org/doi/10.1103/RevModPhys.75.281.

[45] S. Earnshaw. On the Nature of the Molecular Forces which Regulate the Constitution of the Luminiferous Ether. *Transactions of the Cambridge Philosophical Society*, 7:97, 1848.

[46] P. L. Knight, E. A. Hinds, M. B. Plenio, S. Gulde, H. HÃdffner, M. Riebe, G. Lancaster, C. Becher, J. Eschner, F. Schmidt-Kaler, I. L. Chuang, and R. Blatt. Quantum information processing with trapped Ca+ ions. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 361(1808):1363–1374, 2003. doi:

10.1098/rsta.2003.1206. URL https://royalsocietypublishing.org/doi/abs/10.1098/rsta.2003.1206.

[47] R. Gerritsma, G. Kirchmair, F. Zähringer, J. Benhelm, R. Blatt, and C. F. Roos. Precision measurement of the branching fractions of the $4p^2p_{3/2}$ decay of ca ii. *The European Physical Journal D*, 50(1):13–19, Nov 2008. ISSN 1434-6079. doi: 10.1140/epjd/e2008-00196-9. URL https://doi.org/10.1140/epjd/e2008-00196-9.

[48] J. J. Bollinger, D. J. Heizen, W. M. Itano, S. L. Gilbert, and D. J. Wineland. A 303-MHz frequency standard based on trapped Be+ ions. *IEEE Transactions on Instrumentation and Measurement*, 40(2):126–128, April 1991. doi: 10.1109/TIM.1990.1032897.

[49] P. T. H. Fisk, M. J. Sellars, M. A. Lawn, and G. Coles. Accurate measurement of the 12.6 GHz "clock" transition in trapped 171Yb+ ions. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 44(2):344–354, March 1997. doi: 10.1109/58.585119.

[50] D Heinrich, M Guggemos, M Guevara-Bertsch, M I Hussain, C F Roos, and R Blatt. Ultrafast coherent excitation of a 40Ca ion. *New Journal of Physics*, 21(7):073017, jul 2019. doi: 10.1088/1367-2630/ab2a7e. URL https://doi.org/10.1088%2F1367-2630%2Fab2a7e.

[51] Efrat Shimshoni, Giovanna Morigi, and Shmuel Fishman. Quantum zigzag transition in ion chains. *Phys. Rev. Lett.*, 106:010401, Jan 2011. doi: 10.1103/PhysRevLett.106.010401. URL https://link.aps.org/doi/10.1103/PhysRevLett.106.010401.

[52] J. P. Schiffer. Phase transitions in anisotropically confined ionic crystals. *Phys. Rev. Lett.*, 70:818–821, Feb 1993. doi: 10.1103/PhysRevLett.70.818. URL https://link.aps.org/doi/10.1103/PhysRevLett.70.818.

[53] J. I. Cirac and P. Zoller. A scalable quantum computer with ions in an array of microtraps. *Nature*, 404(6778):579–581, April 2000. ISSN 0028-0836, 1476-4687. doi: 10.1038/35007021. URL http://www.nature.com/articles/35007021.

[54] D. Stick, W. K. Hensinger, S. Olmschenk, M. J. Madsen, K. Schwab, and C. Monroe. Ion trap in a semiconductor chip. *Nature Physics*, 2(1):36âĂŞ39, Nov 2005. doi: 10.1038/nphys171.

[55] M. Kumph, P. Holz, K. Langer, M. Meraner, M. Niedermayr, M. Brownnutt, and R. Blatt. Operation of a planar-electrode ion-trap array with adjustable

RF electrodes. *New Journal of Physics*, 2016. ISSN 13672630. doi: 10.1088/1367-2630/18/2/023047.

[56] Daniel F. V. James. Quantum dynamics of cold trapped ions, with application to quantum computation. *Applied Physics B: Lasers and Optics*, 66(2):181–190, February 1998. ISSN 0946-2171, 1432-0649. doi: 10.1007/s003400050373. URL http://arxiv.org/abs/quant-ph/9702053. arXiv: quant-ph/9702053.

[57] H Haffner, C Roos, and R Blatt. Quantum computing with trapped ions. *Physics Reports*, 469(4):155–203, December 2008. ISSN 03701573. doi: 10.1016/j.physrep.2008.09.003. URL https://linkinghub.elsevier.com/retrieve/pii/S0370157308003463.

[58] C. D. B. Bentley, R. L. Taylor, A. R. R. Carvalho, and J. J. Hope. Stability thresholds and calculation techniques for fast entangling gates on trapped ions. *Physical Review A*, 93(4):042342, April 2016. doi: 10.1103/PhysRevA.93.042342. URL https://link.aps.org/doi/10.1103/PhysRevA.93.042342.

[59] J. J. Sakurai and Jim Napolitano. *Modern quantum mechanics*. Cambridge University Press, 2019.

[60] E.T. Jaynes and F.W. Cummings. Comparison of quantum and semiclassical radiation theories with application to the beam maser. *Proceedings of the IEEE*, 51(1):89–109, 1963. ISSN 0018-9219. doi: 10.1109/PROC.1963.1664. URL http://ieeexplore.ieee.org/document/1443594/.

[61] J. I. Cirac and P. Zoller. Quantum Computations with Cold Trapped Ions. *Physical Review Letters*, 74(20):4091–4094, May 1995. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.74.4091. URL https://link.aps.org/doi/10.1103/PhysRevLett.74.4091.

[62] A. Steane. The ion trap quantum information processor. *Applied Physics B: Lasers and Optics*, 64(6):623–643, June 1997. ISSN 0946-2171, 1432-0649. doi: 10.1007/s003400050225. URL http://link.springer.com/10.1007/s003400050225.

[63] Klaus Mølmer and Anders Sørensen. Multiparticle entanglement of hot trapped ions. *Phys. Rev. Lett.*, 82:1835–1838, March 1999. doi: 10.1103/PhysRevLett.82.1835. URL https://link.aps.org/doi/10.1103/PhysRevLett.82.1835.

[64] D.F.V. James. Quantum dynamics of cold trapped ions with application to quantum computation. *Applied Physics B: Lasers and Optics*, 66(2):181–190,

February 1998. ISSN 0946-2171, 1432-0649. doi: 10.1007/s003400050373. URL http://link.springer.com/10.1007/s003400050373.

[65] D. Jonathan, M. B. Plenio, and P. L. Knight. Fast quantum gates for cold trapped ions. *Phys. Rev. A*, 62:042307, Sep 2000. doi: 10.1103/PhysRevA.62.042307. URL https://link.aps.org/doi/10.1103/PhysRevA.62.042307.

[66] F Schmidt-Kaler, S Gulde, M Riebe, T Deuschle, A Kreuter, G Lancaster, C Becher, J Eschner, H Häffner, and R Blatt. The coherence of qubits based on single Ca ions. *Journal of Physics B: Atomic, Molecular and Optical Physics*, 36(3):623–636, February 2003. ISSN 0953-4075. doi: 10.1088/0953-4075/36/3/319. URL http://stacks.iop.org/0953-4075/36/i=3/a=319?key=crossref.1cdc3c8cbb8a6724d7b8a523a68109b0.

[67] Anders Sørensen and Klaus Mølmer. Entanglement and quantum computation with ions in thermal motion. *Phys. Rev. A*, 62:022311, July 2000. doi: 10.1103/PhysRevA.62.022311. URL https://link.aps.org/doi/10.1103/PhysRevA.62.022311.

[68] Alexander K. Ratcliffe, Lachlan M. Oberg, and Joseph J. Hope. Micromotion-Enhanced Fast Entangling Gates For Trapped Ion Quantum Computing. *arXiv Quantum Physics*. doi: arXiv:1902.06364v1. URL http://arxiv.org/abs/1902.06364.

[69] Alexander K. Ratcliffe, Richard L. Taylor, Joseph J. Hope, and Andre R.âĂŁR. Carvalho. Scaling Trapped Ion Quantum Computers Using Fast Gates and Microtraps. *Physical Review Letters*, 120(22), May 2018. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.120.220501. URL https://link.aps.org/doi/10.1103/PhysRevLett.120.220501.

[70] Mahmood Irtiza Hussain, Matthew Joseph Petrasiunas, Christopher D. B. Bentley, Richard L. Taylor, André R. R. Carvalho, Joseph J. Hope, Erik W. Streed, Mirko Lobino, and David Kielpinski. Ultrafast, high repetition rate, ultraviolet, fiber-laser-based source: application towards Yb+; fast quantum-logic. *Opt. Express*, 24(15):16638–16648, Jul 2016. doi: 10.1364/OE.24.016638. URL http://www.opticsexpress.org/abstract.cfm?URI=oe-24-15-16638.

[71] C. D. B. Bentley, A. R. R. Carvalho, D Kielpinski, and J J Hope. Fast gates for ion traps by splitting laser pulses. *New Journal of Physics*, 15(4):043006, apr 2013. doi: 10.1088/1367-2630/15/4/043006. URL https://doi.org/10.1088%2F1367-2630%2F15%2F4%2F043006.

[72] R. Blatt and C. F. Roos. Quantum simulations with trapped ions. *Nature Physics*, 8(4):277–284, April 2012. ISSN 1745-2473, 1745-2481. doi: 10.1038/nphys2252. URL http://www.nature.com/articles/nphys2252.

[73] J. Casanova, A. Mezzacapo, L. Lamata, and E. Solano. Quantum Simulation of Interacting Fermion Lattice Models in Trapped Ions. *Physical Review Letters*, 108(19):190502, May 2012. doi: 10.1103/PhysRevLett.108.190502. URL https://link.aps.org/doi/10.1103/PhysRevLett.108.190502.

[74] Lucas Lamata, Antonio Mezzacapo, Jorge Casanova, and Enrique Solano. Efficient quantum simulation of fermionic and bosonic models in trapped ions. *EPJ Quantum Technology*, 1(1), 2014. doi: 10.1140/epjqt9.

[75] E. Y. Loh, J. E. Gubernatis, R. T. Scalettar, S. R. White, D. J. Scalapino, and R. L. Sugar. Sign problem in the numerical simulation of many-electron systems. *Phys. Rev. B*, 41:9301–9307, May 1990. doi: 10.1103/PhysRevB.41.9301. URL https://link.aps.org/doi/10.1103/PhysRevB.41.9301.

[76] Andrew Tranter, Peter J. Love, Florian Mintert, and Peter V. Coveney. A comparison of the Bravyi Kitaev and Jordan Wigner transformations for the quantum simulation of quantum chemistry. *Journal of Chemical Theory and Computation*, 14(11):5617–5630, Jun 2018. doi: 10.1021/acs.jctc.8b00450.

[77] H. F. Trotter. On the product of semi-groups of operators. *Proceedings of the American Mathematical Society*, 10(4):545–545, Jan 1959. doi: 10.1090/s0002-9939-1959-0108732-6.

[78] Naomichi Hatano and Masuo Suzuki. *Finding Exponential Product Formulas of Higher Orders*, pages 37–68. Springer Berlin Heidelberg, Berlin, Heidelberg, 2005. ISBN 978-3-540-31515-5. doi: 10.1007/11526216_2. URL https://doi.org/10.1007/11526216_2.

[79] J. Hubbard and Brian Hilton Flowers. Electron correlations in narrow energy bands. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 276(1365):238–257, 1963. doi: 10.1098/rspa.1963.0204. URL https://royalsocietypublishing.org/doi/abs/10.1098/rspa.1963.0204.

[80] M. Schüler, M. Rösner, T. O. Wehling, A. I. Lichtenstein, and M. I. Katsnelson. Optimal hubbard models for materials with nonlocal coulomb interactions: Graphene, silicene, and benzene. *Phys. Rev. Lett.*, 111:036601, Jul 2013. doi: 10.1103/PhysRevLett.111.036601. URL https://link.aps.org/doi/10.1103/PhysRevLett.111.036601.

[81] Wayne W. Lukens, Nicola Magnani, and Corwin H. Booth. Application of the hubbard model to Cp*2Yb(bipy), a model system for strong exchange coupling in lanthanide systems. *Inorganic Chemistry*, 51(19):10105–10110, 2012. doi: 10.1021/ic300037q. URL https://doi.org/10.1021/ic300037q. PMID: 22988887.

[82] P W Anderson, P A Lee, M Randeria, T M Rice, N Trivedi, and F C Zhang. The physics behind high-temperature superconducting cuprates: the plain vanilla version of RVB. *Journal of Physics: Condensed Matter*, 16(24):R755–R769, jun 2004. doi: 10.1088/0953-8984/16/24/r02. URL https://doi.org/10.1088%2F0953-8984%2F16%2F24%2Fr02.

[83] Panos Aliferis and Andrew W. Cross. Subsystem fault tolerance with the Bacon-Shor code. *Phys. Rev. Lett.*, 98:220502, May 2007. doi: 10.1103/PhysRevLett.98.220502. URL https://link.aps.org/doi/10.1103/PhysRevLett.98.220502.

[84] Andrew M. Steane. Overhead and noise threshold of fault-tolerant quantum error correction. *Phys. Rev. A*, 68:042322, Oct 2003. doi: 10.1103/PhysRevA.68.042322. URL https://link.aps.org/doi/10.1103/PhysRevA.68.042322.

[85] Peter Brooks and John Preskill. Fault-tolerant quantum computation with asymmetric Bacon-Shor codes. *Phys. Rev. A*, 87:032310, Mar 2013. doi: 10.1103/PhysRevA.87.032310. URL https://link.aps.org/doi/10.1103/PhysRevA.87.032310.

[86] I. V. Inlek, C. Crocker, M. Lichtman, K. Sosnova, and C. Monroe. Multispecies trapped-ion node for quantum networking. *Phys. Rev. Lett.*, 118:250502, Jun 2017. doi: 10.1103/PhysRevLett.118.250502. URL https://link.aps.org/doi/10.1103/PhysRevLett.118.250502.

[87] Dong C. Liu and Jorge Nocedal. On the limited memory bfgs method for large scale optimization. *Mathematical Programming*, 45(1-3):503âĂŞ528, 1989. doi: 10.1007/bf01589116.

[88] J. J. Garcia-Ripoll, P. Zoller, and J. I. Cirac. Fast and robust two-qubit gates for scalable ion trap quantum computing. *Physical Review Letters*, 91 (15), October 2003. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.91.157901. URL http://arxiv.org/abs/quant-ph/0306006. arXiv: quant-ph/0306006.

[89] Rainer Blatt. Private communication, 2016.

[90] S. Dasgupta, T. Kushwaha, and D. Goswami. Dependence of adiabatic popula-

tion transfer on pulse profile. *Pramana*, 66(6):999, Jun 2006. ISSN 0973-7111. doi: 10.1007/BF02708454. URL https://doi.org/10.1007/BF02708454.

[91] Yu Guo, Daoyi Dong, and Chuan-Cun Shu. Optimal and robust control of quantum state transfer by shaping the spectral phase of ultrafast laser pulses. *Phys. Chem. Chem. Phys.*, 20:9498–9506, 2018. doi: 10.1039/C8CP00512E. URL http://dx.doi.org/10.1039/C8CP00512E.

[92] Michael Hush. Private communication, October 16 2019.

[93] Alok Sharan and Debabrata Goswami. Prospects of ultrafast pulse shaping. *Current Science*, 82(1):30–37, 2002. ISSN 00113891. URL http://www.jstor.org/stable/24105924.

[94] Thomas Oksenhendler and Nicolas Forget. *Pulse-Shaping Techniques Theory and Experimental Implementations for Femtosecond Pulses*. 02 2010. ISBN 978-953-7619-80-0. doi: 10.5772/7958.

[95] W. C. Campbell, J. Mizrahi, Q. Quraishi, C. Senko, D. Hayes, D. Hucul, D. N. Matsukevich, P. Maunz, and C. Monroe. Ultrafast gates for single atomic qubits. *Phys. Rev. Lett.*, 105:090502, Aug 2010. doi: 10.1103/PhysRevLett.105.090502. URL https://link.aps.org/doi/10.1103/PhysRevLett.105.090502.

[96] L. M. K. Vandersypen and I. L. Chuang. NMR techniques for quantum control and computation. *Rev. Mod. Phys.*, 76:1037–1069, Jan 2005. doi: 10.1103/RevModPhys.76.1037. URL https://link.aps.org/doi/10.1103/RevModPhys.76.1037.

[97] Kaveh Khodjasteh and Lorenza Viola. Dynamically error-corrected gates for universal quantum computation. *Phys. Rev. Lett.*, 102:080501, Feb 2009. doi: 10.1103/PhysRevLett.102.080501. URL https://link.aps.org/doi/10.1103/PhysRevLett.102.080501.

[98] Chingiz Kabytayev, Todd J. Green, Kaveh Khodjasteh, Michael J. Biercuk, Lorenza Viola, and Kenneth R. Brown. Robustness of composite pulses to time-dependent control noise. *Phys. Rev. A*, 90:012316, Jul 2014. doi: 10.1103/PhysRevA.90.012316. URL https://link.aps.org/doi/10.1103/PhysRevA.90.012316.

[99] J. True Merrill and Kenneth R. Brown. Progress in compensating pulse sequences for quantum computation. *Advances in Chemical Physics Quantum Information and Computation for Chemistry*, page 241âĂŞ294, 2014. doi: 10.1002/9781118742631.ch10.

[100] C. L. Edmunds, C. Hempel, R. J. Harris, V. M. Frey, T. M. Stace, and M. J. Biercuk. Dynamically corrected gates suppress spatio-temporal error correlations as measured by randomized benchmarking. *arXiv:1909.10727 [quant-ph]*, September 2019. URL `http://arxiv.org/abs/1909.10727`. arXiv: 1909.10727.

[101] Nikolay V Vitanov, Thomas Halfmann, Bruce W Shore, and Klaas Bergmann. Laser-induced population transfer by adiabatic passage techniques. *Annual Review of Physical Chemistry*, 52(1):763–809, 2001. doi: 10.1146/annurev.physchem.52.1.763. URL `https://doi.org/10.1146/annurev.physchem.52.1.763`. PMID: 11326080.

[102] Emil Zak. Optimal control theory for rapid-adiabatic passage techniques in inhomogeneous external fields. *arXiv:1712.06998 [quant-ph]*, arXiv:1712.06998 [quant-ph], 12 2017.

[103] V.S. Malinovsky and J.L. Krause. General theory of population transfer by adiabatic rapid passage with intense, chirped laser pulses. *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, 14(2): 147–155, May 2001. ISSN 1434-6079. doi: 10.1007/s100530170212. URL `https://doi.org/10.1007/s100530170212`.

[104] S. Chelkowski and A. D. Bandrauk. Raman chirped adiabatic passage: a new method for selective excitation of high vibrational states. *Journal of Raman Spectroscopy*, 28(6):459–466, 1997. doi: 10.1002/(SICI)1097-4555(199706)28:6<459::AID-JRS124>3.0.CO;2-Y. URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/%28SICI%291097-4555%28199706%2928%3A6%3C459%3A%3AAID-JRS124%3E3.0.CO%3B2-Y`.

[105] T. Watanabe, S. Nomura, K. Toyoda, and S. Urabe. Sideband excitation of trapped ions by rapid adiabatic passage for manipulation of motional states. *Phys. Rev. A*, 84:033412, Sep 2011. doi: 10.1103/PhysRevA.84.033412. URL `https://link.aps.org/doi/10.1103/PhysRevA.84.033412`.

[106] Yu Chen, C. Neill, P. Roushan, N. Leung, M. Fang, R. Barends, J. Kelly, B. Campbell, Z. Chen, B. Chiaro, A. Dunsworth, E. Jeffrey, A. Megrant, J. Y. Mutus, P. J. J. O'Malley, C. M. Quintana, D. Sank, A. Vainsencher, J. Wenner, T. C. White, Michael R. Geller, A. N. Cleland, and John M. Martinis. Qubit architecture with high coherence and fast tunable coupling. *Phys. Rev. Lett.*, 113:220502, Nov 2014. doi: 10.1103/PhysRevLett.113.220502. URL `https://link.aps.org/doi/10.1103/PhysRevLett.113.220502`.

[107] D. Hayes, D. N. Matsukevich, P. Maunz, D. Hucul, Q. Quraishi, S. Olmschenk, W. Campbell, J. Mizrahi, C. Senko, and C. Monroe. Entanglement of atomic qubits using an optical frequency comb. *Phys. Rev. Lett.*, 104:140501, Apr 2010. doi: 10.1103/PhysRevLett.104.140501. URL `https://link.aps.org/doi/10.1103/PhysRevLett.104.140501`.

[108] Marlan O. Scully and M. Suhail. Zubairy. *Quantum optics*. Cambridge University Press, 1997.

[109] Herbert Goldstein. *Classical mechanics*. Pearson, 2014.