Distributed Quantum Computing across a Two-Node Trapped-Ion Quantum Network

Dougal Main

St Anne's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Michaelmas Term, 2024

For my mum

Abstract

Distributed quantum computing (DQC) combines the computing power of multiple networked quantum processing modules, ideally enabling the execution of large quantum circuits without compromising performance or connectivity. Photonic networks are well-suited as a versatile and reconfigurable interconnect layer for DQC; remote entanglement shared between matter qubits across the network enables all-to-all logical connectivity via quantum gate teleportation (QGT). For a scalable DQC architecture, the QGT implementation must be deterministic and repeatable; until now, no demonstration has satisfied these requirements. In this thesis, we report the first experimental demonstration of distributed quantum computing across a quantum network of two photonically interconnected trapped-ion modules.

Our experimental apparatus comprises two trapped-ion modules, separated by $\sim 2\,\mathrm{m}$, each co-trapping one $^{43}Ca^{+}$ ion and one $^{88}Sr^{+}$ ion. The $^{43}Ca^{+}$ ions provide a robust quantum memory for the modules, while the $^{88}Sr^{+}$ ions provide an interface to the optical quantum network, enabling the remote entanglement of $^{88}Sr^{+}$ ions in separate modules. Using local mixed-species operations, we demonstrate the enhanced functionality of the mixed-species trapped-ion quantum network by creating more complex remotely entangled states. These include remote $^{88}Sr^{+}$-$^{43}Ca^{+}$ and $^{43}Ca^{+}$-$^{43}Ca^{+}$ entanglement, as well as 3- and 4-qubit mixed-species GHZ states. Additionally, we demonstrate the storage of remote entanglement in the $^{43}Ca^{+}$ ions for up to 10 s.

Utilising these techniques, we demonstrate the distribution of quantum computations across the mixed-species trapped-ion quantum network. Using heralded remote entanglement between the $^{88}Sr^{+}$ ions, we deterministically teleport controlled-Z gates between two $^{43}Ca^{+}$ ions in separate modules, achieving a fidelity of 86 %. We then implement distributed iSWAP and SWAP circuits, compiled with 2 and 3 instances of QGT, respectively, demonstrating the ability to distribute arbitrary two-qubit operations. Finally, we execute Grover's search algorithm – the first implementation of a distributed quantum algorithm comprising multiple non-local two-qubit gates – and measure a success rate of 71 %. These results represent a significant step towards realising a large-scale distributed quantum computer.

Acknowledgements

I am extremely grateful for all the guidance and support given to me by my supervisor, Prof. David Lucas. Thank you for providing the opportunity to focus on incredibly interesting physics problems without the burden of administrative or financial concerns. Despite leading such a large research group, you always made time to chat, whether it was to impart your incredible ion-trapping wisdom, or to reassure me that setbacks are a normal part of experimental physics.

I have been privileged to work alongside some fantastic people in the quantum networking team–Ellis Ainley, Ayush Agrawal, Dr. Gabriel Araneda, Dr. Chris Ballance, Dr. Jacob Blackmore, Dr. Peter Drmota, Dr. Péter Juhász, Adam Martínez, Dr. Laurent Stephenson, Dr. David Nadlinger, Dr. Bethan Nichol, and Dr. Raghavendra Srinivas–without whom the work presented in this thesis would not have been possible. In particular, thank you to Ellis for all your support in the lab and your willingness to track down mysterious phase offsets at 5 a.m.; to Gabriel for your invaluable guidance and for making my DPhil journey so enjoyable; to Peter for teaching me so much about the experiment and hunting magnetic field signals from the roof of the physics department; to David for always sharing your incredible knowledge about all areas of physics; to Beth for not letting me take anything too seriously and making each day fun; and to Raghu for your incredible support throughout some of the toughest points in my DPhil. Finally, thank you to everyone who read through this thesis–any remaining errors are my own.

I am also incredibly grateful to the wider Oxford Ion Trap group–five years is too short a time to spend among such amazing people. In particular, thank you to my Lab 1 co-inhabitants, Dr. Oana Băzăvan and Dr. Sebastian Saner, for your immense patience during data-taking sessions. Thanks also to Dr. Ana Sotirova for the late-night cryogenics lessons.

I would like to thank my undergraduate tutors at St Anne's College–Prof. Pat Irwin, Prof. Neville Harnew, and Prof. Norman McCubbin–who taught me not only all I know about physics but also that it could be fun. Thanks also to Dr. Tom Hird who introduced me to the world of research and has supported me ever since.

I have been extremely fortunate to have been surrounded by some incredible friends who have been unbelievably supportive over the years – to you all, I am deeply indebted. Furthermore, thanks to everyone at Horspath Cricket Club and Oxford City Chess Club for providing many welcome distractions from work.

Thank you so much to my Dad, David, and my brother, Ruairidh, for the endless support. And finally, to Lowri–thank you for being a constant source of joy, motivation, and cups of tea; I could not have done this without you.

Contents

List of Figures

List of Tables

1

Introduction

*

The turn of the twentieth century saw two scientific revolutions: the development of quantum mechanics and the theory of relativity. These theories transformed our understanding of the universe and are often regarded as the foundations of modern science. While relativity provided insights into the large-scale structure of the universe, quantum mechanics offered a framework for developing physical theories at the microscopic scale. One foundational concept from this quantum revolution, wave-particle duality, led to significant breakthroughs, including the explanation of the photoelectric effect, the understanding of quantum tunnelling, and the development of the quantum model of atomic structure.

As seen throughout history, new fundamental insights often lead to rapid technological development. Indeed, the quantum revolution led to the creation of transformative technologies such as lasers, superconductors, and semiconductors, leading to the invention of the transistor in 1947, an innovation that heralded the beginning of the *information age*.

We are now in the *second quantum revolution* [1]. In contrast to the first quantum revolution, which used quantum mechanics to explain macroscopic phenomena, such

as the properties of materials, from their microscopic quantum origins, the second quantum revolution focuses on directly manipulating individual quantum systems according to quantum mechanical principles. This has enabled the development of a new class of technologies, known as *quantum technologies*. These include atomic clocks, quantum cryptography, and, the focus of this thesis, the quantum computer.

1.1 Quantum computing

In 1982, Feynman [2] pointed out that it is not possible to efficiently simulate a quantum mechanical system using a classical computer, where "efficiently" means that the resources required by the classical computer grow polynomially with the size of the quantum system. To illustrate this, consider a system of N two-level systems. A classical computer would need to store an amplitude for each of the 2^N possible states of the system. As N increases, the required memory grows exponentially, making classical simulation infeasible for large quantum systems. Therefore, to efficiently simulate a quantum mechanical system, Feynman [2] proposed that a new kind of computer would be needed – one that is itself quantum mechanical.

Deutsch [3] proposed the idea of a *universal quantum computer* – a quantum analogue of the Turing machine, capable of efficiently simulating any arbitrary physical system, assuming the system obeys the laws of quantum mechanics. This concept opened up a new computational paradigm, one in which a quantum computer could perform computations where each step could be any operation allowed by quantum mechanics. Deutsch and Jozsa [4] provided an example of how such a quantum computer could outperform its classical counterpart in a specific computational task. Shortly after, Shor's algorithm [5], for efficient integer factorisation, and Grover's algorithm [6], for searching through an unstructured set of items, provided concrete evidence of the computational advantages of quantum computers. Around the same time, the development of quantum error correction demonstrated that large-scale quantum computers could be realised even in the presence of noise and experimental imperfections, paving the way for *fault-tolerant* quantum computing [7, 8].

We now turn to the practical challenges of building a quantum computer. While classical computers store information using binary states (0 or 1), known as *bits*, a quantum computer uses *qubits*, which are two-level quantum systems with states often denoted $|0\rangle$ and $|1\rangle$. Unlike classical bits, qubits can exist in a superposition of states $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where α and β are complex numbers subject to $|\alpha|^2 + |\beta|^2 = 1$.

To construct a quantum computer from a physical system, certain requirements must be met, as outlined by DiVincenzo [9]. These criteria are:

- a scalable physical system with well-characterised qubits
- the ability to prepare qubits in a simple fiducial state, such as $|0\rangle$
- long coherence times, much longer than the gate operation time
- a universal set of quantum gates
- qubit-specific measurement capabilities

DiVincenzo [9] mentions two additional criteria that enable quantum computers to communicate by sending *flying* qubits between systems containing *stationary* qubits. The additional criteria are:

- the ability to create an interface between stationary and flying qubits,
- the ability to faithfully transmit flying qubits between specified locations.

Numerous physical platforms are currently being explored for the realisation of a quantum computer, including trapped ions, neutral atoms, superconducting qubits, photons, and diamond colour centers.

The work in this thesis was performed using the trapped-ion platform, one of the leading platforms for quantum computing. The underlying atomic physics of trapped ions is well-understood, and all ions with the same number of protons, neutrons, and electrons will exhibit identical properties, making the qubits encoded in the energy level structure of the ions free from variations that could arise from,

e.g., manufacturing processes. Trapped-ion qubits have been used to demonstrate state-of-the-art qubit state preparation and measurement [10, 11, 12]. Additionally, trapped ions have been used to demonstrate universal sets of single- and two-qubit gates, also with state-of-the-art fidelities [13, 14, 15, 16, 17, 18]. Finally, coherence times exceeding $\sim 1\,\mathrm{h}$ have been observed [19], far exceeding typical gate durations of $\sim 1 - 100\ \mu\mathrm{s}$.

1.2 Scaling up a quantum computer

Regardless of the physical platform used to realise a quantum computer, scaling up the number of qubits while maintaining precise control and interconnectivity between the qubits comes with major technical challenges [20, 21, 22]. One prominent example is the *wiring problem*. Suppose we construct a quantum computer containing N qubits such that the number of *wires* – that is, input/output signals – scales linearly with N. Theoretically, this scaling does not impose a fundamental limit on the size of the quantum computer. However, in practice, this inevitably leads to engineering constraints when scaling to large qubit numbers. For high-performance operation, quantum computers must be sufficiently isolated from their environment, often achieved using ultra-high vacuum or cryogenic systems. As the density of wires increases, physically engineering the interfaces between the external control system and the isolated quantum system, e.g., via feedthroughs in a vacuum chamber, becomes significantly more challenging. Moreover, a larger number of wires increases the amount of thermal energy that is dissipated in the quantum system, leading to a degradation in performance.

Tackling this wiring problem is a serious technical challenge for realising large-scale quantum computers. Malinowski, Allcock, and Ballance [23] recently proposed a control architecture for the trapped-ion platform capable of wiring a quantum computer with up to 1000 qubits. While this represents a scaling improvement of one to two orders of magnitude over current systems, the authors acknowledge that beyond this point, the wiring problem remains a significant barrier to scalability.

To summarise, although there is no fundamental limitation on the size of a quantum computer in principle, technical limitations arising from increased wire density and associated thermal dissipation present major challenges in practice.

Classical computing faces similar issues; for example, the volume occupied by a classical processor can be limited by the efficiency of heat extraction from its surface area. To overcome these obstacles, distributed computing architectures are frequently employed, whether in the form of multi-core processors or large-scale super-computing clusters. Analogously, adopting modular or distributed architectures in quantum computing may offer a practical pathway to scale up qubit numbers while managing control complexity and mitigating thermal issues.

1.2.1 Distributed quantum computing

Just as distributed architectures have proven essential in classical computing, a quantum computer could similarly benefit from a distributed quantum computing (DQC) architecture, depicted in Figure 1.1, where large quantum computations may be executed by an array of quantum processing modules [24, 25]. By preserving the reduced complexity of the individual modules and transforming the scaling challenge into the task of building more modules and establishing an interface between them, this architecture provides a scalable approach to fault-tolerant quantum computing [26, 27]. The interface between modules could be realised by directly transferring quantum information between modules via, e.g., the transmission of flying qubits as in the DiVincenzo criteria above. However, losses in the interconnecting quantum channels would lead to a loss of quantum information.

Quantum teleportation offers a lossless alternative interface, using only bipartite entanglement (e.g., Bell states) shared between modules, together with local operations and classical communication (LOCC) to effectively replace the direct transfer of quantum information across quantum channels [28, 29]. Quantum gate teleportation (QGT) efficiently implements non-local entangling gates between qubits in separate modules, consuming only one Bell pair and the exchange of two classical

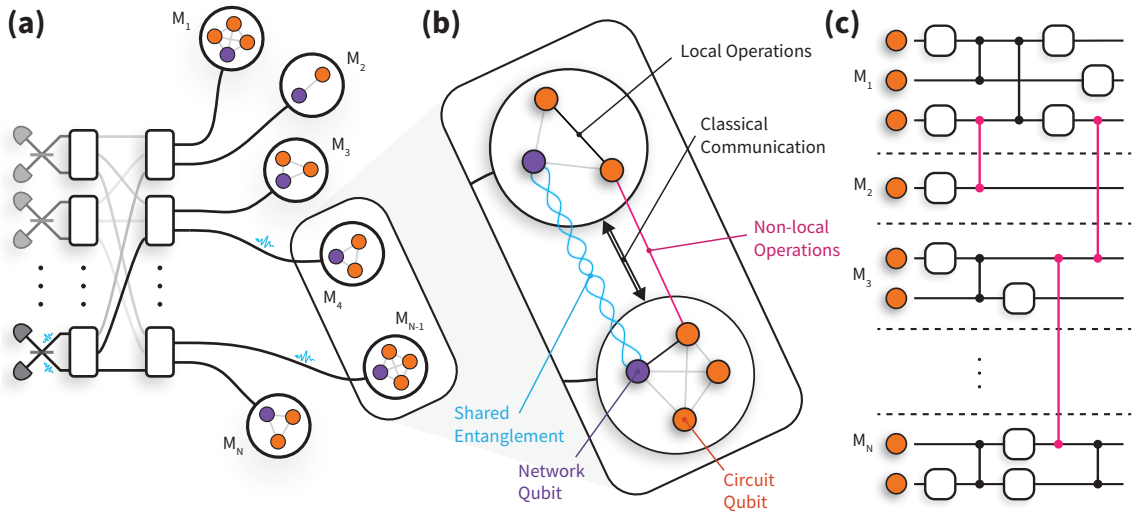

Figure 1.1: **(a)**, Schematic of a DQC architecture comprising photonically interconnected modules. Entanglement is heralded between network qubits via the interference of photons on beamsplitters. A photonic switchboard provides a flexible and reconfigurable network topology. **(b)**, The modules consist of at least one network qubit (purple) and at least one circuit qubit (orange), which may directly interact via local operations. Quantum gate teleportation mediates non-local gate interactions (pink) between circuit qubits in separate modules. These protocols require the resources of shared entanglement, local operations, and classical communication. **(c)**, A quantum circuit distributed across a network of small quantum processing modules that function together as a single, fully intra-connected quantum computer.

bits [30, 31], as depicted in Figure 1.1(b). Given arbitrary single- and two-qubit operations within each node, QGT completes a universal gate set for the distributed quantum computer [29]. The primary advantage of teleportation-based schemes over direct transfer is the exclusive use of the quantum channel for generating identical Bell states; channel losses can be overcome by repetition without losing quantum information, and the distance between modules can be increased by inserting quantum repeaters [32]. Additionally, channel noise may be suppressed using entanglement purification [33]. Since teleportation protocols are executed strictly after entanglement has been established, they enable continuous deterministic operation even if the entanglement is generated non-deterministically. This deterministic nature eliminates the need for post-selection of singular successful outcomes out of an exponentially large set of undesired ones.

Teleportation protocols are agnostic to the physical implementation of the quantum channels, making them a versatile tool for DQC across different platforms. In

the trapped-ion quantum charge-coupled device (QCCD) architecture, qubits can be dynamically transported between modules within a single chip [34] – or even across chips [35] – and thus be used to mediate entangling gates between different trap zones [36, 37]. Photons, however, make natural carriers of quantum information since they can travel long distances without significant degradation of their quantum state. Photonic interconnects enable all-to-all connectivity between qubits distributed across the network whose topology can be dynamically reconfigured without the need to open up complex vacuum or cryogenic systems. Moreover, optical components are widely available and can be operated under ambient conditions. These properties make photonic interconnects particularly appealing for networking quantum computing modules, as shown in Figure 1.1(a).

As depicted in Figure 1.1(b), we consider modules containing "network" and "circuit" qubits with full interconnectivity via local quantum operations. Remote entanglement of network qubits in separate modules is generated by the interference of photons, where reconfigurability and flexibility could be provided via a photonic switchboard. This entanglement can then be used to mediate entangling gates between the circuit qubits in different modules via QGT, enabling the network to function as a single fully connected quantum processor, as shown in Figure 1.1(c). Quantum circuits can be partitioned freely in this architecture, down to a minimum of one circuit qubit per module in the fully distributed case. Heralded entanglement between spatially separated qubits has been achieved experimentally in a variety of platforms including diamond colour centres [38, 39], superconducting qubits [40], neutral atoms [41, 42], and trapped-ions [43, 44, 45, 46].

QGT has been implemented probabilistically in purely photonic systems, requiring passive optical elements and post-selection to perform the conditional rotations that complete the gate teleportation [47, 48]. Chou *et al.* [49] demonstrated deterministic teleportation of a controlled-NOT gate between two qubits encoded in the modes of two superconducting cavities on the same device, separated by $\sim 2\,\mathrm{cm}$, while a third cavity enabled the deterministic generation of entanglement between

two transmon network qubits. Recently, there have been demonstrations of QGT between superconducting qubits within a single device, demonstrating the viability of QGT to overcome nearest neighbour constraints in this architecture [50, 51]. In the trapped-ion QCCD architecture, Wan *et al.* [36] demonstrated QGT in which the entanglement was deterministically generated between two "network" qubits via local operations before being transported $\sim 840\,\mu m$ to two separate locations within the same trap. Additionally, there have been demonstrations of heralded non-local entangling gates across a photonic quantum network in which photons are used to directly transfer quantum information between modules [52, 53]. However, in these demonstrations, unavoidable photon loss destroyed the states of the circuit qubits, rendering these schemes non-deterministic. Until now, there has been (i) no demonstration of deterministic QGT across a quantum network, and (ii) no demonstration of distributed circuits comprising multiple non-local entangling gates.

1.3 Thesis outline

The key result of this thesis is the demonstration of distributed quantum computing between two trapped-ion modules across an optical quantum network link. Within each module, we delegate specific roles to different species, choosing $^{88}Sr^{+}$ for the communication across the network, while $^{43}Ca^{+}$ is used for the storage of quantum information within each module. This thesis first presents the key components of the mixed-species trapped-ion modules, then demonstrates the enhanced capabilities of this mixed-species trapped-ion quantum network, before culminating in the results from the demonstration of distributed quantum computing.

The remainder of this thesis is structured as follows:

Chapter 2 discusses the theoretical tools used throughout this thesis, covering the representations of quantum processes, normal mode analysis and the quantisation of the motion of ions in a Paul trap, the manipulation of qubits encoded in the energy levels of ions via the laser-ion interaction, and the implementation of single- and

two-qubit gates in trapped-ion quantum processors.

Chapter 3 outlines a number of important calibration and characterisation protocols, including an adaptive Bayesian phase estimation protocol, quantum state and process tomography protocols, and randomised benchmarking.

Chapter 4 outlines the experimental apparatus used to obtain the experimental results presented later in this thesis. In particular, we discuss the choice of ion species, provide an overview of the trapped-ion modules, and present a detailed discussion and characterisation of the Paul traps used to confine the ions. Additionally, we cover the Raman laser system, which is integral to many of the operations performed in each module, outline the range of dissipative operations such as ion loading, laser cooling, state preparation, and qubit readout, and conclude with a discussion of our magnetic field stabilisation system.

Chapter 5 provides a discussion and characterisation of several crucial local operations performed in each module, covering single-qubit gates, mixed-species two-qubit gates, hyperfine qubit conversion, and mid-circuit measurements.

Chapter 6 begins with a characterisation of our quantum networking apparatus, starting with the characterisation of entanglement between a $^{88}Sr^{+}$ ion and a photon, followed by the remote entanglement of two $^{88}Sr^{+}$ ions across the quantum network. Following this, we present the results from our mixed-species trapped-ion quantum network, including the generation of remote entanglement between a $^{88}Sr^{+}$ and a $^{43}Ca^{+}$ ion, two $^{43}Ca^{+}$ ions, and 3-qubit ($^{88}Sr^{+}$-$^{88}Sr^{+}$-$^{43}Ca^{+}$) and 4-qubit ($^{88}Sr^{+}$-$^{88}Sr^{+}$-$^{43}Ca^{+}$-$^{43}Ca^{+}$) GHZ states. We then demonstrate the robust storage of remote entanglement using the $^{43}Ca^{+}$ ions for up to 10 s.

Chapter 7 presents results from the first-ever demonstration of distributed quantum computing across an optical network link. In particular, we demonstrate the deterministic teleportation of controlled-Z, iSWAP, and SWAP gates between two

macroscopically separated "circuit" qubits, mediated using remote entanglement shared between two "network" qubits. Finally, we execute the two-qubit Grover's algorithm using this distributed quantum computer.

Chapter 8 concludes this thesis, summarising our results and discussing an outlook for the experiment and the field of quantum networking.

2

Theory

*

While this thesis primarily presents experimental results, it is essential to first introduce the theoretical framework that enables the generation, analysis, and interpretation of these results. This chapter therefore presents established theoretical background for trapped-ion quantum computing.

I begin by reviewing the quantum-mechanical formalism underpinning all analyses in this thesis, following Wood, Biamonte, and Cory [54]. Next, I summarise the dynamics and quantisation of the motion of trapped-ions, largely based on Home [55] and Drmota [56]. Finally, I outline the manipulation of ion internal states via laser-based interactions, drawing from Ballance [57] and Hughes [58].

2.1 Quantum theory

We begin with a brief discussion of the quantum mechanical theory that is used throughout this thesis. In addition to providing theoretical context for the subsequent chapters, the purpose of this section – in particular, Section 2.1.2 – is to provide the tools used in a range of calculations and analyses throughout this thesis.

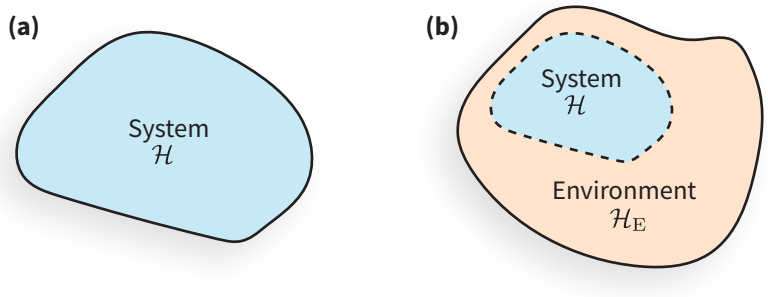

Figure 2.1: Graphical representation of **(a)** closed and **(b)** open quantum systems. **(a)** The closed quantum system, $\mathcal{H}$, does not interact with any environment, and thus quantum states within $\mathcal{H}$ undergo unitary evolution. **(b)** The open quantum system, $\mathcal{H}$, is now able to interact with its surrounding environment, $\mathcal{H}_{\mathrm{E}}$. The total system, $\mathcal{H} \otimes \mathcal{H}_{\mathrm{E}}$, is considered to be closed and thus undergoes unitary evolution; however the coupling of the system to its environment means that evolution within the subspace $\mathcal{H}$ is non-unitary.

2.1.1 Closed quantum systems

A *closed* quantum system is one that is isolated from, and cannot interact with, any other system – this is represented in Figure 2.1(a). The quantum system has a particular space of states, $\mathcal{H}$, such that the state of the quantum system is specified by an element from this space, $|\psi\rangle \in \mathcal{H}$. We consider more general states through the use of the density operator, $\hat{\rho} \in L(\mathcal{H})$, where $L(\mathcal{H}) \equiv L(\mathcal{H}, \mathcal{H})$ and $L(\mathcal{H}_1, \mathcal{H}_2)$: $\mathcal{H}_1 \to \mathcal{H}_2$ is the set of linear operators between $\mathcal{H}_1$ and $\mathcal{H}_2$. The density operator enables the description of statistical mixtures of quantum states, where if p_i is the probability for the quantum system being in the state $|\psi_i\rangle$, then the state of the quantum system can be written as $\hat{\rho} = \sum_i p_i |\psi_i\rangle \langle\psi_i|$. In this thesis, we consider two specific Hilbert spaces. The first is the Hilbert space of a *qubit*, $\mathcal{Q}$, i.e., a two-level quantum system spanned by the orthonormal computational states $|0\rangle$ and $|1\rangle$. The second relevant space is the Fock space, $\mathcal{A}$, which is an infinite-dimensional space used for describing the excitations of a harmonic oscillator.

The time-evolution of a closed quantum system is governed by the Schrödinger equation,

$$i\hbar \frac{\mathrm{d}\hat{\rho}}{\mathrm{d}t} = \left[\hat{H}, \hat{\rho}\right], \tag{2.1}$$

where $\hat{H}$ is the Hamiltonian for the system. This equation describes a unitary

evolution of the quantum system; the evolution over a discrete time interval may be described by the unitary transformation,

$$\hat{\rho}(t_2) = \hat{U}\hat{\rho}(t_1)\hat{U}^\dagger, \tag{2.2}$$

where $\hat{U} = \hat{U}(t_2; t_1)$ is a unitary operator.

2.1.2 Open quantum systems

In general, no system – except perhaps the Universe itself – can be considered perfectly isolated. A quantum system that interacts with an environment, the state of which we cannot directly access, is referred to as an *open* quantum system. Consider an open quantum system with a state space $\mathcal{H}$ that interacts with an environment, as depicted in Figure 2.1(b). We denote the Hilbert space of the environment by $\mathcal{H}_\mathrm{E}$. The total state of the system and its environment is represented by $\hat{\rho}_\mathrm{T} \in L\left(\mathcal{H} \otimes \mathcal{H}_\mathrm{E}\right)$.

By extending the environment to encompass the entire Universe, we can consider the combined system of the quantum system and its environment as a closed quantum system. Suppose the Hamiltonian for the total system is given by

$$\hat{H}_\mathrm{T} = \hat{H} + \hat{H}_\mathrm{E} + \hat{H}_\mathrm{int}, \tag{2.3}$$

where $\hat{H}$ and $\hat{H}_\mathrm{E}$ are the Hamiltonians for the quantum system and the environment, respectively, and $\hat{H}_\mathrm{int}$ describes their interaction. Since the total system is closed, it will therefore evolve according to the Schrödinger equation,

$$i\hbar\frac{\mathrm{d}\hat{\rho}_T}{\mathrm{d}t} = \left[\hat{H}_\mathrm{T}, \hat{\rho}_\mathrm{T}\right]. \tag{2.4}$$

However, as the state of the environment is inaccessible, we focus on the state of the quantum system alone. By tracing over the environmental subsystem in Eq. (2.4), we obtain the master equation for the state of the open quantum system,

$\hat{\rho} = \mathrm{tr_E}\left[\hat{\rho}_\mathrm{T}\right] \in L\left(\mathcal{H}\right)$, as

$$i\hbar\frac{\mathrm{d}\hat{\rho}}{\mathrm{d}t} = \left[\hat{H}, \hat{\rho}\right] + \sum_i \gamma_i \left(\hat{L}_i \hat{\rho} \hat{L}_i^\dagger - \frac{1}{2}\left\{\hat{L}_i^\dagger \hat{L}_i, \hat{\rho}\right\}\right), \tag{2.5}$$

where $\hat{L}_i$ are known as "Lindblad operators", and γ_i are the relaxation rates corresponding to the coupling with the environment [59]. In Section 5.2, we will use this master equation extensively in the investigation of various error mechanisms in our local entangling gates.

As before, it is often useful to consider the evolution of a quantum system over a discrete time interval. For a closed system, this evolution is described by a unitary transformation. In contrast, for an open quantum system, the interaction with the environment leads to non-unitary dynamics. Despite this, the evolution must still preserve the positivity of the density matrix and should maintain its trace[1]. Such an evolution can be described by completely positive trace-preserving (CPTP) maps, $\mathcal{E}$, which act on the space of density operators[2], $\mathcal{E} : L\left(\mathcal{H}_1\right) \to L\left(\mathcal{H}_2\right)$. These maps describe general quantum *processes* and are a valuable tool used throughout this thesis.

Vectorisation

Before discussing the different representations of quantum processes, it is useful to take a moment to briefly introduce vectorisation – a useful tool for calculating and manipulating quantum processes [54, 60].

Vectorisation involves the reshaping of a matrix into a vector in a higher-dimensional space by stacking the columns of the matrix on top of each other[3] [54, 60]. The notation for this is $\hat{\rho} \in L\left(\mathcal{H}\right) \to |\hat{\rho}\rangle\rangle \in \mathcal{H} \otimes \mathcal{H}$. This can be illustrated with the simple

[1] Although non-trace-preserving processes can also be considered.
[2] Note that the input and output spaces need not be the same.
[3] The choice of stacking the columns, rather than the rows, is purely a convention choice. We will use the column-stacking convention throughout.

example,

$$\hat{\rho} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \rightarrow |\hat{\rho}\rangle\rangle = \begin{pmatrix} a \\ c \\ b \\ d \end{pmatrix}. \tag{2.6}$$

Among a number of useful properties of vectorisation, which are discussed in [60] and [54], is Roth's lemma [61], which states

$$|\hat{A}\hat{B}\hat{C}\rangle\rangle = \left(\hat{C}^T \otimes \hat{A}\right) |\hat{B}\rangle\rangle. \tag{2.7}$$

Representations of quantum processes

There are several different representations for a CPTP map, each with its own advantages for interpreting, calculating, and processing the represented map. In this section, we provide an overview of the representations used in this thesis.

Kraus representation: Let us start by first assuming the system and environment are initially in a separable state, such that $\hat{\rho}_T(t_1) = \hat{\rho}(t_1) \otimes |\phi\rangle \langle\phi|$ where $|\phi\rangle \in \mathcal{H}_\mathrm{E}$ is some arbitrary initial state of the environment. We would like to find a representation for the quantum process, $\mathcal{E}$, that describes the evolution of the system state under the discrete time interval between t_1 and t_2, such that $\hat{\rho}(t_2) = \mathcal{E}[\hat{\rho}(t_1)]$. First consider the unitary transformation of the total (closed) system,

$$\hat{\rho}_T(t_2) = \hat{U}\hat{\rho}_T(t_1)\hat{U}^\dagger \in \mathcal{H} \otimes \mathcal{H}_\mathrm{E}. \tag{2.8}$$

Tracing out the environment, we find that the state of the system is given by

$$\mathcal{E}(\hat{\rho}(t_1)) = \sum_i \hat{K}_i \hat{\rho}(t_1) \hat{K}_i^\dagger, \tag{2.9}$$

where

$$\hat{K}_i = \langle i| \hat{U} |\phi\rangle \tag{2.10}$$

are operators acting on states in $\mathcal{H}$, known as *Kraus* operators, and $|i\rangle \in \mathcal{H}_{\mathrm{E}}$ are states of the environment system.

Superoperator representation: Here, we make use of vectorisation, introduced briefly in the last section. The superoperator, $\hat{\mathcal{S}} : \mathcal{H}_1 \otimes \mathcal{H}_1 \to \mathcal{H}_2 \otimes \mathcal{H}_2$, implements

$$\hat{\mathcal{S}}|\hat{\rho}\rangle\rangle = |\mathcal{E}\left(\hat{\rho}\right)\rangle\rangle. \tag{2.11}$$

This superoperator representation can be calculated from the Kraus representation using Roth's lemma [54, 61],

$$\hat{\mathcal{S}} = \sum_i \hat{K}_i^* \otimes \hat{K}_i. \tag{2.12}$$

Choi matrix representation: We define the Choi matrix, $\hat{\Lambda} : \mathcal{H}_1 \otimes \mathcal{H}_2 \to \mathcal{H}_1 \otimes \mathcal{H}_2$, as

$$\hat{\Lambda} = \sum_{i,j} |i\rangle \langle j| \otimes \mathcal{E}\left(|i\rangle \langle j|\right). \tag{2.13}$$

In this form, we can write the action of the process as

$$\mathcal{E}\left(\hat{\rho}\right) = \mathrm{tr}_{\mathcal{H}_1}\left[\left(\hat{\rho}^T \otimes \hat{\mathbb{I}}_{\mathcal{H}_2}\right)\hat{\Lambda}\right]. \tag{2.14}$$

As shown in Wood, Biamonte, and Cory [54], this is equivalent to

$$\hat{\Lambda} = \sum_i |\hat{K}_i\rangle\rangle\langle\langle \hat{K}_i|, \tag{2.15}$$

where $|K_i\rangle\rangle$ are the vectorised Kraus operators. From Eq. (2.15), we see that the eigenvectors of the Choi matrix are the Kraus operators.

Process matrix representation: The final relevant representation is the process matrix, or χ matrix, representation. This representation is simply the Choi matrix under a basis transformation. Consider an operator basis comprising the set of orthonormal operators, σ_α, satisfying $\mathrm{tr}\left[\hat{\tau}_\alpha \hat{\tau}_\beta\right] = \delta_{\alpha\beta}$. If $\hat{T}$ represents the basis

transformation from the computational basis to the new basis $\{|\hat{\tau}_\alpha\rangle\rangle\}$, then we have that the process matrix in this basis is

$$\hat{\chi} = \hat{T}\hat{\Lambda}\hat{T}^\dagger. \tag{2.16}$$

We therefore have that

$$\hat{\Lambda} = \sum_{\alpha,\beta} \chi_{\alpha\beta} |\hat{\tau}_\alpha\rangle\rangle\langle\langle\hat{\tau}_\beta|. \tag{2.17}$$

Finally, we can write the action of the quantum process as

$$\mathcal{E}(\hat{\rho}) = \sum_{\alpha,\beta} \chi_{\alpha\beta} \hat{\tau}_\alpha \hat{\rho} \hat{\tau}_\beta^\dagger. \tag{2.18}$$

Since the processes we deal with all typically act on systems of N-qubits, we usually choose the N-qubit Pauli basis as the operator basis. This basis is constructed from products of the 1-qubit Pauli matrices, $\mathcal{P} = \{\hat{\mathbb{I}}/\sqrt{2}, \sigma_x/\sqrt{2}, \sigma_y/\sqrt{2}, \sigma_z/\sqrt{2}\}$. The Pauli operator basis for the N-qubit system is then $\mathcal{P}^{\otimes N}$. Throughout this thesis, we present representations of quantum processes by plotting the process matrices, $\chi_{\alpha,\beta}$, in the N-qubit Pauli basis, such that the Choi matrix can be written as

$$\hat{\Lambda} = \sum_{\alpha,\beta} \chi_{\alpha\beta} |\hat{\sigma}_\alpha\rangle\rangle\langle\langle\hat{\sigma}_\beta|. \tag{2.19}$$

2.2 Confinement of trapped-ions

With some of the more abstract theoretical framework in place, we now turn our attention to the theoretical tools that are essential for the trapped-ion quantum processor. In this section, we outline how ions are confined using a surface Paul trap and how this confinement enables precise control over their motion. Following a similar approach to Drmota [56], we begin with a classical treatment of ion motion before explaining how this motion can then be quantised.

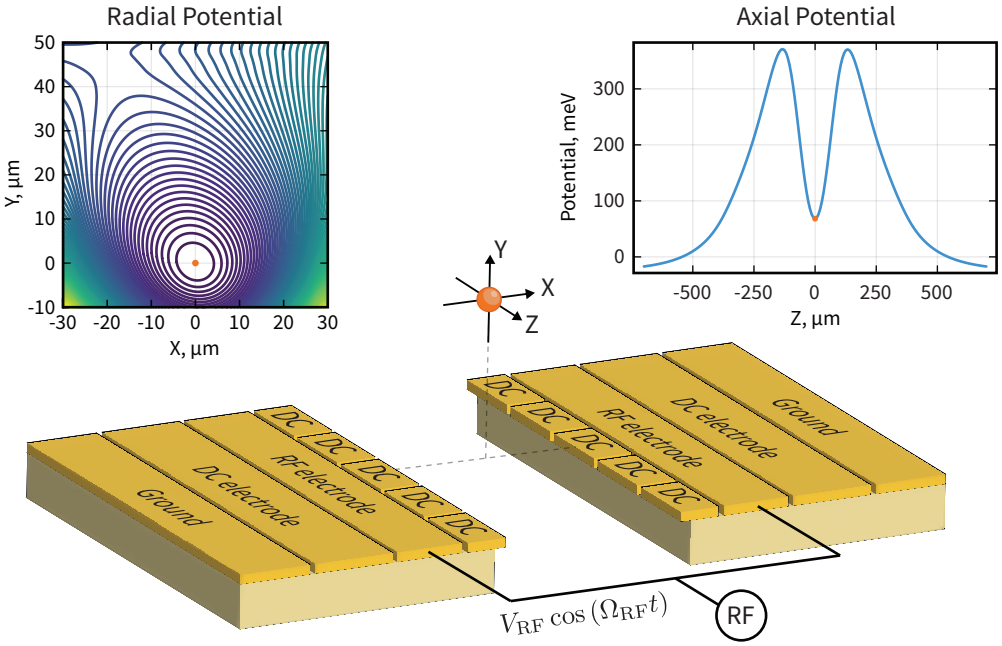

Figure 2.2: Geometry of a surface Paul trap used for ion confinement, representative of the traps employed in this thesis. An oscillating (RF) voltage is applied to the RF electrodes, while static (DC) voltages are applied to the DC electrodes, generating a time-averaged harmonic potential at the trap center. The resulting pseudo-potential is shown in the radial plane (left) and along the trap axis (right). Note that the origin of the coordinate system is at $\sim 70\,\mu$m above the trap surface.

2.2.1 Classical dynamics

We consider an ion to be a point particle of mass m and charge $q = +e$. We begin by noting that no configuration of static electric fields can realise a potential that is confining in all three directions. To show this, let us consider an electric field potential, $\phi(\vec{r})$, that is confining, such that expanding about the confining point yields

$$q\phi(\vec{r}) = \frac{1}{2}\left(\alpha x^2 + \beta y^2 + \gamma z^2\right). \tag{2.20}$$

For the field to be confining, we would need $\alpha, \beta, \gamma > 0$, however, the field must obey the Laplace equation, $\nabla^2\phi = 0$, and so we find that $\alpha + \beta + \gamma = 0$. Since there are no solutions with $\alpha, \beta, \gamma > 0$, we conclude that there is no configuration of static electric fields that is confining. In order to confine the ions, we use a combination of static and oscillating (radio-frequency (RF)) electric fields – known as a Paul trap

– such that the *time-averaged* field experienced by the ion is confining [55, 62]. In particular, for the work in this thesis, we make use of a surface Paul trap geometry, depicted in Figure 2.2.

Let us assume a trap geometry such that the oscillating RF field is applied with a cylindrical symmetry. Applying the pseudo-potential approximation, the time-averaged oscillating field gives rise to the effective potential given by

$$q\varphi_{\mathrm{RF}}(\vec{r}, m) = \frac{\kappa^2}{2m}(x^2 + y^2), \tag{2.21}$$

where the cylindrical axis is along z, and κ is a parameter which depends on the trap geometry and applied RF voltage [56]. In addition to the oscillating RF field, we also apply a static field

$$q\phi_{\mathrm{DC}}(\vec{r}) = \frac{1}{2}\left(-ax^2 - by^2 + (a+b)z^2\right), \tag{2.22}$$

where a and b are coefficients which depend on the trap geometry and voltages applied to various electrodes, such that $a + b > 0$, and the form of this potential is to ensure that it satisfies the Laplace equation.

Suppose we have N ions each with mass m_i confined to such a potential well. We can write down the Lagrangian for the system as $L = T - V$ where

$$T = \sum_{i=1}^{N} \frac{1}{2} m_i \left|\dot{\vec{r}}_i\right|^2, \tag{2.23}$$

$$V = \sum_{i=1}^{N} \left[q\phi_{\mathrm{DC}}(\vec{r}_i) + q\varphi_{\mathrm{RF}}(\vec{r}_i, m_i)\right] + \sum_{i,j=1, i\neq j}^{N} \frac{q^2}{4\pi\epsilon_0 |\vec{r}_j - \vec{r}_i|}, \tag{2.24}$$

are the kinetic and potential energies, respectively, and the last term in the potential energy describes the electrostatic repulsion between the ions. The goal is to find the $3N$ normal modes for this system.

We define $\vec{r}_{i,0}$ to be the equilibrium position of the i^{th} ion such that

$$\vec{\nabla} V \Big|_{\vec{r}_i = \vec{r}_{i,0}} = 0. \tag{2.25}$$

Furthermore, let us define a set of mass-weighted coordinates for each ion about their equilibrium positions,

$$\vec{q}_i = \sqrt{m_i} \left(\vec{r}_i - \vec{r}_{i,0} \right), \tag{2.26}$$

and concatenate them into a single vector $\vec{q} = (\vec{q}_1, \vec{q}_2, \ldots, \vec{q}_N)^T$. Note that q_i now denotes the i^{th} element of $\vec{q}$, which has $3N$ elements in total. By expanding the potential, V, about the equilibrium positions and applying the Euler-Lagrange equation, we obtain the equations of motion for the coordinates $\vec{q}$ as

$$\rightarrow \ddot{q}_i \approx -\sum_{j=1}^{3N} H_{ij} q_j, \tag{2.27}$$

where $H_{ij} = \frac{1}{\sqrt{m_i m_j}} \frac{\partial^2 V}{\partial r_i \partial r_j} \Big|_{\vec{q}=0}$ is the Hessian of the potential evaluated at the equilibrium positions [55]. Finally, we make use of the ansatz $\vec{q} = e^{i\omega t} \vec{\zeta}$, and obtain the eigenvalue equation

$$\omega^2 \zeta_i = \sum_{j=1}^{3N} H_{ij} \zeta_j. \tag{2.28}$$

Once the normal modes, $\{\alpha\}$, are found, we can decompose the motion of the i^{th} ion about its equilibrium position into a linear combination of the normal modes

$$\vec{q}_i(t) = \sum_{\alpha} a_\alpha e^{i\omega t} \vec{\zeta}_{i,\alpha} \tag{2.29}$$

where a_α denotes the amplitude of the normal mode, and $\zeta_{i,\alpha}$ describes the participation of the i^{th} ion in the mode α.

One ion

For the single ions case, the situation is relatively straightforward. We have that the equilibrium position is simply at $\vec{r}_0 = 0$, and the potential has the form

$$V(\vec{r}) = \frac{1}{2m}\left[\left(\frac{\kappa^2}{m} - a\right) q_1^2 + \left(\frac{\kappa^2}{m} - b\right) q_2^2 + (a+b)\, q_3^2\right]. \tag{2.30}$$

Immediately, we see that the Hessian takes the form $H = \mathrm{diag}\left(\omega_x^2, \omega_y^2, \omega_z^2\right)$, where

$$\omega_x = \sqrt{\frac{1}{m}\left(\frac{\kappa^2}{m} - a\right)}, \tag{2.31}$$

$$\omega_y = \sqrt{\frac{1}{m}\left(\frac{\kappa^2}{m} - b\right)}, \tag{2.32}$$

$$\omega_z = \sqrt{\frac{1}{m}(a+b)}. \tag{2.33}$$

The eigenvectors, $\vec{\zeta}_\alpha$, are then trivially aligned along the coordinate axes.

Two ions

We now come to the system of two ions of masses m_1 and m_2. This is of interest in this thesis, since the majority of the results presented here make use of crystals comprising two ions of different masses. Typically, we operate with the axial confinement weaker than the radial confinement, such that the ions form an axial crystal, and thus the equilibrium positions of the ions will lie along the z-axis. Evaluating Eq. (2.25), we find that the ions are located at $z = \pm\frac{d}{2}$, where d is the spacing of the ions given by

$$d = \left[\frac{q^2}{2m_1\omega_{1,z}^2\pi\epsilon_0}\right]^{1/3}, \tag{2.34}$$

where

$$\omega_{1,z} = \sqrt{\frac{1}{m_1}(a+b)} \tag{2.35}$$

is the axial frequency for a crystal comprising only the $i = 1$ ion.

The next step is to solve the eigenvalue equation $\omega_\alpha^2\vec{\zeta}_\alpha = H\vec{\zeta}_\alpha$, where in the

Mode α	**Eigenvalues**	**Eigenvectors**					
	ω_α (MHz)	$\zeta_{1,x}$	$\zeta_{1,y}$	$\zeta_{1,z}$	$\zeta_{2,x}$	$\zeta_{2,y}$	$\zeta_{2,z}$
1, axial IP	1.72	0	0	0.892	0	0	0.453
2	2.61	-0.999	0	0	0.053	0	0
3	3.19	0	-0.999	0	0	0.048	0
4, axial OOP	3.34	0	0	-0.453	0	0	0.892
5	6.14	0.054	0	0	0.999	0	0
6	6.67	0	0.048	0	0	0.999	0

Table 2.1: Eigenvectors and eigenvalues for a $^{88}Sr^{+}$-$^{43}Ca^{+}$ crystal. We assume typical single-$^{88}Sr^{+}$ mode frequencies of $\omega_x = 2.84\,\text{MHz}$, $\omega_y = 3.38\,\text{MHz}$, and $\omega_z = 1.52\,\text{MHz}$.

two-ion case, the eigenvectors are $\vec{\zeta}_\alpha = \left(\vec{\zeta}_{1,\alpha}, \vec{\zeta}_{2,\alpha}\right)^T$. This can be done analytically using symbolic programming.

While the analytical solutions can become quite messy, it is instructive to explore the mode structure of the two-ion crystal using some typical parameters. We consider the case where ion 1 is a $^{88}Sr^{+}$ ion and ion 2 is a $^{43}Ca^{+}$ ion and assume some typical single-$^{88}Sr^{+}$ mode frequencies of $\omega_x = 2.84\,\text{MHz}$, $\omega_y = 3.38\,\text{MHz}$, and $\omega_z = 1.52\,\text{MHz}$. The eigenvalues and eigenvectors are presented in Table 2.1. From these calculations, we see that the modes in the xy-plane will not be suitable for mediating interactions between the two ions since in all the normal modes in this plane, the participation is dominated by one of the ions. As a result, coupling the ions via these modes will result in a very weak interaction. We will therefore make use of the axial mode in which the two ions move out-of-phase – this is referred to as the axial out-of-phase (OOP) mode. Both ions participate relatively strongly in this mode, and the "breathing" motion of the mode means that the normal mode is relatively protected from heating due to common mode electric field noise.

The two eigenvalues for the axial modes can be calculated analytically as

$$\omega_{\text{oop}} = \omega_{1,z}\sqrt{\frac{1+\mu+\sqrt{\mu^2-\mu+1}}{\mu}}, \tag{2.36}$$

$$\omega_{\text{ip}} = \omega_{1,z}\sqrt{\frac{1+\mu-\sqrt{\mu^2-\mu+1}}{\mu}}, \tag{2.37}$$

where $\mu = m_2/m_1$. As we find in Section 5.2, the mass ratio $\mu \sim 0.49$ results in the second harmonic of the in-phase (IP) mode becoming near degenerate with the OOP mode, which can result in unwanted off-resonant excitation of this mode.

2.2.2 Quantisation of the motion

So far, we have treated the motion of the ions classically. However, in order to achieve coherent quantum control over the motion of the ions – thereby enabling the mediation of entangling interactions via this motion – we must move to a quantum mechanical picture of the motion.

We therefore treat each normal mode as a quantum harmonic oscillator with oscillator frequency ω_α. Each mode has an associated pair of creation and annihilation operators, $\hat{a}_\alpha$ and $\hat{a}_\alpha^\dagger$, respectively, which satisfy the bosonic commutation relations, $[\hat{a}_\alpha, \hat{a}_\alpha^\dagger] = 1$. The free Hamiltonian is then

$$\hat{H}_0 = \hbar \sum_{\alpha=1}^{3N} \omega_\alpha \left(\hat{a}_\alpha^\dagger \hat{a}_\alpha + \frac{1}{2} \right), \tag{2.38}$$

which has eigenvectors $|n\rangle$ corresponding to the energies $\hbar\omega_\alpha \left(n + \frac{1}{2}\right)$, for $n \in \mathbb{Z}$. Note that henceforth, we will neglect the zero-point energies, $\hbar\omega_\alpha/2$.

The eigenvectors, $|n\rangle$, span the infinite-dimensional Hilbert space for the quantum harmonic oscillator, $\mathcal{A} := \{|n\rangle\}_{n=0}^{\infty}$. Recall Eq. (2.29) in which we decompose the motion of the i^{th} about its equilibrium position into the normal motional modes. To quantise this motion, we promote $\vec{q}_i(t)$ to the operator

$$\hat{\vec{q}}_i = \sum_{\alpha=1}^{3N} \vec{\zeta}_{i,\alpha} l_{i,\alpha} \left(\hat{a}_\alpha + \hat{a}_\alpha^\dagger \right) \tag{2.39}$$

where $l_{i,\alpha} = \sqrt{\frac{\hbar}{2m_i\omega_\alpha}}$ is a characteristic length scale for the ground state wavefunction. Letting $\vec{r}_{0,i}$ be the equilibrium position for the i^{th} ion, we have that the position operator for the i^{th} ion is given by $\hat{\vec{r}}_i = \vec{r}_{0,i} + \hat{\vec{q}}_i$.

2.3 Ion-Light Interactions

The internal electronic states of atomic systems provide a rich setting for many quantum technologies, including quantum information processing. However, in order to utilise this structure, we must be able to coherently manipulate these states. In this section, we discuss the ion-light interaction through which we coherently manipulate the internal electronic states of trapped ions. For the work in this thesis, these interactions are realised using lasers.

We consider a single ion in a harmonic potential well. In addition to the 3 motional modes, discussed in the previous section, the ion has an internal electronic energy level structure, where the energies are $E_i = \hbar\omega_i$ and correspond to the orthogonal energy eigenstates $|i\rangle$. The total free Hamiltonian is therefore

$$\hat{H}_0 = \sum_i \hbar\omega_i |i\rangle \langle i| + \sum_{\alpha=1}^{3} \hbar\omega_\alpha \hat{a}_\alpha^\dagger \hat{a}_\alpha. \tag{2.40}$$

Note that often, we will denote the energy difference between two states as $\hbar\omega_{ij} = E_j - E_i$.

Here, we are not so concerned with experimental implementation of the laser systems; we simply assume that a laser is a classical electromagnetic plane wave, with an electric field given by

$$\vec{E}(\vec{r}, t) = \xi\vec{\epsilon}\cos\left(\vec{k}\cdot\vec{r} - \omega t + \phi\right), \tag{2.41}$$

where ξ is the amplitude of the electric field, $\vec{\epsilon}$ is the polarisation vector, describing the direction of electric field oscillation, $\vec{k}$ is the wavevector of the laser, $\omega = c\left|\vec{k}\right|$ is the frequency of the laser, and ϕ is the phase of the laser.

The laser interacts with the multipole moments of the ion. This interaction can be expanded in terms of the multipole moments [58, 63], such that

$$\hat{H}_{\mathrm{IL}} = \hat{H}_{\mathrm{ED}} + \hat{H}_{\mathrm{MD}} + \hat{H}_{\mathrm{EQ}} + \ldots, \tag{2.42}$$

where $\hat{H}_{\mathrm{ED}}$ describes the coupling to the electric dipole moment, $\hat{H}_{\mathrm{MD}}$ describes the coupling to the magnetic dipole moment, $\hat{H}_{\mathrm{EQ}}$ describes the coupling to the electric quadrupole moment, and so on.

Note that, for the work in this thesis, we do not consider magnetic dipole interactions. Instead, we focus on electric dipole interactions, which are often the dominant term in the expansion. When electric dipole interactions are forbidden, electric quadrupole interactions are considered as the next leading order term.

2.3.1 Electric dipole interactions

We begin with the electric dipole interaction. Here, the electric field component of the laser couples to the ion's electric dipole moment, described by the Hamiltonian

$$\hat{H}_{\mathrm{ED}} = -\hat{\vec{d}} \cdot \vec{E}(\hat{\vec{r}}, t) \tag{2.43}$$

$$= -\hat{\vec{d}} \cdot \vec{\epsilon}\, \xi \cos\left(\vec{k} \cdot \hat{\vec{r}} - \omega t + \phi\right), \tag{2.44}$$

where $\hat{\vec{d}} = -e\hat{\vec{r}}$ is the ion's dipole moment operator, and the position of the ion appears as the operator, $\hat{\vec{r}}$. Expanding in the energy eigenbasis, we obtain

$$\hat{H}_{\mathrm{ED}} = \sum_{i,j} \frac{\hbar\Omega_{ij}}{2} |i\rangle \langle j| \left[e^{i\left(\vec{k}\cdot\hat{\vec{r}} - \omega t + \phi\right)} + \mathrm{h.c.}\right], \tag{2.45}$$

where we have defined the Rabi frequency

$$\Omega_{ij} = -\frac{\xi \langle i| \hat{\vec{d}} \cdot \vec{\epsilon} |j\rangle}{\hbar}, \tag{2.46}$$

which determines the strength of the coupling of the laser to the transition $|i\rangle \leftrightarrow |j\rangle$.

Due to various symmetries of the dipole operator and the energy eigenstates, the matrix elements from Eq. (2.46) will vanish for particular transitions. This gives rise to selection rules on the transitions that can be driven in the electric dipole regime, typically depending on the quantum numbers representing the states involved in the transition. Transitions with vanishing electric dipole matrix elements are referred

to as electric-dipole *forbidden*, while those with non-vanishing matrix elements are electric-dipole *allowed*. In this thesis, we do not consider magnetic dipole transitions, and we frequently refer to electric-dipole-allowed transitions as dipole transitions.

Excited states connected to lower energy states via dipole-allowed transitions typically undergo rapid spontaneous emission, resulting in larger transition linewidths – or alternatively shorter natural lifetimes of the excited states – compared to dipole-forbidden transitions. The short natural lifetime of the excited states makes these transitions less suitable for coherently manipulating quantum information, since the spontaneous emission process destroys the encoded state. However, such transitions are extremely important for a variety of incoherent operations which require high photon scattering rates, such as laser cooling and state preparation.

AC Stark shifts

Let us consider a scenario in which a laser is far detuned from any transitions within the electronic structure of the ion. In this context, far detuned implies that the detuning of the laser from a nearest transition is large, compared to the strength of the coupling to the transition, i.e., $|\omega - \omega_{ij}| \gg \Omega_{ij}$.

For simplicity, let us begin by considering a single energy level, $|0\rangle$, which is connected to an excited state $|e\rangle$ via a dipole transition. Furthermore, let us neglect the movement of the ion, such that $\vec{k} \cdot \hat{\vec{r}} = \hat{\mathbb{I}}$ up to a constant complex phase, which we absorb into the laser phase, and then set $\phi = 0$. Under these simplifications, the total Hamiltonian becomes

$$\hat{H} = -\frac{\hbar\omega_{0e}}{2}\left(|e\rangle\langle e| - |0\rangle\langle 0|\right) + \frac{\hbar\Omega_{0e}}{2}\left(|0\rangle\langle e| + |e\rangle\langle 0|\right)\left[e^{-i\omega t} + \text{h.c.}\right]. \tag{2.47}$$

Entering the frame rotating with the laser frequency, ω, and applying the rotating wave approximation (RWA) to neglect rapidly oscillating terms, the total Hamiltonian becomes

$$\hat{\tilde{H}} = -\frac{\hbar\Delta}{2}\left(|e\rangle\langle e| - |0\rangle\langle 0|\right) + \frac{\hbar\Omega_{0e}}{2}\left(|0\rangle\langle e| + |e\rangle\langle 0|\right), \tag{2.48}$$

where $\Delta = \omega - \omega_{0e}$ denotes the detuning of the laser from the transition $|0\rangle \leftrightarrow |e\rangle$.

For $\Omega_{0e} = 0$, the unperturbed states, $|0\rangle$ and $|e\rangle$ correspond to the energy eigenvalues $+\frac{\Delta}{2}$ and $-\frac{\Delta}{2}$, respectively. However, in the far-detuned regime, i.e., $\Delta \gg \Omega_{0e}$, the second term results in a perturbation $\sim \mathcal{O}\left(\frac{\Omega_{0e}^2}{\Delta^2}\right)$, such that the perturbed eigenvalues are

$$\mp\frac{\Delta}{2}\left(1 + \frac{\Omega_{0e}^2}{2\Delta^2}\right). \tag{2.49}$$

We therefore find that the interaction with a far-detuned laser results in a shift of the energy levels by

$$\mp\hbar\frac{\Omega_{0e}^2}{4\Delta}. \tag{2.50}$$

This shift is known as the *AC Stark-shift*, although is frequently referred to as a *light-shift*.

Of course, to reach this result, we considered the ion to be a two-level system. However, typically many transitions will contribute to the shift. Therefore, suppose we are considering the state, $|i\rangle$, which is connected to the manifold $\{|e_j\rangle\}$ via dipole transitions. The total light shift of the state $|i\rangle$ will be the sum of the contributions from each transition $|i\rangle \leftrightarrow |e_j\rangle$, such that

$$\delta E_i = -\frac{\hbar\Omega_i}{2}, \tag{2.51}$$

where

$$\Omega_i = \xi^2 \sum_{\{|e_j\rangle\}} \frac{\langle i|\,\hat{\vec{d}}\cdot\vec{\epsilon}\,|e_j\rangle\,\langle e_j|\,\hat{\vec{d}}\cdot\vec{\epsilon}\,|i\rangle}{2\hbar\Delta_{ij}}, \tag{2.52}$$

where $\Delta_{ij} = \omega - \omega_{ij}$ is the detuning of the laser from the $|i\rangle \leftrightarrow |e_j\rangle$ transition.

2.3.2 Quadrupole interactions

Ignoring the magnetic dipole interaction, as it is not used in the work in this thesis, the next term of interest in the multipole expansion is the electric quadrupole

interaction. This interaction is described by the Hamiltonian

$$\hat{H}_{\mathrm{EQ}} = -\frac{1}{6}\sum_{m,n}\hat{Q}_{mn}\left.\frac{\partial E_n}{\partial x_m}\right|_{\vec{r}=\vec{0}}\cos\left(\vec{k}\cdot\hat{\vec{r}}-\omega t+\phi\right), \tag{2.53}$$

where $\hat{Q}_{ij}$ is a rank-2 tensor operator associated with the electric quadrupole moment [58, 63]. As before, we can expand this in terms of the energy eigenbasis as

$$\hat{H}_{\mathrm{EQ}} = \sum_{i,j}\frac{\hbar\Omega_{ij}}{2}\left|i\right\rangle\left\langle j\right|\left[e^{i\left(\vec{k}\cdot\hat{\vec{r}}-\omega t+\phi\right)}+\mathrm{h.c.}\right], \tag{2.54}$$

where we have defined the Rabi frequency

$$\Omega_{ij} = -\frac{1}{6\hbar}\sum_{m,n}\left\langle i\right|\hat{Q}_{mn}\left|j\right\rangle\left.\frac{\partial E_n}{\partial x_m}\right|_{\vec{r}=\vec{0}}. \tag{2.55}$$

Typically, electric quadrupole interaction strengths are smaller than electric dipole interaction strengths by a factor of the fine structure constant, $\alpha \approx 1/137$, [58].

Transitions that are forbidden in the electric dipole approximation but have non-zero electric quadrupole matrix elements are known as *quadrupole transitions*. Because electric quadrupole interactions are typically much weaker than electric dipole couplings, quadrupole transitions exhibit much narrower linewidths. Excited states connected to lower energy states by quadrupole transitions exhibit significantly longer natural lifetimes, often on the order of seconds, and are referred to as *metastable* states. The long lifetimes associated with quadrupole transitions make them more suitable for coherent manipulation of electronic states.

2.3.3 Raman interactions

So far, we have considered electric dipole and quadrupole interactions where transitions are directly driven using a single field. Here, we introduce Raman transitions, a two-photon process that allows us to indirectly address transitions that are forbidden in the electric dipole regime.

Consider two states, $|0\rangle$ and $|1\rangle$, connected by a transition of frequency ω_{01}, that

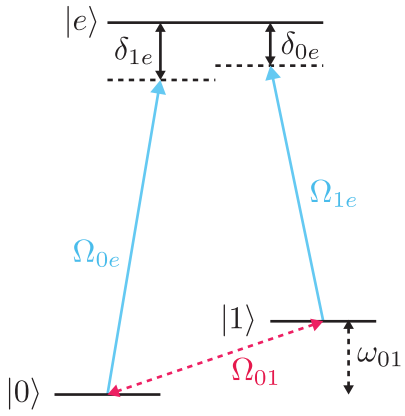

Figure 2.3: Energy level diagram for the implementation of a Raman interaction, utilising two lasers. The first (second) laser has a dipole Rabi frequency denoted Ω_{0e} (Ω_{1e}), wavevector $\vec{k}_0$ ($\vec{k}_1$), is detuned from the $|0\rangle \leftrightarrow |e\rangle$ ($|1\rangle \leftrightarrow |e\rangle$) transition by δ_{0e} (δ_{1e}). This configuration gives rise to an effective coupling between the states $|0\rangle$ and $|1\rangle$, with an effective coupling strength Ω_{01}.

is electric dipole forbidden. These states are, however, both connected to a third excited state, $|e\rangle$, via dipole-allowed transitions. In Raman interactions, depicted in Figure 2.3, two lasers are used to couple the states $|0\rangle$ and $|1\rangle$ via a two-photon process.

In the interaction picture, defined by the unitary transformation $\hat{U} = e^{-i\hat{H}_0 t/\hbar}$, the Hamiltonian for this interaction becomes

$$\hat{\tilde{H}} = \frac{\hbar\Omega_{0e}}{2} |e\rangle\langle 0| e^{i\left(\vec{k}_0\cdot\vec{r}-\delta_{0e}t+\phi_0\right)} + \frac{\hbar\Omega_{1e}}{2} |e\rangle\langle 1| e^{i\left(\vec{k}_1\cdot\vec{r}-\delta_{1e}t+\phi\right)} + \text{h.c.}. \tag{2.56}$$

For $|\delta_{1e} - \delta_{0e}| \ll \delta_{1e}, \delta_{0e}$, we can apply the James-Jerke approximation [57, 64], which yields the effective Hamiltonian

$$\begin{aligned}\hat{\tilde{H}} = &-\frac{\hbar\Omega_{0e}^2}{4\delta_{0e}}\left(|e\rangle\langle e| - |0\rangle\langle 0|\right) - \frac{\hbar\Omega_{1e}^2}{4\delta_{1e}}\left(|e\rangle\langle e| - |1\rangle\langle 1|\right) \\ &+ \frac{\hbar\Omega_{0e}\Omega_{1e}}{4\bar{\delta}_e}\left[|0\rangle\langle 1| e^{i\left(\delta\vec{k}\cdot\vec{r}-\Delta t+\Delta\phi\right)} + \text{h.c.}\right],\end{aligned} \tag{2.57}$$

where $\Delta\vec{k} = \vec{k}_0 - \vec{k}_1$, $\Delta = \delta_{0e} - \delta_{1e}$, $\Delta\phi = \phi_0 - \phi_1$, and

$$\bar{\delta}_e = \tfrac{1}{2}\left(\tfrac{1}{\delta_{0e}} + \tfrac{1}{\delta_{1e}}\right)^{-1} \tag{2.58}$$

The first two terms describe light-shifts of the two dipole transitions. The last term describes a coherent coupling of the two states $|0\rangle$ and $|1\rangle$. Note that by detuning the two Raman beams far enough away from the two dipole transitions, i.e., $\delta_{ie} \ll \Omega_{ie}$, we can assume that the resonant absorption of the Raman light via the dipole transitions becomes negligible.

The Raman interaction allows us to coherently drive the $|0\rangle \leftrightarrow |1\rangle$ transition, mediated via the $|e\rangle$ level. As in the case of the light-shift, the Raman beams will not just couple to the state $|e\rangle$, but all states that connect to $|0\rangle$ and $|1\rangle$ via dipole transitions. Denoting these states $\{|e_j\rangle\}$, we find that the effective Hamiltonian describing the Raman interaction between the states $|0\rangle$ and $|1\rangle$ is given by

$$\hat{\tilde{H}} = \frac{\hbar}{2}\Omega_{01}\left[|0\rangle\langle 1| e^{i\left(\Delta\vec{k}\cdot\vec{r}-\Delta t+\Delta\phi\right)} + \text{h.c.}\right], \tag{2.59}$$

where the total effective Rabi frequency is given by

$$\Omega_{01} = \xi_1\xi_2 \sum_{\{|e_j\rangle\}} \frac{\langle 0|\,\hat{\vec{d}}\cdot\vec{\epsilon}_1\,|e_j\rangle\langle e_j|\,\hat{\vec{d}}\cdot\vec{\epsilon}_2\,|1\rangle}{2\hbar\bar{\delta}_j}, \tag{2.60}$$

where ξ_m and $\vec{\epsilon}_m$ is the amplitude and polarisation vector of the m^{th} laser beam, respectively. Note that we have excluded the expressions for the light-shifts, but these are treated similarly.

2.4 Quantum information processing with ions

So far, we have discussed how linear Paul traps can be used to confine ions in harmonic potential wells, allowing their motion to be treated as that of quantum harmonic oscillators, and discussed the coherent manipulation of the electronic state of the ions using ion-light interactions. In this section, we explain how trapped ions can be utilised to realise a quantum information processor.

2.4.1 Qubits

We encode a qubit in the internal electronic state of a trapped ion by selecting a pair of energy levels to form the qubit states. We denote the qubit states as $\mathcal{Q} := \{|0\rangle, |1\rangle\}$. These states are separated by an energy $\hbar\omega_0$, where ω_0 is referred to as the qubit transition frequency. The qubit follows the formalism of a spin-$\frac{1}{2}$ system, and therefore the projections of the qubit state along the x, y, and z axes are associated with the Pauli matrices $\hat{\sigma}_x$, $\hat{\sigma}_y$, and $\hat{\sigma}_z$, respectively. The z-axis is referred to as the quantisation axis and is defined experimentally in the laboratory through an applied magnetic field.

The Hamiltonian describing the free evolution of a system of N qubits encoded in ions confined to a single harmonic well is given by

$$\hat{H}_0 = \sum_{i=1}^{N} \frac{\hbar\omega_0^{(i)}}{2}\hat{\sigma}_z^{(i)} + \sum_{\alpha=1}^{3N} \hbar\omega_\alpha \hat{a}_\alpha^\dagger \hat{a}_\alpha, \tag{2.61}$$

where the first term describes the free evolution of the qubits, and the second term describes the free evolution of the shared motional modes of the ions. Note that the total space on which this Hamiltonian acts is the product of the N qubit spaces, $\mathcal{Q}^{\otimes N}$, and the $3N$ Fock spaces, $\mathcal{A}^{\otimes 3N}$, with operators acting on their respective subspaces. The operators acting on individual subspaces are therefore implied to be embedded into this larger space; for example, $\hat{\sigma}_z^{(i)} = \hat{\mathbb{I}}_{\mathcal{Q}} \otimes \cdots \otimes \hat{\sigma}_z \otimes \hat{\mathbb{I}}_{\mathcal{Q}} \otimes \hat{\mathbb{I}}_{\mathcal{A}}^{\otimes 3N}$.

2.4.2 Single-qubit operations

In Section 2.3, we considered manipulation of the internal states of the ion using a single laser, via electric dipole and quadrupole interactions, as well as via Raman interactions using a two-photon process. Here, we discuss how such interactions can be used to mediate single-qubit rotations.

For the dipole and quadrupole interactions, we have shown that both interactions are described by similar Hamiltonians, Eq. (2.45) and Eq. (2.54), where the key difference is in the coupling strengths, Ω_{ij}. For now, let us ignore the physical

realisation of the interaction, and consider these interactions to take the form

$$\hat{H}_{\mathrm{IL}} = \frac{\hbar\Omega}{2}\left[\hat{\sigma}_+ + \hat{\sigma}_-\right]\left[e^{i(\vec{k}\cdot\hat{\vec{r}}-\omega t+\phi)} + e^{-i(\vec{k}\cdot\hat{\vec{r}}-\omega t+\phi)}\right], \tag{2.62}$$

where $\hat{\sigma}_+ = |1\rangle\langle 0|$ and Ω is the Rabi frequency describing the coupling of the laser to the qubit states encoded in the internal energy level structure of the ion.

Let us also neglect motional effects, i.e., assume that $e^{i\vec{k}\cdot\hat{\vec{r}}} = \hat{\mathbb{I}}$ up to a complex phase, which we can simply absorb into ϕ. Entering the frame rotating at the laser frequency, defined by the unitary transformation $\hat{U} = e^{-i\omega t\hat{\sigma}_z/2}$, the total Hamiltonian becomes

$$\hat{\tilde{H}} = -\frac{\hbar\Delta}{2}\hat{\sigma}_z + \frac{\hbar\Omega}{2}\left[e^{i\omega t}\hat{\sigma}_+ + e^{-i\omega t}\hat{\sigma}_-\right]\left[e^{-i(\omega t-\phi)} + e^{i(\omega t-\phi)}\right], \tag{2.63}$$

where $\Delta = \omega - \omega_0$ is the relative detuning of the laser and qubit transition. Applying the RWA, such that rapidly oscillating terms can be neglected, we are left with

$$\hat{\tilde{H}} = -\frac{\hbar\Delta}{2}\hat{\sigma}_z + \frac{\hbar}{2}\Omega\left(e^{i\phi}\hat{\sigma}_+ + e^{-i\phi}\hat{\sigma}_-\right). \tag{2.64}$$

Note that this form in Eq. (2.64) is also valid for the Raman transitions, described by the interaction in Eq. (2.59). The subsequent analysis therefore applies to dipole, quadrupole, and Raman interactions, such that the physical mechanism by which we realise the interaction, e.g., a single laser in the case of the dipole and quadrupole interactions, or a pair of lasers in the case of the Raman interactions, is generically referred to as a laser.

Now, when the laser is tuned to the qubit transition frequency, i.e., $\Delta = 0$, the frame rotating with the laser is also the frame rotating with the qubit transition. In this frame, the evolution over this discrete time interval is given by the unitary operator

$$\hat{R}(\theta,\phi) = \begin{pmatrix} \cos\left(\frac{\theta}{2}\right) & i\sin\left(\frac{\theta}{2}\right)e^{-i\phi} \\ i\sin\left(\frac{\theta}{2}\right)e^{i\phi} & \cos\left(\frac{\theta}{2}\right) \end{pmatrix}, \tag{2.65}$$

where $\theta = \Omega t$. Therefore, by "switching on" the laser for some time t with some phase ϕ, we can implement single-qubit rotations by the angle θ about the axis at an angle ϕ to the X-axis in the XY-plane of the Bloch sphere. To implement single-qubit rotations experimentally, we calibrate the durations that implement $\theta = \pi/2$, $\theta = \pi$, and $\theta = 2\pi$.

In practice, the absolute phase of the laser at the position of the ion is not known. However, we can define this phase to be zero at the start of an experimental sequence, and thus all rotations after this point will be referenced to this phase. We can therefore apply rotations about different axes by shifting the laser phase by ϕ from this absolute value for the duration of the rotation. Additionally, arbitrary rotations about the Z axis can be trivially implemented by simply adjusting this referenced value. Since the laser provides access to the rotating frame of the qubit, it is critical that the laser and qubit transition maintain a constant phase relation over the duration of an experimental sequence.

It is important to note that throughout this thesis we make use of a number of notations for single qubit gates. We denote rotations by θ about the axis i as $\hat{R}_i(\theta)$. Pulses with $\theta = \pi$, i.e., π-pulses, are written as a Pauli matrix, as $\hat{R}_i(\pi) = \hat{\sigma}_i$. In circuit diagrams, we depict such gates by the axis about which the flip is performed, i.e., X, Y, and Z. We also make use of S and H (Hadamard) gates, corresponding to the operations,

$$\hat{S} = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}, \tag{2.66}$$

$$\hat{H} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}. \tag{2.67}$$

Finally, in circuit diagrams we often make use of the compact notation $\theta|_i$ to depict a rotation of θ about some axis, i.

Qubit dephasing

As discussed above, it is critical that the laser and qubit transition remain phase coherent over the duration of an experimental sequence. To understand why this is important, let us consider the case in which the laser and qubit transition experience fluctuations, denoted by $\beta(t)$. Furthermore, we consider the simple act of storing the state of the qubit over some interval $[t_0, t_0 + \tau]$. In the frame rotating with the laser, the Hamiltonian over this interval is given by

$$\hat{\tilde{H}} = \frac{\hbar\beta(t)}{2}\hat{\sigma}_z. \tag{2.68}$$

Since at the end of this interval, we will typically want to use the laser to manipulate the stored information, these fluctuations will result in the *dephasing* of the qubit with respect to the laser. We assume that $\beta(t)$ describes Gaussian fluctuations with the time-average $\langle\beta(t)\rangle = 0$. Therefore, the statistical properties of the fluctuations are completely described by the two-point correlation function, $S(t_1 - t_2)$, or equivalently, the power spectral density [65]

$$\tilde{S}(\omega) = \int_{-\infty}^{\infty} \mathrm{d}t S(t) e^{-i\omega t}. \tag{2.69}$$

As outlined in Appendix A, the action of storing a quantum state while the system undergoes these fluctuations can be described by the dephasing process

$$\mathcal{E}(\hat{\rho}) = (1-p)\,\hat{\rho} + p\hat{\sigma}_z\hat{\rho}\hat{\sigma}_z. \tag{2.70}$$

The phase-flip probability, p, can be calculated from the power spectral density as

$$p = \frac{1}{2}\left[1 - \exp\left(-\int_0^{\infty} \frac{\mathrm{d}\omega}{\pi}\tilde{S}(\omega)\frac{F(\omega\tau)}{\omega^2}\right)\right], \tag{2.71}$$

where $F(\omega\tau)$ is the *filter function* [65], given by

$$F(\omega\tau) = \frac{\omega^2}{2}\left|\tilde{\Pi}\left(\omega;\tau\right)\right|^2, \tag{2.72}$$

where $\left|\tilde{\Pi}\left(\omega;\tau\right)\right|^2$ is the mod-squared value of the Fourier transform of the rectangle function, defined as $\Pi\left(t;t_0,t_0+\tau\right)=1$ if $t\in[t_0,t_0+\tau]$ and 0 otherwise.

We see from Eq. (2.71) that the filter function determines how noise at different frequencies contributes to dephasing. Filter functions can be engineered that suppress noise at different frequencies, using spin-flips to effectively invert the sign of the fluctuations – a technique known as *dynamical decoupling*. The simplest form of this technique is the Hahn spin-echo sequence, which applies a single π-pulse at the midpoint of the free evolution. This pulse effectively reverses the sign of any static detuning, $\beta(t)=\beta_0$, in the second half of the storage period, thereby cancelling out its effect on the quantum state. More advanced multi-pulse sequences, such as Carr-Purcell-Meiboom-Gill (CPMG) sequences [66], offer additional suppression of low-frequency noise. Moreover, by carefully varying the phases of the pulses, sequences have been developed that are robust to pulse errors and can protect arbitrary quantum states from decoherence. Particular examples of such pulse sequences used in this thesis are Knill dynamical decoupling [67], and universally robust dynamical decoupling [68]. Dynamical decoupling techniques are widely used across a variety of quantum technologies and are essential to the work presented in this thesis.

2.4.3 Two-qubit gates

The ability to mediate entangling operations between qubits encoded in the spin states of ions is a critical resource for the realisation of universal trapped-ion quantum processors. For ions situated in a potential well separated by $\sim\mu$m, the direct spin-spin coupling between two ions is not strong enough to mediate such interactions. However, the Coulomb repulsion between ions within a harmonic potential

well strongly couples the motion of ions. Therefore, by coupling the spin states of the ions to their motion, we can indirectly mediate the interaction between spin states of multiple ions via a shared motional mode.

In the Cirac-Zoller scheme [69], one of the first proposed entangling gates for trapped-ions, a laser maps the internal state of the i^{th} ion to a shared motional mode, α, by coherently driving a π-pulse on the red sideband of the motional mode:

$$|1_i\rangle |n_\alpha = 0\rangle \leftrightarrow |0_i\rangle |n_\alpha = 1\rangle . \tag{2.73}$$

Here, $|1_i\rangle$ and $|0_i\rangle$ represent the qubit states of the i^{th} ion, while $|n_\alpha = 0\rangle$ and $|n_\alpha = 1\rangle$ are the ground and first excited states of the shared motional mode, α. Next, a 2π-pulse is applied to the j^{th} ion on the red sideband transition between the $|1_j\rangle$ qubit state and an auxiliary state $|a_j\rangle$. If the motional mode is in the excited state $|n_\alpha = 1\rangle$, the j^{th} ion will acquire a phase of π in the $|1_j\rangle$ state. Finally, another π-pulse is applied to the i^{th} ion on the red sideband to transfer the motional state back to the qubit state of the i^{th} ion, disentangling the motional mode from the spin states. The overall result is a controlled-Z gate, where the state of the j^{th} ion undergoes a phase flip if and only if the i^{th} ion is in the state $|1_i\rangle$.

This method relies on the the shared motional mode being initially prepared in the ground state $|0_m\rangle$; unwanted population in higher Fock states results in a gate error. A more commonly employed class of gate – known as *geometric phase gates* – do not, within the Lamb-Dicke approximation, depend on the initial state of the harmonic oscillator, and are thus more robust to experimental imperfections [70]. Geometric phase gates have been used to demonstrate state-of-the-art two-qubit gates in trapped-ion quantum processors, including both laser-based [13, 14, 71] and laser-free mechanisms [16, 72, 73].

A key component of geometric phase gates is the generation of a spin-dependent force (SDF), which coherently drives the ions' motion depending on their collective spin state. The SDF drives a motional mode of the ions along closed trajectories in phase space, causing them to acquire a phase proportional to the area enclosed by

the trajectory – known as a *geometric phase*. Since the force depends on the ions' spin states, different spin configurations traverse different trajectories, leading to acquisition of different phases. By carefully tuning this interaction, this geometric phase mechanism can be used to implement entangling gates on the spins of the ions.

The light-shift spin-dependent force

In this thesis, we employ a specific realisation of this SDF, known as the light-shift force. Geometric phase gates utilising this SDF – known as light-shift gates – have been widely used to implement entangling gates in trapped-ion processors [13, 70, 74], including the current state-of-the-art for mixed-species entangling gates [15]. Accordingly, we chose this mechanism to implement the mixed-species entangling gates in our work. In this section, we discuss the theoretical framework underlying this force.

To understand this gate mechanism, first explore the interaction of a single ion in a harmonic potential well and two lasers, as in Schäfer [75]. For now, we consider just two states of the ion, one qubit state, $|0\rangle$, and an excited state $|e\rangle$, which are connected by a dipole transition. The i^{th} laser is detuned from this dipole transition, $|0\rangle \leftrightarrow |e\rangle$, by $\delta_{0e}^{(i)}$, such that the relative detuning of the two lasers is $\Delta\omega = \delta_{0e}^{(2)} - \delta_{0e}^{(1)}$. The Hamiltonian governing this interaction can be constructed from two instances of the interaction in Eq. (2.45), such that in the interaction picture, we obtain

$$\begin{aligned}\hat{\tilde{H}} = \; & \frac{\hbar}{2}\Omega_{0e}^{(1)} |0\rangle\langle e| \left[e^{i\left(\vec{k}_1\cdot\hat{\vec{r}} - \delta_{0e}^{(1)}t + \phi_1\right)} + \text{h.c.}\right] \\ & + \frac{\hbar}{2}\Omega_{0e}^{(2)} |0\rangle\langle e| \left[e^{i\left(\vec{k}_2\cdot\hat{\vec{r}} - \delta_{0e}^{(2)}t + \phi_2\right)} + \text{h.c.}\right],\end{aligned} \tag{2.74}$$

where $\Omega_{0e}^{(i)}$ is the Rabi frequency for the coupling of the i^{th} laser to the $|0\rangle \leftrightarrow |e\rangle$ dipole transition, $\vec{k}_i$ is the wavevector of the i^{th} laser, and ϕ_i is the phase of the i^{th} laser.

This is a similar scenario to the Raman interaction in Section 2.3.3, however

we do not choose the relative detuning to match the energy difference of two ion states, rather, we choose a smaller detuning that we will later choose to be close to a motional mode frequency. As in Section 2.3.3, we apply the James-Jerke approximation [64], and obtain

$$\hat{H} = \frac{\hbar\Omega_0}{2} |0\rangle \langle 0| \left[e^{i\left(\Delta\vec{k}\cdot\hat{\vec{r}} - \Delta\omega t + \Delta\phi\right)} + \text{h.c.} \right], \tag{2.75}$$

where $\Delta\vec{k} = \vec{k}_2 - \vec{k}_1$ is the relative wavevector, $\Delta\phi = \phi_2 - \phi_1$ is the relative phase of the two lasers, and Ω_0 is given by

$$\Omega_0 = \frac{\Omega_{0e}^{(1)}\Omega_{0e}^{(2)}}{2\bar{\delta}_{0e}} \tag{2.76}$$

$$= -\xi_1\xi_2 \frac{\langle 0| \hat{\vec{d}} \cdot \vec{\epsilon_1} |e\rangle \langle e| \hat{\vec{d}} \cdot \vec{\epsilon_2} |0\rangle}{2\hbar^2\bar{\delta}_{0e}}, \tag{2.77}$$

where ξ_i and $\vec{\epsilon_i}$ are the electric field amplitudes and polarisation vectors of the two lasers, and in similar fashion to Section 2.3.3, we define

$$\bar{\delta}_{0e} = \tfrac{1}{2}\left(\tfrac{1}{\delta_{0e}^{(1)}} + \tfrac{1}{\delta_{0e}^{(1)}} \right)^{-1}. \tag{2.78}$$

Now, we wish to a system of N ions, such that the qubit states of the i^{th} ion are labelled $|s_i\rangle \in \{|0\rangle, |1\rangle\}$. These qubit states are connected to the excited states $\{|e_j\rangle\}$ via dipole transitions. Following the same process as above, we can therefore construct the full Hamiltonian for this system as

$$\hat{H} = \sum_{i=1}^{N} \sum_{s=0,1} \frac{\hbar\Omega_{i,s}}{2} |s_i\rangle \langle s_i| \otimes \left[e^{i\left(\Delta\vec{k}\cdot\hat{\vec{r}}_i - \Delta\omega t + \Delta\phi\right)} + \text{h.c.} \right], \tag{2.79}$$

where $\Omega_{i,s}$ is the light-shift Rabi frequency for the spin state $|s_i\rangle$ of the i^{th} ion, given by

$$\Omega_{i,s} = -\frac{\xi_1\xi_2}{\hbar^2} \sum_{\{|e_j\rangle\}} \frac{\langle s_i| \hat{\vec{d}}_i \cdot \vec{\epsilon_1} |e_j\rangle \langle e_j| \hat{\vec{d}}_i \cdot \vec{\epsilon_2} |s_i\rangle}{2\bar{\delta}_{s_i e_j}}, \tag{2.80}$$

where $\hat{\vec{d}}_i$ is the electric dipole operator for the i^{th} ion.

Physically, this Hamiltonian describes the two lasers interfering at the position of the ions, giving rise to a spatially and temporally varying polarisation interference pattern, which has an envelope given by

$$\cos\Big(\Delta\vec{k}\cdot\vec{r}-\Delta\omega t+\Delta\phi\Big). \tag{2.81}$$

As the light-shift experienced by the qubit states depends on the polarisation at the position of the ion, the light-shifts experienced by the qubit states will also be spatially and temporally varying. Since the spin-states experience an energy shift that depends on their position in the interference pattern, this interaction realises a spin-dependent force.

Notice that the position, $\vec{r}_i$, of the ions appears as the operator, $\hat{\vec{r}}$. Recall from Section 2.2.2 that we can write this operator in terms of the creation and annihilation operators for the different motional modes, and thus the term $\Delta\vec{k}\cdot\hat{\vec{r}}_i$ becomes

$$\Delta\vec{k}\cdot\hat{\vec{r}}_i=\Delta\vec{k}\cdot\vec{r}_{0,i}+\sum_{\alpha=1}^{3N}\eta_{i,\alpha}\left(\hat{a}_\alpha+\hat{a}_\alpha^\dagger\right), \tag{2.82}$$

where we have defined the Lamb-Dicke parameter,

$$\eta_{i,\alpha}=\Delta\vec{k}\cdot\vec{\zeta}_{i,\alpha}l_{i,\alpha}, \tag{2.83}$$

for the coupling of the laser to the i^{th} ion participating in the motional mode α. We can Taylor expand the term $e^{i\vec{k}\cdot\hat{\vec{r}}_i}$, to obtain

$$e^{i\vec{k}\cdot\hat{\vec{r}}}=e^{\vec{k}\cdot\vec{r}_{0,i}}\left(1+i\sum_{\alpha=1}^{3N}\eta_{i,\alpha}\left(\hat{a}_\alpha+\hat{a}_\alpha^\dagger\right)-\left[\sum_{\alpha=1}^{3N}\eta_{i,\alpha}\left(\hat{a}_\alpha+\hat{a}_\alpha^\dagger\right)\right]^2+\dots\right). \tag{2.84}$$

In the case where $\eta_{i,\alpha}\ll 1$ and the motional modes are near their ground state, we may truncate this expansion after the second term – this is known as the Lamb-Dicke approximation.

Applying this approximation, and moving to the interaction picture, defined by

the unitary transformation $\hat{U} = e^{-i\hat{H}_0 t/\hbar}$, the total Hamiltonian becomes

$$\begin{aligned}\hat{\tilde{H}} =& \sum_{i=1}^{N} \sum_{s=0,1} \frac{\hbar\Omega_{i,s}}{2} e^{-i(\Delta\omega t - \Delta\phi_i)} \left|s_i\right\rangle \left\langle s_i\right| \\ & \otimes \left[1 + i \sum_{\alpha}^{3N} \eta_{i,\alpha} \left(e^{-i\omega_\alpha t} \hat{a}_\alpha + e^{i\omega_\alpha t} \hat{a}_\alpha^\dagger\right)\right] + \text{h.c.},\end{aligned} \tag{2.85}$$

where $\Delta\phi_i = \Delta\phi + \vec{k} \cdot \vec{r}_{0,i}$ is the relative phase of the two lasers at the equilibrium position of the i^{th} ion. Assuming that the detuning of the lasers is close to one of the mode frequencies, such that $\delta = \Delta\omega - \omega_\alpha$ where $|\delta| \ll |\Delta\omega|$, we can use the RWA to eliminate rapidly oscillating terms and obtain

$$\hat{\tilde{H}} = i \sum_{i=1}^{N} \sum_{s=0,1} \frac{\hbar\Omega_{i,s}\eta_{i,\alpha}}{2} e^{i\Delta\phi_i} e^{-i\delta t} \left|s_i\right\rangle \left\langle s_i\right| \otimes \hat{a}_\alpha^\dagger + \text{h.c.}. \tag{2.86}$$

We have used the slight abuse of notation to imply

$$\left|s_i\right\rangle \left\langle s_i\right| = \hat{\mathbb{I}}_{\mathcal{Q}} \otimes \cdots \otimes \left|s_i\right\rangle \left\langle s_i\right| \otimes \ldots \hat{\mathbb{I}}_{\mathcal{Q}}. \tag{2.87}$$

However, since it is diagonal in the computational basis, it is clear that the spin-eigenstates of the spin-dependent force will be eigenstates of $\bigotimes_{i=1}^{N} \hat{\sigma}_z$. We denote these eigenstates $|\psi_n\rangle$, where

$$\left|\psi_n\right\rangle = \bigotimes_{i=1}^{N} \left|s_{i,n}\right\rangle, \tag{2.88}$$

where $|s_{i,n}\rangle \in \{|0\rangle, |1\rangle\}$ is the spin-state of the i^{th} ion in the n^{th} eigenstate, $|\psi_n\rangle$. We can therefore write the Hamiltonian as

$$\hat{\tilde{H}} = i\hbar \sum_{n} \left|\psi_n\right\rangle \left\langle \psi_n\right| \otimes \left[\Omega_n(t) \hat{a}_\alpha^\dagger + \text{h.c.}\right]. \tag{2.89}$$

where

$$\Omega_n(t) = \lambda_n e^{-i\delta t}, \tag{2.90}$$

$$\lambda_n = \frac{1}{2}\sum_{i=1}^{N} \Omega_{i,s_{i,n}} \eta_{i,\alpha} e^{i\Delta\phi_i}. \tag{2.91}$$

Physically, this Hamiltonian corresponds to the ions in the collective spin-state $|\psi\rangle_n$ being coherently driven by a time-dependent force given by $\Omega_{\alpha,n}(t)\hat{a}_\alpha^\dagger$.

The light-shift geometric phase gate

With the derivation of the light-shift spin-dependent force, we would now like to understand how to apply this force to realise a geometric phase gate.

We begin by using a Magnus expansion [76] to exactly calculate the unitary propagator, obtaining

$$\hat{U}(t) = \sum_n |\psi_n\rangle\langle\psi_n| \otimes \hat{D}\left(\int_0^t \Omega_n(t_1)\,\mathrm{d}t_1\right) \exp\left[i\mathrm{Im}\left(\int_0^t \mathrm{d}t_1 \int_0^{t_1} \mathrm{d}t_2 \Omega_n^*(t_2)\Omega_n(t_1)\right)\right], \tag{2.92}$$

where $\hat{D}$ is the displacement operator, defined as

$$\hat{D}(\beta) = \exp\left(\beta\hat{a}_\alpha^\dagger - \beta^*\hat{a}_\alpha\right). \tag{2.93}$$

We therefore see that the ions in the collective spin state $|\psi_n\rangle$ will be driven along the phase space trajectory

$$\beta_n(t) = \int_0^t \Omega_n(t_1)\,\mathrm{d}t_1. \tag{2.94}$$

Examples of such phase space trajectories are shown in Figure 2.4.

Evaluating these trajectories for the light-shift force as given in Eq. (2.90), we obtain

$$\beta_n(t) = \frac{2\lambda_n}{\delta} e^{-i\delta t/2} \sin\left(\frac{\delta t}{2}\right). \tag{2.95}$$

For durations $t = t_g = 2m\pi/\delta$ for $m \in \mathbb{Z}$, the trajectory of each spin-state returns to

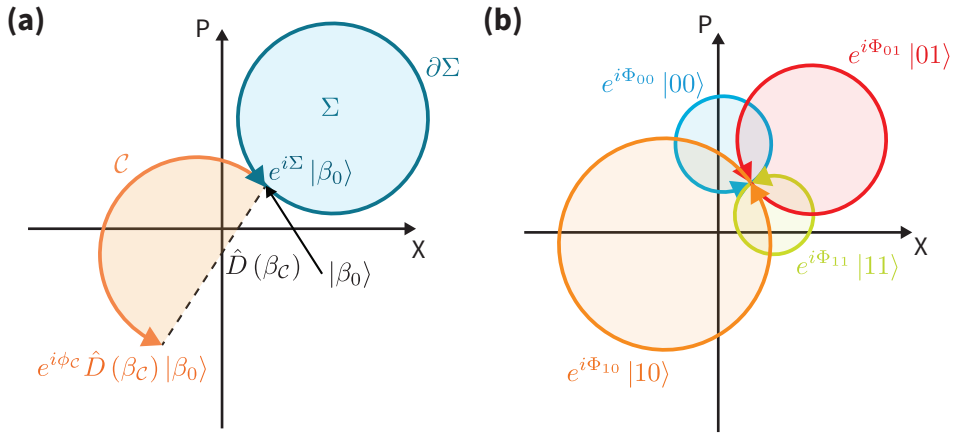

Figure 2.4: **(a)** Examples of closed and open phase space trajectories The closed trajectory $\partial\Sigma$ encloses the area Σ and therefore does not result in a net displacement of the harmonic oscillator. The harmonic oscillator state acquires a geometric phase equal to the area enclosed by its trajectory, $|\beta_0\rangle \rightarrow e^{i\Sigma}\,|\beta_0\rangle$. The open trajectory $\mathcal{C}$ does not return the harmonic oscillator to its original state and results in an overall displacement of $\beta_\mathcal{C}$ in addition to a phase shift. **(b)** Example of a geometric phase gate for two-qubits. The four spin states traverse different, but closed, trajectories, and thus each acquire different geometric phases, thus leading to an entangling interaction.

the initial state of the motional mode, and thus all trajectories are closed, as depicted in Figure 2.4(b). Along these phase space trajectories, the collective spin state of the ions is entangled with the shared motional mode, α, however by returning the motional mode to its original state, we disentangle the spin states from the motion. At these times, the propagator for the light-shift force becomes

$$\hat{U}(t_g) = \sum_n |\psi_n\rangle\langle\psi_n| \otimes \hat{D}\,(0) \exp\left[i\mathrm{Im}\left(\int_0^{t_g} \mathrm{d}t_1 \int_0^{t_1} \mathrm{d}t_2 \Omega_n^*(t_2)\Omega_n(t_1)\right)\right]. \quad (2.96)$$

We therefore see that by traversing the closed trajectory, the spin state $|\psi_n\rangle$ acquires a phase, equal to the area enclosed by its trajectory, given by

$$\Phi_n(t_g) = \mathrm{Im}\left(\int_0^{t_g} \mathrm{d}t_1 \int_0^{t_1} \mathrm{d}t_2 \Omega_n^*(t_2)\Omega_n(t_1)\right) \quad (2.97)$$

$$= \frac{|\lambda_n|^2}{2m\pi} t_g^2, \quad (2.98)$$

which is referred to as a geometric phase. However, since the operator $\hat{D}(0)$ corresponds to no displacement, the operator $\hat{U}(t_g)$ leaves the motional state in its

starting state. Thus, we can trace out the motional state to obtain the unitary evolution of the spin states

$$\hat{U}(t_g) = \sum_n |\psi_n\rangle \langle\psi_n| \, e^{i\Phi_n(t_g)}. \tag{2.99}$$

The collective spin eigenstates, $|\psi_i\rangle$, therefore acquire different phases, and, with appropriate choice of parameters t_g, δ, and $\Omega_{i,s}$, this can be used to generate entanglement between the spin states. With appropriate choices of t_g, δ, and the light-shifts $\Omega_{i,s_{i,n}}$, these interactions can be used to implement entangling gates.

As an example, let us consider the example of $N = 2$ ions and a single mode of motion. We consider further simplifications, such as the light-shifts experienced by the the ions are $\Omega_0 = -\Omega_1 = \Omega$, and the Lamb-Dicke parameters for the ions are the same, i.e., $\eta_{1,\alpha} = \eta_{2,\alpha} = \eta$. Writing each term explicitly, the Hamiltonian in Eq. (2.89) can be written as

$$\begin{aligned}\hat{\tilde{H}}_{\mathrm{LS},\alpha} = i\hbar\Omega\eta e^{-i\delta t} \Bigg[& \cos\left(\frac{\Delta\phi_{12}}{2}\right) |00\rangle\langle 00| \\ & + i\sin\left(\frac{\Delta\phi_{12}}{2}\right) |01\rangle\langle 01| \\ & - i\sin\left(\frac{\Delta\phi_{12}}{2}\right) |10\rangle\langle 10| \\ & - \cos\left(\frac{\Delta\phi_{12}}{2}\right) |11\rangle\langle 11| \Bigg] \otimes \hat{a}_\alpha^\dagger + \mathrm{h.c.},\end{aligned} \tag{2.100}$$

where $\Delta\phi_{12} = \Delta\phi_1 - \Delta\phi_2 = \Delta\vec{k} \cdot \Delta\vec{r}_{i,0}$ is the phase difference between the spin-dependent forces as the equilibrium positions of the two ions. By spacing the ions such that $\Delta\vec{k} \cdot \Delta\vec{r}_{i,0} = \pi$, we obtain

$$\hat{\tilde{H}}_{\mathrm{LS},\alpha} = i\hbar\Omega\eta e^{-i\delta t} \left(|01\rangle\langle 01| \otimes \hat{a}_\alpha^\dagger - |10\rangle\langle 10| \otimes \hat{a}_\alpha^\dagger \right). \tag{2.101}$$

By applying this interaction for the time $t_g = 2\pi/\delta$ that closes the trajectories, we

obtain the propagator

$$\hat{U}(t_g) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{-i\Phi} & 0 & 0 \\ 0 & 0 & e^{-i\Phi} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{2.102}$$

where

$$\Phi(t_g) = \frac{\Omega^2 \eta^2}{2\pi} t_g^2. \tag{2.103}$$

With appropriate choice of gate time, and corresponding δ, we can obtain $\Phi(t_g) = \frac{\pi}{2}$ at which point the unitary propagator in Eq. (2.102) becomes equivalent to the controlled-Z gate, up to single-qubit rotations.

So far in this section, we have considered the implementation of these interactions in the absence of any experimental imperfections. We discuss such imperfections in the discussion of our experimental implementation in Section 5.2.

3

Calibration and Characterisation Routines

*

This chapter presents the calibration and characterisation methods used in our experiment. I begin with the adaptive Bayesian phase-estimation protocol developed as part of my MPhys thesis [77], which builds on existing approaches [78, 79, 80, 81]. This protocol is employed extensively for gate calibration and magnetic-field stabilisation. Next, I review the well-established techniques of quantum state and process tomography [82, 83, 84]. Finally, I outline the randomised benchmarking protocol – another widely adopted characterisation method – following Magesan, Gambetta, and Emerson [85].

3.1 Phase estimation

Quantum systems are highly sensitive to their environment. For applications such as quantum computing, isolating the quantum system from unwanted environmental interactions is crucial for achieving high-precision control. Conversely, this inherent

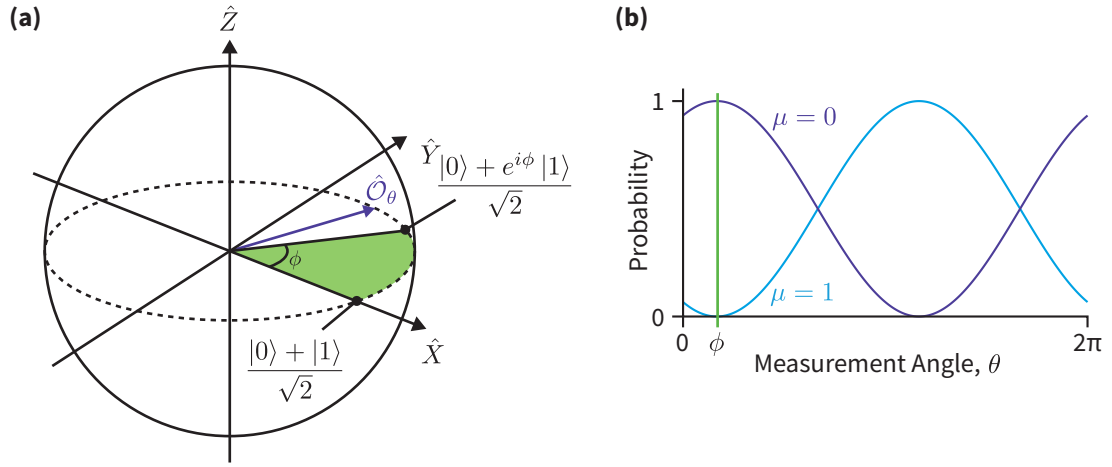

Figure 3.1: **(a)** Bloch sphere representation of a single Ramsey measurement. A qubit is prepared in the superposition, $(|0\rangle + |1\rangle)/\sqrt{2}$, and undergoes free evolution, in which time it acquires the phase, ϕ, about the Z axis, such that the qubit state becomes $\left(|0\rangle + e^{i\phi}|1\rangle\right)/\sqrt{2}$. After this evolution, a measurement is made in the XY plane of the Bloch sphere, along the axis at an angle θ from the X axis. **(b)** Example of a scan over the measurement angle θ. The probability for $\mu = 0$ ($\mu = 1$) is maximal (minimal) for $\theta = \phi$.

sensitivity makes quantum systems well-suited for realising quantum sensing technologies. Various applications of quantum sensing have already been demonstrated, with applications including the detection of electric [86] and magnetic fields [87]. Generally, quantum sensing protocols involve preparing the quantum system in an initial state, exposing it to a target signal, and then measuring the system to extract information about the system's response to the signal. A particularly important class of these protocols is *phase estimation*. In phase estimation, exposure to the target signal induces a phase shift, ϕ, between two orthogonal quantum states, which is proportional to the integral of the target signal over the exposure time. By accurately determining this phase, one can extract information about the target signal.

Ramsey experiments form the building block of many phase estimation schemes. A single Ramsey experiment, or shot, is depicted in Figure 3.1 and consists of preparing the qubit in a superposition state at some time t_0,

$$|\psi(t_0)\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \tag{3.1}$$

and allowing the qubit to evolve between t_0 and $t = t_0 + \tau$, after which the state has acquired some phase $\phi(t_0; t)$,

$$|\psi(t; t_0)\rangle = \frac{|0\rangle + e^{i\phi(t;t_0)}|1\rangle}{\sqrt{2}}. \tag{3.2}$$

A measurement is then made in the XY-plane of the Bloch sphere at some angle θ, yielding a measurement outcome $\mu \in \{0, 1\}$. We describe the Ramsey shot by this measurement outcome, μ, the measurement phase angle, θ, and the exposure time, τ – we combine this into the notation $u = (\mu, \theta, \tau)$. The probability of getting the measurement outcome μ, given the phase acquired by the system is ϕ, is given by

$$\mathrm{P}(u|\phi) = \frac{1 + (-1)^{\mu} C(\tau) \cos(\phi(\tau) - \theta)}{2}, \tag{3.3}$$

where $C(\tau)$ describes a reduction in contrast due to decoherence.

One typical method of estimating the phase acquired by the qubit would be to perform N Ramsey shots at different phase angles and fit Eq. (3.3) to the data to extract an estimate of the phase $\phi(\tau)$. However, for a given τ, the precision of the phase estimate, σ_ϕ, is limited by the standard quantum limit, i.e., $\sigma_\phi \sim \mathcal{O}\left(N^{-1/2}\right)$.

3.1.1 Adaptive Bayesian phase estimation

Adaptive Bayesian phase estimation (ABPE) protocols have been proposed as an efficient way of obtaining the phase of a qubit [78, 79, 80, 81]. These protocols make use of Bayes' theorem to adaptively choose the measurement settings of each Ramsey shot, based on information obtained by the preceding shots. By choosing measurement settings at each shot which would provide the most information, the protocols achieve Heisenberg scaling of the precision, i.e., $\sigma_\phi \sim \mathcal{O}(N^{-1})$. The ABPE protocol we implement is based on elements from these protocols.

Bayesian phase estimation

Consider some base Ramsey duration, τ_0, such that the phase acquired by the qubit over this time is denoted ϕ. Therefore, if we run a Ramsey shot for some integer multiple M of this base duration, $M\tau_0$, the phase acquired will be $M\phi$. We can therefore denote the i^{th} Ramsey shot by $u_i = (\mu_i, \theta_i, M_i)$. Furthermore, we denote the history of a set of l Ramsey experiments as

$$n_l = (u_1, u_2, \ldots, u_l)\,. \tag{3.4}$$

Let us assume that we have performed $l-1$ Ramsey shots, after which our knowledge of the qubit phase is described by the probability distribution $\mathrm{P}\left(\phi|n_{l-1}\right)$. We then perform another Ramsey shot and obtain u_l with probability given by Eq. (3.3). We can then use Bayes' theorem to update our probability distribution representing our knowledge of the qubit phase according to

$$\mathrm{P}\left(\phi|n_l\right) \propto \mathrm{P}\left(u_l|\phi\right)\mathrm{P}\left(\phi|n_{l-1}\right). \tag{3.5}$$

Following this update, we ensure normalisation of the distribution such that $\int \frac{\mathrm{d}\phi}{2\pi}\mathrm{P}(\phi) = 1$. The probability distribution can be used to extract information about our knowledge of the qubit phase, such as the estimation qubit phase

$$\hat{\phi} = \arg\left\langle e^{i\phi}\right\rangle, \tag{3.6}$$

and the Holevo variance [80, 88]

$$\sigma_\phi^2 = \left\langle \cos\left(\phi - \hat{\phi}\right)\right\rangle^{-2} - 1, \tag{3.7}$$

which is frequently used to calculate variances of probability distributions over cyclic coordinates.

In order to simplify many of the calculations involved, we express the probability

distribution as a discrete Fourier series [80, 81],

$$\mathrm{P}\left(\phi|n_l\right) = \sum_{n=-\infty}^{\infty} p_n^l e^{in\phi}. \tag{3.8}$$

Our distribution can therefore be described by the set of complex Fourier coefficients p_n^l, which must satisfy $p_{-n}^l = (p_n^l)^*$ to ensure the distribution is real. The Bayesian update in Eq. (3.5) can then be performed according to

$$p_n^l \to \frac{p_n^{l-1}}{2} + \frac{(-1)^{\mu_l} C(\tau)}{4}\left[e^{-i\theta_l} p_{n-M_l}^{l-1} + e^{i\theta_l} p_{n+M_l}^{l-1}\right], \tag{3.9}$$

followed by

$$p_n^l \to p_n^l \times \left|p_n^l\right|^{a^2-1}, \tag{3.10}$$

where $a \geq 1$ parameterises an artificial broadening of the probability distribution. This artificial broadening protects against numerical instabilities occurring when the probability distribution becomes narrow compared to the resolution of the truncated Fourier series. Larger values of a lead to a more extreme broadening of the distribution. Empirically, we found that $a = 1.002$ was a suitable parameter for most cases. This instability could be avoided by using a different representation of the probability distribution $\mathrm{P}\left(\phi|n_l\right)$, such as a particle filter approach [79]. However, these methods can be computationally expensive, particularly if one would like to perform the entire protocol on an field programmable gate array (FPGA).

Finally we ensure normalisation of the distribution by performing

$$p_n^l \to \frac{p_n^l}{p_0^l}. \tag{3.11}$$

We model decoherence in our system by using a Gaussian model,

$$C(\tau) = e^{-\tau^2/\tau_c^2}, \tag{3.12}$$

where τ_c is some timescale decoherence in the system. With this inclusion of deco-

herence into our model, we are able to balance the increased sensitivity associated with probing at long Ramsey durations, with the reduced information extracted at these Ramsey durations due to the loss of contrast. This discrete Fourier representation of our probability distribution also allows us to find simple expressions for the phase estimator and Holevo variance,

$$\hat{\phi} = \arg\left\langle e^{i\phi}\right\rangle = \arg(p_{-1}) \tag{3.13}$$

$$\sigma_H^2 = \left\langle \cos\left(\phi - \hat{\phi}\right)\right\rangle^{-2} - 1 = |p_{-1}|^{-2} - 1. \tag{3.14}$$

Adaptive measurements

The adaptive nature of the protocol comes from the ability to choose the next measurement settings, based on the information gained from preceding measurements. Each shot has two measurement settings to be chosen: the Ramsey duration, $M\tau_0$, and the measurement phase, θ. When probing for the base Ramsey duration, τ_0, the phase can take a value between 0 and 2π. By increasing the Ramsey duration in a subsequent shot, the phase acquired will now be a narrower window of width $2\pi/M$. It is therefore important not to increase the Ramsey duration until the phase being estimated is known to be within a specific window (within some tolerance), before further narrowing that window. We choose the next Ramsey duration according to

$$\tau_l = M_l\tau_0, \text{ where } M_l = \left\lfloor \frac{\beta}{\sigma_H} \right\rfloor, \tag{3.15}$$

where $\beta = 0.4$ is chosen empirically. Adaptively choosing the Ramsey duration in this way means that once enough measurements have been made to sufficiently constrain the qubit phase at a particular duration, this duration can then be increased to increase the sensitivity of the Ramsey probe.

In addition to the Ramsey duration, we also adaptively choose the measurement phase, θ, such that the next shot will reveal the maximum amount of information about the phase, ϕ. In the work by de Neeve [81], the next measurement phase was chosen as the one that resulted in the biggest gain in Shannon entropy. Here, we

take a different approach by choosing the measurement angle that is expected to minimise the Holevo variance after the measurement.

Recalling the form of the Holevo variance in Eq. (3.7), it is clear that this quantity will be minimised when the quantity $\left\langle \cos\left(\phi - \hat{\phi}\right)\right\rangle$ is maximised. Therefore, given a measurement history n_{l-1}, choosing the next measurement phase, θ_l, can be mapped to the optimisation problem

$$\theta_l = \underset{\theta}{\operatorname{argmax}} \left[\operatorname{Re}\left\{g_l\left(\theta; n_{l-1}\right)\right\}\right], \tag{3.16}$$

where

$$g_l\left(\theta; n_{l-1}\right) = \left\langle e^{i\left(\phi-\hat{\phi}\right)}\right\rangle_l. \tag{3.17}$$

Here, $\langle \dots \rangle_l$ denotes the expectation value after the l^{th} measurement, which has not yet been made. Therefore, we must also average over the two possible measurement outcomes, $\mu_l \in \{0, 1\}$ and take into account the effect that each outcome would have on the probability distribution – and therefore the Holevo variance.

In order to calculate the expectation value in Eq. (3.17), we average over the possible outcomes for the next Ramsey measurement, which we describe by $u = (\mu, \theta, M\tau_0)$. The objective function can then be written as

$$g_l\left(\theta; n_{l-1}\right) = \sum_{\mu\in\{0,1\}} \mathrm{P}\left(u\right) \int_0^{2\pi} \frac{\mathrm{d}\phi}{2\pi}\, \mathrm{P}\left(\phi|u n_{l-1}\right)\; e^{i(\phi-\hat{\varphi})}, \tag{3.18}$$

where $\mathrm{P}\left(u\right)$ is the probability that the l^{th} Ramsey experiment is u, and $\mathrm{P}\left(\phi|u n_{l-1}\right)$ and $\hat{\varphi}\left(u, n_{l-1}\right)$ are the probability distribution and phase estimate for ϕ, respectively, given the set of Ramsey experiments $n_l = u n_{l-1}$. We can obtain an expression for $\hat{\varphi}\left(u, n_{l-1}\right)$ by considering how the Fourier coefficients would be updated after a Ramsey experiment u. Using Eq. (3.9), we obtain

$$p_{-1} \xrightarrow{u} \frac{p_{-1}^{l-1}}{2} + \frac{e^{-(M\tau_0)^2/\tau_c^2}}{4} (-1)^{\mu} \left[e^{-i\theta} p_{-1-M}^{l-1} + e^{i\theta} p_{-1+M}^{l-1}\right], \tag{3.19}$$

and finally $\hat{\varphi} = \arg\left(p_{-1}\right)$.

We have from Bayes' theorem that

$$\mathrm{P}\left(u\right)\mathrm{P}\left(\phi|un_{l-1}\right)=\mathrm{P}\left(u|\phi\right)\mathrm{P}\left(\phi|n_{l-1}\right), \tag{3.20}$$

and thus we find that Eq. (3.18) can be written as

$$g_l\left(\theta;n_{l-1}\right)=\sum_{\mu\in\{0,1\}}e^{-i\hat{\varphi}}\int_0^{2\pi}\frac{\mathrm{d}\phi}{2\pi}\,\mathrm{P}\left(u|\phi\right)\mathrm{P}\left(\phi|n_{l-1}\right)e^{i\phi}. \tag{3.21}$$

Making use of the Fourier expansion, this becomes

$$\begin{aligned}g_l\left(\theta;n_{l-1}\right)=&\sum_{\mu\in\{0,1\}}e^{-i\hat{\varphi}}\sum_{n=-\infty}^{\infty}p_n^{l-1}\times\\&\times\int_0^{2\pi}\frac{\mathrm{d}\phi}{4\pi}\,e^{i(n+1)\phi}\left[1+\frac{(-1)^{\mu}\,e^{-(M\tau_0)^2/\tau_c^2}}{2}\left(e^{i(M\phi-\theta)}+e^{-i(M\phi-\theta)}\right)\right].\end{aligned} \tag{3.22}$$

The integral over ϕ can be trivially implemented due to the orthogonality of the exponential functions, and thus we arrive at the final form of the objective function,

$$g_l\left(\theta;n_{l-1}\right)=\sum_{\mu\in\{0,1\}}\frac{e^{-i\hat{\varphi}}}{2}\left[p_{-1}^{l-1}+\frac{(-1)^{\mu}\,e^{-(M\tau_0)^2/\tau_c^2}}{2}\left(e^{-i\theta}p_{-M-1}^{l-1}+e^{i\theta}p_{M-1}^{l-1}\right)\right]. \tag{3.23}$$

The solution to the optimisation problem Eq. (3.16) is then obtained using a golden-section search algorithm [89]. Since the objective function only contains a sum over the two possible measurement outcomes, the optimisation is relatively inexpensive to calculate compared to, e.g., the Shannon entropy used by de Neeve [81]. While a typical desktop computer could perform this optimisation in $\sim 40\,\mu\mathrm{s}$, it would be interesting to investigate the performance of such an algorithm on an FPGA.

A generic algorithm for ABPE is presented in Algorithm 1, and depicted in Figure 3.2. This algorithm accepts an *Experiment* subroutine as an argument; the subroutine performs the experimental shot according to the given measurement settings and returns the description, u, of the shot. This provides a flexible framework for ABPE, allowing the protocol to be implemented in a range of scenarios.

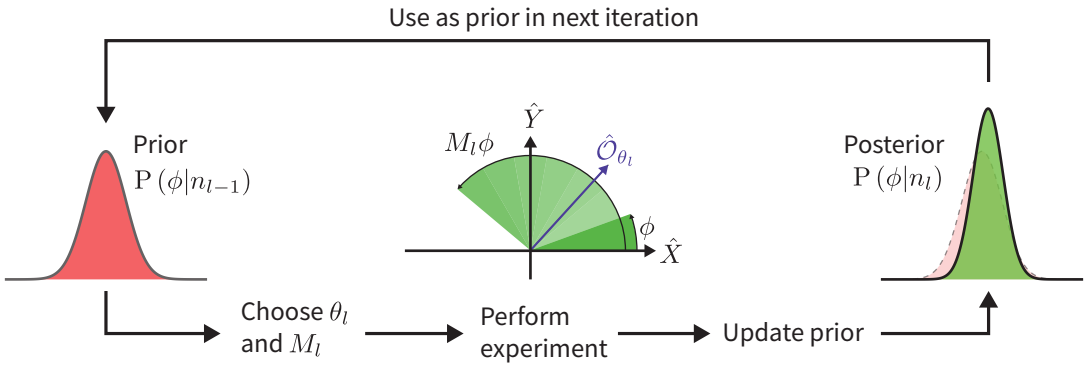

Figure 3.2: Adaptive Bayesian phase estimation protocol. Each iteration begins with a Bayesian prior, $P(\phi|n_{l-1})$. This prior is used to adaptively choose the measurement settings θ_l and M_l. The phase estimation experiment is performed using these settings, yielding the measurement outcome μ. This measurement outcome is used to update the prior distribution according to Eq. (3.9). This posterior distribution is then used as the prior for the next iteration.

Algorithm 1 Adaptive Bayesian Phase Estimation

Input:

N: The total number of experimental shots to perform, where $N \in \mathbb{N}$

a: Parameter for artificially broadening the distribution, where $a \geq 1$

τ_0: Base exposure duration, where $\tau_0 > 0$

Experiment: Subroutine that executes the experiment with the chosen measurement settings, returning the shot description $u = (\mu, \theta, M\tau_0)$

Output:

$\hat{\phi}$: Estimate of the phase acquired after a duration τ_0

σ_H^2: Holevo variance of the final probability distribution of ϕ

Initialise p_n^0 with uniform prior
for $l = 1$ **to** N **do**
  if $l = 1$ **then**
    $M_l \leftarrow 1$
    $\theta_l \leftarrow 0$
  else
    choose M_l - Eq. (3.15)
    choose θ_l - Eq. (3.16)
  end if
  $u_l \leftarrow$ EXPERIMENT(θ_l, M_l, τ_0) ▷ Perform experimental shot
  Update p_n^l - Eq. (3.9)
  Artificially broaden distribution - Eq. (3.10)
end for
Calculate $\hat{\phi}$ - Eq. (3.13)
Calculate σ_H^2 - Eq. (3.14)
return $(\hat{\phi}, \sigma_H^2)$

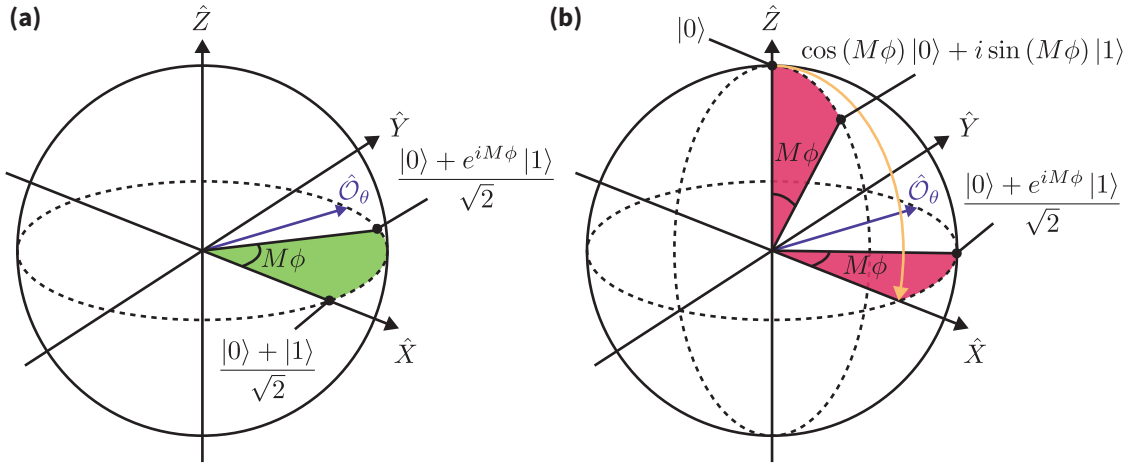

Figure 3.3: **(a)** Ramsey subroutine. The qubit is prepared in the superposition $(|0\rangle + |1\rangle)/\sqrt{2}$ and evolves for a time $M\tau_0$, after which time it acquires the phase $M\phi$. A measurement is made along the axis at an angle θ to the X axis in the equator of the Bloch sphere, yielding the outcome μ. **(a)** Rabi subroutine. The qubit is prepared in the state $|0\rangle$. M pulses of duration τ_0 and phase $\varphi = \pi$ are applied, rotating the qubit about the $-X$ axis by an angle $M\phi$. A $\pi/2$ rotation about the Y axis maps this rotation to a rotation in the XY plane. A measurement is made along the axis at an angle θ to the X axis in the equator of the Bloch sphere, yielding the outcome μ.

3.1.2 Optimal Ramsey experiments

As discussed previously, Ramsey experiments are used extensively in a range of quantum technologies. One particular use is in measuring the detuning between an oscillator and a qubit transition. Suppose the oscillator has some frequency, ω, and sets the reference frame for the qubit, which has a transition frequency ω_0. In the reference frame of the oscillator, the qubit evolves according to the propagator

$$\hat{U} = e^{-i\frac{\delta t}{2}\hat{\sigma}_z}, \tag{3.24}$$

where $\delta = \omega - \omega_0$ is the detuning between the oscillator and the qubit. As a result, over a time τ_0 the qubit will acquire a phase ϕ relative to the oscillator, given by $\phi = \delta\tau_0$. Therefore, by determining the phase acquired by the qubit over a time τ, one can find the detuning between the oscillator and the qubit transition as $\delta = \phi/\tau_0$. In our apparatus, this is used extensively to measure offsets from – and therefore calibrate – qubit transition frequencies.

Algorithm 2 Ramsey Subroutine

Input:
θ: Measurement angle, where $\theta \in [0, 2\pi]$
M: Ramsey exposure scale-factor, where $M \in \mathbb{N}$
τ_0: Base Ramsey exposure, where $\tau_0 > 0$

Output:
u: Measurement description

Prepare qubit in superposition $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$
Let qubit evolve freely for time $M\tau_0$
Measure qubit along axis at θ from the X-axis in the equator of the Bloch sphere
$\mu \leftarrow$ Measurement outcome
$u \leftarrow (\mu, \theta, M\tau_0)$
return u

We perform optimal Ramsey experiments using the ABPE algorithm outlined in Algorithm 1. The *Experiment* subroutine used for the Ramsey experiments is given in Algorithm 2, and depicted in Figure 3.3(a).

3.1.3 Optimal Rabi experiments

Another use for ABPE is for gate-time calibrations. Suppose we resonantly drive a qubit transition with some Rabi frequency Ω and phase φ, the qubit will evolve according to the propagator

$$\hat{U} = e^{-i\frac{\Omega t}{2}\hat{\sigma}_\varphi}. \tag{3.25}$$

Notice that this is very similar to the propagator from the Ramsey case, except now the rotations are about a different axis, with $\hat{\sigma}_\varphi = \cos(\varphi)\hat{\sigma}_x + \sin(\varphi)\hat{\sigma}_y$. We can therefore perform ABPE to efficiently measure the rotation angle, Ωt, for a particular gate duration by first performing gate pulses with phase $\varphi = \pi$, such that the qubit will undertake a circular trajectory in the ZY-plane of the Bloch sphere, followed by a $\pi/2$ rotation[1] about the Y-axis to map the ZY-plane to the XY-plane. With this gate rotation then effectively mapped to a rotation about the Z-axis, we

[1]Note that there is a circular dependency on the $\pi/2$-time, since we are using $\pi/2$ gates to implement the protocol. However, we have found that the protocol is relatively robust to miscalibration of this gate time.

perform the rest of the ABPE protocol as normal – thereby allowing measurement of the gate angle for a particular gate duration. The *Experiment* subroutine used for these Rabi experiments is given in Algorithm 3, and depicted in Figure 3.3(b). We use this technique frequently for efficient calibration of single-qubit gates.

Algorithm 3 Rabi Subroutine

Input:
θ: Measurement angle, where $\theta \in [0, 2\pi]$
M: Number of pulses, where $M \in \mathbb{N}$
τ_0: Base pulse duration, where $\tau_0 > 0$

Output:
u: Measurement description

Prepare qubit in state $|0\rangle$
Apply M pulses with duration τ_0 with pulse phase $\varphi = \pi$
Perform $\pi/2$ rotation at phase $\varphi = \pi/2$
Measure qubit along axis at θ from the X-axis in the equator of the Bloch sphere
$\mu \leftarrow$ Measurement outcome
$u \leftarrow (\mu, \theta, M\tau_0)$
return u

3.2 Quantum state tomography

A density matrix, $\hat{\rho}$, provides a complete description of the quantum state of a system. Quantum state tomography is the process by which measurements of a number of copies of a state are used to reconstruct the density matrix describing the state of a system.

Let us now consider what measurements are sufficient to gain complete information about an unknown quantum state. Consider the case of a single qubit, prepared in some unknown, normalised state $\hat{\rho} \in L(\mathcal{Q})$. The space of possible density matrices, $L(\mathcal{Q})$, is spanned by the set of Hermitian Pauli operators, $\hat{P}_\alpha \in \mathcal{P} = \{\hat{\mathbb{I}}/\sqrt{2}, \hat{\sigma}_x/\sqrt{2}, \hat{\sigma}_y/\sqrt{2}, \hat{\sigma}_z/\sqrt{2}\}$, which are orthogonal under the Hilbert-Schmidt inner product

$$\mathrm{tr}\left[\hat{P}_\alpha \hat{P}_\beta\right] = \delta_{\alpha\beta}. \tag{3.26}$$

Therefore, we can write $\hat{\rho}$ in terms of this basis as

$$\hat{\rho} = \frac{1}{2} \sum_{\alpha} \mathrm{tr} \left[\hat{P}_{\alpha} \hat{\rho} \right] \hat{P}_{\alpha}. \tag{3.27}$$

By measuring along the X, Y, and Z axes in the Bloch sphere, we can therefore determine the expectation values $\left\langle \hat{P}_{\alpha} \right\rangle = \mathrm{tr} \left[\hat{P}_{\alpha} \hat{\rho} \right]$ for $\alpha = 1, 2, 3$. Note that $\left\langle \hat{P}_0 \right\rangle = \mathrm{tr} \left[\hat{\mathbb{I}} \hat{\rho} \right] = 1$ from the normalisation of $\hat{\rho}$. Thus, the measurements along these axis of the Bloch sphere correspond to a *tomographically complete* set of measurements for a single qubit.

The generalisation to the N-qubit case is straightforward, since the space of N-qubit density matrices, $L\left(\mathcal{Q}^{\otimes N}\right)$, is spanned by tensor products of the Pauli operators, $\hat{P}_{\alpha} \in \mathcal{P}^{\otimes N}$. Thus, we can construct a tomographically complete set of measurements for the N-qubit case from measurements of each qubit individually; we need only single-qubit rotations and measurements to completely characterise the state $\hat{\rho}$ which, in general, may be an entangled state.

3.2.1 Maximum-likelihood estimation

Suppose we have a number of copies of some unknown quantum state that we can measure using some measurement apparatus. Now, suppose our measurement apparatus has a number of different measurement settings, labelled by the index i, and each measurement setting has a set of possible outcomes, labelled by the index j. The measurement outcomes for a particular measurement setting are associated with the positive operator-valued measures (POVMs), $\hat{M}_{ij}$, satisfying

$$\sum_{j} \hat{M}_{ij} = \hat{\mathbb{I}}, \tag{3.28}$$

such that the probability that a measurement of $\hat{\rho}$ using the measurement setting i yields the outcome j is given by

$$p_{ij} = \mathrm{tr} \left[\hat{M}_{ij} \hat{\rho} \right]. \tag{3.29}$$

Suppose we make measurements of our N copies of our state using different measurement settings and count how many times each measurement outcome occurs for each measurement setting, n_{ij}. Thus, the likelihood of obtaining a particular dataset, $\{n_{ij}\}$, given the state $\hat{\rho}$, can be written as

$$\mathcal{L}\left(\hat{\rho}\right) = \prod_{ij} \left(\mathrm{tr}\left[\hat{M}_{ij} \hat{\rho} \right] \right)^{n_{ij}} . \tag{3.30}$$

Therefore, we can find an estimation for the density matrix describing the state of the system by finding the density matrix, $\hat{\rho}$, which maximises the likelihood in Eq. (3.30).

After performing the tomographic measurements and collecting the raw data, we find the density matrix that maximises the likelihood by following an iterative diluted maximum-likelihood estimation algorithm developed by Řeháček *et al.* [83].

3.2.2 Tomographic measurements

Here, we consider the tomographic measurements and their respective POVMs which are utilised in this work. In particular, we consider ideal and nonideal measurements of ionic qubits. By modelling experimental imperfections of the measurements into the tomography, we can isolate features occurring in the reconstructed quantum state from those occurring as a result of nonideal measurement.

Ideal measurements

For a quantum system comprising N qubits, we perform a tomographic measurement by randomly selecting a measurement setting, i, for the each qubit. We perform a measurement with this setting by applying the single-qubit rotation

$$\hat{U}_i \in \left\{ \hat{\mathbb{I}},\ \hat{R}_X\left(\frac{\pi}{2}\right),\ \hat{R}_Y\left(\frac{\pi}{2}\right),\ \hat{R}_X\left(-\frac{\pi}{2}\right),\ \hat{R}_Y\left(-\frac{\pi}{2}\right) \right\} \tag{3.31}$$

to the n^{th} qubit and measuring in the computational basis, $\hat{\sigma}_z$. This corresponds to a measurement of the observable $\hat{U}_i^{\dagger} \hat{\sigma}_z \hat{U}_i$. Note that the set of rotations in Eq. (3.31)

corresponds to a tomographically overcomplete set of measurements, as we measure along the $\pm X$ and $\pm Y$ directions to improve robustness to systematic errors. Note that when selecting a measurement setting, we assign twice the weight to the $\hat{\mathbb{I}}$ element, ensuring that we acquire equal statistics along each axis.

Assuming perfect single-qubit rotations and qubit readout, we have that the POVMs for the n^{th} qubit can be written as

$$\hat{M}_{i_n,0} = \hat{U}_i^\dagger \left|0\right\rangle \left\langle 0\right| \hat{U}_i \tag{3.32}$$

$$\hat{M}_{i_n,1} = \hat{U}_i^\dagger \left|1\right\rangle \left\langle 1\right| \hat{U}_i. \tag{3.33}$$

Thus, the N-qubit POVMs may be constructed as

$$\hat{M}_{ij} = \bigotimes_{n=1}^{N} \hat{M}_{i_n,j_n}, \tag{3.34}$$

where the overall measurement setting, i, is the collection of single-qubit measurement settings, $(i_1, i_2, \ldots, i_N)$, and the overall measurement outcome, j, is the collection of the individual single-qubit measurement outcomes, $(j_1, j_2, \ldots, j_N)$.

Imperfect measurements

One of the key motivations for characterising various quantum states using state tomography is to calculate how errors in the creation of a state propagate through a circuit. This is particularly relevant for calculations where our remote entanglement serves as a resource state. Therefore, to obtain a better description of the created states, it is important to mitigate the effect of imperfections in the tomographic measurements. Otherwise, these imperfections may compound during such calculations. We account for the imperfect tomographic measurements by modifying the POVMs used to calculate the likelihood in Eq. (3.30).

For our trapped-ion processor, we discriminate between the two qubit states using fluorescence detection such that one of the qubit states scatterings cooling light, while the other does not. Let ϵ_0 and ϵ_1 denote the readout errors of qubit

states, such that ϵ_0 (ϵ_1) corresponds to the probability that a measurement of a qubit prepared in the state $|0\rangle$ ($|1\rangle$) will yield the measurement outcome 1 (0). This is encapsulated by the POVMs

$$M_0 = (1-\epsilon_0)\,|0\rangle\langle 0| + \epsilon_1\,|1\rangle\langle 1| \tag{3.35}$$

$$M_1 = (1-\epsilon_1)\,|1\rangle\langle 1| + \epsilon_0\,|0\rangle\langle 0|\,. \tag{3.36}$$

We therefore construct the single-qubit tomographic measurement POVMs as

$$M_{i,0} = (1-\epsilon_0)\hat{U}_i^\dagger\,|0\rangle\langle 0|\,\hat{U}_i + \epsilon_1\hat{U}_i^\dagger\,|1\rangle\langle 1|\,\hat{U}_i \tag{3.37}$$

$$M_{i,1} = (1-\epsilon_1)\hat{U}_i^\dagger\,|1\rangle\langle 1|\,\hat{U}_i + \epsilon_0\hat{U}_i^\dagger\,|0\rangle\langle 0|\,\hat{U}_i. \tag{3.38}$$

The readout errors, ϵ_0 and ϵ_1, for the different qubits are determined in Section 4.4.4. We neglect errors in the single-qubit rotations that implement $\hat{U}_i$, since these are typically smaller than the measurement errors, and we do not believe that we have a complete understanding of the structure of these errors.

3.2.3 Entanglement fidelity

Once we have reconstructed a state, $\hat{\rho}$, we can calculate the *fidelity* of the state to some target pure state, $|\psi\rangle$, as

$$\mathcal{F}\left(\hat{\rho}, |\psi\rangle\langle\psi|\right) = \langle\psi|\,\hat{\rho}\,|\psi\rangle\,. \tag{3.39}$$

Typically, the states we would like to characterise using state tomography will be bipartite entangled states. In such cases, we are often interested in quantifying the "amount" of entanglement in a state. We therefore make use of the *fully entangled fraction*, $\mathcal{F}\left(\hat{\rho}\right)$, which, for two qubits, is defined as the fidelity to a maximally entangled Bell-state allowing for arbitrarylocal rotations,

$$\mathcal{F}_{\mathrm{E}}\left(\hat{\rho}\right) = \max_{\hat{U}_1,\hat{U}_2} \left\langle\Phi^+\right|\left(\hat{U}_1\otimes\hat{U}_2\right)\hat{\rho}\left(\hat{U}_1\otimes\hat{U}_2\right)^\dagger\left|\Phi^+\right\rangle, \tag{3.40}$$

where we are performing the maximisation over the local rotations $\hat{U}_1$ and $\hat{U}_2$. Often, we will refer to this metric as the *entanglement fidelity*.

3.2.4 Resampling

Resampling of the measurement outcomes is used to generate new data sets which are analysed in the same way as the original data set, and are used to determine the sensitivity of the analysis to the statistical fluctuations in the input data. Error bars on the fidelities of reconstructed states are quoted as the standard deviation of the fidelities of the resampled data sets.

3.3 Quantum process tomography

As discussed in Section 2.1, closed quantum systems undergo unitary evolution. However, when a quantum system is open – that is, when it interacts with an environment – the evolution is generally not unitary. We represent a general quantum process by the CPTP map, $\mathcal{E} : L(\mathcal{H}_1) \rightarrow L(\mathcal{H}_2)$, which maps the density matrices $\hat{\rho} \in L(\mathcal{H}_1)$, to $\mathcal{E}(\hat{\rho}) \in L(\mathcal{H}_2)$. There are a number of representations for such maps, as discussed in Section 2.1, however we extensively make use of the process matrix representation, where the process $\mathcal{E}$ is represented by the matrix $\chi_{\alpha\beta}$ such that

$$\mathcal{E}(\hat{\rho}) = \sum_{\alpha,\beta} \chi_{\alpha\beta} \hat{P}_\alpha \hat{\rho} \hat{P}_\beta^\dagger, \tag{3.41}$$

where $\hat{P}_\alpha \in \mathcal{P}^{\otimes N}$ are operators from the set of N-qubit Pauli operators.

Quantum process tomography (QPT) is a technique by which we can reconstruct the process matrix representing a particular process. This is done by acting the process upon a set of input states, performing tomographic measurements on the output states, and thus using the outcomes to reconstruct the process matrix. Since the process matrix provides complete information about the action of a process on a particular quantum system, QPT is an important bench-marking tool for quantum

information processing (QIP) applications.

Let us consider the measurements we need to perform to implement QPT. Again, consider the case of a single qubit, $\mathcal{Q}$, which undergoes some process $\mathcal{E} : L(\mathcal{Q}) \to L(\mathcal{Q})$ which we would like to characterise. Suppose we take some set of input states, $\hat{\rho}_k \in \mathcal{I}$, which spans $L(\mathcal{Q})$, and reconstruct the image of these states, $\mathcal{E}(\hat{\rho}_k)$. Due to the linearity of the CPTP maps, we can deduce the action of the process on any arbitrary input via

$$\mathcal{E}\left(\sum_k a_k \hat{\rho}_k\right) = \sum_k a_k \mathcal{E}(\hat{\rho}_k). \tag{3.42}$$

Thus, the data acquisition for QPT will comprise tomographic measurements on the images of a set of input states which span the input space, $\mathcal{E}(\hat{\rho}_k)$. As with state tomography, generalising to the N-qubit case is straightforward because the set of tensor products of states from $\mathcal{I}^{\otimes N}$ spans the space $L\left(\mathcal{Q}^{\otimes N}\right)$.

3.3.1 Maximum-likelihood estimation

A single shot of the QPT protocol consists of choosing some input state from the spanning set, $\hat{\rho}_k$, performing the process on the state, $\mathcal{E}(\hat{\rho}_k)$, and then performing a tomographic measurement on the output state. As in Section 3.2.1, the measurement is associated with the POVMs $\hat{M}_{ij}$, where i indexes the measurement setting and j indexes the measurement outcome. Thus, the probability for measuring a particular outcome is given by

$$p_{ijk} = \mathrm{tr}\left[\hat{M}_{ij}\mathcal{E}(\hat{\rho}_k)\right]. \tag{3.43}$$

Let us now perform N shots and count the number of times each measurement outcome occurs for each input state and measurement setting, n_{ijk}. The likelihood of obtaining a particular dataset, $\{n_{ijk}\}$, given the process $\mathcal{E}$, can be written as

$$\mathcal{L} = \prod_{ijk}\left(\mathrm{tr}\left[\hat{M}_{ij}\mathcal{E}(\hat{\rho}_k)\right]\right)^{n_{ijk}}. \tag{3.44}$$

Therefore, our best estimate for the process is the one which maximises the likelihood in Eq. (3.44). Once we have performed the tomographic measurements and collected the raw data, to find the process which maximises the likelihood we perform an interactive diluted-maximum likelihood estimation algorithm [82, 83, 84].

3.3.2 State preparation and measurements

Here, we consider the preparation of the input states and the tomographic measurements used to execute the QPT protocol. As with state tomography, we consider the ideal case, followed by the treatment of experimental imperfections.

Ideal state preparation and measurements

The states $\hat{\rho}_k = |\psi_k\rangle\langle\psi_k| \in \mathcal{I}$ where

$$|\psi_k\rangle \in \left\{ |0\rangle, |1\rangle, \frac{|0\rangle + |1\rangle}{\sqrt{2}}, \frac{|0\rangle + i|1\rangle}{\sqrt{2}} \right\} \tag{3.45}$$

span the space of single-qubit states $L(\mathcal{Q})$, and thus form a set of input states for QPT. We prepare these states by first preparing the state $|0\rangle$, then applying a single qubit rotation $|\psi_k\rangle = \hat{V}_i|0\rangle$, where

$$\hat{V}_k \in \left\{ \hat{\mathbb{I}},\ \hat{R}_X(\pi),\ \hat{R}_Y\left(\frac{\pi}{2}\right),\ \hat{R}_X\left(-\frac{\pi}{2}\right) \right\}. \tag{3.46}$$

We therefore prepare the N-qubit input states by choosing an input state, k_n, for each qubit, such that

$$\hat{\rho}_k = \bigotimes_{n=1}^{N} \hat{V}_{k_n} |0\rangle\langle 0| \hat{V}_{k_n}^\dagger \in \mathcal{I}^{\otimes N}, \tag{3.47}$$

where $k = (k_1, k_2, \ldots, k_N)$ indexes the N-qubit input state.

The tomographic measurements are performed as for state tomography, with the N-qubit POVMs constructed as in Eq. (3.2.2).

Imperfect state preparation and measurements

Similar to state tomography, we want to suppress unwanted features occurring as a result of imperfect state preparation and measurement. We do this by incorporating these imperfections into the input states, $\hat{\rho}_i$, and POVMs, $\hat{M}_{ij}$, that occur in the likelihood in Eq. (3.44). Imperfections in the tomographic measurements are treated in the same as for state tomography, see Section 3.2.2.

We model the imperfect state preparation with some state preparation error parameter, ϵ, such that instead of the pure $|0\rangle$ state, we prepare

$$\hat{\rho}_{\text{prep}} = (1-\epsilon)\,|0\rangle\langle 0| + \epsilon\,|1\rangle\langle 1|\,. \tag{3.48}$$

We therefore have that the imperfect input states take the form

$$\hat{\rho}_i = (1-\epsilon)\,\hat{V}_i\,|0\rangle\langle 0|\,\hat{V}_i^\dagger + \epsilon\hat{V}_i\,|1\rangle\langle 1|\,\hat{V}_i^\dagger. \tag{3.49}$$

The state preparation errors, ϵ, for the different qubits are determined in Section 4.4.3. This model assumes that any imperfections in state preparation remain confined to the qubit subspace; in general, however, imperfect state preparation typically results in population outside that subspace. Nevertheless, for our purposes, this simplified model is a sufficient approximation that allows us to suppress features in the reconstructed process arising due to imperfect state preparation errors. As in the case of state tomography, we also neglect errors in the single-qubit rotations that implement $\hat{V}_i$ and $\hat{U}_i$, since these are typically smaller than the state preparation and measurement (SPAM) errors.

3.3.3 Average gate fidelity

Typically, the process we are characterising is ideally described by a unitary transformation, $\hat{U}$. Often, we would like to know how close the reconstructed "noisy" process, $\mathcal{E}$, is to the ideal unitary process. To quantify this, we make use of the

average gate fidelity as defined by Nielsen [90]

$$\bar{F}_{\mathcal{E},U} = \int \mathrm{d}\psi \, \langle\psi| \, U^\dagger \mathcal{E} \left(|\psi\rangle \langle\psi| \right) U \, |\psi\rangle \,, \tag{3.50}$$

where $\mathrm{d}\psi$ corresponds to the Haar measure.

Defining the process $\mathcal{E}'$ as the application of the process $\mathcal{E}$, followed by the inverse ideal process, $U^\dagger$, such that

$$\mathcal{E}'(\rho) = U^\dagger \mathcal{E}(\rho) U. \tag{3.51}$$

If the process, $\mathcal{E}$, perfectly implemented the unitary, $\hat{U}$, then the inverse ideal rotation would return the system to the state, such that $\mathcal{E}' \left(|\psi\rangle \langle\psi| \right) = |\psi\rangle \langle\psi|$. However, generally the resulting state has less than unit fidelity to the initial state,

$$F = \langle\psi| \, \mathcal{E}'(|\psi\rangle \langle\psi|) \, |\psi\rangle \leq 1, \tag{3.52}$$

where equality is only if $\mathcal{E}'$ perfectly implements the identity, i.e., if $\mathcal{E}$ perfectly implements $\hat{U}$. The average gate fidelity is obtained by then averaging over all possible pure input states,

$$\bar{F}_{\mathcal{E}'} = \int \mathrm{d}\psi \, \langle\psi| \, \mathcal{E}'(|\psi\rangle \langle\psi|) \, |\psi\rangle \,. \tag{3.53}$$

As shown by Wood, Biamonte, and Cory [54], this can be written in terms of the different representations of the process $\mathcal{E}'$,

$$\bar{F} = \frac{d + \mathrm{tr}\left[\hat{\mathcal{S}}'\right]}{d(d+1)} \tag{3.54}$$

$$= \frac{d + \langle\langle \hat{\mathbb{I}} | \hat{\Lambda}' | \hat{\mathbb{I}} \rangle\rangle}{d(d+1)} \tag{3.55}$$

$$= \frac{d + \sum_i \left| \mathrm{tr}\left[\hat{K}'_i\right] \right|^2}{d(d+1)} \tag{3.56}$$

$$= \frac{d + d\chi'_{00}}{d(d+1)}, \tag{3.57}$$

where $\hat{\mathcal{S}}'$, $\hat{\Lambda}'$, $\{\hat{K}_i'\}$, and $\hat{\chi}'$ are the superoperator, Choi, Kraus, and χ matrix representations of the process $\mathcal{E}'(\hat{\rho}) = \hat{U}^\dagger \mathcal{E}(\hat{\rho})\,\hat{U}$.

3.3.4 Resampling

Resampling of the measurement outcomes is used to generate new data sets which are analysed in the same way as the original data set, and are used to determine the sensitivity of the analysis to the statistical fluctuations in the input data. Error bars on the fidelities of reconstructed processes are quoted as the standard deviation of the fidelities of the resampled data sets.

3.4 Randomised benchmarking

While QPT provides a complete characterisation of a quantum process, it often requires a prohibitively large number of measurements to achieve a desired level of precision. To address this, characterisation protocols have been developed that provide a compromise between completeness and measurement efficiency. Randomised benchmarking (RBM) allows the characterisation of a gate set by extracting an average gate fidelity from significantly fewer measurements compared to QPT. Although RBM does not provide detailed information about the structure of the gate errors, it serves as a fast and efficient tool for benchmarking a quantum processor. Here, we provide a brief overview of the RBM protocol used in this thesis.

We follow the protocol proposed by Magesan, Gambetta, and Emerson [85], considering a system of N qubits. The protocol begins with the register of qubits prepared in the state $|0\rangle^{\otimes N}$. Next, we generate a sequence of $m-1$ gates,

$$\hat{S}_{m-1} = \hat{C}_{m-1} \ldots \hat{C}_1, \tag{3.58}$$

where each gate, $\hat{C}_i$, is uniformly sampled from the Clifford group. A final gate, $\hat{C}_m$,

is then chosen to invert the sequence, such that

$$\hat{C}_m = \hat{S}_{m-1}^{-1}. \tag{3.59}$$

Since the Clifford gates form a group, $\hat{C}_m$ also belongs to the Clifford group. For an ideal gate sequence, the resulting quantum channel, $\mathcal{E}$, corresponds to the identity operation, $\mathcal{E}(\hat{\rho}) = \hat{\rho}$, ensuring that the register returns to the state $|0\rangle^{\otimes N}$ with certainty.

However, in practice, the gates are imperfect. As a result, the process $\mathcal{E}$ deviates from the identity operation, and the probability of measuring the register in the state $|0\rangle^{\otimes N}$, known as the *survival probability* S, is given by

$$S = \mathrm{tr}\left[\left(|0\rangle\langle 0|\right)^{\otimes N} \mathcal{E}\left(\left(|0\rangle\langle 0|\right)^{\otimes N} \right) \right]. \tag{3.60}$$

We now make use of the fact that the Clifford group is a unitary 2 design, and thus the noisy process $\mathcal{E}$ can be "twirled" over the Clifford group. This twirling process effectively transforms the average noise process $\bar{\mathcal{E}}$ into a depolarising channel:

$$\bar{\mathcal{E}}\left(\hat{\rho}\right) = \left(1 - p(m)\right)\hat{\rho} + p(m)\frac{\hat{\mathbb{I}}}{d}, \tag{3.61}$$

where $d = 2^N$ is the dimension of the Hilbert space, and $p(m)$ is the depolarising probability for a sequence of m Clifford operations.

Let p denote the average depolarising probability for a single Clifford operation. Then the survival probability for different sequence lengths m can be modelled as

$$S(m) = A_1 + B_1 p^m, \tag{3.62}$$

where the parameters A_1 and B_1 account for SPAM errors. The depolarising probability p is related to the average error per Clifford operation, $\bar{\epsilon} = 1 - \bar{\mathcal{F}}$, by

$$\bar{\epsilon} = \left(1 - \frac{1}{d}\right) p. \tag{3.63}$$

By measuring the survival probability for different sequence lengths m and fitting the model in Eq. (3.62) to the data, we can extract p. In experimental implementations, the final Clifford gate applies additional Pauli operators $\hat{P}_i$, sampled from $\{\hat{\mathbb{I}}, \hat{\sigma}_x, \hat{\sigma}_y, \hat{\sigma}_z\}$, to each qubit, ideally leaving the register in the state

$$\left(\bigotimes_{i=1}^{N} \hat{P}_i\right) |0\rangle^{\otimes N}, \tag{3.64}$$

and thus the survival probability is calculated as the probability of measuring the register to be in this state. This allows us to avoid systematic errors from SPAM imperfections.

The RBM protocol thus provides an efficient way to determine an average gate fidelity, $\bar{\mathcal{F}} = 1 - \bar{\epsilon}$, for a set of Clifford operations. Sampling from the Clifford group, rather than all possible unitary operations, allows efficient classical simulation of the gate sequences, with the classical overhead associated with generating the gate sequence and calculating the final inverting gate scaling polynomially with the number of qubits. This makes RBM a scalable and practical technique for characterising large quantum systems.

4

Apparatus

*

In this chapter, I outline the apparatus used for our mixed-species quantum networking experiments. With two ion traps – each capable of trapping and manipulating two ion species – considerable effort has gone into the design and engineering, making a completely detailed description impractical in a single thesis. Three theses already cover the core setup in depth [91, 92, 93], and the reader is referred to those for full details. However, I led several subsequent projects not documented in these works – specifically, the characterisation of the Phoenix trap, the implementation of magnetic-field stabilisation, and the characterisation of our Raman laser system. I therefore present detailed accounts of these contributions here and provide a concise overview of the remaining components covered in the existing theses.

4.1 Choice of ion species

The choice of ion species for a trapped-ion apparatus depends on the specific application. For quantum information processing, it is important to select a species that supports high-fidelity entangling gates with minimal fundamental error sources. High-fidelity entangling gates have been demonstrated with various ion species, in-

cluding $^{9}Be^{+}$ [14], $^{25}Mg^{+}$ [16], and Ca^{+} [13, 17, 72]. In contrast, for atomic clocks, the ability to mediate entangling gates is less important. Instead, precise characterisation and control over systematic shifts that limit frequency measurements is critical. Ion species with magnetic-field-insensitive transitions, known as *clock* transitions, are particularly well-suited for this purpose. Other relevant factors include second-order Doppler shifts, black-body radiation shifts, and Stark shifts. Species such as $^{27}Al^{+}$ [94] and $^{171}Yb^{+}$ [95, 96], which exhibit such clock transitions, have been used to realise highly accurate atomic clocks. Practical factors, such as the availability of suitable laser systems and atomic sources for the ion species, also play a key role in the selection process.

For our quantum networking experiments, a number of properties are important to us. We require:

- the ability to interface with a quantum network – discussed in Section 6.1 – for which ions with a simple energy level structure are most suitable,
- the ability to store quantum information while we make use of this network interface; ions with hyperfine structure, thus exhibiting magnetic-field-insensitive transitions, are most suitable for this purpose,
- and finally, the ability to perform quantum logic between the ions within a module.

The requirements of both a simple level structure and a hyperfine structure are somewhat contradictory. $^{171}Yb^{+}$, with a nuclear spin of $\frac{1}{2}$, is a promising candidate that is being actively used for quantum networking [44, 97]. Additionally, $^{133}Ba^{+}$ exhibits a nuclear spin of $\frac{1}{2}$, however, there is still extensive work needed to develop the experimental techniques for this species [98, 99].

For the work in this thesis, instead of attempting to satisfy these requirements with a single species of ion, we delegate roles to two different species of ion: $^{88}Sr^{+}$, which has a simple level structure for realising the network interface, and $^{43}Ca^{+}$ for storing quantum information and executing quantum logic. Please note that

henceforth, we will drop the mass numbers from this notation and assume that Sr^+ and Ca^+ refer to these isotopes.

4.1.1 Calcium-43

The Calcium-43 ion has long been a favourite of the Oxford Ion-Trap Group for executing quantum logic, and has been central in achieving state-of-the-art results for quantum computing primitives, such as SPAM, single-, and two-qubit gates [10, 13, 72, 100]. It therefore makes a natural choice for performing the role of storage/logic ion in our modules. This isotope has a nuclear spin of $\frac{7}{2}$, which gives rise to hyperfine structure, as shown in Figure 4.1. The ground $S_{1/2}$ manifold is split into two hyperfine manifolds, corresponding to the quantum numbers $F = 3$ and $F = 4$, separated by $\sim 3.2\,\text{GHz}$. We apply a $0.496\,\text{mT}$ magnetic field to define the quantisation axis and lift the degeneracies of the Zeeman states. In total, we have 16 states in the ground $S_{1/2}$ manifold. We then choose pairs of states from this hyperfine manifold to form our qubits.

One choice of qubit from the hyperfine manifold, which is relevant to this thesis, is the *circuit* qubit, defined as $\mathcal{Q}_C := \{|0_C\rangle, |1_C\rangle\}$ where $|0_C\rangle = \left|S_{1/2}, F = 4, m_F = 0\right\rangle$ and $|1_C\rangle = \left|S_{1/2}, F = 3, m_F = 0\right\rangle$. At exactly zero magnetic field, the transition frequency connecting the qubit states has no first-order sensitivity to magnetic field fluctuations. Due to the application of the $0.5\,\text{mT}$ magnetic field to set the quantisation axis, this transition acquires a small, but non-zero first-order sensitivity[1] of $122\,\text{kHz}\,\text{mT}^{-1}$. In addition to the circuit qubit, we define an *auxiliary* qubit, defined as $\mathcal{Q}_X := \{|0_X\rangle, |1_X\rangle\}$ where $|0_X\rangle = \left|S_{1/2}, F = 4, m_F = +4\right\rangle$ and $|1_X\rangle = \left|S_{1/2}, F = 3, m_F = +3\right\rangle$. The auxiliary qubit has a significantly higher sensitivity to magnetic field fluctuations of $-24.5\,\text{MHz}\,\text{mT}^{-1}$, and thus it is not suitable

[1] Note that at this field, there is a field-insensitive optical transition which could be used. However, at the time of writing, the 729 nm laser used to address this transition is currently under construction. Additionally, the storage time of this optical qubit would then be limited by the $\sim 1\,\text{s}$ lifetime of the Ca^+ $D_{5/2}$ manifold. Alternatively, at approximately $14.6\,\text{mT}$, the ground hyperfine manifold of Ca^+ has another field-insensitive qubit between the states $\left|S_{1/2}, F = 4, m_F = 0\right\rangle$ and $\left|S_{1/2}, F = 3, m_F = +1\right\rangle$, however operating at this field would have drawbacks for the implementation of the network interface.

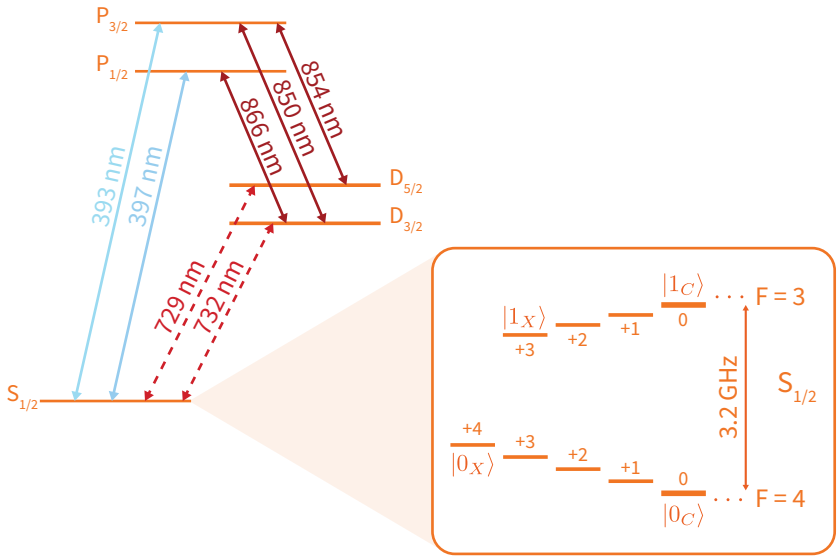

Figure 4.1: Energy level structure of Calcium-43, showing the relevant transitions for manipulating the electronic state of the Ca^+ ions. Transitions drawn with solid lines indicate the transitions which we address with lasers in our apparatus; transitions drawn with dashed lines are not addressed in our apparatus. The box shows the hyperfine structure of the ground $S_{1/2}$ manifold. This manifold is split into two sub-levels, labelled by the quantum numbers $F = 3, 4$, and split by $\sim 3.2\,\text{GHz}$. The application of a magnetic field splits the sub-levels into Zeeman states, labelled by the m_F quantum numbers. Within the ground state manifold, we define the circuit qubit as the pair of states $|0_\text{C}\rangle = \left|S_{1/2}, F = 4, m_F = 0\right\rangle$ and $|1_\text{C}\rangle = \left|S_{1/2}, F = 3, m_F = 0\right\rangle$. We also define the auxiliary qubit as the pair of states $|0_\text{X}\rangle = \left|S_{1/2}, F = 4, m_F = +4\right\rangle$ and $|1_\text{X}\rangle = \left|S_{1/2}, F = 3, m_F = +3\right\rangle$.

for storing quantum information. However, it is used in the state preparation and measurement of the circuit qubit, as well as for mediating mixed-species quantum logic (see Section 5.2).

4.1.2 Strontium-88

For realising a network interface, we choose the Strontium-88 ion; the energy level structure of this ion is shown in Figure 4.2. This species of ion does not have any nuclear spin, and thus doesn't exhibit hyperfine structure. Once again, the applied 0.496 mT magnetic field lifts the degeneracies in the fine structure via the Zeeman effect. In particular, the ground manifold is split into the two Zeeman states, $\left|S_{1/2}, m_J = -\frac{1}{2}\right\rangle$ and $\left|S_{1/2}, m_J = +\frac{1}{2}\right\rangle$, separated by a $\sim 14\,\text{MHz}$ transition. We define the *network* qubit as $\mathcal{Q}_\text{N} := \{|0_\text{N}\rangle, |1_\text{N}\rangle\}$ where $|0_\text{N}\rangle = \left|S_{1/2}, m_J = -\frac{1}{2}\right\rangle$

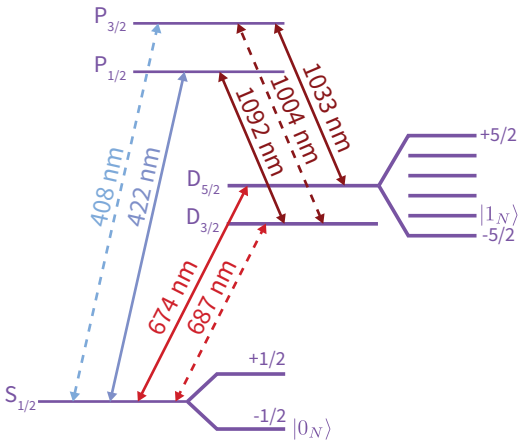

Figure 4.2: Energy level structure of Strontium-88, showing the relevant transitions for manipulating the electronic state of the Sr^+ ions. Transitions drawn with solid lines indicate the transitions which we address with lasers in our apparatus; transitions drawn with dashed lines are not addressed in our apparatus. Sr^+ does not have hyperfine structure. The application of a magnetic field splits the gross structure into Zeeman states, labelled by the m_J quantum numbers. We define the network qubit as the pair of states $|0_\mathrm{N}\rangle = \left|\mathrm{S}_{1/2}, m_J = -\frac{1}{2}\right\rangle$ and $|1_\mathrm{N}\rangle = \left|\mathrm{D}_{5/2}, m_J = -\frac{3}{2}\right\rangle$, which is connected by the optical 674 nm quadrupole transition.

and $|1_\mathrm{N}\rangle = \left|\mathrm{D}_{5/2}, m_J = -\frac{3}{2}\right\rangle$, which is separated by the optical 674 nm transition. The transition frequency of this qubit has a magnetic field sensitivity of $-11.2\,\mathrm{MHz\,mT^{-1}}$, which is the lowest of the accessible 674 nm quadrupole transitions[2], but is significantly higher than the sensitivity of the circuit qubit. While this species is therefore not suited to storing quantum information, the $\mathrm{P}_{1/2}$ level has a relatively short natural lifetime of ~ 7 ns, which, in addition to the simple level structure, enables the generation of ion-photon entanglement at a high rate. Therefore, the Sr^+ ion is well-suited for realising an interface to the optical quantum network.

Of course, Sr^+ is not the only species of ion which can be used to realise a network interface. Such interfaces have also been realised using $^{40}Ca^+$ [101], $^{111}Cd^+$ [102], $^{138}Ba^+$ [45, 46], and $^{171}Yb^+$ [43]. The choice of Sr^+ to provide our network qubit is motivated by a number of factors. Firstly, the similarity in mass to the Ca^+ ions means that Sr^+ can be efficiently used as a sympathetic coolant [103]. This enables

[2]The quadrupole transitions with $\Delta m_J = 0$ transitions exhibit a lower magnetic field sensitivity, $\pm 5.6\,\mathrm{MHz\,mT^{-1}}$. However, geometric constraints result in low coupling strengths for these transitions.

the cooling of ion crystals comprising Sr^+ and Ca^+ ions without destroying quantum states stored in the Ca^+ circuit qubits. Secondly, there is sufficient spectral isolation between the transitions of these two species of ion. This enables the addressing of each ion species with global beams, without needing to pay careful attention to cross-talk between the species. Finally, while the transitions are sufficiently different to provide the spectral isolation, the 408 nm and 397 nm dipole transitions in Sr^+ and Ca^+, respectively, are close enough that we can mediate mixed-species entangling gates using a single pair of Raman beams tuned to between these transitions [15], as outlined in Chapter 5.

4.2 Trapped-ion modules

Our quantum networking apparatus comprises two trapped-ion modules, Alice and Bob, which are connected via optical fibre links. Each trapped-ion module is capable of co-trapping Sr^+ and Ca^+ ions; in this work, we trap one ion of each species, in each module. An overview of a trapped-ion module is shown in Figure 4.3. In this section, we describe the hardware that is used to trap, cool, manipulate, and image the ions in each node.

4.2.1 Ion traps

Each module comprises a room-temperature micro-fabricated surface Paul trap fabricated by Sandia National Laboratories, housed in an ultra-high vacuum chamber ($< 1 \times 10^{-11}$ mbar) [91]. In Alice, we use the HOA2 trap [104]; in Bob, we use the Phoenix trap [105] – the successor to the HOA2 trap. The traps feature two RF electrodes and ~ 100 segmented DC control electrodes, enabling precise control over the trapping potentials and the execution of dynamic waveforms used for ion shuttling, splitting and merging of ion crystals, and for micromotion compensation. We use the Julia library SURF [106] to calculate trapping solutions and transport waveforms that are applied to the DC electrodes. The traps also feature a ~ 60 µm

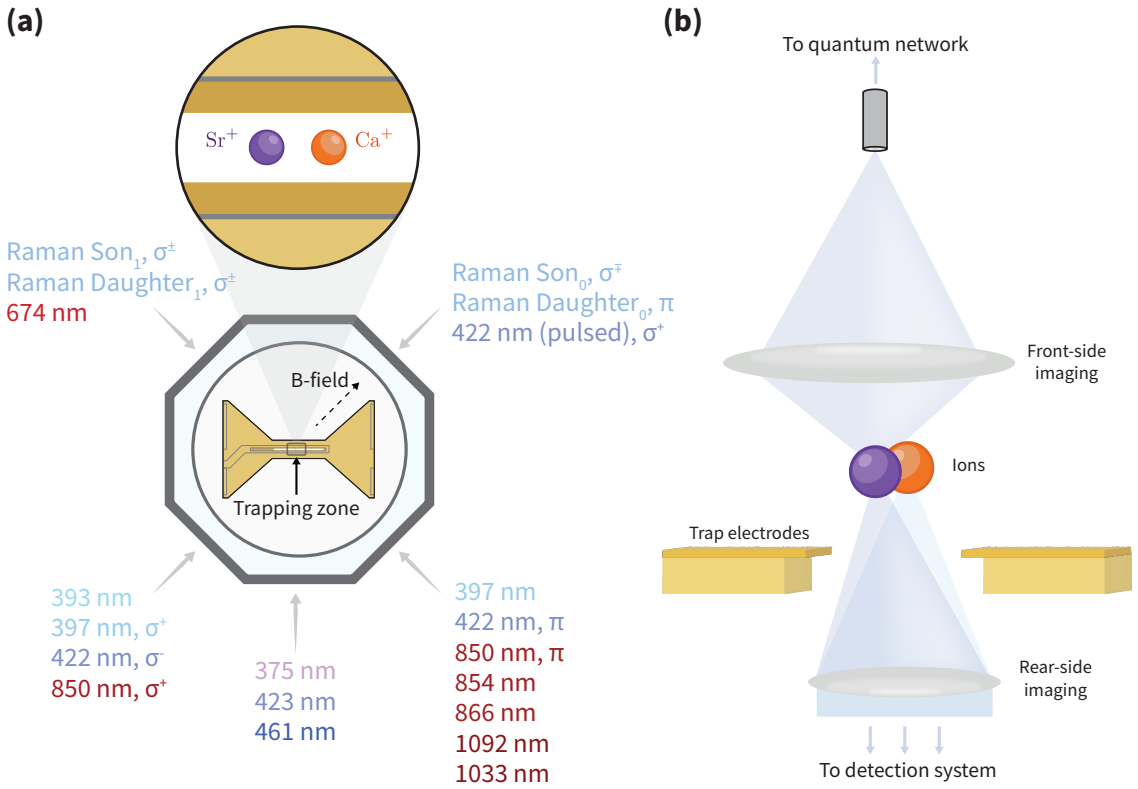

Figure 4.3: Overview of the trapped-ion modules. Ions are not to scale. **(a)** Top-view of one of the trapped-ion modules. A vacuum system contains a surface Paul trap which is used to co-trap a Ca^+ ion and a Sr^+ ion. The direction of the applied magnetic field is indicated. The various laser beam paths, and if necessary their polarisations, are listed, with arrows indicating the view-port by which the beams enter the vacuum chamber and impinge on the ions. **(b)** Side-view of the trapped-ion modules. A rear-side imaging system is used to collect light from both species of ion through a slot in the surface of the trap. This light is then sent to a detection system, where it is used to perform fluorescence detection of the ions. A front-side imaging system is used to collect single photons from the Sr^+ ions and couple them into a single-mode, non-polarisation maintaining fibre, which constitutes the optical quantum network.

slot in the trap surface, enabling optical access from the rear of the trap.

In each module, the RF electrodes are impedance matched to 50 Ω at a chosen frequency, Ω_{rf}, and are driven with a voltage of amplitude, V_{rf}. We apply voltages to the DC electrodes to provide an axial confinement. Additionally, we use the DC electrodes to lift the degeneracy in the radial mode frequencies, and ensure that the radial modes have a sufficient projection onto the Doppler cooling beams, travelling parallel to the trap surface. Typical trap parameters chosen for each module, as well as typical mode frequencies, are given in Table 4.1. Note the higher RF frequency and lower RF voltage in Bob, which results in an overall shallower trap. The higher

Parameter		Alice	Bob
Trap parameters	Ω_{rf}	50 MHz	60 MHz
	V_{rf}	180 V	135 V
Single Sr^+ ion secular frequencies	Axial	1.52 MHz	1.19 MHz
	Lower radial	2.84 MHz	1.91 MHz
	Upper radial	3.38 MHz	2.10 MHz

Table 4.1: Typical trapping parameters and single Sr^+ mode frequencies for each module.

RF frequency was chosen since empirically it was found to improve the lifetime of Sr^+-Ca^+ crystals. Ideally, the reduction in the RF pseudo-potential caused by this increase in RF frequency would be compensated for by an increase in the RF voltage. However, we found an interesting, yet extremely annoying, feature of the particular trap chip in Bob, such that when using voltages exceeding that given in Table 4.1, the trap would begin to charge. As a result, while it would be desirable to increase Bob's trap depth (thereby reaching lower heating rates and increasing separation between motional modes), it was not possible with this particular trap.

Any stray electric fields at the position of the ion will result in the ion being displaced from the RF null – this would result in unwanted excess micromotion. We compensate for these stray fields by applying DC voltages in order to return the ion to the null. We find the compensation voltages that need to be applied using a micromotion compensation technique based on parametric excitation of the radial modes [107].

Characterising the ion motion

As we discuss in Section 5.2, one of the key tools we have at our disposal is the ability to mediate entangling gates between the Sr^+ and Ca^+ ions. Since entangling gates between ions typically utilise the shared motion to mediate the interaction, we require precise control over the quantum state of the motion. It is therefore essential to understand and characterise effects such as motional heating and dephasing of this motion.

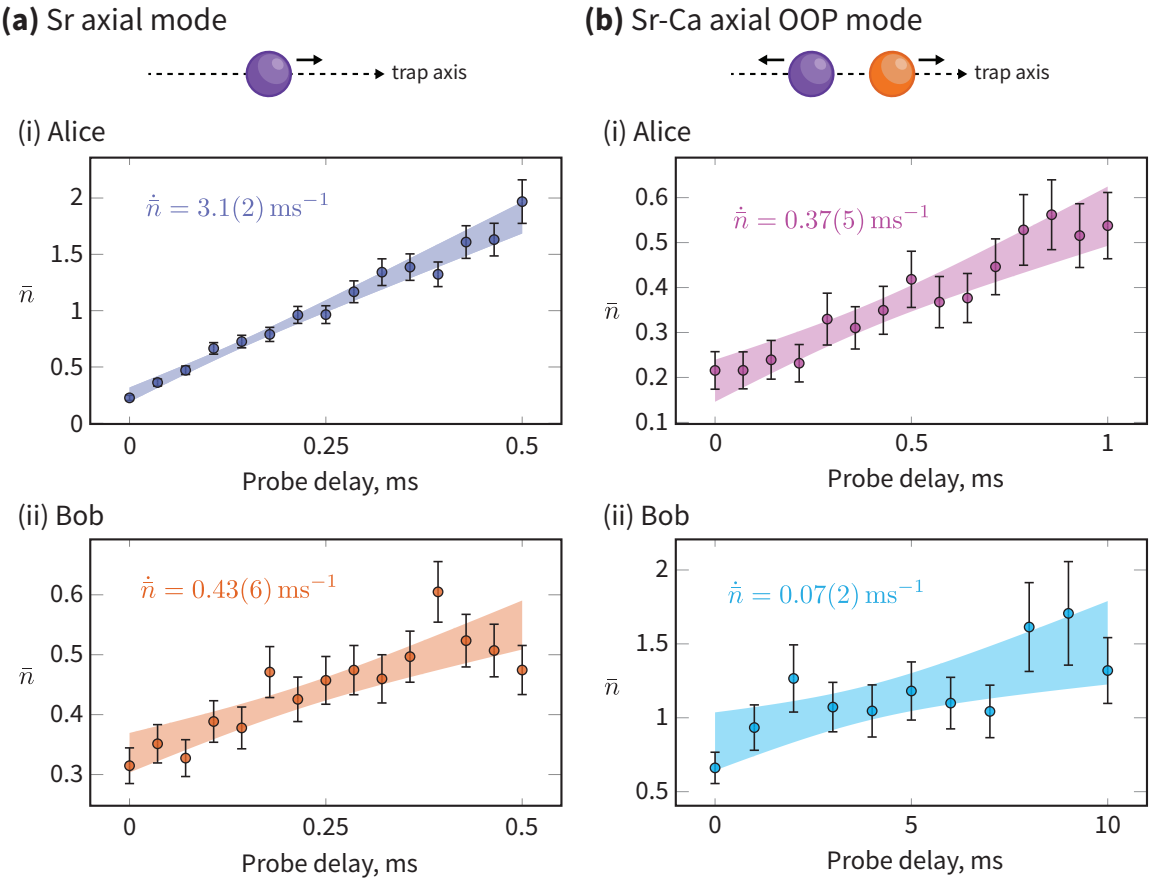

Figure 4.4: Heating rate measurements for **(a)** the single-Sr^+ axial motional mode, **(b)** and the Sr^+-Ca^+ axial OOP motional mode. For the two modes, we measure heating rates in (i) Alice and (ii) Bob.

We probe the motional heating in each module, measuring heating rates for the single-Sr^+ axial mode and the Sr^+-Ca^+ axial OOP mode. The heating rates are measured by cooling the mode, waiting for some delay, before Rabi flopping on the red and blue sidebands of the mode to estimate the mean excitation number, $\bar{n}$ [108]. The results from the heating rate measurements in each module are shown in Figure 4.4. For the single-Sr^+ axial mode, we observe heating rates of $3.1(2)\,\mathrm{ms}^{-1}$ and $0.43(6)\,\mathrm{ms}^{-1}$ for Alice and Bob, respectively. Bob exhibits a significantly lower heating rate than Alice, highlighting the improved performance of the Phoenix trap over the HOA2 trap[3]. The axial OOP mode of the Sr^+-Ca^+ crystal is of particular interest since this mode is used to mediate the mixed-species entangling gate. We

[3]While this could also be explained by some large amount of noise at the Sr^+ axial frequency in Alice, we have observed heating rates of this magnitude in Alice over a range of axial frequencies.

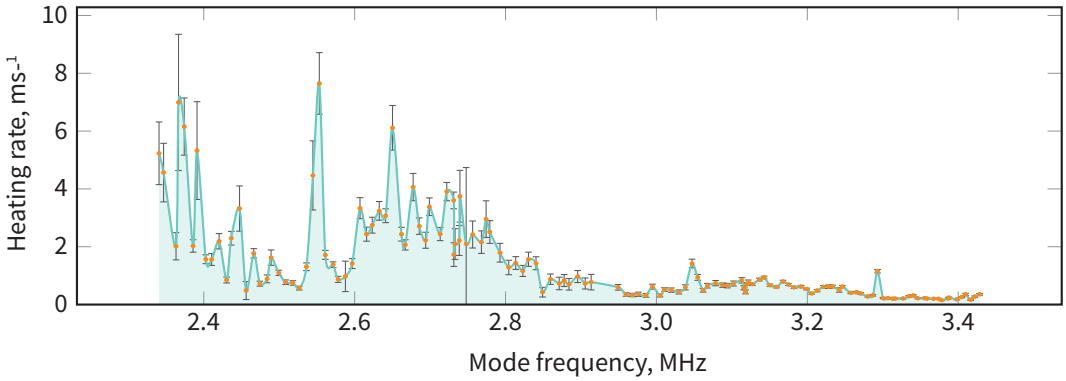

Figure 4.5: Heating rate spectrum of the single Sr^+ radial modes in the Phoenix trap. The spectrum was obtained by varying the radial confinement to modify the radial frequencies, thereby allowing the heating rate to be probed at different frequencies.

observe heating rates of $0.37(5)\,\mathrm{ms}^{-1}$ and $0.07(2)\,\mathrm{ms}^{-1}$ for Alice and Bob, respectively. Once again, we observe a significantly better performance of the Phoenix trap compared to the HOA2 trap. The Sr^+-Ca^+ axial OOP mode exhibits a lower heating rate than the single-Sr^+ axial mode since the approximately asymmetric nature of the OOP mode[4] provides some protection from common mode electric field noise.

Another interesting test of the Phoenix trap in Bob was to map out the radial heating rate spectrum for a single-Sr^+ ion. We implemented this by applying DC voltages to the electrodes to modify the radial potential, and thus the radial mode frequencies. By measuring the heating rate over a range of mode frequencies, we constructed the heating rate spectrum of the radial modes, shown in Figure 4.5. Note that the measurements were made with both radial modes rotated at $\sim 45°$ to the trap plane. The observed spectrum clearly does not follow a power law, which might be expected from, e.g., anomalous heating [109]. In particular, we observe a high density of narrow features for frequencies below 3 MHz. This likely indicates high-frequency coherence signals being coupled into the trap, possibly as a result of ground-loops[5]. The presence of such features indicates that with a lot of attention to technical noise sources, we could likely eliminate a large amount of the

[4]The ions move exactly out of phase, however the motional amplitudes are different due to the mass disparity.

[5]Of which we likely have many in our lab.

electrical field noise at the ions. It is worth pointing out that the frequency range shown in Figure 4.5 is higher than the mode frequencies we operate with in Bob (see Table 4.1). This is a result of this data being taken prior to the observation that the higher RF voltages led to charging of the trap[6]. In order to further diagnose sources of noise in the system, it may be of interest to measure the radial heating rates over a larger range, in addition to repeating the experiment in Alice.

A similarly important metric for the performances of the traps is the motional coherence time. Previously, we have observed a motional coherence time for the Sr^+-Ca^+ axial OOP mode of 3.4(4) ms in Alice [56]. We anticipate the motional coherence time to be limited by technical noise being coupled into the system via the control electronics. While we have not made a recent measurement of the motional coherence time in Alice, or any measurement yet in Bob, we anticipate that both systems will exhibit motional coherence times roughly on the same order of magnitude as each other.

4.2.2 Lasers

We use lasers for photo-ionisation during loading, cooling, SPAM, and manipulation. Since we use two different species of ion, we require two sets of beams per module. The laser light is delivered to the trap table using optical fibres, before free-space optics direct the beams through view-ports located around the vacuum chamber and onto the ion. All beams are delivered parallel to the trap surface. Figure 4.3 displays the different beam paths by wavelength, showing the geometry by which the light is incident on the ions relative to the trap axis and the magnetic field. Acousto-optical modulators (AOMs) are used for fast switching and precise control of the frequency of all beams, with the exception of the photo-ionisation lasers, which are controlled using mechanical shutters.

The laser systems themselves are common to both modules, and are split between the two modules, and other trapped-ion experiments, prior to the AOM network.

[6]Since this data was taken, we have increased the RF frequency and reduced the RF voltage.

The majority of these lasers are standard Toptica diode lasers, however there are a number of more complex laser systems. Firstly, for the remote entanglement scheme discussed in Chapter 6, we use a mode-locked Ti:Sapphire laser with a custom pulse-picking module to produce ~ 5 ps pulses of 422 nm light, at a repetition rate of up to 1 MHz [92, 110]. Secondly, the 674 nm light used for coherently addressing the optical quadrupole transitions in Sr^+ is produced by a SolsTis Ti:Sapphire laser, manufactured by M Squared Lasers Ltd, and locked to a high-finesse cavity system. Finally, a Raman laser system manufactured by Toptica is used for the coherent manipulation of the Ca^+ hyperfine qubits and the generation of the SDF used to mediate entangling gates between the Sr^+ and Ca^+ ions. This laser system is discussed in more detail in Section 4.3.

4.2.3 Imaging systems

Each module has two separate imaging systems for collecting light from the ions, as depicted in Figure 4.3(b). A rear-side imaging system collects fluorescence light from the ions through the slot in the trap chip. This imaging system is designed to collect light at 422 nm and 397 nm, corresponding to the dipole transitions in Sr^+ and Ca^+, respectively. A dichroic mirror is used to separate the two wavelengths, and direct the fluorescence onto photomultiplier tubes (PMTs), thereby enabling the simultaneous detection of single photons from each ion species. The apparatus was also designed to allow the detection via an Andor EMCCD camera; however this was found to produce a significant amount of magnetic field noise at around 110 Hz. The outputs from the PMTs are sent to the control system, which registers the detection event and time-stamps the event with 1 ns precision. Typical values for experimentally measured collection efficiencies, i.e., the probability that a photon emitted from an ion is detected by the corresponding PMT, for Sr^+ (Ca^+) rear-imaging system are 0.34(1) % (0.30(1) %) and 0.27(1) % (0.45(2) %) for Alice and Bob, respectively.

The second imaging system is central to the generation of ion-photon entangle-

ment and, consequently, the generation of remote entanglement between the two modules, as discussed in Chapter 6. A 0.6 NA lens[7] collects spontaneously emitted 422 nm photons from a Sr^+ ion located at the focus of the lens, and couples the photons into a single-mode optical fibre. This imaging system is carefully aligned to minimise optical aberrations, which is essential for avoiding polarisation mixing in the generation of ion-photon entanglement.

4.2.4 Control system

Each module is assigned a "host" computer, which is used to submit and schedule experiments, communicate with the experimental control system, and collect, analyse, and save experimental data. Real-time execution of experimental sequences are handled using the ARTIQ control system [111]. The majority of the hardware comprising this control system is from the Sinara ecosystem, which is well-supported by ARTIQ. Central to this control system is the Kasli – a FPGA that executes programs, also known as kernels, in real-time with nanosecond precision. In addition to the Kasli, the Sinara family includes a number of extension modules that can be controlled by the Kasli. These modules are designed to perform various functions: transistor-transistor logic (TTL) input/output, direct digital synthesis (DDS), digital-to-analogue conversion (DAC), and analogue-to-digital conversion (ADC). These modules therefore provide the interface between the programs executed on the Kasli and the physical experiment. As a simple example, qubit readout, which we discuss in Section 4.4.4, is performed by switching on DDS channels connected to AOMs which control the relevant fluorescence lasers. The output of the detection PMT is then connected to a TTL input, which counts the number of fluorescence photons detected within a certain window. The Kasli can then determine whether this number of counts corresponds to a bright or dark state, and use this information accordingly.

While the general functionality of the Sinara/ARTIQ control system is usu-

[7]Note that this lens is located outside of the vacuum chamber.

ally sufficient for our purposes, the tight timing requirements and high-density of real-time input/output (RTIO) events involved in our ion-photon and ion-ion entanglement generation schemes mean that additional extensions to the Kasli gateware were developed. These extensions enable the execution of our entanglement generation sequences and real-time decision branching at a $\sim 1\,\mathrm{MHz}$ attempt rate. We will not go into detail here, however interested readers are referred to the DPhil thesis of Drmota [56].

4.3 Raman lasers

We use a pair of 402 nm Raman lasers to coherently address transitions within the ground hyperfine manifold of Ca^+. The Raman lasers are detuned by $\Delta \sim 10\,\mathrm{THz}$ from the 397 nm dipole transition, and have a relative detuning of $\sim 3.2\,\mathrm{GHz}$, which is fine-tuned with AOMs to address individual hyperfine transitions. Recall from Section 2.4.2 that in order to coherently implement single qubit rotations between qubit states connected by a particular transition, the driving "laser" must remain coherent with the qubit over the course of the experimental sequence. For the hyperfine qubits driven by the pair of Raman lasers, the "laser" in this context refers to the $\sim 3.2\,\mathrm{GHz}$ beat-note from the interference of the Raman lasers. Thus, to ensure that the hyperfine qubits remain coherent over an experimental sequence, we require that the *relative* phase between the two Raman lasers remain stable. The experimental sequences presented in Chapters 6 and 7 have durations on the order of 10 ms to 1 s; we therefore require that the total integrated phase noise down to frequencies $\mathcal{O}(\mathrm{Hz})$ is significantly less than 1 rad.

4.3.1 Laser system

Typically, a pair of Raman beams is produced by phase-locking two separate lasers. However, typical bandwidths for the phase-locking electronics are $\sim 1\,\mathrm{MHz}$, and thus fluctuations in the relative phase at higher frequencies than this cannot be

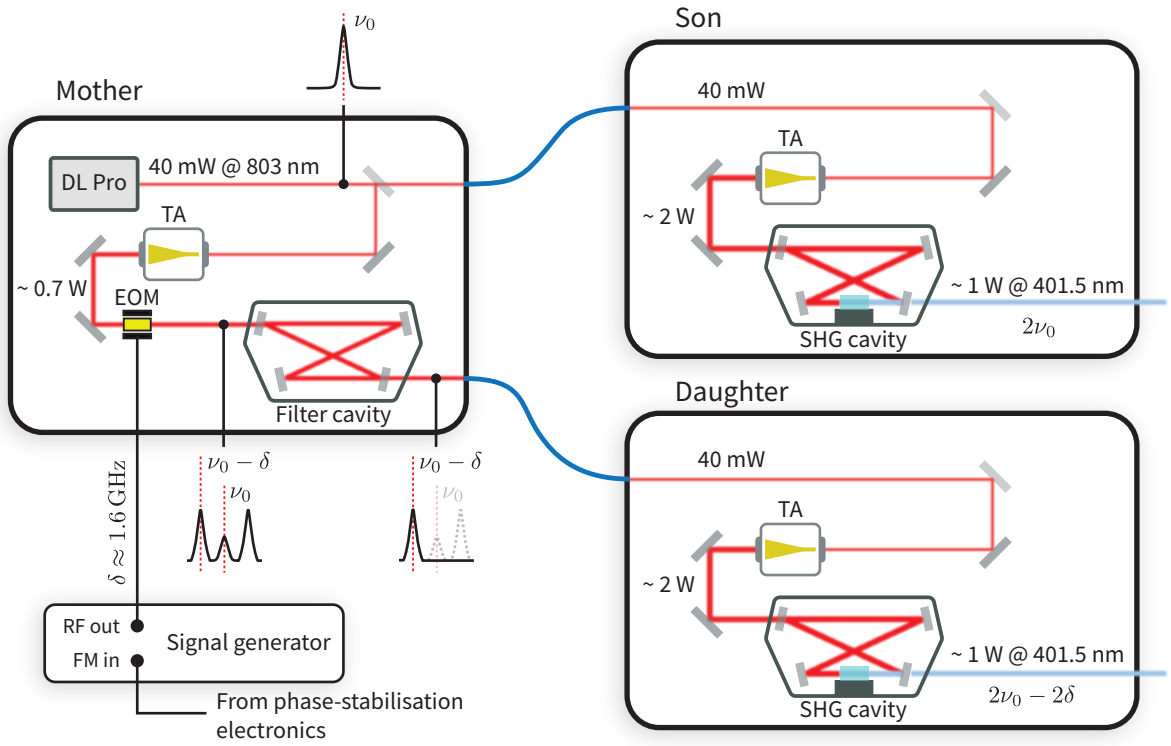

Figure 4.6: Schematic for Toptica Raman laser system. The diode in the Mother laser produces light at 803 nm which is split into two beams. One beam is immediately fibre coupled and sent to seed the Son laser system. The other beam is amplified using a tapered amplifier (TA). This light is then modulated using an electro-optic modulator (EOM) driven at $\delta \approx 1.6$ GHz to produce sidebands. The signal generator driving the EOM has an input for a frequency modulation signal, which we use to modulate the relative phase of the lasers to perform phase-stabilisation. A filter cavity is then used to filter out the carrier and one of the sidebands, leaving only one of the sidebands. This light is hence shifted by δ with respect to the light sent to the Son. The Son and Daughter laser systems are nominally identical. The seed light is amplified using a TA and then frequency doubled using a second harmonic generation (SHG) crystal, located within a second cavity. The result is two beams, each with ~ 1 W power and with a frequency difference of $2\delta \approx 3.2$ GHz, with an intrinsically stable relative phase.

suppressed. We opt for a fundamentally different approach; a common laser is used to optically seed the two laser systems, resulting in two beams that have an intrinsically low relative phase noise at all frequencies without any active locking.

Figure 4.6 shows the setup for our Raman laser system. We refer to the seeding laser as the *Mother* and the pair of seeded Raman lasers as the *Son* and *Daughter*. The diode laser in the Mother emits 803 nm light which is amplified using a tapered amplifier (TA) and split into two beam paths which seed the Son and Daughter systems. One of these beams is immediately fibre-coupled and sent to the Son laser

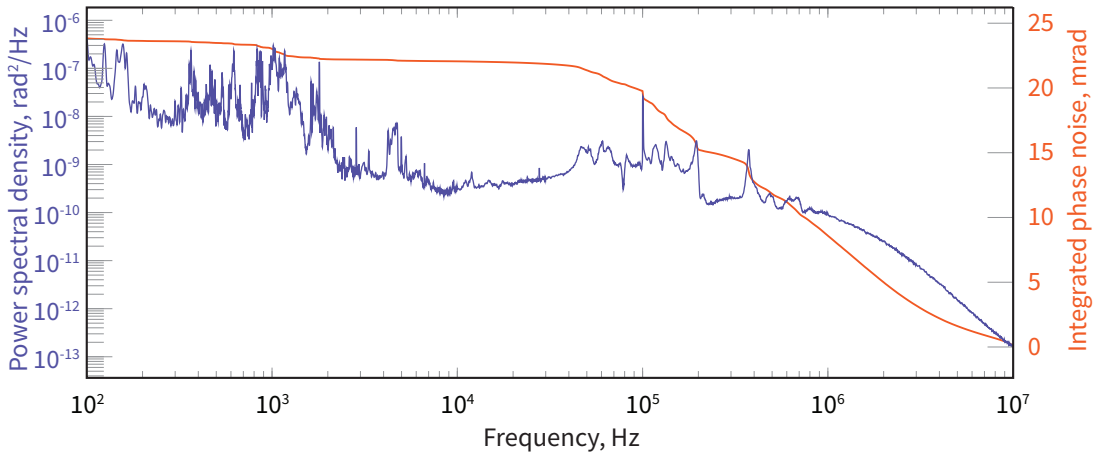

Figure 4.7: Power spectral density (blue) and reverse-integrated phase noise (red) for the relative phase noise between the two laser systems. The total integrated phase noise over 100 Hz to 10 MHz is 24 mrad.

system. The beam for the Daughter is modulated using an electro-optic modulator (EOM)[8] to produce sidebands at 1.6 GHz. A filter cavity is used to lock to one of the sidebands and filter out the other sideband and the carrier, resulting in light that is shifted in frequency by $\sim$1.6 GHz, with respect to the light sent to the Son. This light is fibre-coupled and sent to the Daughter laser system.

The Son and Daughter laser systems are identical. The light coming from the mother is used to seed a TA, which amplifies the light from $\sim$40 mW to $\sim$2 W. The amplified light is then coupled into a cavity containing a second harmonic generation (SHG) crystal which doubles the frequency of the light to 401.5 nm. After the frequency doubling the Son and Daughter beams are then separated by $\approx$ 3.2 GHz, and retain an intrinsically stable relative phase.

To characterise the relative phase noise of the Raman lasers, we superpose the two blue laser beams and direct them onto a photodiode to observe the beat note at 3.2 GHz. The DC-offset is removed using a bias-tee, leaving only the oscillating beat note signal. This signal is amplified and then analysed using a phase noise analyser[9], which measures the power spectral density of the phase noise of the beat note. The power spectral density and the reverse-integrated phase noise between

[8]The modulation signal driving the EOM is produced by an Rohde & Schwarz SMA 100 A signal generator and amplified using a Qubig QDG8 RF driver.

[9]Rhode and Schwarz FSWP phase noise analyzer.

the two Raman lasers is shown in Figure 4.7. We observe a total integrated phase noise of 24 mrad, over 100 Hz to 10 MHz. For comparison, a different Raman laser system in our group exhibited a total integrated phase noise of $\approx$ 9 mrad after phase-stabilisation, over 1 kHz to 1 MHz [75].

4.3.2 Phase noise stabilisation

The light from each Raman laser is delivered separately to the trap table via ≈ 7 m of polarisation-maintaining photonic crystal fibres (PCFs), where the beams are directed to an AOM network, as shown in Figure 4.8(a), and then on to the traps. Since the light from the Son and Daughter systems is delivered through separate fibres, the relative phase between the two beams is highly susceptible to fluctuations caused by differential path-length changes in the fibres. These changes, typically induced by thermal and acoustic stress, result in significant relative phase noise in the acoustic regime (<20 kHz). To suppress this noise, we implement a phase stabilisation circuit, also depicted in Figure 4.8(a).

After the fibres deliver the light to the trap table, a pick-off is used to sample light from each beam; the remaining light is sent to the AOM network. The sampled light from each beam is superimposed and the resulting interference signal measured on a photodiode, which produces a voltage. This voltage comprises a DC offset, v_0, and a signal corresponding to the beat note signal,

$$v(t) \propto \sin\left(2\delta t + \phi(t)\right) \tag{4.1}$$

where $2\delta \approx 3.2$ GHz and $\phi(t)$ corresponds to the time-dependent relative phase fluctuations. The DC component is removed using a bias-tee with a 50 Ω terminator at the DC output. The beat note signal, at this point ~ -70 dBm, is sent through a band pass filter centered at 3.2 GHz, and amplified using a series of three amplifiers, yielding a total amplification of ~ 60 dB. This signal is then passed to a frequency mixer. A signal generator, phase locked to the signal generator which drives the

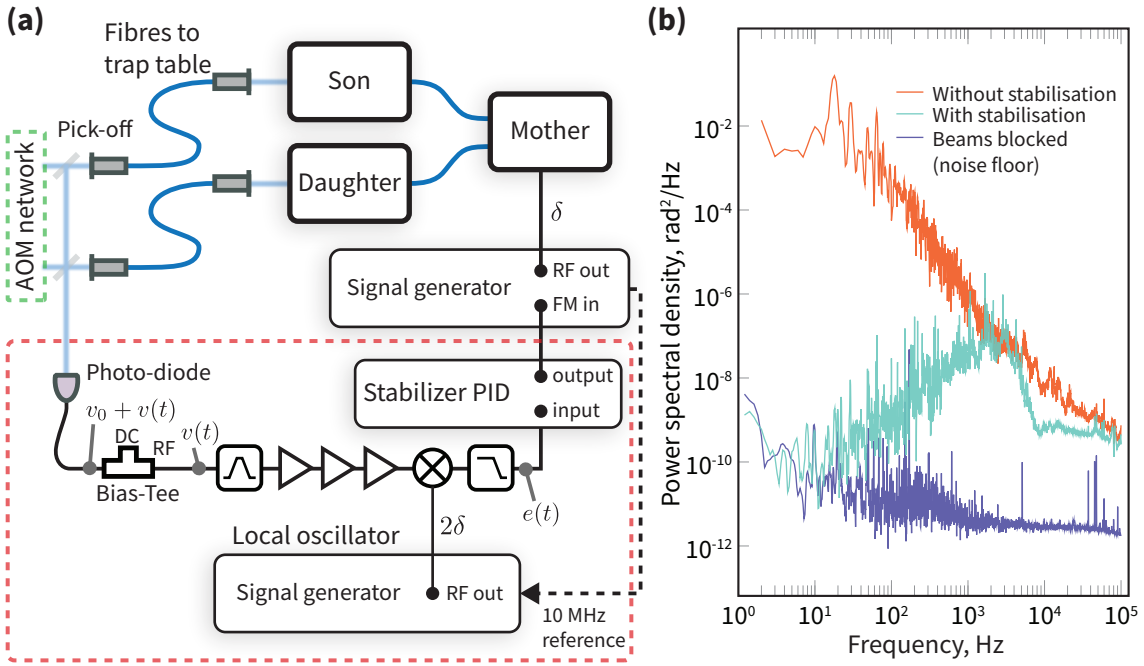

Figure 4.8: Raman phase noise stabilisation system. **(a)** The light from each Raman laser is delivered separately to the trap table. On the trap table, a small fraction of each beam is sampled at a pick-off, the remaining light is sent to the AOM network (see Figure 4.9). The two sampled beams are superimposed and interfere on a photodiode, which produces a voltage comprising a DC offset and an oscillating voltage proportional to the amplitude of the beat note between the two Raman beams. A bias-tee removes the DC component of the signal, and a band-pass filter centered at 3.2 GHz is used to isolate the beat note signal from noise. A series of amplifiers amplifies the beat note signal, before the signal is sent to a mixer, where a local oscillator at $2\delta \approx 3.2\,\text{GHz}$ mixes the signal down to DC. The local oscillator is phase-locked to the signal generator used to drive the EOM in the Mother laser. The down-mixed signal passes through a low-pass filter which is used as the error signal for a PID controller. The output of the PID controller is used to frequency modulate the EOM signal, such that the relative phase of the Raman lasers at the trap table is stable. **(b)** In-loop measurements of the power spectral density of the phase noise with (turquoise) and without (orange) phase noise stabilisation. The power spectral density of the phase noise with both beams blocked (blue) indicates the noise floor for the feedback circuit.

EOM in the Mother, provides a local oscillator signal at 2δ. The output signal from the mixer therefore contains terms at the sum frequency (4δ), and at DC. A low-pass filter is used to remove all but the DC component, such that the signal we are left with is given by

$$e(t) = e_0 \sin\left(\phi(t)\right). \tag{4.2}$$

For small phase fluctuations, we have

$$e(t) \approx e_0 \phi(t), \tag{4.3}$$

which corresponds to a voltage signal proportional to the phase fluctuations.

The voltage signal, $e(t)$, is used as an error signal for a PID controller – for this we deploy a Sinara stabilizer. The output of the PID controller is then used to frequency modulate the signal generator driving the EOM in the Mother, thereby allowing us to adjust the relative phase of the Raman lasers and thus stabilise this phase at the trap table. We characterise the performance of the phase stabilisation circuit by using a voltage pick-off to observe the error signal $e(t)$. The measurements with and without the phase stabilisation circuit are shown in Figure 4.8(b). Without phase stabilisation, we see a significant increase in the relative phase noise for frequencies $< 1\,\mathrm{kHz}$. This is likely due to slow thermal and acoustic noise inducing differential path length fluctuations in the two fibres, as discussed above. When we activate the phase stabilisation, we see a clear reduction in relative phase noise over this range.

4.3.3 Geometry

The Raman laser light is divided between the two modules, such that each module receives two beams from the Son laser, labelled $\mathrm{Son}_{\parallel}$ and $\mathrm{Son}_{\perp}$, and two beams from the Daughter laser, labelled $\mathrm{Daughter}_{\parallel}$ and $\mathrm{Daughter}_{\perp}$, where the subscripts, $\parallel$ and $\perp$, indicate the alignment of the beams relative to the applied magnetic field. The AOM switching network, shown in Figure 4.9(a), enables individual switching of these beams, as well as fine-tuning of their frequency and phase. The polarisation

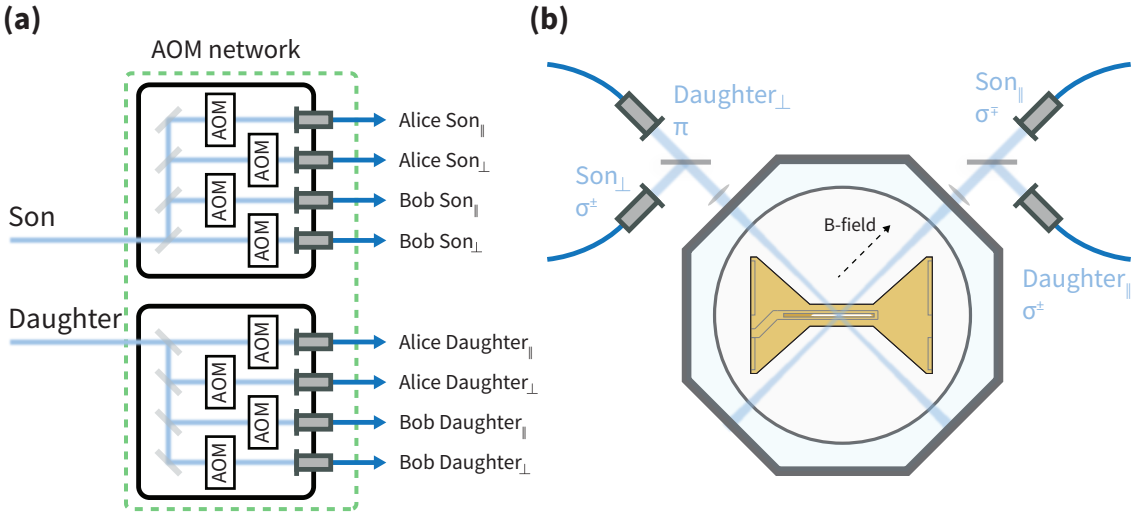

Figure 4.9: **(a)** AOM network for the Raman beams. For each ion trap module, we have a total of four Raman beams, two from each Raman laser. We label these beams $\text{Son}_{\parallel}$, $\text{Son}_{\perp}$, $\text{Daughter}_{\parallel}$, and $\text{Daughter}_{\perp}$, where the subscript denotes the alignment with respect to the magnetic field. Each beam has its own AOM for individual switching as well as fine tuning of the frequency and phase. **(b)** Geometry of the beam delivery for the four Raman beams to one of the modules. Each beam is linearly polarised, with the polarisation indicated next to the respective beam path.

and geometry of the beam delivery is depicted in Figure 4.9(b). This setup provides the flexibility to select different pairs of beams for various applications. For example, for applications that involve carrier interactions, such as single-qubit rotations (Section 5.1) and interconverting between hyperfine qubits (Section 5.3), we can minimise the coupling to the ions' motion by selecting a pair of co-propagating Raman beams to minimize the relative wavevector. Conversely, for applications where coupling to the ions' (axial) motion is desired, such as driving sideband interactions and implementing entangling gates (Section 5.2), we can select pairs of Raman beams from different view-ports, thereby achieving a large relative wavevector.

4.4 Incoherent operations

So far, we have discussed the hardware comprising the modules. Here, we outline the implementation of a number of important incoherent operations that enable the manipulation of the ions. These operations include the loading of and laser cooling

of ions, as well as the preparation and measurement of qubits.

4.4.1 Loading

Inside each vacuum chamber, we have one atomic oven containing calcium and another containing strontium. The ovens are resistively heated to around 300 °C, producing a flux of neutral atoms across the trapping zone [112]. The neutral atoms are ionised via an isotope-selective two-step photo-ionisation process [113]. This is implemented using a pair of photo-ionisation lasers. For ionisation of calcium (strontium), the first laser is tuned to the $4s^2\ {}^1S_0 \leftrightarrow 4s4p\ {}^1P_1$ ($5s^2\ {}^1S_0 \leftrightarrow 5s5p\ {}^1P_1$) transition at 423 nm (461 nm). A second laser at 378 nm provides enough energy to ionise either the neutral calcium or strontium. The isotope shifts for the two species are significantly larger than, e.g., the natural of the ionisation transitions or the residual Doppler broadening, and thus this laser can be tuned select the desired isotope to ionise. Once the atom has been ionised, it will see the confining potential and will be trapped.

4.4.2 Laser cooling

Ions confined to a trapping potential experience heating and occasional background gas collisions. These processes cause the ions to gain kinetic energy, which, if left unchecked, will eventually be sufficient for the ions to escape the trapping potential. It is therefore necessary to remove this kinetic energy, i.e., to *cool* the ions. Additionally, as discussed in Section 2.4.3, to coherently manipulate the motional state of the ions within the Lamb-Dicke regime, we must ensure that the motional state remains close to its motional ground state, $|0\rangle$. Therefore, for QIP applications, the ability to cool the motional modes to their ground state is critical.

Due to the conservation of momentum, the absorption and emission of photons necessarily induce a change in the momentum of an ion. This has enabled the development of a variety of laser cooling techniques, which are critical to many trapped-ion and neutral atom experiments. In this work, we use two types of laser cooling:

Doppler cooling and electromagnetically induced transparency (EIT) cooling.

Doppler cooling

The most common laser cooling technique is Doppler cooling [114, 115], see Figure 4.10(a). This technique typically involves red-detuning a laser from a dipole transition. Due to the conservation of momentum, absorption of a photon from this laser will result in a momentum change of $\hbar k$ along the direction of the laser propagation. The ion will then emit a photon; however, this emission will be isotropic, such that the average change in momentum from the emission will be zero. Now consider an ion confined to a potential well. Suppose the motion of the ion has a component along the direction of the incoming laser. Due to the harmonic motion of the ion, this component alternates between moving towards and away from the laser. If the ion is moving towards the incoming laser, then in the rest frame of the ion, the frequency of the laser will be Doppler-shifted towards the dipole transition, thus making absorption more likely. Conversely, if the ion is moving away from the incoming laser, the laser will be Doppler-shifted away from the dipole transition, making absorption less likely. This results in the ion being more likely to receive a momentum kick in the opposite direction to its motion, leading to a damping of its motion. As the amplitude of the ion's motion is being reduced, this process lowers the temperature of the ion. Any motional mode with a projection along the Doppler cooling laser can then be cooled in this manner.

The minimum temperature that can be reached using Doppler cooling is limited by the random walk resulting from the momentum kicks from the isotropic spontaneous emission, and from photon absorption. Therefore, we cannot use Doppler cooling to attain the ground state of the ion's motion. Given a motional mode with a secular frequency ω, and a cooling transition with a linewidth Γ, the smallest mean excitation number that can be achieved using Doppler cooling is $\bar{n} \sim \Gamma/2\omega$. For Sr^+, we cool using the 422 nm dipole transition, while using a 1092 nm beam to repump any population that decays to the $D_{3/2}$ manifold, as shown in Figure 4.10(a)(i).

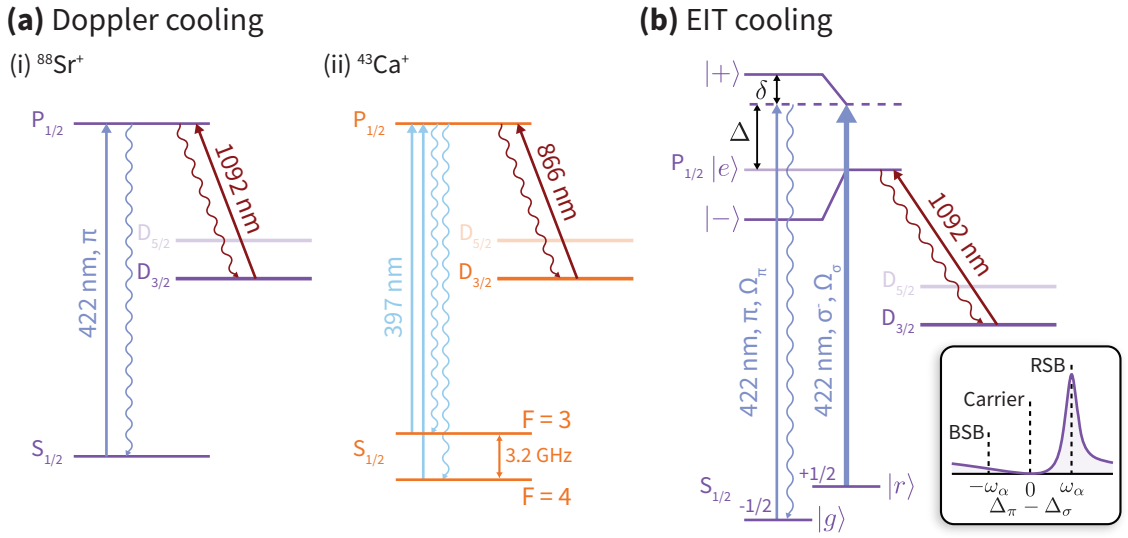

Figure 4.10: Energy level diagrams for **(a)** Doppler cooling of (i) Sr^+ and (ii) Ca^+, and **(b)** EIT cooling of Sr^+. **(a)** Doppler cooling involves red-detuning a laser from the $S_{1/2} \leftrightarrow P_{1/2}$ dipole transition, along with an infrared repumping beam to depopulate the $D_{3/2}$ level. For (i) Sr^+ [(ii) Ca^+], a 422 nm [397 nm] dipole laser and a 1092 nm [866 nm] repumper are used. The 397 nm laser is modulated to produce a sideband at the ground hyperfine splitting, thereby enabling us to address both hyperfine manifolds, $F = 3, 4$. **(b)** In EIT cooling, an intense coupling field, Ω_σ, is blue-detuned from the 422 nm dipole transition by Δ, creating a pair of dressed states, $|\pm\rangle$. A weaker cooling laser with $\Omega_\pi \ll \Omega_\sigma$ is also applied. The coupling field modifies the ion's absorption profile, resulting in a Fano profile, as shown in the inset. By setting the detunings of both lasers to Δ, the EIT condition is achieved, suppressing resonant absorption. When the light shift is set to $\delta = \omega_\alpha$, where ω_α is the relevant motional mode frequency, the red sideband overlaps with the peak of the Fano profile, enhancing red-sideband absorption while suppressing resonant and blue-sideband absorption.

The 422 nm transition has a linewidth of ≈ 20 MHz, and thus the Doppler limit for a 1.5 MHz motional mode is $\bar{n} \sim 6$. For Ca^+, we cool using the 397 nm dipole transition, together with a 866 nm beam to repump any population that decays to the $D_{3/2}$ manifold, as shown in Figure 4.10(a)(ii). Since the hyperfine splitting in the ground $S_{1/2}$ manifold, ~ 3.2 GHz, is significantly larger than the linewidth of the 397 nm laser, we use an EOM to modulate the 397 nm laser light to produce a sideband at the hyperfine splitting, thereby enabling coupling to both ground hyperfine levels. The 397 nm transition in Ca^+ has a similar linewidth to Sr^+, thus we expect a similar Doppler cooling limit.

The broad linewidth of the dipole cooling transition allows us to scatter photons

at a high rate, and thus achieve a high cooling rate. However, this linewidth is significantly larger than the typical motional mode frequencies of the trap, and thus we get significant resonant absorption, thereby limiting the temperature achievable using Doppler cooling. To achieve cooling below this Doppler limit, one can use a transition which has a much narrower linewidth, thereby allowing the individual motional sidebands to be resolved. In this regime, one can coherently drive these sidebands to remove motional excitations – a technique known as *resolved sideband cooling*. However, the need to resolve the motional sidebands limits the cooling rates that can be achieved.

Electromagnetically-induced transparency cooling

As discussed above, the broad linewidth of the dipole cooling transition allows us to scatter photons at a high rate; however, resonant absorption of the cooling light limits the temperatures achievable with Doppler cooling. To overcome this limitation, we can exploit the electromagnetically induced transparency (EIT) effect to coherently suppress absorption on the carrier transition while enhancing absorption on the red sideband transition. This approach allows us to maintain rapid cooling rates similar to Doppler cooling but, by suppressing resonant absorption, we can surpass the Doppler cooling limit and cool the ions significantly closer to the motional ground state [116, 117].

We perform EIT cooling using the Sr^+ ions, and the energy level diagram for this technique is shown in Figure 4.10(b). Following Morigi, Eschner, and Keitel [116], we consider three relevant levels within the Sr^+ ion: the two ground Zeeman states, labelled $|g\rangle$ and $|r\rangle$, and an excited state, $|e\rangle$, in the $P_{1/2}$ manifold. We apply a "coupling laser", consisting of a σ^--polarised 422 nm beam that is blue-detuned from the $S_{1/2} \leftrightarrow P_{1/2}$ dipole transition by Δ_σ with a Rabi frequency Ω_σ. This coupling field creates a pair of dressed states, $|\pm\rangle$, with the state $|+\rangle$ shifted from Δ_σ by a light shift δ, as shown in Figure 4.10(b). Next, we apply a "cooling laser", comprising a π-polarised 422 nm beam, which is blue-detuned from the $S_{1/2} \leftrightarrow P_{1/2}$

dipole transition by Δ_π, and has a Rabi frequency $\Omega_\pi \ll \Omega_\sigma$. The coupling laser modifies the absorption profile of the ion for the cooling laser, resulting in a Fano-like absorption profile, as depicted in the inset of Figure 4.10(b). We see that by setting $\Delta_\pi = \Delta_\sigma = \Delta$, the EIT effect suppresses resonant absorption of the cooling laser.

However, the ion may still absorb light via its motional sidebands. We consider a single mode of motion with secular frequency ω_α. By tuning the Rabi frequency of the coupling field, Ω_σ, such that $\delta = \omega_\alpha$, the peak of the Fano absorption profile overlaps with the red sideband transition, $|g\rangle |n\rangle \leftrightarrow |+\rangle |n-1\rangle$, while absorption on the blue sideband is significantly reduced, and resonant carrier absorption is suppressed. As a result, we can drive the red sideband without heating from resonant absorption, enabling us to achieve much lower temperatures. Typically, we set $\Delta/2\pi \sim 150\,\text{MHz}$ and use coupling and cooling laser powers on the orders of $1.5\,\mu\text{W}$ and $0.5\,\mu\text{W}$, respectively. As discussed in the next section, we make extensive use of EIT cooling for ground-state cooling of mixed-species ion crystals.

Ground-state cooling of mixed-species crystals

For the work in this thesis, the modules each co-trap one Sr^+ ion and one Ca^+ ion, and local entangling gates are mediated using the axial OOP mode of motion. For high-fidelity entangling gates, it is crucial that we are able to efficiently cool this mode to close to its motional ground state. Additionally, as we discuss in Section 5.2, we also need to be able to cool the axial IP mode, as this has a significant impact on the gate fidelity. Finally, we would like to be able to sympathetically cool the Sr^+-Ca^+ crystal using only the Sr^+ ion, thus enabling this cooling to be performed mid-circuit while the Ca^+ ion stores quantum information in the long-lived circuit qubit.

To satisfy these requirements, we perform a cooling sequence, using only the Sr^+ ion, that consists of $\sim 1\,\text{ms}$ of Doppler cooling, followed by short pulses ($\sim 10\,\mu\text{s}$) of EIT cooling which alternatingly target the two axial modes. Using this cooling

sequence, we typically achieve mean occupation numbers for the axial OOP (IP) mode of $\bar{n} = 0.15$ ($\bar{n} = 2.0$) and $\bar{n} = 0.1$ ($\bar{n} = 0.8$) for Alice and Bob, respectively.

4.4.3 State preparation

The ability to prepare qubits in a particular state is critical for any quantum information processing platform. Here, we outline the processes by which we prepare the computational $|0\rangle$ state for the network, auxiliary, and circuit qubits. For each state preparation process, we estimate a state preparation error, ϵ, which corresponds to the probability that the state preparation process leaves the ion in any state other than the intended $|0\rangle$ state. Making the (relatively crude) assumption that the population remains in the qubit subspace of the ion, we model imperfect state preparation as the state

$$\hat{\rho} = (1 - \epsilon) |0\rangle \langle 0| + \epsilon |1\rangle \langle 1| . \tag{4.4}$$

This model is then used when performing process tomography, as outlined in Section 3.3.2. Note that we do not directly measure the state preparation error. Instead, we measure the combined SPAM error and estimate the contributions from imperfect state preparation and imperfect measurement.

Strontium

We prepare the computational $|0_{\mathrm{N}}\rangle$ state of the network qubit using a polarisation-selective optical pumping scheme, shown in Figure 4.11(a). A σ^--polarised 422 nm beam pumps population from the $\left|\mathrm{S}_{1/2}, m_J = +\frac{1}{2}\right\rangle$ state to the $\left|\mathrm{P}_{1/2}, m_J = -\frac{1}{2}\right\rangle$. This excited state then either decays to one of the ground Zeeman states, or to the $\mathrm{D}_{3/2}$ manifold which is depopulated using a 1092 nm repumper. Since angular momentum selection rules prevent the σ^- polarised 422 nm beam from coupling to the $|0_{\mathrm{N}}\rangle = \left|\mathrm{S}_{1/2}, m_J = -\frac{1}{2}\right\rangle$ state, the population is trapped in this state. Due to the simple energy level structure of Sr^+, this state preparation sequence can be

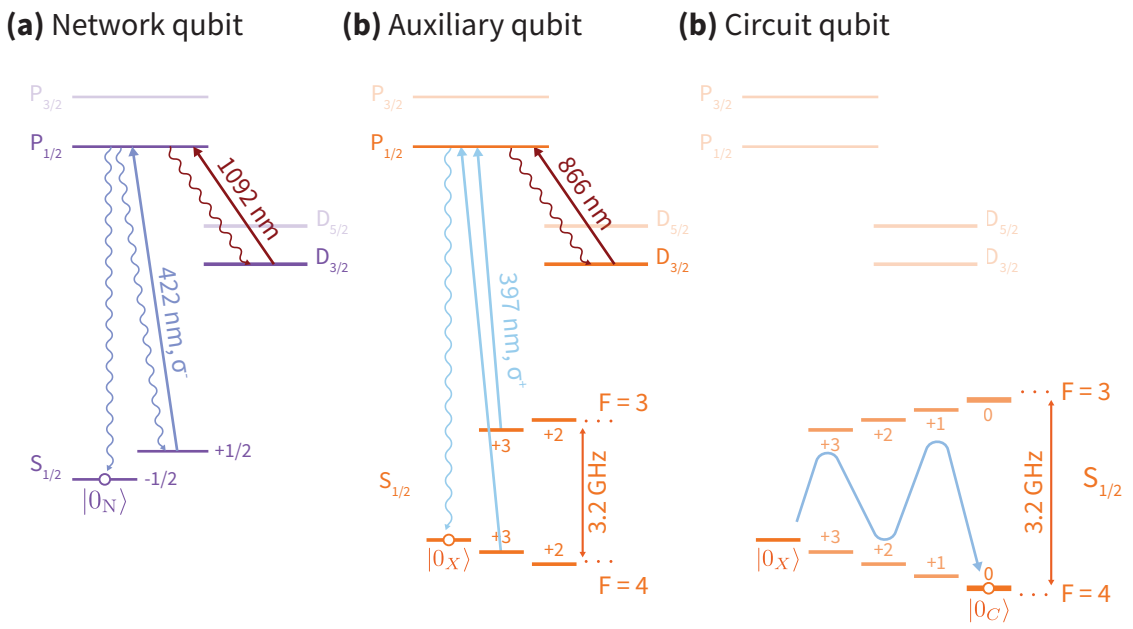

Figure 4.11: Energy level diagrams for state preparation of the **(a)** network qubit in Sr^+, and the **(b)** auxiliary and **(c)** circuit qubits in Ca^+. **(a)** A σ^--polarised 422 nm laser, together with a 1092 nm repumper, pumps population from the $\left|S_{1/2}, m_J = +\frac{1}{2}\right\rangle$ state into the $|0_N\rangle = \left|S_{1/2}, m_J = -\frac{1}{2}\right\rangle$ state. **(b)** A σ^+-polarised 397 nm laser, in addition to a 866 nm repumper, pumps population in the ground hyperfine manifold into the $|0_X\rangle = \left|S_{1/2}, F = 4, m_F = +4\right\rangle$ state. As in Doppler cooling, the 397 nm laser is modulated to produce sidebands at the hyperfine splitting, enabling the addressing of both hyperfine levels, $F = 3, 4$, with a single laser. **(c)** State preparation of the circuit qubit is performed by first preparing the auxiliary qubit in the state $|0_X\rangle$, before using a sequence of Raman π-pulses to coherently transfer the population to the $|0_C\rangle$ state.

performed much faster than for ions with a more complex level structure such as, e.g., $^{43}Ca^+$. This is particularly important for achieving high remote entanglement attempt rates; in our entanglement generation scheme discussed in Chapter 6, we optically pump to this state with $\sim 1\,\%$ error in ~ 300 ns.

For the network qubit in Alice (Bob), we estimate a state preparation error of $\epsilon_N = 4.7(5) \times 10^{-3}$ ($\epsilon_N = 5.0(5) \times 10^{-3}$). We attribute this error to imperfect polarisation of the (nominally) σ^- polarised 422 nm beam.

Calcium

Preparation of the auxiliary qubit to the state $|0_X\rangle$ via a polarisation-selective optical pumping scheme is slightly more complicated than for the network qubit due to

the hyperfine structure of the Ca^+ ion, as shown in Figure 4.11(b). A σ^+-polarised 397 nm beam pumps population towards the ground hyperfine state with the largest m_F quantum number via the $P_{1/2}$ manifold; here this state is one of the computational states of the auxiliary qubit, $|0_X\rangle = \left|S_{1/2}, F = 4, m_F = 4\right\rangle$. As with Doppler cooling, the EOM modulating the 397 nm laser enables the addressing of both hyperfine sub-levels. Again, angular momentum selection rules prevent the σ^+ polarised 397 nm beam from coupling to the $\left|S_{1/2}, F = 4, m_F = 4\right\rangle$ state, thus the population becomes trapped in this state. An 866 nm repumper is used to depopulate the $D_{3/2}$ manifold. Since the ground hyperfine manifold of the Ca^+ ion has a larger number of states, and since the scattering of a single σ^+-polarised 397 nm photon increases the m_F quantum number by at most one unit, state preparation in the Ca^+ ion takes $\sim 300\,\mu s$, significantly longer than for the Sr^+ ion.

Unlike the network and auxiliary qubit, we do not prepare the circuit qubit by direct optical pumping. Instead, we first prepare the Ca^+ ion in the $|0_X\rangle$ state of the auxiliary qubit and use a sequence of Raman π-pulses to coherently transfer this state to the $|0_C\rangle$ state, depicted in Figure 4.11(c).

For the auxiliary qubit in Alice (Bob), we estimate a state preparation error of $\epsilon_X = 3.1(4) \times 10^{-3}$ ($\epsilon_X = 3.8(4) \times 10^{-3}$). For the circuit qubit in Alice (Bob), we estimate a state preparation error of $\epsilon_C = 4.1(3) \times 10^{-3}$ ($\epsilon_C = 4.6(3) \times 10^{-3}$). We attribute the error in the auxiliary qubit state preparation to imperfect polarisation of the (nominally) σ^+ polarised 397 nm beam; this mechanism also contributes to the circuit qubit error, but here we have the additional error arising from imperfect π-pulses which transfer the $|0_X\rangle$ state to the $|0_C\rangle$ state.

4.4.4 Readout

The ability to extract information from a quantum computer via the measurement of qubits is of fundamental importance for universal quantum computing, both for the final output of any algorithm and for the error correction essential for the computer's successful operation. For qubits encoded in energy levels of trapped-ions,

such a measurement in the computational basis, $\{|0\rangle, |1\rangle\}$, is typically performed by shelving one of the qubit states outside of the Doppler cooling cycle, followed by illuminating the ion with cooling light. The qubit state that has been moved outside of the cooling cycle will therefore not scatter any photons, while the qubit state left within the cooling cycle will scatter the cooling light. We perform this fluorescence detection for some duration, in which time we count the number of photon detection events recorded by the relevant PMT. If the number of photons detected is found to exceed some threshold, the ion is determined to be in the bright state; otherwise, the ion is determined to be in the dark state.

As for the state preparation, we anticipate imperfections in the readout of the qubits. For each qubit, we estimate the readout errors, ϵ_0 and ϵ_1, where ϵ_0 (ϵ_1) corresponds to the probability that a measurement of a qubit prepared in the state $|0\rangle$ ($|1\rangle$) will yield the measurement outcome 1 (0), as per Chapter 3. Such experimental imperfections result from imperfect shelving of one of the qubit states, the finite lifetime of the shelved state, and the statistical uncertainty from the detection of scattered photons.

Strontium

Readout of the network qubit is straightforward, since the $|1_N\rangle$ state already lies in the $D_{5/2}$ manifold, outside of the Doppler cooling cycle, as shown in Figure 4.12(a). Therefore, we perform a measurement of the network qubit in the computational basis by performing fluorescence detection, without the need to perform any shelving. This measurement is performed by illuminating the Sr^+ ion with 422 nm and 1092 nm beams for 500 µs, and counting the number of scattered photons detected. For an ion in the bright (dark) state, typical photon detection rates are $4.850(3) \times 10^4\,s^{-1}$ ($36(2)\,s^{-1}$) and $3.943(3) \times 10^4\,s^{-1}$ ($17(1)\,s^{-1}$) for Alice and Bob, respectively. From these values, and the ≈ 391 ms lifetime of the $D_{5/2}$ manifold, we estimate network qubit readout errors, given in Table 4.2.

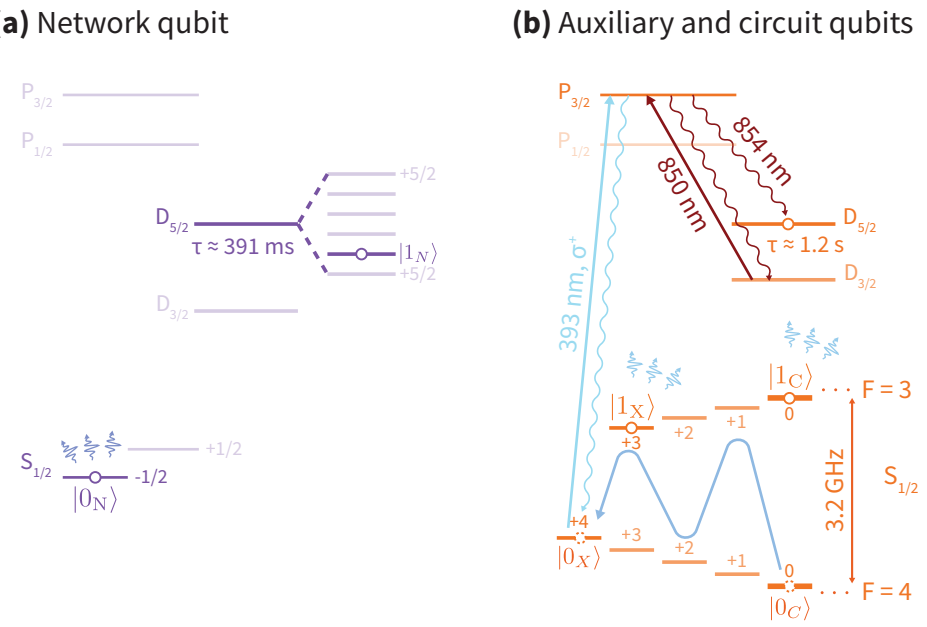

Figure 4.12: Energy level diagrams for readout of the **(a)** network qubit in Sr^+, and the **(b)** auxiliary and circuit qubits in Ca^+. **(a)** The $|1_N\rangle$ state of the network qubit is already in the $D_{5/2}$ manifold, and thus we can perform readout by simply applying the cooling lasers and counting the number of photons scattered in a given time window. The $|0_N\rangle$ state will scatter photons and thus be observed as bright, while the $|1_N\rangle$ state will appear dark. **(b)** We shelve the $|0_X\rangle$ state to the $D_{5/2}$ manifold using an optical pumping sequence, described in the main text, involving a σ^+-polarised 393 nm beam, and π- and σ^+- polarised 850 nm beams. When illuminating with the cooling lasers, the state remaining in the $F = 3$ manifold will appear bright, while the shelved state will appear dark. To shelve the $|0_C\rangle$ state, we coherently transfer this state to the $|0_X\rangle$ state, before performing the dissipative shelving sequence described above. The $S_{1/2} \leftrightarrow P_{1/2}$ and $D_{3/2} \leftrightarrow P_{1/2}$ Doppler cooling transitions, which are driven to excite fluorescence, are not shown in these diagrams.

Calcium

For the readout of the auxiliary and circuit qubits in Ca^+, all qubit states lie in the ground hyperfine manifold, which participates in the Doppler cooling cycle. Therefore, we must first shelve one of the qubit states such that it is outside of this cycle. Note that we do not shelve the $|0_C\rangle$ state directly; we first use Raman π-pulses to coherently transfer to the $|0_X\rangle$ state, before shelving the $|0_X\rangle$ state[10]. The $|0_X\rangle$ state is shelved using a polarisation-selective optical pumping scheme [119][11].

[10] Direct shelving of the $|0_C\rangle$ state is possible, although the lowest achievable error is limited to 4×10^{-3} [118]. In practice, we found that direct shelving caused a significantly higher error than this, although we did not investigate the cause of this thoroughly.

[11] The shelving of this state should be much once we integrate the 729 nm laser.

Module	Qubit	State-preparation ($\times 10^{-3}$)	Readout ($\times 10^{-3}$)		
		ϵ	ϵ_0	ϵ_1	$\frac{\epsilon_0+\epsilon_1}{2}$
Alice	$\mathcal{Q}_\mathrm{N}$	4.7(5)	< 0.001	1.068(4)	0.534(2)
	$\mathcal{Q}_\mathrm{C}$	4.1(3)	4.1(4)	2.7(4)	3.4(4)
	$\mathcal{Q}_\mathrm{X}$	3.1(4)	1.77(4)	0.357(4)	1.06(2)
Bob	$\mathcal{Q}_\mathrm{N}$	5.0(5)	< 0.001	1.085(4)	0.543(2)
	$\mathcal{Q}_\mathrm{C}$	4.6(3)	2.7(5)	1.7(5)	2.2(5)
	$\mathcal{Q}_\mathrm{X}$	3.8(4)	1.02(2)	< 0.001	0.51(1)

Table 4.2: Summary of state preparation and measurement (SPAM) errors.

The shelving of the $|0_\mathrm{X}\rangle$ state comprises repetitions of the following pulse sequence. Firstly, a $\sim 100\,\mathrm{nW}$ pulse of σ^+-polarised 393 nm pulse excites the $|0_\mathrm{X}\rangle = \left|\mathrm{S}_{1/2}, F=4, m_F=4\right\rangle$ state to $\left|\mathrm{P}_{3/2}, F=5, m_F=5\right\rangle$. This state then has three decay channels: to the desired $\mathrm{D}_{5/2}$ manifold via the 854 nm transition with a branching ratio of $\sim 6\,\%$, to the $\mathrm{D}_{3/2}$ manifold via the 850 nm transition with a branching ratio of $\sim 0.7\,\%$, or back to initial state in the $\mathrm{S}_{1/2}$ manifold. Secondly, since the $\mathrm{D}_{3/2}$ manifold is inside the Doppler cooling cycle, we must recover any population that has decayed to this manifold. The 850 nm decay may occur via emission of a σ-polarised photon, or a π-polarised photon. The σ decay channel leaves population in either the $\left|\mathrm{D}_{3/2}, F=4, m_F=4\right\rangle$ or $\left|\mathrm{D}_{3/2}, F=5, m_F=4\right\rangle$ states. We recover this population using a pulse of σ^+-polarised 850 nm light. The π decay channel leaves population in the $\left|\mathrm{D}_{3/2}, F=5, m_F=5\right\rangle$ state. This population is then recovered using a pulse of π-polarised 850 nm light. We find that after $\sim$3-4 repetitions of this pulse sequence, the population in the $|0_\mathrm{X}\rangle$ state is transferred to the $\mathrm{D}_{5/2}$ manifold, outside of the Doppler cooling cycle with sufficiently high probability for our purposes.

With the $|0\rangle$ state shelved to the $\mathrm{D}_{5/2}$ manifold, we perform fluorescence detection by illuminating the Ca^+ ion with 397 nm and 866 nm light for 1.2 ms, while counting the number of scattered photons detected. For an ion in the bright (dark) state, we typically detect photons at a rate of $1.615(1) \times 10^4\,\mathrm{s}^{-1}$ ($440(4)\,\mathrm{s}^{-1}$) and $3.060(1) \times 10^4\,\mathrm{s}^{-1}$ ($331(3)\,\mathrm{s}^{-1}$) for Alice and Bob, respectively. From these values,

and the $\approx 1.16\,\mathrm{s}$ lifetime of the $D_{5/2}$ manifold, we estimate the auxiliary qubit readout errors, given in Table 4.2.

For the circuit qubit, we also have error due to the transfer of the state $|0_C\rangle$ to the state $|0_X\rangle$, prior to the fluorescence detection, resulting in a higher readout error. We estimate circuit qubit readout errors, given in Table 4.2.

4.5 Magnetic-field stabilisation

The applied 0.496 mT magnetic field lifts the degeneracies of the Zeeman states in the ions, and provides a quantisation axis for the qubits. However, the transition frequencies for the network and auxiliary qubits have relatively large sensitivities to fluctuations of this magnetic field; the circuit qubit has a much lower, but still non-zero, first-order sensitivity. As a result, we must suppress fluctuations of the magnetic field experienced by the ions. There are three main sources of magnetic field noise. Firstly, there are ambient sources of magnetic fields, not synchronised with the AC mains cycle. We passively suppress these sources by identifying possible sources of magnetic field fluctuations from around the lab and either removing them or moving them as far from the ion traps as possible. Examples of such sources include the movement of metallic lab chairs, which were replaced with wooden chairs[12], fans used to cool various RF amplifiers and the electron multiplying charge-coupled device (EMCCD) cameras, and a lift nearby in the Clarendon laboratory[13]. Secondly, the coils providing the quantisation field in each module are driven by a power supply unit operating in a (nominally) constant current mode. Fluctuations in the current supplying these will directly lead to magnetic field fluctuations at the ion. Finally, wires carrying AC mains current to the various instruments in the lab will induce magnetic field signals that are phase-locked to the 50 Hz AC mains cycle.

The latter two sources of magnetic field noise – the current supply noise and

[12]Many thanks to my colleagues in lab 1 – in particular Sebastian – for putting up with this request!

[13]Conveniently, this lift broke down around the time of taking the distributed quantum computing data. I had nothing to do with this.

the ambient AC mains signals – account for the majority of the magnetic field noise in our apparatus. To suppress these sources, we deploy an active magnetic field stabilisation circuit.

4.5.1 Stabilisation electronics

The magnetic field stabilisation circuit was designed by Dr. D. P. Nadlinger and Dr. T. P. Harty, and is based on the design employed by Merkel *et al.* [120]. The circuit, outlined in Figure 4.13, operates by shunting current away from the coils in order to maintain a constant magnetic field at the ion. A feedback loop enables the stabilisation of the coil current, while a feedforward circuit allows modulation of the coil current at harmonics of 50 Hz to induce a magnetic field that destructively interferes with the ambient field induced by mains wires near the trap.

The stabilisation electronics comprise a versatile stabilizer PCB board [121] connected to a current-sense mezzanine board [122]. The stabilizer board was designed such that it could be interfaced with a variety of mezzanine boards, allowing adaptation to tasks such as laser frequency locking, temperature stabilisation, and, as here, magnetic field stabilisation. The board features a central processing unit (CPU) connected to two ADCs and two DACs via serial peripheral interface (SPI) buses. The mezzanine board is interfaced with the stabilizer board via an ADC and both DACs, as well as an extra DAC integrated into the CPU and a set of general-purpose input/output (GPIO) pins.

The current-sense board features a 10 Ω precision sense resistor[14] that is connected in series with the magnetic field coils. The voltage across the sense resistor, denoted V_{sense}, is measured using a low-noise instrumentation amplifier so as not to draw any current from the coils. A coarse offset voltage V_{offset}, set by an auxiliary DAC, is subtracted from V_{sense} and the resulting signal is amplified and passed to the ADC. This is used as the process variable, x, for the feedback loop. The feedforward is then implemented by digitally modulating the set point of the feedback

[14] Part number Y160710R0000D9R

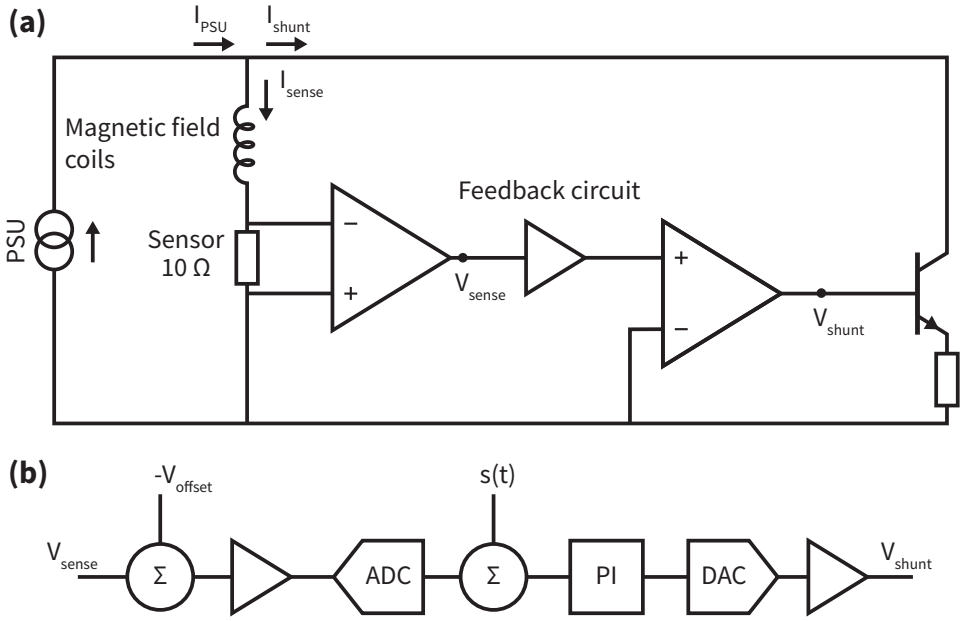

Figure 4.13: Schematic of the magnetic field stabilisation electronics. **(a)** A PSU provides a nominally constant current, I_{PSU}, to the magnetic field coils (depicted as an inductor). The stabilisation circuit dynamically shunts current from the coils based on a voltage, V_{shunt}, applied to a transistor. A $10\,\Omega$ sense resistor is connected in series with the coils, such that the voltage across it, V_{sense}, is proportional to the coil current, I_{sense}. This voltage is fed to a feedback circuit (shown in **(b)**), which controls the shunt voltage, thus allowing the current through the coils to be adjusted. **(b)** The feedback circuit first subtracts a constant offset voltage, V_{offset}, from V_{sense} and amplifies the remaining voltage. An ADC converts this voltage into a digital signal, where the feedforward signal $s(t)$ is added digitally. This combined signal serves as the error signal for a PI filter. The filter's output is then converted to an analogue voltage using a DAC, amplified, and used as V_{shunt} to adjust the coil current.

by the feedforward signal, $s(t)$, resulting in the error signal $e(t) = x(t) - s(t)$. The feedforward signal is phase-locked to the mains line cycle. A standard PI-loop is implemented digitally by the CPU with an update rate of 500 kHz. The output of the PI-loop is passed to a DAC, the output of which drives a transistor which shunts current away from the coils. The stabilizer PCB is connected to the laboratory network via an Ethernet connection over which the stabilizer settings can be configured, and the output of the ADC can be streamed for logging of the error signal. Although the current-sense board was designed with a separate feedforward circuit that was driven by the second DAC, due to concerns about the DC stability of that circuit, the digital option described above was chosen.

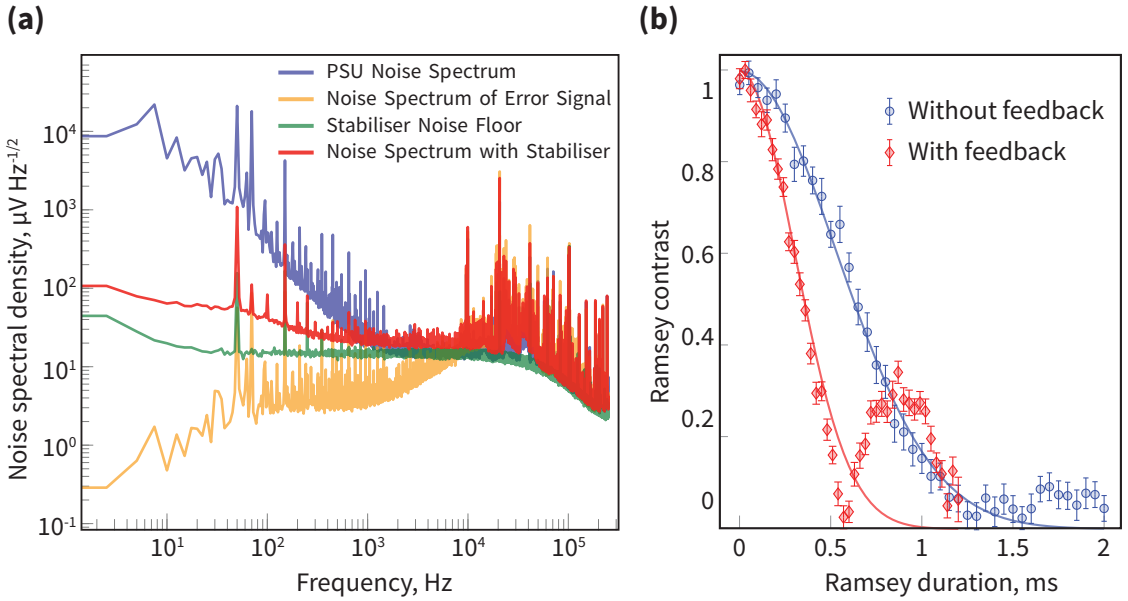

Figure 4.14: **(a)** Noise spectral densities for the magnetic field coil current with (red) and without (blue) the current feedback. The red data was measured using a second Stabiliser unit. Also shown is the noise spectral density of the in-loop error signal (orange), and the noise floor of the stabilisation electronics (green). **(b)** Ramsey contrast of the Sr^+ network qubit with (red) and without (blue) the current feedback, as a function of Ramsey duration. The solid lines are Gaussian fits to the decay in contrast.

4.5.2 Feedback circuit

Let us first consider the action of the feedback circuit only, i.e., we set the feedforward signal to zero, $s(t) = 0$. We measure the current fluctuations by logging the stabilizer ADC input over a period of time. A Fourier transform of this data yielded the noise spectral density of the current fluctuations, as measured by the stabilizer board.

Figure 4.14 shows the noise spectra of the following: the coil current without any feedback (i.e., the power supply unit (PSU) noise spectrum), the coil current as measured by the Stabiliser feeding back on the current supply, the noise floor of the Stabiliser circuit, and the coil current as measured by a second Stabiliser board connected in series with the Stabiliser feeding back on the current supply. Notice that the noise spectrum as measured by the Stabiliser feeding back on the current (orange) goes below the measured noise floor. This is, of course, unphysical since the Stabiliser cannot stabilise the current to better than its noise floor. A second

Stabiliser was connected in series to be able to characterise the true effect of the feedback circuit (red). This noise spectrum does not quite reach the noise floor, suggesting that the performance of the feedback circuit could be further improved, however it does demonstrate a significant reduction of the coil current fluctuations over a 1 kHz bandwidth.

From this data, the rms fluctuations of the magnetic field due to the PSU noise was calculated to be $\sigma_B = 22.52(3)$ nT without any stabilisation, and $\sigma_B = 1.873(2)$ nT with feedback. Using a filter function approach, discussed in Section 2.4.2, the noise spectra in Figure 4.14 were used to estimate the coherence times of the network qubit with and without stabilisation. Coherence times on the order of ~ 8 ms and ~ 1 ms were predicted with (red) and without (blue) stabilisation, respectively.

The coherence of the network qubit was measured with and without feedback, see Figure 4.14(b). The coherence datasets are fitted to Gaussian functions, $e^{-\tau^2/\tau_c^2}$. It is clear that a Gaussian decay is not a good fit for the (red) data – the coherent revival is a result of the 50 Hz oscillations of the ambient magnetic field. However, the fit allows an estimate of the decoherence time, τ_c. Without feedback, the coherence time was estimated to be 0.75(1) ms, while with feedback, the estimated coherence time was 0.43(2) ms.

While the data in figure Figure 4.14(b) shows a longer coherence time without feedback than with feedback, the reason for this is attributed to a difference in amplitudes of the ambient magnetic field at the times when these datasets were taken. This could be due to a change in the electronics operating in the environment of the ion trap. The presence of the coherent revival around 900 µs in the (red) data indicates that the dominant magnetic field decohering mechanism was the 50 Hz oscillations of the ambient AC field. The (blue) data does not exhibit such a revival, indicating that the power supply noise is the dominant source of decoherence. We therefore conclude that while the ambient field amplitude was larger when the (red) data was taken, the Stabiliser effectively reduced the power supply noise.

4.5.3 Feedforward circuit

We now turn our attention to the suppression of the AC ambient magnetic field using the feedforward circuit. Let us denote the ambient AC magnetic field as $\beta(t)$ where $t \in [0, 20\,\text{ms}]$ is the offset from the line trigger. This is the field that results from the mains current-carrying wires, and comprises a 50 Hz signal and harmonics thereof. We decompose this signal into a discrete Fourier series, where $\beta(t)$ can be described by the set of complex amplitudes $\{\beta_n\}$ for $n = 1, \ldots, 5$, such that

$$\beta(t) = \sum_{n=1}^{5} \text{Re}\left\{\beta_n e^{in\Omega t}\right\} \tag{4.5}$$

where $\Omega = 2\pi \times 50\,\text{Hz}$ is the primary harmonic. For our purposes, it is sufficient to consider only the first five harmonics.

The feedforward circuit allows us to apply a signal, $s(t)$, to the coils, which is phase-locked to the 50 Hz line cycle, and thus is phase-coherent with the ambient field from the mains current. As with the ambient field, we can decompose this signal current using the set of complex coefficients $\{s_n\}$, such that

$$s(t) = \sum_{n=1}^{5} \text{Re}\left\{s_n e^{in\Omega t}\right\}. \tag{4.6}$$

The complex amplitudes, $\{s_n\}$, are stored in the stabilizer CPU and define the feedforward signal that is applied to the set-point of the feedback loop. This modulation induces a modulation in the current applied to the coils, which in turn induces a modulation in the magnetic field experienced by the ion. For small signals, we assume that the induced magnetic field depends linearly on the feedforward signal as

$$\sum_{n=1}^{5} \text{Re}\left\{\alpha_n s_n e^{in\Omega t}\right\} \tag{4.7}$$

where $\alpha_n \in \mathbb{C}$ to allow for phases acquired by the Fourier modes of the feedforward in the conversion from current to magnetic field. The magnetic field as measured by the ion, referred to as the response $r(t)$, can then be written as a superposition

of the ambient field and the feedforward field, where a Fourier decomposition yields

$$\sum_{n=1}^{5} \mathrm{Re}\left\{r_n e^{in\Omega t}\right\} = \sum_{n=1}^{5} \mathrm{Re}\left\{(\alpha_n s_n + \beta_n) e^{in\Omega t}\right\}. \tag{4.8}$$

Thus, we obtain

$$r_n = \alpha_n s_n + \beta_n. \tag{4.9}$$

We therefore see that we can suppress the ambient magnetic field arising from the AC mains current by finding the feedforward signal such that the induced field deconstructively interferes with the ambient field, thus minimising the field experienced by the ion. The optimal feedforward coefficients are found by applying a number of stimuli feedforward signals and measuring the response. For qubits with a first-order magnetic field sensitivity, a magnetic field shift will induce a detuning of the qubit frequency proportional to the shift. In order to maximise the sensitivity to the magnetic field shifts, we use a qubit in Sr^+ defined by the states $\left|\mathrm{S}_{1/2}, m_J = -\frac{1}{2}\right\rangle$ and $\left|\mathrm{D}_{5/2}, m_J = \frac{3}{2}\right\rangle$, which has an optical qubit transition with a magnetic field sensitivity of $39.2\,\mathrm{MHz\,mT^{-1}}$. We use Ramsey experiments to map this detuning to a qubit phase, which can be efficiently determined using ABPE, described in Section 3.1. By performing ABPE at different time offsets from a mains line trigger, we can determine the magnetic field shifts as experienced by the ion at different points along the line cycle, i.e., the response $r(t)$. The ABPE is performed twice per line trigger, once at the line trigger itself, and then again at the offset from the line trigger. By taking the difference of the two phase estimates, it is then possible to approximately cancel out any slow drifts in the ambient magnetic field, leaving only the magnetic field induced by the mains wires.

We calibrate the feedforward by applying M stimuli feedforward signals, specified by $\{s_n^i\}$ where $n = 1, \ldots, 5$ and $i = 1, \ldots, M$, measuring the responses $\{r_n^i\}$ on the ion, and solving for $\{s_n^{(\mathrm{opt})}\}$ such that

$$0 = \alpha_n s_n^{(\mathrm{opt})} + \beta_n. \tag{4.10}$$

We follow the method employed in [81], making use of a complex linear regression. For each harmonic, n, we define the response vector, $\vec{r}_n$, and the matrix of regressors, S_n, as

$$\vec{r}_n = \begin{pmatrix} r_n^1 \\ r_n^2 \\ \vdots \\ r_n^M \end{pmatrix}, \text{ and } S_n = \begin{pmatrix} 1 & s_n^1 \\ 1 & s_n^2 \\ \vdots & \vdots \\ 1 & s_n^M \end{pmatrix}. \tag{4.11}$$

Furthermore, by defining the parameter vector, $\vec{p}_n = (\beta_n, \alpha_n)$, we can write

$$\vec{r}_n = S_n \vec{p}_n + \vec{\epsilon}_n \tag{4.12}$$

where ϵ_n are the residuals of the model. This allows us to write the sum of the squared residuals as,

$$R_n = \vec{\epsilon}_n^{\dagger} \vec{\epsilon}_n = (S_n \vec{p}_n - \vec{r}_n)^{\dagger} (S_n \vec{p}_n - \vec{r}_n) \tag{4.13}$$

which is just the sum of the residuals of each trial i. The best estimate of the parameters, $\hat{\vec{p}}_n = (\hat{\beta}_n, \hat{\alpha}_n)$, are such that

$$\frac{\partial R_n}{\partial \vec{p}_n} = 0. \tag{4.14}$$

Computing this derivative, we obtain find

$$\hat{\vec{p}}_n = (S_n^{\dagger} S_n)^{-1} S_n^{\dagger} \vec{r}_n. \tag{4.15}$$

The optimal feedforward amplitudes are then given by

$$s_n^{(\text{opt})} = -\frac{\hat{\beta}_n}{\hat{\alpha}_n}. \tag{4.16}$$

Figure 4.15(a) shows the results from the measurements of the magnetic field over the line cycle with (red) and without (blue) the feedforward applied. Without

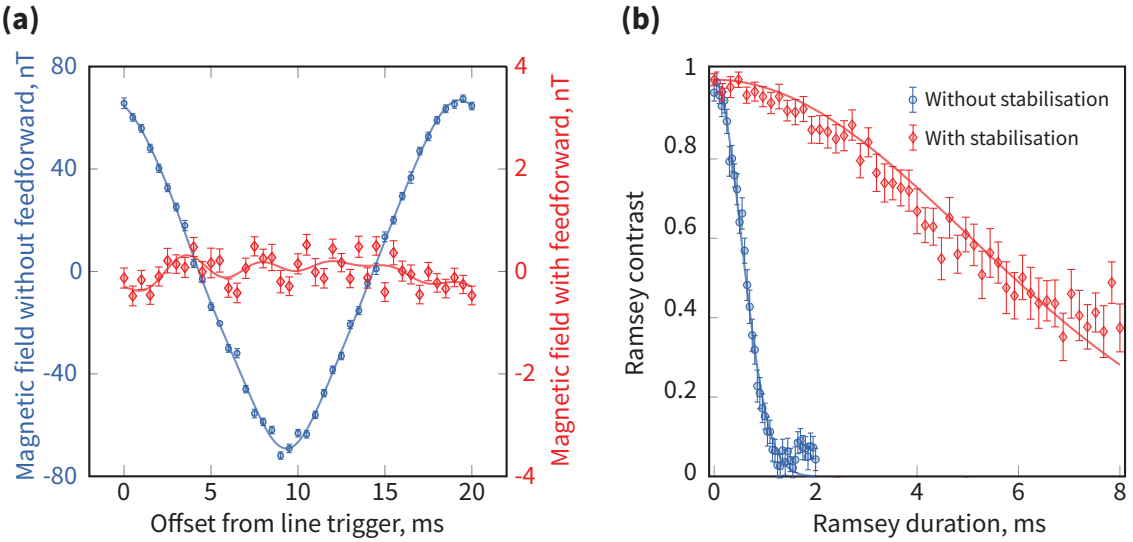

Figure 4.15: **(a)** Deviation of magnetic field from nominal value at different points along the 50 Hz AC line-cycle with (red) and without (blue) magnetic field stabilisation (comprising both feedback and feedforward). The deviations were measured using ABPE. **(b)** Ramsey contrast of the Sr^+ network qubit with (red) and without (blue) magnetic field stabilisation, as a function of Ramsey duration. The solid lines are Gaussian fits to the decay in contrast.

the feedforward, the peak-to-peak value of the ambient field is 136(2) nT [rms of 46.0(7) nT], which is reduced to 0.7(4) nT [rms of 0.299(4) nT] with the application of the feedforward. The Bayesian phase estimation scheme involves performing Ramsey experiments at the same point in the line cycle, so this rms value does not account for ambient field fluctuations that are asynchronous with the line cycle. Indeed, we expect the magnetic field noise to be dominated by the residual noise from the feedback circuit, which has $\sigma_B \approx 2\,\text{nT}$.

Measurement of the Ramsey contrast of the network qubit with feedforward is shown in Figure 4.15(b). The data without any magnetic field stabilisation is also plotted for reference. From this data, the network qubit coherence time is estimated to be 7.1(1) ms. This is comparable to the coherence time of 8.24(3) ms, predicted in section Section 4.5.2, which assumes only residual noise from the feedback circuit. Combining both sources, we estimate an rms for the total magnetic field fluctuations of 2.0(1) nT. This represents over an order of magnitude reduction in the magnetic field fluctuations from the noisy coil current supply and ambient fields induced by wires carrying 50 Hz mains current (and harmonics thereof).

5

Local Operations

*

This chapter presents the local control operations essential for our trapped-ion modules: single-qubit rotations on Ca^+ qubits, mixed-species two-qubit entangling gates, coherent mapping between hyperfine qubits, and mid-circuit measurement techniques. These tools form the building blocks for the complex networked experiments described in later chapters.

The results in this chapter reflect contributions from across the quantum networking team, many of which I directly supported. I implemented and characterised high-fidelity single-qubit rotations on our Ca^+ qubits using the Raman laser system characterised in Chapter 4. I contributed to the implementation of mixed-species entangling gates based on work by Hughes *et al.* [15]. Finally, I performed randomised benchmarking of the coherent hyperfine-qubit mapping pulse developed by P. Drmota.

5.1 Single-qubit gates

The ability to coherently manipulate individual qubits is a critical tool for universal quantum computing. Here, we outline the mechanisms by which we mediate ar-

bitrary single-qubit rotations for the network qubits encoded in Sr^+ ions, and the auxiliary and circuit qubits encoded in the Ca^+ ions. We benchmark the performance of our single-qubit gates using RBM, as outlined in Section 3.4.

For all relevant transitions that we coherently address, we utilise the ABPE protocols outlined in Section 3.1 to perform efficient and precise calibration of the qubit transition frequencies, in addition to the pulse durations corresponding to $\pi/2$-, π-, and 2π-rotations. Combined with arbitrary Z-rotations – trivially implemented by adjusting the tracked qubit phases – we are able to implement any arbitrary single-qubit rotation.

5.1.1 Strontium quadrupole gates

We coherently manipulate the electronic state of the Sr^+ ions using a narrow-linewidth 674 nm laser to address electric quadrupole transitions between sublevels of the $S_{1/2}$ manifold and sublevels of the $D_{5/2}$ manifold. In particular, we are interested in the addressing of the transition between the $\left|S_{1/2}, m_J = -\frac{1}{2}\right\rangle$ and $\left|D_{5/2}, m_J = -\frac{3}{2}\right\rangle$ states forming the network qubit. We benchmark the performance of the single-qubit rotations of the network qubit using RBM. The RBM results for the network qubit in each module are given in Table 5.1. Notice that while the auxiliary qubit exhibits a magnetic field sensitivity similar to the network qubit, it has an order of magnitude smaller single-qubit gate error; thus, we can conclude that fast magnetic field noise is not the dominant error mechanism. Additionally, since the error for Alice is almost twice as large as that in Bob, this suggests that laser phase noise does not fully account for the observed errors. We therefore assume that the off-resonant driving of nearby motional sidebands is a significant source of error, with the larger effect in Bob resulting from the lower secular motional frequencies.

5.1.2 Calcium Raman gates

We coherently manipulate states in the ground hyperfine manifold of the Ca^+ ions using a pair of Raman beams – see Section 4.3 for a description of the Raman laser

Qubit	Error per Clifford ($\times 10^{-4}$)	
	Alice	Bob
network ($\mathcal{Q}_{\mathrm{N}}$)	4.8(3)	9.8(3)
auxiliary ($\mathcal{Q}_{\mathrm{X}}$)	1.4(3)	1.2(3)
circuit ($\mathcal{Q}_{\mathrm{C}}$)	1.0(3)	1.2(4)

Table 5.1: Errors per single-qubit Clifford operation for network, auxiliary, and circuit qubits, measured using RBM.

system. For transitions between states where $\Delta m_F = 0$, such as in the case of the circuit qubit, we use the $\mathrm{Son}_{\parallel}$ and $\mathrm{Daughter}_{\parallel}$ beams; for transitions between states where $\Delta m_F = \pm 1$, such as in the case of the auxiliary qubit, we use the $\mathrm{Son}_{\perp}$ and $\mathrm{Daughter}_{\perp}$ beams. By choosing pairs of co-propagating beams to drive the transitions, we minimise the relative wavevector of the two beams, and therefore the coupling to the ions' motion.

We benchmark the performance of the single-qubit rotations of the circuit and auxiliary qubits using RBM. The RBM results for the auxiliary and circuit qubits in each module are given in Table 5.1. Since the Raman beam geometry suppresses potential off-resonant motional coupling effects, we anticipate the dominant sources of single-qubit gate error to be from magnetic field noise, and relative phase noise between the two Raman lasers. However, given the similarity in the gate errors between the auxiliary qubit and the field-insensitive circuit qubit, we assume that the dominant error mechanism is relative phase noise between the Raman lasers, as this affects both qubits equally.

5.2 Mixed-species entangling gates

The ability to perform logical entangling gates between ions of different species allows us to delegate roles for performing tasks which have diametric requirements, such as storing quantum information and realising a network interface. We perform these mixed-species entangling gates using geometric phase gates. As discussed in Section 2.4.3, the implementation of geometric phase gates requires the generation

of a SDF. In this section, we discuss the implementation and calibration of our entangling gates, and the characterisation of our mixed-species controlled-Z (CZ) and iSWAP gates. These particular gates are deployed frequently in the more complex experiments presented in Chapters 6 and 7.

5.2.1 The spin-dependent force

In Section 2.4.3, we discussed the theoretical implementation of a light-shift SDF; here, we outline our experimental realisation of this SDF. As depicted in Figure 5.1, the SDF is created using a single pair of intersecting 402 nm Raman beams, detuned from one another by $\Delta\omega$, which induce differential light-shifts of the ions' spin states. The Raman beams we use are the $\text{Son}_{\parallel}$ and $\text{Son}_{\perp}$ beams, as defined in Chapter 4. At this wavelength, we couple most strongly to the 397 nm $S_{1/2} \leftrightarrow P_{1/2}$ dipole transition in Ca^+ and the 408 nm $S_{1/2} \leftrightarrow P_{3/2}$ dipole transition in Sr^+, as shown in Figure 5.1(a). The $\sim$ THz detuning from these transitions suppresses unwanted single-photon scattering, while still providing a sufficiently strong driving force with experimentally feasible beam intensities. This configuration is advantageous since we can address both ion species with only a single pair of beams.

The interference of the two beams results in a polarisation travelling-standing wave propagating along the direction of the wavevector differential $\Delta\vec{k}$, as depicted in Figure 5.1(b). The spin states of an ion sitting in this travelling-standing wave will experience a light-shift oscillating at $\Delta\omega$ with a phase depending on the position of the ion in the travelling-standing wave. The ions experience this spatially-dependent light-shift as a dipole force which drives motion along $\Delta\vec{k}$, with a sign dependent on the spin state of the ion. By choosing the differential detuning of the Raman beams to be close to one of the motional modes, i.e., $\Delta\omega = \omega_m + \delta$ where $|\delta| \ll \omega_m$, we off-resonantly drive the motional mode and coherently displace the spin states around circular trajectories in the phase space of the motional mode.

As discussed in Section 2.2.1, the mass disparity between Sr^+ and Ca^+ results in large asymmetries in the radial mode participations of the ions, making these

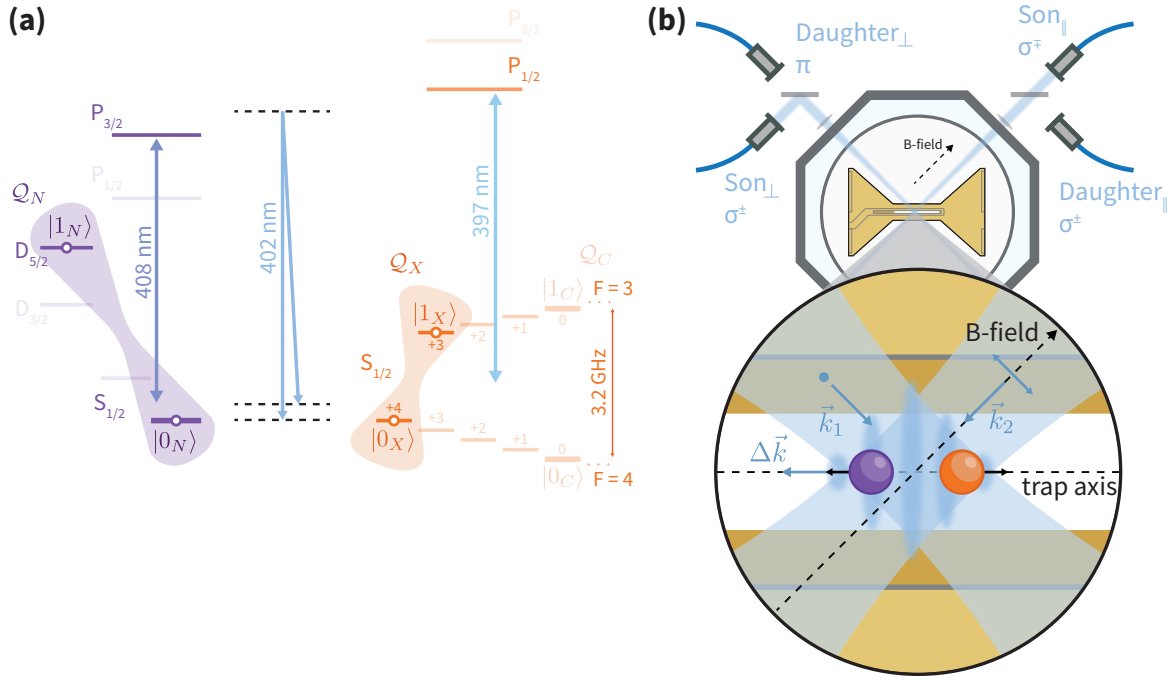

Figure 5.1: **(a)** Energy level diagram for Sr^+ and Ca^+ showing states used to realise the mixed-species SDF. A pair of 402 nm Raman lasers with a relative detuning of $\Delta\omega$ induce a light-shift of the auxiliary ($\mathcal{Q}_X$) and network ($\mathcal{Q}_N$) qubit states. By tuning the Raman lasers to 402 nm, we couple to both the 397 nm $S_{1/2} \leftrightarrow P_{1/2}$ dipole transition in Ca^+ and the 408 nm $S_{1/2} \leftrightarrow P_{3/2}$ dipole transition in Sr^+. **(b)** Geometry of our SDF implementation. The Raman lasers are aligned orthogonal to one another, such that their relative wavevector is along the trap axis. The interference of the Raman beams gives rise to a polarisation travelling-standing wave which induces spin-dependent light-shifts of the spin states oscillating at $\Delta\omega$ with a phase depending on the position of the ions in the interference pattern. The detuning of the Raman lasers, $\Delta\omega$, is tuned close to the secular frequency of the OOP normal mode, ω_{OOP}. The ions experience a SDF off-resonantly driving the spin states around closed trajectories in phase space.

modes unsuitable for mediating entangling gates. Instead, we choose the axial OOP mode since each ion has a significant participation in the mode, and the "breathing" character of the ions' motion provides some suppression of the motional heating due to common-mode electric field noise. We therefore align the Raman beams in an orthogonal geometry such that their differential wavevector points in the axial direction and we can maximally drive the axial motion.

As discussed in Section 2.4.3, the SDF Hamiltonian for two ions in the travelling-

standing wave is given by

$$\hat{H}_{\text{SDF}} = i \sum_{i=1}^{N} \sum_{s=0,1} \frac{\hbar \Omega_{i,s} \eta_{i,\text{OOP}}}{2} e^{i\Delta\phi_i} e^{-i\delta t} \left|s_i\right\rangle \left\langle s_i\right| \otimes \hat{a}^{\dagger}_{\text{OOP}} + \text{h.c.}, \tag{5.1}$$

where $\Omega_{i,s}$ determines the size of the light-shift of the spin state $s \in \{0, 1\}$ of the i^{th} ion, $\eta_{i,\text{OOP}}$ is the Lamb-Dicke parameter for the coupling of the Raman beams to the component of the i^{th} ion's motion in the OOP mode, $\Delta\phi_i$ is the phase of the travelling standing wave at the position of the i^{th} ion, $\hat{a}^{\dagger}_{\text{OOP}}$ is the creation operator the OOP motional mode, δ is the detuning of the travelling-standing wave from the OOP mode frequency, and $\left|s_i\right\rangle \left\langle s_i\right|$ is the projector for the spin state s of the i^{th} ion embedded into the total spin Hilbert space.

The coupling of the SDF to the spin states depends on the ability to generate differential light shifts between the qubit states, i.e., that we can realise $\Omega_{i,1} \neq \Omega_{i,0}$. For the Ca^+ circuit ($\mathcal{Q}_{\text{C}}$) qubits, we have $\Omega_{\text{C},1} = \Omega_{\text{C},0}$ and thus the circuit qubit states will be driven along the same trajectories in phase space. As a result, the circuit qubits will only acquire a global geometric phase and so cannot be used to generate entanglement. We therefore drive the entangling gates using the auxiliary ($\mathcal{Q}_{\text{X}}$) qubit instead.

We apply the gate mechanism directly to the network qubit in Sr^+ – rather than the Zeeman ground state qubit, as done in Hughes *et al.* [15] and Drmota *et al.* [123] – which has the advantage of eliminating mapping pulses to convert between the Zeeman qubit and the optical qubit. However, a consequence of this is an asymmetry between the matrix elements, $|\Omega_{\text{N},1}| \ll |\Omega_{\text{N},0}|$, which is not present for the Zeeman qubit. Due to the hyperfine structure of $^{43}Ca^+$, the matrix elements for the auxiliary qubit also exhibits an asymmetry, $\Omega_{\text{X},1} = -\Omega_{\text{X},0} + \epsilon$. This will lead to an asymmetry in the trajectories taken by the different qubit states. As we discuss in Section 5.2.2, we mitigate the effect of these asymmetries through the deployment of Walsh modulation, consisting of a $\hat{\sigma}_x$ rotation bisecting the SDF pulse [124, 125]. This modulation symmetrises the trajectories of the qubit states,

thereby enabling the execution of the entangling gate at the cost of a reduced gate efficiency that is compensated for by the use of higher laser powers (the laser powers used are summarised in Table 5.2).

The phase of the travelling standing wave at the position of the i^{th} ion, $\Delta\phi_i$, is given by

$$\Delta\phi_i = \Delta\phi_r + \Delta\vec{k} \cdot \vec{r}_{i,0}, \tag{5.2}$$

where $\Delta\phi_r$ is the relative phase of the two Raman beams and $\vec{r}_{i,0}$ is the equilibrium position of the i^{th} ion. The relative phase of the driving force acting on each of the ions is then $\Delta k \Delta z$ where Δz is the axial separation of the two ions. We may then set the relative phase of the driving force by tuning the ion spacing. With an ion spacing such that $\Delta k \Delta z = 2\pi(n + 1/2)$, the driving force experienced by the ions in the even-parity states $|1_{\text{N}}1_{\text{X}}\rangle$ and $|0_{\text{N}}0_{\text{X}}\rangle$ will be maximally out-of-phase and so we most efficiently drive the OOP mode. For our beam geometry, we have $\Delta k = \sqrt{2} \times 2\pi/\lambda$ and thus in each module we choose axial confinements such that

$$\Delta z_n = \left(n + \frac{1}{2}\right) \frac{\lambda}{\sqrt{2}} \tag{5.3}$$

for $n \in \mathbb{Z}$. The ion spacings used in each module are summarised in Table 5.2, and the method by which we find the correct axial confinement to realise these ion spacings is described in Section 5.2.3.

5.2.2 Walsh-modulated entangling gate

As discussed in Section 2.4.3, we implement unitary entangling gates between spin states through careful application of the light-shift SDF. This SDF couples the spins to the shared motion, and so the force will, in general, generate entanglement between the spin states and the state of the motional mode. In order to mediate a unitary interaction between the spin states, we require that once we have completed the interaction, the spin states and motional states are fully decoupled. Residual spin-motion entanglement will manifest as non-unitary evolution of the spin states,

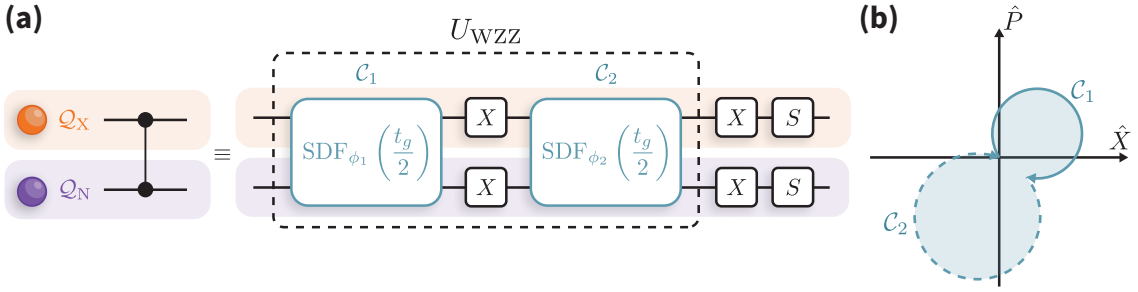

Figure 5.2: **(a)** Circuit used to implement the Walsh-modulated CZ gate. The Walsh modulation comprises two pulses of the spin-dependent force of length $t_g/2$ and detuning $\delta_g = 4\pi/t_g$, separated by $\hat{\sigma}_x$ pulses on each qubit. The gate parameters are chosen such that the acquired geometric phase is $\Phi = \pi/2$, resulting unitary process is the U_{WZZ} entangling gate. Finally, single-qubit rotations are then used to transform this to the desired CZ-gate. **(b)** Phase space trajectory for one of the spin states during the Walsh-1 modulated gate for a mis-set gate detuning and a Rabi frequency asymmetry. The mis-set gate detuning results in the first trajectory ($\mathcal{C}_1$) not closing. The phase of the second SDF pulse is chosen to be $\phi_2 = \phi_1 + \delta_g\tau$ such that the second loop ($\mathcal{C}_2$) closes the total trajectory. The asymmetry in the matrix elements of Ω results in trajectories of different sizes. Since the spin states travel both trajectories, the total geometric phase acquired is the same.

when considering only the reduced density matrix of the spin states, and thus a gate error. The decoupling of the spins and the motion occurs when the trajectories taken by the spin states around the phase space are closed; in the absence of external perturbations, this is achieved when the SDF is applied for integer multiples of the duration $t = 2\pi/\delta$.

To make the entangling gate robust to coherent error processes, we employ a dynamical error suppression protocol known as Walsh modulation [124, 125]. Walsh modulation consists of dividing the interaction into a sequence of shorter interactions with different phases. Similar to dynamical decoupling schemes used to suppress qubit decoherence processes, this modulation of the gate interaction allows us to coherently suppress error processes that occur during the gate.

In our experiment, we use 1$^{\mathrm{st}}$-order Walsh modulation. Shown in Figure Figure 5.2, this consists of dividing the gate interaction into two applications of the SDF separated by $\hat{\sigma}_x$ rotations of each qubit. The Walsh modulation ensures that if the phase space trajectories are not closed after the first SDF pulse, the relative phase of the two SDF pulses may be selected such that the second loop always closes

the total trajectory. To achieve this, we require that the phase of the SDF on each spin state is the same at the start of each loop. We therefore choose $\phi_2 = \phi_1 + \delta_g \tau$, where the $\delta_g \tau$ term accounts for the free evolution of the motional state over the duration, τ, between the start of the first SDF pulse and the start of the second. Additionally, the $\hat{\sigma}_z \otimes \hat{\sigma}_z$-type interaction realised by the SDF commutes with qubit frequency shifts, and so the $\hat{\sigma}_x$ rotation implements a spin-echo on each qubit, making the gate robust to qubit frequency detunings. Finally, the modulation mitigates the asymmetries in the Rabi frequencies, $\Omega_{i,1}$ and $\Omega_{i,0}$, since the spin states will undergo both trajectories.

In the absence of errors, this implementation of the Walsh-modulated ZZ (WZZ) gate yields the unitary transformation

$$\hat{U}(\Phi) = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & e^{-i\Phi} & 0 \\ 0 & e^{-i\Phi} & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \tag{5.4}$$

where Φ is the acquired geometric phase. We choose gate parameters such that $\Phi = \pi/2$, i.e., $\hat{U}_{\mathrm{WZZ}} = \hat{U}(\pi/2)$. Single-qubit rotations then transform this gate into the CZ-gate, given by the matrix

$$\hat{U}_{\mathrm{CZ}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}. \tag{5.5}$$

This mixed-species CZ gate is an important resource in the subsequent experiments.

5.2.3 Gate calibration

In order to realise an entangling gate that is as close to the ideal unitary as possible, there are a number of parameters of that must be carefully calibrated. This includes

Parameter	Alice	Bob
Mode frequency, ω_{OOP}	3.341 MHz	2.614 MHz
Ion-spacing, Δz_n	$\left(11+\frac{1}{2}\right)\lambda/\sqrt{2} \sim 3.3\,\mu\text{m}$	$\left(13+\frac{1}{2}\right)\lambda/\sqrt{2} \sim 3.8\,\mu\text{m}$
Raman powers, $(P_\parallel, P_\perp)$	(20 mW, 19 mW)	(16 mW, 24.5 mW)
Gate time, t_g	62 µs	58 µs

Table 5.2: Typical gate parameters for each module. At the ions, the Raman beams are focused to a beam waist of $\sim 15\,\mu\text{m}$.

the secular frequency of the axial OOP mode, ω_{OOP}, the ion spacing, Δz, and the gate parameters $\mathcal{G} = \{P_\parallel, P_\perp, t_g, \delta_g\}$ where $P_\parallel$ and $P_\perp$ are the powers in the $\text{Son}_\parallel$ and $\text{Son}_\perp$ Raman beams, respectively. Typical values of these parameters used in each module are given in Table 5.2.

Mode frequency

We calibrate the frequency of the axial OOP mode, ω_{OOP}, such that at $\delta_g = 0$ the force resonantly drives the motion. In this scenario, the spin states do not undergo circular trajectories; instead, they are coherently displaced along a straight line in phase space[1].

We calibrate this frequency by reducing the power in each Raman laser beam to $\sim 1\,\text{mW}$ and applying the SDF for $\sim 500\,\mu\text{s}$. We then probe on the red sideband to detect the coherent displacement. By scanning the detuning $\Delta\omega$ of the Raman beams, we are able to identify the point at which the SDF becomes resonant with the OOP motional mode. Results from a typical calibration scan in Alice are shown in Figure 5.3. Typical axial OOP mode frequencies for each module are given in Table 5.2.

Ion spacing

As discussed in Section 5.2.1, the ion spacing of the Sr^+–Ca^+ crystal is chosen such that $\Delta z_n = \left(n+\frac{1}{2}\right) \times \lambda/\sqrt{2}$, for $n \in \mathbb{Z}$, so as to maximise the coupling of the spin states to the axial OOP mode. We tune the ion spacing by either

[1] In the frame rotating with the secular frequency ω_{OOP}.

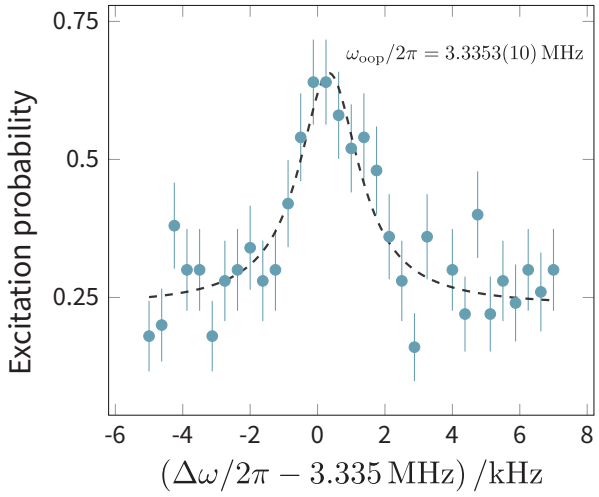

Figure 5.3: Typical scan for the calibration of the mode frequency in Alice. A weak SDF is applied for $\sim 500\,\mu s$ with a variable frequency $\Delta\omega$. This frequency is scanned to identify the resonance of the OOP mode used to mediate the spin-dependent force. At $\Delta\omega = \omega_{OOP}$, the SDF resonantly drives the OOP mode, therefore increasing the excitation probability when probing the red sideband. A Lorentzian fit yields a mode frequency of $\omega_{OOP}/2\pi = 3.3353(10)\,\mathrm{MHz}$.

increasing or reducing the trap axial confinement. We identify the optimal ion spacing experimentally by finding the points $\Delta z_{n,\pm} = \left(n + \frac{1}{2} \pm \frac{1}{2}\right) \lambda/\sqrt{2}$ either side of the desired ion spacing that minimise the coupling of the SDF to the OOP mode. The chosen ion spacing for each module is given in Table 5.2.

The coupling of the force to the motion is determined by resonantly driving the OOP mode before probing on the red motional sideband of a qubit transition. If the force couples to the motion, then the motional state is coherently displaced in phase space and we observe an increase in excitation probability on the red sideband. At the points where the coupling is minimised, we do not observe a signal on the red sideband since the motional state remains close to the ground state.

Since the axial confinement sets the axial frequencies, we measure the single-ion axial frequencies for the axial confinements that minimise this coupling either side of the desired ion spacing. We then set the single-ion axial frequency between the two points[2]. At this point, the coupling between the SDF and the OOP mode should be

[2]Of course, it would be simpler to simply find the point at which the coupling to the IP mode was minimised. However, the cooling of the IP mode was not good enough to clearly see a signal.

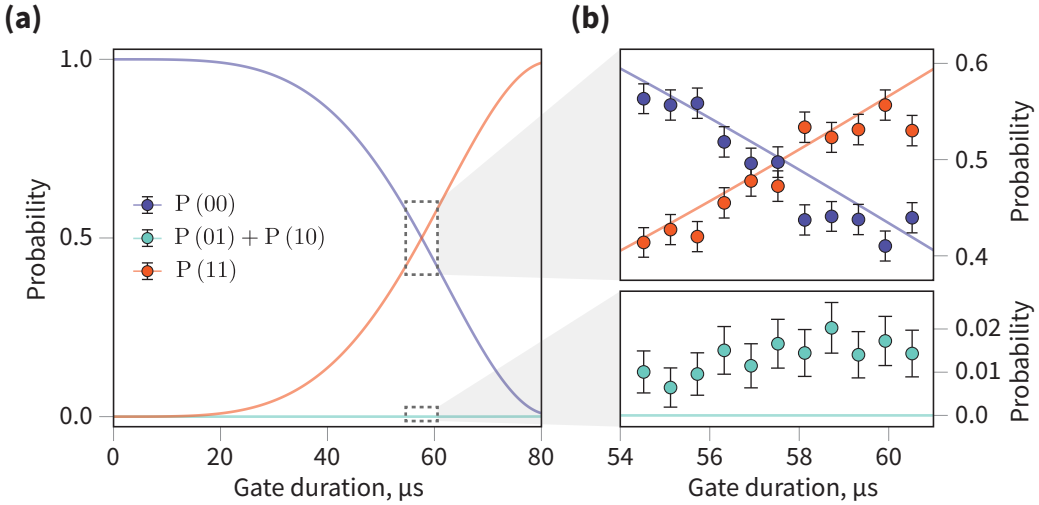

Figure 5.4: **(a)** Simulation of ideal gate dynamics for typical gate parameters, showing population as a function of gate duration. At each duration, δ is set to $4\pi/t_g$ to ensure that the phase space trajectories are closed, allowing observation of the dynamics from Eq. (5.6). The duration at which $\mathrm{P}(00) = \mathrm{P}(11) = \frac{1}{2}$ indicates the creation of the maximally entangled $|\Phi^+\rangle$ state and is thus used as the gate duration for the CZ gate. **(b)** Experimental data from Bob (circles) overlaid with simulated ideal dynamics (solid lines) for a $\sim 8\,\mu$s window around the calibrated gate time. The crossing point of P(00) and P(11) determines the gate duration for the CZ gate.

maximised and thus we most efficiently drive gates on the OOP mode.

Gate parameters

In order to leave the spin states decoupled from the motional states at the end of gate interaction, we must close the phase space trajectories. This is achieved when the time taken to traverse the closed trajectory is equal to the inverse detuning of the force from the motion. Thus, for an m-loop gate with force detuning δ_g, the total gate duration will be $t_g = 2\pi m/\delta_g$. In the case of our Walsh-1 modulated gate, we use two loops and thus our gate detuning is set to $\delta_g = 4\pi/t_g$.

We then have two free gate parameters that we may vary to achieve the desired geometric phase $\Phi = \pi/2$: the Raman laser powers, and the gate duration t_g. We choose to fix the Raman beam powers, $P_\parallel$ and $P_\perp$, and vary the duration t_g (and gate detuning δ_g accordingly) to achieve the desired dynamics.

We calibrate the gate duration by preparing the Sr^+-Ca^+ crystal in the state $|0_N 0_X\rangle$ and applying a sequence consisting of the Walsh-modulated gate unitary,

given by Eq. (5.4), interposed between single-qubit $\hat{R}_X\left(\frac{\pi}{2}\right)$ rotations on each ion. In the absence of experimental imperfections, this yields

$$\hat{R}_X\left(\frac{\pi}{2}\right)^{\otimes 2}\ \hat{U}(\Phi)\ \hat{R}_X\left(\frac{\pi}{2}\right)^{\otimes 2}\ |0_{\mathrm{N}}0_{\mathrm{X}}\rangle = \cos(\Phi)\,|0_{\mathrm{N}}0_{\mathrm{X}}\rangle - i\sin(\Phi)\,|1_{\mathrm{N}}1_{\mathrm{X}}\rangle\,, \quad (5.6)$$

where Φ is the acquired geometric phase and, for fixed Raman intensities, depends only on the gate duration, t_g. We measure in the computational basis to estimate the populations of the resulting state.

A numerical simulation of this sequence in the absence of experimental imperfection is shown as the solid lines in Figure 5.4. In the ideal case, we find no population in the odd-parity states at any duration[3], i.e., $P_{01} + P_{10} = 0$. The points at which $P_{11} = P_{00} = 0.5$ indicates the creation of a maximally entangled state $\frac{1}{\sqrt{2}}(|0_{\mathrm{N}}0_{\mathrm{X}}\rangle - i\,|1_{\mathrm{N}}1_{\mathrm{X}}\rangle)$, and thus that $\Phi = \pi/2$, yielding the desired entangling gate. Figure 5.4(b) shows typical experimental results of such a scan in Alice. The deviations from the ideal simulation are discussed in Section 5.2.5.

5.2.4 Characterisation of the mixed-species CZ gate

We characterise our mixed-species gates using QPT, outlined in Section 3.3. To suppress the effects of imperfect SPAM on the reconstructed process matrix, we incorporate these imperfections into the model used to reconstruct the process matrix, as outlined in Section 3.3.2. This allows us to use the characterisations presented here in error budgets for more complex experiments, without worrying about compounding SPAM errors.

As discussed in Section 5.2.2, due to the Walsh modulation, the gate we implement is (in the absence of experimental imperfections) the unitary U_{WZZ}. However, since this gate is equivalent to the CZ-gate up to single-qubit rotations, we consider the CZ-gate as the base gate in our local mixed-species gate set. Therefore, the characterisation of our mixed-species entangling gate, and any subsequent circuits utilising these operations, will be in terms of the CZ-gate.

[3]Note that for each gate duration, t_g, we adjust the gate detuning, δ_g, accordingly.

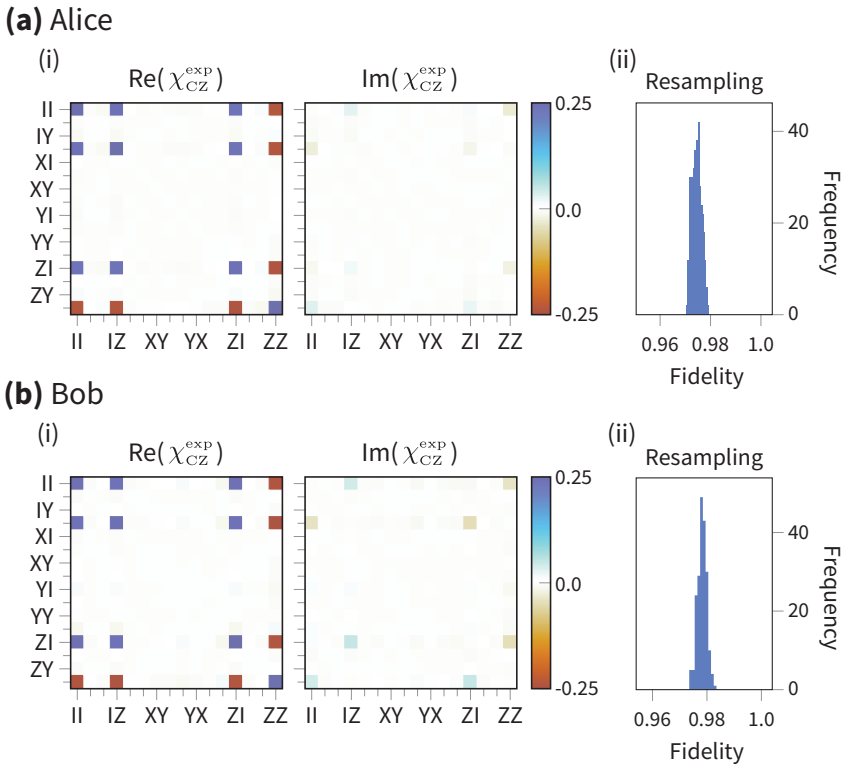

Figure 5.5: **(a)**(i) and **(b)**(i) show the process matrices for the mixed-species CZ gates in Alice and Bob, respectively, reconstructed from QPT using a total of 32 400 measurements. From the process matrices, we observe CZ gate fidelities of 97.6(2) % and 98.0(2) % for Alice and Bob, respectively. **(a)**(ii) and **(b)**(ii) show the histograms of the gate fidelities of process matrices reconstructed from resampled datasets.

The process matrices, reconstructed from 32 400 tomographic measurements, for the mixed-species CZ-gates from each module are shown in Figure 5.5. Compared to the ideal CZ-gate, we measure average gate fidelities of 97.6(2) % and 98.0(2) % for Alice and Bob, respectively. The experimental imperfections to which we attribute the loss in fidelity for the CZ gate are discussed in the following section.

5.2.5 Experimental imperfections

There are a number of mechanisms which result in a deviation from the desired unitary gate, $\hat{U}_{CZ}$. The experimental results from a gate dynamics scan in Alice, shown in Figure 5.4(b), show population in the odd-parity states, in contrast to the numerical simulation of the ideal dynamics. Here, we consider the deviations from the desired unitary dynamics which we expect to affect the performance of our

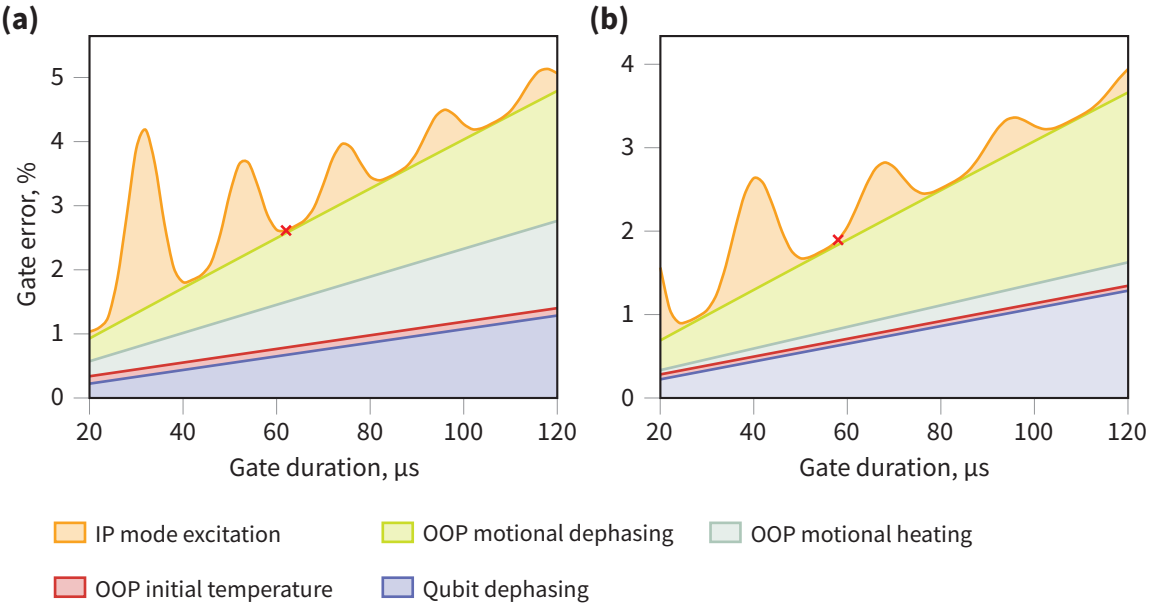

Figure 5.6: Simulated contributions to the CZ gate error as a function of gate duration for **(a)** Alice and **(b)** Bob. At each gate duration, the optimal gate detuning and gate powers are chosen. For both modules, we assume qubit dephasing times of $\tau_N = 14\,\text{ms}$ and $\tau_X = 7\,\text{ms}$ for the network and auxiliary qubits, respectively, and assume a motional dephasing time of $\tau_{OOP} = 3.4\,\text{ms}$. For Alice (Bob), we assume the following parameters: $\bar{n}_{OOP} = 0.15$ ($\bar{n}_{OOP} = 0.1$), $\dot{\bar{n}}_{OOP} = 350\,\text{s}^{-1}$ ($\dot{\bar{n}}_{OOP} = 80\,\text{s}^{-1}$), $\bar{n}_{IP} = 2.0$ ($\bar{n}_{IP} = 0.8$), and $\dot{\bar{n}}_{IP} = 2700\,\text{s}^{-1}$ ($\dot{\bar{n}}_{IP} = 500\,\text{s}^{-1}$). The red markers indicate typical gate durations of 62 µs and 58 µs for Alice and Bob, respectively. This visualisation of the error contributions was inspired by Fig. 4.21 in Drmota [56].

mixed-species entangling gate. The effects of the experimental imperfections on the gate dynamics are modelled by numerically solving the master equation, Eq. (2.5), while including the different error models. The estimated error contributions as a function of gate duration for the different error mechanisms discussed in this section are shown in Figure 5.6.

Miscalibration of gate parameters

The spin states must be completely disentangled from the motion after the gate interaction; residual spin-motion entanglement will manifest as a non-unitary evolution of the spin states. In the gate dynamics scans, this non-unitary evolution will lead to an increased population in the odd-parity states. One mechanism by which residual spin-motion entanglement may occur is if the phase space trajectories taken by the spin states are not closed due to miscalibration of the gate parameters.

While we always set $\delta_g = 2\pi m/t_g$ for an m-loop gate, any deviations of the mode frequency from its calibrated value will result in the trajectories not being closed. As discussed in Section 5.2.2, Walsh modulation suppresses the effects of most mis-set gate parameters. However, this modulation relies on precise calibration of the mode frequency, ω_{OOP} to calculate the phase of the second SDF pulse, and so this gate mechanism is sensitive to deviations in the mode frequency. There are a number of mechanisms that lead to significant deviation in the mode frequency.

Firstly, we observe that the Raman lasers charge the trap chips and cause a significant miscompensation of the trapping potential. This in turn affects the mode structure of the crystal and thus the OOP mode frequency. To mitigate this effect, we automatically calibrate the micromotion compensation in the out-of-plane direction every $\sim$ 5-10 minutes, keeping the mode frequency stable.

Secondly, the two possible ion orders of the Sr^+-Ca^+ crystal have slightly different mode frequencies due to experimental asymmetries. We observe reordering events of the ion crystal on the $\sim$10 minute timescale. In order to detect reordering events, before every experimental shot we look for Sr^+ fluorescence through the high-numerical aperture lens used for collection of photons in the ion-photon interface. Since the coupling of the collected light to the detection system is extremely sensitive to ion position, we only collect fluorescence for one ion order. Upon detection of a reordering event, we automatically reorder the crystal, thus returning it to the correct state.

Qubit dephasing

Typical gate durations in our modules are $\sim 60\,\mu$s, and typical coherence times for the network and auxiliary qubits are $\sim 10\,$ms and $\sim 1\,$ms, respectively. As a result, we expect some non-negligible error simply due to dephasing of the qubits over the gate interaction. Since Z-rotations commute with the gate interaction, the Walsh modulation also acts as a spin-echo on the qubits, thereby suppressing low-frequency dephasing processes such as slow magnetic field drifts. We model qubit dephasing

into our gate simulation using the Lindblad operator $\hat{L} = \hat{\sigma}_z^i$, i.e., the Pauli Z-operator acting on the i^{th} qubit, with an associated relaxation rate $\gamma = 1/2\tau_i$, where τ_i is the coherence time of the i^{th} qubit.

By modelling qubit dephasing in this way, we are assuming a white noise source for the dephasing process. This assumption is not particularly great; in reality, qubit dephasing occurs due to laser phase noise and magnetic field noise, which will generally exhibit a non-white noise spectrum. The effectiveness of the Walsh modulation for suppressing qubit dephasing strongly depends on the noise spectrum of the dephasing processes; this modulation is least effective for a white-noise spectrum due to the significant presence of high-frequency noise. We therefore overestimate the qubit coherence times to balance the pessimistic white noise model. For both modules, we choose coherence times of $\tau_{\text{N}} = 14\,\text{ms}$ and $\tau_{\text{X}} = 7\,\text{ms}$ for the network and auxiliary qubits, respectively. Simulations using a more realistic noise models could be performed using a stochastic solver, however this was not explored in this work.

Heating of the axial OOP mode

The OOP mode is cooled to near its ground state at the start of each experimental sequence using Doppler cooling, followed by EIT cooling. In Alice (Bob), we achieve a mean occupation number of $\bar{n} \approx 0.15$ ($\bar{n} \approx 0.1$). While the geometric phase gate is nominally insensitive to the initial state of the motional mode, this is only true in the Lambe-Dicke regime; deviations from the ground motional state will result in dynamics beyond the Lambe-Dicke regime, which may lead to gate errors. It is clear from the simulations presented in Figure 5.6, however, that this effect should not be significant in either module.

While the gate is resilient to non-zero initial temperatures of the OOP mode, heating of the OOP mode over the gate interaction leads to incoherent evolution of the motional state [126]. Since the motional mode is then not returned exactly to its initial state, we are left with a residual spin-motion entanglement, and thus

a departure from the desired unitary gate dynamics. The mass asymmetry of the mixed-species crystal results in a significantly higher heating rate of the OOP motion than in a single-species crystal where the symmetry suppresses coupling of the OOP mode to electric field noise.

We frequently measured heating rates for the axial OOP mode of the Sr^+-Ca^+ crystal in each module, typically observing heating rates of $\dot{\bar{n}}_{\mathrm{OOP}} \approx 350\,\mathrm{s}^{-1}$ and $\dot{\bar{n}}_{\mathrm{OOP}} \approx 80\,\mathrm{s}^{-1}$, in Alice and Bob, respectively. We model the effects of heating during our gate by including the Lindblad operators $\hat{L}_1 = \hat{a}$ and $\hat{L}_2 = \hat{a}^\dagger$, both with the relaxation rates $\gamma_i = \sqrt{\dot{\bar{n}}_{\mathrm{OOP}}}$, into our gate simulations. The estimated contribution to the total gate error as a function of gate duration is shown in Figure 5.6.

Motional dephasing of the OOP mode

Another error mechanism related to the OOP motion is motional dephasing. Just as qubit dephasing can be viewed in terms of fluctuations in the frequency of the qubit transition, motional dephasing can be viewed in terms of a fluctuation of the motional mode frequency. We model motional dephasing using the Lindblad operator $\hat{L} = \hat{a}^\dagger\hat{a}$ with the associated relaxation rate $\gamma = \sqrt{2/\tau_{\mathrm{OOP}}}$, where τ_{OOP} is the motional coherence time of the OOP mode of motion. As outlined in Chapter 4, by performing a Ramsey experiment between the $|0\rangle$ and $|1\rangle$ Fock states of the OOP motion, we estimated a motional coherence time of $\approx 3.4\,\mathrm{ms}$ in Alice. While we have not yet performed the same experiment in Bob, we anticipate that the motional dephasing is a result of technical imperfections, and thus we assume a similar motional coherence time as Alice.

The simulations of the gate error as a function of gate duration, shown in Figure 5.6, indicate that this mechanism is a significant contributor of error in both systems. As such, to push the gate fidelities, further work should be performed to characterise motional dephasing in both systems, and sources of such dephasing, such as noise arising from ground loops, should be identified and eliminated as far as possible.

Excitation of the axial IP mode

From Section 2.2.1, the frequencies of the IP and OOP motional modes for the Sr^+-Ca^+ crystals can be calculated from the single-Sr axial motional frequency, ω_z, as

$$\omega_{\mathrm{IP}} = \omega_z \sqrt{\frac{1+\mu-\sqrt{1+\mu\left(\mu-1\right)}}{\mu}}$$
$$\omega_{\mathrm{OOP}} = \omega_z \sqrt{\frac{1+\mu+\sqrt{1+\mu\left(\mu-1\right)}}{\mu}},$$

where $\mu = m_{\mathrm{Ca}}/m_{\mathrm{Sr}}$. Thus, we find that for $^{43}Ca^+$ and $^{88}Sr^+$, we obtain

$$\frac{\omega_{\mathrm{OOP}}}{\omega_{\mathrm{IP}}} \approx 1.945, \tag{5.7}$$

which is very close to 2. As a result, the second harmonic of the IP mode is close to the first harmonic of the OOP mode; the only way to increase this separation is to increase the axial confinement, however we are currently working at the limit of axial confinement that we can achieve in each module.

Due to the proximity of the second harmonic of the IP mode to the OOP mode, the spin-dependent force will couple to the IP motion. This coupling is described by the Hamiltonian

$$\hat{H}_{\mathrm{IP}} = -\sum_{i=1}^{N}\sum_{s=0,1}\frac{\hbar\Omega_{i,s}\eta_{i,\mathrm{IP}}^2}{4}e^{i\Delta\phi_i}e^{-i\delta_{\mathrm{IP}}t}\left|s_i\right\rangle\left\langle s_i\right|\otimes\left(\hat{a}_{\mathrm{IP}}^{\dagger}\right)^2 + \mathrm{h.c.}, \tag{5.8}$$

where $\eta_{i,\mathrm{IP}}$ is the Lamb-Dicke parameter for the coupling of the Raman beams to the component of the i^{th} ion's motion in the IP mode, $\hat{a}_{\mathrm{IP}}^{\dagger}$ is the creation operator the IP motional mode, and $\delta_{\mathrm{IP}} = \Delta\omega - 2\omega_{\mathrm{IP}}$ is the detuning of the travelling-standing wave from the second harmonic of the IP mode frequency. Note that the tensor product with the identity for the OOP mode Fock space is not shown. This Hamiltonian describes the SDF off-resonantly driving the second harmonic of the IP mode, leading to a spin-dependent interaction that leaves the spin states entangled with the state

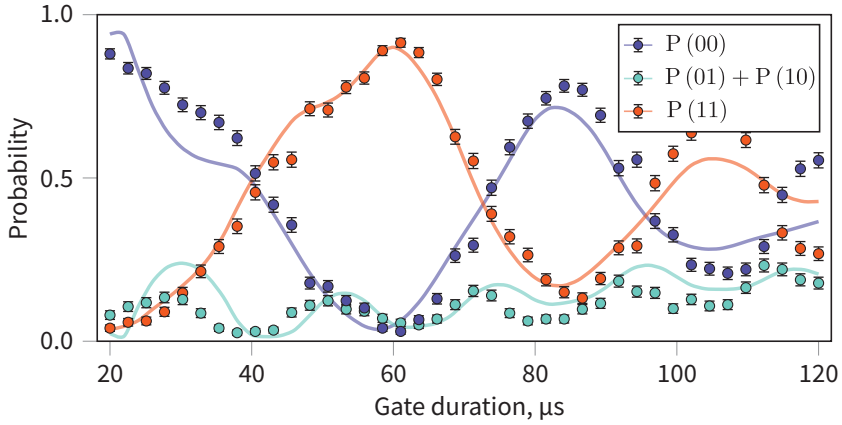

Figure 5.7: Simulated (solid lines) and experimental data from Alice (circles) showing gate dynamics as a function of gate duration, highlighting the characteristic effects of the off-resonant coupling to the second harmonic of the IP mode. The SDF couplings used here are larger than those typically employed in this thesis, resulting in a crossing of the P(00) and P(11) populations at $\sim 40\,\mu s$. Additionally, suboptimal cooling of the IP mode – assumed to be $\bar{n} \approx 3.0$ in the simulation – led to significant modulation of the gate dynamics due to the IP coupling.

of the IP motional mode.

We model the effect of this coupling by combining the Hamiltonian in Eq. (5.8) with the Hamiltonian for the desired spin-dependent force on the OOP mode, Eq. (5.1). Figure 5.7 compares experimental data taken for a gate duration scan to the simulated dynamics including both of these terms. For this scan, we increased the power as far as possible[4], achieving a crossing between the P(00) and P(11) populations around 40 µs. In the simulation, we included all the error sources described in this section, in order to achieve a reasonable match between simulation and experiment[5]. We observe particular gate durations where the odd-parity population, $P_{01} + P_{10}$, deviates significantly from zero, indicating non-unitary evolution and thus degradation of the entangling gate fidelity.

Empirically, we have found that the effect of this off-resonant excitation increases with the temperature of the IP mode. Interestingly, we found that this effect does not have a significant direct dependence on the heating rate of the IP mode. However, a

[4]The experimental data here was taken using the Zeeman qubit in Sr^+, rather than the network qubit.

[5]Note that due to the time taken to simulate this scan, fitting the simulation parameters to the data was not attempted.

high heating rate of the IP mode makes it challenging to cool the mode to its ground state, hence indirectly affecting the gate dynamics. We observe heating rates for the IP mode of $\dot{\bar{n}}_{\mathrm{IP}} \approx 2700\,\mathrm{s}^{-1}$ and $\dot{\bar{n}}_{\mathrm{IP}} \approx 500\,\mathrm{s}^{-1}$ in Alice and Bob, respectively. As a result, we typically achieve initial temperatures of $\bar{n}_{\mathrm{IP}} \approx 2.0$ and $\bar{n}_{\mathrm{IP}} \approx 0.8$ for Alice and Bob, respectively, using EIT cooling.

Using this model for the off-resonant excitation of the IP mode, we simulate the predicted effects on the gate fidelity as a function of gate duration, shown in Figure 5.6. We observe a similar oscillating effect in the gate error that we observed in the $P_{01} + P_{10}$ population. Interestingly, we find that there are gate durations that we can choose at which the error arising from this error mechanism is very small.

Summary

Here, we have provided an overview of our current understanding of the error mechanisms affecting our mixed-species entangling gates. These gates are a critical tool for our mixed-species quantum networking experiments, presented in Chapters 6 and 7, and thus understanding and mitigating these mechanisms is crucial. The simulated error contributions as a function of gate duration, depicted in Figure 5.6, are broadly consistent with the observed performance of the entangling gates. This suggests that, to further reduce gate error, efforts should focus on minimising qubit dephasing and motional dephasing. For future work, it will be important to invest time in characterising and suppressing these sources of noise.

5.2.6 Characterisation of the mixed-species iSWAP gate

The ability to coherently transfer quantum information from Sr^+ to Ca^+ is an important resource in our trapped-ion modules, enabling the remote entanglement generated across the quantum network to be transferred from the Sr^+ network qubit to the long-lived Ca^+ circuit qubit.

The obvious choice for implementing this transfer is the SWAP gate, which exchanges the quantum states of the participating qubits. The SWAP gate requires

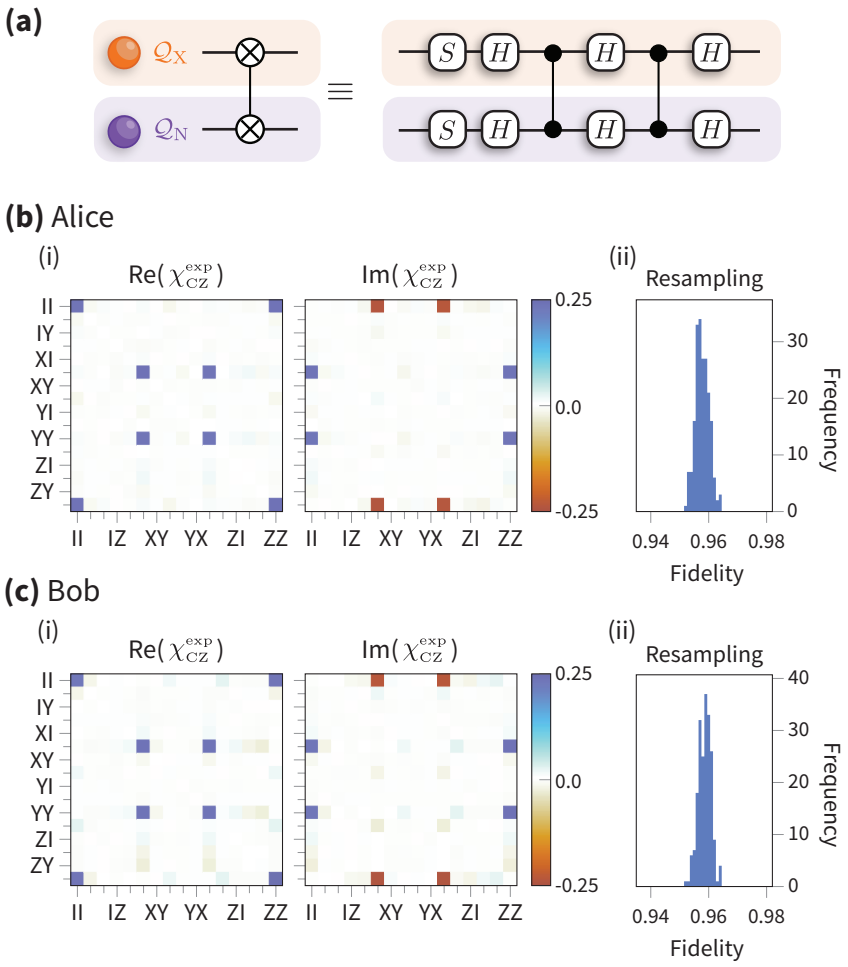

Figure 5.8: **(a)** Circuit for implementing the mixed-species iSWAP gate, compiled with two CZ gates. **(b)**(i) and **(c)**(i) show the process matrices for the mixed-species iSWAP gates in Alice and Bob, respectively, reconstructed from QPT using a total of 32 400 measurements. From the process matrices, we observe CZ gate fidelities of 95.9(2) % and 96.0(2) % for Alice and Bob, respectively. **(b)**(ii) and **(c)**(ii) show the histograms of the gate fidelities of process matrices reconstructed from resampled datasets.

3 instances of the CZ-gate and therefore can build up a significant error contribution, particular when the transfer is implemented in each module. However, in typical scenarios, we are only looking to map the state of the network qubit onto one of the Ca^+ qubits, and thus the Ca^+ ion will initially be in a known state, e.g., $|0_X\rangle$. Therefore, we only need to implement the iSWAP gate, which requires only 2 instances of the CZ-gate.

The iSWAP circuit is shown in Figure 5.8(a), and implements the unitary

$$U_{\text{iSWAP}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}.$$

Thus, if we let the Sr^+ network qubit be in some arbitrary state $|\psi_N\rangle$ and the Ca^+ auxiliary qubit be prepared in the state $|0_X\rangle$, the iSWAP gate therefore implements the mapping

$$|\psi_N\rangle |0_X\rangle \xrightarrow{\text{iSWAP}} |0_N\rangle \left(\hat{S}^\dagger |\psi_X\rangle\right) \in \mathcal{Q}_N \otimes \mathcal{Q}_X. \tag{5.9}$$

We can then trivially apply an S-gate to the auxiliary qubit, completing the transfer of the state $|\psi\rangle$ from the Sr^+ ion to the Ca^+ ion.

As with our mixed-species CZ-gate, we characterise the iSWAP gate using QPT with SPAM errors included into the tomography, see Section 3.3. We perform the tomography with a total of 32 400 measurements. The process matrices for the mixed-species iSWAP-gates from each module are shown in Figure 5.8. Compared to the ideal iSWAP-gate, we measure average gate fidelities of 95.9(2) % and 96.0(2) % for Alice and Bob, respectively.

Network to auxiliary qubit state transfer

An important metric is the fidelity with which we transfer quantum information from the network qubit to the auxiliary qubit. The circuit implementing this transfer is shown in Figure 5.9(a)(i), and the graphical calculation of the superoperator representation of the transfer process [54] is shown in Figure 5.9(a)(ii). Using this graphical calculus, we can predict the action of the iSWAP-based state transfer, taking into account experimental imperfections. Such imperfections include imperfect state preparation of the auxiliary qubit and a noisy iSWAP gate.

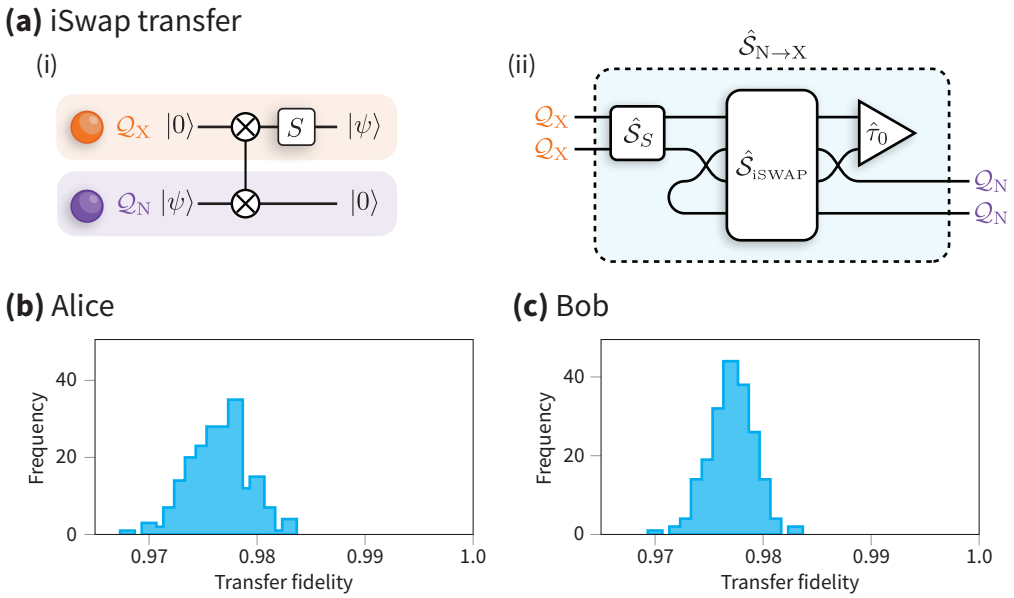

Figure 5.9: **(a)** (i) Quantum circuit for the state transfer, in which the state $|\psi\rangle \in \mathcal{Q}_N$ is transferred to the auxiliary qubit, initially in the state $|0\rangle$, using an iSWAP gate. (ii) Graphical calculation, using the calculus proposed by Wood, Biamonte, and Cory [54], of the superoperator, $\hat{\mathcal{S}}_{N\to X}$, representing the action of implementing the state transfer. The noisy iSWAP process and the imperfect state preparation of the auxiliary qubit are represented by the superoperator $\hat{\mathcal{S}}_{iSWAP}$ and the state $\hat{\tau}_0$, respectively. The superoperator $\hat{\mathcal{S}}_S$ represents an ideal S-gate. **(b)** and **(c)** show histograms for the transfer fidelities, calculated as in Eq. (5.15) using iSWAP superoperators obtained from resampling, for Alice and Bob, respectively. From these histograms, we obtain average gate fidelities of 97.8(3) % and 97.9(2) % for Alice and Bob, respectively.

Suppose imperfect state preparation of the auxiliary qubit yields the state

$$\hat{\tau}_0 = (1-\epsilon)\,|0\rangle\langle 0| + \epsilon\,|1\rangle\langle 1| \in L(\mathcal{Q}_X)\,. \tag{5.10}$$

Additionally, let the noisy iSWAP process be represented by the superoperator $\hat{\mathcal{S}}_{iSWAP}$. The transfer process, described by

$$\mathcal{E}_{N\to X} : L(\mathcal{Q}_N) \to L(\mathcal{Q}_X)\,. \tag{5.11}$$

is represented by the superoperator $\hat{\mathcal{S}}_{N\to X}$. This superoperator is calculated according to Figure 5.9(a)(ii), yielding

$$\hat{\mathcal{S}}_{N\to X} = \hat{\mathcal{S}}_S \left(\hat{\mathbb{I}}^{\otimes 2}_{\mathcal{Q}_X} \otimes \langle\langle \hat{\mathbb{I}}_{\mathcal{Q}_N}|\right) \hat{\mathcal{V}}_2 \hat{\mathcal{S}}_{iSWAP} \hat{\mathcal{V}}_2^\dagger \left(|\hat{\tau}_0\rangle\rangle \otimes \hat{\mathbb{I}}^{\otimes 2}_{\mathcal{Q}_N}\right), \tag{5.12}$$

where $\hat{\mathcal{V}}_2$ is the unravelling operation, as in Wood, Biamonte, and Cory [54]. The action of the unravelling operation on the iSWAP superoperator,

$$\hat{\mathcal{V}}_2 \hat{\mathcal{S}}_{\mathrm{iSWAP}} \hat{\mathcal{V}}_2^{\dagger}, \tag{5.13}$$

can be implemented by swapping the ordering of the second and third subspaces[6].

The ideal transfer process is effectively the identity operation

$$\mathcal{E}_{\mathrm{N}\to\mathrm{X}} \left(\hat{\rho} \in L\left(\mathcal{Q}_{\mathrm{N}}\right)\right) = \hat{\rho} \in L\left(\mathcal{Q}_{\mathrm{X}}\right), \tag{5.14}$$

and thus we can calculate the transfer fidelity as in Eq. (3.54),

$$\mathcal{F}_{\mathrm{N}\to\mathrm{X}} = \frac{d + \mathrm{tr}\left[\hat{\mathcal{S}}_{\mathrm{N}\to\mathrm{X}}\right]}{d\left(d+1\right)}. \tag{5.15}$$

Histograms showing the spread of transfer fidelities, calculated from resampled iSWAP gates, are shown for Alice and Bob in Figure 5.9(b) and (c), respectively. We infer transfer fidelities of 97.8(3) % and 97.9(2) % for Alice and Bob, respectively.

Error-detected network to auxiliary qubit state transfer

Recalling Eq. (5.9), since the auxiliary qubit was (nominally) initially prepared in the computational state $|0_{\mathrm{X}}\rangle$, the network qubit will (nominally) be left in the same state, $|0_{\mathrm{N}}\rangle$. As a result of errors in the iSWAP gate or imperfect state preparation of the auxiliary qubit, the network qubit will occasionally be left in the "wrong" state. Therefore, following the iSWAP-gate, we perform a mid-circuit measurement of the network qubit; if we measure the network qubit in the state $|0_{\mathrm{N}}\rangle$, we have projected the system into the $|0_{\mathrm{N}}\rangle\langle 0_{\mathrm{N}}| \otimes \hat{\mathbb{I}}_{\mathrm{X}}$ subspace, and thus have effectively projected out part of the iSWAP gate error. If the measurement outcome yields the state $|1_{\mathrm{N}}\rangle$, we simply restart the protocol. We can perform this measurement and decision branching in real-time.

[6] In practice, this can be implemented computationally using, e.g., the `permutesystems` function from QuantumOptics.jl [127].

The circuit for this error-detected iSWAP transfer is shown in Figure 5.10(a)(i), and the graphical calculation of the superoperator representation of the transfer process is shown in Figure 5.10(a)(ii). As before, we can use this calculation to predict the action of the error-detected transfer in the presence of experimental imperfections, such as imperfect state preparation of the auxiliary qubit, noise in the iSWAP gate, and imperfect measurement of the network qubit.

As in the case of the non-error-detected transfer, we assume the imperfect state preparation of the auxiliary qubit is described by the state in Eq. (5.10) and that the noisy iSWAP gate is represented by the superoperator $\hat{\mathcal{S}}_{\text{iSWAP}}$. After the transfer, we make a measurement of the network qubit, and abort the experiment if we do not observe the qubit to be in the state $|0\rangle$. The POVM corresponding to a measurement of the network qubit yielding the outcome $|0\rangle$ is given by

$$\hat{M}_0 = (1-\epsilon_0)\,|0\rangle\langle 0| + \epsilon_1\,|1\rangle\langle 1|\,. \tag{5.16}$$

We calculate the superoperator representing the error-detected transfer process according to Figure 5.10, obtaining

$$\hat{\mathcal{S}}'_{\text{N}\to\text{X}} = \hat{\mathcal{S}}_S\left(\hat{\mathbb{I}}^{\otimes 2}_{\mathcal{Q}_\text{X}} \otimes \langle\!\langle \hat{M}_0|\right)\hat{\mathcal{V}}_2\hat{\mathcal{S}}_{\text{iSWAP}}\hat{\mathcal{V}}_2^\dagger\left(|\hat{\tau}_0\rangle\!\rangle \otimes \hat{\mathbb{I}}^{\otimes 2}_{\mathcal{Q}_\text{N}}\right). \tag{5.17}$$

Note that since we have introduced an abort condition, the calculated process is no longer trace-preserving. For example, suppose the network qubit is prepared in some state $\hat{\rho}$ that we would like to transfer to the auxiliary qubit, then the probability that an error is detected, denoted p, can be calculated as

$$1 - p = \text{tr}\left[\mathcal{E}'_{\text{N}\to\text{X}}\left(\hat{\rho}\right)\right]. \tag{5.18}$$

Note that, in general, this probability depends on the input state, $\hat{\rho}$, and thus it is not so straightforward to normalise the process[7].

Due to the difficulty in normalising the process, we calculate the average gate

[7] If the reader knows of a good way to normalise such processes, please do let me know!

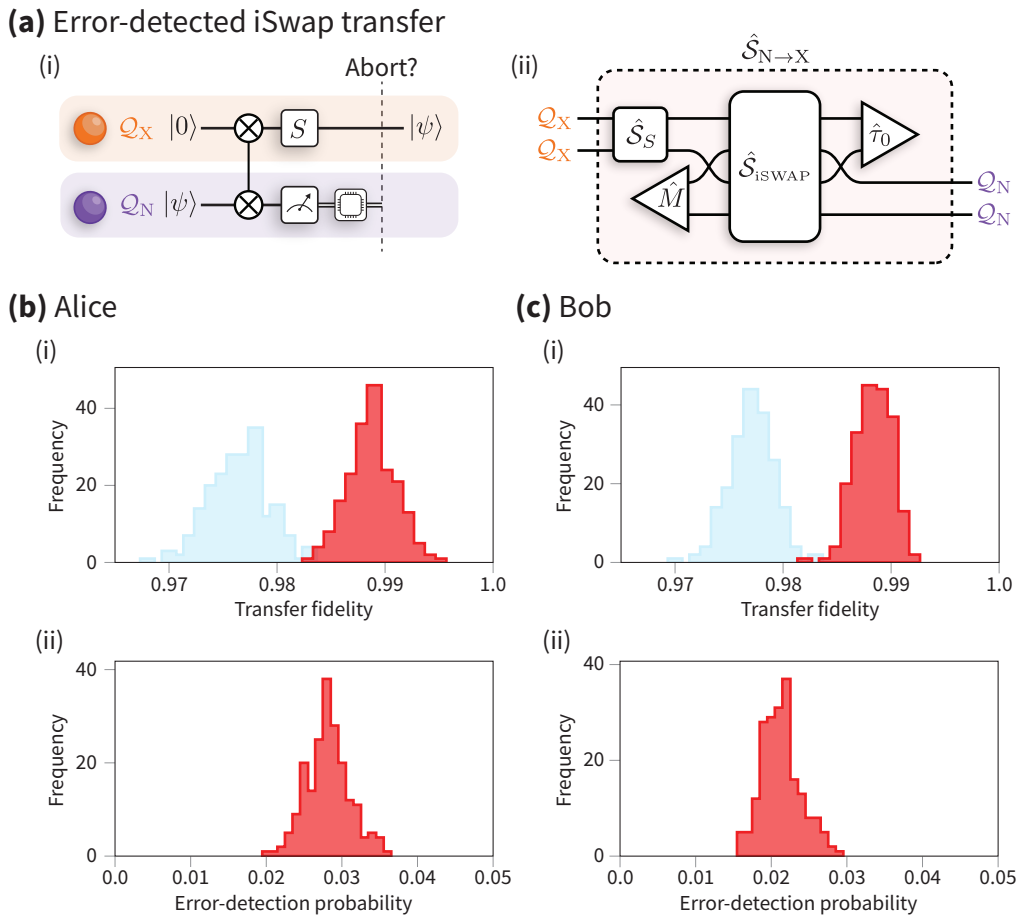

Figure 5.10: **(a)** (i) Quantum circuit for the error-detected state transfer. (ii) Graphical calculation of the superoperator, $\hat{\mathcal{S}}'_{N\to X}$, representing the action of implementing the error-detected state transfer. The measurement of the network qubit is represented by the POVM $\hat{M}_0$ **(b)**(i) and **(c)**(i) show histograms for the error-detected transfer fidelities (red), calculated as in Eq. (5.19) using iSWAP superoperators obtained from resampling, for Alice and Bob, respectively. From these histograms, we obtain average gate fidelities of 99.0(2) % and 99.0(2) % for Alice and Bob, respectively. The histograms for the non-error-detected transfer fidelities (blue) are also shown for comparison. **(b)**(ii) and **(c)**(ii) show histograms for the error-detection probabilities, calculated as in Eq. (5.20) using iSWAP superoperators obtained from resampling, for Alice and Bob, respectively. From these histograms, we obtain average error detection probabilities of 2.2(3) % and 2.8(3) % for Alice and Bob, respectively.

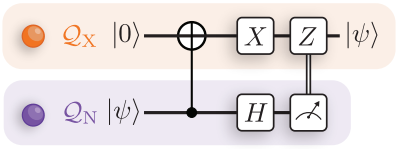

Figure 5.11: Quantum circuit for the CNOT-based state transfer. The state $|\psi\rangle \in \mathcal{Q}_N$ is transferred from the network qubit to the auxiliary qubit – initially prepared in $|0\rangle$ – using a single CNOT gate. A mid-circuit measurement of the network qubit then conditions a single-qubit rotation on the auxiliary qubit, completing the state transfer.

fidelity for the error-detected transfer using a Monte-Carlo approach to calculating the integral in Eq. (3.53). This is performed by sampling N input states $|\psi_i\rangle$ from the Haar measure and computing the average transfer fidelity as

$$\bar{\mathcal{F}}_{N\to X} = \frac{1}{N}\sum_{i=1}^{N} \frac{\langle\psi_i|\,\mathcal{E}'_{N\to X}\left(|\psi_i\rangle\langle\psi_i|\right)|\psi_i\rangle}{\mathrm{tr}\left[\mathcal{E}'_{N\to X}\left(|\psi_i\rangle\langle\psi_i|\right)\right]}. \tag{5.19}$$

We calculate the average error-detection probability, $\bar{p}$, in similar fashion, as

$$\bar{p} = 1 - \frac{1}{N}\sum_{i=1}^{N} \mathrm{tr}\left[\mathcal{E}'_{N\to X}\left(|\psi_i\rangle\langle\psi_i|\right)\right]. \tag{5.20}$$

Histograms showing the spread of average transfer fidelities and error-detection probabilities, as calculated from iSWAP gates reconstructed from resampled datasets, are shown in Figure 5.10(a) and (b) for Alice and Bob, respectively. The results for for Alice (Bob) indicate an average transfer fidelity of $\bar{\mathcal{F}}_{N\to X} = 99.0(2)\,\%$ ($\bar{\mathcal{F}}_{N\to X} = 99.0(2)\,\%$) and an average error-detection probability of $\bar{p} = 2.2(3)\,\%$ ($\bar{p} = 2.8(3)\,\%$).

Alternative state transfer protocol

An alternative method for transferring quantum information from the network qubit to the auxiliary qubit, explored by D. P. Nadlinger and illustrated in Figure 5.11, employs a single CNOT gate followed by a mid-circuit measurement of the network qubit to condition a local rotation on the auxiliary qubit. This scheme uses fewer gates than the iSWAP-based transfer and has no abort-on-failure condition. How-

ever, process tomography of this circuit yielded a transfer fidelity of 98.5(2) %, which is lower than that observed for the error-detected iSWAP circuit.

5.3 Hyperfine transfer

Since the circuit qubit does not participate in the mixed-species gate, the gate interaction is performed on the network and auxiliary qubits. Consequently, we require the ability to coherently inter-convert between the circuit and auxiliary qubit before and after the local operations. This mapping is performed using the Raman beams to coherently address the transitions within the ground hyperfine manifold of Ca^+. The pulse sequence we use to perform this transfer is outlined in Figure 5.12.

The transfer of the circuit qubit to the auxiliary qubit, depicted in Figure 5.12(a), begins with the mapping of the state $|0_C\rangle$ to the state $|0_X\rangle$. However, due to the near degeneracy of the transition $\mathcal{T}_0 : |0_C\rangle \leftrightarrow |F = 3, M_F = +1\rangle$ and the transition $\mathcal{T}_1 : |1_C\rangle \leftrightarrow |F = 4, M_F = +1\rangle$ (see Figure 5.12(c)(i)), separated by only $\Delta f \sim 15\,\text{kHz}$, it is not possible to map the $|0_C\rangle$ state out of the circuit qubit without off-resonantly driving population out of the $|1_C\rangle$ state. We suppress this off-resonant excitation using a composite pulse sequence, shown in Figure 5.12(c)(i), comprising three pulses resonant with the $\mathcal{T}_0$ transition, with pulse durations equal to the 2π-time of the $\mathcal{T}_1$ transition, and phases optimised to minimise the off-resonant excitation. This pulse sequence allows us to simultaneously perform a π-pulse on the $\mathcal{T}_0$ transition and the identity on the off-resonantly-driven $\mathcal{T}_1$ transition. Raman π-pulses are then used to complete the mapping to the $|0_X\rangle$ state. Another sequence of Raman π-pulses coherently maps $|1_C\rangle \rightarrow |1_X\rangle$, thereby completing the transfer of the circuit qubit to the auxiliary qubit, $\mathcal{Q}_C \rightarrow \mathcal{Q}_X$. To implement the mapping $\mathcal{Q}_X \rightarrow \mathcal{Q}_C$, depicted in Figure 5.12(b), the same pulse sequence is applied in reverse.

This off-resonant suppression pulse sequence, developed by Dr. P. Drmota, is an improvement over the previously used Ramsey-like sequence [123, 128]. Since the pulse sequence utilises the difference in Rabi frequencies of the two near-degenerate transitions, rather than in the Ramsey-like case, where the duration of the sequence

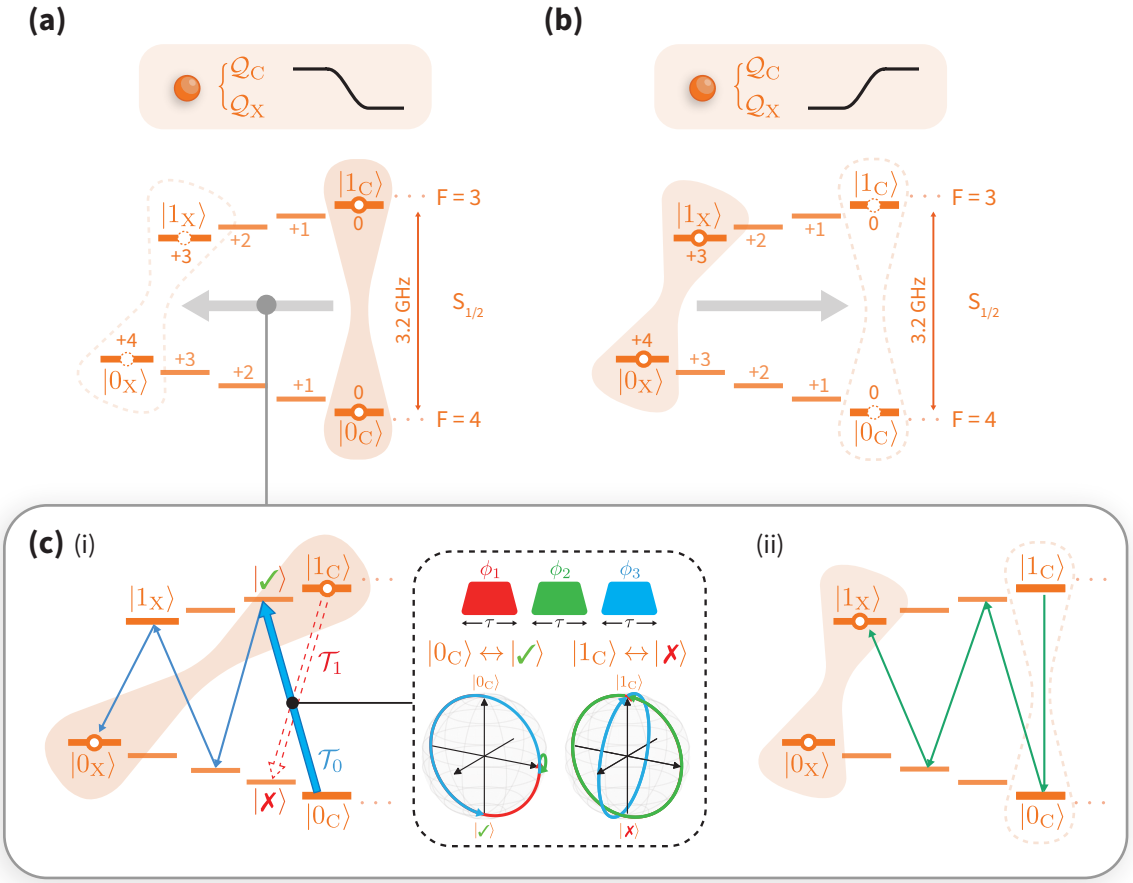

Figure 5.12: **(a)** Circuit element and level diagram depicting the coherent transfer of quantum information from the $\mathcal{Q}_C$ qubit to the $\mathcal{Q}_X$ qubit. **(b)** The inverse transfer is implemented by performing the same steps in reverse. **(c)** The $\mathcal{Q}_C \to \mathcal{Q}_X$ transfer pulse sequence comprises two steps. (i) The first step maps the state $|0_C\rangle$ to $|0_X\rangle$. Due to the near-degeneracy of the intended transition $\mathcal{T}_0 : |0_C\rangle \leftrightarrow |✓\rangle$ (thick blue arrow) and the unwanted transition $\mathcal{T}_1 : |1_C\rangle \leftrightarrow |✗\rangle$ (red dashed arrow), separated by only $\sim 15\,\text{kHz}$, we employ a composite pulse sequence to suppress off-resonant coupling to the $\mathcal{T}_1$ transition. The composite pulse sequence, shown in the dashed box, comprises 3 pulses of duration τ resonant with the $\mathcal{T}_0$ transition with differing phases ϕ_i. The pulse duration, τ, is equal to the 2π-time of the $\mathcal{T}_1$ transition, $\phi_1 = \phi_3 = 0$, and $\phi_2 \approx 2\pi \times 0.231$ is optimised experimentally. The subsequent transfer pulses (thin blue arrows) are π-pulses on the relevant transitions. This sequence therefore performs the mapping $|0_C\rangle \to |0_X\rangle$, while leaving the state $|1_C\rangle$ unaffected. (ii) The second step comprises a sequence of π-pulses which maps $|1_C\rangle \to |1_X\rangle$. This completes the coherent transfer $\mathcal{Q}_C \to \mathcal{Q}_X$.

is Fourier-limited by the spectral separation of the transitions. For example, in Drmota *et al.* [123], the Ramsey-like sequence had a duration of $2/\Delta f \sim 157\,\mu s$; the full sequence also required the time for the additional remaining π-pulses. In the Rabi-like sequence used here, the off-resonant suppression has a duration of three 2π-durations of the $\mathcal{T}_1$ transition. Thus, this duration can be made arbitrarily short by simply increasing the Raman power used to drive the transitions, thus increasing the Rabi frequency mis-match between the near-degenerate transitions. For comparison, in Alice, the Ramsey-like sequence had a duration of $\sim 157\,\mu s$, while the Rabi-like sequence had a duration of $\sim 45\,\mu s$. The advantage of the reduced duration is that it reduces the exposure to magnetic field fluctuations and hence reduces the transfer error.

5.3.1 Characterisation

We characterise our $\mathcal{Q}_C \leftrightarrow \mathcal{Q}_X$ transfer sequence by performing a modification of single-qubit RBM, in which we alternate Clifford operations on the $\mathcal{Q}_C$ and $\mathcal{Q}_X$ qubits, as illustrated in Figure 5.13(a). We assume that (i) the single-qubit gate errors for the $\mathcal{Q}_C$ and $\mathcal{Q}_X$ qubits are negligible compared to the $\mathcal{Q}_C \leftrightarrow \mathcal{Q}_X$ transfer infidelity (we typically observe single-qubit Clifford errors $\sim 1 \times 10^{-4}$ for the Ca^+ hyperfine qubits), and (ii) the fidelity of the transfer $\mathcal{Q}_C \rightarrow \mathcal{Q}_X$ is similar to $\mathcal{Q}_X \rightarrow \mathcal{Q}_C$. We therefore we model the survival probability as

$$S(m) = \frac{1}{2} + Bp^m, \tag{5.21}$$

where m is the number of hyperfine transfers, B accounts for SPAM error offsets, and p is the depolarising probability for a single transfer, related to the error per transfer as

$$\epsilon_{C\leftrightarrow X} = \frac{1-p}{2}. \tag{5.22}$$

The RBM results are shown in Figure 5.13(b); we measure an error per transfer of $3.8(2) \times 10^{-3}$ ($2.6(1) \times 10^{-3}$) for Alice (Bob). Each transfer comprises 7 Raman

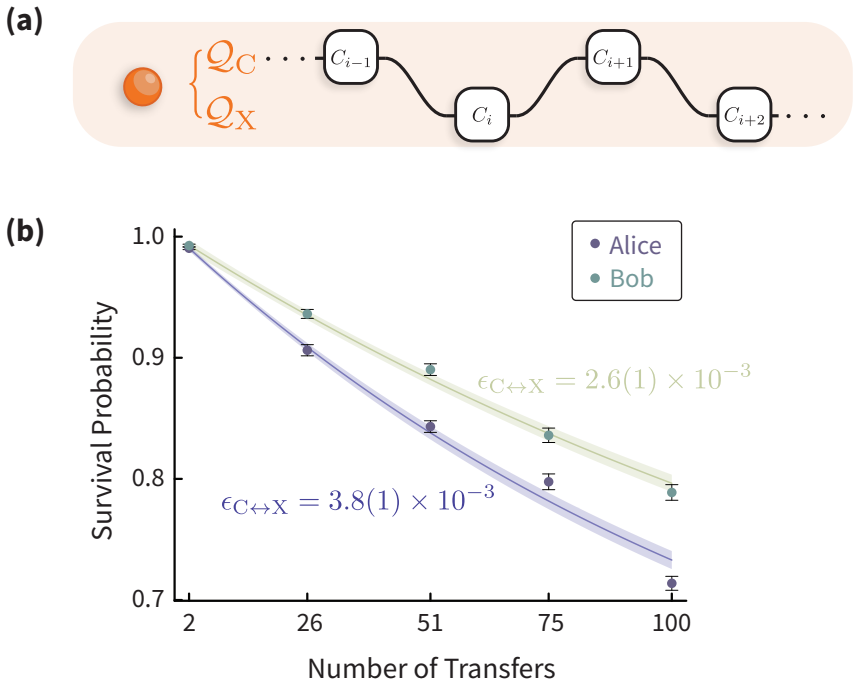

Figure 5.13: The performance of the transfer sequence is characterised using a modified version of RBM in which we alternately perform Clifford operations on the $\mathcal{Q}_C$ and $\mathcal{Q}_X$ qubits. By measuring the survival probability for different numbers of transfers, and neglecting the errors of the single-qubit gates C_i (which are $\sim 1 \times 10^{-4}$), we extract the error per transfer, $\epsilon_{C\leftrightarrow X}$, yielding $3.8(2) \times 10^{-3}$ and $2.6(1) \times 10^{-3}$ for Alice and Bob, respectively. All error bars indicate one standard deviation.

π-pulses on transitions within the hyperfine manifold, in addition to three 2π-pulses comprising the off-resonant suppression step. Assuming each π-pulse incurs an error of order $\sim 1 \times 10^{-4}$, and assuming that the contribution from each 2π-pulse is equivalent to four π-pulses (since these pulses drive two transitions simultaneously), we would expect an error per transfer of order $\sim 2 \times 10^{-3}$ from single-qubit gate errors alone. We attribute the additional error to a number of mechanisms. Firstly, during the transfer $\mathcal{Q}_X \to \mathcal{Q}_C$ ($\mathcal{Q}_C \to \mathcal{Q}_X$), the qubit acquires a phase $\phi_{X\to C}$ ($\phi_{C\to X}$). We calibrate these phases using a Bayesian phase estimation scheme[8], and apply a Z rotation to remove it. Typical values for the phase acquired over the $\mathcal{Q}_X \to \mathcal{Q}_C$ ($\mathcal{Q}_C \to \mathcal{Q}_X$) transfers are $\phi_{X\to C} = 0.565(5)$ turns ($\phi_{C\to X} = 0.346(5)$ turns) and $\phi_{X\to C} = 0.542(5)$ turns ($\phi_{C\to X} = 0.327(6)$ turns) for Alice and Bob, respectively.

[8]Note that we cannot adaptively vary M, so we only adaptively choose the next measurement phase.

Miscalibration of these phases contributes to the error per transfer. Secondly, we have observed drifts in the beam pointing of the Raman lasers which manifests as a drift in the gate durations. This causes an increase in the error contributions from the individual π-pulses, but also in the composite pulse sequence, which relies on the calibration of the 2π-time of the $\mathcal{T}_1$ transition.

5.4 Mid-circuit measurement

Mid-circuit measurements are, as their name suggests, measurements made within a circuit, rather than simply at the end of a circuit. The outcomes of these mid-circuit measurements can then be used in some later part of a circuit, for example, to condition a set of operations based on the outcome of the measurement, or to restart the circuit if the measurement indicated the presence of some un-correctable error. Mid-circuit measurements have applications in a wide range of quantum information processing applications: in error correction protocols, mid-circuit measurements of stabilizer qubits enable the detection and correction of errors occurring within a quantum computer [7, 8]; in teleportation protocols, the focus of Chapter 7, the classical exchange of mid-circuit measurement outcomes is essential, otherwise such protocols would be in violation of causality; and finally, we employ mid-circuit measurements for our $\mathcal{Q}_\mathrm{N} \rightarrow \mathcal{Q}_\mathrm{X}$ transfer, using a measurement to project out some of the iSWAP gate error, as outlined in Section 5.2.6.

One obviously desirable property of a mid-circuit measurement is that the measurement of one qubit does not degrade the encoded quantum information of any other qubit. However, isolating these qubits from the measurement can become challenging, regardless of the physical platform used to realise the quantum processor. As an example, consider a trapped-ion processor, in which a measurement is typically performed using a laser tuned close to a dipole transition, such that the presence or lack of fluorescence enables discrimination between qubit states. It would only require the unwanted scattering of a single photon by another ion in the processor to completely depolarise the state of that ion.

In our apparatus, we use two separate species of ion, Ca^+ and Sr^+. The large spectral separation of the dipole transitions of these ions – 393 nm and 397 nm for Ca^+, and 408 nm and 422 nm for Sr^+ – means that we can perform fluorescence detection on one species, and unwanted scattering in other species will be massively suppressed. A similar effect can also be achieved with a single species of ion using an *omg* architecture [129], however this is not pursued in our apparatus. In all circuits explored in this thesis, we employ mid-circuit measurement of the Sr^+ network qubits only.

6

Quantum Networking with Mixed-Species Modules

*

This chapter outlines the process by which remote entanglement between network qubits in separate modules is established and integrated with our mixed-species logic. First, the ion–photon interface – generating entanglement between a Sr^+ ion and a 422 nm photon – is presented, highlighting the Sr^+ ion's simple level structure and the use of photons as flying qubits. Next, the creation of remote entanglement between Sr^+ qubits in different modules is described. Building on the local operations from Chapter 5, complex networked states are then constructed: Sr^+–Ca^+ entanglement, Ca^+–Ca^+ entanglement, and mixed-species multi-partite Greenberger-Horne-Zeilinger (GHZ) states. Finally, quantum-memory capabilities are demonstrated by mapping Sr^+–Sr^+ entanglement into long-lived Ca^+ circuit qubits and observing its preservation for up to 10 s.

Our group has previously published extensive work on the generation of remote Sr^+–Sr^+ entanglement [91, 130], which was key to our demonstrations of device-independent QKD [131] and entanglement-enhanced clock comparisons [132].

The integration of a Ca^+ quantum memory qubit into one of our modules has been demonstrated in Refs. [56, 123] and enabled our recent demonstration of blind quantum computing [128]. In contrast, the generation and characterisation of mixed-species entanglement across separate modules is a novel contribution of this thesis and is currently being prepared for publication [133]. Specifically, I led the experimental demonstrations of remote entanglement between Sr^+ and Ca^+ qubits, between two Ca^+ qubits, and the generation of the mixed-species multi-partite GHZ states across the network. I also led the demonstration of long-lived remote entanglement by mapping Sr^+–Sr^+ entanglement into long-lived Ca^+ memory qubits.

6.1 Ion-photon entanglement

The ability to interface our trapped-ion modules with a photonic quantum network is a critical component of our apparatus. This interface is realised through the generation of entanglement between the "stationary" Sr^+ network qubits and the "flying" photonic qubits. In this section, we discuss the process by which we generate ion-photon entanglement (IPE) and the characterisation of this interface.

6.1.1 Generation of ion-photon entanglement

The process by which we generate entanglement between the Sr^+ network qubit and the polarisation degree of freedom of an 422 nm photon is depicted in Figure 6.1. This process is try-until-success, i.e., we repeatedly attempt to generate this entanglement until we detect a photon. Each entanglement attempt comprises (i) optical pumping of the Sr^+ ion into the lower ground Zeeman state, $\left|5s\ S_{1/2},\ m_J = -\frac{1}{2}\right\rangle$, using σ^--polarised 422 nm light, and a 1092 nm repumper, (ii) excitation of the Sr^+ ion to the $\left|5p\ P_{1/2},\ m_J = +\frac{1}{2}\right\rangle$ level using a ~ 5 ps pulse of σ^+-polarised 422 nm light, and (iii) rapid decay (~ 7 ns) to one of the two ground Zeeman states. Decay via the 1092 nm decay channel to the $D_{3/2}$ manifold may also occur with a branching ratio $\sim 5\,\%$; however, these photons will not be detected by our apparatus and thus

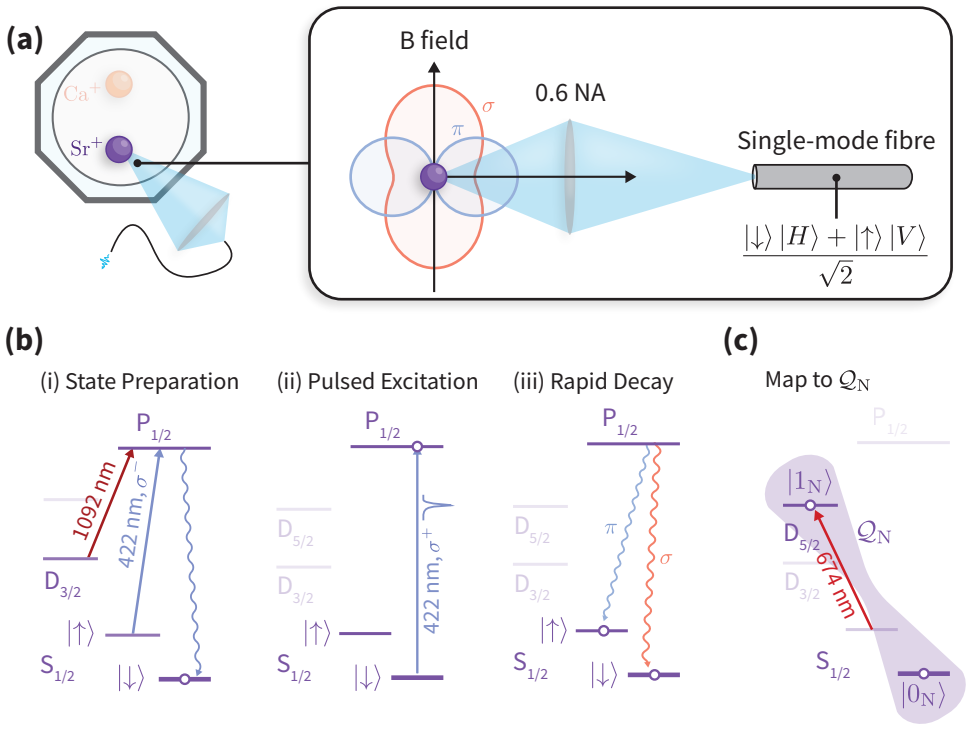

Figure 6.1: **(a)** Geometry of the ion-photon interface enabling entanglement between the Sr^+ ion and the polarisation of a spontaneously emitted 422 nm photon. A high-numerical aperture lens collects photons perpendicular to the quantisation axis and couples them into a single-mode, non-polarisation-maintaining fibre. **(b)** Energy level diagrams for a single excitation attempt. (i) A σ^--polarised 422 nm beam and 1092 nm repumper optically pump the Sr^+ ion into $|\downarrow\rangle$. (ii) A $\sim$5 ps σ^+-polarised 422 nm pulse excites the ion to the $P_{1/2}$ manifold. (iii) The ion decays to the ground Zeeman states via π and σ decay channels, entangling the photon polarisation with the ion's Zeeman states. **(c)** Upon photon detection, the $|\uparrow\rangle$ state is mapped to $|1_N\rangle$, transferring quantum information into the network qubit.

can be treated simply as an unsuccessful attempt – we henceforth ignore this decay channel.

The spontaneously emitted photons will be entangled with ground Zeeman state of the ion [130], and can be described by the state

$$\sqrt{\frac{2}{3}}\,|\downarrow\rangle\,|\sigma^+\rangle + \sqrt{\frac{1}{3}}\,|\uparrow\rangle\,|\pi\rangle\,, \tag{6.1}$$

where $|\downarrow\rangle = \left|5s\ \mathrm{S}_{1/2},\ m_J = -\frac{1}{2}\right\rangle$ and $|\uparrow\rangle = \left|5s\ \mathrm{S}_{1/2},\ m_J = +\frac{1}{2}\right\rangle$ denote the Zeeman states of the Sr^+ ion, $|\sigma^+\rangle$ and $|\pi\rangle$ denote the quantum states of the spontaneously emitted photon, and the amplitudes are given by the Clebsch-Gordon coefficients.

We collect the photons in free space using a high numerical aperture (0.6 NA) lens perpendicular to the quantisation axis. Exactly perpendicular to the quantisation axis, the electric fields of the π and σ^+ fields are orthogonal and thus the polarisation states of the photons, labelled $|H\rangle, |V\rangle \in \mathcal{Q}_{\text{ph}}$, are orthogonal. Additionally, along this axis, the π field has double the intensity of the σ^+ field, which exactly cancels the imbalanced amplitudes in Eq. (6.1). Thus, for the photons collected exactly perpendicular to the quantisation axis, we can write the entangled ion-photon state as the maximally entangled Φ^+ Bell state

$$\left|\Phi^+_{\text{IP}}\right\rangle = \frac{|\downarrow\rangle |H\rangle + |\uparrow\rangle |V\rangle}{\sqrt{2}}. \tag{6.2}$$

However, the high numerical aperture of the collection optics results in the collection of a relatively large amount of the emitted light off-axis, where the electric fields are no longer orthogonal. As a result, the polarisation states of the photons become non-orthogonal, $\langle H|V\rangle \neq 0$, and hence the ion-photon state in Eq. (6.2) no longer describes a maximally entangled Bell state. This polarisation mixing effect was one of the dominant error sources in early experiments demonstrating ion-photon entanglement [102], and the higher numerical-aperture used in our apparatus would result in an even larger effect.

We mitigate this polarisation mixing effect by coupling the collected light into a single-mode, non-polarisation-maintaining fibre. The single-mode fibre acts as a spatial mode filter, coupling only a single spatial mode of the emitted light into the fibre. The spatial mode of the fibre has two possible polarisations, which have orthogonal electric fields. Therefore, by coupling the emitted photons into this fibre, we preserve the orthogonality of the polarisation fields and cancel out the amplitude imbalance in Eq. (6.1), independent of the numerical aperture of the collection optics. This allows us to increase the numerical aperture of the lens while achieving, in principle, unit fidelity of the ion-photon state to the maximally entangled Bell state in Eq. (6.2). However, this improvement comes at a cost, as photons with a polarisation component along the fibre axis are filtered out and do not couple into

the single-mode fibre [130].

Once coupled into the single-mode fibre, the photon is brought to a polarisation analyser[1], shown in Figure 6.3(a), which enables the measurement of the polarisation state of the photon. A detection event, or detector "click", from one of four avalanche photodiodes (APDs) indicates the successful generation of IPE and the subsequent collection of the single-photon. By choosing different settings of the waveplates in the polarisation analyser, we can perform arbitrary unitary rotations of the photon's polarisation, and the combination of polarising beamsplitters (PBSs) and APDs provides a polarisation measurement in the computational basis, $\{|H\rangle, |V\rangle\}$. This enables polarisation measurements of the photon from a tomographically-complete set. Upon successful detection of the photonic part of the IPE, we map the ionic part from the ground Zeeman states, $\{|\downarrow\rangle, |\uparrow\rangle\}$, to the optical network qubit, $\mathcal{Q}_{\mathrm{N}} = |0_{\mathrm{N}}, 1_{\mathrm{N}}\rangle$, for any further manipulation.

6.1.2 Experimental implementation

Since the generation of IPE, followed by the collection and detection of the photon, is a probabilistic process, we repeatedly attempt to generate IPE until we successfully detect a photon. Each entanglement attempt, shown in Figure 6.2, has a duration of 1168 ns. Due to the density of RTIO events, the execution of this sequence is handled by the *Entangler*, a custom gateware extension to the Kasli – for more information about the Entangler, the reader is referred to the DPhil thesis by Drmota [56]. The Entangler has a set of timings which determine the on/off times of TTL outputs, which are then used to switch the 422 nm and 1092 nm state preparation pulses, and request a single pulse from the pulsed-laser. This framework allows us to perform state preparation, pulsed excitation, photon detection, and decide whether the attempt was successful or not, within the duration of a single entanglement attempt, as shown in Figure 6.2.

As shown in Figure 6.2, a significant fraction of the excitation attempt is con-

[1]This polarisation analyser is the same apparatus used to perform a projective Bell state measurement which creates the remote entanglement.

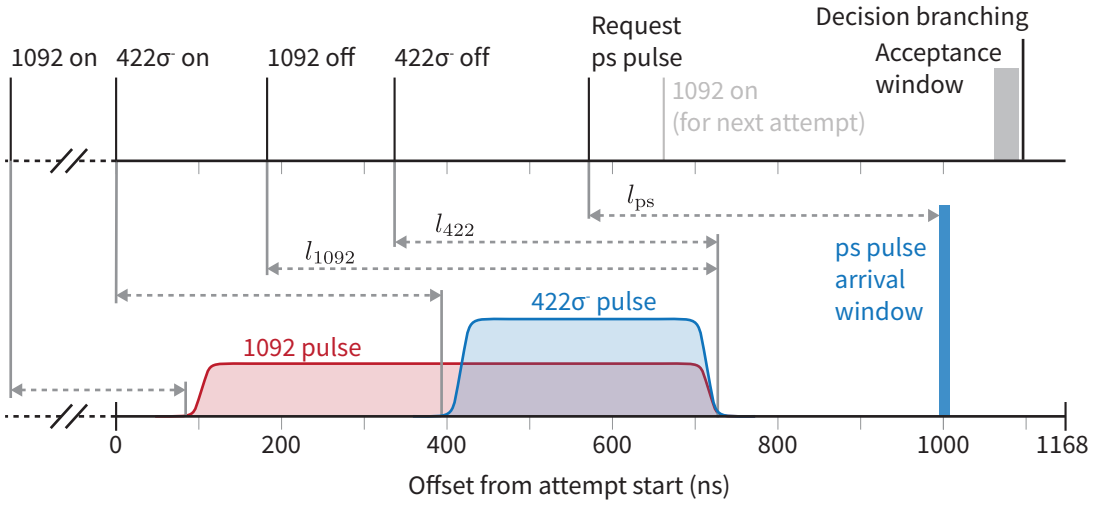

Figure 6.2: The lower timeline depicts the pulse sequence for a single entanglement generation attempt. Each attempt lasts 1168 ns and, as shown on the lower timeline, includes: $\sim$320 ns of state preparation using a σ^--polarised 422 nm pulse and a 1092 nm repumper, the request and arrival window of the picosecond pulse, a 30 ns photon detection window, and a decision step to terminate the loop upon successful detection. The upper timeline shows the entangler's event timings, corresponding to TTL signals that trigger the fast switching of the state preparation beams, and issue a request for the picosecond 422 nm pulse. To minimise latency, the 1092 nm pulse is requested during the previous attempt to arrive at the start of the current attempt. The latencies l_i are given in Table 6.1

sumed by latency, i.e., the delay between the Entangler switching a TTL output and the time taken for the response to arrive at the ion. The latencies, denoted l_i, determined for each module are summarised in Table 6.1. The σ^--polarised 422 nm and 1092 nm pulses implement state preparation for $\sim$320 ns. Since we do not want to destroy the state of the ion once the photon has been detected, we cannot issue requests for the 422 nm state preparation or excitation pulses until we have decided that the previous attempt was unsuccessful – the earliest we can issue these pulses is at the start of the attempt. However, since the 1092 nm light does not couple to the ground Zeeman states, we can issue a request for this pulse in advance, such that the pulse arrives towards the start of the next attempt; we choose for it to arrive $\sim$100 ns into the next attempt. If we successfully detect a photon, the incoming 1092 nm pulse will not affect the ion state. As shown in Figure 6.2, we therefore request the 1092 nm pulse for the next attempt as part of the current attempt. The Entangler then sends the signal for the 422 nm and 1092 nm pulses to be turned off,

Module	**Latencies**		
	l_{422}	l_{1092}	l_{ps}
Alice	414 ns	559 ns	485 ns
Bob	456 ns	500 ns	469 ns

Table 6.1: State preparation "switch-off" latencies, l_{422} and l_{1092}, and the picosecond pulse request latency, l_{ps}. These correspond to the time between an event being issued by the entangler, and the corresponding effect reaching the ion.

such that after the desired state preparation duration, both beams are off at the ion.

There is also a latency associated with requesting a picosecond pulse. When the Entangler requests a pulse, the pulse picker picks a pulse from an 80 MHz pulse train, resulting in a non-deterministic timing of the picked pulse, uniformly distributed across a 12.5 ns window. The timing of the pulse within the window is determined using a fast photodiode trigger, which is sent to the Entangler and enables the determination of the pulse timing to 1 ns. The Entangler uses the timestamp of this trigger as a reference event to gate the APDs inputs to the entangler, and records which detector (APDs) "clicked" during the acceptance window. Finally, the click pattern is compared to the desired click patterns to determine whether the attempt was successful or not. For example, for IPE generation, a successful click pattern would be one and only one detector clicking.

Once an entanglement attempt is deemed successful, the Entangler exits the attempt loop and we set the timestamp of the detection event as the phase reference of the 674 nm DDS – thereby ensuring that we enter the rotating frame of the Sr^+ ion with the same phase, regardless of the timing detection event. As discussed above, we then map the ionic part of the IPE state into the optical network qubit using a 674 nm pulse. If we do not successfully detect a photon within 1 ms, we recool the crystal and try again; however, this only occurs very infrequently when generating ion-photon entanglement.

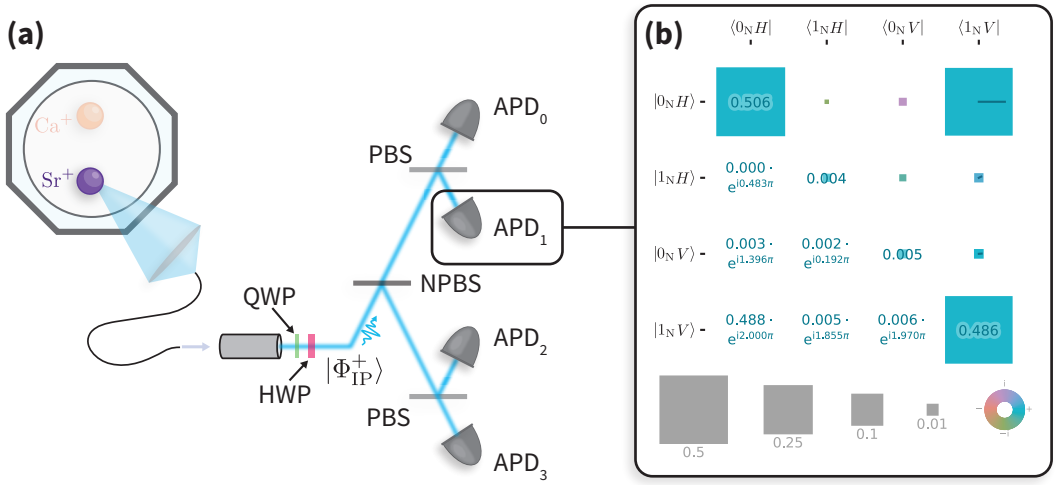

Figure 6.3: **(a)** The single-mode, non-polarisation-maintaining fibre delivers the photon from the module to the central station for tomographic polarisation measurements. A quarter waveplate (QWP) and half waveplate (HWP) at the fibre output enable arbitrary single-qubit rotations before the photon impinges upon a non-polarising beamsplitter (NPBS). At each output port of the non-polarising beamsplitter (NPBS), a PBS and two APDs allow computational basis measurements, with the ability to choose different waveplate settings providing a tomographically complete set of measurements. **(b)** Reconstructed density matrix of the entangled ion-photon state from photon detection events at APD_1. Fidelities and inferred photon collection efficiencies measured for each APD are listed in Table 6.2.

6.1.3 Characterisation

We characterise the entanglement generated between the Sr^+ ion and the polarisation state of the 422 nm photon using state tomography. As discussed in Section 6.1.1, the polarisation analyser provides a tomographically-complete set of measurements of the photonic part of the IPE, while manipulation of the network qubit using the 674 nm laser followed by fluorescence detection provides a tomographically-complete set of measurements of the ionic part. The tomography algorithm differs slightly from the algorithms discussed in Section 3.2, since the measurements of the photonic qubits have different POVMs from those used for the ionic qubits [56]. It is also worth noting that the ground Zeeman states of the Sr^+ ion are only mapped to the optical network qubit after the measurement of the photon, and thus the photon no longer exists when we perform this mapping to the network qubit. However, for consistency with subsequent sections, we characterise this state using the network

Module		APD_0	APD_1	APD_2	APD_3	**Average**
Alice	$\mathcal{F}_{IP}$	98.6(2) %	97.8(2) %	97.8(2) %	98.4(1) %	98.16(8) %
	η_{IP}	1.533(4) %	1.540(4) %	1.873(4) %	1.994(4) %	1.736(2) %
Bob	$\mathcal{F}_{IP}$	98.6(2) %	97.4(2) %	97.3(2) %	98.6(2) %	97.96(8) %
	η_{IP}	1.361(3) %	1.376(3) %	1.814(4) %	1.892(4) %	1.565(2) %

Table 6.2: Entanglement fidelities, $\mathcal{F}_{IP}$, of the ion-photon states reconstructed from tomographic measurements at each APD, and the inferred collection efficiencies, η_{IP}, for each APD.

qubit anyway.

For each module, we took 9 tomographic datasets over a period of $\sim 2\,\mathrm{h}$, each dataset comprising 83 200 tomographic measurements of the ion-photon state. Each dataset was analysed individually, enabling reconstruction of the entangled ion-photon state, $\hat{\rho} \in L\left(\mathcal{Q}_N \otimes \mathcal{Q}_{ph}\right)$ for the photons detected at each APD, in addition to the inferred photon collection efficiency, η_{IP}. The averaged entanglement fidelities of the reconstructed states, and inferred photon collection efficiencies, are given in Table 6.2. Averaging over the four APDs, we observe ion-photon fidelities of 98.16(8) % and 97.96(8) % for Alice and Bob, respectively. Additionally, we observe average photon collection efficiencies of 1.736(2) % and 1.565(2) % for Alice and Bob, respectively.

In principle, this scheme for generating IPE has no fundamental limitation on the achievable fidelity; we expect that the errors observed result from technical imperfections only. From Table 6.2, we see that the entanglement fidelities of the ion-photon states measured by APD_0 and APD_3 are significantly higher than those measured by APD_1 and APD_2. This would suggest that the reflected modes of the PBSs in the Bell state analyser are not as pure as the transmitted modes. Therefore, it appears that for both Alice and Bob, the ion-photon entanglement produced has an entanglement fidelity of at least 98.6 %, before the Bell state analyser. An extensive theoretical treatment of the error sources arising during the generation of IPE was presented by Nadlinger [92]; however, a full experimental characterisation of the error sources has not yet been performed. We expect the leading sources of

error to be polarisation mixing resulting from misalignment of the collection optics, imperfections in the optical elements, and errors in the 674 nm pulse mapping the ionic part of the entanglement into the network qubit.

6.2 Remote entanglement of strontium ions

The ability to generate entanglement between a Sr^+ ion and a 422 nm photon provides each module an interface with the optical quantum network. The next step is to make use of this interface to establish remote entanglement between the two modules. We do this by making use of *entanglement swapping* [134]. In this scheme, we begin by synchronously generating IPE in each module, then performing a projective measurement of the two photons in the Bell basis. This measurement then projects the ions into a maximally entangled state, thereby *swapping* the entanglement from the ion-photon pairs, onto the ions themselves. The IPE generation process is probabilistic, and thus the remote entanglement generation process is also probabilistic, with particular measurement events heralding the creation of remote entanglement.

The projective Bell measurement can be implemented using only linear optics, as shown in Figure 6.4(a). In this apparatus, the photons first undergo unitary rotations of their polarisation – implemented with motorised quarter waveplates (QWPs) and half waveplates (HWPs) – which allow us to account for unwanted birefringence of the optical fibres, such that after the waveplates the ion-photon pairs are in the Φ^+ Bell state. The total state of the two ion-photon pairs can then be written as

$$\left|\Phi^+_{\mathrm{IP}}\right\rangle_A \left|\Phi^+_{\mathrm{IP}}\right\rangle_B = \frac{|\downarrow\rangle|H\rangle + |\uparrow\rangle|V\rangle}{\sqrt{2}} \otimes \frac{|\downarrow\rangle|H\rangle + |\uparrow\rangle|V\rangle}{\sqrt{2}} \tag{6.3}$$

$$= \frac{1}{2}\left[\left|\Phi^+_{\mathrm{II}}\right\rangle\left|\Phi^+_{\mathrm{PP}}\right\rangle + \left|\Phi^-_{\mathrm{II}}\right\rangle\left|\Phi^-_{\mathrm{PP}}\right\rangle + \left|\Psi^+_{\mathrm{II}}\right\rangle\left|\Psi^+_{\mathrm{PP}}\right\rangle + \left|\Psi^-_{\mathrm{II}}\right\rangle\left|\Psi^-_{\mathrm{PP}}\right\rangle\right], \tag{6.4}$$

where to get to the second line, we have reordered the tensor products such that the ion states and photon states are together. From Eq. (6.4), we see that measuring the

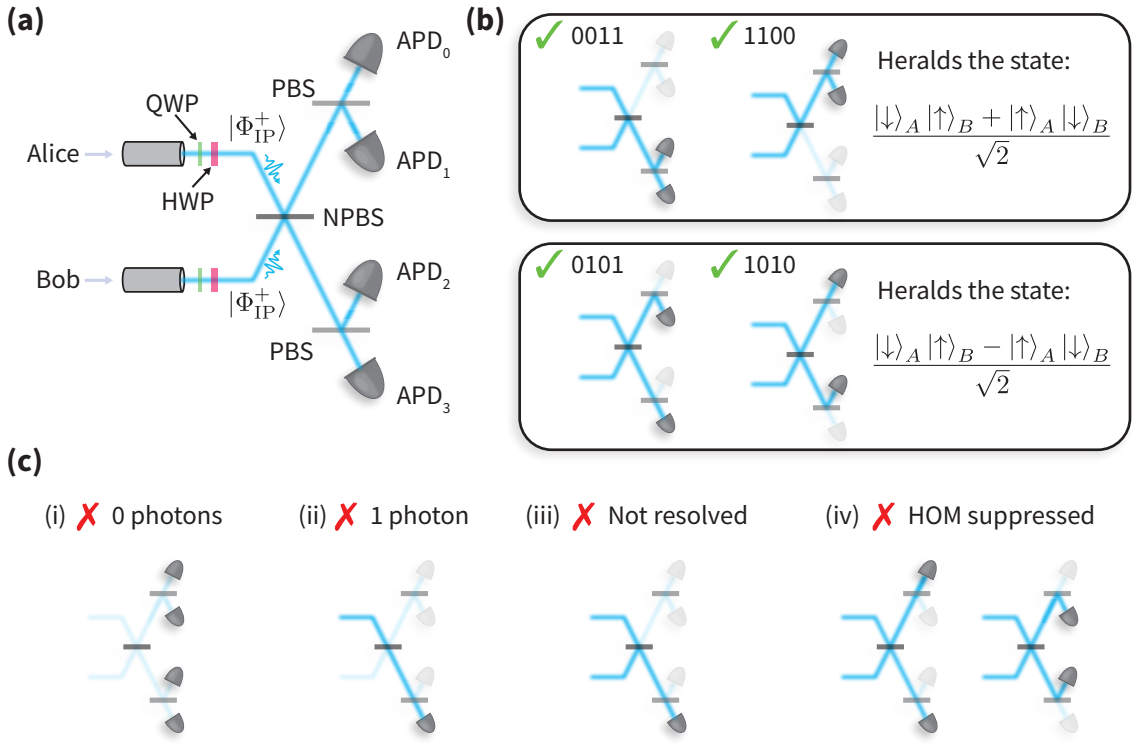

Figure 6.4: **(a)** The Bell-state measurement apparatus, also known as the central heralding station, is the same setup used for tomographic measurements of the ion-photon states. A QWP and HWP at the fibre outputs compensate for fibre birefringence, ideally creating a maximally entangled $|\Phi^+\rangle$ state. The two photons interfere at a NPBS, where the Hong-Ou-Mandel (HOM) effect erases the "which-path" information. PBSs and APDs measure the polarisation of the two-photon state, with different detector click events corresponding to different measurement outcomes. **(b)** Detection patterns that herald successful entanglement generation. Same-side detections indicate a symmetric Bell state, while opposite-side detections indicate an anti-symmetric Bell state. **(c)** Examples of detection patterns that do not herald entanglement generation, including (i) zero and (ii) single-photon detection, (iii) two-photon signals allowed by the HOM effect but are not resolved by the APDs, and (iv) patterns resulting from imperfect photon distinguishability or central station imperfections.

photons in the Bell basis will project the ions into a maximally entangled Bell state, with the measurement outcome determining which state the ions are projected into.

The photons are interfered on a non-polarising beamsplitter (NPBS) which, for perfectly indistinguishable photons, erases the "which path" information via the Hong-Ou-Mandel (HOM) effect [135]. This erasure of which path information manifests as a symmetrisation of two-photon state, with the total wavefunction obeying bosonic exchange symmetry. Without the polarisation degree of freedom, the resulting spatial wavefunction of the photons must obey bosonic statistics, and thus the photons will leave via the same port of the beamsplitter – a phenomenon known as *bunching*, which is typical for HOM interference. Now considering the polarisation states of the photons, we find additional options for the spatial and polarisation degrees of freedom, still with requirement that the overall state be symmetric under exchange. For the symmetric spatial wavefunction, the polarisation state of the photons must also be symmetric and thus will be one of the Bell states $\{\left|\Phi^+_{\mathrm{PP}}\right\rangle, \left|\Phi^-_{\mathrm{PP}}\right\rangle, \left|\Psi^+_{\mathrm{PP}}\right\rangle\}$; for the anti-symmetric spatial wavefunction, the polarisation state must also be anti-symmetric to ensure overall symmetric exchange symmetry, which is only satisfied by the singlet polarisation state $\left|\Psi^-_{\mathrm{PP}}\right\rangle$.

We make a measurement of the two-photon state using a PBS and two APDs at each output port of the NPBS. If the measurement reveals the photons had opposing polarisations, then we know that the photons had been in one of $\left|\Psi^+_{\mathrm{PP}}\right\rangle$ or $\left|\Psi^-_{\mathrm{PP}}\right\rangle$. We can further distinguish between these two states by considering the spatial wavefunction of the photons: if the photons left by the same port, i.e., the two-photon coincidence occurred on the same side of the NPBS, then the photons must have been in the symmetric $\left|\Psi^+_{\mathrm{PP}}\right\rangle$ state and thus the ions are projected into the $\left|\Psi^+_{\mathrm{II}}\right\rangle$ state; if the photons left by different ports, then the photons must have been in the anti-symmetric $\left|\Psi^-_{\mathrm{PP}}\right\rangle$ state and thus the ions are projected into the $\left|\Psi^-_{\mathrm{II}}\right\rangle$ state. Observation of these two-photon coincidences, depicted in Figure 6.4, thus herald the creation of entanglement between the Sr^+ ions. Note that non-linear optics would be required to be able to resolve the other two Bell states, and so even

with perfect photon collection efficiency, we can only resolve half of the possible heralds. Therefore, given ion-photon entanglement probabilities of η_A and η_B in Alice and Bob, the remote entanglement success probability can be estimated as

$$\eta = \frac{1}{2}\eta_{\mathrm{A}}\eta_{\mathrm{B}}. \tag{6.5}$$

The above discussion relied on the underlying assumption that the photons are indistinguishable; any distinguishability between the photons will thus lead to a reduction in fidelity of the heralded state to the ideal maximally entangled state. In particular, care must be taken to ensure the complete overlap of the photonic wavepackets at the NPBS.

6.2.1 Experimental implementation

As discussed above, the remote entanglement generation process is try-until-success. Each entanglement generation attempt involves synchronously generating IPE in each module, as described in Section 6.1, before performing the Bell measurement on the photons. The pulse sequence used for the entanglement attempts is the same as in Figure 6.2; however, the click patterns which denote a success are now the two-photon detection coincidences outlined in Figure 6.4. From Eq. (6.5), it is clear that a significantly higher number of attempts will be required. Therefore, to avoid excess heating while generating the entanglement, we interleave 500 µs of entanglement attempts with ≈ 320 µs of Doppler cooling and 500 µs of EIT cooling. As with the IPE, once we have successfully generated the entanglement, we immediately map the ground Zeeman states, $\{|\downarrow\rangle, |\uparrow\rangle\}$, to the optical network qubit, $\mathcal{Q}_{\mathrm{N}}$, for further manipulation.

As outlined in Figure 6.4, the different herald patterns give rise to maximally entangled states with different phases. While Figure 6.4 presented the two resulting states as the Bell states $|\Psi^+\rangle$ and $|\Psi^-\rangle$, in practice, residual fibre birefringence leads to an additional phase offset for the entangled states. While this phase should, in

principle, be common to all herald patterns, we assume that each pattern has a different entangled state phase, ϕ_i, such that

$$\left|\Psi^i_{\mathrm{NN}}\right\rangle = \frac{|0_{\mathrm{N}}1_{\mathrm{N}}\rangle + e^{i\phi_i}\,|1_{\mathrm{N}}0_{\mathrm{N}}\rangle}{\sqrt{2}}. \tag{6.6}$$

These phases are calibrated, thereby allowing us to perform the necessary Z-rotation to remove this phase for each herald pattern. Once we've mapped the remote entanglement into the optical network qubit and accounted for this phase, each herald pattern corresponds to the same entangled state,

$$\left|\Psi^+_{\mathrm{NN}}\right\rangle = \frac{|0_{\mathrm{N}}1_{\mathrm{N}}\rangle + |1_{\mathrm{N}}0_{\mathrm{N}}\rangle}{\sqrt{2}}. \tag{6.7}$$

Note that we can trivially rotate this Bell state to any of the other Bell states using single-qubit $\hat{\sigma}_x$ and $\hat{\sigma}_z$ rotations in one of the modules. In particular, the $\left|\Phi^+_{\mathrm{NN}}\right\rangle$ is used frequently in this chapter and is created by performing a $\hat{\sigma}_x$ rotation on the second qubit in Eq. (6.7).

6.2.2 Characterisation

We characterise the "raw" remote entanglement between the Sr^+ network qubits using state tomography, as outlined in Section 3.2. The tomography comprises tomographic measurements on 200 000 copies of the $\left|\Psi^+_{\mathrm{NN}}\right\rangle$ state. The reconstructed density matrix representing the entangled state of the network qubits, shown in Figure 6.5(b), has an entanglement fidelity of 96.94(9) %, and a fidelity of 96.89(8) % to the desired $|\Psi^+\rangle$ Bell state.

As with the ion-photon entanglement, a detailed theoretical treatment of the possible error mechanisms of our remote entanglement generation scheme is presented by Nadlinger [92]; however, an experimental determination of our remote entanglement error budget has not yet been determined. If we consider the Bell state analyser to be perfect and the photons to be perfectly indistinguishable when impinging on the NPBS, then we would expect a 95.81(9) % remote entanglement

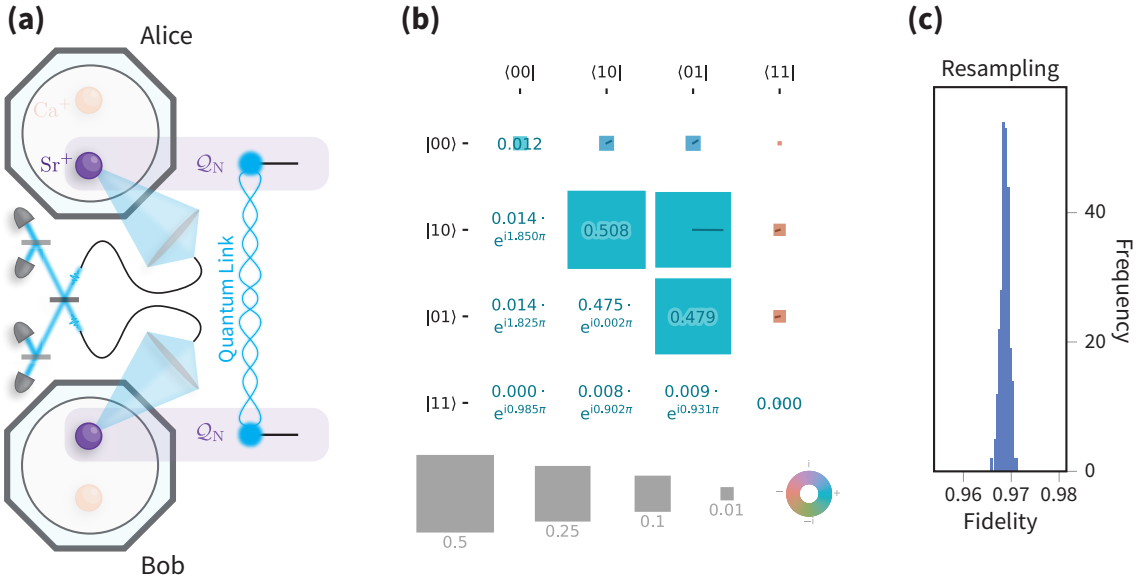

Figure 6.5: Generation of the "raw" Sr^+-Sr^+ remote entanglement. **(a)** Remote entanglement is generated between the network qubits in separate modules. **(b)** Density matrix $\hat{\rho} \in \mathcal{Q}_N{}^{\otimes 2}$, reconstructed from measurements of 200 000 copies using quantum state tomography. The reconstructed state has an entanglement fidelity of 96.94(9) %, and a fidelity of 96.89(8) % to the $|\Phi^+\rangle$ state. **(c)** Histogram of the maximally-entangled state fidelities of states reconstructed from resampled datasets.

fidelity[2], which is significantly lower than the observed fidelity. Due to the large number of measurements that formed the tomographic measurement presented here (alternatively, the relatively small spread in fidelities acquired from resampling), we do not believe this measurement to be a statistical anomaly. A full characterisation of the Bell state analyser should be performed so that we can better understand the imperfections of the Bell state analyser, and compare to theoretical predictions [92]; however, this was not done for this thesis.

Additional error mechanisms that do not appear in the ion-photon entanglement – such as an imperfect NPBS, distinguishability of photons interfering on the NPBS, and detector timing resolution and jitter – will also contribute. However, these mechanisms would result in the remote entanglement error exceeding that predicted from the ion-photon errors. As we do not observe this, we believe these errors are relatively small.

[2]Calculated by averaging over the fidelities of the entangled states that would be produced using the observed ion-photon density matrices and a perfect Bell state analyser

The average number of attempts before successful generation of entanglement was $m = 8089(18)$, which corresponds to an average of success probability of $1.236(3) \times 10^{-4}$. Including recooling, the average time taken to generate entanglement was 25.44(6) ms, which corresponds to an entanglement generation rate of $39.31(9)\,\mathrm{s}^{-1}$. This recooling consumes a significant fraction of the total entanglement generation duty-cycle. O'Reilly *et al.* [46] recently demonstrated that the entanglement generation rate could be increased using sympathetic cooling of a second species of ion; however, while the co-trapping of Ca^+ would enable a continuous attempt loop, this was not explored in this work[3]. Previously in our apparatus, Stephenson *et al.* [130] reported an average success probability of 2.18×10^{-4}, and an average entanglement generation rate of $182\,\mathrm{s}^{-1}$, in the same apparatus. We attribute the lower success probability reported here due to losses through the Bell state analyser due to misalignment and degradation of optical components. These losses, in addition to the extra cooling, account for the lower entanglement generation rate reported here.

6.3 Remote entanglement of strontium and calcium ions

Up to this point, we have demonstrated the key building blocks of our mixed-species trapped-ion quantum network: local control of the mixed-species trapped-ion modules, and the ability to share entanglement between network qubits in the separate modules. We now reach the point where we can start to piece these together to construct more complex experimental sequences. In this section, we explore a variety of remotely entangled states that can be created with two trapped-ion modules, each holding a Sr^+ ion and a Ca^+ ion. The states that we will consider will be of the

[3]While certainly of interest, this was not explored since the additional time required to perform state preparation of the Ca^+ ion after acting as a sympathetic cooling during the generation of entanglement would lead to dephasing of the Sr^+ ion, and hence errors in the mixed-species entanglement. Additionally, in the DQC demonstrations, the Ca^+ ion provides a robust quantum memory, and thus cannot be utilised as a sympathetic coolant.

form

$$|\psi\rangle = \frac{|0\rangle^{\otimes N} + |1\rangle^{\otimes N}}{\sqrt{2}} \in \mathcal{Q}^{\otimes N}, \tag{6.8}$$

where for $N = 2$ this is the Φ^+ Bell state, while for $N > 2$, this is the N-qubit GHZ state. We characterise these states using two techniques: parity/population measurements, and state tomography.

6.3.1 Characterisation via state tomography

Sr^+-Ca^+ entanglement

We first consider the remote entanglement of a Sr^+ ion and a Ca^+ ion across the quantum network. The circuit with which we create this state, shown in Figure 6.6(a), begins with the generation of remote entanglement between the network qubits, as in Eq. (6.7), which is then rotated to the state $\left|\Phi^+_{\mathrm{NN}}\right\rangle$. To mitigate the effect of heating on the local entangling operations, we interleave 200 µs of entanglement attempts with 2.25 ms of cooling. This cooling comprises 1.05 ms of Doppler cooling, followed by 1.2 ms of EIT cooling, which alternates between cooling the axial IP and OOP modes, which are critical for the mixed-species entangling operations. With the raw entanglement generated, we perform an error-detected iSWAP-gate in Bob to map the state of the network qubit to the auxiliary qubit in Ca^+, as outlined in Section 5.2.6. The outcome of the error detection is sent to Alice in real-time via a TTL link – allowing the protocol to be aborted in real-time if an error is detected. If no error is detected, we have successfully created the state

$$\left|\Phi^+_{\mathrm{NX}}\right\rangle = \frac{|0_{\mathrm{N}}0_{\mathrm{X}}\rangle + |1_{\mathrm{N}}1_{\mathrm{X}}\rangle}{\sqrt{2}}. \tag{6.9}$$

The density matrix for the created state, $\hat{\rho}_{\mathrm{NX}}$, is reconstructed from 10 000 tomographic measurements using state tomography, and is shown in Figure 6.6(b). The reconstructed state has an entanglement fidelity of 94.1(6) %. To create these states, we generated a total of 10 442 raw Sr^+-Sr^+ remotely entangled states; the iSWAP error detection detected an error in 442 cases, corresponding to an iSWAP

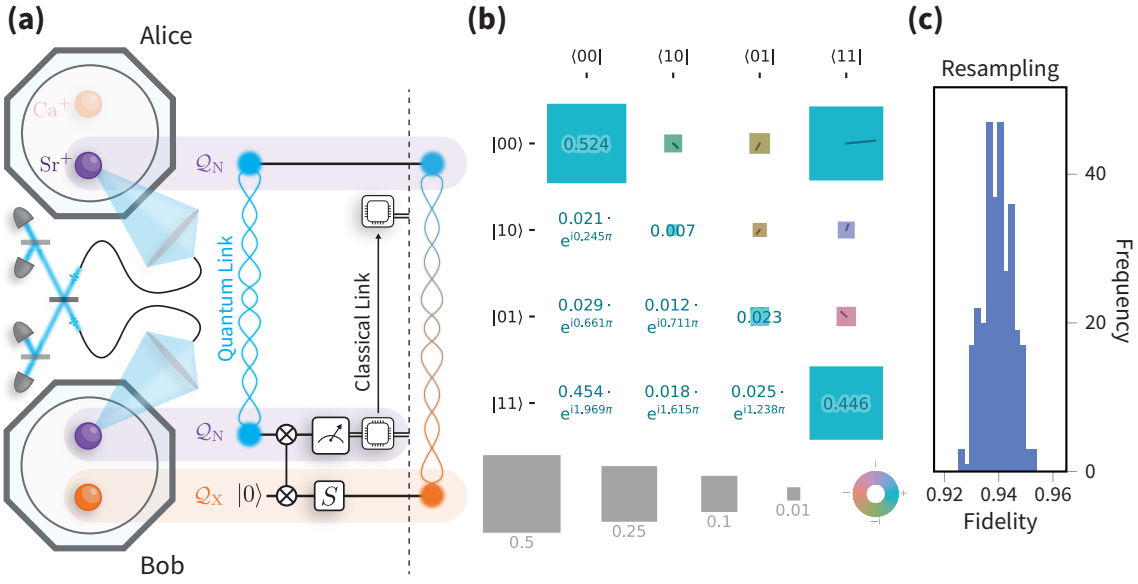

Figure 6.6: Generation of Sr^+-Ca^+ remote entanglement. **(a)** Remote entanglement is first generated between the network qubits in separate modules. Subsequently, an error-detected iSWAP gate is used to transfer the state of the network qubit in Bob to the auxiliary qubit. The error-detection measurement outcome is exchanged via the classical (TTL) link, such that both modules only progress passed the dashed line if no error is detected. **(b)** Density matrix $\hat{\rho} \in \mathcal{Q}_N \otimes \mathcal{Q}_X$, reconstructed from measurements of 10 000 copies using quantum state tomography. The reconstructed state has an entanglement fidelity of 94.1(6) %. **(c)** Histogram of the maximally-entangled state fidelities of states reconstructed from resampled datasets.

error rate of 4.2(2) %. Including the iSWAP error detection, we observe a total success probability of $1.03(1) \times 10^{-4}$, and a generation rate of $7.14(7)\,\mathrm{s}^{-1}$.

Ca^+-Ca^+ entanglement

We now consider the remote entanglement of Ca^+ ions across the quantum network. The circuit for creating this state is shown in Figure 6.7(a). As for the Sr^+-Ca^+ state, we begin by generating remote entanglement between the network qubits, given by Eq. (6.7), followed by rotation to the $\left|\Phi^+_{NN}\right\rangle$ Bell state. The same cooling sequence is used in each module to ensure the Sr^+-Ca^+ crystals are close to their ground state in preparation for the local entangling operations. We perform error-detected iSWAP gates in both modules to map the state of the network qubits to the auxiliary qubits. The modules exchange their respective error detection outcomes via the TTL link; if either module detects an error, the protocol is aborted in real-time. If no error is

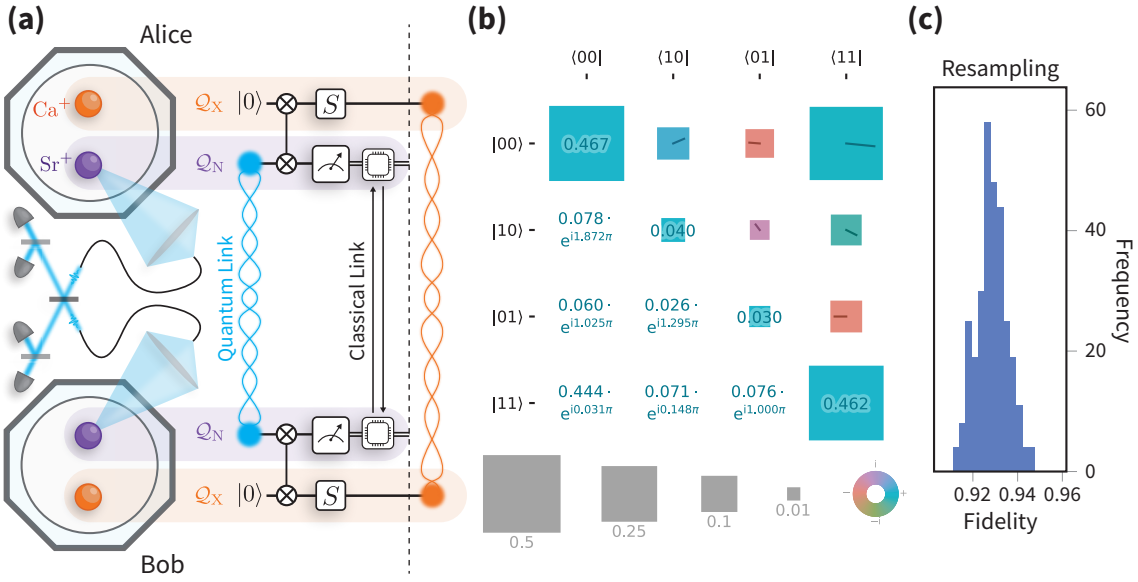

Figure 6.7: Generation of Ca^+-Ca^+ remote entanglement. **(a)** Remote entanglement is first generated between the network qubits in separate modules. Subsequently, error-detected iSWAP gates are used to transfer the states of the network qubits to the auxiliary qubits. The error-detection measurement outcomes are exchanged via the classical (TTL) link, such that both modules only progress passed the dashed line if no error is detected. **(b)** Density matrix $\hat{\rho} \in \mathcal{Q}_X^{\otimes 2}$, reconstructed from measurements of 10 000 copies using quantum state tomography. The reconstructed state has an entanglement fidelity of 93.1(7) %. **(c)** Histogram of the maximally-entangled state fidelities of states reconstructed from resampled datasets.

detected, we have successfully created the state

$$\left|\Phi^+_{XX}\right\rangle = \frac{|0_X 0_X\rangle + |1_X 1_X\rangle}{\sqrt{2}}. \tag{6.10}$$

The reconstructed state of the entangled auxiliary qubits, $\hat{\rho}_{XX}$, is shown in Figure 6.7(b) and has an entanglement fidelity of 93.1(7) %. As with the Sr^+-Ca^+ state, we performed state tomography using tomographic measurements on 10 000 copies of $\hat{\rho}_{XX}$. To create these states, we generated a total of 10 922 raw Sr^+-Sr^+ entangled states; the iSWAP error detection detected an error in 922 cases, corresponding to an iSWAP error rate of 8.4(3) %. Including the iSWAP error detection, we observe a total success probability of $1.24(1) \times 10^{-4}$, and a generation rate of $8.6(1)\,s^{-1}$.

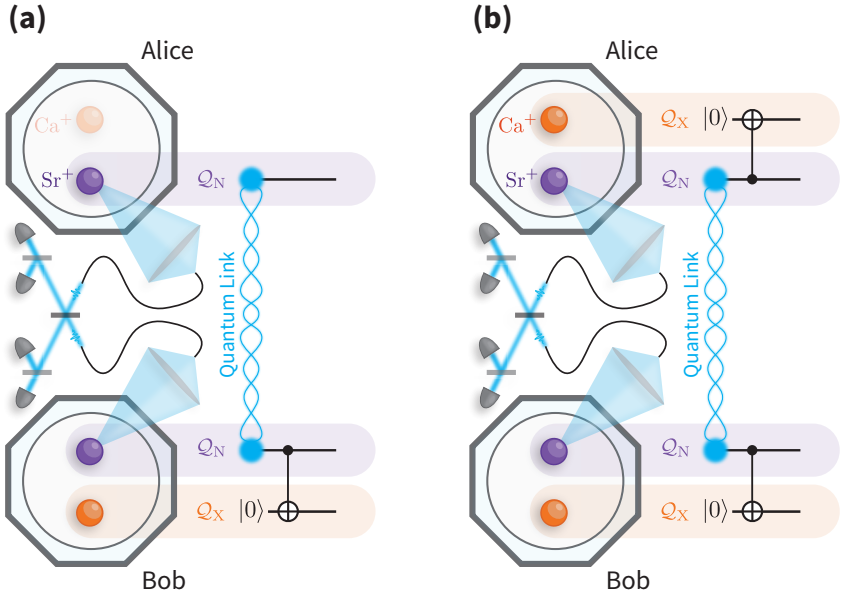

Figure 6.8: Circuits for creating mixed-species **(a)** 3-qubit and **(b)** 4-qubit GHZ states. Both circuits begin with the generation of remote entanglement between the network qubits. **(a)** A mixed-species CNOT gate in Bob create the 3-qubit GHZ state. **(b)** Mixed-species CNOT gates in both modules create the 4-qubit GHZ state.

3- and 4-qubit GHZ states

So far, we have considered only bipartite entanglement – using the error-detected iSWAP gate to transfer states of the network qubits to the auxiliary qubits. However, since have have more than two qubits at our disposal, we can create multipartite entanglement.

The circuits for creating the 3- and 4-qubit GHZ states are shown in Figure 6.8. The circuits start with the auxiliary qubits prepared in the $|0_X\rangle$ states. As before, we generate remote entanglement between the network qubits, followed by single-qubit rotations mapping this entanglement to the Φ^+ Bell state. To create the 3-qubit GHZ state, we apply a mixed-species controlled-NOT (CNOT) gate between the network and auxiliary qubits in Bob, creating the state

$$|\mathrm{GHZ}_3\rangle = \mathrm{CNOT}_B \left[\frac{|0_N 0_N\rangle + |1_N 1_N\rangle}{\sqrt{2}} \otimes |0_X\rangle \right] \tag{6.11}$$

$$= \frac{|0_N 0_N 0_X\rangle + |1_N 1_N 1_X\rangle}{\sqrt{2}}, \tag{6.12}$$

where $\mathrm{CNOT_i}$ is the CNOT gate acting on the qubits in module i. Similarly, the 4-qubit GHZ state is created by applying the mixed-species CNOT gates in each module, such that

$$|\mathrm{GHZ}_4\rangle = \mathrm{CNOT_A CNOT_B}\left[\frac{|0_\mathrm{N}0_\mathrm{N}\rangle + |1_\mathrm{N}1_\mathrm{N}\rangle}{\sqrt{2}} \otimes |0_\mathrm{X}0_\mathrm{X}\rangle\right] \tag{6.13}$$

$$= \frac{|0_\mathrm{N}0_\mathrm{N}0_\mathrm{X}0_\mathrm{X}\rangle + |1_\mathrm{N}1_\mathrm{N}1_\mathrm{X}1_\mathrm{X}\rangle}{\sqrt{2}}. \tag{6.14}$$

We reconstruct the created GHZ states using state tomography with tomographic measurements performed on 10 000 copies of the states. The reconstructed density matrices for the 3- and 4-qubit GHZ states are shown in Figure 6.9 and Figure 6.10, respectively. In calculating the fidelity of the reconstructed states, $\hat{\rho}_\mathrm{AB}$, to the ideal GHZ states, we allow arbitrary rotations about the Z-axis, such that

$$\mathcal{F}_{\mathrm{GHZ}_N} = \max_\phi \langle \mathrm{GHZ}_N | \, \hat{R}_Z(\phi)\, \hat{\rho}_\mathrm{AB} \hat{R}_Z(\phi)^\dagger \, | \mathrm{GHZ}_N \rangle \,, \tag{6.15}$$

where $\hat{R}_Z(\phi)$ denotes a single-qubit Z-rotation by the angle ϕ on each qubit. Using this definition, we calculate GHZ fidelities of 93.1(7) % and 91.9(8) %, for the 3- and 4-qubit GHZ states, respectively.

Since the circuit for creating the GHZ states does not involve any error detection, the success probability for creating these states is simply the success probability for generating the raw remote entanglement. For the 3-qubit GHZ state, we measured a success probability of $1.20(1) \times 10^{-4}$, and a generation rate of $8.26(8)\,\mathrm{s}^{-1}$; for the 4-qubit GHZ state, we measured a success probability of $1.44(2) \times 10^{-4}$, and a generation rate of $9.9(1)\,\mathrm{s}^{-1}$. We note the variability in the success probabilities, not just for the GHZ states, but also for the other remote entanglement experiments. We attribute this variability to drifts in the beam pointing of the excitation laser, and alignment of the collection optics.

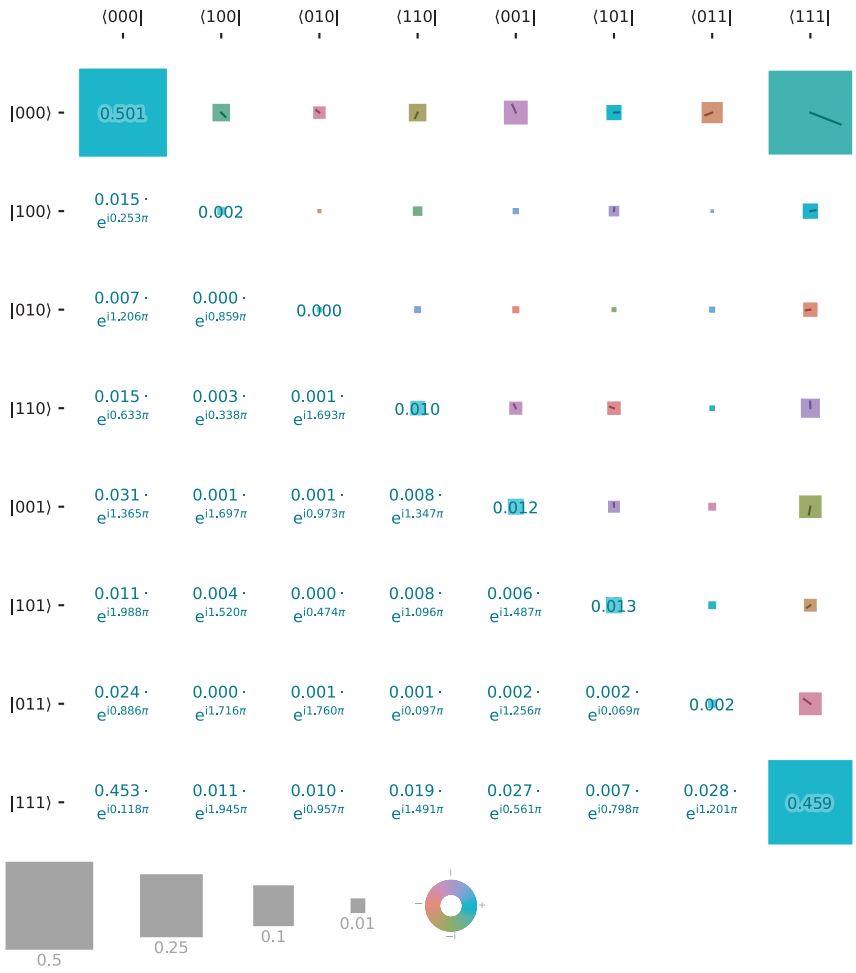

Figure 6.9: Density matrix of the 3-qubit (Sr^+-Sr^+-Ca^+) GHZ state, $\hat{\rho} \in \mathcal{Q}_N{}^{\otimes 2} \otimes \mathcal{Q}_X$, reconstructed from measurements of 10 000 copies using quantum state tomography. The reconstructed state has a fidelity of 93.1(7) % to the 4-qubit GHZ state, after allowing for arbitrary Z rotations.

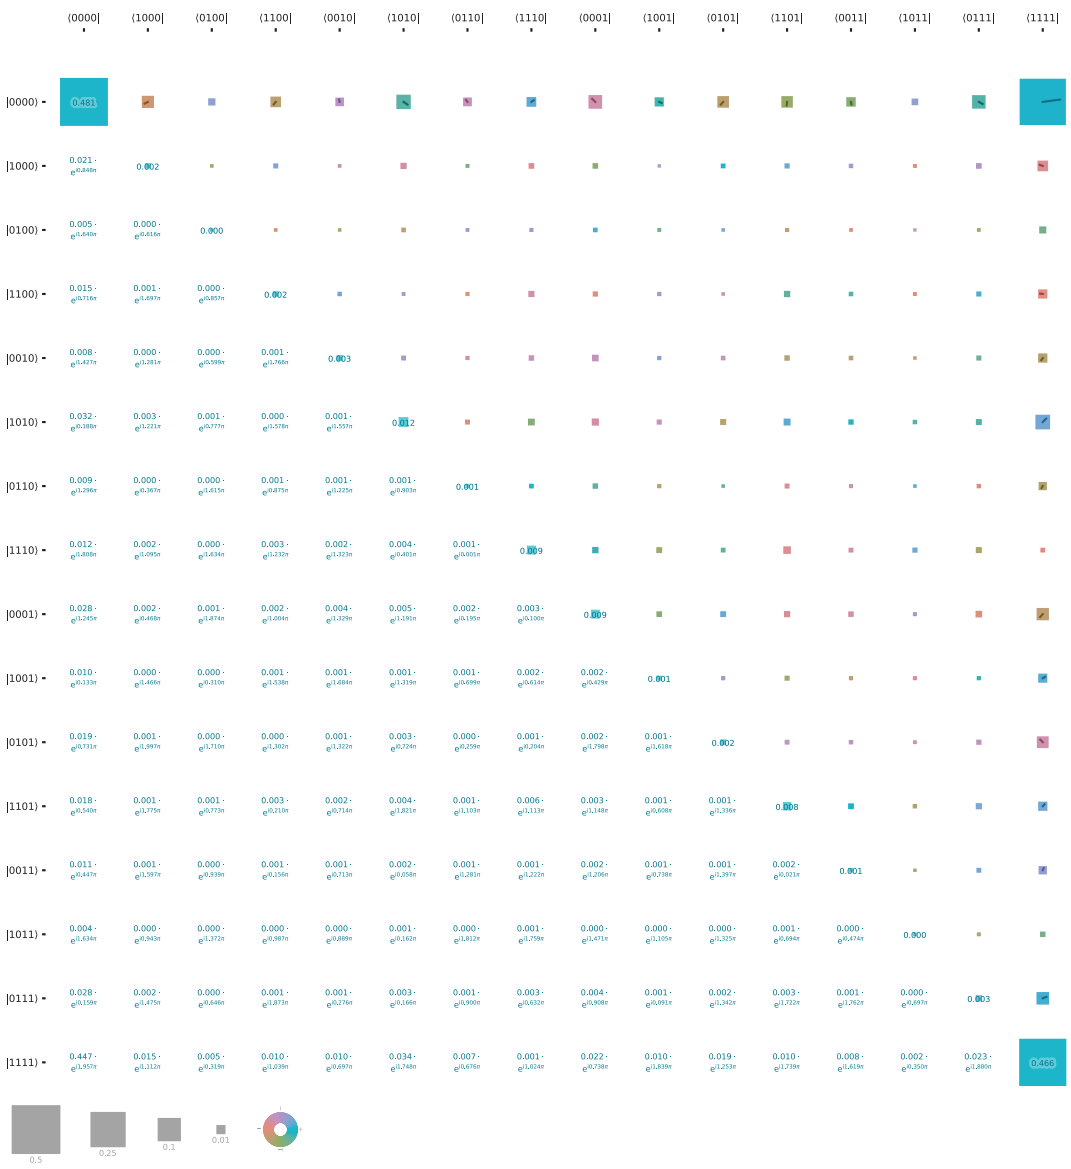

Figure 6.10: Density matrix of the 4-qubit (Sr^+-Sr^+-Ca^+-Ca^+) GHZ state, $\hat{\rho} \in {\mathcal{Q}_N}^{\otimes 2} \otimes {\mathcal{Q}_X}^{\otimes 2}$, reconstructed from measurements of 10 000 copies using quantum state tomography. The reconstructed state has a fidelity of 91.9(8) % to the 4-qubit GHZ state, after allowing for arbitrary Z rotations.

6.3.2 Characterisation via parity/population measurements

As discussed, quantum state tomography provides a complete characterisation of an unknown quantum state; however, it is expensive in the number of measurements required to achieve a particular precision. However, given prior knowledge about what form the state should take, parity/population measurements can be an inexpensive way to extract information about various properties of the state.

Let us write the states that we create as the density matrix

$$\hat{\rho} = \sum_{i,j} \rho_{i,j} \left|i\right\rangle \left\langle j\right| \in L\left(\mathcal{Q}^{\otimes N}\right), \tag{6.16}$$

where $|i\rangle \in \mathcal{Q}^{\otimes N}$. The fidelity of the created state to the desired state, Eq. (6.8), is given by

$$\begin{aligned}\mathcal{F} &= \langle\psi|\,\hat{\rho}\,|\psi\rangle \\ &= \frac{1}{2}\left[\rho_{0\ldots0,0\ldots0} + \rho_{1\ldots1,1\ldots1} + 2|\rho_{0\ldots0,1\ldots1}|\cos\left(\varphi\right)\right],\end{aligned}$$

where $\varphi = \arg\left(\rho_{0\ldots0,1\ldots1}\right)$. Let us now define the population observable, $\hat{P} = \hat{\Pi}_0^{\otimes N} + \hat{\Pi}_1^{\otimes N}$, such that

$$P = \left\langle \hat{P} \right\rangle = \rho_{0\ldots0,0\ldots0} + \rho_{1\ldots1,1\ldots1}. \tag{6.17}$$

Furthermore, define the parity observable, $\hat{\sigma}_\phi^{\otimes N}$, where $\hat{\sigma}_\phi = \cos\left(\phi\right)\hat{\sigma}_x + \sin\left(\phi\right)\hat{\sigma}_y$. Given quantum states of the form of Eq. (6.8), we anticipate that the expectation values for the parity measurements will be of the form

$$\left\langle \hat{\sigma}_\phi^{\otimes N} \right\rangle = C\cos\left(N\phi - \varphi\right), \tag{6.18}$$

where $C = 2|\rho_{0\ldots0,1\ldots1}|$. Thus, if we define the fidelity of the created state to the state in Eq. (6.8) up to an arbitrary Z-rotation, we find that the fidelity is given by

$$\mathcal{F} = \frac{1}{2}\left(P + C\right). \tag{6.19}$$

Experimentally, the population measurement can be performed by simply measuring the created state in the computational basis, and computing the expectation value $\left\langle \hat{P} \right\rangle$, i.e., the probability that the qubits were all measured to be in the same state, e.g., $\mathrm{P}(00) + \mathrm{P}(11)$ for the 2-qubit case. The parity measurements can be implemented by first performing $\hat{R}_{\phi+\pi/2}(\pi/2)$ rotations on each qubit, before measuring in the computational basis. This corresponds to a measurement of the qubits along the axis at an angle ϕ to the X-axis in the XY-plane of the Bloch sphere, associated with the operator $\hat{\sigma}_\phi$. We thus use these measurement outcomes to calculate the expectation values, $\langle \hat{\sigma}_\phi^{\otimes N} \rangle$.

The results for the parity/population measurements for all remotely entangled states discussed in this chapter are shown in Figure 6.11. These results are generally consistent with the state tomography characterisations, but required significantly fewer measurements to obtain. We note that in both the tomography and the parity/population measurements, a number of the states exhibit phase offsets. We attribute this to miscalibration of the Sr^+-Sr^+ herald pattern phases.

6.4 Remote entanglement of long-lived quantum memories

One of the key reasons for integrating Ca^+ into our modules, is that the hyperfine structure gives rise to magnetic field-insensitive transitions, which provide ideal qubits for storing quantum information due to their relatively long coherence times. In particular, the ability to faithfully store quantum information for durations significantly longer than the time taken to establish remote entanglement between modules is particularly important for applications requiring multiple instances of remote entanglement generation. As discussed earlier in this chapter, we can generate remote entanglement at a rate $\sim 40\,\mathrm{s}^{-1}$; however, this drops to $\sim 10\,\mathrm{s}^{-1}$ when additional cooling to the ground is required for executing the mixed-species logic. We have observed coherence times of the optical network qubit of $\sim 8\,\mathrm{ms}$, which,

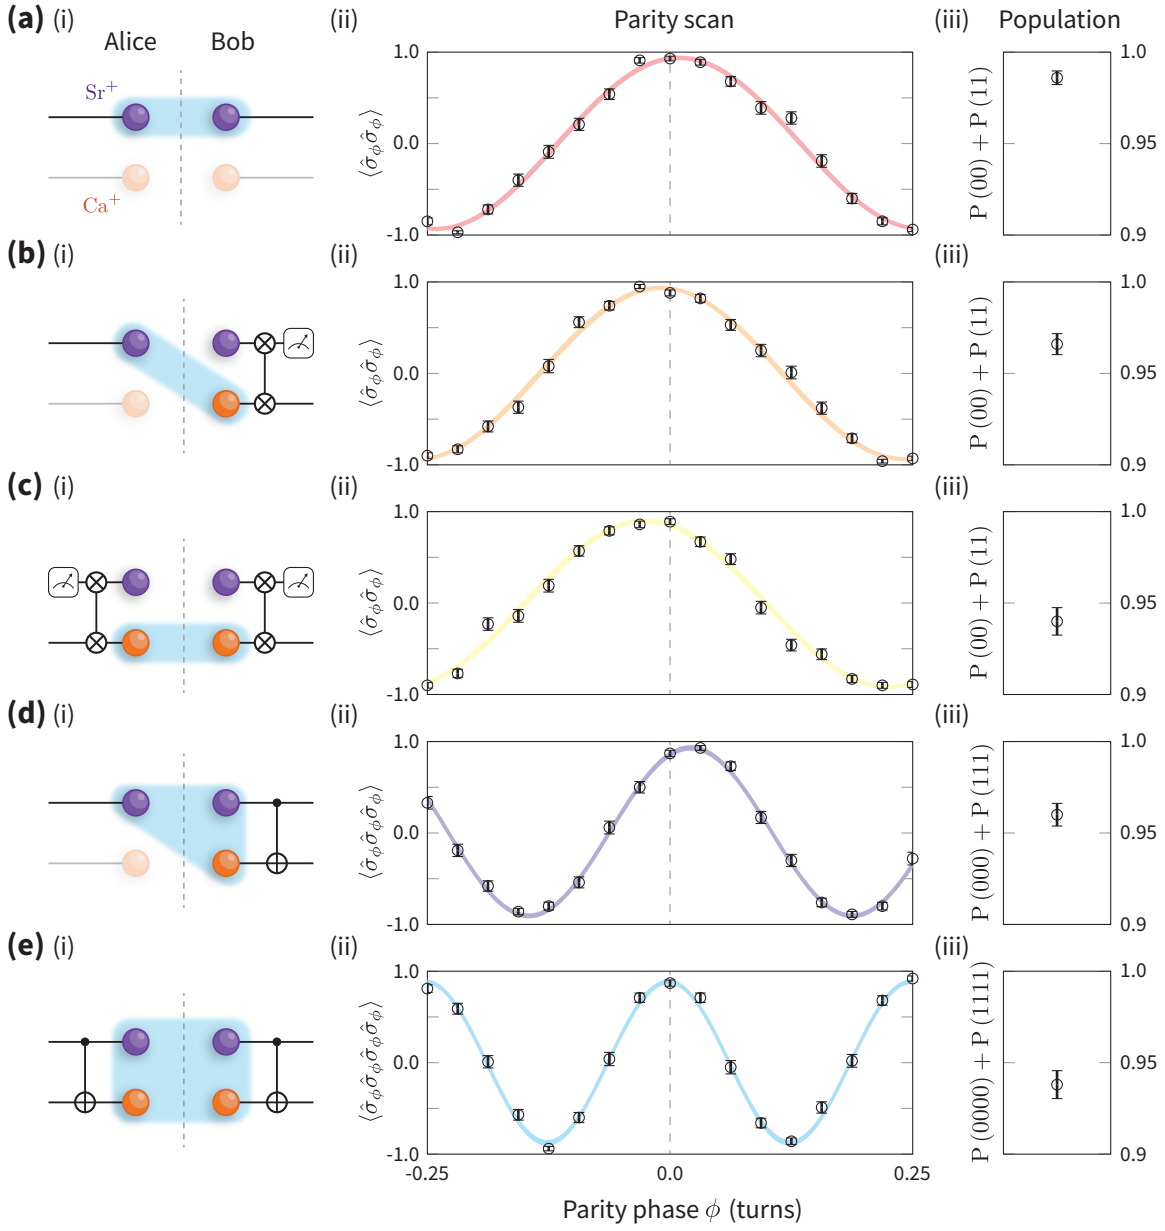

Figure 6.11: Parity and population measurements for mixed-species states, showing (i) a minimal depiction of circuit for generating the entangled state, (ii) a parity scan, and (iii) a population measurement for each state. **(a)** Remotely entangled state of network qubits, with a calculated fidelity of 96.0(7) % to the $\left|\Psi^+_{\mathrm{NN}}\right\rangle$ state. **(b)** Remotely entangled state of a network and an auxiliary qubit, with a calculated fidelity of 95.1(8) % to the $\left|\Psi^+_{\mathrm{NX}}\right\rangle$ state. **(c)** Remotely entangled state of auxiliary qubits, with a calculated fidelity of 92(1) % to the $\left|\Psi^+_{\mathrm{XX}}\right\rangle$ state. **(d)** GHZ state with one auxiliary and two network qubits, with a calculated fidelity of 94(1) % to the $|\mathrm{GHZ}_3\rangle$ state. **(d)** GHZ state with two auxiliary and two network qubits, with a calculated fidelity of 91(1) % to the $|\mathrm{GHZ}_4\rangle$ state.

while this can be extended a little further using dynamical decoupling, is not sufficient for practical applications. Additionally, the spectral isolation between the two species of ions makes the Ca^+ qubits robust to network activity [123].

The circuit qubit, $\mathcal{Q}_C$, takes on the role of the long-lived quantum memory for our modules. Recall that this qubit is only first-order insensitive to magnetic field fluctuations at zero magnetic field; at 0.5 mT the circuit qubit still exhibits a relatively low first-order sensitivity of $122\,\mathrm{kHz\,mT^{-1}}$, however, this is significantly lower than, e.g., the network qubit, which has a sensitivity of $-11.2\,\mathrm{MHz\,mT^{-1}}$. In previous work, we demonstrated the integration of this circuit qubit into one of our network modules [123]; we stored the ionic part of an ion-photon pair for up to 10 s, and demonstrated that the storage was robust to network activity. Here, we extend this to both modules, thereby enabling the long-lived storage of remote entanglement.

In order to demonstrate the storage enhancement provided by our circuit qubits, we compare the abilities of the network and circuit qubits to store remote entanglement, measuring the fidelity as a function of storage duration. The experimental circuits used to characterise the storage capabilities of the qubits are shown in Figure 6.12. The first step is to generate entanglement between the desired pair of qubits. For the network qubits, this is generated as discussed in Section 6.2, before single-qubit rotations are used to map the state to the Φ^+ state. To generate the remote entanglement between the circuit qubits, we first use error-detected iSWAP-gates to map the entanglement from the network to the auxiliary qubits, as in Section 6.3, before using the hyperfine transfer sequence to coherently map the entanglement into the circuit qubits, as shown in Figure 6.12(b). Note that the error detection is performed in real-time, thus enabling the protocol to be restarted if an error is detected prior to the storage part of the circuit.

With the entanglement created, we store the states for a variable storage duration Δt. We deploy universally robust (UR) dynamical decoupling [68] to suppress sources of dephasing, with the pulse sequence lengths optimised for each storage

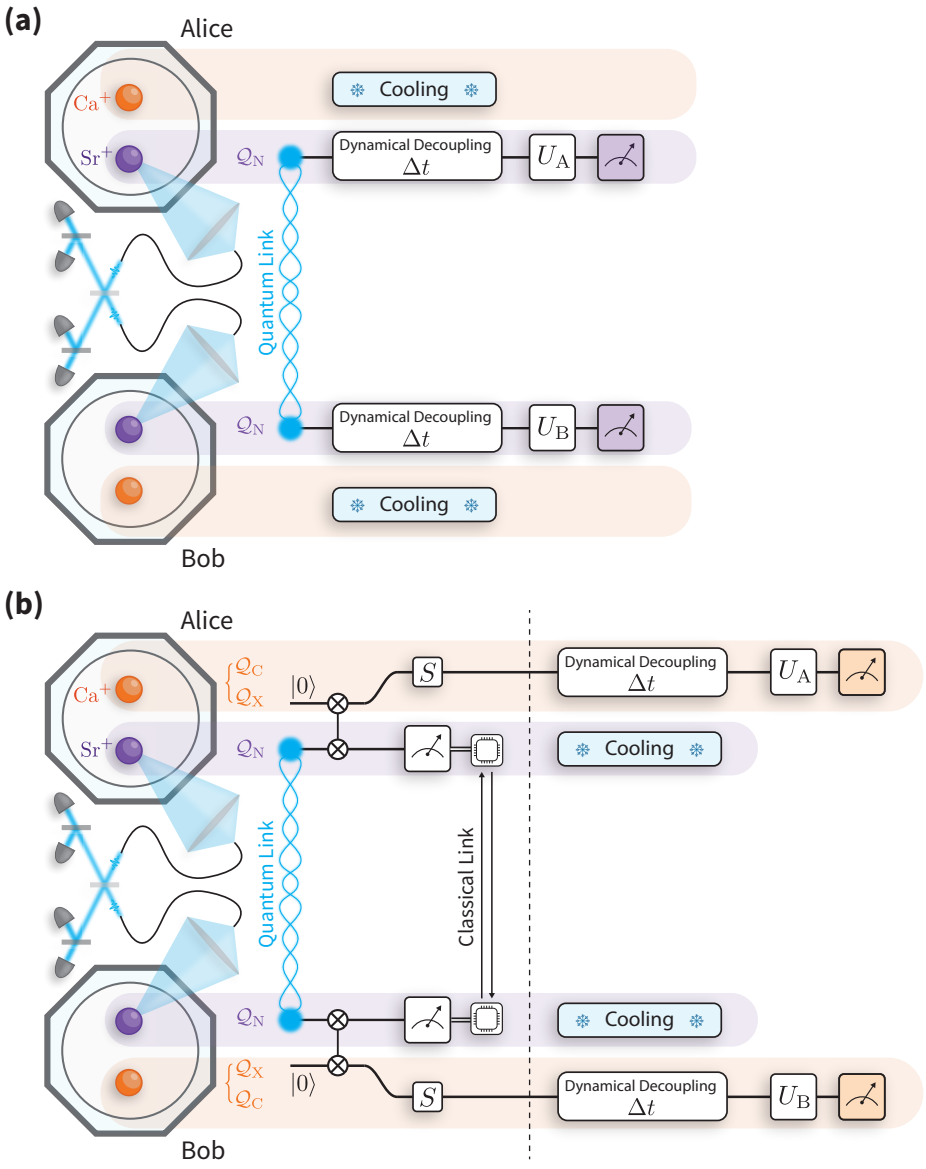

Figure 6.12: Experimental sequences for measuring the remote entanglement fidelity as a function of storage duration between **(a)** network qubits and **(b)** circuit qubits. Both sequences start by generating entanglement between two network qubits. In **(b)**, this entanglement is transferred to the circuit qubits via error-detected iSWAP gates to the auxiliary qubits, followed by hyperfine qubit transfers. The error detection outcomes are exchanged via a classical (TTL) link, and if no error is detected, the modules proceed to the point marked by the dashed line. Once remote entanglement is established between the **(a)** network [**(b)** circuit] qubits, it is stored for a variable duration Δt. During this time, sympathetic cooling is applied to the ion species not used for storage, and dynamical decoupling suppresses dephasing. Finally, a parity/population measurement is used to estimate the fidelity of the stored entanglement after Δt.

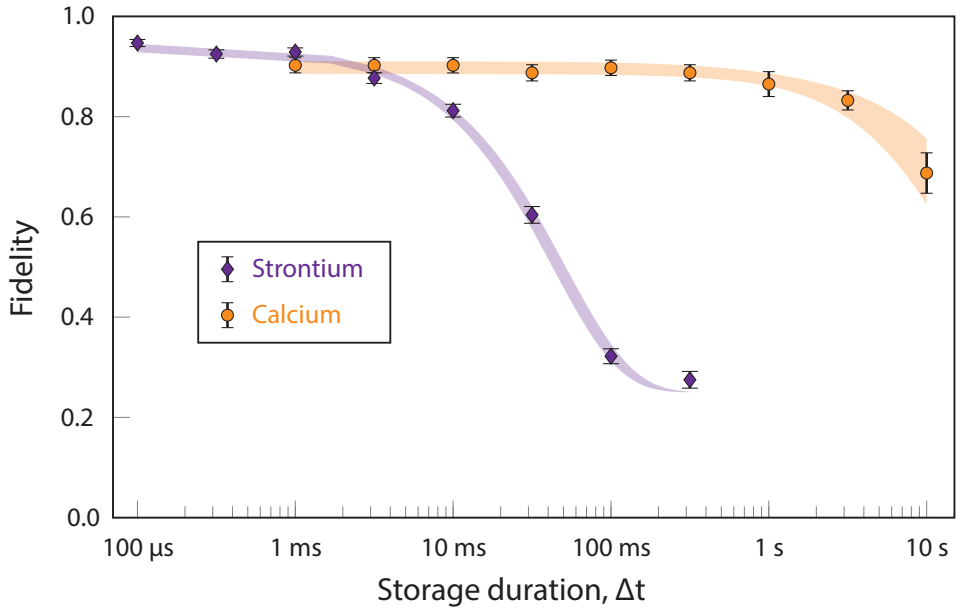

Figure 6.13: Remote entanglement fidelity as a function of storage duration for entanglement between the network qubits (purple) and circuit qubits (orange). For the circuit qubits, we observe a remote entanglement fidelity of 69(4) % at 10 s storage duration. The shaded areas indicate 95 % confidence intervals for exponential decay fits. For storage of entanglement between the network (circuit) qubits, we extract an exponential decay time-constant of 46(2) ms (25(6) s).

duration, ranging from 4 pulses (UR4) to 48 pulses (UR48). Additionally, over this duration, the ion species not used to store the entanglement is used as a sympathetic coolant to prevent ion loss and mitigate single-qubit rotation errors due to heating. Finally, we perform parity and population measurements to estimate the fidelity of the stored entanglement after storage.

The fidelities of the stored entanglement as a function of storage duration for the two qubits is shown in Figure 6.13. For the network qubit, we observe a significant degradation in the stored entanglement for storage durations exceeding 10 ms. We attribute this to the relatively high sensitivity of the network qubit to magnetic field fluctuations ($-11.2\,\mathrm{MHz\,mT^{-1}}$). Additionally, we observe that the fidelity of the network qubit entanglement decays to ~ 0.25 at storage durations exceeding 100 ms, indicating the presence of amplitude damping processes such as the decay of the $|1_N\rangle$ state due to the natural lifetime (≈ 390 ms) of the 674 nm transition[4]. We fit

[4]Note that the Φ^+ Bell state will decay twice as fast.

an exponential decay[5] to the fidelities, extracting a decay time-constant of 46(2) ms. As discussed previously, it takes $\mathcal{O}(10\,\text{ms})$ to generate remote entanglement between network qubits (see Section 6.2), and this increases to $\mathcal{O}(100\,\text{ms})$ when additional cooling is required for the mixed-species experiments (see Section 6.3). It is therefore clear that in our apparatus, the network qubit alone is not sufficient for applications required multiple instances of remote entanglement generation.

For the circuit qubits, we observe a marked enhancement of the storage capabilities, observing a remote entanglement fidelity of 69(4) % after 10 s. Using the same exponential decay model as above, we extract a decay time-constant of 25(6) s. This significantly exceeds the average time taken to generate entanglement in the mixed-species experiments ($\mathcal{O}(100\,\text{ms})$), thereby demonstrating the utility of the circuit qubits for storing quantum information for applications involving multiple instances of remote entanglement generation. This ability is key in our demonstration of DQC, presented in the next chapter.

6.5 Summary

In the first part of this chapter, we outlined the process of generating entanglement between a Sr^+ ion and a 422 nm photon, thereby establishing a mechanism for the modules to interface with the optical quantum network. A projective Bell measurement of these photons swaps the entanglement from two ion-photon pairs to the ions, heralding the creation of remote entanglement between Sr^+ network qubits in two separate modules. This remote entanglement, with an entanglement fidelity of 96.89(8) % to the $|\Psi^+\rangle$ state, is a crucial resource for our quantum networking experiments.

Previously, we integrated a robust quantum memory, provided by a Ca^+ ion, into one of our networking modules [123]. In the second part of this chapter, we discuss the integration of Ca^+ ions into both modules, using tools discussed in Chapter 5.

[5]Note that the exponential decay is a greatly simplified model and does not encapsulate the various error channels, each with their own decay profiles and time-constants.

This enabled the creation of more complex entangled states, such as the remotely entangled state of two ions of different species, the transfer of remote entanglement into the Ca^+ auxiliary qubits, and the creation of mixed-species GHZ states. The fidelities of these created states all exceeded 90 %

While the primary motivation for generating these states was to demonstrate the capabilities of our mixed-species trapped-ion quantum network, such states would have application in fields such as sensing and metrology. For example, in previous work, we demonstrated entanglement enhanced frequency comparisons of two Sr^+ clocks [132]. However, the precision of these comparisons was limited by magnetic field-induced dephasing. By mapping this entanglement to Ca^+, we gain access to a magnetic field-insensitive optical (729 nm) transition, allowing us to probe close to the lifetime limit of the Ca^+ $D_{5/2}$ manifold (~ 1 s), significantly enhancing the measurement precision. This experiment would have exciting implications, paving the way for a quantum network of atomic clocks [136], and thus is being actively pursued in our group. The ability to generate the mixed-species entangled state, such as the Sr^+-Ca^+ Bell state and the 3- and 4-qubit GHZ states, could have similar metrological applications, such as frequency comparisons within a heterogeneous clock network.

Finally, we demonstrated the storage capabilities of the long-lived circuit qubits, observing the storage of remote entanglement for durations up to 10 s. This extends previous work [123], where these circuit qubits were integrated into one module and shown to be robust against network activity. By preserving the state of the circuit qubits for durations significantly exceeding the time taken to generate the remote entanglement, we open the door to more complex circuits which require multiple rounds of remote entanglement generation.

Combining the tools for controlling mixed-species trapped-ion modules, as presented in Chapter 5, with the ability to generate remote entanglement between the two modules, marks a key milestone in quantum networking experiments. This development enables exciting applications, such as entanglement-enhanced frequency

comparisons of Ca^+ clock transitions, entanglement distillation, and, as will be discussed in the next chapter, distributed quantum computing.

7

Distributed Quantum Computing

*

This chapter presents the main results of this thesis: the demonstration of DQC between two mixed-species trapped-ion modules. I led the experimental work, performed the data analysis, and served as primary author on the related publication [137]; where appropriate, text and figures from that publication have been reproduced verbatim.

The chapter begins with a brief overview of quantum teleportation, followed by a detailed description of the experimental implementation and characterisation of a teleported CZ gate between circuit qubits in separate modules. Together with arbitrary single-qubit operations, this teleported CZ completes a universal gate set for a two-qubit distributed quantum computer. These capabilities are then used to realise more complex distributed circuits—the iSWAP and SWAP gates—and to implement Grover's algorithm on the distributed system.

7.1 Quantum teleportation

As discussed in Section 1.2.1, quantum teleportation protocols – characterised by the resources of shared entanglement and LOCC – are critical for the DQC architecture,

enabling a network of smaller quantum processing modules to function as a single, fully intra-connected quantum computer. It is an interesting, and indeed counter-intuitive, result that one can used pre-shared entanglement in order to effectively replace quantum communication channels. To explore how quantum teleportation can be used to connect two modules, let us first discuss the original teleportation protocol – quantum state teleportation.

7.1.1 Quantum state teleportation

Proposed in 1993 by Bennett *et al.* [28], quantum state teleportation enables the faithful transmission of quantum information between two modules, even when the quantum channels interconnecting the modules are lossy and noisy. The modules, Alice and Bob, can perform ideal local operations, exchange classical information, and share a pair of network qubits that are in the maximally entangled Bell state[1],

$$\left|\Phi^+\right\rangle_{\mathrm{AB}} = \frac{1}{\sqrt{2}} \sum_{i=0,1} \left|i\right\rangle_{\mathrm{A}} \left|i\right\rangle_{\mathrm{B}} . \tag{7.1}$$

Notice that we can write the other Bell states in terms of the $|\Phi^+\rangle$ state as $\hat{\sigma}_z^b \hat{\sigma}_x^a \left|\Phi^+\right\rangle$, such that the bit string $ab \in \{00, 01, 10, 11\}$ encodes a particular Bell state.

Suppose now Alice has a circuit qubit in some unknown state, $|\psi\rangle$, and would like to send it to Bob using quantum state teleportation, depicted in Figure 7.1(a). Alice and Bob each share half of an entangled state, such that the total state can be written as $\left|\psi\right\rangle_{\mathrm{A}} \left|\Phi^+\right\rangle_{\mathrm{AB}}$. Alice now measures her qubits in the Bell basis, yielding the bit string ab which corresponds to the Bell state measurement outcome $\hat{\sigma}_z^b \hat{\sigma}_x^a \left|\Phi^+\right\rangle$. Alice therefore projects Bob's half of the Bell pair into the (unnormalised) state

$$\left[\left(\left\langle\Phi^+\right|_{\mathrm{AA}} \hat{\sigma}_x^a \hat{\sigma}_z^b\right) \otimes \hat{\mathbb{I}}_{\mathrm{B}}\right] \left|\psi\right\rangle_{\mathrm{A}} \left|\Phi^+\right\rangle_{\mathrm{AB}} . \tag{7.2}$$

Expanding $|\Phi^+\rangle$ in terms of the basis states, as in Eq. (7.1), it is straightforward to

[1] Of course, if the interconnecting quantum channel is noisy, some form of entanglement distillation would be required to reach high fidelities.

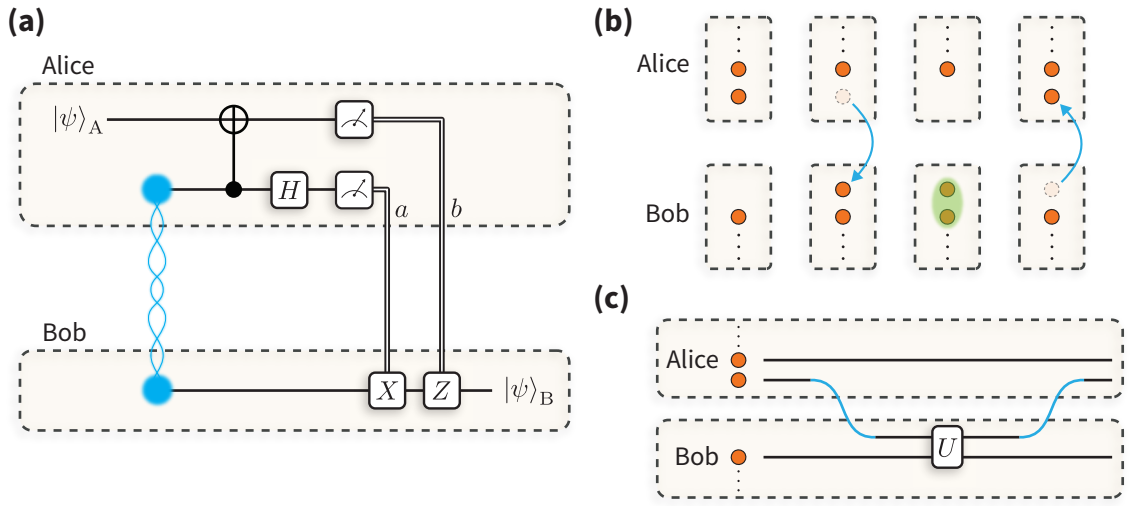

Figure 7.1: **(a)** Circuit for quantum state teleportation. Alice teleports the state $|\psi\rangle_A$ to Bob using shared entanglement, in addition to local operations and classical communication. **(b)** Visualisation of a distributed quantum computing architecture based on state teleportation. Two modules, Alice and Bob, can perform an arbitrary two-qubit unitary operation, $\hat{U}$, between qubit from each module by teleporting a qubit from Alice to Bob (indicated by the blue arrow), locally implementing the desired unitary, and teleporting the qubit back to Alice. **(c)** Equivalent circuit diagram for **(b)**, highlighting the use of teleportation to provide connectivity between qubits in separate modules.

find that the state we project into can be written as

$$\frac{1}{2}\sum_{i=0,1} \langle i|_A \, \hat{\sigma}_x^a \hat{\sigma}_z^b \, |\psi\rangle_A \, |i\rangle_B = \frac{1}{2}\hat{\sigma}_x^a \hat{\sigma}_z^b \, |\psi\rangle_B \, . \tag{7.3}$$

This tells us that when Alice performs the Bell state measurement, Bob's network qubit is projected into the state $\hat{\sigma}_x^a \hat{\sigma}_z^b \, |\psi\rangle_B$, with a probability of $\frac{1}{4}$. Alice then transmits the measurement outcome, i.e., the classical bit string ab, to Bob. Bob can then use this information to perform the correction operation $\left(\hat{\sigma}_x^a \hat{\sigma}_z^b\right)^\dagger$, and thus obtain the unknown state $|\psi\rangle$ sent by Alice.

This result has a number of interesting implications. Firstly, each measurement outcome ab occurs with equal probability, and crucially does *not* depend in any way on the state $|\psi\rangle$ that Alice was trying to send. If this measurement outcome did depend on $|\psi\rangle$ in some way, then Alice would learn some information about the state, which would in turn affect the unknown state she was trying to send. Secondly, before Bob receives the classical bit string ab from Alice, he has no idea

what the measurement outcome was. Therefore, from Bob's point of view, his half of the Bell pair is described by the statistical mixture

$$\hat{\rho}_{\mathrm{B}} = \frac{1}{4} \sum_{a,b=0,1} \hat{\sigma}_x^a \hat{\sigma}_z^b \left|\psi\right\rangle \left\langle\psi\right| \hat{\sigma}_z^b \hat{\sigma}_x^a, \tag{7.4}$$

which is a maximally mixed state. This means that until Bob receives the classical information, he cannot gain any information about the state $|\psi\rangle$. In other words, quantum teleportation does not allow superluminal communication.

This result illustrates how shared entanglement between two modules can enable the transfer of quantum information between them. One can then consider applying this to the distributed quantum computing architecture, as outlined in Figure 7.1(b) and (c), where teleportation could be used to bring two separated qubits into the same local register, locally implementing the desired two-qubit operation, and finally teleporting one of the qubits back into the original module. Thus, two instances of teleportation can mediate any arbitrary two-qubit operation between qubits in separate modules. Gottesman and Chuang [29] proposed a technique for implementing universal quantum computation based on quantum teleportation – demonstrating a circuit for teleporting two qubits through a CNOT gate, requiring two Bell pairs.

7.1.2 Quantum gate teleportation

Entanglement is an important resource for quantum computers, so we must ask the question: what is the most *efficient* use of this resource? Eisert *et al.* [30] answered this by showing that one Bell pair and the exchange of one classical bit in each direction are necessary and sufficient resources for implementing a non-local CZ[2] gate on two separated qubits. This optimal implementation of a non-local gate is referred to as *quantum gate teleportation (QGT)* – where instead of the teleportation of a quantum state, we can view this process as the teleportation of the *interaction* between the two qubits.

[2]Note that Eisert *et al.* [30] considered this problem in terms of a CNOT gate. However, the CNOT gate is equivalent to the CZ gate up to single-qubit rotations.

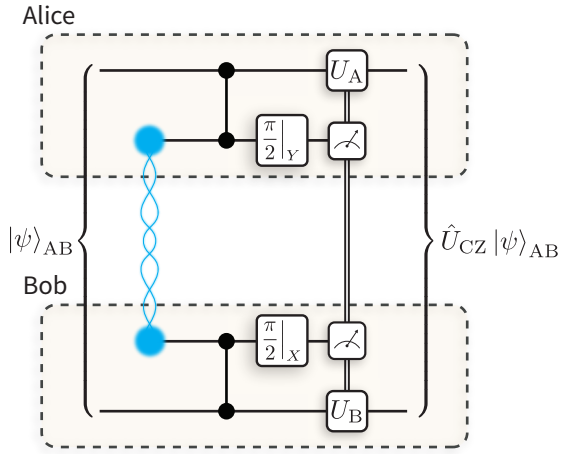

Figure 7.2: Circuit for quantum gate teleportation between two modules, Alice and Bob. Using only one shared Bell state, in addition to local operations and classical communication, the two modules mediate a CZ gate between a qubit from each module.

The fact that at least one Bell pair is necessary for the implementation of a non-local CZ gate can be seen from the fact that a CZ gate can be used to create one maximally entangled Bell state from an unentangled product state. Since the amount of entanglement between the two separated qubits cannot be increased using only LOCC, we must consume at least one Bell state in the realisation of the non-local gate. The requirement of at least one bit of classical communication in each direction once again comes down to the fact that if one could implement a non-local CZ gate with less than this, it would enable superluminal communication. Thus we have a minimal set of resources that are required to implement a non-local CZ gate.

To show that these resources are also sufficient, we construct a circuit which implements a non-local CZ gate, as shown in Figure 7.2. Let Alice and Bob each hold a circuit qubit, such that the state of the two circuit qubits is written generally as $|\psi\rangle_{AB}$. Now suppose that they also share a pair of network qubits in a maximally entangled Bell state, $|\Phi^+\rangle_{AB}$, as in Eq. (7.1). The total state is given by the product state $|\Phi^+\rangle_{AB}|\psi\rangle_{AB}$. Alice and Bob now perform a CZ gate between their half of the Bell state and their circuit qubit. The action of these bilateral CZ gates can be

written as

$$\frac{1}{\sqrt{2}}\sum_{j=0,1}|j\rangle_{\mathrm{A}}\,|j\rangle_{\mathrm{B}}\,|\psi\rangle_{\mathrm{AB}}\rightarrow\frac{1}{\sqrt{2}}\sum_{j=0,1}|j\rangle_{\mathrm{A}}\,|j\rangle_{\mathrm{B}}\left(\hat{\sigma}_z^j\otimes\hat{\sigma}_z^j\,|\psi\rangle_{\mathrm{AB}}\right). \tag{7.5}$$

Alice now makes a measurement of her network qubit in the X-basis and obtains the outcome a, and Bob makes a measurement of his qubit in the Y-basis and obtains the outcome b. This measurement projects the circuit qubits into the (unnormalised) state

$$\frac{1}{2\sqrt{2}}\left[\hat{\mathbb{I}}+i(-1)^{a\oplus b}\hat{\sigma}_z^j\otimes\hat{\sigma}_z^j\right]|\psi\rangle_{\mathrm{AB}}\,. \tag{7.6}$$

This can then be written in terms of a CZ gate acting on the two circuit qubits as

$$\frac{1}{2}\left[\hat{U}_{\mathrm{A}}(a\oplus b)\otimes\hat{U}_{\mathrm{B}}(a\oplus b)\right]\hat{U}_{\mathrm{CZ}}\,|\psi\rangle_{\mathrm{AB}}\,, \tag{7.7}$$

where $\hat{U}_i\,(x)$ are single-qubit rotations on the circuit qubit in module i, conditioned on the bit x, given by

$$\hat{U}_i\,(x)=\begin{cases}\hat{S},\text{ if }x=0,\\ \hat{S}^\dagger,\text{ otherwise,}\end{cases} \tag{7.8}$$

where $\hat{S}=\mathrm{diag}\,(1,i)$, as defined in Section 2.4.2. Therefore, in order to complete the teleportation of the CZ gate, Alice and Bob must exchange their measurement outcomes, and use this information to condition single-qubit rotations to invert $\hat{U}_{\mathrm{A}}(a\oplus b)\otimes\hat{U}_{\mathrm{B}}(a\oplus b)$. By applying these single-qubit rotations, Alice and Bob complete the teleportation of the CZ gate, such that the final state of the circuit qubits is given by $\hat{U}_{\mathrm{CZ}}\,|\psi\rangle_{\mathrm{AB}}$.

As in the case of quantum state teleportation, the four possible measurement outcomes ab all occur with equal probability. If this were not the case, then sending the measurement outcomes would correspond to the exchange of less than one bit in each direction[3] which, as discussed in Eisert *et al.* [30], is not sufficient for

[3]For example, consider the extreme case that only one measurement outcome occurs. Here, exchanging the outcome does not provide the modules with any new information, since they already had information about what the measurement outcome would be in advance.

implementing a non-local CZ gate. Additionally, the teleportation protocol is not completed until the modules have each received, and made use of, the classical information sent from the other module.

This gate teleportation protocol realises the optimal implementation of a non-local CZ gate. Of course, this is just one class of entangling gate, with up to three CZ gates, in addition to arbitrary single-qubit rotations, required to implement any arbitrary two-qubit unitary operation [138]. Therefore, one would assume that performing an arbitrary non-local two-qubit operation, would require up to shared entangled states. However, as pointed out by Eisert *et al.* [30], one always has the option of using two rounds of teleportation and a local implementation of the arbitrary operation, as discussed at the end of the last section. We therefore conclude that at most two Bell states are sufficient for realising arbitrary two-qubit operations. There has been ongoing work developing techniques for compiling distributed circuits that are efficient in their use of entangled resource states; however, this is beyond the scope of this thesis.

7.2 Teleportation of a controlled-Z gate

In the previous section, we illustrated how entanglement could be used to logically connect qubits in separate modules through QGT. QGT therefore represents a critical component of a distributed quantum computer. In this section, we present our experimental implementation and characterisation of a teleported CZ gate between two circuit qubits located in separate trapped-ion modules.

7.2.1 Experimental implementation

The QGT protocol used here to mediate CZ gates between the circuit qubits in separate modules is depicted in Figure 7.3. We allow the circuit qubits to start in an arbitrary state $\left|\psi_{\text{in}}^{\text{AB}}\right\rangle \in {\mathcal{Q}_{\text{C}}}^{\otimes 2}$, which could be part of a larger, long-running computation. We begin the QGT protocol by generating the remotely entangled

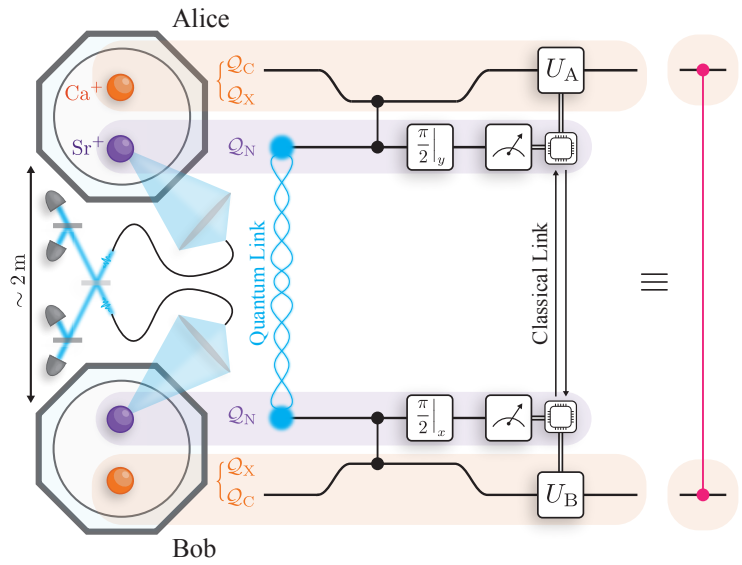

Figure 7.3: Experimental sequence for the teleportation of a CZ gate. The two modules, Alice and Bob, each hold a Sr^+ ion (purple) and a Ca^+ ion (orange). Sr^+ provides a network qubit, $\mathcal{Q}_N$, while Ca^+ provides both a long-lived circuit qubit, $\mathcal{Q}_C$, and an auxiliary qubit, $\mathcal{Q}_X$. Prior to the protocol, the circuit qubits are in some arbitrary state. The protocol begins by generating entanglement between the network qubits via a photonic link. Upon heralding entanglement, each module applies a local CZ gate between the network and circuit qubits, using the auxiliary qubit temporarily to mediate the gate mechanism. The outcomes of mid-circuit parity measurements of the network qubits are exchanged in real-time via a classical (TTL) link connecting the control systems of the two modules. This information is used to condition local feed-forward operations, U_A and U_B, on the circuit qubits – completing the teleportation of the CZ gate.

Bell state,

$$|\Psi^+\rangle = \frac{|10\rangle + |01\rangle}{\sqrt{2}} \in \mathcal{Q}_N^{\otimes 2},$$

between the network qubits. Note that this is different to the $|\Phi^+\rangle$ state used in the discussion presented in Section 7.1, however, these states are equivalent up to single-qubit rotations and our remote entanglement generation scheme naturally generates $|\Psi^+\rangle$ states. Since this entanglement is generated via a try-until-success process, the circuit qubits provide a robust quantum memory until a herald indicates a success (see Section 7.2.3).

At this stage, we map the state stored in the circuit qubits ($\mathcal{Q}_C$) to the auxiliary qubits ($\mathcal{Q}_X$) in preparation for the local entangling operations, see 5.3. In each module, we perform local CZ gates between the network and auxiliary qubits (see Section 5.2), before transferring the auxiliary qubit back to the circuit qubit. We

then perform mid-circuit measurements of the network qubits in the X and Y bases in Alice and Bob, respectively. The mid-circuit measurement outcomes, $m_{\mathrm{A}}, m_{\mathrm{B}} \in \{0, 1\}$, are exchanged in real-time via a classical communication channel between the modules – in our demonstration, this is a TTL link connecting the control systems of the two modules. Following the exchange of the measurement outcomes, the modules, Alice and Bob, perform the conditional rotations $\hat{U}_{\mathrm{A}}$ and $\hat{U}_{\mathrm{B}}$, respectively, where

$$\hat{U}_{\mathrm{A}} \otimes \hat{U}_{\mathrm{B}} = \begin{cases} \hat{S}^{\dagger} \otimes \hat{S} & \text{if } m_{\mathrm{A}} \oplus m_{\mathrm{B}} = 0, \\ \hat{S} \otimes \hat{S}^{\dagger} & \text{if otherwise.} \end{cases} \tag{7.9}$$

(7.10)

This completes the non-local gate $\left|\psi_{\mathrm{in}}^{\mathrm{AB}}\right\rangle \rightarrow \hat{U}_{\mathrm{CZ}}^{\mathrm{AB}} \left|\psi_{\mathrm{in}}^{\mathrm{AB}}\right\rangle$.

This QGT protocol is completely self-contained: the input states of the circuit qubits are set prior to the execution of the non-local gate, and the output states are available for further computation. In addition to arbitrary single-qubit rotations of the circuit qubits, this teleported CZ gate is a key element of a gate set for DQC, enabling the modules to act as a single, fully-connected universal quantum processor.

7.2.2 Characterisation

We characterise the action of the teleported CZ gate using QPT, as outlined in Section 3.3. From a total of 16 000 tomographic measurements, we reconstruct the process matrix for the teleported gate, as shown in Figure 7.4(a). From this process matrix, we calculate an average gate fidelity of 86.2(9) %, compared to the ideal CZ gate. The average gate duration was ~ 103 ms, which was almost entirely dominated by the time taken to generate the remote entanglement, as we discuss in the next section. This result marks the first deterministic demonstration of a non-local two-qubit gate across a quantum network.

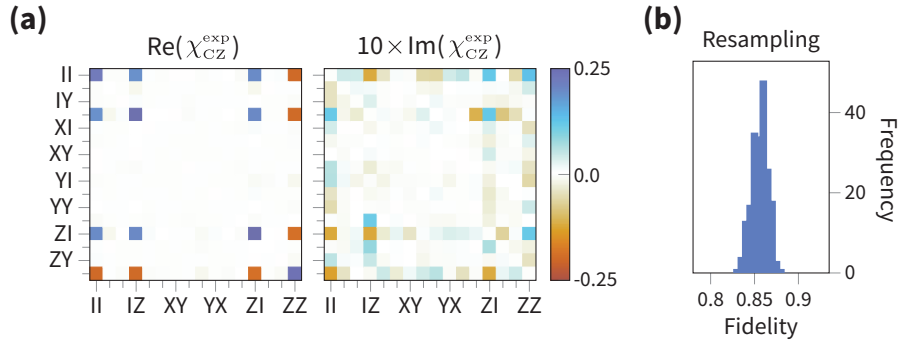

Figure 7.4: Characterisation of the teleported CZ gate **(a)** Process matrix for the CZ gate, reconstructed from 16 000 tomographic measurements. The reconstructed process matrix represents an average gate fidelity of 86.2(9) %, with respect to the ideal CZ gate. **(b)** Histogram of gate fidelities for process matrices calculated from resampled datasets. The uncertainty quoted for the gate fidelity represents one standard deviation of the resampled gate fidelities.

7.2.3 Error budget

The average gate fidelity we achieve is significantly below any threshold that would be required for, e.g., fault tolerant distributed quantum computing. Here, we perform a systematic discussion of the various error sources that contribute to the degradation of the gate fidelity. A summary of the estimated error contributions is given at the end in Table 7.1.

Raw entanglement

The heralded generation of remote entanglement between network qubits in separate modules, outlined in Chapter 6, is central to our QGT protocol and is depicted in Figure 7.5(a). Successful generation of entanglement is heralded by particular detector click patterns and, after subsequent local rotations, indicates the creation of the maximally entangled Ψ^+ Bell-state,

$$\left|\Psi^+\right\rangle = \frac{|10\rangle + |01\rangle}{\sqrt{2}} \in \mathcal{Q}_{\mathrm{N}}^{\otimes 2}. \tag{7.11}$$

This process is executed while simultaneously storing quantum information in the circuit qubits which, as we have shown in Section 6.4, can faithfully store quantum information for significantly longer than it takes to generate remote entanglement.

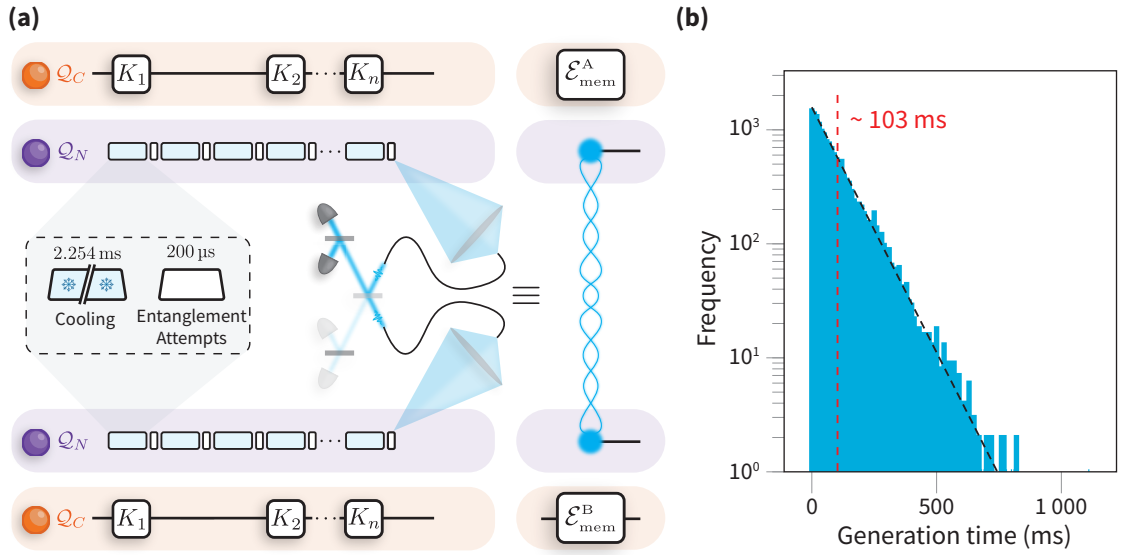

Figure 7.5: **(a)** Overview of entanglement generation with circuit qubit storage. Entanglement is generated between the network qubits using 200 µs of entanglement attempts interleaved with 2.254 ms of sympathetic recooling using the Sr^+ ion. This is repeated until the entanglement is successfully heralded by a particular detector click pattern. While attempting to generate entanglement between the network qubits, Knill dynamical decoupling pulses, K_i, are used to preserve the state of the circuit qubits. **(b)** Histogram of the time taken to successfully herald the generation of entanglement. The red dashed line at 103 ms indicates the average time taken to successfully generate entanglement between the network qubits.

On average it took 7084 attempts to successfully herald entanglement, corresponding to a success probability of 1.41×10^{-4}. To mitigate heating of the ion-crystal, we interleave 200 µs of entanglement generation attempts with 2.254 ms of sympathetic re-cooling of the Sr^+-Ca^+ crystal with the Sr^+ ion. The sympathetic recooling comprises 1.254 ms of Doppler cooling, followed by 1 ms of EIT cooling. The time taken to generate the remote entanglement for each tomographic measurement is shown in Figure 7.5(b). On average, it took 103 ms to generate entanglement between the network qubits, corresponding to an entanglement generation rate of $9.7\,s^{-1}$, although this rate could be increased by optimising the interleaved cooling sequence. As discussed in Section 6.2, we are able to generate Ψ^+ Bell states between the two modules with a fidelity of 96.89(8) %.

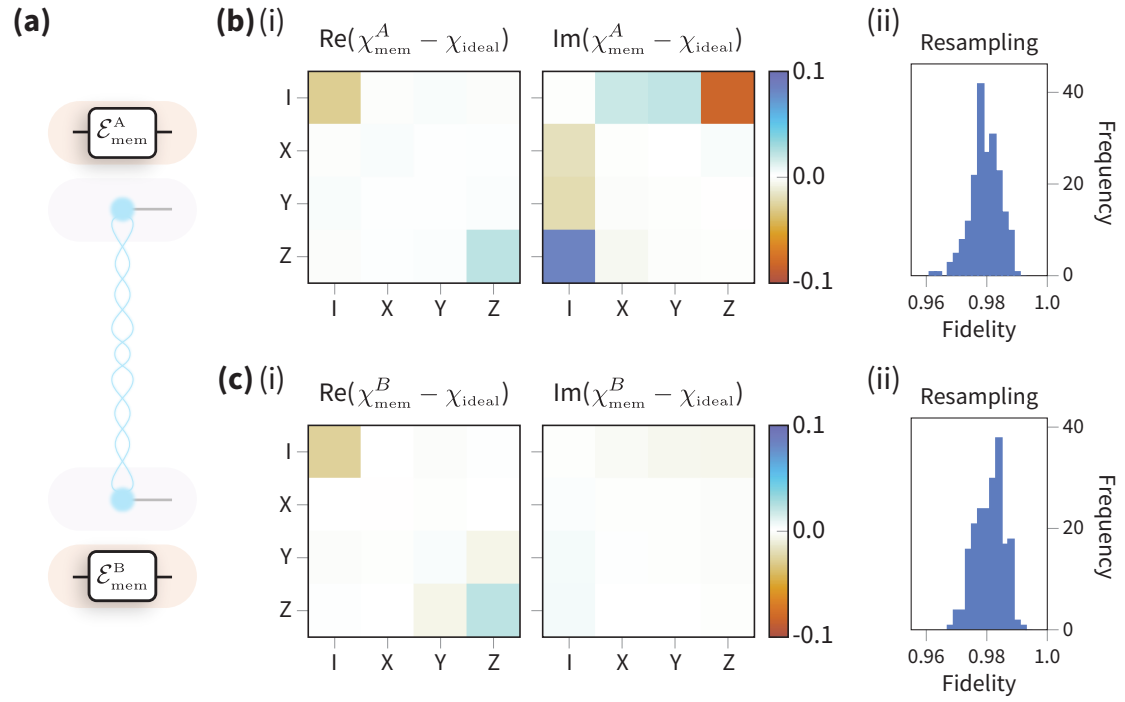

Figure 7.6: **(a)** The quantum channels, $\mathcal{E}^{\mathrm{A}}_{\mathrm{mem}}$ and $\mathcal{E}^{\mathrm{B}}_{\mathrm{mem}}$, are the processes of storing the states of the circuit qubits in Alice and Bob, respectively, while generating remote entanglement between the two modules. **(b)**(i) and **(c)**(i) show the residuals of the process matrices representing the storage processes $\mathcal{E}^{\mathrm{A}}_{\mathrm{mem}}$ and $\mathcal{E}^{\mathrm{B}}_{\mathrm{mem}}$, in Alice and Bob, respectively. The residuals are taken with respect to the ideal storage process matrix, i.e., $\hat{\chi}_{\mathrm{mem}} - \hat{\chi}_{\mathrm{ideal}}$, where $\hat{\chi}_{\mathrm{ideal}} = \mathrm{diag}\,(1,0,0,0)$. **(b)**(ii) and **(c)**(ii) show histograms of the fidelities of resampled datasets, for Alice and Bob, respectively. Compared to the ideal identity process, the reconstructed process matrices yield average gate fidelities of 98.1(4) % and 98.2(5) % for Alice and Bob, respectively.

Circuit qubit storage

Since each instance of QGT requires the generation of entanglement between network qubits, it is necessary to ensure that the circuit qubits preserve their encoded quantum information during this process. Due to their low sensitivity to magnetic field fluctuations, the circuit qubits have exhibited ~ 100 ms coherence times, and in previous work we demonstrated these qubits to be robust to network activity [123]. As in Section 6.4, we further suppress dephasing through dynamical decoupling.

Typically, dynamical decoupling is implemented over a fixed period of time, however the success of the entanglement generation process is non-deterministic and would therefore leave the dynamical decoupling sequence incomplete. One solution would be to complete the dynamical decoupling pulse sequence once the entan-

glement has been generated, however it is desirable to minimise the time between heralding the entanglement generation and performing the QGT protocol, in order to prevent dephasing of the network qubits. Instead, we make use of the fact that the action of a dynamical decoupling pulse on one of the circuit qubits can be propagated through the teleported CZ gate as

$$(\hat{\sigma}_x \otimes \hat{\mathbb{I}})\hat{U}_{\mathrm{CZ}} = \hat{U}_{\mathrm{CZ}}(\hat{\sigma}_x \otimes \hat{\sigma}_z). \tag{7.12}$$

We therefore perform the dynamical decoupling pulses on the circuit qubits until successful herald of remote entanglement, at which point we immediately perform the QGT sequence – implementing a CZ gate on the state of the circuit qubits at the point of interruption. Once this gate is completed, we perform the remaining dynamical decoupling pulses (without any inter-pulse delay), and use Eq. (7.12) to apply the appropriate Z rotations required to correct for the propagation through the CZ gate. With this method, we suppress the dephasing errors in the circuit qubits during entanglement generation, while minimising the time between successfully heralding the entanglement and consuming it for QGT.

We deploy a Knill dynamical decoupling pulse sequence [19, 67] with a $\approx 7.4\,\mathrm{ms}$ inter-pulse delay (corresponding to a pulse every three rounds of interleaved entanglement attempts and re-cooling). This dynamical decoupling sequence was found to perform marginally better in this scenario than the universally robust dynamical decoupling sequence which was used for the static storage duration in Section 6.4. We use QPT to reconstruct the process of storing the quantum information while generating entanglement; ideally, this process would not alter the quantum information stored in the circuit qubit. QPT is implemented by choosing input states for the circuit qubits from a tomographically complete set, generating remote entanglement between the network qubits while dynamically decoupling the circuit qubits, then upon successful herald, completing the dynamical decoupling sequence and performing tomographic measurements of the circuit qubits. The reconstructed process matrices for each module, representing to the action of storing quantum information

during entanglement generation, are shown in Figure 7.6. We observe fidelities to the ideal operation of 98.1(4) % and 98.2(5) % for Alice and Bob, respectively.

Local entangling gates

The ability to perform logical entangling gates between ions of different species allows us to delegate the roles of network and circuit ions which is key to our demonstration of DQC. These local entangling gates are characterised in Section 5.2, where we measure average CZ gate fidelities of 97.6(2) % and 98.0(2) % for Alice and Bob, respectively.

Hyperfine transfer

Since the circuit qubit does not participate in the mixed-species gate, the gate interaction is performed on the network and auxiliary qubits. Consequently, we require the ability to coherently inter-convert between the circuit and auxiliary qubit before and after the local operations. The process by which we implement this mapping is discussed in see Section 5.3; we measure average error per transfers of $3.8(2) \times 10^{-3}$ and $2.6(1) \times 10^{-3}$ for Alice and Bob, respectively.

Mid-circuit measurements

To complete the QGT protocol, the two modules make mid-circuit measurements of the network qubits, exchange the measurement outcomes, and perform a local rotation of their circuit qubits conditioned on the outcomes of these measurements. Utilising the spectral isolation between the two species of ions, mid-circuit measurements of the network qubits can be made without affecting the quantum state of the circuit qubits. However, imperfect mid-circuit measurements of the network qubits would lead to the application of the wrong conditioned gates and thus contribute to the infidelity of the teleported gate.

From the form of the conditional rotations, Eq. (7.9), we see that an error in performing the parity measurement will result in a joint phase-flip on each qubit.

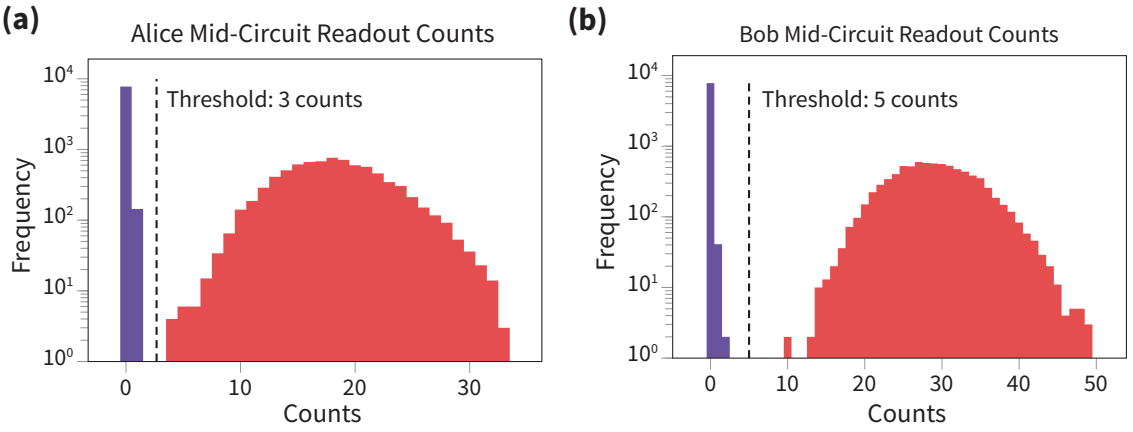

Figure 7.7: Photon detection histograms for the mid-circuit measurement of the network qubits in **(a)** Alice and **(b)** Bob. The mid-circuit measurement detection window was 500 µs. The counts below [above] the readout threshold (dashed line) were deemed to be the dark [bright] state. From these measurements, and the ≈ 391 ms lifetime of the $D_{5/2}$ manifold, we estimate fluorescence detection errors for the measurement of a dark (bright) state of $1.071(4) \times 10^{-3}$ $(1.44(6) \times 10^{-7})$ and $1.059(4) \times 10^{-3}$ $(5.0(2) \times 10^{-8})$ for Alice and Bob, respectively.

The process for a teleported CZ gate where the QGT was executed with imperfect mid-circuit measurements, can be described by the Kraus operators

$$\hat{K}_1 = \sqrt{1-p}\,\hat{U}_{CZ}$$
$$\hat{K}_2 = \sqrt{p}\;(\hat{\sigma}_z \otimes \hat{\sigma}_z)\,\hat{U}_{CZ},$$

where we define p as the average parity measurement error, i.e., the probability that a measurement of a state of a particular parity will yield an outcome of opposite parity. Using the fact that the average gate fidelity of a process represented by a set of Kraus operators, K_i, to the ideal CZ gate, $\hat{U}_{\mathrm{CZ}}$, can be written as [54]

$$\mathcal{F} = \frac{d + \sum_j \left|\mathrm{tr}\left(\hat{U}_{\mathrm{CZ}}^{\dagger}\hat{K}_j\right)\right|^2}{d(d+1)}, \tag{7.13}$$

we find that the teleported gate error due to imperfect mid-circuit measurements is

$$1 - \mathcal{F} = \frac{dp}{d+1}. \tag{7.14}$$

From Section 5.1, the single-qubit gate errors for the network qubits are $4.8(3) \times 10^{-4}$ and $9.8(3) \times 10^{-4}$ for Alice and Bob, respectively. The error in the fluorescence detection is estimated from the observed photon scattering rates of $\mathcal{Q}_N$ states during the mid-circuit measurement. The readout counts acquired in the 500 µs readout window is shown in Figure 7.7; by taking into account the $\approx$ 390 ms lifetime of the $|1_N\rangle$ state [139], we estimate fluorescence detection errors for the measurement of a dark (bright) state of $1.071(4) \times 10^{-3}$ ($1.44(6) \times 10^{-7}$) and $1.059(4) \times 10^{-3}$ ($5.0(2) \times 10^{-8}$) for Alice and Bob, respectively. Combining these error mechanisms, we estimate average mid-circuit measurement errors of $1.02(3) \times 10^{-3}$ and $1.51(3) \times 10^{-3}$. These measurement errors correspond to a contribution to the teleported CZ gate error of 0.082(3) % and 0.121(3) % for Alice and Bob, respectively.

Summary

We summarise the leading error sources affecting our teleported CZ gate in Table 7.1. The measured fidelity of our gate is slightly lower than that predicted by the error budget, which we attribute to drifts in the calibration of various components over the duration of the data acquisition. The majority of identified errors occur during local operations in each module. Our local errors do not represent the state-of-the-art for trapped-ion processors; however, local operations exceeding the fidelity threshold for fault-tolerant quantum computing have been demonstrated in this platform [13, 14, 15, 16, 17, 18, 140]. Relevant to our implementation, Hughes *et al.* [15] demonstrated mixed-species two-qubit gates between Sr^+ and Ca^+ ions with a gate error of 0.2(1) %. We therefore conclude that the technical limitations in our implementation can be overcome. The other significant source of error is the remote entanglement of the network qubits across the photonic quantum network; we observe a fidelity of the remotely entangled network qubits to the desired $|\Psi^+\rangle$ state of 96.89(8) %. Unlike the local operations, the performance of our remote entanglement is at the state-of-the-art. To improve this, and hence enable the teleportation of high-fidelity entangling gates between modules, entanglement distillation could

Source	Error	
	Alice	Bob
Raw entanglement	3.11(8) %	
Mixed-species gate	2.4(2) %	2.0(2) %
$\mathcal{Q}_C$ decoherence	1.9(4) %	1.8(5) %
$\mathcal{Q}_X \leftrightarrow \mathcal{Q}_C$ transfer	0.76(3) %	0.52(1) %
Mid-circuit measurement	0.082(3) %	0.121(3) %
$\mathcal{Q}_C$ rotations	0.010(3) %	0.012(4) %
Predicted total error	12.1(6) %	

Table 7.1: Error budget for CZ gate teleportation.

be used to distribute high-fidelity entangled states from a number of lower-fidelity entangled states [33, 141].

7.3 Universal distributed quantum computing

Almost 30 years ago, it was shown by DiVincenzo [142] that universal quantum computation could be achieved using only quantum gates that acted on up to two qubits. Following this result, Barenco *et al.* [143] showed that the CNOT gate, in addition to single-qubit operations, was sufficient for constructing arbitrary n-qubit unitary operations, and thus provided a universal gate-set for quantum computing. Since the CNOT gate and the CZ gate are equivalent up to single-qubit rotations, we construct a universal gate-set for our distributed quantum computer from the teleported CZ gate and single-qubit rotations. Of course, if more circuit qubits were available in each module, local CZ gates between these qubits would be required to complete the gate-set.

Until now, demonstrations of QGT have been limited to only a single instance of the non-local gate. In photonic platforms, this has been prevented by the inability to store the photons between interactions [47, 48]. For the QCCD demonstration by Wan *et al.* [36], the input states of the of the circuit qubits were prepared *after* the generation of entanglement between the network qubits; the $^9Be^+$ ions provid-

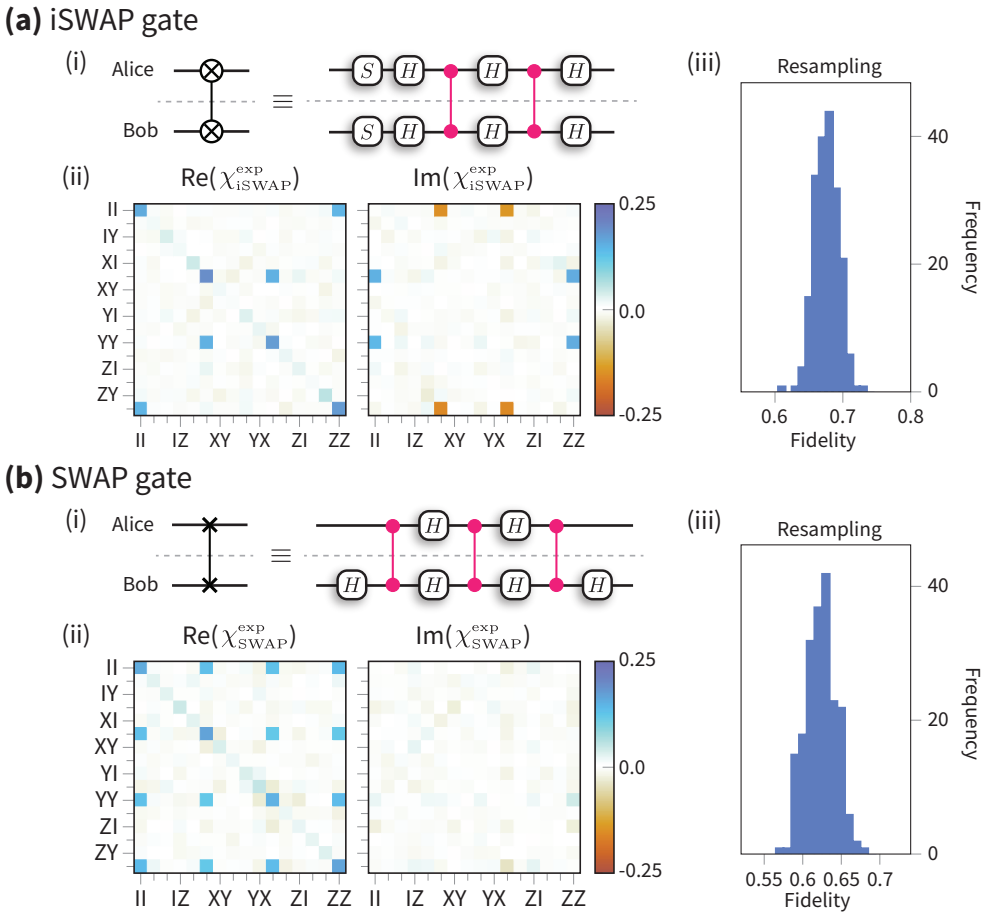

Figure 7.8: CZ decompositions of the distributed **a**(i) iSWAP and **b**(i) SWAP circuits, comprising two and three instances of QGT, respectively. The reconstructed process matrices for the **a**(ii) iSWAP and **b**(ii) SWAP gates indicate average gate fidelities of 70(2) % and 64(2) %, respectively. The histograms for the fidelities of resampled datasets are shown in **(a)**(iii) and **(b)**(iii).

ing the circuit qubits were used as a sympathetic coolant after the establishment of entanglement between trap zones. As a result, while an excellent proof-of-principle demonstration of QGT, this demonstration was unable to execute circuits comprising multiple teleported gates. In our QGT implementation, we prepare the states of the circuit qubit prior to the execution of the protocol and the output states are available for further computation. The self-contained nature of our implementation thus makes the extension to circuits comprising multiple instances of QGT straightforward.

We demonstrate our ability to perform sequential rounds of QGT by executing the CZ decompositions of the iSWAP and SWAP gates, which are shown in

Figures 7.8(a)(i) and 7.8(b)(i), and comprise two and three instances of QGT, respectively. As with the teleported CZ gate, we characterise these circuits via QPT, reconstructing the process matrices shown in Figures 7.8(b)(ii) and 7.8(b)(ii). We measure average gate fidelities of 70(2) % and 64(2) % for the iSWAP and SWAP gates, respectively. In general, any arbitrary two-qubit unitary operation can be decomposed into at most three CZ gates [138]. Therefore, by constructing circuits with up to 3 instances of QGT – enabled by our ability to perform QGT deterministically and on-demand – we demonstrate the ability to construct arbitrary two-qubit unitary operations, and thus perform universal DQC.

7.4 Grover's algorithm

Grover's algorithm, developed by Lov Grover in 1996 [6], was one of the first quantum algorithms conceived. Today, it is one of the classic examples of the capabilities of a quantum computer. As such, this algorithm is a exciting candidate for demonstrating the capabilities of our distributed quantum computer.

Otherwise known as the *quantum search algorithm*, this algorithm concerns searching through an unstructured set of items to identify some solution to a given problem. As framed by Grover, this could be looking up a name in a phone directory containing N names in a random order. Classically, the optimal strategy would be to compare the desired name to the entries in the phone directory one by one until the desired entry is found. On average, this would take $\mathcal{O}(N/2)$ accesses to the directory before the desired entry was found. The optimal quantum strategy – Grover's algorithm – makes use of the fact that a quantum system may be placed in superpositions, therefore enabling simultaneous accesses to the phone directory. While the problem given here is more illustrative than practical – indeed the typical example of database searching is unlikely to be a genuinely practical application – quantum search algorithms have potential application in a range of problems such as determining graph connectivity and pattern matching [144].

Let us first consider a formal statement of the problem. We consider items from

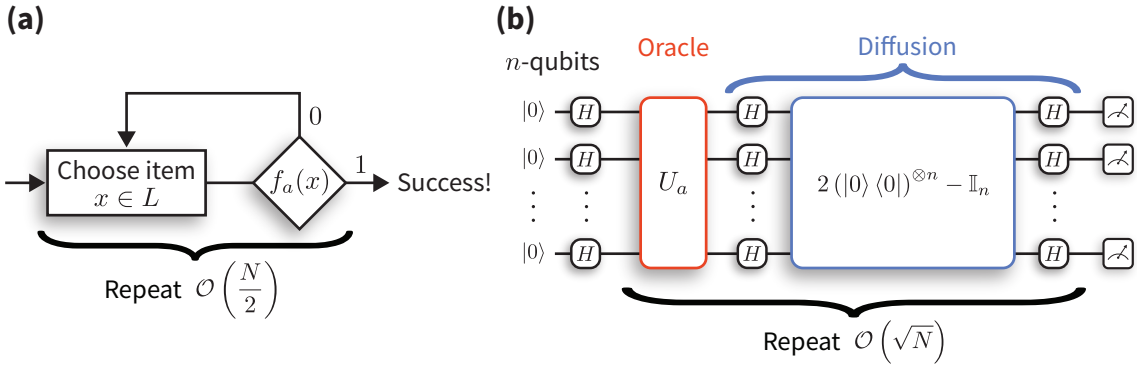

Figure 7.9: **(a)** Classical and **(b)** quantum algorithms for unstructured database searches. **(a)** The classical search algorithm comprises choosing an item, $x \in L$, and computing $f_a(x)$ to see if it is the solution. This is repeated until the solution is found, which occurs after $\mathcal{O}((N/2))$ repetitions, on average. **(b)** The quantum search algorithm, aka Grover's algorithm, comprises repetitions of the Grover oracle and diffusion circuits. After $\mathcal{O}\left(\sqrt{N}\right)$, a measurement of the register reveals the solution with probability $\mathcal{O}(1)$.

an unstructured set $x \in X$. We would like to identify some unique solution, $a \in X$, that satisfies a condition encoded by the function

$$f_a(x) = \begin{cases} 1 & \text{if } x = a, \\ 0 & \text{otherwise.} \end{cases} \tag{7.15}$$

The optimal classical strategy, outlined in Figure 7.9 (a), comprises evaluating this function for each input one by one until the solution is found; on average, this would require $\mathcal{O}(N/2)$ evaluations of f_a.

Grover's algorithm for identifying the solution, a, is shown in Figure 7.9 (b). For N items, the algorithm requires $n = \lceil \log_2 N \rceil$ qubits. The algorithm begins by preparing a superposition of all possible inputs with parallel Hadamard gates, implementing

$$|0\rangle^{\otimes n} \rightarrow \frac{1}{\sqrt{N}} \sum_{x=0}^{N-1} |x\rangle \,. \tag{7.16}$$

The function f_a can then be evaluated over this quantum register using the Grover *oracle*, $\hat{U}_a$, that implements the unitary operation

$$\hat{U}_a : |x\rangle \rightarrow (-1)^{f_a(x)} |x\rangle \,, \tag{7.17}$$

thereby marking the state $|a\rangle$ representing the solution to the problem. The Grover *diffusion* circuit comprises the unitary operation $2\left(|0\rangle\langle 0|\right)^{\otimes n} - \hat{\mathbb{I}}_n$ interleaving parallel Hadamard gates. An iteration of the oracle and diffusion circuits has the effect of increasing the amplitude of the marked state $|a\rangle$ by $\sim\mathcal{O}\left(1/\sqrt{N}\right)$; by performing $\mathcal{O}\left(\sqrt{N}\right)$ iterations, the amplitude of the state $|a\rangle$ therefore approaches $\mathcal{O}(1)$. A measurement of the register then reveals the marked state, $|a\rangle$, with probability $\mathcal{O}(1)$. This algorithm therefore enables the identification of the solution, a, with $\mathcal{O}\left(\sqrt{N}\right)$ implementations of the oracle $\hat{U}_a$, providing a polynomial speed-up over the best classical strategy which requires $\mathcal{O}(N/2)$ evaluations of the function $f_a(x)$.

7.4.1 Distributed Grover's algorithm

Our distributed quantum computer comprises a two-qubit register, distributed between the modules Alice and Bob. Using this register, we can perform Grover's algorithm for the set of items $X = \{00, 01, 10, 11\}$. For the two-qubit case, we require just a single iteration of the Grover oracle and diffusion circuits to achieve unit probability of revealing the marked state, assuming ideal operations; in the classical case, we would require 2 evaluations of f_a, on average. The circuit used to execute Grover's algorithm on our distributed computer, shown in Figure 7.10(a), comprises two instances of QGT. The first instance implements the oracle, while the second implements the diffusion circuit.

The results of Grover's algorithm – executed on our distributed quantum processor – are shown in Figure 7.10(b). For the marked states $a \in \{00, 01, 10, 11\}$, we obtain the correct result with an average success rate of $71(1)\,\%$. To our knowledge, this represents the first deterministic execution of any algorithm on a distributed quantum computer.

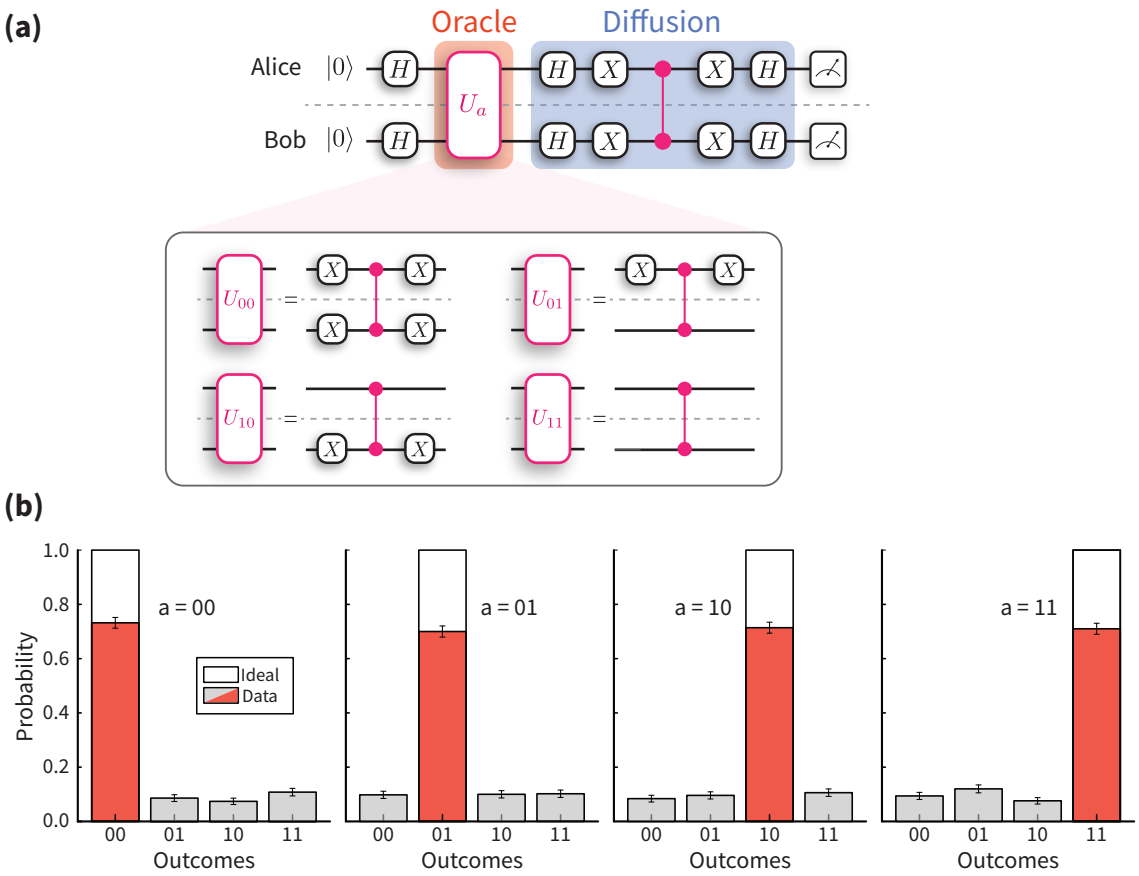

Figure 7.10: **(a)** Circuit for the two qubit Grover's algorithm, executed on our two-qubit distributed quantum computer. The circuit is compiled with two instances of QGT: the first implements the Grover oracle call which marks a particular state, a, while the second implements the diffusion circuit. The four different oracle circuits are shown in the box. **(b)** Measurement outcomes from 500 repetitions of Grover's algorithm per marked state. The average success probability is 71(1) %.

7.5 Summary

In summary, we have demonstrated the deterministic teleportation of CZ gates between two trapped-ion qubits across an optical quantum network. By leveraging storage capabilities of the long-lived circuit qubits, we are able to faithfully store quantum information while generating remote entanglement between network qubits. This enabled the implementation of distributed quantum circuits compiled with multiple teleported CZ gates, such as the iSWAP and SWAP gates, culminating in the execution of Grover's algorithm.

Our DQC implementation features a single circuit qubit in each module; how-

ever processors with larger numbers of qubits have been realised. With only 3 circuit qubits (and one network qubit) per module, the purification of arbitrary quantum channels would be possible [141]. The capabilities of the individual modules may be extended even further by deploying the QCCD architecture. With recent demonstrations in both academic research [145] and industry [37] highlighting the power of this approach, embedding these systems in a quantum network would combine their power with the reconfigurability and flexibility of the DQC architecture. Conversely, computational bottlenecks associated with ion transport overheads observed in the QCCD architecture [37] could be mitigated using photonic interconnects integrated into a single device [146].

While the results presented here were achieved using trapped-ion quantum processing modules, photons may be interfaced with a variety of systems. The connectivity and reconfigurability enabled by photonic networks provides a scalable approach for other quantum computing platforms such as diamond colour centres and neutral atoms. Additionally, modules of different platforms could be connected via wavelength conversion, enabling a hybrid DQC platform. Furthermore, teleportation protocols can be extended to higher-dimensional quantum computing paradigms, such as qudits [147] and continuous variable quantum computing (CVQC) [148, 149], allowing these platforms to benefit from the DQC architecture. Quantum repeater technology [32] would enable large physical separation between the quantum processing modules, thereby paving the way for the development of a *quantum internet* [150]. The scope of these networks extends beyond quantum computing technologies; the ability to control distributed quantum systems, as enabled by this architecture, to engineer complex quantum resources has applications in multi-partite secrete sharing [151], metrology [136], and probing fundamental physics [152].

8

Conclusion

*

This thesis presents experimental results from our mixed-species trapped-ion quantum network. The experiments, for the first time, combined three critical components developed in our group: high-rate, high-fidelity remote entanglement between Sr^+ network qubits in macroscopically separated modules [130]; the integration of long-lived Ca^+ circuit qubits into each module [123]; and mixed-species quantum logic operations between Sr^+ and Ca^+ qubits [15, 56]. In this chapter, we summarise the key results of the thesis and explore potential technical improvements to our apparatus that could enhance its performance. We also present an outlook on future experiments that could be conducted using our apparatus and briefly explore the broader implications for the future of quantum networking.

8.1 Summary

In Chapters 2, 3, 4, and 5, we outlined the theoretical and experimental tools that enabled the development of our quantum network. This network consists of two mixed-species trapped-ion modules, each containing one Sr^+ network qubit and one Ca^+ circuit qubit. This work lays the foundation for the key results of this thesis: the

integration of mixed-species techniques into quantum network modules (Chapter 6) and the first-ever demonstration of distributed quantum computing (Chapter 7).

Mixed-species remote entanglement: In Chapter 6, we first described the procedure for generating heralded remote entanglement between Sr^+ network qubits in the two modules. We then combined this "raw" remote entanglement with mixed-species quantum logic to create more complex remotely-entangled states, including the remote entanglement of an Sr^+ ion and a Ca^+ ion, the remote entanglement of two Ca^+ ions, a 3-qubit (Sr^+-Sr^+-Ca^+) GHZ state, and a 4-qubit (Sr^+-Sr^+-Ca^+-Ca^+) GHZ state. Finally, we demonstrated the storage capabilities of the Ca^+ circuit qubits by transferring the remote entanglement generated between the network qubits to the long-lived circuit qubits, demonstrating the storage of remote entanglement for up to 10 s.

At the time of writing, the results from this chapter are being prepared for publication.

Distributed quantum computing: In Chapter 7, we present the results from the first-ever demonstration of distributed quantum computing across an optical network link. In this demonstration, remote entanglement generated between network qubits is used to deterministically mediate CZ-gates between circuit qubits in macroscopically separated modules via quantum gate teleportation. This represents a key milestone in the construction of a distributed quantum computer. By utilising the robust memory of the long-lived circuit qubits, we deterministically execute quantum circuits compiled with up to three CZ-gates. In particular, we execute the iSWAP and SWAP gates, compiled with two and three teleported CZ-gates, respectively. This demonstrates our ability to teleport arbitrary two-qubit gates and, consequently, to perform universal distributed quantum computing. Finally, we execute the distributed two-qubit Grover's algorithm, which represents, to our knowledge, the first deterministic execution of any algorithm on a distributed quantum computer.

At the time of writing, the results from this chapter have been published on arXiv [137].

8.2 Outlook

To conclude this final chapter, we discuss possible next steps for our apparatus and provide an outlook on the future of quantum networking.

8.2.1 Potential improvements

Several potential improvements could be implemented to enhance the performance of our apparatus and enable the demonstration of more complex experiments. Here, we identify a number of these improvements that we did not have time to implement over the course of my DPhil.

Reducing motional heating and dephasing: Motional effects, such as heating rates and motional dephasing, are responsible for a significant portion of the infidelity in the mixed-species entangling gates within each module. Since these gates are a crucial component of the mixed-species modules, mitigating these effects would substantially improve the performance of the entire network. Furthermore, lowering the heating rates would reduce the cooling time required during remote entanglement generation, which constitutes a considerable portion of the overall entanglement generation success rate.

As discussed in Section 4.2.1, the heating rates observed with the Phoenix trap in Bob are significantly lower than those observed in Alice. Therefore, when building the next iteration of Alice, we will replace the older HOA2 trap with the Phoenix trap. This upgrade should improve the performance of the mixed-species entangling gates in this module.

Additionally, although we have not characterised the motional dephasing rate in Bob, the observed mixed-species gate fidelity suggests that the motional coherence time is comparable to that of Alice. The error budget simulations in Section 5.2.5

indicate that motional dephasing accounts for a significant portion of the gate infidelity. It would therefore be valuable to conduct a full characterisation of the motional dephasing in each trap. We believe this limitation is caused by technical noise coupled into the trap electrodes, likely through ground loops. Identifying and eliminating these noise sources would be a crucial step toward achieving local mixed-species entangling gates with fidelities exceeding 99 %.

Error analysis of the remote entanglement generation: Nadlinger [92] performed a comprehensive theoretical analysis of the possible error mechanisms in our remote entanglement scheme. However, we have not systematically compared this theoretical analysis with experiment. To better understand and mitigate these error mechanisms, it would be valuable to invest time in mapping out the error budget for the remote entanglement generation process. By comparing experimental results with theoretical predictions, we would gain a better understanding of the existing error mechanisms and could likely identify previously unconsidered sources of error.

Addition of 729 nm quadrupole laser for Ca^+: A 729 nm laser system has recently been set up, allowing us to address the $^{43}Ca^+$ quadrupole transitions between the ground $S_{1/2}$ manifold and the meta-stable $D_{5/2}$ manifold. Specifically, at 0.496 mT, there is an optical clock transition with no first-order sensitivity to magnetic field fluctuations. In contrast, the circuit qubit used in this thesis exhibited a small, but non-zero, first-order sensitivity. While the ~ 1 s natural lifetime of the meta-stable $D_{5/2}$ manifold would likely make this qubit unsuitable for storing quantum information during the entanglement generation process, we intend to use this clock transition for entanglement-enhanced frequency comparisons, as discussed in the next section.

The quadrupole laser also provides access to a wider range of entangling gate schemes. The gate mechanism used in this thesis did not couple directly to the circuit qubit, requiring temporary transfer to the auxiliary qubit. The addition of the quadrupole laser enables the exploration of gate schemes that could be applied

directly to the circuit qubit.

Increasing the number of ions per module: The results in this thesis were achieved with one Sr^+ ion and one Ca^+ ion in each trap. Increasing the number of ions per module would significantly enhance the quantum information processing capabilities of each module. Adding another Ca^+ ion per module would allow us to realise a distributed quantum computer capable of performing both local and non-local logical gate operations between circuit qubits. Moreover, with two Sr^+ ions and three Ca^+ ions, the architecture proposed by Nigmatullin *et al.* [141] could be realised.

Unfortunately, the lifetime of the Sr^+-Ca^+ crystal used in this thesis was frustratingly short, approximately 5-10 min. Despite extensive investigation, the primary cause of this issue remains unknown, although we suspect a combination of increased background pressure from the atomic ovens and sub-optimal trapping parameters. Further exploration of the Sr^+-Ca^+ crystal's stability and techniques to extend its lifetime would be a valuable investment of time[1].

8.2.2 Future experiments

Our apparatus combines a state-of-the-art Sr^+ quantum network with long-lived memory qubits provided by Ca^+ ions. The significant experimental complexity involved means that few experiments worldwide operate with these capabilities. Consequently, we are now able to explore a wide range of novel applications.

In this section, we will briefly outline several potential experiments that could be conducted using our apparatus.

Entanglement Distillation: Entanglement distillation is the process by which multiple noisy entangled states are transformed into a smaller number of higher-fidelity entangled states [141]. Such protocols will likely be a critical component

[1]For instance, splitting the Sr^+-Ca^+ crystal into separate wells while the experiment is idle has proven to be quite effective.

of mature quantum network technologies, enabling networks to share high-fidelity entangled resource states even when quantum channels interconnecting the modules are noisy. Due to its importance for quantum networks, demonstrating entanglement distillation has been one of the long-term goals of our experiment.

There have been several notable experimental demonstrations of entanglement distillation, implemented across various platforms, including photonics [153, 154], trapped ions [155], and diamond colour centres [156]. However, among these, only the experiment involving diamond colour centres was conducted across a quantum network [156]. While an excellent proof-of-concept, they only considered the distillation of a specific type of "raw" state and encountered issues with memory qubit dephasing and low remote entanglement fidelities. As a result, they were limited to a distilled fidelity of approximately 65 %.

Given the capabilities of our apparatus, we have an excellent opportunity to explore entanglement distillation schemes. With a single distillation round – requiring two rounds of entanglement generation, an LOCC – with perfect local operations we would, in theory, be able to distill our remote entanglement from $\sim 97\,\%$ to $> 99\,\%$.

Error detection in non-local logical qubits: Another interesting experiment, proposed by Prof. S. Benjamin, involves phase-flip error detection for a non-local logical qubit. In this scheme, a logical qubit is encoded in the subspace of two macroscopically separated circuit qubits, with $|0_{\mathrm{L}}\rangle = |{+}{+}_{\mathrm{C}}\rangle$ and $|1_{\mathrm{L}}\rangle = |{-}{-}_{\mathrm{C}}\rangle$. Phase-flip errors can then be detected using a $|\Phi^{+}\rangle$ resource state shared between two network qubits, enabling a non-local parity measurement of the circuit qubit. This would allow the detection of phase-flip errors in the encoded logical qubit.

As with the entanglement distillation experiment, this demonstration requires the same experimental resources as the DQC demonstration.

Entanglement-enhanced frequency comparisons of atomic clocks: Outside of quantum information processing applications, quantum networks of atomic clocks have been proposed as an important resource for metrology [136]. Using

our quantum networking apparatus, Nichol *et al.* [132] recently demonstrated that entangling two Sr^+ optical clocks enhances the precision of frequency comparisons between atomic clocks. However, the magnetic field sensitivity of the Sr^+ transition limited the duration for which we could perform the spectroscopic measurements.

Ongoing work led by Dr. R. Srinivas aims to use an optical clock transition in $^{43}Ca^+$, which is first-order insensitive to magnetic field fluctuations and could enable measurement durations approaching the 1 s natural lifetime of the $D_{5/2}$ manifold. As demonstrated by the results in this thesis, we are already able to generate remote entanglement between the Ca^+ ions. At the time of writing, the 729 nm laser system that will be used to probe the optical quadrupole transitions in Ca^+ is being set up.

Entanglement of remote harmonic oscillators: Another application of quantum networks is in quantum sensing. Quantum systems are generally highly sensitive to environmental fluctuations, making them excellent sensors when the exposure to these fluctuations can be controlled precisely. In particular, the harmonic motion of the ion is sensitive to electric fields and can be controlled with high precision. This suggests the potential for a network of harmonic oscillators for specific sensing applications. Just as with atomic clock comparisons, entanglement could provide a powerful resource for such a sensing network.

We are exploring the possibility of generating remotely-entangled motional states of two-ion crystals by first generating remote entanglement between the spin states of two Sr^+ ions, and then using a SDF to map this entanglement onto their motional states.

8.2.3 The next generation of trapped-ion quantum networks

The work presented in this thesis was conducted using a quantum networking apparatus that is at the state-of-the-art for these technologies. However, there is substantial room for improvement in many key metrics. For example, while the $\approx 97\,\%$ remote entanglement fidelity reported here is state-of-the-art, it is signifi-

cantly below the fidelities achieved for local entangling operations. Similarly, while an entanglement generation rate of $\sim 100\,\mathrm{s}^{-1}$ is close to the state-of-the-art, it could quickly become a bottleneck in a large-scale quantum network. Therefore, we conclude this thesis by discussing ongoing efforts to address these metrics and advance the field of quantum networking, with a focus on trapped-ion quantum networking.

Entanglement fidelity: The quantum channels interconnecting a quantum network will always, in general, be noisy – particularly when considering quantum networks on metropolitan or even international scales. Entanglement distillation will likely play a critical role in achieving the fidelities necessary for applications such as fault-tolerant distributed quantum computing. However, the distillation of high-fidelity states will require high-fidelity local operations. The integration of a quantum network interface into a quantum processing module capable of high-fidelity operations remains a significant technical challenge. Consequently, the distillation of entangled states with fidelities comparable to those of local operations has yet to be demontrated.

Entanglement rate: Our apparatus is capable of generating remote entanglement at a rate of approximately $100\,\mathrm{s}^{-1}$, which is primarily limited by the low efficiency of our free-space collection of single photons. To overcome this challenge, optical cavities are being pursued as a way to exceed the collection efficiencies achievable with free-space collection [157, 158, 159]. The presence of the cavity forces the ion to emit a photon into the cavity mode, enabling the photon to be extracted from the system with an efficiency approaching unity. In principle, this could enable ion-photon entanglement generation rates approaching the photon scattering limit, i.e., on the order of tens to hundreds of megahertz.

In the context of long-distance quantum networks, the entanglement attempt rate may be limited by the time-of-flight of the photons. As highlighted by You *et al.* [160], a photon travelling across an optical network link of 10 km would take approximately 100 µs, thereby limiting the entanglement attempt rate to 10 kHz—

two orders of magnitude lower than the attempt rate achieved in our apparatus. To explore methods of mitigating this limitation, temporal multiplexing of the photonic interface has recently been demonstrated [101, 160]. In these demonstrations, ions are sequentially transported into the ion-photon interface, before being transported away for storage until either the relevant heralding signal is measured or no herald is detected, allowing the ion to be reused. This technique offers a linear increase in the entanglement generation rate when the main limiting factor is the network link distance.

Finally, another approach to increase the entanglement generation rate is to increase the number of quantum interfaces per module, thereby increasing the surface area of the module's network interface. Realising such a device would require a significantly more compact design than the macroscopic lens system used in our apparatus. Ideally, these interfaces would be integrated into the quantum processing chip in a compact, scalable design. Focusing on trapped ions, there is ongoing work developing trap chips featuring integrated photonics, enabling on-chip free-space collection of photons [146]. This design could enable many interface zones per module, providing a scalable architecture for these trapped-ion modules.

Network link distance: As the distance of the optical network links increases, there will be an exponential reduction in the entanglement generation success probability due to photon loss. For example, at the 422 nm wavelength used in this work, optical fibres typically exhibit an attenuation of approximately $33\,\mathrm{dB\,km^{-1}}$ [161]. Even at metropolitan scales, this would result in a prohibitively low entanglement generation success probability.

However, by utilising down-conversion [42] or a different choice of optical transition to realise the interface, it would be possible to operate using infrared photons where fibre attenuation is significantly lower, thereby enabling the interconnection of quantum processors over metropolitan distances. For international distances, free-space transmission via satellite links [162] and quantum repeater technologies [32] could be employed to extend the reach of quantum networks.

A

Dephasing of a Two-Level Quantum System

*

In this brief appendix, we consider the effect of fluctuations in the qubit transition frequency[1] on the action of storing some quantum state, $\hat{\rho}$, between the times t_0 and $t_0 + \tau$. This mechanism is particularly important for our system, where the ions are exposed to magnetic field fluctuations and laser phase noise. Here, we derive the Kraus representation for the process of storing a quantum state while the system is exposed to fluctuations in the qubit transition frequency.

Suppose that the qubit frequency fluctuates in time as $\omega_0 - \beta(t)$, where ω_0 is the nominal transition frequency. The Hamiltonian describing evolution of the qubit in the frame rotating at the nominal frequency of the laser, ω_0, is

$$\hat{H} = \frac{\hbar}{2}\beta(t)\hat{\sigma}_z, \tag{A.1}$$

where $\beta(t)$ is the functional form of the frequency fluctuations [65]. The evolution

[1]Equivalently, one could consider fluctuations in the laser frequency, since what we actually care about are fluctuations relative to each other.

of the qubit is then described by the propagator

$$\hat{U}\left[\beta(t)\right] = e^{-i\Phi[\beta(t)]\hat{\sigma}_z/2} = \cos\left(\frac{\Phi\left[\beta(t)\right]}{2}\right)\hat{\mathbb{I}} - i\sin\left(\frac{\Phi\left[\beta(t)\right]}{2}\right)\hat{\sigma}_z, \tag{A.2}$$

where $\Phi\left[\beta(t)\right]$ is the functional

$$\Phi\left[\beta(t)\right] = \int_{-\infty}^{\infty} \mathrm{d}t\beta(t)\Pi(t; t_0, t_0+\tau), \tag{A.3}$$

and $\Pi(t; t_0, t_0+\tau)$ is the rectangle function which returns $\Pi(t; a, b) = 1$ if $t \in [a, b]$ and $\Pi(t; a, b) = 0$ otherwise.

Since the fluctuations are stochastic, the functional form of $\beta(t)$ is unknown – equivalently, the mechanism responsible for these fluctuations is part of the environment, which we cannot access. As such, we can write the process for storing the quantum state $\hat{\rho}$ by integrating over all possible realisations of $\beta(t)$

$$\mathcal{E}\left(\hat{\rho}\right) = \int \mathscr{D}\beta(t)\mathcal{P}\left[\beta(t)\right]\hat{U}\left[\beta(t)\right]\hat{\rho}\hat{U}^{\dagger}\left[\beta(t)\right], \tag{A.4}$$

where $\mathcal{P}\left[\beta(t)\right]$ is the probability density that the function $\beta(t)$ is realised. Making use of Eq. (A.2), we can write Eq. (A.4) as

$$\begin{aligned}\mathcal{E}\left(\hat{\rho}\right) =& \frac{1+\left\langle\cos\left(\Phi\left[\beta(t)\right]\right)\right\rangle}{2}\hat{\rho} \\ &+ \frac{1-\left\langle\cos\left(\Phi\left[\beta(t)\right]\right)\right\rangle}{2}\hat{\sigma}_z\hat{\rho}\hat{\sigma}_z \\ &- \frac{i}{2}\left\langle\sin\left(\Phi\left[\beta(t)\right]\right)\right\rangle\left[\hat{\sigma}_z, \hat{\rho}\right],\end{aligned} \tag{A.5}$$

where

$$\left\langle f\left(\Phi\left[\beta(t)\right]\right)\right\rangle = \int \mathscr{D}\beta(t)\mathcal{P}\left[\beta(t)\right] f\left(\Phi\left[\beta(t)\right]\right). \tag{A.6}$$

Let us write the Taylor expansion of f as $f(x) = \sum_{n=0}^{\infty} a_n x^n$. Then we find

$$\langle f\left(\Phi\left[\beta(t)\right]\right)\rangle = \sum_{n=0}^{\infty} a_n \int \mathscr{D}\beta(t)\mathcal{P}\left[\beta(t)\right]\Phi\left[\beta(t)\right] \tag{A.7}$$

$$= \sum_{n=0}^{\infty} a_n \int_{-\infty}^{\infty} \mathrm{d}t_1 \cdots \int_{-\infty}^{\infty} \mathrm{d}t_n \left\langle \beta(t_1)\ldots\beta(t_n)\right\rangle \Pi(t_1)\ldots\Pi(t_n). \tag{A.8}$$

Therefore, we have reduced the problem to one of calculating n-point correlation functions for the noise, $\beta(t)$.

It is useful at this point to define the characteristic functional, $\varphi\left[k(t)\right]$, as

$$\varphi\left[k(t)\right] = \frac{\int \mathscr{D}\beta(t)\mathcal{P}\left[\beta(t)\right] e^{i\int \mathrm{d}t k(t)\beta(t)}}{\int\int \mathscr{D}\beta(t)\mathcal{P}\left[\beta(t)\right]}, \tag{A.9}$$

which has the properties $\varphi[0] = 1$ and is related to the n-point correlation function via the functional derivatives

$$\left\langle \beta(t_1)\ldots\beta(t_n)\right\rangle = (-i)^n \left.\frac{\delta\varphi\left[k(t)\right]}{\delta k(t_1)\ldots\delta k(t_n)}\right|_{k(t)=0}. \tag{A.10}$$

We now make the assumption that the statistics of the fluctuations are Gaussian, and thus is completely defined by the two-point correlation function $S(t_1 - t_2) = \langle\beta(t_1)\beta(t_2)\rangle$ [65]. The general form of a characteristic functional for Gaussian noise that is time-independent can be written as [163]

$$\varphi\left[k(t)\right] = \exp\left(-\frac{1}{2}\int \mathrm{d}t \int \mathrm{d}t' k(t)k(t')S(t-t')\right). \tag{A.11}$$

From this, we find that for even n, the n-point correlators can be written as

$$\left\langle \beta(t_1)\ldots\beta(t_n)\right\rangle = \sum_{P(1,\ldots,n)} S(t_{P_1} - t_{P_2})\ldots S(t_{P_{n-1}} - t_{P_n}), \tag{A.12}$$

where the sum is over all possible pairings of $\{1,\ldots,n\}$ – this is known as *Wick's theorem*. For odd n, the correlators vanish.

Returning to Eq. (A.5), we see that since the n-point correlation functions for

odd n vanish, then the $\langle \sin(\Phi[\beta(t)]) \rangle$ term will also vanish. Eq. (A.5) now takes the appearance of a dephasing quantum channel,

$$\mathcal{E}(\hat{\rho}) = (1-p)\hat{\rho} + p\hat{\sigma}_z\hat{\rho}\hat{\sigma}_z, \tag{A.13}$$

where

$$p = \frac{1 - \langle \cos(\Phi[\beta(t)]) \rangle}{2}. \tag{A.14}$$

Using Eq. (A.8), we find that

$$\langle \cos(\Phi[\beta(t)]) \rangle = \sum_{n=0}^{\infty} \frac{(-1)^n}{(2n)!} \int_{-\infty}^{\infty} \mathrm{d}t_1 \cdots \int_{-\infty}^{\infty} \mathrm{d}t_n \, \langle \beta(t_1) \dots \beta(t_n) \rangle \, \Pi(t_1) \dots \Pi(t_n). \tag{A.15}$$

Using Wick's theorem, Eq. (A.12), we can write this in terms of the two-point correlation functions

$$\begin{aligned} \langle \cos(\Phi[\beta(t)]) \rangle = & \sum_{n=0}^{\infty} \frac{(-1)^n}{(2n)!} \sum_{P(1,\dots,2n)} \int_{-\infty}^{\infty} \mathrm{d}t_1 \cdots \int_{-\infty}^{\infty} \mathrm{d}t_{2n} \\ & \times S(t_{P_1} - t_{P_2}) \cdots S(t_{P_{2n-1}} - t_{P_{2n}}) \, \Pi(t_{P_1}) \cdots \Pi(t_{P_{2n}}). \end{aligned} \tag{A.16}$$

We can simplify this expression further by introducing the spectral noise density, $\tilde{S}(\omega)$, defined by the Fourier transform

$$S(t) = \int_{-\infty}^{\infty} \frac{\mathrm{d}\omega}{2\pi} \tilde{S}(\omega) e^{-i\omega t}. \tag{A.17}$$

In terms of the spectral noise density, Eq. (A.16) becomes

$$\langle \cos(\Phi[\beta(t)]) \rangle = \sum_{n=0}^{\infty} \frac{(-1)^n}{(2n)!} \mathcal{N}_{2n} \left[2 \int_0^{\infty} \frac{\mathrm{d}\omega}{2\pi} \tilde{S}(\omega) \left| \tilde{\Pi}(\omega, \tau) \right|^2 \right]^n, \tag{A.18}$$

where $\left| \tilde{\Pi}(\omega, \tau) \right|^2$ is the mod-squared value of the Fourier transform of the rectangle function,

$$\left| \tilde{\Pi}(\omega, \tau) \right|^2 = \frac{4}{\omega^2} \sin^2\left(\frac{\omega\tau}{2} \right), \tag{A.19}$$

and $\mathcal{N}_{2n}$ is the number of terms in the sum over the possible pairings of $\{1, \ldots, 2n\}$,

$$\mathcal{N}_{2n} = \frac{(2n)!}{n!2^n}. \tag{A.20}$$

We therefore find that

$$\langle \cos\left(\Phi\left[\beta(t)\right]\right)\rangle = \sum_{n=0}^{\infty} \frac{1}{n!}\left[-\int_0^{\infty} \frac{\mathrm{d}\omega}{2\pi}\tilde{S}(\omega)\left|\tilde{\Pi}(\omega,\tau)\right|^2\right]. \tag{A.21}$$

In terms of the quantum channel Eq. (A.13), we see that the phase-flip probability, p, can be written as

$$p = \frac{1}{2}\left[1 - \exp\left(-\int_0^{\infty} \frac{\mathrm{d}\omega}{2\pi}\tilde{S}(\omega)\left|\tilde{\Pi}(\omega,\tau)\right|^2\right)\right], \tag{A.22}$$

This result means that given some spectral noise density of the frequency fluctuations, $\tilde{S}(\omega)$, we can calculate the quantum channel for storage of a quantum state.

Acronyms

CNOT controlled-NOT

CZ controlled-Z

ABPE adaptive Bayesian phase estimation

ADC analogue-to-digital conversion

AOM acousto-optical modulator

APD avalanche photodiode

CPMG Carr-Purcell-Meiboom-Gill

CPTP completely positive trace-preserving

CPU central processing unit

CVQC continuous variable quantum computing

DAC digital-to-analogue conversion

DDS direct digital synthesis

DQC distributed quantum computing

EIT electromagnetically induced transparency

EMCCD electron multiplying charge-coupled device

EOM electro-optic modulator

FPGA field programmable gate array

GHZ Greenberger-Horne-Zeilinger

GPIO general-purpose input/output

HOM Hong-Ou-Mandel

HWP half waveplate

IP in-phase

IPE ion-photon entanglement

LOCC local operations and classical communication

NPBS non-polarising beamsplitter

OOP out-of-phase

PBS polarising beamsplitter

PCF photonic crystal fibre

PMT photomultiplier tube

POVM positive operator-valued measure

PSU power supply unit

QCCD quantum charge-coupled device

QGT quantum gate teleportation

QIP quantum information processing

QPT quantum process tomography

QWP quarter waveplate

RBM randomised benchmarking

RF radio-frequency

RTIO real-time input/output

RWA rotating wave approximation

SDF spin-dependent force

SHG second harmonic generation

SPAM state preparation and measurement

SPI serial peripheral interface

TA tapered amplifier

TTL transistor-transistor logic

UR universally robust

WZZ Walsh-modulated ZZ

Bibliography

[1] A. G. J. MacFarlane, J. P. Dowling, and G. J. Milburn. "Quantum technology: the second quantum revolution". Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences. **361**, pp. 1655–1674. 2003.

[2] R. P. Feynman. "Simulating physics with computers". International Journal of Theoretical Physics. **21**, pp. 467–488. 1982.

[3] D. Deutsch. "Quantum theory, the Church-Turing principle and the universal quantum computer". Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences. **400**, pp. 97–117. 1985.

[4] D. Deutsch and R. Jozsa. "Rapid solution of problems by quantum computation". Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences. **439**, pp. 553–558. 1992.

[5] P. W. Shor. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer". SIAM J. Comput. **26**, pp. 1484–1509. 1997.

[6] L. K. Grover. "A fast quantum mechanical algorithm for database search". In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. STOC '96. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, pp. 212–219.

[7] P. W. Shor. "Scheme for reducing decoherence in quantum computer memory". Phys. Rev. A. **52**, R2493–R2496. 1995.

[8] A. M. Steane. "Error Correcting Codes in Quantum Theory". Phys. Rev. Lett. **77**, pp. 793–797. 1996.

[9] D. P. DiVincenzo. "The Physical Implementation of Quantum Computation". Fortschritte der Physik. **48**, pp. 771–783. 2000.

[10] T. P. Harty *et al.* "High-Fidelity Preparation, Gates, Memory, and Readout of a Trapped-Ion Quantum Bit". Phys. Rev. Lett. **113**, p. 220501. 2014.

[11] F. A. An *et al.* "High Fidelity State Preparation and Measurement of Ion Hyperfine Qubits with $I > \frac{1}{2}$". Phys. Rev. Lett. **129**, p. 130501. 2022.

[12] A. S. Sotirova *et al.* "High-fidelity heralded quantum state preparation and measurement". Preprint at `https://arxiv.org/abs/2409.05805`. 2024.

[13] C. J. Ballance *et al.* "High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits". Phys. Rev. Lett. **117**, p. 060504. 2016.

[14] J. P. Gaebler *et al.* "High-Fidelity Universal Gate Set for $^{9}Be^{+}$ Ion Qubits". Phys. Rev. Lett. **117**, p. 060505. 2016.

[15] A. C. Hughes *et al.* "Benchmarking a High-Fidelity Mixed-Species Entangling Gate". Phys. Rev. Lett. **125**, p. 080504. 2020.

[16] R. Srinivas *et al.* "High-fidelity laser-free universal control of trapped ion qubits". Nature. **597**, pp. 209–213. 2021.

[17] C. R. Clark *et al.* "High-Fidelity Bell-State Preparation with $^{40}Ca^{+}$ Optical Qubits". Phys. Rev. Lett. **127**, p. 130505. 2021.

[18] M. A. Weber *et al.* "Robust and fast microwave-driven quantum logic for trapped-ion qubits". Preprint at `https://arxiv.org/abs/2402.12955`. 2024.

[19] P. Wang *et al.* "Single ion qubit with estimated coherence time exceeding one hour". Nat. Commun. **12**, p. 233. 2021.

[20] C. D. Bruzewicz *et al.* "Trapped-ion quantum computing: Progress and challenges". Appl. Phys. Rev. **6**, p. 021314. 2019.

[21] S. Bravyi *et al.* "The future of quantum computing with superconducting qubits". J. Appl. Phys. **132**, p. 160902. 2022.

[22] S. S. Gill *et al.* "Quantum Computing: Vision and Challenges". Preprint at `https://arxiv.org/abs/2403.02240`. 2024.

[23] M. Malinowski, D. Allcock, and C. Ballance. "How to Wire a 1000-Qubit Trapped-Ion Quantum Computer". PRX Quantum. **4**, p. 040313. 2023.

[24] L. K. Grover. "Quantum Telecomputation". Preprint at `https://arxiv.org/abs/quant-ph/9704012`. 1997.

[25] J. I. Cirac *et al.* "Distributed quantum computation over noisy channels". Phys. Rev. A. **59**, pp. 4249–4254. 1999.

[26] L. Jiang *et al.* "Distributed quantum computation based on small quantum registers". Phys. Rev. A. **76**, p. 062323. 2007.

[27] C. Monroe *et al.* "Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects". Phys. Rev. A. **89**, p. 022317. 2014.

[28] C. H. Bennett *et al.* "Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels". Phys. Rev. Lett. **70**, pp. 1895–1899. 1993.

[29] D. Gottesman and I. L. Chuang. "Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations". Nature. **402**, pp. 390–393. 1999.

[30] J. Eisert *et al.* "Optimal local implementation of nonlocal quantum gates". Phys. Rev. A. **62**, p. 052317. 2000.

[31] D. Collins, N. Linden, and S. Popescu. "Nonlocal content of quantum operations". Phys. Rev. A. **64**, p. 032302. 2001.

[32] H.-J. Briegel *et al.* "Quantum Repeaters: The Role of Imperfect Local Operations in Quantum Communication". Phys. Rev. Lett. **81**, pp. 5932–5935. 1998.

[33] W. Dür and H.-J. Briegel. "Entanglement Purification for Quantum Computation". Phys. Rev. Lett. **90**, p. 067901. 2003.

[34] D. Kielpinski, C. Monroe, and D. J. Wineland. "Architecture for a large-scale ion-trap quantum computer". Nature. **417**, pp. 709–711. 2002.

[35] M. Akhtar *et al.* "A high-fidelity quantum matter-link between ion-trap microchip modules". Nat. Commun. **14**, p. 531. 2023.

[36] Y. Wan *et al.* "Quantum gate teleportation between separated qubits in a trapped-ion processor". Science. **364**, pp. 875–878. 2019.

[37] J. M. Pino *et al.* "Demonstration of the trapped-ion quantum CCD computer architecture". Nature. **592**, pp. 209–213. 2021.

[38] P. C. Humphreys *et al.* "Deterministic delivery of remote entanglement on a quantum network". Nature. **558**, pp. 268–273. 2018.

[39] C. M. Knaut *et al.* "Entanglement of nanophotonic quantum memory nodes in a telecom network". Nature. **629**, pp. 573–578. 2024.

[40] S. Storz *et al.* "Loophole-free Bell inequality violation with superconducting circuits". Nature. **617**, pp. 265–270. 2023.

[41] S. Ritter *et al.* "An elementary quantum network of single atoms in optical cavities". Nature. **484**, pp. 195–200. 2012.

[42] T. van Leent *et al.* "Entangling single atoms over 33 km telecom fibre". Nature. **607**, pp. 69–73. 2022.

[43] D. L. Moehring *et al.* "Entanglement of single-atom quantum bits at a distance". Nature. **449**, pp. 68–71. 2007.

[44] D. Hucul *et al.* "Modular entanglement of atomic qubits using photons and phonons". Nature Physics. **11**, pp. 37–42. 2015.

[45] S. Saha *et al.* "High-fidelity remote entanglement of trapped atoms mediated by time-bin photons". Preprint at `https://arxiv.org/abs/2406.01761`. 2024.

[46] J. O'Reilly *et al.* "Fast photon-mediated entanglement of continuously-cooled trapped ions for quantum networking". Preprint at `https://arxiv.org/abs/2404.16167`. 2024.

[47] Y.-F. Huang *et al.* "Experimental Teleportation of a Quantum Controlled-NOT Gate". Phys. Rev. Lett. **93**, p. 240501. 2004.

[48] W.-B. Gao *et al.* "Teleportation-based realization of an optical quantum two-qubit entangling gate". Proc. Natl. Acad. Sci. U.S.A. **107**, pp. 20869–20874. 2010.

[49] K. S. Chou *et al.* "Deterministic teleportation of a quantum gate between two logical qubits". Nature. **561**, pp. 368–373. 2018.

[50] E. Bäumer *et al.* "Efficient Long-Range Entanglement using Dynamic Circuits". Preprint at `https://arxiv.org/abs/2308.13065`. 2023.

[51] A. Hashim *et al.* "Efficient Generation of Multi-partite Entanglement between Non-local Superconducting Qubits using Classical Feedback". Preprint at `https://arxiv.org/abs/2403.18768`. 2024.

[52] P. Maunz *et al.* "Heralded Quantum Gate between Remote Quantum Memories". Phys. Rev. Lett. **102**, p. 250502. 2009.

[53] S. Daiss *et al.* "A quantum-logic gate between distant quantum-network modules". Science. **371**, pp. 614–617. 2021.

[54] C. J. Wood, J. D. Biamonte, and D. G. Cory. "Tensor networks and graphical calculus for open quantum systems". Preprint at `https://arxiv.org/abs/1111.6950`. 2015.

[55] J. P. Home. "Quantum science and metrology with mixed-species ion chains". Preprint at `https://arxiv.org/abs/1306.5950`. 2013.

[56] P. Drmota. "Blind quantum computing with trapped-ions and single-photons". DPhil thesis. University of Oxford, 2023.

[57] C. J. Ballance. "High-fidelity quantum logic in Ca^+". DPhil thesis. University of Oxford, 2014.

[58] A. C. Hughes. "Benchmarking memory and logic gates for trapped-ion quantum computing". DPhil thesis. University of Oxford, 2021.

[59] H.-P. Breuer and F. Petruccione. "The Theory of Open Quantum Systems". Oxford University Press, 2007.

[60] A. Gilchrist, D. R. Terno, and C. J. Wood. "Vectorization of quantum operations and its use". Preprint at `https://arxiv.org/abs/0911.2539`. 2011.

[61] C. R. Johnson and R. A. Horn. "Matrix analysis". Cambridge university press Cambridge, 1985.

[62] D. Leibfried *et al.* "Quantum dynamics of single trapped ions". Rev. Mod. Phys. **75**, pp. 281–324. 2003.

[63] R. Loudon. "The quantum theory of light". eng. 3rd ed. Oxford science publications. Oxford: Clarendon Press, 2000.

[64] D. F. James and J. Jerke. "Effective Hamiltonian theory and its applications in quantum information". Canadian Journal of Physics. **85**, pp. 625–632. 2007.

[65] Ł. Cywiński *et al.* "How to enhance dephasing time in superconducting qubits". Phys. Rev. B. **77**, p. 174509. 2008.

[66] S. Meiboom and D. Gill. "Modified spin-echo method for measuring nuclear relaxation times". Review of scientific instruments. **29**, pp. 688–691. 1958.

[67] A. M. Souza, G. A. Álvarez, and D. Suter. "Robust Dynamical Decoupling for Quantum Computing and Quantum Memory". Phys. Rev. Lett. **106**, p. 240501. 2011.

[68] G. T. Genov *et al.* "Arbitrarily Accurate Pulse Sequences for Robust Dynamical Decoupling". Phys. Rev. Lett. **118**, p. 133202. 2017.

[69] J. I. Cirac and P. Zoller. "Quantum Computations with Cold Trapped Ions". Physical Review Letters. **74**, pp. 4091–4094. 1995.

[70] D. Leibfried *et al.* "Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate". Nature. **422**, pp. 412–415. 2003.

[71] C. H. Baldwin *et al.* "Subspace benchmarking high-fidelity entangling operations with trapped ions". Physical Review Research. **2**, p. 013317. 2020.

[72] T. P. Harty *et al.* "High-Fidelity Trapped-Ion Quantum Logic Using Near-Field Microwaves". Phys. Rev. Lett. **117**, p. 140501. 2016.

[73] C. M. Löschnauer *et al.* "Scalable, high-fidelity all-electronic control of trapped-ion qubits". Preprint at `https://arxiv.org/abs/2407.07694`. 2024.

[74] C. H. Baldwin *et al.* "High-fidelity light-shift gate for clock-state qubits". Phys. Rev. A. **103**, p. 012603. 2021.

[75] V. Schäfer. "Fast gates and mixed-species entanglement with trapped ions". DPhil thesis. University of Oxford, 2018.

[76] W. Magnus. "On the exponential solution of differential equations for a linear operator". Communications on Pure and Applied Mathematics. **7**, pp. 649–673. 1954.

[77] D. Main. "Magnetic Field Stabilisation in Ion Traps". MPhys thesis. University of Oxford, 2020.

[78] D. W. Berry *et al.* "How to perform the most accurate possible phase measurements". Phys. Rev. A. **80**, p. 052114. 2009.

[79] N. Wiebe and C. Granade. "Efficient Bayesian Phase Estimation". Phys. Rev. Lett. **117**, p. 010503. 2016.

[80] C. Bonato and D. W. Berry. "Adaptive tracking of a time-varying field with a quantum sensor". Phys. Rev. A. **95**, p. 052348. 2017.

[81] B. de Neeve. "Calibrating an Ion-Trap Quantum Computer". MSc thesis. ETH Zürich, 2017.

[82] J. Fiurášek and Z. Hradil. "Maximum-likelihood estimation of quantum processes". Phys. Rev. A. **63**, p. 020101. 2001.

[83] J. Řeháček *et al.* "Diluted maximum-likelihood algorithm for quantum tomography". Phys. Rev. A. **75**, p. 042108. 2007.

[84] G. C. Knee *et al.* "Quantum process tomography via completely positive and trace-preserving projection". Phys. Rev. A. **98**, p. 062336. 2018.

[85] E. Magesan, J. M. Gambetta, and J. Emerson. "Scalable and Robust Randomized Benchmarking of Quantum Processes". Phys. Rev. Lett. **106**, p. 180504. 2011.

[86] K. A. Gilmore *et al.* "Quantum-enhanced sensing of displacements and electric fields with two-dimensional trapped-ion crystals". Science. **373**, pp. 673–678. 2021.

[87] S. Kolkowitz *et al.* "Probing Johnson noise and ballistic transport in normal metals with a single-spin qubit". Science. **347**, pp. 1129–1132. 2015.

[88] A. S. Holevo. "Covariant measurements and imprimitivity systems". In: *Lecture Notes in Mathematics, Berlin Springer Verlag*. **1055**. 1984, p. 153.

[89] J. Kiefer. "Sequential Minimax Search for a Maximum". Proceedings of the American Mathematical Society. **4**, pp. 502–506. 1953.

[90] M. A. Nielsen. "A simple formula for the average gate fidelity of a quantum dynamical operation". Physics Letters A. **303**, pp. 249–252. 2002.

[91] L. Stephenson. "Entanglement between nodes of a quantum network". DPhil thesis. University of Oxford, 2019.

[92] D. P. Nadlinger. "Device-independent key distribution between trapped-ion quantum network nodes". DPhil thesis. University of Oxford, 2022.

[93] B. Nichol. "An elementary quantum network of entangled optical atomic clocks". DPhil thesis. University of Oxford, 2022.

[94] S. M. Brewer *et al.* "$^{27}Al^+$ Quantum-Logic Clock with a Systematic Uncertainty below 10^{-18}". Phys. Rev. Lett. **123**, p. 033201. 2019.

[95] N. Huntemann *et al.* "Single-Ion Atomic Clock with 3×10^{-18} Systematic Uncertainty". Phys. Rev. Lett. **116**, p. 063001. 2016.

[96] A. Tofful *et al.* "$^{171}Yb^+$ optical clock with 2.2×10^{-18} systematic uncertainty and absolute frequency measurements". Metrologia. **61**, p. 045001. 2024.

[97] H.-X. Yang *et al.* "Realizing coherently convertible dual-type qubits with the same ion species". Nature Physics. **18**, pp. 1058–1061. 2022.

[98] D. Hucul *et al.* "Spectroscopy of a Synthetic Trapped Ion Qubit". Phys. Rev. Lett. **119**, p. 100501. 2017.

[99] J. E. Christensen *et al.* "High-fidelity manipulation of a qubit enabled by a manufactured nucleus". npj Quantum Information. **6**, p. 35. 2020.

[100] M. Smith *et al.* "Single-qubit gates with errors at the 10^{-7} level". Preprint at `https://arxiv.org/abs/2412.04421`. 2024.

[101] M. Canteri *et al.* "A photon-interfaced ten qubit quantum network node". Preprint at `https://arxiv.org/abs/2406.09480`. 2024.

[102] B. B. Blinov *et al.* "Observation of entanglement between a single trapped atom and a single photon". Nature. **428**, pp. 153–157. 2004.

[103] M. D. Barrett *et al.* "Sympathetic cooling of $^9Be^+$ and $^{24}Mg^+$ for quantum logic". Phys. Rev. A. **68**, p. 042302. 2003.

[104] P. L. W. Maunz. "High Optical Access Trap 2.0." 2016.

[105] M. C. Revelle. "Phoenix and Peregrine Ion Traps". Preprint at `https://arxiv.org/abs/2009.02398`. 2020.

[106] M. Weber. "High-Fidelity, Near-Field Microwave Gates in a Cryogenic Surface Trap". DPhil thesis. University of Oxford, 2022.

[107] D. P. Nadlinger *et al.* "Micromotion minimisation by synchronous detection of parametrically excited motion". Preprint at `https://arxiv.org/abs/2107.00056`. 2021.

[108] F. Diedrich *et al.* "Laser Cooling to the Zero-Point Energy of Motion". Phys. Rev. Lett. **62**, pp. 403–406. 1989.

[109] M. Brownnutt *et al.* "Ion-trap measurements of electric-field noise near surfaces". Rev. Mod. Phys. **87**, pp. 1419–1482. 2015.

[110] D. Nadlinger. "Entanglement between Trapped Strontium Ions and Photons". MSc thesis. ETH Zürich, 2016.

[111] S. Bourdeauducq *et al.* "m-labs/artiq: 6.0 (Version 6.0)". 2021.

[112] T. G. Ballance *et al.* "A short response time atomic source for trapped ion experiments". Review of Scientific Instruments. **89**, p. 053102. 2018.

[113] D. M. Lucas *et al.* "Isotope-selective photoionization for calcium ion trapping". Phys. Rev. A. **69**, p. 012711. 2004.

[114] D. J. Wineland and W. M. Itano. "Laser cooling of atoms". Phys. Rev. A. **20**, pp. 1521–1540. 1979.

[115] J. Eschner *et al.* "Laser cooling of trapped ions". J. Opt. Soc. Am. B. **20**, pp. 1003–1015. 2003.

[116] G. Morigi, J. Eschner, and C. H. Keitel. "Ground State Laser Cooling Using Electromagnetically Induced Transparency". Phys. Rev. Lett. **85**, pp. 4458–4461. 2000.

[117] C. F. Roos *et al.* "Experimental Demonstration of Ground State Laser Cooling with Electromagnetically Induced Transparency". Phys. Rev. Lett. **85**, pp. 5547–5550. 2000.

[118] D. Szwer. "High fidelity readout and protection of a 43Ca+ trapped ion qubit". DPhil thesis. Oxford University, 2009.

[119] A. H. Myerson *et al.* "High-Fidelity Readout of Trapped-Ion Qubits". Phys. Rev. Lett. **100**, p. 200502. 2008.

[120] B. Merkel *et al.* "Magnetic field stabilization system for atomic physics experiments". Review of Scientific Instruments. **90**, p. 044702. 2019.

[121] Quartiq. "Stabilizer". `https://github.com/quartiq/stabilizer`.

[122] O. I. T. Group. "Analogue front-end mezzanine board for Stabilizer". `https://github.com/OxfordIonTrapGroup/stabilizer_current_sense`.

[123] P. Drmota *et al.* "Robust Quantum Memory in a Trapped-Ion Quantum Network Node". Phys. Rev. Lett. **130**, p. 090803. 2023.

[124] D. Hayes *et al.* "Reducing sequencing complexity in dynamical quantum error suppression by Walsh modulation". Phys. Rev. A. **84**, p. 062323. 2011.

[125] D. Hayes *et al.* "Coherent Error Suppression in Multiqubit Entangling Gates". Phys. Rev. Lett. **109**, p. 020503. 2012.

[126] R. T. Sutherland *et al.* "One- and two-qubit gate infidelities due to motional errors in trapped ions and electrons". Phys. Rev. A. **105**, p. 022437. 2022.

[127] S. Krämer *et al.* "QuantumOptics. jl: A Julia framework for simulating open quantum systems". Computer Physics Communications. **227**, pp. 109–116. 2018.

[128] P. Drmota *et al.* "Verifiable Blind Quantum Computing with Trapped Ions and Single Photons". Phys. Rev. Lett. **132**, p. 150604. 2024.

[129] D. T. C. Allcock *et al.* "omg blueprint for trapped ion quantum computing with metastable states". Applied Physics Letters. **119**, p. 214002. 2021.

[130] L. J. Stephenson *et al.* "High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network". Phys. Rev. Lett. **124**, p. 110501. 2020.

[131] D. P. Nadlinger *et al.* "Experimental quantum key distribution certified by Bell's theorem". Nature. **607**, pp. 682–686. 2022.

[132] B. C. Nichol *et al.* "An elementary quantum network of entangled optical atomic clocks". Nature. **609**, pp. 689–694. 2022.

[133] D. Main *et al.* "Multi-Partite Mixed-Species Entanglement over a Quantum Network". In preparation.

[134] J.-W. Pan *et al.* "Experimental Entanglement Swapping: Entangling Photons That Never Interacted". Phys. Rev. Lett. **80**, pp. 3891–3894. 1998.

[135] C. K. Hong, Z. Y. Ou, and L. Mandel. "Measurement of subpicosecond time intervals between two photons by interference". Phys. Rev. Lett. **59**, pp. 2044–2046. 1987.

[136] P. Kómár *et al.* "A quantum network of clocks". Nat. Phys. **10**, pp. 582–587. 2014.

[137] D. Main *et al.* "Distributed Quantum Computing across an Optical Network Link". Preprint at `https://arxiv.org/abs/2407.00835`. 2024.

[138] G. Vidal and C. M. Dawson. "Universal quantum circuit for two-qubit transformations with three controlled-NOT gates". Phys. Rev. A. **69**, p. 010301. 2004.

[139] V. Letchumanan *et al.* "Lifetime measurement of the metastable $4d\,^2\mathrm{D}_{5/2}$ state in $^{88}\mathrm{Sr}^+$ using a single trapped ion". Phys. Rev. A. **72**, p. 012509. 2005.

[140] S. A. Moses *et al.* "A Race-Track Trapped-Ion Quantum Processor". Phys. Rev. X. **13**, p. 041052. 2023.

[141] R. Nigmatullin *et al.* "Minimally complex ion traps as modules for quantum communication and computing". New J. Phys. **18**, p. 103028. 2016.

[142] D. P. DiVincenzo. "Two-bit gates are universal for quantum computation". Phys. Rev. A. **51**, pp. 1015–1022. 1995.

[143] A. Barenco *et al.* "Elementary gates for quantum computation". Phys. Rev. A. **52**, pp. 3457–3467. 1995.

[144] A. Montanaro. "Quantum algorithms: an overview". npj Quantum Information. **2**, p. 15023. 2016.

[145] J. Hilder *et al.* "Fault-Tolerant Parity Readout on a Shuttling-Based Trapped-Ion Quantum Computer". Phys. Rev. X. **12**, p. 011032. 2022.

[146] F. W. Knollmann *et al.* "Integrated photonic structures for photon-mediated entanglement of trapped ions". Preprint at `https://arxiv.org/abs/2401.06850`. 2024.

[147] Y.-H. Luo *et al.* "Quantum Teleportation in High Dimensions". Phys. Rev. Lett. **123**, p. 070505. 2019.

[148] S. Lloyd and S. L. Braunstein. "Quantum Computation over Continuous Variables". Phys. Rev. Lett. **82**, pp. 1784–1787. 1999.

[149] B. W. Walshe *et al.* "Continuous-variable gate teleportation and bosonic-code error correction". Phys. Rev. A. **102**, p. 062411. 2020.

[150] S. Wehner, D. Elkouss, and R. Hanson. "Quantum internet: A vision for the road ahead". Science. **362**, eaam9288. 2018.

[151] M. Hillery, V. Bu žek, and A. Berthiaume. "Quantum secret sharing". Phys. Rev. A. **59**, pp. 1829–1834. 1999.

[152] D. M. Greenberger *et al.* "Bell's theorem without inequalities". Am. J. Phys. **58**, pp. 1131–1143. 1990.

[153] P. G. Kwiat *et al.* "Experimental entanglement distillation and hidden non-locality". Nature. **409**, pp. 1014–1017. 2001.

[154] J.-W. Pan *et al.* "Experimental entanglement purification of arbitrary unknown states". Nature. **423**, pp. 417–422. 2003.

[155] R. Reichle *et al.* "Experimental purification of two-atom entanglement". Nature. **443**, pp. 838–841. 2006.

[156] N. Kalb *et al.* "Entanglement distillation between solid-state quantum network nodes". Science. **356**, pp. 928–932. 2017.

[157] V. Krutyanskiy *et al.* "Entanglement of Trapped-Ion Qubits Separated by 230 Meters". Phys. Rev. Lett. **130**, p. 050803. 2023.

[158] S. Gao *et al.* "Optimization of Scalable Ion-Cavity Interfaces for Quantum Photonic Networks". Phys. Rev. Appl. **19**, p. 014033. 2023.

[159] E. Kassa *et al.* "How to integrate a miniature optical cavity in a linear ion trap: shielding dielectrics and trap symmetry". Preprint at `https://arxiv.org/abs/2409.05123`. 2024.

[160] B. You *et al.* "Temporally multiplexed ion-photon quantum interface via fast ion-chain transport". Preprint at `https://arxiv.org/abs/2405.10501`. 2024.

[161] E. M. Ainley. "Characterising an Optical Fibre Transmitting Visible Light for Applications in Long Quantum Networks". MPhys thesis. University of Oxford, 2022.

[162] Y.-A. Chen *et al.* "An integrated space-to-ground quantum communication network over 4,600 kilometres". Nature. **589**, pp. 214–219. 2021.

[163] R. P. Feynman, D. F. Styer, and A. R. Hibbs. "Quantum mechanics and path integrals". Emended ed. Dover books on physics. Mineola, N.Y: Dover, 2010.