

Novel Entangling Gates and Scalable Trap Designs for Trapped-Ion Quantum
Computing

by

ALEXANDER D. QUINN

A dissertation accepted and approved in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
in Physics

Dissertation Committee:

Hailin Wang	Chair
David Allcock	Advisor & Core member
Jens Nöckel	Core Member
James Prell	Institutional Representative

University of Oregon
Summer 2024

© 2024 Alexander D. Quinn
All rights reserved.

DISSERTATION ABSTRACT

Alexander D. Quinn

Doctor of Philosophy in Physics

Title: Novel Entangling Gates and Scalable Trap Designs for Trapped-Ion Quantum Computing

Trapped ions have received much attention as a platform for quantum computing, a purpose they may be well-suited for on account of their natural features: Ions of the same species are inherently identical; they can be manipulated using electric and magnetic fields via both their net charge, which allows external fields to couple to ions' center-of-mass motion, and via internal electronic transitions; and when they are isolated from the environment, their states can remain coherent for long time spans, at least by the standards of quantum information experiments. Despite these features, the ability to carry out useful quantum computing in trapped ions is limited, with two major practical constraints being 1) the difficulty of coherently controlling ions' states and 2) the limited physical scale of existing trapped-ion quantum computers, which typically hold no more than tens of ions. This document presents a set of projects meant to help address these limitations through different tracks, including the development and testing of control techniques for trapped ions and the design of new types of traps. Firstly, we present an entangling gate carried out on quantum bits (qubits) encoded in a set of electronic energy levels that have been relatively unexplored until recently and whose viability for quantum information processing may enable more efficient architectures for trapped ion quantum computing. Specifically, we entangled a pair of qubits encoded in two Zeeman sublevels of the $D_{5/2}$ metastable excited state of a pair of trapped $^{40}\text{Ca}^+$ ions using Raman laser

beams 10s of THz detuned from resonance to limit scattering rates (a fundamental error source in Raman gates). We demonstrate that high-fidelity gates (98.6(1)% subtracting state preparation and measurement error, and 99.1(1)% subtracting both SPAM and erasure) can be performed with this encoding scheme, show that the main source of error is technical noise, and employ a leakage detection scheme that allows decay or deshelling from the metastable level to be heralded, potentially making correcting this class of errors easier. After this, we shift focus from control techniques and discuss scalable trap design, focusing on a project to fabricate ion traps using 3D printing, a technique that could potentially enable microfabricated traps with high harmonicity, power efficiency, and depth of confinement relative to the 2D (planar) microfabricated traps widely used in efforts to scale up trapped-ion quantum computing. We design, simulate, fabricate, and carry out preliminary electrical testing on metallized trap prints, demonstrating some minimum viability of the technique, and show computationally that traps produced this way could have trapping characteristics similar to those of other 3D microfabricated designs. Finally, we consider scalable trap design for continuous-variable quantum computing (CVQC), a quantum computing scheme where, in trapped ions, information is encoded in the states of the vibrational modes of an ion crystal. A key requirement for universal CVQC is the ability to perform non-Gaussian operations, which can be difficult to carry out electronically. In this work, we present a 2D ion trap design that can perform non-Gaussian motional operations in an all-electronic way, with the design process accounting for the geometric limitations imposed by 2D traps. We consider some of the motional operations possible with this trap and estimate their associated coupling rates, finding that under a reasonable set of assumptions about operating parameters, coupling rates for non-Gaussian operations could be achieved that are comparable

to those previously achieved with all-electronic Gaussian operations could likely be achieved with this trap design. This dissertation includes co-authored, published material.

ACKNOWLEDGEMENTS

None of the work presented here would have been possible without the independent contributions (e.g. the fabrication/assembly and testing of subsystems of the experiment apparatus as the lab was first being built up), the collaboration (e.g. in the slow-converging feedback loop between simulation and experiment), and the patient guidance of the many members, past and present, of the Oregon Ions lab and those the lab has worked with.

TABLE OF CONTENTS

Chapter	Page
I. INTRODUCTION	27
1.1. A conceptual overview of this document	27
1.2. Trapped ions - a brief guide	29
1.2.1. Traps	29
1.2.2. Ions	31
1.3. Quantum computing	32
1.3.1. Building blocks for QC	33
1.3.2. Technical requirements	36
1.4. TIQC architecture	37
1.5. The <i>omg</i> architecture	40
1.5.1. Optical, metastable, and ground qubits	40
1.5.2. Using <i>omg</i> qubits together	41
1.6. Continuous-variable quantum computing	42
1.6.1. Building blocks for CVQC	42
1.6.2. Implementing CVQC in trapped ions	43
1.7. Scalable ion trap design	44
1.7.1. The QCCD architecture and surface electrode traps.	45
II. EXPERIMENTAL APPARATUS	47
2.1. Ion trap	47
2.1.1. Trap design	47
2.1.2. Trapping RF drive	48
2.1.3. Trapping DC potential	50

Chapter	Page
2.1.4. Shim potentials	51
2.1.5. UHV system	51
2.2. Magnetic field control	54
2.2.1. Magnetic field coils	54
2.2.2. Line noise mitigation	55
2.3. Imaging and state detection	57
2.3.1. Imaging system overview	58
2.3.2. Imaging system performance	60
2.4. Laser system	61
2.4.1. Rack-mounted laser boards	62
2.4.2. Laser frequency locking	65
2.4.3. AOM boards	67
2.5. Far-detuned lasers	69
2.5.1. 976 nm Raman lasers	70
2.5.2. 854 nm AC Stark shift laser	72
2.6. Beam delivery	72
2.6.1. Beam pointing	75
2.6.2. Beam focusing	76
2.6.3. Beam power control	78
2.6.4. Polarization control	80
2.7. Electronic control fields	81
2.7.1. Magnetic Rabi drive	81
2.7.2. Displacement drive	82
2.7.3. Squeeze/beamsplitter drives	83
2.7.4. Pulse shaping	83

Chapter	Page
2.8. Experiment control system	86
2.8.1. Control hardware	86
2.8.2. Clocking	86
2.8.3. DDSes and phase-coherent RF control	87
2.8.4. Experimental control codebase	89
2.9. Ion loading	91
2.9.1. Calcium oven	92
2.9.2. Neutral calcium detection	92
2.9.3. Photoionization	94
III. EXPERIMENTAL METHODS	97
3.1. Trapping	98
3.2. Cooling	101
3.2.1. Doppler cooling	102
3.2.2. EIT cooling	104
3.2.3. Pulsed sideband cooling	105
3.2.4. Cooling testing with thermometry	106
3.2.5. Micromotion compensation	107
3.3. m qubit preparation	109
3.3.1. Optical pumping	109
3.3.2. Preparation fidelity and limits	110
3.3.3. Isolating qubit states	112
3.4. State detection	114
3.4.1. State-selective deshelving	114
3.4.2. Distinguishing shelved and unshelved ions - Fluorescence	115
3.5. Coherent control - Spin	116

Chapter	Page
3.5.1. RF magnetic Rabi flopping	117
3.5.2. Raman Rabi flopping	121
3.5.3. Pulse types and coherent errors	123
3.5.4. Composite pulses	124
3.5.5. AC Stark shifts	127
3.6. Coherent control - Motion	128
3.6.1. General operations	130
3.6.2. Displacement	131
3.6.3. Single-mode squeezing and phase shifting	131
3.6.4. Two-mode squeezing and beamsplitter	132
3.7. Coherent control - Spin-motion coupling	133
3.7.1. Raman sidebands	133
3.7.2. Spin-dependent forces	136
3.8. Motional state characterization	138
3.9. Coherence time measurement	140
3.9.1. Ramsey experiment	140
3.9.2. Ramsey with spin echoes	141
3.9.3. Fock state Ramsey experiment	143
3.10. Routine calibration techniques	143
3.10.1. Coherent drive calibrations	143
3.10.2. Light shift calibrations	145
3.10.3. Optical frequency calibrations	146
3.10.4. Two-point calibrations	148
IV. RAMAN GATES IN METASTABLE QUBITS	150
4.1. Erasure errors in m qubits	151

Chapter	Page
4.1.1. Erasure conversion with fluorescence	152
4.1.2. Leakage pathways in m qubits	153
4.2. Geometric phase gates	154
4.2.1. Fidelity estimation	155
4.3. Gate errors	158
4.3.1. Error budget	158
4.3.2. Addressing dominant error sources	159
4.3.3. Errors from leakage detection	161
4.3.4. Outlook	164
V. 3D-PRINTED ION TRAPS	166
5.1. Introduction	167
5.2. Quantifying Trap Performance	169
5.3. Trench Trap Geometries	172
5.3.1. Simple Trench	174
5.3.2. Stacked Trench	175
5.3.3. Comparison Geometries - SET and Wafer Traps	176
5.3.4. Summary	177
5.4. 3D Printing	178
5.4.1. Dielectric Material	179
5.4.2. Initial Fabrication	180
5.4.2.1. Print	180
5.4.2.2. Metalization	182
5.4.2.3. Integrated Optics	182
5.5. Conclusion	182

Chapter	Page
VI. SURFACE ELECTRODE TRAPS FOR CVQC	189
6.1. Introduction	189
6.2. Single ion in a cubic potential	190
6.2.1. Hamiltonian	190
6.2.2. The $\omega_p = 0$ case (cubic phase gate)	191
6.2.3. The $\omega_p = \omega_x$ case	192
6.2.4. The $\omega_p = 3\omega_x$ case (trisqueezing)	192
6.3. Two ions in a cubic potential	193
6.3.1. Hamiltonian	195
6.3.2. The $\omega_p = 0$ case	198
6.3.3. The $\omega_p = \omega_c$ case	199
6.3.4. The $\omega_p = \omega_c + 2\omega_r$ case	200
6.3.5. The $\omega_p = \omega_c - 2\omega_r$ case (parametric oscillation)	200
6.3.6. The $\omega_p = 3\omega_c$ and $\omega_p = 3\omega_r$ cases (trisqueezing)	200
6.4. Trap design for generating cubic potentials	201
6.4.1. Multipole expansion - the hexapole potential	201
6.4.2. Surface-electrode design	203
6.4.2.1. Equivalent cylindrical electrode systems	203
6.4.2.2. Designing electrode systems to produce hexapoles	203
6.4.2.3. Integrating hexapole electrodes into surface trap	205
6.4.3. Axial cubic potentials - design limitations	209
6.5. Practical implementation	209
6.5.1. Trisqueezing	209
6.5.2. Parametric oscillation	211
6.6. Conclusion	212

Chapter	Page
VII. CONCLUSION	213
APPENDIX: GATE ERROR CALCULATIONS	214
A.1. Error budget	214
REFERENCES CITED	221

LIST OF FIGURES

Figure		Page
1.1.	A schematic illustration of the energy levels of an ion in a trap, showing two independent modes: A harmonic oscillator mode (left) and the internal electronic state of the ion (right). Line thicknesses in the right plot correspond to the number of Zeeman/hyperfine states in a given energy level. Solid arrows denote dipole transitions, while dashed arrows denote quadrupoles. Splittings between energy levels are not to scale.	32
1.2.	Examples of commonplace one- and two-qubit gates written in circuit notation (with each line representing an individual qubit over time) and the operation of the gate on a specific input state. a) The Hadamard gate, which rotates a single qubit. b) The controlled-NOT (CNOT) gate, which carries out a bit flip on the target qubit conditional on the state of the control qubit, generating entanglement.	35
1.3.	An illustration of the conceptual reasoning for using mixed-species ion chains. a) A pair of ions sharing a potential well can be entangled by driving both ions with laser fields. b) For carrying out auxilliary functions (e.g. sympathetic cooling of the ions being used for gates), additional ions need to be added to the chain. If these ions are of the same species as those used for gates, when driven with laser fields, they will emit photons that the other ions can absorb, destroying the information they contain. c) If the auxilliary functions are performed using ions of a different species with different optical transition frequencies, then photons scattered by one species will not be absorbed by the other.	38
1.4.	The three qubit types considered by the <i>omg</i> architecture.	41

Figure	Page
1.5. A cartoon illustration of a possible QCCD “unit cell,” which can be tiled to produce arbitrarily large trap networks. Dark blue represents DC electrodes, while pink represents trapping RF. Ions are shown in light blue. The illustration on the left shows the unit cell of a 2D/surface electrode implementation of a QCCD architecture, while the one on the right shows how the concept can be extended to three-dimensional traps.	46
2.1. Photographs of our trap, showing (left) the trap attached to a vacuum flange and connected to electrical feedthroughs and (middle) a close-up shot of the unmounted trap, with the rough location of the trapping zone marked with a blue box. The associated inset in a picture taken with a CMOS camera of a five-ion chain, with ion spacing marked. Finally, a COMSOL render (right) shows trapping RF electrodes marked in red.	48
2.2. Schematic illustration of the hardware setup used for estimating the thickness of a layer of electroplated metal by measuring the total charge passed through the system during electroplating.	49
2.3. Block diagram of the RF chain used to supply the RF trapping voltage. The trapping RF itself is supplied by the chain in red. Quadrupole motional drives (discussed further in Section 2.7.3 are combined with the trapping RF at the resonator.	51
2.4. Circuit diagram for the electrical connections to the DC trap electrodes (including AC drives applied to these electrodes), with endcap needles in blue, DC trapping electrodes in white, and compensation electrodes in light yellow.	52
2.5. a) CAD rendering of the vacuum chamber, associated support hardware, and components directly attached to this hardware, including magnetic field coils (Section 2.2.1 and b) mounts for beam delivery optics (Section 2.6).	53

Figure	Page
2.6. a) A schematic of the magnetic field coil arrangement used to set the quantization axis in the trap. Current is marked with white arrows, with arrow thickness corresponding the current magnitude. b) A photograph showing the five magnetic field coils used on the trap. Worth noting are the two smaller coils (white and blue) wrapped around the top left coil, which are used for line noise feedforward (Section 2.2.2).	55
2.7. Schematic illustrations of our two line noise mitigation schemes, explained in greater detail in text.	56
2.8. Schematic illustration (and photograph below) of our imaging system, with light emitted by the ion (a) being captured by an objective (b) , passed through a slit (c) to reduce stray light and a narrow-band filter (d) to prevent the system from registering light far away from the wavelength of the cycling transition of Ca^+ , collimated with a lens (f) in an adjustable barrel mount to be redirected by a flip-mounted mirror (g) , which can direct light to a focusing lens leading to the PMT (h) and (i) or to a focusing lens (j) and (k) heading to the camera (l)	59
2.9. Schematic of our overall laser system (exluding far-detuned lasers, discussed in Section 2.5) shown in terms of the output of a single laser breadboard. Free-space laser beams are shown with red and blue lines. Fiber optics are shown with solid black lines. Data connections are shown with dashed black lines.	62
2.10. Photograph of our laser rack, showing the drawers containing laser breadboards as well as an additional rack for housing laser controllers and a third for the wavemeter and OSA board. Inset shows a sample laser breadboard mounted in a drawer.	64
2.11. Schematic illustrations of the three main AOM configurations we use in the lab. These schematics omit the focusing/recollimating lenses that focus the beams into the AOMs.	69

Figure	Page
2.12. Schematics of the two 976 nm Raman laser setups we have used in our experiments, showing the diode laser sources (marked in dark red) all the way through to the ion (marked in light blue). These schematics omit details of the beam delivery optics.	70
2.13. Schematic of the laser system used to supply our high-power 854 nm beam used for generating AC Stark shifts on the non-qubit levels of the $D_{5/2}$ manifold.	72
2.14. Schematic of the optics leading from the output collimators on our optics table to our trap. Definitions of abbreviations are as follows - <i>PD: photodiode, GTP: Glan-Thompson polarizer, PM: piezo motorized mirror, PS: piezo stack mirror, GP: Glan polarizer.</i>	74
2.15. A photograph with labels of one of our outcoupler bases, with the front plate of the photodiode box removed to show internal hardware.	79
2.16. Schematic summary of the RF voltages applied to the trap for coherent driving, marking the trap electrodes (shown in cross-section) that they are connected to.	82
2.17. Block diagram showing the RF chain for three motional drives used in our experiments. <i>PS: pulse shaper, BS: beamsplitter.</i>	84
2.18. Schematic illustration of our two pulse-shaping schemes. a) Pulse shaping by multiplying a fixed-amplitude RF pulse with a nearly-trapezoidal DC envelope. b) Pulse shaping by running the DDS in RAM mode and generating arbitrarily shaped pulses, with the DC output on the AWG being unchanged.	85
2.19. Schematic illustration of our experiment control system. a) Block diagram showing how our experiment control hardware (blue) interfaces with other blocks of our experiment, including beam control (red), the trap/imaging system (yellow), and the host PC allowing users to interface with the system (purple). b) A schematic of the beam power control system we use (SU-Servo).	87

Figure	Page
2.20. Class inheritance diagram showing the hierarchy of classes in our in our experiment codebase, down to the level of code for running individual experiments (e.g. an experiment for driving Rabi flopping on a qubit, RabiFlopping). NDScan classes are shown in yellow, in-house support classes are shown in pink, and experiments are shown in blue. The inset in purple shows the sequence for a generic scan point. In both portions, the main block of code that differs between individual experiments is shown in dark blue.	90
2.21. a) Illustration of the relevant photoionization geometry, showing the neutral calcium beam coming from the oven and the PI beam relative to the trap (electrodes shown in black). The neutral calcium and PI beams are both in the plane of the page. b) Energy level diagram of the two-photon PI process.	91
2.22. a) Fluorescence spectrum data (black dots) on the 423 nm line for our neutral calcium beam with a manual fit including three calcium isotopes (individual isotopes shown with dashed colored lines, total fluorescence shown with a solid black line). b) Relative probability of loading $^{43}\text{Ca}^+$ at different 423 nm beam detunings and with different amounts of Doppler broadening.	94
3.1. A level diagram showing a Ca^+ ion trapped in a harmonic oscillator potential (only one dimension of which is shown). Transitions between internal electronic energy levels are marked with the wavelengths and polarizations that we use to drive them. Energy levels are not to scale.	98
3.2. A schematic illustration of how the RF electrodes of a linear Paul trap generate a harmonic effective potential.	101
3.3. A set of diagrams showing how a) Doppler cooling, b) EIT cooling, and c) pulsed sideband cooling are implemented in our setup.	102

Figure	Page
3.4. Schematic illustrations of techniques that we use to minimize micromotion. The offset direction each technique allows us to optimize is marked with a black double-headed arrow. The direction of the micromotion is marked with a double-headed red arrow. These techniques are a) minimization of position change with RF confinement change, b) minimization of Doppler broadening when driving with the 397π beam (blue arrow), and c) minimization of the micromotion sideband associated with our two-beam (red arrows) Raman Rabi drive.	108
3.5. a) The generic sequence for characterizing the effect of some operation (“Expt.”) on some known state, performing repeated experimental “shots” and calculating the probability of different outcomes. b) The specific pulse sequences of each experimental step (not showing cooling, which either happens before or is interleaved with state preparation).	110
3.6. Sample Rabi flopping, driven with a resonant RF magnetic field, showing two-level sinusoidal dynamics when the 854 LS beam is on (blue) and precession through the $D_{5/2}$ manifold before returning to the starting state when the 854 LS beam is off (red).	113
3.7. A sample count histogram for two-ion fluorescence detection, with the “zero bright” (blue), “one bright” (red), and “two bright” (green) Poisson distributions labelled above.	116
3.8. Schematic illustration of our RF magnetic Rabi drive, showing trap electrodes (pink for trap RF, blue for Rabi drive current), the quantization axis magnetic field \vec{B} (in the xz plane), and the magnetic Rabi drive \vec{B}_{drive} (in the xy plane).	120
3.9. Schematic illustration of the dynamics of a Λ system, showing a) the relationship between the three-level dynamics of the actual system and the effective two-level dynamics and b) the three-level system we use in our experiments.	122
3.10. The results of experiments testing composite a) $\pi/2$ and b) π pulses. Different colors denote time scans at different ϕ_{BB1} , showing that a value of ϕ_{BB1} can be selected such that the output of the pulse sequence is made insensitive to change in pulse times.	126

Figure	Page
3.11. A reference table for the motional operations performed in our experiments, showing in the phase space picture a) the Wigner function of the motional ground state, b) the effect of a displacement operation, c) the effect of squeezing, d) the effect of a phase shift on a displaced state, e) the effect of a beamsplitter operation on a pair of modes (with the Wigner function of the second mode shown in red), and f) the effect of two-mode squeezing on the correlations between the quadratures of a pair of modes.	130
3.12. Two techniques of achieving spin-motion coupling. a) Raman sideband transitions, showing the $\Delta n = +1$ (BSB) and $\Delta n = -1$ (RSB) transitions, with the spin states being represented as a pair of harmonic oscillator energy ladders separated by the spin state splitting. b) Spin-dependent forces, showing how two interfering beams with the same polarizations will generate periodic intensity gradients of different amplitudes (and different force magnitudes) on the two spin states.	134
3.13. Pulse sequences for Ramsey sequences for characterizing spin (left) and motion (right) coherence, with (bottom) and without (top) a spin echo pulse. Carrier pulses are shown in dark blue, blue sideband pulses in light blue. To generate fringes at a given t , the phase ϕ of the final pulse can be scanned.	142
3.14. Experiments for measuring the AC Stark shifts caused by a given beam (red). a) Carrying out Rabi spectroscopy while the beam is on and comparing this result to the result when the beam is off. b) Performing a Ramsey sequence with a spin echo, with the beam only being on during one delay so as to cause asymmetrical accumulation of phase.	146
3.15. Pulse sequences (with time running from left to right) of our four laser line calibrations, with the state of the ion at the end of the experiment (during a fluorescence check, FC) shown assuming that laser associated with the line being characterized is far off-resonance.	147
3.16. Illustration of the two-point calibration servo loop.	148

Figure	Page
4.1. Schematic illustration of our erasure error detection scheme in the context of a gate, showing how checks for 397 fluorescence can be used both to confirm successful state preparation prior to running a gate and check for leakage errors after a gate.	153
4.2. a) Gate pulse sequence showing beams relevant to the gate, with SDF phase ($0, \pi$) denoted by color (white, blue). b) A scan over SDF detuning from the LF OOP mode resonance, showing the crossing point in $ \downarrow\downarrow\rangle$ and $ \uparrow\uparrow\rangle$ populations at which the gate was performed. Inset shows high-shot time series data at the gate operating point. c) A sample parity fringe, with insets showing high-shot time-series data at the parity fringe extrema.	157
4.3. <i>a)</i> A flowchart illustrating the possible outcomes of a one-ion gate sequence with an FC carried out at the end, considering only leakage errors. Outcomes corresponding to an erasure error are labelled in yellow, while outcomes corresponding to a missed leakage error that is missed are labelled in pink. <i>b)</i> An accounting of the error probabilities in each category with and without leakage checks.	163
5.1. Cross-section of a trap fabricated using our proposed method of 3D printing an electrodes onto a wafer containing electrode routing and some of the potential integrated ion control elements that have been demonstrated.	168
5.2. Cross-sectional views of of the trap designs simulated. The ions' location is shown as an orange dot and the trap axis is out of the page. DC electrodes are blue, RF electrodes are red. Numerical apertures are shown in lavender.	173

- 5.3. Simulated parameters of surface-electrode traps vs the symmetric and anti-symmetric ‘simple trench’ trench geometries. (a) Trap depth at a fixed radial frequency. (b) The quadrupole component (C_2) of the trapping potential at the ion. (c) The hexapole component (C'_3) of the trapping potential. (d) The octopole component (C'_4) of the trapping potential. (e) Electrode dimensions at a constant ion-electrode separation for the symmetric simple trench trap. (f) Electrode dimensions at a constant ion-electrode separation for the anti-symmetric simple trench trap. (g,h) NA above and below the ion. Parameters are plotted against a generic geometric variable w , which is related to an electrode dimension of each trap as shown in the legend. Simulation results are shown with markers. The representative SET, wafer, and simple trench traps whose characteristics are summarized in Table 4 are highlighted with larger markers. Simulation points are connected with dashed lines only as a guide to eye. 184
- 5.4. Simulated parameters of surface-electrode traps and wafer traps vs the symmetric ‘stacked trench’ trap geometry. (a) Trap depth at a fixed radial frequency. (b) The quadrupole component (C_2) of the trapping potential at the ion. (c) The hexapole component (C'_3) of the trapping potential. (d) The octopole component (C'_4) of the trapping potential. (e) The hexapole component of the trapping potential, excluding the SET data to make the other series more clear. (f) Electrode dimensions at a constant ion-electrode separation for the symmetric stacked trench trap. (g,h) NA above and below the ion. Parameters are plotted against a generic geometric variable w , which is related to an electrode dimension of each trap as shown in the legend. Simulation results are shown with markers. The representative wafer and symmetric stacked trench traps whose characteristics are summarized in Table 4 are highlighted with larger markers. Simulation points are connected with dashed lines only as a guide to eye. 185

- 5.5. Simulated parameters of surface-electrode traps and wafer traps vs the anti-symmetric ‘stacked trench’ trap geometry. (a) Trap depth at a fixed radial frequency. (b) The quadrupole component (C_2) of the trapping potential at the ion. (c) The hexapole component (C'_3) of the trapping potential. (d) The octopole component (C'_4) of the trapping potential. (e) The hexapole component of the trapping potential, excluding the SET data to make the other series more clear. (f) Electrode dimensions at a constant ion-electrode separation for the anti-symmetric stacked trench trap. (g,h) NA above and below the ion. Parameters are plotted against a generic geometric variable w , which is related to an electrode dimension of each trap as shown in the legend. Simulation results are shown with markers. The representative wafer and anti-symmetric stacked trench traps whose characteristics are summarized in Table 4 are highlighted with larger markers. Simulation points are connected with dashed lines only as a guide to eye. 186
- 5.6. SEM images of gold coated ormocer trap structures printed on glass. In this structure, many trenches have been connected up with X-junctions, as is envisaged in a large ‘QCCD’ processor (Blakestad et al. (2009)). (a) An overhead view of a grid of symmetric stacked trenches showing segmentation of the top DC electrodes. (b) An oblique view of the same grid. (c) A close up view of an intersection of trenches. (d) A false-color version of (c) showing DC electrodes in blue and RF electrodes in red. (e) A close-up view of a DC electrode segment. 187
- 5.7. Proposed integration of electrode metalization with printing process shown here on a short section of the larger trap shown in Fig 5.6. (a) Silicon wafer with electrode routing and vias up to the surface is fabricated using standard CMOS process (metal shown copper coloured, SiO_2 dielectric shown transparent). (b) Aluminum is deposited and patterned to define gaps between electrodes. (c) DLW used to print ormocer trap structures (blue). (d) Gold is deposition from multiple angles to coat all surfaces of the print. (e) Aluminium is etched to allow liftoff of the gold over the gaps between electrodes. Insets shown the liftoff process in more detail. 188

Figure	Page
6.1. Plot of the Fock state distribution of the state generated by applying the $\omega_p = \omega_x$ operation to the motional ground state for a time $t = 0.23 \times 2\pi/g$. Components $n > 9$ not shown. <i>Inset</i> - Associated Wigner function.	193
6.2. Plot of the Fock state distribution of of the trisqueezed state, obtained by applying the trisqueezing Hamiltonian to the ground state for a time $t = 0.035 \times 2\pi/g$. Components $n > 9$ not shown. <i>Inset</i> - Associated Wigner function.	194
6.3. <i>Left</i> - Illustration of a two-ion chain with coordinates labelled. Equilibrium positions are marked in light blue. <i>Right</i> - Table illustrating two-ion vibrational normal modes.	195
6.4. Trapping potential generated by symmetrical surface-electrode trap around ion (marked in yellow). The decomposition of this potential into overlaid quadrupole, hexapole, and octopole terms are represented schematically.	202
6.5. Cross-section of an asymmetric surface-electrode trap through plane perpendicular to trap axis, with ion position marked by a yellow dot. Cross-section of equivalent cylindrical electrode system (partitioned ring of radius r_0) is overlaid on top, with the relationship between surface and cylindrical electrode positions $\vec{r}(\theta) = -r_0\hat{e}_x + r_0 \tan(\frac{\theta}{2})\hat{e}_y$ illustrated geometrically. Figure adapted from Wesenberg (Wesenberg (2008))	204
6.6. <i>a)</i> . Quadrupole system. <i>b)</i> . Hexapole system. <i>c)</i> . Combined electrode system for generating quadrupole trapping potential and hexapole trisqueezing potential. <i>Red</i> - Quadrupole electrodes. <i>Purple</i> - Hexapole electrodes. <i>Green</i> - Inner ground. <i>Blue</i> - Outer ground. Notably, to produce optimal performance under geometric constraints, $\theta_1 + \theta_2 = 60^\circ$, and the quadrupole electrodes are adjacent to hexapole electrodes.	207
6.7. <i>a)</i> . C_2 for the quadrupole electrodes and C_3 for the hexapole electrodes plotted against θ_1 when $\theta_1 + \theta_2 = 60^\circ$. <i>b)</i> . Widths of different electrodes in the combined electrode system of Fig. 6.6b, normalized to ion height, under the condition $\theta_1 + \theta_2 = 60^\circ$	208

Figure	Page
A.1. a) Motional Ramsey experiment, plotting the contrast between the $ 0\rangle$ and $ 1\rangle$ fock states as a function of the delay time between sideband $\pi/2$ pulses. b) Numerically simulated Bell-state infidelity due to a white motional noise source as a function of the coherence to gate time ratio. Orange lines represent the infidelity expected for measured coherence and gate times used in this work. Dashed lines represent the 68% confidence interval.	215
A.2. a) Bell-state infidelity corresponding to a normally distributed static offset in gate mode frequency with standard deviation σ for different orders of Walsh modulation. b) Distribution of Bell-state infidelities corresponding to a normal distribution of mode frequency offsets with a standard deviation of 100 Hz for a Walsh 1 4-loop gate operating at 11.6 kHz. Dashed orange lines represent 68% confidence interval.	217
A.3. a) Bell-state infidelity corresponding to a normally distributed static offset in gate mode frequency with standard deviation σ for different orders of Walsh modulation. b) Distribution of Bell-state infidelities corresponding to a normal distribution of mode frequency offsets with a standard deviation of 100 Hz for a Walsh 1 4-loop gate operating at 11.6 kHz. Dashed orange lines represent 68% confidence interval.	218
A.4. a) Population scattered out of the $m = +5/2$ and $+3/2$ in the $D_{5/2}$ manifold with either the 976 Rpi or Rsig beam on. b) Population scattered out of the qubit states with the 854 beam on.	220

LIST OF TABLES

Table	Page
1. A summary of projects worked on over the course of the last five years and which areas of trapped-ion quantum computing they are relevant to. My contribution to the projects in the top two rows was on the experimental side (hardware design/construction, writing/running/debugging experiments), while my contribution to the projects on the bottom was on the theoretical/computational side.	29
2. A summary of the components of our control system and what we use them for in our experiments.	88
3. Dominant sources of two-qubit gate infidelity as predicted by our error model. The lower half of the table summarizes the expected contributions due to erasure errors. Values in bold represent infidelity as measured in the gate	165
4. Summary of results for trap parameter calculations for the SET, simple trench trap ('Simple'), stacked trench trap ('Stacked'), and wafer trap. Values shown correspond to a set of geometric parameters selected to balance trap depth and quadrupole strength/purity.	177
5. The dimensions (in microns) of the traps whose parameters are summarized in Table 4.	177
6. Electrical and thermal properties of ormoCer.	179
A.7. Measured scattering rates out of $D_{5/2}$ (in Hz) when initialized in the $m = +3/2$ or $m = +5/2$ qubit states for various beams used in the gate. Measurements for the 854 LS beam were taken with two trapped ions.	219

CHAPTER I

INTRODUCTION

1.1 A conceptual overview of this document

This thesis describes a set of projects, largely independent of each other in terms of the work involved and the specific goals pursued, that may seem disconnected from one another. However, an important common thread runs between them: They are all ultimately focused on developing tools to use trapped ions in technological applications, whether quantum sensing or trapped-ion quantum computing (TIQC). In that light, these projects can be organized according to *how* they relate to the broader project of refining, expanding the capabilities of, and scaling up TIQC. There are two “axes” along which these projects can be organized: The type of TIQC they work to enable and the type of improvement they seek to make. Specifically, these projects can be in support of discrete-variable quantum computing (i.e. quantum computing using two-level Hilbert spaces or “qubits,” which will typically just be called “quantum computing” in the rest of the text) or of continuous-variable quantum computing (CVQC, i.e. quantum computing using infinite Hilbert spaces of evenly-spaced energy states). They also broadly fall under two parallel paths of development: the development and testing of new techniques for controlling the states of trapped ions (mostly carried out on the experimental side), and the design of novel trap geometries (carried out on the theoretical/computational side). These projects are listed specifically in Table 1. Not listed in the table is a fifth project investigating non-Hermitian quantum systems using our qubit platform. While this specific project was less motivated by technological applications, it demonstrates features of non-Hermitian dynamics with potential applications (e.g. the ability to implement fast and deterministic operations) that are worth discussion.

Projects not covered in detail. Of these five projects, two of them (observing of non-Hermitian dynamics in trapped-ion qubits Quinn et al. (2023) and carrying out trapped-ion motional interferometry with phase sensitivity below the standard quantum limit J. Metzner et al. (2024)) are not covered in detail in this thesis. Accounts of these projects can be found in their respective papers and in the thesis of Dr. Jeremy Metzner (see J. M. Metzner (2024)), which discusses them in greater detail (particularly on the theory side). Much of the content laid out in this introduction and in Chapters II and III still lays out the experimental apparatus and methods used to carry out these projects.

Currently unpublished projects. Our project to perform entangling gates in metastable qubits, covered in Chapter IV, is soon to be published. My role in this project was on the experimental side, developing and debugging experiments and setting up hardware (though two very important hardware systems enabling this specific project, a high power injection-locked diode laser system for performing gates 2.5.1 and a patched version of our experimental control system 2.8.3, were built and tested by my colleagues Sean Brudney and Gabe Gregory respectively). The work of simulating our gates and calculating an error budget was performed by Gabe Gregory. Important context for this project is provided by a set of photon scattering rate measurements, taken by me and analyzed and presented in the thesis of Dr. Daniel Moore (see Moore et al. (2023)). An entirely separate theoretical/computational project, developing a surface-electrode trap for carrying out non-Gaussian operations on the motional modes of trapped ions for CVQC applications, was carried out by me alone, with guidance from my advisor David Allcock.

Projects covered in existing publications. Chapter V is based on a paper currently available as a preprint, "Geometries and fabrication methods for 3D

	Discrete-variable quantum computing	Continuous-variable quantum computing (CVQC)
Control techniques	Single- and two-qubit Raman gates in metastable (m) qubits	Coherent control of ion motion for motional interferometry
Trap design	Design of 3D-printable ion traps for achieving scalability and favorable geometries	Design of surface-electrode traps (SETs) for carrying out non-Gaussian operations on ion motion

Table 1. A summary of projects worked on over the course of the last five years and which areas of trapped-ion quantum computing they are relevant to. My contribution to the projects in the top two rows was on the experimental side (hardware design/construction, writing/running/debugging experiments), while my contribution to the projects on the bottom was on the theoretical/computational side.

printing ion traps” (Quinn, Brown, Gardner, and Allcock (2022)). The portions of this paper focusing on theory and simulation were written by me, with David Allcock writing portions of the intro and conclusion and Morgan Brown writing the description of the 3D printing and metallization process. Morgan Brown and David Miller carried out the fabrication and characterization of test prints, and I simulated and analyzed the designs discussed in this chapter.

1.2 Trapped ions - a brief guide

Before discussing their technological applications in quantum information, a short discussion is offered below of what trapped ions are, how they can be modelled, and how, in broad strokes, they can be controlled. These topics, particularly as they relate to our own experiments, are discussed more thoroughly in Chapter III.

1.2.1 Traps. Because they have a net electric charge, ion experience a Lorentz force $\vec{F} = q(\vec{E} + \vec{v} \times \vec{B})$ and this force can be used to confine ions in

an entirely classical way. There are two specific types of traps widely in use for quantum information experiments: Penning traps and Paul traps.

Penning traps. Penning traps confine ions using a combination of static electric and magnetic fields. A pair of endcaps held at a static voltage offer confinement in one direction while a magnetic field parallel to the endcap axis forces ions to undergo cyclotron motion, providing confinement in the other two directions. The ions' continuous cyclotron motion (and the large magnetic fields required to confine them) makes them more difficult to work with in certain ways, and Penning traps are less widely used in research on scalable trapped-ion quantum computing than Paul traps (though there has been work done on microfabricated, scalable Penning traps. See Jain et al. (2024) among others).

Paul traps. Paul traps confine ions using only electric fields. By Earnshaw's theorem, it is impossible to achieve confinement in more than two directions at once using static electric fields (an immediate consequence of the fact that $\nabla \cdot \vec{E} = 0$ in the absence of charge). Paul traps work around this constraint by using dynamic electric fields instead of static ones. A deeper treatment of this is given in Section 3.1. In short though, if an AC multipole potential is applied to a charged particle, in certain parameter regimes, the time-averaged motion of the particle shows the particle experiencing an effective potential (or "pseudopotential") ϕ_{eff} proportional to the energy density of the AC electric field (i.e. $\phi_{\text{eff}} \propto |\vec{E}|^2$). For a "linear" Paul trap (shown schematically in Figure 3.2), the trap electrodes are very long relative to the spacing between them, and they generate a 2D potential in the trapping zone. This potential only confines in two directions, with confinement in the third direction being supplied by a DC electric field, as in Penning traps. Most Paul traps used in trapped-ion work generate, to lowest-order, an AC quadrupole potential, which in turn corresponds to a harmonic pseudopotential. An ion

trapped in this 2D harmonic pseudopotential and the 1D harmonic potential supplied by the DC electric field (which tends to be anti-trapping in the other two directions) is, effectively, a *3D harmonic oscillator*. This feature is important for much of the work done in trapped ions.

1.2.2 Ions. A neutral atom (or molecule) that has had one (or more) electrons removed (or added) is an ion. For all discussions below, “ion” will refer to a more narrow subset of this category: alkaline earth atoms that are singly-ionized, such that they have a single valence electron in a hydrogen-like potential. This accounts for most of the species used in trapped-ion experiments (though other elements like ytterbium have been used in experiments requiring the control and detection of individual ions (Balzer et al. (2006)), as have molecules (Drewsen, Mortensen, Martinussen, Staannum, and Sørensen (2004))).

Electronic level structure. The generic energy level structure of an alkaline earth ion is shown in Figure 1.1. The relevant energy levels are the electronic ground state $S_{1/2}$, the short-lived excited states $P_{1/2}$ and $P_{3/2}$, and the $D_{3/2}$ and $D_{5/2}$ excited states (which are long-lived in most ion species). The S and D levels can both couple to the P levels (where other selection rules are not violated) through electric dipole interactions. On the other hand, transitions between the D levels (also called metastable levels) and the S level are dipole-forbidden and can only take place through electric quadrupole interactions.

Zeeman/hyperfine splitting. Within an energy level (e.g. within $S_{1/2}$), the valence electron of the ion can occupy different angular momentum states based on the orbital angular momentum and its own inherent spin. Uncoupled to other systems, these states would be degenerate in energy. However, applying an external magnetic field will lift this degeneracy via the Zeeman effect, defining a quantization axis (i.e. an axis along which the angular momentum state is more




Figure 1.1. A schematic illustration of the energy levels of an ion in a trap, showing two independent modes: A harmonic oscillator mode (left) and the internal electronic state of the ion (right). Line thicknesses in the right plot correspond to the number of Zeeman/hyperfine states in a given energy level. Solid arrows denote dipole transitions, while dashed arrows denote quadrupoles. Splittings between energy levels are not to scale.

or less aligned with the magnetic field) and a set of non-degenerate states. In addition, if the ion nucleus has an inherent magnetic moment, this will also couple with the electron spin and generate hyperfine splitting, lifting the degeneracy of states based on the (anti-)alignment of the electron and nuclear spins. Both of these processes yield a set of energetically distinct magnetic sublevels in each energy level. In all of the work presented here, we use an ion species with no nuclear spin, meaning degeneracy in our system is lifted only by Zeeman splitting.

1.3 Quantum computing

A quantum computer is a device that processes information that is encoded in quantum states. Typically when people refer to quantum computing, they implicitly mean computing done using information encoded in discrete two-level systems (a quantum bit or “qubit”). Qubits can be encoded in the states of many quantum systems: the current through a superconducting circuit, the spin state of a nitrogen-vacancy center, or the polarization of an individual photon, among

other systems (see Cheng et al. (2023) for a broad overview of the state of field around the time of writing). While as of now, no quantum computing platform has seen the implementation of large-scale “fault-tolerant” quantum computing (described in greater detail later in this section), trapped ions are currently among the leading platforms, with some of the highest reported single-(Löschner et al. (2024)) and two-qubit (Clark et al. (2021); Löschner et al. (2024)) gate performances reported in the literature, as well as recent reports of better-than-break-even error correction in small-scale experiments (Hong, Durso-Sabina, Hayes, and Lucas (2024)). As such, there is currently much interest in scaling up trapped-ion quantum computing.

These next three sections will offer a broad overview of quantum computing. (More detailed discussions can be found in review papers or in books, e.g. Nielsen and Chuang (2010), with this source informing much of the discussion below), Quantum computing’s basic building blocks and the associated technical requirements that must be met for it to be practically useful, its typical implementation in trapped ions, and a recently proposed alternative trapped-ion architecture (the *omg* qubit scheme Allcock et al. (2021)) that motivates much of the work we do in our lab will be discussed.

1.3.1 Building blocks for QC. A quantum algorithm (a set of calculations run on a quantum computer) can be carried out by initializing a quantum register (a set of qubits) to some known state and carrying out a set of state-changing operations and measurements on it. The output of this sequence when the final state of the register is measured will be some tensor product of measurement basis eigenstates (usually written as $|0\rangle$ and $|1\rangle$ or $|\downarrow\rangle$ and $|\uparrow\rangle$). To determine the output state of the algorithm in this basis, the algorithm must (in

general) be repeated many times to build up statistics, though computations can also have deterministic outputs.

The state manipulations performed on qubits during a quantum algorithm are often called “gates,” in analogy to classical computing. It can be shown that a set of single-qubit gates together with a single two-qubit gate are “complete” (i.e. that any quantum algorithm can be carried out using a combination of these gates). The meanings of these gate types are clarified here.

Single-qubit gates. In classical computing, where an individual bit can only be in a 0 or 1 state, the only meaningful “single-bit gate” is a NOT operation, flipping the state of a bit. In quantum computing however, where information is encoded in a coherent superposition of states $|0\rangle$ and $|1\rangle$ ($|\Psi\rangle = C_0|0\rangle + C_1e^{i\phi}|1\rangle$), where $|C_0|^2 + |C_1|^2 = 1$), each qubit has two continuous degrees of freedom to control. In fact, the state of a qubit can be represented as a vector on a unit sphere (the Bloch sphere, discussed in greater detail in Section 3.5), and the general set of single-qubit gates can be regarded as “rotations” of this vector. An example of a single-qubit gate is the Hadamard (or H) gate, illustrated in Figure 1.2a.

Multi-qubit (entangling) gates. Entanglement is a key computational resource available to quantum systems. Multi-qubit gates (or entangling gates) generate non-classical correlations between two or more qubits by carrying out an operation that affects the state of one qubit differently depending on the state of the other. A prototypical example of this type of gate, the controlled-NOT (or CNOT) gate, is illustrated in Figure 1.2b.

Gate errors. Like in classical computing, errors can occur during algorithms, leading to incorrect output states. In the classical case, the only single-bit error is an unexpected bit flip. In the quantum case, this idea can be extended to a general set of undesired rotations of the state vector, which can be coherent (i.e. consistent




Figure 1.2. Examples of commonplace one- and two-qubit gates written in circuit notation (with each line representing an individual qubit over time) and the operation of the gate on a specific input state. **a)** The Hadamard gate, which rotates a single qubit. **b)** The controlled-NOT (CNOT) gate, which carries out a bit flip on the target qubit conditional on the state of the control qubit, generating entanglement.

between runs of the algorithm or over time scales allowing coherent cancellation) or incoherent (i.e. random over the same time scales). Other classes of error exist as well, with one notable class (discussed in greater detail in Chapter IV) being leakage errors, i.e. the system entering a state not in the two-level qubit manifold. These different error types can be handled in different ways with different levels of ease. For a gate, the error rate can be quantified in terms of a gate *fidelity*, which is the degree of overlap between the expected and true outputs of the gate, defined as $\mathcal{F} = \langle \Psi_T | \rho | \Psi_T \rangle$, where $|\Psi_T\rangle$ is the target output state and ρ is the density operator of the true output.

Error correction. Error correcting codes exist for the most common types of errors in quantum computing. Typically, these codes require that each qubit encoding data must be entangled with a set of auxiliary qubits whose states can be measured to check whether an error occurred on the data qubit. Together, these qubits are said to make up one “logical qubit” in an error-corrected quantum algorithm. At a given rate of gate errors, in order for a quantum computer to be fault-tolerant (i.e. able to push rates of logical errors to arbitrarily low levels), each logical qubit must consist of some minimum number of physical qubits. In practice, fault-tolerance is difficult to achieve because the number of physical qubits required

for each logical qubit grows quickly with error rate. This problem is discussed further below.

1.3.2 Technical requirements. For quantum computing to be practically useful, two distinct but related technical requirements must be considered: scalability and gate fidelity.

Scalability. A useful quantum computer is one that can perform certain calculations too difficult for classical computers (e.g. the factoring of large prime numbers or the simulation of large molecules). By this benchmark, a quantum computer must be large enough that it cannot be simulated classically. A quantum computer of two logical qubits can be “simulated” with pencil and paper. One made up of 10s can be simulated in a reasonable amount of time on a supercomputer. Beyond this, a quantum computer can perform calculations that are classically infeasible. Quantum computers with 10s to 100s of *physical* qubits have already been implemented (Cheng et al. (2023)). However, as discussed above, error correction requires multiple physical qubits per logical qubit, and in fact, there may need to be orders of magnitude more of the former than the latter, as discussed below. In practice this means that the design of a practical quantum computer must be scalable, i.e. able to generate, store, control, and read out arbitrarily large numbers of qubits with high accuracy. Scalability is the main motivation for the work presented in Chapter V.

Gate fidelity. Related to scalability (through the physical overhead required by error correction) is gate fidelity. Given an error rate of 10^{-4} per gate (and making some assumptions about the distributions of these errors among the different types), fault-tolerance can be achieved using a Steane code for error correction, encoding each logical qubit in seven physical qubits (Nielsen and Chuang (2010)). For higher error rates and different error correcting codes, this

ratio of physical qubits needed for one fault-tolerant logical qubit grows rapidly. Depending on the type of error, this number can potentially be much larger or smaller, but in general, 10^{-4} error rates are regarded as the “threshold” for technically feasible fault-tolerant quantum computing, and much of the current work being done focuses on reducing these rates, whether by improving coherence times (the length of time for which the phase between qubit states remains well defined), improving qubit lifetimes, developing coherent control techniques more robust to decoherence and/or calibration errors, etc. Gate fidelity is a motivating factor and important consideration for the work presented in Chapter IV.

Though these requirements may seem to set very stringent standards on the size and performance of quantum computers, there has been much work on how NISQ (noisy intermediate-scale quantum) trapped-ion devices could be optimally designed for scales/performances expected in the near term (Murali, Debroy, Brown, and Martonosi (2020)) and how NISQ devices in general could be used for practical applications (Bharti et al. (2022)).

1.4 TIQC architecture

In trapped-ion quantum computing (TIQC), quantum information is encoded in the Hilbert space available to an ion (or set of ions) in a trap. As discussed above, a single ion has a large set of long-lived electronic energy levels as well as three harmonic oscillator modes and their associated Fock spaces. In principle, any pair of electronic states can be selected as the $|0\rangle$ and $|1\rangle$ states of a qubit (which we will from here on out write as $|\downarrow\rangle$ and $|\uparrow\rangle$ to distinguish them from Fock states in discussions where these are relevant). Typically however, qubits are encoded in Zeeman/hyperfine sublevels of the $S_{1/2}$ state (ground-state qubits, or just “qubits” for the remainder of this section), which has an effectively infinite lifetime.

To explain the standard TIQC architecture and some of its core techniques (techniques for generating entanglement, handling error correction, and reducing unwanted cross-talk between qubits with different functions), we can consider the process of building up a chain of ions in a shared potential well and the operations that need to be performed on them for carrying out a quantum algorithm. This building-up process is represented schematically in Figure 1.3.




Figure 1.3. An illustration of the conceptual reasoning for using mixed-species ion chains. **a)** A pair of ions sharing a potential well can be entangled by driving both ions with laser fields. **b)** For carrying out auxiliary functions (e.g. sympathetic cooling of the ions being used for gates), additional ions need to be added to the chain. If these ions are of the same species as those used for gates, when driven with laser fields, they will emit photons that the other ions can absorb, destroying the information they contain. **c)** If the auxiliary functions are performed using ions of a different species with different optical transition frequencies, then photons scattered by one species will not be absorbed by the other.

Qubit coupling through shared motion. Two ions sharing a harmonic potential well will be driven together by the well and pushed apart by their electrostatic repulsion. The upshot of this is that the ions effectively make up a two-atom solid. (In general, this solid can be called an ion crystal, or an ion chain if the crystal is one-dimensional). This effective solid has a set of shared vibrational

modes. Because both ions couple to the same modes, these modes can serve to mediate entanglement between the ions if there is a way of driving the modes that depends on the state of each qubit. In other words, if there is a way of entangling the qubit state with the state of the motional mode, then there is a way of entangling two qubit states sharing a motional mode. This entanglement is usually carried out with lasers but can also be performed with strong magnetic field gradients (e.g. Löschnauer et al. (2024)).

Ancilla ions. As discussed in Section 1.4 above, fault-tolerant quantum computing requires extra error-correcting qubits. We can imagine adding additional ions to the two-ion chain discussed above to serve this purpose. An issue arises with the fact that, to carry out certain error correction schemes, the ancilla ions need to be read out “mid-circuit” (i.e. before the data qubits are to be measured). This readout process involves driving an optical transition and scattering photons. While this optical drive can be targeted to a single ion (i.e. ions can be individually addressed), some scattered photons will inevitably be absorbed by data qubits if the ancilla and data ions couple to the same optical transition (e.g. if they are both ground-state qubits in the same species of ion). This absorption and spontaneous decay destroys any information in the qubit and induces an error. Similar issues arise when trying to use ancilla ions for other tasks (e.g. sympathetically cooling the chain, as discussed in later chapters).

Multi-species ion chains. To work around the problem of cross-coupling between ancilla and data ions, common practice is to use different ion species (e.g. different alkaline earth metals) for different purposes. These elements have ground state transitions at very different optical wavelengths (e.g. the $S_{1/2} \rightarrow P_{1/2}$ transition being 397 nm in Ca^+ and 422 nm in Sr^+), meaning that photons

scattered during ancilla readout or sympathetic cooling will not be absorbed by the data ions.

1.5 The *omg* architecture

While dual-species ion chains resolve the issue of separating the functions of data and ancilla qubits, they have their own set of drawbacks. Dual-species chains require sources for two types of ion species, two separate laser systems for addressing a different set of optical transitions, and the ability to detect different wavelengths of emitted light. In addition to this increased technical overhead, mixed-species chains see less coupling through common motional modes (due to mass differences between the species), making both mixed-species entangling gates and sympathetic cooling less efficient, and the mixed masses makes the chain more vulnerable to distortion by stray electric fields. To work around these issues and achieve separation of ancilla/data functions with a single ion species, the *omg* qubit scheme was proposed (Allcock et al. (2021)), building off a body of previous work done using optical (Monz et al. (2016); Roos et al. (2004)) and metastable (Sherman et al. (2013)) qubits for multi- and single-qubit operations respectively and proposed in parallel with efforts to implement a similar scheme in ytterbium ions (Yang et al. (2022)). This scheme suggests using the fuller range of electronic states available for encoding for elimination of cross-talk between qubit types.

1.5.1 Optical, metastable, and ground qubits. The *omg* scheme draws its name from the three qubit types it identifies: optical (*o*), metastable (*m*), and ground (*g*) qubits. These qubit encodings are shown in Figure 1.4. As discussed above, *g* qubits, encoded in Zeeman/hyperfine splitting in the electronic ground manifold, are the typical choice for encoding information in trapped ions. In contrast, *o* qubits, encoded in the narrow-linewidth quadrupole transition between ground and metastable states, are sometimes used for information processing, but

typically see more use as tools for state detection in g qubits (since this transition can be used to selectively “shelve” one of the g qubit states where it will not couple to the transition used for readout). Finally, m qubits, encoded in the Zeeman/hyperfine splitting of the metastable manifold, were not widely explored at the time that the work described in this document was done, but can be useful in conjunction with these other qubit types as described below.

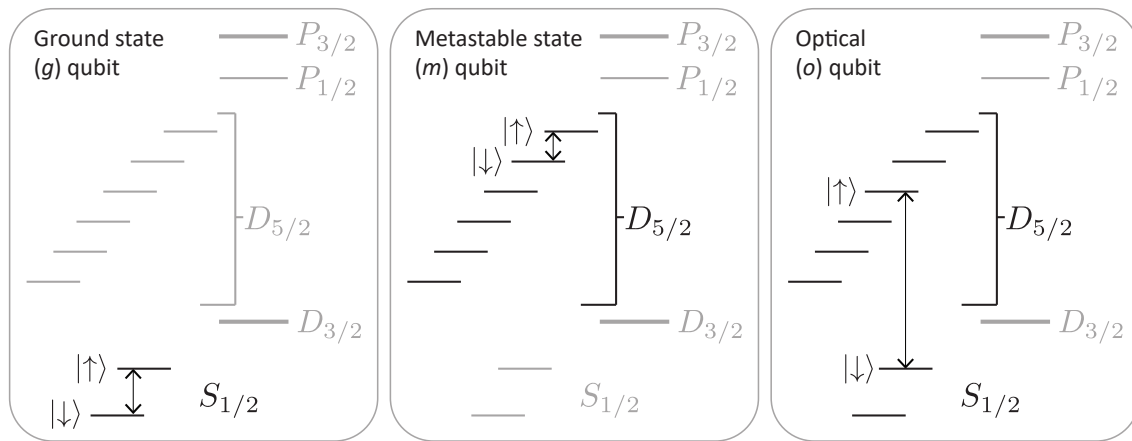


Figure 1.4. The three qubit types considered by the *omg* architecture.

1.5.2 Using *omg* qubits together. To illustrate how the *omg* scheme might be used, an example will be given which specifically motivates the need to investigate the prospect of doing gates in m qubits. One can consider two ions (one in g , the other in m) sharing a potential well. The g qubit could be used for some sort of dissipative operation (e.g. sympathetic cooling of the chain or mid-circuit readout of some ancilla qubit previously encoded in the g ion) while information is stored (or is potentially even being actively manipulated) in the m qubit. In a setup like this, the separation of functions discussed in Section 1.4 is achieved in a single ion species, where g qubits are used for dissipative operations and m qubits are used for encoding and processing information.

1.6 Continuous-variable quantum computing

In typical quantum computing schemes, as discussed above, information is encoded in two-level Hilbert spaces. Often, as in the electronic state of an ion, this Hilbert space is a subspace of some larger but still functionally finite space with varied splitting between levels. In continuous-variable quantum computing (CVQC), information is instead encoded in an infinite Hilbert space of states evenly spaced in energy. Much of the literature on CVQC focuses on its implementation in photons, where photon number states (i.e. the number of excitations of some bosonic mode) are the suggested encoding space (see Fukui and Takeda (2022) for a recent review of the state of photonic CVQC). The harmonic oscillator modes of a trapped ion can be used for much the same purpose, with phonon number state used as a computational basis (Chen, Gan, Zhang, Matuskevich, and Kim (2021); Ortiz-Gutiérrez et al. (2017)), and approaches employing hybrid discrete-variable/continuous-variable approaches using spin and motional modes together have also been proposed (Sutherland and Srinivas (2021)).

1.6.1 Building blocks for CVQC. In the way that a discrete-variable quantum computing scheme requires both single-qubit gates and multi-qubit entangling gates to be computationally complete, CVQC requires a set of Gaussian operations *and at least one* non-Gaussian operation that can be carried out on a bosonic mode (Braunstein and Van Loock (2005); Lloyd and Braunstein (1999)), as well as operations for entangling modes. The theory behind these operations are discussed in greater detail in Chapter III, but a brief overview of these techniques is given below.

Gaussian and non-Gaussian operations. The terms “Gaussian” and “non-Gaussian” refer to the shape of the Wigner function associated with a continuous-variable mode’s state. The Wigner function (described more fully in Section 3.6)

is a quasiprobability distribution in the two quadratures of a mode. In a trapped ion, these quadratures may be the position and momentum associated with a motional mode, potentially in some rotating frame. A Gaussian state is one for which the Wigner function is a two-dimensional Gaussian function, and operations which preserve this feature are said to be Gaussian. Notably, the bosonic ground state (the zero-photon/phonon state) and the familiar coherent state (the quantum analog to classical mechanical oscillators / coherent light fields) are both Gaussian, as is the operation for promoting one to the other. Operations which are first- or second-order in the raising and lowering operators of a mode will be Gaussian (since, in the case of mechanical oscillators, they preserve the parabolic shape of the harmonic potential well, just changing its curvature or position). Higher-order operations (corresponding in light to interactions of two fields in e.g. nonlinear optical media) will distort the state's Gaussian shape and are therefore called non-Gaussian. Non-Gaussian operations can also be implemented by coupling a mode to some non-bosonic system (e.g. through spin-motion entanglement in ions).

Entangling operations. Like two-level systems, continuous-variable modes can be entangled with each other in ways useful for quantum information processing. In photonic systems, two modes can be entangled through interference (e.g. combining two modes on a beamsplitter), or entangled modes can be generated through correlated generation of photons in different spatial modes (as in e.g. spontaneous parametric downconversion). Analogues for both types of operations exist for motional modes in trapped ions, as discussed briefly below and in greater detail in Section 3.6.

1.6.2 Implementing CVQC in trapped ions. A trapped ion chain has an associated set of harmonic oscillator modes. Because ions are charged, their motion will couple to electric fields, and this presents a direct way of controlling

these modes. This means that the operations described above can be achieved in ions in a deterministic, all-electronic fashion. Gaussian operations are routinely performed on ion motion for mode/trap characterization and are sometimes used for sensing applications (Burd et al. (2019)), and electronic operations that generate entanglement between modes by exchanging energy between them have also been demonstrated (Gorman, Schindler, Selvarajan, Daniilidis, and Häffner (2014); Hou et al. (2024); Toyoda, Hiji, Noguchi, and Urabe (2015)). (The ways that we implement these operations in our setup are described in Section 2.7). The final piece required for all-electronic trapped-ion CVQC, the ability to perform non-Gaussian motional operations efficiently, is relatively difficult to achieve, as it requires strong electric potentials that are third-order or higher in ion position, which are generally not desirable in a trap. Overcoming this difficulty and implementing non-Gaussian operations motivates the work presented in Chapter VI.

1.7 Scalable ion trap design

As discussed in Section 1.3, if a quantum computer carried out gates perfectly, even a small one (e.g. 100 qubits) could be used to solve useful problems. However, error rates are likely to remain high (e.g. $> 10^{-4}$ for two-qubit gates) for the near future, meaning that to carry out useful computations, much larger qubit numbers will be needed for performing error correction. Because of this, a practical hardware architecture is one that is scalable, e.g. one that can be expanded from controlling hundreds of physical qubits to hundreds of thousands of them with minimal increase in technical overhead and minimal reduction in quality of control. One key consideration in scalable quantum computer design is all-to-all connectivity, or the ability to perform entangling gates between arbitrary pairs of qubits. In trapped ions, there are two widely-investigated systems for achieving this

all-to-all connectivity. In architectures relying on quantum networking (e.g. the modular universal scalable ion trap quantum-computer or MUSIQ system, Monroe et al. (2014)), ions in entirely separate traps are entangled via single photons. Most proposed architectures however rely on the shuttling of ions through a trap network using electric fields. The discussions in the rest of this document will be focused on this second type of trap architecture, though many of the same design considerations apply for both.

Ideally, a scalable trap architecture is one that can be microfabricated (e.g. with CMOS or MEMS techniques) and integrated easily with components for ion control and state detection (which themselves could be microfabricated, potentially in ways such that integration is inherent to the fabrication process) and that can be driven with off-trap resources (e.g. laser power, trapping power, etc.) for which demand scales in a manageable way with the number of ions trapped. These considerations are discussed more thoroughly in Chapter V, but a brief account is given below of the trap architecture most widely in use and the trap designs most commonly used in research seeking to implement scalable TIQC.

1.7.1 The QCCD architecture and surface electrode traps..

A quantum charge-coupled device (QCCD, Kielpinski, Monroe, and Wineland (2002), first described conceptually in D. Wineland et al. (1998)) is a set of traps, individually controllable segments (generally sharing a trapping RF voltage), making up interconnected “channels” in which ions can be confined and through which they can be transported by adjusting the DC voltage on each segment. Using these DC voltages, pairs of ions can be brought together in short crystals, made to interact, and pulled apart for further operations. A schematic view of this type of trap is shown in Figure 1.5. The scalability of this architecture comes from the fact that a block like the one shown in Figure 1.5 represents a sort of “unit

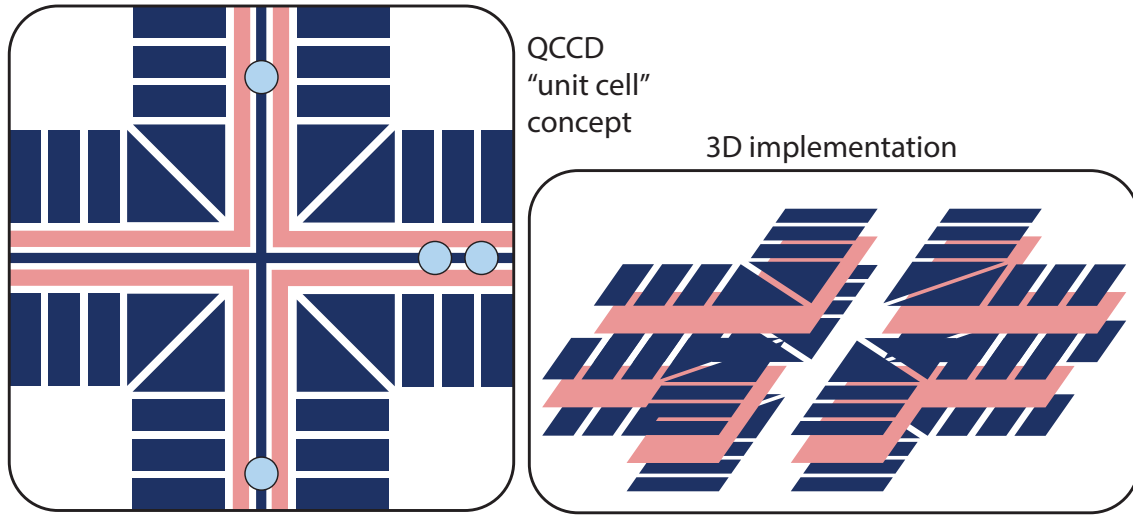


Figure 1.5. A cartoon illustration of a possible QCCD “unit cell,” which can be tiled to produce arbitrarily large trap networks. Dark blue represents DC electrodes, while pink represents trapping RF. Ions are shown in light blue. The illustration on the left shows the unit cell of a 2D/surface electrode implementation of a QCCD architecture, while the one on the right shows how the concept can be extended to three-dimensional traps.

cell” which can be repeated endlessly. Each of these unit cells introduces some fixed overhead in terms of power requirements and number of control lines, but this overhead can be reduced by engineering e.g. new approaches for ion shutting and swapping (Malinowski, Allcock, and Ballance (2023)). Implementing the QCCD architecture, or something like it, motivates much of the work carried out in Chapter V, where we also discuss the typical hardware approach for seeking to implement this architecture, the surface electrode trap (or SET Chiaverini et al. (2005)), which adapts the three-dimensional Paul trap to a set of electrodes in a two-dimensional plane.

CHAPTER II

EXPERIMENTAL APPARATUS

2.1 Ion trap

To trap ions, we need to subject them to electric fields, and to generate these electric fields, we use an arrangement of electrodes to which appropriate voltages can be applied. This arrangement, shown in photographs before and after mounting on a vacuum flange (Figure 2.1), is our trap. The details of this trap’s design, geometric features, electrical drives, and vacuum environment are presented in greater detail below.

2.1.1 Trap design. We run experiments using a macroscopic Paul trap. Our trap electrodes are 0.5 mm dia. rods, with the end cap electrodes (American Probe 72BE-E3/1000x1.25”) having conical tips with $\approx 11^\circ$ internal angle and 200 μm dia. spherical ends. These ten electrodes are arranged in an X-configuration, with four RF trapping electrodes surrounding a pair end cap electrodes, with a set of DC compensation electrodes ringing the outside, with uniform spaces to maximize numerical aperture (NA) for both imaging and laser access. This trap has an 0.75 mm ion-electrode separation, and separation between end cap needle tips is ≈ 3 mm.

Simulation. The design process for this trap was informed by simulations performed using the AC/DC module of COMSOL Multiphysics to generate electrostatic potentials associated with voltages applied to individual sets of rods. These electrostatic potentials/fields were further analyzed in the Electrode simulation suite from NIST (NIST (2017)), allowing calculation of expected trap parameters. From simulations, the trap design used should have a geometric efficiency of 0.90 (defined as the ratio of the mode frequency of a trapped ion at a




Figure 2.1. Photographs of our trap, showing (left) the trap attached to a vacuum flange and connected to electrical feedthroughs and (middle) a close-up shot of the unmounted trap, with the rough location of the trapping zone marked with a blue box. The associated inset is a picture taken with a CMOS camera of a five-ion chain, with ion spacing marked. Finally, a COMSOL render (right) shows trapping RF electrodes marked in red.

given set of parameters to the mode frequency at the same parameters given ideal, i.e. hyperbolic, electrodes).

Fabrication. The electrode rods and needles are made of beryllium-copper (98% Cu 2% Be with a hard temper) electroplated (to prevent oxidation of the electrode surface) with a ~ 100 nm thick layer of gold plated on top of a palladium diffusion barrier (which is non-magnetic). To ensure that the electroplated gold was sufficiently thick, we used the scheme shown in Figure 2.2, measuring the number of gold atoms deposited onto the rod by measuring the total charge passed through the electroplating circuit. To electrically insulate the trap, the rods are held in a pair of end plates machined from Shapal Hi-M Soft, a machinable ceramic.

2.1.2 Trapping RF drive. To generate an effective potential (see Section 3.1) to confine the ions along the trap axis we need an RF voltage that can be applied to the two RF trap electrodes. This voltage is supplied by four-stage RF chain shown schematically in Figure 2.3 which can be summarized as follows: 1) an AWG supplies a sinusoidal voltage with high amplitude noise to 2) a Squareatron, which generates a stable-amplitude signal by turning the AWG input

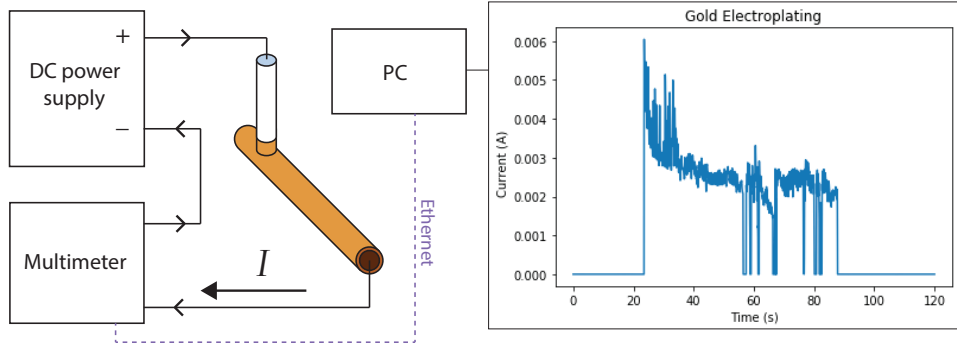


Figure 2.2. Schematic illustration of the hardware setup used for estimating the thickness of a layer of electroplated metal by measuring the total charge passed through the system during electroplating.

into an effectively digital signal and stripping off all higher harmonics, with this signal in turn being fed into 3) an amplifier (Minicircuits ZX60-100VH+) which drives 4) our resonator cavity. These components are discussed in greater detail below, but for a full accounting, see J. M. Metzner (2024).

AWG. A Rigol DG1022Z arbitrary waveform generator (AWG) generates an RF sinusoid voltage. In the context of the rest of the chain, this signal acts as a “frequency reference,” and can be controlled over ethernet to change the trap RF frequency when needed (e.g. for automatically reducing radial confinement for ion chain recrystallization). The output of this AWG has a large amount of amplitude noise, so we do not use it to drive our resonator directly, since fluctuations in the trap RF voltage translate to fluctuations in the mode frequencies of the ions. Instead, we take a 0 dBm AWG output and feed it into a second stage of RF conditioning.

Squareatron. The Squareatron (the design for which originates at the ion trapping group at Oxford with modifications developed in our own lab) is a device that takes in a frequency reference with large amounts of amplitude noise and outputs a sinusoid with a stable amplitude. The basic operating principle is that

the input signal saturates an amplifier, generating an effectively digital signal, from which the higher harmonics are stripped with filters. More complete accounts of how this device works are offered in the theses of Tom Harty (Harty (2013)) and Jeremy Metzner (J. M. Metzner (2024)). The Squareatron allows ≈ 1 dB of digital tuning of its output amplitude, which is fed into an amplifier before being sent to our resonator. The Squareatron and amplifier are both attached to a temperature-stabilized aluminum plate in a box intended to reduce airflow.

Helical resonator. To get a factor of ≈ 80 step up in the RF voltage going to the trap, we use a helical resonator with a resonant frequency of ≈ 14.5 MHz with Q-factor of up to 150. Coils are mechanically stabilized with low-loss polyethylene foam supports to prevent pickup of acoustic-frequency vibrational noise. The step-up coil includes a connection for applying other signals to the trap electrodes, discussed in greater detail in Section 2.7. The output of the resonator is connected to the trap electrodes through a single wire in a vacuum feedthrough to which the RF trap electrodes are connected.

Under typical operating conditions, we see radial center-of-mass (COM) mode frequencies of ~ 1.8 MHz, corresponding to RF voltages at the trap of ≈ 0.8 kV.

2.1.3 Trapping DC potential. To generate confinement along the trap axis, an equal DC voltage is supplied to the two endcap needles using a DC source (Stanford Research Systems DC205), which can output up to ± 100 V. The output of this voltage source is connected to the trap through a pair of wires on a vacuum feedthrough, with these connections being grounded at RF through 0.68 nF capacitors, as shown in Figure 2.4. Under typical operating conditions (e.g. ≈ 60 V DC), we see ≈ 0.55 MHz axial mode frequencies.

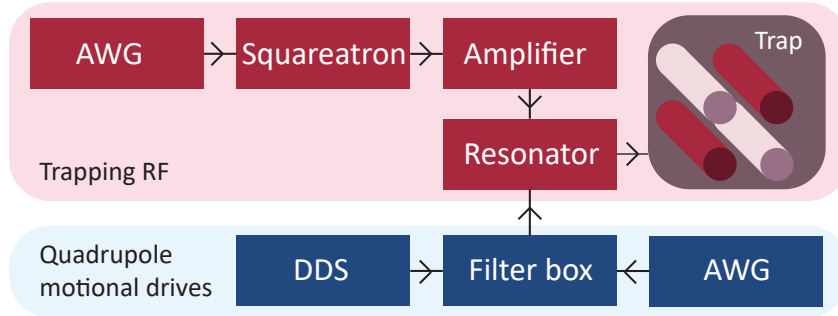


Figure 2.3. Block diagram of the RF chain used to supply the RF trapping voltage. The trapping RF itself is supplied by the chain in red. Quadrupole motional drives (discussed further in Section 2.7.3) are combined with the trapping RF at the resonator.

2.1.4 Shim potentials. Unwanted electric fields in the trap (typically called “stray fields” and coming from e.g. charging of nearby dielectric materials by stray UV laser light) will tend to move the ion off of the trap center (the “RF null” discussed further in Section 3.2.5). To counteract these fields, we use our compensation electrodes to provide “shim potentials,” small DC potentials that shift the trap equilibrium point. We drive these electrodes using digital-to-analog converters (DACs) incorporated into our ARTIQ control system, described in Section 2.8 below. These DACs can supply up to ± 10 V DC, and each compensation electrode can generate an electric field of magnitude 19.9 (V/m)/V at the ion. Connections are through four wires on a vacuum feedthrough, each connection being grounded at RF through a 2.2 nF capacitor, as shown in Figure 2.4.

2.1.5 UHV system. The trap is kept in an ultrahigh vacuum (UHV) chamber held at 10^{-11} Torr. Low pressure is required in order to minimize rates of collision between trapped ions and background gasses (mainly hydrogen) that can cause decrystallization, dephasing, etc. The system is designed to be held under UHV continuously and never brought up to atmosphere, and the hardware used to

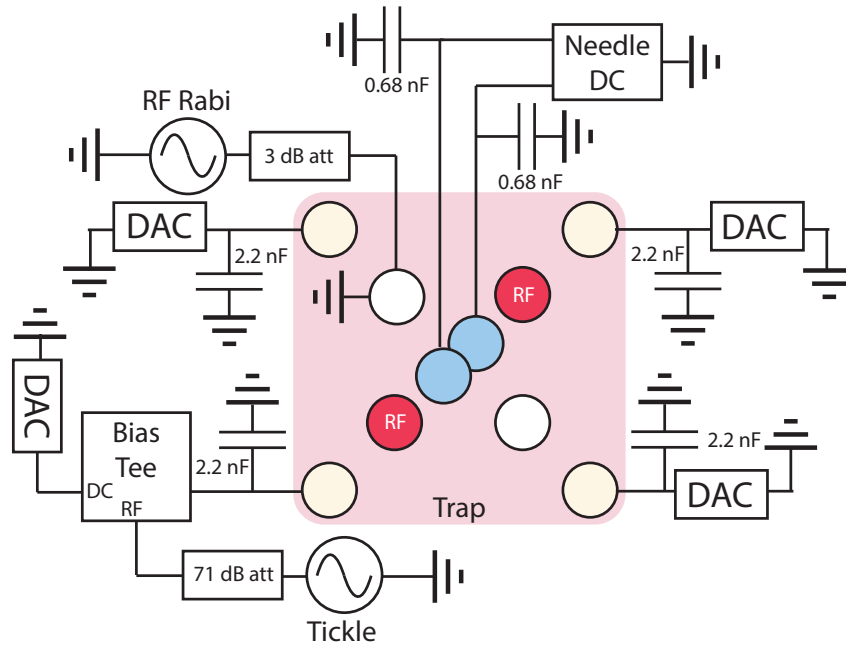


Figure 2.4. Circuit diagram for the electrical connections to the DC trap electrodes (including AC drives applied to these electrodes), with endcap needles in blue, DC trapping electrodes in white, and compensation electrodes in light yellow.

maintain vacuum is not connected to air. A render of this system, with components labelled, is presented in Figure 2.5.

The chamber. The ion trap sits in an octagonal chamber with six viewports on 2-3/4" flanges for laser entry, one viewport (used for imaging/state detection) on a 6" flange with electrical feedthroughs for the trap, and one custom made high-NA viewport for future ion-photon entanglement experiments. The top of the chamber is connected to a custom-fabricated spherical manifold attaching to hardware (described below) for maintaining/measuring UHV as well as a valve for pumping the system down with mechanical pumps. The bottom of chamber is covered with a 2-3/4" blank flange. The chamber is mounted in stainless steel bracketing designed in-house. Chamber components are held together with silver-

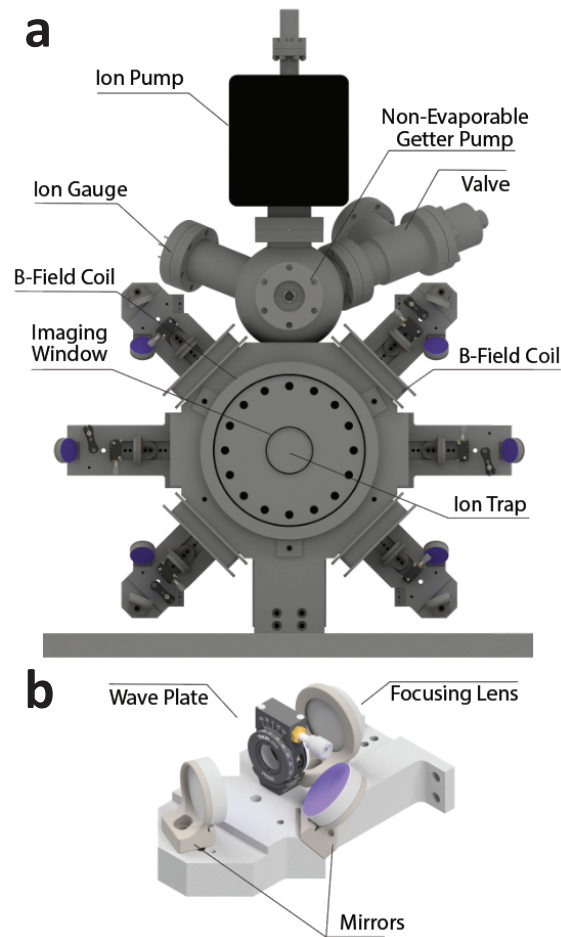


Figure 2.5. **a)** CAD rendering of the vacuum chamber, associated support hardware, and components directly attached to this hardware, including magnetic field coils (Section 2.2.1 and **b)** mounts for beam delivery optics (Section 2.6).

coated screws to prevent parts from seizing during bakeout. Conflat (CF) gaskets used at the interfaces of flanges are silver-coated for the same reason.

Maintaining and measuring vacuum. Vacuum is maintained using a passive non-evaporable getter pump (a NEG, SAES CapaciTorr Z 100) and an active ion pump (Agilent StarCell, 20). These parts hold the system at UHV without being connected to any mechanical pumps. Low pressures can be measured using an ion gauge (Agilent UHV-24P), which measures the current generated by ionized gasses.

Preparing chamber hardware. Before assembly and pumpdown of the vacuum system, most of the stainless steel components were baked at 450°C under atmosphere for several days. This baking process causes the formation of a bronze-colored oxide layer that reduces the amount of hydrogen outgassing in the interior of the system (Sefa, Fedchak, and Scherschligt (2017)). Hardware was washed in isopropanol and acetone before and after pre-assembly baking.

2.2 Magnetic field control

Qubits are encoded in the magnetic moments of our trapped ions. In $^{40}\text{Ca}^+$, which has no nuclear spin, the quantization axis and energies associated with these moments are set by an external magnetic field, which induces Zeeman splitting and breaks the degeneracy of spin states in each energy level. In our system, this magnetic field is at a 45° angle to the trap axis and is a combination of the magnetic field generated by a set of wire coils and the local magnetic field of the earth. All in all, the magnitude of our magnetic field is ≈ 1.56 G, producing a Zeeman splitting (between spin states, e.g. $m=+3/2$ and $m=+5/2$) of 2.63 MHz in the $D_{5/2}$ metastable level of $^{40}\text{Ca}^+$. Noise in this magnetic field is the limiting factor in the coherence of spins encoded in these Zeeman sublevels (since their splitting is first-order sensitive to magnetic field strength). Techniques for characterizing the strength and (to some extent) spectrum of this noise will be discussed in Section 3.9, but an important noise type to discuss here is line noise (noise coming from the 60 Hz mains), which we cancel using a feed-forward system. This system is described in greater detail below in Section 2.2.2.

2.2.1 Magnetic field coils. To control the magnitude and direction of the magnetic field at the ion, we use a set of five wire coils, wound with 20 gauge copper wire. Opposite pairs (spaced about 8.6” apart) of 3.6”x3.0” rectangular coils with ~ 300 windings per coil provide magnetic field control in the xz -plane (in the

trap coordinate system defined above), while a single circular 6.1” dia. coil (about 1.4” from the ion) is used for magnetic field cancellation in the y -direction. These coils are driven with a set of three DC currents (with opposing coil pairs wired in series) supplied by a regulated DC power supply (QL355TP), which can provide up to 500 mA in its low-noise mode. This maximum current is used to drive the coils along the quantization axis. An image of the coil setup is shown in Figure 2.6.




Figure 2.6. **a)** A schematic of the magnetic field coil arrangement used to set the quantization axis in the trap. Current is marked with white arrows, with arrow thickness corresponding the current magnitude. **b)** A photograph showing the five magnetic field coils used on the trap. Worth noting are the two smaller coils (white and blue) wrapped around the top left coil, which are used for line noise feedforward (Section 2.2.2).

2.2.2 Line noise mitigation. Line noise refers to the magnetic fields produced by the 60 Hz (in America) alternating current used to supply electric power to the lab. The noise itself consists of a stable, narrow-band 60 Hz component and its harmonics (with the leading harmonic ordinarily being the first odd harmonic, 180 Hz, produced by e.g. clipping of a 60 Hz sinusoid). Essentially, line noise causes the frequency of our Zeeman qubits to oscillate on a 60 Hz cycle. If we are running an experiment that is not itself in-phase with

this line noise (i.e. shots of the experiment are being performed at randomly sampled points over a cycle), then the line noise will look like a broadening of our qubit frequency / a reduction in spin coherence time. There are two ways of addressing this issue, discussed below: 1) running experiments in sync with the 60 Hz line, and 2) cancelling out 60 Hz noise using feed-forward. These techniques are represented schematically in Figure 2.7a and b respectively. (As an important note, both techniques involve the use of a Line Trigger (Sinara (2020)), a component developed as part of the Sinara open source hardware project. This device generates a 5 V digital output in-phase with an input from a 60 Hz power supply connected to mains).

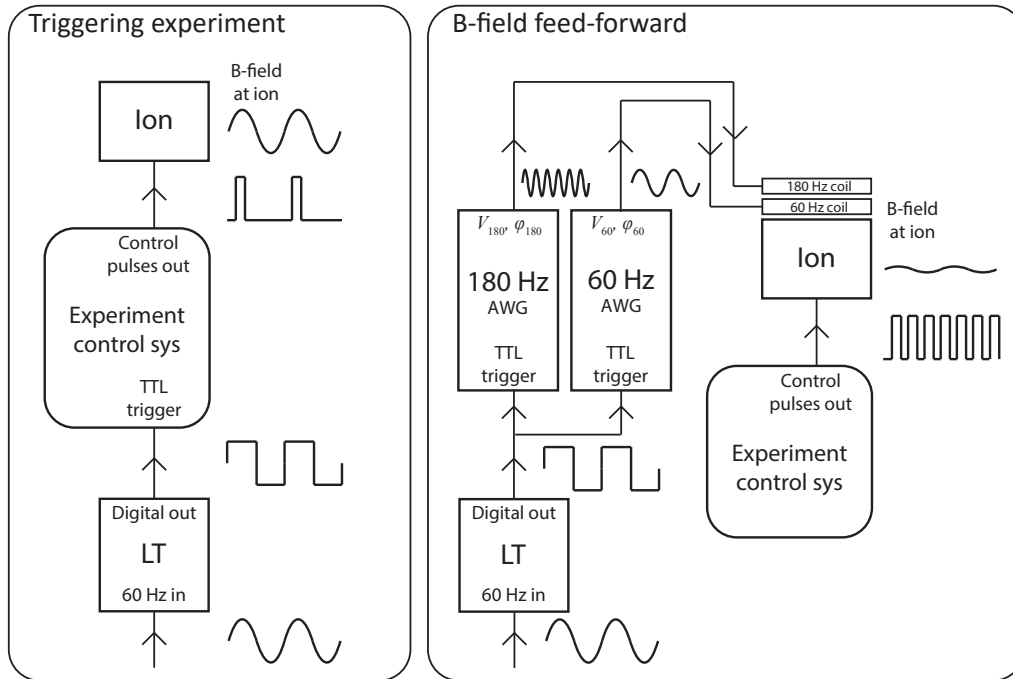


Figure 2.7. Schematic illustrations of our two line noise mitigation schemes, explained in greater detail in text.

Line triggering. A simple solution to the problem of line noise is to run experiments in sync with the 60 Hz line. This can be done by taking a Line Trigger output, reading it in with a TTL channel on the ARTIQ crate (described in more

detail below), and using that input to time experiments to run at consistent points on the line cycle, ideally selecting a time where the line noise is at a crest/trough, so that the qubit frequency will be first-order insensitive to experiment run time, for experiments $\ll 16.7$ ms. This technique introduces little technical overhead, but it can make experiments substantially slower, since it forces the time between experimental shots to be 16.7 ms (or multiples of this). In addition, for experiments lasting an appreciable portion of the line cycle, the qubit frequency will still be shifted over the duration of the experiment.

Line noise feed-forward. A technique that requires more technical overhead, but which we generally use in our experiments, is to cancel out line noise with a feed-forward loop. In this setup, we drive small coils (~ 5 windings each) with 60 Hz and 180 Hz voltages generated by a pair of AWG channels triggered using the Line Trigger. (As shown in Figure 2.6b, we only wind feed-forward coils around one of the quantization axis coils. In our case, our line noise magnetic field is small relative to our quantization field, so our qubit frequency will only be first-order sensitive to line noise along the quantization axis). We calibrate the amplitudes and phases of the 60 and 180 Hz voltages by selecting values which maximize the contrast of our spin Ramsey fringes at long delay times (the details of this technique are explained further in Section 3.9.1). While this technique requires substantially more calibration than line triggering the experiment, it allows experiments to be run much more quickly and ensures there will be no systematic effects from coherent changes in qubit frequency over the course of an experiment.

2.3 Imaging and state detection

If we have the ability to generate and trap ions, before we can use them for experiments, we need to be able to 1) confirm when ions are in the trap and how many of them we have and 2) detect the states of these ions during/after

experiments. For both of these purposes, we use fluorescence on the 397 nm $S_{1/2} \leftrightarrow P_{1/2}$ “cycling transition” in $^{40}\text{Ca}^+$. Illuminating an ion with 397 nm laser light, it can scatter photons at a rate proportional to $1/\tau_{P_{1/2}}$, where $\tau_{P_{1/2}}$ is the lifetime of the $P_{1/2}$ excited state. This means that an ion can scatter $\sim 10^8$ photons per second. These photons can be detected and used to infer that ions are in the trap and that they are in the electronic ground state (since if they were e.g. shelved in a metastable state, the cycling transition could not be driven). This being the case, we built a high-NA (numerical aperture) microscope for detecting 397 nm photons, both allowing ions in a chain to be imaged on a camera in a spatially resolved way and for photons to be counted using a photomultiplier tube (PMT). The hardware details and performance specifications of this imaging system are discussed below.

2.3.1 Imaging system overview. This discussion of the imaging system will start with a hardware overview and an accounting of the design process. A schematic of the imaging system, along with a photograph of the system itself in our experimental setup, is presented in Figure 2.8. The design of the objective was based on one used at NIST, with parameters optimized for 397 nm light using Zemax OpticStudio. OpticStudio was used for optimizing the placement and selection of optical elements and for predicting performance factors like magnification and spot size.

Objective. The objective consists of, in order of nearest- to farthest-from-ion components, a 2” dia. aspheric convex lens (Asphericon AFL50-60-S-U), a 2” dia. aspheric convex lens (Asphericon AFL50-60-S-U), and a 1” dia. planoconvex lens (Edmund 48-321). At the end of the objective, an iris (not shown in schematic) allows the NA to be restricted. By itself, this objective should provide -3.8x magnification and a 2-3 micron spot size.

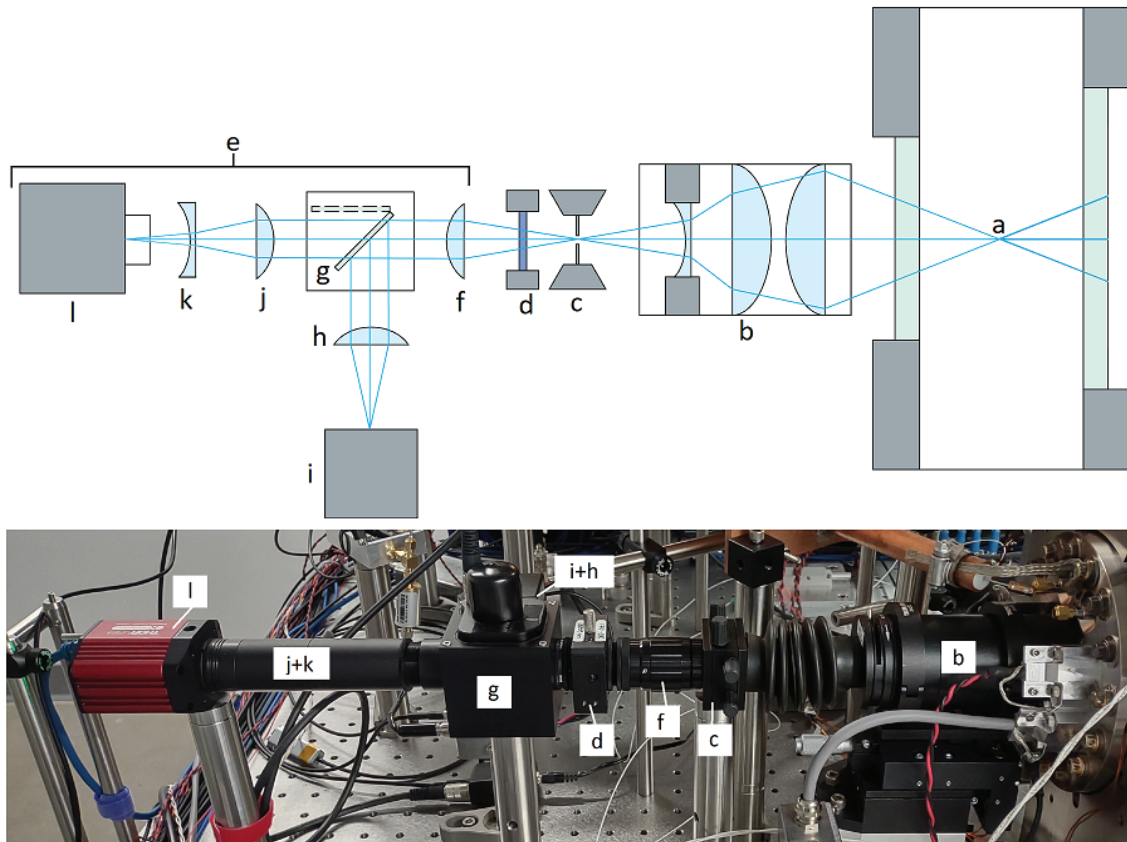


Figure 2.8. Schematic illustration (and photograph below) of our imaging system, with light emitted by the ion (**a**) being captured by an objective (**b**), passed through a slit (**c**) to reduce stray light and a narrow-band filter (**d**) to prevent the system from registering light far away from the wavelength of the cycling transition of Ca^+ , collimated with a lens (**f**) in an adjustable barrel mount to be redirected by a flip-mounted mirror (**g**), which can direct light to a focusing lens leading to the PMT (**h**) and (**i**) or to a focusing lens (**j**) and (**k**) heading to the camera (**l**).

Stray/unwanted light reduction. After the objective, a slit (Owis SP-40) is used to reduce stray light (i.e. light entering the system from sources other than the ions at the objective focus). After this, a narrow-band optical filter (Semrock FF01-390/18-25) is used to screen out light at wavelengths other than 397 nm.

Collimation and reimaging. Focused light is collimated using a convex lens (Thorlabs AC127-030-A) and sent along a path with a mirror (Edmund Optics 85-311) mounted in a motorized flippable mount (Owis KSHM 40) which can

direct light down two paths: To a PMT (Hamamatsu H10682-210) with just a 1" dia. lens (Thorlabs LA1540-A) used to focus the collimated light onto the PMT's active area, or to a CMOS camera (Kiralux CS505MU) with a combination convex/concave lens pair (Thorlabs AC127-075-A and ACN127-050-A respectively) to get reimaging with a longer effective focal length than would otherwise be possible with the available space.

2.3.2 Imaging system performance. Below, a brief account is given for how the imaging system was tested (before and after setting it up in the experimental apparatus) and a summary of imaging system performance.

Testing imaging. The system's ability to focus an image and resolve features was first tested with a (negative) 1951 USAF resolution test chart backlit with 400 μm light. On the apparatus, we established that we could visualize the rod electrodes by adjusting the position of our objective and used our ability to bring the DC needle tips into focus to locate the rough imaging plane in which the ions would likely sit, as well as where in that plane they would be (i.e. exactly midway between the needle tips). Finally, we were able to tune up objective position on actual ions (which we could detect through 397 nm photon counts without being able to image, since the PMT is less sensitive to depth-of-field considerations) and ion chains in the trap, allowing determination of ion number and ion spacing (e.g. on the image in Figure 2.1).

Testing click efficiency. We tested the click efficiency of our imaging system (the total fraction of emitted photons that are recorded as PMT clicks) by shelving an ion in the $D_{3/2}$ level by driving the 397 nm $S_{1/2} \leftrightarrow P_{1/2}$ transition and looking for PMT clicks while deshelving by driving the $D_{3/2} \leftrightarrow P_{1/2}$ transition. This deshelving process generates a single 397 nm photon, so a click efficiency can be estimated by comparing the number of clicks generated by the PMT to the number

of photons that should have been emitted (i.e. the number of shelving/deshelving cycles) wound up being 1.18%, consistent with an $\approx 4\%$ collection efficiency and a ~ 0.3 PMT quantum efficiency.

Performance. The system's performance can be summarized in a few numbers as follows: the whole setup has an NA of 0.41 (based on a 50.8 mm objective lens dia. and a lens-to-ion separation of 57 mm), meaning that it should collect at most $\approx 4\%$ of scattered 397 nm photons. The camera path has a total magnification of -19x and a measured(predicted) spot size of 35(15) μm (along the path to the camera). The reason that the spot sizes are larger than expected has not been established, but together with our magnification, it means that we can reliably resolve ions down to $\approx 2 \mu\text{m}$ of ion-ion separation. As mentioned above, the overall click efficiency of the PMT path is 1.18%.

2.4 Laser system

While there are all-electrical techniques for cooling, controlling the spins of, and reading out charged particles in a trap (Yu et al. (2022)), the use of laser fields is typical for all of these applications. In our lab, we drive ions using a set of CW diode lasers tuned to near-resonance with the optical transitions of $^{40}\text{Ca}^+$.

Conceptually speaking, our laser system can be broken up into three parts, separated by function: 1) laser breadboards (which generate and divide optical power), 2) the Wavelength Analysis and Display laser diagnostics suite (WAnD, Oxford (2019)) (which measures laser frequencies and enables frequency locking), and 3) AOM boards (which act as both fast optical switches and a means of modulating laser frequencies, phases, and amplitudes at the ion). A schematic representation of this system for a single laser is shown in Figure 2.9. It should be noted that, for the 854 nm and 976 nm Raman lasers discussed separately from the rest, the WAnD is used only for diagnostic purposes and never for laser locking.

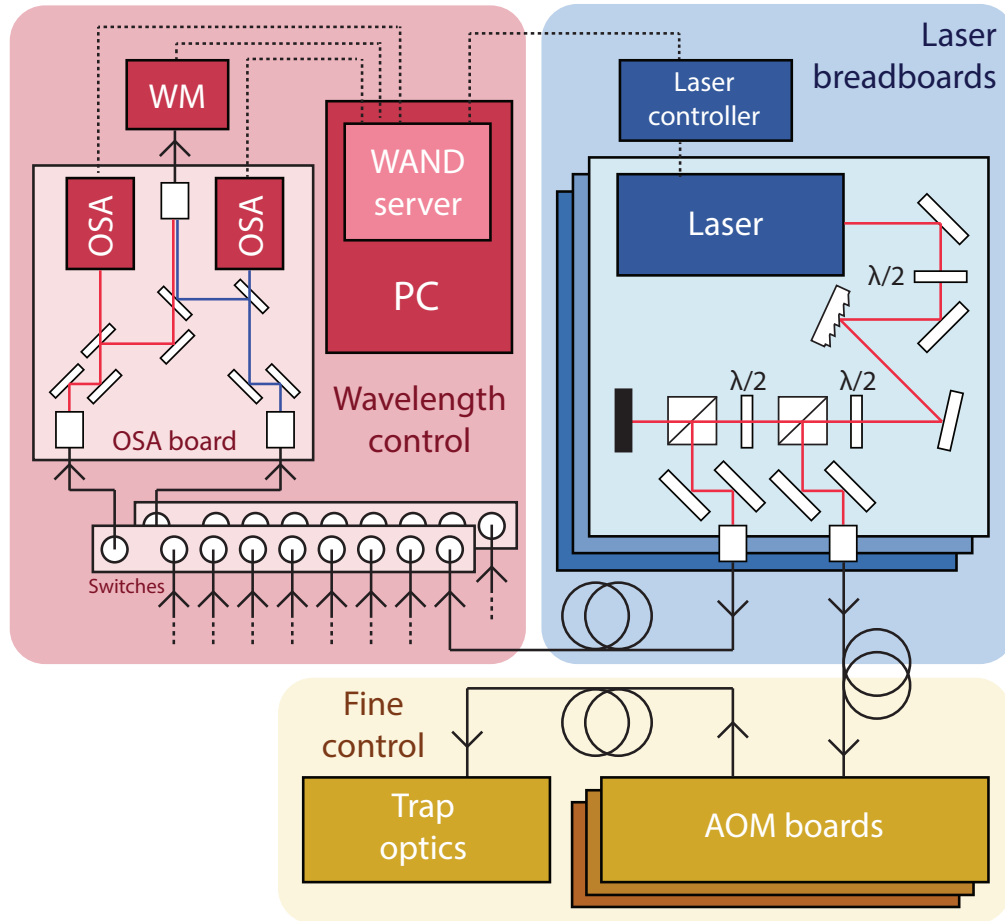


Figure 2.9. Schematic of our overall laser system (excluding far-detuned lasers, discussed in Section 2.5) shown in terms of the output of a single laser breadboard. Free-space laser beams are shown with red and blue lines. Fiber optics are shown with solid black lines. Data connections are shown with dashed black lines.

2.4.1 Rack-mounted laser boards. For most of the laser wavelengths we need, we have a set of tunable external cavity diode lasers (ECDLs) mounted on breadboards placed in drawers on a standard 19" rack. Photographs of this rack and a representative breadboard are shown in Figure 2.10. The function of these breadboards is 1) to clean up laser outputs with blazed gratings, which spatially separate the broad-spectrum power in a spontaneous amplified emission (ASE) “pedestal” from the desired narrow-band output mode of the laser and 2) to

divide laser power along different paths, where it can be sent into couplers and routed to wherever this power is needed. Power is divided with a series of $\lambda/2$ waveplates (WPs) and polarizing beamsplitters (PBSs). The amount of power passing along a specific path can be adjusted by changing the orientation of the waveplate before this path splits. A schematic for this setup is shown in the top right block of Figure 2.9.



Figure 2.10. Photograph of our laser rack, showing the drawers containing laser breadboards as well as an additional rack for housing laser controllers and a third for the wavemeter and OSA board. Inset shows a sample laser breadboard mounted in a drawer.

2.4.2 Laser frequency locking. A portion of the laser power on each breadboard is sent to the WAnD for analysis and, potentially, feedback on the laser frequency. The four key hardware components of the WAnD are 1) optical switches for selecting the laser to be characterized (since the system can only analyze one laser at a time), 2) optical spectrum analyzers (OSAs) for ensuring that the laser is not lasing on multiple modes (as well as sometimes being used for characterizing rates of laser frequency drift/fluctuation on a faster time scale than the wavemeter would be capable of), 3) a wavemeter (WM) that provides high-precision absolute measurement of a laser’s frequency, and 4) a WAnD server running on a host PC (“PC” in Figure 2.9) capable of communicating with the laser controller and adjusting both the current to the laser diode and the voltage on a piezo stack that fine-tunes the characteristics of the external cavity, changing the laser frequency. In practice, we have two optical switches and two OSAs, one each for the UV-blue (just called “blue”) and NIR (“red”) frequency ranges, and a single WM. These components and their operation are described in greater detail below.

Optical switches. We use optical switches to select one of eight inputs (per switch) to direct to an output port. These outputs in turn are coupled directly onto the OSA board. These optical switches are directly controlled by the WAnD.

OSA board. The OSA board (schematic shown in the left panel of Figure 2.9) is a breadboard sitting in a mount on a 19” rack. This board takes the outputs from the red and blue optical switches and splits each of them, with one path going to the OSA (the back mirror piezo stack for which is driven by a linear ramp voltage supplied by an AWG) and the other path going to the WM. The red and blue beams are directed to separate OSAs (Thorlabs SA200-8B and SA200-3B respectively), but are combined before being directed to a photonic crystal fiber going to the WM.

Wavemeter. The wavemeter (WM, HighFinesse WS8) measures laser frequency at a resolution of 0.2 MHz and returns this value to the WAnD server running on the laser controller PC. The wavemeter is held in a temperature-stabilized box to prevent fluctuations in temperature-dependent systematics.

WAnD server and frequency locking. A WAnD server (Oxford (2019)) monitors and reports laser wavelengths as measured by the WM and transmission lines measured by the OSAs. In addition, the WAnD has an option to lock lasers to a given frequency set-point (i.e. a given detuning from some absolute reference). If a lock is active, whenever the WAnD collects a frequency measurement from the WM, it will calculate an error signal (based on a PI loop) and, from this error signal, calculate some change to apply to the voltage on the external cavity piezo stack, which it passes along to the laser controller. In our setup, this loop enables a slow (~ 0.1 Hz) lock, good to within ~ 2 MHz, for six or more lasers at once. The quality of this lock (and absolute frequency of each laser) can be measured using the optical transitions of the ions themselves.

Testing WM/lock performance with ions. The optical transitions of our $^{40}\text{Ca}^+$ for our purposes serve as an absolute frequency reference, and by looking at drifts in line positions relative to the reference frequencies being used by the WAnD, we can both identify calibration errors in our WM and track how these errors drift over time. Over several years, we have seen 10s of MHz drifts in WM measurements at UV frequencies and comparatively little drift in the measurement of IR frequencies. In addition to looking for long-term drift in our WM calibrations, we can also assess lock performance using ions by, e.g. sitting on the side of a spectral line and looking for changes in scattering rates, or by direct, rapid measurement of the line with e.g. a two-point calibration (discussed

more completely in Section 3.10.4). Section 3.10.3 also outlines the experiments we perform to measure laser lines.

2.4.3 AOM boards. The laser light coming from laser breadboards is CW, and the frequency of it can be adjusted by changing the cavity piezo voltage on the associated laser with a corresponding change to the laser diode current. (In practice, laser diode current alone could also be changed, but this approach tends to make the laser more likely to jump to a different mode or “mode hop”). However, to run experiments, we must be able to modulate the frequency/phase of our beams on time scales useful for a single experimental shot (i.e. $\ll 1$ ms), as well as pulsing our laser beams at the trap and regulating beam powers. To achieve this, we use acousto-optic modulators (AOMs) mounted on custom-fabricated breadboards stored in drawers on a 19” rack. A brief discussion follows of the operating principle behind AOMs and how they can be used for switching / fast frequency control of lasers, along with a description of our AOM boards and typical AOM configurations used in our experiments.

AOMs. An AOM is a device that, when driven with an RF voltage at frequency f , splits the beam passing through it into diffraction orders, with the n th orders frequency shifted by nf . Inside of an AOM, the beam passes through a crystal cemented to a piezoelectric actuator. When this actuator is driven, this sets up a travelling wave in the crystal that is essentially a walking diffraction grating. An AOM can be oriented such that interference causes most of the power to go into a specific diffraction order, with 70-80% efficiencies typical for a single pass through.

AOM boards and typical configurations. We operate our individual AOMs as either single-pass (SP) or double-pass (DP), and we use combinations of these in three configurations, shown schematically in Figure 2.11. Most of our beams

are operated with a DP+SP configuration, where we first pass a beam (brought in from one of our laser breadboards by fiber optics) through a DP and then split this beam power along different SP paths to supply different beams at the trap. (For example, the 397 nm laser is split after a DP and provides power for a near-detuned π -polarized beam, a far red-detuned π -polarized beam, and a near-detuned σ^+ -polarized beam, each of which have their own SP AOMs). In the DP+SP configuration, frequency/phase tuning is carried out on the DP (since the beam deflection caused by changing RF drive frequency will matter less for a DP) while power servoing is carried out using the SP. The lone SP configuration has the advantage of lower power loss, but compared to lone DP configuration has 1) an inherent factor of two smaller beam tuning bandwidth (since the beam is only being diffracted once), 2) sensitivity of beam pointing to RF frequency further reducing effective bandwidth, and 3) higher optical power leakage (with only ≈ 30 dB attenuation) when the RF drive is off, which can be problematic if leakage is occurring on a beam used for optical pumping. As a result, the configuration used for a particular beam depends on power and tuning requirements, as well as sensitivity to leakage. A fourth configuration, used only in one experiment (touched on in Chapter IV), is a “noise eater.” In this configuration, the 0th-order beam is sent to the trap, and power is directed to the 1st-order beam only for the sake of power regulation. This configuration allows the full available beam power to be sent to the ion, but suffers the disadvantage that the beam can never be turned off.

RF drive and beam power servoing. RF voltages for driving AOMs are supplied by direct digital synthesizers (DDSes) in our experiment control system. These voltages are passed through 1 W amplifiers before being sent to our AOMs, and beam powers are controlled by adjusting the DDS output powers going into these amplifiers. This setup is discussed in greater detail in 2.6.3 and 2.8.1 below.

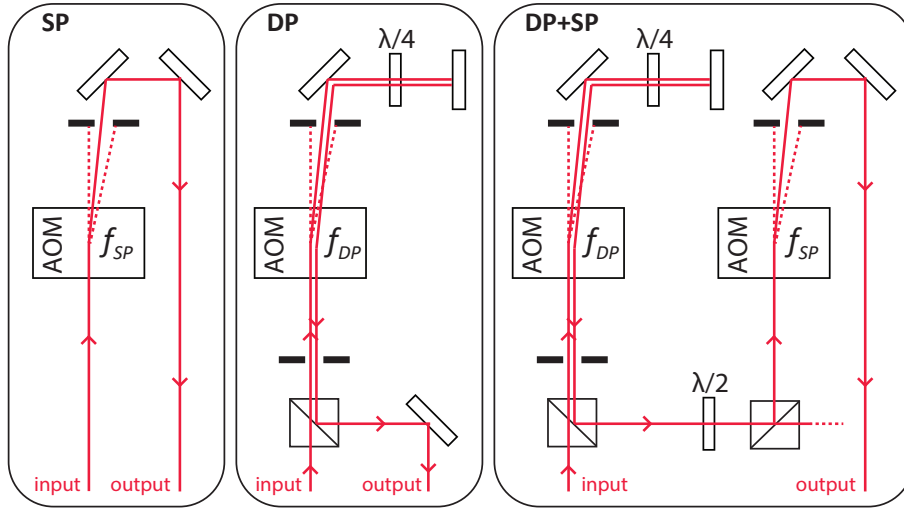


Figure 2.11. Schematic illustrations of the three main AOM configurations we use in the lab. These schematics omit the focusing/recollimating lenses that focus the beams into the AOMs.

2.5 Far-detuned lasers

The laser system described above accounts for all of our near-resonant lasers. When operating near resonance on lines with widths $O(10 \text{ MHz})$, the tuning of a laser to within MHz is significant, and a system like the WANd, in conjunction with a tuning knob like piezo voltage fast enough for the lock (i.e. one that can be run at $> 0.1 \text{ Hz}$), is necessary to prevent problematic laser drifts. For our lasers operating at detunings orders of magnitude larger than the transition linewidths, these requirements are far less stringent. Conversely, since these lasers are operating so far off-resonance, large beam intensities are needed to achieve substantial coupling to our ions. To achieve these high intensities, we use laser systems that enable us to get very large powers at the ions (i.e. ≈ 10 to $\approx 200 \text{ mW}$ per beam, rather than the $\sim 1\text{-}100 \mu\text{W}$ range at which we operate our near-resonant beams). We operate two wavelengths of far-detuned laser: 976 nm (tuned -43 THz from the $D_{5/2} \leftrightarrow P_{3/2}$ transition, used for carrying out coherent operations on our qubits, see e.g. 3.5.2) and 854 nm (tuned -21 GHz from the same transition, used for

providing an AC Stark shift to states outside of our qubit manifold, see 3.3.3). The hardware setups for these laser systems are outlined below.

2.5.1 976 nm Raman lasers. We have used two 976 nm laser setups in different experiments: A one-laser free-space setup and a three-laser injection-locked setup. In both cases, the setup is driving two separate beams at the ion.

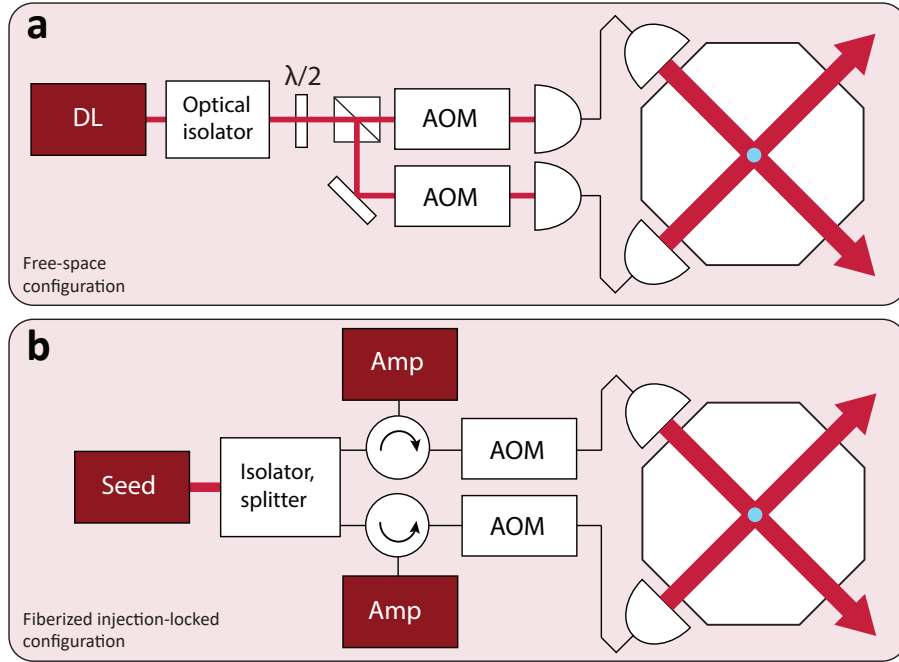


Figure 2.12. Schematics of the two 976 nm Raman laser setups we have used in our experiments, showing the diode laser sources (marked in dark red) all the way through to the ion (marked in light blue). These schematics omit details of the beam delivery optics.

Free-space configuration. One of our early 976 nm laser setups was a single-laser free-space setup, consisting of just an external cavity diode laser (IPS I0976SB0750B), an optical isolator (Newport ISO-04-980-MP), a WP/PBS pair for splitting power along two paths, and a pair of SP AOMs for independently tuning two beam paths. This setup gave us about 70 mW per beam at the ion. For taking data on scattering rates out of the $D_{5/2}$ level in $^{40}\text{Ca}^+$ (see Chapter IV), we temporarily switched this setup to the “noise eater” configuration (with a

mechanical shutter after the isolator to allow the beam at the ion to be turned off on slow time scales) and would put all power along one path at a time, which gave us upwards of 200 mW in a single beam. A schematic of this setup is shown in Figure 2.12a.

Fiberized injection-locked configuration. To increase our available 976 nm beam power, we switched over to an injection-locked configuration (Shimasaki, Edwards, and DeMille (2019)). In this configuration, we split power from a free-space 976 nm diode laser (the “seed laser,” IPS 10976SB0700B) and send it, through optical circulators (DK Photonics PMOC-F-3-98-5-PM980-90-05-FA), into a pair of fiberized FP 976 nm diode lasers (the “amp lasers,” Aerodiode 976LD-2-0-0), which will frequency lock to the seed given sufficient seed power and a free-running output frequency sufficiently close to the seed laser frequency. With this setup, we were able to achieve powers up to ≈ 220 mW per beam at the ion (a roughly 3x improvement over the free-space setup). This setup notably uses fiberized SP AOMs (Castech) instead of free-space ones. A schematic of this setup is shown in Figure 2.12b.

Phase stability at ion. The 976 nm beams are used for coherent operations, so it is important that the optical phase difference between them at the trap remains stable on experimental time scales. Since the beams are travelling first through separate fibers (subject to e.g. changes in refractive index due to mechanical stress or change in temperature) and then along separate free space paths, changes in relative beam phases are expected. Measurements of the width of an optical beatnote signal generated between the two beams close to the trap suggested that the phase stability of the fiberized injection-locked setup at the ion is ≈ 4 Hz (measured at the power FWHM).

2.5.2 854 nm AC Stark shift laser. The configuration for our 854 nm AC Stark shift laser is shown in Figure 2.13. This setup makes use of a fiberized diode laser (IPS, 10854.2SB0100PA) driven with a CTL300E-1-400 laser controller, passed through a fiberized optical isolator (DK Photonics, PMISO-F-850-L-C-PM780HP-90-10-FA) and a fiberized 99:1 splitter (DK Photonics PMFBTC-P-102-85-L-01-1-90-10-FA) directing 1% of the total power to the WAnD for wavelength measurement. The remaining 99% of power goes to a DP AOM. In total, this system gives us an ≈ 8 mW beam at the ion.




Figure 2.13. Schematic of the laser system used to supply our high-power 854 nm beam used for generating AC Stark shifts on the non-qubit levels of the $D_{5/2}$ manifold.

2.6 Beam delivery

After a beam is generated by a laser and passed through an AOM setup, it is delivered by fiber to the optical table where the trap is located, being brought into free space by a fiber collimator. The optics and electronics between this collimator and the ion are responsible for four things: 1) directing the beam onto the ion, 2) focusing the beam so as to achieve intensities high enough to couple reasonably well to the ion, 3) measuring beam powers for power servoing, and 4) adjusting beam polarizations for addressing specific optical transitions. These

four tasks are explained in detail below, and a full-picture schematic of our beam delivery optics is shown in Figure 2.14.

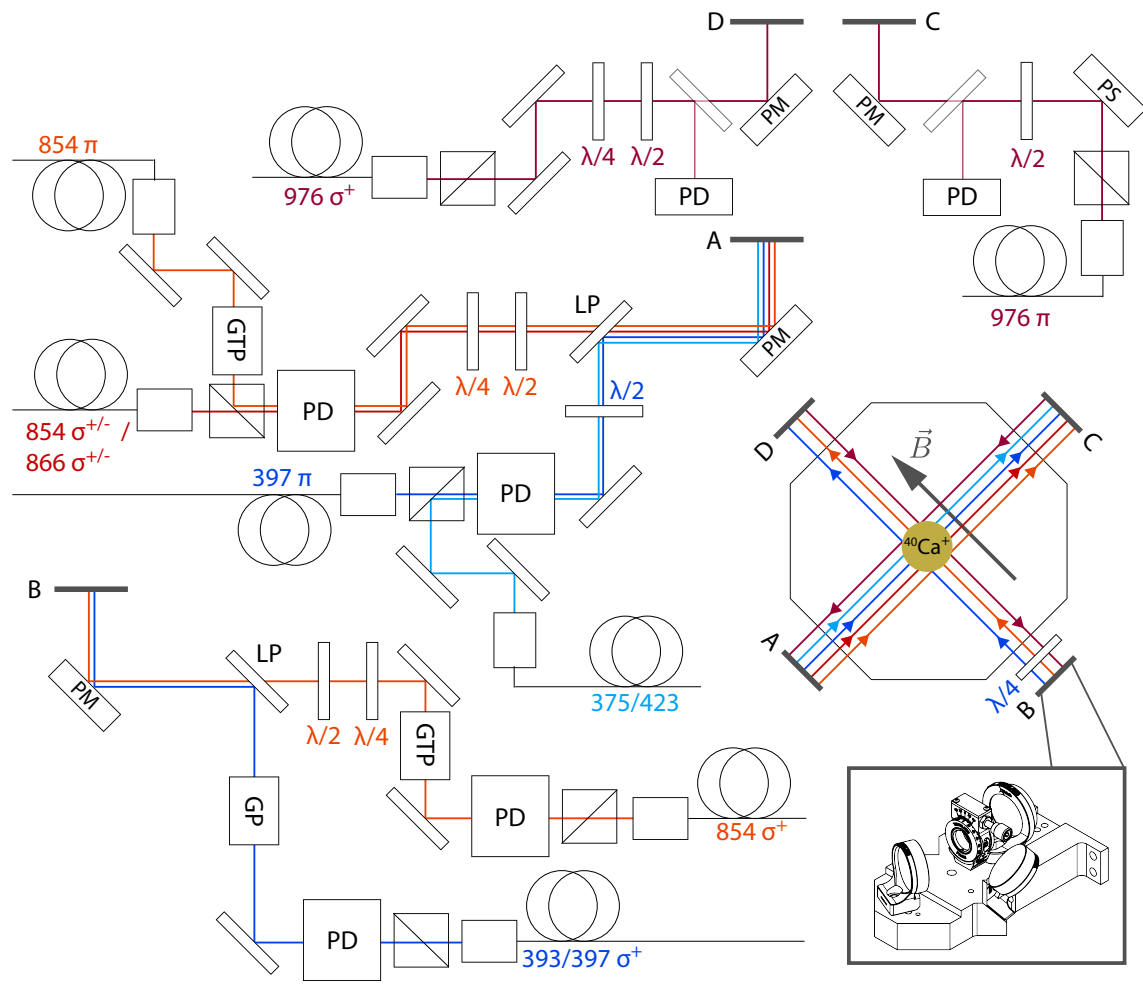


Figure 2.14. Schematic of the optics leading from the output collimators on our optics table to our trap. Definitions of abbreviations are as follows - PD: photodiode, GTP: Glan-Thompson polarizer, PM: piezo motorized mirror, PS: piezo stack mirror, GP: Glan polarizer.

2.6.1 Beam pointing. Most of the elements in our beam delivery system are mirrors, either on fixed or movable mounts, which allow beams to be directed to the trap. We use different types of movable mounts to achieve different levels of fineness in pointing adjustment, as described below.

Coarse adjustment of beam pointing. Before fine adjustment of beam pointing (i.e. trying to center the beam exactly on the ion) makes sense, the beam waist must be close enough to the ion that it is generating some sort of detectable signal (e.g. inducing fluorescence, optically shelving/deshelving, or causing quantum jumps). This level of beam alignment is usually achieved using manual adjustable mirror mounts alone (POLARIS-K05G4 and K1G4).

Techniques for beam alignment. In cases where a beam that needs to be aligned is co- or counter-propagating with another beam that is already on the ion, getting the unaligned beam onto the ion as well is relatively straightforward. Whether co- or counter-propagating, the new beam will be brought roughly onto the ion by adjusting its pointing until it is well aligned with the on-ion beam at two points on the optical path. In the co-propagating case, a mirror can be put after the shared focusing lens, and the focused beams can be visualized on a beam camera (e.g. the Thorlabs DCC1545M that we use in our lab), allowing the focus of the unaligned beam to be aligned with the focus of the on-ion beam directly. In the counter-propagating case, the unaligned beam can be back-aligned into the output coupler of the on-ion beam. In cases where there are no co- or counter-propagating beams, coarse alignment can be attempted by finding the RF trap electrodes and needle tips (i.e. by finding where the beam is “shadowed” on the other side of the vacuum chamber) and positioning the beam in the middle between them.

Piezo motor-driven mirror mounts. For fine adjustment and compensation for slow day-to-day beam drifts, we use piezo motor-driven mirror mounts

(Newport AG-M100N). These mounts move the mirror holder by using vibrations of a piezo motor (driven by a motor controller AG-UC8PC) to move the mount by discrete steps. This means that the mounts have a large tuning range ($\pm 2^\circ$), though the fact that the mechanism of action relies on the surface characteristics of the mount holder means that behavior is not consistent across the range, and trying to use small steps (i.e. ones that would deflect a beam a few microns at the ion) can be difficult due to friction.

Piezoelectric mirror mounts. For one of our 976 nm Raman beams, we need a level of fine control that the piezo motor-driven mirror mount is not able to achieve. For this, we use a piezoelectric mirror mount (POLARIS-K1S2P), which deflects the mirror using a piezoelectric stack driven by a DC voltage. In our setup, this voltage is supplied by a DAC channel in our experimental control system, described in more detail in Section 2.8. This DAC channel can supply up to +10 V DC, corresponding to an angular deflection of 50-100 μrad (translating to a shift at the ion of 15-30 μm). Since the DAC channel used can supply negative voltages, potentially damaging the piezo stack, the voltage is applied through a diode and across a resistor.

2.6.2 Beam focusing. The beams coming from output collimators at the trap have fairly large diameters (at the $1/e^2$ point in intensity), typically 1-2 mm. These are large relative to our trap geometries, so getting light onto our ions with these beams would be fairly easy. However, in addition to clipping on the trap electrodes, at the powers readily available to us, the intensities of these beams are far too low to appreciably couple to our ions. We increase these beams' intensities by orders of magnitude by focusing them down. Because we are sending multiple beams of very different wavelengths down the same paths (e.g. the 397 nm and 866 nm), we use 1" dia. achromatic focusing lenses (Edmund Optics 49-362)

with a focal length of $f = 150$ mm. The relevant features of our focused beams are described below.

Spot size. The diameter of the beam waists x_0 at the focus depends on the collimated beam diameter D according to the relationship $x_0 = 4\lambda f/\pi D$, where λ is the wavelength of the light and f is the focal length of the lens. Of course, for a given beam power P , average beam intensity over the spot $I = P/\pi x_0^2$ grows rapidly as x_0 decreases. However, the need for a tighter focus to maximize intensity must be considered alongside the need to drive multiple ions. All of our beams are global, coupling to all of our ions at once, and more tightly-focused beams (i.e. ones where the beam diameter is within a factor of a few of the ion-ion spacing in a crystal) will be prone to intensity inhomogeneities in multi-ion chains. In two-ion chains, these inhomogeneities arise if the beam is not centered on the chain, while for chains of three or more ions, these inhomogeneities will be inherent. In general, our red beams have spot sizes of $\sim 80 \mu\text{m}$, with $\sim 50 \mu\text{m}$ more typical for our blue beams, while our 976 nm Raman beams are focused to a tighter $\sim 30 \mu\text{m}$. A more minor consideration introduced by spot size is how easy it becomes to get the ion in the focal/Rayleigh range.

Rayleigh range. The Rayleigh range is the range around the focus in which the intensity is within a factor of two of its maximum (or, in other words, the region where the beam diameter is within a factor of $\sqrt{2}$ of its minimum). For a circular Gaussian beam, the range z_R is $z_R = 4\pi x_0^2/\lambda$, where the definitions of these values are the same as above. When we are focusing our beams down (either by adjusting their collimation or by moving our focusing lenses), we aim to ensure that the ion is within the Rayleigh range. Since adjusting a focus has the same goal as optimizing beam pointing (i.e. maximizing intensity), the same types of on-ion experiments can be used for both purposes.

2.6.3 Beam power control. For certain applications (i.e. operations where we are using lasers for coherent operations), we want to be able to continuously stabilize beam power over small time spans (10s to 100s of μs). For other applications (i.e. optical pumping), we only need to counteract drifts in average beam power over long time spans. Though precise control matters much more in the first case than the second, in all cases, we need some way of managing beam power fluctuations. This is done by measuring picked-off beam power with a photodiode and using this measurement to feed back on the RF amplitude of the associated AOM, changing its diffraction efficiency. Feedback here is carried out using the SU-Servo system built into our ARTIQ experiment control system, described in greater detail below.

Pickoff photodiodes. Most of our pickoff photodiode setups are of a design from the ion trapping group at Oxford University. This setup consists of a photodiode connected to a low-noise amplifier board, mounted in an anodized aluminum box with a filter for keeping room light off of the photodiode and a beam pickoff in the beam path, all mounted on an “outcoupler base” which has space for an output collimator, a PBS, a photodiode box, and an adjustable mirror. A labelled photograph of this setup is shown in Figure 2.15. Placing a PBS immediately before the photodiode serves two purposes: 1) acting as a first stage of polarization cleanup and 2) converting polarization fluctuations (which we cannot correct) into beam power fluctuations (which we can correct).

976 nm Raman beam photodiodes. For our high-power 976 nm Raman beams, we do not need low-noise amplification of our photodiode signals, and in fact, at these high powers, we see very long rise/fall times ($\sim 50\text{-}100 \mu\text{s}$) that effectively acts as a low-pass filter on laser power noise, reducing servo bandwidth. For these beams, we instead use 5% beam pickoffs directed towards fast (35 ns response time




Figure 2.15. A photograph with labels of one of our outcoupler bases, with the front plate of the photodiode box removed to show internal hardware.

with no additional loads) photodiodes (Thorlabs DET100A2) which, with 50Ω termination, give $\sim 1\text{-}2\text{ V}$ output signals with a $\sim 1\ \mu\text{s}$ rise time given a $\sim 10\text{ mW}$ pickoff power.

SU-Servo. Our experiment control system, as described further in Section 2.8.1, has DDS channels that we use to drive our AOMs as well as ADC (analog-to-digital converter) channels that can be used to read in analog voltages. This system has a built-in utility, the SU(Sampler-Urukul)-Servo, that allows these devices to be used in tandem. An ADC can read in a pickoff photodiode voltage, the FPGA board managing the experiment control system can compare this input reading to some set point and, using a PI loop, generate an error signal which it can use to feed back on the RF amplitude on the corresponding DDS. This loop is illustrated schematically in Figure 2.19b.

Calibrating power servo parameters. SU-Servo regulates AOM RF power using a PI loop, meaning that its behavior is defined by a proportional servo parameter k_p and an integral parameter k_i . To find optimum values of these parameters, we want to find what pair gives the smallest percent RMS error ϵ in

beam power / photodiode voltage ($\epsilon = (V - V_0)_{rms}/V_0$, where V is measured photodiode voltage and V_0 is photodiode voltage set point, calculated from desired beam power) in the time/frequency ranges relevant to the experiment in question. There are two techniques we generally use for acquiring this RMS error. In the first technique, we turn the power-servoed beam on and measure the photodiode voltage at fixed time intervals of length t (where t is necessarily longer than the $\approx 1.2 \mu s$ SU-Servo cycle time), giving measurements sensitive to noise in the $f > 1/t$ frequency range. This limited bandwidth is not a problem for most of our beams so long as t is less than or comparable to operation times, since these beams are being used for either optical pumping or for coherent operations (i.e. Rabi flopping) dependent only on total pulse energy, and because this process can be carried out with the limited sampling readout speeds achievable with the experimental control ADC, it can be performed while rapidly scanning k_i or k_p to try to minimize ϵ . In cases where limited bandwidth of this “manual” sampling can be a problem, we can instead measure photodiode voltages at higher resolution using an oscilloscope and analyze the RMS error in the generated trace after the fact.

2.6.4 Polarization control. In general, we want our beams to drive transitions between specific sets of Zeeman sublevels. This means that we must be able to set the polarization of our beams at the ions.

Polarization purification. As discussed above, the first stage of polarization purification is the passing of a beam through a PBS prior to power measurement with a photodiode pickoff. For most of our beams, this is the only polarization purification we carry out. For our 854 LS and 854 π beams, we use an additional Glan-Thompson polarizer (and a Glan-Laser polarizer for our 393/397 σ^+ beams).

Polarization adjustment and optimization. To change the polarizations of our beams, we use waveplates. Waveplates are plates of birefringent material with

thicknesses such that, at a defined wavelength, the polarization components along the two axes of the material accumulate a $\pi/2$ phase difference (in the case of a $\lambda/2$ waveplate) or a $\pi/4$ phase difference (in the case of a $\lambda/4$ waveplate). We use $\lambda/2$ waveplates to rotate linear polarizations and $\lambda/4$ waveplates to convert linear polarization to circular. For beams where our ion is highly sensitive to polarization impurities (e.g. on the 854LS beam, where polarization impurities will tend to cause deshelving of the qubit), we can minimize these impurities by minimizing the associated signal on the ion while adjusting the associated waveplates.

2.7 Electronic control fields

The state of an ion can be controlled without using optical fields. The fact that an ion is electrically charged means that the motion of its center of mass can be driven using electric fields, and the fact that the ion has a magnetic moment means that its magnetic orientation (its spin state) can be made to precess using magnetic fields. The physics of these processes and coupling rates under normal operating parameters are described further in Chapter III. This section describes how all-electronic operations are implemented in our hardware. In every case, we apply electronic control fields to the ion by putting some control voltage on our trap electrodes. These systems are represented schematically in Figure 2.16.

2.7.1 Magnetic Rabi drive. To generate an RF magnetic dipole that will couple between the Zeeman levels of our ion, we connect a DDS to one of our DC trap electrodes through an SMA vacuum feedthrough (MPF A1908-2-W). This DC trap electrode is shorted on the other end. Typically, the DDS channel supplying this drive is operated at ≈ 10 dBm with an ≈ 13 dB internal attenuation applied to it before the output. It is then passed through a 20 dB coupler, with the lower-power output connected to a patch panel (a configuration shared with all of

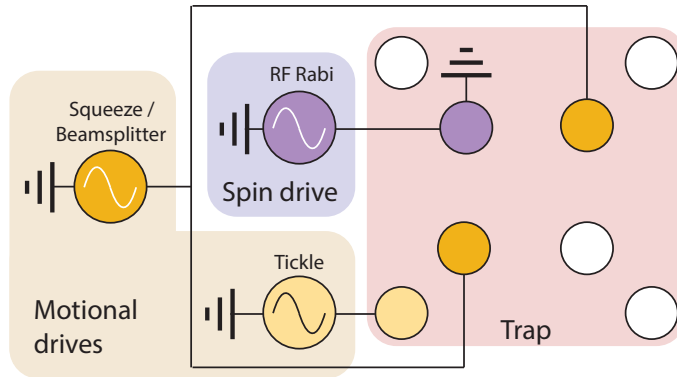


Figure 2.16. Schematic summary of the RF voltages applied to the trap for coherent driving, marking the trap electrodes (shown in cross-section) that they are connected to.

our DDS outputs) and the other connecting to the DC trap electrode through a 3 dB external attenuator meant to reduce back-reflection to the DDS.

2.7.2 Displacement drive. We can generate displaced motional states (analogous to classical harmonic oscillator states) of an ion’s motional modes by applying a dipole electric field resonant with these motional modes (a “tickle”). To generate these fields, we apply an RF voltage generated by a DDS to one of our DC compensation electrodes. The DC voltage and displacement RF voltage are combined on a bias-tee (Mini-Circuits ZFBT-4R2G+), before which the displacement RF is attenuated 73 dB. This attenuation is applied in order to prevent ion driving/heating through leakage on the associated DDS channel, which only has 70 dB of isolation on its switch. As a side note, because of how the motional modes of the trap will tend to align when an RF bias voltage is applied to split the radial mode frequencies, the electric dipole from the DC rod drive will tend to have a 10-100x higher coupling strength on one mode than the other.

Axial displacement. Normally, we only carry out motional operations on our “radial modes” (motional modes perpendicular to the trap axis). However, when initially trying to characterize our trap, it was necessary to measure our axial mode

frequency. In order to do this, we connected a displacement drive to one of our DC needles through a bias-tee made in-house to accommodate large DC voltages.

2.7.3 Squeeze/beamsplitter drives. To perform motional squeezing/beamsplitter operations, we need to apply quadratic potentials at the ion. Since this is the same type of field being used to trap the ion, we apply these potentials directly to our RF trap electrodes, combining them with the trap voltage as a bias. We do this through a set of tank circuits (just labelled “Filter box” in Figure 2.17), which also allows the application of a DC bias to the trap RF voltage. In our motional control experiments, we end up operating at squeeze drive frequencies of ~ 3.6 MHz and beamsplitter frequencies of ~ 33 kHz. This relatively low beamsplitter frequency means that this drive cannot be supplied by one of the DDS channels in our experiment control system, which experience fast roll-offs in power below ~ 100 kHz. Instead of using a DDS channel for this drive, we use an AWG channel triggered off of a TTL input. The squeeze and beamsplitter drives are combined before the resonator box with a simple tee intersection. A block diagram of the overall motional drive setup is shown in Figure 2.17.

2.7.4 Pulse shaping. By default, the DDSes we use for our electronic drives operate at a fixed amplitude and are simply turned on and off with a digital switch. This arrangement generates rectangular pulses, which can cause issues with off-resonant driving due to the long-tailed sinc^2 power spectrum associated with this pulse shape. These issues are especially acute in many of our motional control experiments, which involve running at small mode frequency splittings relative to our operation times (e.g. $100 \mu\text{s}$ operations performed on a mode 10s of kHz separated from the next nearest mode). To address this, we have two options for reducing off-resonant driving. One option is simply to turn down coupling rates g such that $g \ll \Delta\omega$, where $\Delta\omega$ is the mode frequency difference. This is often

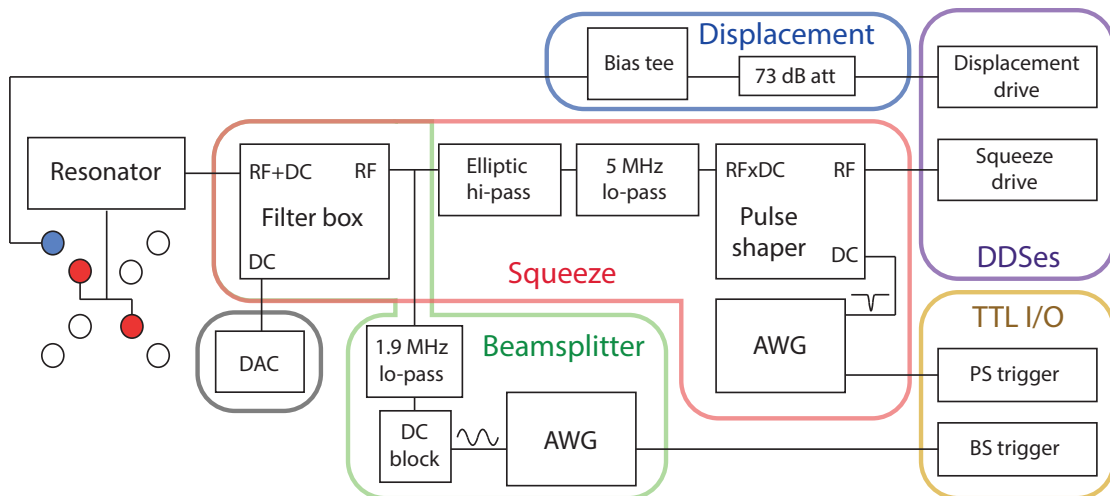


Figure 2.17. Block diagram showing the RF chain for three motional drives used in our experiments. *PS*: pulse shaper, *BS*: beamsplitter.

problematic because, at sufficiently low g (such that operations must be very long to generate the same state parameters), the operation in question becomes more sensitive to decoherence. The other option is to shape pulses to change their power spectra and reduce components that induce off-resonant driving. In our setup, we use two types of pulse shaping, performed different ways in hardware. These setups and their outputs are represented schematically in Figure 2.18. It should be noted that these methods, described below, were only ever applied to our motional drives. RF magnetic field drives are always operated as rectangular pulses.

Blackman pulses, DDS RAM mode. Our DDS channels use AD9910 DDS chips (“1 GSPS, 14-Bit, 3.3 V CMOS Direct Digital Synthesizer” (2016)). These chips can be run in “RAM mode,” in which the DDS can be triggered to generate a pre-programmed pulse with arbitrary time-varying amplitude, frequency, and/or phase. In our experiments, we used this mode to generate pulses with a Blackman amplitude profile, described by the function $y(t) = y_0[a_0 - a_1 \cos(2\pi t/T) + a_2 \cos(4\pi t/T)]$, where T is total pulse time, y_0 is a measure of the RF amplitude (in

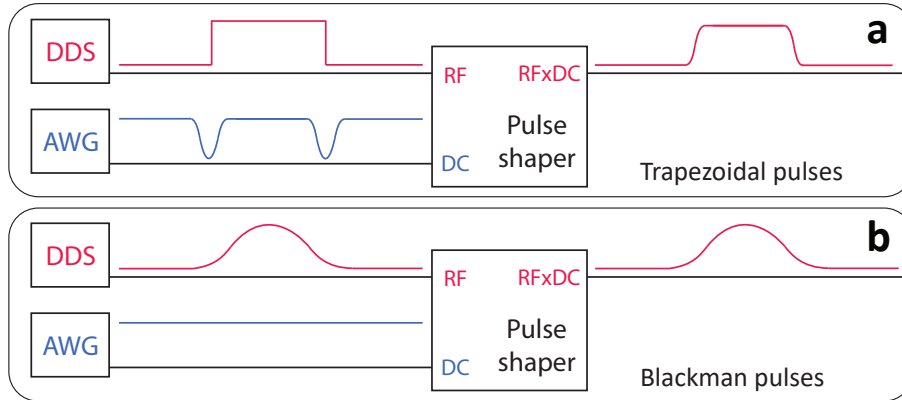


Figure 2.18. Schematic illustration of our two pulse-shaping schemes. **a)** Pulse shaping by multiplying a fixed-amplitude RF pulse with a nearly-trapezoidal DC envelope. **b)** Pulse shaping by running the DDS in RAM mode and generating arbitrarily shaped pulses, with the DC output on the AWG being unchanged.

our control system, usually units of full scale, ranging from 0 to 1), and a_0, a_1, a_2 are a set of constant coefficients such that the amplitudes at $t = 0$ and $t = T$ are zero.

Trapezoidal pulses, RFxDC multiplication. In contexts where there is less need for aggressive pulse shaping (i.e. cases where adjacent modes are further away from each other, but it is still desirable to suppress far-off-resonant driving, ringing/power reflection from tank circuits, etc), we instead use (approximately) trapezoidal pulses. Despite having a power spectrum with more far-off-resonant power, this profile has advantages over Blackman windows. Because trapezoidal pulses spend more of their time at full amplitude, they are more efficient. In addition, for operations extending beyond twice the rise time, there is some window of constant power during which RF amplitude can be servoed to control beam power, if shaping is applied to an RF drive for an AOM. The way that this shaping hardware is set up, this shaping can be applied to both servoed and non-servoed DDS channels (as opposed to RAM mode pulse shaping, which can only be used on non-servoed channels). This shaping is done using an RFxDC multiplication

circuit. In this setup, the RF input to be shaped is multiplied by a constant voltage output from an AWG. When triggered, this AWG channel will generate a single Blackman pulse going down to a voltage of zero. If these pulses are correctly timed with the RF input, the output will be a RF pulse with sloped sides, as sketched in Figure 2.18a.

2.8 Experiment control system

For controlling the low-level functions of our experiments, we use ARTIQ (Advanced Real-Time Infrastructure for Quantum physics, Kasrowicz et al. (2020)), a software/hardware system specifically designed for running atomic physics experiments. A schematic overview of the ARTIQ control system and how it interfaces with both our experiment hardware and our local computer network is presented in Figure 2.19a. The individual components of the ARTIQ system are discussed below.

2.8.1 Control hardware. ARTIQ is designed to control a set of experiment control cards (EEMs, or Eurocard Extension Modules) designed as part of the Sinara open-source hardware project and fabricated by Technosystem and Creotech. During an experiment, every action taken that requires exact timing (which excludes tasks that require communicating with the host PC, like reading/writing data from/to this PC, controlling remote devices like AWGs, etc.) is performed by one of these cards. A reference for these cards, with names and descriptions (serving as a glossary for future sections) is shown in Table 2.

2.8.2 Clocking. Though each Kasli FPGA board has its own onboard 125 MHz quartz oscillator clock, we clock the control system to an external source. The ultimate timing reference for the system is a 10 MHz rubidium clock (SRS FS725), to which we lock a 125 MHz phase-lock oscillator (Wenzel VHF PLO). This PLO is powered with a +12 V voltage regulator (TPS7A4701EVM-094). The

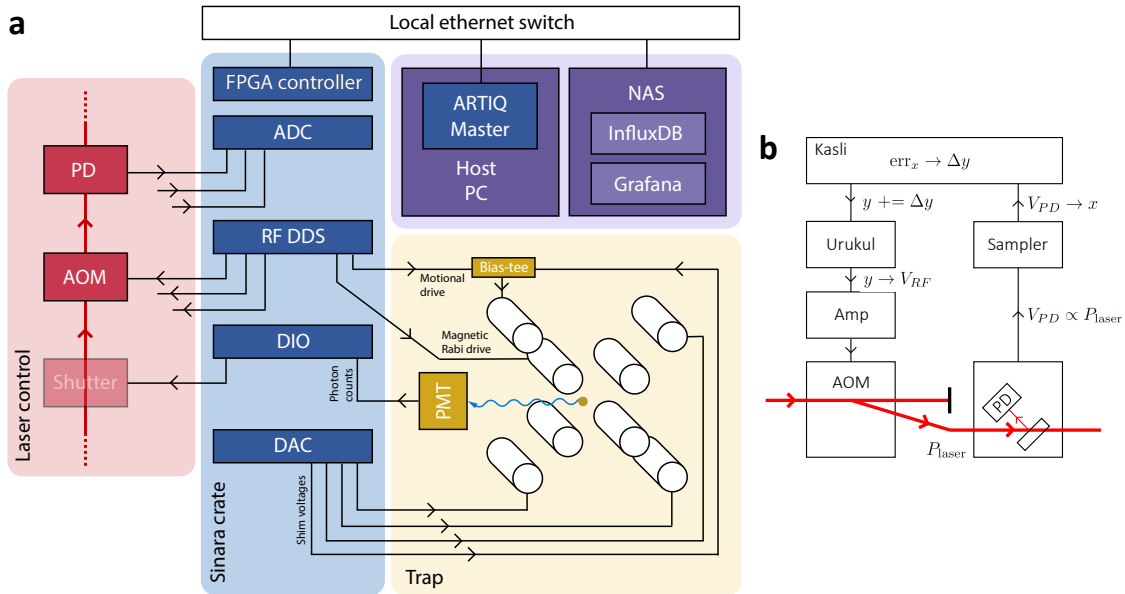


Figure 2.19. Schematic illustration of our experiment control system. **a)** Block diagram showing how our experiment control hardware (blue) interfaces with other blocks of our experiment, including beam control (red), the trap/imaging system (yellow), and the host PC allowing users to interface with the system (purple). **b)** A schematic of the beam power control system we use (SU-Servo).

PLO output is split four ways and directed to each of our labs, along with the raw 10 MHz Rb clock output. Within the control system crate, the four Urukuls are all clocked to the Kasli through local connections.

2.8.3 DDSes and phase-coherent RF control. Compared to other components of our control system hardware, our direct digital synthesizer (DDS) boards (Urukuls, Kasprowicz et al. (2022)) assume more varied functions (i.e. applying electrical, magnetic, and optical drives in a phase-coherent way) and have multiple modes of operation. These functions of the DDSes have been outlined above in Sections 2.4.3, 2.6.3, and 2.7, so discussion below will discuss some of the basic principles of these DDSes and the modes of operation that we use in our experiments.

Board	Description	Purpose
Kasli	FPGA carrier boards	Experiment control
Urukul	DDS boards	AOM driving for laser control motional E-field drives spin B-field drive
Sampler	ADC boards	Photodiode voltage read-in for SU-Servo
Zotino	DAC boards	Micromotion compensation DC bias on trap RF piezo mirror control
DIO	TTL input/output boards	PMT photon counting line-trigger timing measurement mechanical shutter control AWG triggering Squareatron output amplitude control

Table 2. A summary of the components of our control system and what we use them for in our experiments.

Free-running mode. By default, a DDS channel in our setup can be thought of as an always-running phase accumulator. If the DDS is set to output a signal with a frequency f , then at each Δt long step of the clock, a total phase is incremented by $\Delta\phi = f\Delta t$ (where ϕ is phase in turns, one turn being 2π radians). The output at a given time t is then calculated $y = y_0 \sin(2\pi(\phi(t) + \phi_0))$, where ϕ_0 is a fixed phase offset (also in turns) and y_0 is a fixed amplitude in units of full scale, both of which are values set by the user. The total phase $\phi(t)$ depends on the full history of the DDS channel, even when the channel is not putting out a signal, which can be problematic in contexts where the initial phase of a DDS channel at the start of an experiment matters. In these contexts, the total phase of the DDS can be reset at the start of each experiment, effectively setting an absolute time reference at the point where this reset happens. Notably, this feature is only available by default on DDS channels that are *not* part of the SU-Servo system. This issue and its solution are described in greater detail below.

SU-Servo and the Oxford patch. We regulate the power of our laser beams at the ion by measuring the voltage generated on photodiodes by beam power picked off before the trap and feeding these measurements back on the RF voltage driving our AOMs, which is ultimately provided by DDSes in our experiment control setup. The ARTIQ system supports this sort of feedback loop natively with the SU-Servo system, the loop for which is illustrated schematically in Figure 2.19b. Many of the relevant details of this system (how we implement it in hardware and how we optimize the proportional/integral parameters of the servo) have already been outlined in Section 2.6.3 above. Something important to note here though is that the SU-Servo DDS channels, though driven by the same DDS chips as the non-SU-Servo channels, are in certain respects less flexible, trading the ability to regulate beam power in real time with the ability to control DDS phase in an absolute way. This is worked around in a software/firmware patch developed by the ion trapping group at Oxford which allows absolute time references for DDS phase to be set on channels that are part of the SU-Servo system. The features enabled by this patch are used in our experiments requiring phase-coherent control of laser operations performed at different frequencies (e.g. in entangling gates).

2.8.4 Experimental control codebase. The structure of our experiment codebase can be summarized as follows: The control code for each experiment is an object which in turn contains a set of objects and methods that are called when the experiment is run. Each experiment inherits from some higher-level experiment class which comes from our patched version of ARTIQ and which allow ARTIQ utilities to compile the experiment and process user inputs / experiment outputs. Most experiments also inherit from some generic experiment class (written in-house) which defines a set of commonly-used objects, methods, and procedures that are used in nearly every experiment, defining a

generic “experiment loop” (represented schematically in Figure 2.20b) for which the only experiment-specific code that needs to be written is the middle block (“Experiment”), where some set of operations is carried out on a qubit (or set of qubits) prepared in a set state prior to state characterization. A block diagram of our codebase reflecting the discussion above is shown in Figure 2.20a.

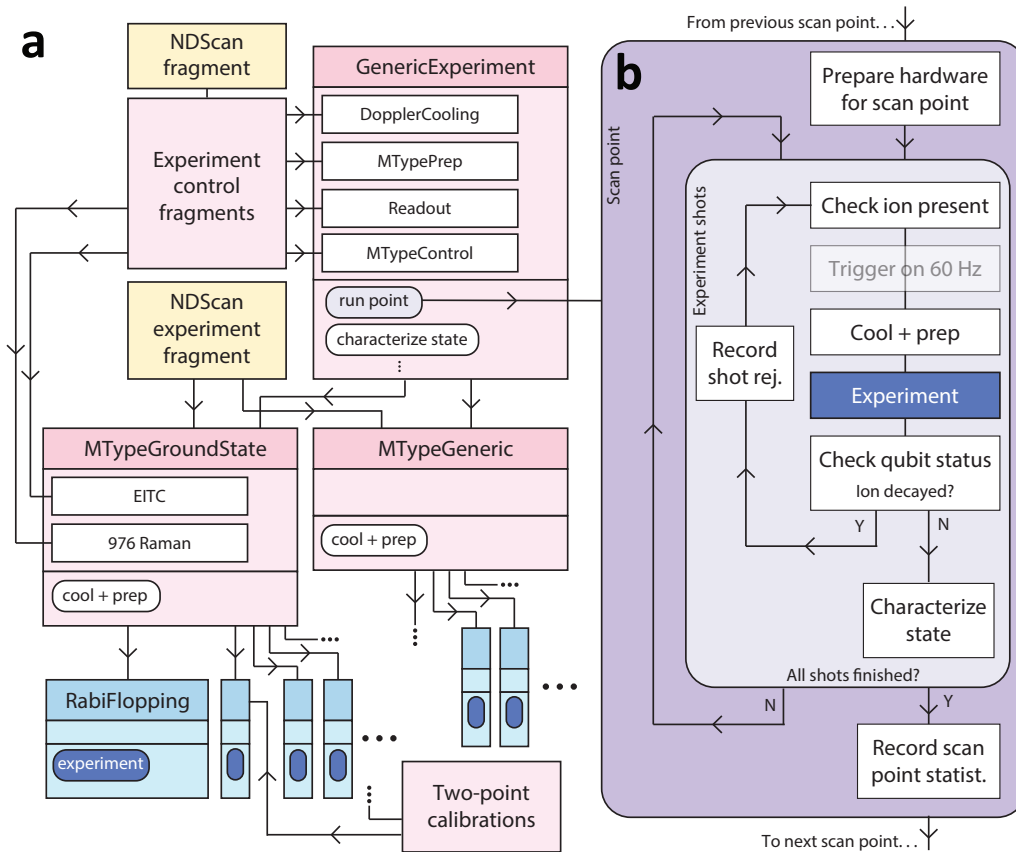


Figure 2.20. Class inheritance diagram showing the hierarchy of classes in our in our experiment codebase, down to the level of code for running individual experiments (e.g. an experiment for driving Rabi flopping on a qubit, RabiFlopping). NDScan classes are shown in yellow, in-house support classes are shown in pink, and experiments are shown in blue. The inset in purple shows the sequence for a generic scan point. In both portions, the main block of code that differs between individual experiments is shown in dark blue.

The platform. Our codebase is written in Python 3 using ARTIQ classes and utilities, patched with an add-on system out of the Oxford ion trap group called

NDScan (standing for “N-dimensional scan,” Oxford (2024)), which helps to enable a more object-oriented approach through the inclusion of “fragments” (objects reusable across experiments with their own associated arguments and result) in experiment code as well as allowing experiments to perform scans over any set of experiment arguments. Python code is sent from the ARTIQ Master and compiled on the Kasli using a Python compiler with a reduced feature set designed to be able to run on the Kasli FPGA.

2.9 Ion loading

There are four essential requirements for loading an ion in our trap: 1) a trapping potential with parameters such that calcium can be stably trapped, 2) a beam of neutral calcium atoms passing through the trapping region, 3) a photoionization (PI) beam to promote a valence electron in a neutral calcium atom to the continuum, leaving the calcium atom singly-ionized, and 4) a laser cooling beam capable of cooling very hot ($\sim 400^\circ\text{C}$) ions to a point where they can be detected. Requirement 1 was discussed in Section 2.1, and requirement 4 will be discussed next chapter in Section 3.2. Requirements (2) and (3), satisfied respectively by a calcium oven and a 423+375 nm PI beam, are discussed in greater detail below. The overall geometry of the oven/trap/PI setup is shown in Figure 2.21.




Figure 2.21. **a)** Illustration of the relevant photoionization geometry, showing the neutral calcium beam coming from the oven and the PI beam relative to the trap (electrodes shown in black). The neutral calcium and PI beams are both in the plane of the page. **b)** Energy level diagram of the two-photon PI process.

2.9.1 Calcium oven. We generate an uncollimated beam of neutral calcium in our vacuum chamber using an oven. The oven itself is a 1.65 mm dia. stainless steel tube (Goodfellow FF217200) with 0.13 mm thick walls and a hole 0.5 mm dia. drilled in the side. The ends are pinched off and spot-welded to a pair of 2 mm dia. stainless steel rods that are connected to an electrical feedthrough. The oven is filled with a 7.1 mg sample of calcium with much higher fraction of ^{43}Ca ($\approx 20.5\%$) than natural abundance (0.135%) to enable easier loading of this isotope. To heat it up, the oven is driven with a 4.4 A DC current provided by a Rigol DP711 power supply. Typical output voltages in this current range (including voltage drop across cabling and connectors) is $\approx 0.9 - 1.0$ V.

2.9.2 Neutral calcium detection. To confirm that the oven is putting out a beam of neutral calcium in the trapping zone, we perform spectroscopy on the $423\text{ nm } ^1S_0 \rightarrow ^1P_1$ transition in neutral calcium and look for fluorescence at this wavelength. Though the imaging system is designed for use at 397 nm, if the 397 nm narrow-band filter is replaced with one for 423 nm (Edmund Optics 34-497), we can measure 423 nm photon counts with our PMT. As discussed below, our 423 nm beam does not have an associated AOM, so we cannot use an RF drive to modulate the beam frequency. Instead, we adjust the cavity grating piezo voltage of the 423 nm diode laser itself and correlate this voltage with an absolute frequency using our wavemeter.

In addition to verifying the presence of calcium in the trap, 423 nm spectroscopy also enables us to 1) verify the relative quantities of different calcium isotopes in our enriched sample and 2) measure the Doppler shifts/broadening of our beam.

Isotope shifts and abundances. Differences in the mass and volume of the nucleus between different isotopes of the same element cause isotope shifts, or small

($O(100\text{ MHz})$) shifts in the frequencies of optical transitions (Lucas et al. (2004); Tanaka, Morita, and Urabe (2007)). Given that these shifts are larger than the 34.5 MHz linewidth of the 423 nm line (Salumbides et al. (2011)), it is in principle possible to resolve these transition for every isotope in our sample. We can further use the strengths of these transitions (as measured by peak photon count rates) to estimate the relative abundance of each isotope, with relative abundance and relative scattering rates being directly proportional.

Doppler effects. If the calcium beam (i.e. calcium travelling along the line between the oven opening and the trapping zone) were orthogonal to the PI beam, the temperature of the beam (via the velocities of the calcium atoms) would have no effect on the absorption of 423 nm photons. However, because the calcium beam is only $\approx 70^\circ$ from the PI beam, the hot calcium atoms have a component of their velocity parallel to this beam. This means that the 423 nm line will be Doppler shifted (due to the non-zero average atom velocity parallel to the 423 beam) and Doppler broadened (due to the non-zero spread of atom velocities). Assuming a Maxwell-Boltzmann velocity distribution, we can calculate the expected Doppler shift and broadening. Based on earlier neutral atom work in our lab using similar ovens, we estimated an initial oven temperature of $\approx 400^\circ\text{C}$, corresponding to an average atom speed of $\approx 650\text{ m/s}$, from which can calculate the expected Doppler shift to be 526 MHz, and the Doppler broadening to be $\sigma = 292\text{ MHz}$, where σ is the standard deviation of a Gaussian line shape.

Measured abundances and linewidths. A sample 423 nm fluorescence line, including manual fits of Doppler-broadened Gaussian lines for ^{40}Ca , ^{43}Ca , and ^{44}Ca (accounting for the shifts of the three hyperfine levels of ^{43}Ca and taking the heights of these lines to be approximately equal) is presented in Figure 2.22a. These fits suggest relative abundances of approximately 56%, 26%, and 17% for

^{40}Ca , ^{43}Ca , and ^{44}Ca respectively, roughly consistent with expectations. The Doppler broadened linewidth, ≈ 300 MHz, matches with the calculated estimate almost exactly. The measured Doppler shift, ≈ 300 MHz, is smaller than expected, but this is consistent with systematics of the wavemeter in this frequency range. As seen from the data, the lines for specific isotopes are not well-resolved. This problem is discussed in greater detail below.

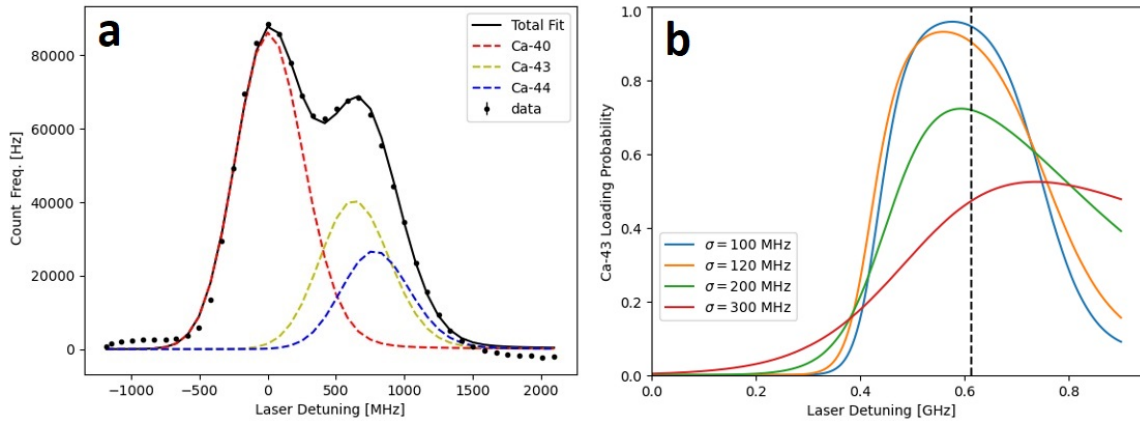


Figure 2.22. **a)** Fluorescence spectrum data (black dots) on the 423 nm line for our neutral calcium beam with a manual fit including three calcium isotopes (individual isotopes shown with dashed colored lines, total fluorescence shown with a solid black line). **b)** Relative probability of loading $^{43}\text{Ca}^+$ at different 423 nm beam detunings and with different amounts of Doppler broadening.

2.9.3 Photoionization. Photoionization is a technique for generating ions by optically exciting a valence electron to an unbound state. In calcium, this is a two-photon process. The first photon (at 423 nm) excites the electron from the 1S_0 state to 1P_1 state, and a second photon (at 375 nm) excites from the 1P_1 state to the continuum, illustrated schematically in Figure 2.21b. Unlike older ionization techniques like bombarding calcium with hot electrons, this process allows ionization without the introduction of substantial ion heating or trap charging and allows loading with a lower calcium flux, reducing calcium deposition on and near the trap. In addition, as discussed above, because the 423 nm line sees

isotope shifts larger than the natural linewidth, this technique allows for isotope-selectivity, or the ability to control with a high level of certainty what isotope will be loaded into the trap.

423/375 nm laser breadboard. The power for the two frequencies in our photoionization beam comes from a 423 nm diode laser of the same type as the other lasers (and on the same type of breadboard) discussed in Section 2.4.1 combined on a second breadboard (the 423/375 breadboard) with power from a 375 nm laser (Toptica iBeam smart). This combined power is then split, fiberized, and directed to different labs. Because we do not frequency-modulate these beams with AOMs, and because we do not need turn-on/off speeds faster than milliseconds, we use mechanical shutters placed before the output couplers to block power from entering the experiment setup. These mechanical shutters are driven by a box triggered using a TTL output of our experiment control system.

Isotope-selectivity. In principle, the >100 MHz isotope shifts on the $^1S_0 \rightarrow ^1P_1$ line should allow us to control what isotope is photoionized by simply tuning the 423 nm laser. In practice, the Doppler broadening we observe makes this selectivity more challenging, particularly for loading the low-abundance ^{43}Ca , whose $^1S_0 \rightarrow ^1P_1$ line sits between those of ^{40}Ca and ^{44}Ca . Based on the fits $F_i(f)$ for each isotope i , where f is the tuning of the 423 nm laser, we can write an expression for the probability of a photoionization event being of ^{43}Ca (assuming identical efficiency of the 375 nm beam for each isotope)

$$P_{43}(f) = \frac{F_{43}(f)}{F_{40}(f) + F_{43}(f) + F_{44}(f)}.$$

This relative probability $P_{43}(f)$ is plotted for several Doppler broadening parameters in Figure 2.22b, showing that for the $\sigma \approx 300$ MHz value observed, we expect no better than a $\sim 50\%$ selection probability.

Improving isotope-selectivity. The issue with isotope selectivity discussed above could be addressed in part by changing the angle of the PI beam through the viewports such that the PI beam is closer to orthogonal with the neutral calcium beam (though we can only rotate the PI beam from 4.2° , assuming that the PI must still pass through the focusing lens it shares with other beams on that path). Additionally, running the oven at a lower temperature so as to reduce the Doppler broadening is also likely to help. Finally, because the cycling transition on the ionized calcium also experiences an isotope shift, and because the $^{43}\text{Ca}^+$ lines are higher-frequency than the $^{40}\text{Ca}^+$ line, the 397 nm beam for driving the cycling transition can be tuned such that it is red of the $^{43}\text{Ca}^+$ lines, cooling this isotope, while it is blue of the $^{40}\text{Ca}^+$ line and will tend to heat it out of the trap.

CHAPTER III

EXPERIMENTAL METHODS

With our specific apparatus explained in Chapter II, it makes sense to step back and look at the more general techniques we seek to implement with this hardware. Most are common throughout ion trapping (or even QIP in general) and some (e.g. certain types of electronic motional control) are more specific to our own work.

The organization of this section roughly follows the “order of operations” for controlling an ion. In other words, to run an experiment with trapped ions, one must 1) trap the ion in an effective potential well (as a precursor to everything else), 2) cool it to a point that Doppler broadening is negligible relative to the linewidths of the control lasers (or for some applications, as far down as the motional ground state), 3) initialize the ion into some known state, 4) read out the state of the ion following an experiment, and 5) carry out some set of coherent operations that put the ion into some state different than the one it started in. The sequence for a generic experiment in our setup is presented in Figure 3.5a, showing how experiments are performed repeatedly (i.e. several “shots” of the experiments are run) to build up statistics that allow the final state of the system to be characterized. The specific details of these experimental steps are shown in Figure 3.5b (omitting the details of cooling, which vary between experiments but which generally take place before/during state preparation).

For reference in discussions below, an electronic/motional level diagram (showing just one of three modes) for a single $^{40}\text{Ca}^+$ atom is presented in Figure 3.1, showing the transitions addressed and the lasers (or other control fields) used to address them.




Figure 3.1. A level diagram showing a Ca^+ ion trapped in a harmonic oscillator potential (only one dimension of which is shown). Transitions between internal electronic energy levels are marked with the wavelengths and polarizations that we use to drive them. Energy levels are not to scale.

3.1 Trapping

While a brief discussion was offered in Section 1.2.1 of how a Paul trap works, this section presents a more thorough and quantitative explanation. The motion of a charged particle in a time-dependent potential is a potentially complex dynamical problem, but it can be simplified greatly by thinking in terms of the ion experiencing some static effective potential (pseudopotential, ϕ_{eff}) arising from the time-averaging of an RF electric field. Considering the shape of this pseudopotential and how this relates to the shape of the underlying field in turn offers a way of thinking about (and calculating) the degree of anharmonicity in an ion trap.

A brief derivation of the pseudopotential, directly following a treatment by Dehmelt (1968), with some additional comments to offer greater intuitive clarity, is presented below.

Force on a charged particle in a uniform electric field. An ion with charge $+e$ sitting in a uniform, one-dimensional electric field (i.e. one that does not depend on the coordinate of the particle z) $E_z(t) = E_0 \cos \Omega t$ (where E_0 is the electric field amplitude and Ω is, in the context of an actual trap, the drive frequency) experiences a corresponding one-dimensional force

$$F_z(t) = m\ddot{z} = eE_0 \cos \Omega t. \quad (3.1)$$

A solution satisfying this equation of motion has two parts: a slow component (the average position \bar{z} of the ion, which in this case is static) and a fast component $\zeta(t)$ matched in frequency to the drive (“micromotion,” discussed more thoroughly in Section 3.2.5). The full solution is

$$z(t) = \bar{z} + \zeta(t) \quad (3.2)$$

$$= \bar{z} - \zeta_0 \cos \Omega t, \quad (3.3)$$

where the amplitude of the fast component ζ_0 can be shown to be $\zeta_0 = eE_0/(m\Omega^2)$. (Important to note is that $\zeta(t)$ is π out of phase with the driving field. The in-phase solution is not steady-state and would cause the amplitude of the ion’s motion to grow indefinitely. This of course is not possible for an unbounded ion, which has no “restoring force” other than the drive itself to keep its motion in-phase. This fact can be thought of as the DC limit of the more general idea that drives which are higher-frequency than the oscillator they are driving tend to drive out-of-phase motion. Consequences of this are touched on briefly in some of the discussion below).

Time-averaging force, adding non-uniformity. If we are interested in the motion of the ion on time scales longer than a period $2\pi/\Omega$ of the motional drive, we can consider time-averaging the force $F_z(t)$ over this period. In the case where E_0 is uniform over the range travelled, this averaged force is trivially zero.

However, if some position-dependence is introduced, such that $E_0 = E(\bar{z})$ and the the electric field during the ion's motion is $\approx E(\bar{z}) + \zeta(t)(\partial E/\partial z)|_{z=\bar{z}}$, we can consider time-averaging this force over a period, such that

$$\langle F_z(z, t) \rangle = F_{\text{ave}}(\bar{z}) = e \langle [E(\bar{z}) + \zeta(t)(\partial E/\partial z)] \cos \Omega t \rangle. \quad (3.4)$$

Recognizing that the first term (since it depends only on $\cos \Omega t$) will time average to zero, assuming that the electric field gradient is sufficiently small such that $\zeta(t) \approx -\zeta_0 \cos \Omega t$ as in the uniform case, and substituting in the value of ζ_0 , we get an average force over one cycle of

$$F_{\text{ave}}(\bar{z}) = e \langle -\zeta_0 (\partial E/\partial z)|_{z=\bar{z}} \cos^2 \Omega t \rangle \quad (3.5)$$

$$= -e \zeta_0 (\partial E/\partial z)|_{z=\bar{z}} \langle \cos^2 \Omega t \rangle \quad (3.6)$$

$$= -e \zeta_0 (\partial E/\partial z)|_{z=\bar{z}}/2 \quad (3.7)$$

$$F_{\text{ave}}(\bar{z}) = -\frac{1}{2} \frac{e^2 E_0}{m \Omega^2} \frac{\partial E}{\partial z} \Big|_{z=\bar{z}} = -\frac{1}{2} \frac{e^2}{m \Omega^2} E(\bar{z}) \frac{\partial E}{\partial z} \Big|_{z=\bar{z}}. \quad (3.8)$$

The pseudopotential. The effective force above is proportional to $E(\bar{z})$ and its gradient. If we take $F_{\text{ave}}(\bar{z})$ to be the gradient $-(\partial \psi/\partial z)|_{z=\bar{z}}$ of some potential $\psi(\bar{z})$, a potential of the form $E^2(\bar{z})$ satisfies this condition, i.e.

$$\frac{1}{e} F_{\text{ave}}(\bar{z}) = -\frac{\partial \psi}{\partial z} \Big|_{z=\bar{z}} \rightarrow \psi(\bar{z}) = \frac{e}{4m\Omega^2} E^2(\bar{z}). \quad (3.9)$$

This $\psi(\bar{z})$ is the pseudopotential (in one dimension), which can be generalized to three dimensions as $(1/e)\vec{F}_{\text{ave}} = -\nabla \psi$, where $\psi = (e/4m\Omega^2)|\vec{E}|^2(\vec{r})$.

Pseudopotential in a linear Paul trap. A linear Paul trap, as shown schematically in Figure 3.2, generates an approximate quadrupole potential $\psi \propto (x^2 - y^2) \cos \Omega t$, which yields an approximately harmonic pseudopotential $\psi_{\text{eff}} \propto x^2 + y^2$. Higher-order multipole potentials correspond to pseudopotentials that are anharmonic, as discussed further in Chapter V.

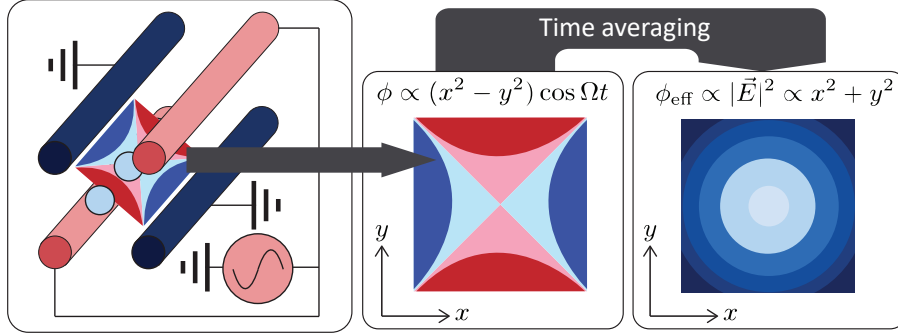


Figure 3.2. A schematic illustration of how the RF electrodes of a linear Paul trap generate a harmonic effective potential.

3.2 Cooling

To cool selected motional modes of our ions to the motional ground state (i.e. $\bar{n} \ll 1$, where \bar{n} is the expected number of phonons in a mode), we employ a three-stage cooling process: 1) Doppler cooling, 2) EIT cooling, and 3) pulsed sideband cooling (carried out on the motional sidebands of Raman transitions in our m qubits). These are all laser cooling techniques and work on the same fundamental principle. On a broad, qualitative level, cooling comes about as a result of a difference in the frequencies of photons absorbed by a spin-oscillator system (e.g. an atom or ion in a trap) and the frequencies of photons emitted by this system, with the emitted photons on average carrying away excess motional energy. The difference between these types is that Doppler cooling is a dissipative process that relies on an *average* difference in absorbed/emitted photon frequency, while resolved sideband cooling (as we do it, using a stimulated Raman process) relies on a coherent exchange of photons between a pair of well-defined optical modes followed by a dissipative step that resets the electronic state of the system without adding more motional energy. EIT cooling is dissipative, though it also relies on coherent effects to partially suppress unwanted absorption pathways. The

specific operating principles behind these techniques, their implementation in our setup, and the level of cooling we can achieve with them are described below.

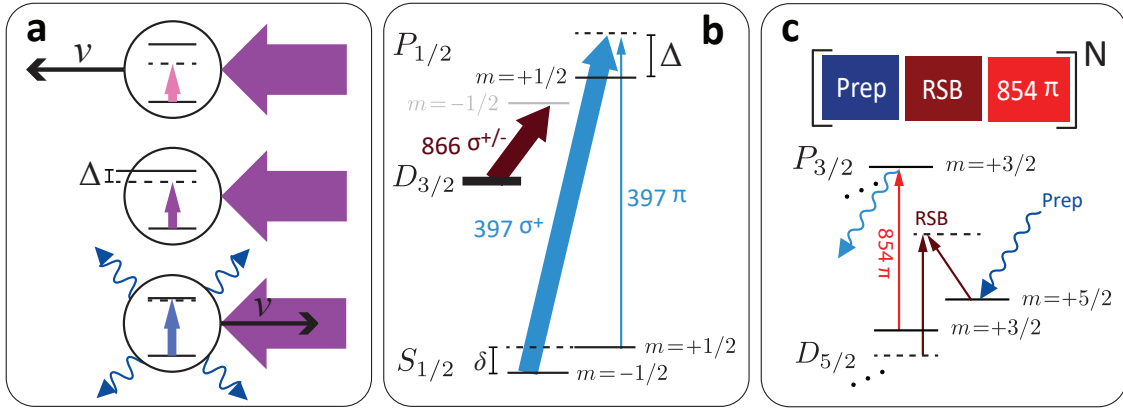


Figure 3.3. A set of diagrams showing how a) Doppler cooling, b) EIT cooling, and c) pulsed sideband cooling are implemented in our setup.

3.2.1 Doppler cooling. Doppler cooling, a longstanding technique in ion trapping (D. J. Wineland, Drullinger, and Walls (1978)) for which theoretical treatments can be found in many textbooks on atomic physics and quantum optics, relies on the Doppler effect (as suggested by the name) to reduce the energy of a bound system. In this technique, the difference in absorbed/emitted photon frequency comes from the fact that photon absorption is a motion-dependent process, while photon emission is not. This fact is illustrated schematically in Figure 3.3a. When a beam is red-detuned from the cycling transition in the lab frame, it will be closer to the transition resonance and more likely to be absorbed when the ion is moving against the beam. In this instance, the energy supplied to drive the optical transition is coming partially from the laser field and partially from the motion of the ion itself. This means that, when the ion spontaneously decays back to the ground state, the photon emitted (which will on average be less Doppler shifted than the absorbed photon, since the direction

of emission is random) will be carrying away some amount of the ion's motional energy.

Limits. Each spontaneous emission event “kicks” the ion, subtracting the momentum $\hbar\vec{k}$ of the emitted photon (where \vec{k} is the photon's associated wavevector). On average, the momentum changes from these kicks will sum to zero, but the same is not true of the energy changes $(\hbar k)^2/2m$ imparted to the ion (where k is wavevector magnitude and m is ion mass). Doppler cooling will reach equilibrium when the rate at which photon kicks add energy matches the rate at which the motionally-biased absorption process subtracts energy. Doppler cooling is a widely used technique, and full theoretical accountings of it can be found in many books on atomic physics or quantum optics, but an important result to note here is the Doppler cooling limit, or

$$k_B T = \frac{\hbar\Gamma}{4} \frac{1 + (2\Delta/\Gamma)^2}{2|\Delta|/\Gamma} \quad (3.10)$$

where k_B is Boltzmann's constant, T is the mode temperature, Γ is the linewidth of the cycling transition used for cooling, and Δ is the detuning of the cooling laser. Mode temperature T is minimized when $|\Delta| = \Gamma/2$, giving a minimum temperature $T_{\min} = \hbar\Gamma/2k_B$ limited by the natural linewidth of the cycling transition. In general, we are less interested in talking about our mode temperatures as we are in talking about average phonon numbers. Assuming a thermal distribution, this average phonon number will be $\bar{n} = 1/(e^{\hbar\omega_m/k_B T} - 1)$, and assuming optimum cooling laser detuning such that $T = T_{\min}$, this expression can be written $\bar{n} = 1/(e^{2\omega_m/\Gamma} - 1)$.

Our Doppler cooling scheme. In our setup, we carry out Doppler cooling on the 397 nm $S_{1/2} \leftrightarrow P_{1/2}$ transition. The Doppler cooling performed in experiments is two stages, each being applied for about 1 ms: The first stage employs a high-

power (50-100 μW) 397 nm beam red-detuned ≈ 500 MHz of resonance and is intended to cool ions hot enough for Doppler broadening to significantly affect the 397 nm line profile (e.g. ions ejected from the $\approx 400^\circ\text{C}$ calcium oven, or already trapped ions that have decrystallized due to e.g. motional driving or background gas collisions). The second stage employs a near-resonant (≈ 10 -15 MHz) 397 nm beam operated at lower powers (4 μW). Given an ≈ 35 MHz transition linewidth and optimized cooling beam detuning, this process should take the ~ 1.8 MHz radial modes down to average occupancies of $\bar{n} \sim 10$ and the ~ 0.6 MHz axial mode down to $\bar{n} \sim 20$. To prevent shelving in the $D_{3/2}$ level, we leave on a high-power (i.e. $\approx 150 \mu\text{W}$), $\sigma^{+/-}$ polarized 866 nm beam which optically pumps from the full $D_{3/2}$ manifold. To prevent the formation of three-level dark resonances (which could yield higher than desired sensitivity to small laser frequency drift, potentially even causing the Doppler cooling lasers to heat the ion), the 866 beam is tuned $\approx +60$ MHz from resonance.

3.2.2 EIT cooling. A limitation of Doppler cooling is the fact that, with a broad cycling transition like $S_{1/2} \rightarrow P_{1/2}$, we are unable to resolve motional sidebands (i.e. transitions which, when a photon is absorbed, add ($n \rightarrow n+1$, where n is the number of phonons in a motional state) or subtract ($n \rightarrow n-1$) a phonon from the motional state, described in greater detail in Section 3.7.1 in the context of Raman sidebands), since the width Γ of the transition will be much greater than the motional mode frequency ω_m , $\Gamma \gg \omega_m$ (in our case, by about a factor of twenty), meaning that it is impossible to drive only the phonon-subtracting sideband, a requirement for most common ground state cooling techniques.

Electromagnetically-induced transparency (EIT) cooling (Morigi, Eschner, and Keitel (2000)), a technique that has seen recent widespread adoption for near ground-state cooling, is a three-level process that relies on destructive interference

between absorption pathways to selectively suppress the carrier ($n \rightarrow n$) and phonon-adding transitions. This means that the photon-subtracting transition will dominate (and as long as the ion is in the Lamb-Dicke regime, i.e. as long as the motional wavefunction extent is very small relative to the wavelength of the cooling lasers, emission is unlikely to change the phonon number of the state). In terms of the absorption spectrum, this suppression of unwanted transitions generates narrow, asymmetrical features in the absorption profile of the cycling transition that can be used to achieve sub-Doppler cooling in a way that is fast and broadband (i.e. able to cool several closely-spaced motional modes at once). The laser setup and states used to perform EIT cooling are shown in Figure 3.3b where a strong 397 nm σ^+ -polarized “pump” beam suppresses fluorescence from a weak 397 nm π -polarized “probe” beam. In our typical setup, we operate at an average detuning of $\Delta = +150$ MHz and a relative detuning of $\delta \approx 4$ MHz (though the optimum value depends on mode frequencies and the magnitude of AC Stark shifts and occasionally needs to be recalibrated, which is done by characterizing the mode temperature using phonon-subtracting sideband of our Raman transitions, as discussed below in Section 3.2.4). With $2.7 \mu\text{W}$ of power in the pump beam and $0.15 \mu\text{W}$ in the probe, we are able to achieve cooling from the Doppler limit to $\bar{n} < 1$ on our radial modes in 1 ms.

During EIT cooling, the 866 nm laser must be on to prevent shelving in the $D_{3/2}$. As in the Doppler cooling case, care must be taken that during EIT cooling, this beam is sufficiently far detuned that it does not itself interfere with the absorption pathways used in this operation.

3.2.3 Pulsed sideband cooling. To achieve full ground state cooling (or GSC, $\bar{n} \ll 1$), we employ resolved sideband cooling (Diedrich, Bergquist, Itano, and Wineland (1989)), where the linewidths of the phonon-adding and phonon-

subtracting transitions are much smaller than the mode frequency such that these transitions can be resolved and driven separately. Specifically, we make use of our Raman sidebands (discussed in greater detail in Section 3.7.1 below) and employ a pulsed cooling scheme. We rely on the fact that we can drive our red sideband (RSB) to induce spin flips conditional on the removal of one phonon from the motional state. As shown in Figure 3.3c, a single sideband cooling (SBC) pulse consists of three steps: 1) initialization of the qubit into the $D_{5/2}$, $m = +5/2$ state (using the preparation sequence outlined in Section 3.3), 2) driving the qubit on the 976 nm Raman RSB, which will remove one phonon from the motional state if a spin flip occurs (with no flip being possible if the ion is already in the motional ground state), and 3) deshelving of the spin-flipped/cooled $m = +3/2$ population for reinitialization on the next SBC pulse (a process which will usually add no motional energy, and if it does, this energy can be taken out on the next cycle). This sequence removes up to one phonon per cycle and has allowed us to achieve ground state cooling on a single radial mode down to $\bar{n} \approx 0.05$. Parameters for the preparation and deshelving pulsed are similar to those used in state preparation and detection. Optimum drive times for the RSB depend on how cool the mode is when SBC starts. In cases where $\bar{n} \ll 1$ to start with, such that most of the motional population outside of the ground state is in the $n = 1$ state, fastest ground state cooling is achieved by driving each RSB pulse for the $n = 1$ RSB t_π . However, since RSB π -time scales like \sqrt{n} , longer drive times will overshoot the t_π for higher Fock states. In general, since EIT cooling puts us at $n \sim 0.5$, we perform the first half of RSB pulses at half the $n = 2$ RSB π -time and the next half at the full $n = 1$ π -time. We have operated with π -times of $\approx 20\text{-}40 \mu\text{s}$.

3.2.4 Cooling testing with thermometry. To measure the quality of our GSC, we use a standard motional thermometry technique. When $\bar{n} \ll 1$

(and assuming the mode is approximately in a thermal state), the “height” of the red and blue sidebands (i.e. the probability of seeing a spin flip after driving on these sidebands) after a given time can be compared, the ratio of these heights P_{RSB}/P_{RSB} giving the average phonon number \bar{n} according to $P_{RSB}/P_{RSB} = \bar{n}/(1 + \bar{n})$. On a day-to-day basis, cooling effectiveness can be monitored just by looking at the height of the red sideband after GSC.

3.2.5 Micromotion compensation. Minimizing micromotion is not cooling in a strict sense, but it falls under the broad category of reducing undesired ion motion and will be included in this section. Unlike the other methods described in this section, micromotion compensation is a entirely passive.

Micromotion. The motion of an ion in a trap can be divided into two portions: secular motion (amounting to a set of quantum harmonic oscillators with frequencies determined by the curvatures of the pseudopotential and any additional DC trapping potentials) and micromotion (driven motion parallel to and 180° out of phase with the RF electric field seen by the ion). Because micromotion is driven by the electric field generated by the trap electrodes, it will be minimized where the electric field is minimized (which, in a large, symmetrical trap, will be at the “RF null,” a spot where the electric field is zero by symmetry). Because the pseudopotential is minimized at the RF null, this will ideally be where the ion is trapped. However, stray DC electric fields (stray E-fields) in the trap (e.g. from dielectric charging by UV lasers) will tend to push the equilibrium point off of the null. “Compensating micromotion” means applying DC offsets (shim voltages) to trap electrodes to move the equilibrium back onto the null. There are three techniques we have used for calibrating these compensation voltages: 1) Minimizing dependence of ion position on trap RF voltage, 2) minimizing Doppler broadening on the 397 nm fluorescence line, and 3) minimizing the height of the micromotion

sideband of our 976 nm Raman transition. These techniques are described briefly below, and Figur 3.4 presents a comparison of these techniques and the associated geometries / directions of sensitivity. Notably, we do not compensate micromotion in the z -direction (by e.g. applying differential voltages on the two endcaps).




Figure 3.4. Schematic illustrations of techniques that we use to minimize micromotion. The offset direction each technique allows us to optimize is marked with a black double-headed arrow. The direction of the micromotion is marked with a double-headed red arrow. These techniques are **a)** minimization of position change with RF confinement change, **b)** minimization of Doppler broadening when driving with the 397π beam (blue arrow), and **c)** minimization of the micromotion sideband associated with our two-beam (red arrows) Raman Rabi drive.

Minimizing micromotion - Ion position vs. trap RF voltage. If there is a stray E-field driving the ion off of the RF null, then in principle, increasing trap RF voltage should drive the ion closer to the null, while reducing RF voltage should do the opposite. Based on this, the first technique we used for minimizing micromotion was looking on our camera at how much the ion would move in the x -direction when trap RF voltage was reduced and adjusting the compensation field E_x until the ion's position on camera no longer visibly changed with trap RF voltage. This minimization technique is relatively imprecise, but it allows us to minimize

micromotion that our lasers are not sensitive to (which is therefore also the type of micromotion that will have the least appreciable effect on experiments).

Minimizing micromotion - Doppler broadening. Micromotion parallel to our 397π beam will tend to cause Doppler broadening of the 397 nm absorption line. As a result, by sitting at some point red-detuned from the 397 nm resonance (e.g. at the expected quarter-fluorescence point given this transition’s natural linewidth of ≈ 35 MHz) and adjusting The y-direction compensation field E_y , an optimum E_y for micromotion compensation can be found by minimizing scattering seen at that 397π tuning.

Minimizing micromotion - Micromotion sideband. By sitting on our red micromotion sideband (located at our Raman Rabi carrier frequency plus our ≈ 14.5 MHz trap RF frequency), we can detect whether or not there is appreciable micromotion parallel to the wave-vector difference $\Delta\vec{k}$ of the Raman beams, since if there is no micromotion (i.e. no energy in the classical “micromotion mode”), we will not be able to drive this transition. An optimum E_y for micromotion compensation here will be the one which minimizes amplitude/frequency of Rabi flopping on the red micromotion sideband.

3.3 m qubit preparation

By default, an ion will tend to be found in its electronic ground state (e.g. the $S_{1/2}$ manifold). We encode our qubits in the $m = +5/2$ and $m = +3/2$ states of the metastable $D_{5/2}$ manifold. While shelving states in $S_{1/2}$ to states in $D_{5/2}$ can be done using narrow-linewidth lasers addressing the associated quadrupole transition, in all of the experiments described in the rest of this thesis, we initialize our qubits with an optical pumping scheme described below.

3.3.1 Optical pumping. We initialize our qubit in the $m = +5/2$ level of the $D_{5/2}$ manifold. To do so, we cycle between a set of three laser pulses,

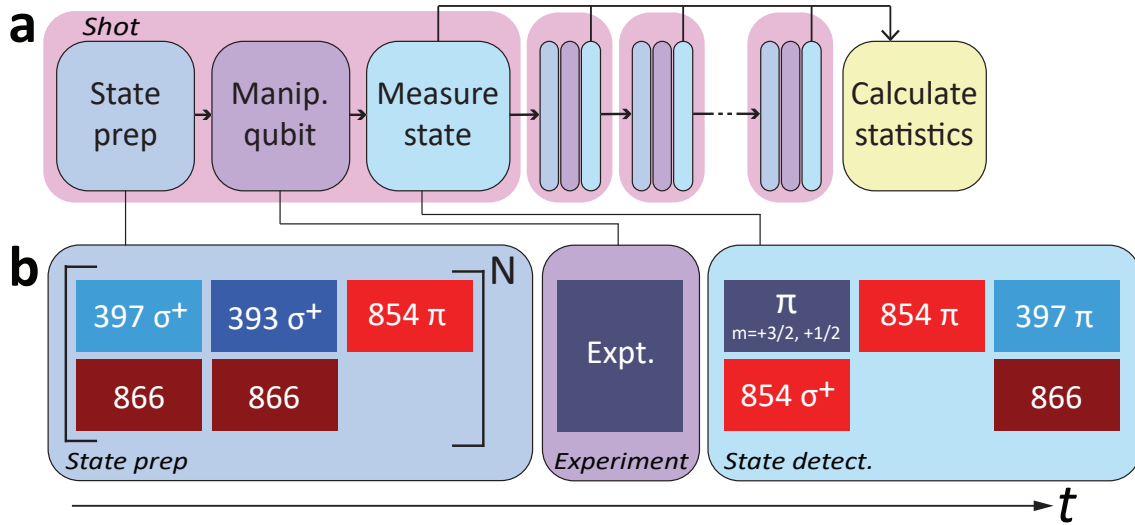


Figure 3.5. **a)** The generic sequence for characterizing the effect of some operation (“Expt.”) on some known state, performing repeated experimental “shots” and calculating the probability of different outcomes. **b)** The specific pulse sequences of each experimental step (not showing cooling, which either happens before or is interleaved with state preparation).

shown (along with the rest of a generic, schematic experiment in m qubits, excluding cooling) in the leftmost block of Figure 3.5b. Pulsing a σ^+ -polarized, 397 nm ($397 \sigma^+$) beam shelves the ion in the $S_{1/2}, m = +1/2$ state. Pulsing a 393 σ^+ beam (resonant with the $S_{1/2} \leftrightarrow P_{3/2}$ transition) shelves the ion in the $D_{5/2}, m \geq -1/2$ states. Given the $P_{3/2}$ to $D_{5/2}$ branching ratios, we expect these two pulses to leave approximately 67% of the ion population in the $m = +5/2$ state. By pulsing an 854 π beam, we deshelve all the remaining population in $D_{5/2}, m < +5/2$ back into the ground state. This pulse sequence is repeated several times (typically $N=7$ times in our experiments), with preparation fidelity bottoming out as per-cycle error rates (discussed below) balance out fidelity gains.

3.3.2 Preparation fidelity and limits. There are three limiting factors in our preparation fidelity: 1) accidental shelving of population in $D_{5/2}, m = -5/2$, 2) deshelving of $m = +5/2$ population by polarization impurities in the

854 π beam, and 3) the likelihood of natural decay from $D_{5/2}$ during one repeat of the optical pumping sequence. These errors are roughly quantified below.

Accidental shelving in $m = -5/2$. This error comes about as a result of population that ends up in $m = -1/2$ during shelving by the 393 σ^+ beam being scattered twice by the 854 π beam, first into $m = -3/2$ and then into $m = -5/2$, where it will no longer couple to 854 π or any of the other beams in state preparation, only being deshelled by pulsing a resonant 854 $\sigma^{+/-}$ beam at the end of an experiment shot. Assuming pessimistically that the 393 σ^+ beam is always driving out of the $S_{1/2}$, $m = -1/2$ state, the probability of shelving in $m = -5/2$ on a given repeat of the optical pumping sequence is the product of three $P_{3/2}$ to $D_{5/2}$ branching ratios: $m = +1/2 \rightarrow m = -1/2$, $m = -1/2 \rightarrow m = -3/2$, and $m = -3/2 \rightarrow m = -5/2$. Pessimistically estimating all of these branching ratios to be $\approx 5\%$, the probability of shelving in $m = -5/2$ on a given repeat of the shelving sequence is then $< 0.02\%$.

Deshelving by 854 π polarization impurities. Ideally, the 854 π beam would only deshelve states between $m = -3/2$ and $m = +3/2$. However, $\sigma^{+/-}$ polarization impurities in this beam (caused in part by factors that are functionally unavoidable, like stress-induced birefringence in our viewports) will tend to deshelve the $m = +5/2$ state. Once the relative populations in the $m = +5/2$ state and the rest of the manifold are such that equal $m = +5/2$ and $m < +5/2$ populations are deshelled by an 854 π pulse, no further preparation fidelity is gained. While there is not a simple expression for this population ratio in terms of the polarization purity of the 854 π beam, a simple order-of-magnitude proxy is the ratio of deshelving rates for the $m = +5/2$ and $m = +1/2$ states. In our setup, we are able to achieve a ratio of $\approx 1:5000$.

Decay due to natural lifetime. During the $40\ \mu\text{s}$ $854\ \pi$ pulses at the end of one optical pumping cycle, there is some chance that the qubit will decay back to the ground state on its own. This effect is only important in the last optical pumping cycle, since otherwise the decayed ion will most likely be reshelved on the next pumping cycle. The probability of a decay event is simply the length of the $854\ \pi$ pulse divided by the natural lifetime $\tau=1.16\ \text{s}$ (Benhelm (2008)). This amounts to a decay probability of $<0.0034\%$.

Heralding decay to reduce preparation error. While preparation errors from shelving in the $m = -5/2$ state can only be resolved with a near-resonant σ^+ -polarized $854\ \text{nm}$ beam, which we do not use, other two error sources discussed above can be mitigated by heralding shelving by checking for fluorescence on the $397\ \text{nm}$ transition after performing state preparation. If fluorescence is observed, then the ion is not shelved, and the state preparation process can be repeated. This concept of using $397\ \text{nm}$ fluorescence to herald deshelling is elaborated on more thoroughly in Chapter IV, where it is used to detect leakage errors that occur during qubit operations.

Given the procedures and limitations outlined above, we in practice measure an $\approx 99.8\%$ $m = +5/2$ preparation fidelity in a single qubit.

3.3.3 Isolating qubit states. The $1.565\ \text{G}$ field defining our quantization axis lifts the degeneracy of the spin states in the $D_{5/2}$ manifold of $^{40}\text{Ca}^+$, giving a Zeeman splitting between adjacent spin states (e.g. between $m = +5/2$ and $m = +3/2$) of about $2.63\ \text{MHz}$. The fact that this splitting is uniform means that we cannot selectively drive between one pair of Zeeman sublevels (so that a drive between the $m = +5/2$ and $m = +3/2$ states will also drive the $m = +3/2$ to $m = +1/2$ transition). To break this symmetry, we apply a beam that generates a large AC Stark shift (also called a light shift or LS, the

physics of which are discussed below in Section 3.5.5) on the $m < +3/2$ Zeeman sublevels and does not couple to the qubit sublevels (Sherman et al. (2013)). Specifically, we apply a σ^+ -polarized beam -21 GHz detuned from the $D_{5/2} \leftrightarrow P_{3/2}$ transition (the “854 LS” beam). With 8 mW of power (supplied by the fiberized diode laser described in Section 2.5.2) focused to an $\approx 70 \mu\text{m}$ beam diameter, we observe an ≈ 0.50 MHz AC Stark shift on the $m = +1/2$ state, allowing us to perform coherent operations on our qubit with minimal leakage for Rabi frequencies ~ 20 kHz (discussed further below). Examples of Rabi flopping (driven with an RF magnetic field) with and without the 854 LS applied are shown in Figure 3.6.




Figure 3.6. Sample Rabi flopping, driven with a resonant RF magnetic field, showing two-level sinusoidal dynamics when the 854 LS beam is on (blue) and precession through the $D_{5/2}$ manifold before returning to the starting state when the 854 LS beam is off (red).

Caveats of qubit isolation via AC Stark shifts. We operate the 854 LS beam at moderately large detunings (10s of GHz) despite the fact that we could generate larger light shifts by tuning closer to resonance. This is because ~ 0.1 - 0.2% polarization impurities in the 854 LS beam (similar to those discussed for the 854 π above) allow this beam to couple to the qubit states. Notably, this caveat is not fundamental, in that 1) if polarization purity could be made perfect, this coupling would not be there and 2) the photon scattering caused by this coupling can be

made arbitrarily low by increasing detuning at the cost of increasing power. In practice however, the 854 LS generates a relative light shift between the qubit states (ordinarily ≈ 1 kHz at our operating settings, when polarization is made maximally pure) that makes the qubit coherence sensitive to 854 LS power stability (even if to a small degree). More importantly for certain experiments, this coupling means that the 854 LS can deshelve the qubit states. At the maximally pure polarization we can achieve, we generally see a deshelving rate on both qubit states with a time constant ~ 0.4 s, not accounting for decay due to the natural lifetime of the qubit.

3.4 State detection

In an experiment, after the ion is cooled and a qubit is initialized as described above, we perform some set of operations that in general changes the qubit state (i.e. that changes the qubit from being purely in the $m = +5/2$ state to being in some superposition/mixture of $m = +5/2$ and $m = +3/2$). At the end of these operations, we want to perform a projective measurement in the qubit basis. This measurement / state detection process has two stages: 1) state-selective deshelving from the $D_{5/2}$ manifold and 2) driving/collection of fluorescence on the 397 nm cycling transition to check for deshelving. These steps (with the associated pulse sequence illustrated schematically in Figure 3.5) are described in greater detail below.

3.4.1 State-selective deshelving. In our state-selective deshelving procedure, we aim to deshelve the $m = +3/2$ state from the $D_{5/2}$ manifold while leaving the $m = +5/2$ state untouched. In principle, this could be done with a single pulse of our 854 π beam. However, because of the branching ratios from the $P_{3/2}$, $m = +3/2$ state, if we drove directly out of the $m = +3/2$ level with π -polarized light, $\approx 3.9\%$ of the time, the ion would scatter into the $m = +5/2$ level, giving us a readout error. Since this branching ratio is non-negligible, before

deshelving with the 854π , we perform a π -pulse between the $m = +3/2$ and $m = +1/2$ levels. The upshot of this is that, when driving this state with 854π , the probability of optically pumping into the $m = +5/2$ state goes from being a single branching ratio, $\approx 3.9\%$, to a product of this branching ratio and the $P_{3/2, m = +1/2} \rightarrow D_{5/2, m = +3/2}$ branching ratio.

3.4.2 Distinguishing shelved and unshelved ions - Fluorescence.

To detect deshelving, we drive the 397 nm cycling transition and look for fluorescence. If an ion is deshelved, it will scatter 397 nm photons. Of these photons, $\approx 1.2\%$ will generate a click on a PMT which can be detected as a digital input in our experiment control system. If an ion is still shelved, it will not scatter 397 nm photons, and the PMT will only record some background of photons scattered elsewhere in the system. Because this “bright count rate” (r_b , typically $\sim 60,000$ counts per second per ion) and “dark count rate” (r_d , typically ~ 2000 counts per second) are constant over a detection period, the associated distributions of photon counts observed in this detection period are described by a pair of Poisson distributions $P(n) = (rt)^n e^{-rt} / n!$ where $P(n)$ is the probability of counting n photons in the cycling transition drive time t and where r is the count rate. On the basis of these distributions, we can decide on a photon count threshold for classifying an ion as deshelved ($m=+3/2$) or shelved ($m=+5/2$). The optimum count threshold will be the one for which the risk of state misclassification is the same for both states. A sample set of photon count distributions are shown in Figure 3.7, illustrating how photon counts can be used to clearly distinguish states, including for multiple ions, as discussed below.

Fluorescence detection with multiple ions. The same principle of count thresholding applies for identifying the state of a pair of ions, with the important

note that, if ion 1 is shelved and ion 2 is deshelled, we cannot distinguish this case from the the opposite, since the photon count rate for both is the same.




Figure 3.7. A sample count histogram for two-ion fluorescence detection, with the “zero bright” (blue), “one bright” (red), and “two bright” (green) Poisson distributions labelled above.

3.5 Coherent control - Spin

The state of a single, unentangled qubit (a qubit whose state can in general be written $|\psi\rangle = c_{\uparrow}|\uparrow\rangle + c_{\downarrow}e^{i\phi}|\downarrow\rangle$, where c_{\uparrow} and c_{\downarrow} are real) can be represented using a Bloch sphere (Feynman, Vernon Jr, and Hellwarth (1957)). In this representation, the populations c_{\uparrow}^2 and c_{\downarrow}^2 are encoded in the polar angle (with the poles representing the spin eigenstates) while the phase ϕ between the states is encoded in the azimuthal angle around the sphere.

A typical quantum information experiment will involve manipulating the state of an individual spin without entangling this state with other modes in the system (i.e. other spins, motional modes). These operations can be thought of as translations of the spin on the surface of the Bloch sphere. Rabi flopping, population exchange driven by the coupling of qubit states through a (near-)resonant external field, is discussed at length below (though treatments of Rabi

flopping can be found in many textbooks on quantum mechanics or quantum optics). For completeness, AC Stark shifts are also briefly touched on as a form of spin operation, since these will change the phase ϕ between the states, but they are not generally used as such.

A general Hamiltonian that will cause Rabi flopping can be written as follows, the first term being for the free evolution of the qubit and the second term being the external drive:

$$H = \frac{\hbar\omega_0}{2}\sigma_z + \hbar\Omega \cos(\omega t)\sigma_x \quad (3.11)$$

where Ω is the Rabi frequency, ω is the frequency of the drive field, and ω_0 is the qubit resonance (typically zero in the absence of calibration errors). On the Bloch sphere, a Rabi drive amounts to rotation at a fixed rate Ω about an axis defined by the basis of the Rabi drive (which in general is some combination of σ_x and σ_y and which can be changed physically by changing the phase of the drive field). We have two physical mechanisms for driving Rabi flopping in our qubits: An oscillating magnetic dipole that causes the spin of the ion to precess and a three-level Raman process that flips the qubit state by exchanging photons between a pair of optical modes. The details of these processes and how they are implemented in our system are discussed in greater detail below. Also discussed are the “standard” operations implemented with Rabi flopping (π - and $\pi/2$ -pulses), coherent errors associated with these operations, and techniques for mitigating these coherent errors.

3.5.1 RF magnetic Rabi flopping. A straightforward way of changing the state of a Zeeman qubit is to induce precession of the ion’s magnetic moment using an oscillating magnetic field. Considering a two-level, spin-1/2 system for which the magnetic moment can be aligned with or against the quantization magnetic field B_z , we can write the unperturbed Hamiltonian of this

system as $H = \hbar\mu B_z \sigma_z / 2 = \hbar\omega_z \sigma_z / 2$, where μ is the magnetic moment of the ion and ω_z is the qubit splitting. If we then apply an AC magnetic field oscillating at a frequency ω perpendicular to the quantization field (e.g. in the x -direction), then the Hamiltonian can be written

$$H = \frac{\hbar\mu B_z}{2} \sigma_z + \frac{\hbar\mu B_x}{2} \cos \omega t \sigma_x. \quad (3.12)$$

Rewriting this Hamiltonian in the interaction representation (i.e. switching from considering the dynamics of the coefficients c_\uparrow and c_\downarrow of a general state $|\psi\rangle = c_\downarrow |\downarrow\rangle + c_\uparrow |\uparrow\rangle$ to considering the dynamics of the coefficients in a frame co-rotating with the freely evolving qubit states, $\bar{c}_n = c_n e^{i\omega_n t}$), we get an “interaction Hamiltonian” \bar{H} (defined as the Hamiltonian governing the time evolution of these rotating coefficients via Schroedinger’s equation, $i\hbar(\partial/\partial t)[\bar{c}_\uparrow, \bar{c}_\downarrow]^T = \bar{H}[\bar{c}_\uparrow, \bar{c}_\downarrow]^T$), which can in this case be written

$$\bar{H} = \hbar\Omega \cos \omega t \begin{bmatrix} 0 & e^{-i\omega_0 t} \\ e^{i\omega_0 t} & 0 \end{bmatrix} \quad (3.13)$$

(rewriting $\mu B_x / 2$ as a Rabi frequency Ω , which is taken to be real here but which can in general be complex, representing the phase of the drive). The time-dependence of the drive (via $\cos \omega t$) can be multiplied with the inherently time-dependent matrix elements of \bar{H} to give matrix elements with two terms each: a “fast-rotating” term (oscillating at the sum of frequencies $\omega + \omega_0$) and a “slow-rotating” term (oscillating at the difference of frequencies $\delta = \omega - \omega_0$), i.e.

$$\bar{H} = \frac{\hbar\Omega}{2} \begin{bmatrix} 0 & e^{i(\omega-\omega_0)t} + e^{-i(\omega+\omega_0)t} \\ e^{i(\omega+\omega_0)t} + e^{-i(\omega-\omega_0)t} & 0 \end{bmatrix} \quad (3.14)$$

To simplify this picture, an important approximation can be applied. So long as $\omega + \omega_0 \gg |\delta|$ and $\omega + \omega_0 \gg |\Omega|$, the contribution of the fast-rotating term to the dynamics will be negligible. On an intuitive level, this is because in this regime,

the fast-rotating drive will go out of phase with the qubit and start driving the opposite way before it has time to drive any substantial population between the qubit states. If we drop this fast-rotating term (i.e. if we apply the rotating-wave approximation, or RWA), this yields

$$\bar{H} \stackrel{RWA}{\approx} \frac{\hbar\Omega}{2} \begin{bmatrix} 0 & e^{i\delta t} \\ e^{-i\delta t} & 0 \end{bmatrix}. \quad (3.15)$$

Given this interaction Hamiltonian, the dynamics of \bar{c}_\downarrow and \bar{c}_\uparrow can be considered in a straightforward way. In the limiting case where the magnetic drive is on resonance (i.e. $\delta = 0$), via Schroedinger's equation, the equations of motion governing the system are simply

$$\frac{\partial \bar{c}_\uparrow}{\partial t} = -i \frac{\Omega}{2} \bar{c}_\downarrow \quad (3.16)$$

$$\frac{\partial \bar{c}_\downarrow}{\partial t} = -i \frac{\Omega}{2} \bar{c}_\uparrow, \quad (3.17)$$

which, given the initial conditions $\bar{c}_\uparrow(t=0) = 0$, $\bar{c}_\downarrow(t=0) = 1$, have the solutions

$$\bar{c}_\downarrow(t) = \cos \frac{\Omega}{2}t \rightarrow P(\downarrow) = \cos^2 \frac{\Omega}{2}t \quad (3.18)$$

$$\bar{c}_\uparrow(t) = -i \sin \frac{\Omega}{2}t. \rightarrow P(\uparrow) = \sin^2 \frac{\Omega}{2}t \quad (3.19)$$

where $P(n)$ corresponds to the probability of finding the qubit in state n at a given time. These solutions correspond to state population moving back and forth between $|\downarrow\rangle$ and $|\uparrow\rangle$ at a frequency Ω . This behavior is Rabi flopping. On-resonance, Rabi flopping can move a qubit entirely from one state into another. It can be shown however that if $\delta \neq 0$, the probability of seeing a spin flip will only be as high as $\Omega^2/(\Omega^2 + \delta^2)$ (as will be touched on in Section 3.5.3), with these reduced-contrast oscillations having a higher Rabi frequency than the resonant case, the “generalized Rabi frequency” being $\Omega' = \sqrt{\Omega^2 + \delta^2}$.

In our setup, we apply an RF magnetic field an angle to the quantization axis, as shown in Figure 3.8. We generate this magnetic field using a voltage

generated by one of our DDSes. Under ordinary operating conditions, we drive spin flips with RF at a Rabi frequency of ≈ 10 kHz, with headroom to increase this number by about a factor of four. We limit our Rabi frequency to avoid unwanted off-resonant driving outside of the qubit manifold, since in our case, rather than driving a two-level spin system, we are driving between two levels of a larger manifold that have been separated from the rest using an AC Stark shift. Without this Stark shift, we are driving population through a uniformly spaced ladder of six states. The dynamics associated with this driving differ, as the ion will precess through the entire manifold before returning to its starting state, as shown in Figure 3.6.




Figure 3.8. Schematic illustration of our RF magnetic Rabi drive, showing trap electrodes (pink for trap RF, blue for Rabi drive current), the quantization axis magnetic field \vec{B} (in the xz plane), and the magnetic Rabi drive \vec{B}_{drive} (in the xy plane).

The discussion above describes magnetically-driven Rabi flopping between a pair of spin states, but the same general derivation can apply to any two-level system driven by a single-mode field. In addition, many of the concepts described in brief above (rotating frames and interaction Hamiltonians, in addition to the

RWA and other resonance-like conditions) will recur throughout this section and be referenced more shorthandedly.

3.5.2 Raman Rabi flopping. While AC magnetic fields are a straightforward way of driving Zeeman/hyperfine qubits, they have certain limitations: 1) AC magnetic drives are hard to target to specific ions compared to laser beams, which allow easy individual addressing, and 2) as described below, generating spin-motion coupling requires a substantial intensity gradient over an ion’s motional wavefunction (i.e. the Rabi frequency of the drive must be correlated with the ion’s position), which is difficult to achieve with the long wavelengths associated with RF/microwave frequencies. While much work has been done to circumvent both limitations (on both individual addressing and on generating spin-motion coupling with AC magnetic fields, e.g. Weber et al. (2024) or Löschnauer et al. (2024)), the standard technique for achieving these ends in Zeeman/hyperfine qubits is to use two-photon Raman processes. These processes allow qubits encoded in spin states in the same manifold to be controlled using optical fields. The actual principle behind spin-motion coupling (which is also in principle applicable to magnetic and resonant, single-mode optical drives) is explained in Section 3.7. To introduce Raman transitions, the driving of three-level systems and how this is achieved in our system are discussed below.

Three-level (Λ) systems. Two-photon Raman processes can occur in Λ systems, or three-level systems where two of the levels are far closer to each other than they are to the third level (e.g. in trapped ions, a pair of low-lying Zeeman levels and another spin state above separated from these levels by an optical transition). Such a system, and the associated drives for carrying out a Raman process, are shown in Figure 3.9a. In the discussion below, the bottom levels are labelled $|\uparrow\rangle$ and $|\downarrow\rangle$, and the top (excited) level is labelled $|e\rangle$. The basic principle

of Raman processes is that population transfer between $|\uparrow\rangle$ and $|\downarrow\rangle$ can be achieved without directly driving between these levels. Instead, if the $|\uparrow\rangle \leftrightarrow |e\rangle$ and $|\downarrow\rangle \leftrightarrow |e\rangle$ transitions are driven simultaneously, population transfer can take place using $|e\rangle$ as an “intermediate” state, and in the limit that the drive is very far-detuned ($\Delta \gg \Omega_{\uparrow}, \Omega_{\downarrow}$) the occupation probability for $|e\rangle$ approaches zero, and the dynamics effectively become those of a two-level system, with the relationships between the three-level and effective two-level system represented schematically in Figure 3.9b. Notably, for the optical dipole transitions we use in our work, the Rabi frequencies Ω_{\downarrow} and Ω_{\uparrow} are proportional to the amplitude E_{\uparrow} of the field used to drive them, e.g. $|\Omega_{\uparrow}| = |\mu_{\uparrow e} E_{\uparrow}|/\hbar$, where $\mu_{\uparrow e}$ is the dipole matrix element for the associated optical transition. This means that the Rabi frequencies of individual transitions are proportional to \sqrt{I} , where I is the beam intensity, but that the combined effective Rabi frequency is proportional to $\sqrt{I_{\downarrow} I_{\uparrow}}$, meaning that e.g. doubling the power of the Raman beams used to drive the effective two-level system will double the effective Rabi frequency. More thorough discussions of three-level systems in general can be found in any textbook on quantum optics.

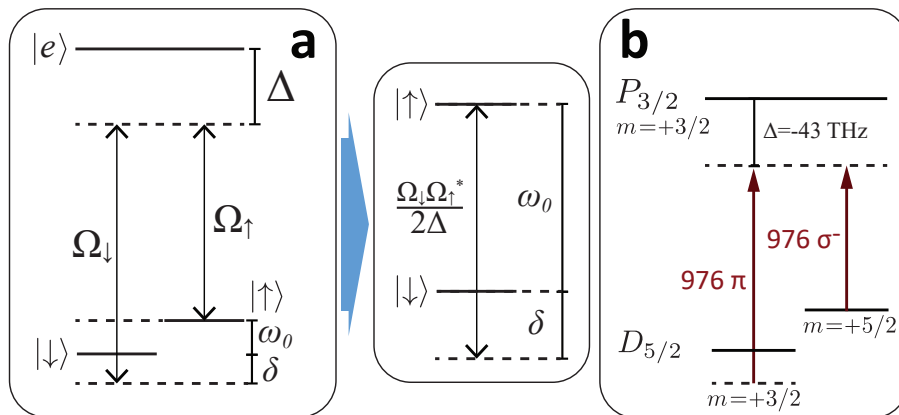


Figure 3.9. Schematic illustration of the dynamics of a Λ system, showing **a)** the relationship between the three-level dynamics of the actual system and the effective two-level dynamics and **b)** the three-level system we use in our experiments.

Our Λ system. The Λ system we use, represented schematically in Figure 3.9b, includes our two qubit levels and the $m = +3/2$ level of the $P_{3/2}$ manifold. These levels are coupled to each other using two optical modes (at 976 nm, tuned ≈ -43 THz from the 854 nm $D_{5/2} \rightarrow P_{3/2}$ resonance) at different polarizations and tunings, and we can drive resonant “Raman Rabi flopping” when the π -polarized mode is tuned ω_0 closer to the optical resonance than the σ^- -polarized one, where ω_0 is the qubit frequency. As shown in Figure 2.12, we apply two Raman beams (≈ 200 mW each) with different polarization components to our ions, giving us two options for how we can drive Raman Rabi flopping. We can drive Raman Rabi flopping by tuning the two beams such that the σ^- -polarized beam and the π -polarized component of the other beam drive the transition. Alternatively, if this second beam has both of the polarization components needed, two tones can be applied to it such that a set of polarization/frequency modes exist in the same beam, two of which can couple to drive Raman Rabi flopping. These sorts of single-beam operations end up throwing away beam power (since most of the modes generated do not contribute to the dynamics), but are used in cases where it is advantageous being able to drive Raman Rabi flopping without worrying about AC Stark shifts, since these can be cancelled out with a single beam (as discussed below in Section 3.5.5). With these systems, we can achieve Raman Rabi frequencies of 100s of kHz.

3.5.3 Pulse types and coherent errors. Rabi flopping is a continuous process. However, certain drive pulse times / rotation angles are used frequently in QIP and are worth noting. A Rabi pulse that tends to rotate a state π radians about the drive axis (e.g. one that takes the qubit from $|\uparrow\rangle$ to $|\downarrow\rangle$) is called a π -pulse, while one that rotates a state $\pi/2$ radians (e.g. one that takes the qubit from $|\uparrow\rangle$ to $(|\uparrow\rangle + |\downarrow\rangle)/\sqrt{2}$) is called a $\pi/2$ -pulse.

Calibrating pulse times. The drive times associated with the pulses described above (called the π - and $\pi/2$ -times, or t_π and $t_{\pi/2}$ respectively) can be calibrated by measuring the Rabi frequency Ω and calculating $t_\pi = \pi/\Omega$, $t_{\pi/2} = \pi/2\Omega$. In practice however, systematic effects (e.g. pulse shaping that does not scale with pulse length, AOM rise times) often mean that the Ω will not be constant over the first Rabi flop. Because of this, calibrating π -times is typically a matter of initializing the qubit in $|\uparrow\rangle$ and measuring the pulse time that gives maximum transfer to $|\downarrow\rangle$, while calibrating $\pi/2$ -times can be done by chaining $\pi/2$ -pulses, with π pulses in between them to counteract the effects of phase offsets/drive detuning.

Coherent errors. Rabi flopping experiences two types of coherent error (i.e. errors that are consistent between shots of an experiment): detuning error and timing error. Detuning error ($\delta \neq 0$, in terms of the Rabi drive Hamiltonian), caused by uncalibrated offsets to the qubit frequency which come either from drift or from systematic effects (e.g. AC Stark shifts) that have not been otherwise accounted for. Timing error is a miscalibration of π - and/or $\pi/2$ -times, leading to consistent over- or under-rotation (generally labelled ϵ) on the Bloch sphere. Timing errors can be caused by slow drifts in control field strength (e.g. changes in Raman beam pointing), which can be addressed by simply recalibrating $\pi(/2)$ -times. A more “fundamental” timing error can arise with Raman Rabi flopping in two ions. Specifically, if each ion sees a different Raman beam intensity, they will have different associated π -times. The effects of such errors can be greatly mitigated by using composite pulse sequences, as described below.

3.5.4 Composite pulses. In general, a single-spin operation (i.e. an arbitrary displacement of the state vector on the surface of the Bloch sphere) can be made more robust against the coherent errors described above by converting this operation from a single Rabi pulse to a chain of Rabi pulses, with the trade-off

that doing this conversion increases operation time (sometimes substantially) and introduces sensitivity to the relative phases of pulses in the chain (which can be difficult to control precisely due to e.g. AC Stark shifts). In our experimental work, we investigated two types of “compositing”: BB1 composite pulses (which we did not end up using in other experiments) and composite Ramsey sequences (which we use routinely).

BB1 composite pulses. Following Mount et al. (2015), we investigated using BB1 composite pulses in our setup. A general BB1 sequence takes a single *targeted* rotation θ_t about rotation axis ϕ_t and converts it into a sequence of four rotations about different axes, in the process converting a rotation error ϵ to an error $O(\epsilon^2)$. In the language of Mount et al., a target rotation $R(\theta_t, \phi_t)$ is converted according to

$$R(\theta_t, \phi_t) \rightarrow R(\theta_t, \phi_t)R(\pi, \phi_t + \phi_{BB1})R(2\pi, \phi_t + 3\phi_{BB1})R(\pi, \phi_t + \phi_{BB1}), \quad (3.20)$$

$$\phi_{BB1} = \cos^{-1} \left(\frac{-\theta_t}{4\pi} \right). \quad (3.21)$$

Using these BB1 sequences, we were able to demonstrate high robustness to timing errors for both our π - and $\pi/2$ -pulses, using both RF B-field and Raman Rabi drives. The results of these tests and associated pulse sequences are shown in Figure 3.10.

While we were able to greatly reduce sensitivity to timing errors using BB1 sequences, the fact that we need to drive the equivalent of four π -pulses on top of the pulse we are intending to do means that we are increasing our pulse times by a factor of five(nine) for a $\pi/2$ -pulse. Given that, in the experiments described in future chapters, we generally operate at $t_\pi \approx 20 \mu\text{s}$, these BB1 sequences end up being $\approx 100 \mu\text{s}$, and given the sensitivity of these sequences to relative pulse phases coupled with the large amount of low-frequency spin-phase noise in our system (described further later in this chapter), it was not obvious that BB1 sequences

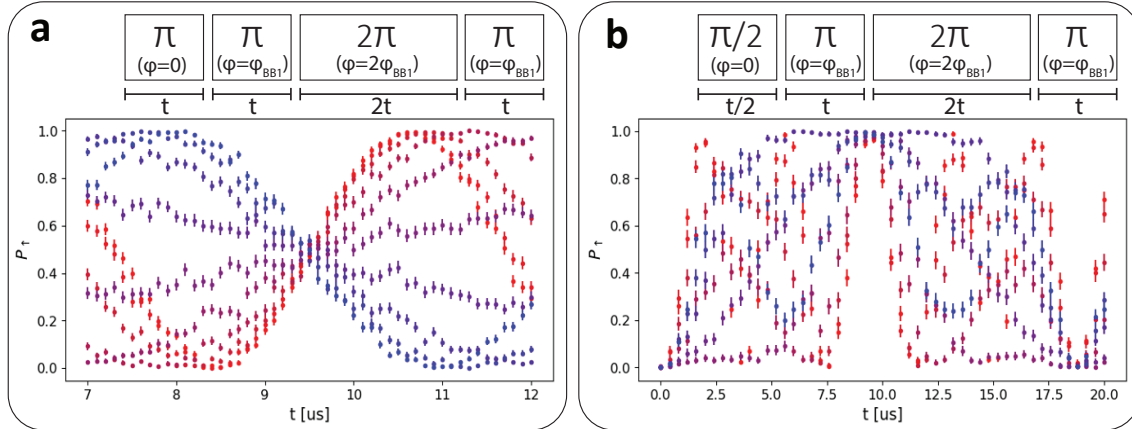


Figure 3.10. The results of experiments testing composite **a)** $\pi/2$ and **b)** π pulses. Different colors denote time scans at different ϕ_{BB1} , showing that a value of ϕ_{BB1} can be selected such that the output of the pulse sequence is made insensitive to change in pulse times.

would offer appreciable advantage compared to other approaches to addressing timing error.

Composite Ramsey sequences. Many of our experiments are effectively Ramsey sequences (a more complete description of which is offered below in Section 3.9), and we are interested in ensuring that the contrast of these sequences is as high as possible, particularly for experiments like our two-qubit entangling gates (see Chapter IV). To that end, we want to make these sequences relatively robust to coherent errors. In a Ramsey sequence with spin echoes, this can be done by selecting pulse phases so as to cancel out correlated rotation errors. For example, we can consider a Ramsey sequence with a single spin echo pulse, with timing errors such that each $\pi/2$ -pulse has an ϵ rotation error and the spin echo π -pulse has a 2ϵ error. In this situation, if the spin echo is in-phase with the initial $\pi/2$ -pulse, then the $|\uparrow\rangle$ and $|\downarrow\rangle$ fringes will both see rotation errors $O(\epsilon)$. On the other hand, if the phase of the spin echo is offset $\pi/2$ from the initial $\pi/2$ -pulse (such that, if the drive for the $\pi/2$ -pulse is in σ_x , the drive for the spin echo is

in σ_y), then the $|\downarrow\rangle$ fringe will see an error $O(\epsilon^2)$. Effectively then, this fringe is acting as a composite π -pulse. In terms of average rotation error on both fringes, the sequence where the spin echo pulse is in the σ_y basis sees an error a factor of two smaller (1.5ϵ , ignoring contributions to the error that go like ϵ^2 or higher) than the sequence where the spin echo is in phase with the initial $\pi/2$ -pulse.

Compared to BB1 sequences, composite Ramsey sequences have the advantage that they add no additional overhead to experiments already using Ramsey sequences with spin echoes (since changing the phase of a Rabi drive is only a matter of changing a phase on a DDS). However, the improvement they offer is less general in the sense that e.g. a Ramsey sequence performed with two BB1 $\pi/2$ -pulses should see $O(\epsilon^2)$ rotation errors on both spin fringes rather than just one. In addition, because Ramsey sequences are in general operated with delays between the pulses, there is more time for phase offsets to accumulate (i.e. from detuning errors or from AC Stark shifts) that can undermine the performance of these sequences. In our two-qubit entangling gate for example, our ability to use a composite Ramsey sequence was compromised by large (~ 1 kHz, applied for $\sim 200 \mu\text{s}$), drifty differential AC Stark shifts between the two ions making it impossible to drive e.g. a $\sigma_x\sigma_y\sigma_x$ sequence on both ions. Nonetheless, because there is no associated technical overhead, when running Ramsey sequences with spin echoes, we in general use these sorts of alternating drive phases to improve fringe contrast.

3.5.5 AC Stark shifts. Driving a two-level system far off-resonance (i.e. at a detuning $\Delta \gg \Omega$, where Ω is the Rabi frequency of the drive) will drive Rabi flopping at vanishingly small amplitudes. However, far-detuned fields still perturb the system, rotating the associated eigenstates (with the new eigenstates being called “dressed states”) and shifting the energy levels (an AC Stark shift).

By analyzing Rabi flopping in the field interaction representation (analogous to the interaction representation used to derive Rabi flopping in Section 3.5.1, but where the frame is co-rotating with the drive field rather than the unperturbed qubit states), it can be shown (and is shown in many standard textbooks on quantum optics, e.g. Steck (2007)) that the rotation angle θ of the eigenstates can be written $2\theta = \text{atan}\Omega/\Delta \approx \Omega/\Delta$ (for very large Δ), while the energy shifts are $\Delta E = \pm\hbar|\Omega|^2/4\Delta$ on the upper and lower levels respectively. (Notably, this means that when Δ is positive, such that the transition is being driven at a frequency higher than resonance, the levels are driven closer together, while the opposite is true when Δ is negative). The rotation angle rapidly approaches zero as Δ grows, such that the eigenstates of the system are unaffected by the drive.

The detuned beams we apply to our ions (e.g. the 976 nm beams) generate AC Stark shifts on the scale of 10s to 100s of kHz on the levels we use in operations. These shifts are extremely small relative to the optical transitions they are driving, which sit at 100s of THz. However, on the few MHz scale of our Zeeman splittings, these shifts are able to meaningfully affect dynamics in the $D_{5/2}$ manifold (as they do when e.g. the 854 LS beam is used to ensure two-level dynamics by generating an AC Stark shift on the $m = +1/2$ level).

Measurement of AC Stark shifts (usually via Rabi spectroscopy probing differential AC Stark shifts between levels in the $D_{5/2}$ manifold), along with knowledge of beam polarization and detuning, gives us a method of determining beam intensities via the relations above. The ability to measure (and attempt to maximize) AC Stark shifts also serves as a useful tool for adjusting beam pointing.

3.6 Coherent control - Motion

Each ion in a chain brings with it three harmonic oscillator modes. The state of each of these modes can be represented in the energy basis as an infinite

set of evenly-spaced Fock states. In other words, the state of the motional mode i can be written as

$$|\psi_i\rangle = \sum_{n=0}^{\infty} C_n e^{in\omega_i} |n\rangle \quad (3.22)$$

where ω_i is the motional frequency of mode i .

Wigner functions. Oftentimes, it is useful to visualize the state of a motional mode in terms of its associated Wigner function (Steck (2007); Wigner (1932)), a picture of the state in terms of its distribution in phase space. The Wigner function $W(x, p)$ of a motional state $\psi(x)$,

$$W(x, p) = \frac{1}{2\pi\hbar} \int_{-\infty}^{\infty} \psi(x + x'/2) \psi^*(x - x'/2) e^{ipx'/\hbar} dx',$$

is a quasi-probability distribution, only returning a true probability distribution along one quadrature (in the case of a motional mode, position or momentum in some potentially rotating frame) when the other quadrature is integrated over. Because they offer a clear visual picture of the state, the discussions below of our motional operations will include plots (given in Figure 3.11) of how they transform the state's Wigner function.

Ground state Wigner function and Gaussian operations. An important Wigner function to note is that of the motional ground state $|n = 0\rangle$, which is a two-dimensional Gaussian curve (a direct consequence of the fact that the ground state wavefunction of a harmonic oscillator is Gaussian) with minimal uncertainty $\Delta x \Delta p = \hbar/2$. In most of our motional experiments, we are preparing one or more modes in the motional ground state and carrying out operations that shift, deform, or transfer (i.e. between modes) the Wigner function. These operations preserve the Gaussian shape of the Wigner functions (and the parabolic shape of the harmonic potential well) and are therefore called ‘‘Gaussian operations.’’ Non-Gaussian operations are discussed in Chapter VI.

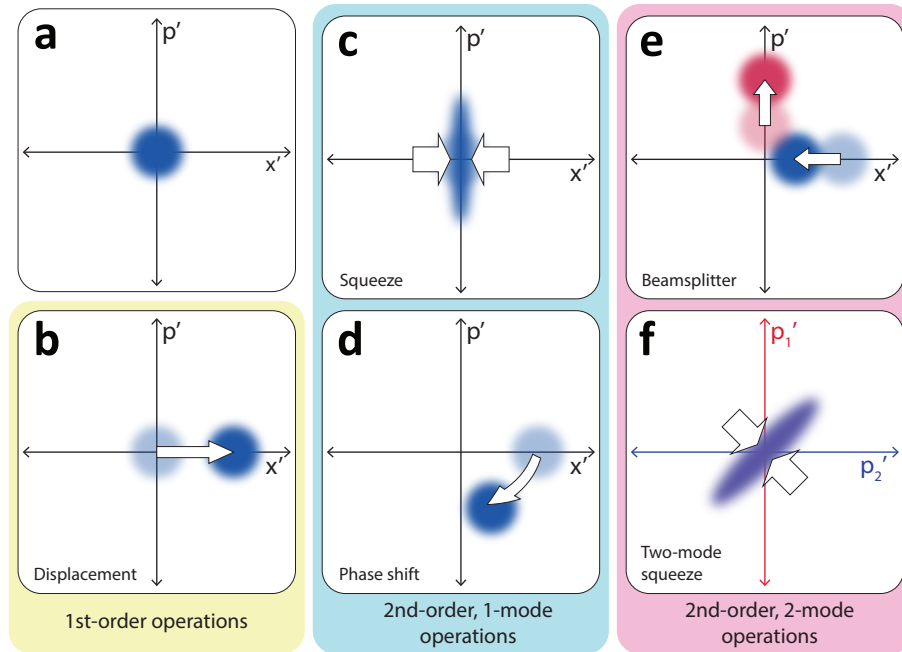


Figure 3.11. A reference table for the motional operations performed in our experiments, showing in the phase space picture **a**) the Wigner function of the motional ground state, **b** the effect of a displacement operation, **c**) the effect of squeezing, **d**) the effect of a phase shift on a displaced state, **e**) the effect of a beamsplitter operation on a pair of modes (with the Wigner function of the second mode shown in red), and **f**) the effect of two-mode squeezing on the correlations between the quadratures of a pair of modes.

3.6.1 General operations.

In the Fock state picture, the Hamiltonian associated with a given motional operation can be represented as a general combination of raising and lowering operator(s) for the mode(s) being driven. These operations can be carried out physically by applying a position-dependent electric potential. For example, an operation on a single mode aligned with the x -axis can be performed by applying a potential $V(x)$ at the ion, which can be Taylor expanded in x to $V(x) = \sum x^n (\partial^n / \partial x^n) V(x) / n!$, with $(\partial^n / \partial x^n) V(x)$ evaluated at the equilibrium point of the ion. If the raising and lowering operators of the mode to be controlled are a^\dagger and a , then the position operator x can be expressed as $x = x_0(a + a^\dagger)$, where $x_0 =$ is the ground state wavefunction extent of

the mode. In this representation, terms in $V(x)$ proportional to x^n will give terms n th order in the raising and lowering operators. Similar to the analysis of Rabi flopping, these terms will only contribute to dynamics if they do not time-average to zero in the interaction picture (i.e. if the drive is sufficiently close to resonance). The operations we carry out each have their own coupling rate, generally in the range of $2\pi \times 1-10$ kHz.

3.6.2 Displacement. Applying a first-order potential $V(x) \propto x$ yields an interaction Hamiltonian first-order in a and a^\dagger , $\bar{H} = \hbar g(e^{-i\delta t}a + e^{i\delta t}a^\dagger)$, where $\delta = \omega - \omega_m$ is detuning of the drive from motional resonance ω_m and g represents a coupling rate. If $\delta = 0$, then the effect of this operation on a mode is to generate a displacement $\alpha = gt$ after a time t . An intuitive picture of what this operation (displacement) does to the motional state can be inferred from the Wigner function in Figure 3.11b. By applying a dipole potential resonant with the ion's motional frequency, the ion is made to oscillate at increasingly large amplitudes in a coherent, phase-defined way. The displaced state is analogous to the classical harmonic oscillator state, with a minimally uncertain position/momentum tracing out closed loops in a (non-rotating) phase space at the mode frequency ω_x .

3.6.3 Single-mode squeezing and phase shifting. If instead of applying a first-order potential, a second-order one $V(x) \propto x^2$ is applied to the ion, a pair of operations can be performed. The x^2 operator can be factored into four terms which can be suggestively grouped as $H_{\text{drive}} \propto (a^\dagger a + aa^\dagger) + (a^{\dagger 2} + a^2)$.

Phase shift operation. The first grouping corresponds to cross-terms that do not change the population of the motional state population (since both of these operators commute with the harmonic oscillator Hamiltonian) but which, when the potential is held static, will shift the energy levels (i.e. applying a DC second-order potential changes the ion's confinement and motional frequency). Applied for a

finite time, this operation amounts to a phase shift. The effect of this operation on a displaced state is shown in Figure 3.11d.

Single-mode squeezing. The second grouping, corresponding to terms that add or remove two phonons at a time, has a less trivial effect than the phase shift. When a second-order potential is applied at a frequency of $2\omega_x$, a mode in the ground state will be converted into a superposition of even Fock states. This is a single-mode squeezed state. In the phase space picture, this amounts to a reduction in the uncertainty of one quadrature by a factor of e^{-r} (where $r = gt$ is the squeezing parameter, g and t being the squeeze operator coupling rate and drive time respectively) while increasing it by a factor of e^r in the other. In this way, single-mode squeezed states are still minimum uncertainty states. These states have been used in trapped ions in the past for phase-sensitive amplification of displacements (e.g. Burd et al. (2019)). The effect of single-mode squeezing on the ground state is shown in Figure 3.11c.

3.6.4 Two-mode squeezing and beamsplitter. If we have two orthogonal modes, one along the x -direction (with associated raising and lowering operators a and a^\dagger) and one along the y -direction (b and b^\dagger), and a second-order potential is applied which lies at a 45° angle in between them, such that, in this plane, $V(x, y) \propto xy$, a pair of operations can be performed which entangle the motional states of the two modes. As in the The xy operator can be factored into four terms, $H_{\text{drive}} \propto (a^\dagger b + ab^\dagger) + (a^\dagger b^\dagger + ab)$.

Beamsplitter. The first pair of terms, when the potential is applied at the difference frequency $\omega_x - \omega_y$ between the two modes, enacts the beamsplitter operation, which coherently exchanges phonons between modes. This operation can be used to transfer entire motional states between modes, e.g. moving a displaced state from one mode to the other as shown in Figure 3.11e. This operation is

analogous to an optical beamsplitter, where the two modes are spatial modes of light, and has been implemented in trapped ions before (Gorman et al. (2014)) for implementing motional interferometry (Toyoda et al. (2015)) and for the cooling of motional modes not easily accessed with laser cooling (Hou et al. (2024)).

Two-mode squeezing. The second pair of terms, when the potential is applied at the sum of the mode frequencies $\omega_x + \omega_y$, causes correlated generation/removal of phonons in the two modes, generating entangled states whose correlated nature are only visible when plotting quadratures from the two modes against each other. The output traced out over a single mode will appear to be a thermal state. This operation is analogous to the process of spontaneous parametric down-conversion in optics, with a key difference that two-mode squeezing on the vibrational modes of trapped ions can be carried out deterministically using electric potentials.

3.7 Coherent control - Spin-motion coupling

The ability to generate entanglement between an ion’s internal spin state and the state of its associated motional modes is needed for many of our experimental purposes, including ground state cooling (as discussed above in Section 3.2.3), motional state characterization (as discussed in Section 3.8), and gates (as discussed in Chapter IV). In our setup, we generate spin-motion entanglement in two ways (in both cases using our Raman beams): Driving Raman sideband transitions and driving spin-dependent forces (SDFs). These two techniques are illustrated schematically in Figure 3.12.

3.7.1 Raman sidebands. The treatment of Raman Rabi flopping given above in Section 3.5.2 implicitly considered the phase of two-mode optical field driving the ion to depend only on time (according to the tuning difference $\omega = \omega_0 + \delta$ of the optical fields). However, the phase of the field will also vary with

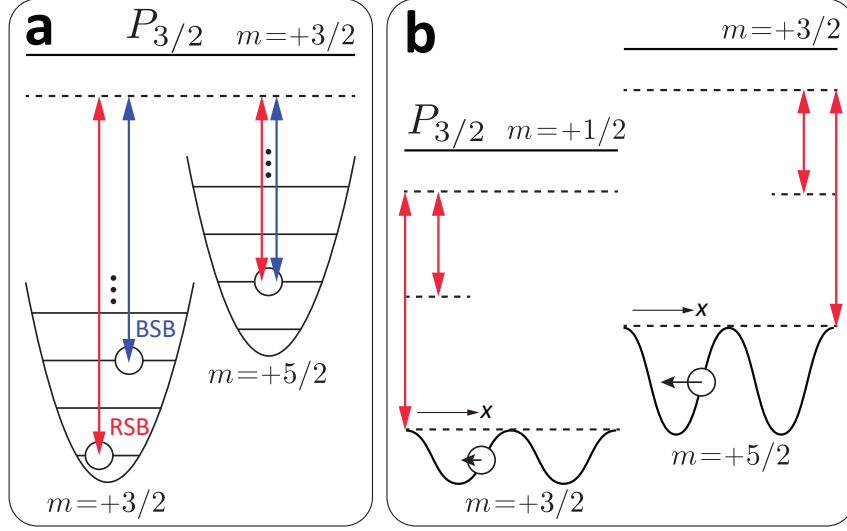


Figure 3.12. Two techniques of achieving spin-motion coupling. **a)** Raman sideband transitions, showing the $\Delta n = +1$ (BSB) and $\Delta n = -1$ (RSB) transitions, with the spin states being represented as a pair of harmonic oscillator energy ladders separated by the spin state splitting. **b)** Spin-dependent forces, showing how two interfering beams with the same polarizations will generate periodic intensity gradients of different amplitudes (and different force magnitudes) on the two spin states.

position (according to the wavevector difference $\Delta\vec{k} = \vec{k}_\downarrow - \vec{k}_\uparrow$. If we consider only the direction along which $\Delta\vec{k}$ lies, calling this the x -direction (which coincidentally aligns with the $\Delta\vec{k}$ of our two-beam Raman drive and the coordinate system of our trap), we can write the optical field phase at the ion as $\phi = \omega t - \Delta kx$, where x is the position of the ion. If the ion position was defined at an exact point, Δkx would be a fixed optical phase. Importantly however, the spatial wavefunction of the ion has a finite width on the scale of the ground state wavefunction extent $x_0 = \sqrt{2\hbar^2/m\omega_x^2}$ of the mode in the x -direction [[double check this]]. If the extent of this harmonic oscillator mode is taken into account, considering the position operator in terms of the raising and lowering operators such that $x = x_0(a + a^\dagger)$, then the position-dependent portion of the field phase becomes $\Delta kx = \Delta kx_0(a + a^\dagger) = \eta(a + a^\dagger)$, where η is the Lamb-Dicke parameter, which defines the ratio of the

ion's wavefunction extent to the steepness of the intensity gradient defined by the interfering optical fields. From this coupling between field phase and the motional operators, spin-motion coupling can be derived as follows.

If we consider the effective two-level interaction Hamiltonian \bar{H} of a Raman Rabi drive with effective Rabi frequency Ω , we can write (after applying the RWA and ignoring any fixed AC Stark shifts, which are simply calibrated out in real operations)

$$\bar{H} = \frac{\hbar\Omega}{2} \begin{bmatrix} 0 & e^{i(\delta t - \Delta kx)} \\ e^{-i(\delta t - \Delta kx)} & 0 \end{bmatrix}, \quad (3.23)$$

which can be rewritten in terms of the qubit raising and lowering operators σ_- and σ_+ (just the complementary off-diagonal terms in a 2x2 matrix) as

$$\bar{H} = \frac{\hbar\Omega}{2} (\sigma_- e^{i(\delta t - \Delta kx)} + \sigma_+ e^{-i(\delta t - \Delta kx)}), \quad (3.24)$$

which, for $|\eta| = |\Delta kx_0| \ll 1$ (which for small average mode occupations \bar{n} puts the system in the Lamb-Dicke regime), can be expanded to first order in η , such that $e^{i\Delta kx} \approx 1 + i\eta(ae^{-i\omega_m t} + a^\dagger e^{i\omega_m t})$ in the interaction picture, where ω_m is the motional mode frequency. Expanding the interaction Hamiltonian above gives a set of terms

that can be grouped into three distinct operations,

$$\bar{H} = \underbrace{\frac{\hbar\Omega}{2}(\sigma_- e^{-i\delta t} + \sigma_+ e^{i\delta t})}_{\text{carrier}} \quad (3.25)$$

$$- \underbrace{\frac{i\hbar\eta\Omega}{2}(\sigma_- a e^{-i(\delta-\omega_m)t} + \sigma_+ a^\dagger e^{i(\delta-\omega_m)t})}_{\text{BSB}} \quad (3.26)$$

$$- \underbrace{\frac{i\hbar\eta\Omega}{2}(\sigma_- a^\dagger e^{-i(\delta+\omega_m)t} + \sigma_+ a e^{i(\delta+\omega_m)t})}_{\text{RSB}} \quad (3.27)$$

$$\xrightarrow{\text{on resonance}} \quad (3.28)$$

$$\bar{H}_{\text{carrier}} = \frac{\hbar\Omega}{2}(\sigma_- + \sigma_+) \quad (3.29)$$

$$\bar{H}_{\text{BSB}} = \frac{i\hbar\eta\Omega}{2}(\sigma_- a + \sigma_+ a^\dagger) \quad (3.30)$$

$$\bar{H}_{\text{RSB}} = \frac{i\hbar\eta\Omega}{2}(\sigma_- a^\dagger + \sigma_+ a) \quad (3.31)$$

The first operation in the equation above describes carrier Raman Rabi flopping (as in Section 3.5.2), on-resonance when $\delta = 0$. If $\delta = \omega_m$, the blue sideband (BSB) will be driven on-resonance. The effect of the BSB will be to drive from the lower to upper qubit state while adding a phonon via the $\sigma_+ a^\dagger$ term. Alternatively, if $\delta = -\omega_m$, the RSB will be driven on-resonance, and driving from the lower to upper qubit state will subtract a phonon. Notably, the RSB and BSB are “relative.” If the qubit is initialized in the upper state, the BSB Hamiltonian above will serve to subtract a phonon while generating a spin flip. These sorts of sideband operations can be used for cooling and the preparation/detection of motional states, as laid out elsewhere in this document, or for carrying out entangling gates, i.e. the Mølmer-Sørensen gate (Mølmer and Sørensen (1999)).

3.7.2 Spin-dependent forces. The sidebands described above generate spin-motion entanglement by changing the state of the ion’s spin and motion in a correlated way. In contrast, spin-dependent forces (SDFs) will tend

not to change the spin state but will, as the term suggests, produce different effects on the motional state depending on the spin state. This technique has been used for generating cat states (citemonroe1996schrodinger) and for performing gates (Leibfried et al. (2003)). The conceptual logic for generating SDFs using Raman beams can be laid out as follows.

State-dependent AC Stark shifts. A single Raman beam generates an AC Stark shift on the ion, with the size of this shift depending on the Clebsch-Gordon coefficient of the transition being addressed via the optical Rabi frequency. Since a given polarization will address different transitions (or no transition at all) in different spin states, this means that the AC Stark shift is state-dependent.

Spatially periodic AC Stark shifts. If a second Raman beam with the same polarization is added to this picture, it will interfere with the first beam and generate interference fringes, spatially periodic in $2\pi/\Delta k$. These interference fringes represent areas of high and low intensity and, therefore, high and low AC Stark shifts.

State-dependent forces. Because these AC Stark shifts represent a shift in the energy associated with a given spin state, if this shift spatially varies, this amounts to a force on the ion that will tend to drive the ion to a spot of minimum energy (either minimum or maximum AC Stark shift, depending on the direction of the shift). For an ion in a superposition of spin states, the magnitude (and potentially direction) of this force will differ by the ratio of the AC Stark shifts on each of these states, meaning that, if the interference pattern and the force are static, the ion's potential well will be shifted by different amounts for each of these states. The ion's motional state is now entangled with its spin state, though likely in a way that isn't large (in the sense that these static shifts are likely to be small) or useful.

Time-varying SDFs. To generate appreciable differences in the motional states associated with different spin states, an SDF can be made to vary in time. This can be achieved by detuning the two Raman beams from each other such that the interference fringes, rather than being static with a phase depending on the path length difference between the two beams, will travel, with a speed $\omega\Delta k$, where ω is the tuning difference between the beams, generating a time-varying force at the ion which can be used to generate spin-dependent displacements if ω is tuned close to the mode frequency ω_m .

3.8 Motional state characterization

As described in Section 3.4 above, we can directly measure the spin state of an ion at the end of an experiment through state-selective shelving and fluorescence. We have no analogous way of directly measuring the motional state of the ion. However, by using the techniques for spin-motion coupling described above (specifically, use of sidebands), we can map the state of a selected motional mode onto the spin. The mismatch between these two Hilbert spaces means that, to fully characterize the occupation of different Fock states, we in general need to analyze sideband Rabi flopping (Meekhof et al. (1997)) as follows.

BSB motional state characterization. To see how sideband Rabi flopping could be used to characterize the Fock state distribution of a motional state, we can consider the interaction Hamiltonian for on-resonant BSB Rabi flopping

$$\bar{H}_{BSB} = \frac{i\hbar\eta\Omega}{2}(\sigma_- a + \sigma_+ a^\dagger), \quad (3.32)$$

and the effect that this Hamiltonian will have on spin/motion systems in different Fock states. Operating on the n -phonon Fock state $|n\rangle$ and taking the spin to be

initialized in the $|\downarrow\rangle$ state, the Hamiltonian yields

$$\bar{H}_{BSB} |\downarrow\rangle |n\rangle = \frac{i\hbar\eta\Omega}{2} (\sigma_- a + \sigma_+ a^\dagger) |\downarrow\rangle |n\rangle \quad (3.33)$$

$$= \frac{i\hbar\eta\Omega}{2} \sqrt{n+1} |\uparrow\rangle |n+1\rangle. \quad (3.34)$$

So if the motional mode is in the n -phonon Fock state, this Fock state will, along with the spin, Rabi flop at a frequency $\eta\Omega\sqrt{n+1}$. What this means is that, if the motional mode is in a combination of Fock states, the spin will Rabi flop with different frequency components, the amplitudes of which are the occupation probability $P(n) = |\langle n|n\rangle|^2$ of each Fock state. By carrying out this Rabi flopping and fitting to it the expression

$$P_\downarrow = \sum_{n=0}^{\infty} P(n) (\cos \eta\sqrt{n+1}\Omega t)^2, \quad (3.35)$$

we are able to infer $P(n)$, given η and Ω (the carrier Raman Rabi frequency) are well known.

Motional state characterization at small \bar{n} . The technique of looking at the frequency components of SB Rabi flopping works for general motional states but is time consuming. Motional state characterization can be performed more easily in the regime where the average mode occupation $\bar{n} \ll 1$ and where occupation of Fock states $n = 2$ and higher are negligible. In this case, by driving the RSB for the t_π of the $n = 1$ Fock state, the motional state can be mapped directly onto the spin, and we are able to “count” our phonon numbers directly. We use this technique of driving the RSB for a fixed span of time in other contexts as well where \bar{n} is not $\ll 1$. In these cases, we are typically looking less to characterize the exact motional state but to detect the presence or absence of phonons (i.e. whether or not the mode has been driven from the motional ground state).

3.9 Coherence time measurement

Decoherence is the main effect limiting our ability to control our ion for arbitrary amounts of time without reinitializing its state (aside from the fact that our qubit lifetime is limited to the natural lifetime of the $D_{5/2}$). Spin decoherence comes about through fluctuations in our qubit frequency translating to fluctuations in the phase between qubit states over the run time of the experiment. If these frequency fluctuations are consistent between runs of the experiment (e.g. the result of the application of AC Stark shifts that are consistent between runs, or the result of a 60 Hz line current that the experiment is being run in sync with, as discussed in Section 2.2.2), then the phase shift will also be consistent, and the state will still be coherent. If however the frequency fluctuations are random relative to the experiment, then the phase shifts that they produce will also be random. The effect of this can be most clearly illustrated by considering the most typical technique for characterizing these fluctuations and measuring coherence time: Ramsey experiments.

3.9.1 Ramsey experiment. The simplest type of Ramsey experiment is shown in the top left quadrant of Figure 3.13. In this experiment, after initializing the qubit in $|0\rangle$ we use a $\pi/2$ pulse to prepare our qubit in a state $|\psi\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ (taking the initial phase to be arbitrary). It will be helpful later to consider the density matrix $\rho = |\psi\rangle\langle\psi|$ of this state, which is initially

$$\rho = \frac{1}{2} \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}. \quad (3.36)$$

If we apply an identical $\pi/2$ pulse immediately, we will move the qubit to $|1\rangle$ (i.e. we will have done a π pulse). If however we allow some time t to pass between the $\pi/2$ pulses, fluctuations in the qubit frequency will cause some relative phase ϵ_ϕ to accumulate between the qubit states, with ϵ_ϕ effectively drawn from some

distribution $P(\epsilon_\phi, t)$ dependent on the frequency fluctuations seen in the system. Then the effective (experimentally observed) density matrix of the state at time t after the $\pi/2$ pulse can be written

$$\rho(t) = \frac{1}{2} \int_{-\infty}^{+\infty} d\epsilon_\phi P(\epsilon_\phi, t) \begin{bmatrix} 1 & e^{-i\epsilon_\phi} \\ e^{i\epsilon_\phi} & 1 \end{bmatrix}. \quad (3.37)$$

As t grows, so will the width of P in ϵ_ϕ , since random phase shifts from frequency fluctuations have had more time to accumulate. As this width increases, the off diagonal terms of the total density matrix become increasingly small until they vanish entirely, i.e.

$$\rho(t \rightarrow \infty) = \frac{1}{2} \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad (3.38)$$

at which point the qubit is now in a mixed state and can no longer be written as a coherent superposition of kets (a pure state). At this point, by the time of the second $\pi/2$ pulse, the qubit is randomly distributed around the equator of the Bloch sphere, and the $\pi/2$ pulse will, on average, not change the populations at all. At times between $t = 0$ and $t = \infty$, the qubit state will be some combination of mixed and pure states. To characterize the state's level of "mixedness" at different times, we can generate Ramsey fringes.

Ramsey fringes. To generate a single Ramsey fringe, we sit at a delay time t and adjust the phase ϕ of the last $\pi/2$ pulse such that, ideally, at certain phases we see full passage to $|1\rangle$ or full return to $|0\rangle$. Plotting the spin-flip probability will show a sinusoid (a Ramsey fringe) that has some upper and lower bound (the contrast) quantifying how mixed the state is at time t . This contrast can be taken at multiple values of t , and these values serve as a measurement of coherence time.

3.9.2 Ramsey with spin echoes. In the discussion above, frequency fluctuations (which in the rest of this section can just be called "frequency noise" or

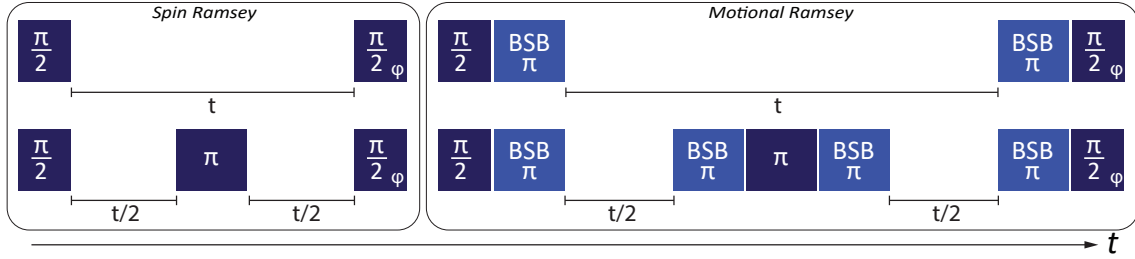


Figure 3.13. Pulse sequences for Ramsey sequences for characterizing spin (left) and motion (right) coherence, with (bottom) and without (top) a spin echo pulse. Carrier pulses are shown in dark blue, blue sideband pulses in light blue. To generate fringes at a given t , the phase ϕ of the final pulse can be scanned.

“noise”) were discussed with no particular consideration of their spectral character. In other words, no consideration was made of whether these fluctuations are fast or slow, the relative strengths of these components, etc. The spectrum of frequency noise can in principle be teased out using Ramsey sequences that include π -pulses (called spin echoes), timed so as to cancel out noise in certain frequency domains. This allows a qubit to serve in principle as a spectrum analyzer e.g. of magnetic field noise causing fluctuations in the splitting of a Zeeman/hyperfine qubit (Wei et al. (2022)). We typically do not seek to use Ramsey sequences to characterize our noise this carefully. However, the principle of using spin echo pulses (explained below) to cancel out low-frequency noise and extend coherence time is one that we use often in our work.

Spin echoes. In a situation where the frequency noise is concentrated in the low-frequency to DC range, qubit frequency may differ on a shot-to-shot basis, but it does not change appreciably over a single shot (i.e. between the first and last $\pi/2$ -pulses). This means that, effectively, after the first pulse, the qubit state is simply precessing along the Bloch sphere at a rate δ , where δ is the fixed tuning error in a given shot. If halfway in between the $\pi/2$ -pulses (at $t/2$, a π -pulse is performed, it will (in the Bloch sphere picture) effectively mirror the qubit

state about the plane defined by the poles and the axis of rotation of the π -pulse. Considering for concreteness the case where the π -pulse is about the same axis as the $\pi/2$ -pulses, the qubit phase is taken from $\delta t/2$ to $\pi - \delta t/2$. Allowing the qubit to precess for another $t/2$ then brings the qubit phase to π , regardless of the size of δ (so long as δ is sufficiently small that it does not affect the contrast of the carrier pulses). The fact that this π -pulse causes a qubit's state to reconverge to the same value after diverging is why it is called a spin echo pulse. Since prior to the last $\pi/2$ -pulse, the qubit phase is always brought to the same value, Ramsey contrast in a sequence with a spin echo pulse will not start to decay until t is large enough that δ is no longer constant over a single shot.

3.9.3 Fock state Ramsey experiment. Just like the stability of the qubit splitting can be measured at different time scales using a spin Ramsey sequence, the stability of a motional mode frequency can be measured in the same way using a motional Ramsey sequence, using the $|n = 0\rangle$ and $|n = 1\rangle$ Fock states rather than the two qubit levels. There is no analogue to π - and $\pi/2$ -pulses for a motional mode, so a combination of qubit carrier and BSB pulses are used, as shown in the right of Figure 3.13. Using this technique requires good ground state cooling to properly initialize the “motional qubit” to $|n = 0\rangle$.

3.10 Routine calibration techniques

To conclude this chapter, a summary will be given of the techniques used in our lab to calibrate our coherent (spin and motional drives), our light shifts, and our laser frequencies. Finally, a brief description will be given of a technique (two-point calibrations/two point servoing) for performing calibrations in real-time, such that a parameter can be tracked as another experiment is run.

3.10.1 Coherent drive calibrations. Applying all of the techniques discussed in this chapter in the context of an actual experiment requires a way

of measuring the relevant parameters (e.g. t_π for some type of Rabi flopping). Broadly, there are two things that need to be known when trying to produce a certain effect with a coherent drive: 1) at what frequency should it be set so as to drive the relevant transition on-resonance (or off-resonance by a known amount), and 2) for how long (or, at a fixed drive time, how hard) should this drive be turned on. These needs suggest two broad classes of calibration experiment: calibrating resonant frequencies (i.e. performing spectroscopy) and calibrating coupling rates.

As discussed in the sections above, we have two classes of drive to consider: spin (Rabi drives) and motion (motional drives). Of the drives that generate spin-motion coupling, Raman sideband flopping is a Rabi drive and SDF is a motional drive for the sake of the discussion below.

Rabi drives. Rabi drives can be calibrated using Rabi spectroscopy (i.e. driving at different frequencies and looking for where spin flips occur) and Rabi flopping (i.e. driving on-resonance for different amounts of time to see how long spin flips take).

Motional drives. While similar calibrations need to be done for motional drives, a key difference from the Rabi drive is that there is no direct way of measuring the motional states. Instead, we need to map the motional state onto the spin using a Raman sideband. For frequency calibrations, after applying the motional drive, we pulse the red sideband, such that if the motional mode is excited, we see spin flips, while if the motional mode is still in the ground state, we see none. In the case of coupling rates, we only need to calibrate these rarely (since our motional drives are all-electronic and tend not to drift beyond any dependence on small mode frequency changes. When we do calibrate them however, we use

blue sideband Raman Rabi flopping to characterize the motional state generated, as discussed in Section 3.8.

With one motional drive, our squeeze drive, we typically use time-reversal protocol as a way of calibrating the resonant frequency. In this protocol, we perform the drive and then attempt to undo it by running it again with a π phase shift. If the drive is resonant, the squeezing will be undone and the motional mode will be returned to the ground state. If it is near-resonant, the second pulse will be phase-shifted, and the squeezed state will not be undone and could in fact be made larger. Because of the exponential growth of \bar{n} with the squeezing parameter r , it is possible to achieve small feature widths by this method with relatively short interrogation times (~ 1 kHz or narrower lines with $\sim 100 \mu\text{s}$ interrogation times).

3.10.2 Light shift calibrations. Oftentimes, it is important to know the size of a light shift generated by e.g. a Raman beam, either to know how much the tuning of a Rabi drive needs to be changed by or as a way of inferring beam intensity. There are two techniques we use for measuring light shifts on our qubit states, both represented in Figure 3.14. The first is to perform RF Rabi spectroscopy while the beam we are interested in is turned on (Figure 3.14a). This technique returns the light shifted qubit frequency directly, which can be compared to the un-shifted frequency to determine the light shift. An alternative technique for light shift measurement is to use a Ramsey sequence with a spin echo (Figure 3.14b). Here, the two “segments” of the Ramsey sequence act like the arms of an interferometer, with a phase difference between the two of $\delta t/2$, where δ is the magnitude of the light shift and t is the total sequence time. If either t or the beam power are scanned, Ramsey fringes will result, the frequency of which can be used to infer the lightshift at a given power.




Figure 3.14. Experiments for measuring the AC Stark shifts caused by a given beam (red). **a)** Carrying out Rabi spectroscopy while the beam is on and comparing this result to the result when the beam is off. **b)** Performing a Ramsey sequence with a spin echo, with the beam only being on during one delay so as to cause asymmetrical accumulation of phase.

3.10.3 Optical frequency calibrations. The optical transitions of a trapped ion serve as an absolute frequency reference against which we can both calibrate our wavemeter at different frequency ranges and monitor long-term stability (wavemeter drift) and short-term stability (laser lock quality) of our frequency locking system, discussed in Section 2.4.2. We can perform direct spectroscopy on our 397 nm cycling transition by simply driving it at different frequencies and looking for scattered photons. However, to prevent shelving in the $D_{3/2}$ level, we must have the 866 nm beam on during this scattering as well. While the 866 nm beam is generally tuned ~ 50 MHz blue of resonance, it is run at sufficiently high powers that the 397 nm line can still be distorted by dark resonances. In order to probe our transitions without distortion by other beams, we look for shelving or deshelling using the sequences shown in Figure 3.15 and described below.

393 nm line. The 393 nm line is the simplest to probe. The 393 beam is pulsed on for a short time and readout is done to check for shelving to the $D_{5/2}$ level.

397 nm line. To probe the 397 nm line, we pulse the 397 beam on to shelve in the $D_{3/2}$ level. Because this level will be deshelled by the 866 beam, before

readout, we shelve any population that remains in $S_{1/2}$ in the $D_{5/2}$ level with a pulse of the 393 beam.

854 nm line. To probe the 854 nm line, we shelve our entire population in $D_{5/2}$, then pulse the 854 beam and check for desheling in readout.

866 nm line. For the 866 nm line, we shelve the entire population in $D_{3/2}$ using the 397 beam, pulse the 866 beam to deshelve into $S_{1/2}$, then reshelve any population here into $D_{5/2}$ with the 393 beam.

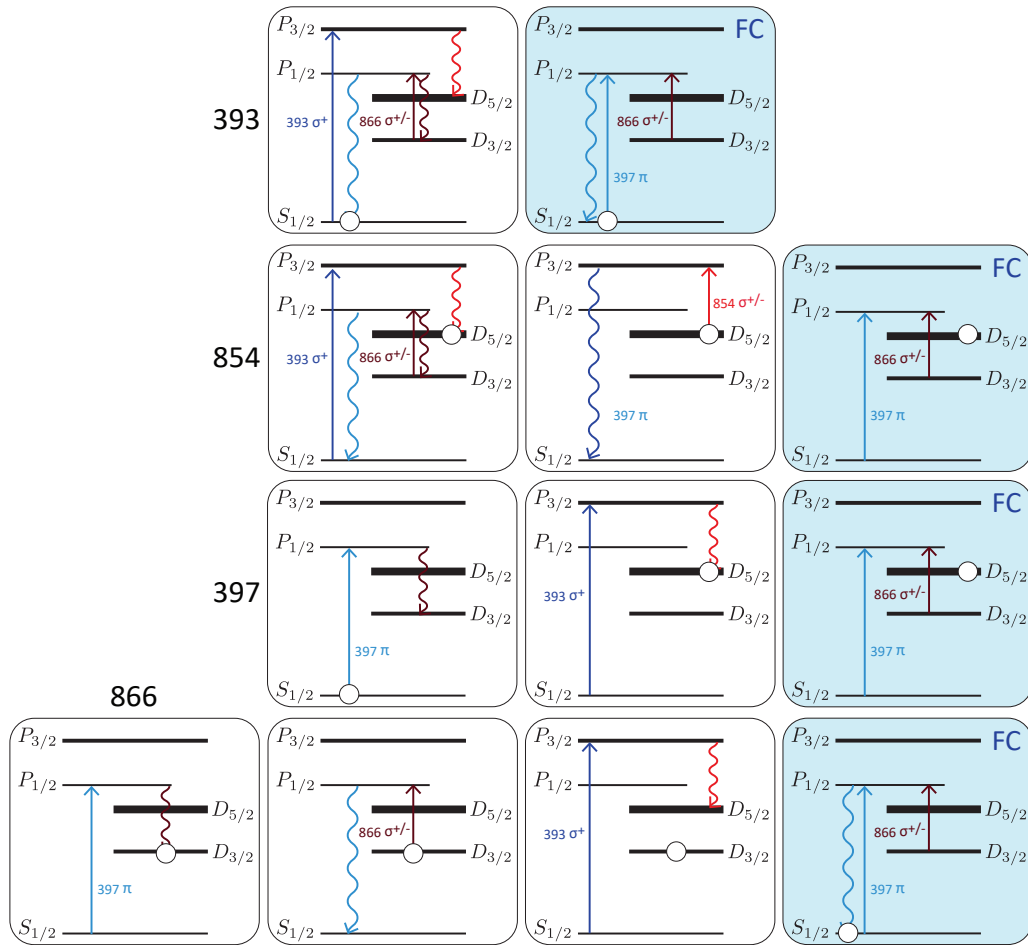


Figure 3.15. Pulse sequences (with time running from left to right) of our four laser line calibrations, with the state of the ion at the end of the experiment (during a fluorescence check, FC) shown assuming that laser associated with the line being characterized is far off-resonance.

Using these calibrations, we can do spectroscopy on all of our available dipole transitions. An important point is that, because we are measuring bright state probabilities rather than counting scattered photons, the lineshapes we measure will only be approximately Lorentzian (with the approximation improving at smaller shelving/deshelving probabilities). However, even with a distorted lineshape, we can still accurately measure the line centers.




Figure 3.16. Illustration of the two-point calibration servo loop.

3.10.4 Two-point calibrations. All of the calibration described above are typically carried out as scans over some parameter, i.e. they require a set of tens of points to be taken to which a curve can be fitted. This is a long process, and there are contexts where we want to track calibrated parameters in real time (e.g. during other experiments that would be very sensitive to drift in a certain parameter). To do this, we use a two-point calibration method. This method can be explained by considering a generic spectroscopic line with a center x_o , a width

(however defined) of $2\Delta x$, and a height y_0 . To check the position of this line, we measure y (which is usually spin-flip probability) at points we expect to be centered around x_0 ($x_0 \pm \Delta x$). If the value has not changed, we expect these two points to have the same value of y (to within shot noise). If there is some difference Δy between these points, it indicates that x_0 has moved and needs to be updated, with the direction of the update determined by Δy and the sign of the peak. This feedback on x_0 is carried out with a PI loop. The cycle for a two-point calibration is shown in Figure 3.16.

CHAPTER IV

RAMAN GATES IN METASTABLE QUBITS

Trapped ions are one of the leading platforms for quantum information processing (QIP), with some of the highest reported one- and two-qubit gate fidelities (Clark et al. (2021); Löschnauer et al. (2024)). Traditionally, trapped-ion qubits are encoded in electronic ground states, and prevention of photon crosstalk between ions used for information processing and ions used for ancillary functions (cooling, mid-circuit readout, etc.) is achieved by using different ion species. However, compared to a single-species setup, dual-species schemes see higher technical overhead (due to e.g. the need for multiple laser systems) and issues with motional control (e.g. reduced efficiency of motional coupling between ions). Recently, there has been interest in achieving dual-species functionality with a single ion species by using the more complete set of electronic states available for encoding quantum information (the so-called *omg* qubit scheme, Allcock et al. (2021)). In this vein, work has been done to verify that ground-state (g) and metastable-state (m) qubits can be operated independently without photon crosstalk (Feng et al. (2024); Yang et al. (2022)) and to null unwanted light shifts on given qubit types (Vizvary et al. (2024)). Much recent work has also been done on entangling gates in optical (o) qubits (Aksenov et al. (2023); Băzăvan et al. (2023); Clark et al. (2021); Manovitz, Shapira, Gazit, Akerman, and Ozeri (2022)). Parallel work has been done investigating potential uses of the *omg* scheme in neutral atom QIP, including as a tool for midcircuit measurement (Huie et al. (2023); Lis et al. (2023)) and erasure error conversion (Kobayashi and Nagayama (2024); Ma et al. (2023); Wu, Kolkowitz, Puri, and Thompson (2022)). However, the full advantages of *omg* (e.g. being able to perform dissipative and coherent operations simultaneously on a single ion chain) require coherent operations in

m qubits. While some work has been done on single m qubits (e.g. DeBry et al. (2023)), entangling gates in m qubits have not been explored until recently (Băzăvan et al. (2023)), and these gates relied on coupling between m and g states. Recent models of far-detuned Raman scattering (Moore et al. (2023)) suggest that two-qubit Raman gates in m qubits should be able to achieve spontaneous Raman scattering (SRS) errors $< 10^{-5}$ in calcium and barium with Raman beams tuned far-red from resonance. We experimentally verified in our own setup that the SRS rates in m qubits are consistent with this theory, and an accounting of this work can be found in Daniel Moore’s thesis (Moore (2023)). Given that Raman scattering represents a fundamental limit to Raman gate fidelity, this suggests that metastable trapped-ion qubits are likely a viable platform for QIP and that entangling gates in this platform are a promising area of investigation. This is especially true since, as recent investigations of erasure error suggest, SRS errors in m qubits may be uniquely easy to correct (Kang, Campbell, and Brown (2023)).

This chapter is organized as follows: First, leakage errors in m qubits and their conversion into erasure error is discussed. Next, experimental work is presented, with this erasure conversion technique implemented in the context of a two-qubit entangling gate (one of the first carried out in m qubits, and the first to break 98 – 99% estimated fidelity) and its associated error budget broken into erasure and non-erasure components. Finally, the estimated error rate of this gate is discussed, with particular emphasis placed on errors arising due to the erasure conversion scheme used and the fundamental limits for this approach.

4.1 Erasure errors in m qubits

As discussed in Chapter I, the types of errors most commonly considered in quantum computing are those which leave the system in its qubit subspace (e.g. Pauli errors, dephasing). However, classes of error (*leakage* errors) which leave the

system outside of this subspace have received much recent attention. Undetected leakage error is a serious problem from an error correction standpoint, as these errors amount to qubits being destroyed (such that an attempt to read them will always produce the same result). However, if they leave the system in states that can be probed without disturbing the qubit subspace, they can be detected in ways impossible with other error types. When a leakage error is detected, it becomes an erasure, a form of heralded error where the qubit is in (or can be put into) a known state outside of the qubit manifold. Compared to Pauli or spin-flip errors, erasure errors have higher error rate thresholds for efficient correction (see e.g. Kang et al. (2023) for a theoretical analysis of m qubits in trapped ions or the analysis offered in Ma et al. (2023) of neutral atoms), so there is currently interest in both 1) effectively converting leakage error to erasure error (erasure conversion) in cases where it naturally occurs and 2) finding ways of converting non-leakage error into (detectable) leakage error. These points, as they relate to m qubits in trapped ions, are discussed below.

4.1.1 Erasure conversion with fluorescence. All leakage errors in m qubits can in principle be converted into erasures. The two places a leakage error in an m qubit can put the ion are 1) outside of the $D_{5/2}$ manifold or 2) in one of the other Zeeman/hyperfine levels of $D_{5/2}$. If a leakage of the first type occurs and we attempt to drive the $S_{1/2} \leftrightarrow P_{1/2}$ cycling transition (a “fluorescence check” or FC) while preventing shelving in the $D_{3/2}$ manifold by driving the $D_{3/2} \rightarrow P_{1/2}$ transition, fluorescence will be observed. If the ion is still shelved, it will be dark. In our experiment, most leakage events are of this type and can be detected with an FC. In cases where leakage is to other states in the $D_{5/2}$ manifold, this can in principle be detected by selectively deshelving these states using a narrow-linewidth $S_{1/2} \leftrightarrow D_{5/2}$ beam prior to performing an FC, though we do not perform this

type of deshelling in our work. The way that these leakage detection / erasure conversion fluorescence checks are implemented in the context of a full experimental shot is shown in Figure 4.1.

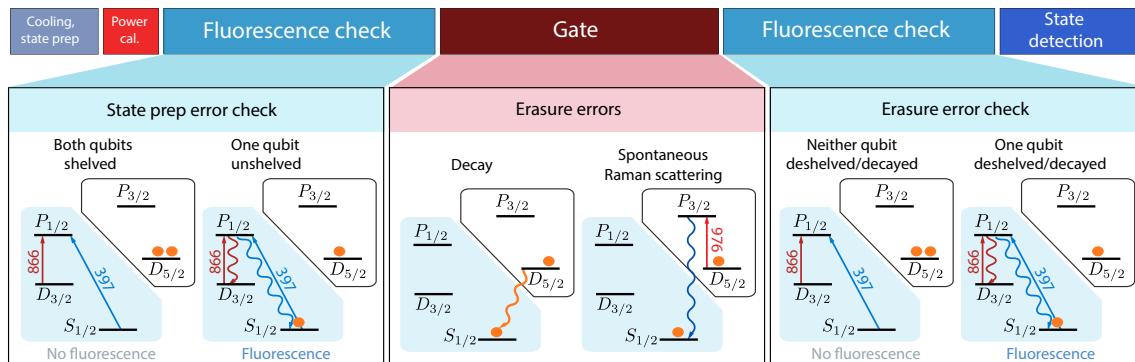


Figure 4.1. Schematic illustration of our erasure error detection scheme in the context of a gate, showing how checks for 397 fluorescence can be used both to confirm successful state preparation prior to running a gate and check for leakage errors after a gate.

4.1.2 Leakage pathways in m qubits. Given that the metastable levels have a finite lifetime before decaying to the electronic ground state, leakage errors are inherent to m qubits. In our experiment, there are two dominant leakage channels. The first is decay due to the natural lifetime τ of the $D_{5/2}$ manifold ($\tau = 1.16$ s in Ca^+ , Benhelm (2008)). In an algorithm, this decay can occur during gates (with a probability per gate t_g/τ , where t_g is the gate time) or during a fluorescence check (with a probability t_{FC}/τ). The latter decay introduces a set of errors, some of which will amount to SPAM (state preparation and measurement) error and some of which contribute to erasure overhead. This error is discussed more thoroughly later on, when discussing the error breakdown for our gate. The other dominant leakage channel is spontaneous Raman scattering (SRS), or the absorption of a Raman beam photon followed by spontaneous decay to a random state different from the starting state. SRS results in both leakage errors and spin

flip errors, the relative probabilities of which are determined by the branching ratios from the excited state. In Ca^+ for example, near-resonant scattering from the $P_{1/2}$ level will put the ion into $S_{1/2} \approx 95\%$ of the time. This means that in a g qubit, accidental shelving into $D_{3/2}$ (a detectable leakage error) occurs only in $\approx 5\%$ of scattering events. This ratio is flipped in m qubits however, where (near-resonance) $\approx 94\%$ of scattering events from $P_{3/2}$ will deshelve the ion from the $D_{5/2}$ manifold (Song et al. (2019)). This means that in m qubits in Ca^+ , the vast majority of SRS errors are leakage errors that can be straightforwardly detected, potentially placing less stringent correction requirements on one of the fundamental error sources for Raman gates in trapped-ion qubits. The picture, where branching ratios favor leakage in SRS events, is similar for m qubits in other ion species.

4.2 Geometric phase gates

Although m qubits may compare favorably to g qubits in terms of erasure conversion, this feature is only useful if high-fidelity gates can be performed in this platform, and the primary goal of our work was to demonstrate that this is the case. To demonstrate entanglement generation in m qubits (without use of the g qubit levels), we perform a geometric phase gate in the σ_z basis (Leibfried et al. (2003)) on the low-frequency out-of-phase (LF OOP) radial mode of a pair of m qubits. As discussed in Chapter III, to generate the spin-dependent force (SDF) used to perform this gate, we intersect our two perpendicular 976 nm beams (R_σ and R_{null}) with frequency difference ω , generating a traveling interference pattern along the radial direction of the two-ion crystal. Pulses of the R_{null} are shaped (as discussed in Section 2.7.4), with a $2\ \mu\text{s}$ rise time, to prevent off-resonant driving of the spin carrier by the motional drive and vice versa. It can be shown that

evolution under the SDF is governed by the Hamiltonian

$$H = \frac{i\eta}{2} \sum_j (\Omega_{\uparrow} |\uparrow_j\rangle \langle \uparrow_j| - \Omega_{\downarrow} |\downarrow_j\rangle \langle \downarrow_j|) e^{-i\delta t + i\phi} a^\dagger + \text{h.c.} \quad (4.1)$$

where j indexes the ion, η is the Lamb-Dicke parameter, $\delta \equiv \omega - \omega_r$ is the detuning of the drive field from the LF OOP mode, ϕ the phase of the drive field, and $\Omega_{\uparrow(\downarrow)}$ is a coupling rate proportional to the amplitude of the AC Stark shift on the \uparrow (\downarrow) spin state. This SDF drives loops in motional phase space, accruing state-dependant geometric phases proportional to the loop area and generating entanglement.

To counteract slow noise and static calibration errors, we use Walsh modulation (Hayes et al. (2012)). Walsh modulation is a technique for cancelling out noise components up to arbitrarily high orders (at the expense of increasing operation time) by modulating the phase of the drive (in the case of motion) or using π -pulses (in the case of spin). The gate sequence we use is Walsh modulated in both spin and motion, with W(1) modulation on the spin (i.e. a spin-echo pulse in the middle of the gate sequence) and W(3) modulation on the motion, performing the gate in four phase space loops. A schematic of the full gate sequence is shown in Figure 4.2a. When characterizing gate output at different SDF detunings δ , we modulate the phases between these loops such that the gate mode ideally returns to the motional ground state, disentangling the motion and the spin in a detuning-insensitive way. A representative scan over δ is shown in Figure 4.2b.

4.2.1 Fidelity estimation. The ideal output of the gate is a Bell state $|\Phi_+\rangle = |\downarrow\downarrow\rangle + |\uparrow\uparrow\rangle$. The fidelity $\mathcal{F} = |\langle \Phi_+ | \psi \rangle|^2$ of the actual gate output $|\Psi\rangle$ can be estimated by rotating the output with a $\pi/2$ pulse performed at different phases ϕ . The output will be a set of states with fringes in parity $\Pi = P(\downarrow\downarrow) + P(\uparrow\uparrow)$

) - $P(\uparrow\downarrow) - P(\downarrow\uparrow)$ ranging from $\Pi = \pm 1$ given an ideal Bell state. A representative parity fringe is shown in Figure 4.2c.

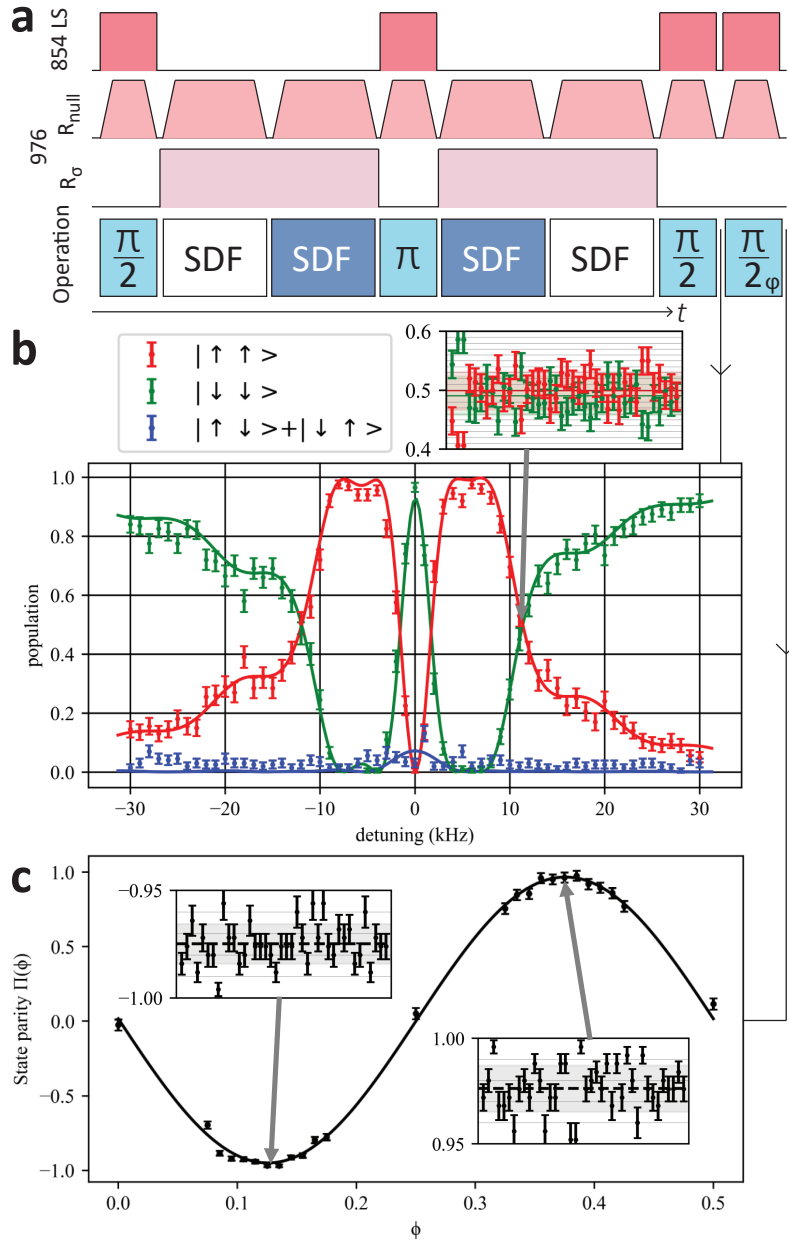


Figure 4.2. **a)** Gate pulse sequence showing beams relevant to the gate, with SDF phase ($0, \pi$) denoted by color (white, blue). **b)** A scan over SDF detuning from the LF OOP mode resonance, showing the crossing point in $|\downarrow\downarrow\rangle$ and $|\uparrow\uparrow\rangle$ populations at which the gate was performed. Inset shows high-shot time series data at the gate operating point. **c)** A sample parity fringe, with insets showing high-shot time-series data at the parity fringe extrema.

To estimate gate fidelity, we measure 1) populations at the end of the gate and 2) the parities at the peaks of positive and negative parity fringes. At these three points, we take 20,000 shots (without erasure errors) in interleaved 500 shot points. The associated time series are shown in insets on Figure 4.2. From these measurements, we calculate gate fidelity according to $\mathcal{F} = (P_{B,\uparrow\uparrow} + P_{B,\downarrow\downarrow} + C)/2$, where $P_{B,\uparrow\uparrow}$ and $P_{B,\downarrow\downarrow}$ are measured populations of the output Bell state and C is the amplitude of the parity fringes. Subtracting out a 0.58% two-ion SPAM error, this gives us a gate error $\epsilon = 1 - \mathcal{F}$ of 1.44(10)%, calculating error bars by bootstrapping the individual sets of 500 shot measurements. If we subtract erasure errors (i.e. in our analysis, ignoring shots where leakage errors were detected), the remaining error is 0.86(10)%. Below, we discuss how these compare to expectations based on ancillary measurements.

4.3 Gate errors

To infer the sources of our estimated infidelity, we calculate the errors that we would expect to see in our gate based on characterization of known error sources (i.e. dephasing rates, scattering rates). These error sources are discussed in greater detail below, both in terms of their causes and avenues for improvement. Finally, errors introduced by our erasure conversion scheme are outlined, and the fundamental limits of those errors are outlined.

4.3.1 Error budget. An error budget for our gate is shown in Table A.7 (with full calculations presented in Appendix A). This table only lists error sources contributing $> 10^{-4}$ to infidelity (aside from Raman scattering back into $D_{5/2}$). Other sources (e.g. mode heating, carrier drive detuning, uneven Raman beam intensities) are not shown, but their contributions are included in

the total infidelity. Error sources can be broadly categorized as erasure and non-erasure.

The bulk of non-erasure error arises from dephasing, which collectively accounts for an infidelity of 0.87(10)%. Adding in smaller coherent errors (e.g. calibration errors and drift) and $< 10^{-4}$ incoherent errors, our predicted gate infidelity excluding erasure is 0.98(-0.09)(+0.18)%, which matches our measured erasure-excluded infidelity of 0.86(10)% to within a 68% confidence interval. For erasure error, the dominant contributor (at 0.29%) is decay from the qubit manifold due to the natural lifetime of the $D_{5/2}$ level. This in turn is mostly (0.17%) due to long (1 ms) fluorescence check times used for leakage detection, with the remainder of decay events occurring during the gate and associated delays/latencies. Predicted erasure error rate from 976 SRS, calculated from direct measurements of scattering rates, are 0.14%. Finally, polarization impurities in the 854 LS beam cause deshelling during carrier operations, contributing an expected 0.06% to the erasure error rate. Altogether, we predict an erasure error rate of 0.48(1)%, which matches the measured 0.55(3)% rate to within 95% confidence intervals.

4.3.2 Addressing dominant error sources. Given the error rates presented above, in the remainder of this chapter, we discuss avenues for reducing both erasure (conceptually divided into erasure caused by technical overhead and erasure inherent to the operation of the gate) and non-erasure errors.

Non-erasure errors. Of the expected 0.98% non-erasure error we calculate, 0.87% is due to dephasing, while nearly all of the remaining 0.11% is due to coherent processes (uncertainty in carrier π -time calibrations, possible beam intensity calibration errors, and mode frequency drift). The first class of error will monotonically decrease with decreased gate time given most typical noise

distributions. This means that this error would be suppressed by improving coherence times or by running faster gates, whether by increasing Raman beam intensities or by improving gate efficiencies, as discussed below. Our dominant coherent errors, with the exception of mode frequency drift, are not directly reduced by increasing gate speed but can be addressed by other means. For example, the effects of π -time calibration uncertainty (0.07%) could in principle be mitigated through the use of composite pulses (Morigi et al. (2000)) or sequences. In our setup, composite sequences were not practically feasible due to drifts in the $\approx 1\%$ differential light shifts on our two ions (a result of large inter-ion spacings relative to the focused beam diameter of our Raman beams).

Erasure errors. Erasure errors in our experiment stem from four factors. Firstly, our gate suffers from scattering errors due to 854 LS polarization impurities. These errors are not fundamental to our gate, since (as discussed in Section 3.3.3) in principle they can be mitigated or eliminated by increasing 854 LS detuning and intensity. More broadly, AC Stark shift beams are not an inherent requirement for m qubits (e.g. in hyperfine clock qubits). Secondly, our gate suffers SRS by the 976 nm gate beams, and this error is fundamentally fixed for a given qubit, gate scheme, and optical setup. Although at this wavelength, we should be able to run gates in Ca^+ with SRS rates $< 10^{-4}$ (Moore et al. (2023)), our rates of SRS events per operation are increased by reduced gate efficiency. These inefficiencies come down to 1) our choice of qubit states giving sub-optimal difference in Clebsch-Gordon coefficients (e.g. for $m = +5/2$ and $m = +3/2$, this difference is $4/15$, while if we encoded in $m = +5/2$ and $m = -1/2$ while keeping everything else the same, we would instead have a difference of $9/15$. The ratio of the squares of these differences, or ≈ 5 , represents the speedup this switch would provide), 2) the need to run a 4-loop gate rather than a 1-loop gate in order to suppress spin/motional

noise (increasing the gate time at a given beam intensity by a factor of two), 3) our use of orthogonal rather than counterpropagating Raman beams combined with the wavevector difference between the beams being at a 45° angle to the driven mode (collectively costing another factor of 2 in efficiency), and 4) polarization components in R_{null} (added in order to allow carrier and sideband operations and to null the differential AC Stark shift on the qubit levels), which contribute to scattering but not to SDF strength (giving only one third of the total power in the σ^- polarization of this beam and reducing efficiency by another factor of $\sqrt{3}$). Taking all of this into account, the gate speed could in principle be increased by a factor of $\sim 30 - 40$ without increasing beam intensities, which would take SRS error rates well below 10^{-4} . Thirdly, our ion decays from the qubit manifold during gate operations. Like SRS error rates, this error would be reduced by increasing gate efficiency and being able to run shorter gates, though this same end can be achieved by increasing Raman beam intensities (e.g. by increasing beam powers). Finally, decay from the qubit manifold during leakage checks (both during the post-state preparation check and the post-gate leakage check) imposes a fixed error overhead on the gate, discussed in greater detail below.

Alternative qubit choices. As a brief aside, many of the errors discussed above can be addressed by switching from a σ_z geometric phase gates in $^{40}\text{Ca}^+$ Zeeman qubits to running σ_x geometric phase gates (i.e. Mølmer-Sørensen gates) in $^{43}\text{Ca}^+$ hyperfine clock qubits would address multiple sources of error at once: 1) SRS could be reduced to $< 10^{-4}$, 2) a separate light shift beam would no longer be necessary for isolating the qubit states, 3) gates could be run a factor of $> 10\times$ more quickly, and 4) spin dephasing could be made negligible.

4.3.3 Errors from leakage detection. To consider the error overhead produced by using an FC to herald leakage (as implemented in

Figure 4.1), we consider the breakdown of errors with and without the final FC, assuming that leakage is the only type of error in the system. If no FC is performed after running gates, there are three possible outcomes: 1) no leakage occurs, 2) shelving during state preparation fails (the probability of this being $P_{dFC,u}$, or the probability of a decay event occurring during an FC and being missed), and 3) deshelling during operations (through decay or SRS, with a total probability P_l). The overall unheralded error then will be $P_l + P_{dFC,u}$. Adding in an FC for heralding leakage expands the space of possible outcomes, as shown in the flowchart in Figure 4.3. In this case, the outcomes where a leakage error occurs can be grouped into erasure error and unheralded leakage error. But assuming that the error complement $\bar{\epsilon}_i = 1 - \epsilon_i \approx 1$ for an error source i , it can be shown that total erasure error should be $P_l + P_{dFC} + \epsilon_d$ (where P_{dFC} is the probability of decay during an FC and ϵ_d is the probability of a dark state being misidentified by an FC, assuming no decay occurs). This ϵ_d component represents the probability of a false positive erasure measurement, which will become an actual erasure if the state of the qubit is then reset to some known location. Unheralded leakage error with an FC still contains a factor of $P_{dFC,u}$ (since unheralded decay during the state preparation FC is caught by the leakage check FC, but the leakage check FC adds a factor of this error of its own), but the unheralded leakage component without the FC is reduced by a factor of ϵ_b (the probability of a bright state being misidentified by an FC). All this gives a total unheralded leakage error of $(P_l + P_{dFC,u})\epsilon_b$. Comparisons between these error components for the two cases are presented in Figure 4.3b.

Given the discussion above, FCs can be said to generate overhead errors of $P_{dFC,u}$ in unheralded leakage error (5.5×10^{-4} in our experiments, with this number being indistinguishable from our SPAM error and included in this number) and

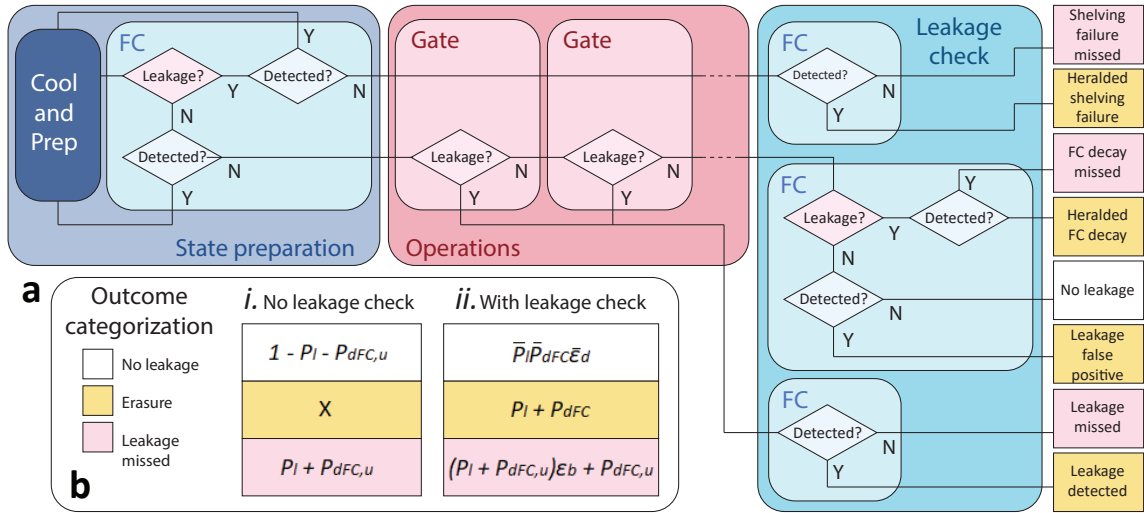


Figure 4.3. a) A flowchart illustrating the possible outcomes of a one-ion gate sequence with an FC carried out at the end, considering only leakage errors. Outcomes corresponding to an erasure error are labelled in yellow, while outcomes corresponding to a missed leakage error that is missed are labelled in pink. b) An accounting of the error probabilities in each category with and without leakage checks.

$P_{dFC} + \epsilon_d$ in erasure error (bearing in mind that $P_{dFC,u}$ of this is state preparation error). As a result, the limit of these errors in the scheme shown in Figure 4.3 is set by 1) how quickly an FC can be run and 2) how accurate state detection can be made. A truly fundamental limit on these is the lifetime of the $P_{1/2}$ level. However, even with this limit, as well as technical limitations (collection efficiency being < 1 , dark counts on detectors) that can be difficult to address, it is possible to high (i.e. $> 99.9\%$) detection fidelities in much shorter time scales than used in our experiment (i.e. $\approx 10 \mu\text{s}$ in both calcium (Myerson et al. (2008)) and other species (Crain et al. (2019))), corresponding to a two orders-of-magnitude reduction in error). In addition, the leakage check scheme presented here is in some ways unoptimized, since the primary focus of this work was as a demonstration of entanglement in m qubits. Beyond the optimization of state detection parameters for FCs, it would in principle be possible to rework the leakage detection scheme

such that the FC and gate operations are carried out in parallel (i.e. the cycling transition being continuously monitored during gate operations such that the ion will begin to fluoresce as soon as a decay event happens). While an optimized FC in this detection scheme would still need to extend past gate operations by some amount to ensure that leakage events towards the end of the gate sequence were detected, an approach of continuous monitoring would allow error overhead from FC to be reduced without needing to improve detection efficiencies.

4.3.4 Outlook. Given the full discussion about gate errors above, there are two points that can be made in summary:

Viability of entangling gates in m qubits? We have shown that reasonably high-fidelity entangling gates can be carried out in metastable trapped-ion qubits, with the main error contribution to our gate being decoherence caused by slow gate speeds, which in turn come down to inefficiencies arising from decisions made to address other experimental constraints. Given the same laser system and the same ion species, this same gate could be run over an order-of-magnitude faster, with a corresponding reductions in leakage error and decoherence.

Vulnerability of m qubits to leakage errors? Compared to g qubits, m qubits are prone to leakage error. However, this leakage error is easily detected by heralding deshelling with fluorescence, and these heralded leakage errors (erasure errors) are in principle much easier to correct than other error types. This means that 1) the limited lifetime of m qubits may be less problematic than it appears and that 2) in fact, the inherent conversion of most SRS errors in m qubits into leakage errors may be a major advantage over g qubits, reducing the impact of an error source fundamental to Raman gates.

Table 3. Dominant sources of two-qubit gate infidelity as predicted by our error model. The lower half of the table summarizes the expected contributions due to erasure errors. Values in bold represent infidelity as measured in the gate

	Error source	Infidelity ($\times 10^{-4}$)	
Non-erasure	Motional dephasing	55	
	Spin dephasing	26	
	π -time calibration	6.8	
	Raman intensity drift	4.8	
	Mode frequency drift	1.9	
	Raman scattering to $D_{5/2}$	0.4	
	Total	97.5	86(10)
Erasure	$D_{5/2}$ lifetime	28.8	
	Raman scattering	13.5	
	854 lightshift scatter	5.6	
	Total	47.9	55(3)

CHAPTER V

3D-PRINTED ION TRAPS

The previous chapter considered one important facet of trapped-ion quantum computing: high-fidelity control of the internal states of trapped ions, particularly in novel qubit encodings with features that may be advantageous for error correction. Control techniques like these are largely independent of the details of the hardware platforms used to implement them (e.g. Raman gates can be performed on ions trapped in 3D macroscopic traps like the ones we used or in 2D microfabricated traps). However, for these techniques to be practically useful, they require hardware that can meet certain performance criteria, as touched on in Section 1.7.1 and discussed more thoroughly below. The purpose of the project discussed in this chapter was to develop trap designs and fabrication methods that can satisfy these criteria without the trade-offs required in the space of existing designs. Notably, in the time since this work was originally put out, additional parallel work has been done on not only fabricating 3D printed traps, but in characterizing their performance as functional traps (Xu et al. (2023)) and beginning to test geometries that could be potentially scalable.

As stated in Section 1.1, this chapter is directly adapted from Quinn et al. (2022), with (to reiterate) the portions of this paper focusing on theory and simulation being written by me, David Allcock writing portions of the intro and conclusion, and Morgan Brown writing the description of the 3D printing and metallization process. Morgan Brown and David Miller carried out the fabrication and characterization of test prints, and I simulated and analyzed the designs discussed in this chapter.

5.1 Introduction

The surface-electrode ion trap (SET, Chiaverini et al. (2005)) was a major step forward for realizing the trapped ion quantum computer when it was first demonstrated in 2006 (Seidelin et al. (2006)). The SET allows for the trap electrodes to be fabricated as a single layer on a wafer. This has allowed for a diverse range of traps to be easily fabricated with rapid turnaround in relatively simple university cleanrooms as well as with commercial MEMS and CMOS processes (Blain et al. (2021)). Underneath this top electrode layer can be many other functional layers containing all the required classical control elements, as shown schematically in Fig. 5.1. These include current-carrying wires for magnetic field-driven gates (Allcock et al. (2013); Warring et al. (2013)), photonic integrated circuits for laser delivery (Ivory et al. (2021); Mehta et al. (2016, 2020)), superconducting nanowires (Slichter et al. (2017); Todaro et al. (2021)) or avalanche photodiodes (“High-fidelity ion state detection using trap-integrated avalanche photodiodes” (n.d.); Setzer et al. (2021)) for photon detection, active CMOS electronics such as DACs for trap voltage generation (Stuart et al. (2019)), and routing layers and thru-wafer vias (Guise et al. (2015)) to external connections. Whilst the complete set of control elements required for QIP has yet to be integrated into a single device, they have all been demonstrated separately and efforts to combine them are underway.

3D ‘wafer’ traps on the other hand typically require thru-wafer machining (Wilpers, See, Gill, and Sinclair (2012)), and usually alignment and stacking of multiple wafers, to achieve the desired geometries. These processes generally preclude the CMOS-like monolithic approach to integrating functional elements, and instead relies on further serial assembly of bespoke and heterogeneous components that requires much new process development (Day et al. (2021));




Figure 5.1. Cross-section of a trap fabricated using our proposed method of 3D printing an electrodes onto a wafer containing electrode routing and some of the potential integrated ion control elements that have been demonstrated.

Ernzer (2018); <https://cordis.europa.eu/project/rcn/218553/factsheet/en> (n.d.); <https://www.translume.com/index.php/item/239-ion-traps-with-integrated-optics> (n.d.)). The recently developed approach of combining wafer stacking and SETs (Auchter et al. (2022)) potentially offers a solution to some of these issues.

Despite these fabrication disadvantages, wafer traps are still under active development because SETs have several major drawbacks (Blakestad et al. (2009)), specifically:

- Trapping potentials are less harmonic. This causes any slight displacement of the trapping location due to stray charges or control voltage inaccuracies to shift the trap frequencies. Without recalibration this will lead to errors in multi-qubit gates. The increased cross-Kerr nonlinearity can also lead to gate errors due to spectator modes (Nie, Roos, and James (2009); Roos et al. (2008)).
- Trap depths are much lower (Chiaverini et al. (2005)), typically close to room temperature. This makes loading ions less efficient and ion losses due to

control field errors, stray charges, and background gas collisions are much more likely. Room temperature operation of more than a couple of ions in lighter ion species is effectively precluded due to short ion lifetimes.

- Junctions between ion traps with favourable trapping strength and depth, and small pseudopotential barriers are harder to achieve than with a 3D geometry (Wesenberg (2009)).
- As the ions reside in an open volume above a surface, they are poorly shielded from electric charges or crosstalk (photons, microwave control fields, or electric trapping fields) from other trapping zones on the SET.

In this paper we propose mitigating these issues by adding a 3D-printed dielectric structure onto the trap wafer before depositing the final metal electrode surface. This allows us to retain the fabrication advantages of the SET whilst also retaining the advantageous trapping geometry of a 3D trap. This requires finding trap geometries that meet these requirements. In section 5.2 we lay out how to quantify the performance of a given geometry. In section 5.3 we propose several feasible geometries and simulate them. Then finally, in section 5.4 we identify a suitable 3D printing process and print demonstration devices.

5.2 Quantifying Trap Performance

As alternatives to existing microfabricated trap designs, we explore several possible ‘trench’ geometries, discussed in greater detail in the following section. To assess the viability of these designs, we simulate them with COMSOL Multiphysics using the boundary element method (BEM). The Electrode software package (NIST (2017)) is used for further post-processing of the electric fields extracted from the simulations.

In each trap simulation we set RF electrodes to unit voltage and DC electrodes to ground and calculate the electric field generated in the vicinity of the ion. From these fields, we can calculate the pseudopotential of the trap, which is the effective potential seen by the ion (Dehmelt (1968)). By finding where this pseudopotential is at a minimum in the plane perpendicular to the trap axis, we locate the equilibrium position of the ion in two dimensions. By also finding the saddle point of the pseudopotential (ie, the ‘escape point’) and comparing the pseudopotential at this point to the pseudopotential at the center of the trap, we have a rough proxy for the depth of the trap (though in reality the situation is more complex due to the dynamical nature of a Paul trap). At the ion’s position, we fit the surrounding electric potential $V = V(r, \theta)$ to a cylindrical harmonic expansion

$$V(r, \theta) = V_0 \sum_{n=2}^{\infty} C_n \left(\frac{r}{r_0} \right)^n \cos(n\theta + \phi_n) + V_{off} \quad (5.1)$$

where r is distance from the trap center, r_0 is the minimum ion-electrode separation, V_{off} is the potential at the ion, ϕ_n is the orientation of the n th-order multipole, and C_n is the expansion coefficient associated with an n th-order multipole. From these expansion coefficients, we define what we call the ‘quadrupole,’ ‘hexapole,’ and ‘octopole’ strengths generated by our traps,

$$\begin{aligned} \text{quadrupole, } C_2, \\ \text{hexapole, } C'_3 = \frac{C_3}{C_2}, \\ \text{octopole, } C'_4 = \frac{C_4}{C_2}, \end{aligned}$$

which we illustrate schematically in Fig. 6.4.

The quadrupole term gives the geometric efficiency of the trap. Greater efficiency is desirable as it reduces the required RF voltage, and thus power

dissipation as well as the possibility of voltage breakdown on the chip. The hexapole and octupole terms give the leading order anharmonicities in the potential and are normalized to the quadrupole term as it is most useful to compare traps at the same ion secular frequency, not the same RF voltage. Anharmonicities increase cross-Kerr nonlinearity and introduce displacement-dependent frequency shifts in the ion's motion, increasing sensitivity of gate fidelity to the temperature of non-gate modes and the presence of stray electric fields. Therefore, we seek designs which minimize higher-order multipole terms around the trap center.

In addition to trap depth and the multipole components of the trapping potential, we also calculate the open numerical aperture (NA) available for fluorescence collection and for laser beam access, as illustrated in Fig. 5.2. A high NA (typically ~ 0.4) is important for fast state readout on the ion. For the novel designs discussed in Sec. 5.3, we give an NA from above and from the substrate, as the ion detectors could be located in the substrate (Setzer et al. (2021); Slichter et al. (2017); Todaro et al. (2021)), or placed above the trap. For SETs the NA is trivially one, while our wafer traps, by their symmetry, have the same NA from both sides. In all cases we consider a circular detector, but a modest improvement could be achieved with a rectangular detector. These NA also gives us an idea of the range of angles available for laser beams perpendicular to the trap axis such that they can exit the trap without striking an electrode.

We scale all of our traps such that the closest electrode to the ion is $75\ \mu\text{m}$ away. Because anomalous heating has a strong $\sim d^4$ dependence (King et al. (1998); Turchette et al. (2000)), the closest electrode will tend to dominate the heating rate. For calculating trap depths we assume a $^{40}\text{Ca}^+$ ion trapped at a secular frequency of 4 MHz by a 40 MHz RF trap drive with no DC potentials applied.

5.3 Trench Trap Geometries

3D printing offers great geometric versatility in trap design, since it can be used to produce arbitrary shapes. In principle, we could 3D print electrodes with hyperbolic cross-sections. Such electrodes would generate pure quadrupole potentials, with lower anharmonicities (as discussed in Sec. 5.2) than other geometries. In practice however, our choices of geometry are limited by other fabrication and operation constraints, namely whether the geometry allows optical access to the ion for laser beam delivery and fluorescence detection and how easily the 3D-printed dielectric can be metalized (a problem discussed further in Sec. 5.4). These constraints naturally suggest what we refer to as a ‘trench’ geometry, illustrated in Fig. 5.1. Trench geometries confine ions between two walls 3D printed on the surface of a wafer. We explore several such ‘trench’ geometries for microfabricated ion traps and compare their performance to SETs and wafer traps of similar dimensions. Cross-sectional views of the geometries simulated are shown in Fig. 5.2.

The trap geometries we considered can be classified as either symmetric or anti-symmetric. Both sets of trap geometries are mirrored around a vertical plane along the trap axis, but in the anti-symmetric traps the role of RF and DC electrode is switched in the mirror half of the trap. Obviously many asymmetric variations on these traps are possible, usually with properties somewhere between the symmetric or anti-symmetric ones. We do not consider them here though as symmetries are generally desirable, and pragmatically we need to keep the parameter space of different traps to simulate tractable.

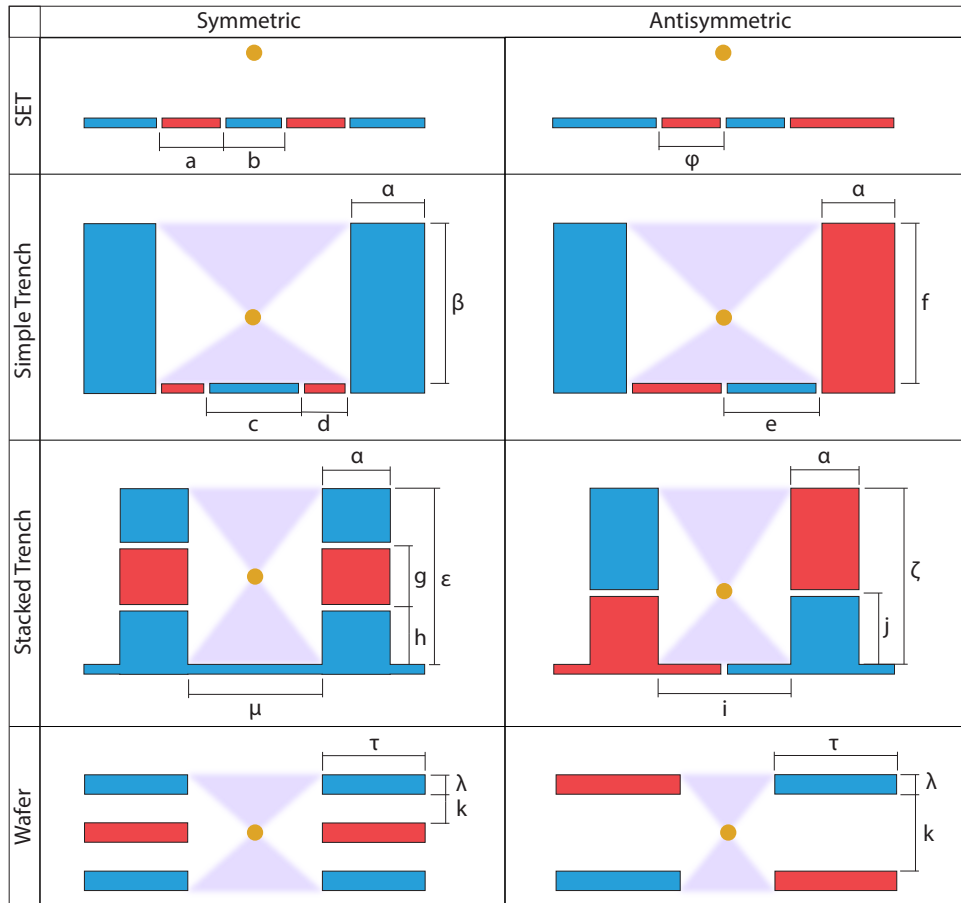


Figure 5.2. Cross-sectional views of the trap designs simulated. The ions' location is shown as an orange dot and the trap axis is out of the page. DC electrodes are blue, RF electrodes are red. Numerical apertures are shown in lavender.

As shown in Fig. 5.2, all simulated traps aside from the SET are characterized by ≥ 3 geometric parameters. We select a subset of parameters (marked with Latin letters) to sweep while fixing the other parameters (marked with Greek letters). For our trench traps, we also investigate how performance changes with the more relevant of the fixed parameters. We do so by making multiple sweeps over the continuous parameters for different values of a fixed parameter.

5.3.1 Simple Trench. The first trap geometry we consider is the ‘simple trench’. In this design, the two outermost electrodes of the SET are extended vertically to above the height of the ion. The symmetric simple trench is primarily characterized by four dimensions, DC and RF electrode widths c and d , trench wall height β , and trench wall thickness α . This last dimension does not strongly affect the potential seen by the ion, since the ion is shielded from all but the inner surfaces of the trench. Because of this, a fixed value $\alpha = 100 \mu\text{m}$, comparable to values used in our test fabrications (see section 5.4), is used in all trench trap simulations.

We simulate a set of symmetric simple trench traps at a fixed ion height, allowing c and d to vary for discrete values of β . Unlike the symmetric design, the anti-symmetric simple trench can be characterized with just two parameters, plane electrode width e and trench wall height f . (We find that ion height in the anti-symmetric trench depends mainly on e , as shown in Fig. 5.3(f). In fact, past $f \approx 200 \mu\text{m}$, ion height shows no dependency on f noticeable beyond numerical noise.) The full simulation results for our symmetric and anti-symmetric simple trench traps are shown in Fig. 5.3.

5.3.2 Stacked Trench. The simple trench geometry discussed above is closer to the ideal quadrupole geometry than a surface trap, but the flexibility of this geometry is limited by the fact that each wall can only contain one electrode. As an alternative, we consider what we call the ‘stacked trench’ geometry. By allowing multiple electrodes to make up the trench wall, we can produce arrangements that more closely match the boundary conditions of a quadrupole, though the fact that the trench must remain open on the top prevents any trench geometry from being fully rotationally symmetrical.

The symmetric stacked trench trap can be characterized by four parameters: trench height ϵ , trench width μ , RF electrode height g , and the height of the RF electrode off of the ground plane h . The anti-symmetric stacked trench trap can be characterized by three parameters: trench height ξ , bottom electrode height j , and trench width i .

These traps have two regimes for which ion-electrode separation is $75\ \mu\text{m}$. In the first regime, the ion is $75\ \mu\text{m}$ from the ground plane and the trench walls are further away, while in the second regime, the ion is $75\ \mu\text{m}$ from the trench walls. We consider both regimes in the antisymmetric case, but only the second in the symmetric case, since performance in the other regime was notably worse.

5.3.3 Comparison Geometries - SET and Wafer Traps. As a point of comparison, we simulate common varieties of microfabricated trap: the SET and the wafer trap. We first simulate the symmetric and anti-symmetric SET (also known as 5-wire and 4-wire SETs respectively). One nice feature of SETs is that under the assumption of an infinitely large trap with no gaps between electrodes, simple closed-form expressions for the potential are available (House (2008)). While we use these closed-form expressions to calculate the parameters shown in Figs. 5.3-5.5, we also simulate our SETs in COMSOL as a way of benchmarking our BEM simulations. We compare the multipole expansions, trap heights, and trap depths of BEM-generated and analytic potentials for the symmetric SET. For most of the free parameter range, agreement across these parameters is within 2%. In the range where $a \ll b$ (see Fig. 5.2), divergence approaches 10% due to the fact that the $1 \mu\text{m}$ electrode spacing present in all BEM simulations is no longer negligible (Schmied (2010)). For a fixed ion height there is one free parameter in the symmetric SET and no free parameters in the anti-symmetric one.

To compare our proposed trench designs against the most prominent set of current 3D microfabricated designs, we also simulate symmetric and anti-symmetric wafer traps, illustrated in the bottom row of Fig. 5.2. A wafer thickness of $\lambda = 50 \mu\text{m}$ of our simulated traps is comparable to those of designs implemented in the lab (Jost (2010); Kienzler (2015)). The depth $\tau = 1 \text{ mm}$ of the wafers is selected to be much larger than wafer thickness and spacing so that the exact depth is unimportant. Therefore, both designs have only one free parameter, the spacing k between wafers in a stack. This parameter is swept over while holding the ion-electrode distance constant.

	SET		Simple		Stacked		Wafer	
	Sym	Anti	Sym	Anti	Sym	Anti	Sym	Anti
Depth (eV)	0.06	0.07	0.08	0.23	0.33	0.22	0.16	0.39
Quadrupole (C_2)	0.17	0.17	0.18	0.24	0.31	0.40	0.35	0.39
Hexapole (C'_3)	1.0	1.0	0.86	0.62	0.020	0.008	0.000	0.001
Octopole (C'_4)	0.75	0.75	0.55	0.36	0.024	0.407	0.344	0.007

Table 4. Summary of results for trap parameter calculations for the SET, simple trench trap (‘Simple’), stacked trench trap (‘Stacked’), and wafer trap. Values shown correspond to a set of geometric parameters selected to balance trap depth and quadrupole strength/purity.

Trap		Dimensions (μm)
SET	Sym	$a = 161.2, b = 59$
	Anti	$\phi = 75$
Simple Trench	Sym	$d = 77.3, c = 210, \beta = 600$
	Anti	$e = 135.2, f = 525$
Stacked Trench	Sym	$g = 140, h = 80, \epsilon = 300$
	Anti	$i = 150, j = 160$
Wafer	Sym	$k = 50$
	Anti	$k = 53$

Table 5. The dimensions (in microns) of the traps whose parameters are summarized in Table 4.

5.3.4 Summary. Having generated the data shown above, we want to compare the simple and stacked trench trap designs with established SET and wafer trap designs and assess their relative merits. For easy comparison between these four sets of simulated geometries, we select a ‘representative’ trap from each category. These traps in some qualitative way optimize the balance between trap depth and quadrupole strength/purity. This selection is somewhat subjective, and optimum parameters may vary by application, but they are unlikely to differ greatly from what we have chosen here. A table of values for these representative traps are shown in Table 4, and the associated dimensions given in Table 5.

From Table 4, we predict that all our trench geometries will outperform SETs. For the symmetric simple trench trap, the improvement is practically

negligible though. The antisymmetric simple trench is a significant improvement over the SET, offering $\sim 3\times$ the depth, $\sim 40\%$ higher quadrupole strength, and $\sim 40\%$ lower hexapole strength. Despite this it is still notably inferior to the wafer traps. The stacked trench traps on the other hand can have very comparable performance to wafer traps, with similar depth, quadrupole strength, and octopole terms. The hexapole term for our representative trap is still larger than in wafer traps (where it can be zero by symmetry), but the factor ~ 50 improvement over SETs will still lead to drastic improvement in trap performance. Further improvement can be gained at the cost of lower NA optical access by making the walls higher. As some point though manufacturing tolerances and curvature from stray fields will set a practical limit for either design.

5.4 3D Printing

Our chosen fabrication process for realizing these trench geometries is 2-photon direct laser write (DLW) 3D printing. The electrode structure is printed directly onto a functional trap wafer using a dielectric material. The printing process uses a focused laser to polymerize this photosensitive material to generate solid structures (Pearre, Michas, Tsang, Gardner, and Otchy (2019)). This can then be metallized to create conductive electrode surfaces.

Whilst not explored here, we note that the simple trench traps could be fabricated using UV-LIGA (Ultraviolet - Lithographie, Galvanoformung, Abformung). This process is already widely used to make electrodes for SETs (Arrington et al. (2013)) and so would be straightforward to adopt. UV-LIGA is a fabrication process where metal is electroplated inside a photoresist mold patterned using UV lithography. As discussed above, the performance advantage of these designs over SETs is fairly modest compared to the stacked trench traps. The stacked trench geometries could also be built using conventional MEMS techniques.

Relative permittivity	2.5 Fröhlich et al. (2002) 2.7 Laine (1992) 3.1 Johansson et al. (2003)
Loss tangent	0.0035 Johansson et al. (2003) 0.00397 Laine (1992) 0.004 Fröhlich et al. (2002),
Dielectric strength (V/cm)	8.7×10^5 Laine (1992)
Bulk resistivity ($\Omega \cdot \text{cm}$)	4.5×10^{16} Laine (1992)
Thermal conductivity (W/m·K)	2.3 Laine (1992)

Table 6. Electrical and thermal properties of ormocer.

However, the 100s of μm vertical heights required are likely challenging to achieve with oxide film deposition.

5.4.1 Dielectric Material. We propose using ormocer (**organically modified ceramic**), a silica-based material, as our printable dielectric material. The high RF frequencies ($\sim 10\text{-}100$ MHz) and voltages ($\sim 10\text{-}1000$ V) applied to ion trap electrodes places stringent requirements on the electrical and thermal properties of any dielectrics used. High dielectric strength is required to avoid breakdown. High resistivity and low loss tangent avoid direct heating of the dielectric. A low dielectric constant reduces parasitic capacitance, requiring less current to drive the electrodes, and thus lower Ohmic losses in the electrode metal. Finally, a high thermal conductivity ensures any heating from losses is minimized. Ormocer properties are somewhat formulation dependent, but in Table 6 we give some representative values measured in the literature to give an idea of what can be achieved.

For room temperature operation, outgassing needs to be low to ensure ultra-high vacuum pressures can be reached and the material must be able to withstand $\sim 200^\circ\text{C}$ vacuum bakes. Low outgassing is also desirable as it reduces the chance of contaminating metal electrode surfaces, which are known to play a

role in anomalous heating (Hite et al. (2013)). The organic cross-linked structure of ormocer reduces outgassing and makes ormocer fit for vacuum processes (*Dielectric ORMOCERs for System-in-Package Electronics and Highly Integrated Systems* (n.d.)), although quantitative outgassing data is so far lacking. Ormocer is stable to $> 300^{\circ}\text{C}$ (*Dielectric ORMOCERs for System-in-Package Electronics and Highly Integrated Systems* (n.d.)). For cryogenic operation, vacuum compatibility is much less of an issue but the structure must be able to survive repeated temperature cycles without structural failure, something that will need testing in future work.

5.4.2 Initial Fabrication.

5.4.2.1 Print. As an initial test, ormocer based trap structures were printed onto a borosilicate glass wafer and metalized with ~ 20 nm of sputtered gold. SEM images of which are shown in Fig. 5.6. The printed structures are $200\ \mu\text{m}$ high with $100\ \mu\text{m}$ wide trenches in between them, giving a $50\ \mu\text{m}$ ion-electrode distance. This is a fairly typical value, with most microfabricated traps historically lying in the few-10s to few-100s μm range. These traps were fabricated via a two-photon direct laser write process. A custom printer was used, based on open source concepts from the Janelia’s open source Modular In vivo Multiphoton Microscopy System (MIMMS) and the printer described in (Pearre et al. (2019)). Recent work has established the viability of mechanically and electrically integrating these prints with traditional silicon wafer processing (Brown et al. (2023)).

This printer was built to address specific needs in terms of print speed and full wafer scalability while retaining the necessary resolution. In the print process, an individual 2D field of view (FOV) is scanned using a galvo for y axis steps paired with a fast resonant imaging mirror operating at 8 kHz in the x axis; allowing for printing at speeds up to ~ 8000 mm/s. A piezoelectric actuator is then

used to sweep this FOV in the normal z axis to form an individual metavoxel. Series of these are then stitched together across the entirety of the wafer via a controllable hexapod (Physik Instrumente) with sub-micron precision. We chose to break each “island” into 4 metavoxels with a $200\ \mu\text{m} \times 200\ \mu\text{m}$ FOV. Each individual island is printed in under a minute with a voxel resolution under $2\ \mu\text{m}$ in the xy print plane (142 x and 512 y voxels) and under $3\ \mu\text{m}$ in z . Demonstrating that, whilst this is an inherently serial process, the overall processing times at the required resolution are tractable for much larger traps than are currently in use.

A hybrid resist is used in this printing process, based on Ormocomp photoresist (Microchem). To this we add a photoinitiator, 2,4,6-trimethylbenzoyl phosphine oxide (TPO) (Sigma), a stabilizing agent, 3,5-Di-tert-butyl-4-hydroxytoluene (BHT), and Fluorescein (Sigma) for in situ imaging during printing. A Chameleon Discovery laser, set to 780 nm, 100 fs pulse width, 80 MHz Rep rate, and $\sim 100\ \text{mW}$ was used to initiate polymerization. Resolution and therefore surface roughness is a function of resin choice as well as the fundamental optics and scan mechanics of the printer. The current design takes advantage of the critical electrical dimensions of the trap being on a vertical surface which minimizes discretization concerns. In addition, a small amount of shrinkage during curing is inevitable. We minimize this by using the highest possible laser power before explosions occur. We do not see delamination from the substrate or noticeable distortion of the structures under SEM imaging.

5.4.2.2 Metalization. The next stage in our fabrication development is metalization. Our trap structure is designed to be ‘self masking’ - the isolation gaps between electrodes can be made deep enough (or with an undercut) such that directional deposition of metal onto the structure will not cause shorting between them. This metalization will connect the trap to vias already pre-fabricated on a planarized wafer using a standard CMOS process (see fig. 5.1). These vias must be off to the side of the printed structures as printing directly onto metal is difficult - the laser heats the metal causing violent disruption of the print. On the substrate around the prints, isolation between electrodes will be achieved using a liftoff process. We envisage using gold as our electrode metal as this is well established in ion trapping, and aluminium for the liftoff material. The full process can be seen in Figure 5.7 and will be described in more detail in a forthcoming publication.

5.4.2.3 Integrated Optics. The final step to producing a working trap will be to integrate optics for laser delivery. This is because trench designs do not offer the required optical access for free-space laser delivery. Waveguides for light routing and grating couplers to focus beams onto the ions have now been demonstrated by several groups (Ivory et al. (2021); Mehta et al. (2016, 2020)). These demonstrations were using SETs, but one of the strengths of the traps proposed here is that the same technique can be used without modification. Other integrated features shown in Fig. 5.1, whilst desirable, aren’t an immediate requirement.

5.5 Conclusion

Additive 3D printing has the potential be a powerful enabling technology for trapped-ion quantum computing. It allows the combination of wafer-integrated ion control with the 3D structures required to obtain optimal trap performance. In order to realise this we have developed a novel trench geometry and shown it can

be printed. We lay out a roadmap to integrate the necessary electrical and optical routing to enable functional devices.




Figure 5.3. Simulated parameters of surface-electrode traps vs the symmetric and anti-symmetric ‘simple trench’ trench geometries. (a) Trap depth at a fixed radial frequency. (b) The quadrupole component (C_2) of the trapping potential at the ion. (c) The hexapole component (C'_3) of the trapping potential. (d) The octopole component (C'_4) of the trapping potential. (e) Electrode dimensions at a constant ion-electrode separation for the symmetric simple trench trap. (f) Electrode dimensions at a constant ion-electrode separation for the anti-symmetric simple trench trap. (g,h) NA above and below the ion. Parameters are plotted against a generic geometric variable w , which is related to an electrode dimension of each trap as shown in the legend. Simulation results are shown with markers. The representative SET, wafer, and simple trench traps whose characteristics are summarized in Table 4 are highlighted with larger markers. Simulation points are connected with dashed lines only as a guide to eye.

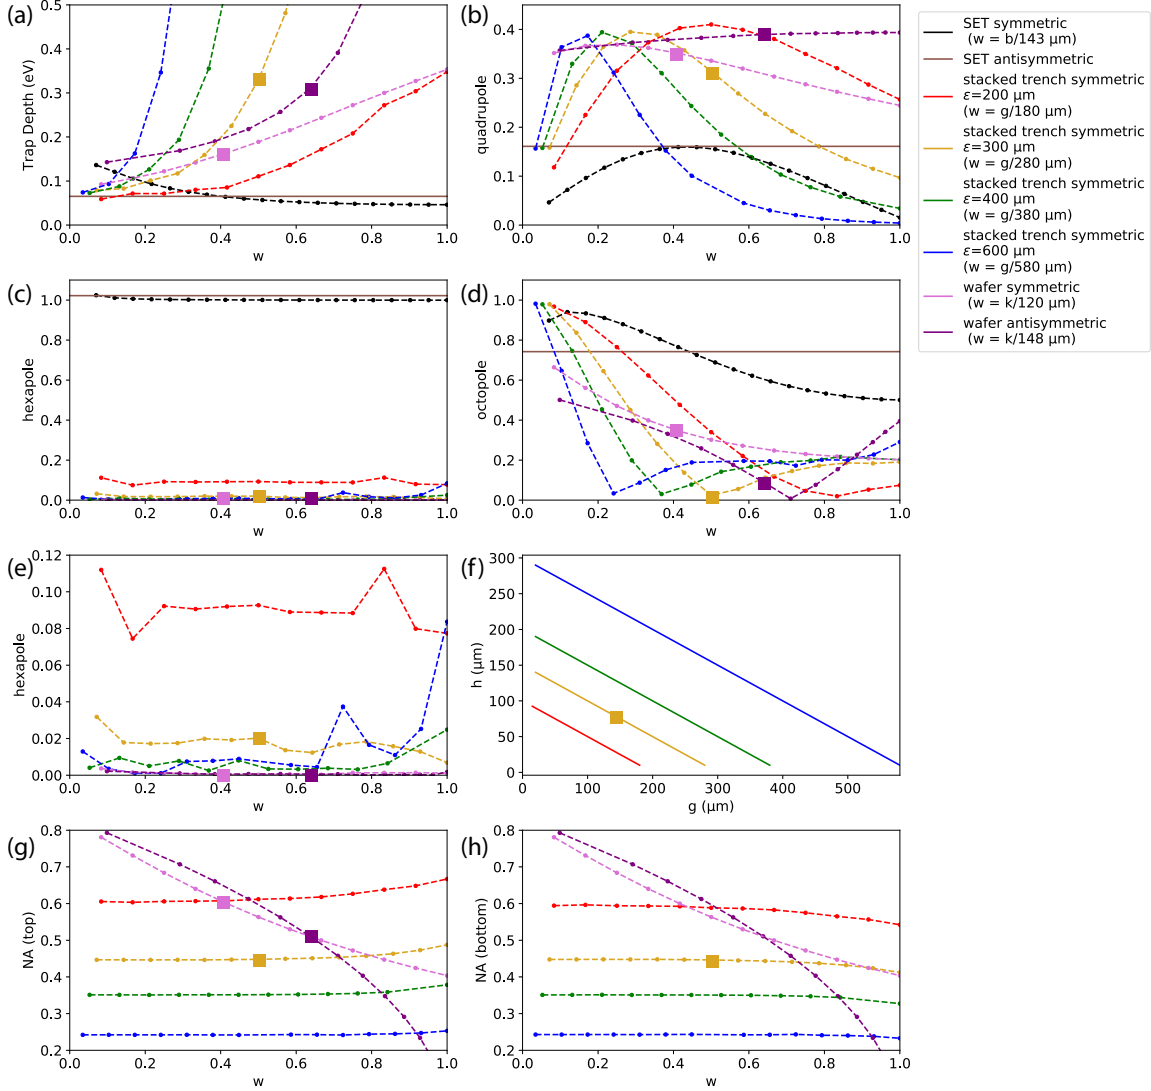


Figure 5.4. Simulated parameters of surface-electrode traps and wafer traps vs the symmetric ‘stacked trench’ trap geometry. (a) Trap depth at a fixed radial frequency. (b) The quadrupole component (C_2) of the trapping potential at the ion. (c) The hexapole component (C_3') of the trapping potential. (d) The octopole component (C_4') of the trapping potential. (e) The hexapole component of the trapping potential, excluding the SET data to make the other series more clear. (f) Electrode dimensions at a constant ion-electrode separation for the symmetric stacked trench trap. (g,h) NA above and below the ion. Parameters are plotted against a generic geometric variable w , which is related to an electrode dimension of each trap as shown in the legend. Simulation results are shown with markers. The representative wafer and symmetric stacked trench traps whose characteristics are summarized in Table 4 are highlighted with larger markers. Simulation points are connected with dashed lines only as a guide to eye.

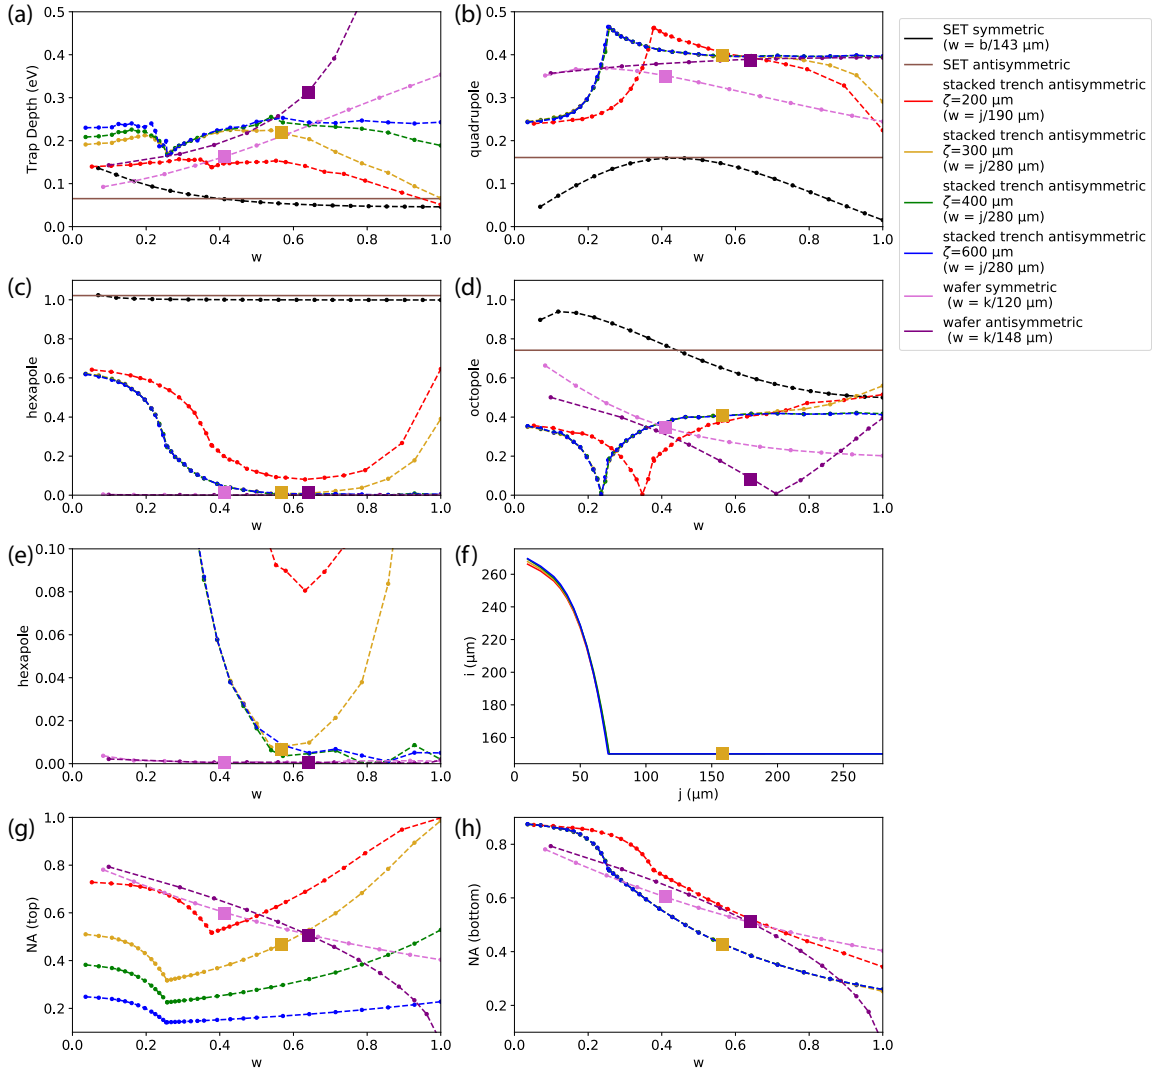


Figure 5.5. Simulated parameters of surface-electrode traps and wafer traps vs the anti-symmetric ‘stacked trench’ trap geometry. (a) Trap depth at a fixed radial frequency. (b) The quadrupole component (C_2) of the trapping potential at the ion. (c) The hexapole component (C'_3) of the trapping potential. (d) The octopole component (C'_4) of the trapping potential. (e) The hexapole component of the trapping potential, excluding the SET data to make the other series more clear. (f) Electrode dimensions at a constant ion-electrode separation for the anti-symmetric stacked trench trap. (g,h) NA above and below the ion. Parameters are plotted against a generic geometric variable w , which is related to an electrode dimension of each trap as shown in the legend. Simulation results are shown with markers. The representative wafer and anti-symmetric stacked trench traps whose characteristics are summarized in Table 4 are highlighted with larger markers. Simulation points are connected with dashed lines only as a guide to eye.




Figure 5.6. SEM images of gold coated ormoer trap structures printed on glass. In this structure, many trenches have been connected up with X-junctions, as is envisaged in a large ‘QCCD’ processor (Blakestad et al. (2009)). (a) An overhead view of a grid of symmetric stacked trenches showing segmentation of the top DC electrodes. (b) An oblique view of the same grid. (c) A close up view of an intersection of trenches. (d) A false-color version of (c) showing DC electrodes in blue and RF electrodes in red. (e) A close-up view of a DC electrode segment.




Figure 5.7. Proposed integration of electrode metalization with printing process shown here on a short section of the larger trap shown in Fig 5.6. (a) Silicon wafer with electrode routing and vias up to the surface is fabricated using standard CMOS process (metal shown copper coloured, SiO₂ dielectric shown transparent). (b) Aluminum is deposited and patterned to define gaps between electrodes. (c) DLW used to print ormozer trap structures (blue). (d) Gold is deposition from multiple angles to coat all surfaces of the print. (e) Aluminium is etched to allow liftoff of the gold over the gaps between electrodes. Insets shown the liftoff process in more detail.

CHAPTER VI

SURFACE ELECTRODE TRAPS FOR CVQC

The final project laid out in this document in some sense connects the themes of the preceding two chapters. The goal of this work was to design ion traps that could be used to perform non-Gaussian motional operations while sticking within the conventional framework of scalable trap design, that being surface electrode traps (SETs). This means that in the design process outlined, both efficient coherent control and trap performance need to be accounted for and balanced against each other. These considerations, and the details of the design process, are all discussed in greater detail below.

6.1 Introduction

Gaussian operations have been applied to the motional modes of a trapped ion simply by applying voltages to the trap electrodes to generate the required potential at the ion. A potential gradient implements operations to the first order in the lowering and raising operators, \hat{a}_i and \hat{a}_i^\dagger . For example, the displacement operation is performed when an oscillating gradient is applied at a frequency resonant with the mode (Alonso et al. (2016); Carruthers and Nieto (1965); Meekhof, Monroe, King, Itano, and Wineland (1996)). Similarly, Gaussian operations second order in \hat{a}_i and \hat{a}_i^\dagger are performed using quadratic potentials modulated at different frequencies: a phase shift is achieved with a static potential (Leibfried et al. (2002)), squeezing with one oscillating at twice the motional frequency (Burd et al. (2019); Heinzen and Wineland (1990)), and the beamsplitter with one oscillating at the difference frequency between two mode (Cornell, Weisskoff, Boyce, and Pritchard (1990); Gorman et al. (2014); D. J. Wineland and Dehmelt (1975)). Non-Gaussian operations, required for universal continuous variable quantum computing (CVQC), must be third order or higher in \hat{a}_i and

\hat{a}^\dagger and consequently are performed using potentials that are cubic or higher-order. A non-Gaussian operation that is third order has yet to be demonstrated in this manner, in part because the cubic potential required is more challenging to generate. In this work, we investigate what kind of non-Gaussian operations can be performed with a cubic potential and design an ion trap optimized to implement them.

This chapter is organized as follows. In Section 6.2, we discuss non-Gaussian operations that can be performed on a single ion with a cubic potential. The analysis is repeated in Section 6.3 for a two-ion system. A novel surface-electrode trap design capable of generating cubic potentials in the radial plane is described in Section 6.4. The viability of these trap designs, i.e. the coupling strengths achievable for different non-Gaussian operations with reasonable trap parameters, is discussed in Section 6.5.

6.2 Single ion in a cubic potential

Non-Gaussian operations can be carried out on the vibrational mode of a single trapped ion by applying a modulated cubic potential. We begin our analysis of this situation by considering the associated Hamiltonian. From there, we demonstrate that three non-Gaussian operations can be performed using this Hamiltonian, including the well-known cubic phase gate and trisqueezing.

6.2.1 Hamiltonian. The Hamiltonian for a single vibrational mode \hat{x} of a trapped ion with an externally applied cubic potential oscillating at modulation frequency ω_p is

$$\hat{H} = \frac{\hat{p}^2}{2m} + \frac{1}{2}m\omega_x^2\hat{x}^2 - \hbar g \sin(\omega_p t - \theta) \left(\frac{\hat{x}}{x_0} \right)^3, \quad (6.1)$$

where ω_x is the selected motional mode frequency, g is the coupling strength (calculated in terms of trap parameters in Section 6.4), θ is the modulation phase,

and $x_0 = \sqrt{\frac{\hbar}{2m\omega_x}}$ is the ground state wavefunction extent of mode \hat{x} . We can rewrite this Hamiltonian in terms of raising and lowering operators $\hat{x} = x_0(\hat{a}^\dagger + \hat{a})$ and $\hat{p} = i\hbar/2x_0(\hat{a}^\dagger - \hat{a})$

$$\hat{H} = \hbar\omega_x \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) - \hbar g \sin(\omega_p t - \theta) (\hat{a}^\dagger + \hat{a})^3. \quad (6.2)$$

The expression $(\hat{a}^\dagger + \hat{a})^3$ can be expanded and rewritten using the commutation relation $[\hat{a}, \hat{a}^\dagger] = 1$ as

$$\begin{aligned} (\hat{a}^\dagger + \hat{a})^3 &= \hat{a}^{\dagger 3} + \hat{a}^3 \\ &\quad + \hat{a}^{\dagger 2} \hat{a} + \hat{a} \hat{a}^{\dagger 2} + \hat{a}^\dagger \hat{a} \hat{a}^\dagger \\ &\quad + \hat{a}^\dagger \hat{a}^2 + \hat{a}^2 \hat{a}^\dagger + \hat{a} \hat{a}^\dagger \hat{a} \\ &= \hat{a}^{\dagger 3} + \hat{a}^3 + 3\hat{a}^\dagger \hat{a} \hat{a}^\dagger + 3\hat{a} \hat{a}^\dagger \hat{a} \end{aligned} \quad (6.3)$$

to give the Hamiltonian

$$\hat{H} = \hbar\omega_x \left(\hat{a}^\dagger \hat{a} + \frac{1}{2} \right) - \hbar g \sin(\omega_p t - \theta) (\hat{a}^{\dagger 3} + \hat{a}^3 + 3\hat{a}^\dagger \hat{a} \hat{a}^\dagger + 3\hat{a} \hat{a}^\dagger \hat{a}). \quad (6.4)$$

This Hamiltonian can be rewritten in the interaction picture as

$$\begin{aligned} \hat{H}_I &= -\hbar g \sin(\omega_p t - \theta) (\hat{a}^{\dagger 3} e^{3i\omega_x t} + \hat{a}^3 e^{-3i\omega_x t} \\ &\quad + 3\hat{a}^\dagger \hat{a} \hat{a}^\dagger e^{i\omega_x t} + 3\hat{a} \hat{a}^\dagger \hat{a} e^{-i\omega_x t}). \end{aligned} \quad (6.5)$$

The interaction Hamiltonian above can be used to carry out three operations on mode \hat{x} , one with a static field ω_p , one at $\omega_p = \omega_x$, and one at $\omega_p = 3\omega_x$. These three cases are discussed below.

6.2.2 The $\omega_p = 0$ case (cubic phase gate). The cubic phase gate $\hat{U}_{cubic} = e^{i\gamma x^3}$ is one possible non-Gaussian resource for universal CVQC. Implementations of the cubic phase gate have been proposed for photonic systems, including putting a squeezed and displaced state through a Kerr nonlinearity (Yanagimoto et al. (2020)) and the squeezing of an intermediate trisqueezed state (Zheng et al. (2021)).

The single-ion Hamiltonian presented in the previous section, when $\omega_p = 0$ (i.e. when the cubic potential is static), can be written in the interaction picture

$$\hat{H}_I = -\hbar g \left(\frac{\hat{x}}{x_0} \right)^3 \quad (6.6)$$

where \hat{x} , being in the interaction picture, is implicitly time-dependent. The effect of applying this Hamiltonian to the motional ground state is shown in Fig.

6.2.3 The $\omega_p = \omega_x$ case. If $\omega_p = \omega_x$, then, after applying the rotating wave approximation (RWA), the interaction Hamiltonian Eq. (6.5) becomes

$$\hat{H}_I = i\hbar \frac{3g}{2} (-\hat{a}^\dagger \hat{a} \hat{a}^\dagger e^{i\theta} + \hat{a} \hat{a}^\dagger \hat{a} e^{-i\theta}). \quad (6.7)$$

To the best of our knowledge, A Hamiltonian of this form has not been described in the literature. In simulations, we found that this interaction Hamiltonian drives the motional ground state into a superposition of odd Fock states, as shown in Fig. 6.1. This is a non-Gaussian state and is therefore a possible resource for universal CVQC. However, this operation may be hard to implement, since any potential proportional to \hat{x} will give a displacement.

6.2.4 The $\omega_p = 3\omega_x$ case (trisqueezing). Trisqueezing is a non-Gaussian operation analogous to squeezing. This operation has been proposed as a tool for generating cubic phase states (Zheng et al. (2021)), a class of states that would enable universal CVQC (Gottesman, Kitaev, and Preskill (2001)). Whilst trisqueezing has been experimentally demonstrated in a superconducting system (Chang et al. (2020)), it has not yet in trapped ions.

If $\omega_p = 3\omega_x$ and the RWA is applied, the interaction Hamiltonian Eq. (6.5) simplifies to

$$\hat{H}_I = i\hbar \frac{g}{2} (-\hat{a}^{\dagger 3} e^{-i\theta} + \hat{a}^3 e^{i\theta}). \quad (6.8)$$

The effect of applying this Hamiltonian to the motional ground state is shown in Fig. 6.2. Analogous to squeezing, which can be accomplished using an

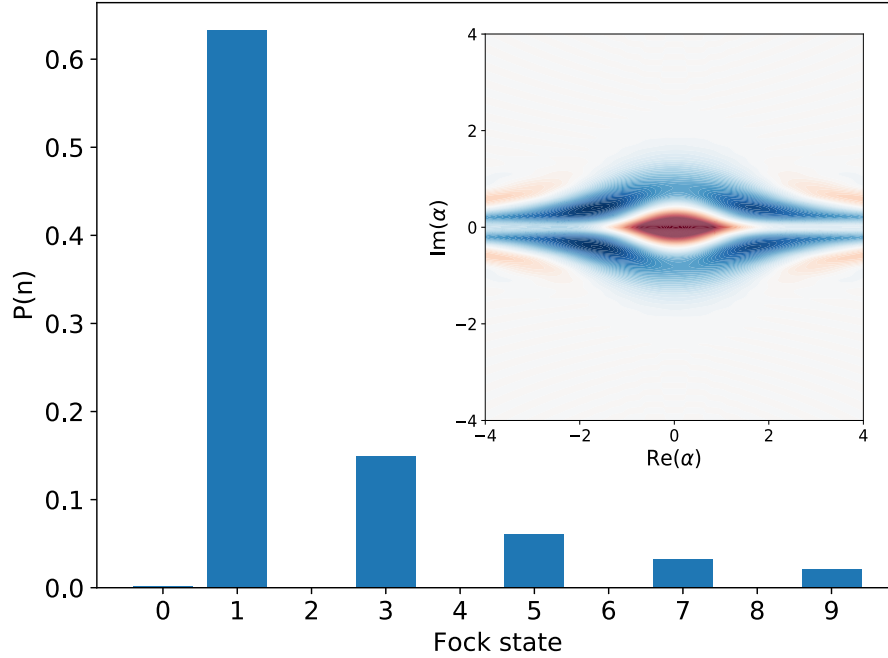


Figure 6.1. Plot of the Fock state distribution of the state generated by applying the $\omega_p = \omega_x$ operation to the motional ground state for a time $t = 0.23 \times 2\pi/g$. Components $n > 9$ not shown. *Inset* - Associated Wigner function.

oscillating quadratic potential (Burd et al. (2019)) and drives the ground state into a superposition of even Fock states ($n = 0, 2, 4, \dots$), trisqueezing drives the ground state into a superposition of every third Fock state ($n = 0, 3, 6, \dots$).

6.3 Two ions in a cubic potential

As established in the previous section, a set of non-Gaussian operations can be carried out on the motional modes of a single trapped ion using a cubic potential. As established in the previous section, a set of non-Gaussian operations can be carried out on the motional modes of a single trapped ion using a cubic potential. With only three modes however, a single ion has limited resources for quantum computing. These resources can be expanded by increasing the number of ions, since a system of N ions has $3N$ normal modes. On top of this, while a single-ion system has only center-of-mass (COM) modes, a multi-ion system




Figure 6.2. Plot of the Fock state distribution of of the trisqueezed state, obtained by applying the trisqueezing Hamiltonian to the ground state for a time $t = 0.035 \times 2\pi/g$. Components $n > 9$ not shown. *Inset* - Associated Wigner function.

has non-COM modes with qualitatively different features. Since operations in practical CVQC will generally be carried out on such systems, we are interested in characterizing the sorts of non-Gaussian operations that can be performed on these systems using a cubic potential.

In this section, we consider a two-ion system ($N = 2$). Such a system will have six normal modes that can be classified into two categories: center-of-mass (COM) modes and out-of-phase (OOP) modes, illustrated in Fig. 6.3. The distinction between COM and OOP modes is relevant to CVQC since, among other reasons, OOP modes have much lower heating rates than COM modes (King et al. (1998); Turchette et al. (2000)). Existing ion traps can hold up to ~ 50 ions. Depending the trap parameters, 1D ion chains or 2D/3D crystals can be formed.

Exploring the third-order operations that can be performed using a cubic potential in the regime of higher N and 2D/3D structures represents an interesting area for future work.




Figure 6.3. *Left* - Illustration of a two-ion chain with coordinates labelled. Equilibrium positions are marked in light blue. *Right* - Table illustrating two-ion vibrational normal modes.

For a single ion, the cubic potential could be applied to any of the three normal modes. Notably, for an ion chain, applying a cubic potential along the chain will introduce displacements. Consequently, we only consider the case of a cubic potential orthogonal to the trap axis.

We demonstrate below that five non-Gaussian operations can be performed on a two-ion system using a radial cubic potential, including parametric oscillation and trisqueezing. Worth noting is that introducing a second ion also introduces a larger set of operations. To investigate these, we begin by considering the Hamiltonian of a two-ion system with a cubic potential applied along the x -axis.

6.3.1 Hamiltonian. For a pair of ions in the same trap, the Hamiltonian can in general be written in terms of these ions' displacement from equilibrium (x_i, y_i, z_i) and separation $(\Delta x, \Delta y, \Delta z)$ as a pair of two three-

dimensional harmonic oscillators coupled by a nonlinear Coulomb interaction

$$\begin{aligned}
\hat{H} &= \frac{\hat{p}_1^2}{2m} + \frac{\hat{p}_2^2}{2m} + \frac{1}{2}m\omega_x^2(\hat{x}_1^2 + \hat{x}_2^2) \\
&+ \frac{1}{2}m\omega_y^2(\hat{y}_1^2 + \hat{y}_2^2) + \frac{1}{2}m\omega_z^2(\hat{z}_1^2 + \hat{z}_2^2) \\
&+ \frac{e^2}{4\pi\epsilon_0} \frac{1}{\sqrt{(\Delta x)^2 + (\Delta y)^2 + (\Delta z)^2}}
\end{aligned} \tag{6.9}$$

where $(\omega_x, \omega_y, \omega_z)$ are the normal mode frequencies for a single ion. Without loss of generality, we can consider the case where the radial frequencies are the same ($\omega_y = \omega_x$). Taking the first- and second-order terms of the Coulomb interaction into account, the above Hamiltonian can be re-expressed in terms of the normal modes of the system, the out of phase (OOP) rocking (q_r) and stretching (q_s) modes and the center of mass (COM) modes, which can be radial (q_c) and axial (q_a)

$$\begin{aligned}
\hat{H} &= 2\frac{\hat{p}_r^2}{4m} + \frac{\hat{p}_s^2}{4m} + 2\frac{\hat{p}_c^2}{4m} + \frac{\hat{p}_a^2}{4m} \\
&+ 2m\omega_r^2\hat{q}_r^2 + m\omega_s^2\hat{q}_s^2 + 2m\omega_c^2\hat{q}_c^2 + m\omega_a^2\hat{q}_a^2 \\
&+ U(\hat{q}_r, \hat{q}_s, \hat{q}_c, \hat{q}_a)
\end{aligned} \tag{6.10}$$

where $\omega_r = \sqrt{\omega_x^2 - \omega_z^2}$, $\omega_s = \sqrt{3}\omega_z$, $\omega_c = \omega_x$, and $\omega_a = \omega_z$ are the mode frequencies for the rocking, stretching, radial COM, and axial COM modes respectively. The term U includes all of the higher-order Coulomb interaction terms. These terms can cause coupling between normal modes, but only at certain normal mode resonances. For example, the nonlinear terms of the Coulomb interaction can generate a parametric oscillator between the stretch and rocking modes at $\omega_s \approx 2\omega_r$ (Ding, Maslennikov, Hablützel, Loh, and Matsukevich (2017)) or a Kerr-type interaction between these same modes (Nie et al. (2009)). These higher-order terms are small when the modes are off resonance, and U can be ignored for most sets of mode frequencies.

If a time-dependent, third-order trisqueezing potential is applied to the system, the Hamiltonian becomes

$$\begin{aligned}
\hat{H} = & 2\frac{\hat{p}_r^2}{4m} + \frac{\hat{p}_s^2}{4m} + 2\frac{\hat{p}_c^2}{4m} + \frac{\hat{p}_a^2}{4m} \\
& + 2m\omega_r^2\hat{q}_r^2 + m\omega_s^2\hat{q}_s^2 + 2m\omega_c^2\hat{q}_c^2 + m\omega_a^2\hat{q}_a^2 \\
& + \hbar g \sin(\omega_p t - \theta) \left[\left(\frac{\hat{x}_1}{x_0} \right)^3 + \left(\frac{\hat{x}_2}{x_0} \right)^3 \right]
\end{aligned} \tag{6.11}$$

where x_0 is the ground state wave function extent for the single-ion \hat{x} mode and g is the single-ion coupling strength (see Section 6.4 for a calculation of g in terms of trap parameters). The radial coordinates x_1 and x_2 of the individual ions can be written in terms of the amplitudes of the radial modes q_c and q_r

$$x_1 = q_c + q_r, \tag{6.12}$$

$$x_2 = q_c - q_r. \tag{6.13}$$

With these relations, the cubic potential that is most naturally expressed in terms of the ion coordinates can be rewritten in terms of the normal mode amplitudes

$$\begin{aligned}
\hat{H} = & 2\frac{\hat{p}_r^2}{4m} + \frac{\hat{p}_s^2}{4m} + 2\frac{\hat{p}_c^2}{4m} + \frac{\hat{p}_a^2}{4m} \\
& + 2m\omega_r^2\hat{q}_r^2 + m\omega_s^2\hat{q}_s^2 + 2m\omega_c^2\hat{q}_c^2 + m\omega_a^2\hat{q}_a^2 \\
& + 2\hbar g \frac{1}{x_0^3} \sin(\omega_p t - \theta) (\hat{q}_c^3 + \hat{q}_c\hat{q}_r^2 + \hat{q}_r\hat{q}_c\hat{q}_r + \hat{q}_r^2\hat{q}_c) \\
\hat{H} = & 2\frac{\hat{p}_r^2}{4m} + \frac{\hat{p}_s^2}{4m} + 2\frac{\hat{p}_c^2}{4m} + \frac{\hat{p}_a^2}{4m} \\
& + 2m\omega_r^2\hat{q}_r^2 + m\omega_s^2\hat{q}_s^2 + 2m\omega_c^2\hat{q}_c^2 + m\omega_a^2\hat{q}_a^2 \\
& + 2\hbar g \frac{1}{x_0^3} \sin(\omega_p t - \theta) (\hat{q}_c^3 + 3\hat{q}_c\hat{q}_r^2),
\end{aligned} \tag{6.14}$$

and the whole expression can be rewritten in terms of the raising and lowering operators of the normal modes

$$\begin{aligned} \hat{H} = & 2\hbar\omega_r(\hat{a}_r^\dagger\hat{a}_r + \frac{1}{2}) + \hbar\omega_s(\hat{a}_s^\dagger\hat{a}_s + \frac{1}{2}) + 2\hbar\omega_c(\hat{a}_c^\dagger\hat{a}_c + \frac{1}{2}) + \hbar\omega_a(\hat{a}_a^\dagger\hat{a}_a + \frac{1}{2}) \\ & + 2\hbar g \sin(\omega_p t - \theta) \left[\frac{q_{c,0}^3}{x_0^3} (\hat{a}_c^\dagger + \hat{a}_c)^3 + 3 \frac{q_{c,0}q_{r,0}^2}{x_0^3} (\hat{a}_c^\dagger + \hat{a}_c)(\hat{a}_r^\dagger + \hat{a}_r)^2 \right]. \end{aligned} \quad (6.15)$$

From this, the interaction Hamiltonian can be written

$$\begin{aligned} \hat{H}_I = & 2\hbar g \sin(\omega_p t - \theta) \left[\frac{q_{c,0}^3}{x_0^3} (\hat{a}_c^\dagger e^{i\omega_c t} + \hat{a}_c e^{-i\omega_c t})^3 \right. \\ & \left. + 3 \frac{q_{c,0}q_{r,0}^2}{x_0^3} (\hat{a}_c^\dagger e^{i\omega_c t} + \hat{a}_c e^{-i\omega_c t})(\hat{a}_r^\dagger e^{i\omega_r t} + \hat{a}_r e^{-i\omega_r t})^2 \right]. \end{aligned} \quad (6.16)$$

Using the relationships between the normal mode frequencies ω_c and ω_r and the single-ion radial and axial frequencies ω_x and ω_z , the coefficients of this interaction Hamiltonian can be simplified,

$$\begin{aligned} \hat{H}_I = & \frac{1}{\sqrt{2}} \hbar g \sin(\omega_p t - \theta) \left[(\hat{a}_c^\dagger e^{i\omega_c t} + \hat{a}_c e^{-i\omega_c t})^3 \right. \\ & \left. + 3 \frac{\omega_c}{\omega_r} (\hat{a}_c^\dagger e^{i\omega_c t} + \hat{a}_c e^{-i\omega_c t})(\hat{a}_r^\dagger e^{i\omega_r t} + \hat{a}_r e^{-i\omega_r t})^2 \right]. \end{aligned} \quad (6.17)$$

Several operations can be performed using this Hamiltonian, as described below.

6.3.2 The $\omega_p = 0$ case. As discussed in Section 6.2.2, the cubic phase gate may be useful tool in universal CVQC. Carrying out this gate on a single ion is trivial in our scheme, since it only requires a application of a static cubic potential. The analogous case for two ions is complicated by interactions between the radial COM and OOP modes. Following the same process as in Section 6.2.2 with the two-ion Hamiltonian gives the propagator \hat{U}_I

$$\hat{U}_I = e^{i2gt(\hat{q}_c^3 + 3\hat{q}_c\hat{q}_r^2)/x_0^3}. \quad (6.18)$$

This propagator does include a term with a cubic dependence on \hat{q}_c . However, for $\omega_c \sim \omega_r$, this propagator will not behave like the cubic phase gate

due to the $\hat{q}_c \hat{q}_r^2$ term coupling \hat{q}_c and \hat{q}_r . However, since $\omega_r = \sqrt{\omega_x^2 - \omega_z^2}$, $\omega_r \rightarrow 0$ for $\omega_x \rightarrow \omega_z$. Given that $\hat{q}_{i,0} \propto 1/\sqrt{\omega_i}$, if the radial frequency is brought close to the axial frequency, $\hat{q}_{r,0}$ becomes much larger than $\hat{q}_{c,0}$. In this case, \hat{U}_I becomes

$$\hat{U}_I = e^{i6gt\hat{q}_c\hat{q}_r^2/x_0^3} \quad (6.19)$$

If $\omega_c \ll \omega_r$, then $\hat{q}_{r,0} \ll \hat{q}_{c,0}$, and the cubic term would become dominant. However, because $\omega_r \leq \omega_x$ always, this situation can never be achieved. This means that the cubic phase gate cannot be carried out on multi-ion chains using a cubic potential.

6.3.3 The $\omega_p = \omega_c$ case. If $\omega_p = \omega_c$, then the interaction Hamiltonian can be written under the RWA as

$$\begin{aligned} \hat{H}_I = & -i\frac{1}{2\sqrt{2}}\hbar g [(-\hat{a}_c\hat{a}_c^\dagger\hat{a}_c e^{i\theta} + \hat{a}_c^\dagger\hat{a}_c\hat{a}_c^\dagger e^{-i\theta}) \\ & + 3\frac{\omega_c}{\omega_r}(-\hat{a}_c e^{i\theta} + \hat{a}_c^\dagger e^{-i\theta})(2\hat{a}_r^\dagger\hat{a}_r + 1)]. \end{aligned} \quad (6.20)$$

Notably, this operator has two terms to it: The first carries out the single-ion $\omega_p = \omega_x$ operator, applied to the COM mode. The second is a displacement in the COM mode whose strength is proportional to the occupation $\hat{a}_r^\dagger\hat{a}_r$ of the rocking mode. As in the two-ion $\omega_p = 0$, case, this second term can be made much stronger than the first by making $\omega_r \ll \omega_c$

$$\hat{H}_I \approx -i\frac{3}{2\sqrt{2}}\hbar g\frac{\omega_c}{\omega_r}(-\hat{a}_c e^{i\theta} + \hat{a}_c^\dagger e^{-i\theta})(2\hat{a}_r^\dagger\hat{a}_r + 1).$$

If the COM and rocking modes are both starting from the ground state, the above operation simply becomes a displacement operator. Population plots generated by

Because the interaction Hamiltonian has no terms that are first-order in \hat{a}_r and \hat{a}_r^\dagger , the $\omega_p = \omega_r$ case does not produce an operation on the OOP mode.

6.3.4 The $\omega_p = \omega_c + 2\omega_r$ case. The two-ion interaction Hamiltonian allows two-mode interactions coupling the COM and OOP modes. If $\omega_p = \omega_c + 2\omega_r$, the interaction Hamiltonian can be rewritten, after the RWA,

$$\begin{aligned}\hat{H}_I &= -i\frac{3}{2\sqrt{2}}\frac{\omega_c}{\omega_r}\hbar g(-\hat{a}_c^\dagger\hat{a}_r^{\dagger 2}e^{i\theta} + \hat{a}_c\hat{a}_r^2e^{-i\theta}) \\ \hat{H}_I &= -i\hbar\xi(-\hat{a}_c^\dagger\hat{a}_r^{\dagger 2}e^{i\theta} + \hat{a}_c\hat{a}_r^2e^{-i\theta})\end{aligned}\quad (6.21)$$

where $\xi = \frac{3}{2\sqrt{2}}\frac{\omega_c}{\omega_r}g$ is the coupling strength for this operation.

6.3.5 The $\omega_p = \omega_c - 2\omega_r$ case (parametric oscillation). Parametric oscillation is a non-Gaussian, multi-mode process that has been investigated as a resource for optical quantum computing (Langford et al. (2011)). This process has been carried out on the vibrational modes of a two-ion system using the non-linearity of the Coulomb interaction (Ding et al. (2017)). We discuss below how parametric oscillation can also be achieved using a cubic potential.

If $\omega_p = \omega_c - 2\omega_r$, then the two-ion interaction Hamiltonian can be written under the RWA as

$$\begin{aligned}\hat{H}_I &= -i\frac{3}{2\sqrt{2}}\frac{\omega_c}{\omega_r}\hbar g(-\hat{a}_c^\dagger\hat{a}_r^2e^{i\theta} + \hat{a}_c\hat{a}_r^{\dagger 2}e^{-i\theta}) \\ \hat{H}_I &= -i\hbar\xi(-\hat{a}_c^\dagger\hat{a}_r^2e^{i\theta} + \hat{a}_c\hat{a}_r^{\dagger 2}e^{-i\theta})\end{aligned}\quad (6.22)$$

where $\xi = \frac{3}{2\sqrt{2}}\frac{\omega_c}{\omega_r}g$ is the parametric oscillator coupling strength, which is identical to the parametric amplifier coupling strength above. This coupling strength ξ is notable in that, if radial and axial frequencies are brought close to each other ($\omega_x \rightarrow \omega_z$), then the rocking mode frequency $\omega_r \rightarrow 0$, and ξ can be made arbitrarily large. Further discussion of ξ can be found in Section 6.5.

6.3.6 The $\omega_p = 3\omega_c$ and $\omega_p = 3\omega_r$ cases (trisqueezing). As discussed in Section 6.2, trisqueezing is a non-Gaussian operation that could allow universal CVQC. Trisqueezing can be carried out on an arbitrary normal mode for

a single ion. For a two-ion chain the choice of modes is more limited. The two-ion interaction Hamiltonian has a \hat{q}_c^3 term, analogous to the single-ion case. This means that $\omega_p = 3\omega_c$ will produce trisqueezing on the radial COM mode. Applying the RWA, the interaction Hamiltonian can be rewritten in this case as

$$\hat{H}_I = -i \frac{1}{2\sqrt{2}} \hbar g (-\hat{a}_c^{\dagger 3} e^{i\theta} + \hat{a}_c^3 e^{-i\theta}).$$

This Hamiltonian is functionally the same as in the single-ion trisqueezing case but with a coupling strength weaker by a factor of $\sqrt{2}$.

Because there are no third-order terms in \hat{a}_r and \hat{a}_r^\dagger , $\omega_p = 3\omega_r$ does not produce trisqueezing. This is in contrast to squeezing where the OOP mode can be squeezed by a quadratic potential with $\omega_p = 2\omega_r$.

6.4 Trap design for generating cubic potentials

In the previous sections, we characterized the third-order operations that can be carried out on one- and two-ion systems using a modulated cubic potential in the radial plane. Cubic potentials are more challenging to generate than the quadratic potentials and gradients needed to carry out Gaussian operations, and electrically-driven third-order operations have not been demonstrated experimentally. We propose a novel surface-electrode trap design that can generate cubic potentials and therefore allow these operations to be carried out in the lab.

6.4.1 Multipole expansion - the hexapole potential. Third-order non-Gaussian operations can be performed on trapped ions by applying a cubic potential along one dimension. Such a cubic potential can be produced, while also satisfying Laplace's equation, by generating a hexapole potential centered about the trap axis. We propose an approach for generating hexapole potentials that relies on the a feature of trap electrodes infinitely long (or effectively so) in the axial direction. Specifically, these electrodes generate potentials that have no dependence

on axial position. Such potentials can be written as a multipole expansion about the equilibrium point (i.e. the trap position) in cylindrical coordinates

$$V(r, \phi) = V_0 \left[\sum_{n=2}^{\infty} C_n \left(\frac{r}{r_0} \right)^n \cos(n\phi) + \sum_{m=2}^{\infty} S_m \left(\frac{r}{r_0} \right)^m \sin(n\phi) \right] + V_{off} \quad (6.23)$$

where V_0 is electrode voltage, V_{off} is the potential at the ion, r_0 is ion-electrode separation, r is distance from the trap center, ϕ is angle from a given axis in the radial plane, and C_n and S_m are expansion coefficients. A schematic representation of the cylindrical harmonic decomposition is shown in Fig. 6.4.




Figure 6.4. Trapping potential generated by symmetrical surface-electrode trap around ion (marked in yellow). The decomposition of this potential into overlaid quadrupole, hexapole, and octopole terms are represented schematically.

It is possible to create an electrode configuration such that the leading term of the expansion is the hexapole (C_3 and S_3) and such that higher-order terms (the octopole, decapole, and so on) are negligible. In what follows, we consider surface-electrode traps symmetric about the plane formed by the trap axis and the normal to the surface, as shown in Fig. 6.6. Therefore if ϕ is defined as the angle from the axis normal to the trap surface, then $S_3 = 0$, and the potential produced by these

electrodes can be written as

$$V(r, \phi) = V_0 C_3 \cos(3\phi) \left(\frac{r}{r_0}\right)^3 + V_{off}. \quad (6.24)$$

If the \hat{x} mode of a single ion is oriented at angle ϕ , a direct comparison can be made between the oscillating cubic potential term in the single-ion Hamiltonian (Eq. 6.1) and the cubic term of the above potential to give

$$-\hbar g \left(\frac{\hat{x}}{x_0}\right)^3 = eV_0 C_3 \cos(3\phi) \left(\frac{\hat{x}}{r_0}\right)^3.$$

This expression gives the following formula for g , the single-ion coupling strength

$$g = -\cos(3\phi) \frac{eV_0 C_3}{\hbar} \left(\frac{x_0}{r_0}\right)^3. \quad (6.25)$$

6.4.2 Surface-electrode design. To implement third-order non-Gaussian operations, we need to design an electrode system that will generate a hexapole potential in the radial plane. We approach this problem using a conformal mapping between cylindrical and planar electrode systems.

6.4.2.1 Equivalent cylindrical electrode systems. A system of surface-electrodes that generates a given multipole can be designed by considering a corresponding cylindrical electrode system related by a conformal mapping (Wesenberg (2008)). For a cylindrical electrode system of radius r_0 , a point on the surface of the cylinder can be mapped to a point on a 2D electrode plane by the relation $\vec{r}(\theta) = -r_0 \hat{e}_x + r_0 \tan\left(\frac{\theta}{2}\right) \hat{e}_y$, where x and y are axes in the radial plane parallel and perpendicular to the electrode plane respectively, and the origin is taken to be at the center of the cylinder, as shown in Fig. 6.5.

6.4.2.2 Designing electrode systems to produce hexapoles.

From the perspective of corresponding cylindrical systems, it is clear from symmetry why surface-electrode trap designs with two parallel RF electrodes embedded in a ground plane can generate strong quadrupoles, as illustrated in




Figure 6.5. Cross-section of an asymmetric surface-electrode trap through plane perpendicular to trap axis, with ion position marked by a yellow dot. Cross-section of equivalent cylindrical electrode system (partitioned ring of radius r_0) is overlaid on top, with the relationship between surface and cylindrical electrode positions $\vec{r}(\theta) = -r_0\hat{e}_x + r_0 \tan\left(\frac{\theta}{2}\right)\hat{e}_y$ illustrated geometrically. Figure adapted from Wesenberg (Wesenberg (2008)

).

Fig. 6.6a. Electrodes that generate a strong hexapole without a quadrupole can be designed on the same basis of symmetry, producing a hexapole electrode system consisting of (in the cylindrical representation) three RF electrodes spaced at 120° from one another on the cylinder (Fig. 6.6b).

The potentials generated by the hexapole electrodes were calculated using the Electrode software developed at NIST (NIST (2017)), and multipole expansions were fit to these potentials. A plot of C_3 for the hexapole electrodes against the parameter θ_1 is shown in Fig. 6.7a. As expected from symmetry considerations, C_3 is maximized when $\theta_1 = 60^\circ$, or when the RF and ground electrodes occupy the same solid angles. The magnitude of C_2 (the quadrupole potential) is zero for these electrodes, since the smooth, one-to-one mapping between the cylindrical and planar representations of the space around the ion cannot change the number of

equipotential lines converging at the ion and, as a result, cannot change the leading term of the harmonic expansion.

6.4.2.3 Integrating hexapole electrodes into surface trap. In order to be useful in carrying out third order operations on trapped ions, the hexapole electrodes must be integrated with a typical surface trap design. We do this by simply superimposing the quadrupole and hexapole electrodes shown in Fig. 6.6a and b to obtain an electrode system with appropriately placed electrodes for trapping and third order operations, shown in Fig. 6.6c.

Superimposing the quadrupole and hexapole electrodes places a constraint on θ_1 and θ_2 . Because the quadrupole RF electrodes are spaced 180° apart and the hexapole electrodes are spaced 120° apart, if the electrodes are arranged symmetrically, the quadrupole electrodes will be separated from the nearest hexapole electrodes by 30° , as shown in Fig. 6.6c. Under this condition, in order for the RF electrodes to be non-overlapping, $\theta_1 + \theta_2 \leq 60^\circ$. If the quadrupole and hexapole electrodes overlap, then in the overlapping regions, another set of electrodes with a voltage equal to the sum of those on the quadrupole and hexapole electrodes must be added to the system. Implementing this third set of RF electrodes introduces technical complications, so to avoid this, we only consider the $\theta_1 + \theta_2 \leq 60^\circ$ range.

The range of optimal θ_1 and θ_2 can be constrained further by recognizing that, for $\theta_1, \theta_2 \leq 60^\circ$, the strength of C_3 for the hexapole electrodes and C_2 for the quadrupole electrodes increase monotonically with θ_1 and θ_2 respectively. This means that, for any value of θ_1 , the optimum value of θ_2 is $\theta_2 = 60^\circ - \theta_1$, and the overall performance of the system is maximized when $\theta_1 + \theta_2 = 60^\circ$. This creates a trade-off between C_2 and C_3 , as shown in Fig. 6.7a. In practical terms, this means there is a trade-off between the voltage needed for efficient third order operations (via C_3) and the voltage needed to produce a given trap frequency (via C_2). Worth noting is that, while this parameter range includes the value of θ_1 that maximizes the hexapole ($\theta_1 = 60^\circ$), the θ_2 maximizing the quadrupole ($\theta_2 = 90^\circ$, by symmetry) is excluded, meaning that including the hexapole electrodes needed for trisqueezing slightly reduces the maximum possible C_2 (in this case, from 0.16 to 0.14). It should be noted that the center hexapole electrode takes up a relatively small fraction of the space between the two rf electrodes (compare sizes of inner hexapole electrode to inner ground in Fig. 6.7b), leaving ample space for DC control electrodes.

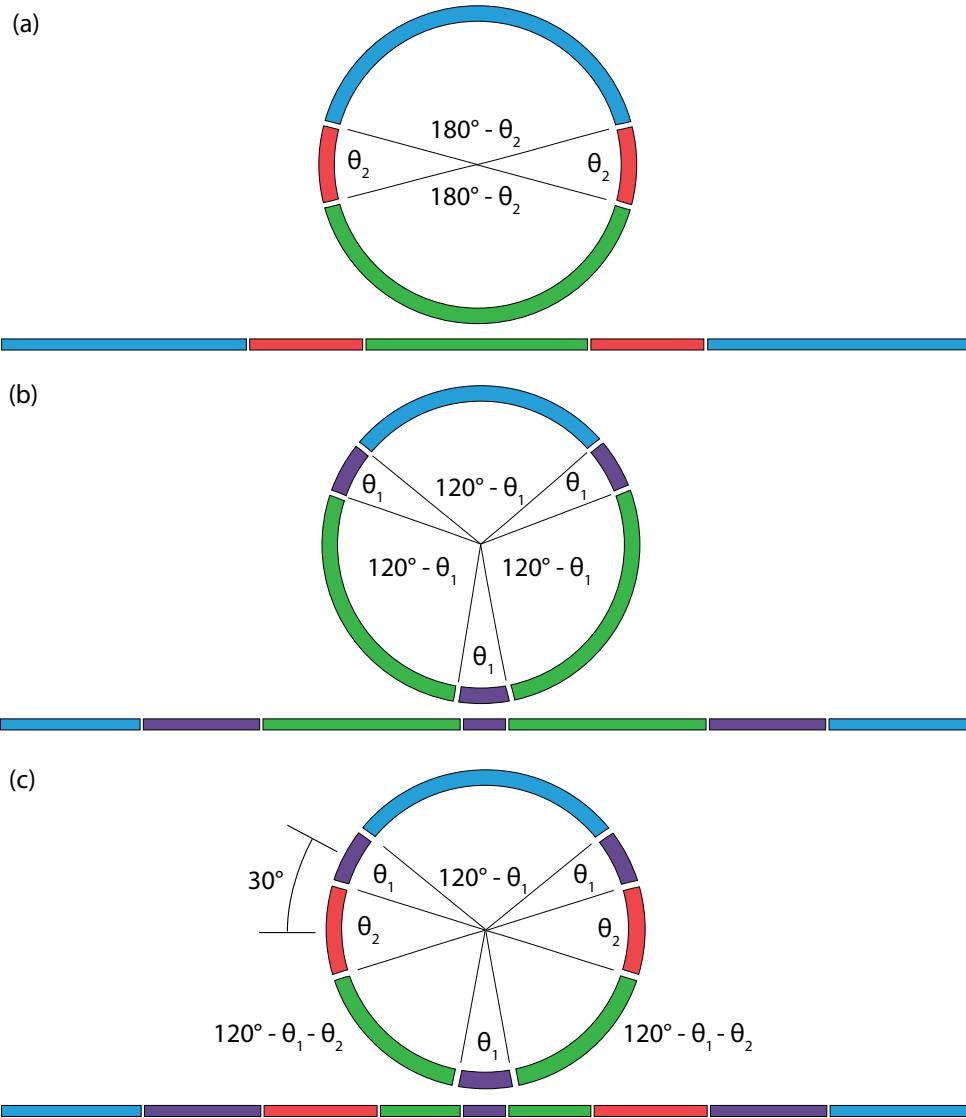


Figure 6.6. a). Quadrupole system. b). Hexapole system. c). Combined electrode system for generating quadrupole trapping potential and hexapole trisqueezing potential. *Red* - Quadrupole electrodes. *Purple* - Hexapole electrodes. *Green* - Inner ground. *Blue* - Outer ground. Notably, to produce optimal performance under geometric constraints, $\theta_1 + \theta_2 = 60^\circ$, and the quadrupole electrodes are adjacent to hexapole electrodes.

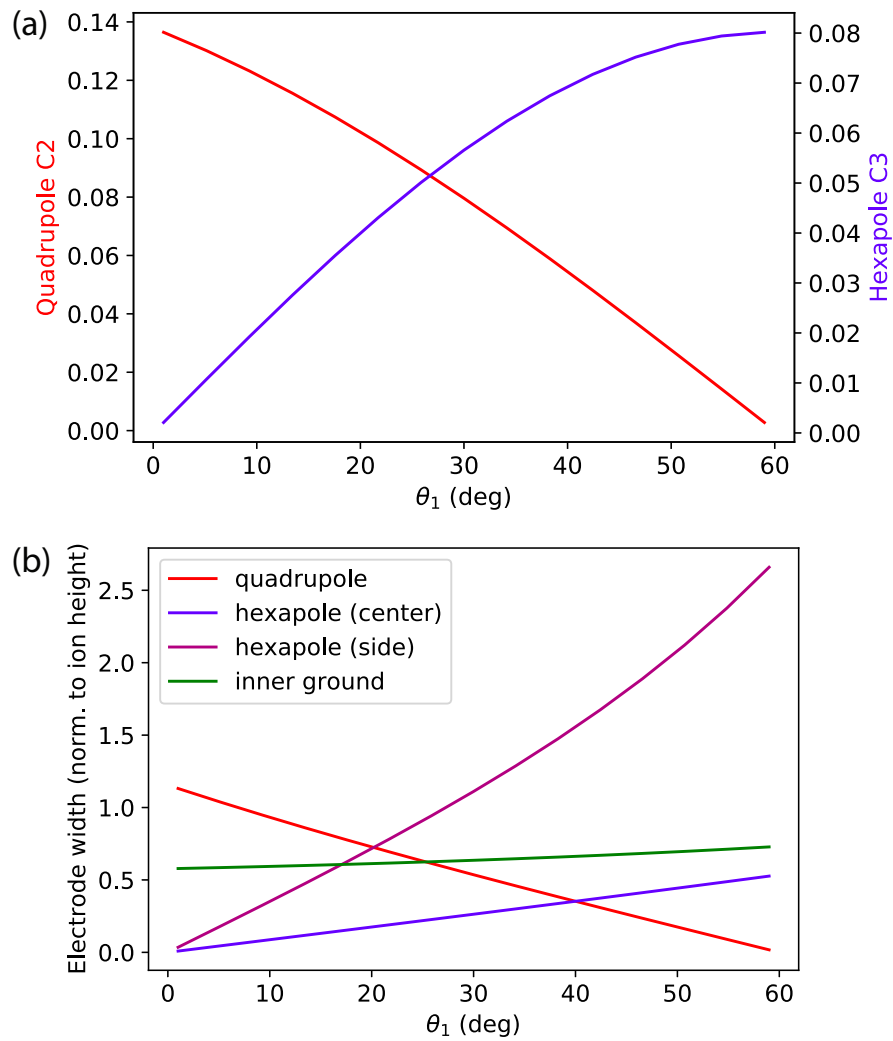


Figure 6.7. a). C_2 for the quadrupole electrodes and C_3 for the hexapole electrodes plotted against θ_1 when $\theta_1 + \theta_2 = 60^\circ$. b). Widths of different electrodes in the combined electrode system of Fig. 6.6b, normalized to ion height, under the condition $\theta_1 + \theta_2 = 60^\circ$.

6.4.3 Axial cubic potentials - design limitations. In principle, it is possible to apply a third order non-Gaussian operations to the axial mode of a single ion using a cubic potential. However, the design approach we describe in this section cannot be applied to axial situation, since this would require rotating the hexapole electrodes to be orthogonal with the trapping electrodes. Making the hexapole and trapping electrodes orthogonal would require breaking the translational symmetry of the hexapole electrodes that allows us to apply the analysis described by Wesenberg. On top of this, an axial system is not practical for working with chains of ions. As discussed in Section 6.3, an axial cubic potential applied to an ion chain will cause unwanted displacements.

6.5 Practical implementation

In Section 6.4, we present a novel design for a surface-electrode trap capable of generating the cubic potentials. We calculate that the hexapole fit coefficient C_3 for this design can range from zero to 0.08. From this range, we can estimate what range of trap parameters would be needed to achieve a practically useful coupling strength.

We will consider the cases of trisqueezing and the parametric oscillator/amplifier. We will compare the coupling strengths achievable for these operations to those that have been used for similar operations in experimental work. Specifically, we compare trisqueezing to second-order squeezing obtained through the application of a quadratic potential, and we compare the rocking/radial COM parametric oscillator to a rocking/stretching parametric oscillator achieved using non-linearity in the ion-ion Coulomb force.

6.5.1 Trisqueezing. To estimate what single-ion coupling strengths g might be achievable for trisqueezing in a practical setting, we will consider the requirements for producing a particular value of g . As shown in Section 6.4, g for

single-ion trisqueezing is

$$g = -\cos(3\phi) \frac{eV_0 C_3}{\hbar} \left(\frac{x_0}{r_0} \right)^3 \quad (6.26)$$

where x_0 is ground state wave function extent for the \hat{x} mode. From Section 6.4, we know that we can expect $C_3 \approx 0.05$. A typical trapping setup might use $^{43}\text{Ca}^+$ with a radial frequency $\omega_r = 2\pi \times 4 \text{ MHz}$ at a distance $r_0 = 100 \mu\text{m}$ above the electrode plane. Taking $\cos(3\phi) \approx 1$, a hexapole electrode voltage of $V_0 \approx 520 \text{ V}$ is needed to generate $g = 2\pi \times 1 \text{ kHz}$. This voltage is potentially feasible as it's similar to the upper end of what can be required for the trap drive. However, a lower voltage would be desirable, as would a higher g (e.g. $g \sim 10 \text{ kHz}$). This kind of coupling strength is more in line with that used in experimentally realized squeezing (Burd et al. (2019)). In the interest of improving coupling strength and decreasing voltage requirements, we consider the following factors:

Distance dependence. The coupling strength g is strongly dependent on ion-electrode separation ($g \propto 1/r_0^3$) and so reducing r_0 from $100 \mu\text{m}$ to $30 \mu\text{m}$ (the smallest r_0 commonly in use) reduces V_0 from 520 V to 14 V .

Frequency dependence. There is a moderate dependence on trap frequency ($g \propto 1/\omega_x^{3/2}$). Low radial frequencies in a small surface trap may be hard to achieve whilst maintaining a reasonable trap depth. Switching to a frequency of 500 kHz (likely easier to achieve using axial modes than radial modes) would reduce the required voltage to 23 V . It should be noted that heating rates tend to scale more strongly than $1/\omega_x^{3/2}$ though so this wouldn't offer an advantage in a scenario where heating rates limited coherence. This caveat applies less strongly to operations that employ OOP modes (e.g. the rocking mode) since these modes experience lower heating rates (King et al. (1998)) while at the same time having frequencies that can be driven arbitrarily low (see discussion in Section 6.5.2).

Mass dependence. The coupling strength g is moderately dependent on ion mass, in the same way as trap frequency ($g \propto 1/m^{3/2}$). Switching from $^{43}\text{Ca}^+$ to $^9\text{Be}^+$ reduces the required voltage to 50 V, while switching to $^{88}\text{Sr}^+$ increases the required voltage to 1520 V.

Other operations. It should be noted that this dependence on ion-electrode separation and mass hold for all third-order operations in both the one-ion and two-ion cases. The dependence on frequency holds for the three single-ion operations and for two-ion trisqueezing.

The coupling ξ for the two-ion, two-mode operations (the parametric oscillator and parametric amplifier) includes both the single-ion g as well as a factor ω_c/ω_r , meaning that $\xi \propto 1/\omega_x \sqrt{\omega_x^2 - \omega_z^2}$.

6.5.2 Parametric oscillation. For the parametric oscillator/amplifier, coupling strength is $\xi = \frac{3}{2\sqrt{2}} \frac{\omega_c}{\omega_r} g$ (see Section 6.3). For a set of representative values (using $^{43}\text{Ca}^+$ with $C_3 \approx 0.05$, radial frequency $\omega_x = 2\pi \times 4$ MHz, axial frequency $\omega_z = 2\pi \times 1$ MHz, ion-electrode separation $r_0 = 100 \mu\text{m}$, hexapole electrode voltage $V_0 = 100$ V and $\cos(3\phi) \approx 1$), the coupling strength for the two-ion parametric oscillator will be $g = 0.4$ kHz. Bringing ion-electrode separation down to $r_0 = 30 \mu\text{m}$ drives coupling strength up to $g = 14.7$ kHz. As discussed in Section 6.3, since $\xi \propto 1/\omega_r$, ξ can be driven arbitrarily high by bringing ω_x to ω_z .

The coupling strength here can be compared with that of an experimentally realized parametric oscillator, as in Ding et al (Ding et al. (2017)), where the rocking and axial stretch modes are coupled via the non-linearity of the Coulomb interaction by adjusting the trap frequency such that $\omega_s \approx 2\omega_r$. The coupling

strength of this process is given as

$$\xi_{coulomb} = \frac{1}{8z_0} \sqrt{\frac{\hbar\omega_s^3}{m\omega_r^2}} \quad (6.27)$$

where z_0 is the equilibrium ion separation. For the masses and frequencies ω_x and ω_z used above, the coupling strength is $\xi_{coulomb} \sim 1$ kHz if $z_0 \sim 1 \mu\text{m}$. Notably, $\xi_{coulomb}$ shares the same $1/\omega_r$ dependence as ξ , meaning that this coupling can be driven up by bring the single-ion frequencies ω_x and ω_z closer together. For the values used in Ding et al, a $\xi_{coulomb}$ of ~ 3 kHz was measured.

The coupling strengths achievable for the parametric oscillator carried out with an external cubic potential seem comparable to those available for similar processes that have been carried out experimentally using between-ion Coulomb interactions. However, there are advantages to using cubic potentials generated by separate trap electrodes. By using external cubic potentials, coupling can be turned on and off without changing trap frequencies to specific ratios which may or may not be stable [should we include citation for statement about stability?]

6.6 Conclusion

We have characterized the operations that can be performed on systems of one and two trapped ions using an externally applied cubic potential, and we have presented a novel trap design that we predict could carry out these operations at efficiencies comparable to those seen in experiments using similar operators.

CHAPTER VII

CONCLUSION

The preceding chapters laid out a set of projects that explore different technological domains of trapped ion quantum computing: New techniques in coherent control (in the form of entangling gates in metastable trapped-ion qubits, with some unique features of this qubit type characterized), the engineering of new trap geometries and fabrication techniques for scalability and favorable trap performance (in the form of the microfabrication of 3D trap networks through 3D printing), and an intersection of the two (in the form of a project to design a 2D, microfabricatable trap that can be used to carry out non-Gaussian operations for continuous-variable quantum computing). Each project suggests further work to be done. In the case of gates in m qubits, deployment of these sorts of gates in a full *omg* context is a natural follow-up, while for the two design projects, the fabrication and testing of functional traps would be the next place to go. The two projects referenced but not discussed in detail, the use of all-electronic motional drives to carry out motional interferometry and the exploration of non-Hermitian quantum systems with metastable qubits, each branch off into sets of projects of both potential practical value (in the applicability of the techniques of motional interferometry to CVQC) and scientific interest (in the investigation of the richly varied physics that arise through non-Hermitian dynamics). Many of these branching paths will continue to be explored, whether here or elsewhere.

APPENDIX

GATE ERROR CALCULATIONS

The discussion of gate errors below was prepared by Gabe Gregory, who calculated these errors based on the results of a suite of auxilliary experiments performed to characterize different error sources in our system (dephasing, heating, calibration errors / drift, and decay/scattering). For reference, these calculations are summarized below.

A.1 Error budget

Here, we detail calculations used to estimate the dominant sources of two-qubit gate infidelity presented in Table 3. We also detail sources of error not presented in the text, including (list stuff), all of which were found to be negligible compared with those presented in the table.

Spin dephasing. The Zeeman qubits in this work suffer from spin decoherence due to the first order sensitivity of qubit splitting to ambient magnetic field fluctuations. To mitigate this, we implemented a Walsh 1 gate sequence in spin, reducing sensitivity to magnetic field noise that is slow on the timescale of the gate. Faster noise components however still degrade the coherence, and we found that loss of spin coherence accounted for a large portion of the total gate infidelity. To estimate the magnitude of this infidelity, we performed a Ramsey spin echo experiment on a single ion, and measured the loss of Ramsey fringe contrast c for a sequence of the same duration as the gate. The corresponding two-qubit gate infidelity is $\frac{3}{2}(1 - c)$ Ballance (2017). We measure a contrast loss $(1 - c)$ of 17.5×10^{-4} after subtracting single ion SPAM, corresponding to a gate infidelity of $26(2) \times 10^{-4}$.

Motional dephasing. To estimate the Bell-state infidelity due to motional dephasing, we performed a motional analogue to the standard Ramsey experiment

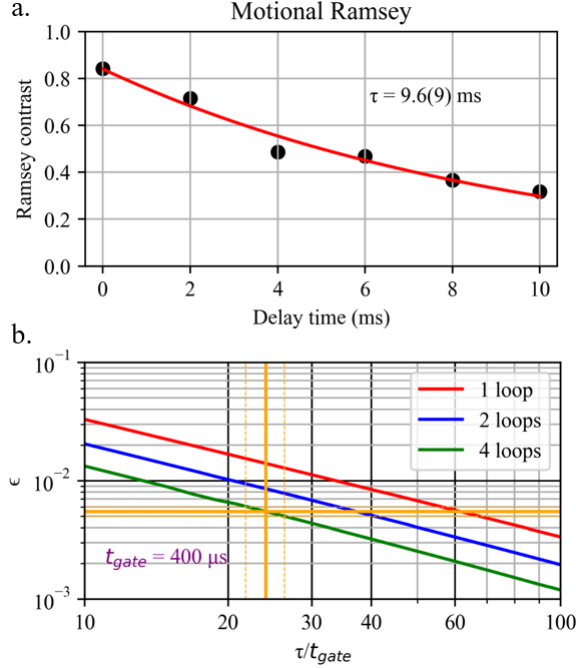


Figure A.1. a) Motional Ramsey experiment, plotting the contrast between the $|0\rangle$ and $|1\rangle$ fock states as a function of the delay time between sideband $\pi/2$ pulses. *b)* Numerically simulated Bell-state infidelity due to a white motional noise source as a function of the coherence to gate time ratio. Orange lines represent the infidelity expected for measured coherence and gate times used in this work. Dashed lines represent the 68% confidence interval.

we utilized to characterize spin dephasing. By driving the red and blue low frequency out of phase sidebands, we first place the motional state of the crystal in a superposition of $|0\rangle$ and $|1\rangle$ Fock states. We then wait a delay time T , and perform the inverse of the motional $\pi/2$ pulse, scanning over the phase of that pulse to obtain a fringe contrast c . We find an exponential decay (with coherence time τ) well fits the Ramsey contrast as a function of the delay time T , suggesting a motional noise source that is white in frequency.

We then numerically simulate the Bell-state infidelity as a function of ratio of the coherence time characterizing the white noise and the total gate duration by solving the master equations in Lindblad form numerically. We fit a motional

coherence time of 9.6(9) ms to the Ramsey contrast data, corresponding to a gate infidelity of $55(5) \times 10^{-4}$.

π -time calibration. Fits to determine Raman driven carrier π -times returned an error of 1%, suggesting that between calibrations, we can expect an error in the π -time drawn from a normal distribution with a standard deviation of 1%. Here we plot numerical simulations of the Bell-state infidelity due to percent offsets in the π -time drawn from a normal distribution characterized by σ , along with distributions of the infidelity to estimate the 68% confidence interval. The degree to which the sequence used is sensitive to π -time miscalibrations depends on the phase of the spin-echo pulse, which is in practice difficult to control due to the fact that we are applying large, inhomogeneous light shifts using a tightly focused beam that can drift over short time scales. To account for this behavior in simulations, we simply average over carrier phases. We find that for an error in the π -time of 1% we expect a Bell-state infidelity of $6.8_{-5.5}^{+12.0} \times 10^{-4}$.

Raman intensity drift. Alignment of the Raman beams generating the spin dependant force is subject to slow drifts over the course of the gate. Drifts of the beam along the axis of the crystal do not contribute any significant infidelity due to the symmetrized nature of the Walsh 1 sequence used in this work. However, drifts orthogonal to the axis of the crystal manifest as errors in the phase accrued to the $|01\rangle$ and $|10\rangle$ spin states. This can be seen as an asymmetry in the $|00\rangle$ and $|11\rangle$ populations after the gate sequence.

Mode frequency drift. The frequency of the mode on which the gate is driven is subject to slow drifts due to several technical reasons (thermal fluctuations, trap anharmonicities, etc.). To estimate the Bell-state infidelity associated with this mode frequency drift, we first measured the width of the low frequency oop mode to have a σ of 100 Hz. We then numerically solved for the Bell-state infidelity

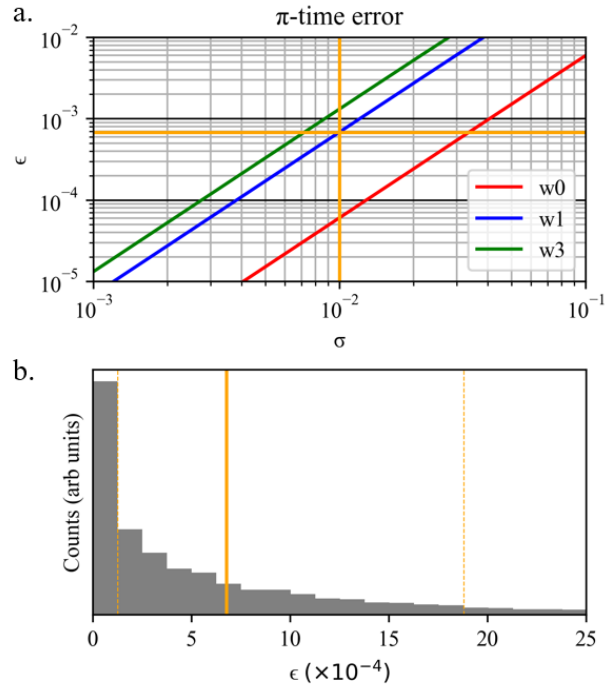


Figure A.2. **a)** Bell-state infidelity corresponding to a normally distributed static offset in gate mode frequency with standard deviation σ for different orders of Walsh modulation. **b)** Distribution of Bell-state infidelities corresponding to a normal distribution of mode frequency offsets with a standard deviation of 100 Hz for a Walsh 1 4-loop gate operating at 11.6 kHz. Dashed orange lines represent 68% confidence interval.

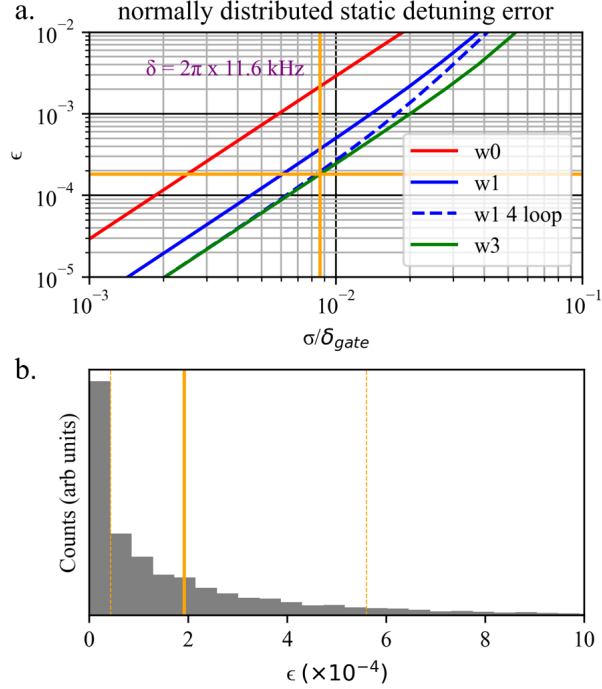


Figure A.3. **a)** Bell-state infidelity corresponding to a normally distributed static offset in gate mode frequency with standard deviation σ for different orders of Walsh modulation. **b)** Distribution of Bell-state infidelities corresponding to a normal distribution of mode frequency offsets with a standard deviation of 100 Hz for a Walsh 1 4-loop gate operating at 11.6 kHz. Dashed orange lines represent 68% confidence interval.

corresponding to an offset error in the gate mode frequency, and performed a weighted average over the infidelities corresponding to a normal distribution of offset errors characterized by a standard deviation σ .

To estimate the uncertainty in the infidelity, we looked at the distribution of infidelities generated by the distribution of mode frequency offsets with standard deviation 100 Hz and simply located the 68% confidence intervals of the infidelity distribution. From this method we found a Bell-state infidelity of $1.9^{+3.7}_{-1.4} \times 10^{-4}$.

Erasure errors. To estimate erasure error type infidelities we simply measured the rate Γ at which each beam scatters population out of the $D_{5/2}$ manifold for a single ion. Together with the duration over which the beam is

on, the Bell state infidelity is $2\Gamma t$ (the factor of two accounts for two ions in the trap during the gate). Because both ions are in an equal spin state superposition throughout the gate, we measure the scattering rate out of each qubit state and consider the average scattering rate when calculating infidelity. The measured scattering rates for each beam and qubit state are presented in Table A.7.

The reported scattering rates in Table A.7 include decay of the natural $D_{5/2}$ lifetime. We subtract the known decay rate of $1/1.16$ Hz from each measured value before calculating the infidelity. We find Bell-state infidelities of $13.5(2) \times 10^{-4}$, $5.6(8) \times 10^{-4}$, and $28.8(1) \times 10^{-4}$ for erasure errors caused by 976 scattering, 854 scattering, and $D_{5/2}$ decay respectively.

Raman/Rayleigh scattering to $D_{5/2}$. We can then use the known branching ratios together with the measured scattering rates (Table A.7) to calculate the scattering rates back into the $D_{5/2}$ manifold and using the same treatment to estimate the Bell-state infidelity as was done for erasure errors. If however a qubit state scatters back to the same state (Rayleigh scattering) a recoil kick occurs and the Bell-state infidelity is instead $\frac{1}{2}\Gamma_{\text{Rayleigh}}t$ Ballance (2017). We find a Raman scattering error (non erasure) of $0.40(1) \times 10^{-4}$ due to the 976 beams, and infidelity $\ll 10^{-5}$ for Rayleigh scattering errors.

Table A.7. Measured scattering rates out of $D_{5/2}$ (in Hz) when initialized in the $m = +3/2$ or $m = +5/2$ qubit states for various beams used in the gate. Measurements for the 854 LS beam were taken with two trapped ions.

Beam	$m = +5/2$	$m = +3/2$
854 LS	5.8(3)	7.9(3)
976 Rpi	1.40(2)	1.75(2)
976 Rsig	2.23(5)	1.38(2)

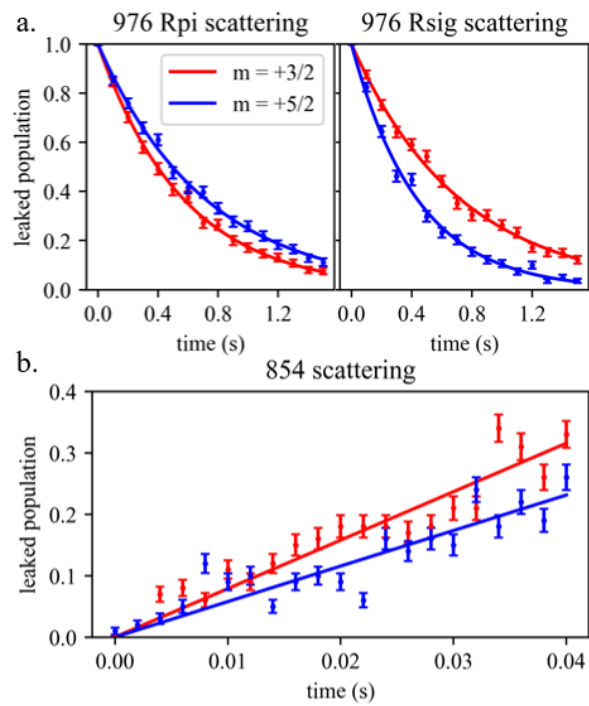


Figure A.4. **a)** Population scattered out of the $m = +5/2$ and $+3/2$ in the $D_{5/2}$ manifold with either the 976 Rpi or Rsig beam on. **b)** Population scattered out of the qubit states with the 854 beam on.

REFERENCES CITED

- 1 gsps, 14-bit, 3.3 v cmos direct digital synthesizer (Computer software manual No. AD9910). (2016). Retrieved from <https://www.analog.com/media/en/technical-documentation/data-sheets/ad9910.pdf> (Rev. E)
- Aksenov, M. A., Zalivako, I. V., Semerikov, I. A., Borisenko, A. S., Semenin, N. V., Sidorov, P. L., ... Kolachevsky, N. N. (2023, May). Realizing quantum gates with optically addressable $^{171}\text{Yb}^+$ ion qudits. *Physical Review A*, *107*, 052612. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevA.107.052612> doi: 10.1103/PhysRevA.107.052612
- Allcock, D. T. C., Campbell, W. C., Chiaverini, J., Chuang, I. L., Hudson, E. R., Moore, I. D., ... Wineland, D. J. (2021). *omg* blueprint for trapped ion quantum computing with metastable states. *Applied Physics Letters*, *119*, 214002. Retrieved from <http://dx.doi.org/10.1063/5.0069544> doi: 10.1063/5.0069544
- Allcock, D. T. C., et al. (2013). A microfabricated ion trap with integrated microwave circuitry. *Applied Physics Letters*, *102*, 044103. Retrieved from <https://doi.org/10.1063/1.4774299> doi: 10.1063/1.4774299
- Alonso, J., Leupold, F., Solèr, M. Z. U., Fadel, M., Marinelli, M., Keitch, B. C., ... Home, J. P. (2016). Generation of large coherent states by bang–bang control of a trapped-ion oscillator. *Nature Communications*, *7*(11243). Retrieved from <https://www.nature.com/articles/ncomms11243> doi: 10.1038/ncomms11243
- Arrington, C. L., et al. (2013). Micro-fabricated stylus ion trap. *Review of Scientific Instruments*, *84*, 085001. Retrieved from <https://doi.org/10.1063/1.4817304> doi: 10.1063/1.4817304
- Auchter, S., Axline, C., Decaroli, C., Valentini, M., Purwin, L., Oswald, R., ... Home, J. (2022). Industrially microfabricated ion trap with 1 eV trap depth. *Quantum Science and Technology*, *7*(3), 035015.
- Ballance, C. J. (2017). *High-fidelity quantum logic in Ca^+* (Unpublished doctoral dissertation). PhD thesis, Oxford.
- Balzer, C., Braun, A., Hannemann, T., Paape, C., Ettlér, M., Neuhauser, W., & Wunderlich, C. (2006). Electrostatically trapped Yb^+ ions for quantum information processing. *Physical Review A—Atomic, Molecular, and Optical Physics*, *73*(4), 041407.

- Băzăvan, O., Saner, S., Minder, M., Hughes, A., Sutherland, R., Lucas, D., . . . Ballance, C. (2023). Synthesizing a σ^z spin-dependent force for optical, metastable, and ground-state trapped-ion qubits. *Physical Review A*, *107*(2), 022617.
- Benhelm, J. (2008). *Precision spectroscopy and quantum information processing with trapped calcium ions*. na.
- Bharti, K., Cervera-Lierta, A., Kyaw, T. H., Haug, T., Alperin-Lea, S., Anand, A., . . . others (2022). Noisy intermediate-scale quantum algorithms. *Reviews of Modern Physics*, *94*(1), 015004.
- Blain, M. G., Haltli, R., Maunz, P., D., N. C., Revelle, M., & Stick, D. (2021). Hybrid MEMS-CMOS ion traps for NISQ computing. *Quantum Science and Technology*, *6*, 034011. Retrieved from <https://doi.org/10.1088/2058-9565/ac01bb> doi: 10.1088/2058-9565/ac01bb
- Blakestad, R. B., et al. (2009). High-Fidelity Transport of Trapped-Ion Qubits through an X-Junction Trap Array. *Physical Review Letters*, *102*, 153002. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevLett.102.153002> doi: 10.1103/PhysRevLett.102.153002
- Braunstein, S. L., & Van Loock, P. (2005). Quantum information with continuous variables. *Reviews of Modern Physics*, *77*(2), 513–577.
- Brown, M. A., Zappitelli, K. M., Singh, L., Yuan, R. C., Bemrose, M., Brogden, V., . . . Gardner, T. J. (2023). Direct laser writing of 3D electrodes on flexible substrates. *Nature Communications*, *14*(1), 3610.
- Burd, S. C., Srinivas, R., Bollinger, J. J., Wilson, A. C., Wineland, D. J., Leibfried, D., . . . Allcock, D. T. C. (2019). Quantum amplification of mechanical oscillator motion. *Science*, *364*(6446), 1163–1165. Retrieved from <http://dx.doi.org/10.1126/science.aaw2884> doi: 10.1126/science.aaw2884
- Carruthers, P., & Nieto, M. M. (1965). Coherent States and the Forced Quantum Oscillator. *American Journal of Physics*, *33*(7), 537-544. Retrieved from <https://doi.org/10.1119/1.1971895> doi: 10.1119/1.1971895

- Chang, C. W. S., Sabín, C., Forn-Díaz, P., Quijandría, F., Vadiraj, A. M., Nsanzineza, I., . . . Wilson, C. M. (2020). Observation of Three-Photon Spontaneous Parametric Down-Conversion in a Superconducting Parametric Cavity. *Physical Review X*, *10*(1), 011011. Retrieved from <http://dx.doi.org/10.1103/PhysRevX.10.011011> doi: 10.1103/physrevx.10.011011
- Chen, W., Gan, J., Zhang, J.-N., Matuskevich, D., & Kim, K. (2021). Quantum computation and simulation with vibrational modes of trapped ions. *Chinese Physics B*, *30*(6), 060311.
- Cheng, B., Deng, X.-H., Gu, X., He, Y., Hu, G., Huang, P., . . . Yu, D. (2023). Noisy intermediate-scale quantum computers. *Frontiers of Physics*, *18*(2), 21308.
- Chiaverini, J., Blakestad, R. B., Britton, J., Jost, J. D., Langer, C., Leibfried, D., . . . Wineland, D. J. (2005). Surface-electrode architecture for ion-trap quantum information processing. *Quantum Information and Computation*, *5*, 419.
- Clark, C. R., Tinkey, H. N., Sawyer, B. C., Meier, A. M., Burkhardt, K. A., Seck, C. M., . . . Brown, K. R. (2021). High-Fidelity Bell-State Preparation with $^{40}\text{Ca}^+$ Optical Qubits. *Physical Review Letters*, *127*, 130505.
- Cornell, E. A., Weisskoff, R. M., Boyce, K. R., & Pritchard, D. E. (1990). Mode coupling in a Penning trap: π pulses and a classical avoided crossing. *Physical Review A*, *41*, 312–315.
- Crain, S., Cahall, C., Vrijsen, G., Wollman, E. E., Shaw, M. D., Verma, V. B., . . . Kim, J. (2019). High-speed low-crosstalk detection of a $^{171}\text{Yb}^+$ qubit using superconducting nanowire single photon detectors. *Communications Physics*, *2*(1), 97.
- Day, M., Choonee, K., Chaboyer, Z., Gross, S., Withford, M. J., Sinclair, A., & Marshall, G. (2021). A micro-optical module for multi-wavelength addressing of trapped ions. *Quantum Science and Technology*, *6*, 024007. Retrieved from <http://iopscience.iop.org/article/10.1088/2058-9565/abdf38>
- DeBry, K., Sinanan-Singh, J., Bruzewicz, C. D., Reens, D., Kim, M. E., Roychowdhury, M. P., . . . Chiaverini, J. (2023). Experimental quantum channel discrimination using metastable states of a trapped ion. *Physical Review Letters*, *131*(17), 170602.

- Dehmelt, H. G. (1968). Radiofrequency Spectroscopy of Stored Ions I: Storage. In D. Bates & I. Estermann (Eds.), (Vol. 3, p. 53). Academic Press. Retrieved from <https://www.sciencedirect.com/science/article/pii/S0065219908601700> doi: [https://doi.org/10.1016/S0065-2199\(08\)60170-0](https://doi.org/10.1016/S0065-2199(08)60170-0)
- Diedrich, F., Bergquist, J. C., Itano, W. M., & Wineland, D. J. (1989). Laser Cooling to the Zero-Point Energy of Motion. *Physical Review Letters*, *62*, 403–406. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevLett.62.403> doi: [10.1103/PhysRevLett.62.403](https://doi.org/10.1103/PhysRevLett.62.403)
- Dielectric ORMOCERs for System-in-Package Electronics and Highly Integrated Systems*. (n.d.). https://www.isc.fraunhofer.de/content/dam/isc/de/documents/Publicationen/Dielectric_ORMOCERs_for_system_in_package_electronics.pdf.
- Ding, S., Maslennikov, G., Hablützel, R., Loh, H., & Matsukevich, D. (2017). Quantum Parametric Oscillator with Trapped Ions. *Physical Review Letters*, *119*(15), 150404. Retrieved from <http://dx.doi.org/10.1103/PhysRevLett.119.150404> doi: [10.1103/physrevlett.119.150404](https://doi.org/10.1103/physrevlett.119.150404)
- Drewsen, M., Mortensen, A., Martinussen, R., Staantum, P., & Sørensen, J. L. (2004). Nondestructive identification of cold and extremely localized single molecular ions. *Physical Review Letters*, *93*(24), 243201.
- Ernzer, M. (2018). *Challenges in Design and Fabrication of a Scalable 3D Ion Trap* (Unpublished master's thesis). ETH Zürich.
- Feng, L., Huang, Y.-Y., Wu, Y.-K., Guo, W.-X., Ma, J.-Y., Yang, H.-X., ... others (2024). Realization of a crosstalk-avoided quantum network node using dual-type qubits of the same ion species. *Nature Communications*, *15*(1), 204.
- Feynman, R. P., Vernon Jr, F. L., & Hellwarth, R. W. (1957). Geometrical representation of the Schrödinger equation for solving maser problems. *Journal of Applied Physics*, *28*(1), 49–52.
- Fröhlich, L., Houbertz, R., Jacob, S., Popall, M., Mueller-Fiedler, R., Graf, J., ... Zychlinski, H. v. (2002). Inorganic-Organic Hybrid Polymers as Photo-Patternable Dielectrics for Multilayer Microwave Circuits. *MRS Online Proceedings Library*, *726*, Q11.4. doi: [10.1557/PROC-726-Q11.4](https://doi.org/10.1557/PROC-726-Q11.4)
- Fukui, K., & Takeda, S. (2022). Building a large-scale quantum computer with continuous-variable optical technologies. *Journal of Physics B: Atomic, Molecular and Optical Physics*, *55*(1), 012001.

- Gorman, D. J., Schindler, P., Selvarajan, S., Daniilidis, N., & Häffner, H. (2014, Jun). Two-mode coupling in a single-ion oscillator via parametric resonance. *Physical Review A*, *89*(6). Retrieved from <http://dx.doi.org/10.1103/PhysRevA.89.062332> doi: 10.1103/physreva.89.062332
- Gottesman, D., Kitaev, A., & Preskill, J. (2001). Encoding a qubit in an oscillator. *Physical Review A*, *64*, 012310.
- Guise, D., N, Fallek, S. D., Stevens, K. E., Brown, K. R., Volin, C., Harter, A. W., ... Youngner, D. W. (2015). Ball-grid array architecture for microfabricated ion traps. *Journal of Applied Physics*, *117*, 174901. Retrieved from <https://doi.org/10.1063/1.4917385> doi: 10.1063/1.4917385
- Harty, T. P. (2013). *High-Fidelity Microwave-Driven Quantum Logic in Intermediate-Field $^{43}\text{Ca}^+$* (Unpublished doctoral dissertation). PhD thesis, University of Oxford.
- Hayes, D., Clark, S. M., Debnath, S., Hucul, D., Inlek, I. V., Lee, K. W., ... Monroe, C. (2012). Coherent Error Suppression in Multiqubit Entangling Gates. *Physical Review Letters*, *109*, 020503.
- Heinzen, D. J., & Wineland, D. J. (1990). Quantum-limited cooling and detection of radio-frequency oscillations by laser-cooled ions. *Physical Review A*, *42*, 2977–2994.
- High-fidelity ion state detection using trap-integrated avalanche photodiodes. (n.d.).
- Hite, D. A., Colombe, Y., Wilson, A. C., Allcock, D. T. C., Leibfried, D., Wineland, D. J., & Pappas, D. P. (2013). Surface science for improved ion traps. *MRS Bulletin*, *38*(10), 826.
- Hong, Y., Durso-Sabina, E., Hayes, D., & Lucas, A. (2024). Entangling four logical qubits beyond break-even in a nonlocal code. *arXiv preprint arXiv:2406.02666*.
- Hou, P.-Y., Wu, J. J., Erickson, S. D., Zarantonello, G., Brandt, A. D., Cole, D. C., ... Leibfried, D. (2024). Indirect Cooling of Weakly Coupled Trapped-Ion Mechanical Oscillators. *Physical Review X*, *14*(2), 021003.
- House, M. G. (2008). Analytic model for electrostatic fields in surface-electrode ion traps. *Physical Review A*, *78*, 033402.
- <https://cordis.europa.eu/project/rcn/218553/factsheet/en>. (n.d.).

- <https://www.translume.com/index.php/item/239-ion-traps-with-integrated-optics>.
(n.d.).
- Huie, W., Li, L., Chen, N., Hu, X., Jia, Z., Sun, W. K. C., & Covey, J. P. (2023). Repetitive readout and real-time control of nuclear spin qubits in ^{171}Yb atoms. *PRX Quantum*, *4*(3), 030337.
- Ivory, M., Setzer, W. J., Karl, N., McGuinness, H., DeRose, C., Blain, M., ... Parazzoli, L. P. (2021). Integrated Optical Addressing of a Trapped Ytterbium Ion. *Physical Review X*, *11*, 041033. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevX.11.041033> doi: 10.1103/PhysRevX.11.041033
- Jain, S., Sägerser, T., Hrmo, P., Torkzaban, C., Stadler, M., Oswald, R., ... Home, J. (2024). Penning micro-trap for quantum computing. *Nature*, *627*(8004), 510–514.
- Johansson, C., Uhlig, S., Tageman, O., Alping, A., Haglund, J., Robertsson, M., ... Frohlich, L. (2003). Microwave circuits in multilayer inorganic-organic polymer thin film technology on laminate substrates. *IEEE Transactions on Advanced Packaging*, *26*(1), 81. doi: 10.1109/TADV.2003.811548
- Jost, J. D. (2010). *Entangled Mechanical Oscillators* (Unpublished doctoral dissertation). University of Colorado.
- Kang, M., Campbell, W. C., & Brown, K. R. (2023). Quantum error correction with metastable states of trapped ions using erasure conversion. *PRX Quantum*, *4*(2), 020358.
- Kasprowicz, G., Harty, T., Bourdeauducq, S., Jördens, R., Allcock, D., Nadlinger, D., ... Nowicka, D. (2022). Urukul – Open-source Frequency Synthesizer Module for Quantum Physics. *International Journal of Electronics and Telecommunications*, *68*(1), 123–128. Retrieved from <https://doi.org/10.24425/ijet.2022.139859> doi: 10.24425/ijet.2022.139859
- Kasprowicz, G., Kulik, P., Gaska, M., Przywozki, T., Pozniak, K., Jarosinski, J., ... Pozniak, K. (2020). ARTIQ and Sinara: Open Software and Hardware Stacks for Quantum Physics. In *Osa quantum 2.0 conference* (p. QTu8B.14). Optica Publishing Group. Retrieved from <https://opg.optica.org/abstract.cfm?URI=QUANTUM-2020-QTu8B.14> doi: 10.1364/QUANTUM.2020.QTu8B.14
- Kielpinski, D., Monroe, C., & Wineland, D. J. (2002). Architecture for a large-scale ion-trap quantum computer. *Nature*, *417*(6890), 709–711.

- Kienzler, D. (2015). *Quantum Harmonic Oscillator State Synthesis by Reservoir Engineering* (Unpublished doctoral dissertation). ETH Zürich.
- King, B. E., Wood, C. S., Myatt, C. J., Turchette, Q. A., Leibfried, D., Itano, W. M., ... Wineland, D. J. (1998). Cooling the Collective Motion of Trapped Ions to Initialize a Quantum Register. *Physical Review Letters*, *81*(7), 1525. Retrieved from <http://dx.doi.org/10.1103/PhysRevLett.81.1525> doi: 10.1103/physrevlett.81.1525
- Kobayashi, F., & Nagayama, S. (2024). Erasure-tolerance protocol for the surface codes on Rydberg atomic quantum computers. *arXiv preprint arXiv:2404.12656*.
- Laine, R. M. (Ed.). (1992). *Inorganic and Organometallic Polymers with Special Properties*. Springer Netherlands. doi: 10.1007/978-94-011-2612-0
- Langford, N. K., Ramelow, S., Prevedel, R., Munro, W. J., Milburn, G. J., & Zeilinger, A. (2011). Efficient quantum computing using coherent photon conversion. *Nature*, *478*(11243). Retrieved from <https://www.nature.com/articles/nature10463> doi: 10.1038/nature10463
- Leibfried, D., DeMarco, B., Meyer, V., Lucas, D., Barrett, M., Britton, J., ... Wineland, D. J. (2003). Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, *422*, 412-415. Retrieved from <https://www.nature.com/articles/nature01492>
- Leibfried, D., DeMarco, B., Meyer, V., Rowe, M., Ben-Kish, A., Britton, J., ... Rosenband, T. (2002). Trapped-Ion Quantum Simulator: Experimental Application to Nonlinear Interferometers. *Physical Review Letters*, *89*(24), 247901. Retrieved from <http://dx.doi.org/10.1103/PhysRevLett.89.247901> doi: 10.1103/physrevlett.89.247901
- Lis, J. W., Senoo, A., McGrew, W. F., Rönchen, F., Jenkins, A., & Kaufman, A. M. (2023). Midcircuit operations using the *omg* architecture in neutral atom arrays. *Physical Review X*, *13*(4), 041035.
- Lloyd, S., & Braunstein, S. L. (1999, Feb). Quantum Computation over Continuous Variables. *Physical Review Letters*, *82*, 1784–1787. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevLett.82.1784> doi: 10.1103/PhysRevLett.82.1784

- Lucas, D., Ramos, A., Home, J., McDonnell, M., Nakayama, S., Stacey, J.-P., . . . Steane, A. (2004). Isotope-selective photoionization for calcium ion trapping. *Physical Review A*, *69*(1), 012711.
- Löschnauer, C. M., Toba, J. M., Hughes, A. C., King, S. A., Weber, M. A., Srinivas, R., . . . Harty, T. P. (2024). Scalable, high-fidelity all-electronic control of trapped-ion qubits. Retrieved from <https://arxiv.org/abs/2407.07694>
- Ma, S., Liu, G., Peng, P., Zhang, B., Jandura, S., Claes, J., . . . Thompson, J. D. (2023). High-fidelity gates and mid-circuit erasure conversion in an atomic qubit. *Nature*, *622*(7982), 279–284.
- Malinowski, M., Allcock, D. T. C., & Ballance, C. J. (2023). How to wire a 1000-qubit trapped-ion quantum computer. *PRX Quantum*, *4*(4), 040313.
- Manovitz, T., Shapira, Y., Gazit, L., Akerman, N., & Ozeri, R. (2022). Trapped-Ion Quantum Computer with Robust Entangling Gates and Quantum Coherent Feedback. *PRX Quantum*, *3*, 010347.
- Meekhof, D. M., Leibfried, D., Monroe, C., King, B. E., Itano, W. M., & Wineland, D. J. (1997). Experimental Creation and Measurement of Motional Quantum States of a Trapped Ion. *Brazilian Journal of Physics*, *27*(2), 15.
- Meekhof, D. M., Monroe, C., King, B. E., Itano, W. M., & Wineland, D. J. (1996). Generation of nonclassical motional states of a trapped atom. *Physical Review Letters*, *76*(11), 1796.
- Mehta, K. K., Bruzewicz, C. D., McConnell, R., Ram, R. J., Sage, J. M., & Chiaverini, J. (2016). Integrated optical addressing of an ion qubit. *Nature Nanotechnology*, *11*, 1066. Retrieved from <https://doi.org/10.1038/nnano.2016.139>
- Mehta, K. K., Zhang, C., Malinowski, M., Nguyen, T.-L., Stadler, M., & Home, J. P. (2020). Integrated optical multi-ion quantum logic. *Nature*, *586*(7830), 533.
- Metzner, J., Quinn, A., Brudney, S., Moore, I. D., Burd, S. C., Wineland, D. J., & Allcock, D. T. C. (2024). Two-mode squeezing and SU (1, 1) interferometry with trapped ions. *Physical Review A*, *110*(2), 022613.
- Metzner, J. M. (2024). *The Sound of Ions: Using Trapped Atomic Ion Motion for Quantum Computation and Sensing* (Unpublished doctoral dissertation). PhD thesis, University of Oregon.

- Mølmer, K., & Sørensen, A. (1999). Multiparticle Entanglement of Hot Trapped Ions. *Physical Review Letters*, *82*, 1835–1838. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevLett.82.1835> doi: 10.1103/PhysRevLett.82.1835
- Monroe, C., Raussendorf, R., Ruthven, A., Brown, K. R., Maunz, P., Duan, L.-M., & Kim, J. (2014). Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects. *Physical Review A*, *89*(2), 022317.
- Monz, T., Nigg, D., Martinez, E. A., Brandl, M. F., Schindler, P., Rines, R., . . . Blatt, R. (2016). Realization of a scalable Shor algorithm. *Science*, *351*(6277), 1068–1070.
- Moore, I. D. (2023). *Easy on the Ions: Photon Scattering Errors from Far-Detuned Raman Beams in Trapped-Ion Qubits* (Unpublished doctoral dissertation). PhD thesis, University of Oregon.
- Moore, I. D., Campbell, W. C., Hudson, E. R., Boguslawski, M. J., Wineland, D. J., & Allcock, D. T. C. (2023). Photon scattering errors during stimulated Raman transitions in trapped-ion qubits. *Physical Review A*, *107*, 032413.
- Morigi, G., Eschner, J., & Keitel, C. H. (2000). Ground State Laser Cooling Using Electromagnetically Induced Transparency. *Physical Review Letters*, *85*, 4458–4461. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevLett.85.4458> doi: 10.1103/PhysRevLett.85.4458
- Mount, E., Kabytayev, C., Crain, S., Harper, R., Baek, S.-Y., Vrijsen, G., . . . Kim, J. (2015). Error compensation of single-qubit gates in a surface-electrode ion trap using composite pulses. *Physical Review A*, *92*, 060301. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevA.92.060301> doi: 10.1103/PhysRevA.92.060301
- Murali, P., Debroy, D. M., Brown, K. R., & Martonosi, M. (2020). Architecting noisy intermediate-scale trapped ion quantum computers. In *2020 acm/ieee 47th annual international symposium on computer architecture (isca)* (pp. 529–542).
- Myerson, A., Szwer, D., Webster, S., Allcock, D., Curtis, M., Imreh, G., . . . Lucas, D. (2008). High-fidelity readout of trapped-ion qubits. *Physical Review Letters*, *100*(20), 200502.

- Nie, X. R., Roos, C. F., & James, D. F. V. (2009). Theory of cross phase modulation for the vibrational modes of trapped ions. *Physical Letters A*, *373*(4), 422. Retrieved from <http://dx.doi.org/10.1016/j.physleta.2008.11.045> doi: 10.1016/j.physleta.2008.11.045
- Nielsen, M. A., & Chuang, I. L. (2010). *Quantum Computation and Quantum Information*. Cambridge university press.
- NIST. (2017). *Electrode*. <https://github.com/nist-ionstorage/electrode>. GitHub.
- Ortiz-Gutiérrez, L., Gabrielly, B., Muñoz, L. F., Pereira, K. T., Filgueiras, J. G., & Villar, A. S. (2017). Continuous variables quantum computation over the vibrational modes of a single trapped ion. *Optics Communications*, *397*, 166–174.
- Oxford. (2019). *Wavemeter analysis and display*. <https://github.com/OxfordIonTrapGroup/wand>. GitHub.
- Oxford. (2024). *Ndscan*. <https://github.com/OxfordIonTrapGroup/ndscan>. GitHub.
- Pearre, B. W., Michas, C., Tsang, J.-M., Gardner, T. J., & Otchy, T. M. (2019). Fast micron-scale 3D printing with a resonant-scanning two-photon microscope. *Additive Manufacturing*, *30*, 100887. Retrieved from <https://www.sciencedirect.com/science/article/pii/S2214860418303944> doi: <https://doi.org/10.1016/j.addma.2019.100887>
- Quinn, A., Brown, M., Gardner, T. J., & Allcock, D. T. C. (2022). Geometries and fabrication methods for 3D printing ion traps. Retrieved from <https://arxiv.org/abs/2205.15892>
- Quinn, A., Metzner, J., Muldoon, J. E., Moore, I. D., Brudney, S., Das, S., . . . Joglekar, Y. N. (2023). Observing super-quantum correlations across the exceptional point in a single, two-level trapped ion. Retrieved from <https://arxiv.org/abs/2304.12413>
- Roos, C. F., Monz, T., Kim, K., Riebe, M., Häffner, H., James, D. F. V., & Blatt, R. (2008). Nonlinear coupling of continuous variables at the single quantum level. *Physical Review A*, *77*, 040302. Retrieved from <http://dx.doi.org/10.1103/PhysRevA.77.040302> doi: 10.1103/physreva.77.040302
- Roos, C. F., Riebe, M., Häffner, H., Hansel, W., Benhelm, J., Lancaster, G. P., . . . Blatt, R. (2004). Control and measurement of three-qubit entangled states. *science*, *304*(5676), 1478–1480.

- Salumbides, E., Maslinskas, V., Dildar, I., Wolf, A., Van Duijn, E.-J., Eikema, K., & Ubachs, W. (2011). High-precision frequency measurement of the 423-nm Ca I line. *Physical Review A—Atomic, Molecular, and Optical Physics*, *83*(1), 012502.
- Schmied, R. (2010). Electrostatics of gapped and finite surface electrodes. *New Journal of Physics*, *12*, 023038. Retrieved from <https://doi.org/10.1088/1367-2630/12/2/023038> doi: 10.1088/1367-2630/12/2/023038
- Sefa, M., Fedchak, J. A., & Scherschligt, J. (2017). Investigations of medium-temperature heat treatments to achieve low outgassing rates in stainless steel ultrahigh vacuum chambers. *Journal of Vacuum Science & Technology A*, *35*(4), 041601.
- Seidelin, S., Chiaverini, J., Reichle, R., Bollinger, J. J., Leibfried, D., Britton, J., . . . Wineland, D. J. (2006). Microfabricated Surface-Electrode Ion Trap for Scalable Quantum Information Processing. *Physical Review Letters*, *96*, 253003.
- Setzer, W. J., Ivory, M., Slobodyan, O., Van Der Wall, J. W., Parazzoli, L. P., Stick, D., . . . McGuinness, H. J. (2021). Fluorescence detection of a trapped ion with a monolithically integrated single-photon-counting avalanche diode. *Applied Physics Letters*, *119*, 154002. Retrieved from <https://doi.org/10.1063/5.0055999> doi: 10.1063/5.0055999
- Sherman, J., Curtis, M., Szwer, D., Allcock, D., Imreh, G., Lucas, D., & Steane, A. (2013). Experimental recovery of a qubit from partial collapse. *Physical Review Letters*, *111*(18), 180501.
- Shimasaki, T., Edwards, E. R., & DeMille, D. (2019). Injection-locking of a fiber-pigtailed high-power laser to an external cavity diode laser via a fiber optic circulator. *Review of Scientific Instruments*, *90*(5).
- Sinara. (2020). *Line trigger*. https://github.com/sinara-hw/Line_trigger. GitHub.
- Slichter, D. H., Verma, V. B., Leibfried, D., Mirin, R. P., Nam, S. W., & Wineland, D. J. (2017). UV-sensitive superconducting nanowire single photon detectors for integration in an ion trap. *Optics Express*, *25*, 8705. Retrieved from <http://www.opticsexpress.org/abstract.cfm?URI=oe-25-8-8705> doi: 10.1364/OE.25.008705

- Song, H.-F., Tang, Y.-B., Chen, S.-L., Du, L.-J., Huang, Y., Guan, H., & Gao, K.-L. (2019, Nov). Combined experimental and theoretical probe of the branching fractions of the $4P_{3/2}$ state in $^{40}\text{Ca}^+$. *Physical Review A*, *100*, 052505. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevA.100.052505> doi: 10.1103/PhysRevA.100.052505
- Steck, D. A. (2007). Quantum and atom optics.
- Stuart, J., Panock, R., Bruzewicz, C. D., Sedlacek, J. A., McConnell, R., Chuang, I. L., ... Chiaverini, J. (2019). Chip-Integrated Voltage Sources for Control of Trapped Ions. *Physical Review Applied*, *11*, 024010. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevApplied.11.024010> doi: 10.1103/PhysRevApplied.11.024010
- Sutherland, R. T., & Srinivas, R. (2021). Universal hybrid quantum computing in trapped ions. *Physical Review A*, *104*, 032609.
- Tanaka, U., Morita, I., & Urabe, S. (2007). Selective loading and laser cooling of rare calcium isotope $^{43}\text{Ca}^+$. *Applied Physics B*, *89*, 195–200.
- Todaro, S. L., Verma, V. B., McCormick, K. C., Allcock, D. T. C., Mirin, R. P., Wineland, D. J., ... Slichter, D. H. (2021). State Readout of a Trapped Ion Qubit Using a Trap-Integrated Superconducting Photon Detector. *Physical Review Letters*, *126*, 010501. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevLett.126.010501> doi: 10.1103/PhysRevLett.126.010501
- Toyoda, K., Hiji, R., Noguchi, A., & Urabe, S. (2015). Hong–Ou–Mandel interference of two phonons in trapped ions. *Nature*, *527*(7576), 74–77.
- Turchette, Q. A., Kielpinski, King, B. E., Leibfried, D., Meekhof, D. M., Myatt, C. J., ... Itano, W. M. (2000). Heating of trapped ions from the quantum ground state. *Physical Review A*, *61*, 063418. Retrieved from <http://dx.doi.org/10.1103/PhysRevA.61.063418> doi: 10.1103/physreva.61.063418
- Vizvary, S. R., Wall, Z. J., Boguslawski, M. J., Bareian, M., Derevianko, A., Campbell, W. C., & Hudson, E. R. (2024). Eliminating Qubit-Type Cross-Talk in the omg Protocol. *Physical Review Letters*, *132*(26), 263201.
- Warring, U., Ospelkaus, C., Colombe, Y., Brown, K. R., Amini, J. M., Carsjens, M., ... Wineland, D. J. (2013). Techniques for microwave near-field quantum control of trapped ions. *Physical Review A*, *87*, 013437. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevA.87.013437> doi: 10.1103/PhysRevA.87.013437

- Weber, M., Gely, M., Hanley, R., Harty, T., Leu, A., Löschnauer, C., ... Lucas, D. (2024). Robust and fast microwave-driven quantum logic for trapped-ion qubits. *Physical Review A*, *110*(1), L010601.
- Wei, W., Hao, P., Ma, Z., Zhang, H., Pang, L., Wu, F., ... Lu, Z. (2022). Measurement and suppression of magnetic field noise of trapped ion qubit. *Journal of Physics B: Atomic, Molecular and Optical Physics*, *55*(7), 075001.
- Wesenberg, J. H. (2008). Electrostatics of surface-electrode ion traps. *Physical Review A*, *78*(6), 063410. Retrieved from <http://dx.doi.org/10.1103/PhysRevA.78.063410> doi: 10.1103/physreva.78.063410
- Wesenberg, J. H. (2009). Ideal intersections for radio-frequency trap networks. *Physical Review A*, *79*, 013416. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevA.79.013416> doi: 10.1103/PhysRevA.79.013416
- Wigner, E. (1932). On the Quantum Correction For Thermodynamic Equilibrium. *Physical Review*, *40*, 749–759. Retrieved from <https://link.aps.org/doi/10.1103/PhysRev.40.749> doi: 10.1103/PhysRev.40.749
- Wilpers, G., See, P., Gill, P., & Sinclair, A. G. (2012). A monolithic array of three-dimensional ion traps fabricated with conventional semiconductor technology. *Nature Nanotechnology*, *7*, 572. Retrieved from <https://doi.org/10.1038/nnano.2012.126>
- Wineland, D., Monroe, C., Itano, W. M., Leibfried, D., King, B. E., & Meekhof, D. M. (1998). *Experimental Issues in Coherent Quantum-State Manipulation of Trapped Atomic Ions* (Vol. 103) (No. 3). National Institute of Standards and Technology.
- Wineland, D. J., & Dehmelt, H. G. (1975). Principles of the stored ion calorimeter. *Journal of Applied Physics*, *46*(2), 919-930. Retrieved from <https://doi.org/10.1063/1.321602> doi: 10.1063/1.321602
- Wineland, D. J., Drullinger, R. E., & Walls, F. L. (1978). Radiation-Pressure Cooling of Bound Resonant Absorbers. *Physical Review Letters*, *40*, 1639–1642. Retrieved from <https://link.aps.org/doi/10.1103/PhysRevLett.40.1639> doi: 10.1103/PhysRevLett.40.1639

- Wu, Y., Kolkowitz, S., Puri, S., & Thompson, J. D. (2022). Erasure conversion for fault-tolerant quantum computing in alkaline earth Rydberg atom arrays. *Nature Communications*, *13*(1), 4657.
- Xu, S., Xia, X., Yu, Q., Khan, S., Megidish, E., You, B., . . . Häffner, H. (2023). 3D-Printed Micro Ion Trap Technology for Scalable Quantum Information Processing. *arXiv preprint arXiv:2310.00595*.
- Yanagimoto, R., Onodera, T., Ng, E., Wright, L. G., McMahon, P. L., & Mabuchi, H. (2020). Engineering a Kerr-Based Deterministic Cubic Phase Gate via Gaussian Operations. *Physical Review Letters*, *124*(24), 240503. Retrieved from <http://dx.doi.org/10.1103/PhysRevLett.124.240503> doi: 10.1103/physrevlett.124.240503
- Yang, H.-X., Ma, J.-Y., Wu, Y.-K., Wang, Y., Cao, M.-M., Guo, W.-X., . . . Duan, L.-M. (2022). Realizing coherently convertible dual-type qubits with the same ion species. *Nature Physics*, *18*(9), 1058–1061.
- Yu, Q., Alonso, A. M., Caminiti, J., Beck, K. M., Sutherland, R. T., Leibfried, D., . . . Häffner, H. (2022). Feasibility study of quantum computing using trapped electrons. *Physical Review A*, *105*(2), 022420.
- Zheng, Y., Hahn, O., Stadler, P., Holmvall, P., Quijandría, F., Ferraro, A., & Ferrini, G. (2021). Gaussian conversion protocols for cubic phase state generation. *PRX Quantum*, *2*(1), 010327.