DISS. ETH NO. 25322

Feedback-stabilised quantum states in a mixed-species ion system

A thesis submitted to attain the degree of
DOCTOR OF SCIENCE of ETH ZÜRICH
(Dr. sc. ETH ZÜRICH)

presented by

VLAD NEGNEVITSKY
BSc (Hons.), BEng. (Hons.), Monash University, 2010

born on 01.10.1987
citizen of Australia

27.09.2018

Accepted on the recommendation of
Prof. Dr. J. P. Home
Prof. Dr. R. Ozeri
Prof. Dr. A. Wallraff
2018

Abstract

Trapped ions are among the leading platforms for realising quantum information processing (QIP). One major challenge in constructing a large-scale QIP device will be to incorporate feedback techniques for performing quantum error correction.

This thesis describes the development of a novel classical control system for ion trap quantum computing incorporating powerful real-time processing, and its use in performing a number of experiments involving such processing which form crucial building blocks for stabilizing large-scale ion trap systems.

A second major component is the demonstration of multi-qubit quantum control in mixed-species ion chains, which allowed low-crosstalk error-check operations to be performed over tens of cycles in a multi-qubit system for the first time. Combined with feedback this allowed the stabilisation of entanglement over extended sequences of operations.

The technical advances in the thesis are a set of control hardware and related firmware and software that is specifically designed to meet the needs of quantum error correction. It enables advanced sequences of measurement, real-time decision making and parameter adjustment needed for scalable experiments, with feedback a core element in its design. Together the feedback-capable system and mixed-species setup were used to test new protocols including a single-qubit adaptive phase estimation scheme relying on rapid real-time classical computation and low-latency parameter updates to optimally extract information, outperforming standard non-adaptive fitting in speed and flexibility.

Single- and mixed-species gates between calcium and beryllium and associated experimental techniques were investigated using registers of two and three ions, leading to the first gates between qubits encoded in optical and hyperfine transitions, which reached two-qubit fidelities above 96% and three-qubit fidelities of 93.8(5)%.

In preparatory steps for further work, a single-species dissipative protocol was used to prepare an entangled steady-state using a new approach devised in our group, while ion transport and separation experiments with up to four single-species and two mixed-species ions into wells 800 µm apart at excitations below ten quanta was implemented and optimised.

The main scientific result of the thesis is the demonstration of the repeated extraction of quantum correlations from a pair of beryllium ions using a calcium ancilla qubit. This type of correlation measurement is critical for performing fault-tolerant algorithms. The measurement was then combined with real-time feedback in order to stabilize beryllium qubits in both subspaces and in entangled states, for sequences including up to fifty rounds of feedback, an order of magnitude more than previous work. Information on the major infidelities in the protocols was extracted from the measurement outcome correlations.

This thesis concludes with an outlook for extending the role of both classical and quantum feedback in trapped-ion QIP experiments.

This is the second edition of the thesis, released on the 27$^{\text{th}}$ of September 2018, incorporating minor corrections. The first edition was released on the 13$^{\text{th}}$ of July 2018.

Zusammenfassung

Gefangene Ionen gehören zu den führenden Plattformen bei der Realisierung von Quanteninformationsverarbeitung (QIP). Eine grosse Herausforderung bei der Konstruktion eines skalierbaren QIP-Systems wird es sein, Rückkopplungstechniken zur Quantenfehlerkorrektur zu implementieren.

Diese Dissertation beschreibt die Entwicklung eines neuartigen klassischen Kontrollsystems für einen Quantencomputer auf Basis gefangener Ionen mit leistungsfähiger Echtzeitverarbeitung und seine Verwendung bei der Durchführung einer Reihe von Experimenten, welche entscheidende Bausteine für die Stabilisierung grosser Ionenfallen-Systeme bilden.

Eine zweite wichtige Komponente ist die Demonstration der Multi-Qubit-Quantensteuerung in Ionen-Ketten gemischter Spezies, die erstmals die Durchführung von Low-Crosstalk-Fehlerüberprüfungsoperationen über Dutzende von Zyklen in einem Multi-Qubit-System ermöglicht. Kombiniert mit Feedback-Techniken ermöglichte dies die Stabilisierung der Verschränkung über ausgedehnte Arbeitsabläufe.

Die technischen Weiterentwicklungen in dieser Dissertation beinhalten eine Reihe von Steuerungshardware sowie zugehörige Firmware und Software, die speziell für die Anforderungen der Quantenfehlerkorrektur entwickelt wurde. Diese ermöglichen fortschrittliche experimentelle Messabläufe, Entscheidungsfindung und Parameteranpassung in Echtzeit, was für skalierbare Experimente erforderlich ist, wobei die Feedback-Funktion ein Kernelement des Designs ist.

Zusammen mit dem rückkopplungsfähigen System und dem Mixed-Spezies-Setup wurden neue Protokolle getestet, einschliesslich eines adaptiven Phasenabschätzungsschemas, das sich auf schnelle klassische Echtzeitberechnungen und Parameter-Updates mit niedriger Latenz stützt, um Informationen optimal zu extrahieren, und gängige, nicht adaptive Anpassungen in Geschwindigkeit und Flexibilität übertrifft.

Einzel- und Mischgatter zwischen Calcium und Beryllium und damit verbundene experimentelle Techniken wurden mit Registern von zwei und drei Ionen untersucht, was zu den ersten Gattern zwischen den in optischen und hyperfeinen Übergängen kodierten Qubits führte, die eine Zuverlässigkeit von über 96% bei zwei Qubits und von 93.5% bei drei Qubits erreichten.

In Vorbereitung künftiger Forschung wurde ein Ein-Spezies-Dissipationsprotokoll verwendet, um einen verschränkten Gleichgewichtszustand unter Verwendung eines neuen, in unserer Gruppe entwickelten Ansatzes herzustellen, während Ionentransport- und Separationsexperimente mit bis zu vier Ein-Spezies- und zwei Misch-Spezies-Ionen in Fallenpotenzialen mit einem Abstand von 800 µm bei Anregungen unter zehn Schwingungsquanten durchgeführt und optimiert wurden.

Das wichtigste wissenschaftliche Ergebnis der Arbeit ist die Demonstration der wiederholten Extraktion von Quantenkorrelationen aus einem Paar Berylliumionen mit Hilfe eines Calcium-Ancillaqubits. Diese Art der Korrelationsmessung ist entscheidend für die Durchführung fehlertoleranter Algorithmen. Die Messung wurde dann mit Echtzeit-Feedback kombiniert, um Beryllium-Qubits in beiden Teilräumen und in verschränkten Zuständen über Sequenzen von bis zu fünfzig Feedback-Runden zu stabilisieren, eine Grössenordnung mehr als bei früheren Arbeiten. Aus den Korrelationen der Messergebnisse wurden Informationen über die wichtigsten Fehlerquellen in den Protokollen extrahiert.

Diese Arbeit schliesst mit einem Ausblick auf die breitere Nutzung von klassischem und Quantenfeedback in QIP-Experimenten mit gefangenen Ionen.

Acknowledgements

Foremost I would like to thank Prof Jonathan Home. His enthusiasm, knowledge on a wide multitude of topics, creativity and at times direct experimental assistance have been absolutely essential to this work. It has been a great pleasure working in the TIQI group thanks largely to his attentive and good-natured supervision and guidance.

I owe the past and present members of the mixed-species team a great deal, especially Matteo Marinelli, who has been my partner in the mixed-species experiments discussed in this thesis over the last few years. Thanks for your tireless maintenance of beryllium lasers and optics, perseverance in getting things running, hard work in writing, testing and debugging experiments, and coaxing more out of Ionizer than I thought was possible. Next time you spend a late night taking data, don't forget to take a break with the Abbott simulator. Karan Mehta has also been instrumental to this work. Thanks for your rigor, insistence on pinning things down and refusal to sweep things under the rug. Your approach has revealed countless experimental problems as well as their solutions. Thanks a lot for all the long brainstorming and planning sessions too, and the dry wit was appreciated more than you know. Hsiang-Yu Lo, thank you for getting the parity measurements started at a time we were on the edge of despair, and setting us on the right path with your otherworldly optics skills and physics insights – we're all still trying to emulate the beryllium master, and the experiments I discuss owe much of their approach to your work. Christa Flühmann, thanks for being the calcium counterpart to Hsiang-Yu, helping keep the mixed-species experiments running behind-the-scenes in many ways – not least, bothering to calibrate the things that we perennially assumed were fine until they weren't. Your bug-finding skills were highly appreciated, and have taught me never to get complacent about code I think is battle-tested. I appreciate your fixing my endless optics mistakes too. Daniel Kienzler, thanks for teaching me the ways of the lab when I was getting started, your help on the original calcium two-qubit gates and rf experiments, useful ideas and feedback on the M-ACTION system, and countless other apparatus (including more or less the lab itself). Ludwig de Clercq, thanks for all your work on lab infrastructure, especially the DEATHs - if it wasn't for your spearheading of their design, we'd probably be using batteries and potentiometers (or worse still, a stack of DAC eval boards). Your work on transport and separation were also crucial. David Nadlinger, thanks for the massive overhaul of much of our codebase, without which many of the experiments in this thesis would never have worked, and bring a wealth of experience to bear on some tricky software engineering just when we needed it most. Ben Keitch, your efforts in getting the control system started on a solid base have paved the way for most of the work we've done since. Thanks also for getting me going back when I started, specifying the DDS board design, and answering my endless questions. And the new generation of students, Tanja Behrle and Francesco Lancellotti, for the big contributions you have already made to the lab apparatus – I think they got our final experiments over the line. I couldn't be leaving things in better hands.

Thanks also to the cryo and fibre trap teams. Frieder, Florian and Joseba, you were excellent testers before I realised I needed them, and have shaped much of the control code we're using today (for better or worse). If there was a prize for constructive and very useful

criticism, you three would win hands down. Thanks for all the company in and away from the lab too.

I'd like to acknowledge all the effort from the excellent stream of semester and Masters students who have contributed to our apparatus over the years – in no particular order, Alex Hungenberg, Lukas Gerster, Martin Sepiol, Karin Fisher, Christoph Fischer, Nelson Darkwah Oppong, Brennan MacDonald-de Neeve and Martin Stadler[1].

Thanks also to Chris Ballance for all the DDS discussions when we were getting our respective systems up and running, as well as your physics learning advice during the week-long period our time at ETH overlapped. John Adair and Reg Fitzer at Enterpoint, for helping debug numerous technical issues. Peter Clements, Marc Oberholzer and Philipp Romann at Enclustra for all your time spent helping us get the new electronics up and running.

I'd like to thank all the group members, past and present, who have helped our collective productivity with a game or three of Kicker[2]. It's been a great leveller, forcing people at all levels of the group to work together in the quest to get the other team under the table. Also, thanks to everyone who provided last-second feedback on this thesis!

Finally I would like to thank my parents, Sveta and Misha, and my grandparents, Bella, Vlad, Sonia and David, for the love and trust that have shaped me, and the opportunities and support that you have always given me.

[1] I'm doubtless forgetting several people, for which I apologise.

[2] Also known as table soccer, foosball, Töggele, or 'spin for the win'.

Contents

1 Introduction

Machines that process information form the backbone of the modern world. We are surrounded by both general-purpose and specialised classical information processors, in the form of computers and digital electronic circuits; they have facilitated many of the technological advances of the last half-century and have become indispensable to our everyday lives. As the exponential growth in performance of silicon-based classical information processing slows down [198], however, technologies such as quantum information processing (QIP) may soon speed up certain tasks handled today by classical devices. QIP offers greater promise than just a speed-up, however: several important technologies, such as quantum computation, simulation and communication, will always be outside the realm of any conceivable classical device [139].

This chapter briefly describes QIP, introduces trapped-ion QIP experiments and their scaling, and discusses the role of error correction and feedback in a trapped-ion quantum information processor, motivating the hardware development and experiments described in later chapters.

1.1 Quantum information processing

The profound insights of information theory form the basis of information processing. Two of its foundational ideas are that the complete state of any conceivable system (classical or quantum) can be encoded to arbitrary precision in a register of classical bits, and that the evolution of the system can be seen as a computation or a simulation acting on these bits. A *universal computer* is an information processor capable in principle of performing an arbitrary computation on the bits stored in its finite memory, given unlimited time. Any such computation can equally be considered a *simulation* of a specific physical system whose state can be encoded into the computer's memory: for the computer there is no practical difference.

QIP relies on the unique properties of quantum superposition, interference, entanglement and projective measurement to carry out operations within a richer paradigm than classical devices. This allows the resources required for a computation (e.g. time, energy, processor size) to scale very favourably compared to classical devices for certain problems; this is known as *quantum speed-up*. An intuitive motivation follows [139]. Just as any classical system can be mapped to a register of bits, any quantum system can be treated as a register of *qubits*. The general state of a qubit can be written as $a\,|g\rangle + b\,|e\rangle$, where $|g\rangle$ and $|e\rangle$ are two discrete states of the system, and a and b are complex numbers readily encoded to arbitrary precision in classical bit strings[1]. Adding qubits to a dynamically interacting quantum system does not linearly increase the classical memory required, however: for every extra qubit, the memory requirements *double*. This occurs due to the possibility of quantum entanglement: we require four complex numbers to represent a general entangled two-qubit state $a\,|gg\rangle + b\,|ge\rangle + c\,|eg\rangle + d\,|ee\rangle$, yet eight are needed for a three-qubit state $a\,|ggg\rangle + b\,|gge\rangle + \ldots + g\,|eeg\rangle + h\,|eee\rangle$. For this reason it becomes infeasible to model

[1] Neglecting normalisation, which removes one degree of freedom.

a general quantum system larger than around 50 qubits' worth on a classical information processor. However, in principle a quantum system can be encoded in a *quantum* memory merely proportional in size to the system state, and simulated by a quantum information processor; this is generally known as *quantum simulation* [52]. The quantum domain also offers computational advantages which provide a dramatic speed-up for certain problems, such as the factoring of integers [181], inverting functions and searching databases [58], and linear algebra operations [61]. Quantum superposition and interference permits 'parallelised' operations in ways that are not possible in a classical machine, although extracting the results can be a subtle task due to the projective nature of quantum measurement. This field of QIP is sometimes referred to as *quantum computation*[2]. Quantum computation and simulation are growing more closely intertwined, and hybrid approaches such as quantum variational algorithms have been proposed to combine the strengths of both in quantum chemistry calculations [120], with proof-of-concept experiments already in progress [62].

Like their classical analogues, quantum information processors may be either specialised, able to run only a single simulation or algorithm, or universal, capable of being applied to any quantum problem that can be encoded in its memory. A large-scale universal device is a long-term ambition of experimental QIP, and the progress towards it has led to many surprising developments such as quantum error correction, discussed in §1.3. The widely-accepted universality requirements are high-fidelity quantum state preparation and measurement, a universal set of single- and multi-qubit quantum gates, and a high ratio of quantum coherence time to gate time [41]. While various platforms such as cold quantum gases, atomic optical lattices and a range of solid-state systems have been used to carry out dedicated quantum simulations, not all currently meet these requirements for universality, which also turn out to be important even for single-purpose devices. Among the universal platforms, the forerunners are superconducting Josephson-junction qubits [40, 51] and trapped atomic ions [204, 122].

1.2 Trapped-ion quantum information

The requirements for an accurate and stable metrological reference are almost the same as for universal QIP, and trapped-ion QIP arose from atomic clock experiments [205]. An ion trap is a set of electrodes, tens of microns to tens of centimetres in size, arranged to produce confining electric or electromagnetic potential wells at particular sites where ions can be trapped and held. A trap is mounted within a vacuum chamber, and the electrode potentials are usually controlled electronically. External lasers are directed at the trapping sites to manipulate the internal and motional states of the ions. The quantum state of a trapped-ion qubit is measured by applying light on a transition which couples to one of the qubit levels, thereby projecting the state either into this level and causing the ion to scatter photons, or into the other qubit level, in which case the ion will remain dark [204]. The photons can be detected using a high-sensitivity camera or photomultiplier tube to determine the ion state.

The most common ion trap design used in QIP and related areas is the Paul trap [148]; others such as the Penning trap are not considered here. The Paul trap uses dc and rf electric fields for confinement; the trap layout, rf and dc parameters may vary widely de-

[2] Another area of QIP research is quantum communication. This takes advantage of quantum principles to communicate information in ways that are beyond classical communication, for example entangling systems remotely or guaranteeing private communication due to the collapse that is caused by quantum measurement. It is less directly relevant to the work in this thesis.

pending on the experimental requirements. Typically the trap is held at a vacuum pressure below 10^{-11} mbar, providing excellent environmental isolation and minimising the effect of background gas collisions [112]; single ions can be trapped for hours. The experimental setup discussed in §2.1 of this thesis, for example, is based around a Paul trap consisting of a stack of four gold-plated alumina wafers, 2 electrodes to which radio-frequency voltages are applied, 28 electrodes to which static voltages are applied and 28 larger electrodes used to null out stray electric fields [81]. Electrodes held at static potentials are often called 'dc' electrodes, although no current flows.

The ions are usually Doppler-cooled and their qubit state read out using dipole-allowed transitions, and the qubits themselves are implemented in a pair of hyperfine ground states, or a ground and excited state separated by an optical dipole-forbidden transition [60] (both are used in this thesis). Hyperfine qubit transitions are driven by microwaves [143, 142] or Raman laser setups [123, 107], and optical qubits by narrow-linewidth lasers [133, 176]; the transitions are used for sideband cooling [125, 129], quantum gates [123, 94] and other coherent internal and motional state manipulation [115]. An advantage for trapped-ion QIP over solid-state approaches that unlike solid-state qubits, every ion of a particular species is identical. Multiple cooled ions naturally form a regular crystal, where tunable trap confinement parameters and mutual repulsion balance out to determine the ions' equilibrium positions. This large spacing between ions (larger than the diffraction limit) facilitates reliable individual addressing with laser beams [30, 204], and reproducible separation and recombination of ion chains [169]. The ions couple to one another via their mutual repulsion, and their shared motional modes are a convenient means of generating entanglement or transferring information [30, 185]. More details on these operations are presented later in this thesis.

The impressive control of trapped ions has facilitated their successful use in QIP, atomic clock and other metrological experiments [174, 19, 202].These have typically used linear strings of ions within ion traps capable of forming only one or a few trapping sites; a scalable QIP platform requires a fundamentally different approach. A single well can hold at most a few tens of ions if high-fidelity entangling gates need to be carried out: the mode spectrum grows more crowded with increasing ion number, and gates relying on the coupled ion motion must run more slowly [122]. An alternative approach that could extend to hundreds of qubits or more is using more complex, microfabricated traps where different operations such as detection, cooling and quantum gates take place in physically separate zones, with ion transport and separation/recombination playing a major role, as well as the use of a second ion species for sympathetic cooling of the qubit-containing ions. This is known as the 'quantum charge-coupled device' (QCCD) architecture [204, 76], with recent proposals including neighbouring arrays of interconnected traps [98] that would feature microwave electronics and photon detectors directly integrated in the trap [2, 183]. Integrated optical beam delivery is also a near-term possibility [116]. A strategy that has attracted recent interest is to use a set of simpler 'unit-cell' traps holding 5-10 ions, with the traps linked by imperfect photonic interfaces, to carry out a large-scale algorithm [122, 137, 124]. This is already feasible with existing trap technology, however photonic interfaces remain too lossy for large entangled states to be prepared in reasonable experimental times, and the entanglement fidelities achieved mean that they do not currently provide a significant resource for complex protocols. The QCCD and photonically-coupled unit cell architectures are largely complementary, with hybrid approaches likely to emerge as photonics, trap fabrication and experimental control technologies advance [24].

1.3 Feedback in quantum information processing

Beyond finding architectures to control and couple tens or hundreds of ions, scaling up trapped-ion QIP requires deep qualitative changes to the way most experiments are operated today [194], particularly in the way infidelities are handled. A QIP algorithm conceptually resembles a series of precise rotations of a state vector in multidimensional space, where the rotation axes and angles are functions of physical experimental parameters such as laser frequency, intensity or trap strength. There are broadly three sources of infidelity in trapped-ion QIP: dc or slowly-varying systematic errors, faster noise such as laser beam intensity fluctuations or magnetic/electric field noise, and decoherence processes due to more fundamental physical effects such as radiative decay in optical qubits, off-resonant excitation or spectator modes in quantum gates, Raman scattering, collisions with background gas, and field emission [204]. Classical feedback in the form of rapid parameter recalibrations and continuous servo loops, and quantum feedback in the form of error correction, are critical to mitigating these.

Systematic errors from experimental parbameter drifts, duty cycle issues or miscalibration may generally be addressed either by engineering the experimental system to be more passively stable, which becomes increasingly difficult as a system grows in size and complexity, or recalibrating the experimental parameters often enough that the error magnitude remains minimal. When dealing with tens of calibrated parameters, as in the experiments discussed in later chapters, this requires the majority of experiment time when carried out manually. Once one round of calibrations is complete, the next must already begin to maximise high fidelities. This can be mitigated by using adaptive Bayesian schemes to measure parameters more rapidly; such schemes adapt the measurement based on incoming data to obtain information in the shortest time. A proof-of-concept example is presented in §6.2, with ongoing work in the group on more advanced protocols [110, 200]. Such schemes also remove human biases and variations from the calibration procedure, though this benefit is shared with all automatic calibrations. At a higher level, an experimental control system with enough built-in automation to intelligently manage most repetitive experimental calibrations is becoming essential as experiments grow in complexity.

Sources of infidelity fluctuating more rapidly than the several-second timescale of a Bayesian calibration, such as acoustic or electrical noise, can often be tackled using continuous-domain classical feedback, with laser frequency and intensity stabilisation being canonical examples. The approaches used in our laboratory to mitigate fast noise are summarised in §3.4.2. In some cases, such as voltage noise on trap electrodes, passive filtering can be used to remove the noise at critical bandwidths.

Residual errors in the calibration, as well as error from rapid noise sources and intrinsic decoherence processes that cannot be dynamically decoupled, accumulate over time during a QIP algorithm, scrambling the state vector direction and limiting the feasible algorithm length[3]. These error sources require a more fundamental form of feedback known as quantum error correction (QEC), which in principle can achieve almost noise-free QIP using noisy components [194]. An analogue occurs in classical digital communication for a noisy channel, where each transmitted bit has a probability of being incorrectly received. To improve the reliability of the channel, a classical error correcting code can be used: n logical bits are

[3] In classical digital logic, a bit has only two states, allowing systems to rely on bistable circuits and positive feedback (the continuous-domain limit of discrete majority-voting error correction) to mitigate errors at a fundamental level. This is less straightforward in universal QIP due to the continuous nature of qubit states, gates and protocols.

encoded redundantly into m physical bits, where $m > n$, such that bit flips on the physical qubits can be detected and corrected[4]. The theoretical minimum ratio of m to n depends on the physical and desired logical bit error probabilities according to the noisy-channel coding theorem [139]; arbitrarily low logical error rates may be achieved by increasing m. QEC is similar; multiple physical qubits are used to encode one or more logical qubits, and individual errors in the physical qubits can be detected and corrected without collapsing the logical qubit state [182, 27, 188]. For a polynomial increase in qubit number, an exponential decrease in logical error rate can be obtained.

Measuring the code qubits directly would collapse their quantum state, so instead they are entangled with an ancillary qubit outside the code. The act of measuring an ancilla projects part or all of the code to a state of 'error' or 'no error', to allow a potential error to be cleanly measured. With several measurements, an error can be localised to a particular physical qubit and manually corrected. For large error codes, this error decoding requires rapid classical computation and feedback to minimise decoherence. QEC can be generalised to implement *fault-tolerant* codes, which can cope with errors in the ancilla measurement and error correction operations as well as the qubit state itself [154, 53, 190].

The form of error correction is dependent on the type of error model. The most common and well-defined models involve uncorrelated noise acting on individual qubits, which leads to feedback-based quantum error correction as described above, and correlated noise acting equally on all qubits, which motivates *decoherence-free subspace* qubit encoding schemes [77] which are resistant to correlated noise in at least one basis at the cost of more physical qubits. Combining both QEC and decoherence-free subspace encoding in two levels can make the logical qubits robust to both correlated and uncorrelated noise (in at least one noise basis). Quantum feedback also plays a major role in QIP aside from QEC, often being used to drive the computation itself, and many protocols such as the widely-known quantum teleportation [139] fundamentally rely on real-time feedback based on measurement outcomes. QEC can be combined with measurement-based schemes to reduce the overall number of measurements required in a quantum computation [190].

In certain cases, dynamical decoupling techniques offer ways to 'engineer' a Hamiltonian interacting with a quantum system such that the effect of a particular noise source is reduced [17, 100]. Unlike passive schemes such as decoherence-free subspaces it typically requires an active drive to be applied, which makes it well-suited for improving the robustness of quantum gates [16, 192, 9, 113]. Although not a universal approach, it addresses one of the major error sources in QIP by reducing quantum gate infidelity.

The partial measurement of a quantum state in QEC and other protocols using ancillary qubits must be carried out in such a way that the code qubits are unaffected. In trapped-ion approaches, detection relies on photon scattering, which heats the motion of the ions. This impairs the ability to perform further multi-qubit gates, which rely on the ions being cooled close to the motional ground state. Sympathetically cooling the code ions using an ion of the same species results in scattered light collapsing the code qubit states. This forms a major motivation for using two different ion species, one for the code and one for the ancilla qubits, whose detection/cooling and qubit transition frequencies are distant from one another, allowing fully independent control of each species. The ancillas can then be used to sympathetically cool the code qubits after readout [68]. Although the experimental setup requires a second set of lasers and detection apparatus, the benefits of mixed-species

[4] A repetition code is a simple (and inefficient) example: 0 and 1 are represented as 000 and 111, so that even if one of the bits is flipped, the correct value can be inferred via a 'majority vote'.

QIP have led to multiple experimental demonstrations in the last decade [67, 9, 193, 50, 26].

This thesis discusses feedback in the various forms outlined above, carried out using mixed-species experiments between calcium and beryllium ions using an experimental control system developed primarily with flexible feedback in mind.

1.4 Thesis outline

Chapter 2 presents the segmented ion trap apparatus and basic single-qubit coherent operations, introduces the relevant level structures of calcium-40 and beryllium-9 qubits, summarises dissipative operations including state detection and preparation, and gives an overview of the vacuum system and the calcium and beryllium lasers. The requirements of an experimental control system able to flexibly feed back on external factors as well as quantum measurement outcomes are closely related; Chapter 3 discusses the design and development of a system meeting these goals, whose features were relied upon in the experiments performed during this thesis. Several auxiliary electronic systems important to laboratory experiments are also discussed therein.

Having established the physical and electronic apparatus, Chapter 4 presents typical experimental operations in more detail, including ion loading, mixed-species detection and associated infidelities, Rabi oscillations and experimental calibration procedures.

Chapter 5 summarises the ion transport and separation operations and the experiments used to characterise and optimise them, as well as several experimental techniques such as transport-based loading, ion reordering and recrystallisation that were important for the mixed-species experiments.

With these building blocks in place, Chapter 6 discusses the various coherent experiments and entangling gates carried out during this thesis. This includes single-ion Ramsey experiments and Bayesian phase estimation, motional coherence measurements on three ions, single- and mixed-species multi-qubit gates and their calibration, Bell state coherence measurements, and a two-qubit dissipative entanglement scheme.

Apart from QEC, quantum feedback is critical to other fundamental protocols such as quantum teleportation and entanglement distillation, and experiments using it to full effect in QIP have only begun to emerge in the last few years [149, 162, 34, 128]. With a full set of mixed-species operations in place, Chapter 7 discusses the calibration and operation of a parity readout of two beryllium ions using a calcium ancilla, and the operation of a three-qubit quantum feedback loop, used to stabilise various subspaces and entangled states in a way that shares many similarities with a full QEC code. Unlike previous work to date [34, 128], the loop is run over up to 50 measurement events, incorporating ancilla detection, recooling and recycling in a fundamentally scalable approach. Two different forms of feedback are investigated, one using global rotations, the other using single-qubit rotations implemented via transport of the crystal, and correlations in the ancilla measurements are investigated to better understand the dynamics of the feedback loop.

Chapter 8 presents the future applications of this thesis work in mixed-species QIP experiments, discusses some of the limitations encountered and suggests potential experimental improvements. The appendices of this thesis cover further details of the experimental control system, several esoteric experimental calibration procedures, a detailed mathematical derivation relevant for Bayesian phase estimation, and a list of abbreviations used throughout the main thesis text.

2 Mixed-species experimental setup

The experiments presented in this thesis were carried out in a segmented three-dimensional Paul trap designed to simultaneously hold beryllium and calcium ions, referred to as the *segmented trap*. This chapter introduces ion trap physics, presents the segmented trap, and briefly reviews the coherent state manipulation, state readout, cooling and state preparation techniques used in the rest of this thesis for beryllium and calcium. It concludes with an overview of the vacuum, laser and imaging apparatus. Chapter 6 will present more details on the experimental operations. Please see the doctoral theses of Daniel Kienzler [81] and Hsiang-Yu Lo [105], who planned and constructed most of the setup, for a more in-depth presentation of the apparatus.

2.1 Segmented trap

The segmented trap used in this work is made from a stack of aluminium oxide (Al_2O_3) wafers. Each wafer was laser-machined to create slots and form electrodes in the desired geometry. A thin layer of gold was evaporated onto the electrode wafers, then a thicker layer was electroplated on top to make the electrodes conducting. The two inner electrode wafers have one rf electrode and 15 dc electrodes each, arranged anti-symmetrically with each other as shown in Figure 2.1a.

The widths of the dc electrodes are nonuniform, to improve the trap confinement in some areas without requiring an excessive number of electrical connections, which would be the case for uniformly narrow electrodes. They are shown in Figure 2.1e. There are two loading zones at either end of the trap, with electrodes 500 µm wide; the maximum axial trap confinement in these regions is lower than in the central regions. Further towards the centre are two experimental/storage zones with 300 µm-wide electrodes, that are useful when multi-well operations are carried out in the trap. Next are two groups of three narrow electrodes 155 µm wide, optimised for maintaining tight axial confinement during ion separation or combining operations [81], and finally a central experimental/storage electrode.

The main experimental zone in the centre of the trap is used for most of the experiments discussed in this thesis. The electrode-ion distance in the trap was chosen to be 184 µm, a compromise between a more compact trap with higher electric field gradients (and thus lower required external voltages) for smaller distances, and lower heating rates and better optical access for larger distances. Custom arbitrary-waveform generators produce the dc electrode voltages, discussed further in Chapter 5.

The voltage for the rf electrodes $v_{\mathrm{rf}} \cos(\Omega_{\mathrm{rf}} t)$ is generated at $\Omega_{\mathrm{rf}} = 2\pi \times 115.2\,\mathrm{MHz}$ using a stable rf source, amplifier and helical resonator [81], with an amplitude of $v_{\mathrm{rf}} \sim 320{-}380\,V$. The trapping principles are briefly summarised below. The electric potential close to a point along the central trap axis that we take as $x, y, z = 0$ is given by

$$V(x,y,z,t) = \frac{1}{2} v_{\mathrm{dc}} \left(\alpha_z z^2 - \alpha_x x^2 - \alpha_y y^2\right) + \frac{1}{2} v_{\mathrm{rf}}(t) \left(\beta_x x^2 - \beta_y y^2\right) \tag{2.1}$$

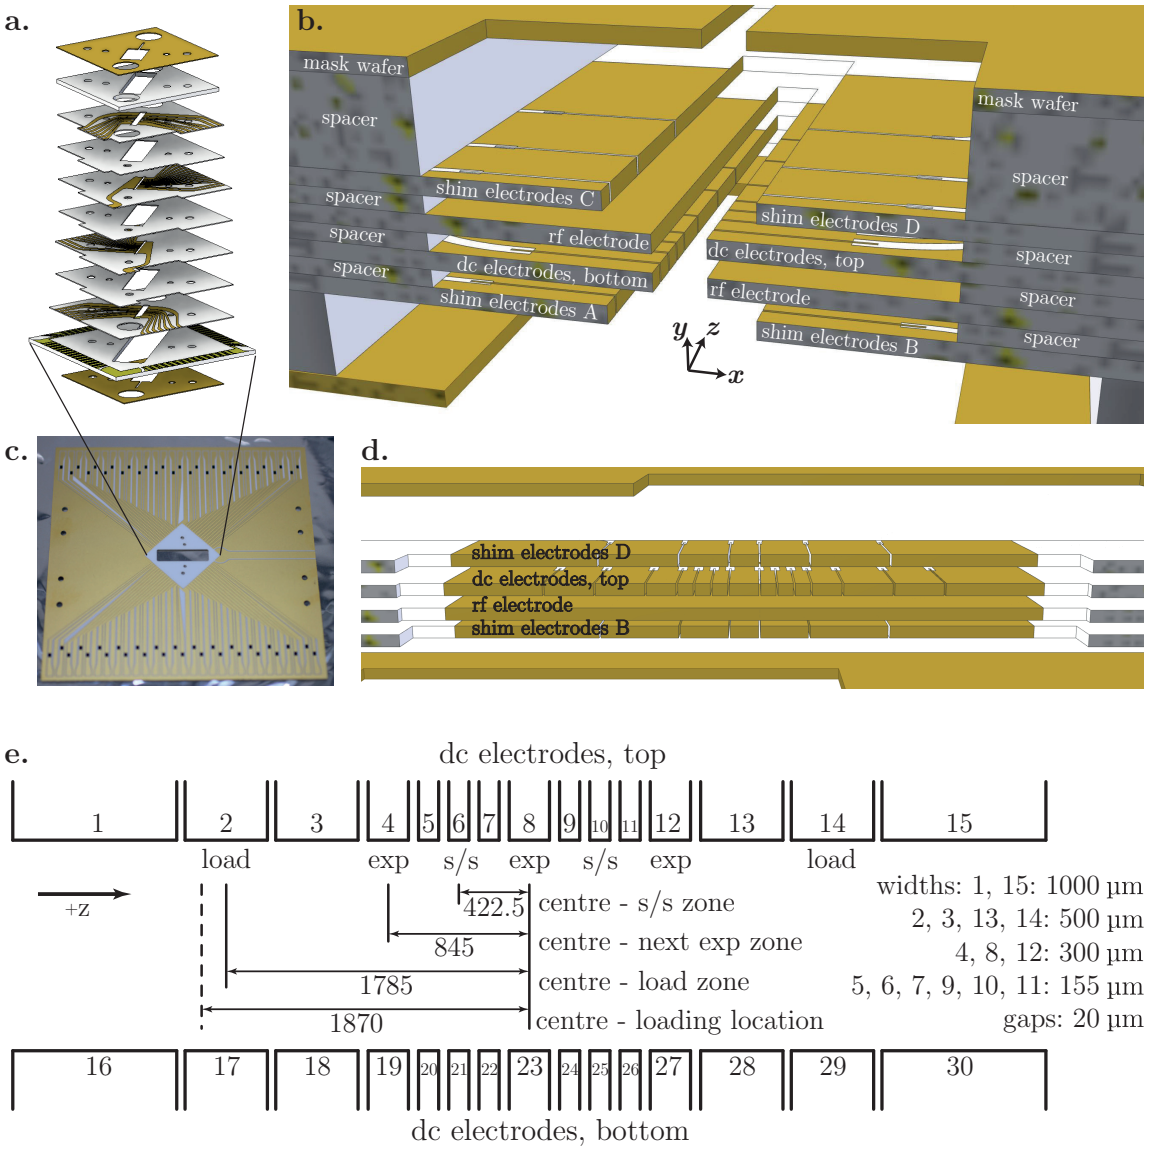

Figure 2.1: Segmented trap structure. **a.** Trap wafer stack-up, showing the wafer ordering and geometry. **b.** Detail of wafer stack when assembled; note the dc, rf and shim electrode positions. **c.** Empty trap filterboard; the wafer stack shown in a. was mounted in the centre, and the filterboard is oriented vertically in the vacuum chamber. **d.** Side view of dc and shim electrodes. **e.** DC electrode numbering and relevant dimensions for this work. All distances are in µm. Adapted from [81], Figures 3.6, 3.7, 3.8, 3.13.

where the curvature v_{dc} is a function of the dc electrode voltages (see §5.1), $\alpha_{x,y,z}$ and $\beta_{x,y}$ are positive geometric factors with $\beta_x = \beta_y$. Higher-order polynomial terms are neglected. Along the z axis the potential is harmonic, and a singly-charged ion oscillates with an angular frequency $\omega_z^2 = v_{\mathrm{dc}}\alpha_z e/m$, where m is the ion mass (cf. Equation 5.4). The radial ion motion can be found by solving $d^2u/dt^2 = -e\partial^2 V(x,y,z,t)/m\,\partial u^2$ with $u = x, y$, which can be cast as Mathieu equations [148, 204] with Mathieu parameters

$$a_u = -\frac{4e\alpha_u v_{\mathrm{dc}}}{m\Omega_{\mathrm{rf}}^2}, \qquad q_x = -q_y = \frac{2e\beta_u v_{\mathrm{rf}}}{m\Omega_{\mathrm{rf}}^2} \tag{2.2}$$

An exponential series can be used to solve the Mathieu equations [95]. We operate the segmented trap in the regime where $|a_u|, q_u^2 \ll 1$, which gives approximate solutions

$$u(t) \simeq u_0 \cos(\omega_{u,\mathrm{rf}} t)\left[1 - \frac{q_u}{2}\cos(\Omega_{\mathrm{rf}} t)\right] \tag{2.3}$$

where x_0, y_0 depend on the initial conditions, and $\omega_{x,\mathrm{rf}}, w_{y,\mathrm{rf}}$ are *secular frequencies* of oscillation along the x and y axes given by

$$\omega_{u,\mathrm{rf}} = \frac{\Omega_{\mathrm{rf}}}{2}\sqrt{a_u + \frac{q_u^2}{2}}. \tag{2.4}$$

Neglecting the effect of the relatively-weak dc potential ($a_u \to 0$), we obtain

$$\omega_{u,\mathrm{rf}} \simeq \frac{\sqrt{2} e \beta_u v_{\mathrm{rf}}}{m \Omega_{\mathrm{rf}}}. \tag{2.5}$$

A shared voltage offset is typically applied to the dc electrodes, which effectively alters the geometric coefficients β_x and β_y in the static quadrupole. This breaks the degeneracy of ω_x and ω_y and rotates the radial mode vectors, which is needed to ensure that all modes overlap with the wavevectors of beams passing through the central slot of the trap along the y axis.

The ion position described by Equation 2.3 oscillates in the radial dimensions at the secular frequencies, modulated by smaller, rapid oscillations at the rf frequency which are called *micromotion*. Micromotion has an intrinsic component, described by the above equation, and an excess component that occurs when there is a stray radial dc electric field (linear potential term in x or y) that shifts the harmonic minima in Equation 2.1 away from $x = y = 0$. This field alters Equation 2.3 to

$$u(t) \simeq u_0 \left[u_e + \cos(\omega_{u,\mathrm{rf}} t)\right]\left[1 - \frac{q_u}{2}\cos(\Omega_{\mathrm{rf}} t)\right], \quad u_e = \frac{m E_u \Omega_{\mathrm{rf}}}{2 e \beta_u v_{\mathrm{rf}}} \tag{2.6}$$

where u_e (x_e, y_e) is the new radial equilibrium position of the ion, and E_u (E_x, E_y) are the excess field components [15, 81]. Note that heavier ions are displaced further by a radial field.

Stray electric fields can be caused by surfaces in or near the trap charging up, often due to UV light, with the charge fluctuating slowly over hours or days. They can cause undesirable effects, especially in mixed-species experiments where the ion masses differ, and are experimentally compensated. As shown in Figure 2.1d the wafers above and below the central dc/rf electrodes have 14 larger dc electrodes each, arranged symmetrically on either side of the trap. These 'shim' electrodes are used to cancel out stray radial dc fields along the trap; the protocol used is discussed in §4.6.1.

Special mask wafers, one above and one below the trap, each cover up the loading zones and neighbouring regions, exposing only a narrow slit along the trap axis to the flux from the atomic Be and Ca ovens. This is to avoid the neutral atoms from the ovens building up on the electrodes, roughening the surfaces and increasing ion heating rates and motional decoherence [204]. The masks limit the area of the trap that can be imaged from each side to slightly more than half, however the masks cover opposite ends of the trap so the imaging systems on either side can be used to image different ends.

The wafer stack is mounted on a larger filterboard, shown in Figure 2.1c, which is connected to the electrode wafers via wirebonds. Most dc (and shim) lines have first-order

RC filters, with $R = 240(20)\,\Omega$ and $C = 820(40)\,\text{pF}$ for a cutoff frequency of 809(78) kHz. Four of the shim lines have only the capacitors present, to allow for rf at several megahertz to be coupled onto certain electrodes for tickling (discussed in §4.5.1). The trap filterboard is mounted inside the vacuum system and connected to external dc and rf voltages.

The discussion of the apparatus is continued in §2.6; to better motivate the system design the theoretical background to the quantum control of ions is presented first.

2.2 Quantum control of ions

Trapped-ion QIP relies on deterministically initialising the internal and motional states of a set of ions, coherently manipulating them using single- and multi-qubit gates, and projectively reading out the state of some or all of them [41]. All three of these must be performed repeatedly to scalably carry out an error-corrected quantum computation. As mentioned in the previous chapter, it is highly advantageous to use different ion species, one for storing the quantum state and one for sympathetic cooling and partial state readout for quantum error correction. An example of this in Chapter 7 of this thesis. Below the general framework of single-ion coherent manipulation is presented, followed by carrier Rabi oscillations and single-qubit gates. Internal state preparation, cooling and state readout are discussed after an introduction to the beryllium and calcium level structures.

2.2.1 Coherent state manipulation

Coherent manipulation is similar for beryllium and calcium, involving two internal levels $|e\rangle$ and $|g\rangle$ that together form a pseudospin or a qubit with state $a\,|g\rangle + b\,|e\rangle$ as discussed in the introduction. Laser or rf fields are used to coherently drive the $|g\rangle \leftrightarrow |e\rangle$ transitions to manipulate both the internal and motional states of the two species.

A single trapped ion can be modelled as a particle confined by quantised harmonic potentials along the axial and two radial directions. The ion Hamiltonian, with the two internal qubit levels $|g\rangle$ and $|e\rangle$ separated by energy $\hbar\omega_0$ and a single harmonic vibrational mode along the z axis with frequency ω_m, is[204, 95, 81]

$$\hat{H}_0 = \hat{H}_{\text{qub}} + \hat{H}_{\text{mot}} = \frac{\hbar}{2}\omega_0\hat{\sigma}_z + \hbar\omega_m\hat{a}^\dagger\hat{a} \tag{2.7}$$

where $\hat{H}_{\text{qub}}$ and $\hat{H}_{\text{mot}}$ give the internal and motional energies of the ion, $\hat{a}^\dagger$ and $\hat{a}$ are the raising and lowering operators acting on the motional state $|n\rangle$, and micromotion effects in the Paul trap are neglected. $\hat{\sigma}_x \equiv |e\rangle\langle g| + |g\rangle\langle e|$, $\hat{\sigma}_y \equiv i\,|g\rangle\langle e| - i\,|e\rangle\langle g|$, $\hat{\sigma}_z \equiv |g\rangle\langle g| - |e\rangle\langle e|$, $\hat{\sigma}_+ \equiv |e\rangle\langle g|$, $\hat{\sigma}_- \equiv |g\rangle\langle e|$ are the Pauli operators on the qubit state. In this thesis the qubit levels which are bright and dark during state readout are labelled $|g\rangle$ and $|e\rangle$ respectively; in beryllium $|g\rangle$ is higher in energy than $|e\rangle$.

The analysis below is carried out for a single motional mode, but can be extended to three modes at different motional frequencies [204]. The position operator of the ion is

$$\hat{z} = z_0\left(\hat{a} + \hat{a}^\dagger\right), \quad z_0 = \sqrt{\frac{\hbar}{2m\omega_m}} \tag{2.8}$$

where m is the ion mass and z_0 is the root-mean-square width of the ion ground-state wavefunction. The wavefunction of a calcium-40 ion in a typical well with $\omega_m = 2\pi\times 1.6\,\text{MHz}$ has a width of 9 nm, whereas a beryllium-9 ion in the same well ($\omega_m = 2\pi \times 3.4\,\text{MHz}$) is 13 nm wide.

When the ion encounters a travelling coherent electromagnetic field at frequency ω, that couples the two states electrically or magnetically and is near-resonant with ω_0, the ion-field interaction Hamiltonian can be written as

$$\hat{H}_{\text{int}}(t) = \frac{\hbar}{2}\Omega\left(\hat{\sigma}_+ + \hat{\sigma}_-\right)\left(e^{i(k\hat{z}-\omega t+\phi)} + e^{-i(k\hat{z}-\omega t+\phi)}\right) \tag{2.9}$$

where Ω is the *resonant Rabi frequency* of the laser-ion interaction[1], and is a function of the $|g\rangle$ and $|e\rangle$ wavefunctions and the exact form of the field, k is the field wavevector projection along the z axis, ω is the field frequency and ϕ is its phase at the ion. Optical and rf transitions can be modelled with this approach, as can Raman transitions if the upper level is adiabatically eliminated. The wavevector k of the Raman transitions discussed in §2.4.1 is $k = (\vec{k_1} - \vec{k_2}) \cdot \tilde{\mathbf{z}}$, where $\tilde{\mathbf{z}}$ is the unit vector along the z axis and k_1, k_2 are the wavevectors of the two Raman beams, and the effective laser frequency becomes $\omega = \omega_1 - \omega_2$.

Expanding $\hat{z}$ according to Equation 2.7 and transforming $\hat{H}_{\text{int}}(t)$ into the interaction picture[2] via $\hat{H}'_{\text{int}}(t) = \hat{U}_0(t)^\dagger \hat{H}_{\text{int}}(t)\hat{U}_0(t)$ where $\hat{U}_0(t) = \exp(i\hat{H}_0 t/\hbar)$, we obtain four terms. Neglecting those containing rapidly-oscillating $e^{\pm i(\omega+\omega_0)}$ factors and keeping those containing $e^{\pm i(\omega-\omega_0)}$, we obtain

$$\hat{H}'_{\text{int}}(t) = \frac{\hbar}{2}\Omega\hat{\sigma}_+ \exp\left[i\eta\left(\hat{a}e^{-i\omega_m t} + \hat{a}^\dagger e^{i\omega_m t}\right)\right] e^{i(\phi-\delta t)} + \text{h.c.} \tag{2.10}$$

where $\delta = \omega - \omega_0$ is the field detuning from the qubit and $\eta = kz_0$ is the *Lamb-Dicke parameter*, proportional to the ratio between z_0 and the field wavelength. For the 729 nm calcium qubit beam $\eta \simeq 0.05$ at 45° to the z axis, while for the beryllium co-90 Raman beam configuration $\eta \simeq 0.4$.

Expanding Equation 2.10 in η results in terms with M $\hat{a}^\dagger$-operators and N $\hat{a}$-operators, and these terms are resonant with the field when $\delta \simeq (M-N)\omega_z = s\omega_z$. These resonances are called the $|s|$-th blue (red) sideband when $s > 0$ ($s < 0$) and the carrier when $s = 0$. They couple the states $|g\rangle|n\rangle$ and $|e\rangle|n+s\rangle$, with resonant Rabi frequencies given by

$$\begin{aligned}\Omega_{n+s,n} = \Omega_{n,n+s} &= \Omega_0\left|\langle n+s|\, e^{i\eta(\hat{a}+\hat{a}^\dagger)}\,|n\rangle\right| \\ &= \Omega_0 \exp\left(-\eta^2/2\right)\eta^{|s|}\sqrt{\frac{n_{\min}!}{n_{\max}!}}L_{n_{\min}}^{(|s|)}(\eta^2), \quad \begin{matrix} n_{\min} = \min(n, n+s) \\ n_{\max} = \max(n, n+s)\end{matrix}\end{aligned} \tag{2.11}$$

where $L_{n_{\min}}^{(|s|)}(\eta^2)$ is the generalised Laguerre polynomial

$$L_{n_{\min}}^{(|s|)}(\eta^2) = \sum_{k=0}^{n_{\min}}(-1)^k \binom{n_{\min}+|s|}{n_{\min}-k}\frac{\eta^{2k}}{k!} \tag{2.12}$$

In this thesis only the carrier and first two sidebands are driven, i.e. $|s| \leq 2$. Figure 2.2b shows the variation of $\Omega_{n+s,n}$ with n and s for calcium and beryllium.

When $\eta\sqrt{\langle(\hat{a}^\dagger + \hat{a})^2\rangle} \ll 1$, the ion is in the *Lamb-Dicke regime*. We can expand Equation 2.10 only to first order in η in this limit to obtain:

$$\hat{H}_{\text{LD}}(t) = \frac{\hbar}{2}\Omega_0\,\hat{\sigma}_+\left[1 + i\eta\left(\hat{a}e^{-i\omega_m t} + \hat{a}^\dagger e^{i\omega_m t}\right)\right] e^{i(\phi-\delta t)} + \text{h.c.} \tag{2.13}$$

[1] In some references Ω is defined as half of this value (i.e. the factor of 1/2 is incorporated into its definition), which changes its interpretation from the oscillation frequency of *probabilities* to that of *probability amplitudes*.

[2] A Hamiltonian $\hat{H}$ can in general be transformed to a different basis using $\hat{H}_{\text{trans}} = \hat{U}\hat{H}\hat{U}^\dagger + i\hbar\frac{\partial\hat{U}}{\partial t}\hat{U}^\dagger$.

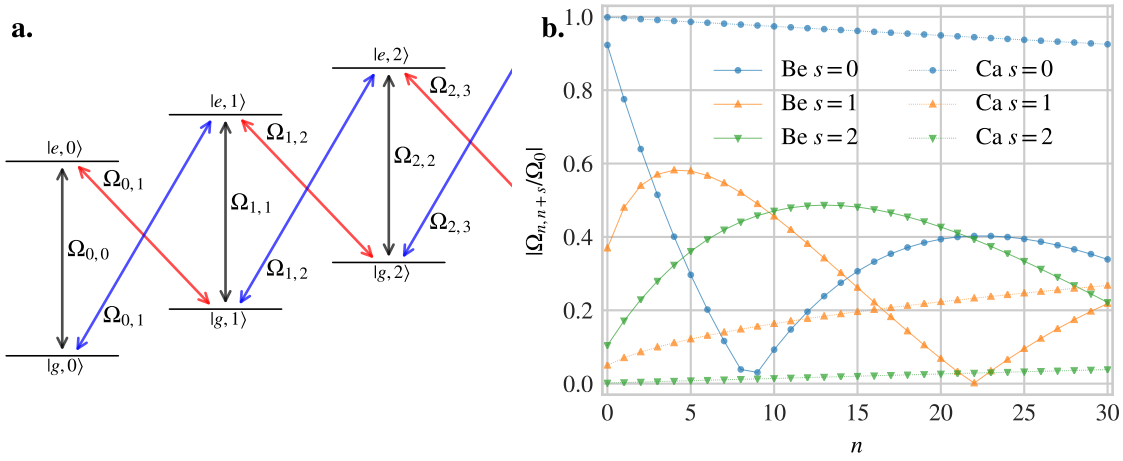

Figure 2.2: Ion spin-motion level structure. **a.** State ladder, showing the transitions driven by the carrier and first red/blue sideband Hamiltonians. In the Lamb-Dicke regime $\Omega_{n,n+1} \simeq \Omega_0\eta\sqrt{n+1}$. For beryllium $|e\rangle$ has lower energy than $|g\rangle$ (see Figure 2.4), so the $|e\rangle$ and $|g\rangle$ levels are swapped. **b.** Dependence of $\Omega_{n+s,n}$ on n and s in Equation 2.11, for calcium ($\eta = 0.05$) and beryllium ($\eta = 0.4$) transitions.

which contains three resonances. The first at $\delta = 0$ is called the carrier resonance, and those at $\delta = \omega_m$ and $\delta = -\omega_m$ are the first blue and first red sidebands respectively:

$$\hat{H}_{\text{car}} = \frac{\hbar}{2}\Omega_0\left(\hat{\sigma}_+ e^{i\phi} + \hat{\sigma}_- e^{-i\phi}\right) \tag{2.14}$$

$$\hat{H}_{\text{RSB}} = \frac{\hbar}{2}\eta\Omega_0\left(\hat{a}\hat{\sigma}_+ e^{i\phi} + \hat{a}^\dagger\hat{\sigma}_- e^{-i\phi}\right) \tag{2.15}$$

$$\hat{H}_{\text{BSB}} = \frac{\hbar}{2}\eta\Omega_0\left(\hat{a}^\dagger\hat{\sigma}_+ e^{i\phi} + \hat{a}\hat{\sigma}_- e^{-i\phi}\right) \tag{2.16}$$

Note that these are time-independent. A similar expansion of Equation 2.11 gives the effective resonant Rabi frequencies of the sideband drives in the Lamb-Dicke regime,

$$\Omega_{\text{car}} = \Omega_0, \quad \Omega_{\text{rsb},n,n-1} = \Omega_0\eta\sqrt{n}, \quad \Omega_{\text{bsb},n,n+1} = \Omega_0\eta\sqrt{n+1} \tag{2.17}$$

$\hat{H}_{\text{RSB}}$ is known as the Jaynes-Cummings Hamiltonian, and is also found in cavity quantum electrodynamics (QED) experiments, and $\hat{H}_{\text{BSB}}$ is known as the anti-Jaynes-Cummings Hamiltonian. The three Hamiltonians couple the state ladder shown Figure 2.2a, and are used to implement the bulk of the coherent state manipulation in QIP, including single- and multi-qubit gates, sideband cooling and motional state manipulation. Coherently driving $\hat{H}_{\text{car}}$, $\hat{H}_{\text{RSB}}$ and $\hat{H}_{\text{RSB}}$ together produces a variety of rich motional dynamics [115, 204, 95, 48], especially when combined with dissipative effects [80, 106, 78].

Equation 2.17 is sufficiently accurate for the majority of the work in this thesis. Because the motion-sensitive Raman beams interacting with the beryllium qubit have a Lamb-Dicke parameter of 0.4, however, accurately modelling carrier and sideband oscillations requires the use of Equation 2.11 when the ion is not ground-state cooled.

2.2.2 Rabi oscillations and single-qubit gates

A general superposition of the states in the state ladder shown in Figure 2.2a can be written as

$$|\Psi(t)\rangle = \sum_{n=0}^{\infty} c_{g,n}(t)\,|g,n\rangle + c_{e,n}(t)\,|e,n\rangle \tag{2.18}$$

and the full Hamiltonian in Equation 2.10 can be solved analytically to obtain the evolution operator

$$\begin{bmatrix} c_{e,n+s}(t) \\ c_{g,n}(t) \end{bmatrix} = \boldsymbol{T}_n^s(t) \begin{bmatrix} c_{e,n+s}(0) \\ c_{g,n}(0) \end{bmatrix} \tag{2.19}$$

$$\boldsymbol{T}_n^s(t) = \begin{Bmatrix} e^{-i\delta' t/2}\left[\cos\frac{f_n^s t}{2} + i\frac{\delta'}{f_n^s}\sin\frac{f_n^s t}{2}\right] & e^{-i\delta' t/2}\left[-i\frac{\Omega_{n+s,n}}{f_n^s}e^{i(\phi+|s|\frac{\pi}{2})}\right]\sin\frac{f_n^s t}{2} \\ e^{i\delta' t/2}\left[-i\frac{\Omega_{n+s,n}}{f_n^s}e^{-i(\phi+|s|\frac{\pi}{2})}\right]\sin\frac{f_n^s t}{2} & e^{i\delta' t/2}\left[\cos\frac{f_n^s t}{2} - i\frac{\delta'}{f_n^s}\sin\frac{f_n^s t}{2}\right] \end{Bmatrix} \tag{2.20}$$

where $\delta' = \delta - s\omega_m$ is the field detuning from the s'th sideband and $f_n^s = \sqrt{\delta'^2 + \Omega_{n+s,n}^2}$ is the 'effective' Rabi frequency at this detuning. Assuming that $c_{e,n+s}(0) = 0$, the state probabilities $p_{g,n}(t) = |c_{g,n}(t)|^2$ evolve according to

$$\begin{aligned} p_{g,n}(t) &= 1 - p_{e,n+s}(t) = 1 - \frac{\Omega_{n+s,n}^2}{(f_n^s)^2}\sin^2\frac{f_n^s t}{2} \\ &= \frac{1}{2} + \frac{1}{2}\frac{\Omega_{n+s,n}^2}{(f_n^s)^2}\cos f_n^s t. \end{aligned} \tag{2.21}$$

If the field is resonant with the carrier or a sideband then $\delta' = 0$, and $\boldsymbol{T}_n^s(t)$ becomes

$$\boldsymbol{T}_{n,\mathrm{res}}^s(t) = \begin{bmatrix} \cos\frac{\theta_{n+s,n}}{2} & -ie^{i(\phi+|s|\frac{\pi}{2})}\sin\frac{\theta_{n+s,n}}{2} \\ -ie^{-i(\phi+|s|\frac{\pi}{2})}\sin\frac{\theta_{n+s,n}}{2} & \cos\frac{\theta_{n+s,n}}{2} \end{bmatrix} \tag{2.22}$$

where $\theta_{n+s,n} = \Omega_{n+s,n}t$ and ϕ can be considered rotation angles for the rotation matrix $\boldsymbol{T}_{n,\mathrm{car}}$. This describes a generalised form of Rabi oscillation between pairs of states of the system, driven on the carrier or the sidebands. It holds beyond the Lamb-Dicke regime. By carrying out a controlled electromagnetic pulse with area $\theta_{n+s,n}$ and phase ϕ, we can coherently manipulate the internal and motional state amplitudes. When the carrier is driven on-resonance we can write

$$\boldsymbol{T}_{n,\mathrm{car}}(t) = \begin{pmatrix} \cos\frac{\theta}{2} & -ie^{i\phi}\sin\frac{\theta}{2} \\ -ie^{-i\phi}\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix} \tag{2.23}$$

where $\theta = \Omega_{n,n}t$. Carrier rotations act only on the internal states of the ion, and in the Lamb-Dicke regime $\theta \simeq \Omega_0$. A continuous drive of the calcium and beryllium qubits are shown in Figure 2.3a; when beryllium is driven with the co-90 beams the Rabi oscillations

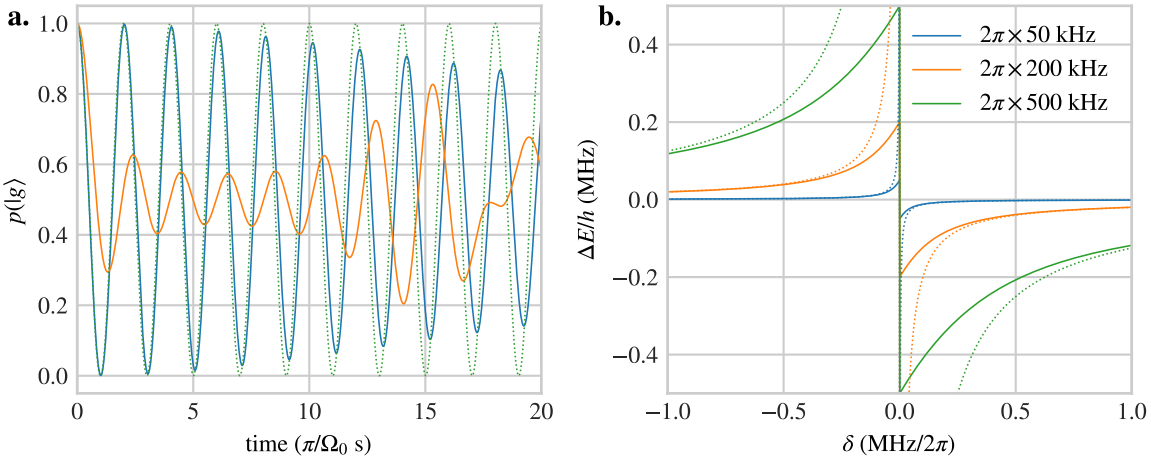

Figure 2.3: Rabi oscillations and ac Stark shifts. **a.** Resonant carrier Rabi oscillations for thermal states using Equation 2.22 and Equation 2.11, where the number state probabilities are $p(|g,n\rangle) = \bar{n}^n/(1+\bar{n})^{n+1}$ initially, at the Doppler limit on calcium (blue, $\eta = 0.05, \bar{n} = 6$) and beryllium (orange, $\eta = 0.4$ using the co-90 lasers, $\bar{n} = 3$), compared with a ground state-cooled calcium ion (dotted, $\eta = 0.05, \bar{n} = 0$). A beryllium ion in the ground state would have the same contrast, but a slightly lower oscillation frequency since from Equation 2.11 $\Omega_{0,0} = \Omega_0 e^{-\eta^2/2}$. Doppler cooling is discussed in §2.5.1. **b.** Plot of ac Stark shift frequency for a two-level system as a function of field detuning using Equation 2.32 (solid lines) and its polynomial approximation Equation 2.33 (dotted lines) for several resonant Rabi frequency values comparable to those used in calcium experiments.

decay due to the superposition of different Rabi frequencies $\Omega_{n,n}$. Sinusoidal oscillations can be attained by sideband-cooling the beryllium ion, discussed in §2.5.3.

Two particularly common rotations used in this work are the carrier $\pi/2$ and π pulses, where $\theta = \pi/2$ and $\theta = \pi$:

$$R_{\pi/2}(\phi) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -ie^{i\phi} \\ -ie^{-i\phi} & 1 \end{pmatrix} \tag{2.24}$$

$$R_{\pi}(\phi) = \begin{pmatrix} 0 & e^{i\phi} \\ e^{-i\phi} & 0 \end{pmatrix} \tag{2.25}$$

Such rotations are also be performed on the sidebands; for a fixed Rabi frequency Ω_0 the first sidebands of a ground state-cooled ion require pulses $1/\eta$ times as long as on the carrier. The qubit drives used for beryllium and calcium are discussed in §2.4.

2.2.3 AC Stark shifts

When the interaction-picture Hamiltonian $\hat{H}'_{\text{int}}(t)$ in Equation 2.10 is applied to the ion with nonzero detuning δ, it shifts the eigenenergies of the levels $|g\rangle$ and $|e\rangle$. To see this we neglect the ion motional energy in the bare Hamiltonian $\hat{H}_0$, treat the ion purely as two-level system with $\eta = 0$, and neglect the field phase ϕ. The total Hamiltonian in the rest frame is then the sum of $\hat{H}_{\text{qub}}$ and $\hat{H}_{\text{int}}(t)$,

$$\hat{H}_{\text{tot}} = \frac{\hbar}{2}\omega_0\hat{\sigma}_z + \frac{\hbar}{2}\Omega\left(\hat{\sigma}_+ + \hat{\sigma}_-\right)\left(e^{i\omega t} + e^{-i\omega t}\right) \tag{2.26}$$

Transforming into the rotating frame of the field using $\hat{U}(t) = e^{i\omega t/2}\left|e\right\rangle\left\langle e\right| + e^{-i\omega t/2}\left|g\right\rangle\left\langle g\right|$, we obtain

$$\hat{H}_{\text{ti}} = \frac{\hbar}{2}\left(\omega_0 - \omega\right)\hat{\sigma}_z + \frac{\hbar}{2}\Omega\left[\hat{\sigma}_+(1+e^{2i\omega t}) + \hat{\sigma}_-(1+e^{-2i\omega t})\right] \tag{2.27}$$

$$= \frac{\hbar}{2}\left(\delta\hat{\sigma}_z + \Omega\hat{\sigma}_x\right) \tag{2.28}$$

where we have discarded the rapidly-oscillating exponentials. The eigenvalues of the rotating-frame Hamiltonian $\hat{H}_{\text{ti}}$ are now

$$E_\pm = \pm\frac{\hbar}{2}\sqrt{\delta^2+\Omega^2}, \tag{2.29}$$

which are the energies of the 'dressed states' of the system in the presence of a field.

By comparison, if we transform the rest-frame 'bare' Hamiltonian $\hat{H}_{\text{ti}} = \hbar\delta\hat{\sigma}_z/2$ into the rotating frame using the same $\hat{U}(t)$, the result has eigenenergies

$$E_{\pm,\text{bare}} = \pm\hbar\delta/2; \tag{2.30}$$

its eigenvalues vary with δ because its eigenstates, which had constant energies in the laboratory frame, are now defined 'relative' to the field frequency and thus vary with the field detuning even though the field has no strength. When $\Omega \to 0$, Equation 2.29 gives $E_\pm = \pm\hbar\left|\delta\right|/2$, which have the same sign as Equation 2.30 for $\delta > 0$ but the opposite sign for $\delta < 0$; this is because the ground and excited dressed states have an avoided crossing in the 'qubit + laser' model whereas the bare states do not [31, Ch. 6]. Correcting the dressed eigenvalues in Equation 2.29 to be consistent with the bare eigenvalues in Equation 2.30 using a $\delta/\left|\delta\right|$ term, and choosing the convention that E_+ is the energy of the dressed excited state, the *differences* between the dressed and bare eigenvalues in the rotating frame are now

$$\Delta E_\pm = \pm\frac{\hbar}{2}\left(\delta - \frac{\delta}{\left|\delta\right|}\sqrt{\delta^2+\Omega^2}\right) \tag{2.31}$$

and the dressed minus the bare eigenvalue splitting is thus

$$\Delta E = \Delta E_+ - \Delta E_- = \hbar\left(\frac{\delta}{\left|\delta\right|}\sqrt{\delta^2+\Omega^2} - \delta\right). \tag{2.32}$$

This is shown in Figure 2.3b. Thus, as the laser detuning δ approaches 0 from below, the energy splitting between the dressed levels grows, increasingly blue-detuning the transition. This detuning is known as the dynamic or *ac Stark shift*. As δ crosses zero the detuning discontinuously changes sign from $+\Omega$ to $-\Omega$ in this model. A more sophisticated treatment includes a spontaneous decay rate γ from the upper level and avoids both the discontinuity and avoided crossing at $\delta = 0$, however if $\gamma \ll \Omega$, as is the case for most qubit transitions, the lineshape remains very similar.

For relatively weak quadrupole or Raman transitions in an ion, such as those typically used for qubits, the dominant ac Stark shifts usually come from far-detuned dipole transitions coupling strongly to the field. For this regime, expanding Equation 2.32 in Ω/δ to first order provides a simpler expression,

$$\Delta E_+ = -\Delta E_- = -\frac{\Omega^2}{4\delta} \quad \text{and} \quad \Delta E = -\frac{\Omega^2}{2\delta}, \quad \left|\delta\right| \gg \Omega \tag{2.33}$$

which is shown with dotted lines in Figure 2.3b.

Dressed states and ac Stark shifts are widely-used experimental tools in quantum optics and cavity QED, however they are also a source of infidelity in QIP, since they alter the qubit resonance frequency and cause a phase evolution of $\Delta Et/\hbar$ while the dressing field is on. They are discussed further in §6.1.2.

This concludes the introduction to coherent state manipulation, which is similar in trapped-ion QIP to other quantum platforms [139]. The techniques described next are specific to atoms and ions.

2.3 Internal state preparation and readout

The internal states of the ions are initialised with high fidelity using optical pumping [203]. Generally the various states of an ion in which undesired population may reside are optically coupled using lasers to short-lived excited states from which the population decays. Over multiple cycles of excitation and decay the population accumulates in those states which are decoupled from the optical fields. The specific techniques used for beryllium and calcium are discussed in §2.4. More details are presented in [81] and [105].

2.3.1 State readout

State readout, also called electron shelving, usually involves driving a dipole-allowed transition that only couples to one of the qubit levels[3]. For a qubit state $a\,|g\rangle + b\,|e\rangle$ this projectively collapses the qubit into the 'bright state' with probability $|a|^2$, and from here the ion fluoresces at a rate of 10^7 – 10^8 photons/s while the transition is driven, which can be observed on a camera or photo-multiplier tube (PMT). If the qubit collapses into the 'dark state' then the electron is 'shelved' and no fluorescence occurs. The contrast in fluorescence allows the qubit states to be distinguished[4], and infidelities around 10^{-4} are routinely achieved for several hundred microseconds of detection time [60]. Generally the detection time used is a tradeoff between collecting more photons and minimising relaxation processes in the ion. Many other ion properties such as motional state can be mapped to the readout outcome [95]. In quantum error correction, detection of a specific ancilla ion may be required without affecting the quantum state of other ions in the trap. This places stringent requirements on crosstalk; typically ~ 20 photons must be collected on a PMT for a detection fidelity below 10^{-4}, and with a detection efficiency of 2% this requires 1000 scattering events from an ion. To reduce the chance of collapsing the state of a neighbouring ion below 10^{-3}, the crosstalk from the scattered photons alone must be below $10^{-3}/1000 = 10^{-6}$, which is another motivation for the use of a separate species for ancilla qubits in QIP.

The calcium and beryllium ions are presented in detail next. Experimental state readout for single and multiple ions is discussed further in §4.3.

[3] Or an auxiliary level into which one of the qubit populations has been transferred.

[4] While mostly synonymous, in this thesis 'detection' refers to the process of measuring the ion fluorescence, and 'readout' to the process of reading its quantum state.

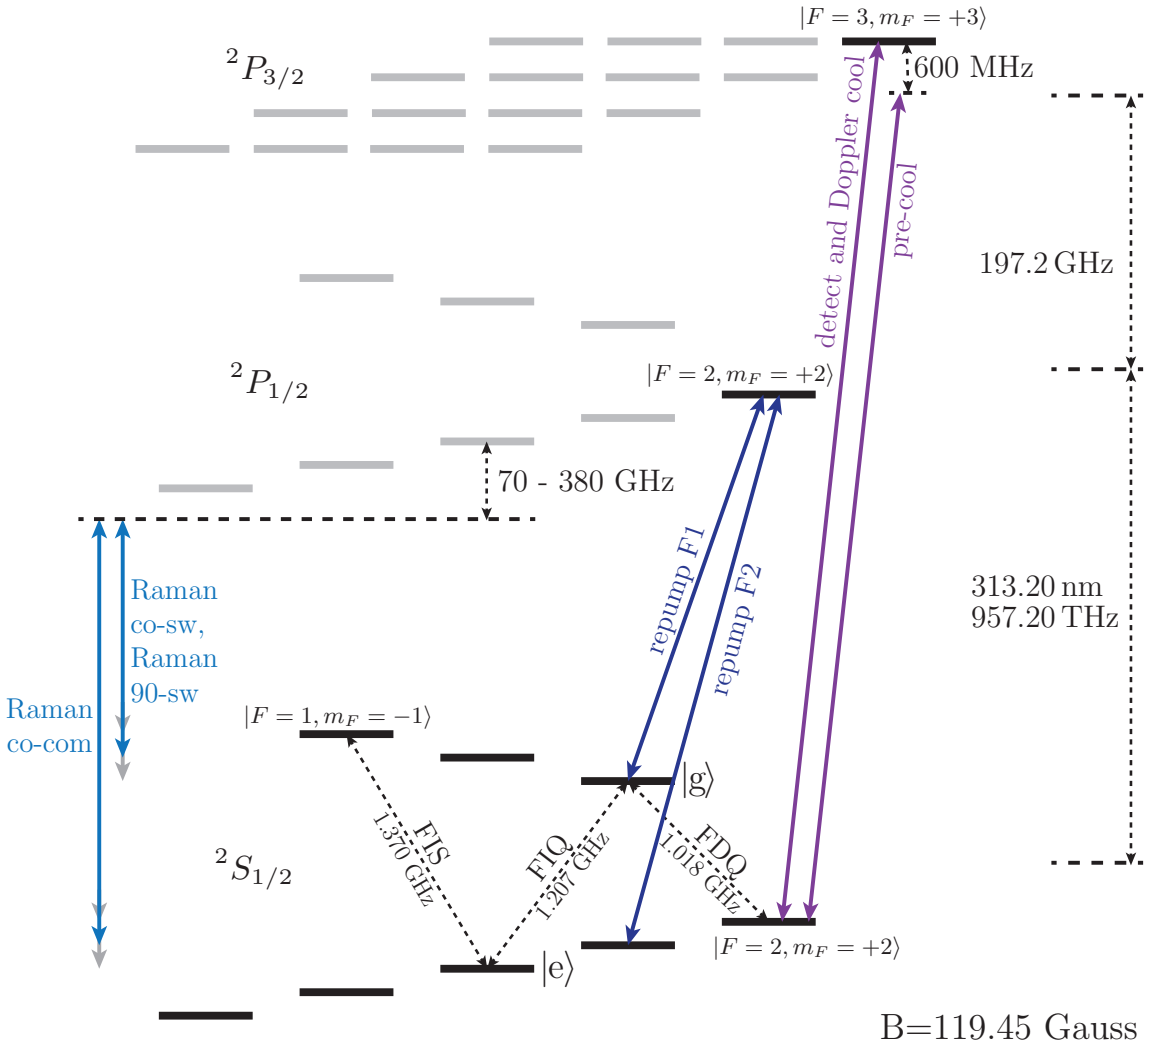

Figure 2.4: Beryllium-9 level structure. Solid arrows between levels represent lasers used in the setup, dashed arrows specify transition properties. The Raman lasers can be tuned to drive either the FDQ, FIQ or FIS transitions. Their common detuning from the $^2P_{1/2}$ manifold is adjustable, and is 230 GHz at the time of writing. For more detail see [86, 87, 105].

2.4 Beryllium and calcium

2.4.1 Beryllium-9

Beryllium-9 (beryllium or Be^+ in this thesis) has an atomic mass of 9.012 amu, and its level structure is shown in Figure 2.4. It is the lightest ion commonly used in QIP [146]. State readout of beryllium uses the cycling transition between $\left|S_{1/2}, F=2, m_F=+2\right\rangle$ and $\left|P_{3/2}, F'=3, m'_F=+3\right\rangle$, which has a natural linewidth of $2\pi\times19.4$ MHz; it is driven with resonant σ_+-polarised laser light. If the ion state is initially in $m_F=+2$, population cycles between this and the excited state and the ion fluoresces, emitting photons at 313 nm that are detected using a camera or PMT. If the ion is in another state in the $\left|S_{1/2}\right\rangle$ manifold, it will initially emit no photons. There are two main sources of infidelity in this readout scheme, which causes misidentification of the bright state as dark and vice-versa. For an

ion starting in the bright state, impure polarisation can excite the ion to a state other than $\left|P_{3/2}, F' = 3, m'_F = +3\right\rangle$, from which there is a chance of decaying to other ground states which are not resonantly coupled. Additionally the detection beam weakly repumps the off-resonant 'dark' $S_{1/2}$ states into the bright state for any polarisation of light over time, causing the photon count rate to rise over time. An intermediate detection time therefore produces optimal results [87, 105].

Near-resonant Doppler cooling is carried out with the same beam as used for detection, red-detuned by around 10 MHz from the resonance. A second beam red-detuned from resonance by 600 MHz is used for far-detuned Doppler cooling, and assists in ion loading and keeping the ions crystallised; see §B.1 and §5.6.4 for further details.

Beryllium has no low-lying D manifolds, reducing the possible decay pathways for the P states used in cooling and detection and simplifying the repumping compared to heavier ions such as calcium. All the transition wavelengths required for manipulating beryllium ions are close to 313.2 nm, which reduces the number of laser sources required. This is also the only wavelength that the beryllium optics need to operate at, unlike the multiple wavelengths needed for the calcium setup.

Repumping and state preparation are carried out using two beams, tuned close to the $\left|S_{1/2}, F = 2, m_F = +1\right\rangle \leftrightarrow \left|P_{1/2}, F = 2, m_F = +2\right\rangle$ and the $\left|S_{1/2}, F = 1, m_F = +1\right\rangle \leftrightarrow \left|P_{1/2}, F = 2, m_F = +2\right\rangle$ transitions. Both serve to optically pump the ion into the $\left|S_{1/2}, F = 2, m_F = +2\right\rangle$ readout state, with the first beam pumping population from the $S_{1/2}, F = 2$ manifold to higher-m_F states, and the second pumping population from the $S_{1/2}, F = 1$ to the $S_{1/2}, F = 2$ manifold. The exact detunings from resonance are optimised experimentally; see [105] for details.

Coherent state manipulation is carried out with Raman beams tuned to couple the $S_{1/2}, F = 1$ and $S_{1/2}, F = 2$ manifolds via the $P_{1/2}$ and $P_{3/2}$ manifolds. They are detuned 230 GHz red of this manifold; the choice of detuning is a tradeoff between Rabi frequency and Raman scattering rate. Three laser beams are used for Raman operations, labelled the *co-com*, the *co-switch* and *90-switch* herein. The co-com and 90-switch beams are arranged at a right-angle such that their difference wavevector $\vec{k}_{\mathrm{co-90}} = \vec{k}_{\mathrm{co-com}} - \vec{k}_{\mathrm{90-sw}}$ is along the axial direction of the trap, and is of magnitude $\left|\vec{k}_{\mathrm{co-90}}\right| = 2.8 \times 10^7$ /m. Collectively they are called the co-90 beams. The co-com and co-switch beams co-propagate and their frequencies differ by at most 1.5 GHz during experiments, hence their difference wavevector magnitude $\left|\vec{k}_{\mathrm{co-co}}\right|$ is only 30 /m. When this beam pair is used to manipulate the internal state of the ion, the effective Lamb-Dicke parameter $\eta \simeq 0$, which decouples the operations from motional imperfections such as intensity fluctuations due to micromotion or thermal occupation[5].

A major advantage of beryllium is that the transition energy between $\left|S_{1/2}, F = 2, m_F = 0\right\rangle$ and $\left|S_{1/2}, F = 1, m_F = +1\right\rangle$ becomes first-order insensitive to magnetic field fluctuations at an external field of 119.45 Gauss [86]. This transition is referred to as the *field-independent qubit* (FIQ), with the qubit states defined as $|g\rangle \equiv |F = 1, m_F = +1\rangle$ and $|e\rangle \equiv |F = 2, m_F = 0\rangle$ throughout this thesis[6]. The transition between the $\left|S_{1/2}, F = 2, m_F = +2\right\rangle$ readout state and $\left|S_{1/2}, F = 1, m_F = +1\right\rangle$ is labelled the *field-dependent qubit* (FDQ). At

[5] Stated differently, any motional Doppler shift is the same to 1 part in 10^6 on both beams, potentially altering their detuning from the $P_{1/2}$ states but not their difference frequency or phase.

[6] Note that $|g\rangle$ is higher in energy than $|e\rangle$; this defininition is chosen so that $|g\rangle$ is the bright state in detection.

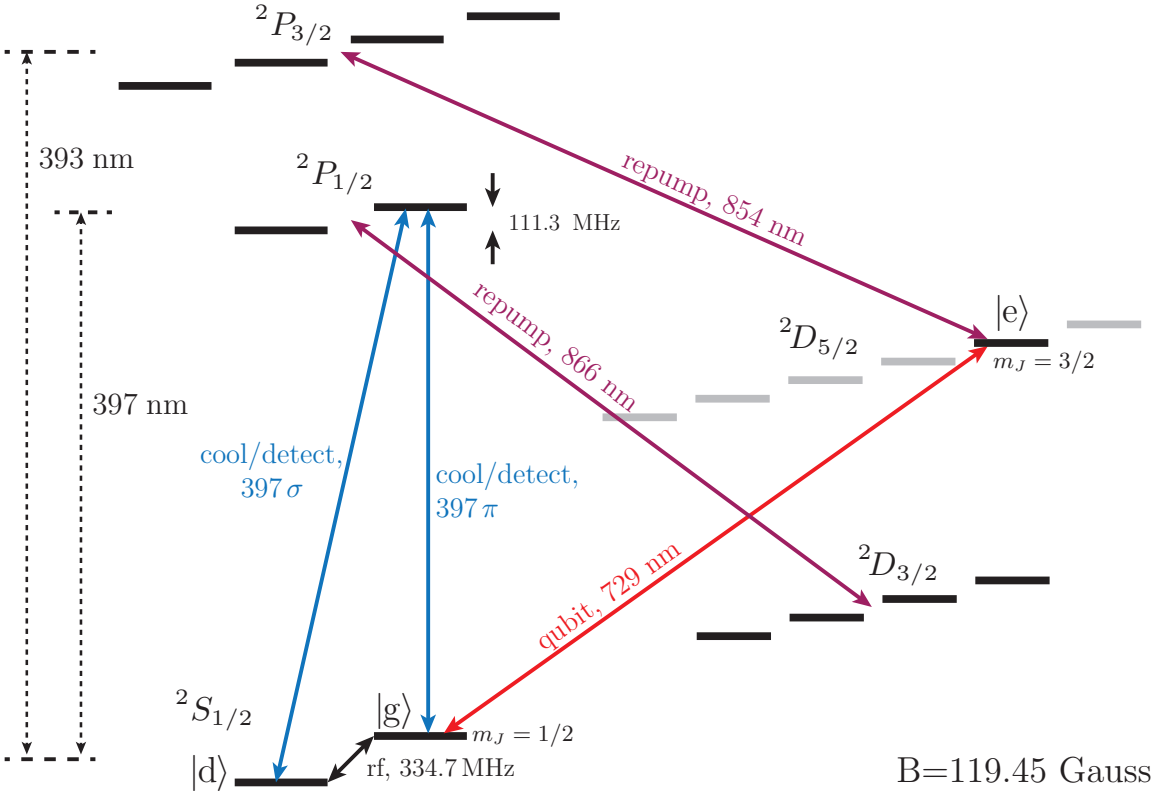

Figure 2.5: Calcium-40 level structure. Solid arrows between levels represent transitions driven by lasers or rf in the setup, dashed arrows specify transition properties.

the beginning of an experiment after cooling and repumping, a π pulse on the FDQ using the co-co beams is used to prepare the qubit in $|g\rangle$. Coherent operations are then carried out on the FIQ using the co-co and co-90 beams, and a final FDQ π pulse is used to transfer population from $|g\rangle$ to $|F = 2, m_F = +2\rangle$ where it can be read out. Note that from $|g\rangle$, the FIQ red sideband is motion-adding and the blue is motion-subtracting, opposite to calcium. The FIQ $|e\rangle$ state is weakly repumped by the detection beam, which contributes significantly to readout infidelity. To mitigate this, a co-co π pulse transfers population from $|e\rangle$ to $|F = 1, m_F = -1\rangle$, which is repumped significantly less. $|e\rangle \leftrightarrow |F = 1, m_F = -1\rangle$ is labelled the *field-independent qubit shelving* (FIS) transition (it is not field-independent itself).

2.4.2 Calcium

Calcium-40 (Ca^+ or calcium in this thesis) has an atomic mass of 39.963 amu and its level structure is shown in Figure 2.5. It is currently used by many ion-trapping groups. State readout uses the four transitions between the $|S_{1/2}, m_J = \pm 1/2\rangle$ and $|P_{1/2}, m_J = \pm 1/2\rangle$ levels, whose natural linewidths are around $2\pi \times 20$ MHz. They are driven by co-propagating π- and σ_+-polarised beams at 396.96 nm. For readout the π beam is tuned to the splitting between $|g\rangle \equiv |S_{1/2}, m_J = +1/2\rangle$ and $|P_{1/2}, m_J = +1/2\rangle$, and the σ_+ is tuned close to that of $|d\rangle \equiv |S_{1/2}, m_J = -1/2\rangle$ and $|P_{1/2}, m_J = +1/2\rangle$. The same beams are 5–10 MHz red-detuned for near-resonant Doppler cooling. Far-detuned Doppler cooling is carried out by red-detuning the π beam by around 40 MHz and operating it at maximum power, with the σ_+ beam off. Electromagnetically-induced transparency (EIT) cooling is carried out using

a third, counter-propagating 397 σ beam as the pump and the π beam as the probe; this is discussed further in §2.5.2.

The $S_{1/2} \leftrightarrow P_{1/2}$ transitions are not closed, and 866.45 nm light at several different frequencies is used to repump population that has decayed into the $D_{3/2}$ manifold. The light is vertically polarised, providing both σ_+ and σ_- components to repump every $D_{3/2}$ state.

The approximate branching ratio from the $P_{1/2}$ to the $S_{1/2}$ versus $D_{3/2}$ manifold is 20:1.

The calcium qubit used in this thesis is between the $|g\rangle \equiv |S_{1/2}, m_J = +1/2\rangle$ and the $|e\rangle \equiv |D_{5/2}, m_J = +3/2\rangle$ states. A narrowband laser coupling these states via an optical quadrupole transition at 729.35 nm is used for coherent manipulation. Unlike the Raman beams in beryllium, its wavevector has a component along the radial modes as well as the axial mode. The $D_{5/2}$ manifold is repumped by coupling to the $P_{3/2}$ states using an 854.44 nm laser with several frequencies separated by several hundred megahertz generated using an acousto-optic modulator (AOM). The ion then rapidly decays from the $P_{3/2}$ to the $S_{1/2}$ states via a 393 nm transition. This is close to the detection wavelength of 397 nm, and the resultant photons can be observed using the same detection system[7]. The transition between $|g\rangle$ and the $|d\rangle \equiv |S_{1/2}, m_j = -1/2\rangle$ level is also directly driven using an rf field in an experiment discussed in §6.6.2.

2.5 Cooling

High-fidelity single- and multi-qubit gates in an ion crystal generally require the motional modes to be cooled close to the quantum ground state, especially the modes used for multi-qubit gates. This requires several steps.

2.5.1 Doppler cooling

The first is Doppler cooling, where a short-lived cycling transition is driven with light red-detuned from the transition resonance. This is used in the regime where the transition lifetime is much shorter than the ion motional frequencies. Considering ion motion classically, an ion is likeliest to absorb photons when its velocity vector is antiparallel to the light wavevector, since in the ion frame the light is then blueshifted closer to resonance. Each photon provides a momentum kick along its wavevector direction when absorbed, slowing the ion down. An ion mode of oscillation whose vector overlaps with the light wavevector will thus be cooled. Unlike atomic Doppler cooling where multiple light beams are required, a single beam can cool the ion along all three motional modes if its wavevector overlaps with every mode vector.

The cooling rate from this process is counteracted by the diffusive heating that occurs due to fluctuations in the absorption time. The recoil from emitted photons causes a random walk in momentum space that also heats the ion, and together these processes impose a temperature limit [95]. It is reached for light below one saturation intensity that is red-detuned by around half the transition linewidth γ. The minimal achievable motional temperature, called the Doppler limit, is $k_B T \simeq \hbar\gamma/2$, at which the cooling and heating rates are in equilibrium [166], which is equivalent to $\bar{n} \simeq (\gamma/\omega_m + 1)/2$ for a motional mode frequency ω_m.

[7] This property has also been used to implement fluorescence at 397 nm in calcium using only the 729 nm, 866 nm and 854 nm lasers, without the use of UV lasers; see Refs [63, 102, 103] for information

For calcium and beryllium the linewidths of the Doppler cooling transitions are both around $2\pi\times$ 20 MHz, thus for a calcium ion in a $\omega_m = 2\pi \times 1.6\,\text{MHz}$ well the Doppler limit is $\bar{n} \simeq 6$. For a beryllium ion in the same well ($\omega_m = 2\pi \times 3.4\,\text{MHz}$) it is $\bar{n} \simeq 3$. See [81] and [105] for further experimental details of Doppler cooling in the mixed-species setup.

2.5.2 Electromagnetically-induced transparency cooling

The next step with a calcium or mixed-species crystal is electromagnetically-induced transparency (EIT) cooling, which operates on the calcium 397 nm $S_{1/2} \leftrightarrow P_{1/2}$ transitions, as do Doppler cooling and readout. It relies on EIT or coherent population trapping between the 397 nm π and σ beams to create a spectral feature much narrower than the 20 MHz $S_{1/2} \leftrightarrow P_{1/2}$ linewidth. It can used to simultaneously cool multiple motional modes to $\bar{n} < 1$ [130, 164, 92]. In our system the $\left|S_{1/2},\, m_J = -1/2\right\rangle \leftrightarrow \left|P_{1/2},\, m_J = +1/2\right\rangle$ transition is driven with strong 397 nm σ_+ light, blue-detuned from resonance by $\Delta\sigma \sim 10$ MHz, counter-propagating with respect to a less strongly-driven 397 nm π beam to create a non-negligible difference wavevector. This creates a sharp Fano-like profile, with an excited-state population given by [91]

$$\begin{aligned}
\rho_{ee} &= \frac{4\Omega_\sigma^2\Omega_\pi^2 P(\Delta_\sigma - \Delta_\pi)^2}{Q}, \qquad (2.34)\\
P &= \Gamma_\sigma + \Gamma_\pi,\\
Q &= (\Delta_\sigma - \Delta_\pi)^2\left[8\Omega_\sigma^2\Omega_\pi^2 P + 16\Omega_\sigma^2\Gamma_\pi\left(P^2 + \Delta_\pi^2\right) + 16\Omega_\pi^2\Gamma_\sigma\left(P^2 + \Delta_\sigma^2\right)\right]\\
&\quad +8\left(\Delta_\sigma - \Delta_\pi\right)\left(\Omega_\sigma^4\Delta_\pi\Gamma_\pi - \Omega_\pi^4\Delta_\sigma\Gamma_\sigma\right) + \left(\Omega_\sigma^2\Gamma_\pi + \Omega_\pi^2\Gamma_\sigma\right)\left(\Omega_\sigma^2 + \Omega_\pi^2\right)^2.
\end{aligned}$$

where $\Delta_\sigma(\Delta_\pi)$, $\Omega_\sigma(\Omega_\pi)$ and $\Gamma_\sigma(\Gamma_\pi)$ are the detunings from resonance, Rabi frequencies and linewidths of the $\sigma(\pi)$ transitions used in EIT (see Figure 2.5). The lineshape as a function of Δ_π for different parameters is shown in Figure 2.6a. For $\Omega_\pi \ll \Omega_\sigma$, there is a zero at $\Delta_\pi = \Delta_\sigma$ (the EIT condition) and a bright resonance at $\Delta_\pi = \Delta_\sigma + \delta$, where $\delta = \left(\sqrt{\Delta_\sigma^2 + \Omega_\sigma^2} - |\Delta_\sigma|\right)/2$ is an ac Stark shift as described earlier. If $\delta \simeq \omega_m$ and $\Delta_\pi = \Delta_\sigma$, no π light is absorbed unless the ion supplies $\hbar\omega_m$ of energy, shifting the bright resonance to $\Delta_\pi = \Delta_s$ and cooling the motion [164]. The spontaneous emission which completes the cycle heats the ion on the order of one recoil energy $E_R = \hbar^2 k_{397}^2/2m_{\text{Ca}} \simeq 1.3\times10^{-10}$ eV per event, whereas the energy of a single motional quantum $\hbar\omega_m = 8\times10^{-9}$ eV is 60× larger. Thus in the Lamb-Dicke regime EIT cooling contributes heating which is much less than one quantum per absorption and emission cycle, although the heating also affects the other motional modes.

EIT cooling is limited by absorption at the $\Delta_\pi \simeq \Delta_\sigma - \omega_m$ which heats the ion in the reverse process to that described above, and the achievable temperature scales with the ratio of these processes. The lineshapes for several different σ-beam detunings are shown in Figure 2.6b. A higher detuning produces a lineshape with a higher ratio between cooling and heating absorption, however requires higher σ power and reduces the EIT cooling rate due to being further off-resonance. In the mixed-species setup $\Delta_\sigma \sim \Gamma/2$ is used for broadband cooling, where $\Gamma = 2\pi \times 21.57\,\text{MHz}$.

Unlike previous EIT experiments using calcium-40 at lower magnetic fields [164, 92], a single impurely-polarised beam consisting of both σ_+ and σ_- components can be used. Because the σ_- impurity is detuned by 20 linewidths from the undesired $\left|S_{1/2},\, m_J = +1/2\right\rangle \leftrightarrow \left|P_{1/2},\, m_J = -1/2\right\rangle$ transition, it contributes negligibly to heating.

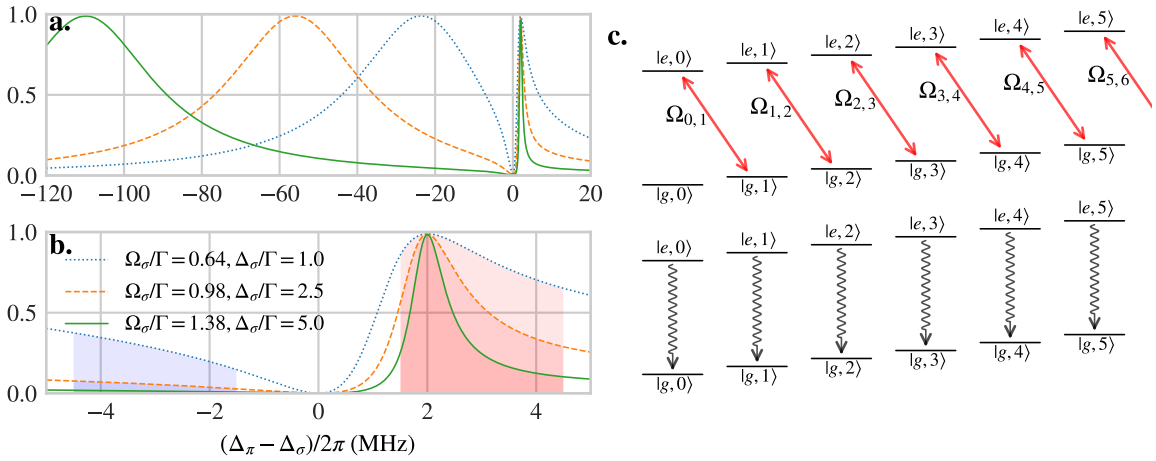

Figure 2.6: EIT and sideband cooling schemes. **a.** EIT cooling lineshapes for calcium $S_{1/2} \leftrightarrow P_{1/2}$ using Equation 2.34 as a function of Δ_π with $\Gamma = 2\pi \times 21.57\,\text{MHz}$, $\Omega_\pi = \Omega_\sigma/100$ and $\Gamma_\sigma + \Gamma_\pi = \Gamma$. Several detunings Δ_σ of the σ beam are shown, with Ω_σ chosen in each case to position the bright resonance maximum at 2 MHz from the EIT condition. **b.** Close-up showing the difference in absorption profile around $\Delta_\pi \sim \Delta_\sigma$ for the different detunings, with the cooling and heating sections of the lineshapes relevant to the mixed-species setup shaded in red and blue. **c.** One cycle of pulsed sideband cooling, consisting of a red sideband pulse followed by repumping, which together cyclically pump population into the $|g,0\rangle$ state.

EIT cooling is used for getting the motional modes of a calcium or mixed-species crystal below $\bar{n} = 1$, however because it is run in the low-detuning regime an extra step is required to ground-state cool the mode used for multi-qubit gates to the ground state, as well as the crystal modes that the calcium ion does not couple to. This requires a final stage of cooling.

2.5.3 Resolved-sideband cooling

Unlike Doppler and EIT cooling, resolved-sideband cooling uses the motional sidebands of the $|g\rangle \leftrightarrow |e\rangle$ qubit transition to carry out a pumping scheme whose dark state is $|g, n = 0\rangle$. This is possible only in the resolved-sideband regime where the natural linewidth $\Gamma_q \ll \omega_m$ for a motional frequency ω_m, which is true by design for qubit transitions. A cycle of the pumping loop consists of driving the $|g,n\rangle \leftrightarrow |e,n-1\rangle$ red sidebands, then effectively repumping the qubit to $|g\rangle$ via the carrier, as shown in Figure 2.6c. The major heating mechanisms are off-resonant carrier and blue sideband excitation [95], and limit the ground-state occupancy to $p_{n=0} \approx 1 - (\Gamma_q/2\omega_m)^2$, which is close to 1. Shaping the sideband pulse amplitude to reduce its spectral power at these frequencies can be used to suppress off-resonant effects. The same technique is used in the context of multi-qubit gates in §6.5.1 and its electronic implementation is in §3.3.6.

As seen earlier the Rabi frequencies for red sidebands vary with n, which lowers the pumping efficiency (for instance, a red sideband π pulse for $|g,1\rangle$ is almost a 2π pulse for $|g,4\rangle$ in calcium). This is partially counteracted by lengthening the sideband pulse time linearly as the cooling loop progresses, with the initial and final lengths optimised experimentally [81, 105]. The pulse power could also be altered for the same effect. For a single-species beryllium crystal the sideband Rabi frequency drops around $n \sim 22$ as shown in Figure 2.2b, and to achieve optimal cooling a loop of second sideband cooling is used initially to pump population from higher n states to below this point. Mixed-species crystals

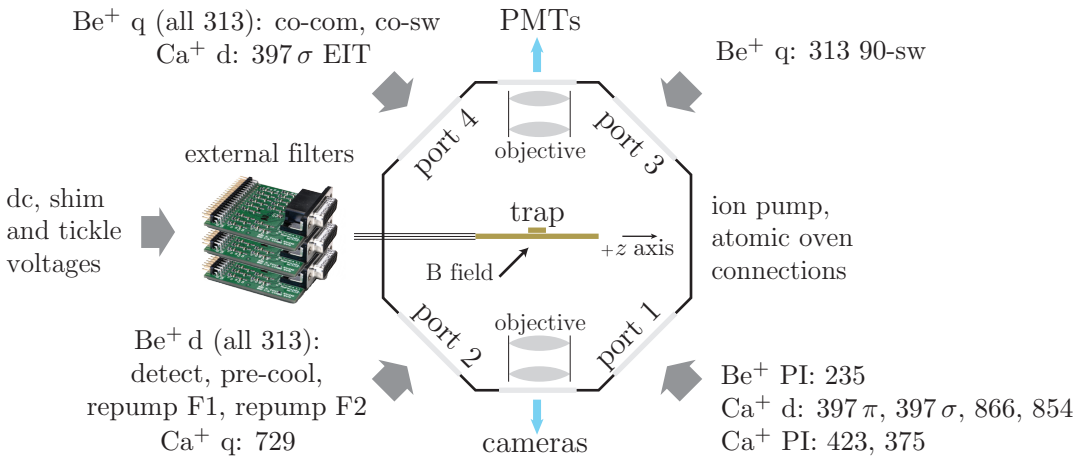

Figure 2.7: Top-down view of the vacuum chamber showing the uses of each viewport. Laser wavelengths are in nm. Beams listed with 'd' are used for dissipative operations (detection, cooling and state preparation), 'q' denotes qubit beams, and 'PI' denotes photo-ionisation (PI) beams used only for loading. The PI beams are directed into the left trap loading zone, the q beams into the central experimental zone, and the d beams are split and directed into both zones.

do not require this since the beryllium can be sympathetically EIT cooled.

The controlled use of dissipation in a scheme like sideband cooling has been applied in the mixed-species setup to prepare coherent and squeezed motional states, where the target state is a dark state of an engineered Hamiltonian created with a combination of carrier and sideband drives [80, 106, 78]. It is also related to the two-qubit technique presented later in §6.6.2.

This concludes the basic description of laser-ion interactions in this work. The experimental apparatus is briefly presented next.

2.6 Vacuum chamber, laser beams and imaging

2.6.1 Vacuum chamber

The vacuum chamber setup was designed and constructed by Daniel Kienzler, and is described in detail in [81]. The trap filterboard and wafer stack introduced in §2.1 are mounted vertically inside an octagonal vacuum chamber with six glass viewports, shown in Figure 2.7. Four of these are used for beam delivery and two for detection and imaging.

The dc feedthroughs on one face of the octagon are used to supply the trap electrode voltages via three 26-pin D-SUB connectors. The voltages are externally filtered through three filterboards shown in Figure 2.7 before the vacuum chamber, which implement a second-order Butterworth lowpass filter with a cutoff frequency of 73 kHz as shown in Figure 2.8 when combined with the in-vacuum RC filters. The filterboards and a backplane into which they plug in were constructed in this work to be modular, and have been replaced multiple times with different component values.

Not shown is the trap rf, which is supplied from the top of the chamber through a helical resonator tuned to 115.2 MHz. RF at ~ 330 MHz was also coupled in through this connection to drive the calcium $|d\rangle \leftrightarrow |g\rangle$ transition; the frequency is almost double the trap rf and the $|d\rangle \leftrightarrow |g\rangle$ splitting was shifted by altering the magnetic field to avoid the rf

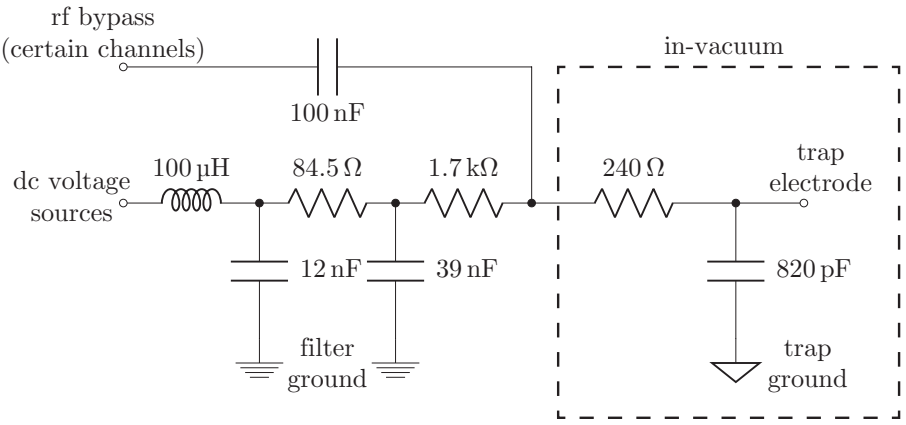

Figure 2.8: Lowpass filter used for dc trap electrodes, with a 73 kHz cutoff.

sources affecting one another. Also omitted are the in-vacuum atomic ovens, which contain metallic calcium and beryllium that can be resistively heated to sublime the metal and produce free atoms. Collimators are used to align the atomic beams with the trap loading zone. The beryllium oven is currently run at 1.3 V and 0.6 A (0.8 W), and calcium at 0.75 V and 4.2 A.

2.6.2 Laser beams

The chamber is surrounded by four optical breadboards to supply the beams at Ports 1, 2, 3 and 4. The beryllium detection, optical pumping and PI beams are raised to their breadboards using free-space periscopes, while all the other beams are brought to the breadboards via optical fibres. The Raman co-com and co-switch beams share a UV-capable photonic crystal fibre [32], and on the breadboard an adjustable piezo mirror[8] is used to finely tune their alignment on the ions. The co-90 beam uses a similar arrangement. Additionally the breadboards in front of vacuum ports 3 and 4 are entirely enclosed by covers reaching to the viewport optics, to minimise beam pointing fluctuations affecting the intensity at the ions. The details will be presented in the upcoming doctoral thesis of Matteo Marinelli [114]. The frequency and amplitude of every dissipative and qubit beam is independently controlled by one or more AOMs. The beryllium beam lines are discussed in [105, Ch. 5] and the calcium in [81, §3.4]. Recently an EOM has been added through which the beryllium detect and pre-cool beams pass[9], which is modulated at the rf frequency to cancel the beryllium frequency modulation caused by axial micromotion. A similar arrangement using two EOMs[10] is used for the Raman 90-switch; their calibration is discussed in §4.6.2.

2.6.3 Imaging

The trap is imaged on one side using two PMT detectors, one for calcium and one for beryllium, as well as two cameras on the other side, one for each ion species. Close to the trap are two in-vacuum objectives, each with a numerical aperture of 0.45, designed for imaging both 397 nm and 313 nm light. The light from these is directed to two imaging

[8] Thorlabs Polaris-K1S3P.
[9] Qubig E0-T110D3-UV.
[10] Qubig EO-T110P3-UV.

boxes, each containing a dichroic beamsplitter[11] with a cutoff wavelength of 347 nm; these reflect 313 nm light into the beryllium imaging devices, and transmit 397 nm light through to the calcium devices; there is another dichroic before these which is used to block the 729, 854 and 866 nm calcium wavelengths, leaving only the UV.

The objectives can be translated in 3D, so that different areas of the trap can be imaged simultaneously. Usually the PMT objective is focused near the central experimental zone of the trap, and the camera objective near the left loading zone in Figure 2.1e. The PMT counts are used for state readout during experiments and sometimes coarse beam alignment in the experimental zone[12], and the cameras are used for optimising the ion fluorescence in the loading zone of the trap. For beam alignment, the camera objective is shifted to image the edges or corners of dc electrodes, upon which laser beams can be aligned to obtain a reference position and calibrate the sensitivity of mirror-mount adjustment screws.

In the next chapter the experimental control system architecture is presented, and the experimental discussion continues in Chapter 4.

[11] Semrock BrightLine beamsplitter, FF347-Di01-25x36.

[12] UV light from the detection and photo-ionisation beams scattered from the ion trap can be detected and minimised using the PMTs.

3 Experimental control

In QIP and related fields such as quantum spectroscopy and metrology, the role of an experimental control system is to automatically and reproducibly carry out the tasks required for state initialisation, manipulation and readout of a quantum system[1]. In previous generations of experiments an overall control system was not essential, however the experiments performed in these fields today would already be unmanageable without a design that systematically addresses experimental requirements. These requirements are discussed in §3.1, and three common control architectures are compared in §3.2.

The 'Modular Advanced Control of Trapped IONs' (M-ACTION) system is discussed in §3.3. It was largely developed during this thesis to meet the requirements of foreseeable experiments in the mixed-species setup while avoiding the tradeoffs of previous systems. A more detailed description of the system libraries, recommended programming style and higher-level techniques is available in §A.1.

This chapter classifies the 'M-ACTION system' as the hardware and software used for real-time control of the experimental setup. Standalone devices necessary to controlling the system that are not synchronously altered during experiments are discussed separately in §3.4.

3.1 System requirements

The system requirements can be broken into the electronic capabilities of the system, the application programming interface (API) which a user builds on to write experiments, and the graphical user interface (GUI) with which the user can interact to run experiments.

3.1.1 Electronic capabilities

Trapped-ion QIP experiments require coordinated rf, dc-coupled analog, and digital outputs with reproducible timing. The rf outputs are most commonly used for controlling acousto-optic modulators (AOMs), through which the frequency, phase and amplitude of laser light can be modulated to produce a sequence of optical pulses on the ions. The rf hardware can also directly drive electromagnetic fields at the trap (see §6.6.2), or modulate the trap electrode voltages (see §4.5.1). Tens of channels capable of rapid frequency and amplitude control are required for complex setups, with the additional requirement of phase coherence and controllability on those channels that control qubit lasers, with a well-defined phase relationship between the channels. 19 modulated rf sources are used in the mixed-species setup, of which 8 must be phase-coherent.

Analog outputs are used primarily to generate the dc trap electrode voltages, however can also be useful for adjusting piezo voltages, analog servo control setpoints or other analog devices. Tens of channels are required for traps with many electrodes; 28 fast and 28 slow

[1] In this chapter, an 'experiment' refers to a particular sequence of these tasks, carried out to perform a measurement or protocol on the quantum system. The experimental apparatus as a whole is referred to as an 'experimental setup' in the literature and elsewhere in this thesis.

Type	Requirements	Timing
rf outputs	100 kHz – 500 MHz frequency (GHz for microwave drive), phase-coherent	< 20 ns resolution, < 1 ns jitter
analog outputs	tens of volts, > 10 MHz bandwidth	< 100 ns jitter
analog inputs	a few volts, > 10 MHz bandwidth	< 100 ns jitter
digital inputs/outputs	LVTTL (0–3.3V) voltage	< 20 ns resolution, < 1 ns jitter
feedback		< 50 µs latency

Table 3.1: Estimated requirements of the experiments in this thesis, an example of typical requirements in trapped-ion QIP.

channels are used in the mixed-species setup (see Chapter 5). Digital outputs are used to control two-state devices such as rf switches, mechanical shutters or stepper motors, to trigger partly-independent external devices such as oscilloscopes and arbitrary waveform generators (AWGs), and to monitor the timing of other outputs or inputs which their switching times coincide with. The mixed-species setup requires 12 outputs at the time of writing, including those used for monitoring other events.

There must also be digital inputs to collect photons and thereby determine the quantum state of the ions, as discussed in §2.3.1. These are single-ended, and are monitored to count digital pulses from photomultiplier tubes (PMTs) within a detection time window. More advanced inputs able to handle digital cameras are also possible [179]. Though less commmon, additional dc-coupled analog inputs are also useful for monitoring or feeding back on experimental parameters such as beam intensity or angle.

Another increasingly important requirement is feedback. The outputs must be linked to the inputs so that the experiment can be altered conditionally based on external events, which could range from measurements of the ions to external alerts such as laser problems. This requires the system to be somewhat 'aware' of the experiment, through computational power either embedded in the hardware or provided by a standalone PC. It also means that the delay (latency) between input and output must be short compared to decoherence processes such as magnetic field noise or motional heating/dephasing, which makes using a PC problematic due to the unpredictable delays in the operating system stack[2]. Outside of feedback, the availability of low-latency computation also enables the real-time monitoring of the experiment, for example in collecting data on the experiment for later analysis offline.

The approximate requirements of the mixed-species setup are summarised in Table 3.1. Outside the concrete specifications the electronics must be reliable, using tried-and-tested approaches, minimally complex, and modular so that the system can scale with experimental demands. These naturally lead to reduced costs.

[2] At least a modern PC running a standard desktop operating system; operating systems such as DOS and real-time Linux have lower and more deterministic latencies and have been used successfully for experimental control [13, 74, and David Lucas, private communication, 2013].

3.1.2 Software stack: API and GUI

Beyond the hardware, the sequence of outputs and measurements in a particular experiment needs to be configurable at a high level of abstraction. Users must be able to rapidly try new experimental sequences, debug the results and iterate their experiment. To achieve this, the system needs to be programmable in a high-level language such as Python, C/C++, MATLAB or NI LabVIEW, by providing a set of functions, objects or LabVIEW VIs that the user can build on. This set is called the application programming interface (API). API design is a broad topic, however one important principle is information hiding, which is the encapsulation of as many lower-level details as possible within higher-level, domain-specific tools, to minimise the knowledge required to program the system. The majority of users in most laboratories are physics students without detailed knowledge of the control system hardware, and a well-structured API helps their productivity.

Once an experiment is written, its high-level execution begins as a largely manual process, where the user sets most parameters by hand and selects the particular scans or plots they wish the system to carry out. Results are checked against intuition or theoretical models, and the experiment is iteratively debugged or improved. The interactive nature of this work, as well as the large number of parameters and settings in complex experiments, mean that a graphical user interface (GUI) is far more efficient than a text-based or programmatic interface for most users. Experimental control systems today virtually all provide a GUI that helps a user by displaying parameters in a logical way, offering options for plotting and fitting experimental results in real time, and alerting them of warnings or errors reported by the lower-level hardware. An important principle in both API and GUI design is that of 'least surprise', whereby they should be designed to match the expectations of a new user where possible to ease their learning curve. Also, altering the GUI to suit new experiments should not require much programming[3]. Finally, experimental parameters and results must be archived for later analysis. Usually software running on a PC, whether it is part of the GUI or standalone, is responsible for this.

3.2 Existing architectures

No single commercial system with all the capabilities listed in Table 3.1 is currently available, and most trapped-ion research groups construct a custom system from a combination of commercial units or based on an existing experimental control system. Three broad architectures, shown in Figure 3.1, are discussed in this section. The most straightforward architecture can be built from commercial devices including arbitrary waveform generators, stable voltage sources, digital counters/cameras and a digital pulse sequencer, all connected to a PC running a custom GUI and software that provides a common interface linking the system together. This architecture is referred to as *PC-centric*. An arbitrary waveform generator (AWG) can store in memory a waveform describing the whole rf sequence, and play it through a digital-analog converter (DAC) when triggered. Such setups are widely and successfully used in ion trapping and other related fields such as atomic physics. The two drawbacks for experimental setups such as ours are the cost for > 20 AWG channels[4] and the difficulty of carrying out fast feedback. The devices usually connect to the PC through protocols such as USB or Ethernet with best-case latencies in the tens of microseconds, and

[3] LabVIEW solves this for example by linking each GUI control to a corresponding object in its visual programming language, however GUIs written in text-based languages require other strategies.

[4] At least 50 000 chf.

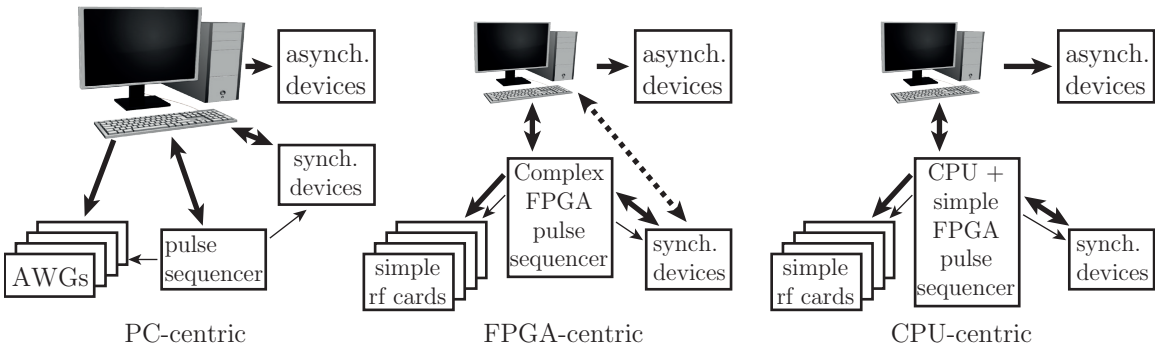

Figure 3.1: Existing trapped-ion experimental control system architectures. Thick arrows represent data communication links, thin arrows show digital trigger lines for synchronised timing. Optional links are shown by dashed lines. Clipart from https://openclipart.org/detail/215108/computer

often no guaranteed *worst-case* latency. Even with low-latency feedback, AWGs are seldom flexible enough for experiments like those performed in Chapter 7, though newer designs show promise [7, 104]. Commercial control systems offering programmable low-latency feedback are also advancing [118] and have been used successfully in NV-centre experiments [149, 34], however their overall capabilities do not yet appear to be sufficient for trapped-ion QIP.

For experimental setups requiring many output channels with greater flexibility at lower cost, such as those in this thesis, much of the electronics must be custom-built. Three notable systems were developed at MIT/University of Innsbruck [150, 175], the NIST quantum computing setup [87] and the NIST aluminium atomic clock setup [168, information in source code] in the 2000s, and remain in use today. Their architectures rely on a central digital pulse sequencer with precise timing, running on a field-programmable gate array (FPGA) with custom-developed gateware (see below). Phase-coherent rf is produced by simpler digital synthesiser PCBs (printed circuit boards) that are controlled by commands from the pulse sequencer, over an addressed parallel bus, which introduces limitations (discussed later) but is capable of driving enough channels to control the coherent qubit lasers in each experiment.

An FPGA consists of thousands of general-purpose logic blocks, memory, hardware multipliers and phase-locked loops, from which almost any digital circuit can be created[5], including a functional computer, by expressing it using a hardware description language (HDL). The major HDLs are VHDL and Verilog[6]; Verilog was used exclusively in this thesis. An HDL description is compiled into low-level FPGA *gateware*, similar to firmware but specifying the physical connections and gates in the FPGA, and loaded onto the FPGA upon power-up like firmware. FPGAs are particularly useful for running optimised circuits that carry out particular operations quickly and with fixed latency; they are widely used in communications, signal processing and high-performance computing. Commercial AWGs and similar devices also increasingly use FPGAs.

In the first two control systems listed above, the FPGA[7] implements a custom-designed

[5] Within the restrictions imposed by FPGA size, connection availability, power and timing.

[6] They are also used in integrated-circuit design; for FPGA work they are equivalent in their expressiveness. In my experience Verilog is less verbose and faster to write/debug but more ambiguous, which causes pitfalls requiring care to avoid.

[7] Altera Cyclone EP1C12Q240C6 on a custom PCB in the first system [150], Xilinx Virtex-4 on an 'XtremeDSP kit' board in the second [87].

processing core that executes hardware instructions from random-access memory (RAM). These describe a digital pulse sequence, with operations for register or memory access (such as reading or writing an address), flow control and feedback (jumping or forking) and pulse input/output (waiting some number of clock cycles, counting digital pulses from PMTs, setting digital outputs, etc). Users write their experimental program in a high-level language, which is compiled into these instructions by custom PC software and sent to the pulse sequencer via Ethernet in the first system and PCI-Express in the second. The FPGA runs through the instructions deterministically when triggered. A GUI on the PC is used to select experiments, load their pulse sequences onto the FPGA board, and process the results. In this architecture, referred to as *FPGA-centric*, the FPGA board runs the pulse sequence and communicates with the PC and peripherals, while the PC does the rest including generating the sequence and processing the experiment results.

In the third system, the FPGA chip[8] has a physical embedded PowerPC CPU (central processing unit), capable of running a compiled C++ program at high speed. This enables a different paradigm where experimental sequences are written directly in C++, and once they are finalised the compiled program is reloaded onto the pulse sequencer board quite infrequently. Instead of a RAM programmed by the PC and read nonlinearly by a custom processing core, the CPU autonomously runs an experiment by feeding pulse input/output instructions into a first-in first-out buffer (FIFO), and they are read and executed sequentially by a finite state machine (FSM). Since off-chip RAM access and experimental flow control are handled by the CPU the FPGA design is much simpler, and in principle pulse sequences of arbitrary complexity can be generated. Tasks such as analysing the results and feeding back on the sequence can be carried out easily in software compared to the FPGA-only design. The manufacturer provides software drivers for Ethernet and many other peripherals, largely shifting the development burden from FPGA firmware to embedded software. The main role of the desktop PC during experiments in this *CPU-centric* architecture is providing a GUI, with no recompilation or significant data transfer required during everyday operation.

While the FPGA-centric and CPU-centric architectures are well-suited for the previous generation of QIP experiments, they each have bottlenecks that limit their medium-term scalability. For experiments requiring fast feedback based on the outcome of an algorithm beyond simple thresholding (integer comparison), an FPGA-centric system requires hand-coding the functionality into the FPGA firmware. This is inflexible, error-prone and time-consuming, and hinders experimental progress; with a CPU-centric system the algorithm can be implemented directly in software and tested/debugged far more quickly and reliably. The high latency between the PC and FPGA memory requires loading the experimental instructions in one block, which can take seconds or more for long sequences, limiting the duty cycle of the setup. Since the CPU-centric system streams pulses out from the embedded CPU in real-time this problem is avoided[9]. While the CPU-centric design avoids these two issues, its major weakness is that the CPU needs to remain unburdened to keep the pulse sequencer supplied with instructions, otherwise timing errors could occur. This problem worsens as more 'simple' peripherals are added to the system, such as rf or dc sources programmed via addressed parallel bus, each requiring a stream of instructions and linearly increasing the demands on the CPU.

[8] Xilinx Virtex-5 on a Xilinx ML-507 board.

[9] Unlimited-length streaming is also used in AWGs such as [7], which are used in the control system of the Siegen ion-trapping group [152, 197, 153], however fast feedback remains difficult due to the latency.

Name	Meaning	Quantity	Typical execution time
Pulse	An optical or rf field applied to the ions for a set time, to carry out an operation on their internal or motional states.	5 – hundreds per shot	1.4 µs – 5 ms
Shot	Single execution of an experimental sequence of cooling, state preparation, manipulation and detection pulses.	10 – 1000 per point	1 – 50 ms
Point	Repeated execution of many shots, to acquire statistics for inferring the state populations.	2 – 1000 per scan	50 ms – 5 s
Scan	Collection of points as a parameter is varied, to plot the experiment results as a function of the parameter.	n/a	1 s – hours

Table 3.2: Definitions of experimental subroutines discussed in this chapter.'Pulse' is not to be confused with *digital pulse*, a rising and falling edge on a digital line.

The M-ACTION system was developed to address the issues discussed above. It is based on the CPU-centric system described above called Ionizer/Aluminizer, which is still used in several NIST atomic-clock setups. It has diverged, however, into a hybrid FPGA-centric and CPU-centric architecture, retaining the benefits of each. The design is presented in the next section, with a comparison to the FPGA- and CPU-centric architectures at the end.

3.3 M-ACTION system

3.3.1 Overview

Like the Ionizer/Aluminizer system, M-ACTION relies on a central embedded CPU to manage the experiment, which is controlled via a PC GUI. It uses a Xilinx Zynq-7000 chip[10] which consists of two ARM CPU cores running at 666 MHz, comparable in speed and features to a smartphone CPU, connected to a medium-sized FPGA. The Zynq also has built-in peripherals that handle Ethernet, USB, and various forms of serial communication. It is mounted on the 'master board' in the M-ACTION system, the Zedboard[11]. Several experimental 'subroutines' referred to in this section are defined in Table 3.2.

M-ACTION reduces demand on the CPU by using standalone direct-digital synthesiser (DDS) boards for rf generation, each with four rf channels and a large FPGA running relatively independent firmware, which does not need a continuous flow of commands from the master as in the CPU- and FPGA-centric designs. Before taking an experiment point, the master sends an rf pulse sequence to each board via a low-latency direct connection. Thereafter the firmware autonomously carries out the sequence for every shot with zero or very low demand on the master CPU. Deterministic single– and multi–channel rf phase coherence is handled on the boards, as well as pulse amplitude shaping. Low-latency feedback is possible due to the firmware architecture and the direct connection to the Zynq.

The system uses custom-designed AWG boards to provide the trap electrode voltages,

[10] Xilinx XC7Z020-CLG484-1.
[11] A commercial Zynq evaluation board, produced by Avnet.

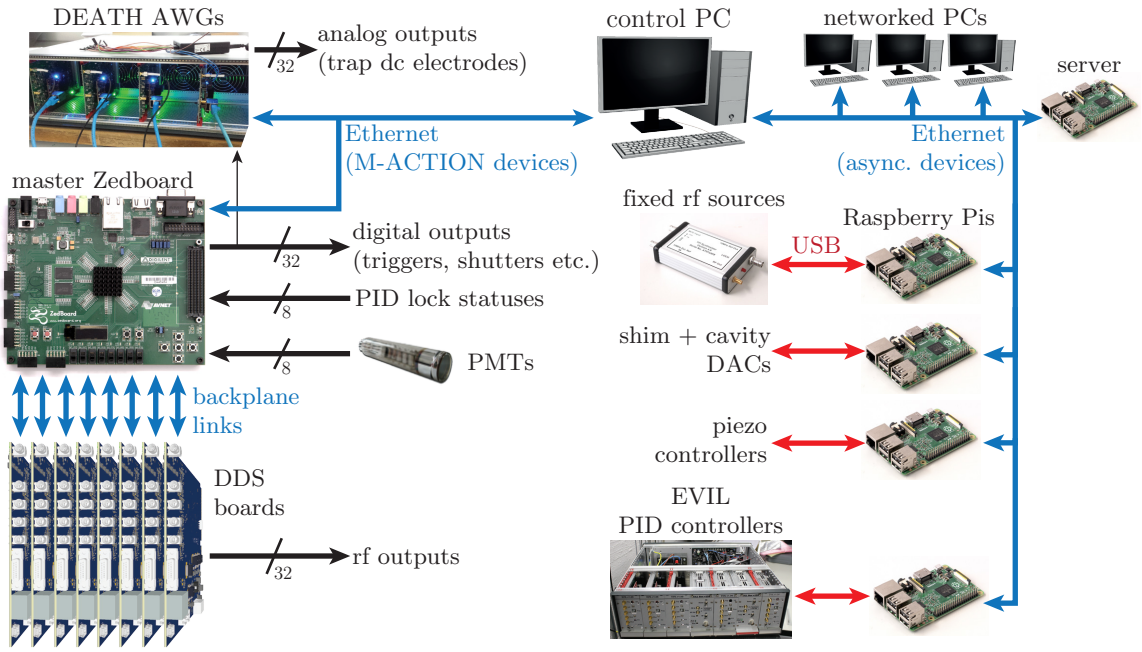

Figure 3.2: Experimental control in the mixed-species setup. The M-ACTION system is on a dedicated network, while the asynchronous devices and other PCs use a separate network that is managed using a Raspberry Pi as the gateway server. Currently only 4 DDS cards (16 rf channels) are used. The central control PC is connected to both. The asynchronous devices are discussed in §3.4.

the Direct Ethernet-Adjustable Transport Hardware (DEATHs). These use smaller versions of the Zedboard, the Microzed, and also run largely independently of the master. The *waveforms*, or time-dependent sequences of voltages, for a set of experiments is loaded by the PC GUI from a waveform file; separate waveforms might carry out transport to or from the loading, separation and experimental zones in the trap, run a separation routine, or alter the trapping confinement without changing the trap location. The master communicates with the DEATHs via Ethernet to select a sequence of waveforms and specify their order, directions and speeds. A shared digital trigger is used to start all the channels in synch with the master. The DEATHs and their usage are discussed in Chapter 5; see the theses of Ludwig de Clercq [37] and David Nadlinger [132] for more details on their design[12].

The M-ACTION system is controlled in two complementary ways: through the software development environment in which experimental code is programmed, run and debugged, the Xilinx software development kit (SDK), and through a C++ GUI called Ionizer2 after the original, which is the front-end to the experiments on the master board and other capabilities. During everyday calibration and data-taking the SDK is rarely used. The M-ACTION and laboratory network topology is shown in Figure 3.2.

3.3.2 Zedboard hardware

The Zedboard PCB holds ports for Ethernet, used for most communication with the PC and DEATHs, USB, used for serial communication and diagnostics as well as debugging

[12] Much of the original software described in [37] has been rewritten; [132] is the more up-to-date reference.

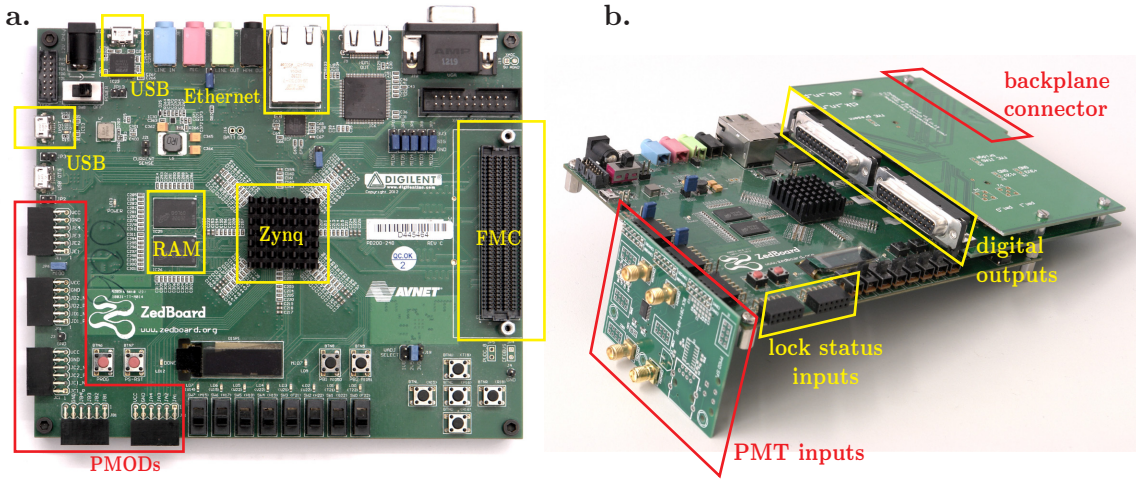

Figure 3.3: Zedboard hardware. **a.** Zedboard PCB and its peripherals used by M-ACTION. **b.** Backplane breakout board and front PMOD adapter card attached.

and programming[13], a FPGA mezzanine card (FMC) connector for high-speed I/O and five 12-pin digital PMOD headers for lower-speed I/O. There is also a pin header connecting to the Zynq XADC peripheral, a built-in 1 mega-sample per second (MSPS[14]) analog-digital converter (ADC) which has been used for real-time laser monitoring; see the theses of Pascal Engeler [44] and Martin Stadler [186] for information. The PCB is shown in Figure 3.3a.

A breakout PCB was designed for mounting on the FMC connector, as shown in Figure 3.3b. This provides 32 single-ended digital outputs on two DB-25 connectors, used for various purposes in the experimental setup. The PCB slots into a backplane, discussed in §3.3.3, and its backplane I/O consists of 16 independent low-voltage differential signaling (LVDS) pairs, two per DDS board; 8 common LVDS pairs driven by a 1:8 clock fanout buffer[15] that provide a global trigger pulse to every DDS board, and a final LVDS pair that provides a global clock. The breakout PCB has several extra I/O lines for future use, such as clocking the Zedboard externally to achieve phase-coherent operation between multiple M-ACTION racks.

The PMOD headers are versatile digital I/O used mainly as single-ended inputs. An external input trigger, up to 8 PMT input lines, and up to 8 experimental status input lines are available (see §3.3.4). An adapter PCB designed by Ben Keitch is used to connect cables to the headers and show several diagnostic LEDs. The PMODs were formerly used to control external DACs over a serial peripheral interface (SPI) link [103]; at the time of writing they are no longer used in any of the TIQI group experimental setups. The PMODs have also been used to interface a quantum random number generator via a two-wire serial protocol to the Zynq, which was used for quantum contextuality experiments [99]; see the upcoming doctoral thesis of Chi Zhang [207] for more details.

[13] Via the JTAG protocol.

[14] I have used MSPS rather than MHz to avoid confusion, since the baseband bandwidth of a 1 MSPS ADC is 0.5 MHz.

[15] Texas Instruments CDCUN1208LP clock buffer.

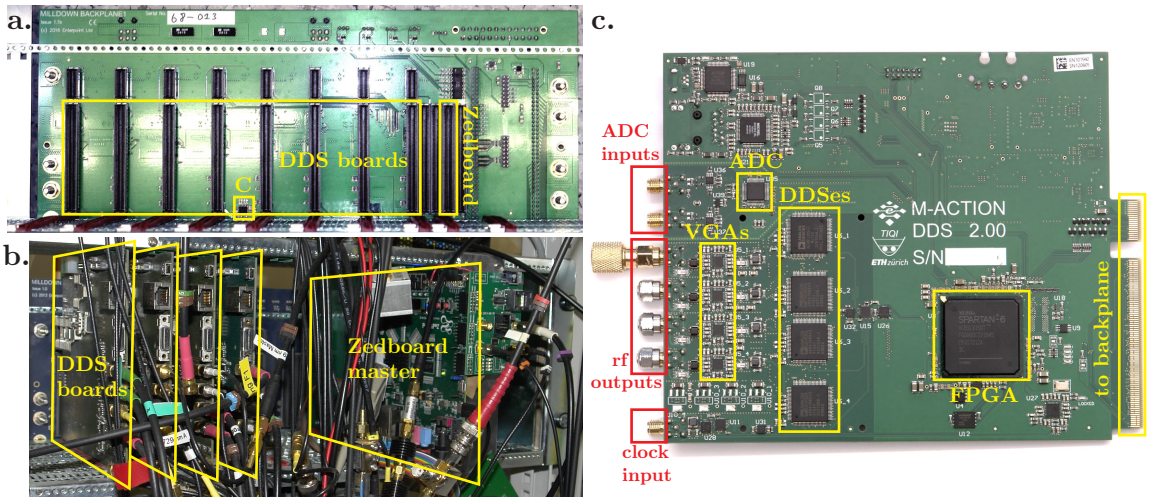

Figure 3.4: M-ACTION rack and DDS board. **a.** Bare backplane, showing backplane interconnections. The global clock is represented by 'C'. **b.** Active rack with a master board and four DDS boards in a Milldown backplane. **c.** Latest revision of DDS board showing relevant hardware highlighted.

3.3.3 Milldown backplane

The master board and DDS boards slot into the Milldown backplane, which was designed[16] to be partially compatible with the MicroTCA form factor, a rack standard designed originally for compact telecommunications [151]. The backplane has a group of 4 neighbouring connectors for an Advanced Mezzanine Card (AMC)-sized master board (74 mm wide × 180 mm deep), and 8 slots for slave boards as shown in Figure 3.4b.

The master board plugs into only the second master slot, which has three dedicated differential lanes to each DDS card and one common lane that connects to a global clock buffer on the backplane[17]. The three dedicated lanes are used for master → slave data transfer, slave → master data transfer, and a trigger common to all 8 slave slots. These lanes are length-matched between all the slave slots to achieve the same propagation delay. The backplane buffer receives a 133 MHz clock from the master board, and distributes it to the DDS boards. Power reaches the backplane through an ATX connection similar to desktop PC motherboards; we use a high-end consumer PC power supply unit which offers relatively stable voltage[18]. The backplane provides 12 V and 3.3 V to each DDS board; the 3.3 V is not part of the MicroTCA standard and is provided via a small additional slot. Several brushless dc fans for cooling the DDS boards are also driven from the power supply.

3.3.4 Master board gateware

The central Zynq chip is separated into a processing system (PS), consisting of two ARM CPU cores, of which only one is currently used, and physical I/O peripherals, and FPGA-style programmable logic (PL), which runs the custom M-ACTION gateware. The PL holds

[16] Design and manufacture were carried out by Enterpoint, a British electronics company, who also sells the backplane and DDS card [45].

[17] Texas Instruments CDCUN1208LP, the same as on the Zedboard breakout PCB.

[18] Corsair AX860i. Lower-noise linear supplies have been tried, however the critical rf circuitry on the DDS boards has several stages of linear regulation and filtering, and we do not see a significant improvement; see [75] for detailed phase noise measurements of the DDS boards.

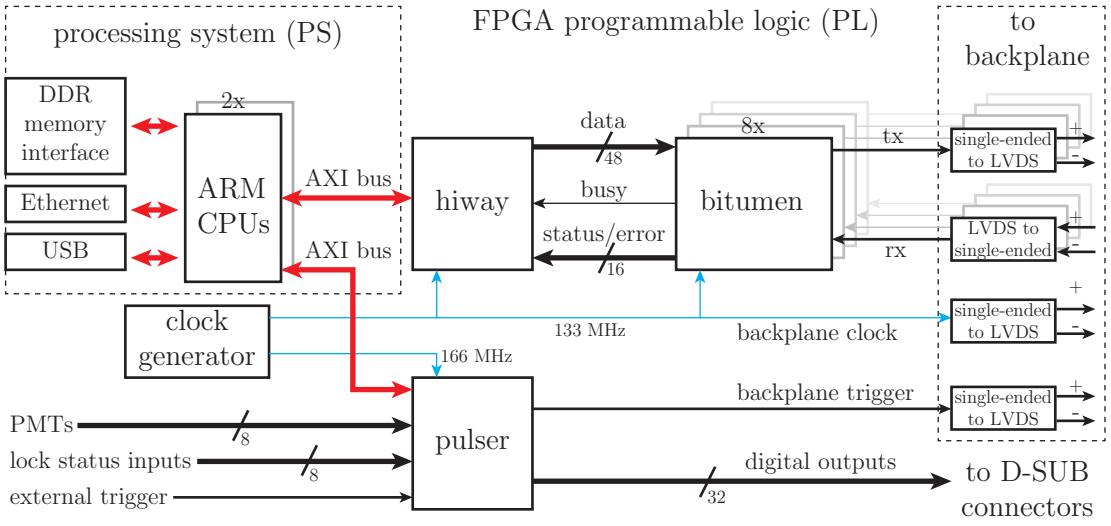

Figure 3.5: Master board gateware architecture. Red lines are bidirectional buses, blue lines are clocks, the remainder are data. `hiway` handles backplane I/O while `pulser` handles the remainder. The `bitumen` cores serialise and transmit 48-bit backplane DDS board instructions, receiving and deserialising the replies.

two major logic cores. `hiway` is responsible for communication with the DDS boards over the backplane, buffering single-channel data and broadcasts, and monitoring status information and errors in the DDS board replies; it is clocked at 133 MHz. `pulser` handles all remaining I/O, including the 32 digital outputs and 8 PMT inputs, trigger input and DDS board trigger output; it is clocked at 166 MHz. The cores communicate with the PS via separate AXI[19] buses; the PS is the bus master and can send or receive data via memory-mapped 32-bit data transfers.

Each transfer takes 30–40 ns, a theoretical throughput of $\approx 800 - 1000$ MBPS (megabits per second) with most of the latency coming from the ARM CPU and PS logic. 'Burst' bus transfers could provide higher data throughput, however this has not yet been required as the buses are far from saturation even in the most demanding experiments.

Pulser

The `pulser` core internally has a 256-slot 'write FIFO' (`wfifo`) to buffer a 32-bit digital pulse sequence, and 8 gated digital counters that collect PMT counts when activated and push their results into corresponding 8-slot 'counter FIFOs' (`cfifo`). The CPU can write to or read from particular addresses in `pulser`, and the act of reading or writing causes the core to carry out specific operations. The main elements are shown in Figure 3.6. Two 32-bit writes to Address 0 then 1 transfers a 64-bit instruction from the CPU and pushes it into the `wfifo`, filling one slot. The instruction contains a 32-bit digital output pattern, how many clock cycles to hold it before moving to the next instruction, whether to count pulses from the PMTs channels during the instruction duration, and whether to allow an external trigger to end the current instruction and/or output a trigger to the backplane after the instruction is finished. The last two features make it possible to trigger the M-ACTION system from an external line. Instructions from the `wfifo` are read by a FSM that outputs

[19] Advanced eXtensible Interface, a bus standard used by ARM.

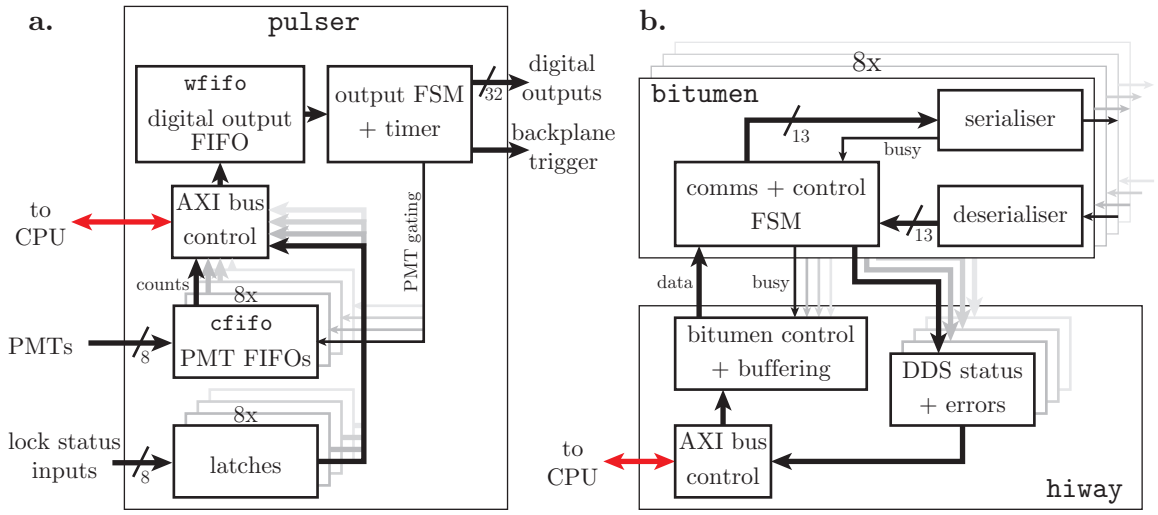

Figure 3.6: Key elements in the `pulser`, `hiway` and `bitumen` cores.

the digital pattern from `pulser` and holds it for the number of cycles requested, and carries out the other optional tasks, before reading the next instruction. See §A.2 for more detail.

Reading from Addresses 0 to 7 causes the corresponding counter FIFO to return a PMT result if one is available, emptying a FIFO slot. If not, it returns an error code to the CPU. Several other features are available in `pulser`. Reading Address 10 provides information on how many slots are available in the `wfifo`, and whether it has been empty at any point since the previous read; this would indicate that the core has paused for an unknown time and is shown as a timing error in the GUI. Reading Address 11 shows which of the counter FIFOs are full or empty.

Finally, a set of latches are used to monitor 8 digital inputs, and if any input goes high at any point then the latch remains set high. Reading the latch status on Address 14 returns the latch statuses and resets them. Although this has not been used in experiments so far, it is designed to check whether lasers remained stably locked during data taking, since the remotely-controlled lock boxes (EVILs) output a digital lock status signal (see §3.4.2). If a laser unlocked during an experiment shot, however briefly, the shot could be automatically retaken or the user could be alerted.

Hiway and Bitumen

The `hiway` core is an interface between the CPU and the `bitumen` cores, which handle low-level communication with the DDS boards. The communication uses a custom protocol rather than common ones such as Ethernet or USB to minimise latency and keep the gateware simple. `hiway` receives 64-bit instructions from the CPU; the lower 48 bits make up the data payload to be transferred to the DDS boards, and the lower 8 of the upper 16 bits specify which DDS boards to send the payload to. The upper 8 bits are unused. Before a valid instruction can be received, `hiway` checks whether the relevant `bitumen` cores are busy with an existing transfer; those cores that are idle are sent the data immediately. If any are busy `hiway` waits for them to become idle, and blocks the AXI bus if the CPU sends too

much data until it can be transmitted[20].

Each `bitumen` core serialises its 48-bit payload into four 12-bit words along with 'special' bits used for status, forming a total of 13 bits, and transmits them serially in four groups to a DDS board. The protocol is similar to Universal Asynchronous Receiver-Transmitter (UART). The output level is high when idle, and set low for one clock cycle when a transmission begins. Once the transmission ends the core receives a reply from the DDS board, which is running a mirrored `bitumen` core as a slave. This reply is sent back to `hiway`, letting it know whether the DDS board received the data correctly and its internal status. The CPU can read specific addresses in `hiway` to obtain the status and errors for each DDS board during operation. See §A.3 for more detail on the communications protocol.

3.3.5 DDS board hardware

The DDS boards are designed to meet the rf requirements discussed in §3.1.1. They consist of four Analog Devices AD9910 DDS chips, connected to a large Xilinx FPGA[21] as shown in Figure 3.4a, which is linked to the master through the backplane. On-board voltage regulation and filtering has been optimised to reduce phase noise, digital spurs and inter-channel crosstalk below levels that would affect our QIP experiments. Measurements show that the phase noise is currently limited by the clock source used, with a single-sideband power spectral density of below -80 dBc/Hz from 0.01–1 kHz and below 100 dBc/Hz above this.

Other than the backplane there is USB, Ethernet and general-purpose digital I/O for flexible communications, with a version of the gateware available that uses only USB for standalone operation. The boards were manufactured by Enterpoint [45] based on specifications provided by the TIQI group; the rf chain was mostly based on an earlier design developed at the University of Innsbruck [175]. We have published the board details and performance [75].

The heart of an AD9910 DDS is a 32-bit binary counter called the *phase accumulator*, which is incremented once per clock cycle (a 1 GHz clock is used) by a 32-bit frequency control word F. The accumulator wraps around past 0 every $2^{32}/F$ clock cycles; for example if $F = 2^{29}$ the counter will wrap at a 125 MHz frequency. The frequency resolution is thus $1\,\text{GHz}/(2^{32}) = 233\,\text{mHz}$. A 16-bit phase offset control word is added to the most significant bits of the phase accumulator [39], providing a phase resolution of $2\pi/2^{16} = 96\,\mu\text{rad}$. The result is sent into an angle-to-amplitude converter, equivalent to a look-up table (LUT) with a sinusoidal pattern. Its output is multiplied by an amplitude control word. Finally the scaled digital sinusoid is converted to an rf current by a 14-bit DAC, which is output from the AD9910 chip differentially [6]. The block diagram is shown in Figure 3.7.

By modifying the frequency, phase and amplitude words at precise times, the AD9910 can be used to generate arbitrary rf patterns. The words are stored in single-tone profile registers, which can be written to via an SPI link; it currently takes 1.296 µs to rewrite a profile register from the FPGA. A pulse on a dedicated IOUPDATE input switches the DDS core to the new parameters. We configure the DDS to reset the phase accumulator to 0 upon switching. The words can also be rapidly modulated using a 16-bit parallel bus; this

[20] This also pauses the CPU; this is not detrimental to the system performance under the transmission loads we have tested.

[21] Xilinx Spartan-6 XC6SLX150T.

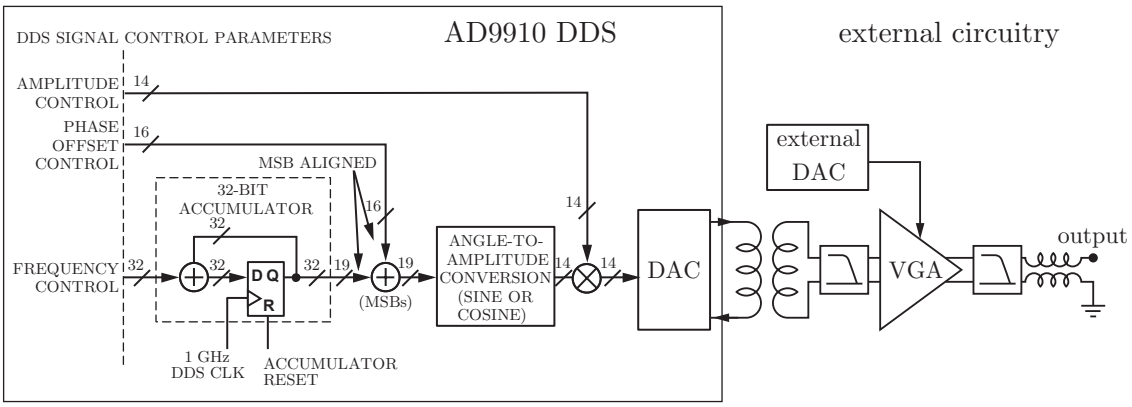

Figure 3.7: DDS internal logic and rf chain. Block diagram of AD9910 DDS core, adapted from [6], and the output rf chain. The DDS current output passes through a transformer, variable-gain amplifier and two 400 MHz lowpass filters.

is used for automated phase coherence as discussed in §3.3.6. Additional AD9910 options such as RAM modulation or fast profile switching are not currently used in experiments.

The DDS output goes to a variable-gain amplifier (VGA)[22], which is used for pulse amplitude shaping. Its gain is set by by one channel of a two-channel 14-bit DAC[23]. The VGA response is exponential in gain voltage, which requires the digital DAC input to be the logarithm of the desired signal. Pulse shaping is discussed in the next section. The rf chain including DDS, VGA and filtering has a passband of 10–400 MHz.

There are several requirements for a deterministic phase relationship between DDS chips. The 1 GHz clock must arrive in phase at each chip, which is carried out using a clock buffer[24] and length-matched PCB tracks. The DDSes have 16-state FSMs that control their internal clock phases and rf generation; these are random upon power-on and must be synchronised. The AD9910 has a sync circuit for this purpose, which resets the FSM state whenever a rising edge is detected on a differential sync input. It can also output a differential 1 GHz/16=62.5 MHz sync signal for synchronising other AD9910s to itself. On the DDS boards, DDS 1 provides the sync signal and it is sent to all four chips (including itself) via another clock buffer and length-matched tracks. Within the AD9910 both the sync input and output delays are finely tuned to ensure the sync is robust and does not fail due to temperature drifts. Finally the `ioupdate` pulses and parallel data from the FPGA must both be in synch with the DDSes. A 250 MHz sync clock from DDS 1 (not to be confused with the 62.5 MHz *sync signal*) is used for this purpose; the next section discusses how.

In the latest DDS board revision, a dual-channel 62.5 MSPS ADC[25] has been added for carrying out low-latency feedback on pulse amplitude or potentially other parameters. The two channels use pseudo-differential inputs and input buffers to reduce the risk of ground loops. The system is still under testing; see [186] for more information.

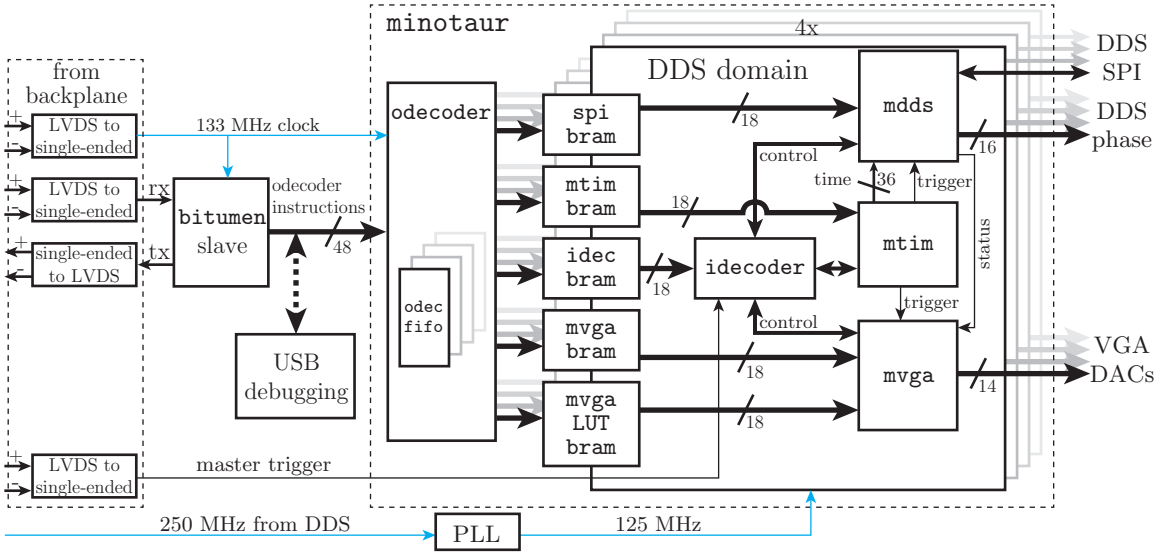

Figure 3.8: DDS board gateware. Data from the master is received and decoded by `odecoder`, transferred into BRAMs and executed by `idecoder`, which controls `mdds`, `mtimer` and `mvga`.

3.3.6 DDS board gateware

The FPGA gateware on the DDS board is shown in Figure 3.8. It consists of two clock domains, the 'backplane domain' driven by the backplane clock from the master board, and the 'DDS domain' driven by the sync clock from DDS 1. The backplane domain is run by the same clock as the `hiway` and `bitumen` cores in the master board gateware, currently 133 MHz. It holds a `bitumen` slave that communicates with the master board, and outputs the 48-bit instructions to a top-level core known as `minotaur`. Within `minotaur` the instructions enter a core known as `odecoder`, which decodes them and either forwards them to a set of dual-port block RAMs (BRAMs) on the FPGA, or to 2048-slot FIFO buffers `odec_fifo`, one per DDS channel.

The 250 MHz sync clock from DDS 1 is divided to 125 MHz using a phase-locked loop (PLL) on the FPGA, which provides the DDS domain clock and imposes a minimum timing resolution of 8 ns. The DDS domain has four channels, one per DDS, each holding four main cores, `mdds`, `mvga`, `mtimer` and `idecoder`, which together produce an rf pulse sequence. The BRAMs bridge the backplane and DDS domains, with each BRAM storing instructions or parameters for the cores in the DDS domain. They have separate read and write interfaces, which allows them to be written at any time by `odecoder`, even as a core is reading them from the DDS domain during a running sequence.

Upon power-up, the AD9910s by default divide their 1 GHz input clock by 2, and produce a 125 MHz sync clock that causes the DDS domain to run at half its operational speed until the DDSes are initialised over SPI; additionally if the DDSes are reset or their configuration registers are mis-programmed, the DDS domain clock can halt. For these reasons `odecoder` handles top-level resetting of the cores in the DDS clock domain as well as the DDSes

[22] Analog Devices ADL5330.
[23] Texas Instruments DAC5672A.
[24] Analog Devices ADCLK846 clock fanout buffer.
[25] Analog Devices AD9248.

themselves, since its own clock depends only on the master board being programmed.

As the current gateware was being developed, an earlier version[26] programmed over USB without real-time pulse sequencing capabilities or backplane communication was used for early experiments [80, 106]. USB capabilities were also implemented in the current gateware, primarily for debugging without the master, which allows the board to be used outside the M-ACTION system for simple experiments.

DDS domain cores

The `mdds` core handles most communication between the on-board FPGA and the DDS. It reads frequencies, phases and amplitudes from the 512×18-bit `mdds_spi` BRAM, in which the 14-bit amplitudes and 16-bit phases take a single address each and the 32-bit frequencies take two. It then serialises a frequency, phase and amplitude and transfers them to a DDS profile word at 62.5 MHz using a fast SPI core. The SPI core was optimised to have minimal delay. Its frequency is half of the DDS domain clock, which is close to the maximum DDS data rate of 70 Mbps. While an SPI transfer is running, the data for the next transfer is read in parallel, so that the core itself adds no latency to DDS communications. A reply from the DDS could also be read from the core, which is useful for debugging when it is connected to a gateware debugging interface such as Xilinx Chipscope; the read data is not otherwise used. `mdds` also controls the 16-bit DDS parallel bus, with more details in §3.3.8.

`mvga` handles the DAC channel controlling the VGA, and is responsible for pulse shaping. The shaping speed is read from the `mvga` BRAM. The core generates a linear ramp based on this, whose output value is used as the input to a 1024×18-bit BRAM (`mvga_lut`). The logarithm of the pulse shape is stored here; each value consists of 13 bits for the amplitude and 5 bits for the gradient between the current and the next point in the BRAM. This plays a similar role to the DDS angle-to-amplitude converter as shown in Figure 3.9. The currently-used shaping function is a Nuttall window[27], according to $w(\tau) = 0.3558 - 0.4874\cos(\pi\tau) + 0.1442\cos(2\pi\tau) - 0.0126\cos(3\pi\tau)$, where τ represents the normalised shape time ($0 \leq \tau \leq 1$) and $w(\tau)$ is the fraction of the maximum VGA gain. The core linearly interpolates between values and is able to efficiently produce smooth amplitude shapes, which is important especially for pulses shaped over tens of µs (which admittedly were not used in this thesis work) due to the exponential response of the VGA.

`mtimer` manages the timing of the rf pulse sequence, coordinating the other cores in the DDS domain. Internally it has a 36-bit counter that counts up from the start of a pulse sequence, and a 36-bit comparison register which is always incremented by the differential times of pulses or pauses in the sequence. This allows for a single $8\,\text{ns} \times 2^{36} = 9$-minute phase-coherent sequence to be run on the DDS boards. Whenever the counter exceeds the value in the comparison register, the core waiting for a timer event is triggered and carries out its operation. Figure 3.9b shows its operational principle. The most complex example of its use is in shaped pulses, where `mdds` and `mvga` are tightly coordinated. The timer also forwards the lower 32 bits of its counter, minus a reference time, to `mdds` for calculating the absolute phase. The reference time is 0 by default, and can be set to the current time by a special command; the former reference times are stored in an 8-level stack so that previous reference times can be recovered. This is used for phase-insensitive blocks of pulses, discussed in §3.3.8.

[26] Called the BRIQI gateware, based largely on the University of Innsbruck gateware [175].

[27] Similar to the more common Blackman window but with a sharper spectral rolloff; see https://en.wikipedia.org/wiki/Window_functions. Implemented in `minotaur/crete/maths.py`.

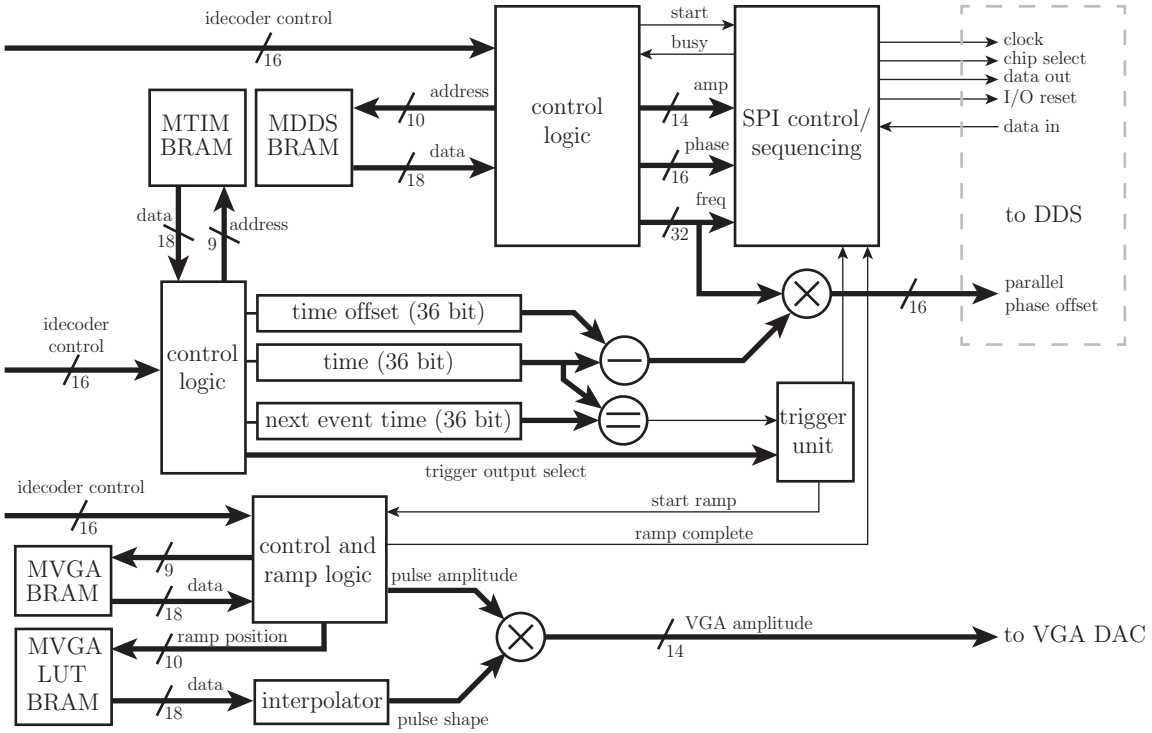

Figure 3.9: `mdds`, `mtimer` and `mvga` collective operation. `mddds` controls the DDS outputs, `mvga` controls the rf pulse amplitude shaping, and `mtimer` controls the timing of the pulse sequence.

`idecoder` is a 'pulse processor' that executes a sequence of instructions from the `idec` BRAM describing an rf pulse sequence. It is pipelined and executes one instruction per clock cycle in ideal circumstances. This performance is not yet required in the M-ACTION system due to the bandwidth limits of the AD9910 communication. There are two main instruction categories: flow control and control of the other cores in the DDS domain. Flow control includes instructions for conditional or unconditional program jumps, setting or decrementing internal registers, simple subroutines, and waiting for other cores to finish their task/s. Instructions to control the other cores include setting the DDS frequency/phase/amplitude, sequencing an event in the timer, or specifying a VGA ramp. Importantly, most of these specify the BRAM addresses from where a core must obtain the instruction parameters, rather than providing the parameters directly in the instruction. An SPI transfer instruction to `mdds`, for example, specifies a *pointer* to the `mdds_spi` BRAM addresses that it should read to obtain the frequency, phase and amplitude, instead of providing the parameters themselves. Much of the power of the `minotaur` design comes from this 'parameter pointer' principle, discussed next.

3.3.7 Real-time feedback, loops, and conditional forking

A simple rf pulse to carry out for example a π rotation consists of a frequency, phase, amplitude and duration, and all the parameters except the duration might be shared by a $\pi/2$ pulse. Because the `idecoder` instructions carrying out both of these pulses will refer to the same BRAM addresses except for the duration, a single write to the BRAM by `odecoder` will change the pulse properties for the entire sequence. The pulse is thus encoded to minimise communication from the master for routine pulse property changes.

This is used in protocols where parameters are varied not only between experimental points but between shots or even within a single shot, such as the scheme in §6.2 or the experiments in Chapter 7, since the feedback delay imposed by the ion is typically short (for cooling) and much of it is taken up by computation on the CPU.

Another use of this feature is to bypass BRAM size limitations, in protocols such as randomised benchmarking where an arbitrary sequence of pulses must be carried out. Instead of a random series of pulses, which would use a lot of memory, randomised benchmarking is carried out by looping over a *single* pulse, whose parameters are altered synchronously with the loop itself to create a random series. Here the CPU writes parameter updates into an `odec_fifo` instead of directly to a BRAM, and a special `idecoder` instruction flag causes the FIFO to write to the relevant BRAM at a well-controlled time; for example at the end of a single execution of a loop. This in principle allows for a near-infinite randomised sequence, as long as the CPU can continue feeding parameters into the FIFO more quickly than they are used.

The `idec` BRAM can similarly be written in real-time to conditionally fork the pulse sequence, for example to run one of several mutually exclusive sub-sequences based on an experimental outcome. The sub-sequences are written into sequential regions of the BRAM, and an unconditional jump operation precedes the first. Its jump target is the start of the default sub-sequence to execute. When M-ACTION reacts to an experimental outcome, it can choose the sub-sequence that will run by overwriting the instruction with a new jump targeted at the beginning of the desired sub-sequence[28]. This is not very efficient, as every `minotaur` channel on every board would in the worst-case require a different jump target address, causing around 36 µs of latency for a global fork. A 'conditional jump' instruction could easily be added to `idecoder`, which would read a global fork flag updated by a single backplane broadcast, reducing the latency to < 5 microseconds; however this would not be advantageous for many of our experiments because the fork delay occurs during slow experimental operations such as cooling.

3.3.8 Phase coherence and reference time shifts

When driving qubit rotations with an optical or microwave field controlled by the DDSes, the DDS rf phase determines the qubit axis of rotation. Two rf pulses generated from the DDSes must be in phase with one another, even if the same DDSes have generated pulses on other qubits/transitions (i.e. at other frequencies) in between. The AD9910 resets its phase accumulator in our configuration; this is needed for phase-coherent operation to set a predictable phase at the beginning of every pulse. By default a phase offset is calculated inside `mdds`, using the pulse frequency about to be applied and the time since the sequence was begun, according to $\phi = \omega(t - t_{\text{ref}})$. This is done continuously using a hardware multiplier, and the output is sent to the AD9910 via its parallel bus precisely when its frequency is switched. The result is phase-coherent operation as shown in Figure 3.10. Other hardware schemes such as calculating the phase using a set of parallel accumulators on the FPGA are also possible [175, 10].

An additional requirement is shifting the reference time t_{ref} for phase calculations; this can be required if the qubit phase evolves differently with the qubit drive on than with it off, due for example to ac Stark shifts (see the discussion in §6.1.2). This effectively causes the phase *difference* between pairs of pulses in the qubit frame to evolve as a function of

[28] This is a slight simplification of the instructions used for forking; see §A.4 for full details.

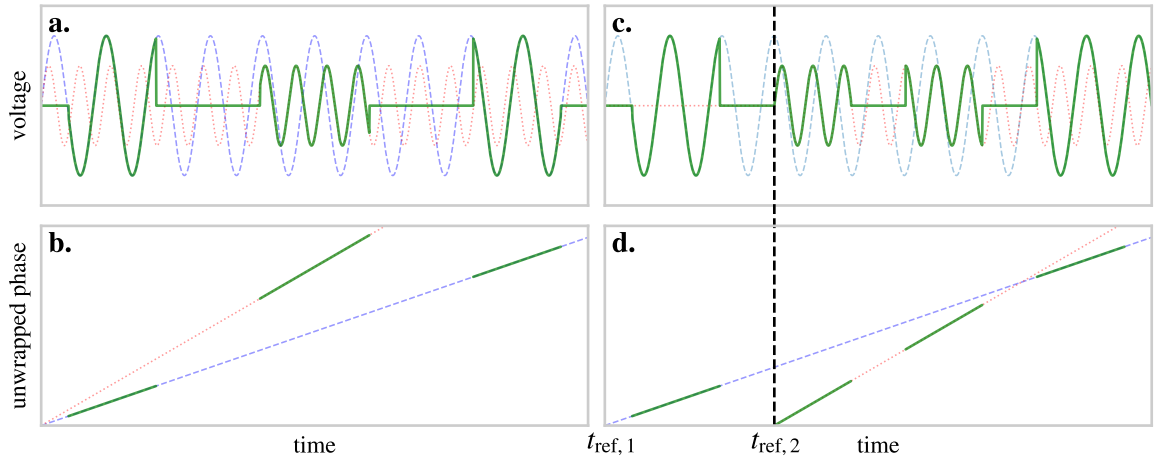

Figure 3.10: Phase-coherent operation and shifting of the phase reference time t_0. **a.** Frequency switching while phase coherence is maintained between Pulses 1 and 3. **b.** The unwrapped phase ramp in the pulses shown in a. **c.** Pulses 1 and 4 use $t_{\mathrm{ref},1}$ as a time reference for their phase coherence, whereas pulses 2 and 3 use a different reference $t_{\mathrm{ref},2}$, slightly before the beginning of pulse 2. **d.** The unwrapped phase ramp in the pulses shown in c. Dashed lines show the background sinusoids and phases for the two frequencies.

time t. By setting $t_{\mathrm{ref}} = t_p$ for the starting time t_p of the first pulse in a pulse block that needs to be phase-coherent, the time dependence of the phase difference is removed.

Having covered the main hardware and firmware features of the M-ACTION system, the software is presented next.

3.3.9 M-ACTION software and application programming interface

The C++ software running on the master board is hierarchical, with each stage building on functions and classes provided by lower-level code as explained in §3.1.2. A block diagram is shown in Figure 3.11a. The lowest *driver* level consists of drivers for physical hardware including Ethernet, serial and USB, and FPGA gateware such as the `pulser` and `hiway` cores. The second *API* level wraps the drivers and provides classes and functions for creating rf and digital pulse sequences, reading out the PMTs, communicating with the PC and DEATH boards, and creating *experiments* and *remote parameters* which provide a link to the PC GUI. The API level also contains mathematical libraries and algorithms for fitting, Bayesian estimation, randomised benchmarking and so on, relying on the EIGEN linear algebra library [59].

The third *global operation and experiment* level is unique to each M-ACTION setup, and its files are stored in a separate folder (e.g. `segtrap` for the mixed-species setup). It is written and maintained largely by experimenters rather than the people developing the hardware and software, and significantly differs between setups. It implements the trapped-ion 'toolbox', with functions for cooling, repumping/state preparation, qubit carrier and sideband pulses, multi-qubit single- and mixed-species gates, transport and separation, and ion readout. These operations are used by multiple experiments.

The top and final level consists of an Ethernet server handling requests from the PC. The PC may query the server about what remote parameters or experiments are available, modify remote parameters, execute experiments, or run auxiliary functions such as set-

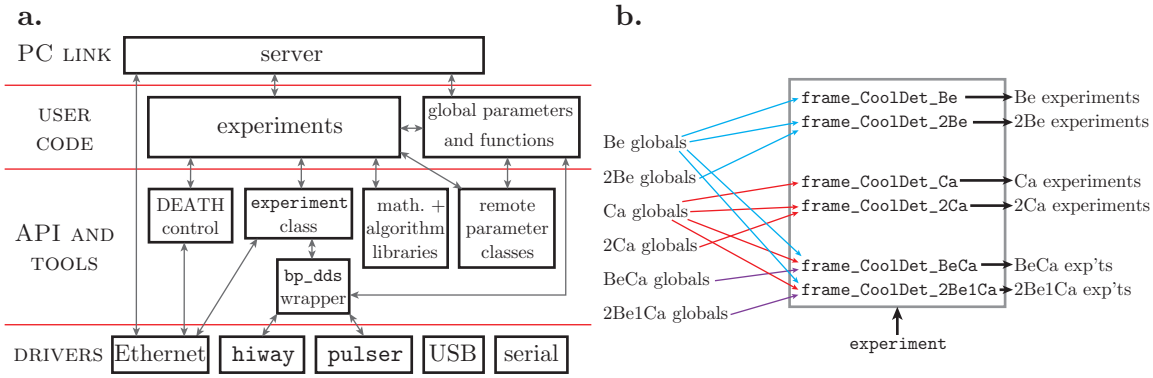

Figure 3.11: M-ACTION software hierarchy. **a.** A server handles requests from the PC, executing experiments or altering remote parameters. The experiments rely on global functions and classes, which in turn use the M-ACTION API that wraps the low-level drivers. **b.** Simplified inheritance of experiments and globals files in the mixed-species setup. Transport/separation and 2-Be-2-Ca frames are omitted.

ting the digital outputs manually or testing or resetting hardware. Requested information, experimental results, and any system warnings or errors are sent back to the PC.

Some key aspects of the software design are discussed next; §A.1 presents more information.

Pulser and hiway driver level

The `pulser` driver is straightforward, with functions to read and write its registers, carry out digital pulses of arbitrary length, read out the PMT FIFOs, and check the status of the core, including whether timing errors due to an empty `wfifo` have taken place.

The `hiway` driver is more complex, to handle the interlinked structure of the DDS board `minotaur` gateware. At the lowest level are commands for board communication and management, such as sending instructions or data and checking the board status or resetting them. The data and functions belong to a `DdsManager` class that centrally tracks the driver and DDS board state.

A `DdsBramObject` class stores arbitrary data for the `minotaur` BRAMs; it has a BRAM index, address, size, a 32-bit 'channel mask' specifying which of the 32 DDS channels use it, and Boolean flags indicating whether an address is assigned and the data has been sent to the DDS. Various `Setting` classes inherit from `DdsBramObject`, including frequency, phase, amplitude, time and shaped pulse settings.

A pulse sequence is constructed from `Pulse` classes, most importantly the `Edge`, `Cap`, `Shaped` and `Wait`.Each stores a list of `Setting`s that it requires to run, for example a frequency, phase, amplitude and time for an `Edge`. A `Pulse` provides the `idecoder` instructions to execute a DDS pulse, obtaining the parameter addresses from its `Setting` list. Importantly, multiple `Pulse`s may refer to a single `Setting`, allowing the sequence parameterisation discussed in §3.3.7.

After the basic `Pulse`s are specified, they are put into a `Sequence` object according to their order of execution. A `Sequence` handles flow control such as triggering and jumps, and along with `Pulse`s it can store sub-`Sequence`s for loops, forks or subroutines. The top-level pulse sequence for a single DDS channel is held in a `LoopSequence`. It consists of DDS

configuration that is run once per experimental point, then a *main loop* that is run once per DDS trigger (i.e. once per shot).

The main loop contains all the pulses, sub-loops and other flow control to run the shot. When a `Sequence` is written to the DDS, it checks the internal consistency of all its subordinate `Sequences` and `Pulses`, for example to make sure they have valid BRAM addresses and their channel masks cover its own; it intelligently assigns them addresses or masks by default, and alerts the user if an inconsistency arises. The same checks are then performed by its subordinates on *their* subordinates and so on; this assigns valid addresses and masks to every `Setting`, `Pulse` and `Sequence` used by the DDS channel. Once every `Sequence` has completed this process, all the data and instructions in the hierarchy are sent to the DDS boards, which takes hundreds of microseconds to milliseconds depending on the pulse sequence complexity, and the experiment can begin.

A child class of `DdsManager` called `Boss` is used to automate the consistency checks and manual management of `Sequences` in most M-ACTION experiments. It has a list of 32 top-level `LoopSequences`, one per channel, and provides functions to add `Pulses` to one or multiple channels and begin or end loops or forks, handling the details internally. It tries to enforce equal pulse times on all the channels; for example if one channel is running a `ShapedPulse`, the other channels must either run pulses or waits of equal length (adding equal-length waits is automated). This allows the user to think of the pulse sequence as a global unit, although each DDS channel is fully independent in reality. This independence may be required in future experiments where QIP operations must be carried out in parallel on different trap zones, at which point the `Boss` code will require changes.

M-ACTION API level

The frequently used libraries in the API layer are summarised next; for more information see §A.1.

The `bp_dds` library wraps the lower-level `pulser` and `hiway` drivers, providing functions to create synchronised digital and rf pulses. Sets of frequencies, phases and amplitudes are stored in objects labelled `FPAs`, and times in `Times`. There are several commonly-used pulse constructs. An `Edge` consists of a wait time followed by a change of digital output and rf frequency, phase and amplitude. A `Cap`[29] is two sequential `Edges`; usually the first turns a DDS channel and/or digital output on after an initial time $t_{\text{on\_delay}}$ and the second turns them off after $t_{\text{off\_delay}}$. A `Shaped` pulse first sets the VGA to minimum gain, turns on the DDS channel, and ramps the gain smoothly to a maximum; after a delay (t_{plateau}) it ramps the gain down and turns the DDS channel off. A `Wait` pulse is simply a delay followed by a change of digital output; this is most often used for controlling external devices such as the DEATHs. Any of these pulses, usually the `Caps`, can be set to record PMT events during their execution; this is most useful for detection pulses, however can also be used to monitor ion fluorescence during Doppler cooling.

`bp_dds` also provides a convenient interface for loops using the DDS FIFOs, which is used most often for pulsed sideband cooling where the pulse length is swept during a cooling loop, or randomised benchmarking where pulse amplitudes, phases and lengths are altered arbitrarily. A vector of `Setting` objects, each associated with a vector of parameters that it must take during the loop, is supplied to a `run_loop()` function; the relevant parameters are thereafter transferred to the correct DDS FIFOs before every experimental shot, and as

[29] So-called since its shape looks like a baseball cap.

explained in §3.3.7 they are written to the DDS BRAMs in synch with the loop execution. Fork operations, including opening a fork in the pulse sequence execution, adding pulses to each path, and closing the fork, are also handled. Nested forks, loops within forks, and forks within loops up to a depth of 4 loops are possible; no experiment has yet required more than 2 nested loops however (the most complex M-ACTION experiments performed so far are discussed in Chapter 7).

Much of the `bp_dds` code is responsible for synchronising digital and rf pulses. Due to the fundamental differences in the `pulser` and `minotaur` architectures, this has become quite elaborate; for instance to carry out nested forks and loops, `bp_dds` implements a nested linked list of digital pulse objects. Because of this, `bp_dds` has to process the loops in the sequence twice; once to construct the TTL sequence and send the DDS pulse sequence to the cards, and a second time to update the DDS FIFOs with loop data in real-time. Most of the complexity could be removed if `pulser` were replaced by a simple BRAM-and-processor structure capable of looping and forking similar to `minotaur`.

The API also supports the DEATHs, providing an interface to create and run sequences of voltage waveforms, real-time forking, and applying static offsets to individual voltage channels. It was written by David Nadlinger and is discussed in [132]. A sequence of waveforms is created using a `DeathSequence` object, and is transferred to the DEATHs at the beginning of an experimental point. A digital trigger from the master causes the DEATHs to advance through one waveform in the sequence, which is cyclical and repeated in every shot. The DEATHs are discussed further in §5.2.

Remote parameters are objects that can be used in the C++ code as regular variables, however their values are controlled by the PC GUI. They are used to alter settings such as rf pulse properties and times as well as higher-level experiment properties such as what kind of pulse sequence to carry out, whether to include various pulses or not, and so on. The API provides remote parameters for integer and floating-point numbers, Booleans (tickboxes in the GUI), matrices and 'combo-boxes', which are integers which appear in the GUI as drop-down menus with user-defined menu items. Floating-point remote parameters, `rp_double`s, can be scanned from the GUI to produce 1d or 2d plots. Mathematical operations can be carried out on remote parameters; for example the phase offset of an rf pulse is often a sum of several remote parameters specifying times, multiplied by one specifying a frequency.

Finally the API provides an `experiment` class, which is the major link through which the GUI can run experiments. An `experiment` is shown as a graphical 'page' whose remote parameters can be modified or scanned; a scan repeatedly runs the `experiment` while altering the parameter for each point, producing one or multiple live plots. When an `experiment` is run, it acquires an experimental point. The user must define several functions: `init_vars()`, `init_pulse_sequence()`, `read_out_pmts()` and `calc_results()`. The first two initialise the system and prepare for acquiring multiple shots, clearing experimental variables and sending the rf and digital pulse sequences. Next, the sequence is run multiple times (often 50) to acquire statistics. `read_out_pmts()` is run once per shot, reading out the PMTs and performing any real-time operations such as thresholding the PMT counts, status monitoring or forking. Once the shots have all been finished, `calc_results()` often simply calculates the average of the thresholded PMT outcomes, however the calculations may be more elaborate for experiments with multiple PMT readouts per shot, where for example the correlations may be extracted. Finally the experiment returns the user-defined set of results for the point as well as the raw PMT data for each shot. Experiments can also have user-defined *remote actions*, which appear as clickable buttons in the GUI; they run tasks

such as turning off the DDS outputs or manually transporting the ion.

Experiments and global operations

New experiments are usually built on simpler experiments via C++ object inheritance, where a child experiment extends the features of a parent. Experiments targetting a particular ion crystal (e.g. two-Be^+ , or Be^+-Ca^+) usually use identical cooling, state preparation and detection sequences. 'Frame' experiments provide a common set of these operations; experiments inheriting from a frame need only to implement a custom `QubitManipulation()` function, which is a few lines long for simple sequences.

Using inheritance is inconvenient when experiments require different, partially-overlapping subsets of a shared set of operations. This is handled with *globals* files containing many different calibrated operations, which experiments can individually use as required. For example, a `globals_qip_Ca` file implements carrier and sideband π and $\pi/2$ pulses, Ramsey sequences and randomised benchmarking for a single calcium ion, and can be used by whichever experiments require these functions.

The experiment inheritance in the mixed-species setup is shown in Figure 3.11b. At the time of writing there are around 160 experiments for the mixed-species setup, although most are seldom used, as well as 18 globals files. Future expansion of the system will require streamlining and optimising the experimental code; currently it forms a 90 MB executable when compiled. A more efficient scheme where experiments or experiment groups can be dynamically loaded will be useful in the medium-term.

3.3.10 M-ACTION graphical user interface

The M-ACTION GUI, called Ionizer2, is an evolution of the NIST Ionizer GUI with many additional features. Written in C++/Qt, it communicates with the master board, DEATHs and multiple asynchronous devices (see §3.4) using a standard protocol based on the MsgPack library, discussed in [132]. The protocol can be used in Python and many other languages, allowing the replacement of Ionizer2 if necessary[30].

When Ionizer2 is started, it queries the master board about the available experiments, the remote parameters of each (which may be shared by multiple experiments), and the current digital output status. From this information it dynamically generates a control page for every experiment, with a GUI element for each remote parameter. Global parameters are shown with a grey background, and altering them on one page will affect them on every other page where they are used. Ionizer2 also has static pages for asynchronously controlling hardware including the 397 nm, 866 nm and 854 nm cavity piezo voltages, trap shim electrode voltages, piezo mirrors and the DEATH calibration settings and waveforms. Parameters from the experimental or static pages may be scanned from any experimental page, with controls for setting the start, stop and number of points to acquire. A scan over multiple parameters can be used to create a rectangular colour plot.

Ionizer2 displays the results from each experimental point in plot windows, with the number of windows and plotted data determined by the experiment on the master board. Single-ion calibrations require only a single plot window, whereas complex mixed-species experiments require up to six. Both the result data and the PMT outcomes for each shot are saved to hard disk as timestamped text files, along with all the parameter settings used

[30] Michael Meth and Philipp Schindler, private communication, 2017.

for the scan. A completed scan can be fitted to calibrate remote parameter values such as frequencies or times.

Each experiment has a priority, and when several are started in parallel their priorities determine how often each is executed. A continuous or 'scrolling' scan at a low priority is often used to run an experiment in the background, such as cooling and detection of the ion crystal.

A Python-based scripting interface can be used to automate much of the experimental operations. A script can set parameters, carry out scans and perform complicated logic on the results. Scripts are often used to keep experimental parameters calibrated without human intervention, with functions that allow a script to intelligently decide if a parameter scan looks 'normal', and if not to check and calibrate other parameters that the calibration experiment relies on.

More details about Ionizer2 will be presented in the upcoming doctoral thesis of Matteo Marinelli [114].

3.3.11 Future M-ACTION upgrades

The M-ACTION system has been designed with the primary goal being executing the current generation of trapped-ion QIP experiments. Given the limited time available for engineering work in a physics setting, deliberate compromises were made that may affect the long-term system capabilities. Foreseeable issues are discussed here to guide the planning of hardware, gateware and software upgrades and maximise the longevity of the design.

A new master board will be required once production of the Zedboard is halted, which may occur in the coming years. A replacement board will not be difficult to design or buy, since there are no especially sensitive or high-speed elements in the Zedboard design. Although not yet required in foreseeable experiments, high-speed serial links at 1 Gbps or more may become useful; the downsides of increased complexity and latency will require some consideration. The master gateware is unlikely to require any major changes, though the `bitumen` communication protocol could be improved for higher bandwidth and the backplane clock speed could be increased with little effort; altering the design of `pulser` could simplify the M-ACTION API however.

The DDS board hardware is based on well-tested elements that will not require replacement any time soon; at the time of writing the FPGA and DDS chips are widely used and will remain in production for years[31]. It should be capable of meeting the needs of trapped-ion QIP until the experimental paradigm changes fundamentally, such as optical control electronics integrated in the trap.

Although it can currently handle rf parameter changes as short as 100 ns, in limited circumstances, the DDS board gateware was designed with microsecond-length pulses in mind and will need updates to arbitrarily modulate waveforms over 100 ns timescales or faster. If this becomes required, for example to carry out ultrafast entangling gates, the memory bandwidth will grow significantly and the FPGA BRAMs in the Minotaur design may no longer be enough. The DDS boards come with 128 MB of separate DDR memory, which could be used in an upgrade. The `idecoder` instruction set could also be made

[31] Analog Devices still lists the AD9910 as 'recommended for new designs', the category for chips at the beginning of their production cycle, although it has been in production for over a decade. Additionally Analog Devices notifies users of impending obsolescence two years in advance to allow designs to be migrated. Xilinx has stated that they will continue producing the Spartan-6 FPGAs for at least another decade.

significantly more efficient per clock cycle. Improved DDS gateware may also involve a pulse and sequence 'library', consisting of subroutines read from DDR memory and cached in the BRAMs. Some feedback flexibility would be lost in such an architecture, so two types of parameters could be used. Those that never require real-time changes could be stored in DDR, while those requiring real-time alteration could remain in BRAMs as in the current design. These gateware changes could probably be made within several weeks, since most of the DDS board gateware complexity lies in the interface to the DDS chips, which could largely remain unchanged.

Another important area that will need improvement is the M-ACTION API and experiment code. While the low-level hardware drivers were carefully designed to detect foreseeable inconsistencies, edge cases and user errors, their wrappers such as `bp_dds` have been shaped over time by user demands and currently require more user knowledge to avoid making mistakes. After a testing period of a few months, the time required by users to write and debug experiments turned out to be significantly less than expected; for this reason the API documentation was never polished beyond a basic level. Because the experimental requirements, structure and typical usecases of the API have been clearly established over several years, its streamlining and documentation could be easily carried out.

With these potential upgrades in mind, it is hoped that M-ACTION will remain useful for trapped-ion QIP experiments in the foreseeable future.

3.3.12 M-ACTION compared to other experimental control systems

A new experimental control system similar in architecture to M-ACTION, known as the Advanced Real-Time Infrastructure for Quantum physics (ARTIQ), has been in development [108]. At the time of writing its production hardware is similar to M-ACTION, and higher-performance designs are being developed. It is a promising platform with the backing of multiple experimental groups.

ARTIQ consists of a hybrid architecture, with the central 'master' implemented on a fast FPGA rather than a physical CPU as is available on the Zynq chip. The CPU is programmed by the user in a subset of Python rather than C++, which allows experimenters to use a single language at all levels of the system. At the time of writing the real-time computational performance of the CPU is significantly lower than the ARM CPUs used in M-ACTION, however, and because existing numerical libraries and algorithms cannot yet be used directly on the CPU, more effort is required to implement real-time computations as needed for adaptive Bayesian schemes or rapid feedback. These are important considerations in the continued use of M-ACTION. Being based around open hardware, gateware and software, the ARTIQ-compatible boards could easily be interfaced with M-ACTION or vice-versa, since the risk-prone step of developing and debugging the hardware has been completed for both M-ACTION and most ARTIQ PCBs.

3.4 Asynchronous devices

Many devices in the experimental setup are not synchronously controlled by the M-ACTION system during experimental execution. These include the laser frequency, intensity and fibre noise stabilisation setups, the magnetic field stabilisation and feedforward, fixed-frequency sources, shim electrode voltages, and mirror piezo controllers. These devices are briefly summarised in this section.

Reliably communicating with these tens of USB and serial devices has been problematic from a single desktop computer due to the variation of port names, long USB cables, driver issues and so on. The strategy used instead is discussed next.

3.4.1 Controlling devices using the Raspberry Pi

The Raspberry Pi (RPi) is a single-board computer [158] that has greatly simplified our control of standalone always-on devices. Eight RPis are used in the mixed-species setup, all running Arch Linux ARM, a version of the Arch Linux distribution ported to ARM CPUs[32], which offers the standard Linux benefits including a reliable operating system, C++ and Python. Four RPis host servers that control the digital PID boards, fixed-frequency rf sources and piezo voltages for the Raman beam mirror angles. Two more RPis communicate with slow serial DACs[33] controlling shim electrode voltages and the voltages of piezos used to adjust the lengths of cavities used as optical frequency references. The piezo and shim voltages are controlled directly from Ionizer2, and their optimisation is discussed in the next chapter. One RPi is used exclusively for the magnetic field feedforward, discussed below. A final RPi acts as the network gateway and firewall to the internet. Each RPi and it software can be remotely configured, rebooted and upgraded via secure shell (SSH) login. The network topology is shown in Figure 3.2.

3.4.2 Frequency and intensity stabilisation

Together with Ludwig de Clercq a mixed-signal board featuring two 10-bit ADCs, one 12-bit slow DAC, one 14-bit fast DAC and an FPGA was developed, called the electronically-variable interactive lockbox (EVIL); see [37] for details. An EVIL rack is shown in Figure 3.12a. This is used for most of the analog feedback tasks in the mixed-species setup, which are outlined here.

The frequency of laser light must often be stabilised to an optical reference, and vice versa. For example the frequencies of the 397 nm, 866 nm and 854 nm diode lasers in the calcium setup are referenced to medium-finesse cavities using the Pound-Drever-Hall (PDH) scheme [43, 18], as are the doubling cavities used in the beryllium laser setup, which must be resonant with the incoming light to operate. This relies on sending frequency-modulated light into the cavity and using its differential effect on the optical sidebands to probe the frequency difference between the light and cavity resonance, then beating the reflected light on a photodiode, demodulating the beatnote and using the dc-coupled result as a frequency error signal to feed back on either the cavity or the laser frequency. A standalone board, the Keitch Integrated Laser Lock (KILL) produces the modulation signal, demodulates the beatnote with a mixer and filters the output [74]. A firmware version of the EVIL containing built-in digital demodulation was tested, however the signal-to-noise was marginal [46]. Six PDH locks are controlled by EVILs in the setup, all with standard EVIL firmware except the beryllium Raman frequency-doubling lock. This feeds back simultaneously on the doubler cavity piezo and the seed laser current for better lock stability.

A related technique known as modulation transfer spectroscopy is used to stabilise the 626 nm beryllium laser frequencies before they are doubled, by using iodine vapour in place of a cavity and counter-propagating beams to eliminate room-temperature Doppler broad-

[32] Arch Linux ARM can also be run on the Zedboard. It is smaller and has more up-to-date software than more commonly used RPi operating systems such as Raspbian.

[33] Analog Devices AD5371.

ening [28, 180]. A 'pump' beam of several mW is frequency modulated by an AOM at 125 kHz [105, §5.4], and the modulation is transferred to a weaker 'probe' beam of a few hundred µW by the vapour through modulated hole burning [135], where the differential effect on the sidebands is similar to PDH. A commercial lock-in amplifier is used to demodulate the beatnote and produce an error signal. The iodine locks have a lower signal-noise ratio than the PDH locks and their dc setpoint is laser power-dependent.

Another use of the EVILs is laser intensity noise cancellation (INC), where the dc input of a variable-gain rf amplifier or analog multiplier is adjusted to vary the rf power into an AOM, or the gain of a tapered amplifier is altered by adjusting its current. The error signal is the dc output of a photodetector elsewhere in the setup. Three INC locks are used, one on the beryllium detection beam and two in the 729 nm setup.

The final use is fibre noise cancellation (FNC). Due to acoustic or other forms of noise, the optical path length of a fibre can change over kilohertz timescales, introducing significant phase noise in the output beam. This can be cancelled by shifting the beam frequency (and phase) using an AOM driven by an analog VCO, passing the beam through the fibre and back, optically beating it with the un-diffracted AOM beam on a photodetector to produce an rf tone at twice the VCO frequency, and electronically beating this tone with a stable local oscillator [109, 141]. If the phase of the beat signal is within $\pm\pi/2$ of the reference oscillator phase, the resultant signal is proportional to the phase shift caused by the fibre, and can be used to feed back on the VCO, lock the optical phase to the local oscillator and cancel the fibre noise. A standalone board with a VCO and attenuator, mixer and filtering is used to produce the error signal [141]. Two FNC locks are used in the 729 nm laser setup.

Each EVIL board has a programmable digital output that can be used to alert the M-ACTION system when a lock fails using a lock status input (as mentioned in §3.3.4; the primary purpose of the lock detect algorithm is automatic relocking, discussed in [69]), although this has not yet been used during experiments. The EVIL parameters are controlled by C++ servers running on the RPis that also stream the EVIL lock signals to a live Python dashboard on the PC; see [132] for details.

3.4.3 Fixed-frequency sources

Many rf sources are used in the setup that only need to be remotely switched on or off, with their frequency seldom altered after initial tuning. These include sources for permanently-on EOMs used for PDH locks and AOMs used in intensity servos, and phase references for fibre noise cancellation. For these applications the mixed-species setup uses 6 'TPI' rf sources [206][34], which produce up to 9 dBm of rf from 35 MHz to 4.4 GHz and are powered and controlled over USB. Internally they use a flexible PLL with an integrated wideband VCO[35] configured via a USB-to-serial chip[36]. Although a Windows GUI is available, a direct interface was written in Python to allow control of multiple sources from an always-on Raspberry Pi through a simple webpage, along with saving and loading PLL parameters. Sources must be connected to a Raspberry Pi through an externally-powered USB hub, due to their variable and occasionally high current requirements.

[34] At the time of writing we have not tested the latest version commercially available.
[35] Analog Devices ADF4351.
[36] FTDI FT232H.

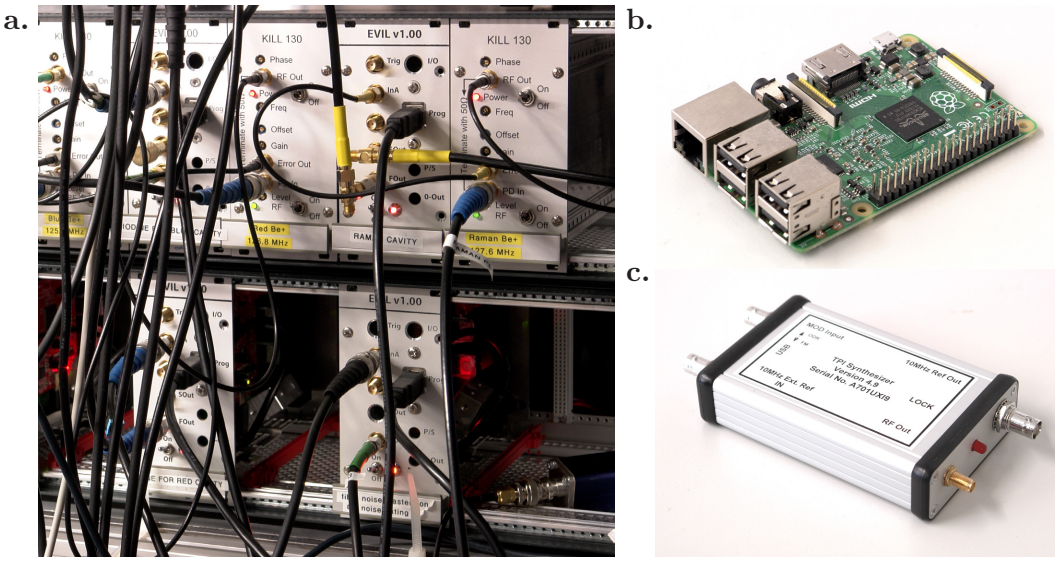

Figure 3.12: Asynchronous devices commonly used in the mixed-species setup. **a.** Eurocard rack holding 4 active EVIL PID boards. **b.** RPi single-board computer used as a laboratory device controller. **c.** 'TPI' rf source.

3.4.4 Magnetic field stabilisation and feedforward

While we usually operate the beryllium ion on a field-insensitive transition at 119.45 G, the calcium coherence time is limited by magnetic field noise. This arises both from noise in the current supply for the magnetic field coils[37], which runs at 115.2 A, and the external field fluctuations caused by ac currents at harmonics of the 50 Hz mains frequency. Current noise is minimised by an analog loop that feeds back on the supply current. Its error signal is provided by a fluxgate sensor. See the upcoming doctoral thesis of Christa Flühmann [49] for more details.

Mains noise is tackled by generating phase-locked harmonics at 50 Hz, 100 Hz etc. using a Raspberry Pi controlling a DAC. These are applied through two additional field compensation coils, with their amplitudes and phases chosen to cancel out the noise as measured by the ion. These are controlled from the Ionizer2 GUI, and can be calibrated using an Ionizer2 script written by Brennan MacDonald-de Neeve [110]. This iteratively applies a series of randomly-chosen amplitudes and phases at each feedforward harmonic, and measures their effect on the ion using a Bayesian scheme (see §6.2). From this data, it estimates both the background magnetic field noise and the transfer function between the feedforward and the ion for each harmonic, using a linear complex regression. The amplitudes and phases to counteract the background noise are calculated from these parameters.

[37] Agilent 6682A.

4 Experimental operation

This chapter presents the methods used to run experiments, optimise the mixed-species setup and calibrate parameters for the more complex protocols discussed in later chapters. The basic experimental sequence, loading and ion state detection and readout are presented, before simple coherent experiments and calibrations are discussed. The chapter concludes with a summary of micromotion compensation for mixed-species experiments.

There are multiple calibrations that rely on similar techniques to those presented here, including repumping, EIT and precooling beam optimisation and alignment, rf drive calibration and dark→bright and bright→dark leakage rate estimation in beryllium; these are discussed in Appendix B.

4.1 Basic sequence

A single experimental shot is similar in all the experiments presented in this thesis. The following steps are carried out.

1. Several milliseconds of intense far-detuned Doppler 'pre-cooling' then $\sim 600\,\mu s$ of near-resonant Doppler cooling ensure that the ions are cooled after any strong motional excitation due to background gas collisions, dipole lasers that have drifted blue of resonance, or miscalibrated transport or separation operations. The level of fluorescence during the near-resonant Doppler cooling can optionally be monitored to check if the ions have heated up (see §5.6.4).

2. EIT and/or sideband cooling are optionally carried out, depending on the ion crystal and the experiment being run. EIT takes 200–400 µs and sideband cooling takes a similar length of time for both beryllium and calcium. Although the calcium 729 nm sideband pulses are 30–50 µs each compared to 10–20 µs for the beryllium FDQ co-90, fewer cycles of calcium sideband cooling are generally needed thanks to the EIT.

3. Internal state preparation, using optical pumping as discussed in §2.4. If the FIQ is being used for beryllium a π pulse on the FDQ to prepare the $|g\rangle$ state follows the regular state preparation.

4. Qubit manipulation, where the coherent quantum protocol is carried out on the internal and motional states of the ions. Often this consists only of coherent pulses on the 729 and Raman lasers, however it includes dissipation, transport and separation, cooling and detection in the experiments presented later in this thesis.

5. The ions are detected and based on their fluorescence the state is estimated, discussed further below. If the FIQ was used for beryllium a π pulse on the FIS before the detection shelves the $|e\rangle$ population to improve readout fidelity.

The shot ends with pre-cooling being switched on indefinitely, or until the next shot is run. Figure 4.1 shows an 'oscilloscope view' of a complete mixed-species pulse sequence. In

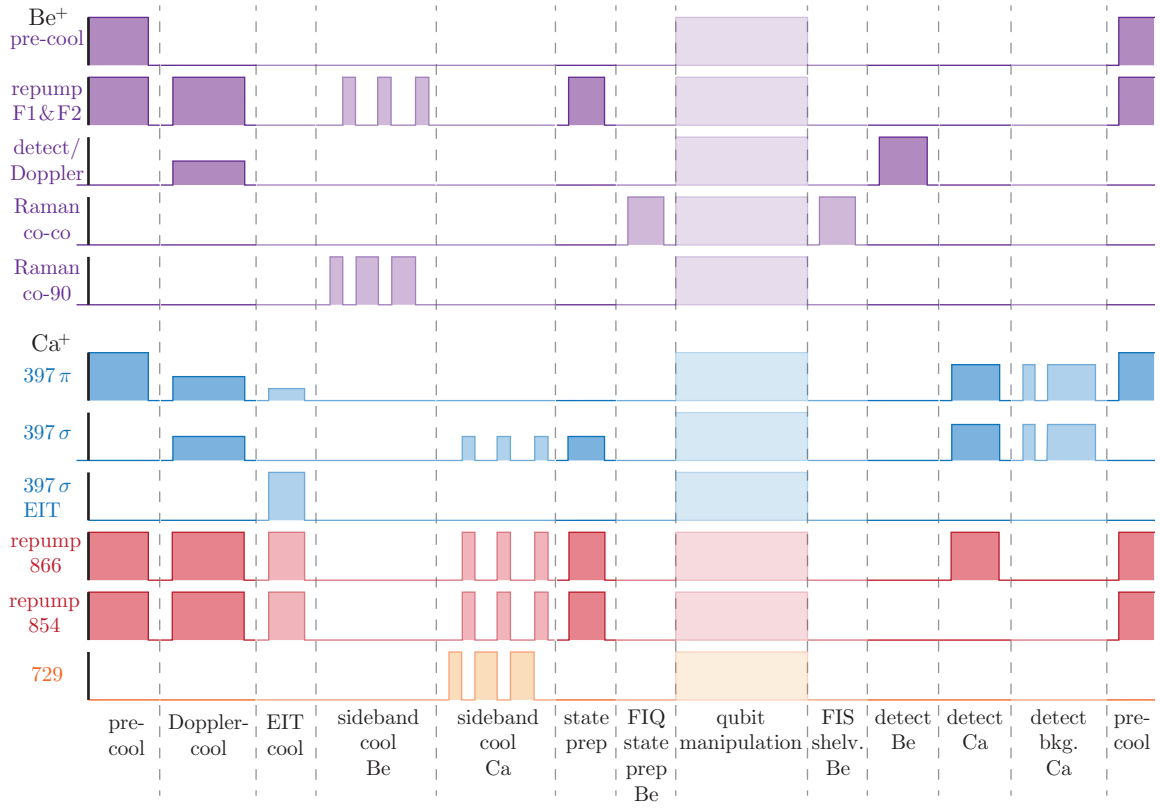

Figure 4.1: Pulse sequence diagram of a mixed-species experimental shot; pulse widths and heights are only a qualitative guide. Lighter sections are optionally run depending on the experiment, and single-species experiments omit the pulses of the other species.

the M-ACTION control code, the sequences are implemented for each ion crystal in an experimental 'frame' (see §3.3.9).

4.2 Ion loading

The basic ion loading sequences for beryllium and calcium are similar. The dc trap electrodes are configured to produce a shallow well in the left trap loading zone at -1870 µm. The sequence described in the previous section is run with only precooling, Doppler cooling, state preparation and detection activated for one or both species. One species is loaded at a time, and current is run through the relevant atomic oven, which directs a beam of neutral atoms through the loading zone. For calcium, photo-ionisation (PI) beams at 422.79 nm and 375 nm aligned on the loading zone are switched on, whereas for beryllium a single beam similarly aligned at 235 nm is used. The two-photon ionisation processes have a Doppler-broadened linewidth, and the 235 nm and 423 nm lasers are not locked to a frequency reference, with a 100 MHz-resolution wavemeter[1] being sufficient. When an atom loses an electron near the trap potential well, it is confined by the electric fields then Doppler cooled, becoming visible on the camera or PMTs. A 'loading conveyor' waveform, which sweeps a well from the loading zone to the experimental zone and merges with the existing well there, is run every few seconds while loading multiple mixed-species ions; it is discussed

[1] HighFinesse WS6-200.

in §5.6.2. Typical loading times when the beam alignment, frequencies, intensities and well properties are optimised are approximately 1 minute for a beryllium ion and 2 minutes for calcium.

When the detection beam alignment or frequencies are poorly optimised the calcium ion fluorescence can be lower than the background count rate on the PMT, making it difficult to distinguish when an ion is present. A two-stage 'background-subtracted detection' can be carried out where the first 397 nm + 866 nm pulse is followed by a short 397 nm–only pulse to shelve the ion into the $D_{3/2}$ manifold, then a second detection pulse without the 866 nm is run for which the ion is certain to be dark. The average PMT count difference between the two detection pulses will thus be due only to the ion. This also provides an unambiguous signal for optimising the detection beams. Since a closed transition is used for beryllium detection this technique cannot be carried out; however the beryllium background count level is < 0.5, low enough that weakly-fluorescing ions nonetheless provide a clear signature. It is significantly lower than calcium partly due to less beam scatter, and partly due to the dichroic filter that only reflects wavelengths below 325 nm to the beryllium PMT.

Further details of the PI lasers and the loading procedures for beryllium and calcium can be found in [105] and [81].

4.3 Detection and state readout

Because ion detection and state readout are fundamental to experimental optimisation and fidelity calculations, the approaches taken in the mixed-species setup are presented below in detail.

4.3.1 Thresholding

The basic aim of readout is to identify how many ions are bright and dark based on the PMT counts. For an n-ion single-species experiment with a set of $n+1$ probabilities $\{p_m\}$ that m ions are bright during state readout, the probability to obtain a total photon count λ from a PMT during the detection pulse is

$$p(\lambda) = \sum_{m=0}^{n} p_m P(\lambda; \mu_m), \quad \text{where} \quad P(\lambda; \mu) = \frac{\mu^\lambda e^{-\mu}}{\lambda!} \tag{4.1}$$

assuming a Poissonian model $P(\lambda; \mu_m)$ of the photon emission[2]. The Poissonian means, $\{\mu_m\}$, are approximately equal to $\mu_m = \mu_0 + m(\mu_1 - \mu_0)$ when the ions are evenly illuminated by the detection beam and imaged by the PMT. The average background count μ_0 arises from stray light, electronic noise and the PMT dark count rate, and $\mu_1 - \mu_0$ is the average bright count measured for a single ion.

We are usually interested in using the set of populations $\{\{p_m(\vec{a})\}\}$ from an experimental scan to estimate some experimental parameters $\vec{a}$, some of which are varied in the scan to obtain $\{\{\lambda(\vec{a})\}\}$ (see the definitions in Table 3.2). In the simplest case this is done by choosing $n+2$ integer thresholds $\tau_0 \ldots \tau_{n+1}$, where the expression

$$I_j = \sum_{\lambda=0}^{\tau_j - 1} P(\lambda; \mu_j) + \sum_{\lambda=\tau_j}^{\infty} P(\lambda; \mu_{j-1}) \tag{4.2}$$

[2]For computation of $P(\lambda; \mu)$ at high λ, the expression $P(\lambda; \mu) = \prod_{j=1}^{\lambda} \mu e^{-\mu/\lambda}/j$, $\lambda > 0$ is more numerically stable.

is minimised for $1 \leq j \leq n$. If $\tau_j \leq \lambda < \tau_{j+1}$ then it is assumed that j ions are bright, with the minimum and maximum thresholds $\tau_0 = 0$ and $\tau_{n+1} = \infty$. This threshold choice is equivalent to minimising the infidelity I_j due to misidentified results assuming uniform p_m populations, which provides a good guide for most scans over experimental parameters. Once S shots have been run for a constant set of experimental parameters $\vec{a}_p$ (an experimental point), thresholding the set of counts $\{\lambda_p(\vec{a}_p)\}$ yields B_m shots for each number of bright ions m. The populations are then estimated according to

$$\{p_m(\vec{a}_p)\} = \{B_m(\vec{a}_p)/S\}\,. \tag{4.3}$$

Each p_m is the result of a binomial process (for each shot, either the number of bright ions was m or not), and its uncertainty comes from the root-mean-square sum of binomial error and Laplace's rule of succession [127, §6.3.2], namely

$$\Delta p_m = \sqrt{\frac{p_m(1-p_m)}{S} + \frac{1}{(S+2)^2}}\,. \tag{4.4}$$

This scales as $S^{-1/2}$ for large S; in this thesis $S \geq 300$ for most data sets. The data $\{\{p_m(\vec{a}), \Delta p_m(\vec{a})\}\}$ from the full scan is then used to estimate $\vec{a}$, usually through a nonlinear weighted least-squares fit in this thesis.

4.3.2 Histogram fitting

For a 200 µs detection time in the mixed-species setup, $\mu_0 \leq 2$ for calcium and $\mu_0 \leq 0.3$ for beryllium. In both cases $\mu_1 \geq 25$, which gives worst-case beryllium and calcium infidelity for a single ion of $I_0 = 1.5 \times 10^{-6}$ and $I_0 = 3 \times 10^{-4}$ for thresholds of $\tau_0 = 6$ and $\tau_0 = 9$ respectively. For two ions, however, $I_1 \simeq 0.06$, which is large enough to dominate the binomial uncertainty Δp_1 and Δp_2 from Equation 4.4. It falls to 0.032 if the fluorescence is raised to $\mu_1 = 30$ yet remains the highest detection infidelity, and makes simple thresholding impractical for high-fidelity experiments with multiple ions at detection times of a few hundred microseconds; this is shown in Figure 4.2a. In beryllium, longer detection times are impeded by off-resonant optical pumping of dark states, while in calcium this is ultimately restricted by finite D state lifetime.

This can be mitigated with more sophisticated state estimation techniques; the simplest is a least-squares fit of Equation 4.1 to $\{\lambda_p(\vec{a}_p\}$, the aggregate histogram of the counts for an experimental point. This is done in two stages for the results in this thesis. A single 1D histogram is created using the counts from a complete scan $\{\{\lambda(\vec{a})\}\}$, and a fit floating both the Poisson means $\{\mu_m\}$ and probabilities $\{p_m\}$ is performed from which $\{\mu_m\}$ are estimated with low uncertainty. Individual fits are then carried out on histograms of $\{\lambda_p(\vec{a}_p)\}$ for each point in the scan, where $\{\mu_m\}$ are fixed and only $\{p_m(\vec{a}_p)\}$ are floated. The uncertainties $\{\Delta p_m(\vec{a}_p)\}$ are estimated as above.

There are known problems with the least-squares approach, especially when

$$p_m + \Delta p_m \approx 1 \quad \text{or} \quad p_m - \Delta p_m \approx 0, \tag{4.5}$$

leading to inaccurate estimates. Bayesian fitting[3] can be used to mitigate these by maximising a multiplicative probability likelihood function instead of minimising an additive error as in least-squares fitting [87, 127]. This can be done in three ways. Each $\{p_m(\vec{a}_p)\}$

[3] In practice these techniques are synonymous with 'maximum-likelihood fitting' in the literature.

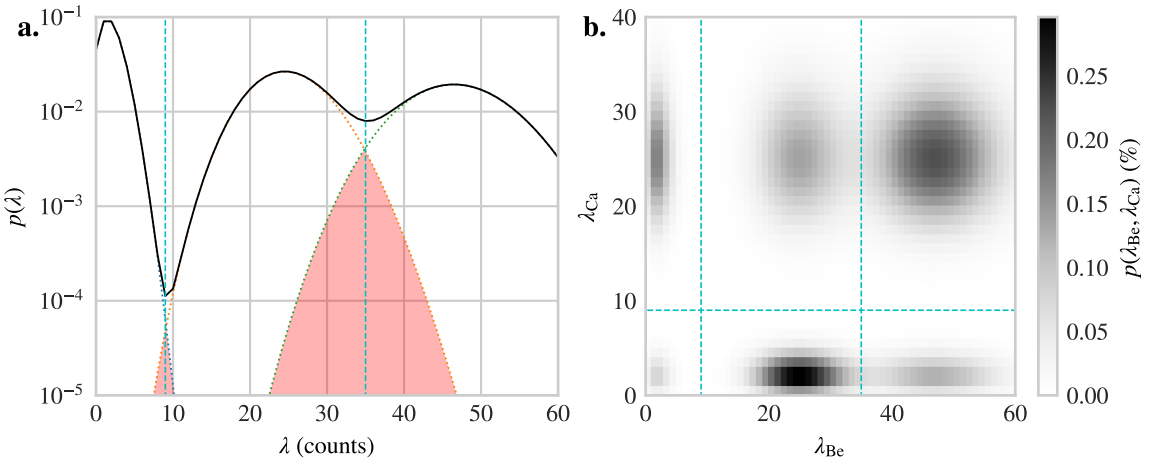

Figure 4.2: Poissonian histograms for single-species and mixed-species detection. **a.** Simulated histogram for two ions with Poisson means $\mu_0 = 2$, $\mu_1 = 25$ and $\mu_2 = 47$, using thresholds $\tau_0 = 9$ and $\tau_1 = 35$. Vertical dashed lines show the thresholds, and dotted lines are the individual Poisson distributions each scaled by $1/3$; the solid line is the sum. The overlap infidelities I_0 and I_1 are filled in red. With reference to Equation 4.7, the red area left (right) of the lower threshold is $p(0|ge)$ ($p(1|ee)$), neglecting state leakage. For the upper threshold the left (right) areas are $p(1|gg)$ ($p(2|ge)$). **b.** 2D histogram for two beryllium and one calcium ion, to which Poissonian distributions are fitted for estimating the correlations between the states. The underlying entangled state is $a_{0,0}\,|gg, g\rangle + a_{0,1}\,|gg, e\rangle + a_{1,0}(|eg, g\rangle + |ge, g\rangle) + a_{1,1}(|eg, e\rangle + |ge, e\rangle) + a_{2,0}\,|ee, g\rangle + a_{2,1}\,|ee, e\rangle$, with $a_{0,0}, a_{0,1} \ldots a_{2,1} = \{0.1, 0.28, 0.37, 0.46, 0.28, 0.69\}$ and the Poisson amplitudes given by $p_{b,c} = |a_{b,c}|^2$.

can be obtained using Bayesian fitting to the histogram $\{\lambda(\vec{a}_p)\}$, with the parameters $\vec{a}$ obtained from weighted least-squares fits to $\{\{p_m(\vec{a})\}\}$ as described earlier [87]. Alternatively $\{p_m(\vec{a}_p)\}$ are obtained from simpler thresholding or least-squares fits, then a Bayesian likelihood function is fitted to $\{\{p_m(\vec{a})\}\}$ assuming a binomial distribution for the p_m values [10]. Third a single likelihood function parameterising the complete set of raw counts $\{\{\lambda(\vec{a})\}\}$ directly in terms of $\vec{a}$ can be used on the entire scan [127]; this makes the fewest assumptions about the underlying data however is the most computationally demanding.

These were all found to yield negligible differences at the working fidelities in the mixed-species setup, and generally sufficient data was taken to ensure that Equation 4.5 was rarely true. The weighted least-squares approach was generally used since it took the shortest time to evaluate. This must be revisited once higher working fidelities are achieved, however.

4.3.3 State leakage during detection

There are two additional contributions to readout infidelity in beryllium that are negligible in calcium below several milliseconds of detection time. These are the chance that an ion initially in the bright state is driven into a dark state, one of the other $S_{1/2}$ states in beryllium, due to impure σ_+ polarisation in the detection beam; or that an initially-dark ion is repumped by the detection beam into the bright state[4]. In each case the cumulative probability that the ion has switched state during the detection increases exponentially with time, leading to histograms formed by the convolution of Poissonian and exponential distributions [1, 87], which can be used in offline data analysis.

[4]Or in calcium, spontaneously decays from the dark ($D_{5/2}$) to bright ($S_{1/2}$)

In our case bright→dark and dark→bright leakage contribute $\sim 0.5\,\%$ each to detection infidelity per beryllium ion on the FIQ in the mixed-species setup. The dark→bright leakage is around 8% for the FDQ, as can be seen in Figure 4.3. The effects of leakage can be reduced by making use of the photon arrival time information rather than only the total counts [87, 131], which helps distinguish cases where the ion has switched state midway through detection. This is implemented in our system by running the 200 µs detection as multiple short intervals of several microseconds, and the likelihoods of the ion being bright or dark are updated in real-time based on the counts in each interval; once the likelihood of bright or dark is above a desired threshold a conclusion can be made. Often this conclusion can be made before the ion state switches due to leakage, which is the source of the fidelity improvement.

This technique relies on rapid computation on the ARM CPU, and this is the first time it has been performed in real-time thanks to the capabilities of the M-ACTION system. It will be discussed in the upcoming thesis of Matteo Marinelli [114].

4.3.4 Mixed-species readout

In mixed-species experiments, the two PMTs return λ_{Be} and λ_{Ca} counts for every shot. For experiments using n_{Be} and n_{Ca} beryllium and calcium ions, there are $(n_{\mathrm{Be}}+1)(n_{\mathrm{Ca}}+1)$ probabilities $\{p_{b,c}\}$ that b beryllium and c calcium ions are bright. The probability of obtaining λ_{Be} and λ_{Ca} beryllium and calcium counts is the sum of $(n_{\mathrm{Be}}+1)(n_{\mathrm{Ca}}+1)$ Poissonian distributions:

$$p(\lambda_{\mathrm{Be}},\lambda_{\mathrm{Ca}}) = \sum_{b=0}^{n_{\mathrm{Be}}}\sum_{c=0}^{n_{\mathrm{Ca}}} p_{b,c} P(\lambda_{\mathrm{Be}};\mu_{b,\mathrm{Be}}) P(\lambda_{\mathrm{Ca}};\mu_{c,\mathrm{Ca}}). \tag{4.6}$$

There are $(n_{\mathrm{Be}}+1)(n_{\mathrm{Ca}}+1)-1$ degrees of freedom for the probabilities of a general state, and Equation 4.6 cannot be factorised into two 1D functions similar to Equation 4.1. The same approaches described in §4.3.1 can be used, however must be carried out in 2D. Estimating $\{p_{b,c}\}$ is straightforward with thresholding, with the 2D histogram split into rectangular regions as shown in Figure 4.2b; fitting Equation 4.6 to the experimental 2D histograms is also similar however it is computationally more demanding since $p(\lambda_{\mathrm{Be}},\lambda_{\mathrm{Ca}})$ is evaluated over a discrete grid instead of a line. Bayesian fitting approaches would suffer a similar increase in complexity, however have not yet been tried with mixed-species data.

4.3.5 Real-time readout and state leakage

During live experimental parameter optimisation, in particular of multi-qubit gates and protocols, reasonable estimates of $\{p_m(\vec{a}_p)\}$ are required for each point during data acquisition since often the calibration procedure consists of experimentally minimising or maximising a particular p_m level. Usually only 50 shots per point are taken to increase the speed of scans and reduce the effect of parameter drifts, however this limits the usability of histogram fitting, which requires hundreds of shots to reliably estimate p_m. The real-time Bayesian technique has been tried with multiple ions however is not yet reliable. Instead, the thresholded populations are corrected to reduce the several-percent error between p_1 and p_2. For two ions, the thresholded results can be expressed as

$$\begin{pmatrix} p_0 \\ p_1 \\ p_2 \end{pmatrix} = \boldsymbol{C}_{\mathrm{thresh}} \begin{pmatrix} \tilde{p}_0 \\ \tilde{p}_1 \\ \tilde{p}_2 \end{pmatrix}, \quad \boldsymbol{C}_{\mathrm{thresh}} = \begin{bmatrix} p(0|ee) & p(0|ge) & p(0|gg) \\ p(1|ee) & p(1|ge) & p(1|gg) \\ p(2|ee) & p(2|ge) & p(2|gg) \end{bmatrix} \tag{4.7}$$

where p_m are the probabilities of inferring m bright ions directly from the thresholded counts using Equation 4.3, and $\tilde{p}_m$ are the corrected population estimates of the quantum state. The parameters $p(m|ee)$, $p(m|eg)$ and $p(m|gg)$ are the probabilities that m bright ions will be inferred given the true ion states of $|gg\rangle$, $|ge\rangle$ or $|eg\rangle$, and $|ee\rangle$. If we neglect state leakage, these parameters are given by only the histogram tails extending beyond the thresholds, shown in Figure 4.2a, which makes $p(0|ge)$, $p(0|gg)$, $p(1|ee)$ and $p(2|ee)$ negligible, $p(0|ee) \simeq 1$, $p(1|ge) + p(2|ge) \simeq 1$ and $p(1|gg) + p(2|gg) \simeq 1$. Applying these equalities and inverting the system of equations, we obtain

$$\begin{pmatrix} \tilde{p}_1 \\ \tilde{p}_2 \end{pmatrix} = \begin{bmatrix} p(2|gg) & p(2|gg) - 1 \\ p(1|ge) - 1 & p(1|ge) \end{bmatrix} / K \begin{pmatrix} p_1 \\ p_2 \end{pmatrix} \tag{4.8}$$
$$\tilde{p}_0 = p_0, \qquad K = p(1|ge) + p(2|gg) - 1.$$

This empirical approach is currently used for live experimental calibrations, with $p(2|gg) = 0.985$ and $p(1|ge) = 0.988$ at the time of writing. It could be extended with the single-ion leakage rates by post-multiplying $\boldsymbol{C}_{\text{thresh}}$ in Equation 4.7 by

$$\boldsymbol{C}_{\text{leak}} = \begin{bmatrix} p(0|e)^2 & 2p(0|g)p(0|e) & p(0|g)^2 \\ 2p(0|e)p(1|e) & 2\left[p(1|g)p(0|e) + p(0|g)p(1|e)\right] & 2p(0|g)p(1|g) \\ p(1|e)^2 & 2p(1|g)p(1|e) & p(1|g)^2 \end{bmatrix} \tag{4.9}$$

and inverting, where $p(0|g)$ ($p(1|e)$) is the bright→dark (dark→bright) leakage probability and $p(1|g) = 1 - p(0|g)$, $p(0|e) = 1 - p(1|e)$. In offline fits to FDQ data the parameter $p(1|e)$ is non-negligible and is included in the fitting equation.

For two-beryllium one-calcium mixed-species readout, no corrections are made on the calcium readout and the input and output vectors in Equation 4.7 simply become two-column matrices, with the resultant system being[5]

$$\begin{pmatrix} p_{0,0} & p_{0,1} \\ p_{1,0} & p_{1,1} \\ p_{2,0} & p_{2,1} \end{pmatrix} = \boldsymbol{C}_{\text{thresh}} \begin{pmatrix} \tilde{p}_{0,0} & \tilde{p}_{0,1} \\ \tilde{p}_{1,0} & \tilde{p}_{1,1} \\ \tilde{p}_{2,0} & \tilde{p}_{2,1}. \end{pmatrix} \tag{4.10}$$

This correction operation is only applied on the averaged populations from multiple shots. For real-time readout carried out *within* a shot, however, currently the only possibilities are the Bayesian scheme and uncorrected thresholding. In the experiments where real-time detection played a role, thresholding was used; the infidelity was not severe since single calcium ions were being read out.

Having presented the detection schemes used in the mixed-species setup, common calibration experiments are discussed next.

4.4 Rabi oscillations

Carrier Rabi oscillations are the simplest and most versatile coherent experiments, and are used in many ways for calibration. The qubit $|g\rangle \leftrightarrow |e\rangle$ carrier transitions are driven near resonance by pulses as discussed in §2.2.2. Using Equation 2.21 for a near-resonant

[5] If multiple calcium ions are used in the future with a separate correction matrix, then instead of two-column matrices, single column vectors will be required, with the total correction matrix formed through a tensor product of the matrices of each species.

carrier pulse, if the ion starts in the bright state p_g then the bright-state population $p_g(t)$ (equivalent to $p_1(t)$ in the previous section) evolves according to

$$\begin{aligned} p_g(t) = 1 - p_e(t) &= 1 - \sum_{n=0}^{\infty} p_{g,n}(0) \frac{\Omega_{n,n}^2}{\left(f_n^0\right)^2} \sin^2 \frac{f_n^0 t}{2} \\ &= \frac{1}{2} + \frac{1}{2} \sum_{n=0}^{\infty} p_{g,n}(0) \frac{\Omega_{n,n}^2}{\left(f_n^0\right)^2} \cos f_n^0 t, \quad f_n^0 = \sqrt{\delta'^2 + \Omega_{n,n}^2} \\ &\simeq \frac{1}{2} + \frac{1}{2} \frac{\Omega_0^2}{\left(f_0^0\right)^2} \cos f_0^0 t \quad \text{in the Lamb-Dicke regime} \end{aligned} \tag{4.11}$$

where $p_g(t)$ is a partial trace over the motional state, i.e. the sum of the diagonal density matrix elements $\sum_{n=0}^{\infty} |c_{g,n}|^2$. If the detuning δ' is scanned for a π pulse with $t = \pi/\Omega_{n,n}$, then the resonance frequency is at the centre of the resultant sinc-like profile. If the time is scanned, a sine curve is observed that is slowest and deepest when the qubit is resonantly driven (see Figure 2.3a). The π time t_π can be estimated from a scan of the resonant oscillations. Iterating the two scans, the qubit frequency and π time can be obtained. The frequency can be calibrated more precisely by scanning the detuning for a $k\pi$ pulse where k is an integer. In the mixed-species setup the maximum k is limited to 30–50 by intensity noise, however; more precise calibrations use the Ramsey sequence discussed in §6.1.

Once the frequency of a transition has been calibrated, the t_π and $t_{\pi/2}$ pulse times are set by repeating k π- or $\pi/2$-pulses and scanning the pulse time. The π and $\pi/2$ times have an offset of between 100 ns and 1 µs on different lasers due to the AOM response times.

The 729 and Raman 90-switch beams each use a double-pass followed by a single-pass AOM, with the double-pass usually being used to modulate the frequency, phase and amplitude of pulses. The single-pass is used to generate two-tone laser light, used for spin-dependent forces and multi-qubit gates (discussed in §6.4), and it is usually driven at a default frequency and amplitude except when two-tone pulses are being run. This minimises thermal effects, which are more severe for the single-pass than the double-pass because the diffraction angle is not cancelled out. The AOM output beam is coupled directly into a fibre for both the 729 and 90-switch, and small variations in angle cause intensity fluctuations. The Raman co-com and co-switch beams use a single double-pass AOM each, so this issue is not as severe, although still present.

4.4.1 Qubit beam alignment

Carrier Rabi oscillations on the calcium qubit and the beryllium FDQ are used to optimise the qubit beam alignments. A single ion is used and the pulse time is set to carry out an odd number of $\pi/2$ rotations such that $p_1 \approx 0.5$ and dp_1/dt is positive[6]. By altering the beam alignment we then attempt to raise p_1, which would imply a Rabi frequency increase. If this is possible the pulse time is shortened to restore $p_1 \approx 0.5$ and the process is repeated iteratively. Care is taken to avoid crossing a sine peak beyond which dp_1/dt turns negative. Using shorter times minimises this risk, however the sensitivity is higher at longer times.

Two ions allow for more sensitive alignment, since they simultaneously probe two points in the beam separated by several microns. The qubits evolve independently and a misaligned beam will cause them to oscillate at different Rabi frequencies Ω_a and Ω_b for ions a and

[6]The routine can similarly be carried out with dp_1/dt initially negative.

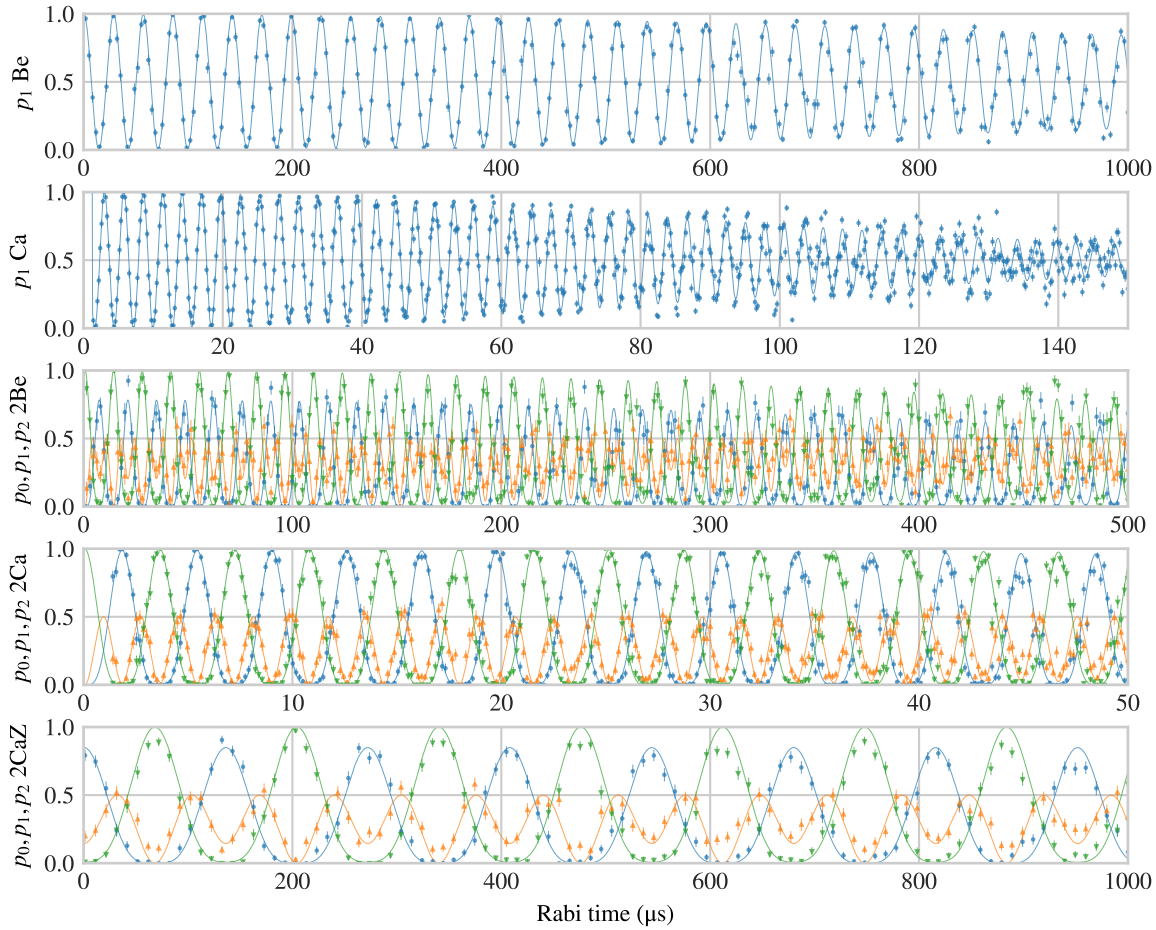

Figure 4.3: Rabi oscillations on various ion crystals, labelled on the y axes. The beams for '2Be' and '2Ca' were aligned using the separation technique. The '2Be' oscillations were on the FDQ, thus their contrast is lower than the FIQ oscillations shown for the single beryllium. '2CaZ' refers to two-ion oscillations on the Zeeman qubit, which begin in the dark state. For the two-ion plots, blue circles, orange upward-pointing and green downward-pointing triangles represent p_0, p_1 and p_2 respectively. The curves have been fitted to equations of the form Equation 4.11 for single ions and Equation 4.12 for two ions, with additional contrast terms that model dark→bright leakage as discussed in §4.3.5, and a decay term $\exp -(\Omega_{n,n}t/\pi\tau)^2$, where τ is a decay constant in units of the fitted π time. For beryllium, calcium and two-beryllium it is 116(6), 88(2) and 102(3) respectively. For two-calcium and the Zeeman qubit there was insufficient data to estimate it with reasonable confidence.

b (the n dependence is suppressed), as well as different detunings due to their different Stark shifts. Separating the ions and detecting them independently (see §5.5.2) produces two sinusoids for a time scan, whose frequency difference is an unambiguous error signal for alignment.

A different method is used for two-beryllium one-calcium crystals, and was used in the past for two-ion crystals before ion separation was implemented. In these cases the qubits are jointly read out, leading to the three populations

$$p_0 = p_{a,0}p_{b,0}, \quad p_1 = p_{a,0} + p_{b,0} - 2p_{a,0}p_{b,0}, \quad p_2 = 1 - p_0 - p_1 \tag{4.12}$$

where $p_{a,0}$ and $p_{b,0}$ are the single-ion dark populations from Equation 4.11 for the different

Rabi frequencies and detunings (time dependence is suppressed). Neglecting Stark shifts and assuming the ions are resonant, their oscillations interfere and p_1 becomes

$$p_1 = \frac{1}{2} - \frac{1}{4}\left\{\cos\left[\left(\Omega_a - \Omega_b\right)t\right] + \cos\left[\left(\Omega_a + \Omega_b\right)t\right]\right\}. \tag{4.13}$$

This describes a slow oscillation at $\Omega_a - \Omega_b$ superimposed on a rapid oscillation at $\Omega_a + \Omega_b$; the beam alignment is optimised by minimising the difference frequency, which is equivalent to minimising the average value of p_1 around long timescales.

The two-ion methods allow the Rabi frequencies to be matched to better than 0.5% for the mixed-species crystal, obtained by fits to the Rabi oscillations.The accuracy is limited by common-mode intensity noise, which destroys the oscillation contrast at long pulse times. This was beyond the angular resolution achievable by hand for the beryllium Raman beams when using a two-beryllium one-calcium crystal[7], which was the main reason for using piezo mirrors before the trap. Raman beam alignment is carried out using the co-co beams first; accurately aligning its piezo mirror simultaneously aligns the co-com and co-switch beams. The co-90 is aligned second, relying on the existing co-com alignment to optimise the 90-switch beam. Although this method is precise, it takes significant experience to carry it out quickly due to piezo hysteresis and the lack of an error signal that is robust to decoherence. Currently an adaptive Bayesian scheme is being investigated to produce an error signal with less ambiguity, and perhaps align the beam entirely without human intervention [200].

4.4.2 Beryllium qubit calibration

Once the beryllium beam alignment, FDQ frequency and FDQ co-co π time are calibrated, an FDQ co-co π pulse is used to prepare $|g\rangle$. Next the FIQ co-co π and $\pi/2$ times are optimised (the FIQ frequency is rarely recalibrated because it is constant by design), according to an overall sequence of (from left to right) $R_\pi^{\mathrm{FDQ}}, R^{\mathrm{FIQ}}, R_\pi^{\mathrm{FDQ}}$. Finally the FIS is calibrated with $R_\pi^{\mathrm{FDQ}}, R_\pi^{\mathrm{FIQ}}, R_\pi^{\mathrm{FIS}}, R_\pi^{\mathrm{FDQ}}$; this sequence highlights any population not transferred out of $|e\rangle$ by the FIS pulse as bright, which eases FIS calibration. The FDQ and FIQ co-90 π and $\pi/2$ pulses are sensitive to the ion temperature, as discussed in §2.2.2, and their times are calibrated once cooling has been optimised.

4.5 Sideband oscillations and thermometry

For a single ion initially in $|g\rangle$ with a motional-state probability distribution $\{p(n)\}$, it can be found from Equation 2.20 that driving a red or blue sideband causes the bright state population to evolve according to

$$\begin{aligned} p_g(t) = 1 - p_e(t) &= 1 - \sum_{n=0}^{\infty} p_{g,n}(0)\frac{\Omega_{n+s,n}^2}{\left(f_n^s\right)^2}\sin^2\frac{f_n^s t}{2} \\ &= \frac{1}{2} + \frac{1}{2}\sum_{n=0}^{\infty} p_{g,n}(0)\frac{\Omega_{n+s,n}^2}{\left(f_n^s\right)^2}\cos f_n^s t, \quad f_n^s = \sqrt{\delta'^2 + \Omega_{n+s,n}^2} \end{aligned} \tag{4.14}$$

where in the Lamb-Dicke regime, $\Omega_{n+s,n} \simeq \eta^{|s|}\Omega_0\sqrt{n+s}$ for $s = \pm 1$ for the first motion-adding and motion-subtracting sidebands respectively, as presented in Chapter 2. This is

[7] Due to the increased separation of 10.5 µm between the beryllium ions in the 3-ion crystal, compared to <4 µm for a typical 2-ion crystal, the angular sensitivity was increased.

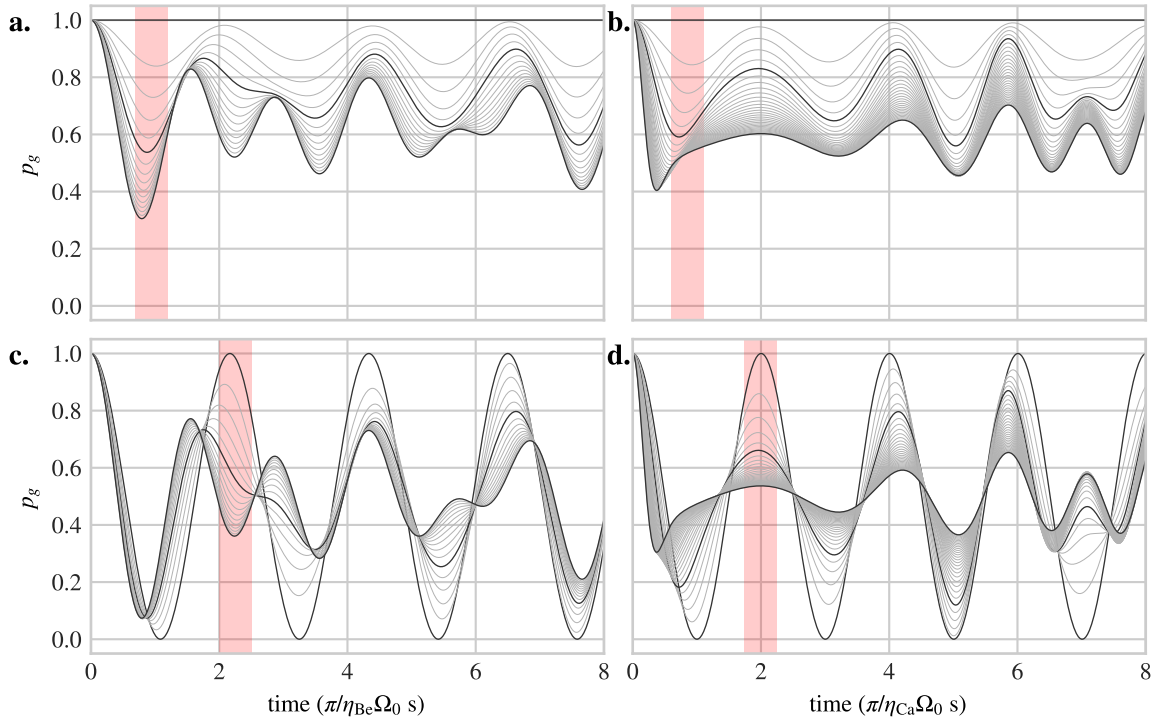

Figure 4.4: Beryllium and calcium sideband oscillations for thermal states with $0 \leq \bar{n} \leq \bar{n}_{\mathrm{Dopp}}$ where $\bar{n}_{\mathrm{Dopp}}$ is the Doppler-cooling limit. Solid lines mark $\bar{n} = 0$, $\bar{n} = 1$ and $\bar{n} = n_{\mathrm{Dopp}}$, with grey lines at $\bar{n} = 0.2$ intervals. The top plots show the motion-subtracting sidebands (red for calcium and the FDQ, blue for the FIQ) and the bottom show the motion-adding sidebands. **a.** and **c.** Beryllium, $\eta_{\mathrm{Be}} = 0.4$ and $n_{\mathrm{Dopp}} = 3$. **b.** and **d.** Calcium, $\eta_{\mathrm{Be}} = 0.05$ and $n_{\mathrm{Dopp}} = 6$. The Lamb-Dicke approximation has not been made. The red areas show the time regions at which the population p_1 can be used as a direct temperature diagnostic for optimising EIT or sideband cooling. The same patterns are visible in p_0 if a carrier π pulse precedes the opposite sideband used each plot (e.g. p_1 for the calcium blue sideband is equivalent to p_0 for a π pulse followed by the calcium red sideband).

a sum of Rabi oscillations where both the Rabi frequencies and spectral linewidths grow larger with increasing n, and it can be used for thermometry. Figure 4.4 shows the resonant sideband oscillations for different thermal states, whose occupancies are given by $p_{g,n}(0) = \bar{n}^n/(1+\bar{n})^{n+1}$. Driving sideband pulses for the times shown in red produces sensitive measures for minimising the temperature. For a ground-state cooled ion the motion-subtracting sideband ideally shows no excitation, and the motion-adding sideband is sinusoidal.

The sideband frequencies are asymmetrically influenced by Stark shifts from far-detuned transitions, and it is experimentally more convenient to carry out a red sideband with or without a preceding carrier π pulse to observe the motion-adding dynamics. For non-thermal motional states, the population evolution can be fitted to Equation 4.14 with $\delta' = 0$ to extract $\{p_{g,n}(0)\}$, where an extra exponential decay $\exp(-\gamma t/\Omega_{n+1,n})$ is included to model Rabi-frequency-dependent decoherence processes such as laser intensity noise [95]. The fitting equation is then

$$p_{g,\mathrm{fit}}(t) = \frac{1}{2} + \frac{1}{2}\sum_{n=0}^{n_{\max}} p_{g,n}(0) e^{-\gamma t/\Omega_{n+s,n}} \cos\Omega_{n+s,n} t. \tag{4.15}$$

For a thermal state, the motion-adding and motion-subtracting populations $p_{e,\mathrm{mss}}$ and $p_{e,\mathrm{mas}}$ provide an estimate of $\bar{n}$ independent of time [196], according to $\bar{n}_{\mathrm{est}} \simeq p_{e,\mathrm{mss}}/(p_{e,\mathrm{mss}} + p_{e,\mathrm{mas}})$. at one motion-adding sideband π time and for $\bar{n} < 0.2$ this is approximately $p_{e,\mathrm{mss}}$. This is useful for measuring the motional mode heating rates, since it requires little time for a rough initial estimate.

These sideband techniques are similar for multiple ions, however unlike the carrier oscillations, the ground state-cooled sideband contrast never reaches 1 due to the simultaneous interaction of the ions with a single motional mode. Otherwise the optimisation techniques are the same.

4.5.1 Tickling

Several shim electrodes lack the resistor in their in-vacuum RC filters, allowing an rf drive to be directly applied onto them. On the external filterboards, rf from the DDSes can be capacitively coupled onto the dc electrode lines, creating a coherent rf field in the trap with both a radial and axial component, modulating the potential well position and 'tickling' the ions with an oscillating force. The tickle frequency ω_t can be tuned for a sensitive measurement of the ion motional mode frequencies $\{\omega_m\}$. When $\omega_t \approx \omega_m$, the ions are excited and a loss of fluorescence is observed. This technique permits mode frequency measurements independent of ac Stark shifts.

The ion must be heated strongly to measurably decrease its fluorescence, and is easily lost due to momentary interruptions in cooling or human error. A variation of this technique can be used in which the ion is ground-state cooled, weakly tickled then driven with a motion-subtracting sideband pulse [68]. The sideband pulses use the same frequency parameter as the tickling, thus the laser remains resonant with the tickle as the tickle frequency is scanned. On resonance the weak tickle excites the ion from the motional ground state, which is visible from the sideband drive. Precisions of below 100 Hz can be obtained with this approach.

The tickle can also coherently manipulate the ion motional state. It can be calibrated to provide a precise displacement in phase space, which has been used for cancelling out excitations due to diabatic ion transport [23].

Tickling is currently not used often in the mixed-species setup for multi-ion crystals, with weak sideband pulses used to calibrate the mode frequencies instead.

4.6 Micromotion compensation

4.6.1 Radial compensation

As discussed in §2.1, stray radial electric fields cause excess micromotion by displacing the ion equilibrium position away from the rf potential null. This has adverse effects including altering transition lineshapes and Rabi frequencies [15] and shifting mode frequencies [68], and is particularly detrimental for mixed-species experiments because the displacement is ion mass-dependent.

To minimise radial micromotion we modulate the rf voltage v_{rf} at a modulation frequency ω_{mod} equal to that of a radial motional mode w_m, and observe the effect on the ion fluorescence [70]. This can be considered a 'modulated tickle'. If the ion equilibrium position u_e component along the motional mode vector is 0, this modulation has no effect.

If this is not true, the equilibrium position is modulated at the mode frequency and the radial mode is excited, causing a Doppler shift that reduces the ion fluorescence.

The stray field is compensated using the shim electrodes. Currently the seven electrodes in each radial quadrant of the trap are wired together to form four blocks, however they could be altered individually in the future to compensate different stray fields along the axis. With reference to the quadrant labels and coordinates of Figure 2.1b, the wafer voltages v_{A-D} are parameterised using two field gradients α and β according to

$$v_A = -v_D = \frac{1}{4}(\alpha + \beta), \quad v_C = -v_B = \frac{1}{4}(\alpha - \beta) \tag{4.16}$$

with the α and β degrees of freedom creating orthogonal fields along the trap x and y directions. We scan β against ω_{mod} at a fixed α value[8], and the optimal point is noted. This is carried out for several α values and both radial modes, forming two optimal lines in (α, β) space. Their intersection yields the $(\alpha_{\text{opt}}, \beta_{\text{opt}})$ values where the stray field is compensated in both directions. See [105] for previously-used techniques and further details.

Currently micromotion compensation is run on a single calcium ion once per day, typically with $-1 < \alpha_{\text{opt}}, \beta_{\text{opt}} < 1$, however the optimal compensation varies for different trap wells. This effect has not been systematically studied, and may indicate a miscalibration in the DACs and/or error in the dc electrode models, which is consistent with other anomalies (see Chapter 5). The thesis of Matteo Marinelli will describe investigations into these effects carried out using a mixed-species crystal [114]. The compensation values obtained with calcium are also used for beryllium and mixed-species experiments.

4.6.2 Axial compensation

If there is an axial component of the rf field, it causes axial micromotion that cannot be compensated with dc voltages. The axial position is modulated according to $z(t) = z_{\text{mm}} \cos(\Omega_{\text{rf}} t)$, where z_{mm} is the micromotion amplitude. This scales proportionally to the trap q_x and q_y parameters [81, 105] in Equation 2.2, namely $2e\beta_z v_{\text{rf}}/m\Omega_{\text{rf}}^2$, where β_z is a geometrical factor. Thus z_{mm} is larger for beryllium than calcium by their mass ratio $m_{\text{Ca}}/m_{\text{Be}}$. The micromotion causes a Doppler shift that frequency modulates the laser beams with a non-negligible wavevector component along the trap axis; this includes the beryllium detection and co-90 Raman beams. In the ion rest frame this introduces sidebands to the laser spectrum at integer multiples of Ω_{rf}, reducing the carrier Rabi frequency Ω. The new Rabi frequencies can be parameterised by a beam-dependent modulation index β, according to

$$\Omega_{\text{mm},n} = \Omega \, |J_n(\beta)| \tag{4.17}$$

where $n = 0$ is the micromotion sideband index and J_n are Bessel functions of the first kind. β depends on the beam wavevector component along the axis according to [4] $\beta = k_z z_{\text{mm}}$. Figure 4.5b shows the variation of $\Omega_{\text{mm},0}$, $\Omega_{\text{mm},1}$ and $\Omega_{\text{mm},2}$ with β. This is especially detrimental to the co-90 Raman beams, since the co-90 carrier Rabi frequency is reduced while the Raman scattering rate is unchanged, which requires using longer pulses and worsening the fidelity (discussed in §6.5.1).

Applying a counteracting axial rf field on an electrode has been used in an experiment in the Ozeri ion trapping group [117], however this was unsuccessful in our trap. Instead

[8] The 2D scan is required because the mode frequency changes as a function of β.

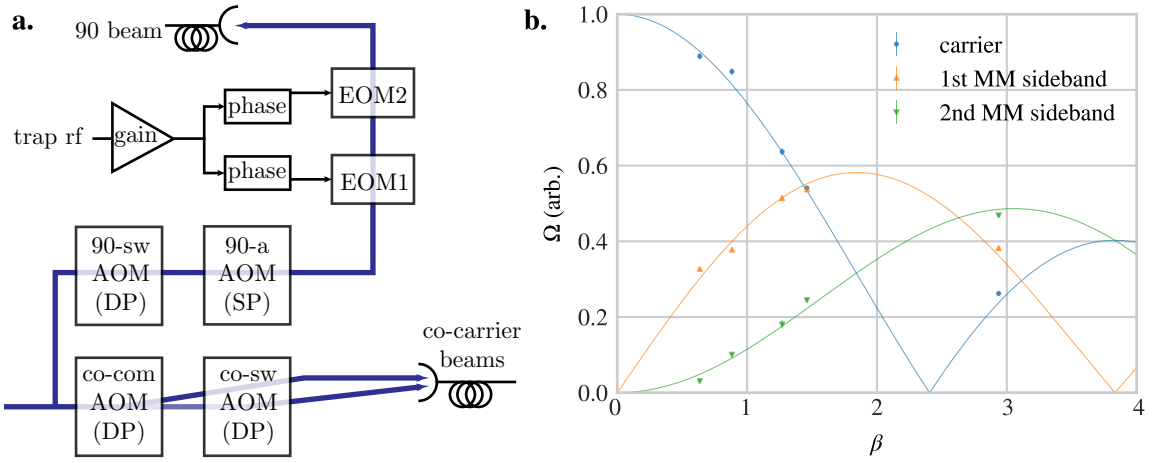

Figure 4.5: Axial micromotion compensation on the Raman beams. **a.** Location of the EOMs in the Raman beam setup. They are driven with an amplified and phase-shifted version of the trap rf, with the calibration procedure described in the text. DP and SP refer to double- and single-pass optical layouts of the AOMs. **b.** Fitted modulation index for various EOM settings. Each group of points shows the estimated Rabi frequencies from fits to Rabi time scans on the carrier, first and second micromotion sidebands using the Raman co-90 beam; the three frequencies are collectively fitted to the theoretical curve by floating β and Ω in Equation 4.17. The values at $\beta = 2.94$ were measured with the EOMs off, and the values at $\beta = 0.64$ were after optimisation. The intermediate points were recorded during the optimisation process.

in the mixed-species setup the beryllum detection and Raman co-90 beams are themselves frequency-modulated at Ω_{rf} using two EOMs with an adjustable amplitude and phase[9] as shown in Figure 4.5a. The micromotion sidebands are probed by detuning the Raman co-90 beam by $n\Omega_{\mathrm{rf}}$, and based on the Rabi frequency ratios the amplitude and phase are adjusted to minimise the sidebands. Initially EOM2 is disconnected, EOM1 is driven at a low amplitude and its phase is adjusted to maximally reduce β. EOM1 is then connected and its phase similarly adjusted Thereafter the rf amplitude is gradually increased until the first and second sideband Rabi frequencies are minimised, at which point the amplitude and EOM2 phase are iterated further. Using this approach β was reduced from 3 to 0.7. The sidebands could not be entirely nulled out for unknown reasons; we speculate that a nonlinearity or higher-order modulation could be the cause. In the Ozeri group, 50 Hz electric field noise was found to be the major reason why this occurred[10].

Each EOM is temperature-stabilised using a Peltier cooler, with a heatsink and small fan to dissipate the heat. The resonance frequency is highly sensitive to temperature and input rf power. If the rf is interrupted for more than 5 minutes, the EOM cools, and due to the shifted resonance it reflects too large a proportion of the rf power to warm up again; the resonance must then be manually re-tuned and tracked over around 20 minutes to ensure the EOM absorbs enough rf power to reach thermal equilibrium. Two sequential EOMs are used because around 35 dBm of rf power is used in total, and the resonance frequency of a single EOM driven at full power is unstable even with strong cooling, presumably due to internal temperature gradients.

The beryllium detection beam is modulated with a similar EOM at lower power, with

[9] This scheme was suggested by Ting-Rei Tan, private communication, 2015.
[10] Roee Ozeri, private communication, 2018.

the carrier detection fluorescence itself providing a diagnostic signal for optimisation. Its amplitude and phase rarely require adjustment, and were not characterised in detail during this thesis.

5 Ion transport and separation

Reliable positioning, transport, separation and recombination of ion crystals are core requirements for scalable trapped-ion QIP architectures based on the 'quantum CCD' [204, 76] introduced in §1.2, which relies on transporting small ion crystals around a multi-zone trap with dedicated regions for different operations such as readout or quantum gates. A quantum CCD additionally requires well-controlled transport through trap junctions.

The first transport experiments in the mixed-species setup were carried out by Ludwig de Clercq, who built custom dc-coupled AWGs to drive the dc trap electrodes. These were used to perform parallel transport gates [97, 38] and, with Robin Oswald, ion velocimetry for better controlling the velocities required for such gates [36, 144]. During this thesis transport was further investigated, optimised and more tightly integrated into various experiments. It was used for profiling and optimising beams (see §5.6.3) and implementing 'addressed' operations on a mixed-species crystal (see §7.4.2). Ion separation was implemented and optimised for various single-species crystals and investigations on mixed-species beryllium-calcium and beryllium-calcium-beryllium separation were begun. This chapter briefly reviews the theory behind transport and separation, describes the numerical process used for generating the waveforms and how it was optimised, presents suggestions for further improvement, and concludes by introducing several related auxiliary operations used in mixed-species experiments.

5.1 Ion position, frequency and dc offset

A trap electrode $j \in 1, 2, \ldots, J$ at a voltage v_j creates an electric potential $V_j(x, y, z)$ in the surrounding space, dependent on the geometry of the electrode and the environment around it. This potential can be expressed as

$$V_j(x, y, z) = v_j \phi_j(x, y, z) \tag{5.1}$$

where $\phi_j(x, y, z)$ is the potential at a point x, y, z created by the electrode for a 1 V test voltage, or the dimensionless *potential moment*. The total potential for a set of voltages is then

$$V(x, y, z) = \sum_{j=1}^{J} v_j \phi_j(x, y, z) \tag{5.2}$$

and the electric field is

$$\vec{E}_u(x, y, z) = \sum_{j=1}^{J} v_j \frac{\partial \phi_j(x, y, z)}{\partial u}, \qquad u \in \{x, y, z\}\,. \tag{5.3}$$

In the segmented trap $J = 30$. We are most interested in the ion behaviour along the trap axis. Electrodes on the top and bottom dc wafers are symmetric across the axis, and when opposing pairs such as electrodes 8 and 23 (see Figure 2.1e) have the same voltage,

the potential gradient in the radial directions is ideally zero everywhere along the axis. There is thus no radial electric field, and the ion is centred in the radial pseudopotential 'tube' along the axis[1], where $x = y = 0$. We can thus suppress the x and y dependence of ϕ_j and V. Positive singly-charged ions experience a force in the $+z$-direction of $F_z = -eE_z(z) = -e\,\partial V(z)/\partial z$, where e is the electric charge. Hence given some dissipation (i.e. laser cooling), the ions will settle in minima of $V(z)$.

Near a minimum at $z = z_0$, $V(z)$ is quite smooth over a range of tens of µm, and can be well-approximated by a quadratic potential $V_q(z)$. A short string of equal-mass singly-charged ions experiences this as a harmonic well [65], with

$$V_q(z) = \frac{m\omega^2}{2e}(z-p)^2 + d \tag{5.4}$$

where m is the ion mass, ω is the angular frequency of the axial common mode of motion, and d is a constant dc offset. Hence the angular frequency for $V(z)$ at a local minimum $z = p$ is

$$\omega^2 = \frac{e}{m}\frac{\partial^2}{\partial z^2}V(z)\bigg|_{z=p}. \tag{5.5}$$

The dc offset d alters neither ω nor the position of p. It is, however, correlated with the radial curvature, and affects the radial mode directions and their Lamb-Dicke parameters with respect to the cooling laser beams. It is chosen primarily to optimise cooling. Thus a trapping location can be parameterised by p, ω and d, where ω is specified with respect to the calcium ion mass.

Figure 5.1a shows the electrode moments $\phi_1(z)$ to $\phi_{10}(z)$ along the axis. They are obtained from an electrostatic boundary-element simulation of the segmented trap and its surroundings[2] at points on a 3D grid with 5 µm spacing, along and around the trap axis. By applying voltages to the electrodes, as shown in Figure 5.1b, the positions, frequencies and dc offsets of the minima of $V(z)$ can be controlled to within around 100 nm, within 1 kHz, and within 10 mV in the mixed-species setup. It is worth noting that, using the electrode numbering in Figure 2.1e, $\phi_j(z) = \phi_{15+j}(z)$ because the top and bottom electrode wafers are symmetric, and $\phi_j(z) = \phi_{16-j}(-z)$ because the trap is symmetric around $z = 0$.

5.2 DEATHs and their usage

The AWGs used to drive the inner dc trap electrodes are called the Direct Ethernet-Adjustable Transport Hardware (DEATHs). Each DEATH consists of a Microzed board[3] mounted on a motherboard hosting two dual-channel 16-bit DACs[4] that are run at 100 MSPS. Four analog amplification and filtering circuits produce dc-coupled voltages between -9.6 V and 9.6 V with a 50 MHz cutoff bandwidth [37]. The two channels of each DAC are used to drive matching electrodes on the top and bottom wafers, to ensure that temperature-dependent gain or offset errors between DACs do not cause voltage differences that displace the ions radially.

[1] Neglecting the effect of stray radial dc fields, which in any case shift the ions by at most several µm, not far enough to significantly alter the value of ϕ_j.

[2] COMSOL Multiphysics was used for the current moments, and the results agree to 0.1% with the NIST BEM package [47, 37].

[3] Avnet Microzed, a compact version of the Zedboard using a similar Zynq chip.

[4] Maxim MAX5898.

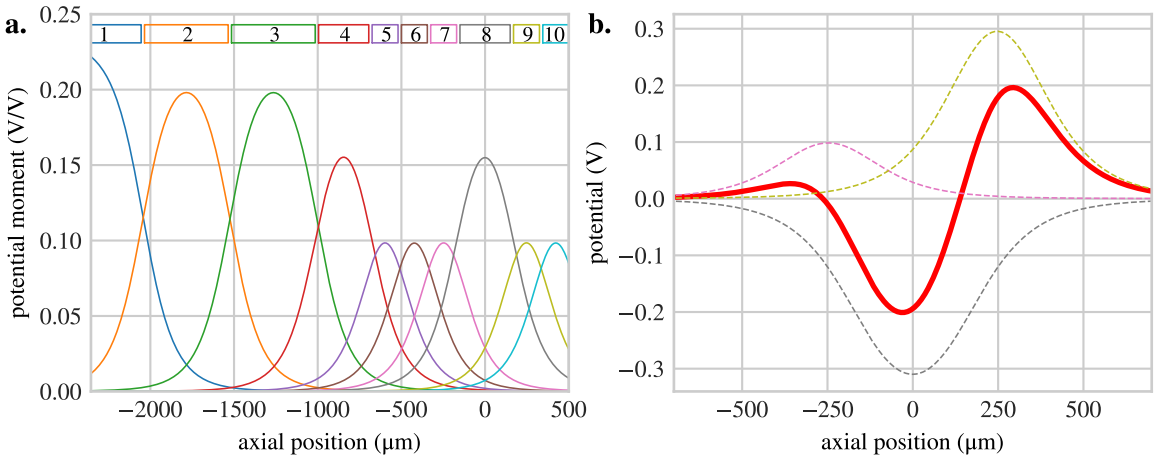

Figure 5.1: Electrode moments and potentials. **a.** The electrode moments ϕ_j of the top wafer, from the left loading zone of the segmented trap to the central experimental zone are shown, with electrode widths drawn to scale above. The bottom wafer moments are the same, and the right side of the trap is a mirror image around 0 µm. Note that the bottom wafer moments contribute equally in ion motion simulations. **b.** Creation of an example potential well by applying 1 V, -2 V and 3 V to electrodes 7, 8 and 9; their contributions to the total potential are shown by dashed lines. The resulting minimum is centred at -30.7 µm with a dc offset $d = -200$ meV and a second derivative of 1.44×10^7 V/m^2, resulting in an axial frequency $\omega = 2\pi \times 940$ kHz for Ca^+. Note that in practical waveforms, the same voltages would be applied symmetrically to the electrodes on the bottom wafer, leading to a doubled trap potential and second derivative (a $\sqrt{2}$ increase in ω).

The FPGA gateware to drive the DACs consists of BRAMs which store *voltage waveforms*, made up of voltage samples that are sent sequentially to the DACs in forward or reverse order, with a dynamically-adjustable delay between each sample. Currently 16 384 total samples are supported, which can be divided between up to 256 waveforms. The delay is an integer number of clock cycles called the *slowdown*; a slowdown of 0 leads to a waveform being played at the full sample rate of 100 MSPS, while a value of 1 halves this to 50 MSPS. A slowdown of 19 is commonly used, corresponding to a 200 ns sample period or a 5 MSPS update rate.

The software running on the ARM CPU is programmed in C++ similar to the Zedboard, however it is loaded from a FLASH chip rather than over USB from a PC and has hardly ever been modified since installation. It runs a server that receives *waveform sets* from the Ionizer2 PC GUI, where a waveform set consists of the voltage waveforms to be loaded to BRAM along with their metadata such as their descriptions, unique IDs, and starting and ending voltages. A waveform set is a JSON-formatted[5] text file storing the floating-point voltages of multiple waveforms [132], generated by numerical techniques presented in the next sections. For a complex mixed-species experiment the waveform set may contain single-well loading, loading conveyor, recrystallisation, reordering, transport and separation, and addressed rotations; all are discussed in this chapter except the addressed rotations (see §7.4.2). Ionizer2 handles the distribution of a waveform set between the different DEATHs; currently 8 DEATH boards are used for 32 voltage outputs. 28 are connected to electrodes, one is used for monitoring the waveform timing and 3 are spare.

The Zedboard master communicates with the DEATHs via the API discussed in §3.3.9,

[5] JavaScript Object Notation, http://json.org/.

specifying a cyclical sequence of waveforms to play through. For each waveform the master specifies whether to begin playing it immediately upon the end of the previous waveform or to wait for a digital trigger from the Zedboard, whether to play it forwards or backwards, and what slowdown to use. There is also a real-time forking capability which has not yet been used in experiments [132]. A feature was added to allow the master to apply a constant offset to any combination of electrode voltages, and alter this in between experiment shots; this was highly useful for investigating ion separation.

Having described the hardware used to supply the dc electrode voltages, the current processes for generating waveforms are discussed.

5.3 Generating transport waveforms

5.3.1 Finding the optimal electrode voltages

For a set of N desired static wells $\{\{p_1, \omega_1, d_1\}, \ldots, \{p_N, \omega_N, d_N\}\}$, we wish to find the set of electrode voltages $\{v_j\}$ that minimises a cost function C_{static} according to

$$C_{\text{static}} = \int_{z_\text{L}}^{z_\text{R}} \sum_{n=1}^{N} G_n(z') D_n(z') dz' = \int_{z_\text{L}}^{z_\text{R}} \sum_{n=1}^{N} \exp\left[-\frac{(z'-p_n)^2}{2R_n^2}\right] \left[V(z') - V_{q,n}(z')\right]^2 dz' \tag{5.6}$$

where for a desired well n with $\{p_n, \omega_n, d_n\}$, $V_{q,n}$ is defined in Equation 5.4, and z_L, z_R are the trap boundaries. The Gaussian term $G_n(z)$ defines a region of interest of width R_n centred around each p_n, where the local shape of $V(z)$ is the most important. It ensures that the cost function is unaffected by large $D_n(z)$ values far from the $z = p_n$ locations. It is parameterised such that its width coincides with a chosen minimum well depth V_R, according to $R_n^2 = 2V_R e/m\omega_n^2$. By this method R_n is widened for shallow wells. Usually V_R is set to a few mV.

Obtaining a set of voltages $\{v_j\}$ by minimising C_{static} is sufficient for positioning ions, and manually combining the voltages obtained from a series of positions into a single waveform is enough for slow adiabatic transport. For faster transport, some other effects must be included. We must treat the electrode voltages as time-dependent, namely $\{v_1(t), \ldots, v_J(t)\}$. The desired wells $\{p_n(t), \omega_n(t), d_n(t)\}$ are now time-dependent and follow *trajectories* through parameter space. With explicit t-dependence for clarity, Equation 5.2, Equation 5.4 and Equation 5.6 become

$$V(t,z) = \sum_j v_j(t)\phi_j(z), \qquad V_{q,n}(t,z) = \frac{m}{2e}\omega_n^2(t)\left[z - p_n(t)\right]^2 + d_n(t) \tag{5.7}$$

$$G_n(t,z) = \exp\frac{-\left[z-p_n(t)\right]^2}{R_n^2(t)}, \qquad D_n(t,z) = \left[V(t,z) - V_{q,n}(t,z)\right]^2 \tag{5.8}$$

$$\begin{aligned} C = \int_{t_0}^{t_f} \Bigg[& \int_{z_\text{L}}^{z_\text{R}} \sum_{n=1}^{N} G_n(t',z') D_n(t',z') dz' \\ & + \sum_{j=1}^{J} \left\{ r_0 \left[v_j(t') - v_{\text{def}}\right]^2 + r_1 \left[\frac{dv_j(t')}{dt'}\right]^2 + r_2 \left[\frac{d^2 v_j(t')}{dt'^2}\right]^2 \right\} \Bigg] dt' \end{aligned} \tag{5.9}$$

where t_0 and t_f are the initial and final times for the transport routine. There are additional cost terms, weighted by r_0, r_1 and r_2. The first penalises the deviation of $\{v_j(t)\}$ from a

default voltage v_{def}, and tends to stabilise the voltages chosen for electrodes far from any well locations, which are otherwise unconstrained by the cost function. The second and third terms penalise the first and second time derivatives of the waveform, which serves to reduce its spectral content at high frequencies and limit the slew rate demands on the DEATH output buffers [144].

Apart from the cost terms, two rigid constraints on the voltages are usually set. $\{v_j(t_0)\}$ and $\{v_j(t_f)\}$ are fixed to the values obtained using C_{static} in Equation 5.6 instead of C, so that if two waveforms begin or end with the same set of $\{p_n, \omega_n, d_n\}$, the derivative terms in Equation 5.9 do not shift their voltages and the waveforms will be continuous. Also, $|v_j(t)| \leq v_{\max}$ to stay within the DEATH output limits. Usually $v_{\max} = 8.9\,\text{V}$, to provide a small amount of overhead for waveform interpolation or filtering (discussed later). A third constraint on the maximum voltage slew rate, or $|dv_j(t)/dt|$, has been used in the past [36, 144] but not for the waveforms in this thesis, since it was found to cause unpredictable effects; instead a warning is printed when a slew rate above $5\,\text{V}/\mu\text{s}$ is detected.

In practice C is minimised using a numerical algorithm called *quadratic programming*, which out of several methods was found to be the most useful for solving the transport problem [37, 144]. Originally the code was implemented in MATLAB and relied on preparing a set of sparse matrices for the solver to use, however the algorithm was ported to Python to be compatible with the M-ACTION scripting (see §3.3.10) for automated waveform generation. It uses the CVXPY library[6], which automates much of the work of specifying the problem and tuning the solver.

The minimisation is carried out on a discrete grid of K z locations (usually $K = 943$, spaced at 5 µm intervals, from -2355 µm to +2355 µm) and M timesteps (usually $1 \leq M \leq 2000$), currently spaced at 200 ns intervals to match the typical DAC sampling rate during transport. Each term in Equation 5.9 can be discretised. For each electrode j the COMSOL simulation provides a vector of $\phi_j[k]$ moment values, where k indexes the axial grid. At a particular timestep we write the J electrode voltages as a column vector $\vec{v} = \{v_1, v_2, \ldots, v_J\}$. Combining the $\phi_j[k]$ column vectors in a $K \times J$ matrix $\boldsymbol{\Phi}$, the discretised axial potential at this timestep is the K-element vector $\vec{V} = \boldsymbol{\Phi}\vec{v}$. Similarly combining the $\vec{v}[m]$ column vectors in a $J \times M$ matrix $\boldsymbol{v}$, where m indexes the timesteps, we discretise the first part of Equation 5.7 as $\boldsymbol{V} = \boldsymbol{\Phi v}$, where $\boldsymbol{V}$ is $K \times M$ in size. For the second part we can express the n desired well trajectories as $\{p_n[m], \omega_n[m], d_n[m]\}$ at discrete times, writing them as row vectors $\left\{\vec{p}_n, \vec{\omega}_n, \vec{d}_n\right\}$ with M elements each, and tiling them vertically into $K \times M$ matrices $\{\boldsymbol{p_n}, \boldsymbol{\omega_n}, \boldsymbol{d_n}\}$. Similary we discretise z as a column vector and tile it horizontally for a $K \times M$ matrix $\boldsymbol{z}$. Now $\boldsymbol{V_{q,n}} = (m/2e)\boldsymbol{\omega_n^2}\,(\boldsymbol{z} - \boldsymbol{p_n})^2$, where all multiplications are elementwise. Using these definitions and elementwise operations, Equation 5.8 can similarly be discretised to obtain $\boldsymbol{G_n}$ and $\boldsymbol{D_n}$. Thus the discretised cost function becomes

$$C_d = \sum_{\text{mat}} \left(\sum_{n=1}^{N} \boldsymbol{G_n D_n} \right)^2 + r_0 \sum_{\text{mat}} (\boldsymbol{v} - v_{\text{def}})^2 + r_1 \sum_{\text{mat}} (\Delta_m \boldsymbol{v})^2 + r_2 \sum_{\text{mat}} \left(\Delta_m^2 \boldsymbol{v}\right)^2 \tag{5.10}$$

where all operations are elementwise, the $\sum_{\text{mat}}$ symbol indicates a cumulative sum of all matrix elements, and $\Delta_m \boldsymbol{v}$ is the discrete difference operator on the time axis of V (producing a $J \times (M-1)$ matrix). Each term in Equation 5.10 is a function of $\boldsymbol{v}$ and can easily be specified to a convex optimisation solver, to calculate $\boldsymbol{v}$ such that C_d is minimised.

[6] With the MOSEK solver backend; https://www.mosek.com.

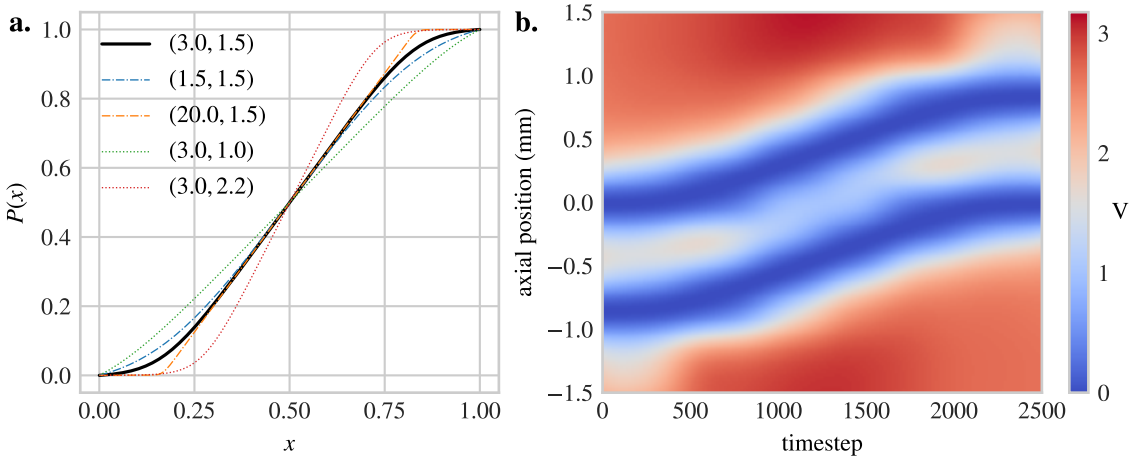

Figure 5.2: Two-well transport. **a.** $P(x)$ profiles using Equation 5.14 for different (a, b) values; the bold black line is used for most routines. **b.** Two-well transport with $p_0 = 0 \rightarrow 845\,\mu\text{m}$, $p_1 = -845 \rightarrow 0\,\mu\text{m}$, $\omega = 2\pi\times 1.6\,\text{MHz}$ (Ca^+), $d = 0\,\text{V}$.

5.3.2 Well trajectories and solver parameters

Avoiding motional excitation when transporting ions is desirable, reducing how frequently a crystal needs to be recooled. Adiabatic transport is straightforward to carry out since the exact trajectories of $\{p_n(t), \omega_n(t), d_n(t)\}$ do not matter, however the waveform may take hundreds of microseconds, comparable to the recooling timescales and negating the benefit of low excitation. Faster diabatic transport over tens of microseconds is thus desirable. If ω_n is held constant, diabatic transport produces a coherent motional state $\alpha_m(t)$ according to

$$\alpha_m(t_f) = \sqrt{\frac{m\omega_n}{2\hbar}}\left[-e^{-i\omega_n t}\int_{t_0}^{t_f}\frac{dp_n(t')}{dt'}e^{i\omega_n t'}\,dt'\right] \tag{5.11}$$

where the terms are defined as in Equation 5.9. If a linear ramp in $p_n(t)$ for a constant ω_n is used [89, 23], then $dp_n(t)/dt$ is constant and $\alpha_m(t_f)$ becomes

$$\alpha_m(t_f) = \sqrt{\frac{m\omega_n}{2\hbar}}\left[\frac{i}{\omega}\frac{dp_n(t)}{dt}\left(1 - e^{-i\omega_n(t_f - t_0)}\right)\right] \tag{5.12}$$

which can be precisely cancelled for a velocity $dp_n(t)/dt$ by choosing $\omega_n(t_f - t_0) = 2\pi k$, k an integer. This approach has been used to carry out transport over 370 µm in 8 µs [23], with a similar technique used to transport 280 µm in 4 µs [199]. It requires precise control of the waveform timing, something not yet tried with the DEATHs, and the compensation of distortion caused by the trap filters (see §2.6). Such fast timescales were unnecessary for the experiments in this thesis, and instead the well position $p_n(t)$ profiles were chosen to be smooth functions minimising the impulse on the ion [169, 12], which minimised the motional excitation for transport taking 30–100 µs without the need for careful calibration. Position profiles take the form

$$p_n(t) - p_n(t_0) = [p_n(t_f) - p_n(t_0)]\,P\left(\frac{t - t_0}{t_f - t_0}\right) \tag{5.13}$$

where $P(x)$ is a profile function whose domain and range are $0 \le x \le 1$ and $0 \le P(x) \le 1$. A common choice is $P(x) = \sin^2(\pi x/2)$, however this was found to be somewhat inflexible.

Instead $P(x)$ was defined according to

$$P(x) = \frac{\rho(x) - \rho(0)}{\rho(1) - \rho(0)}, \quad \rho(x) = \ln\left|\frac{\zeta(x) - ie^{-a}}{\zeta(x) - ie^{a}}\right|, \quad \zeta(x) = \exp\left[ab(2x-1)\right] \tag{5.14}$$

with various a and b values, where a determines the sharpness at the ends of the profile and b determines the slope in the middle; usually $a = 3$ and $b = 1.5$ were used. Figure 5.2a shows profiles for several sets of values. The $\omega_n(t)$ and $d_n(t)$ trajectories between the starting and ending values in a transport waveform were also usually set to follow Equation 5.14. With these settings, single-well transport has been carried out largely without any further adjustment.

Additional optimisation was carried out for multi-well transport, since the solver cost function grows more complex and often one electrode moment may contribute significantly to multiple $D_n(t, z)$ terms in Equation 5.9. A significant issue is maintaining high axial frequencies. This is desirable for many single- or two-ion QIP experiments, since Doppler and EIT cooling rates are improved if the axial sidebands are further from the carrier, so that a larger differential between carrier and sideband absorption can be achieved. Additionally the heating and motional decoherence rates from both technical noise and anomalous heating sources scale inversely on frequency [5, 65]. Generally the maximum ω at a trap location z_0 is limited by the local second derivatives of the moments $d^2\phi_j(z)/dz^2|_{z=z_0}$, the maximum achievable electrode voltages, and the desired dc offset (especially if it is positive). Problems arise when creating a multi-well transport waveform with axial frequencies of above $\sim$ 1.6 MHz, and the waveform results are usually checked against the expected trajectories to avoid heating or losing the ions.

Some tuning of the solver weights was required for robust results over a range of desired trajectories; Table 5.1 summarises the tradeoffs involved. Figure 5.2b shows a typical two-well transport routine.

5.4 Ion separation

5.4.1 Separation process

In the mixed-species setup, ion separation is designed to occur at the centres of electrodes 6 and 10. To understand the separation process, the potential around the separation location z_{sep} can be modelled as a 4th-order polynomial

$$V_{\text{sep}}(z_s) = \alpha z_s^2 + \beta z_s^4 + \gamma z_s + \delta + g z_s^3 \tag{5.15}$$

where $z_s = z - z_{\text{sep}}$ (i.e. $z_s = 0$ at the separation location) and z_{sep} is -422.5 (+422.5) µm for the left (right) separation zone, as shown in Figure 2.1e. The gz_s^3 term is assumed negligible if the electrode geometry and applied voltages are locally symmetric around z_s.

The separation distance s and axial centre-of-mass angular frequency ω of a two-ion crystal in this potential, including the effect of their Coulomb repulsion, can be calculated according to [64]

$$\beta s^5 + 2\alpha s^3 = \frac{e}{2\pi\epsilon_0} \tag{5.16}$$

$$\omega^2 = 2\alpha + 3\beta s^2 \frac{e}{m} \tag{5.17}$$

where ϵ_0 is the free-space permittivity and m is the ion mass; s is readily calculable with a root-finding algorithm for a numerical α and β using Equation 5.16.

An ion separation waveform consists of three stages. The first brings a potential well holding a multi-ion crystal to the location of the separation. Next a potential 'wedge' is driven through the centre of the well, splitting the crystal and dividing it between the two wells formed on either side of the centre. Finally the two wells are further separated until the ion crystals are hundreds of microns apart and can be handled independently. Recombination is simply the reverse of this process.

At the beginning of the second, 'splitting' stage, α is positive and β is small. Next β is increased, and α is decreased through 0 and made negative to create a double-well potential. It is desirable to keep the well frequency ω as high as possible while α is lowered, since this minimises the motional excitation of the ions due to imperfect waveforms and external noise; as discussed in §2.6 the electrode filters have a cutoff of a few hundred kHz, the same range ω reaches during splitting. It can be seen from Equation 5.17 that a high β results in the highest ω for any α value, thus β should be kept as high as possible during this stage. The time when $\alpha \simeq 0$ is especially important, and called the *critical point*. Note that in this chapter, *splitting* refers to this specific stage of separation, and not to ion separation as a whole (which also involves transport).

Early experiments notwithstanding [38], the requirements of splitting could not be handled optimally by the transport solver described in §5.3, as it uses a constant number of fixed-parameter wells and is not structured to favour certain well parameters at the expense of others. Another approach was developed to handle the splitting stage, while the first and third stage were still carried out using the transport solver.

5.4.2 Generating splitting waveforms

Several steps are involved in generating splitting waveforms. The first is fitting 4th-order polynomials centred around z_{sep} to the moments of the electrodes involved in separation. This is done for Electrodes 4 – 8, and their bottom-wafer counterparts, for the left splitting zone over a 200 µm-wide region of interest (ROI). The fits are shown in Figure 5.3a. The fits for electrode j provide $\tilde{\alpha}_j$, $\tilde{\beta}_j$, $\tilde{\gamma}_j$ and $\tilde{\delta}_j$ according to

$$\tilde{\phi}_j(z_s) = \tilde{\alpha}_j z_s^2 + \tilde{\beta}_j z_s^4 + \tilde{\gamma}_j z_s + \tilde{\delta}_j + \tilde{g}_j z_s^3 \tag{5.18}$$

for $j = \{4, 5, 6, 7, 8\}$. For a set of voltages, Equation 5.2 gives $\tilde{V}_{\text{sep}}(z_s) = \sum_{j=4}^{8} v_j \phi_j(z_s)$, where $\tilde{V}_{\text{sep}}(z_s) \simeq V_{\text{sep}}(z_s)$ in the ROI. Relating this to Equation 5.15, a linear relationship holds between electrode voltages and polynomial coefficients of the potential

$$\begin{bmatrix} \alpha \\ \beta \\ \gamma \\ \delta \end{bmatrix} = \begin{bmatrix} \tilde{\alpha}_4 & \dots & \tilde{\alpha}_8 \\ \tilde{\beta}_4 & \dots & \tilde{\beta}_8 \\ \tilde{\gamma}_4 & \dots & \tilde{\gamma}_8 \\ \tilde{\delta}_4 & \dots & \tilde{\delta}_8 \end{bmatrix} \begin{bmatrix} v_4 \\ \vdots \\ v_8 \end{bmatrix} \tag{5.19}$$

which can be inverted to calculate the voltages for a given set of $\{\alpha, \beta, \gamma, \delta\}$. This is not an optimal approach, however, since it does not consider voltage limits or necessarily maximise β. Instead, we minimise a discrete cost function

$$C = \sum_{j=4}^{8} \left[q_1 \left(\alpha - \tilde{\alpha}_j v_j\right)^2 + q_2 \left(\gamma - \tilde{\gamma}_j v_j\right)^2 + q_3 \left(\delta - \tilde{\delta}_j v_j\right)^2 \right] - \sum_{j=4}^{8} \tilde{\beta}_j v_j \tag{5.20}$$

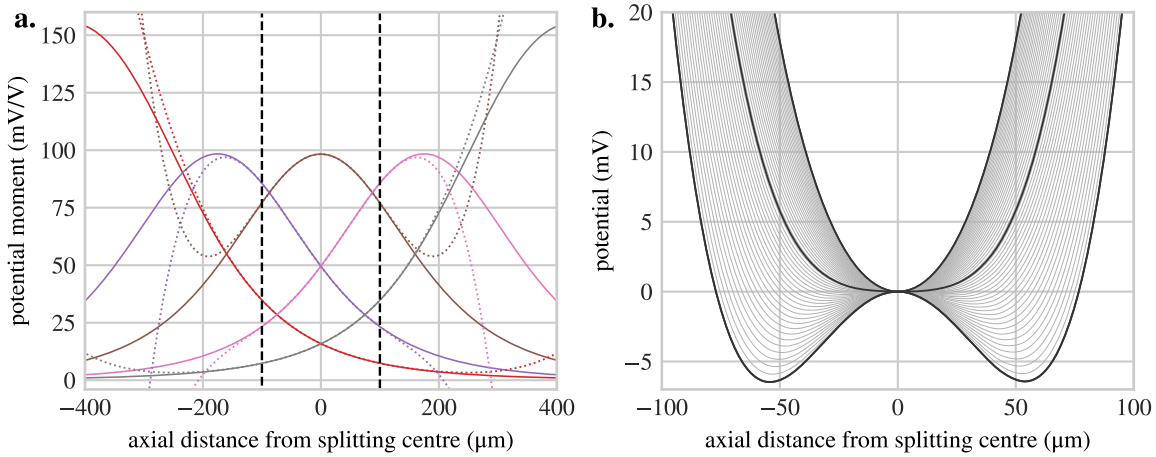

Figure 5.3: Separation electrode model and potential. **a.** Polynomial fits to Electrodes 4–8 within the ROI marked by vertical dashed lines. Electrodes 5 and 7 mirror each other around z_s, as do 4 and 8. **b.** Potential during successive stages of splitting from $\alpha = 5 \times 10^6$ to $\alpha = -5 \times 10^6$, the same range as used for Figure 5.4. The dc offsets have been subtracted for clarity. Thick lines show, from top to bottom, the potentials for $(\alpha, \beta) = (5 \times 10^6, 5.64 \times 10^{14})$, $(0, 6.33 \times 10^{14})$ and $(-5 \times 10^6, 7.02 \times 10^{14})$ (units of V/m^2, V/m^4). The well parameters can be seen in Figure 5.4.

where q_1, q_2 and q_3 specify how closely to meet the desired α, γ and δ values compared to maximising β. All terms are linear or quadratic in v_j, and can be handled by the same solver used for transport. The solution is subject to the constraint $|v_j| \leq v_{\max}$, to keep within the DEATH output limits.

Trajectory of α

The next question is how to choose the trajectory of $\alpha(t)$ between the initial and final times t_0 and t_f. For a single well with $\beta \approx 0$, we can equate Equation 5.15 and Equation 5.4 to obtain $\alpha = m\omega^2/2e$. When the splitting is begun, the initial value α_0 is chosen so that the well frequency ω is around halfway between the usual transport well frequency and the minimum frequency reached during splitting; currently $\alpha_0 = 5 \times 10^6\,\mathrm{V/m^2}$ corresponding to $\omega = 2\pi \times 780\,\mathrm{kHz}$ for calcium, although this choice is arbitrary. The final value α_f is chosen to be the negative of α_0; currently $\alpha_f = -5 \times 10^6\,\mathrm{V/m^2}$. Note that the time taken between α_0 and α_f is referred to as the *splitting speed* later in this chapter. In between, the approach taken is similar to that carried out in the Mainz ion trapping group [73, 172]. As with transport, a primary goal of separation is to minimise the motional excitation of the ions, which arises from heating, uncontrolled stray fields, and impulsive acceleration that occurs near the critical point when α is close to 0 and the confinement is weak. Heating is reduced by keeping β high, as discussed earlier. A simple approach would be to simply ramp α linearly from α_0 to α_f, however this causes large fluctuations in the ion acceleration, as shown in Figure 5.4a.

Instead, the waveform voltages are re-parameterised in terms of the ion separation distance s instead of α. First the solver is run over a set of desired α values sampling the range between α_0 and α_f, with denser sampling around $\alpha = 0$, to produce a discrete set of voltages $\vec{v}[k]$ (where k is a discrete index). 50–100 values is more than sufficient. The solver also returns the attained $\alpha[k]$[7] and $\beta[k]$ values, which are sufficient to determine $s[k]$

[7] Attaining the desired α is only part of Equation 5.20 and there is a tradeoff with the other terms,

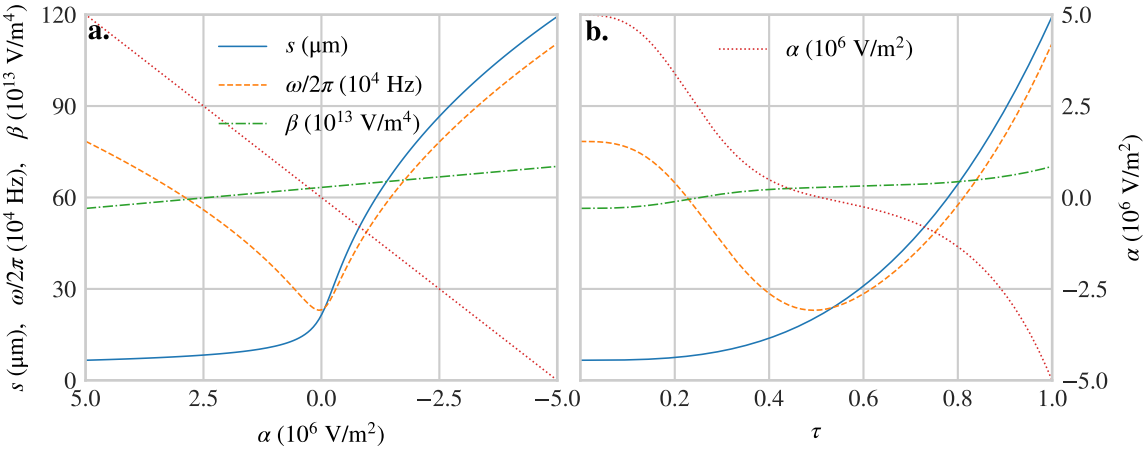

Figure 5.4: Ion separation distance s, motional frequency ω (Ca^+ is assumed), and quadratic and quartic terms α and β during a 300-step splitting waveform as a function of **a.** α and **b.** τ. The y axes and legends are shared, with s, ω and β referenced to the left axis and α to the right axis. **a.** A linear ramp of α from 5×10^6 to -5×10^6 V/μm^2; note that the x axis is inverted. The motional frequency dips to 230 kHz at the critical point where $\alpha \simeq 0$, and the ions sharply move apart in the weakly confining well as α decreases further. **b.** The waveform is re-parameterised in terms of separation s, according to $s(\tau) = s_0 + \tau^3(s_f - s_0)$. This slows down the separation around the critical point and lowers the ion acceleration. Throughout the splitting, β is held almost constant by the solver.

and $\omega[k]$ from Equation 5.16 and Equation 5.17. $s[k]$ is monotonic with respect to $\alpha[k]$, which is necessary for the final step. The voltages are re-parameterised using fourth-order spline interpolation with $(s[k], \vec{v}[k])$ treated as (x, y) coordinate sets, numerically creating the function $\vec{v}(s)$. Between the initial and final separations s_0 and s_f, a splitting trajectory $s(t)$ can now be chosen, and $\vec{v}(t)$ calculated directly.

It is convenient to use a normalised time $\tau = (t - t_0)/(t_f - t_0)$. A number of $s(t)$ choices are suitable for low-excitation splitting [73]; power-law functions such as

$$s(t) = s_0 + \tau^\epsilon (s_f - s_0) \tag{5.21}$$

were found to perform well with $2 \leq \epsilon \leq 3.5$. Currently $\epsilon = 3$ is used in the mixed-species setup, and the resulting trajectories are shown in Figure 5.4b. There is a tradeoff between minimising time spent in a shallow well, during which anomalous heating and external noise will excite the ions, and minimising ion acceleration; this was not investigated in detail during this thesis. In this work, 300 evenly-spaced τ values were used to generate a voltage matrix $\boldsymbol{v}_{\text{split}}$. The usual DAC output rate was 5 MSPS (200 ns sampling period), carrying out $\boldsymbol{v}_{\text{split}}$ in 60 μs.

Trajectories of γ and δ

The parameters γ and δ in Equation 5.15 specify the potential gradient and dc offset at $z_s = 0$; γ effectively specifies an axial electric field. Although ideally $\gamma = 0$ for symmetric separation, in practice the deviations of the true potential moments from the model, as well as fluctuating axial stray fields, require a correction field to be applied during the

especially β, which causes roughly constant discrepancies between the desired and attained α of $\sim 10^4$ V/m^2, or 1 % of α_0. Increasing q_1 reduces this, however numerical instability can arise when $q_1 \geq 1000$.

critical point to minimise motional excitation. Currently γ is set to a constant value γ_{comp} for the entire splitting waveform. In earlier tests of separation there was no method to alter waveform voltages during experiments, thus many waveforms were generated between $-300 \leq \gamma_{\text{comp}} \leq 300\,\text{V/m}$ to find a working range, which was then probed using a finer γ_{comp} spacing. Altering individual waveform voltages was implemented soon afterwards, and the axial field is currently adjusted by differentially raising and lowering the left (Electrodes 5 and 20) and right (Electrodes 7 and 22) voltages, discussed later in §5.5.2, and γ_{comp} is used only for coarse adjustments of $> 10\,\text{V/m}$. From the $\tilde{\gamma}_j$ values in Equation 5.18, i.e. the linear terms of the polynomial fits at the splitting centre, the axial field moment magnitude is $\sim 343\,(\text{V/m})/\text{V}$ for Electrodes 5, 7, 20 and 22, resulting in a net field of $1372.8\,\text{V/m}$ when a unit voltage is manually applied.

δ specifies a dc voltage for the splitting waveform. It was not constrained during the experiments in this thesis (the cost term q_3 in Equation 5.20 was set to 0), though it may be useful for mixed-species separation where the radial mode directions of the crystal need to be controlled.

5.4.3 Auxiliary transport operations for separation

Having created the splitting waveform, the transport solver is used to produce a separate waveform bringing the ions from the experimental zone to the separation location, carrying out the first stage of separation discussed earlier. The starting axial potential is calculated from the splitting voltages $\vec{v}(t_0)$, and a localised parabola is fitted to the well at the splitting site to estimate the precise position p, frequency ω and dc offset d. These are used as end points for the transport solver, with the starting points usually simply the default experimental zone well. Because the splitting solver maximises β, the voltages differ significantly between its beginning and the end of the transport, and to avoid discontinuities linear interpolation is used to smooth the transition. In waveform sets this transport waveform is labelled 'trans from start $\rightarrow$ split start'. It is labelled a in Figure 5.5.

The end of the splitting waveform is similarly spliced with a two-well transport waveform using linear interpolation, however the initial values of the two wells and their rates of change are found by extrapolating the derivatives of the well parameters in the final timesteps of the splitting; the timesteps required for interpolation are taken into account. $ds(t)/dt|_{t=t_f}$ in Equation 5.21 determines the initial well velocities of this waveform, thus higher exponents ϵ in τ^ϵ can allow the ions to be transported apart more quickly relative to the splitting speed, however this has the potential to cause undesired additional excitation. The end of the transport usually places the two wells in the centres of the experimental/storage electrodes around the separation zone, at $z = -845\,\mu\text{m}$ and $z = 0\,\mu\text{m}$. The presence of ions in the rightmost well in the trap centre is then verified, and cooling or temperature measurement are carried out. The combined splitting and transport separation waveform is labelled 'split apart'.

To complete the separation toolbox, a third waveform moves the left well from -845 µm to the trap centre, simultaneously moving the right well to +845 µm. This allows the operations carried out on the right well to be repeated on the left well. The transport requires no particular constraints, and the ordinary transport routines are used. It is labelled '-far to centre, centre to +far'. Figure 5.2b shows a waveform of this type, and the potentials for the three waveforms are shown in Figure 5.5. The optimised parameters in current use are listed in Table 5.2.

term	meaning	notes	default value	transport value	reasonable range
V_R	potential difference in harmonic well to include in the ROI	Higher → better ω agreement, lower → better p, d agreement but shallower ω results if requested ω is high.	4 mV	2 mV	$1 - 20$ mV
r_0	default electrode voltage weight	Higher → less electrode voltage fluctuation, lower → solution may improve if there are multiple competing cost terms.	10^{-15}	5×10^{-10}	10^{-15}–10^{-5}
v_{def}	default electrode voltage	Higher → deeper potential wells, however p, ω and d agreement worsens.	5 V	5 V	0 – 8 V
r_1	voltage derivative weight	Higher → smoother waveforms, p, ω and d agreement worsens during transport (start/end agreement is unaffected).	10^{-5}	5×10^{-4}	$0 - 10^{-3}$
r_2	voltage second derivative weight	Similar to r_1, more susceptible to numerical noise in $\phi_j(z)$.	0	2×10^{-6}	$0 - 10^{-5}$

Table 5.1: Solver coefficients used in ion transport waveform generation. The positive effects of using higher values must be weighed against how well the resultant waveform will agree with the $\{p(t), \omega(t), d(t)\}$ specifications.

term	notes	default value	reasonable range
α_0	Higher → separation solver begins at a higher ω, may require more interpolation between a and b_2 sections	5×10^6 V/m^2	2×10^6 to 2×10^7 V/m^2
α_f	Lower → separation solver ends at higher ω and s, requires more interpolation between b_2 and b_4 sections	-7×10^6 V/m^2	-2 $\times 10^6$ to -2 $\times 10^7$ V/m^2
ϵ in Equation 5.21	Higher → fewer timesteps in section b_4 required, possibly higher excitation	3	1.5 to 4
b_1 and b_3 timesteps	Higher → smoother interpolation waveform segment, less excitation, slower waveform	100	20 – 200
b_2 timesteps	Higher → slower splitting (not necessarily lower excitation), slower waveform	300	100 – 500
c timesteps	Higher → slower transport (not necessarily lower excitation), slower waveform	1000	500 – 2500
DEATH sampling rate	Higher → slower transport and separation (not necessarily lower excitation, due to noise), slower waveforms	200 ns, Be^+ 500 ns, Ca^+	10 ns – 2 µs

Table 5.2: Constants used in ion separation waveform generation.

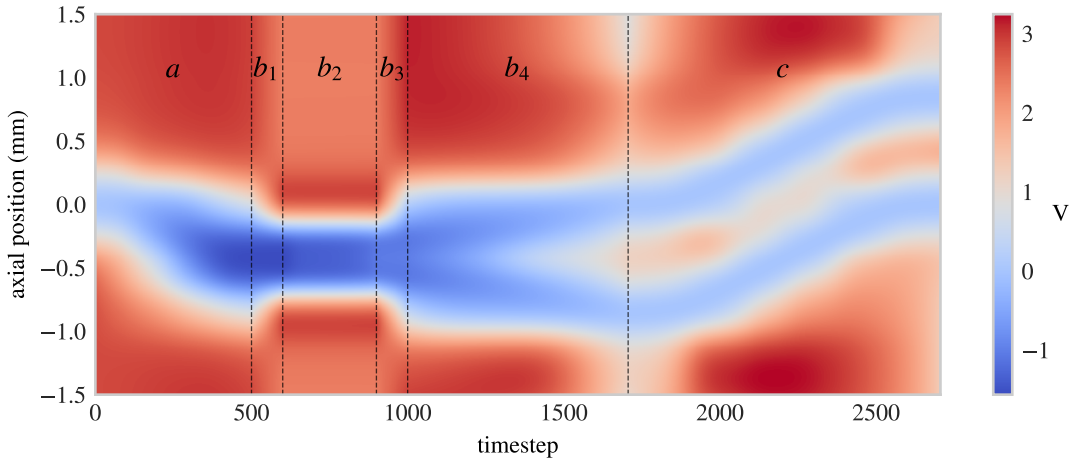

Figure 5.5: Potentials for the set of waveforms used for separation. Waveform a transports the ions to the separation zone. Waveform b separates them into two wells in the experimental and left storage zones. It has several distinct sections as described in the text; b_2 is dedicated to low-excitation splitting and calculated using the polynomial solver whereas b_4 uses the regular transport solver, with regions b_1 and b_3 interpolating between the solvers. Waveform c shifts the wells to the experimental and right storage zones.

5.5 Experimental optimisation

The steps used in generating transport and separation waveforms in the mixed-species setup have been described. Next the procedures that were used to optimise them are discussed.

5.5.1 Transport optimisation

The transport waveform generation required substantial parameter tuning but no fundamental alterations from the approach of Ludwig de Clercq and Robin Oswald [37, 144]. After the solver described earlier in this chapter was implemented with smoothed position trajectories, the parameters listed in Table 5.1 were empirically tuned by estimating the motional excitation of a single ion. This was done by cooling to the ground state, applying a transport waveform, re-applying it in reverse, then driving a motion-adding sideband[8] (MAS) oscillation to estimate the populations in each n level as described in §4.5. Alternative techniques relying on higher-order sidebands are also possible for faster data acquisition [199].

The motional excitation for transport waveforms whose axial frequency remained above $\sim 1\,\mathrm{MHz}$ was less than 0.2 quanta, for mean velocities both below and above $4\,\mathrm{m/s}$. Heating of several quanta was seen for early waveforms with velocities below $4\,\mathrm{m/s}$ whose motional frequencies fell below $500\,\mathrm{kHz}$, however, such as for waveform a in Figure 5.5. The experiments were carried out using trap electrode filters with a cutoff of $260\,\mathrm{kHz}$, thus it may have been related to external noise. The dc offset also varied from $+1\,\mathrm{V}$ to $1\,\mathrm{V}$ during these waveforms, which may have contributed. For waveforms with average velocities above $4\,\mathrm{m/s}$ this effect was weaker, and it was not seen at all for beryllium ions, whose axial frequencies were $\sim 2\times$ higher for the same potentials. The filter cutoff frequency has since been lowered to $73\,\mathrm{kHz}$, and at the time of writing the effects on transport and separation are

[8] A blue sideband for Ca^+, a red sideband for the Be^+ FIQ.

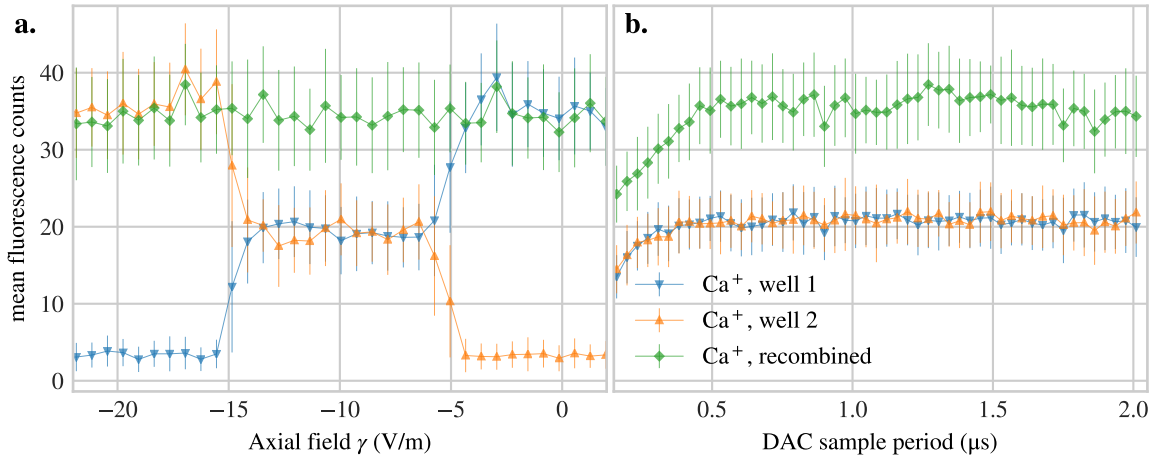

Figure 5.6: Separation of a calcium-calcium crystal. **a.** Mean PMT counts in Well 1 and 2 (right and left wells respectively), with the region between $-15 \leq \delta \leq -5\,\mathrm{V/m}$ showing one ion in each well, at a DAC sampling rate of 500 ns (2 MSPS, total of 150 µs for $5 \times 10^6 \rightarrow \alpha \rightarrow -5 \times 10^6$). Left (right) of this region, the axial field drives both ions into Well 2 (Well 1). The recombined counts show the fluorescence for a repeated detection of the ions in a single well. **b.** Fluorescence as a function of DAC sampling rate. Below 200 ns (5 MSPS), the ions are strongly excited and the fluorescence falls. The scan was recorded on a different day when the fluorescence was higher than in **a**, and the traces have been scaled by 0.64 for better comparison with **a**. For sampling rates above 1.5 µs, the recombined fluorescence begins to fall slightly due to the slowness of the waveform. Error bars are the standard deviations of the counts.

being investigated[9].

As mentioned earlier, for transport speeds above a few tens of m/s diabatic effects contribute significantly to the excitation and must be compensated by adjusting the transport distance or waveform timing [23, 199]. This is not straightforward with the DEATHs due to the pre-computed waveforms and the quantised sampling rates, however it could nonetheless be investigated by loading multiple incrementally-varying waveforms. Diabatic transport is inherently less robust to parameter fluctuations, however, and is unlikely to be a high priority in the mixed-species setup until separation is fully optimised and transport delays become a limiting factor in experiments.

5.5.2 Separation optimisation

Optimising the separation required several further steps. For a given waveform sequence, the ions were first Doppler or sideband cooled. Next the separation waveforms were applied (waveforms a and b in Figure 5.5), splitting the ions and bringing those in Well 1 to the experimental zone. Here they were optionally Doppler cooled or probed with a MAS pulse, then detected. The parallel-transport waveform c was applied next, and the same operations were repeated on the ions in Well 2. Finally the wells were recombined by running c, b and a in reverse, allowing the combined ion crystal to be cooled or probed then detected.

Using this sequence, the axial field γ was scanned over several tens of V/m to identify a range where the ions would separate into two wells. Initially the Doppler cooling was necessary for un-optimised waveforms, which would cause such severe heating after sep-

[9] While a careful study has not been performed, experiments using transport seem to perform at least as well as before the filter change.

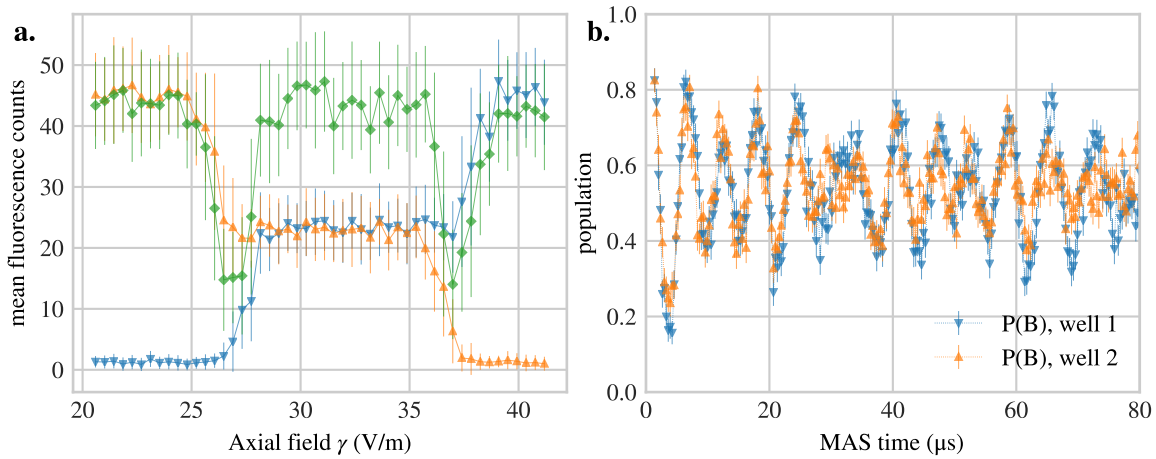

Figure 5.7: Separation of a beryllium-beryllium crystal. **a.** Similar scan to Figure 5.6a, at a DAC sampling rate of 1 µs (1 MSPS). Triangles represent the fluorescence in each well, and diamonds represent the fluorescence after recombination. **b.** Ion populations as a function of MAS (blue sideband) probe time, corresponding to $\bar{n} \sim 7$.

aration that the ions would be lost. Several milliseconds of cooling were required before ion fluorescence was recovered. Once this basic routine was operational, the parameters in Table 5.2 were iteratively altered until the cooling time necessary to regain fluorescence could be reduced or disabled entirely without losing ions.

The next stage relied on optimising the fluorescence as a function of γ and waveform speed, with the aim of minimising its sensitivity to both. A well-optimised waveform shows constant fluorescence over a range of waveform speeds, and clear 'steps' for a scan of γ; examples are shown in Figure 5.6a and Figure 5.7a. Characteristics such as non-uniform steps, gradients in the fluorescence, asymmetry and other artifacts identify sub-optimal waveforms.

Final optimisation involved probing the MAS for the single ions in each well, with the temperature being used for finer adjustment of γ and the waveform properties in Table 5.2. The splitting dc offset δ can also be optimised[10], however this has not yet been systematically carried out.

Several approaches to the separation solver were tried before the final version discussed in §5.4 was implemented. An initial approach where the voltages were manually chosen and interpolated performed poorly. An approach using the solver cost function Equation 5.20 where the α profile was smoothly ramped but not parameterised in terms of s was robust, however it swept rapidly through the critical point. With additional filtering[11] the fluorescence remained high and the waveforms were usable for the beryllium separation experiments discussed in §4.4 and §6.1.2. There was a more severe loss in fluorescence for calcium due to its lower axial frequencies, however. For two beryllium ions initialised in the ground state, this solver approach gave an excitation of approximately 7 quanta for each ion, with a MAS scan shown in Figure 5.7b.

This approach caused too much heating for mixed-species separation, prompting the re-parameterisation of the solver in terms of separation distance. This greatly improved

[10] Thomas Ruster, private communication, 2016.
[11] Savitzky-Golay filtering, of order 2 and a window of 101 timesteps.

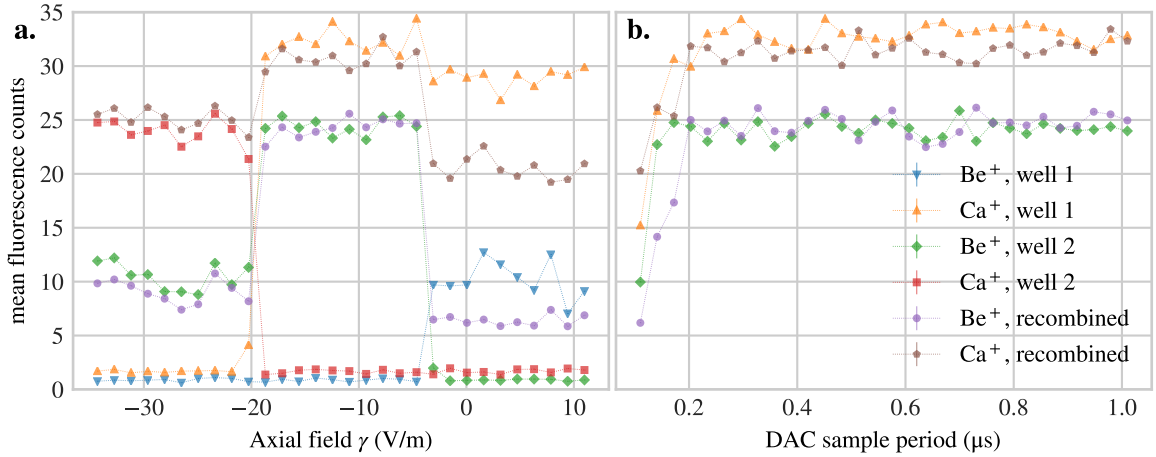

Figure 5.8: Separation of a beryllium-calcium crystal. **a.** The mean PMT counts in Well 1 and Well 2 (right and left wells respectively), with the region between $-19 \leq \delta \leq -5\,\mathrm{V/m}$ showing Ca^+ in Well 1 and Be^+ in Well 2, at a DAC sampling rate of 200 ns (5 MSPS, total of 60 µs for $5 \times 10^6 \to \alpha \to -5 \times 10^6$). Left (right) of this region, the axial field drives both ions into Well 2 (Well 1), however the Be^+ heats up strongly and its counts are more than halved. The recombined counts show the fluorescence for a repeated detection of the ions in a single well; $-19 \leq \delta \leq -5\,\mathrm{V/m}$ shows optimal separation and recombination, resulting in fluorescence comparable to the single-well levels. Error bars are not displayed for clarity. **b.** Fluorescence as a function of DAC sampling rate. Below 200 ns (5 MSPS), the ions are both strongly excited and the fluorescence falls sharply.

separation of a beryllium and calcium crystal, shown in Figure 5.8. The separate PMTs allowed the ion order to be directly distinguished. Initially, the ions frequently swapped positions due to heating causing them to decrystallise and recrystallise in a random order, and a waveform to obtain a deterministic order was applied before every shot to compensate (see §5.6.1). This was reduced to once per experimental point after the waveform was further optimised. The solver re-parameterisation also allowed separation of longer single-species ion chains, as shown in Figure 5.9. Although these crystals were not used in experiments, the current separation approach is robust enough for future single-species use on longer ion chains.

Separation was also investigated for beryllium-calcium-beryllium crystals. Manual shots were successful, however an axial field for automated deterministic separation could not be found. In early experiments some structure similar to Figure 5.8a was visible in a similar scan, however the ion fluorescence levels were less than half of their stationary values and the calcium ion was lost every few thousand shots despite intermediate Doppler cooling. It is suspected that an unknown heating source affects this crystal more strongly than the others discussed earlier; one possibility is a nonuniform radial potential causing the crystal to deform.

When a radial offset was applied using a differential voltage on Electrodes 6 and 21, calcium could be deterministically split into one well with the beryllium ions in the other; the fluorescence after separation was also comparable to normal levels. This could be explained by a slight misalignment between the top and bottom trap wafers leading to a nonuniform axial potential as a function of radial displacement. Similar to the beryllium-calcium reordering discussed below in §5.6.1, the different radial displacements of calcium and beryllium lead them to encounter different axial fields. Such effects may be corrected by experimentally characterising the electrode moments rather than relying exclusively on

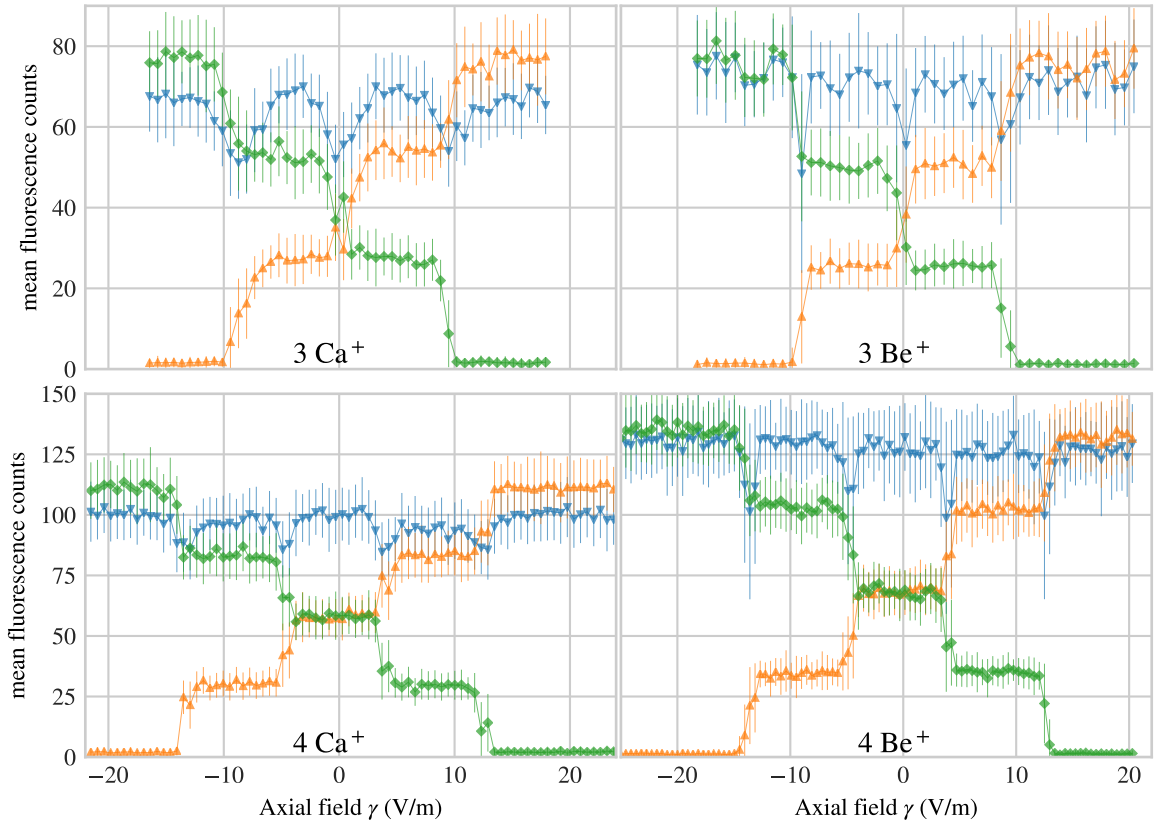

Figure 5.9: Separation of multi-ion single-species crystals. The top row shows three-ion Ca^+ and Be^+ separation respectively, and the bottom shows four-ion separation. The axial field offset has been subtracted from the x axes to facilitate comparison. Green diamonds and orange triangles represent the ion fluorescence in each well, blue triangles represent the fluorescence after recombination.

the simulation model, as was carried out in previous low-excitation single-species separation work [73, 172].

5.5.3 Future steps

Transport is currently reliable and well-controlled at moderate velocities. It is regularly used in loading, the beam profiling experiments discussed in §5.6.3 as well as the operations in §7.4.2, and to speed it up diabatic effects will have to be considered. This is not an immediate priority as transport delays do not yet limit experimental duty cycles, especially compared to separation.

The separation performance is approaching the level of [23] and it is reliably used in everyday operations such as beam alignment, ac Stark shift calibration and MS gate optimisation, as well as separating the ions to remove hydrides. It is not yet as robust as transport, however, largely because the operations involved are more sensitive to noise and parameter variations, especially in the electrode moment models. The immediate goal is to optimise separation for beryllium-calcium-beryllium crystals. The performance of the existing waveforms must first be optimised further using sideband spectroscopy, however this may be insufficicent: previous work has required experimental adjustment of the electrode moment models to minimise excitation [172]. Many issues, such as strong stray fields

or modelling errors, can probably be diagnosed using beryllium-calcium splitting, which is simpler to optimise because ions already are not lost.

As the separation approaches diabatic timescales, the distortion of the trap electrode filters grows more significant. This was modelled early in the separation experiments, found to have a negligible impact, and removed to improve computation times; it may need to be re-introduced since the new filters have a lower cutoff frequency that may already be affecting the waveforms when using 200 ns sampling rates.

Currently γ is set to a constant value for the whole splitting routine, which is not ideal because each γ_{comp} requires a slightly different transport waveform before and after the splitting operation, taking extra time to generate. The trajectory of $\gamma(t)$ could easily be ramped from 0 at t_0 to γ_{comp} at the critical point, then back to 0 at t_f, ensuring that the start and end voltages remain unchanged by γ. This is not yet necessary, but may become useful if the waveform calculation time becomes the limiting factor in experiments or scripts. Modifications to the DEATH embedded software are being tested that will allow real-time weighted sums of waveforms to be calculated on the DEATH CPUs, such that a waveform containing a $\gamma(t)$ ramp could be added to a standard separation waveform, with its weight in the sum altered between experimental points as γ is currently. This will also allow separate γ values in the separation and recombination, which will probably be necessary once diabatic speeds are used.

Beyond improved modelling and stray field compensation, the separation approach itself can be further optimised. Optimal trajectories for $\alpha(t)$ and $\beta(t)$ have been calculated using a 'shortcut-to-adiabaticity' approach in the Muga group [147], which could easily be applied to produce a separation waveform. A notable difference in these trajectories is that unlike in this work, $\beta(t)$ is not maximised at all times[12] , but only near $\alpha = 0$. This could be directly implemented using the current cost function in Equation 5.20 by simply altering the final $\tilde{\beta}_j$ cost term to $\sum_{j=4}^{8} q_4 \left(\beta - \tilde{\beta}_j v_j\right)^2$, and using the optimally calculated trajectories for $\alpha(t)$ and $\beta(t)$.

5.6 Auxiliary operations

This concludes the transport and separation discussion. A mixed-species experiment also requires several auxiliary operations with their own specialised waveforms, described here.

5.6.1 Mixed-species crystal reordering

Maintaining a deterministic ion crystal order is required for mixed-species experiments. For beryllium-calcium crystals it is carried out in the central experimental zone using the six surrounding electrodes [68]. First a radial displacement is applied by increasing the voltages on Electrodes 7, 8, 9 by V_{rad} on the top wafer, shifting both ions off the trap axis towards the bottom wafer, with calcium displaced further than beryllium due to its weaker radial confinement (see Equation 2.6). At this point a tight axial confinement is helpful, to maximise the radial separation of the ions. Next a 'twist' is applied to the dc potential by increasing the voltages of Electrodes 7 and 24 by V_{tw}, and decreasing those of 9 and 22. This axially displaces both ions leftwards, with the calcium ion shifted further since it is closer to Electrodes 22 and 24 and encounters a stronger net axial field. The radial

[12] The units of β used in some references are N/m^3, whereas in other references and this thesis they are V/m^4; they differ by a factor of e, with $10^{-4}\,\text{N/m}^3$ corresponding to $6.2 \times 10^{14}\,\text{V/m}^4$.

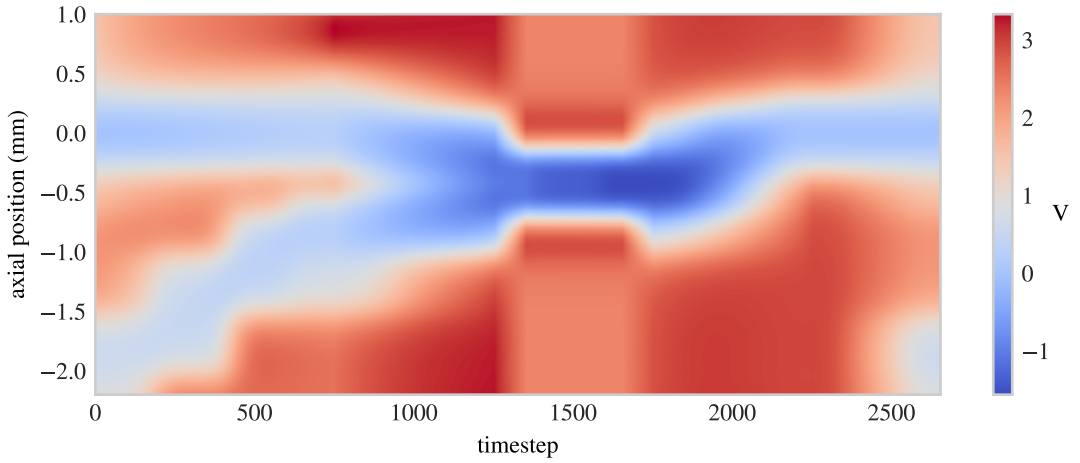

Figure 5.10: Cyclical 'loading conveyor' waveform potential. Two wells in the loading zone at -1.870 mm and experimental zone at 0 mm are merged into a single well in the experimental zone. The loading zone well is then recreated; the potentials at the initial and final timestep are identical. The loading zone uses a shallow 600 kHz well (Ca^+).

displacement is removed, returning the ions to the axis with calcium on the left. Finally the twist is removed and the well returned to the normal axial frequency. The current values are $V_{\text{rad}} = 0.8\,\text{V}$ and $V_{\text{tw}} = 0.9\,\text{V}$, calibrated by distinguishing the ion order from the PMT counts (the different axial positions cause PMT counts to differ slightly) and identifying the minimum voltages to set it deterministically.

For beryllium-calcium-beryllium crystals as used in Chapter 7, the twist is not required since the crystal forms an isosceles triangle when displaced far enough radially, with the calcium ion forming the outer apex. Once the displacement is relaxed, the triangle flattens into a chain with calcium in the centre. Here $V_{\text{rad}} = 2.0\,\text{V}$ is currently used. Reordering can be carried out on larger mixed-species crystals by generalising these techniques [68]

5.6.2 Loading conveyor waveform

As discussed in §4.2 the ions are loaded in the dedicated left loading zone, with a shallow well located at -1870 µm. Currently the PMT counts in the experimental zone are used to tell the number of ions loaded; the imaging cameras of the mixed-species setup are usually used only to optimise beam pointing in the loading zone. There are two loading techniques; the first is to keep a well in the loading zone, and periodically transport it to the experimental zone to check for the presence of ions. This has the drawbacks that while the well is in the experimental zone no ions can be loaded, and that a user must manually run the transport waveform forwards and backwards. It was also found that once an ion was already present in the loading zone, it seemed to affect the likelihood of loading single additional ions versus pairs.

Instead a single 'loading conveyor' waveform was calculated that began with simultaneous wells in the loading and experimental zones, and when run would merge them together, transport the combined well to the experimental zone, and recreate the loading well. This cyclic 'pump' routine, shown in Figure 5.10, allowed simultaneous loading and inspection of the experimental zone fluorescence, and significantly eased the loading process[13]. The mer-

[13]Especially for a single person simultaneously managing the photoionisation lasers. Currently the wave-

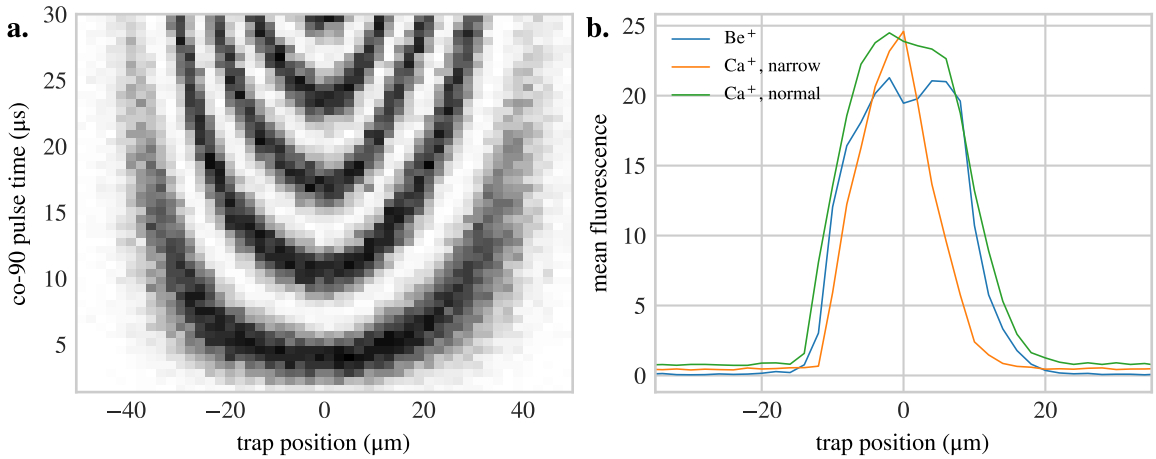

Figure 5.11: Beam profiling. **a.** Co-90 Raman beam profile, where darker regions indicate lower ion fluorescence. The beam full-width-half-maximum is around 45 µm. It has since been more tightly focused to ∼ 20 µm. **b.** Beryllium and calcium PMT aperture beam profiles for wide and narrow calcium aperture settings. For a narrow aperture the background counts are lower, however the optimal ion position is restricted.

ging of distant wells also helped 'scoop up' ions trapped in local potential minima between the loading and experimental zones due to heating or un-optimised separation.

5.6.3 Beam profiling

The position and shape of laser beams directed on the ions, as well as the PMT apertures, can be profiled to sub-micron resolution by transporting a single ion. For a qubit beam in calcium or beryllium, the experiment consists of observing how the Rabi frequency scales as a function of axial position. Using a waveform with an array of target wells, in every shot the ion is transported to a position, the beam is pulsed, then the ion is returned to the centre for detection. A 2D map of Rabi oscillations against position is produced, from which the centre, width and symmetry of the beam can be estimated. Figure 5.11a shows the Raman co-90 beam profile after optimisation.

To optimise the detection beams and PMT apertures, the detection beam is pulsed and photons are collected, before the ion is detected a second time back in the centre; this way the two levels of fluorescence can be observed simultaneously and the optimal aperture positions can be set. Narrowing the PMT apertures without sacrificing ion fluorescence is required to minimise background photon counts. This is shown in Figure 5.11b. A similar technique is being implemented to automatically identify the ordering of a three-ion mixed-species crystal for separation experiments.

5.6.4 Recrystallisation

In experiments using beryllium or mixed-species crystals, stray collisions with background gas or transient electronic noise cause the ions to heat up and lose fluorescence, at a rate of roughly one event per 30 s per beryllium ion. When this occurs the crystal must be 'recrystallised', which is carried out by 'transporting' using a waveform whose well frequency is lowered to ∼ 400 kHz (for calcium) at a constant position and a dc offset chosen to optimise

form is run manually with a mouse click once every few seconds, however this could be trivially automated.

cooling in the shallow well. The system carries out recrystallisation automatically when the fluorescence counts measured during Doppler cooling of beryllium fall below a threshold for more than several shots in a row. Once the shallow well is applied, successive cycles of far-detuned and near-detuned Doppler cooling are carried out until the fluorescence returns, or an adjustable timeout of several seconds is reached. In the former case, the waveform is run in reverse to return to the original well, and the experimental point is retaken; in the latter the shallow well is maintained and the user is alerted.

At frequencies below 400 kHz, a single ion seemed to be displaced by several microns axially and radially compared to frequencies above 1 MHz. This was visible through a drop in fluorescence even when micromotion compensation had been carried out recently, and extra static axial and radial fields similar to those used for reordering were required to compensate it. Apart from the mixed-species separation behaviour this may be another sign of asymmetric electrode potential moments.

6 Quantum operations and protocols

This chapter presents the key single- and multi-qubit operations and protocols carried out in the mixed-species setup to characterise spin and motional decoherence and Stark shifts, and the real-time adaptive Bayesian calibration of the phase of a single qubit. Multi-qubit entangled states and the Mølmer-Sørensen gate are introduced, and gate calibration techniques for single-species gates are presented. Two multi-qubit protocols are demonstrated, making use of the properties of entangled states and multi-qubit gates, including a novel dissipative entanglement scheme. The chapter concludes with a discussion of multi-species two- and three-qubit gates, their calibration, and potential contributions to the current gate infidelity in the mixed-species setup.

6.1 Ramsey spectroscopy and related techniques

The Ramsey sequence and its variations are widely-used techniques in quantum metrology and QIP [156, 157, 204, 161]. They are applied in the mixed-species setup for precise laser frequency calibration, spin and motional coherence measurements, and optimising the magnetic field feedforward coefficients. This section first introduces the standard Ramsey sequence and its applications, followed by an adaptive Bayesian phase estimation scheme that uses the Ramsey sequence as a subroutine. Motional coherence measurements are discussed in §6.7.2.

In a standard Ramsey sequence the qubit is initialised in state $|g\rangle$, and a $R_{\pi/2}(\frac{\pi}{2})$ pulse prepares $(|g\rangle + |e\rangle)/\sqrt{2}$. The qubit is exposed to a Hamiltonian

$$H_R = \hbar\delta(t)\sigma_z/2, \tag{6.1}$$

over an 'interrogation time' or 'wait time' τ that causes phase evolution between the states, leading to a state $e^{i\Phi(\tau)/2}|g\rangle + e^{-i\Phi(\tau)/2}|e\rangle$. A second $R_{\pi/2}(\frac{\pi}{2}+\theta)$ pulse maps the relative phase $\Phi(\tau)$ into the final $|g\rangle$ and $|e\rangle$ population, resulting in

$$p_g(\theta) = \sin^2\frac{\theta+\Phi}{2} \quad = \frac{1}{2} - \frac{1}{2}\cos(\theta+\Phi), \quad \text{where} \quad \Phi = \int_0^\tau \delta(t)dt. \tag{6.2}$$

where we have assumed the pulses acted instantaneously, a valid assumption if $\tau \gg t_{\pi/2}$. In the mixed-species setup $t_{\pi/2} < 6\,\mu s$ for all qubits, while in most experiments $100\,\mu s < \tau < 100\,ms$.

6.1.1 Frequency calibration and spin coherence

The sequence is used in several ways, with the most direct being laser frequency calibration. If $\delta(t) = \delta_0$, as is the case for a frequency detuning between the $R_{\pi/2}$ pulse laser and the qubit, the cosine argument is $\theta + \delta_0\tau$. By scanning τ and minimising the oscillation in $p_g(\theta,\tau)$, the laser frequency can be matched to the qubit. The gradient of p_g with respect to δ_0 is $\frac{\tau}{2}\sin(\delta_0\tau+\theta)$, thus longer wait times linearly improve the calibration precision. For

the beryllium FIQ the frequency was calibrated to within 1 Hz using this technique, and has remained stable over a year.

The Ramsey contrast on the other transitions in the ions is worsened due to a fluctuating phase argument $\theta + \Phi$ in Equation 6.2 between experimental shots, which is caused by a varying qubit energy splitting as well as laser frequency/phase noise. The dominant contribution to this dephasing in the mixed-species setup comes from magnetic field noise.

If Φ is sampled from a normalised probability distribution $D(\Phi)$, representing the various slow noise sources contributing to Φ such that the fluctuation in Φ is slow compared to the Ramsey wait time, the Ramsey curve averaged over many shots will approximate the form $D(\Phi)(1 - \cos(\theta + \Phi))/2$. The contrast C is equal to the curve offset from 0.5 at zero phase shift, integrated over possible Φ values, namely[1]

$$C(D) = \int_{-\infty}^{\infty} D(\Phi') \cos(\Phi') d\Phi'. \tag{6.3}$$

We assume that $D(\Phi)$ forms a Gaussian distribution, given by

$$D_G(\Phi) = \frac{1}{\sigma\sqrt{\pi}} e^{-\frac{\Phi^2}{\sigma^2}}, \tag{6.4}$$

where σ is the $1/e$ Gaussian half-width. Thus Equation 6.3 becomes [87]

$$C_G = e^{-\frac{\sigma^2}{4}}. \tag{6.5}$$

This is approximately valid for other symmetric distributions as well. The Ramsey contrast can be used to estimate the decoherence rate γ of the system, using $\sigma = \tau/\gamma$; namely by observing the contrast as a function of wait time and fitting to one of the expressions in Equation 6.5 depending on the assumed noise model, σ can be estimated. In this thesis we are concerned primarily with relative decay rates, thus Equation 6.5 is sufficient [66]. A more sophisticated treatment of Ramsey contrast decay curves can be found in [87].

This approach was used to characterise the decoherence in beryllium and calcium for single ions [81, 105], and is used on two-ion entangled states in §6.6.1. At the time of writing the calcium $1/e$ decay time is 1.32(3) ms with the magnetic field feedforward calibrated, limited by magnetic field noise; this is compared with two-ion entangled states in §6.6.1. The beryllium FDQ is several hundred microseconds, and beryllium FIQ is several seconds [105].

6.1.2 ac Stark shift calibration

A qubit laser pulse causes an effective ac Stark shift ΔE, as discussed in §2.2.3, which is a sum of the shifts from multiple contributing transitions. In the limit of long wait times and short $\pi/2$ pulses, a Ramsey sequence allows precise calibration of the bare qubit frequency ω_0. While the laser is on, however, the qubit frequency is $\omega_0 + \omega_{ss}$, where $\omega_{ss} = \Delta E/\hbar$ is the Stark shift detuning.

There are two ways of handling this systematic shift in the mixed-species setup. For low Stark shifts relative to the pulse Rabi frequency, the laser frequency ω can be tuned directly to ω_0, which is done for beryllium and facilitates automatic calculation of the required rf phases $\phi = \omega(t - t_0)$ in the M-ACTION system for phase-coherent pulses (see §3.3.8). This

[1] Note that Φ is in principle unbounded, thus the integral is taken over $\{-\infty, \infty\}$.

beam	$t_{\pi/2}$ (µs)	t_π (µs)	ω_{ss} ($2\pi\times$kHz)	laser ω set to:	infidelity
Ca 729	1.4	2.6	2 (est.)	$\sim\omega_0$	10^{-4}
Be FDQ co-co	1.9	3.3	-	$\omega_0+\omega_{ss}$	
Be FDQ co-90	-	2.3	-	$\omega_0+\omega_{ss}$	
Be FIQ co-co	6.2	12.3	2.0 (co-com) + 1.1 (co-sw)	ω_0	6×10^{-3}
Be FIQ co-90	2.1	4.2	2.0 (co-com) + 2.2 (90-sw)	ω_0	10^{-3}
Be FIS co-co	-	32.8	-	$\omega_0+\omega_{ss}$	

Table 6.1: Pulse times and ac Stark shifts for calcium and beryllium on 18.01.2018, as well as how the carrier frequency is chosen for each transition. The beryllium FDQ and FIS frequencies are calibrated via Rabi oscillations, yielding $\omega_0+\omega_{ss}$. For the beryllium FIQ, the qubit frequency calibrated via a Ramsey experiment is used. For calcium, a Rabi oscillation at low power is used. Dashes indicate parameters not used in experiments. The calcium Stark shift is estimated from $\Delta\phi$ used in Ca^+ experiments, due to lack of measurement data.

drives the qubit off-resonantly however, and according to Equation 2.21 the qubit population will only be inverted to a depth of $p_{\text{inv}}=\Omega^2/(\Omega^2+\omega_{ss}^2)$ by a π pulse. The deviation of p_{inv} from 1 scales as $1-\omega_{ss}^2/\Omega^2$. Currently for beryllium $\Omega\gg\omega_{ss}$ for the FIQ, leading to infidelities below 1% (see Table 6.1). For calcium the values are similar.

This approach is not as straightforward when the Rabi frequency is increased, since from Equation 2.32, $\omega_{ss}\propto\Omega^2$ for the far-detuned transitions that dominate the calcium Stark shift. This leads to a larger infidelity; currently this is reduced by operating the calcium qubit laser closer to $\omega=\omega_0+\omega_{ss}$ and calibrating the required pulse phase offsets $\Delta\phi=-(\omega-\omega_0)(t-t_0)$ experimentally[2].

If $\omega\neq\omega_0+\omega_{ss}$, each laser pulse of length t_p causes a phase precession of $\phi_p\simeq(\omega_0+\omega_{ss}-\omega)t_p$ on the bare qubit, in the limit where $\Omega\gg\omega_{ss}$. This precession must be known to calculate a-priori the rf phases in longer QIP protocols. This can be measured independently of the ω choice for the beryllium lasers using a variation of the Ramsey sequence known as a Hahn echo or a spin-echo. If in the Ramsey sequence a π pulse is applied at time $\tau/2$, whose phase does not matter, the population becomes

$$p_g(\theta)=\frac{1}{2}-\frac{1}{2}\cos\left(\theta+\Phi_2-\Phi_1\right),\quad \Phi_1=\int_0^{\frac{\tau}{2}}\delta(t)dt,\quad \Phi_2=\int_{\frac{\tau}{2}}^{\tau}\delta(t)dt. \tag{6.6}$$

As long as the three Ramsey pulses are close to resonance, this cancels out the dc offset of $\delta(t)$, which is equal to $\omega-\omega_0$, and allows the quantity $\Phi_2-\Phi_1$ to be measured, which is nonzero if $\delta(t)$ is odd around $\tau/2$. This could be caused by a linear ramp in time, for instance. Forms of the spin-echo sequence with multiple π pulses, sensitive to frequency components in $\delta(t)$ with periods of $(\tau+1)/m$ where m is the number of π pulses, can be used to profile the $\delta(t)$ spectrum [85].

Running a pulse on the co-com, co-switch or 90-switch Raman beam alone, before the spin-echo π pulse, causes a qubit phase shift ϕ_p due to the Stark shift without altering the qubit populations. Comparing the Ramsey curves with and without the pulse allows ω_{ss} for that beam to be calculated; the Stark shift during a co-co or co-90 pulse can be estimated by adding those of its constituent beams. Alternatively, if n discrete $\pi/2$ or π pulses are carried out on a single beam, then $\Phi_2-\Phi_1=n\phi_p$, and the Stark shift phase from a calibrated

[2] A planned third way will involve operating at $\omega=\omega_0+\omega_{ss}$ and automatically calculating the rf phase according to $\phi=\omega_0(t-t_0)$, however this will require changes to the M-ACTION DDS board gateware.

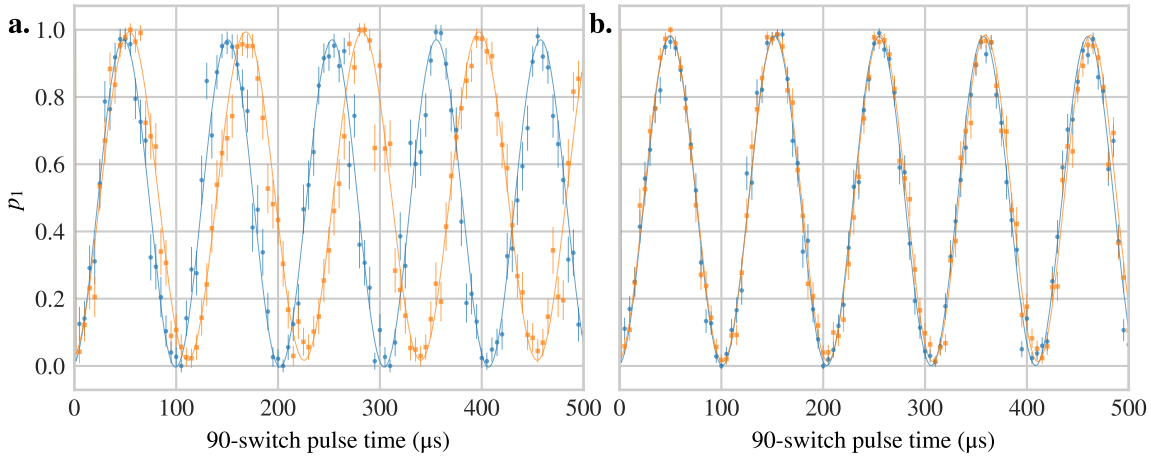

Figure 6.1: 90-switch beam alignment using a spin-echo Ramsey sequence and readout of the ions in separate wells. **a.** Population oscillations in the Ramsey sequence at frequencies $\omega_{ss,1} = 2\pi \times 9.77(2)\,\text{kHz}$ and $\omega_{ss,2} = 2\pi \times 8.74(2)\,\text{kHz}$. **b.** Population after equalisation, at frequencies $\omega_{ss,1} = 2\pi \times 9.75(2)\,\text{kHz}$ and $\omega_{ss,2} = 2\pi \times 9.69(2)\,\text{kHz}$. Because $\omega_{ss} \propto \Omega^2$, this corresponds to a difference in Ω of 0.3 %. Note that recently-measured ω_{ss} values are lower by 3×, it is believed due to beam polarisation changes.

gate can be directly measured, already taking pulse transients such as switching times into account. Note that the phase shift observed in the Ramsey experiment is independent of τ.

The shift in phase due to Stark shift is also used to independently optimise the beryllium 90-switch beam alignment on the ions without requiring a Rabi sequence, which is useful since the ion must be ground-state cooled using the 90-switch beam itself to observe coherent co-90 Rabi oscillations. The sensitivity is also higher than the Rabi technique, since according to Equation 2.33, ω_{ss} scales as Ω^2 for the far-detuned Raman beams. Separating the ions and equalising the differential phase shifts for the FIQ 90-switch in the Ramsey sequence is shown in Figure 6.1. A ω_{ss} agreement of 0.6 % can be obtained between the ions.

Finally, when $\omega = \omega_0$ there is a small fixed Ramsey phase offset due to the Stark shifts of the Ramsey pulses themselves, which cannot be removed without changing the second pulse phase θ, and must be considered when comparing Ramsey phases. Table 6.1 lists recently measured experimental parameters.

6.1.3 Spectral isolation

The decoherence and Stark shift on one ion species caused by detection pulses on the other is another source of infidelity; the Stark shifts can be compensated with a phase shift however the decoherence is difficult to mitigate. Many potential mixed-species QIP protocols, including those discussed in the next chapter, rely on a quantum state being encoded in the beryllium FIQ while the calcium is detected and re-cooled multiple times.

To characterise the effect of the 397 nm beams on the FIQ, a Ramsey sequence was carried out with a wait time of 1 s. Both its contrast and phase were compared with and without a 500 ms pulse on the 397 nm lasers at the typical Doppler-cooling settings, around 10^3 longer than the typical Doppler cooling time. When the lasers were on the contrast altered from 0.51(3) to 0.53(3), consistent with no detectable change. The phase shift on

the FIQ was 11(5) degrees, which implies a negligible shift of 0.01 degrees for a single round of Doppler cooling.

The ac Stark shift on the FIQ from the 729 nm beam was also measured, and the contrast and amplitude shift were both consistent with 0.

The ac Stark shift on the calcium 729 nm transition from the co-90 beams was approximately 3 kHz, with a shift of several degrees for a single co-90 π pulse.

6.2 Bayesian phase estimation

6.2.1 Introduction

The Ramsey experiments described so far involve either estimating the Ramsey sine phase or its contrast, and have been carried out by collecting statistics over uniformly-spaced scans of frequency, phase or wait time, then estimating parameters from the set of aggregate data. However, a more efficient scheme is possible in which each successive experimental measurement outcome is used to update a probabilistic model, or *likelihood function*, describing the state of knowledge of a parameter to be estimated. This is known as a *Bayesian scheme*, and its underpinnings are related to the Bayesian fitting techniques discussed earlier in §4.3.2. Such schemes are widely used in engineering and computer science, especially in robotics, artificial intelligence and machine learning.

A particular benefit of Bayesian estimation over standard measurement techniques is that the likelihood function and its associated conditional update rules can be used not only to incorporate new measurements from the system, but also to predict which next measurement would maximally constrain the model, which is equivalent to maximising the information gain for each measurement. Adaptive *Bayesian experimental design* can be carried out, in which the scheme acts in two stages: it incorporates a measurement result, then predicts the parameters that would be informationally optimal for the next round of measurement.

A Bayesian scheme was developed by Andrey Lebedev [90] that allows us to estimate the phase ϕ in Equation 6.2 more rapidly than with a fit. It consists of a likelihood function, an estimator for ϕ, a rule to update the likelihood based on the outcome of an experimental shot, and a method to choose the optimal Ramsey pulse phase θ for the next shot. This scheme is computationally intensive compared to typical real-time QIP protocols, nonetheless it can be run in real-time on the M-ACTION CPU in the current shot time of several milliseconds, modifying the Ramsey pulse phase not between experimental points but shot-to-shot. The Lebedev scheme is presented in the next section.

6.2.2 Phase estimation protocol

We begin with Bayes' theorem; the posterior probability distribution $p(a|b)$ for a parameter a given an observation b is

$$p(a|b) = \frac{p(b|a)p(a)}{p(b)}, \tag{6.7}$$

where $p(a)$ is the prior probability distribution for a, representing our knowledge of the system before the measurement. In our case we have a Ramsey experiment, that from Equation 6.2 yields a bright or dark outcome, which we label as $\sigma = 1$ and $\sigma = -1$ with

probability

$$p(\sigma, \theta, \phi) = \frac{1}{2} + \frac{\sigma}{2}\cos(\theta + \phi), \tag{6.8}$$

where ϕ is the qubit phase we wish to estimate and θ is the rotation angle of the second $\pi/2$ pulse that we can experimentally adjust. The first step is constructing a likelihood function; we model the knowledge we have of ϕ at an estimation step s using

$$L_s(\phi) = \sum_{n=-\infty}^{\infty} c_n^{(s)} e^{in\phi}, \quad \text{where } c_n^{(s)} = c_{-n}^{*(s)} \text{ and } c_0^{(s)} = \frac{1}{2\pi} \tag{6.9}$$

Note that this should be interpreted as a normalised *probability density* function, unlike the probability amplitude in Equation 6.8. $c_0^{(s)} = 1/2\pi$ is derived from the condition $\int_0^{2\pi} L_s(\phi)d\phi = 1$.

Next we define a phase estimator $\hat{\phi}$ that acts on the likelihood function, according to

$$\hat{\phi}_s = \arg \int_0^{2\pi} L_s(\phi) e^{i\phi} d\phi = \arg\left(c_{-1}^{(s)}\right) \tag{6.10}$$

which we use to predict the likeliest phase given the knowledge L_s at estimation step s; all the terms cancel in the integral except $c_{-1}^{(s)}$. This is equivalent to calculating the expectation value of ϕ from $L_s(\phi)$, and can further be justified from the fact that $\arg(c_{-1}^{(s)})$ gives the phase at which a symmetric distribution in L_s would be maximised[3].

We wish to update the likelihood function based on experimental observations. At shot s for an outcome σ_s at a measurement angle θ_s, the probability that the qubit phase was ϕ is simply $p(\sigma_s, \theta_s|\phi) = p(\sigma_s, \theta_s, \phi)$, given by Equation 6.8. We relate this to the term $p(b|a)$ in Equation 6.7, with $p(a)$ and $p(a|b)$ the likelihood function before and after the measurement b. Expressing Equation 6.7 in terms of the distributions we have defined, the likelihood function update at shot s is

$$L_s(\phi|\sigma_s, \theta_s) = \frac{1}{K(\sigma_s, \theta_s)} p(\sigma_s, \theta_s|\phi) L_{s-1}(\phi) \tag{6.11}$$

$$K(\sigma_s, \theta_s) = \int_0^{2\pi} p(\sigma_s, \theta_s|\phi) L_{s-1}(\phi) d\phi = \pi + \frac{\pi}{2}\sigma_s \left[c_{-1}^{(s-1)} e^{i\theta_s} + c_1^{(s-1)} e^{-i\theta_s}\right] \tag{6.12}$$

where the denominator $K(\sigma_s, \theta_s)$ is proportional to the prior probability of measuring σ_s for a measurement angle θ_s over all possible ϕ values, and was calculated by expanding $p(\sigma_s, \theta_s|\phi)$ and $L_{s-1}(\phi)$ in powers of $e^{i\phi}$. It represents $p(b)$ in Equation 6.7, and is used as a normalisation term independent of ϕ.

Using Equation 6.8 and Equation 6.9 to expand Equation 6.11 and equating powers of $e^{i\phi}$, we obtain

$$\sum_{n=-\infty}^{\infty} c_n^{(s)} e^{in\phi} = \frac{1}{2K(\sigma_s, \theta_s)} \left\{1 + \frac{\sigma_s}{2}\left[e^{i(\theta_s+\phi)} + e^{-i(\theta_s+\phi)}\right]\right\} \sum_{n=-\infty}^{\infty} c_n^{(s-1)} e^{in\phi} \tag{6.13}$$

$$\text{hence,} \quad c_n^{(s)} = \frac{1}{2K(\sigma_s, \theta_s)} \left\{c_n^{(s-1)} + \frac{\sigma_s}{2}\left[c_{n-1}^{(s-1)} e^{i\theta_s} + c_{n+1}^{(s-1)} e^{-i\theta_s}\right]\right\}. \tag{6.14}$$

[3]Note that L_s is real, and can also be written as $c_0 + \sum_{n=1}^{\infty} 2\left[\mathbb{R}\mathrm{e}(c_n^{(s)})\cos n\phi - \mathbb{I}\mathrm{m}(c_n^{(s)})\sin n\phi\right]$.

From Equation 6.14, $L_s(\phi)$ can be calculated from $L_{s-1}(\phi)$ based on σ_s and θ_s. At the beginning of the protocol we initialise the coefficients of L_0 to $c_0 = 1/2\pi$, $c_{n\neq 0} = 0$, to reflect the lack of prior knowledge about ϕ. For s shots, according to the propagation rule Equation 6.14 the coefficients from $c_{-s}^{(s)}$ to $c_s^{(s)}$ are nonzero. Thus if $s_{\max}$ total shots are to be run, the protocol requires $2s_{\max}+1$ coefficients to be stored, and the computation time scales as $\mathrm{O}(s_{\max})$. This also determines the maximum Fourier frequency in $L_s(\phi)$, given by $e^{i\phi s_{\max}}$, which governs the uncertainty in $\hat{\phi}_s$.

We now have the ingredients for a Bayesian estimation of ϕ using a series of measurements $\{\sigma_s\}$ at angles $\{\theta_s\}$, which can be chosen non-adaptively. In this case, the update rules in Equation 6.14 amount to a maximum-likelihood estimation, and $L_s(\phi)$ is insensitive to the order in which the measurements $\{\sigma_s, \theta_s\}$ were conducted.

Any non-adaptive choice runs the risk of measuring the Ramsey curve close to an extremum, where the marginal information gain is low if previous measurements have already bracketed the extrema locations, because the measurement is likely to be bright or dark. Instead it is desirable to optimise the information gained from each shot; we wish to choose θ_s based on L_{s-1}. We calculate the information gain using the Shannon entropy, defined as

$$H\left[L(\phi)\right] = -\int_0^{2\pi} L(\phi)\ln\left[L(\phi)\right]d\phi. \tag{6.15}$$

and the Shannon entropy gain in our knowledge for an outcome σ_s at an angle θ_s is

$$\Delta H\left[\sigma_s, \theta_s\right] = H\left[L_s(\phi|\sigma_s, \theta_s)\right] - H\left[L_{s-1}(\phi)\right]. \tag{6.16}$$

We wish to choose an angle θ_s that maximises the *expected entropy gain*, given by

$$\langle \Delta H\left[\theta_s\right]\rangle = \sum_\sigma K(\sigma, \theta_s)\Delta H\left[\sigma, \theta_s\right] \tag{6.17}$$

where the term $\sum_\sigma K(\sigma, \theta_s)$ is included to take into account the probabilities of observing the outcomes σ. For example, although the information gain would be high if a bright outcome was observed at a phase likely to produce a dark outcome, the chance of this occurring is low and should be suitably weighted.

We now seek to maximise $\langle \Delta H\left[\theta_s\right]\rangle$. Expanding Equation 6.17, we can obtain [111]

$$\langle \Delta H\left[\theta_s\right]\rangle = \sum_\sigma \left\{\int_0^{2\pi} L_{s-1}(\phi)p(\sigma, \theta_s|\phi)\ln p(\sigma, \theta_s|\phi)d\phi - K(\sigma, \theta_s)\ln K(\sigma, \theta_s)\right\} \tag{6.18}$$

which simplifies to

$$\langle \Delta H\left[\theta_s\right]\rangle = 1 - \ln(2) + \sum_{m=1}^{\infty} \frac{\mathbb{R}\mathrm{e}\left(c_{2m}^{s-1}e^{-2im\theta_s}\right)}{m\left(4m^2+1\right)} - \sum_\sigma K(\sigma, \theta_s)\ln K(\sigma, \theta_s). \tag{6.19}$$

There is no analytical way to maximise Equation 6.19 in general, and a numerical search is required to find the θ_s angle where $\langle \Delta H\left[\theta_s\right]\rangle$ is highest. This angle is then fed back to the measurement apparatus.

To summarise, we carry out each shot s of the Bayesian scheme by repeating the following steps:

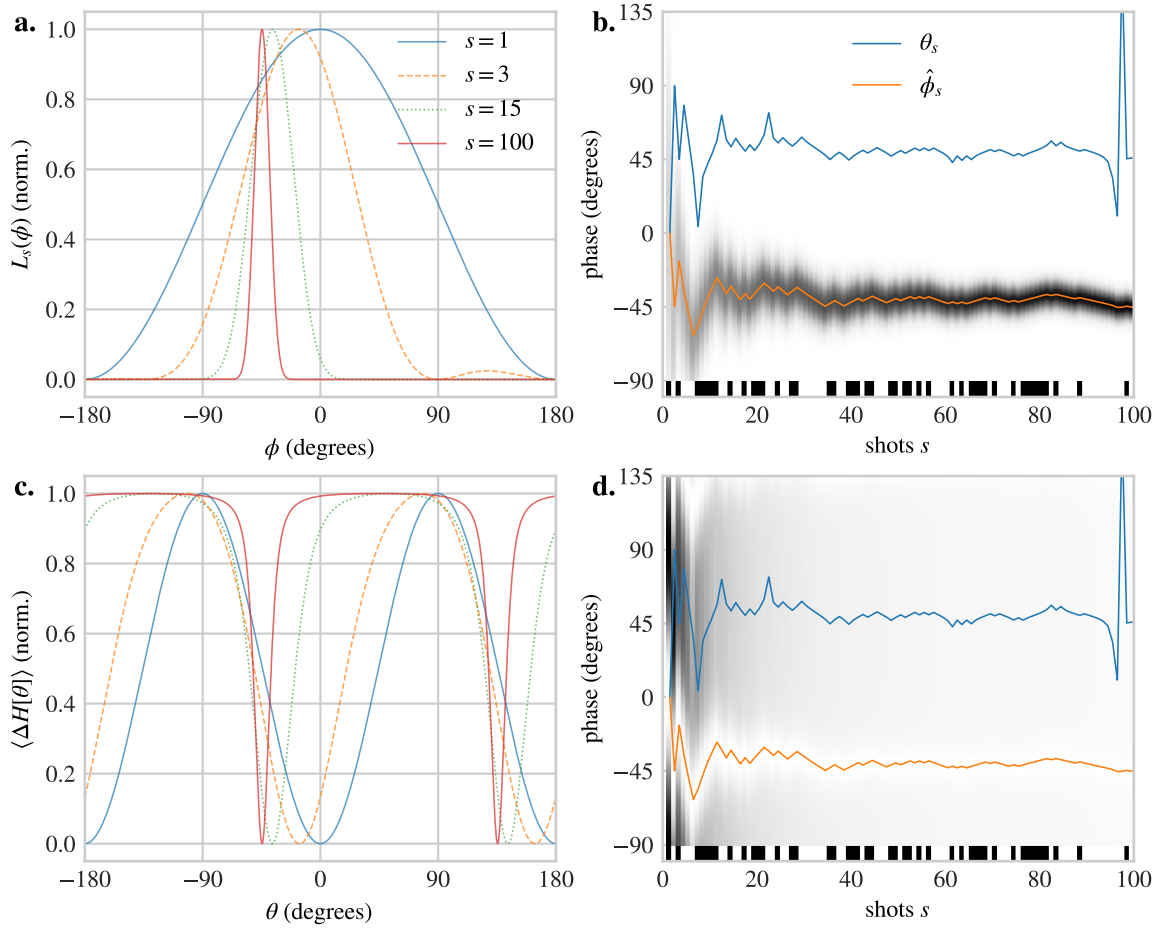

Figure 6.2: Bayesian phase estimation as a function of shot, for $\phi_{\text{real}} = -45°$. The top plots show the evolution of $L_s(\phi)$, which approaches a Gaussian distribution with a width that scales as $1/\sqrt{s}$ as the sequence continues, and the bottom plots show $\langle \Delta H\,[\theta_s] \rangle$, where the measurement angles θ_s are chosen at the maxima of this function. The measurement outcomes σ_s are shown as black and white stripes on the lower edges of **b.** and **d.**; the effect of sequences of dark or bright outcomes is to skew the distribution, growing smaller as the estimation proceeds. As the likelihood converges to a value, the marginal entropy acquired per shot grows smaller, and $\langle \Delta H\,[\theta_s] \rangle$ is almost flat 90° away from $\hat{\phi}_s$.

Choose a measurement angle θ_s by carrying out a numerical maximisation of Equation 6.19 and finding the optimal angle; this is the most computationally-intensive section of the protocol since it requires evaluating Equation 6.19 tens or hundreds of times. Currently a golden-ratio search is used.

Execute the Ramsey sequence to obtain the measurement outcome σ_s for the applied θ_s, with the expected outcome probability given by Equation 6.8; this is the limiting delay if we wish to be duty-cycle-limited.

Update the likelihood $L_s(\phi)$ using σ_s and θ_s in Equation 6.11; after this the system can estimate $\hat{\phi}_s$.

6.2.3 Implementation and results

The protocol has been applied to single calcium and beryllium ions, with the M-ACTION system handling all the steps described above. A single shot takes $\sim 2\,\text{ms}$, limited by pre-cooling time, with the computation taking around $200\,\mu\text{s}$ per shot when 50 coefficients are used; a low-latency phase update is sent from the CPU to the relevant DDS board as discussed in §3.3.7 before the Ramsey sequence begins, taking $\sim 2\,\mu\text{s}$. The Ramsey wait time is up to $600\,\mu\text{s}$, and the protocol copes well with decoherence, which increases the variance of estimates without causing a systematic bias.

Figure 6.2a and b show the evolution of $L_s(\phi)$ and $\langle \Delta H\,[\theta_s] \rangle$ for a typical protocol execution where $\phi \simeq -\pi/4$. After the first shot, $L_1(\phi)$ has a cosine profile whose phase is determined by the shot outcome. With successive shots $L_s(\phi)$ converges to a Gaussian centred on the correct phase. Figure 6.2c and d show the normalised expected entropy gain at various shots, which is what the measurement angle θ_s is set to; it is usually but not always highest $\pi/2$ away from $\hat{\phi}_s$. The system adapts to unexpected strings of bright or dark results with shifts in θ_s, such as at $s = 90 \to 100$ in Figure 6.2.

The Bayesian scheme has been used for several years in the mixed-species experiment as a tool for calibrating the coefficients of the magnetic field feedforward mentioned in §3.4.4. The Ramsey experiment is begun at a fixed delay after the zero-phase point of the 50 Hz ac mains cycle to estimate the frequency shift, and the frequency fluctuations over a scan of the delay are minimised by setting the compensation amplitudes and phases of the feedforward coil current. This is currently automated by fitting the frequency scans to harmonics of the mains; Figure 6.3a shows the improvement from a round of calibration. The automated script was developed by Brennan MacDonald-de Neeve [110], and takes around 5 minutes without human intervention[4].

The Bayesian protocol can be adapted to 'track' a time-varying phase, by 'smearing out' the likelihood function with every shot such that non-negligible probability density is introduced in regions of $L_s(\phi)$ where it had previously been ruled out based on measurements. This is done by multiplying the coefficients $c_n^{(s)}$ in $L_s(\phi)$ by a Gaussian function of n centred on the dc term $n = 0$, which is equivalent to convolving $L_s(\phi)$ with a Gaussian in the phase domain. The approach could be useful for maintaining an optimal measurement angle in spite of experimental drift, which cannot be done with non-adaptive schemes. Originally this approach was planned to track the real-time frequency shifts during an ac mains cycle, however it was not pursued since the algorithm converges rapidly enough even with zero prior information. A simulation is shown in Figure 6.3b, with a smearing function of $e^{-0.003n^2}$; strong smearing reduces the effect of high-Fourier-frequency coefficients on the algorithm, increasing its 'slew rate' at the expense of its precision [90]

Figure 6.3c and d show a Monte-Carlo simulation of 30 000 realisations of three different phase estimation approaches: Bayesian, fitting to data obtained by uniformly sampling the measurement angle θ_s, and using the mean measured outcome $\bar{\sigma}$ according to $\hat{\phi}_s = \arccos(\bar{\sigma}) - \theta$, where θ is always set to $\phi + \pi/2$, the optimal angle as $s \to \infty$ when projection noise and decoherence are considered [71]. The standard deviations of all three methods scale as $1/\sqrt{s}$, the standard quantum limit, however the fit standard deviations are higher than the other two by a near-constant prefactor of 1.25. This is probably due to data taken near the sinusoid extrema contributing less information to the phase estimate, however this requires further investigation. The prefactor causes the fitting approach to require

[4] Beyond making sure the lasers stay locked.

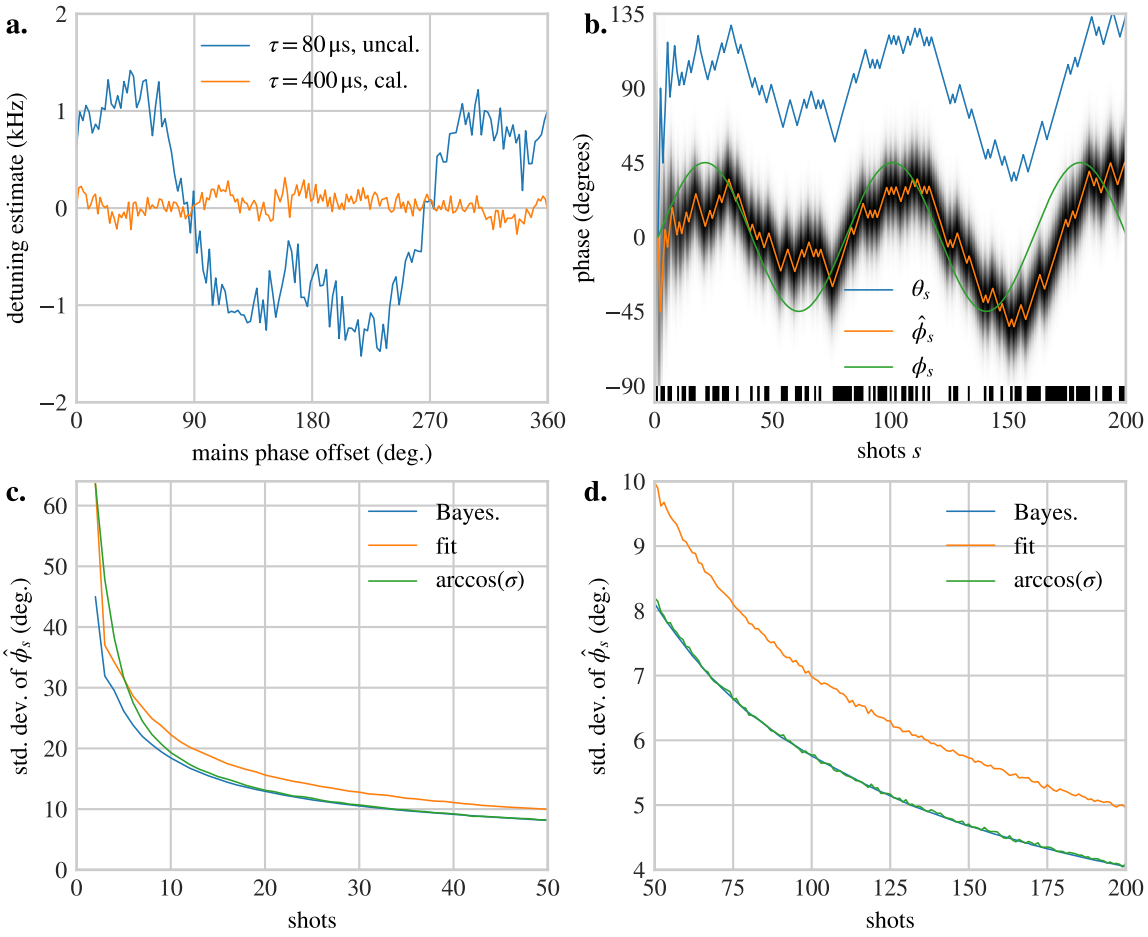

Figure 6.3: Bayesian phase estimation performance. **a.** Frequency detuning estimations at a range of ac mains phases, each using 50 shots. This is used as an error signal for optimising the feedforward to reduce 50 Hz harmonics, with the typical results shown. The pre- and post-optimisation scans use Ramsey wait times τ of 80 and 400 µs, increasing the frequency sensitivity of the latter. **b.** 200 shots for a simulated sinusoidal fluctuation in ϕ, where the coefficients $c_n^{(s)}$ are scaled by $e^{-0.003n^2}$ after each shot and the Bayesian scheme tracks the fluctuation. **c.** and **d.** Monte-Carlo simulation of Bayesian phase estimation, fitting and direct phase estimation from the population, showing the standard deviation of the estimates for 30 000 trials for each shot.

around $1.25^2 \simeq 1.5\times$ more shots to achieve the same precision as the Bayesian and arccos approaches, which is experimentally significant.

The Bayesian and arccos schemes perform identically above $s = 20$, which is expected since as $s \to \infty$ both schemes are optimal. In many experimental calibration scenarios where ϕ is not known initially or drifts between shots, however, the Bayesian scheme offers extra performance because θ is always optimally selected, and it converges more rapidly over the initial shots than the non-adaptive scheme. This is consistent with other adaptive phase estimation experiments [21, 22]. When the arcos scheme was run with sub-optimal θ, the $\hat{\phi}_s$ accuracy suffered considerably compared to the Bayesian and fit approaches.

None of the approaches takes advantage of the $1/\tau$ Heisenberg scaling in accuracy offered by increasing the Ramsey wait time. There are non-adaptive algorithms that feature Heisenberg scaling, such as robust phase estimation (RPE) [82, 170]. The Bayesian phase estim-

ation scheme has been used as a 'subroutine' to speed RPE up further; preliminary results by Brennan MacDonald-de Neeve show that the adaptive scheme improves the scaling prefactor by around 2x. Please see [110] for more detail. The Bayesian scheme can also be applied beyond Ramsey sequences, since its sinusoidal measurement model describes many forms of calibration including Rabi oscillations or parity phase scans (see below).

6.2.4 Discussion

Schemes like Bayesian phase estimation already demonstrate the advantages offered by low-latency computational resources close to the qubits; currently few public experimental control systems are capable of running adaptive algorithms of this complexity without the shot repetition rate being limited by either real-time computation or latency delays. The original algorithm was written and simulated by Andrey Lebedev on a desktop PC in C++ for performance reasons, with no foreknowledge that this would be convenient for us, and was run virtually unchanged[5] on the M-ACTION CPU with an ion.

An equivalent scheme could be implemented on an FPGA, however the nonlinear optimisation would require significant development effort, and changes to the algorithm would take many times as long to prototype and debug than in a purely-software implementation such as ours. Similar adaptive schemes have been implemented in other quantum platforms [21], however they have relied on replacing the optimisation with simpler numerical algorithms, which admittedly do not significantly affect the algorithm scaling for one-dimensional problems such as this. As the dimensionality of a Bayesian estimation problem increases, finding good approximate schemes becomes increasingly difficult, and is an active topic of research [56, 201].

The greatest promise of adaptive Bayesian schemes lies in multi-parameter estimation, in which a multi-dimensional likelihood function is constrained in each dimension simultaneously by measurement outcomes. Analytical approaches such as the one used here become more difficult to formulate, and numerical approaches become useful for describing and updating the likelihood function. One promising scheme is a particle filter or sequential Monte-Carlo model, in which the likelihood function is estimated with a numerical distribution of many 'particles', each with a coordinate in N-dimensional parameter space and a weight [42, 56, 55]. Statistical properties of the likelihood function such as its expectation value, variance or entropy can be computed from the coordinates. Their locations are updated probabilistically in response to new measurements. Although predicting the measurement values with the maximum expected entropy gain is still computationally expensive, approximate techniques approaching optimal performance can be used. A particle filter scheme has been written C++ for the M-ACTION system to be used with multi-parameter calibrations such as micromotion compensation, however so far it has only been used to reproduce Bayesian phase estimation. Its most interesting application, currently under investigation, will be on two- and multi-qubit gates.

This concludes the first half of this chapter. The remainder focuses on the operation, (non-Bayesian) calibration and fidelity estimation of multi-qubit gates.

[5] The largest change was inverting the sign of θ to suit the experiment.

6.3 Multi-qubit entangled states

A general two-qubit pure state in the qubit measurement basis is $|\Psi\rangle = a\,|gg\rangle + b\,|ge\rangle + c\,|eg\rangle + d\,|ee\rangle$, and a general mixed state is described by a density matrix with 16 terms corresponding to $|\Psi\rangle\langle\Psi|$. Three-qubit states are similarly described. The regularly-used states and operations are presented before multi-qubit experiments are discussed.

6.3.1 Bell states and two-qubit rotations

The Bell states for two qubits are

$$|\Phi_+\rangle = \frac{1}{\sqrt{2}}\left(|gg\rangle + |ee\rangle\right), \quad |\Phi_-\rangle = \frac{1}{\sqrt{2}}\left(|gg\rangle - |ee\rangle\right), \tag{6.20}$$

$$|\Psi_+\rangle = \frac{1}{\sqrt{2}}\left(|ge\rangle + |eg\rangle\right), \quad |\Psi_-\rangle = \frac{1}{\sqrt{2}}\left(|ge\rangle - |eg\rangle\right), \tag{6.21}$$

and form a basis spanning the two-qubit Hilbert space [139]. The $|\Phi\rangle$ ($|\Psi\rangle$) states have even (odd) *parity*, namely the probability of measuring an even (odd) number of qubits in state $|g\rangle$ minus that of an odd (even) number. $|\Psi_+\rangle$ and $|\Psi_-\rangle$ are known respectively as the triplet and the singlet state. Using the definitions from §4.3 the parity value is equivalent to $p_0 + p_2 - p_1$, or the probability of detecting 0 or 2 bright ions minus that of detecting 1 bright ion.

The general two-qubit rotation matrix for shared rotations is

$$\begin{aligned}\boldsymbol{T}_{2R}(\theta,\phi) &= \boldsymbol{T}_{n,\mathrm{car}}(\theta,\phi)\otimes\boldsymbol{T}_{n,\mathrm{car}}(\theta,\phi)\\ &= \begin{pmatrix} \cos^2\frac{\theta}{2} & -ie^{i\phi}\sin\frac{\theta}{2}\cos\frac{\theta}{2} & -ie^{i\phi}\sin\frac{\theta}{2}\cos\frac{\theta}{2} & -e^{2i\phi}\sin^2\frac{\theta}{2} \\ -ie^{-i\phi}\sin\frac{\theta}{2}\cos\frac{\theta}{2} & \cos^2\frac{\theta}{2} & -\sin^2\frac{\theta}{2} & -ie^{i\phi}\sin\frac{\theta}{2}\cos\frac{\theta}{2} \\ -ie^{-i\phi}\sin\frac{\theta}{2}\cos\frac{\theta}{2} & -\sin^2\frac{\theta}{2} & \cos^2\frac{\theta}{2} & -ie^{i\phi}\sin\frac{\theta}{2}\cos\frac{\theta}{2} \\ -e^{-2i\phi}\sin^2\frac{\theta}{2} & -ie^{-i\phi}\sin\frac{\theta}{2}\cos\frac{\theta}{2} & -ie^{-i\phi}\sin\frac{\theta}{2}\cos\frac{\theta}{2} & \cos^2\frac{\theta}{2} \end{pmatrix}\end{aligned} \tag{6.22}$$

Transforming $\boldsymbol{T}_{2R}(\theta,\phi)$ into the Bell basis, we obtain

$$\boldsymbol{T}_B = \begin{array}{c} \\ |\Phi_+\rangle \\ |\Phi_-\rangle \\ |\Psi_+\rangle \\ |\Psi_-\rangle \end{array} \begin{array}{c} \begin{array}{cccc} |\Phi_+\rangle & |\Phi_-\rangle & |\Psi_+\rangle & |\Psi_-\rangle \end{array} \\ \begin{pmatrix} 1-2\cos^2\phi\sin^2\frac{\theta}{2} & -i\sin 2\phi\sin^2\frac{\theta}{2} & -i\cos\phi\sin\theta & 0 \\ i\sin 2\phi\sin^2\frac{\theta}{2} & 1-2\sin^2\phi\sin^2\frac{\theta}{2} & -\sin\phi\sin\theta & 0 \\ -i\cos\phi\sin\theta & \sin\phi\sin\theta & \cos\theta & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \end{array}. \tag{6.23}$$

This shows that for suitable θ and ϕ we can transform all the Bell states into each other, except for the singlet $|\Psi_-\rangle$ state, which is invariant to global rotations alone. The transformations can be carried out with the following rotations:

$$|\Phi_+\rangle = \pm iR_\pi\left(\pm\frac{\pi}{4}\right)|\Phi_-\rangle, \quad |\Phi_+\rangle = \pm iR_\pi\left(\pm\frac{3\pi}{4}\right)|\Phi_-\rangle, \tag{6.24}$$

$$|\Phi_+\rangle = iR_{\pi/2}(0)\,|\Psi_+\rangle, \quad |\Phi_+\rangle = -iR_{\pi/2}(\pi)\,|\Psi_+\rangle, \tag{6.25}$$

$$|\Phi_-\rangle = \mp R_{\pi/2}\left(\pm\frac{\pi}{2}\right)|\Psi_+\rangle. \tag{6.26}$$

Note that the triplet $|\Psi_+\rangle$ state is invariant to R_π rotations, and the $|\Phi_+\rangle$ and $|\Phi_-\rangle$ states are invariant to $R_{\pi/2}(0)$ and $R_{\pi/2}(\pi)$ rotations; these invariances are relevant for later experiments. The prefactors before the rotations introduce global phases, which have no effect in two-qubit experiments or multi-qubit states in which the Bell states form a separable part[6].

6.3.2 Parity oscillations

A common method for characterising the coherence of a two-qubit density matrix ρ and its fidelity with respect to a desired Bell state $F = \langle B|\rho|B\rangle$ is measuring parity oscillations [173, 66, 127], where $|B\rangle$ is the target state. State tomography is required to estimate each density matrix element in general, however estimating the coherence and fidelity only requires a subset. From Equation 6.22 it can be shown that an $R_{\pi/2}(\phi)$ *parity pulse* acting on a superposition of the even-parity states $\left(|gg\rangle + e^{i\phi_0}|ee\rangle\right)/\sqrt{2}$, where $\phi_0 = 0$ (π) for $|\Phi_+\rangle$ $(|\Phi_-\rangle)$, maps the off-diagonal coherences $\rho_{gg,ee} = \rho^*_{ee,gg}$ to the diagonal elements of the density matrix. The diagonals are measured during readout, thus by scanning ϕ a sinusoid is obtained in the odd-parity population ($\rho_{ge,eg} + \rho_{eg,ge} = p_1$, i.e. a single bright ion) according to

$$p_{\text{odd}} = \frac{1}{2} + \frac{C}{2}\cos(2\phi + \phi_0), \tag{6.27}$$

where the contrast $0 \leq C \leq 1$ provides an estimate of $|\rho_{gg,ee}|$. C can also be expressed as $C = \left\langle \hat{\sigma}_z^{(1)}\hat{\sigma}_z^{(2)} \right\rangle$ on the two ions. The phase offset ϕ_0 arises from the input Bell state superposition. Thus the states $|\Phi_\pm\rangle$ produce sinusoids 180° out of phase with each other. The overall Bell state fidelity can be estimated according to

$$F = \frac{1}{2}\left(\rho_{gg,gg} + \rho_{ee,ee} + 2|\rho_{gg,ee}|\right) = \frac{1}{2}\left(\rho_{gg,gg} + \rho_{ee,ee} + C\right). \tag{6.28}$$

The fidelity of $|\Psi_+\rangle$ can be estimated by first applying $R_{\pi/2}(\pi/2)$ to produce $|\Phi_-\rangle$ then using the same analysis technique. This cannot be done with $|\Psi_-\rangle$, which must first be mapped to another Bell state with an addressed rotation on one of the qubits. Parity pulses can also be used to prepare particular Bell states, discussed further in §7.2.2.

6.3.3 Greenberger-Horne-Zeilinger states

The Greenberger-Horne-Zeilinger (GHZ) states are a generalisation of the Bell states to n qubits [57, 139], with 2^n basis states. Two of them are

$$|\Phi_{\pm,\text{GHZ}}\rangle = \frac{1}{\sqrt{2}}\left(|g\ldots g\rangle \pm |e\ldots e\rangle\right). \tag{6.29}$$

In this thesis we consider only 3-qubit experiments. Similarly to the Bell states, parity oscillations can be used to characterise a superposition $\left(|ggg\rangle + e^{i\phi_0}|eee\rangle\right)/\sqrt{2}$, with parity curves for $n = 3$ given by

$$p_{\text{odd}} = \frac{1}{2} + \frac{C}{2}\cos(3\phi + \phi_0) \tag{6.30}$$

for an input state where now $p_{\text{odd}} = \rho_{gge,gge} + \rho_{geg,geg} + \rho_{egg,egg} + \rho_{eee,eee}$. Note the period of ϕ is now $2\pi/3$.

[6] However they will affect multi-qubit entangled states where this is not true, e.g. $|g\rangle|\Phi_+\rangle + |e\rangle|\Phi_-\rangle$

6.4 Multi-qubit gates

A universal quantum information processing platform requires methods to generate and control entanglement between qubits and perform conditional operations [139, 41], which are both usually carried out using unitary *multi-qubit gates*. Such gates can be implemented in trapped-ion QIP either deterministically, usually using the Coulomb repulsion, or probabilistically using heralded-entanglement schemes in ion traps coupled via phototonic links [54]. A foundational proposal for Coulomb-based multi-qubit gates[7] was the Cirac-Zoller (CZ) gate [30, 178], which along with state preparation, readout and single-qubit gates showed that trapped ions satisfied the basic QIP criteria.

The CZ gate relies on ions initialised in the ground state of a common mode of motion, with individually-addressed sideband pulses used to entangle their internal states. For the energy ladder discussed in Figure 2.2a, the $|g, n = 0\rangle$ and $|e, n = 0\rangle$ state amplitudes are unchanged by the red sideband (RSB) and blue sideband (BSB) drives respecively; this is used to perform operations where the excitation of the shared motional 'bus' mode is conditional on the state amplitudes. The bus mode excitation in turn affects the internal states of other ions driven by RSB/BSB drives.

Instead of CZ-style techniques, gates used in current experiments usually induce a geometric phase between the two-qubit basis states via a state-dependent drive of the shared motional mode [185, 121, 119, 94, 167, 83]. Such gates are relatively insensitive to motional excitation compared to CZ approaches, and individually-addressed rotations are not required. In the mixed-species setup Mølmer-Sørensen gates [121, 185] are used for both two- and three-ion entangling operations.

6.4.1 Mølmer-Sørensen gate

For a two-ion crystal, the Mølmer-Sørensen (MS) gate is run by off-resonantly driving the red and blue sidebands of a motional mode simultaneously, by applying a bichromatic field with amplitude

$$E_{\mathrm{bsb/rsb}}(t) = E_0 \cos\left\{\left[\omega_0 \pm (\omega_m + \delta_m)\right] t \pm \phi_c/2\right\}, \tag{6.31}$$

where ω_0 and ω_m are the qubit and mode frequencies, δ_m is a detuning of the sideband from the mode frequency, and ϕ_c is a shared phase. The sum frequency of the sideband drives is $2\omega_0$, and the gate carries out two-photon transitions between the states $|gg\rangle \leftrightarrow |ee\rangle$ and $|ge\rangle \leftrightarrow |eg\rangle$. Following the same steps used to derive Equation 2.10, we can write the Hamiltonian describing this process in the interaction picture as [121, 83, 167]

$$\begin{aligned}\hat{H}(t) = \frac{\hbar}{2}\Omega\left(\hat{\sigma}_+^{(1)} + \hat{\sigma}_+^{(2)}\right)\left(e^{i[(\omega_m+\delta_m)t+\phi_c/2]} + e^{-i[(\omega_m+\delta_m)t+\phi_c/2]}\right) \\ \times \exp\left[i\eta\left(\hat{a}e^{-i\omega_m t} + \hat{a}^\dagger e^{i\omega_m t}\right)\right] + \mathrm{h.c.}\end{aligned} \tag{6.32}$$

where $\hat{\sigma}_+^{(n)}$ is the raising operator on ion n. Expanding in the Lamb-Dicke regime and removing off-resonant terms oscillating at or above ω_m, this reduces to

$$\hat{H}(t) = -\frac{\hbar}{2}\eta\Omega\left(\hat{a}e^{-i\delta_m t} + \hat{a}^\dagger e^{i\delta_m t}\right)\hat{S}_{\phi_c}, \quad \hat{S}_{\phi_c} = (\hat{\sigma}_x^{(1)} + \hat{\sigma}_x^{(2)})\cos\phi_c + (\hat{\sigma}_y^{(1)} + \hat{\sigma}_y^{(2)})\sin\phi_c \tag{6.33}$$

[7] The earliest trapped-ion 'multi-qubit' gates were between the spin and motional states of a single ion, similar to the motional coherence experiments described earlier [123].

This can be solved analytically [184, 93, 83] to obtain the propagator

$$U(t) = \hat{D}\left(\alpha(t)\hat{S}_{\phi_c}\right) \exp\left[i\frac{\eta^2\Omega^2}{4\delta_m^2}\left(\delta_m t - \sin\delta_m t\right)\hat{S}_{\phi_c}^2\right], \quad \alpha(t) = \frac{\eta\Omega}{2\delta_m}\left(e^{i\delta_m t} - 1\right) \tag{6.34}$$

where $\hat{D}(\alpha) = e^{\alpha\hat{a}^\dagger + \alpha^*\hat{a}}$ is the motional displacement operator. Critical to the utility of the MS gate, at the time $t = \tau = 2\pi/|\delta_m|$ the displacement becomes the identity, and the gate does not alter the motional state; the propagator is now

$$U(\tau) = \exp\left[i\frac{\eta^2\Omega^2\pi}{2\delta_m}\hat{S}_{\phi_c}^2\right] \tag{6.35}$$

If $\Omega = |\delta_m|/2\eta$, the gate carries out the operation

$$U_{MS} = \exp(i\pi\hat{S}_{\phi_c}^2/8). \tag{6.36}$$

This operator carries out correlated spin-spin inversions on the two ions, and performs the following conditional state mapping:

$$|gg\rangle \to (|gg\rangle - ie^{-i\phi_c}|ee\rangle)/\sqrt{2} \tag{6.37}$$
$$|ge\rangle \to (|ge\rangle - i|eg\rangle)/\sqrt{2} \tag{6.38}$$
$$|eg\rangle \to (-i|ge\rangle + |eg\rangle)/\sqrt{2} \tag{6.39}$$
$$|ee\rangle \to (-ie^{i\phi_c}|ee\rangle + |gg\rangle)/\sqrt{2} \tag{6.40}$$

where a common global phase has been factored out. With the addition of the MS gate, a universal set of quantum operations is thus available in the mixed-species setup, and by initialising in $|gg\rangle$ a Bell state can be created on demand.

For a larger number of ions n, the operator Equation 6.35 is still valid with $\hat{S}_{\phi_c} = \sum_{m=1}^{n}\hat{\sigma}_x^{(m)}\cos\phi_c + \hat{\sigma}_y^{(m)}\sin\phi_c$, and the GHZ states $|g\ldots g\rangle + |e\ldots e\rangle$ can similarly be produced for even n. For odd n the mapping is less straightforward since each qubit is involved in an even number of pairwise interactions with other qubits, however an additional $R_{\pi/2}$ rotation before or after the gate can be used in this case to prepare a GHZ state nonetheless [121].

The gate is experimentally implemented by driving a single-pass AOM with two tones

$$V_{\mathrm{bsb/rsb}}(t) = V_0\cos\left\{\left[\omega_{\mathrm{rf}} \pm (\omega_m + \delta_m) \pm \phi_m/2 + \delta_c\right]t + \phi_c/2\right\}, \tag{6.41}$$

which are produced by electronically adding the rf from two DDS channels, as described in [81] for single-ion spin-dependent force experiments. The optical field amplitude in Equation 6.31 from the AOM can be written as an amplitude modulation, according to

$$E = E_0\cos\left[(\omega_0 + \delta_c)t + \phi_c\right]\cos\left[2(\omega_m + \delta_m)t + \phi_m\right] \tag{6.42}$$

where the expression is the same as above except for δ_c, a common sideband detuning, and and ϕ_m, a differential sideband phase offset. For single-species gates the wavelength of the amplitude modulation $2(\omega_m + \delta_m)/k$ is on the order of metres at our gate parameters, thus the driving force from the field is in-phase for the ions on the CoM mode. In single-species gates the ions share ϕ_c and ϕ_m since they share a common drive, which is not true for mixed-species gates where separate gate lasers are used.

6.5 Single-species gate optimisation

Initially the sidebands are switched on independently and their amplitudes are equalised using a photodiode, then run simultaneously as the photodiode signal is viewed on an oscilloscope. This is used to verify that the amplitude modulation depth in Equation 6.42 reaches close to 100 %; nonlinearity in the AOM when the rf power is high causes intermodulation of the sideband tones that produces a tone at ω_{rf} and harms the gate fidelity.

Figure 6.4a and b show the single-species two-ion populations for calcium and beryllium as the MS drive time is swept. The gates are tuned using a slightly different procedure from previous setups [14, 83]. Initially a detuning δ_m is selected based approximately on the maximum expected Rabi frequency Ω, determined in previous calibrations. If the gate time does not fulfil $2\pi/|\delta_m|$, the phase-space displacement trajectories do not close and there is residual population left in $p_1 = p(|ge\rangle) + p(|eg\rangle)$; this is first-order independent of Ω and is used to optimise the gate time t_{gate}. Because of transient AOM effects, imperfect calibration of ω_m and an offset due to pulse shaping, t_{gate} can be different from τ by several microseconds.

The mapping $|gg\rangle \rightarrow |gg\rangle + e^{i\phi_c}|ee\rangle$ occurs only when $\Omega = |\delta_m|/2\eta$, and Ω is experimentally calibrated by equalising $p_0 = p(|ee\rangle)$ and $p_2 = p(|gg\rangle)$ at t_{gate}. In calcium the gate lasers cause a net ac Stark shift on the sidebands due to far-detuned dipole transitions, as discussed in §2.2.3; this is compensated with a common sideband detuning δ_c of several kHz, which is also tuned to minimise p_1. The Stark shift scales as Ω^2, hence Ω and δ_c are iteratively optimised at t_{gate}. In beryllium the effect is much weaker, however a small detuning is still required. Thus t_{gate}, δ_c and Ω form the key gate parameters. The fidelity of the Bell state given by is estimated using parity oscillations as described above. These are shown in Figure 6.4. The highest fidelities obtained in the mixed-species setup were 97.8(4)% and 99.4(6)% for beryllium and calcium, however more typically achieved fidelities are currently $\sim 96\,\%$ and $\sim 98\,\%$.

6.5.1 Gate error sources

For both ion species, the readout infidelities can cause significant miscalibration if not taken into account. Optimisations where a population is minimised or maximised are less influenced by these than cases where two populations are equalised or otherwise compared quantitatively. This principle is used to calibrate Ω and δ_c by optimising the populations at $2(2n+1)$ gate times, where $n = 0, 1, \ldots$; the gate populations over longer times are shown in Figure 6.10. Duty-cycle effects of the beryllium AOMs were measured and removed by fine realignment of the beams into the AOMs, before this calibration gave a significant improvement.

For ion crystals where η is low, which is the case for calcium, the sideband tones off-resonantly drive the carrier, causing oscillations in the populations that scale according to $\Omega^2/(\Omega^2 + \omega_m^2)$ (see Equation 4.11). Their period is below 1 µs and their phase is highly unstable with respect to beam power, thus uncorrected they contribute significantly to infidelity for faster gates. Pulse shaping, or adiabatically turning the rf on and off, reduces the spectral content of the sideband drives around the carrier frequency [167, 83]. The implementation was discussed in §3.3.6, and for smooth shapes with rise times of several microseconds the effect is negligible at our fidelity levels.

In beryllium MS gates, the gate phase ϕ_c is determined by the relative path difference of the co-com and 90-switch beams of the co-90 drive. This phase fluctuates over timescales

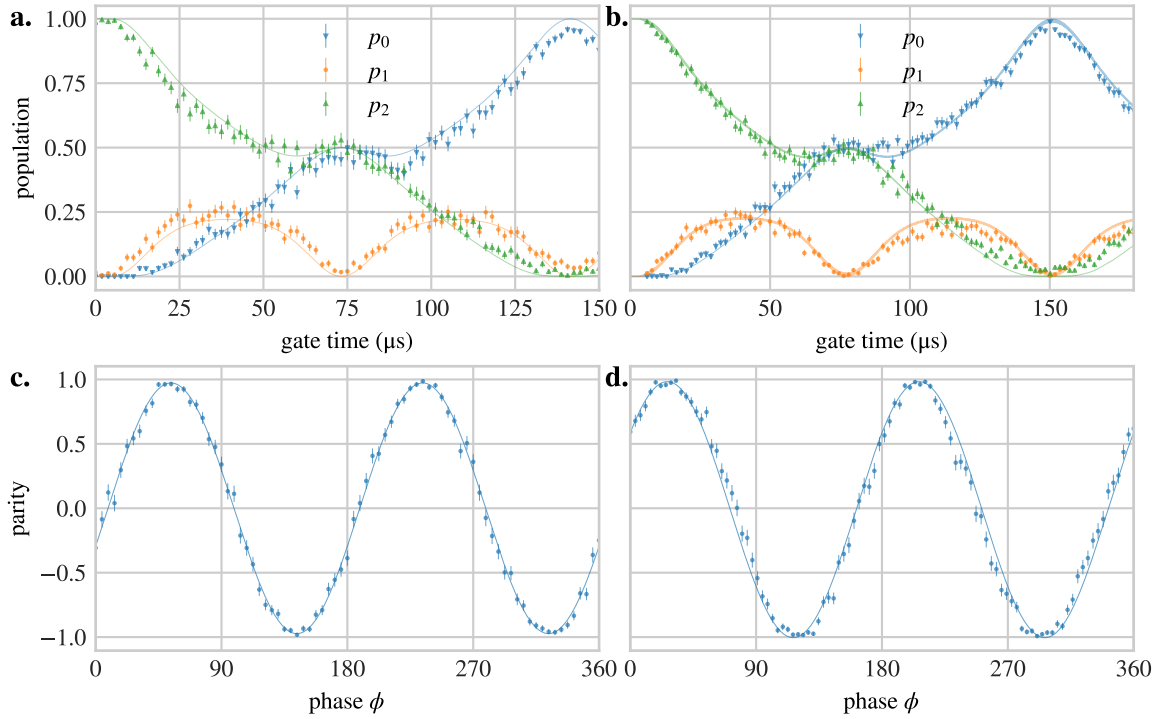

Figure 6.4: Single-species two-qubit gate populations and parity oscillations. **a.** shows beryllium populations as a function of time, with the maximally-entangled state produced in the centre, and **c.** shows parity oscillations at the gate time produced by sweeping the $\pi/2$ analysis pulse phase. A histogram fit to Equation 6.27 using Poissonian distributions gives a contrast of 97.3(4)%, which combined with the populations without the $\pi/2$ pulse at the gate time according to Equation 6.28, gives a fidelity of 97.8(4)%. **b.** and **d.** show the data from a two-calcium gate, with a contrast and fidelity of 99.3(6)% and 99.4(6)% respectively. Continuous lines in **a.** and **b.** were generated by numerically integrating the gate Hamiltonian; the blurring of the lines in **b.** is caused by rapid off-resonant carrier oscillations (see §6.5.1). A fit by eye to the gate detuning was performed.

longer than one experiment shot, which is problematic when phase-coherent rotations using the co-carrier beams are carried out on a superposition such as $|gg\rangle + e^{i\phi_c}|ee\rangle$. This is not a problem for the parity scan because the final $R_{\pi/2}$ rotation also uses the co-90 beams, cancelling common-mode fluctuations, however the fluctuations must be decoupled with extra rotations when part of a longer protocol [93], such as in the sequences used in the next chapter.

The gate is sensitive to collective beam intensity, with deviations also altering the geometric phase ϕ_c; the infidelity scales as $(\delta\Omega/W)^2$ for a small deviation $\delta\Omega$ [10]. A systematic deviation changes the phase but does not significantly affect the populations, thus the gate is somewhat insensitive to Ω miscalibration; however shot-to-shot power and beam pointing fluctuations cause dephasing. Closed-loop stabilisation of the power using the DDS boards is planned to reduce this effect by monitoring it using a photodiode, as discussed in Chapter 3. Beam pointing fluctuations are more difficult to stabilise, and passive stability is achieved in the mixed-species setup using physical enclosures as mentioned earlier, with the ultimate solution being to integrate beam delivery directly into the trap [116].

Generally fluctuations in Ω also contribute to Stark shift fluctuations, causing a secondary dephasing due to a fluctuating δ_c term which from Equation 2.33 also scales as Ω^2. This

has previously been compensated in calcium using an off-resonant beam co-propagating with the gate beam, tuned to cancel out δ_c [60].Since the two beams are affected equally by intensity or beam pointing fluctuations, the compensation is almost independent of power level. This approach is infeasible for beryllium, because a Raman transition is driven, however it is less important since the FIQ Stark shift is significantly lower than for calcium.

Error caused by dephasing, such as magnetic field or small intensity fluctuations, can be reduced using dynamical decoupling techniques with a corresponding increase in gate duration [16, 192, 9], although recently more rapid approaches have been investigated [113].

Although not a current limit, Raman scattering during the gate on beryllium can cause population leakage into hyperfine states outside $|g\rangle$ and $|e\rangle$ [87, 145]. This is currently well below our other error sources, however will become relevant in the future. It is also the major error source in multi-gate feedback protocols, discussed in the next chapter.

6.6 Two-qubit protocols

Using the techniques described so far, two QIP experiments were carried out on a two-calcium crystal, to investigate Bell state coherence and demonstrate a novel dissipative entanglement scheme.

6.6.1 Bell state coherence

The Bell states have several interesting properties with respect to noise. If Hamiltonians similar to Equation 6.1 act independently on both ions, inducing relative phases Φ_1 and Φ_2 between $|g\rangle$ and $|e\rangle$, then a general two-ion input state $|\psi_{\text{in}}\rangle = a\,|gg\rangle + b\,|ge\rangle + c\,|eg\rangle + d\,|ee\rangle$ becomes

$$|\psi_{\text{rot}}\rangle = ae^{i\frac{\Phi_1+\Phi_2}{2}}\,|gg\rangle + be^{i\frac{\Phi_1-\Phi_2}{2}}\,|ge\rangle + ce^{-i\frac{\Phi_1-\Phi_2}{2}}\,|eg\rangle + de^{-i\frac{\Phi_1+\Phi_2}{2}}\,|ee\rangle\,. \tag{6.43}$$

Hence the even Bell states $|\Phi_\pm\rangle$ acquire a differential phase of $\Phi_1 + \Phi_2$ between $|gg\rangle$ and $|ee\rangle$, while the odd Bell states acquire $\Phi_1 - \Phi_2$. Thus, if $\Phi_1 = \Phi_2$ and Φ_1 fluctuates due to different environmental influences as discussed in §6.1.1, the even Bell states acquire a relative phase of $2\Phi_1$, doubling the sensitivity to Φ_1 of a parity phase scan compared to a single-qubit Ramsey scan. For this reason the even Bell and GHZ states can improve the metrological sensitivity of a system to Φ_1, scaling as the number of entangled qubits [20, 96].

Conversely the relative phase cancels out for the odd Bell states, protecting them against common-mode dephasing noise and allowing the construction of a two-qubit decoherence-free subspace [77, 86, 126]. These effects have also been used experimentally to characterise sources of decoherence in ion traps [163, 86]. To probe the decoherence of $|\Phi_\pm\rangle$ in the mixed-species setup, two calcium ions were used and an MS gate was applied, producing an even state according to Equation 6.37. A parity pulse was applied after a wait time; the parity contrasts are shown in Figure 6.5b as orange downward-pointing triangles, with a decay time of 640(44) µs. The decay time of a single ion is indeed 2.0(2)× as long as the even Bell state, as expected.

Similarly the decoherence of $|\Psi_+\rangle$ was probed, by applying an $R_{\pi/2}(-\frac{\phi_c}{2} - \frac{\pi}{4})$ pulse immediately after the gate, followed by a wait time and a second rotation to return to an even state; the phase of a third $R_{\pi/2}$ pulse was scanned to estimate the contrast; shown as blue upward-pointing triangles with a 29(3) ms decay time. This is 20× as long as the

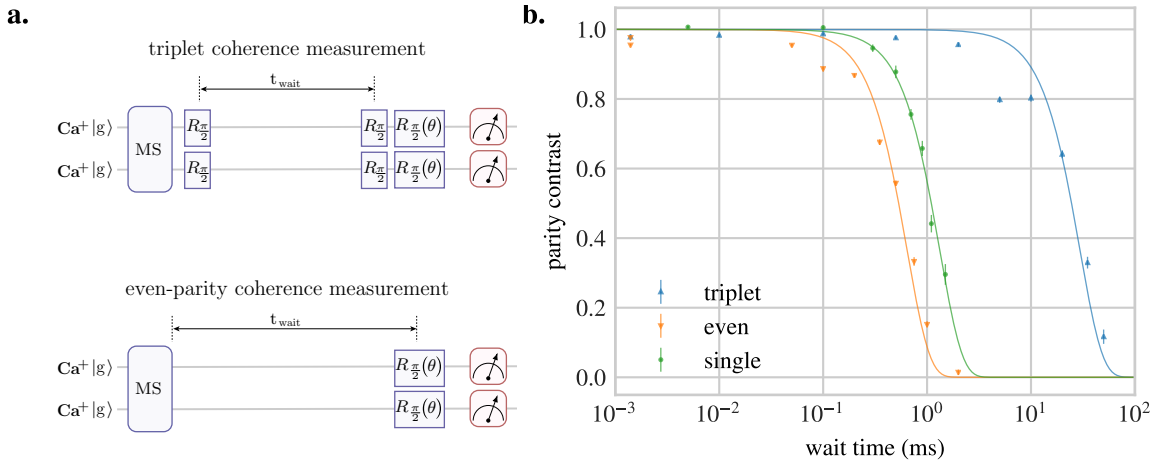

Figure 6.5: Coherence comparison between triplet and even Bell states, and a single ion. **a.** Pulse sequence for preparing the even or triplet states: an MS gate is used to prepare an even state, then either left in this state or rotated to the triplet state. After a wait time, the coherence is analysed. **b.** Parity contrast decay for the even and triplet states, and a single ion for comparison. A Gaussian decay $\exp -(t/\gamma)^2$ was fitted to the curves, resulting in γ values of 29(3) ms for the triplet, 640(44) µs for the even state, and 1.32(3) ms for a single ion.

single-ion decay. The residual decay is due to inhomogeneous dephasing, possibly due to a fluctuating magnetic field gradient, however the field difference between the ions would have to be several percent of the dc field to explain the relatively rapid decoherence, which seems implausible given the 5 µm ion spacing. The overlap of the even-parity populations p_0 and p_2 in the parity curves grew worse for longer wait times, so another possibility is ion heating affecting the Rabi frequency and causing imperfect rotations. Finally a similar decoherence of the odd state observed in a calcium-40 setup was partially caused by residual repumping light [163].

Once other effects are ruled out, field inhomogeneities may be characterised by performing a 'differential Ramsey' experiment, by separating the ions and holding them in different trap locations for a wait time, then recombining them and measuring the parity phase [171]. The phase shift will be proportional to the magnetic field difference at the two locations. Unlike a single-ion transport experiment this technique is immune to common-mode magnetic field noise. Characterising the field inhomogeneities will be required to minimise decoherence in mixed-species experiments involving the transport and separation of multiple calcium ions.

6.6.2 Dissipative entangled state preparation

As mentioned previously, the singlet Bell state is invariant to global rotations. Along with the fact that the MS gate couples only to even-parity states, this makes a dissipative Bell state preparation scheme feasible using a third level in calcium. Using two calcium ions, a scheme was demonstrated where the MS gate was not used to coherently produce a Bell state, but rather conditionally transfer population from $|gg\rangle$ to $|ee\rangle$. It was used along with rotations on the calcium $|g\rangle \leftrightarrow |d\rangle$ Zeeman transition (see Figure 2.5) using the rf drive.

The ion was prepared in $|gg\rangle$, then the MS gate[8] brought the population to $|ee\rangle$. An

[8] A π pulse could also have been used.

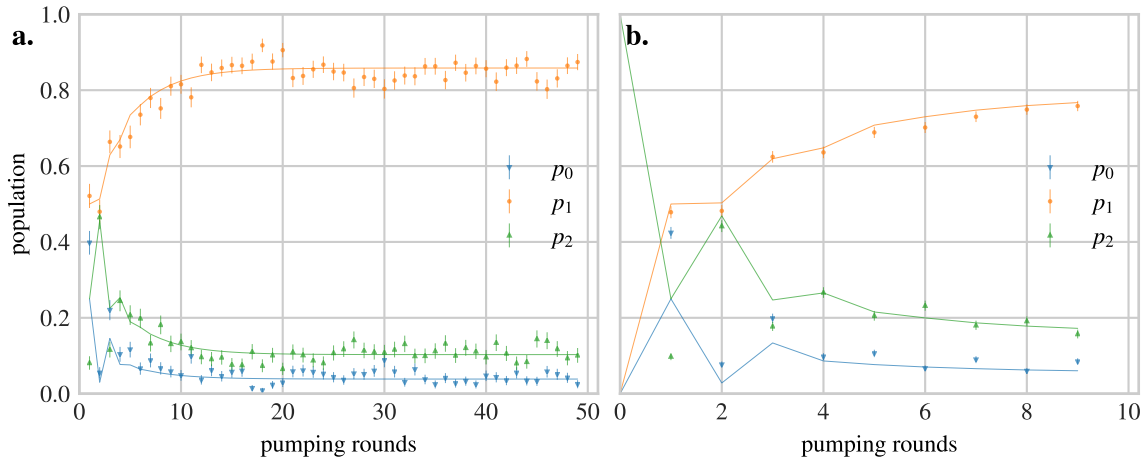

Figure 6.6: Populations for a long-term and short-term scan of the dissipative entanglement **a.** shows the long-term population limits at 300 shots per point, and **b.** shows the short-term dynamics at 1000 shots per point. Solid lines are fits using a model floating the rf rotation parameters and dephasing of the $S_{1/2}$ states.

854 nm repump pulse excited it to the $P_{3/2}$ manifold, from which it decayed into a mixture of $|gg\rangle$, $|gd\rangle$, $|dg\rangle$ and $|dd\rangle$. A $\sim 3\pi/4$ rf pulse was applied on $|g\rangle \leftrightarrow |d\rangle$ to cycle population between $|gg\rangle$, $|dd\rangle$ and $|dg\rangle+|gd\rangle$, while leaving the singlet $|dg\rangle-|gd\rangle$ population unchanged; this completed a single round of the protocol. The calibrated MS gate is a two-photon process, thus its next repetition did not couple to the $|dg\rangle$ or $|gd\rangle$ populations.

Each round of the protocol pumped population into the singlet from the other two-ion states; the evolution is shown in Figure 6.6 for 50 rounds. The single-ion population, representing $p(|dg\rangle - |gd\rangle) + p(|dg\rangle + |gd\rangle)$ rises over time as expected[9]; transport and separation were not yet available to allow state tomography to be performed. The protocol is in principle immune to most infidelities in the MS gate and rf pulses as long as $\tau \simeq 2\pi/\delta_m$, and δ_c are calibrated, and the rf couples evenly to the ions; the parity contrast for a single MS gate was $> 95\%$ both before and after the protocol experiments were concluded, and the rf Rabi flops from the same day were consistent with a $< 5\%$ imbalance (they are shown in Figure 2.3).

A probable cause was realised at the time of writing: the 854 nm repump pulse should have been followed by an 866 nm pulse to repump population decaying into the $D_{3/2}$ manifold. This population would have increased p_2 over time at the expense of p_1, which is indeed visible in Figure 6.6, although it does not fully account for the residual population in p_0.

The populations were fitted to a simple model where the rf rotation angle, phase and a phenomenological dephasing of the $S_{1/2}$ states were floated, with ideal gate operation and repumping. Averaging the results for the two datasets, this yielded an angle and phase of $0.9(1)\pi$ and $0.17(2)\pi$ respectively, and a dephasing factor of 0.3 per round (where the coherences in the density matrix were simply scaled uniformly by 0.7 per round). The fitted curves reproduce the qualitative dynamics, however it is clear from the populations at the first round that the branching ratio of the decay into $|g\rangle$ versus $|d\rangle$ was around 1.8; the model assumed 1. Nonetheless this preliminary investigation demonstrates population

[9]Note that readout is performed by shelving the $|g\rangle$ population using a 729 nm π pulse.

pumping qualitatively similar to that observed in a more complex mixed-species four-ion scheme [101]. Dissipative entanglement could potentially be used to prepare high-fidelity entangled states for future experiments, since in principle it is more robust to many error sources that affect coherent state-preparation methods.

Having presented the single-species gates and protocols performed in this work, the remainder of this chapter discusses mixed-species crystals of beryllium and calcium and gates between them.

6.7 Mixed-species crystals

6.7.1 Cooling

The first mixed-species operations in the setup were performed using beryllium-calcium crystals. A waveform is currently used with an axial CoM and STR frequencies of 2 MHz and 5.1 MHz respectively. Optimising the beryllium-calcium gate relies on initially calibrating the cooling of the crystal using the sidebands as summarised in §4.5. A single round of EIT cooling is optimised to cool both the CoM and the STR modes, with emphasis on the CoM. Next sideband cooling is first carried out on the STR using the FDQ co-90 drive, then on the CoM using the calcium 729 nm. Both modes are cooled to $\bar{n} < 0.1$. Beryllium-calcium crystals have not yet been used in quantum protocols beyond mixed-species MS gates, however their relatively high gate fidelity (discussed below) and independent control of each ion provides a flexible testbed for fundamental QIP experiments.

Beryllium-calcium-beryllium crystals are used in the parity measurement protocols described in the next chapter; the axial motional modes and their parameters are shown in Figure 6.7a. EIT cooling is optimised on the lowest radial mode at 2.5 MHz. This is followed by interleaved sideband cooling on the STR and EGY modes at 4.1 and 4.2 MHz respectively, in which sideband pulses alternate between the STR and EGY modes. Because these modes are narrowly spaced, the sideband pulse power is reduced by $\sim 50\%$ to minimise off-resonant excitation effects. The final step is CoM sideband cooling at 1.56 MHz using calcium. All three modes are cooled to $\bar{n} < 0.1$, and their sideband oscillations are shown in Figure 6.8.

6.7.2 Motional mode coherence

A Ramsey sequence between motional number states can be used to estimate the coherence of different motional modes for various ion crystals [165]. The mode under investigation is ground-state cooled, then the first Ramsey carrier $R_{\pi/2}(\frac{\pi}{2})$ pulse is carried out, preparing the internal-motional state $(|g,0\rangle + |e,0\rangle)/\sqrt{2}$. A motion-subtracting sideband π pulse[10] produces the state $(|g,0\rangle + |g,1\rangle)/\sqrt{2}$. A Hamiltonian $H_M = \hbar\delta_m(t)\hat{a}^\dagger\hat{a}$ representing a fluctuation of the mode frequency $\delta_m(t)$ is applied over a wait time τ, and causes $|g,1\rangle$ to acquire a phase $e^{-i\Phi_m(\tau)}$, where $\Phi_m(\tau) = \int_0^\tau \delta_m(t)dt$. After the wait time another sideband π pulse and subsequent carrier $R_{\pi/2}(\frac{\pi}{2} + \theta)$ pulse produce a state with population

$$p_g(\theta) = \frac{1}{2} - \frac{1}{2}\cos\left(\theta + \Phi_m\right), \tag{6.44}$$

where the form is analogous to the Ramsey sinusoid in Equation 6.2 however the motional mode energy fluctuation $\delta(t)$ is probed. This is used to estimate the mode coherence, which

[10] Red for calcium, blue for the beryllium FIQ.

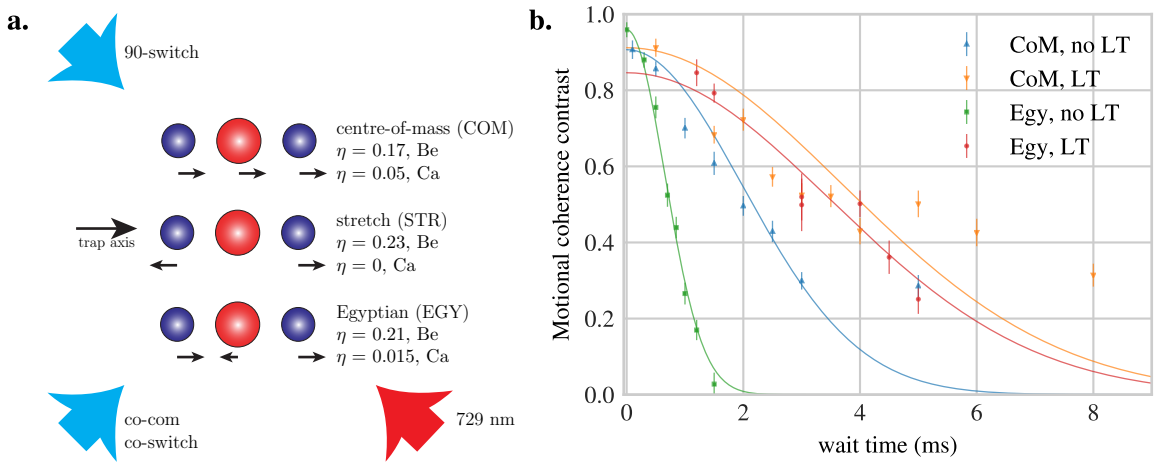

Figure 6.7: Axial motional modes of the mixed-species three-ion crystal and their coherence times. **a.** Motional modes, frequencies, Lamb-Dicke parameters, along with the lasers applied during mixed-species operations discussed later. **b.** Motional coherence contrasts from Equation 6.44, with and without line triggering, fitted to $a \exp\left[-(t/t_0)^2\right]$ where a is the peak value of each dataset The CoM coherence times without and with line triggering are 2.8(2) and 5.2(6) ms; for the EGY they are 0.92(2) and 4.9(2) ms.

is limited by processes causing the confining trap potential to fluctuate, such as unstable electric fields that vary between shots; an investigation of decoherence for different reservoirs can be found in [195]. In the mixed-species setup the radial modes of a single calcium ion have coherence times below 1 ms, however the axial mode coherence time is 15–30 ms as measured in 2014 [81], indicating that the radial rf confinement fluctuates far more than the axial confinement in our system. In similar setups this has been due to rf power fluctuations, which can be reduced using feedback [72].

The coherence time of the stretch (STR) mode was found to be $\sim 300\,\mu$s for calcium compared to ~ 18 ms for the centre-of-mass (CoM) mode[11] The CoM mode was used for multi-qubit gates on each crystal, and for single species the coherence time was roughly half of the single-ion time, as was expected since common-mode noise adds constructively.

Beryllium-calcium-beryllium crystals were also probed, performing the Ramsey on the calcium ion; the three modes are illustrated in Figure 6.7a. A strong improvement from line triggering was observed, lengthening the CoM coherence from 2.8(2) to 5.2(6) ms and the 'Egyptian' mode coherence from 0.92(2) to 4.9(2) ms; fits are shown in Figure 6.7b. The STR mode was not investigated since by symmetry it does not couple to the central calcium ion. The EGY mode is significantly worse than the CoM, potentially due to cross-coupling with the ion radial modes caused by the anharmonicity of the Coulomb potential [165, 138], which are only Doppler-cooled. In addition to their low coherence, the radial modes are currently only Doppler-cooled. This could be investigated using a spin-echo pulse in the motional coherence sequence[12].

The improvement from line-triggering suggests that fluctuating electric fields in phase with the mains are present in the vacuum chamber, either directly in the dc or rf voltages

[11] Although the STR contrast showed revivals at 200 µs intervals when line-triggered. The calcium-calcium coherence was last measured in 2015 and it is unknown if this effect is still present, however.

[12] This has been tried in the setup, however it was noticed at the time of writing that the spin-echo pulse was being carried out on the wrong motional mode.

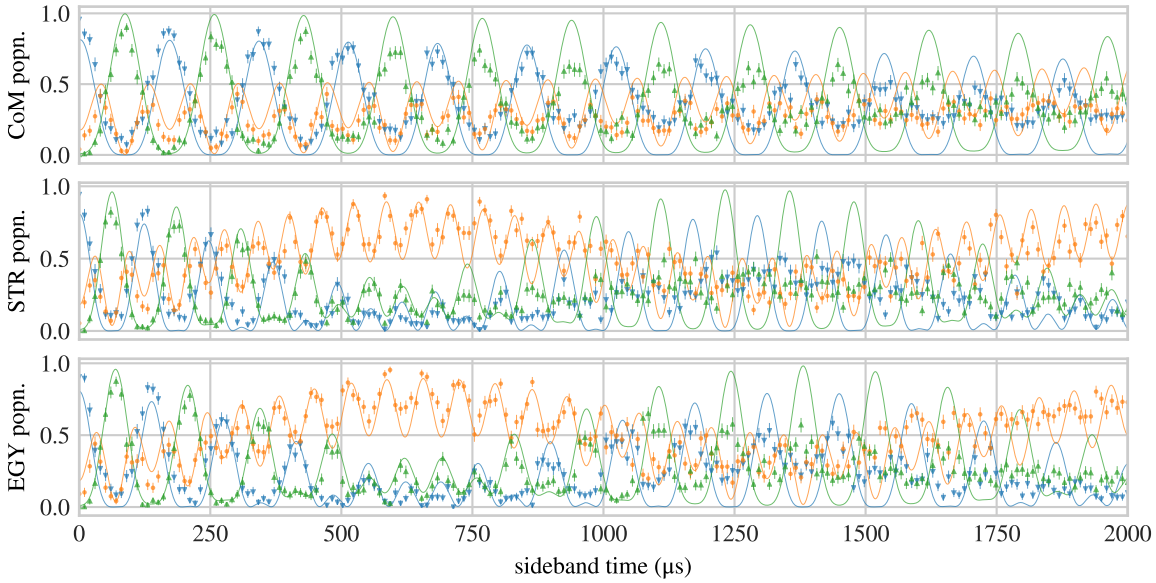

Figure 6.8: Beryllium FDQ red sideband oscillations (after a carrier π pulse) on the axial centre-of-mass, 'stretch' and 'Egyptian' modes of the beryllium-calcium-beryllium crystal. Green upward-pointing triangles, orange circles and blue downward-pointing triangles represent p_2, p_1 and p_0 respectively. They are fitted to Rabi oscillations where the two Rabi frequencies are floated and parameterised as Ω_H and Ω_L, $\Omega_L \leq \Omega_H$, with no decoherence. The Rabi frequency imbalances $\Omega_H/\Omega_L - 1$ for the centre-of-mass, stretch and Egyptian modes are 1.6(9)%, 9.6(2)% and 10.0(2)%.

due to pickup or coupled in via a ground loop. Further experiments are required to identify the source.

6.8 Mixed-species gate optimisation

In mixed-species operation the working principle of the MS gate is the same, however the parameters of both qubit lasers must be tuned. With reference to Equation 6.42 there are two gate detunings $\delta_{c,\mathrm{Be}}, \delta_{c,\mathrm{Ca}}$, intensities $E_{0,\mathrm{Be}}, E_{0,\mathrm{Ca}}$, and sideband phases $\phi_{m,\mathrm{Be}}$ and $\phi_{m,\mathrm{Ca}}$ that must be tuned independently. The differential detunings ω_m and δ_m are set equally for both species.

crystal	t_{gate} (μs)	$\omega_m/2\pi$ (MHz)	$\delta_m/2\pi$ (kHz)	$\delta_{c,\mathrm{Be}}/2\pi$ (kHz)	$\delta_{c,\mathrm{Ca}}/2\pi$ (kHz)	parity contrast (%)	fidelity (%)
Be^+-Be^+	72.5	1.57	14.0	-2.8	-	97.3(4)	97.8(4)
Ca^+-Ca^+	76.9	1.61	13.5	-	-0.1	99.3(6)	99.4(6)
Be^+-Ca^+	58.0	1.96	17.5	-6.0	-1.5	98.9(3)	98.3(9)
Be^+-Ca^+-Be^+	73.3	1.57	14.0	-3.1	-0.3	90.4(6)	93.8(5)

Table 6.2: Typical gate parameters for the various crystals used in the mixed-species setup.

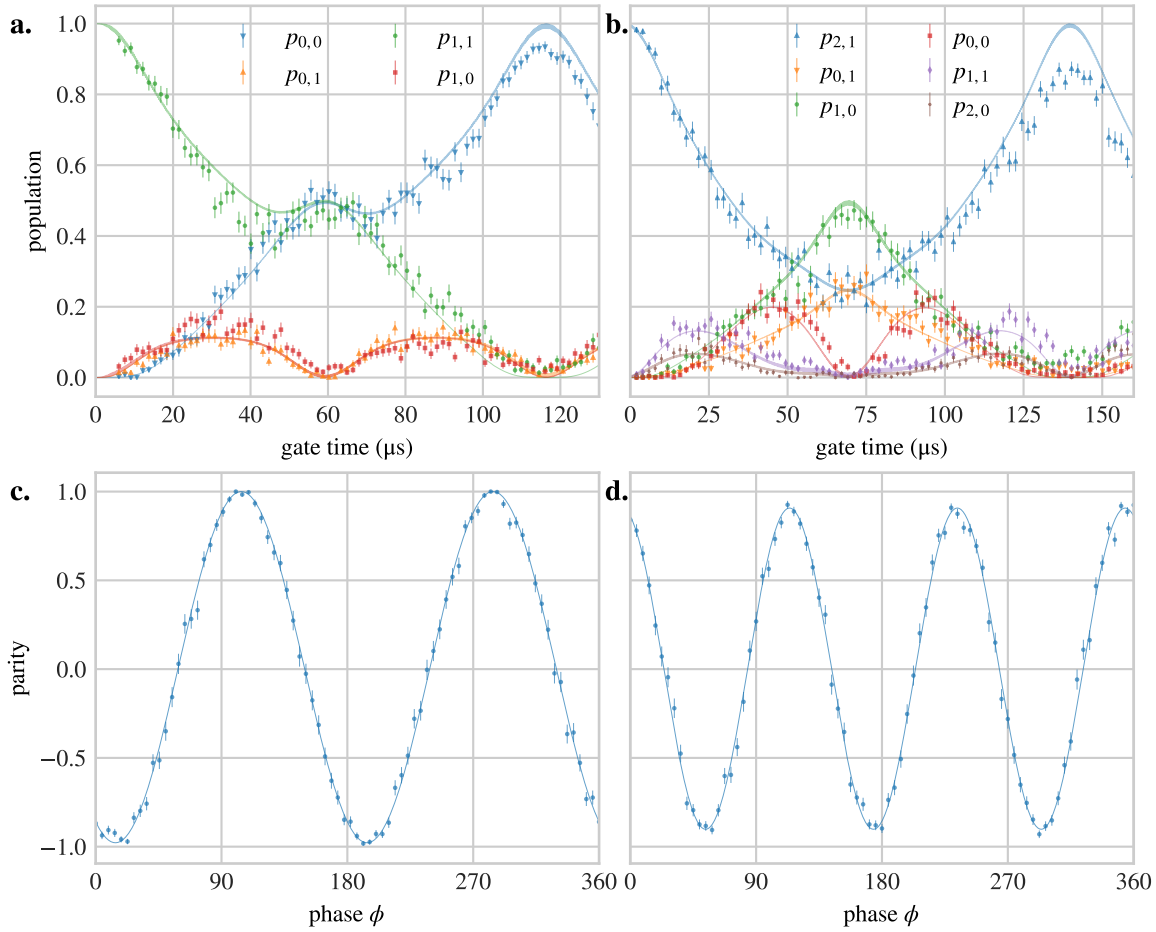

Figure 6.9: Mixed-species two- and three-ion gate populations and parity oscillations. **a.** and **c.** show beryllium-calcium gate time and parity scans, with a parity contrast of 98.9(3)%. **b.** and **d.** show the same for beryllium-calcium-beryllium, with a contrast of 90.4(6)%.

6.8.1 Beryllium-calcium gate

Gate optimisation is begun by adjusting the red and blue sideband amplitudes of the 729 nm laser to be equal in power using an oscilloscope as for single-species gates, with the same repeated for the 90-switch. The relative sideband phases of each drive, ϕ_m from Equation 6.41, are equalised by aligning the amplitude modulation phase of the beatnotes, measured using high-bandwidth photodiodes. Miscalibrating these does not in itself cause gate infidelity, however it misaligns the state-dependent forces of the two species relative to each other, slowing down the gate[13].

The sideband detuning δ_m is coarsely chosen based on the calcium Rabi frequency, which is the limiting factor in the mixed-species gate speed, and initially the common detunings $\delta_{c,\mathrm{Be}}$ and $\delta_{c,\mathrm{Ca}}$ are set to the optimised single-species gate values. The beam powers are tuned to reach a single-ion population of 0.5 for each species at the gate time. This can be achieved over a range of ratios between $E_{0,\mathrm{Be}}$ and $E_{0,\mathrm{Ca}}$, however, and is verified using the overlap of the single-ion $p_{1,\mathrm{Be}}(t)$ and $p_{1,\mathrm{Ca}}(t)$ curves as a function of pulse time.

The gate detunings $\delta_{c,\mathrm{Be/Ca}}$ and beam powers $E_{0,\mathrm{Be/Ca}}$ are fine-tuned using a similar

[13] If $\phi_{m,\mathrm{Be}} = \phi_{m,\mathrm{Ca}} + \pi$ then the forces in fact cancel out.

approach to the single-species gates, with the complication that miscalibration of one species leads to a locally-optimal set of values that causes diminished gate fidelity. This is mitigated by applying single-species π pulses before the gate; when the gate is calibrated near its global optimum, parameter scans with and without the π pulses show the same optimal points. A re-parameterisation to simplify calibration was tried, with $\delta_{c,\mathrm{Be}} = \delta_{c0} + \Delta_c$ and $\delta_{c,\mathrm{Ca}} = \delta_{c0} - \Delta_c$ using common parameters δ_{c0} and Δ_c, such that detunings could be scanned in unison; it requires further testing to establish whether it helps reach the global optimum. As with single-species gates, optimising the populations at multiples of the gate time was useful for fine-tuning. Generally the optimal $\delta_{c,\mathrm{Be}}$ is within several hundred Hz of its single-species value, however $\delta_{c,\mathrm{Ca}}$ shifts by several kHz due to ac Stark shifts from the beryllium Raman beams during the gate.

The optimised beryllium-calcium gate time scan and parity oscillations are shown in Figure 6.9a and c. The parity contrast is 98.9(3)% with an estimated fidelity of 98.3(9)%, which is between the single-species calcium and beryllium gate fidelities.

6.8.2 Beryllium-calcium-beryllium gate

Tuning the beryllium-calcium-beryllium gate is in fact more straightforward than beryllium-calcium, because the beryllium component can be independently calibrated as a standard single-species gate. The calcium parameters ϕ_m, $\delta_{c,\mathrm{Ca}}$ and $E_{0,\mathrm{Ca}}$ are optimised similarly to the beryllium-calcium gate. The time and parity scans are shown in Figure 6.9b and d. A parity contrast of 90.4(6)% was observed, with an estimated fidelity of 93.8(5)%. The multiple populations and associated detection infidelities made fine-tuning at the gate time difficult, and significant improvement was obtained at $2(2n+1)$ gate times, shown in Figure 6.10. Table 6.2 lists the optimal gate parameters for the various crystals.

6.9 Mixed-species gate error sources

The error sources discussed in §6.5.1 also apply to the mixed-species gates. Although mixed-species gate optimisation is similar to the single-species case, systematic miscalibrations or drifts tend to occur more often due to the second set of gate beams and parameters. Currently significant experience is required to maximise the fidelities, as much due to parameter drifts during the optimisation as knowledge of the optimisation procedure. Plans are underway to automate most of the steps, however, to reduce these issues.

Although the beryllium–calcium fidelity is currently between that of single-species beryllium and calcium gates, the three-ion gate fidelity is significantly lower. Because the two-ion beryllium-beryllium fidelity is comparable to the three-ion case, fluctuations in intensity and beam pointing are unlikely to be the source.

Motional decoherence was also investigated as discussed in §6.7.2. The CoM coherence time of 2.8(2) ms is a factor of 40 longer than the MS gate time without line triggering. Using a first-order error model of $\epsilon = 0.686 t_{\mathrm{gate}}/t_{\mathrm{coh}}$ [10], this contributes 1.7% to the infidelity, and will have to be addressed by either line triggering the experiment or removing the source of decoherence. There is no clear reason why the mixed-species gate would be affected more strongly than the single-species case on the same CoM mode, however.

Fluctuations in ϕ_c between beryllium and calcium may be another issue, although the two-ion beryllium-calcium gate is also susceptible to these. Because a three-ion entangled

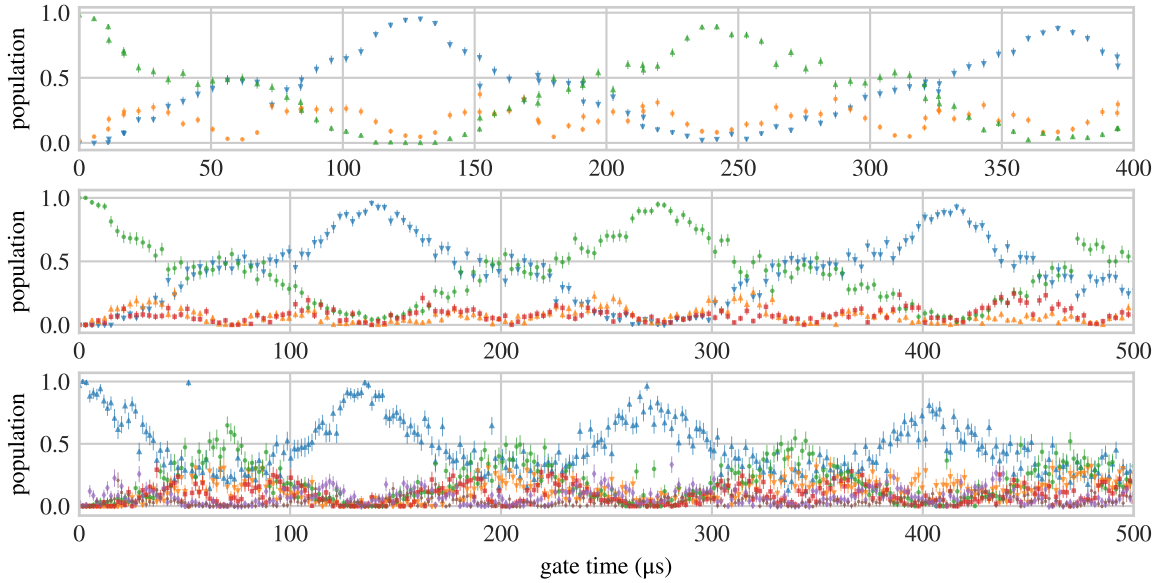

Figure 6.10: Multi-qubit gate populations with the gate time scanned beyond $6t_{\text{gate}}$ for (top to bottom) two-beryllium, beryllium-calcium and beryllium-calcium-beryllium gates. The populations at $6t_{\text{gate}}$ are used for fine-tuning the calibration. Markers have the same meanings as Figure 6.4 and Figure 6.9.

state is created, dephasing may contribute an error 1.5× as large as the two-ion case (cf Equation 6.43). This is still too low to explain the fidelity difference however.

A beat is seen in the sideband Rabi oscillations on the STR and EGY modes of the three-ion crystal, shown in Figure 6.8, that indicates an imbalance of around 10% in sideband Rabi frequencies between the beryllium ions. This suggests either a pseudopotential gradient or anharmonicity in the trap confinement. The pseudopotential has been estimated as a function of trap axial position using the micromotion of a single beryllium ion [105], and the gradient in the experimental zone does not appear significant; the well was moved +20 µm to a region with an even lower gradient and no change in the beat was seen. Anharmonicity also appears too low to explain this effect based on our waveform simulations (which otherwise agree well with observed mode frequencies), however a measurement of the mode frequencies of a beryllium–calcium crystal using the three-ion well ($\omega_{m,\text{CoM}}/2\pi = 1.34\,\text{MHz}$, $\omega_{m,\text{STR}}/2\pi = 3.47\,\text{MHz}$) showed a 4–7 kHz increase in the STR mode and a ~ 0.5 kHz decrease in the CoM when the ion order was reversed. Altering the shim voltages, and thus the radial offset of the crystal in the trap, was found to alter the beat frequency, although it could not be entirely removed, and the shim settings that minimised it were several volts away from the optimal micromotion compensation voltages for a single calcium ion; this suggests that the cause is crystal shape- or motion-related. Additionally, because the beat is not visible on the CoM, it is not certain that the gates themselves are affected by it.

Thus no clear source of infidelity was found for the three-ion gate, although multiple effects were investigated. Ongoing efforts to better characterise the trap, quantitatively estimate the infidelity magnitudes and diagnose the unexplained observations will be discussed in the thesis of Matteo Marinelli [114].

7 Mixed-species parity readout

The previous chapters have introduced and demonstrated many of the building blocks of a mixed-species quantum information processor, including a flexible control system designed for rapid feedback, ion transport and separation, and both single and multi-qubit gates. This chapter presents several mixed-species parity readout and stabilisation experiments, which were carried out by combining these elements to measure and control the parity of two beryllium ions using a calcium ancilla. The role of parity measurement in scalable QIP is introduced, the gate sequence currently used for parity readout is presented, and its calibration, characterisation and fidelity are discussed.

The parity is read out repeatedly multiple times. Since this involves calcium qubit detection which heats the ions, the crystal is cooled sympathetically using the calcium ion. The outcomes are used to feed back on the beryllium using global rotations and stabilise different parity subspaces for up to 50 rounds, as well as preparing the Bell states deterministically. Feedback using only global rotations is not entirely universal, however, thus differential rotations were implemented by transporting the crystal to unevenly illuminate the beryllium ions. With this addition, all four Bell states were deterministically prepared and stabilised.

The chapter concludes with a discussion of the infidelity sources in the multi-round stabilisation experiments, potential diagnostics in the readout correlations and techniques to improve the fidelity.

7.1 Introduction and context

As mentioned in §1.3, quantum error correction (QEC) can be implemented by using multiple physical qubits to encode a smaller number of logical qubits [182, 188]. Many QEC schemes are based on stabiliser codes, which rely on the detection of qubit errors by carrying out multi-qubit Pauli measurements [139, 84]. Although it is not strictly a QEC code, this is illustrated well by the operation of the three-qubit bit-flip code. A logical qubit $|\psi_{LZ}\rangle = a\,|g_{LZ}\rangle + b\,|e_{LZ}\rangle$ is represented by three physical qubits according to $|g_{LZ}\rangle \equiv |ggg\rangle$ and $|e_{LZ}\rangle \equiv |eee\rangle$. If a bit flip corrupts the state to $|\psi_{LZ}\rangle = a\,|gge\rangle + b\,|eeg\rangle$, for example, the parity measurement $S_Z^{(1,2)} \equiv \hat{\sigma}_z^{(1)} \otimes \hat{\sigma}_z^{(2)}$ on the first two ions will yield 1, since the ion states are correlated, whereas $S_Z^{(1,3)}$ will yield -1 since the states are anticorrelated. $S_Z^{(1,2)}$ and $S_Z^{(1,3)}$ form an error *syndrome* whose measurement pinpoints the bit flip error to a single ion, which can be corrected with an $R_\pi(0)$ rotation on a single ion. 'Phase flips', where a qubit acquires a π phase shift between its basis states $|g\rangle$ and $|e\rangle$ due to dephasing noise as introduced in §6.1 in the context of a Ramsey experiment, can similarly be detected by measuring parity in the $\{|+\rangle \equiv (|g\rangle + |e\rangle)/\sqrt{2},\ |-\rangle \equiv (|g\rangle - |e\rangle)/\sqrt{2}\}$ basis using $S_X^{(i,j)} = \hat{\sigma}_x^{(i)}\hat{\sigma}_x^{(j)}$ on a logical qubit encoded using $|\psi_{LX}\rangle$, $|g_{LX}\rangle \equiv |+++\rangle$ and $|e_{LX}\rangle \equiv |---\rangle$, and a correction sequence of $R_{\pi/2}(0)R_\pi(\frac{\pi}{2})R_{\pi/2}(\pi)$ on a single ion can be used. Correcting bit and phase flips is sufficient for a universal error-correcting code that can preserve an arbitrary quantum state, of which the Shor code is a well-known ex-

	$E_X = -1$	$E_X = +1$
$E_Z = -1$	$\|\Psi_-\rangle = \frac{\|ge\rangle - \|eg\rangle}{\sqrt{2}}$	$\|\Psi_+\rangle = \frac{\|ge\rangle + \|eg\rangle}{\sqrt{2}}$
$E_Z = +1$	$\|\Phi_-\rangle = \frac{\|gg\rangle - \|ee\rangle}{\sqrt{2}}$	$\|\Phi_+\rangle = \frac{\|gg\rangle + \|ee\rangle}{\sqrt{2}}$

Table 7.1: Eigenvalues of the Bell states for the operators $S_Z \equiv \hat{\sigma}_z^{(1)} \otimes \hat{\sigma}_z^{(2)}$ and $S_X \equiv \hat{\sigma}_x^{(1)} \otimes \hat{\sigma}_x^{(2)}$. The parity measurement sequences map +1 (-1) eigenvalues to the state $|e\rangle$ ($|g\rangle$) in calcium.

ample [182]; it concatenates the two codes above to implement a single $|\psi_{LZ}\rangle$ qubit, with 9 qubits in total. QEC codes in which individual errors do not multiply within single encoded qubits are known as *fault-tolerant* [53, 190].

The readout of an n-qubit code syndrome can be recast as a series of n-qubit Pauli measurements, of which the two-qubit parity measurements are an example. In practice these are carried out by entangling the code qubits with an ancilla then projectively measuring it, which simultaneously projects the code qubits into an eigenstate of the parity operator. A set of readout results is processed in real-time to pinpoint a specific error, and a suitable correction operation is applied. Limited QEC codes have been demonstrated in several quantum systems [33, 29, 160], which have often relied on protocols where the code qubits were brought out of the code-space for syndrome measurement, however for universal QEC, ideal n-qubit measurements are required [53].

In this work, two-qubit parity stabiliser measurements $S_Z = \hat{\sigma}_z^{(1)} \otimes \hat{\sigma}_z^{(2)}$ and $S_X = \hat{\sigma}_x^{(1)} \otimes \hat{\sigma}_x^{(2)}$ are performed. The Bell states are their eigenstates, with the eigenvalues E_Z and E_X given in Table 7.1. The act of measuring the parity in a single basis projects the qubits into an eigenstate of S_Z or S_X which is a superposition or mixture of Bell states of a particular eigenvalue in that basis, thus collapsing a potential superposition of eigenstates with opposite eigenvalues. Hence S_Z followed by S_X projects the qubits into a single Bell state with particular $\{E_Z, E_X\}$ eigenvalues.

Parity readout has been carried out in several platforms, including trapped ions [11, 177, 140] and nitrogen-vacancy centres [149], including experiments in which up to three rounds of feedback were conditioned on the result of measurement [162, 34]. For indefinite stabilisation of a quantum system, however, dissipative operations such as ancilla measurement must couple minimally to the code qubits. Up to three successive measurements have been performed in trapped ions by 'hiding' qubit states in levels that do not couple to the detection beam using individual ion addressing [128], however this approach has significant overhead and may be difficult to scale. The work presented in this thesis uses a mixed-species paradigm without these disadvantages, whose performance can go beyond previous implementations. Its key features are that the calcium ancilla and beryllium code qubit operations are spectroscopically decoupled from one another; that the quantum information is encoded in the beryllium FIQ, giving it a coherence time which is $> 1000\times$ longer than a single stabiliser readout round; and that the internal and motional degrees of freedom in the system relevant to high-fidelity quantum operations are re-initialised between rounds using the ancilla, such that the parity readout fidelity remains close to its original value even after tens of rounds. Together these elements facilitate two-qubit schemes with the main characteristics of quantum error correction to be investigated experimentally.

7.2 S_Z and S_X measurement protocols

To measure S_Z, the unitary gate sequence U_{S_Z} shown in Figure 7.1a is carried out, which with the addition of a calcium detection forms the M_{S_Z} stabiliser readout operation. The core elements are a three-ion MS gate, π pulses and second MS gate; as discussed in §6.4.1 the MS gates mediate pairwise interactions between the qubits. A $R_{\pi/2}(\pi/2)$ pulse before the first MS gate prepares calcium in $|+\rangle$. For both S_Z and S_X measurements, the MS-π-MS sequence imprints a phase on calcium of $+\pi/2$ ($-\pi/2$) if $E_{Z/X} = +1(-1)$ for the beryllium state, and the second $R_{\pi/2}(0)$ rotation maps the phase to $|e\rangle$ ($|g\rangle$). The central calcium $R_\pi(0)$ pulse is required to ensure the $\pi/2$ phases from each MS gate add constructively, while the beryllium $R_\pi(0)$ pulses decouple the sequence from minor gate miscalibrations.

As discussed in §6.5.1 the beryllium gate phase ϕ_c is not stable due to the path difference fluctuations between the co-com and 90-switch Raman beams, and if we read out the parity over multiple rounds, the beryllium gate basis will fluctuate relative to the calcium basis. This is avoided by rotating the MS-π-MS operation using extra $R_{\pi/2}(-\pi/2)$ and $R_{\pi/2}(\pi/2)$ pulses, rendering it diagonal in the $\{|g\rangle, |e\rangle\}$ measurement basis[1] on the beryllium qubits and thus insensitive to phase fluctuations [93, 50]. The relative phase between the MS gates and single-qubit rotations is still important, however, and must be constant for every shot of the parity measurement even in an extended sequence. As mentioned in §6.5, the red and blue MS gate sidebands have a shared detuning δ_c to cancel out Stark shifts; over time this causes a phase difference $(t - t_{\text{ref}})\delta_c$ between the single-qubit and MS gate pulses due to the automatic phase accumulation described in §3.3.8. This is bypassed by shifting the reference time t_{ref} to the beginning of the first beryllium co-90 pulse, then returning it back to zero when co-co single-qubit pulses are carried out, since these must remain coherent throughout multiple rounds of the sequence.

The combined protocol U_{S_Z} implements the following mapping on input states in the measurement basis:

$$|ggg\rangle \leftrightarrow |gge\rangle\,, \qquad |eeg\rangle \leftrightarrow |eee\rangle \tag{7.1}$$

$$|egg\rangle \leftrightarrow -|egg\rangle\,, \qquad |geg\rangle \leftrightarrow -|geg\rangle \tag{7.2}$$

where the first two qubits are beryllium and the third is calcium, and the omitted states are unaffected by the sequence. U_{S_X} is implemented by adding co-co $R_{\pi/2}$ pulses before and after the U_{S_Z} sequence, rotating the the beryllium parity state into the S_Z basis. Note that the phases chosen in this sequence are not unique, and other combinations can be found to perform a functionally equivalent mapping.

7.2.1 Phase calibration

The U_{S_Z} block is calibrated with $|ggg\rangle$ as the input state, and the phase reference t_{ref} is set to the beginning of the first co-90 $R_{\pi/2}$ pulse as shown in Figure 7.1a. The ac Stark shift on the beryllium FIQ from the calcium 729 laser is negligible as mentioned in §6.1.3, however the calcium qubit experiences a Stark shift of several kHz from the beryllium Raman lasers. For this reason the beryllium gates and phases in U_{S_Z} are calibrated initially, which includes running the three-ion MS gates without the 729 nm beam, using the beryllium populations. Next the calcium pulses are successively introduced, their phases are calibrated using the live beryllium-calcium readout correlations, and U_{S_Z} is verified as described below in §7.2.2.

[1] In other words, a correlated $\hat{\sigma}_z$ rotation on the qubits.

All calcium and beryllium phases are referred to the first MS gate, which is kept at 0 rf phase. The detailed calibration steps are listed in §B.2.2.

U_{S_X} is calibrated by inputting an even-parity state, and adjusting the co-co $R_{\pi/2}$ phases before and after U_{S_Z} to perform a quantum nondemolition operation by maximising the even-parity populations of beryllium at the output. Because U_{S_X} is not diagonal in the measurement basis, it is significantly more sensitive to all of the beryllium $R_{\pi/2}$ phases.

7.2.2 Verification using beryllium–calcium correlations

To verify the parity measurement, parity superpositions are prepared using an initial two-beryllium MS gate U_{MS} with a gate spin phase ϕ_c, as defined in Equation 6.36, which is sandwiched between two co-90 $R_{\pi/2}(\phi_c - \pi/2)$ and $R_{\pi/2}(\phi_c + \pi/2)$ pulses to make the gate diagonal in the measurement basis and thus insensitive to co-90 phase fluctuations between shots [93], as is done for the parity measurement discussed in the previous section. Two co-co $R_{\pi/2}$ pulses before and after this operation, at phases $\phi_g - \pi/2$ and $\phi_g + \pi/2$, rotate the basis back to $\cos(\phi_g)\hat{\sigma}_x + \sin(\phi_g)\hat{\sigma}_y$ as for the original MS gate, however the phase ϕ_g of this decoupled gate is now referenced to the co-co beam. The decoupled gate is followed by a parity pulse with phase ϕ_p to produce a superposition of Bell states $|\psi_{\text{in}}\rangle = R_{\pi/2}(\phi_p)U_{MS}(\phi_g)$ according to

$$|\psi_{\text{in}}\rangle = \frac{|\Phi_+\rangle}{4}\left(G - ie^{-i\phi_g}G^*\right) + \frac{|\Phi_-\rangle}{4}\left(H + ie^{-i\phi_g}H^*\right) - \frac{|\Psi_+\rangle}{2}\left(ie^{i\phi_p} + e^{-i(\phi_p+\phi_g)}\right) \quad (7.3)$$

where $G = 1 - e^{2i\phi_p}$ and $H = 1 + e^{2i\phi_p}$. Setting $\phi_p = -\phi_g/2 - \pi/4$ produces the triplet state $|\Psi_+\rangle$, and setting it to $-\phi_g/2 + \pi/4$ produces the superposition of Bell states $|B\rangle = |\Phi_+\rangle\left(1 - ie^{-i\phi_g}\right)/2 + |\Phi_-\rangle\left(1 + ie^{-i\phi_g}\right)/2$.

Next U_{S_X} or U_{S_Z} were applied, followed by detection on all qubits. We wish to predict the calcium and beryllium populations in the measurement basis when scanning ϕ_p. Verifying U_{S_Z} is straightforward, since $E_Z = 1$ (calcium in $|e\rangle$) for $|B\rangle$ and -1 (calcium in $|g\rangle$) for $|\Psi_+\rangle$. The E_Z expectation values as a function of phase are

$$\langle E_Z = 1\rangle = p(|\Phi_+\rangle) + p(|\Phi_-\rangle) = |\langle\Phi_+|\,\psi_{\text{in}}\rangle|^2 + |\langle\Phi_-|\,\psi_{\text{in}}\rangle|^2 = \frac{1}{2} + \frac{1}{2}\sin(2\phi_p + \phi_g) \quad (7.4)$$

$$\langle E_Z = -1\rangle = p(|\Psi_+\rangle) + p(|\Psi_-\rangle) = |\langle\Psi_+|\,\psi_{\text{in}}\rangle|^2 + 0 = \frac{1}{2} - \frac{1}{2}\sin(2\phi_p + \phi_g), \quad (7.5)$$

which are out-of-phase parity oscillations. For every shot, E_Z is either +1 or -1, and the calcium readouts should match the beryllium state; thus we expect to observe $\langle E_Z = 1\rangle = p_{0,0} + p_{2,0}$, where $p_{m,n}$ is the probability of m beryllium and n calcium ions being bright. Similarly, $\langle E_Z = -1\rangle = p_{1,1}$. Figure 7.1d shows the four correlations between $E_Z = \pm 1$ as measured with beryllium and $M_{S_Z} = \pm 1$ as measured with calcium; the populations $p(E_Z = 1, M_{S_Z} = 1) = p_{0,0} + p_{2,0}$ and $p(E_Z = -1, M_{S_Z} = -1) = p_{1,1}$ are shown by blue circles and red squares respectively, while the other populations are shown as triangles. For comparison, the population $p_0 + p_2 = 1 - p_1$ inferred from reading out beryllium directly after state preparation is shown as violet stars[2].

Verifying U_{S_X} requires further consideration, since S_X does not act in the measurement basis. $|\psi_{\text{in}}\rangle$ only produces $E_X = -1$ when $|B\rangle = |\Phi_-\rangle$ at $\phi_g = \pi/2$, otherwise a pure

[2] This is half the magnitude of the parities plotted in the previous chapter, which show $p_0 + p_2 - p_1$.

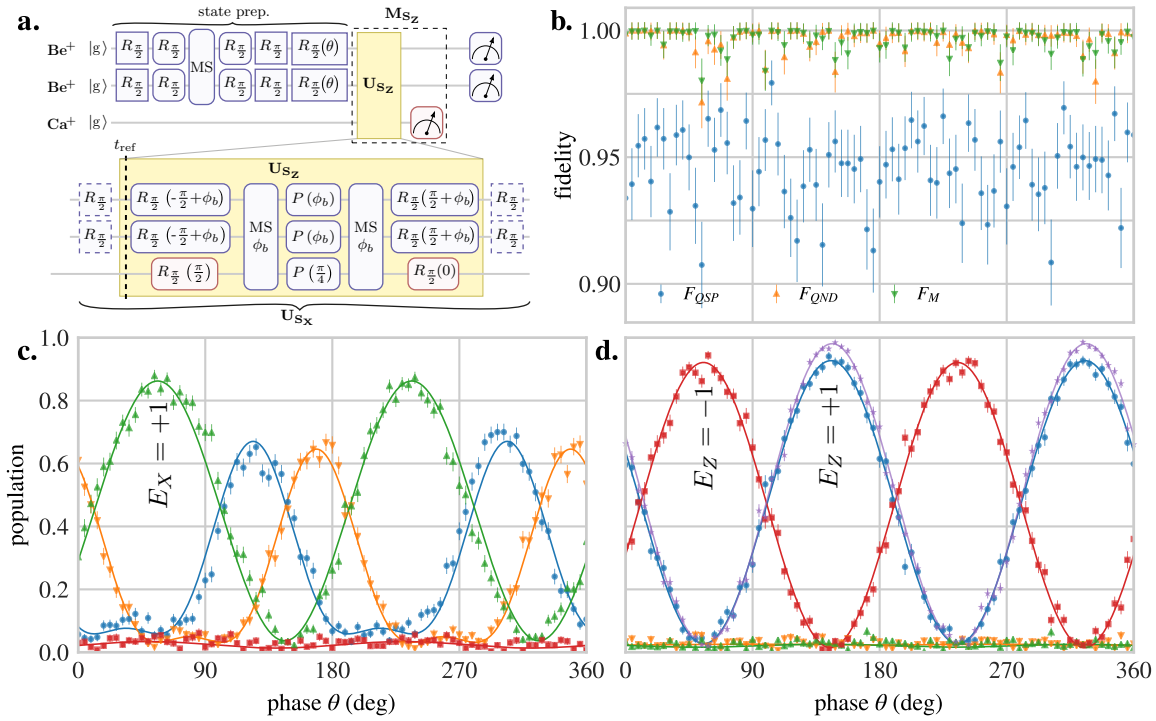

Figure 7.1: Single-round stabiliser readout and verification. **a.** The gate sequence used for the stabiliser readouts M_{S_Z} and M_{S_X}, consisting of unitary operations U_{S_Z} and U_{S_X} followed by a calcium readout. Rounded (sharp-cornered) rectangles indicate pulses using motion-sensitive (insensitive) beams, dashed borders indicate optional pulses. The parity measurement is verified using a phase-insensitive MS gate and $R_{\pi/2}(\phi_p)$ pulse to prepare a two-ion beryllium input state $|\psi_{\text{in}}\rangle$ according to Equation 7.3, where θ in the plots is offset from ϕ_p by an uncompensated ac Stark shift. **d.** Correlations between the calcium result and the parity, measured by detecting the beryllium ions after M_{S_Z}. Blue circles (red squares) show the probability of observing both that calcium is measured in $|e\rangle$ ($|g\rangle$) and that beryllium is measured in the $E_Z = +1$ ($E_Z = -1$) eigenspace, with the anticorrelated populations shown as orange and green triangles at the bottom. Violet stars show the input state parity as measured in a separate experiment where M_{S_Z} is not executed, plotted as $(1 + \langle S_Z \rangle)/2$. **c.** The same measurement, carried out for M_{S_X}. Green upward-pointing triangles indicate $E_X = +1$, corresponding to odd beryllium parity for the $|\Psi_+\rangle$ state, correlated with calcium measured in $|e\rangle$, and blue circles (orange downward-pointing triangles) indicate calcium measured in $|e\rangle$ for the $|\Psi_+\rangle$ state ($|g\rangle$ for the $|\Psi_-\rangle$ state) respectively. The curves are fitted to Equation 7.6, Equation 7.7 and Equation 7.8 respectively with $\phi_g = 0$ and additional contrast and dc offset parameters floated, with contrasts of 82(1)%, 89(2)%, 89(1)% and compared to their ideal values. **d.** Fidelity estimates of M_{S_Z} as a function of phase, using the data from **c** and Equation 7.9 – Equation 7.11.

$E_X = -1$ state cannot be obtained with this scheme. Because $|\Psi_+\rangle$ is an eigenstate of S_X, Equation 7.5 is also valid, thus

$$p(E_X = 1, M_{S_X} = 1) = p_{1,0} = [1 - \sin(2\phi_p + \phi_g)]/2. \tag{7.6}$$

This is shown by the green upward-pointing triangles in Figure 7.1c, which agree with the red squares in the S_Z plot as expected. Unlike S_Z however, $p_{1,0} \neq \langle E_X = 1 \rangle$ because $p(|\Phi_+\rangle)$ also contributes to the $\langle E_X = 1 \rangle$ population.

To predict the other populations, we consider the other two Bell states, which both have

even parity in the measurement basis. From Equation 7.3 their populations are

$$p(|\Phi_+\rangle) = |\langle \Phi_+ | \psi_{\text{in}}\rangle|^2 = \sin^2(\phi_p)\left[\frac{1}{2} + \frac{1}{2}\sin(2\phi_p + \phi_g)\right] \tag{7.7}$$

$$p(|\Phi_-\rangle) = |\langle \Phi_- | \psi_{\text{in}}\rangle|^2 = \cos^2(\phi_p)\left[\frac{1}{2} + \frac{1}{2}\sin(2\phi_p + \phi_g)\right]. \tag{7.8}$$

When calcium is bright, $E_X = -1$, which can only be produced by $|\Phi_-\rangle$. This has even parity in the measurement basis, thus $\langle E_X = -1\rangle = p_{0,1} + p_{2,1} = p(|\Phi_-\rangle)$, which is shown as orange downward-pointing triangles in Figure 7.1c and fitted to Equation 7.8. When calcium is dark, $E_X = +1$, which can be produced by either $|\Phi_+\rangle$ or $|\Psi_+\rangle$, hence we can say $\langle E_X = +1\rangle = p(|\Psi_+\rangle) + p(|\Phi_+\rangle) = p_{0,0} + p_{1,0} + p_{2,0}$, and from Equation 7.6 we can conclude that $p_{0,0} + p_{2,0} = p(|\Phi_+\rangle)$, which is shown as blue circles in Figure 7.1c and fitted to Equation 7.7.

In these experiments $\phi_g \simeq 0$, thus from Equation 7.8 there is no value of ϕ_p at which a pure $E_X = -1$ state is observed[3]; the maximum value of the $p(|\Phi_\pm\rangle)$ curves is ~ 0.728.

7.2.3 Fidelity estimation

An ideal round of parity measurement $M_{S_{Z/X}}$ is a form of quantum nondemolition (QND) measurement, in that it does not cause a back-action on the system, here defined as the beryllium ions, if they are in an eigenstate or superposition of equal-eigenvalue eigenstates of $S_{Z/X}$. A QND measurement should fulfil three conditions: the measurement outcome should match the input state of the system; the measurement should not alter the system state; and repeated measurements should give the same result, in other words the measurement prepares and keeps the system in a measurement eigenstate [155]. These correspond to three fidelities we can estimate for a single round of parity measurement $M_{S_{Z/X}}$. The measurement fidelity F_M quantifies how well the calcium measurement agrees with the input parity state. The quantum nondemolition fidelity F_{QND} estimates the agreement of the beryllium parities with each other before and after M_{S_Z}, i.e. how likely the parity is to be corrupted by the measurement. The quantum state preparation fidelity F_{QSP} is the conditional probability that a given calcium measurement projects the beryllium parity into the correct subspace.

We use the probabilities of measuring the beryllium parity corresponding to $E_{Z/X} = \pm 1$ both before and after $M_{S_{Z/X}}$, given by $p^{\text{in}}_{\pm 1}$ and $p^{\text{out}}_{\pm 1}$, and the probabilities of the calcium outcomes $p_{|g\rangle}$ and $p_{|e\rangle}$. For $M_{S_{Z/X}}$ the fidelities are defined as [11, 155]

$$F_M = \left(\sqrt{p^{\text{in}}_{+1} p^{m}_{|e\rangle}} + \sqrt{p^{\text{in}}_{-1} p^{m}_{|g\rangle}}\right)^2 \tag{7.9}$$

$$F_{QND} = \left(\sqrt{p^{\text{in}}_{+1} p^{\text{out}}_{+1}} + \sqrt{p^{\text{in}}_{-1} p^{\text{out}}_{-1}}\right)^2 \tag{7.10}$$

$$\begin{aligned} F_{QSP} &= p_{|e\rangle} p^{\text{out}}_{+1||e\rangle} + p_{|g\rangle} p^{\text{out}}_{+1||g\rangle} \\ &= p^{\text{out}}_{+1\&|e\rangle} + p^{\text{out}}_{-1\&|g\rangle} \end{aligned} \tag{7.11}$$

where $p^{\text{out}}_{+1||e\rangle}$ ($p^{\text{out}}_{-1||g\rangle}$) is the *conditional* probability of observing $E_{Z/X} = +1$ ($E_{Z/X} = -1$) in beryllium having measured $|e\rangle$ ($|g\rangle$) in calcium, and $p^{\text{out}}_{+1\&|e\rangle}$ ($p^{\text{out}}_{-1\&|g\rangle}$) is the *joint* probability

[3] Unfortunately the possibility of calibrating the co-co $\pi/2$ pulses to obtain $\phi_g = \pi/2$, and thereby testing a pure $E_X = -1$ eigenstate, was investigated only at the time of writing.

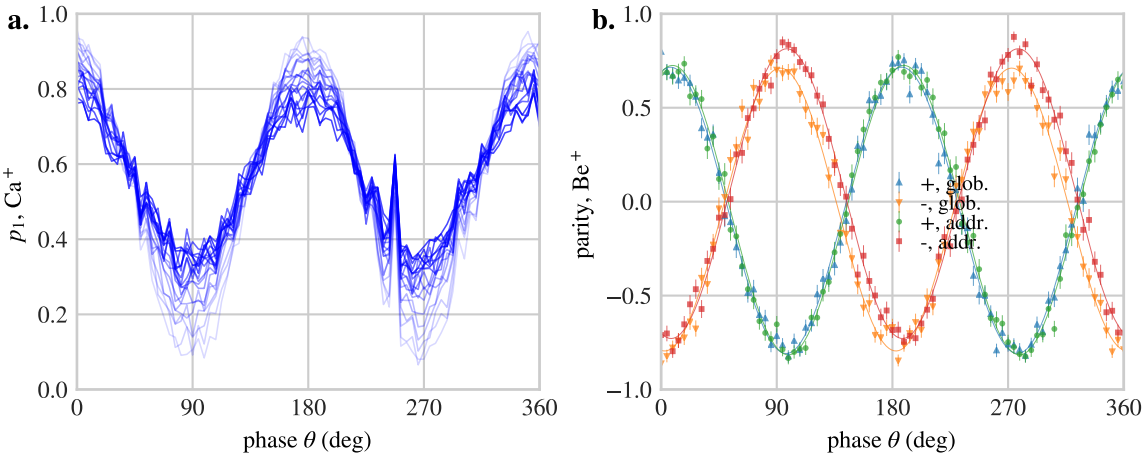

Figure 7.2: Repeated S_Z measurement, and even Bell state preparation using global and addressed rotations. **a.** Calcium population as a function of state-preparation phase ϕ as shown in Figure 7.1a, from 1 to 20 rounds of S_Z measurement. The faintness of the lines represents the shot number, with the faintest lines from the beginning of the measurement loop. The spike is due to a momentary unlocking of the Raman doubling cavity. **b.** Comparison of beryllium parity after preparing even Bell states using global rotations versus addressed rotations. The parity pulse phase is scanned, similar to the plots in Figure 7.1. $|\Phi_+\rangle$ and $|\Phi_-\rangle$ have contrasts of 76.2(8) and 75.3(9)% for global rotations, and 77(8) and 77(7)% for the addressed rotations.

of observing these events. For M_{S_Z} the joint probabilities are simply $p_{0,0} + p_{2,0}$ and $p_{1,1}$, corresponding to the red squares and blue circles in Figure 7.1d respectively.

The fidelity estimates for M_{S_Z} are shown in Figure 7.1b, with the mean fidelities $\bar{F}_M = 99.6(4)\%$, $\bar{F}_{QND} = 99.6(5)\%$ and $\bar{F}_{QSP} = 94.6(2)\%$. The value of F_{QSP} is most relevant for QEC, and Figure 7.1b shows that its level does not vary significantly with phase. As discussed in the previous chapter, two-beryllium and three-ion MS gates have fidelities of 97.8(4)% and 93.8(5)% respectively, and $\bar{F}_{QSP}$ is broadly consistent with these being the major sources of infidelity. The fidelities for M_{S_X}, were not estimated, since unlike for M_{S_Z} the beryllium outcomes do not unambiguously indicate $E_X = \pm 1$, however a lower bound is given by the odd-parity sinusoid contrast of 82(1)%. A more detailed error budget will be presented in the upcoming thesis of Matteo Marinelli [114].

7.3 Repeated measurements

Next the stabiliser measurements M_{S_Z} and M_{S_X} were applied repeatedly on an initial state prepared as above, with a beryllium detection only after the final round. A single round consisted of $M_{S_{Z/X}}$ followed by re-initialisation of the calcium ion, then recooling of the motional modes to mitigate the heating due to photon scattering during calcium detection and maintain high gate fidelities. To weakly Doppler-cool the crystal and reduce this heating, the 397 nm π and σ detection beams were red-detuned by 9 MHz and 8 MHz from resonance respectively. Next two rounds of EIT cooling were applied, the first optimised on the EGY and the second on the CoM mode, for 150 µs and 200 µs respectively; this was followed by 10 cycles of 729 nm sideband cooling on the CoM. Afterwards the CoM sideband oscillations were comparable to those after the cooling in the beginning of the sequence.

For repeated applications of M_{S_Z} on a superposition of $|\Phi_+\rangle$ and $|\Phi_-\rangle$ ($E_Z = +1$) and

feedback type, population	$E_Z = +1$	$E_Z = -1$	$E_X = +1$	$E_X = -1$	fit type
none, Ca	8.0(2)	8.5(3)	-	11.3(2)	exp
global, Be	0.95(1)	0.39(3)	-	-	lin
global, Ca	1.17(1)	0.28(2)	-	-	lin
global, BeCa	1.16(1)	0.51(3)	-	-	lin
addressed, Ca	0.33(1)	0.10(1)	0.22(1)	0.23(1)	lin

Table 7.2: Parity subspace open-loop and stabilised decay rates for different stabilisation subspaces. E_Z (E_X) values were obtained from a sequence of M_{S_Z} (M_{S_X}) measurements as shown in Figure 7.5a. No-feedback values are those obtained from the decay of an initially-prepared subspace without any feedback operations. Population types indicate the data used for the fit; BeCa are the correlated populations plotted in Figure 7.3. Beryllium data were not taken for the addressed stabilisation. Open-loop and closed-loop data were fitted to $a\exp(-\gamma x)+b$ and $-\gamma x+c$ respectively, where $a \simeq 0.5$ and $b \simeq 0.5$. γ values are listed in the table, in units of percent per measurement round.

$|\Psi_+\rangle$ ($E_Z = -1$), the calcium populations are shown as red diamonds (green circles) in the left plot of Figure 7.5b, and green circles for M_{S_Z} acting on $|\Psi_+\rangle$ ($E_X = -1$) in the right plot. These were fitted to exponentials, with decay constants listed in Table 7.2. For multiple rounds both the $E_Z = -1$ and $E_Z = +1$ populations rise slightly, which we attribute mainly to population leakage into the dark $F = 2$ states in the beryllium manifold due to the Raman beams. Figure 7.2a shows the calcium population as a function of parity pulse phase for M_{S_Z} up to 20 measurement rounds for an earlier dataset. The reduction in contrast indicates that the purity of the original Bell state is reduced due to each measurement cycle having a finite chance of corrupting the parity, that is significantly higher than the F_{QND} obtained in §7.2.3. Because a single round has a low probability of corrupting the parity, the mean correlations between the calcium readouts for successive rounds[4] decrease slightly from 95% to 85% over the course of the sequence; this decrease indicates that the fidelity of the M_{S_Z} operation degrades over time. The reasons are not fully understood, however they are likely to be systematic effects such as thermal AOM drifts due to duty cycles.

Similar measurements were carried out for M_{S_X}, however the Stark shift compensation (see below) had not yet been implemented and thus the basis phase precesses with the measurement number, making the plots difficult to interpret.

7.3.1 Stark shift compensation

Because U_{S_X} is not diagonal in the beryllium measurement basis (unlike U_{S_Z}), it is sensitive to the relative phase between $|g\rangle$ and $|e\rangle$ acquired due to ac Stark shifts that occur during its constituent pulses, whose frequencies are set to ω_0 and not $\omega_0 + \omega_{ss}$ (see §6.1.2). During sequential M_{S_X} measurements, the phases of the initial and final co-co $R_{\pi/2}$ rotations are shifted by $m\phi^{ss}_{(S_X)}$, where m is the number of times U_{S_X} has previously been carried out and $\phi^{ss}_{(S_X)}$ is the phase precession caused by a single round of M_{S_X}. $\phi^{ss}_{(S_X)}$ is measured using the procedure in §6.1.2 by carrying out only the co-com or the 90-switch components of U_{S_X} within a Ramsey sequence and measuring the Ramsey phase shift. It is verified empirically by carrying out the scan shown in Figure 7.1c for several rounds of M_{S_X} and adjusting

[4]The correlation between two shots is 1 if calcium was read out in the same state, and -1 otherwise.

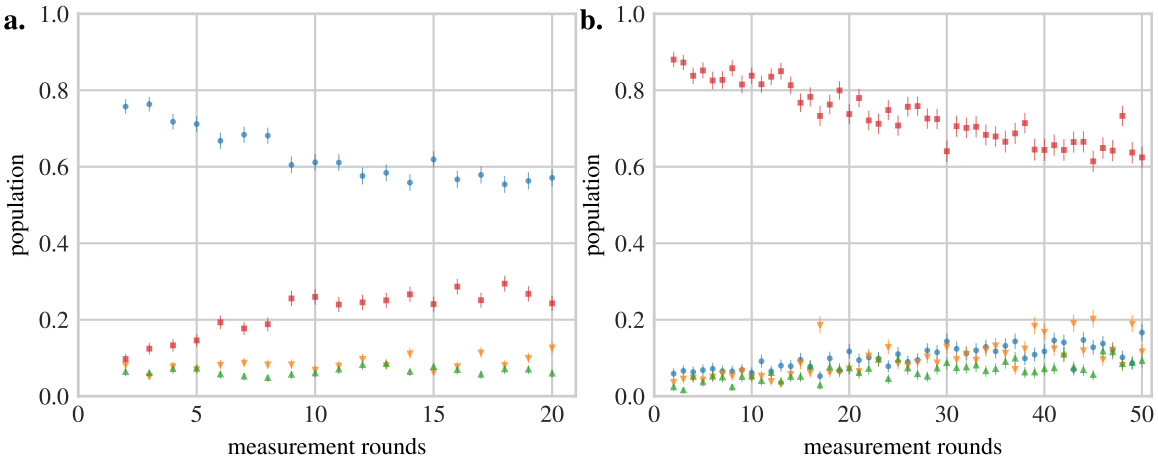

Figure 7.3: Beryllium parities for S_Z subspace stabilisation, with **a.** $E_Z = +1$ and **b.** $E_Z = -1$ the desired states. Blue circles (red squares) show the probability of observing both that calcium is measured in $|e\rangle$ ($|g\rangle$) and that beryllium is measured in the $E_Z = +1$ ($E_Z = -1$) eigenspace, as in Figure 7.1d. The singlet is a dark state of the global correction operations, and its population accumulates over time due mainly to beam pointing fluctuations causing unequal rotations on the beryllium ions; this hastens the decay of $E_Z = +1$ (note the different x axes). Linear fits yield slopes of 1.16(1)% and 0.51(3)% per round for $E_Z = +1$ and $E_Z = -1$ respectively.

$\phi^{ss}_{(S_X)}$ to maximise the correlations for the even-parity input state, while ensuring they do not fall for other input states.

7.4 Conditional feedback

The population in the repeated measurements decays primarily due to infidelity in the quantum operations, especially the multi-qubit gates, rather than fundamental qubit decoherence. This is also likely to be the limiting factor in a larger-scale trapped-ion quantum information processor making use of qubits protected from decoherence, whether in a field-insensitive state or a logical decoherence-free subspace. By correcting the parity errors introduced by the measurement protocol itself, we therefore demonstrate multiple elements required for more general correction of errors introduced by QIP protocols. The correction uses real-time feedback. Unlike the feedback used in the Bayesian scheme of §6.2 where a pulse phase is altered over a near-continuous range of values, discrete parity errors are corrected by carrying out one of several discrete correction operations.

7.4.1 Global rotations

Pulses that applied the same rotation to both beryllium ions, according to the two-qubit rotation matrix Equation 6.22, referred to here as global rotations, were the first kind of feedback investigated. As can be seen from Equation 6.24–Equation 6.26, global rotations can be used to map all the Bell states to one another except for the singlet $|\Psi_-\rangle$. The $E_Z = \pm 1$ parity subspaces were thus stabilised, by repeatedly applying M_{S_Z} and carrying out $R_{\pi/2}(\pm\pi/2)$ rotations to invert the S_Z parity when needed. The populations are shown in Figure 7.3. The $E_Z = -1$ subspace can be stabilised for over 50 measurement rounds, however $E_Z = +1$ decays twice as rapidly to a mixture of parities. This is because the singlet is a dark state of the protocol, unaffected by the global rotations, and its population increases

the $E_Z = -1$ populations while decreasing those of $E_Z = +1$. In one sense, the rise of the $E_Z = -1$ population can be seen as the undesired counterpart to the dissipative pumping scheme described in §6.6.2, where the singlet is also a dark state; here the dissipation is introduced via the calcium detection, and again there is no way to distinguish between the singlet and triplet states.

Global rotations can also be used to deterministically prepare the even Bell states, by first carrying out M_{S_Z} with a conditional $R_{\pi/2}$ pulse to convert the triplet into an even state. Next M_{S_X} is carried out, and if the desired outcome did not occur, E_X is inverted using a $R_\pi(\pm\pi/4)$ or $R_\pi(\pm 3\pi/4)$ pulse[5]. The parity curves for the two Bell states are shown in Figure 7.2b. The Bell states were generated from a product state created with an initial $R_{\pi/2}$ state preparation pulse.

Feedback using conditional global rotations was a large step forward in our experimental capabilities, however it could not be used to alter the singlet population. A more universal approach that overcomes this is discussed next.

7.4.2 Addressed rotations

A second form of feedback was implemented to differentially rotate the ions, carrying out the operations $C_Z = -I^{(1)} \otimes \hat{\sigma}_x^{(2)}$ and $C_X = -I^{(1)} \otimes \hat{\sigma}_z^{(2)}$ to invert the S_Z and S_X parity respectively. To implement C_Z the ion well is shifted rightwards along the trap axis by a distance p of 11–14 µm in several tens of microseconds, then an FIQ co-co pulse is carried out to rotate the right beryllium ion by π and the left ion by 2π. The Rabi frequency of an ion in a Gaussian beam with width[6] w is given by $\Omega(p) = \Omega(0)\exp(-p^2/2w^2)$, thus the ratio for two ions separated by a distance $d \approx 10\,\mu\text{m}$ is

$$\begin{aligned} R = \Omega(p - d/2)/\Omega(p + d/2) &= \exp - \frac{(p - d/2)^2 - (p + d/2)^2}{2w^2} \\ &= \exp \frac{pd}{w^2}, \end{aligned} \tag{7.12}$$

where $R \geq 1$, which yields $p = (\log R)w^2/d$. Thus for this scheme a suitable position p can be found for any ratio, beam width or ion spacing.

C_Z is optimised experimentally by generating multiple transport waveforms with final locations spaced at 50 nm intervals. The optimal well is chosen based on the two-ion Rabi oscillations. For single-ion Rabi frequencies of Ω and $(2+\delta)\Omega$, the two-ion dark population p_0 is

$$p_0 = \cos^2(\Omega t)\cos^2[(2+\delta)\Omega t], \tag{7.13}$$

which has two peaks around $t = 3\pi/\Omega$ whose relative height difference scales with δ (red regions in Figure 7.4a). Experimentally the heights of these peaks were equalised using a 2D scan of waveform index against Rabi pulse time, with an optimal waveform shown in Figure 7.4b. This method is robust to detection infidelity and sensitive enough to achieve $\delta < 0.5\%$. Once the transport distance is calibrated, the co-co pulse time is calibrated

[5] In principle the first $R_{\pi/2}$ phase can be chosen to already prepare $|\Phi_+\rangle$ or $|\Phi_-\rangle$ deterministically, however M_{S_X} is still required to convert $|\Phi_\pm\rangle \leftrightarrow |\Phi_\mp\rangle$.

[6] Neglecting the Raman beam process; namely that the Rabi frequency in fact scales as the product of two beam intensities. In the mixed-species setup the co-com and 90-sw beams were of slightly different widths, thus a more accurate model would consist of the ratio of 4 Gaussian profiles.

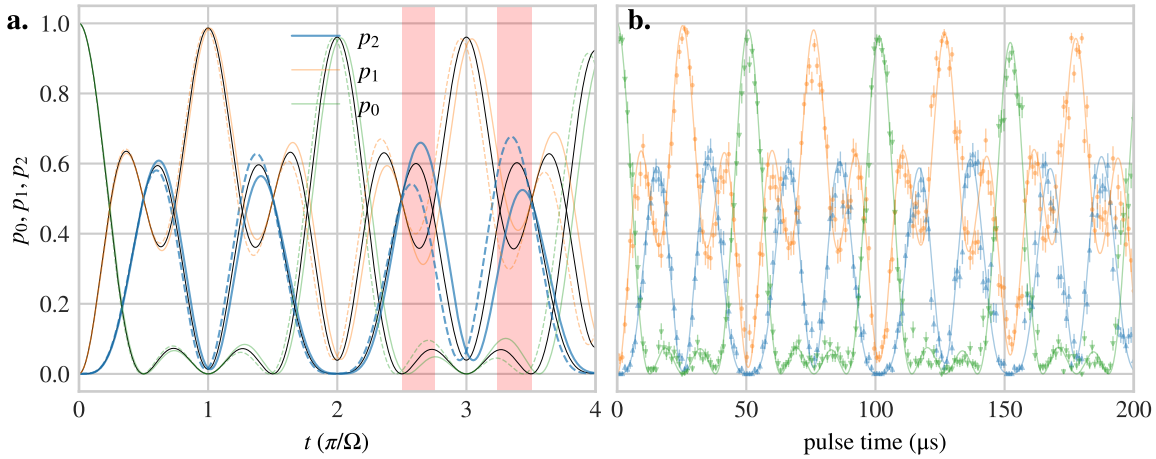

Figure 7.4: Position calibration for differential rotation. **a.** Red regions show the p_0 populations used to calibrate the transport distance; solid (dashed) coloured lines show $\delta = 3.5\%$ (-3.5%) and black lines show $\delta = 0$. **b.** Extended oscillations at an optimised transport distance. Fits to Equation 7.13 with an additional decay term yield $\delta = 0.37(3)\%$.

to maximise p_1 for an odd number of C_Z operations; a peak population above 99.8% is obtained.

C_X consists of $R_{\pi/2}(0)$, $C_Z(\pi/2)$ and $R_{\pi/2}(\pi)$ using the FIQ co-co beams. This carries out an addressed rotation $-I^{(1)} \otimes \hat{\sigma}_z^{(2)}$, which is diagonal in the measurement basis and thus phase-insensitive. It is calibrated by first applying the $R_{\pi/2}$ pulses without C_Z such that they map $|gg\rangle$ to $|gg\rangle$, then adding C_Z and adjusting its phase to create $|ee\rangle$.

7.5 Universal stabilisation

Using addressed rotations, both the S_Z (S_X) parity subspaces were stabilised by applying C_Z (C_X) every time M_{S_Z} (M_{S_X}) yielded an undesired outcome. In addition to the execution fork chosen based on each detection, the new Stark shift compensation phases for the M_{S_X} measurement as described in §7.3.1 are also calculated and sent in real-time. The corrections use the co-co only, thus are applied before re-cooling begins without loss of fidelity; this way any additional excitation due to transport[7] is also cooled. The results with and without feedback are shown in Figure 7.5. The initial state for the stabilisation was produced by applying an $R_{\pi/2}$ beryllium rotation to $|gg\rangle$, to prepare an equal superposition of $E_Z = \pm 1$ eigenstates.

The open-loop decays were fitted to exponentials. Due to the lack of information on the closed-loop stabilised populations in the many-round limit, linear fits were used to obtain approximate decay rates; the rates are listed in Table 7.2, with an average decay of around 0.3% per measurement round. This is lower by over 20× than the average open-loop decay rates.

The solid lines shown in Figure 7.5 are from a simulation incorporating leakage from the beryllium qubit subspace, which was run by Matteo Marinelli and will be discussed further in his upcoming thesis [114]. It does not fully account for the decay, although introducing

[7] This has not been measured but is likely to be minimal, since the ions are transported at a slow speed below 1 m/s over a short distance.

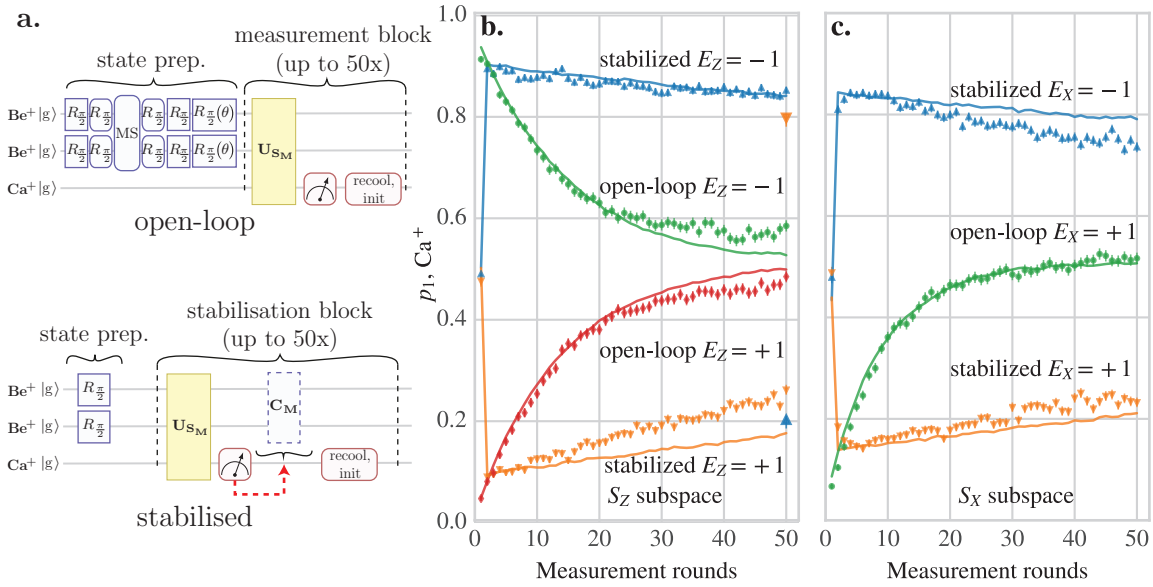

Figure 7.5: Parity subspace stabilisation in **a.** the S_Z and **b.** the S_X basis, using the unitaries U_{S_Z} (U_{S_X}) and the feedback operations C_Z (C_X) respectively. The large upward (downward)-pointing triangles at the right of the S_Z plot show the parity of the Be^+ populations at the end of the $E_Z = -1$ ($E_Z = +1$) stabilisation, defined as $(1 + \langle S_Z \rangle)/2$ or $1 - p_1$. Exponential decays of the form $a \exp(-\gamma t) + b$ are fitted to the open-loop curves, and linear fits are performed on the closed-loop curves since there is insufficient decay to orthogonally estimate the three parameters of the exponential fit. The results are listed in Table 7.2. Solid lines are produced by Matteo Marinelli using a Monte-Carlo simulation, to be discussed in [114].

a gradual decline in the parity readout fidelity of 0.06% per measurement round results in closer agreement. This is within what could be caused by duty cycle effects in the AOMs. The leakage rate is expected to be comparable to the error rate from Raman scattering within the qubit subspace. Once other sources of gate error discussed in §6.5.1 and §6.9 are eliminated, this will become dominant. It cannot be corrected using error-correcting approaches within the qubit subspace, which may motivate the use of ions lacking hyperfine structure in future protocols [194, 25].

Using addressed rotations a single round of Bell state preparation can now deterministically produce any of the four Bell states. It consists of M_{S_Z} followed by M_{S_X}, resulting in the four possible outcomes $\{E_z = \pm 1, E_X = \pm 1\}$. Depending on which Bell state is the target, either no correction, C_Z, C_X or both C_X and C_Z are applied. This sequence can be repeated multiple times to stabilise the Bell states, as shown in Figure 7.6.

The four Stark shift phases $\phi^{ss}_{(S_Z)}$, $\phi^{ss}_{(S_X)}$, $\phi^{ss}_{(C_Z)}$ and $\phi^{ss}_{(C_X)}$ are initially measured as described in §7.3.1, and the co-co $R_{\pi/2}$ phases of U_{S_X} and C_Z are updated before every measurement round based on the number of times each measurement and correction operation has been applied so far. The phases are fine-tuned by preparing Bell states then running the sequence with different combinations of always-on or always-off feedback operations, and adjusting the phases to maximise the expected populations for each kind of sequence[8].

[8] This was not implemented for the single round of feedback based on global rotations, however in that case the correction phases were empirically optimised within the sequence, thereby taking the Stark shifts into account.

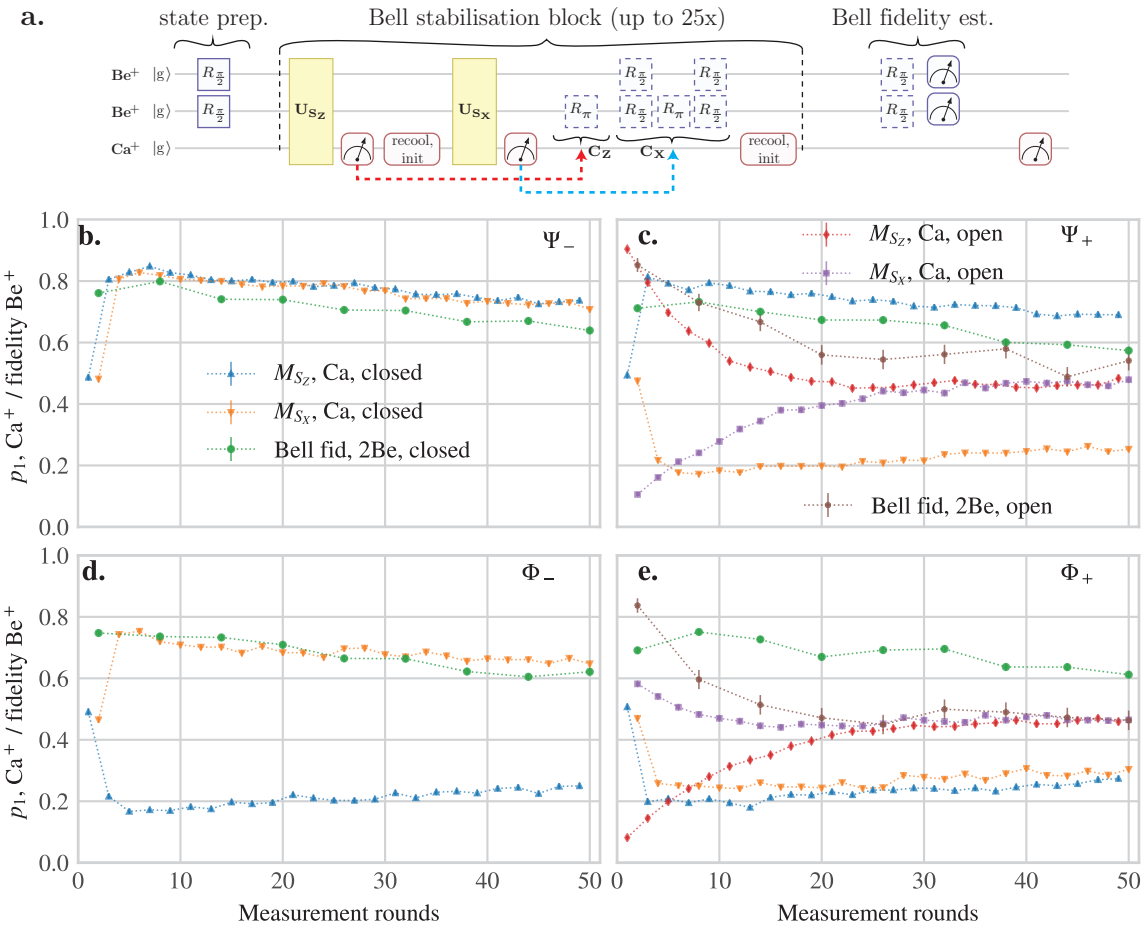

Figure 7.6: Bell state stabilisation using the pulse sequence shown in **a.**, in which the results of S_Z and S_X measurements are used to apply C_Z, C_X to target a particular Bell state. **b** – **e.** Evolution of Ca^+ outcome probabilities over 25 Bell state stabilisation blocks, a total of 50 parity measurements. Beryllium open- and closed-loop fidelities are also shown, obtained by running a certain number of blocks then estimating the resultant Bell state fidelity as described in the text. In **e.** the open-loop M_{S_X} calcium population begins close to 0.5 because as discussed earlier the state preparation could only produce a superposition of $|\Phi_+\rangle$ and $|\Phi_-\rangle$. The behaviour of the open-loop beryllium fidelity in **c.** is not understood; we suspect it is due to a parity pulse phase miscalibration. Open- and closed-loop calcium population and fidelity curves were fitted as described in the caption of Figure 7.5, with values listed in Table 7.3.

Every third round of measurement of the complete syndrome, the Bell state fidelities were estimated by measuring the beryllium populations in three orthogonal bases. For $|\Phi_\pm\rangle$ the first was directly in the measurement basis to obtain $\langle S_Z\rangle$, and the other two were after co-co $R_{\pi/2}$ pulses whose phase was set to prepare either $|\Psi_+\rangle$ or $|\Phi_\mp\rangle$ (i.e. the opposite-parity state), thereby providing $\langle S_X\rangle$ and $\langle S_Y\rangle$ (corresponding to estimating the contrast of parity oscillations from the extremal values). To estimate the fidelities of $|\Psi_\pm\rangle$, a $C_Z = -I_1 \otimes X_2$ operation preceded the analysis to initially convert them to $|\Phi_\pm\rangle$. Although the triplet could have been analysed without C_Z, this scheme minimised the extra calibrations required and was easier to manage and debug. Note that the $R_{\pi/2}$ analysis pulse phases were also dynamically updated based on the preceding operations in the sequence. The fidelity estimates are shown as green and brown circles in Figure 7.6. After a single round of Bell state stabilisation, we observe a mean fidelity of 73.1(4)%, which falls to

feedback type, population	meas. basis	$\|\Phi_+\rangle$	$\|\Phi_-\rangle$	$\|\Psi_+\rangle$	$\|\Psi_-\rangle$	fit type
none, Ca	S_Z	9.1(2)*	9.1(2)*	15.9(6)	-	exp
none, Ca	S_X	-	-	8.3(3)	-	exp
addressed, Ca	S_Z	0.16(2)	0.16(2)	0.26(1)	0.24(2)	lin
addressed, Ca	S_X	0.12(2)	0.18(2)	0.18(2)	0.24(1)	lin

Table 7.3: Bell state open-loop and stabilised decay rates. $|\Phi_\pm\rangle$ and $|\Psi_\pm\rangle$ values were obtained from a sequence of interleaved M_{S_Z} and M_{S_X} measurements as shown in Figure 7.6a No-feedback values are those obtained from the decay of an initially-prepared states without any feedback operations. Stars next to the $|\Phi_\pm\rangle$ values for S_Z indicate that the input state was a superposition of $|\Phi_+\rangle$ and $|\Phi_-\rangle$. Open-loop and closed-loop data were fitted to the same equations as in Table 7.2.

61.3(4)% after 25 rounds (i.e. 50 parity measurements). The calcium parity contrasts and beryllium fidelities slowly decrease over time due to state leakage, similar to the subspace stabilisation, however they have not yet been analysed in detail. It will be interesting to determine whether other error sources such as AOM duty cycle effects contribute to the decay using a more complete model of the protocol.

7.6 Correlations in the ancilla measurements

Currently the Bell state fidelities are significantly lower than expected based on the gate fidelities alone. Each experimental shot of the subspace and Bell state stabilisation protocols provides multiple calcium results, whose correlations were analysed to better understand the stabilisation and its current limitations. The two-point correlations of successive M_{S_Z} and M_{S_X} measurements are categorised and averaged based on which feedback operations had been applied between the measurements, with the results shown in Figure 7.7. Because is calibrated to act in the measurement basis, a correction C_X should commute with M_{S_Z}, with the correlation between M_{S_Z} outcomes in an $M_{S_Z} - M_{S_X} - C_X - M_{S_Z}$ sequence nominally equal to those from $M_{S_Z} - M_{S_X} - M_{S_Z}$. A similar argument can be made for M_{S_X} correlations with C_Z operations. The correlations with a nominally-commuting feedback event are on average $\sim 10\%$ lower than without feedback, implying that the commutation error is worse than the $\sim 1\%$ level expected from calibration experiments. The error mechanism is currently unknown, but could be investigated by monitoring the correlations live during calibration.

We can obtain information about the feedback fidelities by considering the correlations conditional on a feedback event having occurred, labelled as $S_Z : C_Z$ and $S_X : C_X$ in Figure 7.7; here we consider only the S_Z basis. We assume each M_{S_Z} measurement performs perfect parity readout ($F_{QSP} = 1$) however has a nonzero chance of corrupting the parity ($F_{QND} < 1$). Given an undesired outcome followed by a perfect correction (with correction fidelity $F_C = 1$), the next measurement M_{S_Z} will yield the desired outcome, thus producing anticorrelated result and leading to a correlation probability of 0. This is not the case in Figure 7.7, with the probability being around 0.3 to 0.4. In the context of the current fidelities, this can be understood in the following way. Because $F_{QSP} \simeq 0.95$ and $F_{QND} \simeq 0.995$, detecting the undesired parity is much more likely to occur due to a faulty detection than a genuine parity error. In this case the C_Z correction in fact corrupts the parity state, and the next readout is most likely to again show the undesired parity. The original incorrect detection thus causes two feedback events, with a correlated and uncorrelated pair

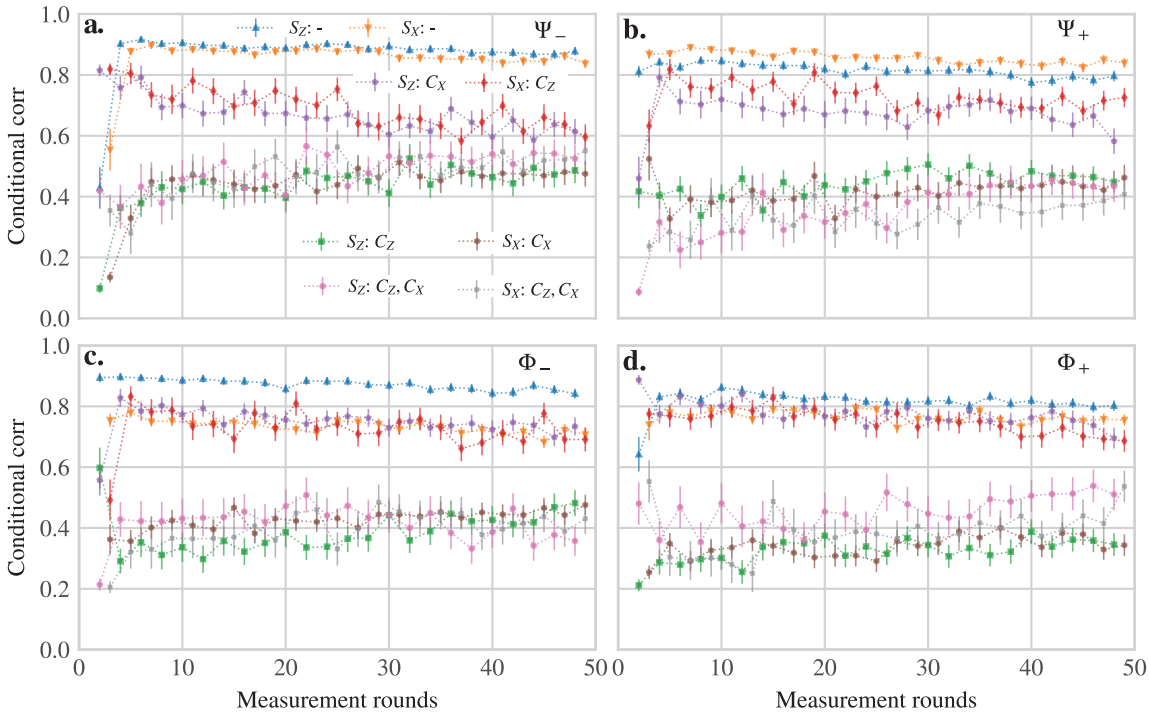

Figure 7.7: Correlations between successive same-basis parity measurements, post-selected based on the feedback that took place between them. Values of 1 (0) represent perfect (anti-) correlation. Upward- (downward-) pointing triangles show the M_{S_Z} (M_{S_X}) correlations where no intermediate feedback took place, i.e. for M_{S_Z} the sequence was $M_{S_Z} - M_{S_X} - M_{S_Z}$. Violet pentagons (red diamonds) indicate the correlations where only a commuting operation was applied between the M_{S_Z} (M_{S_X}) measurements, for example the sequence $M_{S_Z} - M_{S_X} - C_X - M_{S_Z}$; for perfectly commuting operations these would be level with the triangles. Green squares (brown hexagons) indicate where only a non-commuting correction was applied, for example $M_{S_Z} - M_{S_X} - C_Z - M_{S_Z}$, and pink circles (grey crosses) indicate where both C_Z and C_X were applied.

of detections leading to an average correlation of 0.5 instead of 0.

We can analyse a more realistic case based on the previously estimated fidelities. When each round begins, the system is in one of four states; the beryllium parity is either in the desired subspace or not, and the previous M_{S_Z} measurement either gave a desired outcome or not. We can express this as a four-element probability vector $\{P_{00}, P_{01}, P_{10}, P_{11}\}$ where the first index is is 1 if the current parity state is desired, and the second is 1 if the detection outcome was desired. Considering the measurement and feedback process as a Markov chain, we can obtain the steady-state probabilities of each case.Each round has a probability of not corrupting the parity F_{QND}, a probability of measuring the actual parity F_{QSP} and a probability of successful parity inversion, whether or not it is correct to do so, F_C.

We can construct a binary tree of possibilities and sum up the paths by which each P

term is obtained, to obtain the steady-state relation

$$\begin{pmatrix} P_{11} \\ P_{10} \\ P_{01} \\ P_{00} \end{pmatrix} = \begin{bmatrix} A & BF_C + AI_C & B & AF_C + BI_C \\ C & DF_C + CI_C & D & CF_C + DI_C \\ B & CF_C + DI_C & C & DF_C + CI_C \\ D & AF_C + BI_C & A & BF_C + AI_C \end{bmatrix} \begin{pmatrix} P_{11} \\ P_{10} \\ P_{01} \\ P_{00} \end{pmatrix}, \tag{7.14}$$

$$\text{where } A = F_{QSP}F_{QND}, \quad B = F_{QSP}I_{QND}, \quad C = I_{QSP}F_{QND}, \quad D = I_{QSP}I_{QND}$$
$$\text{and } I_{QND} = 1 - F_{QND}, \quad I_{QSP} = 1 - F_{QSP}, \quad I_C = 1 - F_C. \tag{7.15}$$

By applying the matrix repeatedly to a normalised initial condition, the steady-state populations can be obtained. For the fidelities $F_{QSP} = 0.95$, $F_{QND} = 0.995$ and $F_C = 0.99$, these are $P_{11} = 0.89, P_{10} = 0.047, P_{01} = 0.007, P_{00} = 0.052$. In this case P_{11} is the probability of remaining in the target parity subspace. It is highly sensitive to F_{QSP}, and for a value of $F_{QSP} = 0.98$ it rises to 0.95. Because $I_{QSP} \simeq 10 I_{QND}$, a significant gain could be obtained by simply repeating $M_{S_Z/X}$ multiple times and carrying out feedback only in the case of a 'majority vote'. Information such as the probability of feedback events and correlations can also be extracted from such a model, and it could be used to optimise the protocol performance.

This experiment has demonstrated some of the general elements of multiple-round stabilizer readout and correction, as required for performing QEC on a larger-scale system. The major current limitation is the gate infidelities themselves rather than the qubit, which has a coherence time significantly beyond the 60–80 ms that 50 measurement rounds currently take. These arise from motional decoherence during the multi-qubit gates, spin decoherence of calcium, and pulse miscalibrations. The demonstration of conditional feedback opens several opportunities for quantum state control outside QEC, including quantum metrology [96], quantum gate teleportation [54] or measurement-based quantum computing [159, 88].

8 Summary and outlook

8.1 Experimental control and automation

This thesis summarised the mixed-species experimental setup and the basic principles of trapped-ion QIP, then presented the M-ACTION experimental control system, largely developed in this work to perform experiments requiring rapid real-time feedback and/or computation. Its architecture and design choices were discussed in detail. M-ACTION is currently used in all of the setups in the TIQI group, and has proven flexible enough to carry out a wide range of complex experiments [80, 106, 3, 38, 78, 35, 102, 79, 48, 99, 136]. The M-ACTION system was also used to perform the experiments in this thesis, including multiple experiments relying on its real-time features. Additionally, auxiliary elements of the experimental control such as the digital PID controllers, Raspberry Pis and flexible standalone rf sources are important to carrying out robust and repeatable experiments.

8.1.1 Calibration and Bayesian schemes

Laboratory work in trapped-ion QIP has reached a level of complexity where the time needed to manually calibrate the many elements of a multi-ion mixed-species experiment begins to dominate the daily schedule. A short-term goal in the experimental control will be scripting as many mundane calibration steps as possible, including many of the Rabi and Ramsey-based calibrations discussed in Chapter 4 and Chapter 6. Primary candidates for automation are carrier frequency, $\pi/2$ and π time calibrations, as well as motional mode frequencies and possibly cooling using the approaches introduced in §4.5. An already-useful initial step would be automatically carrying out a series of 'sanity checks' to alert the user when parameters are sub-optimal. For example, estimating the bright population after a π pulse calibrated by the user, once every 10 minutes, and showing an alert if it is above a threshold.

Calibrations where heuristics and experience are required, such as multi-qubit gates or micromotion compensation, generally require longer to reach the maximum achievable fidelities. Once automatic calibration is implemented using existing scan-and-fit techniques, Bayesian approaches to calibrating multidimensional parameter spaces such as these are expected to speed these processes up significantly. One promising 'hybrid' approach is applying Bayesian schemes to infer system properties that normally would require one or more manual scans; an example of this is producing a real-time error signal for the frequency detuning or the pi time of a qubit that is updated several times a second to assist in beam alignment or polarisation optimisation. Integrating Bayesian schemes into the overall calibration workflow is likely to be highly productive. Bayesian schemes are not without drawbacks, however; they are often unintuitive to debug or optimise, and are very sensitive to the assumptions made when deriving their models. Nonetheless their use will be necessary as experiments grow more complex, and in fact as experimental fidelities improve, Bayesian models will agree increasingly well with experimental data.

As trapped-ion QIP systems grow, laboratory control must shift to a higher level of ab-

straction, where most of the everyday calibrations are handled by low-level routines without user intervention; the techniques described here will probably all contribute to such a system.

8.1.2 Future of experimental control systems

Every experimental control system has inherent tradeoffs, and several upgrades are already planned for M-ACTION as discussed in §3.3.11. The trapped-ion QIP experimental paradigm will change in the coming years, with technology such as integrated fibres [],beam delivery [116] and in-vacuum electronics [134] coming online. Increasing commercial investment makes it appear likely that there will be a proliferation of incompatible control systems as QIP systems evolve from laboratory experiments to profitable commercial devices in the coming decades, due to the competitive advantages this business model will offer each manufacturer [1].

A metric likely to become more important is the 'cost per qubit' of the classical control; it is too early to anticipate the form of the control system in a scalable QIP architecture, however once thousands of physical qubits become practical, any inefficiencies in classical control systems will rapidly be trimmed away. Trapped ions are a convenient platform for experimental control at their current stage of development, since their relatively slow 'clock speeds' of microseconds [187, 189, 191, 122] allow inexpensive and well-tested control hardware to be used, which lessens the cost per qubit of the control hardware in a scalable architecture compared to other QIP platforms.

8.2 Transport and separation

Ion transport and separation have also been investigated in this work, as well as the use of mixed-species crystals, single- and multi-qubit gates and their calibration. Currently the beryllium-calcium-beryllium crystal offers the most interesting experimental platform in the mixed-species setup. The range of possible protocols would be dramatically expanded with the capability of separating the crystal or managing two mixed-species crystals in the trap simultaneously, however, and the major obstacle to achieving this is currently three-ion mixed-species separation. As discussed in §5.5.3 the separation appears to be fundamentally limited by stray fields in the splitting zone, inaccurate electrode models and/or a sub-optimal separation trajectory.

Experiments at the time of writing have shown that there is a measurably different stray field in the splitting zone from the experimental zone. This fact, along with careful characterisation of the trap electrode moments and the separation performance with single-species crystals, will likely provide insight into the mixed-species separation issues. Additionally, upgrades to the DEATH boards currently in progress will allow more flexible experiments to be carried out; operations such as sweeping the local axial field during separation without affecting it globally will be possible without regenerating the waveforms. In my experience, separation has been straightforward to optimise once the major source of heating is addressed; this is likely to be the case with mixed-species crystals.

[1] This is a common pattern among emerging technologies, with historical examples including the commercial computer systems of the 1960s–1980s, desktop software over recent decades, and recently the various hardware and software making up the 'internet of things'.

Once the three-ion separation is optimised, there are no barriers to performing protocols with 5 or more mixed-species ions using the current mixed-species setup.

8.3 Mixed-species fidelity improvements

Several techniques specific to mixed-species experiments have been discussed in this work, including mixed-species detection and how detection infidelities are handled in the setup as well as various sources of gate infidelity. Currently the three-ion gate infidelity limits the fidelity of the parity measurements discussed in the previous chapter, however an investigation into possible causes has so far not yielded any conclusive results.

The infidelity of the mixed-species two-ion gate is several times lower than the three-ion gate, which rules out many potential error sources. The anharmonicity noticed for a beryllium-calcium crystal, the improvement of motional mode coherence when line-triggered, and the three-ion sideband beat signals discussed in Chapter 6 are currently the three key observations against which we can test different hypotheses. A detailed simulation incorporating noise processes may reveal potential sources. If the source of noise is not found, dynamical decoupling schemes [16, 9, 113] could be investigated.

The next limiting factor after gate infidelity will likely be the laser intensity fluctuations. The calcium and beryllium lasers fluctuate by up to several percent during a gate pulse when their AOMs are slightly misaligned. The analog feedback loop on the DDS boards is being implemented at the time of writing, and will likely reduce this. Additionally, the laser cavities losing lock have been a regular problem in the acquisition of long sequences. Connecting the EVIL lock-status signals to the M-ACTION master will address this problem.

Ultimately multi-qubit gate infidelities below 10^{-3} may be feasible once the error sources are better-understood, as demonstrated in similar experimental setups [50, 8].

8.4 Conclusion

The main result of this work has been the implementation of feedback-based quantum protocols, along with the necessary hardware and constituent operations. Experimentally this area of research is still in its infancy and new schemes and protocols are regularly proposed, waiting to be investigated. This concludes my thesis; I hope that the reader has found it informative. I also hope that the techniques, equipment and experiments it has described will contribute to the ultimate scaling-up of trapped-ion quantum information processing, and the experimental exploration of interesting new systems not yet imagined.

A Further M-ACTION details

A.1 API source details

The control software is compiled and run using the Xilinx Software Development Kit (SDK). This consists of several related projects. `ionpulse_sw` is the primary project, with `sharedlib` and `standalone_bsp_0` containing required libraries. `testing_sandbox` is used only for hardware development and debugging, containing various selftest and diagnostic routines, and can be run independently of the Ionizer2 GUI. Inside the `src` folder of `ionpulse_sw`, there are multiple files that implement a library of experimental control functionality, including the server code, interfaces to the peripherals, and helper functions. There are several experimental folders, corresponding to the various experimental setups in the TIQI group, with `segtrap` usually containing the most up-to-date driver and experimental code. Each folder consists of at least `config_local.cpp/h` files, which define various global properties of the experimental setup, including typical digital-output patterns, instantiated experiment classes, and DDS configuration. There are also several `globals_*.cpp/h` files, which contain globally shared functions and parameters used by multiple experiments; multiple globals files (for example, for beryllium and calcium) may be used in complex experiments.

A final folder, `generic`, implements the `experiment` class. To create an entirely new experiment, the user must write a new class, inheriting from the `experiment` class in `experiment.h`. By using several layers of inheritance, multiple layers of functionality may be implemented. For example, `ionpulse_sw/src/segtrap/ca/frame_cooldet_Ca.h` implements a framework of functions useful for a single calcium ion, including cooling, state preparation and detection; typically a new single-calcium experiment class may be created simply by inheriting from it directly and only redefining the `QubitManipulation()` function. Finally, the experiment class must be instantiated and added to a global list; this is done in `ionpulse_sw/src/segtrap/config_local.cpp/h`.

A.2 Pulser FIFO instruction

The `pulser` instruction is composed of

bits	63	62	61	60 – 55	54 – 48	47 – 32	31 – 0
meaning	cycles $\times 2^{16}$	wait for trigger	send trigger	spare	gate PMT counter/s	cycles	pulse pattern

.

A.3 Bitumen communication protocol

Each `bitumen` transmission (in either direction) has a start bit, special bit, 12 data bits, parity bit and a stop bit. The `bitumen` master send 4 data chunks (with the special bit set to 0) and one chunk with the special bit set to 1 and the data set to 0. It holds a running checksum of the transmitted data, which is compared to the response from the slave for simple error detection.

The slave stores the incoming data, and sends back the 12-bit checksum. In response to the special chunk from the master, any communication errors previously detected by the slave and 8 bits of DDS board status and errors are sent to the master.

A.4 Experimental forking in idecoder

For a series of n possible paths through a fork in the instructions of idecoder, where each path consists of k_n instructions, the system first writes a single jump instruction into the `idecoder` BRAM, followed by n jump instructions, which are in turn followed by each set of k_n instructions for each path. The end of each set is terminated by a jump to the end of all the sets, where the paths rejoin. Each of the n jump instructions is targetted to the beginning of one of the sets, and the single initial jump is targetted to one of the n instructions following it. This way a single initial jump instruction can be broadcast to every DDS channel, and the n instructions following it take care of the variable k_n lengths that arise between different channels.

B Auxiliary experimental calibrations

In this appendix, various calibrations are outlined that are performed infrequently or only for specialised experiments.

B.1 Single-qubit procedures

B.1.1 Pre-cooling and repump beam alignment

The beam alignments of the 854 nm repump for calcium and the repump beams for beryllium are optimised by carrying out a 729 or FDQ π pulse to make the population dark, then running a repumping pulse and scanning its time. An exponential decay to the bright population emerges. The alignment is optimised by setting the repump pulse to a short time (several microseconds) and/or reducing its power, and adjusting the beam to increase the bright population. This is insufficient to optimise the repump frequencies however; see [81] and [105] for details.

The beryllium pre-cooling beam is aligned in a similar way. Instead of repumping a pre-cooling pulse is applied, which weakly repumps the dark states. This method is necessary because the beam is 600 MHz red of the detection resonance and does not induce visible fluorescence.

B.1.2 Beryllium leakage estimation

The dark→bright leakage for a beryllium ion is estimated with a π pulse followed by detection, where the detection time is scanned over several milliseconds. This is less straightforward than the repump scan to interpret because the competing bright→dark process affects the population once a significant proportion is bright. A fit to the initial section of the curve where $p_1 < 0.1$ can provide an estimate.

B.1.3 Calcium EIT σ beam alignment

The 397 EIT σ beam is aligned by using it for state preparation instead of the regular 397 σ beam. The EIT σ beam is pulsed initially, followed by a π pulse on the 729; when it optimally prepares the $|g\rangle$ state the dark population is maximised.

B.1.4 Calcium rf optimisation

The $|g\rangle \leftrightarrow |d\rangle$ rf drive frequency and coupling into the vacuum chamber were optimised by first preparing $|g\rangle$ with the usual 397 σ pulse, applying the rf drive, then a 729 π shelving pulse to transfer the $|g\rangle$ population to $|e\rangle$. Scanning the rf showed the same coherent Rabi dynamics as the other system qubits, which were then used for optimisation as described above.

B.2 Multi-qubit procedures

B.2.1 Beryllium phase-insensitive MS gate

This consists of sandwiching the gate in co-90 $\pi/2$ pulses to rotate it into the $\hat{\sigma}_z$ basis, carrying out a phase gate, which is itself sandwiched between co-co $\pi/2$ pulses to return back to the usual MS $\{\hat{\sigma}_x, \hat{\sigma}_y\}$ basis. It assumes the MS gate and $\pi/2$ pulses are already calibrated. A phase-sensitive beryllium MS gate with spin phase ϕ_c is defined as $MS_{\mathrm{Be}}(\phi_c)$, whereas the phase-insensitive gate is $MS^z_{\mathrm{Be}}(0)$. $R_{\pi/2}$ pulses on the co-co and co-90 beams are defined as $R^{\mathrm{co}}_{\pi/2,\mathrm{Be}}(\phi)$ and $R^{\mathrm{co}}_{\pi/2,\mathrm{Be}}(\phi)$ respectively. The initial state for the calibration is $|gg\rangle$, corresponding to $p_2 = 1$. The following steps are carried out.

Block 1: $MS_{\mathrm{Be}}(0))$ followed by $R^{90}_{\pi/2,\mathrm{Be}}(\phi_1)$. ϕ_1 is set such that p_0 and p_2 are equalised (the maxima and minima of p_0 and p_2 are used to avoid bias due to detection error) and the gradient $dp_0/d\phi_1$ is positive.

Block 2: $R^{90}_{\pi/2,\mathrm{Be}}(\phi_2)$ followed by Block 1. ϕ_2 and ϕ_c are iterated to maximise p_2.

$MS^z_{\mathrm{Be}}(0)$ **sequence:** $R^{\mathrm{co}}_{\pi/2,\mathrm{Be}}(\phi_3)$, Block 2, $R^{\mathrm{co}}_{\pi/2,\mathrm{Be}}(\phi_4)$. ϕ_3 is set to 0 and ϕ_4 is set to minimise p_1.

Once this is complete, the gate operates in the co-co phase basis. A phase scan of an additional $RBe\phi$ pulse should yield the constituent odd and even populations of a parity curve similar to Figure 6.4c.

B.2.2 U_{S_Z} sequence

This relies on the beryllium phase-insensitive MS gate calibration being optimal, as well as the beryllium-calcium-beryllium MS gate. The purpose of the various pulses is discussed in §7.2. A calcium $R_{\pi/2}$ rotation is defined as $R_{\pi/2,\mathrm{Ca}}(\phi)$. A beryllium-beryllium MS gate is defined as $MS_{\mathrm{Be}}(\phi_{c,\mathrm{Be}})$, and a beryllium-calcium-beryllium MS gate as $MS_{\mathrm{Be,Ca}}(\phi_{c,\mathrm{Be}}, \phi_{c,\mathrm{Ca}})$; note that these are not phase-insensitive. p_m refers to the population where m beryllium ions are bright, and $p_{m,n}$ is where m beryllium and n calcium ions are bright respectively.

The initial state for the calibration is $|ggg\rangle$, and the following steps are carried out.

Block 1: $R^{90}_{\pi/2,\mathrm{Be}}(\phi_1)$ followed by $MS_{\mathrm{Be}}(0)$. ϕ_1 is set such that p_0 and p_2 are equalised (the maxima and minima of p_0 and p_2 are used to avoid bias due to detection error) and the gradient $dp_2/d\phi_1$ is positive.

Block 2: $MS_{\mathrm{Be}}(0)$ then $MS_{\mathrm{Be}}(\phi_2)$. ϕ_2 is set to maximise p_0.

Block 3: $MS_{\mathrm{Be}}(0)$, $RpiBee\phi_3$ then $MS_{\mathrm{Be}}(\phi_2)$. ϕ_3 and ϕ_2 are iterated to maximise p_2.

Block 4: Block 1, then $MS_{\mathrm{Be}}(\phi_2)$ and $R^{90}_{\pi/2,\mathrm{Be}}(\phi_4)$. ϕ_4 is set to maximise p_0.

Block 5: $MS_{\mathrm{Be,Ca}}(0,0)$, $R^{90}_{\pi,\mathrm{Be}}(\phi_3)$ then $MS_{\mathrm{Be,Ca}}(\phi_2, \phi_5)$. ϕ_5 is set to maximise $p_{0,1}$.

Block 6: $MS_{\mathrm{Be,Ca}}(0,0)$, $R^{90}_{\pi,\mathrm{Be}}(\phi_3)$, $R_{\pi,\mathrm{Ca}}(\phi_6)$ then $MS_{\mathrm{Be,Ca}}(\phi_2, \phi_5)$. ϕ_6 is set such that $p_{0,0}$ and $p_{2,0}$ are equalised (the maxima and minima of $p_{0,0}$ and $p_{2,0}$ are used to avoid bias due to detection error), and the gradient $dp_{2,0}/d\phi_6$ is positive.

S_Z sequence: $R_{\pi/2,\mathrm{Ca}}(\phi_7)$, $R^{90}_{\pi/2,\mathrm{Be}}(\phi_1)$, Block 6, $R^{90}_{\pi/2,\mathrm{Be}}(\phi_4)$ then $R_{\pi/2,\mathrm{Ca}}(\phi_8)$. ϕ_7 and ϕ_8 are iterated to maximise $p_{2,0}$.

To verify that the S_Z parity measurement is valid, the beryllium phase-insensitive MS gate and $RBe\phi$ are carried out before the S_Z sequence. A scan of ϕ produces the constituent populations of the curves shown in Figure 7.1d.

B.2.3 U_{S_X} sequence

Once the U_{S_Z} sequence is calibrated, the U_{S_X} calibration consists of carrying out $MS^z_{\mathrm{Be}}(0)$, $R^{\mathrm{co}}_{\pi/2,\mathrm{Be}}(\phi)$ then U_{S_Z}, and choosing ϕ such that the even-parity population p_0+p_2 is maximised. Next, the sequence $MS^z_{\mathrm{Be}}(0)$, $R^{\mathrm{co}}_{\pi/2,\mathrm{Be}}(\phi)$, $R^{\mathrm{co}}_{\pi/2,\mathrm{Be}}(\phi_1)$, S_Z and $R^{\mathrm{co}}_{\pi/2,\mathrm{Be}}(\phi_2)$ is carried out. ϕ_1 and ϕ_2 are iterated to minimise $p_{1,1}$ and $p_{1,0}$, in other words maximising the four even-parity populations so that S_X acts as a QND measurement. Several iterations are required, since the populations of interest are quite low. Once an optimum has been reached, a scan of ϕ produces the constituent populations of the curves shown in Figure 7.1c.

C Abbreviations

ADC: *analog-digital converter*, chip used to convert analog voltages or currents into streams of digital numbers.

AOM: *acousto-optic modulator*, used for frequency, amplitude and phase modulation of lasers.

API: *application programming interface*, software library providing users an interface to lower-level system functions.

AWG: *arbitrary waveform generator*, laboratory device to produce an arbitrary time-dependent voltage or current.

BRAM: *block RAM*, physical RAM common on FPGAs, as opposed to a RAM created from FPGA logic cells.

CCD: *charge-coupled device*, device for collecting photons that operates by controlled transporting of charge across a chip.

CoM: *centre-of-mass motional mode*, in-phase axial oscillation of an ion crystal.

co-co: *co-propagating Raman beams*, used for addressing motion-insensitive transitions in beryllium in the mixed-species setup.

co-90: *perpendicularly-propagating Raman beams*, used for addressing motion-sensitive transitions in beryllium in the mixed-species setup.

CPU: *central processing unit*, processor responsible for executing machine code at the heart of PCs or embedded devices.

DAC: *digital-analog converter*, chip used to convert streams of digital numbers into analog voltages or currents.

DDS: *direct-digital synthesiser*, mixed-signal circuit that deterministically produces rf from a digital clock.

DEATH: *direct Ethernet-adjustable transport hardware*, M-ACTION AWGs used to control trap dc electrode voltages.

EGY: *'Egyptian' motional mode*, 3-ion axial oscillation mode; outer ions oscillate out of phase with the central ion.

EVIL: *electronically-variable interactive lockbox*, digital PID controller widely used in the TIQI group.

FDQ: *field-dependent qubit*, beryllium transition between the bright state and $|g\rangle$ used mainly for cooling and beam alignment; see Figure 2.4.

FIFO: *first-in first-out*, address-less memory buffer where data units are read only in the same order they were written in.

FIQ: *field-independent qubit*, beryllium transition between $|g\rangle$ and $|e\rangle$, used for gates and QIP protocols; see Figure 2.4.

FIS: *field-independent qubit shelving*, beryllium transition between $|e\rangle$ and an auxiliary state, used for shelving the $|e\rangle$ population to improve detection fidelity; see Figure 2.4.

FPGA: *field-programmable gate array*, digital logic chip that can be programmed with a pattern of gates and other elements.

FSM: *finite-state machine*, digital circuit often used for controlling other digital logic, with a set of states, conditional transitions between them, and associated behaviours.

GUI: *graphical user interface*, software providing a 'point-and-click' interface for a user.

HDL: *hardware description language*, used for describing digital logic and programming FPGAs.

LUT: *look-up table*, simple random-access data structure that stores the value of a pre-calculated function.

LVDS: *low-voltage differential signalling*, signal standard that encodes a bit in the voltage difference between two digital lines.

M-ACTION: *modular advanced control of trapped ions*, experimental control system used in the TIQI group.

MSPS: *megasamples per second*, measure of the sampling rate of a mixed-signal system such as an ADC or DAC.

PID: *proportional-integral-derivative*, type of closed-loop servo controller which uses a proportional-integral-derivative mathematical law to generate the output from the input.

PLL: *phase-locked loop*, mixed-signal circuit to generate analog rf or digital clocks from a reference signal usually at a different frequency.

PMT: *photo-multiplier tube*, used for counting photons at a high quantum efficiency for ion readout.

QIP: *quantum information processing*, field encompassing quantum computation, simulation and related topics.

QND: *quantum non-demolition*, quantum measurement that in certain circumstances leaves a quantum system in its original state.

RAM: *random-access memory*, digital memory that can be read and written in arbitrary order.

ROI: *region of interest*, a sub-region of some parameter space, typically in data processing or analysis.

SDK: *software development kit*, development environment in which a user can write and compile software for a desktop or embedded computer system.

SPI: *serial peripheral interface*, common digital chip-chip serial communication protocol.

STR: *'stretch' motional mode*, out-of-phase axial mode of oscillation for a two- or three-ion crystal.

VGA: *variable-gain amplifier*, analog amplifier whose gain is externally controlled, often by a voltage.

Bibliography

[1] M. Acton et al. "Near-Perfect Simultaneous Measurement of a Qubit Register". In: *Quantum Info. Comput.* 6.6 (Sept. 2006), pp. 465–482 (cit. on p. 57).

[2] D. T. C. Allcock et al. "A Microfabricated Ion Trap with Integrated Microwave Circuitry". In: *Applied Physics Letters* 102.044103 (2013). DOI: 10.1063/1.4774299 (cit. on p. 3).

[3] J. Alonso et al. "Generation of Large Coherent States by Bang–bang Control of a Trapped-Ion Oscillator". In: *Nature Communications* 7 (Apr. 2016), p. 11243. DOI: 10.1038/ncomms11243 (cit. on p. 131).

[4] J. M. Amini et al. "Micro-Fabricated Chip Traps for Ions". In: *Atom Chips*. Wiley-Blackwell, 2011, pp. 395–420. DOI: 10.1002/9783527633357.ch13 (cit. on p. 65).

[5] J. M. Amini et al. "Toward Scalable Ion Traps for Quantum Information Processing". In: *New Journal of Physics* 12.3 (2010), p. 033031. DOI: 10.1088/1367-2630/12/3/033031 (cit. on p. 74).

[6] Analog Devices. *AD9910 Datasheet, Rev. E*. 2016 (cit. on pp. 37, 38).

[7] M. T. Baig et al. "A Scalable, Fast, and Multichannel Arbitrary Waveform Generator". In: *Review of Scientific Instruments* 84.12 (Dec. 2013), p. 124701. DOI: 10.1063/1.4832042 (cit. on pp. 29, 30).

[8] C. J. Ballance et al. "High-Fidelity Quantum Logic Gates Using Trapped-Ion Hyperfine Qubits". In: *Physical Review Letters* 117.6 (Aug. 2016), p. 060504. DOI: 10.1103/PhysRevLett.117.060504 (cit. on p. 133).

[9] C. J. Ballance et al. "Hybrid Quantum Logic and a Test of Bell's Inequality Using Two Different Atomic Isotopes". In: *Nature* 528.7582 (Dec. 2015), pp. 384–386. DOI: 10.1038/nature16184 (cit. on pp. 5, 6, 106, 133).

[10] C. J. Ballance. "High-Fidelity Quantum Logic in Ca^+". PhD Thesis. Oxford University, 2014 (cit. on pp. 42, 57, 105, 113).

[11] J. T. Barreiro et al. "An Open-System Quantum Simulator with Trapped Ions". In: *Nature* 470.7335 (Feb. 2011), pp. 486–491. DOI: 10.1038/nature09801 (cit. on pp. 116, 120).

[12] M. D. Barrett et al. "Deterministic Quantum Teleportation of Atomic Qubits". In: *Nature* 429.6993 (June 2004), pp. 737–739. DOI: 10.1038/nature02608 (cit. on p. 73).

[13] P. A. Barton et al. "Measurement of the Lifetime of the $3d^2D_{5/2}$ State in $^{40}Ca^+$". In: *Physical Review A* 62.3 (Aug. 2000), p. 032503. DOI: 10.1103/PhysRevA.62.032503 (cit. on p. 27).

[14] J. Benhelm et al. "Towards Fault-Tolerant Quantum Computing with Trapped Ions". In: *Nature Physics* 4.6 (June 2008), pp. 463–466. DOI: 10.1038/nphys961 (cit. on p. 104).

[15] D. J. Berkeland et al. "Minimization of Ion Micromotion in a Paul Trap". In: *Journal of Applied Physics* 83.10 (Apr. 1998), pp. 5025–5033. DOI: 10.1063/1.367318 (cit. on pp. 9, 64).

[16] A. Bermudez et al. "Robust Trapped-Ion Quantum Logic Gates by Continuous Dynamical Decoupling". In: *Physical Review A* 85.4 (Apr. 2012), p. 040302. DOI: 10.1103/PhysRevA.85.040302 (cit. on pp. 5, 106, 133).

[17] M. J. Biercuk et al. "Optimized Dynamical Decoupling in a Model Quantum Memory". In: *Nature* 458.7241 (Apr. 2009), pp. 996–1000. DOI: 10.1038/nature07951 (cit. on p. 5).

[18] E. D. Black. "An Introduction to Pound–Drever–Hall Laser Frequency Stabilization". In: *American Journal of Physics* 69.1 (Dec. 2000), pp. 79–87. DOI: 10.1119/1.1286663 (cit. on p. 50).

[19] R. Blatt and D. Wineland. "Entangled States of Trapped Atomic Ions". In: *Nature* 453.7198 (June 2008), pp. 1008–1015. DOI: 10.1038/nature07125 (cit. on p. 3).

[20] J. J. . Bollinger et al. "Optimal Frequency Measurements with Maximally Correlated States". In: *Physical Review A* 54.6 (Dec. 1996), R4649–R4652. DOI: 10.1103/PhysRevA.54.R4649 (cit. on p. 106).

[21] C. Bonato et al. "Optimized Quantum Sensing with a Single Electron Spin Using Real-Time Adaptive Measurements". In: *Nature Nanotechnology* 11.3 (Mar. 2016), pp. 247–252. DOI: 10.1038/nnano.2015.261 (cit. on pp. 98, 99).

[22] C. Bonato and D. W. Berry. "Adaptive Tracking of a Time-Varying Field with a Quantum Sensor". In: *Physical Review A* 95.5 (May 2017), p. 052348. DOI: 10.1103/PhysRevA.95.052348 (cit. on p. 98).

[23] R. Bowler et al. "Coherent Diabatic Ion Transport and Separation in a Multizone Trap Array". In: *Physical Review Letters* 109.8 (Aug. 2012), p. 080502. DOI: 10.1103/PhysRevLett.109.080502 (cit. on pp. 64, 73, 81, 84).

[24] K. R. Brown, J. Kim and C. Monroe. "Co-Designing a Scalable Quantum Computer with Trapped Atomic Ions". In: *npj Quantum Information* 2 (Nov. 2016), p. 16034. DOI: 10.1038/npjqi.2016.34 (cit. on p. 3).

[25] N. C. Brown and K. R. Brown. "Comparing Zeeman Qubits to Hyperfine Qubits in the Context of the Surface Code: $^{174}Yb^{+}$ and $^{171}Yb^{+}$". In: *Physical Review A* 97.5 (May 2018), p. 052301. DOI: 10.1103/PhysRevA.97.052301 (cit. on p. 126).

[26] C. Bruzewicz et al. "High-Fidelity, Single-Shot, Quantum-Logic-Assisted Readout in a Mixed-Species Ion Chain". In: *arXiv:1706.05102 [physics, physics:quant-ph]* (June 2017). arXiv: 1706.05102 [physics, physics:quant-ph] (cit. on p. 6).

[27] A. R. Calderbank and P. W. Shor. "Good Quantum Error-Correcting Codes Exist". In: *Physical Review A* 54.2 (Aug. 1996), pp. 1098–1105. DOI: 10.1103/PhysRevA.54.1098 (cit. on p. 5).

[28] G. Camy, C. Bordé and M. Ducloy. "Heterodyne Saturation Spectroscopy through Frequency Modulation of the Saturating Beam". In: *Optics Communications* 41.5 (May 1982), pp. 325–330. DOI: 10.1016/0030-4018(82)90406-0 (cit. on p. 51).

[29] J. Chiaverini et al. "Realization of Quantum Error Correction". In: *Nature* 432.7017 (Dec. 2004), pp. 602–605. DOI: 10.1038/nature03074 (cit. on p. 116).

[30] J. I. Cirac and P. Zoller. "Quantum Computations with Cold Trapped Ions". In: *Physical Review Letters* 74.20 (May 1995), pp. 4091–4094. DOI: 10.1103/PhysRevLett.74.4091 (cit. on pp. 3, 102).

[31] C. Cohen-Tannoudji, J. Dupont-Roc and G. Grynberg. *Atom-Photon Interactions: Basic Processes and Applications*. New York: Wiley-VCH, Mar. 1998 (cit. on p. 15).

[32] Y. Colombe et al. "Single-Mode Optical Fiber for High-Power, Low-Loss UV Transmission". In: *Optics Express* 22.16 (Aug. 2014), pp. 19783–19793. DOI: 10.1364/OE.22.019783 (cit. on p. 24).

[33] D. G. Cory et al. "Experimental Quantum Error Correction". In: *Physical Review Letters* 81.10 (Sept. 1998), pp. 2152–2155. DOI: 10.1103/PhysRevLett.81.2152 (cit. on p. 116).

[34] J. Cramer et al. "Repeated Quantum Error Correction on a Continuously Encoded Qubit by Real-Time Feedback". In: *Nature Communications* 7 (May 2016), p. 11526. DOI: 10.1038/ncomms11526 (cit. on pp. 6, 29, 116).

[35] L. E. de Clercq et al. "Estimation of a General Time-Dependent Hamiltonian for a Single Qubit". In: *Nature Communications* 7 (Apr. 2016), p. 11218. DOI: 10.1038/ncomms11218 (cit. on p. 131).

[36] L. E. de Clercq et al. "Time-Dependent Hamiltonian Estimation for Doppler Velocimetry of Trapped Ions". In: *arXiv:1509.07083 [physics, physics:quant-ph]* (Sept. 2015). arXiv: 1509.07083 [physics, physics:quant-ph] (cit. on pp. 68, 72).

[37] L. E. de Clercq. "Transport Quantum Logic Gates for Trapped Ions". Doctoral Thesis. ETH Zürich, 2015 (cit. on pp. 32, 50, 69, 72, 80).

[38] L. E. de Clercq et al. "Parallel Transport Quantum Logic Gates with Trapped Ions". In: *Physical Review Letters* 116.8 (Feb. 2016), p. 080502. DOI: 10.1103/PhysRevLett.116.080502 (cit. on pp. 68, 75, 131).

[39] A. Devices. *Fundamentals of Direct Digital Synthesis (DDS)*. Tutorial MT-085. Analog Devices, Inc., 2009 (cit. on p. 37).

[40] M. H. Devoret and R. J. Schoelkopf. "Superconducting Circuits for Quantum Information: An Outlook". In: *Science* 339.6124 (Mar. 2013), pp. 1169–1174. DOI: 10.1126/science.1231930 (cit. on p. 2).

[41] D. P. DiVincenzo. "The Physical Implementation of Quantum Computation". In: *Fortschritte der Physik* 48.1 (Feb. 2000), pp. 771–783 (cit. on pp. 2, 10, 102).

[42] A. Doucet, S. Godsill and C. Andrieu. "On Sequential Monte Carlo Sampling Methods for Bayesian Filtering". In: *Statistics and Computing* 10.3 (July 2000), pp. 197–208. DOI: 10.1023/A:1008935410038 (cit. on p. 99).

[43] R. W. P. Drever et al. "Laser Phase and Frequency Stabilization Using an Optical Resonator". In: *Applied Physics B: Lasers and Optics* 31.2 (June 1983), pp. 97–105. DOI: 10.1007/BF00702605 (cit. on p. 50).

[44] P. Engeler. "Using the Zynq-7000 XADC and Signal Pre-Conditioning". Internship Thesis. ETH Zürich, 2017 (cit. on p. 33).

[45] Enterpoint Ltd. *Milldown Product Page*. https://www.enterpoint.co.uk/products/spartan-6-development-boards/milldown/. June 2014 (cit. on pp. 34, 37).

[46] C. Fischer. "Implementation of a Digital Lock-in Amplifier on a Field-Programmable Gate Array and Its Remote Control in a Local-Area Network". Semester Thesis. ETH Zürich, 2014 (cit. on p. 50).

[47] C. Fischer. "Optical Trapping in a Build-up Cavity". Masters Thesis. ETH Zürich, 2015 (cit. on p. 69).

[48] C. Flühmann et al. "Sequential Modular Position and Momentum Measurements of a Trapped Ion Mechanical Oscillator". In: *Physical Review X* 8.2 (Apr. 2018), p. 021001. DOI: 10.1103/PhysRevX.8.021001 (cit. on pp. 12, 131).

[49] C. Flühmann. "Title Undecided". in preparation. PhD Thesis. ETH Zürich, 2019 (cit. on p. 52).

[50] J. P. Gaebler et al. "High-Fidelity Universal Gate Set for $^{9}Be^{+}$ Ion Qubits". In: *Physical Review Letters* 117.6 (Aug. 2016), p. 060505. DOI: 10.1103/PhysRevLett.117.060505 (cit. on pp. 6, 117, 133).

[51] J. M. Gambetta, J. M. Chow and M. Steffen. "Building Logical Qubits in a Superconducting Quantum Computing System". In: *npj Quantum Information* 3.1 (Jan. 2017), p. 2. DOI: 10.1038/s41534-016-0004-0 (cit. on p. 2).

[52] I. M. Georgescu, S. Ashhab and F. Nori. "Quantum Simulation". In: *Reviews of Modern Physics* 86.1 (Mar. 2014), pp. 153–185. DOI: 10.1103/RevModPhys.86.153 (cit. on p. 2).

[53] D. Gottesman. "Theory of Fault-Tolerant Quantum Computation". In: *Physical Review A* 57.1 (Jan. 1998), pp. 127–137. DOI: 10.1103/PhysRevA.57.127 (cit. on pp. 5, 116).

[54] D. Gottesman and I. L. Chuang. "Demonstrating the Viability of Universal Quantum Computation Using Teleportation and Single-Qubit Operations". In: *Nature* 402.6760 (Nov. 1999), pp. 390–393. DOI: 10.1038/46503 (cit. on pp. 102, 130).

[55] C. Granade. "Characterization, Verification and Control for Large Quantum Systems". PhD Thesis. University of Waterloo, 2014 (cit. on p. 99).

[56] C. E. Granade et al. "Robust Online Hamiltonian Learning". In: *New Journal of Physics* 14.10 (Oct. 2012), p. 103013. DOI: 10.1088/1367-2630/14/10/103013 (cit. on p. 99).

[57] D. M. Greenberger, M. A. Horne and A. Zeilinger. "Going Beyond Bell's Theorem". In: *Bell's Theorem, Quantum Theory and Conceptions of the Universe*. Fundamental Theories of Physics. Open-access version at https://arxiv.org/abs/0712.0921. Springer, Dordrecht, 1989, pp. 69–72. DOI: 10.1007/978-94-017-0849-4_10 (cit. on p. 101).

[58] L. K. Grover. "From Schrödinger's Equation to the Quantum Search Algorithm". In: *Pramana* 56.2-3 (Feb. 2001), pp. 333–348. DOI: 10.1007/s12043-001-0128-3 (cit. on p. 2).

[59] G. Guennebaud, B. Jacob and and others. *Eigen V3*. http://eigen.tuxfamily.org/, 2010 (cit. on p. 43).

[60] H. Haeffner, C. F. Roos and R. Blatt. "Quantum Computing with Trapped Ions". In: *Physics Reports* 469.4 (Sept. 2008). Phys. Rep. 469, 155-203 (2008), pp. 155–203. DOI: 10.1016/j.physrep.2008.09.003 (cit. on pp. 3, 16, 106).

[61] A. W. Harrow, A. Hassidim and S. Lloyd. "Quantum Algorithm for Linear Systems of Equations". In: *Physical Review Letters* 103.15 (Oct. 2009), p. 150502. DOI: 10.1103/PhysRevLett.103.150502 (cit. on p. 2).

[62] C. Hempel et al. "Quantum Chemistry Calculations on a Trapped-Ion Quantum Simulator". In: *arXiv:1803.10238 [quant-ph]* (Mar. 2018). arXiv: 1803.10238 [quant-ph] (cit. on p. 2).

[63] R. J. Hendricks et al. "Doppler Cooling of Calcium Ions Using a Dipole-Forbidden Transition". In: *Physical Review A* 77.2 (Feb. 2008), p. 021401. DOI: 10.1103/PhysRevA.77.021401 (cit. on p. 20).

[64] J. P. Home and A. M. Steane. "Electrode Configurations for Fast Separation of Trapped Ions". In: *Quantum Information and Computation* 6.4 (July 2006), pp. 289–325 (cit. on p. 74).

[65] J. P. Home et al. "Normal Modes of Trapped Ions in the Presence of Anharmonic Trap Potentials". In: *New Journal of Physics* 13.7 (July 2011), p. 073026. DOI: 10.1088/1367-2630/13/7/073026 (cit. on pp. 69, 74).

[66] J. P. Home. "Entanglement of Two Trapped-Ion Spin Qubits". Doctoral Thesis. Oxford University, 2006 (cit. on pp. 90, 101).

[67] J. P. Home et al. "Complete Methods Set for Scalable Ion Trap Quantum Information Processing". In: *Science* 325.5945 (Apr. 2009), pp. 1227–1230. DOI: 10.1126/science.1177077 (cit. on p. 6).

[68] J. P. Home. "Chapter 4 - Quantum Science and Metrology with Mixed-Species Ion Chains". In: *Advances In Atomic, Molecular, and Optical Physics*. Ed. by P. R. B. a. C. C. L. Ennio Arimondo. Vol. Volume 62. Advances in Atomic, Molecular, and Optical Physics. Academic Press, 2013, pp. 231–277 (cit. on pp. 5, 64, 85, 86).

[69] A. Hungenberg. "Automatic Relocking of an FPGA-Based PID Controller Using a Bandpass-Filtering Approach". Semester Thesis. ETH Zürich, 2013 (cit. on p. 51).

[70] Y. Ibaraki, U. Tanaka and S. Urabe. "Detection of Parametric Resonance of Trapped Ions for Micromotion Compensation". In: *Applied Physics B* 105.2 (Nov. 2011), pp. 219–223. DOI: 10.1007/s00340-011-4463-x (cit. on p. 64).

[71] W. M. Itano et al. "Quantum Projection Noise: Population Fluctuations in Two-Level Systems". In: *Physical Review A* 47.5 (May 1993), pp. 3554–3570. DOI: 10.1103/PhysRevA.47.3554 (cit. on p. 97).

[72] K. G. Johnson et al. "Active Stabilization of Ion Trap Radiofrequency Potentials". In: *Review of Scientific Instruments* 87.5 (May 2016), p. 053110. DOI: 10.1063/1.4948734 (cit. on p. 110).

[73] H. Kaufmann et al. "Dynamics and Control of Fast Ion Crystal Splitting in Segmented Paul Traps". In: *New Journal of Physics* 16.7 (2014), p. 073012. DOI: 10.1088/1367-2630/16/7/073012 (cit. on pp. 76, 77, 84).

[74] B. C. Keitch. "A Quantum Memory Qubit in Calcium-43". Doctoral Thesis. Oxford University, 2007 (cit. on pp. 27, 50).

[75] B. Keitch, V. Negnevitsky and W. Zhang. "Programmable and Scalable Radio-Frequency Pulse Sequence Generator for Multi-Qubit Quantum Information Experiments". In: *arXiv:1710.04282 [physics, physics:quant-ph]* (Oct. 2017). arXiv: 1710.04282 [physics, physics:quant-ph] (cit. on pp. 34, 37).

[76] D. Kielpinski, C. Monroe and D. J. Wineland. "Architecture for a Large-Scale Ion-Trap Quantum Computer". In: *Nature* 417.6890 (June 2002), pp. 709–711. DOI: 10.1038/nature00784 (cit. on pp. 3, 68).

[77] D. Kielpinski et al. "A Decoherence-Free Quantum Memory Using Trapped Ions". In: *Science* 291.5506 (Sept. 2001), pp. 1013–1015. DOI: 10.1126/science.1057357 (cit. on pp. 5, 106).

[78] D. Kienzler et al. "Observation of Quantum Interference between Separated Mechanical Oscillator Wave Packets". In: *Physical Review Letters* 116.14 (Apr. 2016), p. 140402. DOI: 10.1103/PhysRevLett.116.140402 (cit. on pp. 12, 23, 131).

[79] D. Kienzler et al. "Quantum Harmonic Oscillator State Control in a Squeezed Fock Basis". In: *Physical Review Letters* 119.3 (July 2017), p. 033602. DOI: 10.1103/PhysRevLett.119.033602 (cit. on p. 131).

[80] D. Kienzler et al. "Quantum Harmonic Oscillator State Synthesis by Reservoir Engineering". In: *Science* 347.6217 (Jan. 2015), pp. 53–56. DOI: 10.1126/science.1261033 (cit. on pp. 12, 23, 40, 131).

[81] D. Kienzler. "Quantum Harmonic Oscillator State Synthesis by Reservoir Engineering". Doctoral Thesis. ETH Zürich, 2015 (cit. on pp. 3, 7–10, 16, 21–24, 55, 65, 90, 103, 110, 136).

[82] S. Kimmel, G. H. Low and T. J. Yoder. "Robust Calibration of a Universal Single-Qubit Gate Set via Robust Phase Estimation". In: *Physical Review A* 92.6 (Dec. 2015), p. 062315. DOI: 10.1103/PhysRevA.92.062315 (cit. on p. 98).

[83] G. Kirchmair et al. "Deterministic Entanglement of Ions in Thermal States of Motion". In: *New Journal of Physics* 11.2 (Feb. 2009), p. 023002. DOI: 10.1088/1367-2630/11/2/023002 (cit. on pp. 102–104).

[84] E. Knill et al. "Introduction to Quantum Error Correction". In: *arXiv:quant-ph/0207170* (July 2002). arXiv: quant-ph/0207170 (cit. on p. 115).

[85] S. Kotler et al. "Single-Ion Quantum Lock-in Amplifier". In: *Nature* 473.7345 (May 2011), pp. 61–65. DOI: 10.1038/nature10010 (cit. on p. 91).

[86] C. Langer et al. "Long-Lived Qubit Memory Using Atomic Ions". In: *Physical Review Letters* 95.6 (Aug. 2005), p. 060502. DOI: 10.1103/PhysRevLett.95.060502 (cit. on pp. 17, 18, 106).

[87] C. E. Langer. "High Fidelity Quantum Information Processing with Trapped Ions". Doctoral Thesis. University of Colorado, Boulder, 2006 (cit. on pp. 17, 18, 29, 56–58, 90, 106).

[88] B. P. Lanyon et al. "Measurement-Based Quantum Computation with Trapped Ions". In: *Physical Review Letters* 111.21 (Nov. 2013), p. 210501. DOI: 10.1103/PhysRevLett.111.210501 (cit. on p. 130).

[89] H.-K. Lau and D. F. V. James. "Decoherence and Dephasing Errors Caused by the Dc Stark Effect in Rapid Ion Transport". In: *Physical Review A* 83.6 (June 2011), p. 062330. DOI: 10.1103/PhysRevA.83.062330 (cit. on p. 73).

[90] A. Lebedev. *Metrology Protocol Notes*. Tech. rep. 2015 (cit. on pp. 93, 97).

[91] R. Lechner. "Multi-Mode Cooling Techniques for Trapped Ions". PhD Thesis. University of Innsbruck, 2017 (cit. on p. 21).

[92] R. Lechner et al. "Electromagnetically-Induced-Transparency Ground-State Cooling of Long Ion Strings". In: *Physical Review A* 93.5 (May 2016), p. 053401. DOI: 10.1103/PhysRevA.93.053401 (cit. on p. 21).

[93] P. J. Lee et al. "Phase Control of Trapped Ion Quantum Gates". In: *Journal of Optics B: Quantum and Semiclassical Optics* 7.10 (2005), S371. DOI: 10.1088/1464-4266/7/10/025 (cit. on pp. 103, 105, 117, 118).

[94] D. Leibfried et al. "Experimental Demonstration of a Robust, High-Fidelity Geometric Two Ion-Qubit Phase Gate". In: *Nature* 422.6930 (Mar. 2003), pp. 412–415. DOI: 10.1038/nature01492 (cit. on pp. 3, 102).

[95] D. Leibfried et al. "Quantum Dynamics of Single Trapped Ions". In: *Reviews of Modern Physics* 75.1 (Mar. 2003), pp. 281–324. DOI: 10.1103/RevModPhys.75.281 (cit. on pp. 9, 10, 12, 16, 20, 22, 63).

[96] D. Leibfried et al. "Toward Heisenberg-Limited Spectroscopy with Multiparticle Entangled States". In: *Science* 304.5676 (June 2004), pp. 1476–1478. DOI: 10.1126/science.1097576 (cit. on pp. 106, 130).

[97] D. Leibfried et al. "Transport Quantum Logic Gates for Trapped Ions". In: *Physical Review A* 76.3 (Sept. 2007), p. 032324. DOI: 10.1103/PhysRevA.76.032324 (cit. on p. 68).

[98] B. Lekitsch et al. "Blueprint for a Microwave Trapped Ion Quantum Computer". In: *Science Advances* 3.2 (Feb. 2017), e1601540. DOI: 10.1126/sciadv.1601540 (cit. on p. 3).

[99] F. M. Leupold et al. "Sustained State-Independent Quantum Contextual Correlations from a Single Ion". In: *Physical Review Letters* 120.18 (May 2018), p. 180401. DOI: 10.1103/PhysRevLett.120.180401 (cit. on pp. 33, 131).

[100] D. A. Lidar. *Review of Decoherence Free Subspaces, Noiseless Subsystems, and Dynamical Decoupling*. arXiv E-Print 1208.5791. Aug. 2012 (cit. on p. 5).

[101] Y. Lin et al. "Dissipative Production of a Maximally Entangled Steady State of Two Quantum Bits". In: *Nature* 504.7480 (Dec. 2013), pp. 415–418. DOI: 10.1038/nature12801 (cit. on p. 109).

[102] F. Lindenfelser et al. "Cooling Atomic Ions with Visible and Infra-Red Light". In: *New Journal of Physics* 19.6 (2017), p. 063041. DOI: 10.1088/1367-2630/aa7150 (cit. on pp. 20, 131).

[103] F. Lindenfelser. "Broadband Cooling on a Forbidden Transition in a Novel High-Optical-Access Ion Trap". Doctoral Thesis. ETH Zürich, 2017 (cit. on pp. 20, 33).

[104] K. Liu et al. "Precisely Synchronous and Cascadable Multi-Channel Arbitrary Waveform Generator". In: *Review of Scientific Instruments* 88.3 (Mar. 2017), p. 035110. DOI: 10.1063/1.4978067 (cit. on p. 29).

[105] H.-Y. Lo. "Creation of Squeezed Schrödinger's Cat States in a Mixed-Species Ion Trap". Doctoral Thesis. ETH Zürich, 2015 (cit. on pp. 7, 16–18, 21, 22, 24, 51, 55, 65, 90, 114, 136).

[106] H.-Y. Lo et al. "Spin-Motion Entanglement and State Diagnosis with Squeezed Oscillator Wavepackets". In: *Nature* 521.7552 (May 2015), pp. 336–339. DOI: 10.1038/nature14458 (cit. on pp. 12, 23, 40, 131).

[107] D. M. Lucas et al. "Oxford Ion-Trap Quantum Computing Project". In: *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 361.1808 (July 2003), pp. 1401–1408. DOI: 10.1098/rsta.2003.1209 (cit. on p. 3).

[108] M-Labs. *ARTIQ: Advanced Real Time Infrastructure for Quantum Physics*. http://m-labs.hk/artiq/. 2016 (cit. on p. 49).

[109] L.-S. Ma et al. "Delivering the Same Optical Frequency at Two Places: Accurate Cancellation of Phase Noise Introduced by an Optical Fiber or Other Time-Varying Path". In: *Optics Letters* 19.21 (Nov. 1994), pp. 1777–1779. DOI: 10.1364/OL.19.001777 (cit. on p. 51).

[110] B. MacDonald-de Neeve. "Calibrating an Ion Trap Quantum Computer". Masters Thesis. ETH Zürich, 2017 (cit. on pp. 4, 52, 97, 99).

[111] B. MacDonald-de Neeve. *Notes on Adaptive Robust Phase Estimation Gain Derivation*. Tech. rep. 2018 (cit. on p. 95).

[112] F. G. Major, V. N. Gheorghe and G. Werth. *Charged Particle Traps: Physics and Techniques of Charged Particle Field Confinement*. Springer Series on Atomic, Optical, and Plasma Physics. Berlin Heidelberg: Springer-Verlag, 2005 (cit. on p. 3).

[113] T. Manovitz et al. "Fast Dynamical Decoupling of the M\o{}lmer-S\o{}rensen Entangling Gate". In: *Physical Review Letters* 119.22 (Nov. 2017), p. 220505. DOI: 10.1103/PhysRevLett.119.220505 (cit. on pp. 5, 106, 133).

[114] M. Marinelli. "Title Undecided". in preparation. PhD Thesis. ETH Zürich, 2019 (cit. on pp. 24, 48, 58, 65, 114, 121, 125, 126).

[115] D. M. Meekhof et al. "Experimental Creation and Measurement of Motional Quantum States of a Trapped Ion". In: *Brazilian Journal of Physics* 27 (June 1997), pp. 178–192 (cit. on pp. 3, 12).

[116] K. K. Mehta et al. "Integrated Optical Addressing of an Ion Qubit". In: *Nature Nanotechnology* 11.12 (Dec. 2016), pp. 1066–1070. DOI: 10.1038/nnano.2016.139 (cit. on pp. 3, 105, 132).

[117] Z. Meir. "Dynamics of a Single Ground-State Cooled and Trapped Ion Colliding with Ultracold Atoms: A Micromotion Tale". PhD Thesis. Weizmann Institute of Science, 2016 (cit. on p. 65).

[118] J. C. Messtechnik. *ADwin Real-Time Control System*. http://www.adwin.de/index-us.html. 2018 (cit. on p. 29).

[119] G. J. Milburn, S. Schneider and D. F. V. James. "Ion Trap Quantum Computing with Warm Ions". In: *Fortschritte der Physik* 48.9-11 (Sept. 2000), pp. 801–810. DOI: 10.1002/1521-3978(200009)48:9/11<801::AID-PROP801>3.0.CO;2-1 (cit. on p. 102).

[120] N. Moll et al. "Quantum Optimization Using Variational Algorithms on Near-Term Quantum Devices". In: *arXiv:1710.01022 [quant-ph]* (Oct. 2017). arXiv: 1710.01022 [quant-ph] (cit. on p. 2).

[121] K. Mølmer and A. Sørensen. "Multiparticle Entanglement of Hot Trapped Ions". In: *Physical Review Letters* 82.9 (Mar. 1999), pp. 1835–1838. DOI: 10.1103/PhysRevLett.82.1835 (cit. on pp. 102, 103).

[122] C. Monroe and J. Kim. "Scaling the Ion Trap Quantum Processor". In: *Science* 339.6124 (Aug. 2013), pp. 1164–1169. DOI: 10.1126/science.1231298 (cit. on pp. 2, 3, 132).

[123] C. Monroe et al. "Demonstration of a Fundamental Quantum Logic Gate". In: *Physical Review Letters* 75.25 (Dec. 1995), pp. 4714–4717. DOI: 10.1103/PhysRevLett.75.4714 (cit. on pp. 3, 102).

[124] C. Monroe et al. "Large-Scale Modular Quantum-Computer Architecture with Atomic Memory and Photonic Interconnects". In: *Physical Review A* 89.2 (Feb. 2014), p. 022317. DOI: 10.1103/PhysRevA.89.022317 (cit. on p. 3).

[125] C. Monroe et al. "Resolved-Sideband Raman Cooling of a Bound Atom to the 3D Zero-Point Energy". In: *Physical Review Letters* 75.22 (Nov. 1995), pp. 4011–4014. DOI: 10.1103/PhysRevLett.75.4011 (cit. on p. 3).

[126] T. Monz et al. "Realization of Universal Ion-Trap Quantum Computation with Decoherence-Free Qubits". In: *Physical Review Letters* 103.20 (Nov. 2009), p. 200503. DOI: 10.1103/PhysRevLett.103.200503 (cit. on p. 106).

[127] T. Monz. "Quantum Information Processing beyond Ten Ion-Qubits". PhD Thesis. University of Innsbruck, 2011 (cit. on pp. 56, 57, 101).

[128] T. Monz et al. "Realization of a Scalable Shor Algorithm". In: *Science* 351.6277 (Mar. 2016), pp. 1068–1070. DOI: 10.1126/science.aad9480 (cit. on pp. 6, 116).

[129] G. Morigi et al. "Laser Cooling of Two Trapped Ions: Sideband Cooling beyond the Lamb-Dicke Limit". In: *Physical Review A* 59.5 (May 1999), pp. 3797–3808. DOI: 10.1103/PhysRevA.59.3797 (cit. on p. 3).

[130] G. Morigi, J. Eschner and C. H. Keitel. "Ground State Laser Cooling Using Electromagnetically Induced Transparency". In: *Physical Review Letters* 85.21 (Nov. 2000), pp. 4458–4461. DOI: 10.1103/PhysRevLett.85.4458 (cit. on p. 21).

[131] A. H. Myerson et al. "High-Fidelity Readout of Trapped-Ion Qubits". In: *Physical Review Letters* 100.20 (May 2008), p. 200502. DOI: 10.1103/PhysRevLett.100.200502 (cit. on p. 58).

[132] D. Nadlinger. "Experimental Control and Benchmarking for Single-Qubit Trapped-Ion Transport Gates". Semester Thesis. ETH Zürich, 2016 (cit. on pp. 32, 46, 47, 51, 70, 71).

[133] H. C. Nägerl et al. "Investigating a Qubit Candidate: Spectroscopy on the S1/2 to D5/2 Transition of a Trapped Calcium Ion in a Linear Paul Trap". In: *Physical Review A* 61 (Feb. 2000), p. 023405. DOI: 10.1103/PhysRevA.61.023405 (cit. on p. 3).

[134] C. M. Natarajan, M. G. Tanner and R. H. Hadfield. "Superconducting Nanowire Single-Photon Detectors: Physics and Applications". In: *Superconductor Science and Technology* 25.6 (2012), p. 063001. DOI: 10.1088/0953-2048/25/6/063001 (cit. on p. 132).

[135] V. Negnevitsky and L. D. Turner. "Wideband Laser Locking to an Atomic Reference with Modulation Transfer Spectroscopy". In: *Optics Express* 21.3 (Feb. 2013), pp. 3103–3113. DOI: 10.1364/OE.21.003103 (cit. on p. 51).

[136] V. Negnevitsky et al. "Repeated Multi-Qubit Readout and Feedback with a Mixed-Species Trapped-Ion Register". In: *arXiv:1804.09703 [physics, physics:quant-ph]* (Apr. 2018). arXiv: 1804.09703 [physics, physics:quant-ph] (cit. on p. 131).

[137] N. H. Nickerson, J. F. Fitzsimons and S. C. Benjamin. "Freely Scalable Quantum Technologies Using Cells of 5-to-50 Qubits with Very Lossy and Noisy Photonic Links". In: *Physical Review X* 4.4 (Dec. 2014), p. 041041. DOI: 10.1103/PhysRevX.4.041041 (cit. on p. 3).

[138] X. R. Nie, C. F. Roos and D. F. V. James. "Theory of Cross Phase Modulation for the Vibrational Modes of Trapped Ions". In: *Physics Letters A* 373.4 (Jan. 2009), pp. 422–425. DOI: 10.1016/j.physleta.2008.11.045 (cit. on p. 110).

[139] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, Oct. 2000 (cit. on pp. 1, 5, 16, 100–102, 115).

[140] D. Nigg et al. "Quantum Computations on a Topologically Encoded Qubit". In: *Science* 345.6194 (July 2014), pp. 302–305. DOI: 10.1126/science.1253742 (cit. on p. 116).

[141] N. D. Oppong. "Cancelation of Optical Phase Noise Induced by an Optical Fiber". Semester Thesis. ETH Zürich, 2015 (cit. on p. 51).

[142] C. Ospelkaus et al. "Microwave Quantum Logic Gates for Trapped Ions". In: *Nature* 476.7359 (Aug. 2011), pp. 181–184. DOI: 10.1038/nature10290 (cit. on p. 3).

[143] C. Ospelkaus et al. "Trapped-Ion Quantum Logic Gates Based on Oscillating Magnetic Fields". In: *Physical Review Letters* 101.9 (Aug. 2008), p. 090502. DOI: 10.1103/PhysRevLett.101.090502 (cit. on p. 3).

[144] R. Oswald. "Velocity Control of Trapped Ions for Transport Quantum Logic Gates". Masters Thesis. ETH Zürich, 2015 (cit. on pp. 68, 72, 80).

[145] R. Ozeri et al. "Errors in Trapped-Ion Quantum Gates Due to Spontaneous Photon Scattering". In: *Physical Review A* 75.4 (Apr. 2007), p. 042329. DOI: 10.1103/PhysRevA.75.042329 (cit. on p. 106).

[146] R. Ozeri. "The Trapped-Ion Qubit Tool Box". In: *Contemporary Physics* 52.6 (2011), pp. 531–550. DOI: 10.1080/00107514.2011.603578 (cit. on p. 17).

[147] M. Palmero et al. "Fast Separation of Two Trapped Ions". In: *New Journal of Physics* 17.9 (Sept. 2015), p. 093031. DOI: 10.1088/1367-2630/17/9/093031 (cit. on p. 85).

[148] W. Paul. "Electromagnetic Traps for Charged and Neutral Particles". In: *Reviews of Modern Physics* 62.3 (July 1990), pp. 531–540. DOI: 10.1103/RevModPhys.62.531 (cit. on pp. 2, 8).

[149] W. Pfaff et al. "Demonstration of Entanglement-by-Measurement of Solid-State Qubits". In: *Nature Physics* 9.1 (Jan. 2013), pp. 29–33. DOI: 10.1038/nphys2444 (cit. on pp. 6, 29, 116).

[150] P. Pham. "A General-Purpose Pulse Sequencer for Quantum Computing". Masters Thesis. Massachusetts Institute of Technology, 2005 (cit. on p. 29).

[151] PICMG consortium. *MicroTCA Overview*. 2017 (cit. on p. 34).

[152] C. Piltz et al. "A Trapped-Ion-Based Quantum Byte with 10^{-5} next-Neighbour Cross-Talk". In: *Nature Communications* 5 (Aug. 2014), p. 4679. DOI: 10.1038/ncomms5679 (cit. on p. 30).

[153] C. Piltz et al. "Versatile Microwave-Driven Trapped Ion Spin System for Quantum Information Processing". In: *Science Advances* 2.7 (July 2016), e1600093. DOI: 10.1126/sciadv.1600093 (cit. on p. 30).

[154] J. Preskill. "Reliable Quantum Computers". In: *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* 454.1969 (Jan. 1998), pp. 385–410. DOI: 10.1098/rspa.1998.0167 (cit. on p. 5).

[155] T. C. Ralph et al. "Quantum Nondemolition Measurements for Quantum Information". In: *Physical Review A* 73.1 (Jan. 2006), p. 012113. DOI: 10.1103/PhysRevA.73.012113 (cit. on p. 120).

[156] N. F. Ramsey. "A Molecular Beam Resonance Method with Separated Oscillating Fields". In: *Physical Review* 78.6 (June 1950), pp. 695–699. DOI: 10.1103/PhysRev.78.695 (cit. on p. 89).

[157] N. F. Ramsey. "Experiments with Separated Oscillatory Fields and Hydrogen Masers". In: *Reviews of Modern Physics* 62.3 (July 1990), pp. 541–552. DOI: 10.1103/RevModPhys.62.541 (cit. on p. 89).

[158] *Raspberry Pi*. https://www.raspberrypi.org/. 2018 (cit. on p. 50).

[159] R. Raussendorf and H. J. Briegel. "A One-Way Quantum Computer". In: *Physical Review Letters* 86.22 (May 2001), pp. 5188–5191. DOI: 10.1103/PhysRevLett.86.5188 (cit. on p. 130).

[160] M. D. Reed et al. "Realization of Three-Qubit Quantum Error Correction with Superconducting Circuits". In: *Nature* 482.7385 (Feb. 2012), pp. 382–385. DOI: 10.1038/nature10786 (cit. on p. 116).

[161] F. Riehle. *Frequency Standards: Basics and Applications*. John Wiley & Sons, Mar. 2006 (cit. on p. 89).

[162] D. Ristè et al. "Deterministic Entanglement of Superconducting Qubits by Parity Measurement and Feedback". In: *Nature* 502.7471 (Oct. 2013), pp. 350–354. DOI: 10.1038/nature12513 (cit. on pp. 6, 116).

[163] C. F. Roos et al. "Bell States of Atoms with Ultralong Lifetimes and Their Tomographic State Analysis". In: *Physical Review Letters* 92.22 (June 2004), p. 220402. DOI: 10.1103/PhysRevLett.92.220402 (cit. on pp. 106, 107).

[164] C. F. Roos et al. "Experimental Demonstration of Ground State Laser Cooling with Electromagnetically Induced Transparency". In: *Physical Review Letters* 85.26 (Dec. 2000), pp. 5547–5550. DOI: 10.1103/PhysRevLett.85.5547 (cit. on p. 21).

[165] C. F. Roos et al. "Nonlinear Coupling of Continuous Variables at the Single Quantum Level". In: *Physical Review A* 77.4 (Apr. 2008), p. 040302. DOI: 10.1103/PhysRevA.77.040302 (cit. on pp. 109, 110).

[166] C. Roos. "Controlling the Quantum State of Trapped Ions". PhD Thesis. University of Innsbruck, 2000 (cit. on p. 20).

[167] C. F. Roos. "Ion Trap Quantum Gates with Amplitude-Modulated Laser Beams". In: *New Journal of Physics* 10.1 (Jan. 2008), p. 013002. DOI: 10.1088/1367-2630/10/1/013002 (cit. on pp. 102, 104).

[168] T. Rosenband et al. *Ionizer Source Code Repository*. https://github.com/nist-ionstorage/ionizer. 2010 (cit. on p. 29).

[169] M. A. Rowe et al. "Transport of Quantum States and Separation of Ions in a Dual RF Ion Trap". In: *Quantum Info. Comput.* 2.4 (June 2002), pp. 257–271 (cit. on pp. 3, 73).

[170] K. Rudinger et al. "Experimental Demonstration of a Cheap and Accurate Phase Estimation". In: *Physical Review Letters* 118.19 (May 2017), p. 190502. DOI: 10.1103/PhysRevLett.118.190502 (cit. on p. 98).

[171] T. Ruster et al. "Entanglement-Based Dc Magnetometry with Separated Ions". In: *Physical Review X* 7.3 (Sept. 2017), p. 031050. DOI: 10.1103/PhysRevX.7.031050 (cit. on p. 107).

[172] T. Ruster et al. "Experimental Realization of Fast Ion Separation in Segmented Paul Traps". In: *Physical Review A* 90.3 (Sept. 2014), p. 033410. DOI: 10.1103/PhysRevA.90.033410 (cit. on pp. 76, 84).

[173] C. A. Sackett et al. "Experimental Entanglement of Four Particles". In: *Nature* 404.6775 (Mar. 2000), pp. 256–259. DOI: 10.1038/35005011 (cit. on p. 101).

[174] M. Sasura and V. Buzek. "Cold Trapped Ions as Quantum Information Processors". In: *arXiv:quant-ph/0112041* (Dec. 2001). arXiv: quant-ph/0112041 (cit. on p. 3).

[175] P. Schindler. "Frequency Synthesis and Pulse Shaping for Quantum Information Processing with Trapped Ions". Diploma Thesis. University of Innsbruck, 2008 (cit. on pp. 29, 37, 40, 42).

[176] P. Schindler et al. "A Quantum Information Processor with Trapped Ions". In: *New Journal of Physics* 15.12 (2013), p. 123012. DOI: 10.1088/1367-2630/15/12/123012 (cit. on p. 3).

[177] P. Schindler et al. "Experimental Repetitive Quantum Error Correction". In: *Science* 332.6033 (May 2011), pp. 1059–1061. DOI: 10.1126/science.1203329 (cit. on p. 116).

[178] F. Schmidt-Kaler et al. "Realization of the Cirac–Zoller Controlled-NOT Quantum Gate". In: *Nature* 422.6930 (Mar. 2003), pp. 408–411. DOI: 10.1038/nature01494 (cit. on p. 102).

[179] N. Schwegler. "Fast Parallel Readout of Multiple Trapped Ions". in preparation. Masters Thesis. ETH Zürich, 2018 (cit. on p. 27).

[180] J. H. Shirley. "Modulation Transfer Processes in Optical Heterodyne Saturation Spectroscopy". In: *Optics Letters* 7.11 (Nov. 1982), pp. 537–539. DOI: 10.1364/OL.7.000537 (cit. on p. 51).

[181] P. W. Shor. "Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer". In: *SIAM J. Sci. Statist. Comput.* 26.1484 (1997). DOI: 10.1137/S0097539795293172 (cit. on p. 2).

[182] P. W. Shor. "Scheme for Reducing Decoherence in Quantum Computer Memory". In: *Physical Review A* 52.4 (Oct. 1995), R2493–R2496. DOI: 10.1103/PhysRevA.52.R2493 (cit. on pp. 5, 115, 116).

[183] D. H. Slichter et al. "UV-Sensitive Superconducting Nanowire Single Photon Detectors for Integration in an Ion Trap". In: *Optics Express* 25.8 (Apr. 2017), pp. 8705–8720. DOI: 10.1364/OE.25.008705 (cit. on p. 3).

[184] A. Sørensen and K. Mølmer. "Entanglement and Quantum Computation with Ions in Thermal Motion". In: *Physical Review A* 62.2 (July 2000), p. 022311. DOI: 10.1103/PhysRevA.62.022311 (cit. on p. 103).

[185] A. Sørensen and K. Mølmer. "Quantum Computation with Ions in Thermal Motion". In: *Physical Review Letters* 82.9 (Mar. 1999), pp. 1971–1974. DOI: 10.1103/PhysRevLett.82.1971 (cit. on pp. 3, 102).

[186] M. Stadler. "Integrated Laser Amplitude Stabilization for Mixed-Species Trapped-Ion Experiments". in preparation. Masters Thesis. ETH Zürich, 2018 (cit. on pp. 33, 38).

[187] A. Steane. "The Ion Trap Quantum Information Processor". In: *Applied Physics B* 64.6 (June 1997), pp. 623–643. DOI: 10.1007/s003400050225 (cit. on p. 132).

[188] A. M. Steane. "Error Correcting Codes in Quantum Theory". In: *Physical Review Letters* 77.5 (July 1996), pp. 793–797. DOI: 10.1103/PhysRevLett.77.793 (cit. on pp. 5, 115).

[189] A. Steane et al. "Speed of Ion-Trap Quantum-Information Processors". In: *Physical Review A* 62.4 (Sept. 2000), p. 042305. DOI: 10.1103/PhysRevA.62.042305 (cit. on p. 132).

[190] A. M. Steane. "Efficient Fault-Tolerant Quantum Computing". In: *Nature* 399.6732 (May 1999), pp. 124–126. DOI: 10.1038/20127 (cit. on pp. 5, 116).

[191] A. M. Steane. "Overhead and Noise Threshold of Fault-Tolerant Quantum Error Correction". In: *Physical Review A* 68.4 (Oct. 2003), p. 042322. DOI: 10.1103/PhysRevA.68.042322 (cit. on p. 132).

[192] T. R. Tan et al. "Demonstration of a Dressed-State Phase Gate for Trapped Ions". In: *Physical Review Letters* 110.26 (June 2013), p. 263002. DOI: 10.1103/PhysRevLett.110.263002 (cit. on pp. 5, 106).

[193] T. R. Tan et al. "Multi-Element Logic Gates for Trapped-Ion Qubits". In: *Nature* 528.7582 (Dec. 2015), pp. 380–383. DOI: 10.1038/nature16186 (cit. on p. 6).

[194] B. M. Terhal. "Quantum Error Correction for Quantum Memories". In: *Reviews of Modern Physics* 87.2 (Apr. 2015), pp. 307–346. DOI: 10.1103/RevModPhys.87.307 (cit. on pp. 4, 126).

[195] Q. A. Turchette et al. "Decoherence and Decay of Motional Quantum States of a Trapped Atom Coupled to Engineered Reservoirs". In: *Physical Review A* 62.5 (Oct. 2000), p. 053807. DOI: 10.1103/PhysRevA.62.053807 (cit. on p. 110).

[196] Q. A. Turchette et al. "Heating of Trapped Ions from the Quantum Ground State". In: *Physical Review A* 61.6 (May 2000), p. 063418. DOI: 10.1103/PhysRevA.61.063418 (cit. on p. 64).

[197] N. V. Vitanov et al. "Fault-Tolerant Hahn-Ramsey Interferometry with Pulse Sequences of Alternating Detuning". In: *Physical Review A* 91.3 (Mar. 2015), p. 033406. DOI: 10.1103/PhysRevA.91.033406 (cit. on p. 30).

[198] M. M. Waldrop. "The Chips Are down for Moore's Law". In: *Nature News* 530.7589 (Feb. 2016), p. 144. DOI: 10.1038/530144a (cit. on p. 1).

[199] A. Walther et al. "Controlling Fast Transport of Cold Trapped Ions". In: *Physical Review Letters* 109.8 (Aug. 2012), p. 080501. DOI: 10.1103/PhysRevLett.109.080501 (cit. on pp. 73, 80, 81).

[200] N. Widmer. "Adaptive Bayesian Calibration of Two Rabi Frequencies from Collective Two-Qubit Measurements". in preparation. Masters Thesis. ETH Zürich, 2018 (cit. on pp. 4, 62).

[201] N. Wiebe and C. E. Granade. "Efficient Bayesian Phase Estimation". In: *arXiv:1508.00869 [quant-ph]* (Aug. 2015). arXiv: 1508.00869 [quant-ph] (cit. on p. 99).

[202] D. J. Wineland and D. Leibfried. "Quantum Information Processing and Metrology with Trapped Ions". In: *Laser Physics Letters* 8.3 (Mar. 2011), p. 175. DOI: 10.1002/lapl.201010125 (cit. on p. 3).

[203] D. J. Wineland et al. "Double-Resonance and Optical-Pumping Experiments on Electromagnetically Confined, Laser-Cooled Ions". In: *Optics Letters* 5.6 (June 1980), pp. 245–247. DOI: 10.1364/OL.5.000245 (cit. on p. 16).

[204] D. J. Wineland et al. "Experimental Issues in Coherent Quantum-State Manipulation of Trapped Atomic Ions". In: *Journal of Research of the National Institute of Standards and Technology* 103.3 (May 1998), p. 259 (cit. on pp. 2–4, 8–10, 12, 68, 89).

[205] D. J. Wineland. *Nobel Lecture: Superposition, Entanglement, and Raising Schroedinger´s Cat*. https://www.nobelprize.org/nobel_prizes/physics/laureates/2012/wineland-lecture.html. 2012 (cit. on p. 2).

[206] R. Yarbrough. *RF-Consultant*. https://www.rf-consultant.com/. 2017 (cit. on p. 51).

[207] C. Zhang. "Title Undecided". in preparation. PhD Thesis. ETH Zürich, 2019 (cit. on p. 33).