


Benchmarking Memory and Logic Gates for Trapped-Ion Quantum Computing



Amy C. Hughes

Oriel College

Department of Physics, University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Hilary Term 2021

Abstract

Trapped ions are a promising platform for experimental quantum computing, possessing the longest decoherence times and the highest fidelity logic gates of any candidate technology. The challenge remains to scale up ion-trap systems to larger numbers of qubits. This thesis benchmarks two fundamental operations for scalable trapped-ion quantum computing: the memory performance of a $^{43}\text{Ca}^+$ hyperfine qubit, and mixed-element entangling gates between $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$.

Decoherence is usually measured over long timescales, where the memory errors ϵ_m are large compared to state-preparation and measurement errors. Information about the small-error regime relevant for quantum computing is inferred by extrapolation. In this work we use randomised benchmarking to directly measure errors as small as $1.2(7) \times 10^{-6}$ after a storage time of 1 ms, which is around an order of magnitude smaller than would be expected based on the usual model of exponential decay. We find $\epsilon_m < 10^{-4}$ for up to 50 ms with no additional dynamical decoupling, and $\epsilon_m < 10^{-3}$ for up to 2 seconds using a simple CPMG sequence. These timescales exceed previously-demonstrated gate or measurement times in trapped-ion systems by several orders of magnitude — a requirement for quantum error correction, and a highly desirable feature for near-term processors. The qubit is robust to offsets of the external magnetic-field strength, with $\epsilon_m < 10^{-4}$ for 1 ms even at a 50 mG offset, and we identify phase noise on the reference oscillator as the limiting factor on the memory performance.

We provide a comparison of different implementations of mixed-element geometric phase gates in the same experimental system. This includes a light-shift gate which can be implemented on both ion species using a single laser, with a fidelity of 99.8(1)% or 99.7(1)%, measured using partial state tomography or interleaved randomised benchmarking respectively. We also demonstrate several Mølmer-Sørensen gates with measured fidelities of up to 99.6(2)% (by partial state tomography); this mechanism is more susceptible to errors arising from instability of the external magnetic field. For the first time, this pushes mixed-element entangling gate fidelities over the fault-tolerant threshold level (above which error correction is possible), and puts them on par with state-of-the-art single-species gates.

Acknowledgements

I am very grateful to my supervisor, Prof. David Lucas, for giving me the opportunity to play in his labs, and for his atomic physics lectures which inspired me to do so. Thanks for always having time for a chat, for encouraging me, and for checking that the labs were still standing during the first COVID lockdown.

The work in this thesis was only possible thanks to the combined efforts of many members of the Oxford ion trap group. From the Lab 2 crew: Dr. Joe Goodwin first showed me how to run an experiment, and always offered a level-headed perspective or a hand with a ‘back-of-the-paper-plate’ calculation. Dr. Martin Sepiol helped me to trap my first ion, Dr. Tom Harty kindly shared some of his unrivalled knowledge of phase noise, and Dr. Jamie Tarlton took the first memory data and started the all-important ion trappers’ cake club.

In Lab 1, I am lucky to have had the guidance of Dr. Vera Schäfer, who taught me almost everything I know about two-qubit gates and was instrumental in keeping things on track. Dr. Keshav Thirumalai was an integral part of the blade-trap team, and showed me the species-mixing ropes. Thank you, both, for being a joy to work with. Oana Bazavan is a wonderful new addition and is taking the experiment forward masterfully. Dr. Chris Ballance somehow always had a minute to help solve a problem, David Nadlinger offered a wealth of expertise and implemented gate set tomography, and Prof. Guido Pagano provided useful tips for calibrating Mølmer-Sørensen gates.

Much appreciation goes to David, Vera, Joe and Tom for helpful feedback on this thesis, and to the rest of the group with whom I crossed over — Dr. Jochen Wolf, Dr. Laurent Stephenson, Dr. Diana Aude Craik, Dr. Tim Ballance, Dr. Ryan Hanley, Dr. Shuoming An, Dr. Gabriel Araneda, Dr. Shaobo Gao, Dr. Ezra Kassa, Clemens Löschnauer, Beth Nichol, Peter Drmota, Marius Weber, Will Hughes, Ana Sotirova, Dougal Main, Sebastian Saner, Lee Peleg and Prof. Andrew Steane — it’s been a pleasure.

I gratefully acknowledge EPSRC funding via the NQIT and QCS hubs, and a Graduate Teaching and Research Scholarship from Oriel College. Thanks to Prof. Andrew Boothroyd and colleagues at Oriel, and to my students, who were a genuine pleasure to teach and to drink with (not simultaneously).

I have spent countless mornings on the river, breakfasts at Combibos, and evenings in the bar with University College Boat Club, and am proud to have shared an 8+ with so many inspiring women. To Ellie, Aidan and Helena, thanks for all the pizza, games and good times, and also to Florentyna for always being a message away. Thank you so much to my parents and to my brother, Fred, for offering endless support and cups of tea, and for letting me talk to the kittens on Zoom (which improved the thesis-writing process immensely). Finally, to Rory, thank you for still being my person after nine years, three lockdowns and a PhD (each).

Contents

1	Introduction	1
1.1	Quantum computing	1
1.2	Ion-trap quantum computing	3
1.2.1	Mixed-species entanglement	5
1.3	Benchmarking quantum computers	6
1.4	Thesis outline	7
2	Trapped-ion qubits	9
2.1	Linear Paul trap	9
2.2	Ions	13
2.2.1	$^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$	14
2.2.2	Calcium qubits	15
2.2.2.1	Hyperfine structure	15
2.2.2.2	Clock qubit	19
2.2.2.3	Stretch qubit	19
2.2.3	Strontium qubits	20
2.2.3.1	Zeeman qubit	21
2.2.3.2	Optical qubit	21
2.2.4	Photoionisation	22
2.2.5	Cooling	23
2.2.5.1	Doppler cooling	23
2.2.5.2	Dark-resonance cooling	24
2.2.5.3	Resolved sideband cooling	26
2.2.6	State preparation	27
2.2.7	Readout	27
3	Theory	30
3.1	Ion-light interactions	30
3.1.1	The two-level atom	31
3.1.2	Multipole interactions	31
3.1.2.1	Electric dipole interactions	32
3.1.2.2	Magnetic dipole and electric quadrupole interactions	36
3.1.3	Raman interactions	36
3.1.4	Coupling to the ion's motion	38
3.1.4.1	The Lamb–Dicke parameter	39

	3.1.4.2	The Lamb–Dicke limit	40
	3.1.4.3	Outside the Lamb–Dicke limit	43
3.2		Qubit decoherence	43
	3.2.1	Types of decoherence	44
	3.2.1.1	Population decay	44
	3.2.1.2	Dephasing	45
	3.2.2	Qubit phase noise	46
	3.2.2.1	Measuring decoherence	46
	3.2.2.2	Memory error and T_2^*	48
	3.2.2.3	Dephasing due to an arbitrary noise source	48
	3.2.2.4	Common types of noise	50
	3.2.3	Dynamical decoupling	51
	3.2.4	Ion-light interactions in the presence of decoherence	53
	3.2.4.1	Damped Rabi oscillations	53
	3.2.4.2	Scattering	55
3.3		Mixed-species geometric phase gates	57
	3.3.1	Forced quantum harmonic oscillator	58
	3.3.2	Spin-dependent forces	60
	3.3.3	Mølmer–Sørensen / $\sigma_\phi \otimes \sigma_\phi$ gate	62
	3.3.3.1	Beam geometry	67
	3.3.3.2	Carrier coupling	68
	3.3.3.3	Light shifts	69
	3.3.4	Light-shift / $\sigma_z \otimes \sigma_z$ gate	70
	3.3.4.1	Simple same-species case	72
	3.3.4.2	Mixed-species case	75
	3.3.5	Two-qubit gate errors and mitigation	78
	3.3.5.1	Photon-scattering errors	79
	3.3.5.2	Thermal errors	80
	3.3.5.3	Off-resonant excitation errors	82
	3.3.5.4	Errors due to mis-set gate parameters	84
	3.3.5.5	Errors due to qubit frequency offsets	85
3.4		Randomised benchmarking	87
	3.4.1	Standard randomised benchmarking	88
	3.4.1.1	Twirling	88
	3.4.1.2	The Clifford group	89
	3.4.1.3	Standard randomised benchmarking	90
	3.4.2	Interleaved randomised benchmarking	93
	3.4.2.1	Simplified method	94
4		Apparatus	96
4.1		Traps	97
	4.1.1	Apparatus A: planar trap	97
	4.1.2	Apparatus B: blade trap	101
4.2		Vacuum systems	103
4.3		Imaging	104

4.4	Magnetic field	105
4.4.1	Coils	105
4.4.2	Field stabilisation	105
4.5	Reference oscillator	106
4.6	Microwave and RF drive	106
4.7	Lasers	109
4.7.1	Diode lasers	109
4.7.2	Raman lasers	113
4.7.3	674 nm laser	115
4.8	Experimental control	116
5	Memory benchmarking	118
5.1	Preliminary work	118
5.1.1	Ramsey experiments	119
5.1.2	Benchmarking of single-qubit microwave gates	122
5.1.2.1	Standard Cliffords	122
5.1.2.2	BB1-protected Cliffords	123
5.2	Memory benchmarking results	125
5.2.1	Experimental method	125
5.2.2	Results at the clock field	128
5.2.3	Detuned magnetic field	129
5.2.4	Memory performance with dynamical decoupling	131
5.3	Sources of decoherence	131
5.3.1	Phase noise measurements	131
5.3.2	Magnetic-field noise	136
5.3.3	AC Zeeman shift	139
5.4	Summary	140
6	Mixed-species two-qubit gates	142
6.1	Calibration and characterisation	142
6.1.1	Recrystallisation and re-ordering	142
6.1.2	Crystal tilt	143
6.1.3	Mode frequencies	144
6.1.4	Cooling	144
6.1.5	Heating rate	146
6.1.6	Single-qubit randomised benchmarking	146
6.1.7	Pulse shaping	149
6.1.8	SPAM correction	149
6.1.9	Spin dephasing	150
6.2	Light-shift gate	150
6.2.1	Raman beams	150
6.2.2	Ion spacing	152
6.2.3	Gate implementation	154
6.2.3.1	Pulse sequence	154
6.2.3.2	Driving force phase	155

6.2.3.3	Gate parameters	155
6.2.4	Partial state tomography	157
6.2.5	Randomised benchmarking	160
6.2.6	Gate set tomography	164
6.3	Mølmer–Sørensen gates	165
6.3.1	Raman beams	165
6.3.2	Bichromat balance calibration	166
6.3.3	Gate implementation	168
6.3.3.1	Gate parameters	168
6.3.4	MS gate on the clock qubit	168
6.3.5	MS gates on other qubits	173
6.3.6	MS gates using the 674 nm laser	175
6.4	Summary	177
6.5	Comparison	179
7	Conclusions	181
7.1	Memory performance	181
7.1.1	Comparison to other work	182
7.2	Mixed-species two-qubit gates	183
7.2.1	Comparison to other work	185
7.2.1.1	Other mixed-species gates	185
7.2.1.2	Other two-qubit randomised benchmarking	187
7.2.1.3	Scaling up randomised benchmarking	187
7.3	Outlook	189
A	Mølmer–Sørensen gate beam geometries	191
A.1	Phase-insensitive geometry	192
A.2	Phase-sensitive geometry	193
A.3	Phase-insensitive sequence	194
B	Ramsey measurements of small memory errors	195
C	Clifford decompositions	197
C.1	Single-qubit Clifford group	197
C.2	Two-qubit Clifford group	198
	Bibliography	202

List of Figures

2.1	Level structures of $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ at 146 G	16
2.2	Hyperfine structure in the $^{43}\text{Ca}^+$ ground level	17
2.3	Photoionisation	22
2.4	Doppler cooling and state preparation in $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$	25
2.5	Shelving for readout of $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$	28
3.1	The two-level atom model	31
3.2	Driving a Raman transition	37
3.3	Dephasing on the Bloch sphere	45
3.4	CPMG dynamical decoupling	52
3.5	Damped Rabi oscillations	54
3.6	Phase-space trajectories for a general geometric phase gate	62
3.7	MS gate phase-space trajectories	66
3.8	LS gate phase-space trajectories	76
3.9	Phase-space trajectories with Walsh modulation	85
3.10	Standard randomised benchmarking sequence	91
3.11	Interleaved randomised benchmarking sequences	94
4.1	Surface trap	98
4.2	Blade trap	101
4.3	$^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ mode frequencies	103
4.4	Beam geometries	112
4.5	Schematic of apparatus	117
5.1	Ramsey experiments	121
5.2	Benchmarking of standard single-qubit Clifford gates with microwaves	124
5.3	BB1 pulse sequence	125
5.4	Memory benchmarking	127
5.5	Clock qubit magnetic-field dependence	129
5.6	Memory benchmarking with dynamical decoupling	132
5.7	Phase noise measurement setup	133
5.8	Phase detector calibration	134
5.9	Phase noise measurements	135
5.10	Magnetic-field servo corrections	137
5.11	Stretch transition coherence time measurements	138

6.1	Re-ordering a mixed-species crystal	143
6.2	Temperature diagnostics for $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$	145
6.3	$^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ heating rates	146
6.4	Simultaneous randomised benchmarking of single-qubit operations on two ions	147
6.5	Interleaved randomised benchmarking of single-qubit gates with lasers	148
6.6	Light-shift gate beams	151
6.7	Two-qubit gate errors vs. Raman detuning	151
6.8	Ion spacing calibration	153
6.9	Light-shift gate dynamics	156
6.10	Light-shift gate parity scan	159
6.11	Bell-state fidelities using light-shift gate	160
6.12	Interleaved randomised benchmarking of light-shift gate	162
6.13	Mølmer-Sørensen gate beams	166
6.14	Bichromat balance calibration	167
6.15	MS gate dynamics	170
6.16	Bell-state fidelities using MS gate on clock qubit	171
6.17	MS gate parity oscillations	172
6.18	Magnetic-field oscillations	172
6.19	Parity-angle drifts	174
6.20	Bell-state fidelities using MS gate on stretch qubit	174
6.21	MS gate beams with 674 nm laser	176
6.22	MS gate with 674 nm laser: parity scan	177
7.1	Comparison to previous mixed-species gate fidelities	186
A.1	Frequency configurations for the MS gate	192

Chapter 1

Introduction

1.1 Quantum computing

Quantum computers harness quantum-mechanical phenomena in order to process and store information in ways which are impossible to replicate within the confines of classical physics. The origins of quantum computing can arguably be traced back to the early 1980s, when Richard Feynman pointed out the inefficiency of simulating quantum physics with classical computers, instead proposing the use of a quantum system itself for this purpose [55]. A few years later, David Deutsch first proposed the idea of a ‘universal’ quantum computer, which could simulate any other quantum computer with arbitrary precision [48].

Feynman noted that a full description of the collective quantum state of a number of particles N cannot be represented using a classical system with a number of elements proportional to N . Due to non-classical properties such as superposition and entanglement, the amount of information required to represent N two-level quantum-mechanical elements, or ‘qubits’, doubles each time N is increased by one; this exponential scaling means that to describe just 50 qubits would require 10^{15} complex numbers, or 1 petabyte of information at 1 byte per complex number. In

the 1990s, the first quantum algorithms were discovered, which make use of these non-classical properties to deliver more efficient solutions to complex problems than can be achieved in classical systems [49, 157, 64]. With their fundamentally different modes of operation, quantum computers have further potential to outperform their classical counterparts in solving certain computational problems, with applications anticipated in cryptography, drug design and development, quantum chemistry, financial modelling and machine learning.

Storing and processing quantum information, however, comes with additional challenges. The preservation of a quantum state is generally a more delicate issue than the preservation of classical data. For example, a classical bit stored as a low voltage (e.g. 0 – 2 V) or a high voltage (e.g. 3 – 5 V) across an electronic component is robust to noise from external sources, provided the noise does not cause the voltage to fluctuate outside the acceptable range. On the other hand, for a qubit in a general state $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, any change to α or β introduces a finite change to the quantum state, making it more susceptible to errors. Nonetheless, the invention of quantum error correction [167, 158] showed that it is possible for an arbitrarily long quantum computation to succeed as long as error rates are kept below a manageable threshold [99], leading to increased interest in the experimental realisation of a universal quantum computer.

In 1998, David DiVincenzo outlined five criteria which must be met in any physical implementation of a quantum computer [50, 51]. They are:

1. *A scalable physical system with well characterised qubits*
2. *The ability to initialise the state of the qubits to a simple state*
3. *Long relevant decoherence times (much longer than the gate operation time)*
4. *A universal set of quantum gates*
5. *A qubit-specific measurement capability.*

The first quantum algorithms were experimentally demonstrated in the same year, in nuclear magnetic resonance systems [40, 86]. Today, experimental quantum computing is an active area of research in academic and commercial settings, realising the DiVincenzo criteria across qubit platforms including trapped ions [30], superconducting circuits [96], photonic systems [160], and semiconductor materials [35].

1.2 Ion-trap quantum computing

In 1995, Ignacio Cirac and Peter Zoller proposed the trapped-ion quantum processor [41], consisting of a linear array of ions which are cooled close to their ground state of motion. Trapped ions are promising qubit candidates for many reasons. Each qubit is encoded using two of the internal quantum states of a single ion, which are guaranteed to be identical in all ions of the same species. Suspended in ultra-high vacuum using electric fields, they are extremely well isolated from sources of decoherence in their environment. The spacing between ions in the same trap means that their internal states are also well isolated from those of their neighbours, yet they share a motional quantum state; this ‘phonon bus’ can be used to transfer information between qubits when desired. Atomic physics is well understood and can be utilised effectively to prepare, manipulate and read out the qubits via their interactions with electromagnetic fields.

In analogy with the construction of classical circuits from a primitive set of transistor-based logic gates, a universal quantum gate set can be formed from compositions of single-qubit rotations and a two-qubit entangling gate [14]. Mere months after Cirac and Zoller’s original proposal, the first two-qubit quantum logic gate in any physical system was demonstrated at NIST¹ by Chris Monroe, using trapped ions [125]. Over the last quarter of a century, all five of DiVincenzo’s criteria have been demonstrated in trapped-ion systems, with increasing fidelity. The best state-

¹National Institute of Standards and Technology, Boulder, Colorado, USA

preparation and measurement (SPAM) errors are $\sim 10^{-4}$ [67, 43], and trapped ions currently boast the longest coherence times [182] and the lowest single- [67] and two-qubit [12, 60, 166] gate errors ($\sim 10^{-6}$ and $\sim 10^{-3}$ respectively). The systems currently under development that feature the largest numbers of qubits, however, tend to be those based on superconducting qubits [90, 62, 75, 145], where significant improvements have also recently been made to coherence times and to gates [162, 95, 192].

Attention now turns to scaling up trapped-ion processors whilst keeping error rates low. The precise error threshold for fault-tolerant computation depends on the particular error-correcting code to be used, as well as the noise model and the computer architecture. The foremost error-correcting code is the surface code, for which the threshold is $\sim 1\%$ [57]; however, the further below threshold the errors can be kept, the fewer additional resources are required to correct them.

The largest trapped-ion quantum registers to date have a few tens of qubits under coherent control² [59, 190, 54, 33, 74]. For larger linear chains of ions, the motional coupling strengths between ions and to the driving field decrease, and the likelihood of unwanted coupling to spectator motional modes increases, making entangling gates more difficult to implement. Despite impressive results so far [101], the increase to significantly larger numbers of qubits is likely to require a modular approach, with communication between registers containing small numbers of ions. DiVincenzo identifies two additional criteria which are necessary for quantum communication:

6. *The ability to interconvert stationary and ‘flying’ qubits*

7. *The ability to faithfully transmit flying qubits between specified locations.*

The ‘quantum charged-coupled device’ (QCCD) blueprint [188, 91] tackles this by performing entangling gates on small numbers of ions in separate trapping zones,

²Quantum simulations have been performed using more than 50 ions [196, 24], but these systems do not aim for universal quantum computation capabilities.

and transporting ‘communication’ ions between zones. The necessary elements for this architecture — including high-fidelity operations within a single trap zone, fast shuttling of ions between zones, and parallelisation of these operations — have been demonstrated individually in many experimental setups (e.g. [72, 89]), and were recently combined in a single, 6-qubit trapped-ion processor [140].

The QCCD may be limited to ~ 1000 ions due to the increasing complexity of the control voltages required for shuttling, and does not provide scope for the transfer of quantum information across large distances. Another option is to link together registers of ions via photonic interconnects [126], where flying qubits encoded in photons facilitate fast communication across a network. Remote entanglement between two ions in spatially-separated traps connected via optical fibres has been experimentally demonstrated [170, 77]. A fully-scalable trapped-ion quantum computer architecture seems likely to combine the two approaches [28].

1.2.1 Mixed-species entanglement

The simultaneous trapping of two different ion species allows the manipulation of one species without corruption of the electronic state of the other. Mixed-species systems therefore provide access to powerful tools such as sympathetic cooling [92], which combats errors introduced by the heating of ions over long sequences of gates, and during shuttling in a QCCD [72, 140]. They also facilitate low cross-talk, individual addressing without the need to focus beams onto single ions, and quantum non-demolition measurement [80].

Different ions have different properties, which make them well suited for different applications. A high-fidelity entangling gate between two species offers the flexibility to transfer information from one to the other depending on the task at hand. In a single node of an optically-networked quantum processor, quantum states can be swapped from ‘memory’ or ‘logic’ ions to ‘communication’ ions. The communication

ions can then be entangled with photons using intense laser pulses which will not disturb the memory qubits.

The combination and entanglement of mixed species will therefore be an essential ingredient for scalable quantum processors. Other applications include ‘quantum logic’ spectroscopy of molecules or of exotic species for testing fundamental physics, and readout of atomic clocks [152, 37, 27, 153, 161].

1.3 Benchmarking quantum computers

As quantum computing capabilities increase, there is a growing interest in testing the performance of devices in a manner which is comparable across hardware platforms [144, 22]. Classical computers are often benchmarked by running standard programs which closely replicate the demands that would be placed on the system during real-world use-cases. Similar approaches, based on running particular algorithms, have been proposed and implemented in small quantum systems [190, 107, 118]. An increasingly popular measure of the performance of existing devices is the ‘quantum volume’, proposed by IBM [44], and Google recently announced ‘quantum supremacy’ over classical computers, using an approach based on random circuits [5, 25]. However, it is not yet clear what the most useful real-world applications of quantum computers will be, making it difficult to prescribe a single representative algorithm or set of algorithms which would constitute a good benchmark. Even if this were known, the implementation of such a benchmark is likely to be outside the capabilities of near-term ‘noisy, intermediate-scale quantum’ (NISQ) [142] processors, where the number of qubits is modest and error correction is yet to be implemented.

Much of the focus has been on ‘component-level’ benchmarking of individual operations. Methods such as quantum state or process tomography [141, 39] offer full characterisation but are susceptible to errors in the state-preparation and measure-

ment steps. They also scale inefficiently with the system size, which is unfavourable given the potential for cross-talk and context-dependent calibration changes when performing parallel operations on multiple qubits. The most widely used method for quantifying the level of coherent control in a system is randomised benchmarking [113, 112, 114], which provides a measure of gate fidelity that is comparable across platforms, independent of SPAM errors, and more easily scalable.

1.4 Thesis outline

$^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ form a promising pair of ion species for co-trapping in a node of a quantum network, as outlined in [133]. Magnetic-field-independent ‘clock’ qubit states in $^{43}\text{Ca}^+$ make it very well suited for storage and processing of information, whilst the relatively simpler atomic structure of $^{88}\text{Sr}^+$ lends it to photonic entanglement. This work uses randomised benchmarking and partial state tomography of Bell states to characterise the quantum memory performance of a single $^{43}\text{Ca}^+$ ion and the fidelities of several mixed-species two-qubit entangling gates between a $^{43}\text{Ca}^+$ ‘memory’ ion and a $^{88}\text{Sr}^+$ ‘communication’ ion. These operations are fundamental requirements for the implementation of an optically-networked ion-trap quantum computer based on the combination of these two elements. Along with concurrent research on remote entanglement of $^{88}\text{Sr}^+$ qubits [170, 169], this thesis represents progress towards that goal.

Chapter 2: Trapped-ion qubits discusses how we trap ions, the choice of species, and how we initialise and measure qubit states.

Chapter 3: Theory covers the use of ion-light interactions for single-qubit rotations, the role of decoherence, the theory behind the two-qubit gate mechanisms demonstrated in this thesis, and the basics of randomised benchmarking.

Chapter 4: Apparatus describes two different experimental setups — referred to

as ‘apparatus A’ and ‘apparatus B’ — which were used in this work.

Chapter 5: Memory benchmarking presents a characterisation of the memory performance of a hyperfine qubit in $^{43}\text{Ca}^+$, using randomised benchmarking to directly measure errors at the 10^{-6} level. This work is published in [156].

Chapter 6: Mixed-species two-qubit gates covers various experimental implementations of mixed-species entangling gates on $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystals, with high fidelities measured using randomised benchmarking and partial state tomography. Some of these results are published in [79].

Chapter 7: Conclusions summarises the results and compares them to other related work.

Chapter 2

Trapped-ion qubits

This chapter will discuss how we trap ions for use as qubits, and how we satisfy the second and fifth DiVincenzo criteria of initialising and measuring their quantum states. The third and fourth criteria — a universal gate set and a long coherence time — are addressed in chapter 3.

2.1 Linear Paul trap

Charged particles may be trapped using an electric field. As a result of Gauss’s law, with no external charges ($\nabla \cdot \mathbf{E} = 0$), a static electric field cannot form a trap which confines in three dimensions. One way to overcome this is to combine a static field with an oscillating one, providing confinement when averaged over one oscillation period; this is the principle of operation of a ‘Paul’ or ‘radio-frequency’ ion trap.

There are many variations of ‘linear’ Paul trap designs which function similarly. In a common configuration, a static DC voltage is applied to two electrodes in order to confine along a ‘trap axis’, $\hat{\mathbf{z}}$, whilst a radio-frequency (RF) voltage is applied to two more pairs of electrodes to form an oscillating, two-dimensional quadrupole potential in the plane perpendicular to the $\hat{\mathbf{z}}$ axis.

The potential in the trapping region has the form [105]

$$U(x, y, z, t) = \frac{V_{\text{RF}}}{2} \cos(\Omega_T t) [\alpha_z z^2 + \alpha_x x^2 - \alpha_y y^2] + \alpha' \frac{V_{\text{DC}}}{2} [2z^2 - x^2 - y^2], \quad (2.1)$$

where the oscillating voltage has amplitude V_{RF} and frequency Ω_T , the static voltage is V_{DC} , and the origin of the co-ordinate system is the centre of the trapping region. The parameters $\{\alpha_i, \alpha'\}$ are geometric factors which differ from 1 (or 0) because the shapes of the electrodes are not typically hyperbolic, so the resulting potential is only approximately quadrupolar (and end effects are included).

An ion of mass M and charge $+e$ in this potential experiences a force, leading to equations of motion which can be written in the form

$$\frac{d^2 u_i}{d\xi^2} + [a_i + 2q_i \cos(2\xi)] u_i = 0, \quad (2.2)$$

where $u_i = \{x, y, z\}$ and $\xi \equiv \Omega_T t/2$. The parameters a_i and q_i are given by

$$\begin{aligned} a_x &= -\frac{4e\alpha' V_{\text{DC}}}{\Omega_T^2 M}, & a_y &= -\frac{4e\alpha' V_{\text{DC}}}{\Omega_T^2 M}, & a_z &= \frac{8e\alpha' V_{\text{DC}}}{\Omega_T^2 M} \\ q_x &= \frac{2e\alpha_x V_{\text{RF}}}{\Omega_T^2 M}, & q_y &= -\frac{2e\alpha_y V_{\text{RF}}}{\Omega_T^2 M}, & q_z &= \frac{2e\alpha_z V_{\text{RF}}}{\Omega_T^2 M}. \end{aligned} \quad (2.3)$$

Equation (2.2) is known as the Mathieu equation and its analytic solutions are well known (see e.g. [105]). Under the conditions¹ $\{|a_i|, q_i^2\} \ll 1$, an approximate solution is

$$u_i(t) = U_{0i} \cos(\omega_i t) \left[1 - \frac{q_i}{2} \cos(\Omega_T t) \right], \quad (2.4)$$

¹These conditions approximately ensure stability of the ion's motion and may be interpreted physically; $|a_i| \ll |q_i|$ means that the average radial confinement due to the RF voltage is stronger than the radial anti-confinement from the DC, and a small q_i puts a limit on the amplitude of oscillations at frequency Ω_T . For a more thorough discussion, see e.g. [105].

which consists of harmonic oscillations at the ‘secular’ frequencies

$$\omega_i = \frac{\Omega_T}{2} \sqrt{a_i + \frac{q_i^2}{2}}, \quad (2.5)$$

as well as faster, smaller ‘micro-motion’ oscillations at the trap frequency Ω_T .

An ion sitting at the null of the RF potential experiences no excess micro-motion. Traps usually include ‘compensation’ electrodes where extra DC voltages may be applied, to offset the effect of any stray electric fields and ensure that the ions sit as close to the null as possible. We therefore neglect micro-motion and approximate $U(x, y, z, t)$ — as seen by an ion of mass M — as a harmonic pseudo-potential

$$\tilde{U}_M(x, y, z) = \frac{1}{2}M\omega_x^2 x^2 + \frac{1}{2}M\omega_y^2 y^2 + \frac{1}{2}M\omega_z^2 z^2. \quad (2.6)$$

The motion of a single ion in this pseudo-potential may be decomposed into three normal modes of oscillation, along $\hat{\mathbf{x}}$, $\hat{\mathbf{y}}$ and $\hat{\mathbf{z}}$. Two ions in a trap have two normal modes of oscillation along each direction: one ‘in-phase’ (ip) mode and one ‘out-of-phase’ (oop) mode. Expressions for the frequencies of all of these modes, for two ions of masses M_1 and M_2 , are given in [191]. The axial ip and oop mode frequencies are

$$\begin{aligned} \omega_{\text{ax,ip}} &= \sqrt{\frac{1 + \mu - \sqrt{1 - \mu + \mu^2}}{\mu}} \omega_{z,1}, \\ \omega_{\text{ax,oop}} &= \sqrt{\frac{1 + \mu + \sqrt{1 - \mu + \mu^2}}{\mu}} \omega_{z,1}, \end{aligned} \quad (2.7)$$

where $\mu = M_2/M_1$, and $\omega_{z,1}$ is the axial secular frequency for a single ion of mass M_1 in the same trap. For $M_1 = M_2$, this reduces to $\omega_{\text{ax,ip}} = \omega_z$, $\omega_{\text{ax,oop}} = \sqrt{3}\omega_z$.

The motion of the ions may be described using the creation and annihilation

mode	$b_{1,m} (^{43}\text{Ca}^+)$	$b_{2,m} (^{88}\text{Sr}^+)$
ax ip	0.453	0.892
ax oop	0.892	-0.453

Table 2.1: Eigenvector components for the axial motional modes of a $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal.

operators $a_m^\dagger$ and a_m for each normal mode m . The position $\mathbf{r}_i$ of ion i is [127, 71]

$$\mathbf{r}_i = \mathbf{r}_{\text{eq},i} + \sum_{\text{modes } m} \hat{\mathbf{r}}_m l_{i,m} (a_m + a_m^\dagger) b_{i,m}, \quad (2.8)$$

where $l_{i,m} \equiv \sqrt{\hbar/(2M_i\omega_m)}$, $\hat{\mathbf{r}}_m \in \{\hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\mathbf{z}}\}$ is the axis of mode m , and $b_{i,m}$ is the motional mode eigenvector component corresponding to ion i in mode m . The ion oscillates about its equilibrium position $\mathbf{r}_{\text{eq}}$. The root-mean-square ground-state wavefunction spread from $\mathbf{r}_{\text{eq}}$, for ion i in mode m , is $\tilde{l}_{i,m} = l_{i,m}|b_{i,m}|$.

For a single ion, $b_{i,m} = 1$. For two ions of the same species, the motion is perfectly symmetric for all ip modes ($b_{1,m} = b_{2,m} = \frac{1}{\sqrt{2}}$) and perfectly antisymmetric for all oop modes ($b_{1,m} = \frac{1}{\sqrt{2}}$, $b_{2,m} = -\frac{1}{\sqrt{2}}$). The mode eigenvector components $b_{i,m}$ for two-ion mixed-species crystals can be calculated analytically [191]; the components for the axial modes of a $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal are given in table 2.1.

Along the axial direction, the DC confinement is much larger than the confinement from the RF drive ($\alpha_z \approx 0$), such that the axial frequency of a particular ion $\omega_{z,i} \propto \sqrt{V_{\text{DC}}/M_i}$. In the radial directions we find $\omega_{x,y,i} \propto V_{\text{RF}}/M_i\Omega_T$. Therefore, the axially-confining force is independent of the ion's mass, whereas the radial force is approximately proportional to $1/M_i$.

For experiments in this thesis involving more than one ion, we choose trap parameters for which the ratio $\omega_{x,y}/\omega_z$ is large enough that the equilibrium positions of the ions lie along the z -axis, forming a ‘linear crystal’. The spacing between the ions

is determined by the balance of the axial confining force with the Coulomb repulsion due to their charge. Assuming, again, that the DC confinement dominates along the axial direction, the equilibrium z -positions of two singly-charged ions are independent of their masses, and are equal to

$$z_{\text{eq}} = \pm \left(\frac{e}{32\pi\epsilon_0\alpha'V_{\text{DC}}} \right)^{1/3} = \pm \left(\frac{e^2}{16\pi\epsilon_0 M_i \omega_{z,i}^2} \right)^{1/3}. \quad (2.9)$$

These positions may be calculated using any known axial frequency $\omega_{z,i}$ for a single ion of mass M_i in the same trap. The ion spacing is equal to $2|z_{\text{eq}}|$.

2.2 Ions

There are many factors to consider when deciding on an ion species to trap for quantum information processing applications. A table of ions of interest and their properties may be found in [30].

Atomic level structure

Typical choices are group II elements such as beryllium, magnesium, calcium, strontium and barium, as well as ytterbium; when singly ionised, these elements have a ‘hydrogen-like’ atomic energy level structure with a single valence electron. A relatively simple structure is desirable to allow good control over the ion’s internal state, and access to a strong cycling transition for cooling and fluorescence detection. The details of the level structure determine the properties of candidate qubit transitions.

Transition wavelengths

Transitions should be accessible via wavelengths which are convenient to work with and to generate. Optical components for ultra-violet (UV) wavelengths are generally more expensive and more difficult to source, and perform less well than those for

visible light. These shorter wavelengths are more likely to cause charging of trap dielectrics (leading to drift of motional mode frequencies and compensation voltages), and are more readily absorbed in air. Manufacture of on-chip ‘integrated optics’ (such as those in [120]), is also more difficult for UV wavelengths due to increased optical loss [30], which may limit the scalability of trapped-ion processors that require these wavelengths.

Mass

Heavier ions require higher trap voltages and have lower motional mode frequencies, which limits the speed of two-qubit gates. Influencing their motion using lasers requires higher beam powers, leading to increased likelihood of errors due to off-resonant excitation and photon scattering during gate operations. Lighter ions, on the other hand, are more susceptible to heating, and to errors arising from photon recoil momentum.

Compatibility of species

Co-trapping of two different atomic species comes with additional considerations. Ions with similar masses have similar requirements for stable trapping. The closer the masses, the more similar the motional mode eigenvector amplitudes, allowing efficient sympathetic cooling. The transition frequencies of the two species should complement each other; the ability to address one species without disturbing the internal state of the other is an advantage, but one may wish to avoid a large spread of required wavelengths for reasons of practicality and scalability.

2.2.1 $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$

All of the work in this thesis involves either a single $^{43}\text{Ca}^+$ ion or a two-ion mixed-species $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal. $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ are medium-mass group II ions with

similar atomic structures, the relevant fine structure levels being $\{nS_{1/2}, nP_{1/2}, nP_{3/2}, (n-1)D_{3/2}, (n-1)D_{5/2}\}$ in each case, with $n = 4$ for $^{43}\text{Ca}^+$ and $n = 5$ for $^{88}\text{Sr}^+$.

The transition wavelengths for each ion are shown in figure 2.1; they fall in the range 393 – 1092 nm, avoiding complications associated with working in the UV. Those necessary for Doppler cooling may all be produced from commercially-available diode lasers, without the need for frequency doubling. In addition, they may loosely be grouped into ‘blue’ and ‘red’ ranges for each ion, where optical components may be used flexibly within each range.


With a mass ratio of ≈ 2 , $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ may be efficiently sympathetically cooled. Their transition wavelengths are sufficiently well separated to facilitate trivial individual addressing with negligible cross-talk. Nonetheless, the 408 nm $5S_{1/2} \leftrightarrow 5P_{3/2}$ transition in $^{88}\text{Sr}^+$ and the 397 nm $4S_{1/2} \leftrightarrow 4P_{1/2}$ transition in $^{43}\text{Ca}^+$ are close enough in frequency that Raman transitions may be driven on both species using the same wavelength, if desired. This important feature allows for a significant simplification of the apparatus required to perform mixed-species two-qubit gates.

A number of different pairs of qubit states are utilised in this work; their properties are detailed in the following sections, and summarised in table 2.2.


2.2.2 Calcium qubits

2.2.2.1 Hyperfine structure

$^{43}\text{Ca}^+$ is the only stable isotope of calcium which has a non-zero nuclear spin ($I = \frac{7}{2}$). The nuclear spin couples to the angular momentum, J , of the valence electron, resulting in the emergence of hyperfine structure; each fine structure level is split into a number of hyperfine manifolds corresponding to different total angular momenta $F = I + J$, and each hyperfine manifold contains $2F + 1$ degenerate states, labelled by their angular momentum quantum numbers, m_F . This degeneracy is lifted by applying a constant external magnetic field of strength B_0 , which defines a quantisation axis for



(a) $^{43}\text{Ca}^+$: Wavelengths of lasers required for Doppler cooling, state preparation and read-out are shown in blue and red. The ‘clock’ qubit states $\{|\uparrow_C\rangle, |\downarrow_C\rangle\}$ and ‘stretch’ qubit states $\{|\uparrow_S\rangle, |\downarrow_S\rangle\}$, and their transition frequencies, are labelled. Transitions between hyperfine states (shown in green) may be driven using microwaves or Raman beams.



(b) $^{88}\text{Sr}^+$: Solid blue and red arrows show laser wavelengths required for Doppler cooling, state preparation and readout. The $5S_{1/2} \leftrightarrow 5P_{3/2}$ and $4D_{3/2} \leftrightarrow 5P_{3/2}$ transition wavelengths (not used in this work) are also indicated with dashed lines. The 409 MHz Zeeman qubit $\{|\uparrow_Z\rangle, |\downarrow_Z\rangle\}$ may be driven using an oscillating RF field or Raman beams. The optical qubit is formed from $|\downarrow_Z\rangle$ and $|\uparrow_O\rangle$ and is driven using a 674 nm laser.

Figure 2.1: Atomic level structures of $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ in a magnetic field of strength $B_0 = 146 \text{ G}$. Raman beams (wavelength 402 nm) are also used to drive clock and stretch qubit transitions in $^{43}\text{Ca}^+$ and Zeeman qubit transitions in $^{88}\text{Sr}^+$, via the P states; for clarity these beams are not shown, but are discussed further in section 4.7.2.




Figure 2.2: Hyperfine structure in the $4S_{1/2}$ level of $^{43}\text{Ca}^+$ at different external magnetic-field strengths, B_0 ($1 \text{ G} = 10^{-4} \text{ T}$). The $|F, m_F\rangle = |4, 0\rangle \leftrightarrow |3, +1\rangle$ clock qubit at $B_0 = 146 \text{ G}$ is indicated; here, the $4S_{1/2}$ level is still in the low-field regime, where the hyperfine structure dominates and we may label states with F and m_F . This is also the case in the $4P_{1/2}$ level. The $3D_{5/2}$ level is in the strong-field regime at 146 G , and $4P_{3/2}$ and $3D_{3/2}$ are in the intermediate regime (see [66, 1, 85]).

laser polarisations and provides spectral resolution for efficient state preparation and readout.

In the ‘weak-field’ regime, the Zeeman splittings are small compared to the hyperfine structure, and F and m_F remain good quantum numbers. In the ‘strong-field’ or ‘Paschen-Back’ regime, the interaction with the magnetic field dominates, and states are best labelled by the quantum numbers m_I and m_J . In the intermediate regime, the projection of the total angular momentum (equal to m_F at low field and $m_I + m_J$ at high field) remains a good quantum number. The hyperfine structure in the $4S_{1/2}$ level of $^{43}\text{Ca}^+$ is shown, for a range of field strengths, in figure 2.2. Similar plots for the other fine structure levels may be found in [66, 1, 85].

The presence of hyperfine structure is an advantage because it allows access to ‘atomic clock’ transitions, so named because their frequencies are first-order independent of variations in the external magnetic field, making them good candidates

for stable frequency standards. Since magnetic-field noise is usually the dominant source of decoherence in trapped ions, a clock transition is also a very attractive candidate for use as a qubit. Hyperfine splittings also tend to be larger ($\sim$ GHz) than practically-achievable Zeeman splittings ($\sim$ MHz), which offers better frequency selectivity for preparation and readout of hyperfine qubit states.

On the other hand, the more complicated structure can be a disadvantage. With many more spectrally-resolved states involved in the transitions used for cooling and fluorescence detection, a more careful choice of laser parameters is required for these processes (particularly at higher external magnetic-field strengths) [66], and more laser frequency components may be required to access all necessary states. There is more potential for population ‘leakage’ out of the qubit subspace, and the existence of more allowed decay routes from an excited state leads to reduced emission of photons in any particular mode, which is undesirable for ion-photon entanglement.

We make use of hyperfine qubit states in the $4S_{1/2}$ ground level of $^{43}\text{Ca}^+$, because transitions between them are electric-dipole forbidden and have low ($\sim$ GHz) frequencies, meaning that their lifetimes are effectively infinite (see section 3.2).

With transition frequencies in the GHz range, hyperfine qubits may be driven by microwave radiation. The field of RF/microwave electronics is well established and a large range of high-quality components are available off-the-shelf at low cost. Microwave electrodes can straightforwardly be integrated into ion-trap structures, and their use avoids the technical demands of aligning and stabilising laser beams. Driving the qubit with microwaves also eliminates errors due to photon scattering, associated with Raman beams (see section 3.2.4.2). On the other hand, the longer wavelength of microwave radiation means that it couples much less strongly to an ion’s motion (see section 3.1.4), so microwaves are often used in conjunction with an optical frequency for when this is required; the hyperfine qubit transitions in this work are also driven using pairs of Raman beams. Microwaves are also much less easy to

spatially focus than lasers, so individual addressing of ions with microwaves requires more inventive solutions [7, 106] (in this thesis, individual addressing is achieved trivially for two ions through the use of two different atomic elements).

2.2.2.2 Clock qubit

At an external magnetic-field strength of $B_0 = 146.0942$ G, the transition frequency between the states $\{|\uparrow_C\rangle, |\downarrow_C\rangle\} \equiv \{4S_{1/2}^{F=4, m_F=0}, 4S_{1/2}^{F=3, m_F=+1}\}$ (indicated in figure 2.2) becomes first-order independent of deviations of the field from B_0 . The second-order sensitivity is $d^2f/dB^2 = 2.42$ kHz/G². The transition frequency is 3.200 GHz.

Previous work has shown that an exceptionally long coherence time ($T_2^* \approx 50$ s [67]) may be achieved for this qubit, without the use of dynamical decoupling. In chapter 5 we present a thorough characterisation of its memory performance over different timescales and under offsets of B_0 from the field-independent point. All the work in this thesis was performed at 146 G in order to give access to the clock qubit².

A disadvantage is that ‘light-shift’ two-qubit entangling gates cannot easily be implemented directly on a clock qubit. Additionally, neither of the clock qubit states can be prepared directly using optical pumping, so extra microwave pulses are required to transfer population to them (section 2.2.6).

2.2.2.3 Stretch qubit

The $4S_{1/2}^{4,+4}$ state, on the other hand, may be prepared directly using optical pumping. Therefore, another qubit used in this work is the ‘stretch’ qubit: $\{|\uparrow_S\rangle, |\downarrow_S\rangle\} \equiv \{4S_{1/2}^{3,+3}, 4S_{1/2}^{4,+4}\}$. Light-shift gates may be performed on this qubit. At 146 G, its transition frequency is 2.874 GHz and its magnetic-field sensitivity is -2.36 MHz/G.

²There is a ‘zero-field’ clock transition in the ground level between the states $|F=4, m_F=0\rangle$ and $|F=3, m_F=0\rangle$ at $B_0=0$, but the necessity of a small quantisation field means that we cannot operate exactly at the field-independent point of this transition. Furthermore, operating close to $B_0=0$ results in small Zeeman splittings and therefore increased off-resonant excitation of nearby transitions. Previous work did make use of this qubit [12] and demonstrated a coherence time of 1.2(2)s [110].

The stretch qubit is much more sensitive to magnetic-field fluctuations than the clock qubit. We use this property as an advantage, in order to calibrate the field strength B_0 by measuring the stretch transition frequency (section 4.4.2).

2.2.3 Strontium qubits

$^{88}\text{Sr}^+$ has zero nuclear spin and therefore no hyperfine structure. The interaction energy with a weak magnetic field of strength B_0 , for a state with total angular momentum projection m_J , is

$$\Delta E = g_J \mu_B B_0 m_J, \quad (2.10)$$

where

$$g_J = \frac{3J(J+1) - L(L+1) + S(S+1)}{2J(J+1)} \quad (2.11)$$

is the Landé g-factor (see e.g. [20]). This gives rise to Zeeman splittings of levels of different m_J . In this context, the ‘weak-field’ regime applies when the Zeeman splittings are small compared to the fine structure splittings; for practically-achievable field strengths this is always the case. In this regime, transition frequencies always have a linear dependence on the magnetic-field strength, meaning that clock transitions are inaccessible in nuclear-spin-zero ions.

The simpler level structure of $^{88}\text{Sr}^+$ gives straightforward access to a strong cycling transition used for efficient cooling and fluorescence detection. There are only two decay paths from the $5\text{P}_{1/2}^{m_J=+1/2}$ state to the $5\text{S}_{1/2}$ level, meaning that it is much better suited for high-rate, high-fidelity entanglement with a photon for inter-trap communication [170, 169]. For this application, the coherence time of the communication qubit needs only to be longer than the gate operation required to transfer information from it to a memory qubit (such as $^{43}\text{Ca}^+$) for longer term storage, so the lack of clock qubits is not a limitation.

The large majority of the work involving $^{88}\text{Sr}^+$ in this thesis is performed on a Zeeman qubit in the ground level. An optical qubit is an alternative option.

2.2.3.1 Zeeman qubit

The two states in the ground level form the Zeeman qubit $\{|\uparrow_Z\rangle, |\downarrow_Z\rangle\} \equiv \{5S_{1/2}^{m_J=+1/2}, 5S_{1/2}^{m_J=-1/2}\}$. At $B_0 = 146$ G, the frequency of this transition is 409 MHz and its field sensitivity is 2.80 MHz/G (similar in magnitude to that of the calcium stretch qubit). The Zeeman qubit shares many of the desirable qualities of the calcium hyperfine qubits, including an essentially infinite lifetime against spontaneous decay and the option to drive it with an RF field. If stronger motional coupling is required, it may also be driven using the same Raman beams as for $^{43}\text{Ca}^+$.

2.2.3.2 Optical qubit

An alternative qubit is formed from the lower Zeeman qubit state $|\downarrow_Z\rangle = 5S_{1/2}^{m_J=-1/2}$ and the $4D_{5/2}^{m_J=-3/2}$ state, which we label $|\uparrow_O\rangle$. This is an electric-quadrupole-allowed transition which we drive with a 674 nm laser. Of the quadrupole transitions accessible with our geometry, this one is minimally sensitive to magnetic-field noise.

The major disadvantage of optical qubits is that their lifetimes are significantly shorter. The lifetime of the $4D_{5/2}$ level in $^{88}\text{Sr}^+$ is 391 ms. The coherence time of this qubit is limited by the linewidth of the quadrupole laser used to drive it, and because it is a dipole-forbidden transition the laser intensity needs to be appreciable (~ 10 mW); these are significant extra technical considerations. This qubit is rarely used in this work.

On the other hand, the optical qubit offers near-perfect state discrimination for readout. Direct drive with the 674 nm laser provides motional coupling and eliminates photon-scattering errors in two-qubit gates (these are replaced with errors due to spontaneous decay of the upper state — see section 3.3.5.1).

qubit	$\omega_0/2\pi$	drive	field sensitivity
Ca clock $\{ \uparrow_C\rangle, \downarrow_C\rangle\}$	3.200 GHz	microwaves/Raman	$\frac{d^2 f}{dB^2} = 2.42 \text{ kHz/G}^2$
Ca stretch $\{ \uparrow_S\rangle, \downarrow_S\rangle\}$	2.874 GHz	microwaves/Raman	$\frac{df}{dB} = -2.36 \text{ MHz/G}$
Sr Zeeman $\{ \uparrow_Z\rangle, \downarrow_Z\rangle\}$	409 MHz	RF/Raman	$\frac{df}{dB} = 2.80 \text{ MHz/G}$
Sr optical $\{ \uparrow_O\rangle, \downarrow_O\rangle\}$	$c/(674 \text{ nm})$	quadrupole laser	$\frac{df}{dB} = 1.12 \text{ MHz/G}$

Table 2.2: A comparison of the qubits used in this thesis. The frequencies and field sensitivities are given for an external magnetic-field strength of 146 G.

All four qubits and their transition frequencies are indicated in figure 2.1, and a summary of their properties is provided in table 2.2. The rest of this chapter will discuss how we load and cool our ions, and prepare and read out the various qubits.

2.2.4 Photoionisation

To generate ions, granules of neutral calcium and strontium are resistively heated so that a beam of hot atoms passes through the trapping zone. The samples of each atomic species contain a mixture of different isotopes (see section 4.2), so in order to select only the desired isotopes we use a two-stage photoionisation process, illustrated in figure 2.3.




Figure 2.3: Two-stage photoionisation processes for calcium and strontium.

First, an electron is promoted from the ground (1S_0) level to an excited (1P_1) level, using a 423 nm laser for calcium [111] and a 461 nm laser for strontium. A second beam, of wavelength 375 nm, is then used to excite the electron to the continuum. The linewidths of the transitions used in the first stage are smaller than the isotope shifts, allowing for isotope selectivity by tuning the frequencies of the 423 and 461 nm beams.

2.2.5 Cooling

We use three stages of cooling: Doppler cooling, sub-Doppler cooling on a dark resonance, and (if necessary) resolved sideband cooling. Due to its complex structure, previous work has given careful consideration to the parameters required for efficient cooling of $^{43}\text{Ca}^+$ at 146 G. Details can be found in [3, 66, 85].

2.2.5.1 Doppler cooling

Ions are cooled using a laser which is red-detuned from the frequency of a cycling transition, such that photons are preferentially absorbed when the ion is moving towards the beam (due to the Doppler effect). Photons are subsequently emitted in random directions, resulting in a net loss of energy from the ion. The cooling beam has a projection onto all three primary axes of the trap, to cool all motional modes.

The achievable temperature T_{Dopp} is limited by the recoil momentum of the emitted photons. For an ideal two-level atom, it is of the order [105]

$$k_B T_{\text{Dopp}} \approx \frac{\hbar \Gamma}{2}, \quad (2.12)$$

where Γ is the decay rate from the upper state. We use the $S_{1/2} \leftrightarrow P_{1/2}$ transition for Doppler cooling in both ion species, where the $P_{1/2}$ levels have natural linewidths of $\Gamma \approx 2\pi \times 22$ MHz, leading to $T_{\text{Dopp}} \approx 0.5$ mK. For a motional mode of typical secular

frequency $\omega_m \approx 2\pi \times 2$ MHz, this would result in a mean thermal occupation number $\bar{n}_m \approx k_B T_{\text{Dopp}}/(\hbar\omega_m) = 5.5$. In both ions, however (particularly $^{43}\text{Ca}^+$ because of its more complex hyperfine structure), the cooling transition is not a closed, two-level cycling transition, so the achievable temperature is higher than in the simple approximation. Doppler cooling is illustrated for both ions in figure 2.4.

2.2.5.2 Dark-resonance cooling

A consequence of the complex structure of $^{43}\text{Ca}^+$ is the existence of many two-photon dark resonances. A dark resonance occurs when two states are each coupled to a third, excited state via fields which have equal detunings from resonance with the relevant transitions. In this situation, there exist superpositions of the two lower states which no longer couple to the excited state.

One such dark resonance exists between the states $4S_{1/2}^{m_F=+4} \leftrightarrow 4P_{1/2}^{m_F=+4} \leftrightarrow 3D_{3/2}^{m=+5}$. By lowering the 397 nm and 866 nm laser intensities, we may choose frequencies to sit in this narrow feature, where the fluorescence rate is greatly reduced and the gradient of the fluorescence rate with frequency is large and positive. This is advantageous for cooling, because heating from photon recoil is reduced, and the ion is much more likely to emit a photon with larger energy than the one it absorbed.

We use this dark resonance to cool $^{43}\text{Ca}^+$ to below the Doppler limit. For 2 ms of dark-resonance cooling on a single ion, we reach ≈ 0.3 mK in apparatus A (parameters detailed in [3]), and 0.11 mK in apparatus B (parameters detailed in [150])³. For the motional modes of primary interest in each apparatus, this corresponds to $\bar{n}_m \lesssim 2$.

For a $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal (apparatus B only), we Doppler cool $^{88}\text{Sr}^+$ and dark-resonance cool $^{43}\text{Ca}^+$ for a total of 4 ms. For the first $\frac{2}{3}$ of the time, both ions are cooled simultaneously. For the final $\frac{1}{3}$, the $^{88}\text{Sr}^+$ lasers are switched off, to allow fur-

³The cooling performance is slightly better in apparatus B because the σ^+ and π components of the 397 nm laser are delivered via separate beams whose intensities can be varied independently. This allows the choice of more optimal intensities than in apparatus A, where a single beam of mixed σ^+/π -polarisation is used for cooling, and the intensity ratio is fixed by the geometry to $\frac{2}{3} \sigma^+ : \frac{1}{3} \pi$.




Figure 2.4: Doppler cooling and state preparation in $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$. In calcium, the $4S_{1/2}^{4,+4} \leftrightarrow 4P_{1/2}^{4,+4}$ transition is driven, using π -polarised 397 nm light. Population may decay from the upper state to the $3D_{3/2}$ level (with branching ratio 6.0%), so it must be reclaimed from here using an 866 nm repumper. Population may also decay to lower m_F states in both the $F = 4$ and $F = 3$ manifolds of the ground level. We therefore add a ≈ 3 GHz sideband to the 397 nm beam using an electro-optic modulator (EOM), in order to re-excite from both manifolds, and we add a σ^+ polarisation component to ensure that population stays at the higher m_F end of each manifold. To prepare $|\downarrow_S\rangle$, we switch off the EOM and optically pump using purely σ^+ -polarised 397 nm light. Similarly, in strontium we use a 422 nm laser to cool on the $5S_{1/2} \leftrightarrow 5P_{1/2}$ transition, and a 1092 nm beam to repump from the $4D_{3/2}$ level (branching ratio 5.5%). Acousto-optic modulators (AOMs) are used to split the $\sigma^\pm$ -polarised 422 nm beam into a lower frequency (422l) and a higher frequency (422u) component with a frequency difference of 545 MHz, to excite from the $5S_{1/2}^{m_J=+1/2}$ and $5S_{1/2}^{m_J=-1/2}$ states respectively. To prepare $|\downarrow_Z\rangle$, we switch off 422u and optically pump using 422l.

ther sympathetic cooling on the $^{43}\text{Ca}^+$ dark resonance. We reach a mean occupation number of $\bar{n}_{\text{ax,ip}} \approx 3$ for the axial ip mode (frequency 1.49 MHz).

2.2.5.3 Resolved sideband cooling

For many applications, Doppler cooling and dark-resonance cooling are sufficient. However, for high-fidelity two-qubit entangling gates, a mean occupation number of $\bar{n}_m \lesssim 0.1$ is required (see section 3.3.5.2). When necessary, we use the technique of resolved sideband cooling [105]. This is implemented on calcium only; strontium is sympathetically cooled.

The natural linewidths Γ of the Doppler-cooling transitions are large compared to the motional mode frequencies ω_m of the ions. If we instead choose a transition for which $\Gamma \ll \omega_m$, such as a dipole-forbidden transition or a Raman transition with sufficient detuning, we can resolve secular motion sidebands.

We tune the frequency difference of a Raman beam pair to $\omega_0 - \omega_m$, where ω_0 is the frequency of the stretch qubit transition and ω_m is the secular frequency of the mode we wish to cool. This allows us to drive red sideband (RSB) transitions ($|\downarrow_S, n_m\rangle \rightarrow |\uparrow_S, n_m - 1\rangle$), which remove a motional quantum of energy from mode m (see section 3.1.4). A σ^+ -polarised 397 nm beam (with the EOM switched off) is then used to repump population from $|\uparrow_S\rangle$ to $|\downarrow_S\rangle$ indirectly, via excitation to $4\text{P}_{1/2}^{4,+4}$ followed by spontaneous decay. An 866 nm beam is also used to reclaim population which decays to the $3\text{D}_{3/2}$ level.

‘Continuous’ sideband cooling involves turning on the Raman, 397 σ^+ and 866 beams simultaneously. ‘Pulsed’ sideband cooling involves alternating RSB pulses and repumping pulses. We optimise the cooling parameters empirically and choose 2 ms of continuous cooling followed by 20 repetitions of the pulsed sequence for each axial motional mode. The oop mode is cooled first because it has the lower heating rate. We reach $\bar{n}_{\text{ax ip}} = 0.07(1)$ and $\bar{n}_{\text{ax oop}} = 0.030(7)$ (see section 6.1).

2.2.6 State preparation

Calcium

In calcium, we can prepare $|\downarrow_S\rangle = 4S_{1/2}^{4,+4}$ via optical pumping. The process is similar to that used for Doppler cooling (figure 2.4), except we use a purely σ^+ -polarised 397 nm beam, and we switch off the EOM.

The state-preparation fidelity is limited by the polarisation purity of the $397\sigma^+$ beam, which leads to some unwanted population in states in the $4S_{1/2}$, $F = 4$ hyperfine manifold with $m_F < 4$. The fidelity can be improved by applying microwave π -pulses on the $|F, m_F\rangle = |4, +3\rangle \rightarrow |3, +3\rangle$ and $|4, +2\rangle \rightarrow |3, +2\rangle$ transitions [67, 66]. The ‘microwave-enhanced’ state preparation pulse sequence is $[397\sigma^+ \text{ (EOM on)}, \{397\sigma^+ \text{ (EOM off)}, \text{microwaves}\}^N, 397\sigma^+ \text{ (EOM off)}]$. We implement the technique in apparatus A with $N = 2$.

To prepare $|\uparrow_C\rangle$ or $|\downarrow_C\rangle$, three or four more microwave pulses are required: $|F, m_F\rangle = |4, +4\rangle \rightarrow |3, +3\rangle$, $|3, +3\rangle \rightarrow |4, +2\rangle$, $|4, +2\rangle \rightarrow |3, +1\rangle$, and $|3, +1\rangle \rightarrow |4, 0\rangle$.

Strontium

In strontium, we prepare $|\downarrow_Z\rangle$, the lower state of both the Zeeman and the optical qubits, via optical pumping using the 422 nm transition. We use the same beams as for Doppler cooling (figure 2.4) but we switch off the 422u component.

2.2.7 Readout

To read out one of our ions, we first perform a ‘shelving’ process (described in figure 2.5) which maps population in one of the qubit states (the ‘dark’ state) to the $D_{5/2}$ level. Then we turn on the Doppler-cooling beams for 1 ms and detect fluorescence on the 397 nm or 422 nm transition. The $D_{5/2}$ level is outside the Doppler-cooling cycle, so if the ion is in this state the number of photon counts will be small.




Figure 2.5: Shelving processes for readout of $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$. In $^{43}\text{Ca}^+$, a pulse sequence $[(393\sigma^+, 850\sigma_{1,2}^+, 850\pi)^N, 393\sigma^+]$, where $N = 8$, selectively shelves $|\downarrow_S\rangle$: The 393 nm pulses excite to $4P_{3/2}^{5,+5}$, from which population decays with a branching ratio of 5.9% to the $3D_{5/2}$ shelf. The only dipole-allowed decay path to the $4S_{1/2}$ level is back to $|\downarrow_S\rangle$, with a branching ratio of 93.5%. Population may also decay to states in the $3D_{3/2}$ level (branching ratio 0.7%), from which it is repumped using σ^+ -polarised and π -polarised 850 nm light. This readout method is discussed in detail in [171, 31]. To read out the clock qubit, it is first mapped to the stretch qubit using the microwave pulse sequence $[|3, +1\rangle \rightarrow |4, +2\rangle, |4, +2\rangle \rightarrow |3, +3\rangle, |3, +3\rangle \rightarrow |4, +4\rangle, |4, 0\rangle \rightarrow |3, +1\rangle, |3, +1\rangle \rightarrow |4, +2\rangle, |4, +2\rangle \rightarrow |3, +3\rangle]$. This maps $|\downarrow_C\rangle$ to $|\uparrow_S\rangle$, and $|\uparrow_C\rangle$ to $|\downarrow_S\rangle$. We then shelf $|\downarrow_S\rangle$ in the same manner. In $^{88}\text{Sr}^+$, $|\downarrow_Z\rangle$ is shelved more simply using a 674 nm pulse (if we are using the optical qubit, the upper state $|\uparrow_O\rangle$ is already in the $4D_{5/2}$ level so shelving is not required). After shelving, fluorescence is detected on the $S_{1/2} \leftrightarrow P_{1/2}$ transition. The $D_{5/2}$ shelf is outside the fluorescence cycle.

If the number of counts is above a certain threshold, we determine that the ion was in the other (‘bright’) qubit state. Accurate readout relies on the detection of a large number of photons from the bright state before a) the bright state is off-resonantly shelved or b) the dark state decays and joins the fluorescence cycle. Factors which limit readout fidelity therefore include the photon collection efficiency, imperfections in the shelving processes, and spontaneous decay of the $D_{5/2}$ shelf.

$^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ readout

For a $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal, the ions are shelved and read out sequentially. This allows us to determine the state of each ion individually (as opposed to a same-species crystal, where we are unable to distinguish between the states $|\uparrow\downarrow\rangle$ and $|\downarrow\uparrow\rangle$ using only the fluorescence level). $^{88}\text{Sr}^+$ is shelved and read out first, because the required 674 nm shelving pulse is much more sensitive to temperature than the 393 nm optical pumping process in calcium.

Chapter 3

Theory

3.1 Ion-light interactions

We rely on the interaction between an ion and an electromagnetic field in order to prepare, read out and manipulate the ion's electronic energy state, as well as to cool its motional state. This allows us to satisfy the second, fourth and fifth DiVincenzo criteria (initialisation, gates and measurement) in a trapped-ion quantum computing architecture.

We may express the Hamiltonian of an ion in a trap, interacting with a driving field, as

$$H = H_m + H_e + H_I \tag{3.1}$$

where the motional part of the Hamiltonian $H_m = \sum_m \hbar\omega_m(a_m^\dagger a_m + \frac{1}{2})$ represents the excitation in all harmonic oscillator modes m , the electronic part H_e encompasses all of the atomic structure (including the effect of the constant external magnetic field at 146 G), and H_I represents the interaction with the driving field.




Figure 3.1: The idealised two-level atom model. We represent the ‘ z -basis’ states $|\uparrow\rangle$ and $|\downarrow\rangle$ as unit vectors along the $\pm z$ -directions of the Bloch sphere. Any pure state $|\psi\rangle = \cos(\theta/2)|\uparrow\rangle + e^{i\phi}\sin(\theta/2)|\downarrow\rangle$ in this two-dimensional subspace may be represented using a Bloch vector with polar angle θ from the z -axis and azimuthal angle ϕ from the x -axis.

3.1.1 The two-level atom

Often we are interested in only two of the many electronic energy states. We may label the states $|\uparrow\rangle$ and $|\downarrow\rangle$, and the energy splitting between them $\hbar\omega_0$, as shown in figure 3.1a. We define $|\uparrow\rangle$ and $|\downarrow\rangle$ as the eigenstates of the Pauli z -operator σ_z with eigenvalues $+1$ and -1 respectively, so that in the subspace of these two states the electronic Hamiltonian may be written

$$H_e = \frac{\hbar\omega_0}{2}\sigma_z. \quad (3.2)$$

A two-level system such as this is commonly represented geometrically using the Bloch sphere, shown in figure 3.1b.

3.1.2 Multipole interactions

H_I is best expressed as a sum of terms which represent the interaction of the driving field with the electric and magnetic multipole moments of the ion’s charge distribu-

tion:

$$H_I = H_{ED} + H_{MD} + H_{EQ} + \dots \quad (3.3)$$

(see e.g. [109]). The leading terms in this expansion are the electric dipole (ED), magnetic dipole (MD) and electric quadrupole (EQ) contributions

$$\begin{aligned} H_{ED} &= -\mathbf{d} \cdot \mathbf{E}(\mathbf{r}_n) \\ H_{MD} &= -\boldsymbol{\mu} \cdot \mathbf{B}(\mathbf{r}_n) \\ H_{EQ} &= -\frac{1}{6} \sum_{i,j} Q_{ij} \frac{\partial E_j}{\partial x_i} \Big|_{\mathbf{r}_n}, \end{aligned} \quad (3.4)$$

where $\mathbf{d} = -e\mathbf{r}_e$ is the ion's electric dipole moment ($\mathbf{r}_e$ is the position of the electron relative to the position of the nucleus), $\boldsymbol{\mu}$ is the magnetic dipole moment (including contributions from the ion's orbital and spin angular momentum), and the rank-2 tensor Q_{ij} is the electric quadrupole moment. Dipole moments interact directly with the electric and magnetic fields $\mathbf{E}$ and $\mathbf{B}$, which should be evaluated at the position $\mathbf{r}_n$ of the nucleus. The electric quadrupole moment interacts with the gradient of the electric field, which should also be evaluated at $\mathbf{r}_n$.

3.1.2.1 Electric dipole interactions

We drive electric dipole transitions between atomic energy levels using an oscillating field such as a laser. We treat the field classically such that

$$H_{ED} = eE(\mathbf{r}_n)\mathbf{r}_e \cdot \hat{\mathbf{e}} \cos(\mathbf{k} \cdot \mathbf{r}_n - \omega t + \phi_l), \quad (3.5)$$

where the laser frequency is ω , the charge of the electron is $-e$, the electric field has amplitude E along polarisation direction $\hat{\mathbf{e}}$, and ϕ_l is a constant phase offset determined by the laser oscillator.

Typical electric dipole transition wavelengths (~ 400 nm) are long compared to

the size of the ion (~ 0.1 nm). Therefore we may make the ‘dipole approximation’ and assume that $\mathbf{k} \cdot \mathbf{r}_e \ll 1$, so that $\mathbf{E}$ depends only on $\mathbf{r}_n$ and not on $\mathbf{r}_e$. Equivalently, we assume that interaction energies due to gradients in $\mathbf{E}$ or $\mathbf{B}$ will be much smaller than interaction energies with $\mathbf{E}$ and $\mathbf{B}$ themselves. MD and EQ interaction energies have similar magnitudes and are smaller than ED interactions by a factor of α , the fine structure constant ($\approx \frac{1}{137}$). Therefore if the electric dipole interaction is non-zero we may neglect all further contributions to the interaction Hamiltonian.

In the basis $\{|\uparrow\rangle, |\downarrow\rangle\}$, H_{ED} may be written

$$H_{ED} = \frac{\hbar\Omega}{2} \left[e^{i(\mathbf{k} \cdot \mathbf{r}_n - \omega t + \phi_l)} + e^{-i(\mathbf{k} \cdot \mathbf{r}_n - \omega t + \phi_l)} \right] (|\uparrow\rangle\langle\downarrow| + |\downarrow\rangle\langle\uparrow|), \quad (3.6)$$

where $\Omega \equiv \langle\uparrow|e\mathbf{r}_e \cdot \hat{\mathbf{e}}|\downarrow\rangle E(\mathbf{r}_n)/\hbar$ is the Rabi frequency. We know that H_{ED} has no diagonal elements due to the odd parity of the position operator $\mathbf{r}_e$.

Moving to the interaction picture with respect to the electronic Hamiltonian H_e and using the rotating wave approximation (RWA) to neglect fast-oscillating terms at frequency $\omega + \omega_0$, the total Hamiltonian in the dipole approximation becomes

$$\tilde{H}_{ED} = -\frac{\hbar\Omega}{2} \left[e^{i\delta t} e^{-i(\phi_l + \mathbf{k} \cdot \mathbf{r}_n)} |\uparrow\rangle\langle\downarrow| + e^{-i\delta t} e^{i(\phi_l + \mathbf{k} \cdot \mathbf{r}_n)} |\downarrow\rangle\langle\uparrow| \right], \quad (3.7)$$

where $\delta \equiv \omega - \omega_0$.

As discussed in section 2.1, the position $\mathbf{r}_{n,i}$ of the i th ion’s nucleus will oscillate about its equilibrium position $\mathbf{r}_{\text{eq},i}$ in the trap, such that

$$\mathbf{r}_{n,i} = \mathbf{r}_{\text{eq},i} + \sum_{\text{modes } m} \hat{\mathbf{r}}_m l_{i,m} (a_m + a_m^\dagger) b_{i,m}. \quad (3.8)$$

We may absorb $\mathbf{k} \cdot \mathbf{r}_{\text{eq}}$ into a phase $\phi \equiv \phi_l + \mathbf{k} \cdot \mathbf{r}_{\text{eq}}$. In some cases we can neglect motional effects and assume that the ion remains fixed at its equilibrium position ($\mathbf{r}_n = \mathbf{r}_{\text{eq}}$); then we can ignore the second term in $\mathbf{r}_n$.

On resonance

Assuming we may neglect motional effects, if the laser is on resonance with the transition ($\omega = \omega_0$) the Hamiltonian becomes

$$\tilde{H}_{ED}(\delta = 0) = -\frac{\hbar\Omega}{2}\sigma_n, \quad (3.9)$$

where $\sigma_n \equiv \hat{\mathbf{n}} \cdot \boldsymbol{\sigma}$, $\boldsymbol{\sigma} \equiv (\sigma_x, \sigma_y, \sigma_z)$ and $\hat{\mathbf{n}} = (\cos \phi, \sin \phi, 0)$. The propagator on resonance is

$$\tilde{U}(\delta = 0) \equiv e^{-i\tilde{H}_{ED}(\delta=0)t/\hbar} = \begin{pmatrix} \cos\left(\frac{\Omega}{2}t\right) & i \sin\left(\frac{\Omega}{2}t\right) e^{-i\phi} \\ i \sin\left(\frac{\Omega}{2}t\right) e^{i\phi} & \cos\left(\frac{\Omega}{2}t\right) \end{pmatrix}. \quad (3.10)$$

If the ion is in $|\downarrow\rangle$ at $t = 0$, the population in $|\uparrow\rangle$ after a time t is

$$P_{\uparrow}(t) = \sin^2\left(\frac{\Omega}{2}t\right) = \frac{1}{2}(1 - \cos(\Omega t)), \quad (3.11)$$

hence Ω is the frequency of the ‘Rabi oscillations’ in the populations of $|\uparrow\rangle$ and $|\downarrow\rangle$. Geometrically, the Bloch vector rotates through $|\uparrow\rangle$ and $|\downarrow\rangle$ about the axis $\hat{\mathbf{n}}$ in the x - y plane, with the angle ϕ to the x -axis determined by the laser oscillator and the ion’s equilibrium position. If we set $\phi = 0$ we perform coherent x -rotations, and $\phi = \frac{\pi}{2}$ results in y -rotations. We may choose the laser intensity and pulse length to set Ωt for the desired number of oscillations. For example, if $\Omega t = \pi$ the populations are completely inverted; this is referred to as a ‘ π -pulse’.

Off resonance

If we drive the transition away from resonance, the general propagator (again neglecting motional effects) is

$$\tilde{U}(\delta) = \begin{pmatrix} \cos\left(\frac{\Omega_{\text{eff}}}{2}t\right) + i\frac{\delta}{\Omega_{\text{eff}}}\sin\left(\frac{\Omega_{\text{eff}}}{2}t\right) & i\frac{\Omega}{\Omega_{\text{eff}}}\sin\left(\frac{\Omega_{\text{eff}}}{2}t\right)e^{-i\phi} \\ i\frac{\Omega}{\Omega_{\text{eff}}}\sin\left(\frac{\Omega_{\text{eff}}}{2}t\right)e^{i\phi} & \cos\left(\frac{\Omega_{\text{eff}}}{2}t\right) - i\frac{\delta}{\Omega_{\text{eff}}}\sin\left(\frac{\Omega_{\text{eff}}}{2}t\right) \end{pmatrix}. \quad (3.12)$$

This results in rotations about a direction $\hat{\mathbf{n}}'$ which is not in the x - y plane, at a higher frequency $\Omega_{\text{eff}} \equiv \sqrt{\Omega^2 + \delta^2}$. In this case, we no longer reach a full inversion of populations in $|\uparrow\rangle$ and $|\downarrow\rangle$.

Far off resonance

In the limit of large detuning, where $\delta \rightarrow \Delta \gg \Omega$, the off-diagonal terms tend to zero and $\tilde{U}$ becomes

$$\tilde{U}(\Delta) = \begin{pmatrix} e^{i\Omega_{\text{eff}}t/2} & 0 \\ 0 & e^{-i\Omega_{\text{eff}}t/2} \end{pmatrix}. \quad (3.13)$$

Therefore, for large detuning, the probability of transferring population between $|\uparrow\rangle$ and $|\downarrow\rangle$ tends to zero and we no longer see Rabi oscillations. Working backwards from this propagator, the corresponding Hamiltonian is

$$\tilde{H}_{ED}(\Delta) = -\frac{\hbar\Omega_{\text{eff}}}{2}\sigma_z \approx -\left(\frac{\hbar\Delta}{2} + \frac{\hbar\Omega^2}{4\Delta}\right)(|\uparrow\rangle\langle\uparrow| - |\downarrow\rangle\langle\downarrow|), \quad (3.14)$$

where we have used the fact that $\Omega_{\text{eff}} \approx \Delta + \frac{\Omega^2}{2\Delta}$ for large Δ . The first term on the right is simply the electronic Hamiltonian in the rotating frame ($\tilde{H}_e$). The second represents a shift of each of the states $|\uparrow\rangle$ and $|\downarrow\rangle$ by an energy $\frac{\hbar\Omega^2}{4\Delta}$. This is known as a ‘light shift’ or an ‘AC Stark shift’. The directions of the shifts are such that the transition frequency increases (decreases) if the laser frequency is lower (higher) than ω_0 .

3.1.2.2 Magnetic dipole and electric quadrupole interactions

In many cases, the electric dipole moment operator $\langle \uparrow | e\mathbf{r}_e \cdot \hat{\mathbf{e}} | \downarrow \rangle$ between two states $|\uparrow\rangle$ and $|\downarrow\rangle$ vanishes. Such a transition is ‘electric-dipole forbidden’, and we must consider the next interaction terms in the multipole expansion.

Unlike the electric dipole moment operator, the magnetic dipole and electric quadrupole moment operators have positive parity, meaning that the diagonal terms in H_{MD} and in H_{EQ} are, in general, non-zero. This leads to a differential Stark shift (known instead as a ‘Zeeman shift’ in the case of a magnetic field) on the two states. Ignoring these diagonal terms (which may be absorbed into H_e), H_{MD} and H_{EQ} have forms equivalent to that of H_{ED} and give rise to Rabi oscillations with the corresponding Rabi frequencies for each type of interaction

$$\begin{aligned}\Omega_{ED} &= \langle \uparrow | e\mathbf{r}_e \cdot \hat{\mathbf{e}} | \downarrow \rangle E(\mathbf{r}_n)/\hbar, \\ \Omega_{MD} &= -\langle \uparrow | \boldsymbol{\mu} \cdot \hat{\mathbf{b}} | \downarrow \rangle B(\mathbf{r}_n)/\hbar, \\ \Omega_{EQ} &= \sum_{i,j} \langle \uparrow | Q_{ij} | \downarrow \rangle \left(\frac{\partial E_j}{\partial x_i} \right)_{\mathbf{r}_n} = \langle \uparrow | (e\mathbf{r}_e \cdot \hat{\mathbf{e}})(\mathbf{k} \cdot \mathbf{r}_e) | \downarrow \rangle E(\mathbf{r}_n)/2\hbar.\end{aligned}\tag{3.15}$$

Here $\left(\frac{\partial E_j}{\partial x_i} \right)_{\mathbf{r}_n}$ is the gradient of $\mathbf{E}$ at $\mathbf{r}_n$, and we have defined $\mathbf{B} \equiv B\hat{\mathbf{b}} \cos(\mathbf{k} \cdot \mathbf{r}_n - \omega t + \phi_0)$. Matrix elements for transitions relevant to this work are calculated in [150, 11, 66].

3.1.3 Raman interactions

Magnetic dipole transitions may also be driven by a pair of laser beams in a Raman configuration, each of which couples one of the states $\{|\uparrow\rangle, |\downarrow\rangle\}$ to a third, excited state $|e\rangle$ via an electric dipole transition. The beams have wavevectors and initial phases $\{\mathbf{k}_1, \phi_{l,1}\}$ and $\{\mathbf{k}_2, \phi_{l,2}\}$ respectively. $\{\delta_1, \Omega_1\}$ are the frequency detuning from and the Rabi frequency on the $|\downarrow\rangle \leftrightarrow |e\rangle$ transition of the first beam, and $\{\delta_2, \Omega_2\}$




Figure 3.2: Driving a Raman transition between states $|\uparrow\rangle$ and $|\downarrow\rangle$ via a third, excited state $|e\rangle$, using two laser beams. The first beam has frequency ω_1 which is detuned from the $|\downarrow\rangle \leftrightarrow |e\rangle$ transition by Δ so that $\delta_1 = \Delta$, and the second has frequency $\omega_2 = \omega_1 - (\omega_0 + \delta)$ so that its detuning from the $|\uparrow\rangle \leftrightarrow |e\rangle$ transition is $\delta_2 = \Delta + \delta$.

are the detuning from and Rabi frequency on the $|\uparrow\rangle \leftrightarrow |e\rangle$ transition of the second beam.

The Hamiltonian in the interaction picture with respect to H_e is given by the sum of the individual Hamiltonians of the two beams:

$$\begin{aligned} \tilde{H}_R = & -\frac{\hbar\Omega_1}{2} [e^{i\delta_1 t} e^{-i(\phi_{l,1} + \mathbf{k}_1 \cdot \mathbf{r}_n)} |e\rangle\langle\downarrow| + e^{-i\delta_1 t} e^{i(\phi_{l,1} + \mathbf{k}_1 \cdot \mathbf{r}_n)} |\downarrow\rangle\langle e|] \\ & -\frac{\hbar\Omega_2}{2} [e^{i\delta_2 t} e^{-i(\phi_{l,2} + \mathbf{k}_2 \cdot \mathbf{r}_n)} |e\rangle\langle\uparrow| + e^{-i\delta_2 t} e^{i(\phi_{l,2} + \mathbf{k}_2 \cdot \mathbf{r}_n)} |\uparrow\rangle\langle e|] . \end{aligned} \quad (3.16)$$

In the limit of large detunings $\delta_1 \gg \Omega_1$ and $\delta_2 \gg \Omega_2$, the probability of excitation of population to $|e\rangle$ becomes negligible (as we found in section 3.1.2.1 for a single beam).

In this limit we may apply the James–Jerke approximation [84] to time-average over fast dynamics evolving at frequencies $\sim \delta_1, \delta_2$. After doing so, we find an effective Hamiltonian

$$\begin{aligned} \tilde{H}_R^{\text{eff}} = & -\frac{\hbar\Omega_1^2}{4\delta_1} [|e\rangle\langle e| - |\downarrow\rangle\langle\downarrow|] - \frac{\hbar\Omega_2^2}{4\delta_2} [|e\rangle\langle e| - |\uparrow\rangle\langle\uparrow|] \\ & - \frac{\hbar\Omega_2\Omega_1}{4} \frac{1}{2} \left(\frac{1}{\delta_1} + \frac{1}{\delta_2} \right) [e^{i(\Delta\mathbf{k} \cdot \mathbf{r}_n + (\Delta\phi_l) - (\Delta\delta)t)} |\uparrow\rangle\langle\downarrow| + e^{-i(\Delta\mathbf{k} \cdot \mathbf{r}_n + (\Delta\phi_l) - (\Delta\delta)t)} |\downarrow\rangle\langle\uparrow|] , \end{aligned} \quad (3.17)$$

where $\Delta \mathbf{k} \equiv \mathbf{k}_2 - \mathbf{k}_1$, $\Delta \phi_l \equiv \phi_{l,2} - \phi_{l,1}$ and $\Delta \delta \equiv \delta_2 - \delta_1$. In the common configuration shown in figure 3.2 where $\delta_1 = \Delta$ and $\delta_2 = \Delta + \delta$, this becomes

$$\begin{aligned} \tilde{H}_R^{\text{eff}} = & -\frac{\hbar \Omega_1^2}{4\Delta} [|e\rangle\langle e| - |\downarrow\rangle\langle\downarrow|] - \frac{\hbar \Omega_2^2}{4(\Delta + \delta)} [|e\rangle\langle e| - |\uparrow\rangle\langle\uparrow|] \\ & - \frac{\hbar \Omega_R}{2} [e^{i(\Delta \mathbf{k} \cdot \mathbf{r}_n + \Delta \phi_l - \delta t)} |\uparrow\rangle\langle\downarrow| + e^{-i(\Delta \mathbf{k} \cdot \mathbf{r}_n + \Delta \phi_l - \delta t)} |\downarrow\rangle\langle\uparrow|], \end{aligned} \quad (3.18)$$

where $\Omega_R \equiv \Omega_1 \Omega_2 / 2\Delta$ and $1/\Delta_R \equiv (1/\delta_1 + 1/\delta_2)/2 \approx 1/\Delta$. The first two terms in $\tilde{H}_{\text{eff}}^R$ are single-beam light shifts arising from the couplings $|\uparrow\rangle \leftrightarrow |e\rangle$ and $|\downarrow\rangle \leftrightarrow |e\rangle$ as in equation (3.14). The third term represents a two-photon interaction through which we can drive Rabi oscillations between $|\uparrow\rangle$ and $|\downarrow\rangle$ at the Raman Rabi frequency Ω_R .

In a real system we typically couple $|\uparrow\rangle \leftrightarrow |\downarrow\rangle$ via many different excited states $|e_j\rangle$. As long as $\delta_{1,j} \gg \Omega_{1,j}$ and $\delta_{2,j} \gg \Omega_{2,j}$ holds for all j , we can coherently sum the Rabi frequencies

$$\Omega_R = \sum_j \frac{\Omega_{1,j} \Omega_{2,j}}{2\delta_j}. \quad (3.19)$$

Similarly, the total light shifts on $|\uparrow\rangle$ and $|\downarrow\rangle$ due to each beam are $\sum_j \frac{\hbar \Omega_{1,j}^2}{4\delta_j}$ and $\sum_j \frac{\hbar \Omega_{2,j}^2}{4\delta_j}$.

3.1.4 Coupling to the ion's motion

So far we have ignored motional effects and assumed that $\mathbf{r}_n = \mathbf{r}_{\text{eq}}$. This is a good approximation when the wavelength of the light is much longer than the spatial extent of the ion's motional wavefunction from its equilibrium position, such that $\mathbf{k} \cdot (\mathbf{r}_n - \mathbf{r}_{\text{eq}}) \ll 1$. For typical mode frequencies of a few MHz, the ground-state wavefunction spread $\tilde{l}_{i,m}$ for ion i in motional mode m (defined in section 2.1) is ~ 10 nm. Hence this is a good approximation for the microwave and RF wavelengths (~ 10 cm or longer) which are used in this work. However, laser wavelengths (~ 400 nm) are short enough that they do couple to the ion's motion.

This section will discuss motional coupling to a pair of Raman beams, as is most relevant for this work. The discussion is equally applicable to a single beam if $\{\Omega_R, \Delta\mathbf{k}, \Delta\phi_l\}$ are replaced with $\{\Omega, \mathbf{k}, \phi_l\}$.

3.1.4.1 The Lamb–Dicke parameter

Using equation (3.8) for the position $\mathbf{r}_{n,i}$ of the i th ion’s nucleus we may write

$$\Delta\mathbf{k} \cdot \mathbf{r}_{n,i} = \Delta\mathbf{k} \cdot \mathbf{r}_{\text{eq},i} + \sum_m \eta_{i,m} (a_m + a_m^\dagger), \quad (3.20)$$

where

$$\eta_{i,m} \equiv l_{i,m} b_{i,m} \Delta\mathbf{k} \cdot \hat{\mathbf{r}}_m \quad (3.21)$$

is the Lamb–Dicke parameter for the i th ion, for motional mode m , when coupling to a wavevector (difference) $\Delta\mathbf{k}$. The size of the Lamb–Dicke parameter is equal to the ratio of the spread of the motional ground-state wavefunction to the effective wavelength of the light $\lambda_{\text{eff}} = 2\pi/|\Delta\mathbf{k}|$. Alternatively, $\eta_{i,m}^2$ is the ratio of the energy $\hbar^2 k^2 / 2M_i$ transferred to mode m from photon recoil, to the energy spacing $\hbar\omega_m$ of the harmonic oscillator mode. Therefore $|\eta_{i,m}|$ is a measure of the coupling strength of the light to the motion of a particular ion in a particular mode. Note that the sign of $\eta_{i,m}$ depends on the mode eigenvector component $b_{i,m}$; in particular, for a two-ion crystal the sign of $\eta_{i,m}$ for each ion will be the same for an ip mode, but opposite for an oop mode (see section 2.1).

Excluding the single-beam light shift terms and defining $\Delta\phi \equiv \Delta\phi_l + \Delta\mathbf{k} \cdot \mathbf{r}_{\text{eq}}$, the effective Raman Hamiltonian in equation (3.18) may be written in terms of the Lamb–Dicke parameter for a given ion i :

$$\tilde{H}_R^{\text{eff}} = -\frac{\hbar\Omega_R}{2} \left[e^{i\left(\sum_m \eta_{i,m}(a_m + a_m^\dagger) + \Delta\phi - \delta t\right)} |\uparrow\rangle\langle\downarrow| + e^{-i\left(\sum_m \eta_{i,m}(a_m + a_m^\dagger) + \Delta\phi_l - \delta t\right)} |\downarrow\rangle\langle\uparrow| \right]. \quad (3.22)$$

Moving into the interaction picture with respect to the motional Hamiltonian $H_m = \sum_m \hbar\omega_m(a_m^\dagger a_m + \frac{1}{2})$, this becomes

$$\tilde{H}_R^{\text{eff}} = -\frac{\hbar\Omega_R}{2} \left[e^{i\left(\sum_m \eta_{i,m}(a_m e^{-i\omega_m t} + a_m^\dagger e^{i\omega_m t}) + \Delta\phi - \delta t\right)} |\uparrow\rangle\langle\downarrow| + e^{-i\left(\sum_m \eta_{i,m}(a_m e^{-i\omega_m t} + a_m^\dagger e^{i\omega_m t}) + \Delta\phi - \delta t\right)} |\downarrow\rangle\langle\uparrow| \right]. \quad (3.23)$$

3.1.4.2 The Lamb–Dicke limit

The Lamb–Dicke criterion is $\langle\psi_m|\mathbf{k}\cdot(\mathbf{r}_n - \mathbf{r}_{\text{eq}})|\psi_m\rangle \ll 1$ where $|\psi_m\rangle$ is the ion’s motional state. If this criterion is satisfied, the amplitude of the ion’s motion in the direction of the beam is small compared to the laser wavelength, and we are in the Lamb–Dicke regime. For a particular ion i in a particular mode m , this criterion corresponds to $\eta_{i,m}^2(2\bar{n}_m + 1) \ll 1$ where $\bar{n}_m$ is the mean occupation number of the mode. To satisfy this, it is a necessary but not a sufficient condition that $\eta_{i,m} \ll 1$; even if $\eta_{i,m}$ is small, for hot motional modes (high $\bar{n}_m$) we may still be outside the Lamb–Dicke regime unless

- i. The wavelength of the radiation is so large that $\eta_{i,m}$ is extremely tiny (as is the case for our microwave/RF drive), or
- ii. The direction of $\Delta\mathbf{k}$ is such that $\eta_{i,m}$ is extremely tiny, for example
 - (a) $\Delta\mathbf{k} = 0$ (co-propagating Raman beams), or
 - (b) $\Delta\mathbf{k} \perp \hat{\mathbf{r}}_m$ for mode m (as is the case for our Raman beams and the radial motional modes).

In the Lamb–Dicke limit we may approximate

$$e^{i\sum_m \eta_{i,m}(a_m e^{-i\omega_m t} + a_m^\dagger e^{i\omega_m t})} \approx 1 \pm i \sum_m \eta_{i,m} (a_m e^{-i\omega_m t} + a_m^\dagger e^{i\omega_m t}), \quad (3.24)$$

where we have neglected terms of higher order in $\eta_{i,m}$.

The Raman Hamiltonian becomes

$$\begin{aligned} \tilde{H}_R^{\text{eff}} = & -\frac{\hbar\Omega_R}{2} \left[e^{i(\Delta\phi-\delta t)} |\uparrow\rangle\langle\downarrow| + e^{-i(\Delta\phi-\delta t)} |\downarrow\rangle\langle\uparrow| \right] \\ & -\frac{\hbar\Omega_R}{2} \left[i \sum_m (\eta_{i,m} e^{i(\Delta\phi-(\delta+\omega_m)t)} a_m) |\uparrow\rangle\langle\downarrow| - i \sum_m (\eta_{i,m} e^{-i(\Delta\phi-(\delta+\omega_m)t)} a_m^\dagger) |\downarrow\rangle\langle\uparrow| \right] \\ & -\frac{\hbar\Omega_R}{2} \left[i \sum_m (\eta_{i,m} e^{i(\Delta\phi-(\delta-\omega_m)t)} a_m^\dagger) |\uparrow\rangle\langle\downarrow| - i \sum_m (\eta_{i,m} e^{-i(\Delta\phi-(\delta-\omega_m)t)} a_m) |\downarrow\rangle\langle\uparrow| \right]. \end{aligned} \quad (3.25)$$

We can identify three terms:

Carrier

On the first line in equation (3.25) is the carrier term $\tilde{H}_{\text{carrier}}$, which dominates if $\delta = 0$ i.e. the frequency difference of the two Raman beams is on resonance with the $|\uparrow\rangle \leftrightarrow |\downarrow\rangle$ transition. In this case, the other terms are fast-oscillating (at the mode frequencies ω_m) and are suppressed by a factor $\eta_{i,m}$. The carrier term drives transitions between $|\uparrow\rangle$ and $|\downarrow\rangle$ with no change in the motional state.

Red sideband (RSB)

If the frequency difference is tuned to the frequency of a particular motional mode m such that $\delta = -\omega_m$, the carrier term becomes

$$\tilde{H}_{\text{carrier}}^{\delta=-\omega_m} = -\frac{\hbar\Omega_R}{2} \left[e^{i(\Delta\phi+\omega_m t)} |\uparrow\rangle\langle\downarrow| + e^{-i(\Delta\phi+\omega_m t)} |\downarrow\rangle\langle\uparrow| \right], \quad (3.26)$$

and the other terms become

$$\tilde{H}_{\text{RSB}} = -\frac{i\hbar\Omega_R}{2} \eta_{i,m} \sum_{n_m} \sqrt{n_m} \left[e^{i\Delta\phi} |\uparrow, n_m - 1\rangle\langle\downarrow, n_m| - e^{-i\Delta\phi} |\downarrow, n_m\rangle\langle\uparrow, n_m - 1| \right]. \quad (3.27)$$

We have assumed that the motional modes are sufficiently well resolved that we only couple to mode m and we may use the RWA to neglect terms which oscillate at frequencies $\omega_{m'} - \omega_m$ (where $m' \neq m$), as well as at $\omega_{m'} + \omega_m$ and $2\omega_m$.

If the ion starts in the lower energy internal state, $|\downarrow\rangle$, $\tilde{H}_{\text{RSB}}$ drives transitions which both flip the internal state $|\downarrow\rangle \rightarrow |\uparrow\rangle$ of the ion and remove a quantum of motion from mode m . Hence continuous or repeated driving of the RSB can cool the motion. The RSB Rabi frequency for transitions from motional mode $|n_m\rangle$ to $|n_m - 1\rangle$ is

$$\Omega_R^{n_m \rightarrow n_m - 1} = \Omega_R \sqrt{n_m} \eta_{i,m}. \quad (3.28)$$

$\tilde{H}_{\text{carrier}}^{\delta = -\omega_m}$ represents an off-resonant drive of the carrier transition which may not be neglected completely because it is enhanced by a factor of $1/\eta_{i,m}$ compared to the other terms. However, as long as $\omega_m \gg \Omega_R$, it simply results in a light shift $-\frac{\hbar \Omega_R^2}{4(-\omega_m)} [|\uparrow\rangle\langle\uparrow| - |\downarrow\rangle\langle\downarrow|]$.

Blue sideband (BSB)

If $\delta = \omega_m$ we get, similarly,

$$\tilde{H}_{\text{carrier}}^{\delta = +\omega_m} = -\frac{\hbar \Omega_R}{2} \left[e^{i(\Delta\phi - \omega_m t)} |\uparrow\rangle\langle\downarrow| + e^{-i(\Delta\phi - \omega_m t)} |\downarrow\rangle\langle\uparrow| \right] \quad (3.29)$$

and the other terms become

$$\tilde{H}_{\text{BSB}} = -\frac{i\hbar \Omega_R}{2} \eta_{i,m} \sum_{n_m} \sqrt{n_m} \left[e^{i\Delta\phi} |\uparrow, n_m + 1\rangle\langle\downarrow, n_m| - e^{-i\Delta\phi} |\downarrow, n_m\rangle\langle\uparrow, n_m + 1| \right]. \quad (3.30)$$

If the ion starts in the lower energy internal state $|\downarrow\rangle$, $\tilde{H}_{\text{BSB}}$ drives transitions which flip the internal state $|\downarrow\rangle \rightarrow |\uparrow\rangle$ and add a quantum of motion to mode m . The BSB Rabi frequency for transitions from motional mode $|n_m\rangle$ to $|n_m + 1\rangle$ is

$$\Omega_R^{n_m \rightarrow n_m + 1} = \Omega_R \sqrt{n_m + 1} \eta_{i,m}. \quad (3.31)$$

3.1.4.3 Outside the Lamb–Dicke limit

If the Lamb–Dicke criterion is not satisfied we must consider terms of higher order in $\eta_{i,m}$ in equation (3.24). These terms allow for transitions between motional modes $|n_m\rangle$ and $|n'_m\rangle$ where $|n'_m - n_m| > 1$ (higher order sidebands). The general Rabi frequency for transitions from $|n_m\rangle$ to $|n'_m\rangle$ is [188]

$$\Omega_R^{n_m \rightarrow n'_m} = \Omega_R \langle n_m | e^{i\eta_{i,m}(a_m + a_m^\dagger)} | n'_m \rangle = \Omega_R e^{-\eta_{i,m}^2/2} \sqrt{\frac{n_{<}!}{n_{>}!}} \eta_{i,m}^{|n'_m - n_m|} \mathcal{L}_{n_{<}}^{|n'_m - n_m|}(\eta_{i,m}^2), \quad (3.32)$$

where $\mathcal{L}_{n_{<}}^{|n'_m - n_m|}(\eta_{i,m}^2)$ are the generalised Laguerre polynomials. Equation (3.28) and equation (3.31) are special cases of equation (3.32) in the Lamb–Dicke limit.

For carrier transitions ($n_m = n'_m$), the Rabi frequency is

$$\Omega_R^{n \rightarrow n} = \Omega_R e^{-\eta_{i,m}^2/2} (1 - n\eta_{i,m}^2 + \mathcal{O}(\eta_{i,m}^4)). \quad (3.33)$$

The reduction of the carrier Rabi frequency by a factor $e^{-\eta_{i,m}^2/2}$ is known as the Debye–Waller effect. For higher n_m , the carrier Rabi frequency is reduced further, before increasing again. Since different motional transitions are driven with different Rabi frequencies, this can cause decoherence of the ion’s internal state.

3.2 Qubit decoherence

Decoherence may be defined as the loss of a fixed phase relation between qubit states $|\uparrow\rangle$ and $|\downarrow\rangle$. The third of DiVincenzo’s criteria specifies that the time before decoherence occurs should be long compared to the typical duration of a gate operation.

Decoherence arises because trapped ions are not a completely isolated system; there are unwanted couplings to the environment. This means the eigenstates of the idealised Hamiltonian $H_m + H_e$ are not true eigenstates of the full system, including

the environment. This section will discuss decoherence of the ion's electronic state.¹

3.2.1 Types of decoherence

There are two timescales commonly associated with decoherence. The first is the T_1 time, which is associated with unwanted changes in the populations of $|\uparrow\rangle$ and $|\downarrow\rangle$ via exchange of energy with the environment. This is known as ‘longitudinal relaxation’. The second is the T_2^* time, associated with uncontrolled changes in the phase relation between $|\uparrow\rangle$ and $|\downarrow\rangle$. This is known as ‘dephasing’ or ‘transverse relaxation’.

3.2.1.1 Population decay

Section 3.1 showed that an applied electromagnetic field can drive transitions between the eigenstates of the electronic Hamiltonian H_e . After switching off this field, if H_e were the full electronic Hamiltonian then an ion in an eigenstate of H_e would stay in that state indefinitely. However, when the field is treated quantum-mechanically, we note that its ground state has a non-zero energy associated with it [124]. Even in the absence of an intentionally applied field, the ion may interact with fluctuations in the vacuum field, meaning that excited states have a finite lifetime before they spontaneously decay to states which are lower in energy. Spontaneous decay tends to be the dominant longitudinal relaxation process for optical qubits in ions.

As noted in chapter 2, we choose to encode our qubits in the hyperfine or Zeeman states of the ground levels because transitions between them are electric-dipole forbidden and their transition frequencies are so small that their lifetimes against spontaneous decay are effectively infinite ($\sim 10^9$ years)².

¹Environmental couplings which alter the motional Hamiltonian (such as heating) analogously cause decoherence of the motional state.

²In the two-level atom approximation, the decay rate on an electric dipole transition is equal to the Einstein A co-efficient [58] $A_{12} = \frac{\omega_0^3}{3\pi\epsilon_0\hbar c^3}|d_{12}|$, where ω_0 is the transition frequency, ϵ_0 is the permittivity of free space and d_{12} is the dipole moment between states 1 and 2 and is of the order $\sim ea_0$ where a_0 is the Bohr radius. For levels which are split into multiple states, we must sum over all decay routes. Lifetimes for electric-dipole-forbidden transitions are orders of magnitude longer.




Figure 3.3: Dephasing on the Bloch sphere in a frame rotating at the qubit frequency. (i) A Bloch vector points along $+y$. (ii) After a time τ , the Bloch vector acquires a spread in phase $\langle\delta\phi\rangle$. (iii) The resulting average Bloch vector is reduced in length.

The T_1 lifetime in ion-trap systems using these types of qubits is usually limited by other processes, such as collisions with background gas molecules in the vacuum system. In our case, this happens every few hours.

3.2.1.2 Dephasing

Under the Hamiltonian H_e , the Bloch vector of a qubit state on the Bloch sphere rotates about the z -axis at the qubit frequency ω_0 , leading to a linear accumulation of phase $\phi(t) = \phi(t=0) + \omega_0 t$. In a frame rotating at the qubit frequency, the Bloch vector remains stationary and the phase constant.

Unwanted coupling of the ion and the environment — for example, noise in the external magnetic field $\mathbf{B}_0$ — can change the energies of the qubit states and thereby the qubit frequency. The phase accumulated is then given by $\phi(t) = \phi(t=0) + \int_0^t (\omega_0 + \dot{\phi}(t')) dt'$, so this may be thought of as qubit ‘frequency noise’ or ‘phase noise’. It leads to unwanted rotation of the Bloch vector in the rotating frame.

Over many repeats of an experiment, or over an ensemble of many qubits, the Bloch vector rotates by a different amount in each case, due to noise processes which affect the frequency. The collection of Bloch vectors over these experiments or qubits will have some typical spread in phase $\langle\delta\phi\rangle$, as shown in figure 3.3. The ‘polar’ states $|\uparrow\rangle$ and $|\downarrow\rangle$ are insensitive to dephasing.

In order to measure the ion in experiment, we rely on coherence between it and

the field used for probing. Noise in the frequency or phase of the probe field may equally be viewed as noise on the qubit states relative to the probe field, and therefore contributes to the measured dephasing in the same manner.

Longitudinal relaxation processes cause random changes in the populations of the qubit states and therefore also destroy the phase relation between them. They contribute to transverse relaxation at a rate $1/(2T_1)$ [109].

T_2^* is associated with the total measured dephasing rate both from ‘pure dephasing’ processes and from longitudinal relaxation processes. We may write

$$1/T_2^* = 1/(2T_1) + 1/T_2^\perp, \quad (3.34)$$

where $T_2^\perp$ is the pure dephasing contribution.

3.2.2 Qubit phase noise

3.2.2.1 Measuring decoherence

T_1

A simple experimental sequence for measuring the T_1 time of a system is to prepare the higher energy state $|\uparrow\rangle$, wait for a time τ and then read out the state. Although the result will be $|\uparrow\rangle$ or $|\downarrow\rangle$ for each run of the sequence, this may be repeated many times in order to measure the probability $P_\uparrow(\tau)$ of remaining in the excited state after a particular time τ . Measuring $P_\uparrow(\tau)$ for different delay times allows extraction of T_1 by fitting a decay curve to the results. For spontaneous decay, an appropriate model is $P_\uparrow(\tau) = e^{-\tau/T_1}$. Generally, the most appropriate model depends on the underlying noise processes which contribute to changes in the populations.

T_2^*

A similar method, known as a Ramsey experiment, is often used to measure T_2^* . After preparing into $|\downarrow\rangle$, a $\frac{\pi}{2}$ -rotation is performed to put the qubit into an equal superposition of $|\uparrow\rangle$ and $|\downarrow\rangle$ (on the equator of the Bloch sphere), where it is maximally sensitive to dephasing. After a delay time τ , a second $\frac{\pi}{2}$ -pulse is applied with the same phase as the first, such that in the absence of any dephasing the final state would be $|\uparrow\rangle$. The state is read out, and this is again repeated many times for different delay times τ .

The probability of ending in the correct qubit state may be calculated by projecting the actual final state $|\psi\rangle$ onto the ideal final state $|\uparrow\rangle$. If the extra phase accumulated during the delay is equal to $\delta\phi(\tau)$, we find the probability

$$|\langle\uparrow|\psi\rangle|^2 = \cos^2\left(\frac{\delta\phi(\tau)}{2}\right). \quad (3.35)$$

The average over many repeats of the experiment is therefore given by

$$P_{\uparrow}(\tau) \equiv \langle |\langle\uparrow|\psi\rangle|^2 \rangle \approx 1 - \frac{\langle [\delta\phi(\tau)]^2 \rangle}{4}, \quad (3.36)$$

where $\langle [\delta\phi(\tau)]^2 \rangle$ is the mean-square dephasing over time τ , and is assumed to be small.

The phase of the second $\frac{\pi}{2}$ -pulse is usually scanned, in order to see sinusoidal oscillations in the population $P_{\uparrow}(\tau)$, dependent on the relative phase $\tilde{\phi}$ of the two pulses:

$$P_{\uparrow}(\tau) = \frac{C(\tau)}{2} \cos(\tilde{\phi}) + \frac{1}{2}, \quad (3.37)$$

where $C(\tau)$ is the ‘Ramsey fringe contrast’. $C(\tau)$ is equal to 1 in the absence of dephasing and $\approx 1 - \frac{1}{2}\langle [\delta\phi(\tau)]^2 \rangle$ otherwise.

3.2.2.2 Memory error and T_2^*

The loss of Ramsey fringe contrast $1 - C(\tau)$ is not a representative estimate of the average probability of incurring an error due to dephasing over the course of a computation. This is for two reasons:

- i. A contrast loss of 1 (complete decoherence) indicates the qubit is equally likely to be found in $|\uparrow\rangle$ or $|\downarrow\rangle$. For a computation with a given target state, the probability of measuring the incorrect state would then be $\frac{1}{2}$.
- ii. Over the course of a computation, the qubit will spend on average $\frac{1}{3}$ of its time near the poles of the Bloch sphere (which are insensitive to phase drifts), and only $\frac{2}{3}$ of its time near maximally sensitive states on the equator³.

We therefore define a more computationally relevant ‘memory error’ $\epsilon_m(\tau)$ which is related to the contrast loss via

$$\epsilon_m(\tau) = \frac{1}{2} \times \frac{2}{3} \times (1 - C(\tau)) \approx \frac{\langle [\delta\phi(\tau)]^2 \rangle}{6}, \quad (3.38)$$

where the approximation is valid for small mean-square phase deviations $\langle [\delta\phi(\tau)]^2 \rangle$.

An estimate for T_2^* is generally extracted by plotting $C(\tau)$ for different τ and fitting a decay curve. As for T_1 processes, the exact form of the curve depends on the dynamics of the underlying noise processes which contribute to the dephasing.

3.2.2.3 Dephasing due to an arbitrary noise source

For a particular run of the experiment the extra unwanted phase acquired after a time τ is

$$\delta\phi(\tau) = \int_0^\tau \dot{\phi}(t') dt', \quad (3.39)$$

³We assume that an arbitrary computation samples operations uniformly from a unitary two-design (see e.g. section 3.4.1). The result is perhaps apparent if one considers the single-qubit Clifford group (which may be composed from a basis gate set of discrete $\frac{\pi}{2}$ -rotations about the principal axes x , y and z of the Bloch sphere), but may also be extended to a continuous gate set.

where $\dot{\phi} = d\phi/dt$ is the instantaneous qubit frequency difference from ω_0 . Following [117], the mean-square dephasing after a time τ is therefore

$$\begin{aligned}\langle [\delta\phi(\tau)]^2 \rangle &= \left\langle \int_0^\tau \dot{\phi}(t') dt' \int_0^\tau \dot{\phi}(t'') dt'' \right\rangle \\ &= \int_0^\tau dt' \int_0^\tau dt'' \langle \dot{\phi}(t') \dot{\phi}(t'') \rangle,\end{aligned}\tag{3.40}$$

where the angle brackets indicate the average over multiple experimental runs.

The autocorrelation function of $\dot{\phi}$ is a measure of the similarity of $\dot{\phi}$ to a delayed version of itself. It is defined by

$$\begin{aligned}R_{\dot{\phi}}(\tau) &= \lim_{T \rightarrow +\infty} \frac{1}{T} \int_0^T \dot{\phi}(t) \dot{\phi}(t + \tau) dt \\ &= \langle \dot{\phi}(t) \dot{\phi}(t + \tau) \rangle.\end{aligned}\tag{3.41}$$

The Fourier transform of the autocorrelation function $R_{\dot{\phi}}(\tau)$ is equivalent to the power spectral density (PSD) $S(\omega)$ (see e.g. [34]) such that

$$\begin{aligned}R_{\dot{\phi}}(\tau) &= \frac{1}{2\pi} \int_{-\infty}^{\infty} S_{\dot{\phi}}^{(2)}(\omega) e^{-i\omega\tau} d\omega \\ &= \frac{1}{2\pi} \int_0^{\infty} S_{\dot{\phi}}^{(1)}(\omega) \cos(\omega\tau) d\omega,\end{aligned}\tag{3.42}$$

where (1) and (2) indicate the ‘one-sided’ and ‘two-sided’ PSDs respectively. The PSD $S_X(\omega)$ of a noise signal $X(t)$ is commonly used to describe the noise; it represents the noise ‘power’ ($\propto X(t)X^*(t)$) at a particular frequency ω , and has units [units of X]²/[angular frequency]. $S_{\dot{\phi}}(\omega)$ is the noise power, at angular frequency ω , in the qubit angular frequency $\dot{\phi}$, and therefore has units [angular frequency]²/[angular frequency].

Combining equations (3.40), (3.41) and (3.42), we now see that

$$\begin{aligned}\langle [\delta\phi(\tau)]^2 \rangle &= \frac{1}{2\pi} \int_0^\infty d\omega S_{\dot{\phi}}^{(1)}(\omega) \int_0^\tau dt' \int_0^\tau dt'' \cos(\omega(t' - t'')) \\ &= \frac{\tau^2}{2\pi} \int_0^\infty S_{\dot{\phi}}^{(1)}(\omega) \frac{\sin^2(\omega\tau/2)}{(\omega\tau/2)^2} d\omega,\end{aligned}\tag{3.43}$$

in agreement with [137], for an arbitrary noise process.

We may also relate the frequency noise PSD $S_{\dot{\phi}}(\omega)$ to the phase noise PSD $S_\phi(\omega)$ via [81]

$$S_{\dot{\phi}}(\omega) = \omega^2 S_\phi(\omega),\tag{3.44}$$

such that

$$\langle [\delta\phi(\tau)]^2 \rangle = \frac{2}{\pi} \int_0^\infty S_\phi^{(1)}(\omega) \sin^2(\omega\tau/2) d\omega.\tag{3.45}$$

$S_\phi(\omega)$ has units [radians]²/[angular frequency].

3.2.2.4 Common types of noise

This section describes two common types of noise and their effects on the memory error⁴ $\epsilon_m(\tau)$.

White frequency noise

White frequency noise (equivalent to $1/f^2$ phase noise) has equal power at all frequencies, and may therefore be represented by the PSD $S_{\dot{\phi}}(\omega) = S_W$ where S_W is a constant. Using equation (3.43) we calculate the mean-square dephasing for white frequency noise $\langle [\delta\phi(\tau)]^2 \rangle = S_W\tau/2$, leading to a memory error

$$\epsilon_m^W(\tau) = \frac{S_W\tau}{12},\tag{3.46}$$

in the limit of small errors.

⁴For an analysis of the effect of various types of noise on other ion-trap operations, see [66].

For purely white noise, it is therefore usually considered a reasonable approximation to model the Ramsey fringe contrast $C(\tau) = 1 - 3\epsilon_m(\tau)$ as an exponential decay $\approx e^{-\tau/T_2^*}$, in order to characterise the dephasing using the parameter $T_2^* \approx 4/S_W$.

1/f frequency noise

1/f frequency noise — also known as ‘pink’ noise or ‘flicker’ noise, and equivalent to 1/f³ phase noise — has a PSD of the form

$$S_{\dot{\phi}}(\omega) = S_{1/f} \frac{2\pi}{\omega}, \quad (3.47)$$

leading to a memory error (calculated numerically in [137])

$$\epsilon_m^{1/f}(\tau) = \frac{S_{1/f}\tau^2}{6} \ln\left(\frac{0.4007}{f_c\tau}\right), \quad (3.48)$$


where $f_c = \omega_c/2\pi$ is the low-frequency cut-off of the 1/f noise. The cut-off is approximately equal to the inverse of the experiment duration, such that we are insensitive to noise at frequencies lower than this⁵.

Since $\epsilon_m^{1/f}(\tau)$ depends quadratically on τ , a Gaussian model for the decay of Ramsey fringe contrast ($C(\tau) \approx e^{-(\tau/T_2^*)^2}$) is sometimes used to extract T_2^* .


3.2.3 Dynamical decoupling

The effect of dephasing may be mitigated to some extent using ‘dynamical decoupling’ (DD) pulse sequences, which act to undo the spread in phase and ‘refocus’ a Bloch vector. The simplest example of this refocusing is the spin echo, first explained in 1950 by Erwin Hahn in the context of nuclear magnetic resonance (NMR) [65]. Modifications of the spin-echo technique to overcome dephasing have been studied extensively in the field of NMR [180], and include the Carr, Purcell, Meiboom and

⁵Equation (3.48) assumes $f_c\tau \lesssim 0.2$, which is true for relevant experiments in this work.



(a) (i) A Bloch vector initially lies along the $+y$ -axis. (ii) After a time $\frac{\tau_{DD}}{2}$, it has undergone some dephasing. (iii) The dephased Bloch vector is rotated by an angle π about the x -axis. (iv) After another time $\frac{\tau_{DD}}{2}$ the dephasing is perfectly reversed.



(b) The πx pulse may be thought of as flipping the qubit frequency offset $\dot{\phi} \rightarrow -\dot{\phi}$. The phase acquired during a time t is given by integrating $\dot{\phi}$ over that time. For a constant frequency offset, we see that the net area (ii) + (iii) after a time τ_{DD} is zero.

Figure 3.4: Perfect correction of dephasing due to a constant frequency offset, using the Carr-Purcell sequence. The Meiboom-Gill adaptation employs a π -rotation about the y -axis rather than the x -axis, the effect of which is to refocus the spins along their original orientation, $+y$, rather than $-y$.

Gill (CPMG) sequence [32, 122]. CPMG dynamical decoupling consists of repeated π -pulses about an axis on the equator of the Bloch sphere, spaced evenly by a time τ_{DD} .

The success of CPMG relies on the behaviour of the noise $\dot{\phi}(t)$ being similar in successive periods of time τ_{DD} between the DD π -pulses. Perfect cancellation is achieved for any ‘perfectly correlated’ noise; figure 3.4 illustrates the effect for a constant qubit frequency offset $\dot{\phi}$. Errors due to any dephasing process with some degree of correlation over the relevant timescale τ_{DD} may be partially corrected.

The higher the degree of correlation in subsequent delay periods, the better the performance of the sequence. Therefore, under the assumption that the DD pulses

themselves are executed perfectly and instantaneously, decreasing τ_{DD} and increasing the number of pulses will correct the dephasing increasingly well. In reality, imperfections in the DD pulses eventually limit the performance [4, 147, 184], although the sequences can be constructed to increase robustness to these imperfections.

Increasingly complex DD pulse sequences, which work in a similar manner to the CPMG sequence, have since been proposed and implemented. Notable examples include ‘Uhrig’ dynamical decoupling (UDD) [178], in which the delay times τ_{DD} between pulses are varied, and ‘Knill’ dynamical decoupling (KDD) [165], where π -pulses are performed about different axes in the x - y plane. Both of these methods have been successfully used to extend coherence times in ion-trap systems [18, 172, 183, 182]. Errors arising from decoherence during two-qubit gates between ions have also been suppressed using DD techniques [173, 13, 12, 68, 116, 166].

3.2.4 Ion-light interactions in the presence of decoherence

Section 3.1 described the interaction of an ion with a driving field in the limit where there is no decoherence of the internal states. We will now examine the effect of decoherence.

3.2.4.1 Damped Rabi oscillations

We may describe a state $|\psi\rangle = c_{\downarrow}|\downarrow\rangle + c_{\uparrow}|\uparrow\rangle$ using the density matrix

$$\rho = \begin{pmatrix} \rho_{\downarrow\downarrow} & \rho_{\downarrow\uparrow} \\ \rho_{\uparrow\downarrow} & \rho_{\uparrow\uparrow} \end{pmatrix}, \quad (3.49)$$

where $\rho_{\alpha\beta} = c_{\alpha}c_{\beta}^*$. We additionally define $\tilde{\rho}_{\uparrow\downarrow} = e^{-i\delta t}\rho_{\uparrow\downarrow}$ and $\tilde{\rho}_{\downarrow\uparrow} = e^{-i\delta t}\rho_{\downarrow\uparrow}$. where δ is the detuning of the driving field from the transition frequency.

We define Γ as the rate of population decay from $|\uparrow\rangle$ to $|\downarrow\rangle$ and γ as the total dephasing rate. In general, $\gamma = \Gamma/2 + \gamma_{\perp}$, by analogy with equation (3.34). The




Figure 3.5: Solutions to the optical Bloch equations for resonant drive ($\delta = 0$) and initial conditions $|\psi(t=0)\rangle = |\downarrow\rangle$, for different values of the longitudinal and transverse relaxation rates γ and Γ .

relevant equations for the time-evolution of ρ are the ‘optical Bloch equations’ [109]:

$$\begin{aligned}
\dot{\rho}_{\uparrow\uparrow} &= i\frac{\Omega}{2}(\tilde{\rho}_{\uparrow\downarrow} - \tilde{\rho}_{\downarrow\uparrow}) - \Gamma\rho_{\uparrow\uparrow} \\
\dot{\rho}_{\downarrow\downarrow} &= -i\frac{\Omega}{2}(\tilde{\rho}_{\uparrow\downarrow} - \tilde{\rho}_{\downarrow\uparrow}) + \Gamma\rho_{\uparrow\uparrow} \\
\dot{\tilde{\rho}}_{\downarrow\uparrow} &= -i\frac{\Omega}{2}(\rho_{\uparrow\uparrow} - \rho_{\downarrow\downarrow}) - (\gamma + i\delta)\tilde{\rho}_{\downarrow\uparrow} \\
\dot{\tilde{\rho}}_{\uparrow\downarrow} &= i\frac{\Omega}{2}(\rho_{\uparrow\uparrow} - \rho_{\downarrow\downarrow}) - (\gamma - i\delta)\tilde{\rho}_{\uparrow\downarrow}.
\end{aligned} \tag{3.50}$$

To illustrate the effect of decoherence we consider the particular case where the drive is on resonance, and the initial state of the system is $|\psi(t=0)\rangle = |\downarrow\rangle$. In this case, we find an analytic solution for the population in the upper state:

$$P_{\uparrow}(t) = \rho_{\uparrow\uparrow}(t) = \frac{\Omega^2/2}{\Gamma + \Omega^2} \left[1 - \left(\cos(\Omega_{\Gamma}t) + \frac{(\Gamma + \gamma)}{2\Omega_{\Gamma}} \sin(\Omega_{\Gamma}t) \right) e^{-\frac{(\Gamma + \gamma)t}{2}} \right], \tag{3.51}$$

where $\Omega_{\Gamma} \equiv \sqrt{\Omega^2 - (\Gamma - \gamma)^2/4}$, in agreement with [134].

Figure 3.5 shows $\rho_{\uparrow\uparrow}(t)$ for different values of Γ and γ . In the case where $\Gamma = \gamma = 0$ this reduces to $\sin^2(\Omega t/2)$, in agreement with equation (3.11). For larger decoherence rates, the Rabi oscillations are damped and the populations reach an equilibrium value of 0.5 over a timescale $\sim 2/(\Gamma + \gamma)$. In the limit of large damping, the Rabi oscillations are cut off before they reach a peak and we recover the behaviour predicted by the Einstein rate equations (see, for example, [56]), which do not include coherent evolution.

Although in theory a prediction of the behaviour of a multi-level atom in the presence of a driving field and with decoherence would require solution of optical Bloch equations for all of the levels involved, in practice we are usually the ‘low-damping’ or ‘high-damping’ limit, where Rabi oscillations or the Einstein rate equations are good approximations. The former applies during coherent manipulation of qubit states, and the latter during operations such as fluorescence detection or Doppler cooling where the short lifetimes of the P-levels mean that we are very much within the incoherent limit.

3.2.4.2 Scattering

Longitudinal relaxation comes into play slightly differently for Raman interactions, where we must consider decay from the intermediate excited state.

‘Scattering’ involves the absorption of a photon from a Raman beam such that an excited state $|e\rangle$ is occupied, followed by the spontaneous decay of the state $|e\rangle$ and the consequent emission of a photon in a random direction. This two-photon process may be thought of as a Raman transition with the second beam replaced by the vacuum field.

The scattering rate from initial state $|i\rangle$ to final state $|f\rangle$ via multiple excited states $|e_j\rangle$, from a single laser beam with polarisation $\hat{\mathbf{e}}_1$ and electric field amplitude

E at the ion, is given by the Kramers–Heisenberg formula [139]:

$$\Gamma_{i \rightarrow f} = \frac{E^2 A_j}{4\hbar^2} \left| \sum_{j,p} \frac{\langle i | e \mathbf{r}_e \cdot \hat{\mathbf{e}}_1 | e_j \rangle \langle e_j | e \mathbf{r}_e \cdot \hat{\mathbf{e}}_p | f \rangle}{\Delta_j} \right|^2 \quad (3.52)$$

where A_j is the excited-state Einstein A coefficient and we sum over possible polarisations of the emitted photon $\hat{\mathbf{e}}_p = \{\hat{\sigma}^+, \hat{\pi}, \hat{\sigma}^-\}$. The scattering rates for Raman transitions are much lower than the decay rates of the excited states, because the large detuning Δ ensures that the probability of actually populating the excited state remains small. There are two distinct types of scattering [139]:

Elastic/Rayleigh scattering

In some cases the ion may decay from $|e\rangle$ back to the same initial state from which it was excited ($|f\rangle = |i\rangle$). This is known as ‘Rayleigh scattering’ or ‘elastic scattering’. The quantum states of the ion and of the emitted photon are separable (not entangled), so an event of this type does not cause decoherence of the qubit state if the scattering rates from the two qubit states $\{|\uparrow\rangle, |\downarrow\rangle\}$ are the same (this is the case, for example, for clock qubits). If the rates are not the same, the elastic dephasing rate is [179]

$$\Gamma_{\text{el deph}} = \frac{E^2 A_j}{4\hbar^2} \left| \sum_{j,p} \frac{\langle \uparrow | e \mathbf{r}_e \cdot \hat{\mathbf{e}}_1 | e_j \rangle \langle e_j | e \mathbf{r}_e \cdot \hat{\mathbf{e}}_p | \uparrow \rangle - \langle \downarrow | e \mathbf{r}_e \cdot \hat{\mathbf{e}}_1 | e_j \rangle \langle e_j | e \mathbf{r}_e \cdot \hat{\mathbf{e}}_p | \downarrow \rangle}{\Delta_j} \right|^2. \quad (3.53)$$

Inelastic/Raman scattering

Otherwise, the ion may decay to a different state ($|f\rangle \neq |i\rangle$). This is known as ‘Raman scattering’ or ‘inelastic scattering’. In this case, the polarisation and frequency of the emitted photon depend on the initial and final internal states of the ion, i.e. the ion and photon are entangled. This type of scattering event does cause decoherence since the internal state changes with certainty, and the final state $|f\rangle$ may even be outside

the qubit subspace.

Both types of scattering also cause decoherence of the ion’s motional state due to the momentum imparted as the photon is emitted.

3.3 Mixed-species geometric phase gates

Section 3.1 discussed the use of ion-light interactions to drive single-qubit quantum gates on trapped ions. In order to satisfy the fourth DiVincenzo criterion of a universal gate set, the ability to generate entanglement between two qubits is also essential.

Qubits must interact with each other if they are to become entangled. The internal states of trapped ions do interact directly via the magnetic dipole interaction, but the strength of this interaction has a $1/r^3$ dependence on the ion separation and is extremely small (a few mHz) in a typical trap. Entanglement generated via this interaction has been demonstrated [100] but is impractical because of the long interaction times needed and the extreme sensitivity to noise. Alternatively, entanglement may be generated on a much faster timescale by engineering an electric dipole interaction using Rydberg states [195], or, most commonly, by using the fact that the ions’ motion in the trap is strongly coupled via the Coulomb interaction.

We can mediate an interaction between the internal states of different ions by first coupling these states to the shared motional state. As discussed earlier in this chapter, oscillating electromagnetic fields can couple to both the internal and motional states of an ion, so may be used for this purpose. The key ingredient for generating entanglement in this way is that the motion excited by the field must be dependent on the ions’ internal states. The first method for doing this was proposed by Cirac and Zoller in 1995 [41], and proceeds as follows:

1. Cool the ions to their motional ground state
2. Apply a π -pulse, on the RSB of the qubit transition, to the first ion

3. Apply a RSB 2π -pulse, on an auxiliary transition $|\downarrow\rangle \leftrightarrow |e\rangle$, to the second ion
4. Repeat step 2.

The result of this sequence is that if and only if the initial internal state was $|\uparrow\uparrow\rangle$, a phase shift of π is acquired. The operation may be used to perform a controlled NOT (CNOT) gate by sandwiching it between two single-qubit Hadamard gates⁶ on the second ion: $U_{\text{CNOT}} = (\mathbb{1} \otimes H) \cdot U_{\text{CZ}} \cdot (\mathbb{1} \otimes H)$.

The Cirac–Zoller gate requires cooling to the motional ground state and individual addressing of the ions, both of which are experimentally demanding. Other schemes known as ‘geometric phase gates’ — which do not have these requirements but instead only rely on being in the Lamb–Dicke regime — have since been proposed and implemented [164, 104, 17]. In these schemes, the ions’ motion is driven off-resonantly.

3.3.1 Forced quantum harmonic oscillator

We will examine the effect of an off-resonant driving force on a quantum harmonic oscillator (such as ions in a trap), following closely the approach in [103].

Consider a quantum harmonic oscillator subject to a force

$$\mathbf{F}(t) = F_0 \hat{\mathbf{r}}_m \cos(\omega t - \phi), \quad (3.54)$$

where $\omega = \omega_m + \delta_g$ and $\delta_g \ll \omega_m$ such that the force couples to a single motional mode m along the direction of the unit vector $\hat{\mathbf{r}}_m$.

The Hamiltonian for the system is

$$H_{\text{FQHO}} = \sum_{m'} \hbar \omega_{m'} (a_{m'}^\dagger a_{m'} + \frac{1}{2}) - \mathbf{F}(t) \cdot \mathbf{r}. \quad (3.55)$$

⁶ $H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$

Moving into the interaction picture with respect to $H_m = \sum_{m'} \hbar \omega_{m'} (a_{m'}^\dagger a_{m'} + \frac{1}{2})$, we find

$$\tilde{H}_{\text{FQHO}} = -\frac{F_0 \tilde{l}_m}{2} (a_m e^{i\delta_g t} e^{-i\phi} + a_m^\dagger e^{-i\delta_g t} e^{i\phi}), \quad (3.56)$$

where $\tilde{l}_m$ is the spread of the ground-state wavefunction in mode m . We have assumed the motional modes are well separated and used the RWA to ignore fast-oscillating terms at frequencies $2\omega_m$ and $\approx \omega_{m'} \pm \omega_m$ where $m' \neq m$. Note this may be written in the form

$$\tilde{H}_{\text{FQHO}} = \frac{V^*}{2} a_m e^{i\delta_g t} + \frac{V}{2} a_m^\dagger e^{-i\delta_g t}, \quad (3.57)$$

with $V_m = -F_0 \tilde{l}_m e^{i\phi}$.

We can calculate the propagator of $\tilde{H}_{\text{FQHO}}$ using a Magnus expansion [115]:

$$\begin{aligned} U(t) = \exp \Bigg\{ & -\frac{i}{\hbar} \int_0^t \tilde{H}(t') dt' \\ & - \frac{1}{2\hbar^2} \int_0^t dt' \int_0^{t'} dt'' [\tilde{H}(t'), \tilde{H}(t'')] \\ & + \frac{i}{6\hbar^3} \int_0^t dt' \int_0^{t'} dt'' \int_0^{t''} dt''' \left([\tilde{H}(t'), [\tilde{H}(t''), \tilde{H}(t''')]] \right. \\ & \quad \left. + [\tilde{H}(t'''), [\tilde{H}(t''), \tilde{H}(t')]] \right) \\ & + \dots \Bigg\}. \end{aligned} \quad (3.58)$$

The first term in $U(t)$ is equal to

$$D(\alpha(t)) = e^{\alpha a_m^\dagger - \alpha^* a_m}, \quad (3.59)$$

where

$$\alpha(t) \equiv \frac{F_0 \tilde{l}_m}{2\hbar \delta_g} e^{i\phi} (e^{-i\delta_g t} - 1) = \frac{F_0 \tilde{l}_m}{\hbar \delta_g} e^{i(\phi - \frac{\pi}{2} - \frac{\delta_g t}{2})} \sin\left(\frac{\delta_g t}{2}\right). \quad (3.60)$$

Using $[a_m, a_m^\dagger] = 1$ we find that the commutator $[\tilde{H}(t'), \tilde{H}(t'')]$ in the second term in

$U(t)$ is equal to a scalar, meaning that all further terms in $U(t)$ vanish. The second term is equal to

$$e^{-i\Phi(t)}, \quad (3.61)$$

where

$$\Phi(t) \equiv \left(\frac{F_0 \tilde{l}_m}{2\hbar\delta_g} \right)^2 (\delta_g t - \sin(\delta_g t)), \quad (3.62)$$

and the whole propagator is therefore

$$U(t) = e^{-i\Phi(t)} D(\alpha(t)). \quad (3.63)$$

The effect of the displacement operator $D(\alpha(t))$ on a motional state is to displace the state in motional phase space along a trajectory defined by $\alpha(t)$, without altering the shape of the state. The operator $e^{-i\Phi(t)}$ adds a global phase $\Phi(t)$.

In the frame rotating with the harmonic motion, the trajectory defined by $\alpha(t)$ in motional phase space is a circle of radius $\frac{F_0 \tilde{l}_m}{2\hbar\delta_g}$ with an initial phase of $\phi - \frac{\pi}{2}$. Motion is sinusoidally excited and de-excited with a period $t_g \equiv 2\pi/\delta_g$ such that the initial motional state is recovered at points where $\alpha(t) = 0 \Rightarrow t = Kt_g$ where K is the number of loops completed around the circular trajectory. At these points, the global phase acquired is

$$\Phi(Kt_g) = \frac{\pi K (F_0 \tilde{l}_m)^2}{2(\hbar\delta_g)^2}. \quad (3.64)$$

The area inside the circle is equal to $\pi \left(\frac{F_0 \tilde{l}_m}{2\hbar\delta_g} \right)^2 = \Phi(t_g)/2$. The phase $\Phi(Kt_g)$ is therefore proportional to the total area enclosed by the trajectory and is known as a ‘geometric phase’.

3.3.2 Spin-dependent forces

In order to use this off-resonant drive to implement a two-qubit gate for trapped ions, we require that the force felt by an ion is conditional on its internal state or ‘spin

state'. If this is the case, we may write

$$\tilde{H}_{\text{FQHO}} = \sum_{s=\uparrow_{\hat{\mathbf{n}}}, \downarrow_{\hat{\mathbf{n}}}} \left[\frac{V_{s,m}^*}{2} a_m e^{i\delta_g t} + \frac{V_{s,m}}{2} a_m^\dagger e^{-i\delta_g t} \right] |s\rangle\langle s|. \quad (3.65)$$

Here the force couples differently to the eigenstates $\{|\uparrow_{\hat{\mathbf{n}}}\rangle, |\downarrow_{\hat{\mathbf{n}}}\rangle\}$ of $\sigma_n \equiv \boldsymbol{\sigma} \cdot \hat{\mathbf{n}}$, the Pauli spin operator for a direction $\hat{\mathbf{n}}$.

The spin states $\{|\uparrow_{\hat{\mathbf{n}}}\rangle, |\downarrow_{\hat{\mathbf{n}}}\rangle\}$ follow different trajectories in motional phase space and each acquires a phase $\Phi_s(t)$ which depends on the area enclosed by that trajectory. If $V_{\uparrow_{\hat{\mathbf{n}}}} = V_{\downarrow_{\hat{\mathbf{n}}}}$, they both acquire the same global phase (the operator $\tilde{H}_{\text{FQHO}}$ becomes proportional to $\mathbb{1}$). If $V_{\uparrow_{\hat{\mathbf{n}}}} \neq V_{\downarrow_{\hat{\mathbf{n}}}}$, the result is a phase difference between $|\uparrow_{\hat{\mathbf{n}}}\rangle$ and $|\downarrow_{\hat{\mathbf{n}}}\rangle$. For one ion, this is a convoluted way to perform a single-qubit rotation. For two ions, however, the total force is equal to the sum of the forces on each ion:

$$\tilde{H}_{\text{FQHO}}^2 = \sum_{\substack{\text{ions} \\ i=1,2}} \sum_{\substack{s_i= \\ \uparrow_{\hat{\mathbf{n}}_i}, \downarrow_{\hat{\mathbf{n}}_i}}} \left[\frac{V_{i,s_i,m}^*}{2} a_m e^{i\delta_g t} + \frac{V_{i,s_i,m}}{2} a_m^\dagger e^{-i\delta_g t} \right] |s_i\rangle\langle s_i|_i, \quad (3.66)$$

where $|s_i\rangle\langle s_i|_i = |s_1\rangle\langle s_1| \otimes \mathbb{1}$ for $i = 1$ or $\mathbb{1} \otimes |s_2\rangle\langle s_2|$ for $i = 2$. Now each of the four spin states $\{|\uparrow_{\hat{\mathbf{n}}_1}\uparrow_{\hat{\mathbf{n}}_2}\rangle, |\uparrow_{\hat{\mathbf{n}}_1}\downarrow_{\hat{\mathbf{n}}_2}\rangle, |\downarrow_{\hat{\mathbf{n}}_1}\uparrow_{\hat{\mathbf{n}}_2}\rangle, |\downarrow_{\hat{\mathbf{n}}_1}\downarrow_{\hat{\mathbf{n}}_2}\rangle\}$ follows a different trajectory in phase space (as shown in figure 3.6), and the geometric phase acquired for spin state $|s_1, s_2\rangle$ after a closed loop in phase space is

$$\Phi_{s_1, s_2} = \frac{\pi |V_{1,s_1,m} + V_{2,s_2,m}|^2}{2(\hbar\delta_g)^2} = \frac{t_g |V_{1,s_1,m} + V_{2,s_2,m}|^2}{4\hbar^2\delta_g}. \quad (3.67)$$

The propagator associated with this Hamiltonian after K closed loops is

$$U(t_g, K) = \text{diag} \left(e^{-i\Phi_{\uparrow_{\hat{\mathbf{n}}_1}\uparrow_{\hat{\mathbf{n}}_2}}(Kt_g)}, e^{-i\Phi_{\uparrow_{\hat{\mathbf{n}}_1}\downarrow_{\hat{\mathbf{n}}_2}}(Kt_g)}, e^{-i\Phi_{\downarrow_{\hat{\mathbf{n}}_1}\uparrow_{\hat{\mathbf{n}}_2}}(Kt_g)}, e^{-i\Phi_{\downarrow_{\hat{\mathbf{n}}_1}\downarrow_{\hat{\mathbf{n}}_2}}(Kt_g)} \right), \quad (3.68)$$

in the basis $\{|\uparrow_{\hat{\mathbf{n}}_1}\uparrow_{\hat{\mathbf{n}}_2}\rangle, |\uparrow_{\hat{\mathbf{n}}_1}\downarrow_{\hat{\mathbf{n}}_2}\rangle, |\downarrow_{\hat{\mathbf{n}}_1}\uparrow_{\hat{\mathbf{n}}_2}\rangle, |\downarrow_{\hat{\mathbf{n}}_1}\downarrow_{\hat{\mathbf{n}}_2}\rangle\}$. This Hamiltonian may be used to implement a two-qubit geometric phase gate.




Figure 3.6: Phase-space trajectories for a general two-ion geometric phase gate. The sizes of the loops are determined by the values and relative phases of $V_{i,s_i,m}$. The orientations of the loops are determined by the phases of $V_{i,s_i,m}$.

The spin-dependent force may be produced via a variety of methods. In this work we use two different gates known as the the Mølmer–Sørensen gate and the light-shift gate. Their mechanisms are essentially equivalent, with the main difference being the basis in which they are performed.

3.3.3 Mølmer–Sørensen / $\sigma_\phi \otimes \sigma_\phi$ gate

One method of producing a spin-dependent force on the ions is to apply a bichromatic field which simultaneously drives the RSB of a motional mode m at detuning $-\delta_g$ and the BSB of the same mode at detuning $+\delta_g$. This is known as the Mølmer–Sørensen (MS) gate [164, 163].

The Hamiltonians for driving red and blue sidebands using pairs of Raman beams with frequency differences $\{-\omega_m + \delta_r, \omega_m + \delta_b\}$ and relative phases $\{\Delta\phi_{l,r}, \Delta\phi_{l,b}\}$ are⁷

$$\tilde{H}_{\text{RSB}} = -\frac{i\hbar\Omega_R}{2}\eta_m \left[e^{i(\Delta\phi_{l,r} + \mathbf{\Delta k}_r \cdot \mathbf{r}_{\text{eq}})} a_m e^{-i\delta_r t} |\uparrow\rangle\langle\downarrow| - e^{-i(\Delta\phi_{l,r} + \mathbf{\Delta k}_r \cdot \mathbf{r}_{\text{eq}})} a_m^\dagger e^{i\delta_r t} |\downarrow\rangle\langle\uparrow| \right], \quad (3.69)$$

⁷Here we have assumed we are in the Lamb–Dicke limit. If this is not the case, we need to modify Ω_R by a factor $\langle n_m | e^{i\eta_{i,m}(a_m + a_m^\dagger)} | n'_m \rangle$ as in equation (3.32).

$$\tilde{H}_{\text{BSB}} = -\frac{i\hbar\Omega_R}{2}\eta_m \left[e^{i(\Delta\phi_{l,b} + \mathbf{\Delta k}_b \cdot \mathbf{r}_{\text{eq}})} a_m^\dagger e^{-i\delta_b t} |\uparrow\rangle\langle\downarrow| - e^{-i(\Delta\phi_{l,b} + \mathbf{\Delta k}_b \cdot \mathbf{r}_{\text{eq}})} a_m e^{i\delta_b t} |\downarrow\rangle\langle\uparrow| \right]. \quad (3.70)$$

We apply bichromatic driving fields to two ions simultaneously such that the total Hamiltonian is $\tilde{H}_{\text{SB}} = \sum_{\text{ions } i} \left(\tilde{H}_{\text{RSB},i} + \tilde{H}_{\text{BSB},i} \right)$. For the case where $\delta_r = -\delta_g$ and $\delta_b = +\delta_g$, this is

$$\tilde{H}_{\text{SB}} = \sum_{i=1,2} -\frac{i\hbar\Omega_{R,i}}{2}\eta_{i,m} \left[\left(e^{i\phi_{M,i}} a_m e^{i\delta_g t} + e^{-i\phi_{M,i}} a_m^\dagger e^{-i\delta_g t} \right) \left(e^{-i\phi_{S,i}} \sigma_+^{(i)} - e^{i\phi_{S,i}} \sigma_-^{(i)} \right) \right], \quad (3.71)$$

where $\sigma_+^{(i)} = |\uparrow_i\rangle\langle\downarrow_i|$, $\sigma_-^{(i)} = |\downarrow_i\rangle\langle\uparrow_i|$ and we have defined the ‘motional phase’

$$\phi_{M,i} \equiv \frac{(\mathbf{\Delta k}_{r,i} \cdot \mathbf{r}_{\text{eq},i} + \Delta\phi_{l,r,i}) - (\mathbf{\Delta k}_{b,i} \cdot \mathbf{r}_{\text{eq},i} + \Delta\phi_{l,b,i})}{2} \quad (3.72)$$

and the ‘spin phase’

$$\phi_{S,i} \equiv -\frac{[(\mathbf{\Delta k}_{r,i} \cdot \mathbf{r}_{\text{eq},i} + \Delta\phi_{l,r,i}) + (\mathbf{\Delta k}_{b,i} \cdot \mathbf{r}_{\text{eq},i} + \Delta\phi_{l,b,i})]}{2}. \quad (3.73)$$

The second term in brackets in $\tilde{H}_{\text{SB}}$ may be written

$$e^{-i\phi_{S,i}} \sigma_+^{(i)} - e^{i\phi_{S,i}} \sigma_-^{(i)} = i \left[\cos\left(\phi_{S,i} + \frac{\pi}{2}\right) \sigma_x^{(i)} + \sin\left(\phi_{S,i} + \frac{\pi}{2}\right) \sigma_y^{(i)} \right] = i\sigma_{(\phi_{S,i} + \frac{\pi}{2})}^{(i)}, \quad (3.74)$$

where $\sigma_{(\phi_{S,i} + \frac{\pi}{2})}$ is the Pauli operator for a direction $\hat{\mathbf{n}}$ in the x - y plane at an angle $\phi_{S,i} + \frac{\pi}{2}$ from the x -axis. This leads to

$$\tilde{H}_{\text{SB}} = \sum_{i=1,2} \sum_{\substack{s_i = \\ \uparrow(\phi_{S,i} + \pi/2), \\ \downarrow(\phi_{S,i} + \pi/2)}} \frac{\hbar\Omega_{R,i}}{2} \eta_{i,m} \left(e^{i\phi_{M,i}} a_m e^{i\delta_g t} + e^{-i\phi_{M,i}} a_m^\dagger e^{-i\delta_g t} \right) |s_i\rangle\langle s_i|_i, \quad (3.75)$$

which is equivalent to the Hamiltonian $\tilde{H}_{\text{FQHO}}^2$ for a two-ion forced quantum harmonic

oscillator in equation (3.66), with

$$V_{i,s_i,m}^{MS} = \pm \hbar \Omega_{R,i} \eta_{i,m} e^{-i\phi_{M,i}}. \quad (3.76)$$

Hence the effect of $\tilde{H}_{\text{SB}}$ is to drive spin states along different trajectories in motional phase space as described in sections (3.3.1) and (3.3.2). The direction $\hat{\mathbf{n}}_i$ of the spin eigenstates to which the MS force couples differently (the basis of the gate) is defined by the spin phase $\phi_{S,i}$ for each ion, and lies in the x - y plane. The motional phase $\phi_{M,i}$ defines the phase of the motional driving force on each ion.

Using equation (3.67) we find the geometric phase acquired by spin state $|s_1, s_2\rangle$ (proportional to the area enclosed by its phase-space trajectory) after K closed loops is

$$\Phi_{s_1,s_2}^{\text{MS}}(Kt_g) = \frac{Kt_g}{4\delta_g} \left| \text{sgn}(s_1) \Omega_{R,1} \eta_{1,m} e^{-i\phi_{M,1}} + \text{sgn}(s_2) \Omega_{R,2} \eta_{2,m} e^{-i\phi_{M,2}} \right|^2, \quad (3.77)$$

where

$$\text{sgn}(s_i) = \begin{cases} +1 & s_i = \uparrow_{\hat{\mathbf{n}}_i} \\ -1 & s_i = \downarrow_{\hat{\mathbf{n}}_i} \end{cases}. \quad (3.78)$$

From equation (3.68) the general MS gate propagator in the basis $\{|\uparrow_{\hat{\mathbf{n}}_1} \uparrow_{\hat{\mathbf{n}}_2}\rangle, |\uparrow_{\hat{\mathbf{n}}_1} \downarrow_{\hat{\mathbf{n}}_2}\rangle, |\downarrow_{\hat{\mathbf{n}}_1} \uparrow_{\hat{\mathbf{n}}_2}\rangle, |\downarrow_{\hat{\mathbf{n}}_1} \downarrow_{\hat{\mathbf{n}}_2}\rangle\}$ is therefore

$$U_{\text{MS}}^{\hat{\mathbf{n}}}(Kt_g) = \text{diag} \left\{ \begin{aligned} & \exp \left[-i \frac{Kt_g}{4\delta_g} |\Omega_{R,1} \eta_{1,m} e^{-i\phi_{M,1}} + \Omega_{R,2} \eta_{2,m} e^{-i\phi_{M,2}}|^2 \right], \\ & \exp \left[-i \frac{Kt_g}{4\delta_g} |\Omega_{R,1} \eta_{1,m} e^{-i\phi_{M,1}} - \Omega_{R,2} \eta_{2,m} e^{-i\phi_{M,2}}|^2 \right], \\ & \exp \left[-i \frac{Kt_g}{4\delta_g} |-\Omega_{R,1} \eta_{1,m} e^{-i\phi_{M,1}} + \Omega_{R,2} \eta_{2,m} e^{-i\phi_{M,2}}|^2 \right], \\ & \exp \left[-i \frac{Kt_g}{4\delta_g} |-\Omega_{R,1} \eta_{1,m} e^{-i\phi_{M,1}} - \Omega_{R,2} \eta_{2,m} e^{-i\phi_{M,2}}|^2 \right] \end{aligned} \right\}. \quad (3.79)$$

As the forces on the two spin states for each ion have equal and opposite magnitude $\propto \Omega_{R,i}\eta_{i,m}$, the even-parity states acquire the same phase ($\Phi_{\uparrow\hat{n}_1\uparrow\hat{n}_2} = \Phi_{\downarrow\hat{n}_1\downarrow\hat{n}_2} = \Phi_{\text{even}}/2$) and the odd-parity states acquire the same phase ($\Phi_{\uparrow\hat{n}_1\downarrow\hat{n}_2} = \Phi_{\downarrow\hat{n}_1\uparrow\hat{n}_2} = \Phi_{\text{odd}}/2$).

For an optimal gate we require

i. $e^{-i\phi_{M,1}} = \pm e^{-i\phi_{M,2}}$, and

ii. balanced forces on each ion $|\Omega_{R,1}\eta_{1,m}| = |\Omega_{R,2}\eta_{2,m}|$.

Assuming we satisfy (i) with $e^{-i\phi_{M,1}} = +e^{-i\phi_{M,2}}$ (as is the case in this work), for a same-species gate, we may drive both ions with the same Raman beam pairs and we therefore satisfy (ii) ‘for free’ as long as the two ions are equally illuminated in those beams. In the mixed-species case we must use different beam pairs to drive each ion, but we can calibrate the laser powers and polarisations (which fix $\Omega_{R,1}$ and $\Omega_{R,2}$) for each mode such that $|\Omega_{R,1}\eta_{1,m}| = |\Omega_{R,2}\eta_{2,m}|$. Then the MS gate Hamiltonian becomes

$$H_{\text{MS}} = \frac{\hbar\tilde{\Omega}}{2} \left(\sigma_{\phi_{S,1}+\pi/2}^{(1)} \pm \sigma_{\phi_{S,2}+\pi/2}^{(2)} \right) (\hat{a}e^{i\delta_g t} + \hat{a}^\dagger e^{-i\delta_g t}), \quad (3.80)$$

where $\tilde{\Omega} \equiv |\Omega_{R,1}\eta_{1,m}| = |\Omega_{R,2}\eta_{2,m}|$ is the same for each ion. We take the top signs if we are driving the ip mode, where $\eta_{1,m} = \eta_{2,m}$, or the bottom signs if we are driving the oop mode, where $\eta_{1,m} = -\eta_{2,m}$.

The propagators for each mode are

$$\begin{aligned} U_{\text{MS,ip}}^{\hat{n}}(Kt_g) &= \text{diag} \left\{ \exp \left[-i \frac{Kt_g}{4\delta_g} |2\tilde{\Omega}|^2 \right], 1, 1, \exp \left[-i \frac{Kt_g}{4\delta_g} |2\tilde{\Omega}|^2 \right] \right\}, \\ U_{\text{MS,oop}}^{\hat{n}}(Kt_g) &= \text{diag} \left\{ 1, \exp \left[-i \frac{Kt_g}{4\delta_g} |2\tilde{\Omega}|^2 \right], \exp \left[-i \frac{Kt_g}{4\delta_g} |2\tilde{\Omega}|^2 \right], 1 \right\}. \end{aligned} \quad (3.81)$$

Both of these propagators are equal (up to a global phase) to the propagator $\exp \left[-i \frac{\Phi}{2} \sigma_{(\phi_{S,1}+\pi/2)} \otimes \sigma_{(\phi_{S,2}+\pi/2)} \right]$ with $\Phi = \frac{Kt_g}{4\delta_g} |2\tilde{\Omega}|^2$; hence the MS gate is also known




Figure 3.7: Phase-space trajectories for the MS gate assuming $e^{-i\phi_{M,1}} = e^{-i\phi_{M,2}} = 1$ (true in this thesis). If the laser powers are calibrated so that the forces on each ion are balanced ($|\Omega_{R,1}\eta_{1,m}| = |\Omega_{R,2}\eta_{2,m}|$), the smaller loops vanish and we drive an optimal gate.

as the $\sigma_\phi \otimes \sigma_\phi$ gate. We can adjust the Rabi frequencies for a particular gate detuning and gate time $\{m, \delta_g, Kt_g\}$ so that we have $\Phi = \frac{\pi}{2}$, leading to

$$U_{\text{MS,ip}}^{\hat{\mathbf{n}}} = -i \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad U_{\text{MS,oop}}^{\hat{\mathbf{n}}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -i & 0 & 0 \\ 0 & 0 & -i & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (3.82)$$

Using two single-qubit operations — the Hadamard H and the single-qubit $\frac{\pi}{8}$ z -phase gate⁸ T — these two propagators may be converted into a CNOT gate via $U_{\text{CNOT}} = (\mathbb{1} \otimes H) \cdot (T^{\dagger 2} \otimes T^{\dagger 2}) \cdot U_{\text{MS}}^{\hat{\mathbf{n}}} \cdot (\mathbb{1} \otimes H)$. The MS gate may be used in addition to single-qubit operations to form a universal gate set.

⁸ $T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$

In the usual z -basis, up to a global phase, the propagators above are

$$U_{\text{MS,ip,oop}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 0 & \pm i e^{-i(\phi_{S,1} + \phi_{S,2})} \\ 0 & 1 & \mp i e^{i(\phi_{S,2} - \phi_{S,1})} & 0 \\ 0 & \mp i e^{-i(\phi_{S,2} - \phi_{S,1})} & 1 & 0 \\ \pm i e^{i(\phi_{S,1} + \phi_{S,2})} & 0 & 0 & 1 \end{pmatrix}, \quad (3.83)$$

where we again take the top signs for the ip mode and the bottom signs for the oop mode.

3.3.3.1 Beam geometry

In order to satisfy condition (i) for an optimal MS gate we must consider in more detail the motional phase $\phi_{M,i}$ for each ion, given by equation (3.72). It depends on the wavevector differences $\{\Delta \mathbf{k}_{r,i}, \Delta \mathbf{k}_{b,i}\}$ and relative laser phases $\{\Delta \phi_{l,r,i}, \Delta \phi_{l,b,i}\}$ for the red and blue sideband drives, which in turn depend on the geometry of the experimental setup. As noted in [103], the geometry may be chosen to reduce sensitivity to particular experimental calibrations or instabilities.

In this work, we drive the sidebands using beam pairs with frequency differences $\omega_0 \pm (\omega_m + \delta_g)$, where $\Delta \mathbf{k}_r = \Delta \mathbf{k}_b$. In this geometry the motional phase is ≈ 0 for both ions, so we satisfy $e^{-i\phi_{M,1}} = e^{-i\phi_{M,2}}$ automatically. Importantly, the motional phase is independent of the ion spacing, which would otherwise be inconvenient when working with mixed-species crystals⁹. The frequency $\pm(\omega_m + \delta_g)$ is easily bridged using an acousto-optic modulator, so these two tones may originate from the same beam. On the other hand, the spin phases $\phi_{S,i}$ (and therefore the gate basis) are sensitive to fluctuations in the relative optical phases of the Raman beams, on the scale of an optical wavelength. These may arise due to thermal effects or air currents

⁹Generally, the optimal spacing would be different for a Ca–Sr MS gate, for a Ca–Ca MS gate and for a Sr–Sr MS gate, and for efficient operation it is impractical to have to change the trap parameters to alter the ion spacing when switching between these different operations.

which are difficult to control in the lab.

Despite this phase-sensitivity, carrier operations performed using these lasers have a fixed phase relation ($\phi_{S,i}$) to the MS gate phase ($\phi_{S,i} + \frac{\pi}{2}$). Therefore if we use the same lasers for both single- and two-qubit operations, gates within a single run of the experiment will remain coherent, provided the phase is stable over its duration. For longer sequences, we may wish to use microwaves/RF to drive single-qubit operations instead. To do so, the MS gate could be wrapped in single-qubit $\frac{\pi}{2}$ -pulses to eliminate the dependence of the final state on the (random) relative phases between the lasers and the microwaves/RF, as described in [103]. The geometry and this sequence are described in more detail in appendix A.

3.3.3.2 Carrier coupling

The Hamiltonian $\tilde{H}_{\text{SB}}$ is simply the sum of a red and a blue sideband drive. However, the fields required to produce this drive will also result in an unwanted, off-resonant drive of the carrier transition, as in equations (3.26) and (3.29). For each ion i , this drive is

$$\begin{aligned} \tilde{H}_{\text{carrier},i} = & -\frac{\hbar\Omega_{R,i}}{2} \left[\left(e^{i(\Delta\phi_{l,r,i} + \Delta k_r z_{\text{eq},i} + (\omega_m + \delta_g)t)} + e^{i(\Delta\phi_{l,b,i} + \Delta k_b z_{\text{eq},i} - (\omega_m + \delta_g)t)} \right) \sigma_+^{(i)} \right. \\ & \left. + \left(e^{-i(\Delta\phi_{l,r,i} + \Delta k_r z_{\text{eq},i} + (\omega_m + \delta_g)t)} + e^{-i(\Delta\phi_{l,b,i} + \Delta k_b z_{\text{eq},i} - (\omega_m + \delta_g)t)} \right) \sigma_-^{(i)} \right]. \end{aligned} \quad (3.84)$$

This may be written

$$\tilde{H}_{\text{carrier},i} = -\hbar\Omega_{R,i} \cos(\phi_{M,i} + (\omega_m + \delta_g)t) \sigma_{\phi_{S,i}}^{(i)}, \quad (3.85)$$

which is an off-resonant drive of the carrier transition, with phase $\phi_{S,i}$, which oscillates in amplitude. Although strength of this coupling is enhanced by a factor $1/\eta_{i,m}$ relative to the sideband terms, its may be neglected if the laser pulses used to drive the gate are appropriately amplitude-shaped (see section 3.3.5.3).

3.3.3.3 Light shifts

As well as the off-resonant carrier drive, $\tilde{H}_{\text{SB}}$ also neglects light shifts on the qubit states of each ion due to various couplings of the form in equation (3.18). Since there are three beams $\{1, 2r, 2b\}_i$ which couple to each ion, there are three single-beam light shifts for each ion¹⁰ which arise from couplings of the qubit states to the excited states in the P-levels. These are:

$$\begin{aligned}
1 : & \quad \frac{\hbar\Omega_{1,i,\downarrow}^2}{4\Delta_i} |\downarrow_i\rangle\langle\downarrow_i| + \frac{\hbar\Omega_{1,i,\uparrow}^2}{4\Delta_i} |\uparrow_i\rangle\langle\uparrow_i| \\
2r : & \quad \frac{\hbar\Omega_{2r,i,\downarrow}^2}{4(\Delta_i + \omega_{0,i} + \omega_m + \delta_g)} |\downarrow_i\rangle\langle\downarrow_i| + \frac{\hbar\Omega_{2r,i,\uparrow}^2}{4(\Delta_i + \omega_m + \delta_g)} |\uparrow_i\rangle\langle\uparrow_i| \\
2b : & \quad \frac{\hbar\Omega_{2b,i,\downarrow}^2}{4(\Delta_i + \omega_{0,i} - (\omega_m + \delta_g))} |\downarrow_i\rangle\langle\downarrow_i| + \frac{\hbar\Omega_{2b,i,\uparrow}^2}{4(\Delta_i - (\omega_m + \delta_g))} |\uparrow_i\rangle\langle\uparrow_i|.
\end{aligned} \tag{3.86}$$

The differential single-beam light shifts on the two qubit states may be nulled by careful choice of laser polarisations and beam geometry¹¹.

There are also two-beam light shifts due to the RSB and BSB drives coupling off-resonantly to the carrier¹² which, if properly calibrated, will be equal and opposite:

$$\begin{aligned}
1 + 2r : & \quad \frac{\hbar\Omega_{R,i}^2}{-4(\omega_m + \delta_g)} [|\downarrow_i\rangle\langle\downarrow_i| - |\uparrow_i\rangle\langle\uparrow_i|] \\
1 + 2b : & \quad \frac{\hbar\Omega_{R,i}^2}{4(\omega_m + \delta_g)} [|\downarrow_i\rangle\langle\downarrow_i| - |\uparrow_i\rangle\langle\uparrow_i|].
\end{aligned} \tag{3.87}$$

¹⁰We assume there are no light shifts on ion 1 (2) from the beams driving ion 2 (1). For a same-species gate, the gate beams are generally common to both ions.

¹¹When including the qubit frequency splitting into the Raman detuning, a small asymmetry in the σ^+ and σ^- polarisation components is necessary to null the shift. Since the qubit splitting is larger in $^{43}\text{Ca}^+$ than in $^{88}\text{Sr}^+$, the shift cannot be exactly nulled for both species simultaneously. We set the polarisations to null the shift on $^{43}\text{Ca}^+$, and the error resulting from the residual shift on $^{88}\text{Sr}^+$ is negligible.

¹²To a lesser extent the RSB drive will also couple off-resonantly to the BSB, and vice versa. The light shifts on the qubit states due to this coupling will also be equal and opposite if properly calibrated.

3.3.4 Light-shift / $\sigma_z \otimes \sigma_z$ gate

A different method of producing a spin-dependent force on the ions is the light-shift (LS) gate [104]. Two non-copropagating Raman beams, with a frequency difference $\delta = \omega_m + \delta_g$ tuned close to an axial motional mode, form a polarisation travelling standing wave in which the ions sit. Due to the AC Stark effect, the ions experience a periodic potential (i.e. a force) which oscillates at frequency δ and therefore off-resonantly drives the motional mode m . The force felt by an ion depends on its position, its atomic structure, its Lamb–Dicke factor and its spin state.

We may choose the Raman frequency splitting $\delta = \omega_m + \delta_g \ll \omega_0$, and the beam polarisations, such that we do not couple the two qubit states $\{|\uparrow\rangle, |\downarrow\rangle\}_i$ for each ion to each other, but instead we couple to each state individually. The total effective Raman Hamiltonian on two ions $i = 1, 2$ (in the Lamb–Dicke limit) is

$$\begin{aligned} \tilde{H}_R^{\text{LS}} = & \sum_i \left[\frac{\hbar\Omega_{1,i,\downarrow}^2}{4\Delta_i} + \frac{\hbar\Omega_{2,i,\downarrow}^2}{4(\Delta_i + \omega_m + \delta_g)} \right] |\downarrow\rangle\langle\downarrow|_i \\ & + \sum_i \left[\frac{\hbar\Omega_{1,i,\uparrow}^2}{4(\Delta_i - \omega_{0,i})} + \frac{\hbar\Omega_{2,i,\uparrow}^2}{4(\Delta_i - \omega_{0,i} + \omega_m + \delta_g)} \right] |\uparrow\rangle\langle\uparrow|_i \\ & - \sum_{\substack{i, s_i = \\ \uparrow_i, \downarrow_i}} \frac{\hbar\Omega_{R,i,s_i}}{2} |s_i\rangle\langle s_i|_i \left[e^{i(\Delta\phi_i - \delta t)} + e^{-i(\Delta\phi_i - \delta t)} \right. \\ & \left. + i \sum_m \eta_{i,m} (a_m e^{-i\omega_m t} + a_m^\dagger e^{i\omega_m t}) (e^{i(\Delta\phi_i - \delta t)} - e^{-i(\Delta\phi_i - \delta t)}) \right], \end{aligned} \quad (3.88)$$

where $\Delta\phi_i \equiv \Delta\phi_l + \mathbf{\Delta k} \cdot \mathbf{r}_{\text{eq},i}$ is different for the two ions due to their different equilibrium positions in the trap. The first two terms in $\tilde{H}_R^{\text{LS}}$ are single-beam light shifts, which we shall again neglect. Assuming the motional modes are well separated, using

the RWA, and neglecting the single-beam light shifts,

$$\begin{aligned} \tilde{H}_R^{\text{LS}} = & \sum_{\substack{i, s_i = \\ \uparrow_i, \downarrow_i}} -\frac{\hbar \Omega_{R, i, s_i}}{2} |s_i\rangle \langle s_i|_i 2 \cos(\Delta\phi_i - (\omega_m + \delta_g)t) \\ & + \sum_{\substack{i, s_i = \\ \uparrow_i, \downarrow_i}} \left[\left(-\frac{\hbar \Omega_{R, i, s_i} \eta_{i, m}}{2} (-i) e^{-i\Delta\phi_i} a_m e^{i\delta_g t} - \frac{\hbar \Omega_{R, i, s_i} \eta_{i, m}}{2} (+i) e^{i\Delta\phi_i} a_m^\dagger e^{-i\delta_g t} \right) |s_i\rangle \langle s_i|_i \right]. \end{aligned} \quad (3.89)$$

Here, the first term is an unwanted oscillating, time-dependent light shift. If $\Omega_{R, i, s_i} \ll \omega_m + \delta_g$, we may suppress errors caused by this term similarly to the unwanted carrier coupling in the MS gate, by using pulse shaping (see section 3.3.5.3).

The second term is equivalent to $\tilde{H}_{\text{FQHO}}^2$ in equation (3.66), with

$$V_{i, s_i, m}^{\text{LS}} = -i\hbar \Omega_{R, i, s_i} \eta_{i, m} e^{i\Delta\phi_i}. \quad (3.90)$$

The effect of this Hamiltonian is therefore to drive the z -basis spin states $\{|\uparrow_1\uparrow_1\rangle, |\uparrow_1\downarrow_2\rangle, |\downarrow_1\uparrow_2\rangle, |\downarrow_1\downarrow_2\rangle\}$ along different circular trajectories in motional phase space. The geometric phase acquired by spin state $|s_1, s_2\rangle$ after K closed loops is

$$\Phi_{s_1, s_2}^{\text{LS}}(Kt_g) = \frac{Kt_g}{4\delta_g} \left| \Omega_{R, 1, s_1} \eta_{1, m} + \Omega_{R, 2, s_2} \eta_{2, m} e^{i\phi_M} \right|^2, \quad (3.91)$$

where the motional phase $\phi_M \equiv \Delta\phi_2 - \Delta\phi_1$ is the relative phase of the driving force on each ion, and is determined by the ion spacing along the axial direction¹³. The drive phase is the same for both ions if $|\Delta\mathbf{k} \cdot (\mathbf{r}_{\text{eq}, 2} - \mathbf{r}_{\text{eq}, 1})| = 2N\pi$, or opposite if $|\Delta\mathbf{k} \cdot (\mathbf{r}_{\text{eq}, 2} - \mathbf{r}_{\text{eq}, 1})| = 2(N + \frac{1}{2})\pi$; $N \in \mathbb{Z}$.

¹³For light-shift gates in this thesis we drive both ion species with the same Raman beam pair, so the required ion-spacing is the same whether we wish to do a mixed-species or a same-species gate.

By combining equation (3.91) with equation (3.68) we find the general propagator for an LS gate

$$U_{\text{LS}}(Kt_g) = \text{diag} \left\{ \begin{aligned} &\exp \left[-i \frac{Kt_g}{4\delta_g} |\Omega_{R,1,\uparrow_1} \eta_{1,m} + \Omega_{R,2,\uparrow_2} \eta_{2,m} e^{i\phi_M}|^2 \right], \\ &\exp \left[-i \frac{Kt_g}{4\delta_g} |\Omega_{R,1,\uparrow_1} \eta_{1,m} + \Omega_{R,2,\downarrow_2} \eta_{2,m} e^{i\phi_M}|^2 \right], \\ &\exp \left[-i \frac{Kt_g}{4\delta_g} |\Omega_{R,1,\downarrow_1} \eta_{1,m} + \Omega_{R,2,\uparrow_2} \eta_{2,m} e^{i\phi_M}|^2 \right], \\ &\exp \left[-i \frac{Kt_g}{4\delta_g} |\Omega_{R,1,\downarrow_1} \eta_{1,m} + \Omega_{R,2,\downarrow_2} \eta_{2,m} e^{i\phi_M}|^2 \right] \end{aligned} \right\} \quad (3.92)$$

in the z -basis.

3.3.4.1 Simple same-species case

Although all the two-qubit gates in this work are performed on mixed-species crystals, it is instructive to look first at the case where both ions 1 and 2 are of the same species and we use the same qubit states in each ion. In this situation there are many symmetries which simplify $U_{\text{LS}}(Kt_g)$, namely:

- i. $\{|\uparrow_1\rangle, |\downarrow_1\rangle\} = \{|\uparrow_2\rangle, |\downarrow_2\rangle\} = \{|\uparrow\rangle, |\downarrow\rangle\}$,
- ii. $\Omega_{R,1,\uparrow_1} = \Omega_{R,2,\uparrow_2} = \Omega_\uparrow$ and $\Omega_{R,1,\downarrow_1} = \Omega_{R,2,\downarrow_2} = \Omega_\downarrow$
- iii. $|\eta_{1,m}| = |\eta_{2,m}| = \eta_m$.

Noting these simplifications, and noting also that in our experiment we set the ion spacing to be a half-integer multiple of the effective wavelength λ_{eff} so that $e^{i\phi_M} = -1$,

$U_{\text{LS}}(Kt_g)$ becomes

$$U_{\text{LS}}^{\text{same}}(Kt_g) = \text{diag} \left\{ \exp \left[-i \frac{Kt_g \eta_m^2}{4\delta_g} |\Omega_{\uparrow} \mp \Omega_{\uparrow}|^2 \right], \right. \\ \exp \left[-i \frac{Kt_g \eta_m^2}{4\delta_g} |\Omega_{\uparrow} \mp \Omega_{\downarrow}|^2 \right], \\ \exp \left[-i \frac{Kt_g \eta_m^2}{4\delta_g} |\Omega_{\downarrow} \mp \Omega_{\uparrow}|^2 \right], \\ \left. \exp \left[-i \frac{Kt_g \eta_m^2}{4\delta_g} |\Omega_{\downarrow} \mp \Omega_{\downarrow}|^2 \right] \right\}. \quad (3.93)$$

We take the negative signs if we are driving the ip mode, where $\eta_{1,m} = \eta_{2,m}$, or the positive signs if we are driving the oop mode, where $\eta_{1,m} = -\eta_{2,m}$.

No hyperfine structure

Ideally, we would like $\Omega_{\uparrow} = -\Omega_{\downarrow}$. This is true for the Zeeman qubit states $\{|\uparrow_Z\rangle, |\downarrow_Z\rangle\}$ in $^{88}\text{Sr}^+$ — which does not have hyperfine structure — due to the symmetry of the matrix elements. It leads to

$$U_{\text{LS,ip}}^{\text{same}}(Kt_g) = \text{diag} \left\{ 1, \exp \left[-i \frac{Kt_g \eta_m^2}{4\delta_g} |2\Omega_{\uparrow}|^2 \right], \exp \left[-i \frac{Kt_g \eta_m^2}{4\delta_g} |-2\Omega_{\uparrow}|^2 \right], 1 \right\} \\ U_{\text{LS,oop}}^{\text{same}}(Kt_g) = \text{diag} \left\{ \exp \left[-i \frac{Kt_g \eta_m^2}{4\delta_g} |2\Omega_{\uparrow}|^2 \right], 1, 1, \exp \left[-i \frac{Kt_g \eta_m^2}{4\delta_g} |-2\Omega_{\uparrow}|^2 \right] \right\} \quad (3.94)$$

for ip and oop modes respectively. Here we see that for the ip (oop) mode, only the odd (even) parity states are driven in loops in motional phase space and acquire a geometric phase. The other spin states remain unchanged.

Up to a global phase, both of these propagators are equal to the propagator $\exp \left[-i \frac{\Phi}{2} \sigma_z \otimes \sigma_z \right]$ with $\Phi = \frac{Kt_g \eta_m^2}{4\delta_g} |2\Omega_{\uparrow}|^2$; hence the light-shift gate is also known as the $\sigma_z \otimes \sigma_z$ gate.

We can again find a laser power (which fixes $\Omega_{\uparrow}$) for each mode, gate detuning and gate time $\{m, \delta_g, Kt_g\}$ such that $\Phi = \frac{\pi}{2}$, and we recover the same propagators as

we found for the MS gate¹⁴:

$$U_{\text{LS,ip}}^{\text{ideal}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & -i & 0 & 0 \\ 0 & 0 & -i & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad U_{\text{LS,oop}}^{\text{ideal}} = -i \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \quad (3.95)$$

This time, the propagators are already in the z -basis.

Hyperfine structure and the spin-echo sequence

If we perform this gate on two $^{43}\text{Ca}^+$ ions using hyperfine states as our qubit states, due to asymmetries in the matrix elements for hyperfine states it is not necessarily the case that $\Omega_{\uparrow} = -\Omega_{\downarrow}$. In fact, for clock qubits $\{|\uparrow_C\rangle, |\downarrow_C\rangle\}$ it can be shown that $\Omega_{\uparrow_C} = \Omega_{\downarrow_C}$ [102], so U_{LS} can only ever lead to a global phase on all spin states, which does not generate entanglement. Therefore we cannot drive a light-shift gate on a clock qubit¹⁵. In this work we perform the light-shift gate using the stretch qubit states $\{|\uparrow_S\rangle, |\downarrow_S\rangle\}$ of $^{43}\text{Ca}^+$.

There is a small asymmetry in the matrix elements for the stretch states, such that $\Omega_{\downarrow_S} = -\Omega_{\uparrow_S} - \delta\Omega$. The geometric phases $\{\Phi_{\uparrow_S\uparrow_S}, \Phi_{\uparrow_S\downarrow_S}, \Phi_{\downarrow_S\uparrow_S}, \Phi_{\downarrow_S\downarrow_S}\}$ for the ip and oop modes are therefore proportional to

$$\begin{aligned} \text{ip: } & \propto \{0, |2\Omega_{\uparrow_S} + \delta\Omega|^2, |-2\Omega_{\uparrow_S} - \delta\Omega|^2, 0\}, \\ \text{oop: } & \propto \{|2\Omega_{\uparrow_S}|^2, |-\delta\Omega|^2, |-\delta\Omega|^2, |-2\Omega_{\uparrow_S} - 2\delta\Omega|^2\}. \end{aligned}$$

For the ip mode, this again results in a $\sigma_z \otimes \sigma_z$ operator, with a small increase in

¹⁴The gate could equally be performed with integer ion spacing $|(\mathbf{r}_{\text{eq},2} - \mathbf{r}_{\text{eq},1})_z| = N\lambda_{\text{eff}}$, in which case the results $U_{\text{LS,ip}}^{\text{ideal}}$ and $U_{\text{LS,oop}}^{\text{ideal}}$ would swap.

¹⁵This is true if the gate is driven using Raman beams with a detuning Δ from an electric dipole transition, in the limit where Δ is large compared to the qubit frequency. A proposal does exist to demonstrate a light-shift gate on a hyperfine clock qubit by instead using a small detuning from a dipole-forbidden optical transition [8].

laser power needed to account for the asymmetry. For the oop mode, the phase acquired for spin states $|\uparrow_S \uparrow_S\rangle$ and $|\downarrow_S \downarrow_S\rangle$ is no longer equal. We can solve this problem by performing the gate in a ‘spin-echo’ sequence. The gate is driven in two halves (two distinct laser pulses, each driving one phase-space loop) with a resonant π -pulse on both ions between the two halves, flipping the states $|\uparrow_S\rangle \leftrightarrow |\downarrow_S\rangle$. The total geometric phase acquired by each spin state is then equal to the sum of the phases acquired in each loop so that both even-parity spin states now acquire a phase $\Phi_{\uparrow_S \uparrow_S} + \Phi_{\downarrow_S \downarrow_S} = |2\Omega_{\uparrow_S}|^2 + |-2\Omega_{\uparrow_S} - 2\delta\Omega_{\uparrow_S}|^2$.

The effect of the two halves is that all four spin states acquire a global phase

$$\Phi_{\mathbb{1}}^{\text{oop}} = \Phi_{\uparrow_S \downarrow_S} + \Phi_{\downarrow_S \uparrow_S}, \quad (3.96)$$


and the even-parity states acquire an extra, two-qubit phase

$$\Phi_{\sigma_z}^{\text{oop}} = (\Phi_{\uparrow_S \uparrow_S} + \Phi_{\downarrow_S \downarrow_S}) - (\Phi_{\uparrow_S \downarrow_S} + \Phi_{\downarrow_S \uparrow_S}). \quad (3.97)$$


Entanglement is generated due to the two-qubit phase, and the global phase on the final state may be ignored. However, as some of the Rabi frequency is ‘wasted’ on generating this unnecessary global phase, a small increase in laser power is needed to drive the gate. Performing the light-shift gate in this spin-echo sequence has other additional benefits such as reducing errors due to (a) unequal illumination of the two ions and (b) magnetic-field noise over the gate duration (see section 3.3.5.5).

3.3.4.2 Mixed-species case


We implement a light-shift gate on a two-ion, mixed-species crystal of $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$, using the stretch qubit in calcium ($\{|\uparrow_1\rangle, |\downarrow_1\rangle\} = \{|\uparrow_S\rangle, |\downarrow_S\rangle\}$) and the Zeeman qubit in strontium ($\{|\uparrow_2\rangle, |\downarrow_2\rangle\} = \{|\uparrow_Z\rangle, |\downarrow_Z\rangle\}$). For this pair of species we may use the same two Raman beams to drive both ions simultaneously (see section 6.2.1).



(a) Same-species, with $\Omega_{\uparrow} = -\Omega_{\downarrow}$. The gate is maximally efficient.



(b) Same-species with asymmetry $\Omega_{\uparrow} \neq -\Omega_{\downarrow}$. The $|\uparrow_1, \downarrow_2\rangle$ and $|\downarrow_1, \uparrow_2\rangle$ trajectories overlap.



(c) Mixed-species with one species having asymmetry $\Omega_{R,1,\uparrow_1} \neq -\Omega_{R,1,\downarrow_1}$.

Figure 3.8: LS gate phase-space trajectories. Half-integer ion spacing is assumed, and the total Rabi frequency $\sum_{i,s_i} |\Omega_{R,i,s_i}|$ is constant in each plot. For the same-species gates, the $^{43}\text{Ca}^+$ Lamb–Dicke factors in our setup $\{\eta_{\text{ip}}, \eta_{\text{oop}}\} = \{0.112, \pm 0.085\}$ were used, and for mixed-species, $\eta_{\text{ip,oop}}^{\text{Ca}} = \{0.09, 0.127\}$ and $\eta_{\text{ip,oop}}^{\text{Sr}} = \{0.124, -0.045\}$.

For these two qubits, the mixed-species propagator $U_{\text{LS}}^{\text{mixed}}$ takes almost its most general form, except for the simplification that $\Omega_{R,\text{Sr},\uparrow Z} = -\Omega_{R,\text{Sr},\downarrow Z}$ due to the lack of hyperfine structure in $^{88}\text{Sr}^+$. Eliminating $\Omega_{R,\text{Sr},\downarrow Z}$, and with half-integer ion spacing, the propagator is

$$U_{\text{LS}}^{\text{mixed}}(Kt_g) = \text{diag} \left\{ \begin{aligned} &\exp \left[-i \frac{Kt_g}{4\delta_g} |\Omega_{R,\text{Ca},\uparrow S} \eta_{\text{Ca},m} \mp \Omega_{R,\text{Sr},\uparrow Z} |\eta_{\text{Sr},m}|^2 \right], \\ &\exp \left[-i \frac{Kt_g}{4\delta_g} |\Omega_{R,\text{Ca},\uparrow S} \eta_{\text{Ca},m} \pm \Omega_{R,\text{Sr},\uparrow Z} |\eta_{\text{Sr},m}|^2 \right], \\ &\exp \left[-i \frac{Kt_g}{4\delta_g} |\Omega_{R,\text{Ca},\downarrow S} \eta_{\text{Ca},m} \mp \Omega_{R,\text{Sr},\uparrow Z} |\eta_{\text{Sr},m}|^2 \right], \\ &\exp \left[-i \frac{Kt_g}{4\delta_g} |\Omega_{R,\text{Ca},\downarrow S} \eta_{\text{Ca},m} \pm \Omega_{R,\text{Sr},\uparrow Z} |\eta_{\text{Sr},m}|^2 \right] \end{aligned} \right\}. \quad (3.98)$$

We again employ the technique of performing the gate in a spin-echo sequence, driving two loops and flipping the spins of both ions after the first loop. Generalising equations (3.96) and (3.97), the result is a global phase on all four spin states

$$\Phi_{\mathbb{1}} = \begin{cases} \Phi_{\text{even}} & \text{ip} \\ \Phi_{\text{odd}} & \text{oop} \end{cases}, \quad (3.99)$$

as well as the desired two-qubit phase

$$\Phi_{\sigma_z} = \begin{cases} \Phi_{\text{odd}} - \Phi_{\text{even}} & \text{ip} \\ -(\Phi_{\text{odd}} - \Phi_{\text{even}}) & \text{oop} \end{cases}, \quad (3.100)$$

where

$$\begin{aligned} \Phi_{\text{odd}} &\equiv \Phi_{\uparrow S \downarrow Z} + \Phi_{\downarrow S \uparrow Z}, \\ \Phi_{\text{even}} &\equiv \Phi_{\uparrow S \uparrow Z} + \Phi_{\downarrow S \downarrow Z}. \end{aligned} \quad (3.101)$$

We can define the ‘gate efficiency’

$$\Theta \equiv \frac{|\Phi_{\text{odd}} - \Phi_{\text{even}}|}{|\Phi_{\text{odd}} + \Phi_{\text{even}}|}. \quad (3.102)$$

For a smaller gate efficiency, more Rabi frequency is ‘wasted’ on generating the global phase rather than the two-qubit phase and therefore a higher Rabi frequency is needed overall. For a same-species gate, if the ions have no hyperfine structure and are equally illuminated we can theoretically achieve optimal efficiency ($\Theta = 1$). The loss in gate efficiency due to the asymmetric matrix elements of the $^{43}\text{Ca}^+$ stretch qubit is small. However, the loss in efficiency caused by asymmetric Rabi frequencies on $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ in a mixed-species gate is more significant (in our case $\approx 20\%$ of the total acquired geometric phase gives rise to a global phase).

The expression for gate efficiency is equally applicable to a mixed-species MS gate, where achieving optimal efficiency relies on balancing the Rabi frequencies on the two species. Since the different species are driven by different beams, this is more technically demanding than for a same-species MS gate. Mixed-species gates require higher overall Rabi frequency, and the maximum excursion in phase space becomes larger. This increases errors due to heating, photon scattering, and imperfect closure of loops in phase space, compared with more efficient same-species gates.

3.3.5 Two-qubit gate errors and mitigation

Error sources for geometric phase gates are modelled in [11, 150] and discussed in [177]. This section will summarise the main contributions, with a focus on their manifestations in mixed-species gates.

3.3.5.1 Photon-scattering errors

The only fundamental source of error in laser-driven two-qubit gates is photon scattering. As discussed in section 3.2.4.2, after excitation to a P-state via a Raman beam, the excited state may spontaneously decay, emitting a photon. This may cause a loss in coherence of the qubit states, leading to a gate error.

The scattering rates are $\Gamma \propto 1/\Delta^2$ whereas the Rabi frequency used in the gate drive is $\Omega \propto 1/\Delta$. Therefore, in order to reduce photon-scattering errors at the same Rabi frequency it is always advantageous to increase the Raman detuning Δ and increase the laser power. In reality, the amount of laser power available will limit this trade-off.

If we wish to drive a mixed-species $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ gate using shared Raman beams with a wavelength between their respective S $\leftrightarrow$ P transition wavelengths, of course, increasing $|\Delta_{\text{Ca}}|$ decreases $|\Delta_{\text{Sr}}|$, and vice versa, imposing another limit on the scattering error. Fortunately, we find the minimum scattering error achievable for reasonable Rabi frequencies on each ion is still $\sim 10^{-4}$ for this combination of species (see figure 6.7). Reduced gate efficiency means that for a gate of the same duration $t_{\text{tot}} = Kt_g$, higher Rabi frequencies are required, so the photon-scattering error for fixed Raman detuning is increased.

For a gate driven directly rather than using Raman beams (e.g. an MS gate driven by a 674 nm laser on the quadrupole transition in $^{88}\text{Sr}^+$), the equivalent scattering error arises slightly differently. In this case we are concerned with spontaneous decay of the upper qubit state, which has a much shorter lifetime for an optical transition compared to a microwave/RF transition (391 ms in our case). Although this lifetime is clearly unaffected by the laser power, the shortest achievable gate duration — and therefore the probability of spontaneous decay — decreases with increased Rabi frequency. Reduced gate efficiency means that for a fixed Rabi frequency a longer gate time is required, again increasing this error.

3.3.5.2 Thermal errors

Motional amplitude heating

Trapped ions are subject to motional heating. If heating of the gate mode occurs over the duration of the gate, it causes unwanted excursions in motional phase space which lead to gate errors [150]

$$\epsilon_h = \frac{\dot{n}_m t_{\text{tot}}}{2K\Theta}, \quad (3.103)$$

where $\dot{n}_m$ is the heating rate for the gate mode m . For a fixed total gate duration $t_{\text{tot}} = Kt_g$, a smaller number of loops K increases the average excursion in phase space. This means there is a higher level of spin-motion entanglement, and therefore a larger gate error. For mixed-species gates with lower gate efficiency Θ , the gate error increases for the same reason.

For same-species crystals, ip modes are perfectly symmetric and oop modes are perfectly antisymmetric (see section 2.1). Therefore, common-mode noise does not couple to oop modes, leading to much smaller heating rates and gate errors. For mixed-species crystals all modes have some symmetry so the common-mode noise rejection is not perfect, leading to increased oop mode heating rates. For $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystals, the similar masses mean that the ip mode is still significantly more symmetric than the oop mode and therefore more susceptible to heating errors.

A related effect is dephasing of the motional state; similarly to elastic dephasing of the spin state due to Rayleigh scattering, some heating effects may not change the motional state but may instead cause decoherence. For this work the motional dephasing rate is dominated by amplitude heating, and estimates of the motional dephasing error are included in the heating error.

Outside the Lamb–Dicke regime

In equation (3.71) and equation (3.88) for the MS and LS gate Hamiltonians we assumed we were in the Lamb–Dicke limit. If any mode has a sufficiently high thermal

occupation number $\bar{n}_m$, this may not be the case. If a spectator mode m' (not the gate mode) is excited, the amplitude of the ion's motion is increased so that it explores a larger range of the gate beam, leading to a reduction in the average field amplitude it experiences. The effect is a reduction in the carrier Rabi frequency, as we found in equation (3.33), which depends on $n_{m'}$. For a Fock state $|n_{m'}\rangle$, in principle we could compensate for this reduction by increasing the laser power, but in reality the spectator mode is likely to be in a thermal state with an average occupation number $\bar{n}_{m'}$, leading to a spread in carrier Rabi frequencies [188]

$$\frac{\Delta\Omega_R}{\Omega_R} \approx \sqrt{\eta_{m'}^4 \bar{n}_{m'} (\bar{n}_{m'} + 1)}, \quad (3.104)$$

which cannot be corrected and leads to a gate error $\approx \frac{1}{4}\pi^2(\Delta\Omega_R/\Omega_R)^2$ [150]. In this work, since the Raman beams are ideally aligned to null η for the radial modes, cooling of the radial modes is not necessary for gates using these beams. Cooling of the spectator axial mode to close to its ground state is required to reduce this error.

The gate error due to higher initial thermal occupation of the gate mode m is found to be similar [11], although the details are more complex since this mode is displaced in phase space as part of the gate mechanism.

Kerr cross-coupling

Although the motional modes of trapped ions are usually treated as independent, excitation of modes which involve asymmetric motion leads to a change in the mean separation between the ions. This affects the frequencies of other asymmetric modes, leading to a ‘cross-coupling’ interaction of the form $\chi a_m^\dagger a_m a_{m'}^\dagger a_{m'}$. For a same-species crystal, an expression for χ can be found in [131], which shows that to reduce this coupling we should carefully choose trap parameters to avoid the resonance condition $\omega_m = 2\omega_{m'}$.

Even if sensible trap parameters are chosen, this coupling may still not be negli-

gible, in which case thermal occupation of an asymmetric mode m' will change the frequency of an asymmetric gate mode m . Similarly to the Lamb–Dicke error, if mode m' were in a known motional Fock state this would cause a constant frequency offset which could be taken into account. However, a thermal state leads to a spread in mode frequencies and therefore to an effective spread in the gate detuning, causing an error.

In same-species crystals, Kerr cross-coupling does not occur for the ip modes because they are completely symmetric. In order to perform a gate on the axial oop mode with a small Kerr cross-coupling error, we must avoid $\omega_{\text{ax},\text{oop}} = 2\omega_{\text{rad},\text{oop}}$.

In mixed-species crystals there is some coupling between all modes. Again, due to the similar masses of $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$, in this work the ip modes are still significantly more symmetric than the oop modes and are therefore much less susceptible to Kerr cross-coupling errors. To a good approximation, we can still reduce these errors for the axial oop mode by avoiding the same resonance condition with the radial oop modes as for same-species crystals.

3.3.5.3 Off-resonant excitation errors

Carrier coupling

Another error source is unwanted off-resonant coupling of the gate beams to the carrier transition. The nature of this coupling is different for the two types of gate.

In the MS gate, the carrier couplings of the equal intensity red and blue components may be viewed as a single beam that is resonant with the carrier transition but amplitude-modulated with frequency $\omega_m + \delta_g$ (equation (3.85)). The strength of this coupling is enhanced by a factor $1/\eta_{i,m}$ relative to the sideband couplings which drive the gate and, without mitigation, it will cause significant oscillations on top of the desired gate dynamics. Unless the laser pulse durations are carefully controlled to within a small fraction of the motional period $2\pi/\omega_m$, the fidelity of the gate op-

eration will be reduced. For $\phi_{M,i} = 0$, the average reduction in fidelity for a two-ion gate used to generate the Bell state $\frac{1}{\sqrt{2}}(|\downarrow_1\downarrow_2\rangle - i|\uparrow_1\uparrow_2\rangle)$ is calculated in [163] and is equal to

$$\frac{1}{2} \sum_i \left(\frac{\Omega_{R,i}}{\omega_m + \delta_g} \right)^2. \quad (3.105)$$

This is a significant error ($\sim 1\%$ for the gate parameters in this thesis)¹⁶.

However, the effect of this term can be suppressed to an insignificant level ($\lesssim 10^{-5}$ for typical gate times) by amplitude-shaping the laser pulses such that $\Omega_R(t)$ rises smoothly over a few motional periods at the start of the pulse, and falls similarly at the end, leading to a much smaller Fourier component at the carrier frequency [94].

In the LS gate, the coupling to the carrier is in the form of an oscillating light shift of the carrier transition frequency. Since the LS gate operator commutes with σ_z -operations, the propagator becomes

$$U_{\text{LS}} \cdot \exp \left[i \frac{\theta_{\text{carr},1}}{2} \sigma_z^{(1)} \right] \cdot \exp \left[i \frac{\theta_{\text{carr},2}}{2} \sigma_z^{(2)} \right], \quad (3.106)$$

where $\sigma_z^{(1)} \equiv \sigma_z \otimes \mathbb{1}$, $\sigma_z^{(2)} \equiv \mathbb{1} \otimes \sigma_z$ and the extra qubit frequency splitting due to the carrier light shift, integrated over the gate duration, is

$$\begin{aligned} \theta_{\text{carr},i} &= (\Omega_{R,i,\uparrow} - \Omega_{R,i,\downarrow}) \int_0^{Kt_g} \cos(\Delta\phi_i - (\omega_m + \delta_g)t) \\ &= \left(\frac{\Omega_{R,i,\uparrow} - \Omega_{R,i,\downarrow}}{\omega_m + \delta_g} \right) \left[\cos(\Delta\phi_i) - \cos(\Delta\phi_i - (\omega_m + \delta_g)Kt_g) \right]. \end{aligned} \quad (3.107)$$

Its effect is to add random single-qubit phases $\theta_{\text{carr},i}$ (dependent on $\Delta\phi_i$) for each ion. With rectangular laser pulses (constant Ω_R), the gate error caused by this term for

¹⁶For $\phi_{M,i} \neq 0$, the average error is larger [94].

generating a maximally entangled state is [150]

$$\begin{aligned}\epsilon_{\text{carr}} &= \frac{1}{4} (\langle \theta_{\text{carr},1} \rangle^2 + \langle \theta_{\text{carr},2} \rangle^2) \\ &= \frac{1}{2} \left[\sum_i \left(\frac{\Omega_{R,i,\uparrow} - \Omega_{R,i,\downarrow}}{\omega_m + \delta_g} \right)^2 \right] \sin^2 \left(\frac{(\omega_m + \delta_g) K t_g}{2} \right),\end{aligned}\tag{3.108}$$

averaging over a uniform distribution of possible $\Delta\phi_i$ due to instability in the relative laser phases. Again, the size of this error is $\sim 1\%$ for the gate parameters in this work, but again it can be reduced to a negligible level by amplitude-shaping the laser pulses over a few motional periods [11].

Coupling to spectator modes

The gate beams may also couple off-resonantly to other ‘spectator’ motional modes than the gate mode, driving loops in phase space which are, in general, not closed at the end of the gate. Errors resulting from this are smaller than carrier coupling errors [11], not least because they are suppressed by a factor $\eta_{i,m}$. Therefore these, too, may be neglected if pulse shaping is used.

3.3.5.4 Errors due to mis-set gate parameters

Imprecise setting of the gate time t_{tot} , the gate detuning δ_g , or the Rabi frequencies will also cause errors. An incorrect Rabi frequency or δ_g will change the size of the loops in phase space and therefore the geometric phase acquired. An incorrect t_{tot} or δ_g will result in imperfect closure of phase-space loops, leading to an error due to residual spin-motion entanglement. Errors due to imperfect loop closure are larger for less efficient gates because the loops are necessarily larger, so the residual displacement in phase space is larger for the same relative error in t_{tot} or δ_g .

Walsh modulation

To reduce sensitivity of the gate fidelity to imperfect loop closures, it is advantageous to perform a two-loop gate where the phase of the driving force is flipped by π at the beginning of the second pulse, relative to the phase at the beginning of the first pulse. This is known as a first-order Walsh modulation [70].

For perfect gate times at which we complete two loops exactly, it has no effect on the final phase-space displacement. The advantage becomes apparent for mis-set gate times, where we now always return to the initial location in phase space¹⁷, leaving no residual spin-motion entanglement (see figure 3.9).




Figure 3.9: Phase-space trajectories for a particular spin state during a two-loop gate using first-order Walsh modulation. The trajectories begin and end at the origin and the phase of the driving force is flipped by π at the marked points. Shown for fractional offsets $\Delta t/t$ from the ideal gate time of 0.7 (blue), 0.2 (red) and 0 (green).

3.3.5.5 Errors due to qubit frequency offsets

During the gate operation, the qubit frequency may be offset from the nominal value, for example due to residual single-beam light shifts or slow drifts in the external magnetic-field strength. In the interaction picture (rotating at the expected qubit

¹⁷If the spins are also flipped in between the two pulses, for a gate where the forces on opposite spin states are not equal in magnitude, the loops do not return to the initial location. However, the loops for all four spin states end at the same point in phase space, meaning the gate is still robust to errors in t_{tot} (see [177]).

frequency), frequency offsets δ_i^q on each qubit i add an extra term to the Hamiltonian during the gate,

$$H_{\text{err}} = \frac{\hbar}{2} [\delta_1^q \sigma_z^{(1)} + \delta_2^q \sigma_z^{(2)}]. \quad (3.109)$$

For the LS gate, where the gate drive is independent of the absolute qubit frequencies, H_{err} commutes with the gate Hamiltonian ($\sim \sigma_z \otimes \sigma_z$), and a simple spin-echo pulse can protect against these offsets.

The MS gate drive depends directly on the qubit frequencies and is therefore less intrinsically robust to offsets; H_{err} does not commute with the $\sigma_\phi \otimes \sigma_\phi$ operator. If the frequency offsets are small compared to the gate detuning, the gate operator at the gate time t_g may be approximated as

$$U_{\text{MS}}^\delta \approx e^{-i\pi/4} \exp \left[-i \left(\frac{\delta_1^q}{2} \sigma_z^{(1)} + \frac{\delta_2^q}{2} \sigma_z^{(2)} \right) t_g \pm i \frac{\pi}{4} (\sigma_{\phi_{S,1}+\pi/2} \otimes \sigma_{\phi_{S,2}+\pi/2}) \right]$$

$$= e^{-i\pi/4} \exp \left[-i \begin{pmatrix} (\delta_1^q + \delta_2^q) \frac{t_g}{2} & 0 & 0 & \pm \frac{\pi}{4} e^{-i(\phi_{S,1} + \phi_{S,2})} \\ 0 & -(\delta_2^q - \delta_1^q) \frac{t_g}{2} & \mp \frac{\pi}{4} e^{i(\phi_{S,2} - \phi_{S,1})} & 0 \\ 0 & \mp \frac{\pi}{4} e^{-i(\phi_{S,2} - \phi_{S,1})} & (\delta_2^q - \delta_1^q) \frac{t_g}{2} & 0 \\ \pm \frac{\pi}{4} e^{i(\phi_{S,1} + \phi_{S,2})} & 0 & 0 & -(\delta_1^q + \delta_2^q) \frac{t_g}{2} \end{pmatrix} \right]. \quad (3.110)$$

As noted in [108], a change in the average qubit frequency ($\delta_1^q + \delta_2^q \neq 0$) changes the energy gap between $|\uparrow_1 \uparrow_2\rangle$ and $|\downarrow_1 \downarrow_2\rangle$ but leaves $|\downarrow_1 \uparrow_2\rangle \leftrightarrow |\uparrow_1 \downarrow_2\rangle$ unchanged, whilst a change in the frequency difference between the two qubits ($\delta_1^q - \delta_2^q \neq 0$) does the opposite. The lowest-order effect of the offsets on the even-parity input states, $|\uparrow_1 \uparrow_2\rangle$ and $|\downarrow_1 \downarrow_2\rangle$, is to change the phase of the output Bell state if $(\delta_1^q + \delta_2^q)t_g \neq 0$. For $|\downarrow_1 \uparrow_2\rangle$ and $|\uparrow_1 \downarrow_2\rangle$, the phase becomes dependent on $(\delta_1^q - \delta_2^q)t_g$. For larger frequency offsets, phase-space loop closure is also affected and a high-fidelity Bell state can no longer be generated. More complex dynamical decoupling techniques, requiring additional drive frequencies applied continuously during the gate, may be used to protect MS

gates from qubit frequency offset errors [68, 17, 166].

For same-species gates, the frequency offsets on each qubit will be similar. Small differences in δ_1^q and δ_2^q may arise, for example, due to slightly different magnetic-field strengths at each ion or due to unequal illumination of the ions in the gate beams. Qubits in different ion species, however, will in general have different sensitivities to noise sources and can therefore have significantly different frequency offsets.

3.4 Randomised benchmarking

Quantum computing relies on the implementation of unitary operations. Benchmarking how well these operations can be implemented in a particular experimental system, and characterising the underlying noise processes which cause errors in their implementation, are therefore important tasks.

Consider a unitary operation $\hat{U}(\rho) = U\rho U^\dagger$ which acts on a quantum state represented by a density matrix ρ . In order to measure how well this operation can be implemented in an experimental system, one could prepare the system in a known input state ρ_{in} , perform the unitary operation as well as possible, and measure the resulting state ρ_{out} in a particular basis. Repeating the process with measurement in different bases allows us to acquire enough information to construct an estimate of the density matrix ρ_{out} . If the target outcome state ρ_{target} of the unitary $\hat{U}(\rho)$ is known, one can calculate the Uhlmann-Jozsa fidelity of the operation [87]

$$F(\rho_{\text{target}}, \rho_{\text{out}}) = \left[\text{Tr} \left(\sqrt{\sqrt{\rho_{\text{target}}} \rho_{\text{out}} \sqrt{\rho_{\text{target}}}} \right) \right]^2, \quad (3.111)$$

and attempt to correct for known state-preparation and measurement (SPAM) errors. The process of characterising the output density matrix is known as *quantum state tomography* (QST).

QST will only reveal how well the process $\hat{U}$ can be implemented for a particular

input state. A better measure is *quantum process tomography* (QPT) [141, 39], where QST is carried out after the gate process is applied to a set of basis input states, in order to construct the experimentally-implemented superoperator. However, there are some problems with this process. Firstly, the number of measurements required grows exponentially with the number of qubits, making this method intrinsically unscalable to larger systems. In addition, the quantum gates which are needed to prepare the initial states and to measure in the required bases are often not themselves well characterised, and this set can include $\hat{U}$ itself.

Over the last 15 years there has been increasing interest in methods of characterising the degree of coherent control over experimental systems in a scalable manner. In 2005, Emerson, Alicki and Życzkowski [53] proposed a scalable method for noise estimation based on applying unitary operations randomly chosen from the Haar measure (a uniform probability distribution over all unitary operators). This has since evolved into the procedure known as *randomised benchmarking*, variants of which are detailed below.

3.4.1 Standard randomised benchmarking

3.4.1.1 Twirling

The standard randomised benchmarking (SRB) protocol [113, 112] is a method of measuring an average error per operation for a particular set of operations, and is based on the concept of ‘twirling’ a quantum channel [16].

For a channel $\hat{\Lambda}(\rho)$ which acts on an input density matrix ρ , twirling of the channel involves finding the average effect $\hat{\Lambda}^{\text{ave}}(\rho)$ of the composition $(\hat{U} \circ \hat{\Lambda} \circ \hat{U}^\dagger)(\rho)$, where $\hat{U}(\rho) = U\rho U^\dagger$, over unitary operations $\hat{U}$ chosen according to some probability distribution $\mu(\hat{U})$:

$$\hat{\Lambda}^{\text{ave}}(\rho) = \int d\mu(\hat{U}) (\hat{U} \circ \hat{\Lambda} \circ \hat{U}^\dagger)(\rho) = \int d\mu(\hat{U}) U^\dagger \hat{\Lambda}(U\rho U^\dagger) U. \quad (3.112)$$

For a system of n qubits, we can theoretically choose to sample uniformly from the continuous group $\hat{U}(D)$ of unitaries of dimension $D = 2^n$ (the Haar measure). As might be expected due to the high degree of symmetry in this process, the resulting averaged channel is simply a depolarising channel [53]:

$$\bar{\Lambda}(\rho) = p\rho + (1-p)\frac{\mathbb{1}}{2^n}. \quad (3.113)$$

This means that the average effect of $\hat{\Lambda}(\rho)$ is to preserve ρ with probability p , or to ‘depolarise’ ρ into the maximally mixed state, with probability $(1-p)$. If $\hat{\Lambda}(\rho)$ represents a ‘noisy identity’ operation, i.e. one which ideally implements an identity operation but is subject to noise which introduces errors, then the fidelity of this identity operation is completely characterised by the parameter p , known as the ‘depolarising parameter’.

3.4.1.2 The Clifford group

Experimentally, sampling from $\hat{U}(D)$ is inefficient, so we would like to restrict our choice of unitaries to some discrete group which has the property that sampling uniformly from this group is operationally equivalent to sampling from the Haar measure. Such a group is known as a unitary two-design [46], and one example is the Clifford group.

The n -qubit Clifford group is defined as the ‘normaliser’ of the n -qubit Pauli group P_n in $\hat{U}(2^n)$. This is the group of unitaries which transform unitaries made of n -fold tensor products of the single-qubit Pauli matrices $P_1 = \{\mathbb{1} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}, \sigma_x = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \sigma_y = \begin{pmatrix} 0 & i \\ -i & 0 \end{pmatrix}, \sigma_z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}\}$ into other n -fold tensor products of these matrices (not including $\mathbb{1}_n$).

To offer a visual example, the single-qubit Clifford group permutes the axes $\{\pm x, \pm y \pm z\}$ on the Bloch sphere. There are 6 choices of where to fix the z -axis, then

there is a rotational symmetry of 4 about this axis, so 4 options for where to fix the x -axis. Then the structure is completely specified. Therefore the single-qubit Clifford group has $6 \times 4 = 24$ elements. The two-qubit Clifford group has 11,520 elements.

3.4.1.3 Standard randomised benchmarking

The *standard randomised benchmarking* (SRB) protocol provides a measure of the error per Clifford operation in a particular quantum system.

In the SRB protocol, we perform random sequences of Clifford unitaries, or ‘gates’, sampled uniformly from the Clifford group, followed by a final inverting gate C^{-1} whose desired effect is to undo the combined effect of the previous ones. Since the Clifford gates form a group, this final inverting gate will also be a Clifford gate. This process can be viewed as an extension of the Clifford twirling process of a noisy identity operation, so we form a depolarising channel and we can characterise the fidelity of the Clifford operations using a depolarising parameter, p_C .

The full protocol is as follows [113, 112]:

1. Generate a sequence of m Clifford gates, sampled randomly from the Clifford group, and calculate a final ‘inverting’ gate C^{-1} such that the overall effect of the sequence (in the absence of errors) is an identity operation.
2. Prepare the system in a known initial state, perform this sequence and measure the result. Repeat N times and calculate the probability of ending up back in the initial state (i.e. the sequence fidelity).
3. Repeat steps 1 and 2 for k different random sequences of the same length m , and calculate an average sequence fidelity $F(m)$ over different randomisations for that sequence length.
4. Repeat for different sequence lengths m .

The variation of fidelity with sequence length $F(m)$ can then be fitted to an appropriate decay model and an error ϵ_C per Clifford operation can be extracted. An example SRB sequence is shown in figure 3.10.




Figure 3.10: Example SRB sequence of m random Cliffords followed by a final inverting Clifford C^{-1} which, in the absence of errors, returns the system to its initial state $|\psi\rangle$.

The probability that a particular part of the SRB sequence does not cause depolarisation is given by a depolarising parameter for that part of the sequence. We can call the probability of not depolarising during the SPAM and inversion processes $p_{S/I}$, and the probability of not depolarising during each of the m Clifford gates p_C . Therefore, the probability of not depolarising during an SRB sequence of length m is given by

$$P(m) = p_{S/I} \times p_C^m. \quad (3.114)$$

If depolarisation does occur during the sequence, and the state becomes the maximally mixed state, it may not necessarily cause us to measure the wrong outcome at the end of the sequence. For example, in the single-qubit maximally mixed state $\mathbb{1}_1 = \frac{1}{2}(|\uparrow\rangle\langle\uparrow| \otimes |\downarrow\rangle\langle\downarrow|)$, there is only a 50% chance of measuring the wrong outcome even if the identity operation failed. For a system of n qubits, the probability of measuring the wrong outcome, given a depolarisation probability $1 - P(m)$, is equal to

$$E(m) = \frac{1}{\alpha_n} (1 - P(m)), \quad (3.115)$$

where

$$\alpha_n = \frac{2^n}{2^n - 1}. \quad (3.116)$$

In the SRB protocol, we measure the sequence fidelities $F(m)$, where

$$F(m) = 1 - E(m). \quad (3.117)$$

Combining equations equation (3.114) and equation (3.115) leads to

$$P(m) = 1 - \alpha_n E(m) = (1 - \alpha_n \epsilon_{S/I})(1 - \alpha_n \epsilon_C)^m, \quad (3.118)$$

where $\epsilon_{S/I} = (1 - p_{S/I})/\alpha_n$ is the error rate during the SPAM and inversion steps, and $\epsilon_C = (1 - p_C)/\alpha_n$ is the error per Clifford gate. Using equation (3.117), we find that the sequence fidelities we measure are given by

$$F(m) = \frac{1}{\alpha_n} (1 - \alpha_n \epsilon_{S/I})(1 - \alpha_n \epsilon_C)^m + \frac{1}{2^n}, \quad (3.119)$$

or

$$F(m) = A_n p_C^m + B_n, \quad (3.120)$$

where

$$A_n = \frac{1}{\alpha_n} (1 - \alpha_n \epsilon_{S/I}) \quad (3.121)$$

and

$$B_n = \frac{1}{2^n}. \quad (3.122)$$

Therefore by fitting the measured fidelities to the fidelity decay model (equation (3.120)) with parameters A_n , B_n and $p_C = 1 - \alpha_n \epsilon_C$, we can extract the error per Clifford gate ϵ_C .

The restriction to the Clifford group, rather than the full Haar measure, makes this a scalable method of characterising the degree of coherent control over a quantum system. The method scales polynomially with the number of qubits [113], rather than exponentially as for QPT. Unlike in QPT, errors due to SPAM and the inverting

gates are absorbed by the parameter A_n and kept separate from the estimate of ϵ_C . Additionally, due to the symmetry properties of the Clifford group, SRB can provide unbiased information about the underlying noise processes in a quantum system.

Benchmarking of Clifford gates is useful for many reasons. For example, universal quantum computation is possible using only the set of Clifford gates plus any additional unitary gate not in the Clifford group [[113], theorem 6.5 of [128]].

3.4.2 Interleaved randomised benchmarking

SRB allows us to benchmark the full Clifford group, but sometimes we are interested in the performance of one particular unitary operation. The technique of *interleaved randomised benchmarking* (IRB) allows estimation of the error ϵ_G per gate for a particular ‘gate-under-test’ G , as opposed to ϵ_C which is an average error per Clifford gate. The protocol was introduced in [114] and proceeds as follows:

1. Perform the standard (SRB) protocol and extract a depolarising parameter p_C for the Clifford group.
2. Repeat the SRB protocol, but with ‘interleaved’ sequences rather than standard sequences. To form an interleaved sequence of length m , first generate a random sequence of m Clifford gates sampled randomly from the Clifford group, then interleave the gate-under-test G after each Clifford gate C_i (see figure 3.11). In interleaved sequences, the final inverting gate $C^{-1,G}$ should invert the composition of the first m Cliffords *and* the interleaved G gates. Extract a depolarising parameter for the interleaved sequences p_C^{int} .
3. Calculate the error per gate-under-test

$$\epsilon_G = \frac{1}{\alpha_n} \left(1 - \frac{p_C^{\text{int}}}{p_C} \right). \quad (3.123)$$




Figure 3.11: Example pair of SRB and IRB sequences used in measuring the error per gate ϵ_G of the gate-under-test, G . This error is calculated by considering the difference in the fidelities of the two sequence types. The final gate $C^{-1,G}$ in the IRB sequence is chosen to invert the preceding composition of gates, including the interleaved G s.

3.4.2.1 Simplified method

Implementing the IRB protocol as described above would require fitting fidelity decay curves $F(m) = A_n p_C^m + B_n$ for the standard sequences, and $F^{\text{int}}(m) = A_n^{\text{int}} (p_C^{\text{int}})^m + B_n$ for the interleaved sequences, and extracting depolarising parameters in each case. For some of the applications in this work, this would necessitate collecting an impractically large amount of data.

To calculate the error per gate-under-test only requires knowledge of the ratio of the two depolarising parameters

$$p_C = \sqrt[m]{\frac{F(m) - B_n}{A_n}} \quad (3.124)$$

and

$$p_C^{\text{int}} = \sqrt[m]{\frac{F^{\text{int}}(m) - B_n}{A_n^{\text{int}}}}. \quad (3.125)$$

If the experiment can be controlled such that the parameters A_n and A_n^{int} are the same, then the ratio p_C^{int}/p_C may be calculated from the fidelities $F(m)$ and $F^{\text{int}}(m)$ for just *one* sequence length m . The error per gate may be extracted via

$$\epsilon_G = \frac{1}{\alpha_n} \left(1 - \sqrt[m]{\frac{F^{\text{int}}(m) - B_n}{F(m) - B_n}} \right), \quad (3.126)$$

without fitting two full decay curves.

Since A_n depends on the error contribution from SPAM and the inverting Clifford, in order to make use of this simplification it is necessary to ensure that these errors are equal in size for the standard and the interleaved sequences.

Earlier variations of the standard randomised benchmarking protocol were implemented experimentally in a wide range of qubit technologies including ions [98, 19, 67], superconducting qubits [52, 38], neutral atoms [136] and NMR [148].

Chapter 4

Apparatus

The work in this thesis was carried out in two distinct experimental setups, which will be referred to as ‘apparatus A’ and ‘apparatus B’. They were both designed and built by previous Oxford ion trappers and are therefore described in detail elsewhere; this chapter provides an overview of each one. The two setups have many similarities, both having been designed to trap $^{43}\text{Ca}^+$ for quantum computing applications. A schematic representation (relevant to both) can be found at the end of the chapter, in figure 4.5.

Both setups include a linear Paul trap, housed in an octagonal vacuum system with six viewports for optical access. A magnetic field of 146 G is generated by water-cooled coils, in order to give access to the $^{43}\text{Ca}^+$ clock qubit. An imaging system, featuring a photomultiplier tube (PMT) and an electron-multiplying charge-coupled device (EMCCD) camera for fluorescence detection, sits above each vacuum chamber. Lasers are used to cool and read out the ions (and for coherent manipulations in apparatus B). Microwaves are used for coherent drive of hyperfine transitions in $^{43}\text{Ca}^+$, and apparatus B includes additional RF drive for the $^{88}\text{Sr}^+$ Zeeman qubit. Both setups operate at room temperature.

The most significant differences are the designs of the traps and the ion species

used in each; apparatus A features a microfabricated planar trap [2] with integrated microwave electrodes, and was used in this work for memory benchmarking of a single $^{43}\text{Ca}^+$ ion (chapter 5). The setup is identical to that described in [176], with further details in [155, 66, 1].

In apparatus B is a three-dimensional, macro-scale ‘blade’ trap, used in this work for mixed-species entanglement of a $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal (chapter 6). The setup is described in detail in [177, 150]. Changes were made to the Raman AOM network (section 4.7.2) to facilitate the drive of a Mølmer–Sørensen gate using these beams.

4.1 Traps

4.1.1 Apparatus A: planar trap

The trap in apparatus A is a microfabricated ‘surface-electrode’ trap with integrated microwave circuitry, depicted in figure 4.1. It was designed by D. T. C. Allcock, and fabricated and characterised by D. T. C. Allcock and T. P. Harty [1, 2, 66]. In a trap such as this, the trapping electrodes are not arranged in a three-dimensional quadrupolar shape but instead lie in one plane, and reproduce a quadrupolar trapping potential on a smaller scale; this ‘on-chip’ geometry allows the use of microfabrication techniques for trap construction, which is a promising option for scalability. It also allows trapping of ions much closer to the electrodes, leading to larger magnetic-field gradients for driving two-qubit entangling gates with near-field microwaves [138], which was the main purpose of this trap [155, 68].

The surface trap has a ‘five-wire’ geometry [36]. An RF signal with an amplitude of 72 V and a trapping frequency $\Omega_T = 2\pi \times 38.7$ MHz is applied to the RF electrodes, to create a radial quadrupolar potential in the $\hat{\mathbf{x}}$ - and $\hat{\mathbf{y}}$ -directions. Axial confinement along $\hat{\mathbf{z}}$ is provided by DC voltages of $\lesssim 10$ V, applied to the central three electrodes as well as to six further DC control electrodes. These DC voltages also allow for




Figure 4.1: Planar ion trap used in apparatus A. Figure adapted from [155]. Gold is electroplated onto a sapphire substrate to form trap electrodes which double up as half-wave microwave resonators at 3.2 GHz. They are 50 Ω impedance-matched via quarter-wave couplers to the microwave drive circuitry, which is connected via in-vacuum SMAs. DC electrodes tilt the radial trapping field by $\approx 40^\circ$ from $\{\hat{x}, \hat{y}\}$ to $\{\hat{x}', \hat{y}'\}$ ($\hat{y}'$ not shown for clarity). The static magnetic field $B_0 = 146$ G is oriented along $\hat{x}$.

micro-motion compensation in all three directions, and generate an additional, static quadrupole in the radial plane. The purpose of the extra quadrupole is to rotate the principal axes of the radial trapping field away from the $\hat{\mathbf{x}}$ - and $\hat{\mathbf{y}}$ -directions (by $\approx 40^\circ$) and to lift the degeneracy of the two radial modes. This allows cooling of all three motional modes of a $^{43}\text{Ca}^+$ ion using a 397 nm beam which propagates parallel to the trap surface, avoiding scattering of blue light off the surface. Scattered light would otherwise result in a large background count on the PMT and thus a loss of fidelity during ion readout, and could lead to undesirable charging of the trap dielectrics.

For a single $^{43}\text{Ca}^+$ ion, the trapping fields result in an axial frequency $\omega_z \approx 2\pi \times 500$ kHz, and radial frequencies $\omega_{x',y'} \approx 2\pi \times 3$ MHz with a splitting of 330 kHz between them. The estimated trap depth is 59 meV. Ions sit at a distance $d = 75 \mu\text{m}$ above the central electrode. This distance was chosen to allow large magnetic-field gradients at the ion whilst limiting anomalous heating; the origin of anomalous heating in ion traps is poorly understood, but it scales approximately as $1/d^4$ [154, 47, 26].

Microwaves, for driving hyperfine transitions, are applied to the trap chip via in-vacuum SMA cables and 50Ω coplanar waveguides. These are impedance-matched via quarter-wave coupling elements to the three central electrodes, which are microwave half-wave cavities ($Q \approx 5$), resonant at 3.2 GHz (the frequency of the clock qubit in $^{43}\text{Ca}^+$). In previous work, microwaves were applied to the central electrode in order to drive a $^{43}\text{Ca}^+ - ^{43}\text{Ca}^+$ entangling gate (with a fidelity of 99.7(1)% [155, 68]). In this work, microwaves are applied only to one of the outer electrodes, for single-qubit operations.

The trap was fabricated in-house; details of the process can be found in [1]. Gold electrodes ($\sim 5 \mu\text{m}$) were electroplated onto a sapphire substrate (0.5 mm). Sapphire was chosen because its high thermal conductivity allows dissipation of large (~ 5 W) microwave powers, and its reasonably high relative permittivity leads to smaller dimensions of the microwave structures.

Diplexer

Each of the three central trapping electrodes must be connected to the trapping RF voltage (or ground), a DC bias voltage, and the microwave drive. To achieve this, these signals are combined in a custom filter network (diplexer), described in detail in [2, 1].

Trap RF

The 38.7 MHz trapping RF signal is generated by a synthesiser¹, and amplitude-stabilised using a custom circuit designed by T. P. Harty and detailed in [66]. It is then passed through an amplifier² and stepped up by a factor of 9.6 in voltage by a toroidal resonator³, before passing to the diplexer and on to the trap electrodes.

Trap DC

The trap DC voltages are generated by a digital-to-analogue converter (DAC)⁴ and passed through a low-pass filter. The voltages are applied to the three central trap electrodes via the diplexer and to the other six DC electrodes via feedthroughs in the base of the vacuum system.

Heating rate

The heating rate in this trap was previously measured to be $\dot{n}_{\text{rad}} = 27(2)$ quanta/s for the 3.268 MHz radial mode of a single $^{43}\text{Ca}^+$ ion [155]. This is an unusually low heating rate for a room-temperature surface trap.

¹HP/Keysight/Agilent 8656B.

²Spectrum Microwave/API Tech TM3310.

³Micrometals T80-10: Q-factor 65, resonance frequency 38.7 MHz.

⁴Analog devices AD5372 DAC chip on an EVAL-ADUM4160 evaluation board.




Figure 4.2: CAD model of the blade trap in apparatus B; figure adapted from [189]. The endcap electrodes are shown in red. Oscillating RF voltages are applied to the blue blade electrodes. The yellow blades (only one shown) are grounded. The three compensation electrodes are shown in purple. The static magnetic-field direction is in the horizontal plane, at 45° to the $\hat{z}$ -axis.

4.1.2 Apparatus B: blade trap

The trap in apparatus B is a much larger, three-dimensional trap with four ‘blade’-shaped electrodes and two cylindrical endcap electrodes. It was designed by S. R. Woodrow [189] and built by K. Thirumalai [177], and is depicted in figure 4.2.

DC voltages of up to 200 V are applied to the endcap electrodes to provide axial confinement. Two of the blade electrodes are grounded, whilst oscillating RF voltages at frequency $\Omega_T = 2\pi \times 28.0$ MHz with amplitudes up to ~ 1 kV are applied to the other two, to form the two-dimensional radial quadrupole potential. Driving the blade pairs asymmetrically (as opposed to applying equal voltages with a 180° phase difference) removes the technical requirement of stabilising the relative phase between the two pairs, which could otherwise lead to uncompensatable micro-motion. The asymmetric drive, along with the endcap potentials, lifts the degeneracy of the radial modes along $\hat{x}$ and $\hat{y}$ ($\alpha_x \neq \alpha_y$ in equation (2.1)) and results in a small contribution

to the axial confinement from the RF potential ($\alpha_z \neq 0$) [11].

Three compensation electrodes allow adjustment of the ion's position to minimise micro-motion in the horizontal and vertical directions. Axial compensation may be adjusted via the endcap voltages.

Microwave and RF drive for coherent qubit manipulations may be applied either to the grounded blade electrodes or to an in-vacuum antenna mounted to the base of the vacuum system.

Trap RF

The 28.0 MHz trapping RF is generated via a synthesiser⁵ and passes to a variable attenuator⁶ followed by a (fixed) +46 dB power amplifier⁷ and a copper helical resonator [159] with one end connected to two of the trap blades and its ground shared with the remaining two blades. The resonant frequency is 28 MHz, at which the voltage is stepped up by a factor of 35.

The variable attenuator allows us to change the trap depth for loading, recrystallisation, re-ordering (see section 6.1.1) and normal operation. Under normal operating parameters, the trap frequencies for a single $^{43}\text{Ca}^+$ ion are $\omega_z = 2\pi \times 1.92$ MHz, $\omega_{x,y} \approx 2\pi \times 4$ MHz with a splitting of ≈ 300 kHz. The trap frequencies for $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ are given in figure 4.3. The trap depth for the blade trap is approximately 100 times that of the planar trap (~ 1 eV).

Trap DC

The trap DC voltages are supplied by a custom DAC designed by C. J. Ballance [11]. They pass through a filter board to feedthroughs in the vacuum system and on to the endcaps and compensation electrodes. An AC ‘tickle’ voltage, generated

⁵HP/Keysight/Agilent 8656B.

⁶Mini-circuits ZX76-31R5-PP-S+, 0 to 31.5 dB attenuation in 0.5 dB steps.

⁷Frankonia FLL-25.




Figure 4.3: $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ secular frequencies $f_m = \omega_m/2\pi$ at the trap voltages used for normal operation in apparatus B.

by a synthesiser⁸, may also be applied to one of the DC electrodes via an additional connection which bypasses the filter. This is used to deliberately excite resonant motion in order to determine the trap secular frequencies.

Heating rate

For a single $^{43}\text{Ca}^+$ ion in this trap, a heating rate of $\dot{n}_z = 75(10)$ quanta/s was measured in other work [150] for an axial motional mode frequency $\omega_z = 2\pi \times 1.92$ MHz. For a $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal, the heating rates of the axial ip and oop modes are $\dot{n}_{\text{ax ip}} \approx 110$ quanta/s and $\dot{n}_{\text{ax oop}} \approx 30$ quanta/s (measured in section 6.1.5).

4.2 Vacuum systems

Each trap sits in a ‘spherical octagon’-shaped vacuum chamber⁹ with viewports for optical access attached on six sides. A further viewport on the top is used for imaging. An ion pump and getter are attached to one side and maintain the pressure at $\lesssim 1 \times 10^{-11}$ mbar. Electrical feedthroughs on the remaining side and on the base flange are used for the trap RF, trap DC, microwaves, strontium RF (in apparatus B) and ovens. Part numbers for apparatus A can be found in [176, 155, 66] and for apparatus B in [177].

The ovens consist of ~ 10 mm long stainless steel tubes filled with granules of

⁸Keysight 33512B.

⁹Kimball MCF600-SphOct-FC28.

calcium and strontium. They are resistively heated with currents of ~ 5 A. The natural abundance of ^{43}Ca is only 0.14%, so we use an enriched sample¹⁰ of 12% ^{43}Ca and 88% ^{40}Ca . The sample of strontium contains different isotopes in their natural abundances, the most common being ^{88}Sr (82.6%), ^{86}Sr (9.9%) and ^{87}Sr (7.0%).

4.3 Imaging

Both systems collect ion fluorescence (on the 397 nm and 422 nm transitions) via a lens (numerical aperture 0.3 in apparatus A and 0.6 in apparatus B) situated above the top viewport. An aperture is placed in the focal plane of the lens to reduce background scatter. We choose between imaging using an EMCCD camera¹¹ or a PMT¹² by moving a beam splitter (apparatus A) or a mirror (apparatus B) into or out of the beam. A filter selects wavelengths below 440 nm. The camera is used for diagnostics and the PMT for normal operation, including ion readout.

In apparatus A, the PMT count rate for readout of a single $^{43}\text{Ca}^+$ ion is approximately 25 kHz above a background of 5 kHz, with a total photon collection efficiency of $\approx 0.2\%$. In apparatus B we adjust the focus onto the PMT to a ‘compromise’ position which allows collection of fluorescence at 397 nm and at 422 nm with count rates of approximately 200 kHz and 80 kHz respectively, above a background of 180 kHz. The collection efficiency is $\approx 1.5\%$.

¹⁰Supplied by Oak Ridge National Laboratory [1].

¹¹Electron-multiplying charge-coupled device camera. Apparatus A: Andor Technology Ltd, Luca DL-406M-OEM. Apparatus B: Andor Technology iXonEM+ DU-897E.

¹²Photomultiplier tube. Apparatus A: ET Enterprises Limited P25PC-12. Apparatus B: Hamamatsu H10682-02.

4.4 Magnetic field

4.4.1 Coils

The 146 G external magnetic field is generated primarily by a pair of water-cooled, rectangular copper coils in each setup. The coils in apparatus A are supplied¹³ with a total current of 135 A, and those in apparatus B with 56.8 A.

There are additional pairs of ‘trim’ coils in each setup to adjust the magnetic field more finely and in orthogonal directions, ensuring that it is well aligned with the propagation directions of σ -polarised laser beams. The trim coils are supplied¹⁴ with currents of ~ 1 A.

4.4.2 Field stabilisation

In each setup, the current supplying the coils is actively stabilised so that the field strength remains within 1 mG of the desired 146 G, using a circuit designed by B. Merkel. It is described in detail and characterised in [123] (using apparatus B) and [176] (by J. E. Tarlton, using apparatus A).

Feedback

A small amount of current is picked off from the main power supply and measured using a fluxgate sensor¹⁵. Feedback is provided by a transistor-controlled current shunt which diverts a small amount of current from the coils. In apparatus A, an additional sensor¹⁶ on the trap table measures the magnetic-field strength every two seconds and automatically adjusts the set point of the feedback loop, to correct for drifts on this timescale.

¹³Keysight 6671A power supply. Two in parallel in apparatus A. A single unit in apparatus B.

¹⁴Thurlby Thandar Instruments (TTi) Limited QL series power supplies.

¹⁵Apparatus A: LEM IT 200-S ULTRASTAB. Apparatus B: LEM IT 405-S ULTRASTAB.

¹⁶Bartington Instruments Mag670-I-1000.

Feedforward

A feedforward system acting on the coil current suppresses field noise synchronous with harmonics of the 50 Hz mains electricity cycle to a root-mean-square deviation of ~ 0.1 mG [123, 176]¹⁷. The feedforward is calibrated by measuring the mains noise using the ion as a probe.

Ion servo

We use the magnetic-field sensitivity of the stretch qubit transition of $^{43}\text{Ca}^+$ (or of the Zeeman qubit transition when occasionally working with a single $^{88}\text{Sr}^+$ ion) to provide an absolute reference for the field strength. Measurements of the frequency of this transition are taken every few minutes in an automated servo, and are used to adjust the set point of the feedback loop. The statistical precision of this servo is ± 0.04 mG in apparatus A (corresponding to 100 Hz on the stretch qubit frequency), and usually about twice that in apparatus B, depending on the settings chosen.

4.5 Reference oscillator

A rubidium atomic frequency standard¹⁸, shared between both setups, provides a 10 MHz reference which is used to clock direct digital synthesisers (DDSs) for microwave drive, strontium RF drive, and RF drive of AOMs for Raman and 674 nm beams.

4.6 Microwave and RF drive

Apparatus A

The Rb reference clocks a synthesiser¹⁹ which produces a ≈ 1 GHz signal and clocks

¹⁷In apparatus B, the feedforward tones are generated using ARTIQ’s Sinara hardware ‘Stabilizer’, implemented prior to work on Raman-driven MS gates.

¹⁸SRS FS725.

¹⁹Rohde & Schwarz SMA100A.

two DDS boards²⁰.

One DDS provides the clock qubit drive. It has eight different output profiles which are configured with different phases to perform rotations about different axes on the Bloch sphere. The output frequency on all profiles is ≈ 400 MHz, which is octupled to 3.2 GHz using a series of components described in [176]. The resolution of the synthesiser which provides the 1 GHz clock is 10 mHz, which translates to 32 mHz on the qubit frequency. The DDS resolution is 0.23 Hz, or 1.84 Hz on the qubit frequency. The exact output frequencies may be adjusted together, to correct for drifts with a better frequency resolution (2 mHz) than can be achieved by adjusting either setting individually (see [176]).

At the end of the octupler chain are two switches²¹. The first is operated in the usual manner but the second is operated in reverse, with its input connection continuing to the trap via the diplexer. Its second output connection is driven with a synthesiser²² at a frequency of 3.8 GHz and a power similar to that used to drive the qubit, in order to keep it consistently warm. It was found in [176] that this configuration significantly reduced transients on the microwave pulses which could otherwise introduce errors in memory benchmarking. The frequency of 3.8 GHz was chosen to be close enough to 3.2 GHz that the thermal behaviour is similar, but far enough detuned so as not to significantly drive the qubit transition.

The second DDS provides the drive for the other hyperfine transitions needed to prepare and read out the clock qubit. Its profiles are set to frequencies in the range 359 – 387 MHz, which are octupled via standard Mini-Circuits components, described in [66].

The π -times and frequencies of all the microwave transitions driven in this system are listed in table 4.1.

²⁰Analog Devices AD9910 Evaluation Board.

²¹Mini-Circuits ZASWA-2-50DR+.

²²Keysight N5181B.

transition	frequency (GHz)	π -time (μ s)
$ 4, +4\rangle \rightarrow 3, +3\rangle$ (stretch qubit)	2.874	1.5
$ 3, +3\rangle \rightarrow 4, +2\rangle$	2.986	3.5
$ 4, +2\rangle \rightarrow 3, +1\rangle$	3.095	1.4
$ 3, +1\rangle \rightarrow 4, 0\rangle$ (clock qubit)	3.200	18.5
$ 4, +3\rangle \rightarrow 3, +3\rangle$	2.931	2.5
$ 4, +2\rangle \rightarrow 3, +2\rangle$	3.042	3.9

Table 4.1: π -times and frequencies of $^{43}\text{Ca}^+$ microwave transitions driven in apparatus A. The first three transitions are used to prepare the clock qubit. The final two are used in microwave-enhanced optical pumping (see section 2.2.6). States are all in the $4S_{1/2}$ ground level and are labelled using $|F, m_F\rangle$.

Apparatus B

The Rb reference is used to clock a synthesiser²³ which produces a signal at 990 MHz. For the strontium RF drive, this clocks a DDS²⁴ which produces a 409 MHz tone that passes through a switch²⁵, then is amplified²⁶ to ≈ 10 W and sent to the trap.

For the calcium microwaves, it clocks a separate DDS²⁷, one channel of which generates frequencies of a few hundred MHz, depending on the hyperfine transition we wish to drive. The DDS output is mixed with a 2.77 GHz tone from a synthesiser²⁸ in a single-sideband source (described in [11]), passed through a switch²⁹ and amplified³⁰ to ≈ 10 W.

The RF and microwaves may each be applied either to a grounded trap blade or to the antenna. The π -times we achieve for relevant transitions in each case are listed in table 4.2.

²³HP 8656B.

²⁴Arduino-controlled Analog Devices AD9910 evaluation board.

²⁵Mini-Circuits ZASWA-2-50DR.

²⁶Mini-Circuits ZHL-5W-1.

²⁷Enterpoint Milldown: includes an FPGA + 4 Analog Devices AD9910 chips.

²⁸HP E4422B.

²⁹Mini-Circuits ZASWA-2-50DR.

³⁰Mini-Circuits ZHL-16W-43-S+.

drive	transition	π -time (μ s)
RF on antenna	Zeeman qubit	5
RF on blade	Zeeman qubit	18
MW on antenna	stretch qubit	4
	$ 3, +3\rangle \rightarrow 4, +2\rangle$	6.5
	$ 4, +2\rangle \rightarrow 3, +1\rangle$	53
	clock qubit	29
MW on blade	stretch qubit	3.5
	$ 3, +3\rangle \rightarrow 4, +2\rangle$	16
	$ 4, +2\rangle \rightarrow 3, +1\rangle$	11
	clock qubit	19

Table 4.2: π -times of $^{43}\text{Ca}^+$ microwave (MW) transitions and the $^{88}\text{Sr}^+$ RF transition driven in apparatus B.

4.7 Lasers

Across the two experiments we use a combination of diode lasers and Ti:Sapphire lasers. All lasers except the 375 nm photoionisation laser are monitored on a wavelength meter³¹ and an optical spectrum analyser³². Beams are switched using AOMs or shutters, and delivered to the traps via optical fibres so that realignment of preceding optics does not change the alignment onto the trap centre. Most are intensity-stabilised by measuring the beam power using a photodiode and feeding back on the RF drive amplitude for the switching AOM, and focused at the trap centres. The geometry of the beams at each trap is shown in figure 4.4. The intensity-stabilisation process is described in [1] and [11].

4.7.1 Diode lasers

Photoionisation lasers

The photoionisation lasers are shared between both systems. The 423 nm and 461

³¹High Finesse WS7.

³²TOPTICA FPI-100-0980-1 (red lasers), Thorlabs SA200-3B (blue lasers).

nm lasers³³ are external cavity diode lasers (ECDLs). The 375 nm laser³⁴ is a free-running diode, as there are no strict requirements on its frequency stability. They are switched using mechanical shutters.

Calcium lasers

Each setup includes five ECDLs³⁵ at {393, 397, 850, 854, 866} nm, located on separate optical tables from the traps. All except the 854 nm laser are frequency-stabilised to reference cavities³⁶. Diffraction gratings prevent amplified spontaneous emission at unwanted wavelengths in the 397 nm, 866 nm and (apparatus A only) 850 nm beams.

Each beam passes through an AOM³⁷ which is used to add a frequency offset, to stabilise the intensity of the beam via feedback on the amplitude of the RF drive, and to switch the beam on and off; only the diffracted beam is coupled into the fibre to the trap table. To improve extinction, the drive frequency of the AOM is also detuned when it is switched off. These AOMs are used in a double-pass configuration so that when the drive frequency is scanned, the fibre-coupling efficiency does not change.

Sidebands at ≈ 3 GHz are added to the 397 nm lasers via EOMs³⁸ in order to address both hyperfine manifolds of the $4S_{1/2}$ level. In apparatus A, this laser is split into a σ^+ -polarised beam and a ‘397 Doppler’ beam with polarisation $\frac{1}{3}\pi$ and $\frac{2}{3}\sigma^+$. For the 397 Doppler beam, the 400 MHz red-detuned zeroth diffracted order of the corresponding AOM is also coupled into the fibre. A mechanical shutter blocks this order from reaching the fibre apart from during Doppler cooling, where it helps to improve the ion lifetime. In apparatus B, the 397 laser is split into a $397\sigma^+$ and a 397π beam at the same frequency, and a ‘397 load’ beam which is not diffracted by

³³TOPTICA Photonics DL Pro.

³⁴TOPTICA Photonics iBEAM SMART.

³⁵Apparatus A: TOPTICA Photonics DL100. An optical isolator is placed immediately after each laser’s output. Apparatus B: TOPTICA Photonics DL Pro (isolators integrated into laser heads).

³⁶National Physical Laboratory (NPL) Low Drift Etalon.

³⁷Blue lasers: IntraAction Corp. ASM-2001.5B8. Red lasers: IntraAction Corp. ATM-851A2.

³⁸Apparatus A: New Focus 4431M. Apparatus B: New Focus 4435-02.

an AOM and is therefore 400 MHz red-detuned from the others. It is switched using a mechanical shutter, and has a mixed polarisation and a large spot size at the trap centre, to help with cooling of hot ions during loading. The detuned beams in each case therefore perform slightly different tasks due to the different trap depths; in the planar trap, this beam helps to prevent ion loss, whereas in the deeper blade trap it is used to cool ions that have been trapped but still have high motional energy.

The 866 nm lasers are each split into two frequencies approximately 300 MHz apart; one ('866') is used for high fluorescence and the other ('866 dark') for dark-resonance cooling.

The 850 nm lasers are each split into three separate components at different frequencies: 850π , $850\sigma_1^+$ and $850\sigma_2^+$. These are used during shelving to repump from the $3D_{3/2} |F, m_F\rangle = |5, +5\rangle$, $|5, +4\rangle$ and $|4, +4\rangle$ states respectively (see figure 2.5).

The calcium laser beam paths were built by T. P. Harty (apparatus A) and V. M. Schäfer (apparatus B); further details can be found in [66, 150].

Strontium lasers

The strontium laser beam paths (apparatus B only) were built by K. Thirumalai [177]. They include 422 nm, 1092 nm and 1033 nm ECDLs³⁹ which are each filtered with a grating, switched and intensity-stabilised via AOMs⁴⁰, and the 422 nm and 1092 nm lasers are locked to reference cavities⁴¹. The 422 nm beam is split into two separate frequencies using two further AOMs⁴² to form the 422u and 422l components (as seen in figure 2.4).


All the strontium diode lasers and most of the associated optical components are mounted on breadboards and housed in a 19-inch rack. The arrangement was initially developed by R. Nourshargh [135].

³⁹TOPTICA Photonics DL Pro.


⁴⁰Blue: Isomet 1250C-829A. Red: Gooch & Housego FibreQ.

⁴¹Stable Laser Systems, based on a design by NPL.

⁴²Isomet 1250C-829A.



(a) Apparatus A. The 850π beam enters through the top viewport, at 45° to the plane.



(b) Apparatus B

Figure 4.4: Beam geometries at the traps. $1/e^2$ intensity beam radii at the ions are $\approx 20 \mu\text{m}$ (A) / $\approx 50 \mu\text{m}$ (B) for blue diodes (except the 397 load), $\approx 100 \mu\text{m}$ for red diodes, and $\approx 30 \mu\text{m}$ for the Raman lasers and the 674.

4.7.2 Raman lasers

The ‘master’ and ‘slave’ Raman lasers (apparatus B only) are two phase-locked, frequency-doubled Ti:Sapphire lasers⁴³, pumped by 15 W and 10 W of 532 nm light⁴⁴. They are typically operated at 402 nm with an output power of 1 – 2 W each. A small amount of light is picked off from the output of each Ti:Sapph and sent to a photodiode to produce an error signal for a phase lock⁴⁵. This locks the slave laser frequency to 1.6 GHz above the master, to span the $^{43}\text{Ca}^+$ hyperfine splitting after frequency doubling. A more detailed description of the lasers can be found in [150], noting that the pumps were replaced in this work, and further beam paths were added to facilitate drive of Mølmer–Sørensen gates, as described below.

Light from the two lasers is split into multiple beams using AOMs⁴⁶. From the slave we derive ‘ R_H ’ and ‘ R_{H2} ’, and from the master we derive ‘ $R_{||}$ ’, and ‘ R_V ’ or ‘ R_{HSr} ’. All beams are coupled into optical fibres after the AOMs, then directed to the trap via piezo-controlled mirrors. Typical beam powers at the trap are tens of mW. This is much larger than the power required for the diode lasers (tens of μW) due to the large Raman detuning from the $S \leftrightarrow P$ transitions (-9.0 THz for $^{43}\text{Ca}^+$ and $+11.2$ THz for $^{88}\text{Sr}^+$). The powers are measured at the fibre outputs and stabilised using the AOMs.

The polarisations and AOM frequency offsets for all Raman beams are listed in table 4.3. Different pairs of beams are used for different purposes, as shown in table 4.4. The beam geometry at the trap is shown in figure 4.4b; the alignment is chosen so that the wavevector difference of motionally-sensitive beam pairs lies along the trap axis, meaning that the Raman beams only couple to the axial motion (see section 3.1.4).

⁴³2 $\times$ (M Squared SolsTiS + ECDx 400 doubling cavity).

⁴⁴2 $\times$ M Squared Equinox.

⁴⁵provided by M Squared.

⁴⁶ R_H : IntraAction Corp. ASM-851.5B8. R_{H2} : IntraAction Corp. ASM-2001.5B8. $R_{||}$ and R_V/R_{HSr} : Brimrose CQM-200-40-.400/OW.

The R_H AOM is driven via a synthesiser⁴⁷ at a fixed frequency⁴⁸. The RF drive for the other AOMs is via a DDS⁴⁹ (described in detail in [11]) which has four channels, each with eight profiles. The DDS outputs pass through switches⁵⁰ before being amplified⁵¹ to $\lesssim 2.5$ W. For the Mølmer-Sørensen gate, two channels of the DDS are used to drive the $R_{\parallel}$ AOM with two frequency tones simultaneously.

laser	beam	frequency offset (MHz)	polarisation
master	R_V	+217	$\sigma^{\mp}$
	$R_{\parallel}$	$\approx +217$	$\sigma^{\pm}$
	R_{HSr}	-192	π
slave (+3.2 GHz)	R_H	-109	π
	R_{H2}	+217	π

Table 4.3: Frequency offsets and polarisations of Raman beams.

beam pair	drives	motionally sensitive
$R_V + R_H$	stretch qubit	$\times$
$R_V + R_{H2}$	clock qubit	$\times$
$R_{\parallel} + R_H$	stretch qubit	$\checkmark$
$R_{\parallel} + R_{H2}$	clock qubit	$\checkmark$
$R_{\parallel} + R_{HSr}$	strontium	$\checkmark$
$R_V + R_{\parallel}$	light-shift gate	$\checkmark$

Table 4.4: Purposes of pairs of Raman beams. The R_V and all R_H beams are co-propagating, so pairs of these beams do not couple to the ions' motion. $R_{\parallel} + R_H$ is used for sideband cooling, on the stretch qubit. Mølmer-Sørensen gates, where the frequency splittings span the qubit frequencies for each ion, are driven by $R_{\parallel} + R_{HSr}$ in strontium and $R_{\parallel} + R_H/R_{H2}$ in calcium, depending on the qubit. $R_V + R_{\parallel}$ drives the light-shift gate on both species, with a frequency splitting $\approx \omega_{\text{ax ip,oop}}$ ($\lesssim 3$ MHz).

⁴⁷HP 8656B.

⁴⁸For MS gates on the $^{43}\text{Ca}^+$ stretch qubit, the RF drive sources for the R_H and R_{H2} AOMs were swapped, to allow profile-switching and pulse shaping on R_H .

⁴⁹Enterpoint Milldown.

⁵⁰Mini-Circuits ZASWA-2-50DR

⁵¹Mini-Circuits ZHL-03-5WF.

4.7.3 674 nm laser

The 674 nm quadrupole laser (apparatus B only) is a Ti:Sapph laser⁵² with an output of ~ 1 W, pumped with 15 W of 532 nm light⁵³. It is commercially stabilised⁵⁴ via a lock to a reference cavity and a fibre phase noise cancellation system, where light that is retro-reflected from the flat end facet of the fibre leading to the trap table is superposed with incoming light and used to produce an error signal for active phase stabilisation. Since the quadrupole transition has a natural linewidth of only 2.5 Hz, we require a laser with good frequency stability and short term linewidth to avoid decoherence, and a power of ~ 10 mW at the trap in order to drive the transition with a reasonable Rabi frequency ($\Omega \sim 1$ MHz). The specified linewidth is < 10 Hz over 1 s. A detailed description may be found in [177].

On the trap table, the beam passes through a double-pass AOM⁵⁵ for switching and intensity stabilisation, and a second, single-pass AOM⁵⁶, before passing down a second fibre to a piezo-controlled mirror and on to the trap. The single-pass AOM is driven at its centre frequency for single-qubit operations, with frequency offsets usually applied on the double-pass. To drive a Mølmer-Sørensen gate on the strontium optical qubit we apply two tones to the single-pass AOM. The AOMs are driven by the other three channels of the same DDS that is used for the microwave drive. The DDS outputs pass through switches⁵⁷ and are amplified⁵⁸ before reaching the AOMs.

The 674 nm beam enters the vacuum system at a slight angle to the Raman beams, and couples to all motional modes.

⁵²M Squared SolsTiS.

⁵³M Squared Equinox.

⁵⁴Stable Laser Systems.

⁵⁵Isomet 1205C-2-804B.

⁵⁶Isomet M1205-P80L-1.

⁵⁷Mini-Circuits ZASWA-2-DR+.

⁵⁸Mini-Circuits ZHL-03-5WF.

4.8 Experimental control

Experimental control in both setups is done mostly through ARTIQ⁵⁹, a piece of open-source software designed for use in quantum information experiments. A high-level programming language, based on Python, is used to write experiments, which may be scheduled to run at particular times and with different priorities. This allows for a significant amount of modularisation and automation, in order to run complex experiments over many hours without the need for manual recalibration or input; for example, if an ion is lost from the surface trap during an experimental run, this can be detected and another ion loaded automatically before resuming the run.

Nanosecond timing resolution is achieved for time-critical sections of experiments by compiling and running pulse sequences on an FPGA⁶⁰. This may include steps such as switching on and off microwave drive or drive of AOMs, opening and closing shutters, and changing DDS profiles. An input also allows monitoring of counts from the PMT with the same timing resolution. For tasks which are not time-critical, communication with other devices connected to the host computer (or to any other computer with a network connection to the host) is achieved via a remote procedure call mechanism. This includes things like setting the trap voltages, adjusting piezo controllers for cavities or mirrors, configuring DDS profiles, and switching between the PMT and camera.

ARTIQ includes databases for storing experimental parameters and results, and a graphical user interface which can be used to display results and monitor parameters in real time.

An overview of how the different parts of the experiments interact is shown in figure 4.5.

⁵⁹Advanced Real-Time Infrastructure for Quantum physics, M-Labs.

⁶⁰Xilinx KC705.

Chapter 5

Memory benchmarking

The third of DiVincenzo’s criteria for the successful physical implementation of a quantum computer is ‘long relevant decoherence times, much longer than the gate operation time’. Errors caused by decoherence during times when qubits are idle are known as ‘memory’ or ‘storage’ errors. These errors are usually measured in quantum systems in the regime where the error is much larger than the qubit state-preparation and measurement (SPAM) error (typically $\epsilon_{\text{SPAM}} \gtrsim 10^{-3}$). Information about the small-error regime — relevant for quantum computing applications — is inferred by extrapolation of these measurements using a particular model. In this chapter, we characterise the memory performance of a $^{43}\text{Ca}^+$ clock qubit, using a technique based on randomised benchmarking to directly measure decoherence errors down to the 10^{-6} level, without relying on extrapolation. This work is published in [156]. It was carried out in apparatus A.

5.1 Preliminary work

The first measurements of decoherence in this system — which motivated the work in this chapter — as well as characterisations of the single-qubit microwave gates required for benchmarking, were performed by J. E. Tarlton, M. A. Sepiol and J. F.

Goodwin and are described in more detail in [176]. Some key results are reproduced here to lend context to later discussions.

5.1.1 Ramsey experiments

The most widely used technique for measuring qubit decoherence is the Ramsey experiment, as described in section 3.2.2.1. The ion is (dark-resonance) cooled and prepared in $|\downarrow_C\rangle$, then the qubit state is rotated onto the equator of the Bloch sphere using a microwave $\frac{\pi}{2}$ -rotation about the x -axis. The qubit is left idle for a time τ before applying another $\frac{\pi}{2}$ -rotation, with phase $\tilde{\phi}$ relative to the first, and reading out the final state. The probability of measuring the bright state, $|\uparrow_C\rangle$, oscillates sinusoidally with $\tilde{\phi}$:

$$P(|\uparrow_C\rangle) = \frac{C(\tau)}{2} \cos(\tilde{\phi} - \phi_0) + \frac{1}{2}, \quad (5.1)$$

where ϕ_0 is an offset that accounts for any residual detuning between the qubit frequency and the microwave oscillator, and $C(\tau)$ is the Ramsey fringe contrast after a delay time τ . The memory error $\epsilon_m(\tau)$ (as defined in section 3.2.2.2) is

$$\epsilon_m(\tau) = \frac{1}{3} (1 - C(\tau)). \quad (5.2)$$

In previous work [176], the memory error in our system was characterised using Ramsey experiments. The population was only measured at two values of $\tilde{\phi}$, which were chosen to be at the expected peak and trough of the Ramsey fringes ($\tilde{\phi} = \phi_0$ and $\tilde{\phi}' = \phi_0 + 180^\circ$). The required offset, ϕ_0 , was calibrated before (but not during) experimental runs. Extracting the contrast from this two-point measurement ensures that the inferred memory error is sensitive to any drifts in ϕ_0 , as would be the case during a quantum computation.


In order to distinguish the contribution to the measured error due to decoherence from the contribution due to SPAM errors, it is necessary to measure the SPAM

errors with a similar statistical uncertainty as the Ramsey data. Each shot of the Ramsey experiment was therefore followed by ‘control’ measurements of the SPAM errors $\epsilon_{\uparrow_C}(\tau)$ and $\epsilon_{\downarrow_C}(\tau)$ of the bright and the dark state, for the corresponding delay time τ . In this way, any drifts in the total SPAM error $\epsilon_{\text{SPAM}}(\tau) \equiv \frac{1}{2}(\epsilon_{\uparrow_C}(\tau) + \epsilon_{\downarrow_C}(\tau))$, or systematic dependence on τ , could be captured. At the shortest delay time of 20 ms, the SPAM error was $\approx 2 \times 10^{-3}$, but at the longest delay time of 20 seconds the SPAM error was approximately an order of magnitude larger. The measured $\epsilon_{\text{SPAM}}(\tau)$ was subtracted from the measured fringe contrast loss for each delay time, before converting to a memory error.


Figure 5.1 shows the pulse sequences and the results of the Ramsey measurements. Two different decay models are fitted to the data at $\tau > 1$ s, the commonly accessible region where $\epsilon_m(\tau) \gg \epsilon_{\text{SPAM}}(\tau)$. The most commonly used decay model is an exponential decay $\epsilon_m(\tau) = \frac{1}{3}(1 - e^{-\tau/T_2^*})$, which assumes underlying white frequency noise (see section 3.2.2.4). Based on this fit, one would extract a T_2^* time of 22(3) s and would overestimate the error at $\tau = 100$ ms by around an order of magnitude. Also shown is an approximately Gaussian decay model, based on the assumption of $1/f$ frequency noise with a low-frequency cut-off, as described in section 3.2.2.4¹. This model fits the rest of the data much better. This experiment demonstrates that predictions of $\epsilon_m(\tau)$ in the small-error regime which are extrapolated from data at longer times, as is customary in T_2^* measurements, can deviate significantly from the true values unless one has a good understanding of the underlying noise processes which are relevant at these timescales.

The measured memory errors range from $4(9) \times 10^{-5}$ at $\tau = 20$ ms to 0.22(2) at $\tau = 20$ s. Below $\tau = 200$ ms, statistical uncertainty in $\epsilon_{\text{SPAM}}(\tau)$ limits knowledge of $\epsilon_m(\tau)$ to an upper bound only. In this range, the SPAM errors are comparable in size

¹The low-frequency cut-off parameter f_c causes the overturning behaviour and deviation from the form $\epsilon_m(\tau) \approx \frac{1}{3}(1 - e^{-(\tau/T_2^*)^2})$ when the contrast loss approaches 1.



(a) Pulse sequences used for Ramsey experiments and interleaved SPAM measurements.



(b) Decay of Ramsey contrast with time. The SPAM-corrected loss of contrast is plotted, with the conversion to a memory error (a scaling by $\frac{1}{3}$) on the right-hand ordinate. In blue is an exponential (white frequency noise) fit to the data at $\tau > 1$ s (where $\epsilon_m(\tau) \gg \epsilon_{\text{SPAM}}$), fixed at 0 for $\tau = 0$. The fitted T_2^* time is 22(3) s. In red is a fit to the same subset, which assumes underlying $1/f$ frequency noise ($S_{\dot{\phi}}(\omega) = S_{1/f} 2\pi/\omega$) with a low-frequency cut-off f_c , leading to an approximately Gaussian decay of contrast. The fitted parameters are $S_{1/f} = 0.0100(8)$, $f_c = 0.0145(7)$ Hz. Shaded regions indicate the 1σ uncertainties in the fits.

Figure 5.1: Measuring the memory error using two-point Ramsey experiments with interleaved SPAM measurements. The data are reproduced from [176], where more details can be found. The $1/f$ noise fit is new in this work.

to the memory errors. The binomial uncertainty in ϵ_{SPAM} is given by

$$\sqrt{\frac{\epsilon_{\text{SPAM}}(1 - \epsilon_{\text{SPAM}})}{N}}, \quad (5.3)$$

where N is the number of shots (experimental repeats) of the control sequences. Collecting more data would therefore reduce this uncertainty, but only with the square-root of the acquisition time. This experiment already required several days of continuous data acquisition, so it would be practically impossible² to use the Ramsey method for characterising memory errors $< 10^{-4}$.

5.1.2 Benchmarking of single-qubit microwave gates

To circumvent the limitation imposed by the SPAM error, we wish to use a technique based on interleaved randomised benchmarking (IRB), outlined in section 3.4.2.1. This will involve subjecting the qubit to a sequence of dephasing periods of length τ , interleaved with Clifford gates sampled randomly from the single-qubit Clifford group. Before proceeding with memory benchmarking, it is instructive to characterise the randomising Clifford gates themselves.

5.1.2.1 Standard Cliffords

We can implement single-qubit gates via microwave pulses with appropriately-chosen phase, amplitude and duration to perform $\pm\frac{\pi}{2}$ -rotations about the x - and y -axes. The 24 operations in the single-qubit Clifford group can be implemented using just these gates (see appendix C). A standard randomised benchmarking protocol, as described in section 3.4.1.3, was used in [176] to characterise the fidelity of Clifford gates implemented in this way. The same benchmarking sequences are also used as a sensitive method for calibrating the $\frac{\pi}{2}$ -pulse durations ($\approx 10 \mu\text{s}$). The $\frac{\pi}{2}$ -pulses are

²In appendix B, we estimate that to measure a memory error of 10^{-6} to a fractional uncertainty of 10% in our system would require 42 years of continuous data acquisition.


separated by $12\ \mu\text{s}$ to allow the DDS source time to switch between profiles, and there are on average $3.50\ \frac{\pi}{2}$ -pulses per Clifford gate. Results are shown in figure 5.2; the error per Clifford is $\epsilon_C = 1.7(2) \times 10^{-6}$.

5.1.2.2 BB1-protected Cliffords


The memory benchmarking process relies upon measuring the fidelities of random sequences of single-qubit Clifford gates, with interleaved delays and without interleaved delays, and attributing the increased error in the interleaved sequences to the memory error acquired during the delays. In order for this measurement to be reliable, the error contribution from the Clifford gates must remain stable and must not change when delays are introduced. However, it was found in [176] that the microwave power varied with the duty cycle of the microwave chain, due to thermal effects, and that it was necessary to recalibrate the pulse duration every few hours, due to drifts in the microwave power on this timescale. Sequences with and without delays have significantly different microwave duty cycles, so the average gate fidelity is expected to vary. In principle, the microwave power or pulse duration could be recalibrated for each delay length separately. However, for longer delays these calibrations would take an impractically long time. Instead, each $\frac{\pi}{2}$ -pulse was replaced with a composite ‘BB1’ $\frac{\pi}{2}$ -pulse [187], which protects against errors due to over- or under-rotation³. The BB1 pulse sequence is shown in figure 5.3.

Standard randomised benchmarking was used again in [176] to characterise the performance of BB1-protected Cliffords compared to standard Cliffords, at the optimal microwave power and again at a power corresponding to a 1% reduction in Rabi frequency. It was found that the standard pulses perform well if the power remains optimal ($\epsilon_C = 1.4(4) \times 10^{-6}$, consistent with the results in figure 5.2), but the fidelity decreases significantly at the lower power ($\epsilon_C = 3.3(5) \times 10^{-4}$). The

³Many composite pulse sequences exist which protect quantum gates from various sources of error. Most of these were developed in the context of NMR; a summary may be found in [45].



(a) Pulse sequences for standard randomised benchmarking of Clifford gates C_i . The qubit is prepared in $|\downarrow_C\rangle$ and then subjected to m random Cliffords, plus a final Clifford $C^{-1,P}$. The final gate is not always chosen to invert the previous m gates and rotate the qubit state back to $|\downarrow_C\rangle$ (in the absence of errors). Instead, it rotates to a target state of either $|\downarrow_C\rangle$ or $|\uparrow_C\rangle$, with equal probability. This additional ‘Pauli randomisation’ reduces the effect of any asymmetry in the readout errors for the two qubit states. After each benchmarking sequence, a measurement of the SPAM error is made. The target state is chosen to match that of the preceding benchmarking sequence, and the delay between state preparation and measurement is equal to the duration of the sequence.



(b) Measured sequence fidelities, averaged over $k = 50$ randomisations for each sequence length m . A fit to a decay model $F(m) = Ap_C^m + \frac{1}{2}$ (as in equation (3.120), for one qubit) gives an error per Clifford of $\epsilon_C = \frac{1}{2}(1 - p_C) = 1.7(2) \times 10^{-6}$ and a SPAM/inversion error of $\epsilon_{S/I} = \frac{1}{2} - A = 3.0(7) \times 10^{-3}$. The mean concurrently-measured SPAM error is $2.7(4) \times 10^{-3}$, which is consistent with the benchmarking fit intercept. Data reproduced from [176].

Figure 5.2: Standard randomised benchmarking of single-qubit Clifford gates implemented via $\pm\frac{\pi}{2}$ -rotations about the x - and y -axes.




Figure 5.3: To protect a pulse of target rotation angle θ about an axis rotated from the x -axis by an angle ϕ_0 in the x - y plane, the pulse should be followed by a sequence of three pulses about axes at angles $\phi_1 = \phi_0 \pm \arccos\left(-\frac{\theta}{4\pi}\right)$ and $\phi_2 = \phi_0 \pm 3\arccos\left(-\frac{\theta}{4\pi}\right)$ from the x -axis in the x - y plane. In the absence of errors, this sequence performs a net identity operation. In the case of an over-rotation $\delta\theta$ in the target angle, the additional pulses reduce the dependence of the resulting error from $O[(\delta\theta)^2]$ to $O[(\delta\theta)^6]$ [176].

BB1 pulses perform slightly less well than the standard pulses at the optimal power ($\epsilon_C = 6(1) \times 10^{-6}$), but are much more robust to power variations ($\epsilon_C = 6.5(9) \times 10^{-6}$ at the reduced power).

Due to technical limitations of the DDS used for the qubit microwave pulses, we do not have the flexibility to implement $\pm\frac{\pi}{2}$ x - and y -pulses with BB1 protection, so it was necessary to restrict the Clifford decomposition to include a ‘reduced’ gate set of only $+\frac{\pi}{2}$ x - and $+\frac{\pi}{2}$ y -pulses, with an average of 3.58 $\frac{\pi}{2}$ -rotations per Clifford (see appendix C).

5.2 Memory benchmarking results

5.2.1 Experimental method


To extend our measurement of memory errors to smaller scales, we follow the approach introduced by O’Malley et al. [137], which employs IRB sequences where the ‘gate-under-test’ is simply a delay of length τ . In the absence of any decoherence, this would be equivalent to an identity gate. The memory error is amplified relative to the SPAM error by subjecting the qubit to m dephasing periods of length τ , instead of just a single period, whilst ensuring that errors add incoherently by interleaving each delay τ with a Clifford gate, sampled randomly from the full single-qubit Clifford group. To make these Clifford gates robust to drifts in microwave power — which

would otherwise necessitate frequent recalibration — they are implemented via BB1-protected $+\frac{\pi}{2}$ -rotations about the x - and y -axes, as described in the previous section.


We use the simplified method described in section 3.4.2.1, where we do not fit full fidelity decay curves ($F(m)$ versus sequence length m) for each delay length τ ; for the required delay times and numbers of repeats, this would require an impractically large amount of time. Instead we choose a particular sequence length m for each τ and calculate the memory error for that delay via the interleaved and non-interleaved sequence fidelities using equation (3.126), with $B_n = \frac{1}{2}$ and $\alpha_n = 2$ for $n = 1$ qubits. The lengths m range from $m = 1$, for $\tau = 10$ s, to $m = 5000$, for $\tau < 20$ ms, and are chosen such that the average total infidelity of each sequence length is ~ 0.1 . This is larger than the SPAM error, even for long delays where the ion is heated significantly [176], so we avoid the limitation which is encountered using Ramsey experiments.

Figure 5.4a shows the types of sequences that are used. The non-interleaved sequences include a delay of length $m \times \tau$ before readout, to keep the time between qubit initialisation and readout equal to that in the interleaved sequences, thus minimising any systematic differences in the SPAM/inversion error $\epsilon_{S/I}$ and ensuring the validity of the simplified approach. This delay occurs after rotation into an eigenstate of σ_z , so we do not introduce any extra dephasing during this time. The final Clifford gate in the interleaved sequence is chosen to rotate to a target state of either $|\downarrow_C\rangle$ or $|\uparrow_C\rangle$ with equal probability, to reduce the effect of any asymmetry in the readout errors of the two states.

As in the Ramsey experiments, we are sensitive to slow drifts of the detuning of the microwave source relative to the qubit transition frequency. During each data point (which takes up to two days to acquire), we measure this detuning every few hours and correct for these drifts; from these measurements, we estimate that the contribution to $\epsilon_m(\tau)$ due to these detuning errors is negligible, as detailed in [176].



(a) Example pulse sequences, showing decomposition into BB1-protected $+\frac{\pi}{2}$ -pulses. The final gate, $C^{-1,P}$, is chosen to rotate the state to $|\uparrow_C\rangle$ or $|\downarrow_C\rangle$. The non-interleaved sequences include an extra delay $m \times \tau$ before readout, to reduce differences in the SPAM errors due to different amounts of heating in the two sequence types.



(b) Average memory errors measured via IRB, with $k = 50$ randomisations per sequence length. The Ramsey data from section 5.1.1 are overlaid for comparison. We fit a decoherence model (blue line) including white and $1/f$ frequency noise, with fitted parameters $S_W = 0.010(2)$, $S_{1/f} = 0.009(2)$ and $f_c = 25(3)$ mHz. The yellow, green and purple data are from IRB experiments repeated with the magnetic field offset by $\Delta B = (10, 25, 50)$ mG from the optimal (clock) field. The dotted lines show the calculated $\epsilon_m(\tau)$ due to the corresponding frequency offsets $\Delta f = (0.12, 0.76, 3.0)$ Hz.

Figure 5.4: Memory benchmarking results.

5.2.2 Results at the clock field

The results are shown in figure 5.4b. The measured memory errors are consistent with those from the Ramsey experiments, but the superior sensitivity of the IRB approach enables characterisation of $\epsilon_m(\tau)$ for delays as short as 1 ms, where $\epsilon_m = 1.2(7) \times 10^{-6}$. At this point, the limiting factor is the uncertainty on the error introduced by the randomising Clifford gates, which plays a similar role to the SPAM error in the Ramsey measurements.

We find that the memory error remains below 10^{-4} for $\tau \lesssim 50$ ms, which is around three orders of magnitude longer than the typical time required for entangling gates or measurement in ion-trap systems [12, 60, 151, 195, 43]. Depending on the timescale and the decay model used, one could extract an ‘effective T_2^* time’ of between 40 seconds and 400 seconds from this data.

We find that neither an exponential nor a Gaussian decay model fits the whole data set well. Although a model containing only $1/f$ frequency noise seemed appropriate for the Ramsey data, this model in fact underestimates the measured memory error at 1 ms by more than an order of magnitude. Instead, we fit a decoherence model which includes contributions from a source of white frequency noise ($S_{\dot{\phi}}(\omega) = S_W$) and a source of $1/f$ frequency noise ($S_{\dot{\phi}}(\omega) = S_{1/f} 2\pi/\omega$) with a low-frequency cut-off f_c :

$$\epsilon_m(\tau) = \frac{S_W \tau}{12} + S_{1/f} \tau^2 \ln \left(\frac{0.4007}{f_c \tau} \right), \quad (5.4)$$

as discussed in section 3.2.2.4, with S_W , $S_{1/f}$ and f_c as free parameters.

The $1/f$ contribution dominates at the longer timescales that were measured in the Ramsey experiments, and indeed the fitted value of $S_{1/f}$ for the benchmarking data is consistent with the fitted $S_{1/f}$ for the Ramsey data. The white-noise contribution dominates in the regime where $\epsilon_m(\tau) < 10^{-4}$. The typical length of a benchmarking sequence is ≈ 40 seconds, so we expect the experiment to be sensitive to noise down




Figure 5.5: Experimentally-measured dependence of the detuning Δf of the qubit frequency from the clock frequency ($f_0 = 3.2$ GHz) on the magnetic-field offset ΔB from the field-independent point ($B_0 = 146.0942$ G). The fitted second-order field dependence is $d^2 f/dB^2 = 2.40(2)$ kHz/G².

to frequencies of $\approx 1/40$ s = 25 mHz, and the fitted value of f_c is 25(3) mHz.

These results again demonstrate the importance of measuring $\epsilon_m(\tau)$ directly in the region of interest, rather than the customary approach of extrapolating from measurements at longer timescales. They also highlight the complexity which can be overlooked in characterising qubit memory performance simply using a single parameter, the T_2^* time.

5.2.3 Detuned magnetic field

To avoid decoherence, we usually control the external magnetic-field strength (section 4.4.2) in order to remain at the field-independent point of the clock qubit, $B_0 = 146$ G. An offset ΔB of the field from this point results in an offset of the qubit frequency

$$\Delta f = \frac{1}{2} \frac{d^2 f}{dB^2} (\Delta B)^2, \quad (5.5)$$

to lowest order in ΔB . Figure 5.5 shows an experimental measurement of the dependence of the qubit frequency on the magnetic field.

Figure 5.4b includes the results of three additional repeats of the IRB experiment with the magnetic field deliberately detuned away from the optimal operating point, in order to characterise the robustness to field inhomogeneities. This will be an important consideration for scaling up ion-trap processors, which is likely to involve the shuttling of ions across different trapping regions in the device (for example in a ‘quantum CCD’ architecture), where the field is less well controlled.

ΔB is set to three constant offset values of 10 mG, 25 mG and 50 mG, with a similar accuracy as it was previously set to zero. To obtain a ‘worst-case’ estimate of the memory error, the microwave detuning is not adjusted for the known resulting frequency offsets $\Delta f = 0.12$ Hz, 0.76 Hz and 3.0 Hz. We find that the qubit is quite robust to detuning of the field; even with an offset of 50 mG, $\epsilon_m(\tau)$ remains below 10^{-4} for $\tau = 1$ ms, which is still 1 – 2 orders of magnitude longer than typical ion shuttling times⁴ [106, 78].

The field detuning introduces two extra sources of error. Firstly, there is an error due to the (nominally constant) frequency offset Δf , and secondly, there is an extra error due to the increased sensitivity of the qubit frequency to residual field noise at the offset point. The dotted lines in figure 5.4b indicate the predicted memory error accounting for the constant offset, which adds an extra error

$$\epsilon_m^{\Delta B}(\tau) = \frac{2}{3} [1 - \cos^2(\pi\tau\Delta f)] . \quad (5.6)$$

For each offset, it is evident that $\epsilon_m(\tau)$ is dominated by this error, which could, in theory, be corrected for in a quantum processor by calibrating the field across the device.

⁴When shuttling an ion between trap zones, it is preferable that the motional state of the ion is not disturbed. In order to achieve this adiabaticity, shuttling times tend to be long compared to the secular motion period of the ion in the trap. Trap secular frequencies are usually ~ 1 MHz, making reasonable shuttling times 10 – 100 μ s.

5.2.4 Memory performance with dynamical decoupling

As discussed in section 3.2.3, appropriately chosen ‘dynamical decoupling’ (DD) pulse sequences can correct dephasing due to correlated noise. If the underlying noise in our system does indeed consist of a white frequency noise component and a $1/f$ frequency noise component, we expect that applying a DD sequence during the delay periods should suppress the contribution to $\epsilon_m(\tau)$ from the $1/f$ noise, allowing measurement of the uncorrelated noise contribution across a wider range of delay times.

We repeat the IRB experiment, for delay times $\tau \geq 200$ ms, with added CPMG decoupling during the delay periods. This involves repeated π -rotations about the x -axis, with a spacing $\tau_{DD} = 100$ ms between rotations, as shown in figure 5.6a.


The results are shown in figure 5.6b, alongside the previous data for reference. The memory error is now kept below 10^{-3} for up to 2 seconds. The system behaves as one which experiences dephasing due to the previously-fitted white frequency noise component alone, strongly suggesting that the DD scheme is able to effectively suppress one of two separate contributions to the memory error.

5.3 Sources of decoherence


5.3.1 Phase noise measurements

One possible source of decoherence is phase noise on the clock qubit microwave drive. We measure the phase noise at different points in the drive chain, which consists of a 10 MHz rubidium atomic reference oscillator, a 990 MHz clock generated by a synthesiser, and a 400 MHz DDS source which is octupled to 3.2 GHz (as described in section 4.6).

We use the phase-detector method to convert a phase difference into a voltage. This involves setting up a duplicate chain of components with a nominally identical



(a) Sequences used for repeats of the IRB memory experiment with additional CPMG dynamical decoupling pulses (πx) during the delay periods. Each πx -pulse is composed of two BB1-protected $+\frac{\pi}{2}x$ -pulses. In the interleaved sequence, the final gate $C^{-1,P}$ also compensates for the effect of these additional pulses on the qubit state.



(b) The diamond-shaped points show the results of memory benchmarking with DD. At $\tau = 5$ seconds, we measure $\epsilon_m < 0$ (a decrease in error when a delay is added), for unknown reasons. The original data is included for comparison (blue circles). The black dot-dashed line is the predicted contribution to $\epsilon_m(\tau)$ from the white-noise component of the fit to the original data, with the grey shaded region showing the 1σ uncertainty. The red dashed line is the predicted memory error from independently-measured noise on the rubidium reference oscillator used in the microwave drive chain, with the red shaded region representing an estimate of the uncertainty of this prediction due to measurement repeatability.

Figure 5.6: Memory benchmarking with dynamical decoupling. The DD pulses appear to suppress the $1/f$ frequency noise ($1/f^3$ phase noise) component, reducing the memory error to a level consistent with white frequency noise ($1/f^2$ phase noise) from the Rb reference oscillator.




Figure 5.7: To measure phase noise, we use a duplicate chain of components of the same models as those used for our qubit microwave drive. The two chains provide signals at nominally identical frequencies which are combined on a mixer in a phase-detector measurement. To measure the noise on components downstream of the rubidium reference oscillator, we clock both synthesisers from the same reference, as indicated by the blue arrow.

frequency output, as shown in figure 5.7, and combining the two sources on a mixer⁵ followed by a low-pass filter⁶ to remove the sum-frequency component. When the signals are adjusted to be in phase quadrature (a 90° phase difference), the nominal DC output of the phase detector is zero Volts. Any phase fluctuations around the quadrature point are converted to voltage fluctuations, which we monitor on an oscilloscope⁷.

First, it is necessary to calibrate the detector response, which depends on the insertion loss and the drive level used for the mixer. We do so by splitting the signal from the original chain into two arms, recombining them on the mixer and manually stepping the phase difference. The calibration is shown in figure 5.8; we can use this to convert the measured voltage fluctuations to a phase noise power spectral density (PSD) $S_\phi(f)$ in radians²/Hz.

We measure the phase noise on the microwave chain over a period of 1 second with a sampling interval of $2 \mu\text{s}$. We start close to zero voltage, and over the measurement time the phase fluctuations remain in the linear response regime ($\lesssim 0.2$ radians),

⁵Mini-Circuits ZX05-73L-S+

⁶Mini-Circuits ZX75LP-120-S+

⁷Tektronix TDS 5104




Figure 5.8: Phase detector response calibration. The phase is stepped by 18° every 0.25 seconds. The linear gradient around zero Volts is $204.5(1)$ mV/radian.

where the sensitivity is maximised. The measured PSD is shown in figure 5.9; we find clear $1/f^2$ phase noise (white frequency noise), with a magnitude consistent with the magnitude of the white frequency noise component of the fit to the memory benchmarking data.

We can attribute the noise to a particular section of the chain by ensuring that all noise upstream of that section is common-mode rejected from the measurement. To measure the phase noise introduced by all components downstream of the Rb reference oscillator, we choose to clock both synths from the original reference oscillator rather than using the duplicate reference. In this configuration, we use a sampling interval of $2 \mu\text{s}$ over a period of 100 seconds, where the phase excursions remain < 0.1 radians. We find the resulting noise is too low to explain the measured memory error (figure 5.9), allowing us to identify the Rb reference oscillator as the dominant source of phase noise in the microwave chain.

We use the direct results of the 1-second measurement — plus an extrapolation assuming white frequency noise down to 25 mHz — to calculate the predicted contribution to the memory error due to this noise. This is the red dashed line in figure 5.6b. It fits the DD data very well, and lies close to the predicted memory error due to the




Figure 5.9: Power spectral densities of phase noise on the qubit microwave drive. The top panel shows the mean of three measurements of noise on the whole chain. These were taken on an oscilloscope, over a period of 1 second, with a sampling interval of $2\ \mu\text{s}$. The red shaded region represents the uncertainty due to measurement repeatability. The lower panel shows the noise introduced by components downstream of the Rb reference oscillator in the microwave chain. In each case, the noise floor of the oscilloscope (measured using the same settings as for the corresponding phase noise measurement) is shown in grey. The black dot-dashed lines show the $1/f^2$ phase noise (white frequency noise) component of the fit to the memory benchmarking data. We conclude that the Rb reference oscillator is the likely source of this noise.

white noise component of the fit to the original data⁸.

We conclude that phase noise on the Rb reference oscillator is the dominant source of memory error for $\tau \leq 200$ ms, and also for longer times if a simple CPMG dynamical decoupling sequence is implemented. Therefore, the current limit on our memory performance is not imposed by any fundamental property of the qubit, and we expect the performance could be improved by replacing the reference oscillator with a more stable model.

5.3.2 Magnetic-field noise

Another source of qubit decoherence could be magnetic-field noise. As described in section 4.4.2, we stabilise the magnetic field through multiple feedback and feed-forward mechanisms, including a servo routine which uses the ion itself as a probe. During memory benchmarking experiments, this servo is typically run every two minutes. Residual offsets from the clock field can be estimated from the corrections made with each application of the routine, which are shown in figure 5.10. The field remains within ≈ 1 mG of the clock point at all times, with a root-mean-square (RMS) deviation of $\sigma_B = 0.4$ mG. Using equations (5.5) and (5.6), we can calculate the memory error due to an offset of this size, for different delay times: $\epsilon_m(\tau = 1 \text{ ms}) = 2 \times 10^{-13}$, $\epsilon_m(\tau = 0.1 \text{ s}) = 2 \times 10^{-9}$, $\epsilon_m(\tau = 10 \text{ s}) = 2 \times 10^{-5}$. This is several orders of magnitude smaller than the measured memory error.

We can also estimate the contribution of magnetic-field noise to the memory error by considering its effect on the field-sensitive stretch transition. We perform Ramsey experiments on the stretch transition, with and without the addition of a single πx spin-echo pulse half way through the delay period; the results are shown in figure 5.11. The resulting decay of contrast is Gaussian in both cases, indicating correlated noise

⁸To check the validity of our extrapolation, we measure the slower noise on the whole chain by logging the phase-detector output voltage on a digital multimeter (Keysight 34470A) over a period of several hours. The sensitivity is poorer, but we do not find any significant deviation from the PSD which would be extrapolated from the faster measurements.




Figure 5.10: Time series of corrections made to the magnetic field after each application of the servo routine, referenced to the frequency of the stretch transition. The servo runs every two minutes. The RMS deviation of the corrections from zero is 0.4 mG.

$S_B(f) \propto 1/f$ on a timescale longer than the Ramsey delays. Under this condition, the mean-square phase deviation is given by [123, 137]

$$\langle [\delta\phi(\tau)]^2 \rangle = (2\pi)^2 \langle [\delta f_S(\tau)]^2 \rangle \tau^2, \quad (5.7)$$

where f_S is the instantaneous stretch transition frequency. Using $\epsilon_m(\tau) = \frac{1}{6} \langle [\delta\phi(\tau)]^2 \rangle$ and $C(\tau) \approx 1 - 3\epsilon_m(\tau) \approx \exp[-(\tau/T_2^*)^2]$, we can therefore estimate the RMS magnetic field deviation over the relevant delay time from the fitted T_2^* time, via:

$$\sigma_B = \frac{\sqrt{2}}{2\pi |df_S/dB| T_2^*}, \quad (5.8)$$

where $df_S/dB = -2.36$ MHz/G is the magnetic-field sensitivity of the stretch transition.

From the Ramsey contrast decay without spin echo, we extract a T_2^* time of 0.26(2) ms. This is attributable to magnetic-field offsets of ≈ 0.3 mG over the duration of




Figure 5.11: Ramsey contrast at different delay times for the stretch transition, with and without the addition of a single spin-echo π x -pulse. The data are well described by the fitted Gaussian decay curves ($\exp[-(\tau/T_2^*)^2]$), implying a $1/f$ dependence of the underlying frequency noise. The fitted T_2^* is 0.26(2) ms without spin echo, or 2.1(2) ms with spin echo, corresponding to RMS magnetic-field deviations of ≈ 0.3 mG and 47(4) μ G respectively.

the experiment, which is consistent with the deviation of the servo corrections shown in figure 5.10. Adding a spin-echo pulse corrects for offsets, extending the coherence time to $T_2^* = 2.1(2)$ ms, which corresponds to a residual RMS magnetic-field deviation of $\sigma_B = 47(4)$ μ G.

These measurements allow us to put constraints on the memory error due to residual magnetic-field noise at millisecond timescales. From the servo corrections, we assume a typical field offset of $\Delta B = 0.4$ mG. The sensitivity of the clock transition frequency f_C to magnetic-field noise at this point is approximately 1 million times smaller than that of the stretch transition frequency f_S :

$$\begin{aligned}
 \left(\frac{df_C}{dB}\right)_{\Delta B=0.4\text{mG}} &\approx \left(\frac{d^2 f_C}{dB^2}\right)_{\Delta B=0} \times \Delta B \\
 &= 0.97 \text{ Hz/G} \\
 &= 0.4 \times 10^{-6} \times \frac{df_S}{dB}.
 \end{aligned} \tag{5.9}$$

This implies the angular dephasing on the clock qubit is approximately 1 million times smaller than that on the stretch qubit for any given delay time, leading to a memory error $\epsilon_m(\tau = 2.1 \text{ ms}) \sim 10^{-13}$. Assuming the $1/f$ frequency dependence of the magnetic-field noise, implied by the stretch transition contrast decay, extends to lower frequencies, we again find that the predicted memory error remains several orders of magnitude lower than the measured error. In fact, we expect that residual magnetic-field noise will not limit the memory performance until $\Delta B \gg 100 \text{ mG}$ ⁹.

5.3.3 AC Zeeman shift

Having ruled out magnetic-field fluctuations, the source of the $1/f$ frequency noise remains unidentified. We propose variations in the AC Zeeman shift due to the trap RF electrodes as a possible candidate.

The trapping RF frequency Ω_T in apparatus A is $\approx 38 \text{ MHz}$, which is fairly close to the Zeeman splittings ($\approx 50 \text{ MHz}$) between hyperfine states in the ground level of $^{43}\text{Ca}^+$. This results in a shift on the clock qubit frequency, which depends on the polarisation of the RF field and varies over small distances. Previous work [66] measured this shift to be $\sim 5 \text{ Hz}$ at typical powers, and characterised its dependence on the trap RF power ($\sim 0.01 \text{ Hz/mW}$) and on the ion's radial position ($\sim 1 \text{ Hz}/\mu\text{m}$). To account for the $1/f$ frequency noise we observe would require power fluctuations of $\sim 0.5\%$ or displacements of $\sim 100 \text{ nm}$. The trap RF supply is designed to be stable to $< 0.01\%$, and previous measurements of the secular frequency showed that this was approximately true [66]. On the other hand, the required radial displacements correspond to compensation voltages which are smaller than we are able to calibrate for in the out-of-plane direction. We therefore propose movement of the ion in the trap RF field as a possible source of the decoherence at longer times, although we do not have direct evidence of this.

⁹Here, we assume that the nominally constant resulting frequency offset is corrected for, and consider only the residual noise around the offset point.

5.4 Summary

Thanks to the discovery of quantum error correction, the required time before decoherence occurs in a quantum system does not scale with the time required to complete a computation. Instead, as long as the rate at which memory errors are introduced remains below a correctable level, the computation can succeed regardless of its length or complexity. A typical error-correction protocol will require preparation of ancilla qubits, entanglement of these ancilla qubits with logical qubits, and measurement of the ancillas, before any correction can be applied. Therefore, the important metric is the number of gates and measurement steps which can be performed while the memory error remains small [168, 97].

We find the memory error ϵ_m remains $< 10^{-4}$ for up to 50 ms, which exceeds typical gate or measurement times by three orders of magnitude and corresponds to an ‘effective T_2^* time’ of between 40 seconds and 400 seconds, depending on the model used. Our results are consistent with a model which assumes a combination of white frequency noise and $1/f$ frequency noise, where the $1/f$ noise dominates at timescales ≥ 200 ms. A naïve extrapolation of the data in the typically-accessible regime where $\epsilon_m > \epsilon_{\text{SPAM}}$, based on the usual assumption of exponential decay, would significantly overestimate the true memory errors at the shorter times relevant for quantum computation.

Using a simple CPMG dynamical decoupling sequence, we can suppress the effect of the $1/f$ noise and further improve the memory performance such that ϵ_m remains below 10^{-3} for 2 seconds. We identify the remaining limiting noise source in our system as phase noise on the reference oscillator, indicating that the memory performance may be improved further with the use of a more stable reference.

We also provide a measure of the robustness to offsets of the external magnetic field from the ‘clock’ point, which is likely to be relevant when shuttling ions across a larger quantum processor. Even with an offset of 50 mG, ϵ_m remains below 10^{-4}

for 1 ms, which is 1 – 2 orders of magnitude longer than typical ion shuttling times. The error is dominated by the resulting frequency offset of 3.0 Hz, and could be corrected if the magnetic-field strength was characterised across different trap zones. We predict that residual magnetic-field noise will not limit the memory performance until the offset is $\gg 100$ mG.

Chapter 6

Mixed-species two-qubit gates

To achieve universal quantum control, entanglement between qubits is required. The ability to transfer quantum information between different ion species offers significant advantages for the flexibility and scalability of quantum processors, and for a range of other applications such as spectroscopy and optical clocks. This chapter will present various experimental implementations of geometric phase gates between two ions of different species — $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ — with the goal of selecting the most robust mechanism for future applications. This work was carried out in apparatus B.

6.1 Calibration and characterisation

This section will discuss several procedures which are relevant for all two-qubit gates in this chapter.

6.1.1 Recrystallisation and re-ordering

Occasional collisions with background gas molecules can heat the ions. The trap is deep enough that ions are very rarely completely lost, but they can end up in higher orbits rather than in a well defined, linear crystal along the z -axis. Before




Figure 6.1: Re-ordering a mixed-species crystal: (i) Lower the RF confinement until a radial (rather than axial) crystal is formed. (ii) Change the endcap voltages to displace the crystal from the trap centre. The axial confinement due to the RF is mass-dependent, so the crystal tilts, with the lighter ion remaining closer to the centre. (iii) Raise the RF confinement to go back to an axial crystal, but in a well-defined order (determined by the direction of the axial displacement). (iv) Finally, reset the endcap voltages to return to the centre.

each shot of every experiment, therefore, we check the fluorescence rate. If it has fallen below the usual level for the ions we expect to have in the trap, we conclude that a ‘decrySTALLISATION’ event has occurred and we discard the preceding data point. We weaken the trap RF amplitude using the variable attenuator, and switch on the ‘397 load’ beam (a beam with mixed polarisation and a large spot size at the ion, which is red-detuned by 400 MHz from the Doppler-cooling transition in calcium), to facilitate cooling. We then return the trap RF strength to its usual level and check that the fluorescence has recovered.

To ensure consistency between experimental runs and calibrations, we also require that the order of the ions in the trap remains fixed. After recrystallisation, we include a step to re-order the ions into a known configuration. The technique is illustrated in figure 6.1 and is commonly used in mixed-species ion-trap experiments [71].

6.1.2 Crystal tilt

If the trap is poorly compensated such that the ions are displaced radially from the centre, mixed-species crystals tilt away from the z -axis because of the mass-dependence of the radial confinement. This tilt changes the Lamb–Dicke factors of the ions in the gate beams. The Raman beams are aligned such that their differential wavevector $\Delta\mathbf{k}$ lies parallel to the trap axis and they couple only to the axial modes,

on which we drive entangling gates. Tilting of the crystal will increase the projection of $\Delta\mathbf{k}$ on the radial direction, and therefore increase errors due to unwanted coupling to the radial modes.

Since same-species crystals do not tilt, the Raman beams can be carefully aligned on a $^{43}\text{Ca}^+ - ^{43}\text{Ca}^+$ crystal by minimising the excitation of radial modes. Any residual coupling to the radial modes of a $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal must therefore be due to imperfect compensation. We set the radial compensation voltages by minimising this coupling, with final adjustments being made using the entangling gate fidelity itself.

6.1.3 Mode frequencies

The trap parameters were chosen carefully, to avoid certain resonances between the resulting motional mode frequencies of a $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal. All gates in this chapter are carried out on the axial modes ($f_{\text{ax,ip}} = 1.49$ MHz, $f_{\text{ax,oop}} = 2.91$ MHz). As discussed in section 3.3.5.2, we wish to avoid the resonance condition $f_{\text{ax,oop}} \simeq 2f_{\text{rad,oop}}$, to reduce errors due to Kerr cross-coupling for gates on the oop mode. For gates on the ip mode, we also avoid $f_{\text{ax,ip}} \simeq f_{\text{rad,oop}}$, to reduce off-resonant coupling to the radial modes in case of a slight misalignment of the Raman beams. Finally, we avoid $f_{\text{ax,oop}} \simeq 2f_{\text{ax,ip}}$; otherwise, certain choices of gate detuning could lead to higher-harmonic excitation of the spectator axial mode. All of the motional mode frequencies can be found in figure 4.3.

6.1.4 Cooling

Before performing an entangling gate, the ions are dark-resonance cooled and both axial modes are sideband cooled on the $^{43}\text{Ca}^+$ stretch transition, as described in section 2.2.5. The whole process takes ≈ 10 ms. We reach $\bar{n}_{\text{ax,ip}} = 0.07(1)$ and $\bar{n}_{\text{ax,oop}} = 0.030(7)$ (see figure 6.2).




Figure 6.2: Temperature diagnostics for $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$. We can scan the frequency difference of the $R_{\parallel}$ and R_H beams to excite the blue sidebands (BSB) of the axial motional modes on the stretch qubit transition in $^{43}\text{Ca}^+$, as shown in the figure. Both ion species are in the trap but, due to the different qubit frequencies, only $^{43}\text{Ca}^+$ is excited, so we use the calcium bright-state population $p_{\uparrow s} = p_{\uparrow s \downarrow z} + p_{\uparrow s \uparrow z}$ for diagnostics. We measure this population after a π -pulse on the BSB (blue star) and on the red sideband (RSB) (red star) of each mode. The temperature can be estimated from the ratio r of the RSB population inversion to the BSB population inversion: $\bar{n} = r/(1-r)$ [185]. The population measurements are corrected for readout errors, measured concurrently (black stars). The data in the figure correspond to $\bar{n}_{\text{ax,ip}} = 0.07(1)$ and $\bar{n}_{\text{ax,oop}} = 0.030(7)$.

6.1.5 Heating rate

We measure heating rates of $\dot{n}_{\text{ax,ip}} = 120(10)$ quanta/s and $\dot{n}_{\text{ax,oop}} = 30(10)$ quanta/s for a two-ion $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal (figure 6.3).




Figure 6.3: Heating rates of the axial modes of a $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal. We prepare $|\downarrow_S \downarrow_Z\rangle$, wait for a given time, then measure the temperature using the method in figure 6.2. We find $\dot{n}_{\text{ax,ip}} = 120(10)$ quanta/s and $\dot{n}_{\text{ax,oop}} = 30(10)$ quanta/s. Shaded regions indicate the 1σ uncertainties of the fits.

6.1.6 Single-qubit randomised benchmarking

We measure the error associated with single-qubit rotations using randomised benchmarking, generating the single-qubit Clifford group via $\pm\frac{\pi}{2}$ x -, y - and z -rotations. The z -rotations are implemented via software phase shifts on subsequent x - and y -pulses. The decompositions of all 24 elements are listed in appendix C (‘XYZ’ decomposition). We set a target error of $< 10^{-3}$ for single-qubit $\frac{\pi}{2}$ -rotations to ensure that their contributions do not dominate our measured two-qubit gate errors, and remain no larger than state-preparation and measurement errors.

We perform standard (non-interleaved) randomised benchmarking sequences of single-qubit Clifford operations on $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ simultaneously, using microwave and RF pulses. This can be thought of as benchmarking the subset of the two-qubit Clifford group which includes only single-qubit operations on each ion, or — since the




Figure 6.4: Standard randomised benchmarking of simultaneous single-qubit Clifford gates on the $^{43}\text{Ca}^+$ stretch qubit and the $^{88}\text{Sr}^+$ Zeeman qubit, using microwaves on the trap blade and RF on the antenna. The respective $\frac{\pi}{2}$ -pulse times are $1.7 \mu\text{s}$ and $2.61 \mu\text{s}$. The ions are dark-resonance cooled, but not sideband cooled, before preparation into $|\downarrow_S \downarrow_Z\rangle$. There are 15 randomisations per sequence length, for sequence lengths $m = [1, 200, 400, 600, 800, 1000]$. We fit a two-qubit decay model $F(m) = Ap_C^m + \frac{1}{4}$ to the overall sequence fidelity including both qubits (dark blue line). From this, we calculate an error per two-qubit Clifford (restricted to two-qubit Cliffords which consist only of single-qubit operations on each ion) of $\epsilon_C^{(2)} = 1.19(5) \times 10^{-3}$. Since there are, on average, 1.42 $\frac{\pi}{2}$ -pulses per Clifford per ion in the ‘XYZ’ decomposition (appendix C), from this we estimate an average error per $\frac{\pi}{2}$ -rotation of $4.3(2) \times 10^{-4}$. The pink and light-blue lines are fits of the form $F(m) = Ap_C^m + \frac{1}{2}$ to subsets of the data corresponding to the sequence fidelities for each ion individually. We extract an error per calcium Clifford $\epsilon_C^{\text{Ca}} = 5.1(8) \times 10^{-4}$ and, for strontium, $\epsilon_C^{\text{Sr}} = 6.6(2) \times 10^{-4}$, corresponding to $\frac{\pi}{2}$ -pulse errors of $3.6(6) \times 10^{-4}$ and $4.7(1) \times 10^{-4}$ respectively.

two qubits are distinguishable and individually addressable — as simultaneous single-qubit benchmarking of each ion. Correspondingly, we can fit a decay curve either to the overall sequence fidelity or to a subset of the data pertaining to just one of the qubits. The fitted curve in each case has the form $F(m) = Ap_C^m + B_n$, where $B_n = 1/2^n$ depends on the number of qubits n . The error per n -qubit Clifford gate is given by $\epsilon_C = (1 - p_C)/\alpha_n$ where $\alpha_n = 2^n/(2^n - 1)$. The results are shown in figure 6.4; the two-qubit Clifford error agrees well with the sum of the individual $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ errors, and we estimate an average error per $\frac{\pi}{2}$ -rotation of $4.2(6) \times 10^{-4}$.

We also check the performance of single-qubit $\frac{\pi}{2}$ -rotations implemented with the Raman lasers using interleaved randomised benchmarking. Figure 6.5 shows a representative data set where the gate-under-test is a $\frac{\pi}{2} x$ -rotation on the calcium clock qubit. We measure an error of $5.8(6) \times 10^{-4}$ per interleaved $\frac{\pi}{2} x$ -gate, using a motionally-sensitive beam pair, and with a strontium ion also present in the trap.




Figure 6.5: Interleaved single-qubit randomised benchmarking on the calcium clock qubit using the motionally-sensitive Raman beams $R_{\parallel}$ and R_{H2} , whilst a strontium ion is also present in the trap. There are 50 randomisations per sequence length. The $\frac{\pi}{2}$ -pulse length is $5.0 \mu\text{s}$. We fit decay curves $F(m) = Ap_C + \frac{1}{2}$ and $F^{\text{int}}(m) = A^{\text{int}}p_C^{\text{int}} + \frac{1}{2}$, and calculate the error per $\frac{\pi}{2} x$ -gate $\epsilon_X = \frac{1}{2}(1 - p_C^{\text{int}}/p_C) = 5.8(6) \times 10^{-4}$.

6.1.7 Pulse shaping

Amplitude-shaping the gate laser pulses over a few motional periods is necessary to suppress errors due to unwanted carrier couplings (section 3.3.5.3). We shape rising and falling edges for all gates using the function $\sin^2(\pi t/2t_s)$ where $t_s = 2 \mu\text{s}$. Only one driving beam is shaped in each case — the second is left as a rectangular pulse which brackets the shaped beam, to reduce sensitivity to any timing jitter.

6.1.8 SPAM correction

Imperfections in the readout process mean that the measured bright and dark populations do not reflect the true populations $p_\uparrow$ and $p_\downarrow$ in each qubit state, for each ion. The two are related via the matrix M , such that

$$\begin{pmatrix} p_{\text{bright}} \\ p_{\text{dark}} \end{pmatrix} = M \begin{pmatrix} p_\uparrow \\ p_\downarrow \end{pmatrix}, \quad \text{where} \quad M = \begin{pmatrix} 1 - \epsilon_\uparrow & \epsilon_\downarrow \\ \epsilon_\uparrow & 1 - \epsilon_\downarrow \end{pmatrix}, \quad (6.1)$$

and $\epsilon_{\uparrow(\downarrow)}$ is the probability that an ion in the $|\uparrow\rangle$ ($|\downarrow\rangle$) state is wrongly measured as dark (bright). Measured populations can be corrected for known errors $\epsilon_{\uparrow,\downarrow}$ (measured within a few hours of the relevant data) by applying the inverse of M . For a mixed-species crystal, $(M_{\text{Ca}} \otimes M_{\text{Sr}})^{-1}$ is applied. Typical errors are listed in table 6.1.

qubit	$\epsilon_\uparrow$	$\epsilon_\downarrow$
$^{43}\text{Ca}^+$ stretch	0.13(10)%	0.16(5)%
$^{43}\text{Ca}^+$ clock	0.3(1)%	0.35(9)%
$^{88}\text{Sr}^+$ Zeeman	0.35(8)%	0.61(8)%

Table 6.1: Results of a typical SPAM-error measurement. Here, each state is prepared and read out 10,000 times. For correcting populations we combine many of these measurements, to reduce statistical uncertainty.

6.1.9 Spin dephasing

Dephasing of the spin states involved in the two-qubit gates over the gate duration will contribute to the overall error. From measurements of the coherence time of the magnetic-field-sensitive calcium stretch qubit, we estimate that this contribution is negligible ($< 10^{-4}$) for typical gate times. Mølmer-Sørensen gates may be implemented using the calcium clock qubit characterised in chapter 5, thereby suppressing the effect of magnetic-field noise for one of the two ion species.

6.2 Light-shift gate

We perform a light-shift ($\sigma_z \otimes \sigma_z$) gate on the stretch qubit in $^{43}\text{Ca}^+$ and the Zeeman qubit in $^{88}\text{Sr}^+$. These results are published in [79], and are also discussed in [177]. The technique was first used in a previous version of apparatus B to entangle two different isotopes of calcium [13]. A proof-of-principle demonstration with $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ is discussed in [150].

Two Raman beams, with a frequency difference close to an axial motional mode frequency, form a polarisation ‘travelling standing wave’ in which the ions sit. Motion in the chosen mode is off-resonantly excited, resulting in the acquisition of a geometric phase, dependent on the ions’ internal states, and thereby creating entanglement. The theory behind the mechanism is discussed in more detail in section 3.3.4.

6.2.1 Raman beams

The 397 nm $4\text{S}_{1/2} \leftrightarrow 4\text{P}_{1/2}$ transition in $^{43}\text{Ca}^+$ and the 408 nm $5\text{S}_{1/2} \leftrightarrow 5\text{P}_{3/2}$ transition in $^{88}\text{Sr}^+$ are close enough in frequency that Raman transitions may be driven in both species using a single wavelength. Since the required frequency difference between the two beams is independent of the two qubit frequencies ($f_0^{\text{Ca}} = 2.874$ GHz, $f_0^{\text{Sr}} = 409$ MHz), this means we can use the same beam pair to drive the light-shift




Figure 6.6: Light-shift gate beams. R_V and $R_{\parallel}$ come from a single laser with wavelength $c/f_L = 402$ nm, and have a frequency difference equal to the gate detuning δ_g plus the motional mode frequency f_{ax} . Each beam couples to both qubit states in each ion via a light-shift interaction, providing a qubit-state-dependent force on both species. Their differential wavevector $\Delta \mathbf{k}$ lies parallel to the trap z -axis, so they couple only to the axial motional modes. The detunings from $S \leftrightarrow P$ transitions in $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ are $\Delta_{Ca} = -9.0$ THz, $\Delta_{Sr} = +11.2$ THz.




Figure 6.7: Simulated heating and scattering errors vs. Raman detuning for the $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ LS gate on the axial oop mode. The 'X' marks the minimum total error. Adapted from [150] for the approximate parameters used in this thesis.

gate on both species simultaneously. Moreover, the required frequency difference is close to a motional mode frequency ($f_{\text{ax,ip}} = 1.49$ MHz or $f_{\text{ax,oop}} = 2.91$ MHz), which is easily bridged using an AOM, so both beams may be derived from a single laser.

We use the Raman beams R_V and $R_{\parallel}$, described in section 4.7.2, to drive this gate. Their geometries and frequencies are shown in figure 6.6. The detunings Δ_{Ca} and Δ_{Sr} of the Raman beams from $S \leftrightarrow P$ transitions in $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ determine the rate of photon scattering, which contributes to the gate error (section 3.3.5.1). Since the laser power is limited, the detuning also determines the fastest possible gate speed. The slower the gate, the more the ions are heated over its duration, which also contributes to the gate error (section 3.3.5.2). The detunings are therefore chosen to approximately minimise the total error due to Raman photon scattering and due to heating, as shown in figure 6.7.

The Lamb–Dicke parameters for each ion in the Raman beams are given in table 6.2.

ion	η_{ip}	η_{oop}
$^{43}\text{Ca}^+$	0.090	0.127
$^{88}\text{Sr}^+$	0.124	-0.045

Table 6.2: Lamb–Dicke parameters for $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ in the Raman beams.

6.2.2 Ion spacing

As discussed in section 3.3.4, the relative phase ϕ_M of the driving force on each ion is determined by the distance between the two ions in the standing wave, as well as by their internal states. For a maximally efficient gate, the ion spacing should be equal to an integer or half-integer multiple of the effective wavelength $\lambda_{\text{eff}} = 402$ nm/ $\sqrt{2}$. For convenient motional mode frequencies, we choose the latter; in this case,

for a gate on the oop mode, the motional excitation is maximised for the even-parity spin states ($|\uparrow_S\uparrow_Z\rangle$ and $|\downarrow_S\downarrow_Z\rangle$) and minimised for the odd-parity states ($|\uparrow_S\downarrow_Z\rangle$ and $|\downarrow_S\uparrow_Z\rangle$), and vice versa for the ip mode (see figure 3.8).

We calibrate the ion spacing by preparing the state $|\downarrow_S\downarrow_Z\rangle$ and switching on the gate beams, on resonance with each of the axial mode frequencies in turn. Since $|\downarrow_S\downarrow_Z\rangle$ is an even-parity state, we set a half-integer spacing by maximising the excitation of the oop mode and minimising that of the ip mode (figure 6.8).




Figure 6.8: Ion spacing calibration. After preparing $|\downarrow_S\downarrow_Z\rangle$, we excite motion using gate beams resonant with each motional mode. By adjusting the endcap voltages to maximise excitation of the oop mode for this starting state, we set the ion spacing to $3.57 \mu\text{m}$, corresponding to $12.5\lambda_{\text{eff}}$ where $\lambda_{\text{eff}} = 402 \text{ nm}/\sqrt{2}$ is the effective wavelength of the Raman travelling standing wave. The relative excitation of the ip mode is $\sim 1\%$ (note the different x -axis scales).

6.2.3 Gate implementation

6.2.3.1 Pulse sequence

The light-shift gate generates a maximally entangled state from an input state consisting of a superposition of all four z -basis states. Therefore, our first step is to prepare the state $|\downarrow_S \downarrow_Z\rangle$, and to apply single-qubit $\frac{\pi}{2}$ x -rotations to each ion. For the light-shift gate operator, we do not require phase coherence between the gate lasers and single-qubit operations, so these may be driven using microwaves and RF.

We then implement a two-loop gate with a spin-echo π x -pulse between the loops. This spin-echo pulse is also implemented using microwaves and RF. As discussed in section 3.3.4, its purpose is multifold: firstly, it symmetrises the gate against unequal Rabi frequencies. Since the light-shift gate Hamiltonian commutes with z -rotations, the spin-echo sequence also cancels any residual unwanted light shifts due to Raman carrier coupling or unequal illumination, and reduces sensitivity to correlated magnetic-field noise over the gate duration.

Finally, we again apply $\frac{\pi}{2}$ x -rotations, such that the total sequence is

$$\left[\left(\frac{\pi}{2} \right)_{\phi_{\text{RF}}} \otimes \left(\frac{\pi}{2} \right)_{\phi_{\text{MW}}} \right] \cdot \sqrt{U_{\text{LS}}} \cdot [(\pi)_{\phi_{\text{RF}}} \otimes (\pi)_{\phi_{\text{MW}}}] \cdot \sqrt{U_{\text{LS}}} \cdot \left[\left(\frac{\pi}{2} \right)_{\phi_{\text{RF}}} \otimes \left(\frac{\pi}{2} \right)_{\phi_{\text{MW}}} \right], \quad (6.2)$$

where $(\theta)_\phi$ represents a rotation by angle θ around the direction at azimuthal angle ϕ , and

$$U_{\text{LS}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \pm i & 0 & 0 \\ 0 & 0 & \pm i & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad (6.3)$$

is the light-shift gate propagator as defined in section 3.3.4. Starting in $|\downarrow_S \downarrow_Z\rangle$, with perfect gates, this sequence would generate the Bell state $\frac{1}{\sqrt{2}} (|\uparrow_S \uparrow_Z\rangle \pm i |\downarrow_S \downarrow_Z\rangle)$ (where the sign depends on the motional mode).

6.2.3.2 Driving force phase

For calibrating the gate, we set the phase of the driving force on each qubit state to be the same at the beginning of the first loop and at the beginning of the second loop. In this way, we return to the initial position in phase space only if two full, closed loops are completed, which makes diagnostics simpler because the dynamics look the same regardless of the number of loops. To achieve this, the laser phase at the start of the second pulse is shifted by $-(\pi + \delta_g t_{\text{delay}})$ relative to the first, where t_{delay} is the total time between the starts of the two pulses and the shift of $-\pi$ cancels the π phase shift introduced by the spin-echo pulse.

When measuring the gate fidelity we use a first-order Walsh modulation, meaning that the phase of the second pulse is shifted by π relative to the first, so that the total phase offset becomes $\pi - (\pi + \delta_g t_{\text{delay}})$. As discussed in section 3.3.5.4, this reduces errors due to imperfect closure of loops in phase space.

6.2.3.3 Gate parameters

We perform a light-shift gate on the axial oop mode with a gate detuning of $\delta_g = -40$ kHz. There are 60 mW in each of the two Raman beams, which both have a waist radius of $27 \mu\text{m}$ at the trap centre.

Figure 6.9 shows the dynamics we observe for different lengths of the two gate pulses, without Walsh modulation. The dashed line indicates the gate time $t_g = 2/\delta_g = 49.2 \mu\text{s}$, where the length of each gate pulse is $t_g/2$. At this point, we have generated the Bell state $\frac{1}{\sqrt{2}}(|\uparrow_s \uparrow_z\rangle - i|\downarrow_s \downarrow_z\rangle)$, so the populations $p_{\uparrow_s \uparrow_z} = p_{\downarrow_s \downarrow_z} = 0.5$.

Alternatively, we can also implement a gate on the ip mode at a gate detuning of 80 kHz, with 84 mW in each Raman beam. The gate duration in this case is $25 \mu\text{s}$.




Figure 6.9: Light-shift gate dynamics for different Raman pulse lengths. The gate lasers are switched on for a time $t/2$, followed by a spin-echo π -pulse on each qubit and a second gate pulse of duration $t/2$. When $t = t_g$, each gate pulse implements the propagator $\sqrt{U_{LS}}$ and the final measured state is $\frac{1}{\sqrt{2}}(|\uparrow_S \uparrow_Z\rangle - i|\downarrow_S \downarrow_Z\rangle)$, so the populations $p_{\uparrow_S \uparrow_Z}$ and $p_{\downarrow_S \downarrow_Z}$ after this time are both equal to 0.5. The phase of the driving force at the beginning of the second gate pulse matches the phase at the beginning of the first (no Walsh modulation), to improve sensitivity to mis-set parameters for calibration. Single-qubit rotations are implemented using microwaves for calcium and RF for strontium. Solid lines are from a simulation of the corresponding gate parameters, written by V. M. Schäfer, and are not a fit to the data. Populations are corrected for readout errors.

6.2.4 Partial state tomography

A standard and commonly used method for estimating the fidelity of a two-qubit gate operation is to measure the fidelity with which it can be used to create a Bell state, by partial state tomography (PST) [186, 149, 12, 60, 190, 195].

For an output density matrix ρ , the fidelity with which we create a particular Bell state $|\Phi_{\phi_0}\rangle = \frac{1}{\sqrt{2}} (|\uparrow_S \uparrow_Z\rangle + e^{i\phi_0} |\downarrow_S \downarrow_Z\rangle)$ is given by

$$\begin{aligned} F &= \langle \Phi_{\phi_0} | \rho | \Phi_{\phi_0} \rangle \\ &= \frac{1}{2} [\rho_{\uparrow\uparrow, \uparrow\uparrow} + \rho_{\downarrow\downarrow, \downarrow\downarrow} + e^{i\phi_0} \rho_{\uparrow\uparrow, \downarrow\downarrow} + e^{-i\phi_0} \rho_{\downarrow\downarrow, \uparrow\uparrow}] \\ &= \frac{1}{2} [\rho_{\uparrow\uparrow, \uparrow\uparrow} + \rho_{\downarrow\downarrow, \downarrow\downarrow} + 2|\rho_{\downarrow\downarrow, \uparrow\uparrow}| \cos(\phi_0 - \arg(\rho_{\downarrow\downarrow, \uparrow\uparrow}))], \end{aligned} \quad (6.4)$$

where we have dropped the subscripts S and Z for clarity. For some applications it may suffice to produce any generic Bell state, but ideally we should always produce a Bell state with a known phase, $\arg(\rho_{\downarrow\downarrow, \uparrow\uparrow}) = \phi_0$, such that the cosine term is equal to 1.

To measure this fidelity, we can measure the populations in $|\uparrow\uparrow\rangle$ and $|\downarrow\downarrow\rangle$ after the gate operation:

$$\begin{aligned} \text{populations} &= p_{\uparrow\uparrow} + p_{\downarrow\downarrow} \\ &= \rho_{\uparrow\uparrow, \uparrow\uparrow} + \rho_{\downarrow\downarrow, \downarrow\downarrow}. \end{aligned} \quad (6.5)$$

We also measure the ‘parity’ of the state after a single-qubit rotation by angle $\frac{\pi}{2}$ about an axis ϕ in the x - y plane:

$$\begin{aligned} \text{parity}(\phi) &= p'_{\uparrow\uparrow} + p'_{\downarrow\downarrow} - (p'_{\uparrow\downarrow} + p'_{\downarrow\uparrow}) \\ &= \rho_{\uparrow\downarrow, \downarrow\uparrow} + \rho_{\downarrow\uparrow, \uparrow\downarrow} - 2|\rho_{\downarrow\downarrow, \uparrow\uparrow}| \cos(2\phi - \arg(\rho_{\downarrow\downarrow, \uparrow\uparrow})), \end{aligned} \quad (6.6)$$

where p' indicates a population after rotation about ϕ . By scanning the parity angle

ϕ , we measure oscillations in the parity with a periodicity of 2ϕ . The contrast of these oscillations is equal to $2|\rho_{\downarrow\downarrow,\uparrow\uparrow}|$. Therefore, we find the fidelity via

$$F = \frac{1}{2} (\text{populations} + \text{parity contrast}). \quad (6.7)$$

Figure 6.10 shows an example of a fidelity measurement using PST for a light-shift gate on the ip mode. From this particular measurement, we find a fidelity of 99.7(1)% by fitting a sinusoidal function to the parity oscillations. However, it is possible to overestimate the gate fidelity using this method, by trading off between the fitted contrast and the parity angle offset ϕ_0 . In addition, any uncertainty in ϕ_0 would be detrimental during long sequences of single- and two-qubit gates which require phase coherence, so ideally we would like to measure the gate fidelity using a method which is sensitive to this offset. We therefore use full parity scans such as this only to determine ϕ_0 , and not for final fidelity measurements.

For final fidelity measurements, we measure only two points of the scan which we expect to correspond to maximum and minimum parity, extracting the ‘contrast’ from the difference of these two measurements. This ensures that we are sensitive to any changes in ϕ_0 , in a similar manner as for the two-point Ramsey experiments in chapter 5, and is equivalent to measuring the overlap with a particular Bell state (defined by ϕ_0), rather than a general one. A small deviation $\delta\phi$ of the true Bell-state phase from the predicted one will result in a loss of measured fidelity equal to $\frac{1}{4}(\delta\phi)^2$.

Figure 6.11 shows the results of seven such two-point PST fidelity measurements for a gate on the oop mode, taken over two days. Combining all these measurements and correcting for readout errors results in a fidelity of 99.8(1)%. Similar measurements for the ip mode gate also give a fidelity of 99.8(1)%. Errors arising from imperfections in the single-qubit spin-echo π -pulses and the $\frac{\pi}{2}$ -rotations before and after the gate are not corrected for.




Figure 6.10: Light-shift gate parity scan, for a gate on the ip mode. The light-shift gate G is implemented via two Raman pulses separated by a spin-echo π -pulse on each qubit. Single-qubit $\frac{\pi}{2}$ -pulses rotate the state before and after the gate. Another ‘analysis’ pulse, with a variable phase, ϕ , is used for measuring parity oscillations. Single-qubit rotations are implemented via microwaves for calcium and RF for strontium. The solid line is a sinusoidal fit to the parity, $1 - 2(p_{\uparrow_S \downarrow_Z} + p_{\downarrow_S \uparrow_Z})$. The right-hand panel shows a measurement of the populations without the extra analysis pulse. The fidelity is equal to $\frac{1}{2}(\text{populations} + \text{parity contrast})$, where ‘populations’ is the total population in $|\uparrow_S \uparrow_Z\rangle$ and $|\downarrow_S \downarrow_Z\rangle$ in the right-hand panel, and the ‘parity contrast’ is the contrast of the parity oscillations in the left-hand panel. From this measurement we calculate a fidelity of 99.7(1)%.




Figure 6.11: Bell-state fidelity measurements for the oop mode gate, taken over two separate days. The parity is measured in each case with a ‘two-point’ measurement, using analysis pulses with phases $\phi = 45^\circ$ and 135° . Each fidelity is readout-corrected using measurements taken on the same day. The total number of population and parity measurements is 165,000, and the total number of readout-error measurements is 50,000. The mean fidelity is 99.8(1)%.

6.2.5 Randomised benchmarking

Measurement of the gate fidelity by PST, although simple to implement, has a few drawbacks. Firstly, since the gate error is of a comparable size to the SPAM errors, the uncertainty in the measurement is dominated by statistical uncertainty in the SPAM errors, and small variations in the SPAM errors can lead to unphysical gate fidelities of $> 100\%$. Although alternative statistical techniques such as maximum likelihood estimation can account for this to some extent (see e.g. [166]), another drawback is that PST only characterises the performance of the gate for a particular input state, $|\downarrow_S \downarrow_Z\rangle$. A more complete measure of the fidelity can be made using two-qubit interleaved randomised benchmarking.

Two-qubit randomised benchmarking is more technically demanding to implement because it involves long sequences of gates and requires individual addressing of the two ions. However, it provides a measure of the fidelity in the computationally rel-

evant context of a long sequence of operations with random input states, and it is not limited by uncertainties in SPAM errors. In addition, the number of measurements required for randomised benchmarking of larger numbers of qubits scales only polynomially with the number of qubits, rather than exponentially, as is the case for comparable techniques such as process tomography.

The randomising gates are sampled uniformly from the two-qubit Clifford group, which may be implemented using only single-qubit $\pm\frac{\pi}{2}$ x -, y - and z -rotations and the two-qubit light-shift gate. Circuits are optimised for efficiency within Cliffords, leading to an average of 1.5 entangling gates per Clifford ([42], see appendix C). For our decomposition, the average number of single-qubit $\frac{\pi}{2}$ x - and y -rotations per two-qubit Clifford is 6.4 (including both ions) as well as 1.3 $\frac{\pi}{2}$ z -rotations. As for the single-qubit benchmarking earlier in this chapter, z -rotations are implemented via software phase shifts and contribute negligible error. The two-qubit identity operation $\mathbb{1} \otimes \mathbb{1}$ is implemented instantaneously as ‘do nothing’. Single-qubit identity operations are implemented as delays of the length of the coinciding single-qubit Clifford on the other ion.

We implement interleaved two-qubit benchmarking with the light-shift gate, G , as the gate-under-test. G includes the two Raman pulses and the microwave/RF spin-echo pulse, but not the $\frac{\pi}{2}$ -rotations, since they may be absorbed into the surrounding randomising gates. Results and example sequences are shown in figure 6.12.

The sequences are between 16 and 32 ms in duration, including ≈ 16 ms of state-preparation, cooling and readout time. The ions are cooled before, but not during, each sequence. The longest sequences ($m = 60$, interleaved) include a total of 121 Clifford gates, requiring an average of 151.5 entangling gates. On these timescales, gate errors due to heating of the ion and due to duty-cycle effects in amplifiers become non-negligible. The measured error per interleaved light-shift gate, ϵ_G , therefore increases slightly with the sequence length. Including sequence lengths up to 20, we




Figure 6.12: Interleaved randomised benchmarking of the $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ light-shift gate, G . We perform sequences of lengths $m = [1, 10, 20, 30, 40, 50, 60]$, with $N = 100$ shots of $k = [120, 120, 120, 120, 105, 60, 15]$ randomisations per sequence length. The operations C_i are chosen at random from the two-qubit Clifford group. The final gate in each sequence, $C_{(G)}^{-1,P}$, inverts all previous operations and rotates each qubit into a randomly-chosen eigenstate of σ_z , before measurement. The left-hand panel shows the average fidelity $F(m)$ of interleaved and non-interleaved sequences for each length m . We fit a decay curve $F(m) = Ap_C^m + \frac{1}{4}$ to each data set, from which we can calculate the error ϵ_G per interleaved G gate. The right-hand panel shows the inferred ϵ_G as a function of the maximum sequence length used in the fit. We find $\epsilon_G = 2.9(7) \times 10^{-3}$ ($3.8(3) \times 10^{-3}$) from sequences including up to 20 (60) interleaved G gates. Due to thermal effects, the error increases as longer sequences are included. Error bars on ϵ_G are from parametric bootstrapping. Using all sequences up to length 60, we infer an average error per two-qubit Clifford operation $\epsilon_C^{(2)} = 8.3(2) \times 10^{-3}$, consistent with the errors of the constituent single- and two-qubit operations.

measure $\epsilon_G = 2.9(7) \times 10^{-3}$, or a gate fidelity of 99.71(7)%. Including all sequences up to length 60, we measure $\epsilon_G = 3.8(3) \times 10^{-3}$, or a gate fidelity of 99.62(3)%.

From the fitted decay curve we can also extract the error per two-qubit Clifford operation, $\epsilon_C = 8.3(2) \times 10^{-3}$; this is consistent with the error that would be predicted from the average number of constituent light-shift gates and single-qubit rotations, and their associated errors.

The whole benchmarking data set includes a total of 132,000 qubit population measurements with durations between 16 and 32 ms. The PST measurements in figure 6.11 required 265,000 population measurements, including SPAM measurements, each of which take ≈ 16 ms due to the same state preparation, cooling and read-out. Benchmarking results in a lower uncertainty on the gate fidelity in a more time-efficient manner.

The leading contribution to the gate error, for a single gate, is expected to be ion heating during the gate ($\approx 4 \times 10^{-4}$ infidelity). Photon scattering and Kerr cross-coupling to the radial modes contribute $\approx 2 \times 10^{-4}$. The error per single-qubit $\frac{\pi}{2}$ -rotation is $\approx 4 \times 10^{-4}$, as measured by randomised benchmarking in figure 6.4, although this does not necessarily contribute linearly to the measured gate error. We compensate the trap in the vertical direction with a precision of ≈ 0.3 V/m. From measurements of the gate fidelity using randomised benchmarking before and after a correction to the vertical compensation voltage, we find that a change of this size can change the gate error by $2(5) \times 10^{-4}$. For the longer sequences used in randomised benchmarking, heating of the axial ip mode becomes the dominant source of error; a temperature of $\bar{n}_{\text{ax,ip}} \approx 1.8$ after 16 ms contributes 3×10^{-3} error for the final gate in the sequence.

6.2.6 Gate set tomography

Where randomised benchmarking and PST can offer information about the magnitude of gate errors, the technique of gate set tomography (GST)[23] can also provide insight into their nature. This protocol involves state preparation followed by repetitions of sequences of gates, and measurement. Here, the sequences (known as ‘germs’) are not randomly generated but are instead chosen to amplify sensitivity to particular error sources.

Tomography of the gate set $\{G, +\frac{\pi}{2}x, +\frac{\pi}{2}z\}$ was implemented in our system by D. P. Nadlinger. From this characterisation we extract a light-shift gate error of $\epsilon_G = 6(3) \times 10^{-3}$, where the largest contribution is a coherent over-rotation in z . The extra insight offered by GST is invaluable in determining the most appropriate approach to error-correction, but comes at the cost of extra computational complexity in the analysis and experimental resources which scale exponentially with the number of qubits; this implementation required 386,000 population measurements with durations between 16 and 23 ms, and the non-Markovianity of the gate error limits the uncertainty. For more details see [79]. All three fidelity measurements described in this chapter – PST, randomised benchmarking and GST – are consistent at the $\approx 1\sigma$ level.

6.3 Mølmer–Sørensen gates

We also perform mixed-species Mølmer–Sørensen ($\sigma_\phi \otimes \sigma_\phi$) gates on various combinations of qubits; an advantage of the MS gate over the light-shift gate is that it may be used directly on the calcium clock qubit.

To perform an MS gate requires driving red and blue sidebands of each qubit simultaneously. Analogously to the light-shift gate, we drive loops in motional phase-space which result in the acquisition of a geometric phase, depending on the ions' spin states. In this case, however, the spin basis is not along the z -direction, but a direction in the x - y plane which is defined by the phases of the gate-driving lasers. The theory is discussed in more detail in section 3.3.3.

6.3.1 Raman beams

To drive an MS gate we apply two frequency tones simultaneously to the $R_{\parallel}$ AOM, at $\pm(f_{\text{ax}} + \delta_g)$ from the centre frequency, where δ_g is the gate detuning and f_{ax} is the frequency of the chosen motional mode. With the Raman lasers at the same detuning as for the light-shift gate ($\Delta_{\text{Ca}} = -9.0$ THz, $\Delta_{\text{Sr}} = +11.2$ THz), we may use the $R_{\parallel}$ beam for both ion species.

Unlike for the light-shift gate, however, the Raman frequency splitting required for the MS gate depends on the qubit frequency, so only one of the two gate beams may be shared between both species. In conjunction with $R_{\parallel}$, we therefore use the R_{HSr} beam, with a frequency 409 MHz lower than $R_{\parallel}$, to drive the MS gate on the strontium Zeeman qubit, and R_{H2} (3.200 GHz below $R_{\parallel}$) for the calcium clock qubit. Alternatively, we can choose to use R_{H} (2.874 GHz below $R_{\parallel}$) to drive the gate on the calcium stretch qubit. The relevant beams are shown in figure 6.13. For the following work, the $R_{\parallel}$ beam waist radius was reduced to 16 μm and the radii of the R_{H} beams were reduced to 22 μm , unless otherwise specified.




Figure 6.13: Mølmer-Sørensen gate beams. $R_{\parallel}$ includes two tones at $\pm(f_{\text{ax}} + \delta_g)$. $R_{\parallel}$ and R_{HSr} come from the ‘master’ Raman laser and this beam pair drives the gate on strontium, with a frequency splitting equal to the strontium Zeeman qubit frequency $f_0^{\text{Sr},Z} = 409$ MHz. R_H and R_{H2} are from the ‘slave’ laser, in order to span the calcium qubit frequencies $f_0^{\text{Ca},S} = 2.874$ GHz and $f_0^{\text{Ca},C} = 3.200$ GHz. We choose either $R_{\parallel} + R_H$ or $R_{\parallel} + R_{H2}$, depending on which qubit we wish to use for the gate. As for the light-shift gate, the difference wavevector of all gate-driving beam pairs lies parallel to the trap z -axis, and the detunings from $S \leftrightarrow P$ transitions in $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ are $\Delta_{\text{Ca}} = -9.0$ THz, $\Delta_{\text{Sr}} = +11.2$ THz.

6.3.2 Bichromat balance calibration

For each ion, it is important that the red and blue sideband Rabi frequencies are of equal magnitude. Measuring the power in each beam using a photodiode provides a good starting point, but a much more sensitive method is to measure the light shift on the carrier frequency when the gate beams are switched on, and adjust the laser powers accordingly. When the sideband Rabi frequencies are balanced, the two-beam light shifts on the carrier frequency should cancel exactly. An example of this calibration experiment is shown in figure 6.14.




Figure 6.14: Example ‘bichromat balance’ calibration for calcium. A gate pulse of length t is sandwiched in a Ramsey sequence on the calcium qubit. The $\frac{\pi}{2}$ -rotations may be done using microwaves or using the Raman lasers, as long as there is phase coherence between the two pulses. The second $\frac{\pi}{2}$ -rotation is performed with a phase of 90° relative to the first, so we expect the final population in $|\uparrow_c\rangle$ to be constant at 0.5. If there is a light shift of frequency f_{LS} on the calcium qubit during the gate pulse, we will see oscillations at this frequency in the population $p_{\uparrow c}$. The blue curve indicates a light shift of 797(5) Hz. By adjusting the power in each of the two tones in the $R_{\parallel}$ beam, we reduce this light shift such that no Ramsey fringes are seen up to 3 ms (red curve). For this calibration, the gate detuning is set to 100 kHz to avoid any off-resonant drive of the sidebands. If a strontium ion is also in the trap, it must be ‘hidden’ from the gate beams by shelving it in the $D_{5/2}$ level before applying the gate pulse.

For strontium, the calibration is performed in a spin-echo sequence, where the gate beams are only switched on during the first arm, in order to distinguish light shifts from decoherence of the field-sensitive qubit. If a calcium ion is also present in the trap, it may be ‘hidden’ in one of the unused qubit states.

Using this procedure to balance the sideband Rabi frequencies relies on the absence of any single-beam light shifts. In reality, the single-beam shift cannot simultaneously be nulled for the clock and for the stretch qubit, due to the different matrix elements. We choose to null the shift on the stretch qubit (resulting in a ~ 1 kHz shift on the clock), because it happens that zero shift on the clock transition would leave us with no Rabi frequency for sideband cooling on the stretch transition. For gates on the clock qubit, we set the powers based on the gate dynamics and compensate for the single-beam light shift with an extra phase offset on subsequent laser pulses.

6.3.3 Gate implementation

The MS gate generates a maximally entangled state from the starting state $|\downarrow\downarrow\rangle$, so no preceding single-qubit $\frac{\pi}{2}$ -rotations are required for this purpose. As for the light-shift gate, we implement a two-loop gate with a first-order Walsh modulation. To achieve this, the phases of the two $R_{\parallel}$ tones must be adjusted by $\pi \pm \delta_g t_{\text{delay}}$ since they have opposite detunings from resonance. Each of the two pulses implements the propagator $\sqrt{U_{\text{MS}}}$, where

$$U_{\text{MS}} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 0 & 0 & \pm i e^{-i(\phi_{S,1} + \phi_{S,2})} \\ 0 & 1 & \mp i e^{i(\phi_{S,2} - \phi_{S,1})} & 0 \\ 0 & \mp i e^{-i(\phi_{S,2} - \phi_{S,1})} & 1 & 0 \\ \pm i e^{i(\phi_{S,1} + \phi_{S,2})} & 0 & 0 & 1 \end{pmatrix} \quad (6.8)$$

is the MS gate propagator as defined in section 3.3.3. Starting in $|\downarrow\downarrow\rangle$, with perfect gates, this sequence would generate the Bell state $\frac{1}{\sqrt{2}} (\pm i e^{-i(\phi_{S,1} + \phi_{S,2})} |\uparrow\uparrow\rangle + |\downarrow\downarrow\rangle)$ (where the sign depends on the motional mode).

6.3.3.1 Gate parameters

The parameters for MS gates in this section can be found in table 6.3. All gates are two-loop gates with Walsh modulation, on the axial ip mode.

6.3.4 MS gate on the clock qubit

Taking advantage of the fact that the MS gate can be implemented directly on a field-insensitive qubit, we first perform a mixed-species gate on the calcium clock qubit and the strontium Zeeman qubit. Figure 6.15 shows an example of the dynamics we observe for different lengths of the two gate pulses. Again, the gate dynamics are observed without Walsh modulation to increase sensitivity to loop closures for

gate	δ_g (kHz)	Power (mW)				
		R_H	R_{H2}	R_{HSr}	$R_{ }$ (A)	$R_{ }$ (B)
Ca–Sr clock	40 kHz	-	18	8	120	120
Ca–Sr clock (slow)	16 kHz	-	22	11	170	170
Ca–Sr stretch	40 kHz	17	-	3	120	120
Ca–Ca clock	29 kHz	-	17	-	≈ 60	≈ 60
Ca–Ca stretch	29 kHz	5.1	-	-	≈ 60	≈ 60

Table 6.3: Parameters for the MS gates in this chapter, unless otherwise specified. The ‘slow’ Ca–Sr clock gate was performed before reducing the Raman beam waist radii of $R_{||}$ and of all the R_H beams from $27\text{ }\mu\text{m}$ to $16\text{ }\mu\text{m}$ and $22\text{ }\mu\text{m}$ respectively, so a smaller Rabi frequency is achieved for similar laser powers. The $R_{||}$ (A) and (B) powers are measured whilst driving only one of the two frequency tones; the total $R_{||}$ power with both tones on simultaneously is slightly less than the sum of the two, due to non-linearities in the AOM response. We choose to have the most power in $R_{||}$ as this is shared between both ions, and is not overlapped with other beams on beam splitters, so less power is wasted. The amount of photon scattering is slightly larger than for a more symmetric division of power, but the gain in gate speed outweighs this.

calibration and diagnostics.

We measure the fidelity of this gate using partial state tomography as for the light-shift gate, calculating the parity contrast from a two-point measurement. The results are shown in figure 6.16. From a combination of five separate measurements, we find a fidelity of 96.2(3)%. Also shown are the populations and the parity contrast from each measurement, from which it can be seen that the loss in fidelity is due mostly to a loss of parity contrast. The average population, $p_{\uparrow_C\uparrow_Z} + p_{\downarrow_C\downarrow_Z}$, is 99.4(2)%, whereas the inferred average parity contrast from all five measurements is 92.9(6)%.

Figure 6.17 shows three PST measurements, taken approximately 5 minutes apart, where the parity contrast is found by scanning the analysis phase across a full fringe, rather than from two points. Each measurement has a parity contrast which is much greater than that implied by the two-point measurement, but they have different




Figure 6.15: Dynamics of a $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ MS gate on the calcium clock qubit and the strontium Zeeman qubit without Walsh modulation, using the axial ip mode. The gate detuning in this case is $\delta_g = 36$ kHz, corresponding to a total gate time of $55 \mu\text{s}$. The beam powers for R_{H2} and R_{HSr} are 18 mW and 8 mW respectively, and there are 130 mW in each $R_{\parallel}$ tone (these parameters differ slightly from the final parameters used in fidelity measurements). Populations are readout-corrected.

parity offset angles. This indicates that in each case we create a Bell state with high fidelity, but the phase of the Bell state drifts between measurements, so that the phases at which the two-point measurements are made are no longer at the maximum and minimum parity. Unpredictability of this phase means that use of this gate in longer sequences requiring phase coherence would result in a rapid loss of fidelity. PST measurement runs with longer durations also result in a lower measured fidelity.

A change in the Bell-state phase can happen when there is an offset of the qubit frequencies from their nominal values during the gate, for example due to slow drift




Figure 6.16: Two-point PST fidelity measurements of a mixed-species MS gate on the calcium clock qubit and the strontium Zeeman qubit. The gate parameters can be found in the first row of table 6.3; the gate time is $50 \mu\text{s}$. The total number of population and parity measurements is 27,000, corrected using 20,000 measurements of the readout errors. For the MS gate, single-qubit ‘analysis’ pulses for the parity measurement are implemented using the gate lasers. Combining the data gives a mean fidelity of 96.2%. The mean populations are 99.4% and the mean parity contrast is 92.9% (solid lines). The loss in fidelity is due to variation in the parity contrast; the dashed lines indicate the standard deviations between the five experimental runs.




Figure 6.17: Three ‘full-fringe’ PST measurements taken approximately 5 minutes apart, showing the drift of the parity offset angle between measurements. Populations are readout-corrected. The fitted fidelities are 100(3)%, 100(2)% and 100(1)%.




Figure 6.18: Oscillations in the external magnetic-field strength, as measured via offsets on the calcium stretch qubit frequency, and the concurrently measured temperature of the coils. A 1 kHz frequency offset corresponds to a 0.4 mG field offset.

of the external magnetic-field strength (see section 3.3.5.5). For the input state $|\downarrow\downarrow\rangle$, the change depends on the sum of the qubit frequency offsets $\delta_{\text{Ca}}^q + \delta_{\text{Sr}}^q$ and increases for longer gate times. For a small magnetic-field offset, on the calcium clock qubit $\delta_{\text{Ca}}^q \approx 0$, but $\delta_{\text{Sr}}^q \neq 0$ due to the field-sensitivity of the Zeeman qubit. The timescale of the drifts would have to be longer than the duration of a single fidelity measurement ($\lesssim 1$ minute, depending on the number of shots), but comparable to the time between measurements.

Indeed, we see variations in the magnetic field at the ions over this same timescale; figure 6.18 shows the calcium stretch qubit frequency over a period of ≈ 1 hour, logged using the magnetic-field servo that we use to stabilise the magnetic field (section 4.4.2). Also shown are concurrent measurements of the temperature of the high-field coils¹, indicating that the oscillations are due to temperature variations. Plans for future experiments in this apparatus include temperature stabilisation in order to reduce this effect.

6.3.5 MS gates on other qubits

To confirm that the parity angle drifts are caused by the magnetic field, we use the same mechanism to drive same-species, $^{43}\text{Ca}^+ - ^{43}\text{Ca}^+$ gates on the clock qubit and on the stretch qubit, and monitor the angle over time. The results are shown in figure 6.19; as expected, we see no significant drift for the same-species clock qubit gate, where both qubits are field-insensitive². Included for comparison is the drift of the parity phase for mixed-species gates using the strontium Zeeman qubit and the calcium clock qubit. The amount of phase acquired depends on the magnitude of the frequency offset and on the total gate time, so slower gates (with smaller gate detunings) are more susceptible.

More interestingly, we note that the magnetic-field sensitivity of the strontium

¹BetaTHERM 10K3A542I thermistor

²For the same-species $^{43}\text{Ca}^+ - ^{43}\text{Ca}^+$ clock qubit MS gate we measure a fidelity of 99.5(3)%.




Figure 6.19: Parity-angle drifts for various gates. Longer gates (smaller detunings) are expected to be more sensitive to magnetic field offsets, and this is reflected in the results. Table 6.3 lists the parameters used for each gate. The average offset angle is different for each gate due to the different single-beam light shifts on each qubit pair.




Figure 6.20: Two-point PST fidelity measurements of a mixed-species MS gate on the calcium stretch qubit and the strontium Zeeman qubit. The gate parameters can be found in table 6.3. From a total of 30,000 population and parity measurements and 100,000 readout-error measurements, the gate fidelity is 99.6(2)%.

Zeeman qubit is 2.80 MHz/G, and that of the calcium stretch qubit is almost equal and opposite, at -2.36 MHz/G. For an MS gate on these qubits with input state $|\downarrow_S \downarrow_Z\rangle$, we would therefore expect the offsets to almost cancel ($\delta_{\text{Ca}}^q + \delta_{\text{Sr}}^q \approx 0$) and the parity phase to be stable once more. Figure 6.19 also shows this to be the case.

With this in-built robustness, therefore, we can still perform a high-fidelity mixed-species MS gate by using the calcium stretch qubit, despite the magnetic-field instability. However, it should be noted that this cancellation only works for the even-parity input states, $|\uparrow_S \uparrow_Z\rangle$ and $|\downarrow_S \downarrow_Z\rangle$. Figure 6.20 shows the results of partial state tomography using two-point parity measurements, where we find a fidelity of 99.6(2)% for the starting state $|\downarrow_S \downarrow_Z\rangle$.

6.3.6 MS gates using the 674 nm laser

We also implement an MS gate between the calcium clock qubit and the strontium optical qubit, using the 674 nm quadrupole laser to drive the gate on strontium, and $R_{\parallel} + R_{H2}$ for calcium. The single-pass AOM for the 674 nm laser is driven with two frequency tones, as for $R_{\parallel}$, to drive the sidebands in strontium. An outline of the beams is shown in figure 6.21.

For this gate, the lasers for the two species are therefore completely independent, so we tune the wavelength of the Raman lasers closer to the 397 nm transition, choosing $\Delta_{\text{Ca}} = -3.0$ THz.

We find that this method has several drawbacks. Primarily, the Raman lasers cause a significant light shift on the frequency of the strontium optical qubit frequency (-140 kHz for typical laser powers). The centre frequency of the 674 nm bichromat must be shifted during the gate, and the strontium qubit frequency stability becomes dependent on the intensity stability of the Raman lasers. The shift must be recalibrated each time the Raman laser power is adjusted, and necessitates more stringent synchronisation of pulse shaping on each species. Secondly, at this




Figure 6.21: Mølmer-Sørensen gate beams with Raman lasers on calcium and 674 nm quadrupole laser on strontium. $R_{\parallel}$ and the 674 nm beam each include two tones at $\pm(f_{\text{ax}} + \delta_g)$. The Raman detuning from the 397 transition is $\Delta_{\text{Ca}} = -3.0$ THz.

time, the gate speed was limited by the amount of 674 nm power available, making it more susceptible to decoherence. The gate detuning was not much larger than the oscillations in the strontium qubit frequency due to the magnetic-field variations, so it suffered more from the parity-angle drift, not only in the fidelity but also in the practicality of running longer diagnostic experiments. In addition, the need for a third gate laser with a second bichromatic drive increased the number of calibration steps and the amount of daily maintenance required. Finally, in our geometry the 674 nm beam has a non-zero projection onto the radial motional modes, which are not cooled below the Doppler limit. To reduce gate errors due to the temperature of these modes would require implementing sideband cooling of these modes.

This gate mechanism was therefore not fully optimised, in favour of switching to gates driven by the Raman lasers on both qubits. Nevertheless, we measured a fidelity of 94.9(6)% (figure 6.22) for a single-loop gate on the ip mode, although this

fidelity was measured using a four-point parity measurement so it does not take into account the parity-angle wander. The gate detuning was 11 kHz, and the laser powers were 10.0 mW R_{H2} , 17.3 mW $R_{\parallel}$ (total) and 3.75 mW 674 (total). This gate was performed before reducing the Raman spot sizes.




Figure 6.22: Four-point PST fidelity measurement for an MS gate on the calcium clock qubit and the strontium optical qubit. Calcium is driven using the Raman lasers and strontium with the quadrupole laser. Populations are readout-corrected. The inferred fidelity is 94.9(6)%.

6.4 Summary

We perform mixed-species entangling gates between $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ using various different methods, in order to select the most suitable for future experiments. The measured fidelities are summarised in table 6.4.

We reach the highest fidelities (up to 99.8(1)% as measured by PST, or 99.7(1)% by full two-qubit randomised benchmarking) using a light-shift gate between the calcium

mechanism	qubits ($^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$)	fidelity	method
LS	stretch–Zeeman	99.8(1)%	PST (two-point)
		99.7(1)%	interleaved randomised benchmarking
MS	stretch–Zeeman	99.6(2)%	PST (two-point)
	clock–Zeeman	96.2(3)%	PST (two-point)
	clock–optical	94.9(6)%	PST (full-fringe)

Table 6.4: Mixed-species gate fidelities for light-shift (LS) and Mølmer–Sørensen (MS) gates, measured using two-qubit interleaved randomised benchmarking or partial state tomography (PST)

stretch qubit and the strontium Zeeman qubit. This requires just two Raman beams, which are derived from a single laser.

We perform MS gates on the same two ions using three Raman beams from two different lasers. For an MS gate on the calcium clock qubit and the strontium Zeeman qubit we measure a lower fidelity of 96.2(3)% by PST, due to uncertainty in the parity phase arising from slow drifts in the external magnetic-field strength. The average total population $p_{\uparrow_C \uparrow_Z} + p_{\downarrow_C \downarrow_Z}$, however, is 99.4(2)%, so we expect this gate would perform with a fidelity comparable to that of the light-shift gate with appropriate stabilisation of the magnetic field. In our case, for example, this could be done via the coil temperature. More generally, other ion-trap experiments achieve improved stability using a μ -metal shield and/or a permanent or superconducting magnet to generate the field (e.g. [21, 146]). To protect against any remaining magnetic-field noise, more involved dynamical decoupling techniques could be employed during the gate [68, 17, 166], although these are undeniably more complex to implement than the simple spin-echo sequence used for the light-shift gate. If one’s aim is to create a generic Bell state, rather than a specific one, or to perform a gate between two clock qubits, then the drift of the parity phase would not be a limiting factor and the gate could be used with no further modifications. We note, however, that qubits of

different species are extremely unlikely to have clock qubits at the same magnetic-field strength, for $B \neq 0$.

By changing only one of the three Raman beams, we can perform an MS gate between the strontium Zeeman qubit and the calcium stretch qubit, with a Bell-state fidelity of 99.6(2)% as measured by PST starting in the state $|\downarrow_S \downarrow_Z\rangle$. The magnetic-field sensitivities of these two qubits are almost equal and opposite, which allows generation of a Bell state with a stable phase for even-parity input states. With improved field-stabilisation, we would, again, expect this gate to perform with an average fidelity comparable to that of the light-shift gate over all input states. However, this perhaps highlights one of the pitfalls of quoting gate fidelities measured only via PST of a Bell state, and lends strength to the argument for more stringent benchmarking methods; randomised benchmarking of this gate would reveal a lower average fidelity because the cancellation of the field sensitivities does not apply for input states other than $|\uparrow_S \uparrow_Z\rangle$ and $|\downarrow_S \downarrow_Z\rangle$.

Finally, we also demonstrate a gate between the strontium optical qubit and the calcium clock qubit, using two Raman lasers for calcium and an independent 674 nm quadrupole laser for strontium. We measure a fidelity of 94.9(6)% by PST with full-fringe parity-contrast determination (insensitive to changes in the Bell-state phase on timescales longer than the duration of the measurement). This gate was not fully optimised, in favour of pursuing less complex implementations.

6.5 Comparison

A significant advantage of the light-shift gate mechanism is the fact that the required laser frequencies are independent of the qubit frequencies and differ only by the motional mode frequency. Combined with the proximity in wavelength of $S \leftrightarrow P$ transitions in $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$, this means only one laser is required to drive the light-

shift gate. Furthermore, the basis of the gate operation means that phase coherence of the gate lasers with single-qubit operations is not required, and that we may protect against errors due to slow drift of the magnetic-field strength using a simple spin-echo π -pulse implemented via microwave/RF pulses. Both of these features greatly reduce experimental complexity and provide inherent robustness to certain sources of error. However, the light-shift gate cannot straightforwardly be implemented using a magnetic-field-independent qubit. Information may instead be coherently mapped to the calcium clock qubit from the stretch qubit using additional microwave ‘transfer’ pulses with a duration totalling $\lesssim 100 \mu\text{s}$. During these pulses, the qubit state is susceptible to errors due to decoherence and due to the driving of off-resonant transitions.

The MS gate may be used directly on a clock qubit, eliminating the need for these extra pulses. However, due to its dependence on the qubit frequencies, it requires a minimum of one extra laser beam path and two extra AOM frequencies. In turn, this increases the total amount of laser power required for a gate of the same speed. The MS gate operates in an orthogonal basis to that of the light-shift gate, meaning that it is less inherently robust to magnetic-field noise. In addition, the phase of the MS gate is determined by the optical phases of the gate lasers and is therefore subject to variation due to uncontrolled path length changes. However, assuming that the phase is stable over the duration of one entangling gate (true in our case), if necessary, this dependence on the laser phase could be eliminated at the cost of two extra $\frac{\pi}{2}$ -rotations on each qubit before and after the gate operation (see appendix A), allowing the use of microwaves/RF for single-qubit operations.

Chapter 7

Conclusions

7.1 Memory performance

Memory errors are usually measured over long timescales, where $\epsilon_m \gg \epsilon_{\text{SPAM}} \gtrsim 10^{-3}$. Information about shorter timescales is inferred by extrapolation of these measurements, assuming a particular decay model (usually exponential or Gaussian). The level of decoherence is almost universally described using a single parameter from this model, the T_2^* time.

In this thesis, we characterise the memory performance of a $^{43}\text{Ca}^+$ clock qubit in a magnetic field of 146 G. Randomised benchmarking enables us to circumvent the limit imposed by the SPAM error and to directly measure memory errors down to the 10^{-6} level. We find the error ϵ_m remains $< 10^{-4}$ for up to 50 ms, or $< 10^{-3}$ for up to 2 seconds using a simple dynamical decoupling sequence. These timescales are several orders of magnitude longer than previously-demonstrated gate or measurement times, making quantum error correction feasible. Long decoherence times relative to gate times will also be an important feature for near-term applications in devices where error correction is not yet implemented. In addition, they offer more flexibility in the design of computers and the algorithms which can be implemented using them — for

example, if the memory error remains small after the time it takes to move a qubit to a different area of the processor, this increases the scope for connectivity.

Extrapolation of our measured memory errors at longer timescales would overestimate the true error after 1 ms by around an order of magnitude. The inferred T_2^* time depends on the timescale over which memory errors are examined, and on the model used. It is therefore insufficient to use only this parameter to quantify the memory performance, and care should be taken when extrapolating measurements without knowledge of the underlying noise processes.

We find the calcium clock qubit is robust to offsets and residual noise in the external magnetic-field strength, with ϵ_m remaining $< 10^{-4}$ for 1 ms even at an offset of 50 mG. As predicted in [10], the current limit on our memory performance is the stability of the Rb reference oscillator, rather than any fundamental property of the trapped-ion qubit.

7.1.1 Comparison to other work

We can compare our work to that of O'Malley et al. (2015) [137], who first proposed and implemented the technique used in this thesis. In a superconducting quantum system, they measure memory errors over times between 5 ns and 450 ns, finding $\epsilon_m \lesssim 1 \times 10^{-4}$ after 5 ns, with a significant contribution from energy relaxation ($T_1 = 27.6 \mu\text{s}$). Single-qubit rotations in this system have typical durations ≈ 20 ns, and entangling gates ≈ 40 ns, so decoherence contributes non-negligibly to the total two-qubit gate error of $5.6(5) \times 10^{-3}$ [15]. At short timescales, the coherence is limited by telegraph noise.

To the best of our knowledge, this work was the first to directly characterise memory errors in ion-trap systems in the small-error regime. The same method has since been used to measure $\epsilon_m \sim 10^{-4}$ in a $^{171}\text{Yb}^+$ hyperfine qubit [69].

Earlier this year (2021), P. Wang et al. presented the longest measured coherence

time for a single qubit [182]. They achieve $T_2^* = 5500(670)$ seconds for an $^{171}\text{Yb}^+$ hyperfine clock qubit in a magnetic field of 5.8 G, inferred by fitting an exponential decay model to measurements with delay times of up to 960 seconds. A dynamical decoupling sequence [165] with a pulse-spacing of 400 ms is applied during the delays, and the ytterbium ion is continuously sympathetically cooled using a co-trapped $^{138}\text{Ba}^+$ ion. Their coherence time is improved from their previous best of 660 seconds [183] to over an hour. The authors identify magnetic-field fluctuations and instability of their Rb reference oscillator as the previous limiting factors, and improve on these by installing magnetic-field shielding, swapping coils for a permanent magnet, and replacing the Rb reference with a crystal oscillator which has an order of magnitude smaller Allan variance at 1 second.

In 2020, T. R. Tan et al. also identified phase noise of their Rb reference oscillator as the limiting factor in their memory performance [175], improving their coherence time by an order of magnitude to 8.7(7) seconds (without dynamical decoupling) by switching to a 10.6 GHz cryogenic sapphire oscillator. Using a local-oscillator frequency closer to the qubit frequency (in this case 12.6 GHz for $^{171}\text{Yb}^+$) avoids multiplying the phase noise on the 10 MHz reference during upconversion and is a promising approach for reducing memory errors in future.

7.2 Mixed-species two-qubit gates

We present and compare various mechanisms for generating entanglement in a mixed-species, $^{43}\text{Ca}^+ - ^{88}\text{Sr}^+$ crystal, all implemented in the same experimental system.

We demonstrate a light-shift gate between the calcium stretch qubit and the strontium Zeeman qubit, using a single laser. We use partial state tomography to ascertain that we can create a Bell state with 99.8(1)% fidelity, after correction for known read-out errors. This is the highest recorded fidelity of any quantum logic gate between two

different atomic elements to date. We also use full two-qubit interleaved randomised benchmarking to characterise the gate’s performance in the more computationally relevant context of a long sequence of operations with arbitrary input states. This provides a measure which is independent of SPAM errors, and easily comparable across qubit platforms. We find a consistent fidelity of 99.7(1)%, with the inferred error per light-shift gate increasing for longer sequences due to non-Markovianity stemming from thermal effects. The fidelity per two-qubit Clifford operation is 99.17(2)%.

We use two Raman lasers to implement mixed-species MS gates, where one beam contains two frequency tones and is shared between both ion species, again harnessing the advantage of similar transition wavelengths in this ion pair. On the calcium clock qubit, we measure a gate fidelity of 96.2(2)% using partial state tomography of a Bell state; this is lower than expected, due to uncertainty in the Bell-state phase arising from slow drift of the magnetic field. Using the stretch qubit, we create a well-defined Bell state with a fidelity of 99.6(2)%. This qubit has a magnetic-field sensitivity almost equal and opposite to that of the strontium Zeeman qubit, leading to a symmetry of the gate operation against magnetic-field offsets when starting in an even-parity state. In a system with a more stable magnetic field, we predict that this mixed-species MS gate would perform with a fidelity comparable to that of the light-shift gate, on either choice of calcium qubit and for an arbitrary input state, and we plan to implement improved temperature stabilisation for the field-generating coils in our system in future.

A summary of measured fidelities can be found in section 6.4, and a comparison of the various gate mechanisms in section 6.5.

This work was carried out with the application in mind of co-trapping $^{43}\text{Ca}^+$ and $^{88}\text{Sr}^+$ in a node of a quantum network [133], where $^{88}\text{Sr}^+$ would be used for communication between network nodes via photonic entanglement, and the clock qubit in $^{43}\text{Ca}^+$ would be used for long-term storage of quantum information. Concurrently

with the work in this thesis, remote entanglement of two strontium ions in two separate nodes of such a network was demonstrated with a fidelity of 94% at a rate of 182 Hz [170, 169]. Photonic interfacing will be a key ingredient for scaling up trapped-ion quantum processors and distributing computation over a network.

For more general quantum information processing applications, a high-fidelity entangling gate between ions of different species provides the freedom to coherently map information from one to the other to make use of their different strengths. The use of two spectrally-resolved species allows state preparation, cooling and measurement of one species without disturbing the quantum state of the other. Sympathetic cooling using a second ion species is now commonly used in quantum information applications, and mixed-species entanglement may be particularly useful for quantum logic with ancilla qubits for error correction. Other applications include ‘quantum logic’ spectroscopy of molecules or of exotic species for testing fundamental physics, and readout of atomic clocks. In many of these applications, preservation of the phase of superposition states is not strictly required, but the use of an entangling gate has been shown to reduce temperature sensitivity and to scale more favourably to larger numbers of spectroscopy ions [93].

7.2.1 Comparison to other work

7.2.1.1 Other mixed-species gates

Entangling gates for mixed-isotope and mixed-element ion pairs have previously been demonstrated. However, due to the extra technical complications of working with two different species, fidelities have been lower than for the best same-species gates, as shown in figure 7.1. In this work we push mixed-element fidelities up to the level of state-of-the-art same-species gates, and below the $\approx 1\%$ error threshold necessary for error correction using the surface code.




Figure 7.1: Reported mixed-species gate errors in trapped ions (red squares). The best same-species results are shown in blue, for comparison. Squares and crosses indicate laser-driven gates, whilst circles are microwave-driven. The dashed line represents the 1% error threshold necessary for implementation of the surface code.

Mixed-species (fidelities measured by PST):

2015:

- $^{40}\text{Ca}^+ - ^{43}\text{Ca}^+$ LS gate [13]. Fidelity 99.8(6)%.
A previous use of the same gate mechanism as in this thesis, in a previous iteration of the experimental apparatus B.
- $^9\text{Be}^+ - ^{25}\text{Mg}^+$ MS gate [174]. Fidelity 97.9(1)%.

2017:

- $^{171}\text{Yb}^+ - ^{138}\text{Ba}^+$ MS gate [82]. Fidelity $\approx 60\%$.

2018:

- $^9\text{Be}^+ - ^{40}\text{Ca}^+$ MS gate [129]. Fidelity 98.3(9)%.
See also: three-ion $^9\text{Be}^+ - ^{40}\text{Ca}^+ - ^9\text{Be}^+$ GHZ states generated with 93.8(5)% fidelity in [130].

2019:

- $^{40}\text{Ca}^+ - ^{88}\text{Sr}^+$ MS gate [29]. Fidelity 94.3(3)%.
- $^9\text{Be}^+ - ^{25}\text{Mg}^+$ CNOT gate consisting of an optical-phase-insensitive MS gate and single-qubit rotations [181]. Total CNOT fidelity 97(1)%, estimated using PST of the MS gate.

Same-species:

2016: [12, 60, 68], **2018:** [151], **2019:** [54, 194], **2020:** [9], **2021:** [166].

7.2.1.2 Other two-qubit randomised benchmarking

Using a mixed-species crystal of ions with different qubit frequencies allows us to achieve individual addressing with low cross-talk, so that we can explore the full two-qubit Hilbert space. This is less trivial for a same-species ion pair.

Baldwin et al. [9] propose and implement a variation of randomised benchmarking on a symmetric subspace of two $^{171}\text{Yb}^+$ ions, which is achievable without individual addressing. They measure an error per two-qubit gate of $2(1) \times 10^{-3}$, utilising sympathetic cooling via two $^{138}\text{Ba}^+$ ions. Gaebler et al. [61] implement full two-qubit benchmarking on two $^9\text{Be}^+$ qubits by separating the ions into different trap zones for single-qubit rotations, then recombining and sympathetically cooling via $^{24}\text{Mg}^+$ ions before each entangling gate, finding an error per two-qubit gate of 6.9(1.7)%. Similarly, Pino et al. [140] implement full two-qubit benchmarking of an MS gate on two $^{171}\text{Yb}^+$ ions where the ions are spatially separated before each single-qubit rotation, and each ion is sympathetically cooled using a partner $^{138}\text{Ba}^+$ ion. They achieve an error per entangling gate of $7.9(4) \times 10^{-3}$, and confirm that this error does not increase significantly when simultaneous two-qubit randomised benchmarking is performed in a separate trapping zone.

Full two-qubit randomised benchmarking has been more widely implemented in superconducting systems (e.g. [192, 5, 42, 15, 6, 88]) with two-qubit gate errors as low as $5.2(4) \times 10^{-3}$. It has also been used in silicon-based systems [193, 76] to measure two-qubit errors down to 2.0(3)%.

7.2.1.3 Scaling up randomised benchmarking

With more qubits comes the potential for cross-talk, so the performance of larger systems may not necessarily be well predicted from the error rates of isolated single- and two-qubit gates. Operations must therefore be calibrated and assessed in the context in which they will be used.

Although the number of measurements and resources necessary for randomised benchmarking scales more favourably with the number of qubits than for complete characterisation techniques such as process tomography, at the time of writing there has only been one implementation of full three-qubit randomised benchmarking [119], and none on larger numbers of qubits.

Related protocols, however, have been proposed and implemented for more qubits [54, 148, 143, 63]. For example, Erhard et al. [54] note that a single Clifford operation on an n -qubit register, if composed of primitive single- and two-qubit operations, requires $\mathcal{O}(n^2/\ln n)$ entangling gates. With current two-qubit gate errors of $\sim 10^{-3}$, the resulting error for just one 10-qubit Clifford operation would be $\sim 10\%$. This reduces the maximum feasible sequence length, so the number of measurements required to determine these ‘NISQ-era’ error rates can be large. The authors therefore introduce ‘cycle’ benchmarking, similar to a protocol used in [193], where local randomising gates ensure that the number of measurements required to reach a fixed fractional uncertainty in fidelity is approximately independent of the number of qubits. They use the protocol to estimate multi-qubit Mølmer–Sørensen gate fidelities ranging from 99.6(1)% for 2 qubits to 86(2)% for a 10-qubit interaction.

It is also important to acknowledge many other demonstrations of coherent control over larger systems using methods other than randomised benchmarking, with the caveat that they are not all directly comparable since they have different aims and sensitivities to limitations. For example, gate-based quantum algorithms have been run on 10 qubits in an IonQ system [190] and Google implemented random circuits on 53 superconducting qubits in order to claim ‘quantum supremacy’ [5]. IBM’s single-number metric, the ‘quantum volume’ (QV) — which gives equal importance to the number of qubits and to their error rates — has been influential amongst industry players, for benchmarking existing NISQ processors. A QV of 2^D requires a minimum of D qubits and a random-circuit depth of D . The highest QV reported in a research

paper to date is 64 ($= 2^6$), achieved by both IBM [88] and Honeywell [140], although Honeywell have since released data demonstrating a QV of 512 ($= 2^9$) [73].

7.3 Outlook

Trapped ions are a very promising candidate platform for large-scale quantum computing. Qubit state preparation and readout, as well as single- and two-qubit gate operations, have all been demonstrated with fidelities $\gtrsim 99.9\%$, with decoherence times several orders of magnitude longer than typical gate times. All of these features have previously been demonstrated in the two experimental systems used in this thesis.

This work adds mixed-element entangling gates to the repertoire of operations already below the $\sim 1\%$ error threshold required for error correction. However, the required overheads make fault-tolerance with current two-qubit gate errors a challenging task. Reducing errors further below the threshold will reduce these overheads to a more manageable level, and will be a priority for near-term quantum processors. With this in mind, many ion-trap systems are turning to cryogenic cooling, to reduce gate errors due to anomalous heating. Microwave-driven entangling gates also offer the potential to surpass the eventual limit that would be imposed by photon-scattering errors when using lasers. Considering recent progress on decreasing entangling gate times to $\approx 1 \mu\text{s}$ [151, 195], and the availability of more stable reference oscillators, there is also scope to further improve the ratio of decoherence times to gate times.

Another priority will be increasing the numbers of qubits, using the QCCD architecture or photonic interconnects between traps. Mixing ion species is invaluable in either approach (or a combination of the two), because it facilitates sympathetic cooling during shuttling and long gate sequences, and photonic entanglement of ‘communication’ qubits without disturbing the states of logic qubits. Microfabricated surface

traps offer better prospects for scalability, with optics¹ and electronics increasingly being integrated into the trap chips themselves (see e.g. [30] for a review).

The overall performance of a quantum processor will be affected by a number of factors, including the number of qubits and the error rates associated with various computational steps, as well as the connectivity between qubits and the degree to which parallelism of operations is possible. With the rapid development of many differing qubit technologies, benchmarking methods will be a useful tool for assessing progress towards a universal, fault-tolerant quantum computer.


¹Coherent control of calcium and of strontium qubits has been demonstrated using beams delivered via integrated waveguides [120, 132, 121]

Appendix A

Mølmer–Sørensen gate beam geometries

A Mølmer–Sørensen (MS) gate features simultaneous drive of red and blue sidebands in each ion. Following [103], we note that to achieve this requires a minimum of three frequencies for each ion, and each sideband-driving beam pair must have a non-zero differential wavevector $\Delta\mathbf{k}$ in order to couple to the motion. Figure A.1 shows two possible frequency configurations. In each case, two of the frequencies ($\omega_B^\pm$) may be delivered via the same beam path in the laboratory, so that their corresponding wavevectors are parallel. The wavevector of the beam at frequency ω_A must be different. Each sideband is driven by the combination of ω_A with one of $\omega_B^\pm$. The frequency differences driving the sidebands are $\pm(\omega_0 + \omega_m + \delta_g)$ in the ‘phase-insensitive’ geometry and $\omega_0 \pm (\omega_m + \delta_g)$ in the ‘phase-sensitive’ geometry.

The relative phases $\Delta\phi_{l,r,i}$ and $\Delta\phi_{l,b,i}$ would ideally be fixed. However, in reality they are susceptible to uncontrolled fluctuations in the optical path lengths. As suggested by their names, the two different geometries have different sensitivities to these phase instabilities.



(a) Phase-insensitive geometry. The RSB is driven by the frequencies $\{\omega_1, \omega_2\} = \{\omega_B^+, \omega_A\}$ and the BSB is driven by $\{\omega_A, \omega_B^-\}$, where ω_1 is the higher frequency in each case.

(b) Phase-sensitive geometry. The RSB is driven by the frequencies $\{\omega_1, \omega_2\} = \{\omega_A, \omega_B^-\}$ and the BSB is driven by $\{\omega_A, \omega_B^+\}$, where ω_1 is the higher frequency in each case.

Figure A.1: Two different frequency configurations for the MS gate. In this work we use the phase-sensitive geometry.

A.1 Phase-insensitive geometry

In this geometry, the components of $\Delta \mathbf{k}_r$ and $\Delta \mathbf{k}_b$ along the direction of mode m have opposite signs. Instability in the optical path lengths therefore leads to opposite phase variations for each sideband ($\Delta \phi_r = -\Delta \phi_b = \Delta \phi$), so in this configuration we have

$$\begin{aligned} \phi_{M,i} &= \frac{\Delta k_i z_{\text{eq},i} + \Delta \phi_i - (-\Delta k_i z_{\text{eq},i} - \Delta \phi_i)}{2} = \Delta k_i z_{\text{eq},i} + \Delta \phi_i, \\ \phi_{S,i} &= -\frac{[\Delta k_i z_{\text{eq},i} + \Delta \phi_i - (\Delta k_i z_{\text{eq},i} - \Delta \phi_i)]}{2} = 0, \end{aligned} \quad (\text{A.1})$$

where $z_{\text{eq},i}$ is the component of $\mathbf{r}_{\text{eq},i}$ along the axial direction¹. Since typical motional mode frequencies are very small compared to optical transition frequencies, we have made the approximation $|(\Delta \mathbf{k}_r)_z| \approx |(\Delta \mathbf{k}_b)_z| = \Delta k$.

Here we see that the spin phase is always equal to zero for both ions so the gate is always a $\sigma_y \otimes \sigma_y$ operator, independent of the noisy optical phases; this is the

¹We have assumed we are driving the gate on an axial mode, as is the case for all two-qubit gates in this thesis.

advantage of this configuration.

The motional phase for each ion depends on its equilibrium position. The motional phase difference $\phi_{M,2} - \phi_{M,1}$ is equal to $\Delta k_2 z_{\text{eq},2} - \Delta k_1 z_{\text{eq},1}$. For a same-species gate, where we can drive both ions with the same beam pair such that $\Delta k_1 = \Delta k_2$, this means the motional phase depends on the axial separation between the ions in the trap; the drive phase is the same for both ions if $\Delta k(z_{\text{eq},2} - z_{\text{eq},1}) = 2N\pi$ or opposite if $|\Delta k(z_{\text{eq},2} - z_{\text{eq},1})| = 2(N + \frac{1}{2})\pi$; $N \in \mathbb{Z}$. Setting the motional phase to be equal or opposite, as is required for an efficient MS gate, therefore depends on the ion spacing.

A.2 Phase-sensitive geometry

In this geometry, $\Delta \mathbf{k}_r$ and $\Delta \mathbf{k}_b$ are parallel, so instability in the optical path lengths leads to the same phase variations for each sideband ($\Delta \phi_r = \Delta \phi_b = \Delta \phi$) and we have

$$\begin{aligned}\phi_{M,i} &= \frac{\Delta k_i z_{\text{eq},i} + \Delta \phi_i - (\Delta k_i z_{\text{eq},i} + \Delta \phi_i)}{2} = 0, \\ \phi_{S,i} &= -\frac{[\Delta z_i r_{\text{eq},i} + \Delta \phi_i + (\Delta k_i z_{\text{eq},i} + \Delta \phi_i)]}{2} = -(\Delta k_i z_{\text{eq},i} + \Delta \phi_i).\end{aligned}\tag{A.2}$$

In this configuration the motional phase is therefore always equal to zero for both ions, so we satisfy $e^{-i\phi_{M,1}} = e^{-i\phi_{M,2}}$ automatically and independently of the ion spacing. This is the configuration that we choose to use. The phase-sensitive frequency configuration is also more convenient because the difference between frequencies ω_B^+ and ω_B^- is only $\approx 2\omega_m$, which can be bridged using an AOM. The phase-insensitive configuration, on the other hand, requires spanning the qubit frequency; for $^{43}\text{Ca}^+$ this is several GHz.

The disadvantage is that the spin phase (and therefore the gate basis) becomes sensitive to small thermal fluctuations in the optical path lengths, on the scale of an optical wavelength.

A.3 Phase-insensitive sequence

To protect against optical phase fluctuations of the gate beams in longer sequences, comparable to the timescales of these drifts, the MS gate can be wrapped in extra $\frac{\pi}{2}$ -pulses in a ‘phase-insensitive’ sequence, as described in [103] and implemented in e.g. [61, 174, 181, 9, 140].

In order to generate a maximally entangled Bell state from the starting state $|\downarrow_1, \downarrow_2\rangle$, the phase-insensitive sequence is

$$\begin{aligned} & \left[\left(\frac{\pi}{2} \right)_{\phi_{\text{RF}}} \otimes \left(\frac{\pi}{2} \right)_{\phi_{\text{MW}}} \right] \cdot \left[\left(\frac{\pi}{2} \right)_{\phi_{S,1}} \otimes \left(\frac{\pi}{2} \right)_{\phi_{S,2}} \right] \\ & \quad \cdot U_{\text{MS}} \cdot \\ & \left[\left(\frac{\pi}{2} \right)_{\phi_{S,1}+\pi} \otimes \left(\frac{\pi}{2} \right)_{\phi_{S,2}+\pi} \right] \cdot \left[\left(\frac{\pi}{2} \right)_{\phi_{\text{RF}}} \otimes \left(\frac{\pi}{2} \right)_{\phi_{\text{MW}}} \right]. \end{aligned} \tag{A.3}$$

The notation $(\theta)_\phi$ represents a rotation by angle θ around the direction at azimuthal angle ϕ from the x -axis. Pulses with phases ϕ_{MW} or ϕ_{RF} are implemented using microwaves or RF, and those with phases $\phi_{S,i}$ are implemented using the gate lasers. The effect is to remove the dependence of the final Bell-state phase on the phases of the gate lasers; this phase is instead defined by the microwave/RF phases.

Appendix B

Ramsey measurements of small memory errors

Consider using a SPAM-corrected Ramsey measurement (as in figure 5.1) to measure a small memory error $\epsilon_m \ll \epsilon_{\text{SPAM}}$, where the uncertainty on ϵ_m is dominated by the binomial uncertainty on the SPAM measurement,

$$\delta\epsilon_{\text{SPAM}} = \sqrt{\frac{\epsilon_{\text{SPAM}}(1 - \epsilon_{\text{SPAM}})}{N}}, \quad (\text{B.1})$$

for N shots. The fractional uncertainty on the memory error is

$$\frac{\delta\epsilon_m}{\epsilon_m} \approx \frac{\delta\epsilon_{\text{SPAM}}}{\epsilon_m} \approx \sqrt{\frac{\epsilon_{\text{SPAM}}}{N\epsilon_m^2}}, \quad (\text{B.2})$$

so the number of shots required for a fractional uncertainty of 10% is

$$N \approx 100 \frac{\epsilon_{\text{SPAM}}}{\epsilon_m^2}. \quad (\text{B.3})$$

This requires a total time of

$$t_{\text{tot}} = 2 \times N \times t_{\text{shot}}, \quad (\text{B.4})$$

where t_{shot} is the time for one shot, and there is an extra factor of two because each shot of the Ramsey experiment is interleaved with a SPAM measurement.

For a particular Ramsey delay τ , the total time is therefore equal to

$$t_{\text{tot}} \approx 2 \times 100 \frac{\epsilon_{\text{SPAM}}}{(\epsilon_m(\tau))^2} \times (t_{\text{state prep}} + \tau + t_{\text{readout}} + 2t_{\pi/2}). \quad (\text{B.5})$$

According to the memory benchmarking measurements made in chapter 5 (figure 5.4b), the memory error in our system is $\sim 1 \times 10^{-6}$ after a delay of $\tau = 1$ ms. The SPAM error, as measured in figure 5.2, is $\epsilon_{\text{SPAM}} = 2.7 \times 10^{-3}$. The state-preparation time is $t_{\text{state prep}} \approx 430 \mu\text{s}$, the readout time is $t_{\text{readout}} = 1$ ms, and the time required for each Ramsey $\frac{\pi}{2}$ -rotation is $t_{\pi/2} \approx 10 \mu\text{s}$.

The total data acquisition time required to measure a 10^{-6} memory error to a fractional uncertainty of 10% in this system is therefore

$$\begin{aligned} t_{\text{tot}} &\approx 2 \times 100 \frac{2.7 \times 10^{-3}}{(1 \times 10^{-6})^2} \times (430 \mu\text{s} + 1 \text{ ms} + 1 \text{ ms} + 2 \times 10 \mu\text{s}) \\ &= 1.3 \times 10^9 \text{ seconds} \\ &= 42 \text{ years} \\ &\gg 1 \text{ PhD.} \end{aligned} \quad (\text{B.6})$$

Appendix C

Clifford decompositions

C.1 Single-qubit Clifford group

Table C.1 lists the matrix operators of the 24 elements of the single-qubit Clifford group, in the computational (z) basis $\{|\uparrow\rangle, |\downarrow\rangle\}$. The other columns show how these elements can be implemented via different gate sets used in our experiments, where the notation $\pm X, Y, Z$ denotes $\pm\frac{\pi}{2}$ -rotations about the x, y, z axes.

The ‘standard’ gate set, which includes $\pm\frac{\pi}{2}$ -rotations about the x and y axes, and the ‘reduced’ gate set ($+\frac{\pi}{2}$ x - and y -rotations only) are used in chapter 5 (apparatus A). The ‘XYZ’ gate set, which includes $\pm\frac{\pi}{2}$ -rotations about the x , y and z axes, is used in chapter 6 (apparatus B), where the z -rotations are implemented via software phase shifts.

The optimal (fewest physical gates per Clifford gate) decomposition into $\pm X$ and $\pm Y$ gates requires on average 2.17 physical gates per Clifford. However, it was found in apparatus A that the microwave power varied with the duty cycle of the microwave chain due to thermal effects, so extra pairs of $\pm X$ gates were added to some of the decompositions in order to make them more uniform in length — these are indicated in brackets in table C.1. The ‘standard’ decomposition therefore required on average


3.50 physical gates per Clifford. For the same reason, the ‘reduced’ gate set requires on average 3.58 physical gates per Clifford, but the optimal decomposition would require only 3.08 physical gates per Clifford.

The ‘XYZ’ decomposition is the optimal decomposition into $\pm X$, $\pm Y$ and $\pm Z$ gates, as single-qubit gate fidelities in long sequences were less critical for this application. It includes $1.92 \frac{\pi}{2}$ -rotations per Clifford. However, since z -rotations are essentially error-free, a more useful number is the average number of x and y -rotations per Clifford (equal to 1.42).

C.2 Two-qubit Clifford group


The elements $C_i^{(2)}$ of the two-qubit Clifford group may be divided into four classes, depending on the minimum number of entangling gates required to implement them [42].

The first class consists of single-qubit Clifford gates on each qubit, $C_i^{(1)} \otimes C_j^{(1)}$, where $C_i^{(1)}$ is an element of the single-qubit Clifford group as listed in table C.1.



As there are 24 single-qubit Clifford elements, there are $24^2 = 576$ two-qubit Clifford elements of this type.

The second class is the ‘CNOT’ class and contains all combinations of the sequence



i	Matrix operation	‘standard’ (3.50{X,Y}/C)	‘reduced’ (3.58{X,Y}/C)	‘XYZ’ (1.42{X,Y}/C)
1	$\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$	(+X, -X, +X, -X)	(+X, +X, +X, +X)	-
2	$\begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$	+X, -Y, -X	+X, +X, +X, +Y, +X	+Z
3	$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$	+X, +X, +Y, +Y	+X, +X, +Y, +Y	+Z, +Z
4	$\begin{pmatrix} 1 & 0 \\ 0 & -i \end{pmatrix}$	+X, +Y, -X	+X, +Y, +X, +X, +X	-Z
5	$\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$	+X, +X, (+X, -X)	+X, +X	+X, +X
6	$\begin{pmatrix} 0 & i \\ 1 & 0 \end{pmatrix}$	+X, -Y, +X,	+X, +Y, +Y, +Y, +X	+X, +X, -Z
7	$\begin{pmatrix} 0 & -1 \\ 1 & 0 \end{pmatrix}$	+Y, +Y, (+X, -X)	+Y, +Y	+Y, +Y
8	$\begin{pmatrix} 0 & -i \\ 1 & 0 \end{pmatrix}$	+X, +Y, +X	+X, +Y, +X	+X, +X, +Z
9	$\frac{1}{2} \begin{pmatrix} -1+i & 1+i \\ -1+i & 1-i \end{pmatrix}$	+Y, +X, (+X, -X)	+Y, +X	+X, +Z
10	$\frac{1}{2} \begin{pmatrix} -1+i & -1+i \\ -1+i & 1-i \end{pmatrix}$	+X, +X, -Y	+Y, +X, +X	+X, +X, -Y
11	$\frac{1}{2} \begin{pmatrix} -1+i & -1-i \\ -1+i & 1+i \end{pmatrix}$	+Y, -X, (+X, -X)	+Y, +X, +X, +X	+Y, -X
12	$\frac{1}{2} \begin{pmatrix} -1+i & 1-i \\ -1+i & -1+i \end{pmatrix}$	+Y, (+X, -X)	+Y, (+X, +X, +X, +X)	+Y
13	$\frac{1}{2} \begin{pmatrix} 1+i & -1+i \\ -1-i & -1+i \end{pmatrix}$	-Y, -X, (+X, -X)	+X, +X, +Y, +X	-X, +Z
14	$\frac{1}{2} \begin{pmatrix} 1+i & -1-i \\ -1-i & -1-i \end{pmatrix}$	+X, +X, +Y	+X, +X, +Y	+X, +X, +Y
15	$\frac{1}{2} \begin{pmatrix} 1+i & 1-i \\ -1-i & 1-i \end{pmatrix}$	-Y, +X, (+X, -X)	+Y, +Y, +Y, +X	+X, -Z
16	$\frac{1}{2} \begin{pmatrix} 1+i & 1+i \\ -1-i & 1+i \end{pmatrix}$	-Y, (+X, -X)	+Y, +Y, +Y	-Y
17	$\frac{1}{2} \begin{pmatrix} -1-i & -1-i \\ 1-i & -1+i \end{pmatrix}$	-X, -Y, (+X, -X)	+X, +Y, +X, +X	-X, -Y
18	$\frac{1}{2} \begin{pmatrix} -1-i & 1-i \\ 1-i & -1-i \end{pmatrix}$	-X, (+X, -X)	+X, +X, +X	-X
19	$\frac{1}{2} \begin{pmatrix} -1-i & 1+i \\ 1-i & 1-i \end{pmatrix}$	-X, +Y, (+X, -X)	+X, +X, +X, +Y	+Y, +Z
20	$\frac{1}{2} \begin{pmatrix} -1-i & -1+i \\ 1-i & 1+i \end{pmatrix}$	-X, +Y, +Y	+Y, +Y, +X	+X, +Z, +Z
21	$\frac{1}{2} \begin{pmatrix} -1-i & -1-i \\ -1+i & 1-i \end{pmatrix}$	+X, -Y, (+X, -X)	+X, +Y, +Y, +Y	+X, -Y
22	$\frac{1}{2} \begin{pmatrix} -1-i & 1-i \\ -1+i & 1+i \end{pmatrix}$	+X, +Y, +Y	+X, +Y, +Y	+X, +Y, +Y
23	$\frac{1}{2} \begin{pmatrix} -1-i & 1+i \\ -1+i & -1+i \end{pmatrix}$	+X, +Y, (+X, -X)	+X, +Y	+X, +Y
24	$\frac{1}{2} \begin{pmatrix} -1-i & -1+i \\ -1+i & -1-i \end{pmatrix}$	+X, (+X, -X)	+X, (+X, +X, +X, +X)	+X

Table C.1: Decompositions of the 24 single-qubit Clifford group operations into the ‘standard’ ($\pm X, \pm Y$), ‘reduced’ (+X and +Y only) and ‘XYZ’ ($\pm X, \pm Y, \pm Z$) gate sets used in our experiments. The notation ‘ $\pm X, Y, Z$ ’ refers to $\pm \frac{\pi}{2}$ -pulses around the x, y, z axes. The average numbers of $\frac{\pi}{2}$ x - and y -rotations per Clifford are 3.50, 3.58 and 1.42 respectively.

where $S_n^{(1)}$ is an element of the single-qubit exchange group $\{\mathbb{1}, S, S^2\}$. This group permutes the x -, y - and z - axes. S is the operator which maps $(x, y, z) \rightarrow (z, x, y)$, implemented in our decompositions via a $+\frac{\pi}{2}$ y -rotation followed by a $+\frac{\pi}{2}$ x -rotation. S^2 maps $(z, x, y) \rightarrow (y, z, x)$, and is implemented via a $-\frac{\pi}{2}$ x -rotation followed by a $+\frac{\pi}{2}$ y -rotation.


The single-qubit exchange group has 3 elements, so the CNOT class contains $24^2 \times 3^2 = 5184$ elements in total.

The CNOT gate

$$\begin{array}{c} \text{---} \bullet \text{---} \\ | \\ \text{---} \oplus \text{---} \end{array} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix} \quad (\text{C.1})$$

flips the spin of the second qubit if and only if the first ‘control’ qubit is in $|\uparrow\rangle$. It may be implemented using one light-shift gate, G , plus single-qubit rotations.

The third class is the ‘iSWAP’ class




which also contains 5184 elements. The iSWAP gate

$$\begin{array}{c} \text{---} \otimes \text{---} \\ | \\ \text{---} \otimes \text{---} \end{array} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & i & 0 \\ 0 & i & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (\text{C.2})$$


swaps the two qubit states whilst adding a phase i to the odd-parity states, $|\uparrow\downarrow\rangle$ and $|\downarrow\uparrow\rangle$. It may be implemented using two light-shift gates plus single-qubit rotations.

Finally, the ‘SWAP’ class



contains $24^2 = 576$ elements.

The SWAP gate



$$= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (\text{C.3})$$

swaps the two qubit states and may be implemented using three light-shift gates plus single-qubit rotations.

In total, there are 11,520 elements of the two-qubit Clifford group, and the optimal average number of light-shift gates per two-qubit Clifford is 1.5. This is equivalent to the approach in [61].

References

- [1] D. T. C. Allcock. *Surface-electrode ion traps for scalable quantum computing*. PhD thesis. University of Oxford, 2011.
- [2] D. T. C. Allcock, T. P. Harty, C. J. Ballance, et al. “A microfabricated ion trap with integrated microwave circuitry”. In: *Applied Physics Letters* 102.4 (2013), p. 044103.
- [3] D. T. C. Allcock, T. P. Harty, M. A. Sepiol, et al. “Dark-resonance Doppler cooling and high fluorescence in trapped Ca-43 ions at intermediate magnetic field”. In: *New Journal of Physics* 18.2 (2016), p. 023043.
- [4] G. A. Álvarez, A. Ajoy, X. Peng, and D. Suter. “Performance comparison of dynamical decoupling sequences for a qubit in a rapidly fluctuating spin bath”. In: *Phys. Rev. A* 82.4 (2010), p. 042306.
- [5] F. Arute, K. Arya, R. Babbush, et al. “Quantum supremacy using a programmable superconducting processor”. In: *Nature* 574.7779 (2019), pp. 505–510.
- [6] A. Ash-Saki, M. Alam, and S. Ghosh. “Experimental Characterization, Modeling, and Analysis of Crosstalk in a Quantum Computer”. In: *IEEE Transactions on Quantum Engineering* 1 (2020), pp. 1–6.
- [7] D. P. L. Aude Craik. *Near-field microwave addressing of trapped-ion qubits for scalable quantum computation*. PhD thesis. University of Oxford, 2016.

- [8] C. H. Baldwin, B. J. Bjork, M. Foss-Feig, et al. “High-fidelity light-shift gate for clock-state qubits”. In: *Phys. Rev. A* 103.1 (2021), p. 012603.
- [9] C. H. Baldwin, B. J. Bjork, J. P. Gaebler, et al. “Subspace benchmarking high-fidelity entangling operations with trapped ions”. In: *Phys. Rev. Research* 2.1 (2020), p. 013317.
- [10] H. Ball, W. D. Oliver, and M. J. Biercuk. “The role of master clock stability in quantum information processing”. In: *npj Quantum Information* 2.1 (2016), p. 16033.
- [11] C. J. Ballance. *High-Fidelity Quantum Logic in Ca^+* . PhD thesis. University of Oxford, 2014.
- [12] C. J. Ballance, T. P. Harty, N. M. Linke, et al. “High-fidelity quantum logic gates using trapped-ion hyperfine qubits”. In: *Phys. Rev. Lett.* 117.6 (2016), p. 060504.
- [13] C. J. Ballance, V. M. Schäfer, J. P. Home, et al. “Hybrid quantum logic and a test of Bell’s inequality using two different atomic isotopes”. In: *Nature* 528.7582 (2015), pp. 384–386.
- [14] A. Barenco, C. H. Bennett, R. Cleve, et al. “Elementary gates for quantum computation”. In: *Phys. Rev. A* 52.5 (1995), pp. 3457–3467.
- [15] R. Barends, J. Kelly, A. Megrant, et al. “Superconducting quantum circuits at the surface code threshold for fault tolerance”. In: *Nature* 508.7497 (2014), pp. 500–503.
- [16] C. H. Bennett, D. P. DiVincenzo, J. A. Smolin, and W. K. Wootters. “Mixed-state entanglement and quantum error correction”. In: *Phys. Rev. A* 54.5 (1996), pp. 3824–3851.

- [17] A. Bermudez, P. O. Schmidt, M. B. Plenio, and A. Retzker. “Robust trapped-ion quantum logic gates by continuous dynamical decoupling”. In: *Phys. Rev. A* 85.4 (2012), p. 040302.
- [18] M. J. Biercuk, H. Uys, A. P. VanDevender, et al. “Optimized dynamical decoupling in a model quantum memory”. In: *Nature* 458.7241 (2009), pp. 996–1000.
- [19] M. J. Biercuk, H. Uys, A. P. VanDevender, et al. “High-fidelity quantum control using ion crystals in a Penning trap”. In: *Quantum Information and Computation* 9.11 (2009), pp. 920–949.
- [20] J. Binney and D. Skinner. *The Physics of Quantum Mechanics*. Oxford University Press, 2014. ISBN: 9780199688562.
- [21] R. B. Blakestad, C. Ospelkaus, A. P. VanDevender, et al. “Near-ground-state transport of trapped-ion qubits through a multidimensional array”. In: *Phys. Rev. A* 84.3 (2011), p. 032314.
- [22] R. Blume-Kohout and K. Young. “Metrics and Benchmarks for Quantum Processors: State of Play”. In: *Sandia National Laboratories: Technical Report* (2018).
- [23] R. Blume-Kohout, J. K. Gamble, E. Nielsen, et al. “Robust, self-consistent, closed-form tomography of quantum logic gates on a trapped ion qubit”. In: *arXiv* 1310.4492 [quant-ph] (2013).
- [24] J. G. Bohnet, B. C. Sawyer, J. W. Britton, et al. “Quantum spin dynamics and entanglement generation with hundreds of trapped ions”. In: *Science* 352.6291 (2016), pp. 1297–1301.
- [25] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, et al. “Characterizing quantum supremacy in near-term devices”. In: *Nature Physics* 14.6 (2018), pp. 595–600.

- [26] I. A. Boldin, A. Kraft, and C. Wunderlich. “Measuring Anomalous Heating in a Planar Ion Trap with Variable Ion-Surface Separation”. In: *Phys. Rev. Lett.* 120.2 (2018), p. 023201.
- [27] S. M. Brewer, J.-S. Chen, A. M. Hankin, et al. “ $^{27}\text{Al}^+$ Quantum-Logic Clock with a Systematic Uncertainty below 10^{-18} ”. In: *Phys. Rev. Lett.* 123.3 (2019), p. 033201.
- [28] K. R. Brown, J. Kim, and C. Monroe. “Co-designing a scalable quantum computer with trapped atomic ions”. In: *npj Quantum Information* 2.1 (2016), p. 16034.
- [29] C. D. Bruzewicz, R. McConnell, J. Stuart, et al. “Dual-species, multi-qubit logic primitives for Ca^+/Sr^+ trapped-ion crystals”. In: *npj Quantum Information* 5.1 (2019), p. 102.
- [30] C. D. Bruzewicz, J. Chiaverini, R. McConnell, and J. M. Sage. “Trapped-ion quantum computing: Progress and challenges”. In: *Applied Physics Reviews* 6.2 (2019), p. 021314.
- [31] A. H. Burrell. *High Fidelity Readout of Trapped Ion Qubits*. PhD thesis. University of Oxford, 2010.
- [32] H. Y. Carr and E. M. Purcell. “Effects of Diffusion on Free Precession in Nuclear Magnetic Resonance Experiments”. In: *Phys. Rev.* 94.3 (1954), pp. 630–638.
- [33] P. Chapman. *Introducing the World’s Most Powerful Quantum Computer*. 2020. URL: <https://ionq.com/posts/october-01-2020-introducing-most-powerful-quantum-computer> (visited on Apr. 2, 2021).
- [34] C. Chatfield. *The Analysis of Time Series: An Introduction, Sixth Edition*. Chapman & Hall/CRC Texts in Statistical Science. CRC Press, 2016. ISBN: 9781584883173.

- [35] A. Chatterjee, P. Stevenson, S. De Franceschi, et al. “Semiconductor qubits in practice”. In: *Nature Reviews Physics* 3.3 (2021), pp. 157–177.
- [36] J. Chiaverini, R. B. Blakestad, J. Britton, et al. “Surface-Electrode Architecture for Ion-Trap Quantum Information Processing”. In: *Quantum Information and Computation* 5.6 (2005), pp. 419–439.
- [37] C.-w. Chou, C. Kurz, D. B. Hume, et al. “Preparation and coherent manipulation of pure quantum states of a single molecular ion”. In: *Nature* 545.7653 (2017), pp. 203–207.
- [38] J. M. Chow, J. M. Gambetta, L. Tornberg, et al. “Randomized Benchmarking and Process Tomography for Gate Errors in a Solid-State Qubit”. In: *Phys. Rev. Lett.* 102.9 (2009), p. 090502.
- [39] I. L. Chuang and M. A. Nielsen. “Prescription for experimental determination of the dynamics of a quantum black box”. In: *Journal of Modern Optics* 44.11-12 (1997), pp. 2455–2467.
- [40] I. L. Chuang, N. Gershenfeld, and M. Kubinec. “Experimental Implementation of Fast Quantum Searching”. In: *Phys. Rev. Lett.* 80.15 (1998), pp. 3408–3411.
- [41] J. I. Cirac and P. Zoller. “Quantum Computations with Cold Trapped Ions”. In: *Phys. Rev. Lett.* 74.20 (1995), pp. 4091–4094.
- [42] A. D. Córcoles, J. M. Gambetta, J. M. Chow, et al. “Process verification of two-qubit quantum gates by randomized benchmarking”. In: *Phys. Rev. A* 87.3 (2013), p. 030301.
- [43] S. Crain, C. Cahall, G. Vrijsen, et al. “High-speed low-crosstalk detection of a $^{171}\text{Yb}^+$ qubit using superconducting nanowire single photon detectors”. In: *Communications Physics* 2.1 (2019), p. 97.
- [44] A. W. Cross, L. S. Bishop, S. Sheldon, et al. “Validating quantum computers using randomized model circuits”. In: *Phys. Rev. A* 100.3 (2019), p. 032328.

- [45] H. K. Cummins, G. Llewellyn, and J. A. Jones. “Tackling systematic errors in quantum logic gates with composite rotations”. In: *Phys. Rev. A* 67.4 (2003), p. 042308.
- [46] C. Dankert, R. Cleve, J. Emerson, and E. Livine. “Exact and approximate unitary 2-designs and their application to fidelity estimation”. In: *Phys. Rev. A* 80.1 (2009), p. 012304.
- [47] L. Deslauriers, S. Olmschenk, D. Stick, et al. “Scaling and Suppression of Anomalous Heating in Ion Traps”. In: *Phys. Rev. Lett.* 97.10 (2006), p. 103007.
- [48] D. Deutsch. “Quantum theory, the Church-Turing principle and the universal quantum computer”. In: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 400.1818 (1985), pp. 97–117.
- [49] D. Deutsch and R. Jozsa. “Rapid solution of problems by quantum computation”. In: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 439.1907 (1992), pp. 553–558.
- [50] D. P. DiVincenzo and D. Loss. “Quantum information is physical”. In: *Superlattices and Microstructures* 23.3–4 (1998), pp. 419–432.
- [51] D. P. DiVincenzo. “The Physical Implementation of Quantum Computation”. In: *Fortschritte der Physik* 48.9-11 (2000), pp. 771–783.
- [52] J. Emerson, M. Silva, O. Moussa, et al. “Symmetrized Characterization of Noisy Quantum Processes”. In: *Science* 317.5846 (2007), pp. 1893–1896.
- [53] J. Emerson, R. Alicki, and K. Życzkowski. “Scalable noise estimation with random unitary operators”. In: *Journal of Optics B: Quantum and Semiclassical Optics* 7.10 (2005), S347–S352.
- [54] A. Erhard, J. J. Wallman, L. Postler, et al. “Characterizing large-scale quantum computers via cycle benchmarking”. In: *Nature Communications* 10.1 (2019), p. 5347.

- [55] R. P. Feynman. “Simulating physics with computers”. In: *International Journal of Theoretical Physics* 21.6 (1982), pp. 467–488.
- [56] C. J. Foot. *Atomic Physics*. Oxford Master Series in Atomic, Optical and Laser Physics. Oxford University Press, 2005. ISBN: 0198506961.
- [57] A. G. Fowler, M. Mariantoni, J. M. Martinis, and A. N. Cleland. “Surface codes: Towards practical large-scale quantum computation”. In: *Phys. Rev. A* 86.3 (2012), p. 032324.
- [58] M. Fox. *Quantum Optics*. Oxford Master Series in Atomic, Optical and Laser Physics. Oxford University Press, 2006. ISBN: 9780198566724.
- [59] N. Friis, O. Marty, C. Maier, et al. “Observation of Entangled States of a Fully Controlled 20-Qubit System”. In: *Phys. Rev. X* 8.2 (2018), p. 021012.
- [60] J. P. Gaebler, T. R. Tan, Y. Lin, et al. “High-Fidelity Universal Gate Set for $^9\text{Be}^+$ Ion Qubits”. In: *Phys. Rev. Lett.* 117.6 (2016), p. 060505.
- [61] J. P. Gaebler, A. M. Meier, T. R. Tan, et al. “Randomized Benchmarking of Multiqubit Gates”. In: *Phys. Rev. Lett.* 108.26 (2012), p. 260503.
- [62] J. Gambetta. *IBM’s Roadmap For Scaling Quantum Technology*. 2020. URL: <https://www.ibm.com/blogs/research/2020/09/ibm-quantum-roadmap/> (visited on Apr. 10, 2021).
- [63] J. M. Gambetta, A. D. Córcoles, S. T. Merkel, et al. “Characterization of Addressability by Simultaneous Randomized Benchmarking”. In: *Phys. Rev. Lett.* 109.24 (2012), p. 240504.
- [64] L. K. Grover. “A fast quantum mechanical algorithm for database search”. In: *Proceedings of the Twenty-Eighth Annual ACM Symposium on Theory of Computing*. STOC ’96. Philadelphia, Pennsylvania, USA: Association for Computing Machinery, 1996, 212219. ISBN: 0897917855.

- [65] E. L. Hahn. “Spin Echoes”. In: *Phys. Rev.* 80.4 (1950), pp. 580–594.
- [66] T. P. Harty. *High-Fidelity Microwave-Driven Quantum Logic in Intermediate-Field $^{43}\text{Ca}^+$* . PhD thesis. University of Oxford, 2013.
- [67] T. P. Harty, D. T. C. Allcock, C. J. Ballance, et al. “High-Fidelity Preparation, Gates, Memory, and Readout of a Trapped-Ion Quantum Bit”. In: *Phys. Rev. Lett.* 113.22 (2014), p. 220501.
- [68] T. P. Harty, M. A. Sepiol, D. T. C. Allcock, et al. “High-Fidelity Trapped-Ion Quantum Logic Using Near-Field Microwaves”. In: *Phys. Rev. Lett.* 117.14 (2016), p. 140501.
- [69] D. Hayes, D. Stack, B. Bjork, et al. “Eliminating Leakage Errors in Hyperfine Qubits”. In: *Phys. Rev. Lett.* 124.17 (2020), p. 170501.
- [70] D. Hayes, K. Khodjasteh, L. Viola, and M. J. Biercuk. “Reducing sequencing complexity in dynamical quantum error suppression by Walsh modulation”. In: *Phys. Rev. A* 84.6 (2011), p. 062323.
- [71] J. P. Home. “Quantum Science and Metrology with Mixed-Species Ion Chains”. In: *Advances In Atomic, Molecular, and Optical Physics* 62 (2013). Ed. by E. Arimondo, P. R. Berman, and C. C. Lin, pp. 231 –277.
- [72] J. P. Home, D. Hanneke, J. D. Jost, et al. “Complete Methods Set for Scalable Ion Trap Quantum Information Processing”. In: *Science* 325.5945 (2009), pp. 1227–1230.
- [73] Honeywell Quantum Solutions. *Honeywell Sets New Record For Quantum Computing Performance*. 2021. URL: <https://www.honeywell.com/us/en/news/2021/03/honeywell-sets-new-record-for-quantum-computing-performance> (visited on Apr. 12, 2021).

- [74] Honeywell Quantum Solutions. *Quantum Milestone: How We Quadrupled Performance*. 2021. URL: <https://www.honeywell.com/us/en/news/2021/03/quantum-milestone-how-we-quadrupled-performance> (visited on Apr. 2, 2021).
- [75] J. Hsu. *CES 2018: Intel’s 49-Qubit Chip Shoots for Quantum Supremacy*. 2018. URL: <https://spectrum.ieee.org/tech-talk/computing/hardware/intels-49qubit-chip-aims-for-quantum-supremacy> (visited on Apr. 10, 2021).
- [76] W. Huang, C. H. Yang, K. W. Chan, et al. “Fidelity benchmarks for two-qubit gates in silicon”. In: *Nature* 569.7757 (2019), pp. 532–536.
- [77] D. Hucul, I. V. Inlek, G. Vittorini, et al. “Modular entanglement of atomic qubits using photons and phonons”. In: *Nature Physics* 11.1 (2015), pp. 37–42.
- [78] D. Hucul, M. Yeo, S. Olmschenk, et al. “On the transport of Atomic Ions in Linear and Multidimensional Ion Trap Arrays”. In: *Quantum Information and Computation* 8.6 (2008), pp. 501–578.
- [79] A. C. Hughes, V. M. Schäfer, K. Thirumalai, et al. “Benchmarking a High-Fidelity Mixed-Species Entangling Gate”. In: *Phys. Rev. Lett.* 125.8 (2020), p. 080504.
- [80] D. B. Hume, T. Rosenband, and D. J. Wineland. “High-Fidelity Adaptive Qubit Detection through Repetitive Quantum Nondemolition Measurements”. In: *Phys. Rev. Lett.* 99.12 (2007), p. 120502.
- [81] IEEE Standards Coordinating Committee 27 on Time and Frequency, Institute of Electrical and Electronics Engineers, and IEEE-SA Standards Board. *IEEE Standard Definitions of Physical Quantities for Fundamental Frequency and Time Metrology—Random Instabilities*. IEEE Std. IEEE, 2008. ISBN: 9780738158020.

- [82] I. V. Inlek, C. Crocker, M. Lichtman, et al. “Multispecies Trapped-Ion Node for Quantum Networking”. In: *Phys. Rev. Lett.* 118.25 (2017), p. 250502.
- [83] IonQ. *IonQ Unveils World’s Most Powerful Quantum Computer*. 2020. URL: <https://ionq.com/news/october-01-2020-most-powerful-quantum-computer> (visited on Apr. 10, 2021).
- [84] D. F. James and J. Jerke. “Effective Hamiltonian theory and its applications in quantum information”. In: *Canadian Journal of Physics* 85.6 (2007), pp. 625–632.
- [85] H. A. Janacek. *Optical Bloch equations for simulating trapped-ion qubits*. PhD thesis. University of Oxford, 2015.
- [86] J. A. Jones and M. Mosca. “Implementation of a quantum algorithm on a nuclear magnetic resonance quantum computer”. In: *Journal of Chemical Physics* 109.5 (1998), pp. 1648–1653.
- [87] R. Jozsa. “Fidelity for Mixed Quantum States”. In: *Journal of Modern Optics* 41.12 (1994), pp. 2315–2323.
- [88] P. Jurcevic, A. Javadi-Abhari, L. S. Bishop, et al. “Demonstration of quantum volume 64 on a superconducting quantum computing system”. In: *Quantum Science and Technology* 6.2 (2021), p. 025020.
- [89] H. Kaufmann, T. Ruster, C. T. Schmiegelow, et al. “Scalable Creation of Long-Lived Multipartite Entanglement”. In: *Phys. Rev. Lett.* 119.15 (2017), p. 150503.
- [90] J. Kelly. *A Preview of Bristlecone, Google’s New Quantum Processor*. 2018. URL: <https://ai.googleblog.com/2018/03/a-preview-of-bristlecone-googles-new.html> (visited on Apr. 10, 2021).
- [91] D. Kielpinski, C. Monroe, and D. J. Wineland. “Architecture for a large-scale ion-trap quantum computer”. In: *Nature* 417.6890 (2002), pp. 709–711.

- [92] D. Kielpinski, B. E. King, C. J. Myatt, et al. “Sympathetic cooling of trapped ions for quantum logic”. In: *Phys. Rev. A* 61.3 (2000), p. 032310.
- [93] D. Kienzler, Y. Wan, S. D. Erickson, et al. “Quantum Logic Spectroscopy with Ions in Thermal Motion”. In: *Phys. Rev. X* 10.2 (2020), p. 021012.
- [94] G. Kirchmair, J. Benhelm, F. Zähringer, et al. “Deterministic entanglement of ions in thermal states of motion”. In: *New Journal of Physics* 11.2 (2009), p. 023002.
- [95] M. Kjaergaard, M. E. Schwartz, A. Greene, et al. “Programming a quantum computer with quantum instructions”. In: *arXiv* 2001.08838 [quant-ph] (2020).
- [96] M. Kjaergaard, M. E. Schwartz, J. Braumüller, et al. “Superconducting Qubits: Current State of Play”. In: *Annual Review of Condensed Matter Physics* 11.1 (2020), pp. 369–395.
- [97] E. Knill. “Quantum computing with realistically noisy devices”. In: *Nature* 434.7029 (2005), pp. 39–44.
- [98] E. Knill, D. Leibfried, R. Reichle, et al. “Randomized benchmarking of quantum gates”. In: *Phys. Rev. A* 77.1 (2008), p. 012307.
- [99] E. Knill, R. Laflamme, and W. H. Zurek. “Resilient Quantum Computation”. In: *Science* 279.5349 (1998), pp. 342–345.
- [100] S. Kotler, N. Akerman, N. Navon, et al. “Measurement of the magnetic interaction between two bound electrons of two separate ions”. In: *Nature* 510.7505 (2014), pp. 376–380.
- [101] K. A. Landsman, Y. Wu, P. H. Leung, et al. “Two-qubit entangling gates within arbitrarily long chains of trapped ions”. In: *Phys. Rev. A* 100.2 (2019), p. 022332.

- [102] C. Langer, R. Ozeri, J. D. Jost, et al. “Long-Lived Qubit Memory Using Atomic Ions”. In: *Phys. Rev. Lett.* 95.6 (2005), p. 060502.
- [103] P. J. Lee, K.-A. Brickman, L. Deslauriers, et al. “Phase control of trapped ion quantum gates”. In: *Journal of Optics B: Quantum and Semiclassical Optics* 7.10 (2005), S371–S383.
- [104] D. Leibfried, B. DeMarco, V. Meyer, et al. “Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate”. In: *Nature* 422.6930 (2003), pp. 412–415.
- [105] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland. “Quantum dynamics of single trapped ions”. In: *Rev. Mod. Phys.* 75.1 (2003), pp. 281–324.
- [106] B. Lekitsch, S. Weidt, A. G. Fowler, et al. “Blueprint for a microwave trapped ion quantum computer”. In: *Science Advances* 3.2 (2017).
- [107] N. M. Linke, D. Maslov, M. Roetteler, et al. “Experimental comparison of two quantum computing architectures”. In: *Proceedings of the National Academy of Sciences* 114.13 (2017), pp. 3305–3310.
- [108] J. Lishman and F. Mintert. “Trapped-ion entangling gates robust against qubit frequency errors”. In: *Phys. Rev. Research* 2.3 (2020), p. 033117.
- [109] R. Loudon. *The Quantum Theory of Light*. Third. Oxford University Press, 2000. ISBN: 0198501765.
- [110] D. M. Lucas, B. Keitch, J. P. Home, et al. “A long-lived memory qubit on a low-decoherence quantum bus”. In: *arXiv* 0710.4421 [quant-ph] (2007).
- [111] D. M. Lucas, A. Ramos, J. P. Home, et al. “Isotope-selective photoionization for calcium ion trapping”. In: *Phys. Rev. A* 69.1 (2004), p. 012711.
- [112] E. Magesan, J. M. Gambetta, and J. Emerson. “Characterizing quantum gates via randomized benchmarking”. In: *Phys. Rev. A* 85.4 (2012), p. 042311.

- [113] E. Magesan, J. M. Gambetta, and J. Emerson. “Scalable and Robust Randomized Benchmarking of Quantum Processes”. In: *Phys. Rev. Lett.* 106.18 (2011), p. 180504.
- [114] E. Magesan, J. M. Gambetta, B. R. Johnson, et al. “Efficient Measurement of Quantum Gate Error by Interleaved Randomized Benchmarking”. In: *Phys. Rev. Lett.* 109.8 (2012), p. 080505.
- [115] W. Magnus. “On the Exponential Solution of Differential Equations for a Linear Operator”. In: *Communications on Pure and Applied Mathematics* 7.4 (1954), pp. 649–673.
- [116] T. Manovitz, A. Rotem, R. Shaniv, et al. “Fast Dynamical Decoupling of the Mølmer-Sørensen Entangling Gate”. In: *Phys. Rev. Lett.* 119.22 (2017), p. 220505.
- [117] J. M. Martinis, S. Nam, J. Aumentado, et al. “Decoherence of a superconducting qubit due to bias noise”. In: *Phys. Rev. B* 67.9 (2003), p. 094510.
- [118] A. J. McCaskey, Z. P. Parks, J. Jakowski, et al. “Quantum chemistry as a benchmark for near-term quantum computers”. In: *npj Quantum Information* 5.1 (2019), p. 99.
- [119] D. C. McKay, S. Sheldon, J. A. Smolin, et al. “Three-Qubit Randomized Benchmarking”. In: *Phys. Rev. Lett.* 122.20 (2019), p. 200502.
- [120] K. K. Mehta, C. D. Bruzewicz, R. McConnell, et al. “Integrated optical addressing of an ion qubit”. In: *Nature Nanotechnology* 11.12 (2016), pp. 1066–1070.
- [121] K. K. Mehta, C. Zhang, M. Malinowski, et al. “Integrated optical multi-ion quantum logic”. In: *Nature* 586.7830 (2020), pp. 533–537.

- [122] S. Meiboom and D. Gill. “Modified Spin-Echo Method for Measuring Nuclear Relaxation Times”. In: *Review of Scientific Instruments* 29.8 (1958), pp. 688–691.
- [123] B. Merkel, K. Thirumalai, J. E. Tarlton, et al. “Magnetic field stabilization system for atomic physics experiments”. In: *Review of Scientific Instruments* 90.4 (2019), p. 044702.
- [124] P. W. Milonni. “Why spontaneous emission?” In: *American Journal of Physics* 52.4 (1984), pp. 340–343.
- [125] C. Monroe, D. M. Meekhof, B. E. King, et al. “Demonstration of a Fundamental Quantum Logic Gate”. In: *Phys. Rev. Lett.* 75.25 (1995), pp. 4714–4717.
- [126] C. Monroe, R. Raussendorf, A. Ruthven, et al. “Large-scale modular quantum-computer architecture with atomic memory and photonic interconnects”. In: *Phys. Rev. A* 89.2 (2014), p. 022317.
- [127] G. Morigi and H. Walther. “Two-species Coulomb chains for quantum information”. In: *The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics* 13.2 (2001), pp. 261–269.
- [128] G. Nebe, E. M. Rains, and N. J. A. Sloane. “The invariants of the Clifford groups”. In: *Designs, Codes and Cryptography* 24.1 (2001), pp. 99–122.
- [129] V. Negnevitsky. Private communication. 2020.
- [130] V. Negnevitsky, M. Marinelli, K. K. Mehta, et al. “Repeated multi-qubit readout and feedback with a mixed-species trapped-ion register”. In: *Nature* 563.7732 (2018), pp. 527–531.
- [131] X. R. Nie, C. F. Roos, and D. F. James. “Theory of cross phase modulation for the vibrational modes of trapped ions”. In: *Physics Letters A* 373.4 (2009), pp. 422–425.

- [132] R. J. Niffenegger, J. Stuart, C. Sorace-Agaskar, et al. “Integrated multi-wavelength control of an ion qubit”. In: *Nature* 586.7830 (2020), pp. 538–542.
- [133] R. Nigmatullin, C. J. Ballance, N. de Beaudrap, and S. C. Benjamin. “Minimally complex ion traps as modules for quantum communication and computing”. In: *New Journal of Physics* 18.10 (2016), p. 103028.
- [134] H.-R. Noh and W. Jhe. “Analytic solutions of the optical Bloch equations”. In: *Optics Communications* 283.11 (2010), pp. 2353–2355.
- [135] R. Nourshargh. *Miniaturised laser systems for ion trap quantum computing*. Master’s thesis. University of Oxford, 2016.
- [136] S. Olmschenk, R. Chicireanu, K. D. Nelson, and J. V. Porto. “Randomized benchmarking of atomic qubits in an optical lattice”. In: *New Journal of Physics* 12.11 (2010), p. 113007.
- [137] P. J. J. O’Malley, J. Kelly, R. Barends, et al. “Qubit Metrology of Ultralow Phase Noise Using Randomized Benchmarking”. In: *Phys. Rev. Applied* 3.4 (2015), p. 044009.
- [138] C. Ospelkaus, C. E. Langer, J. M. Amini, et al. “Trapped-Ion Quantum Logic Gates Based on Oscillating Magnetic Fields”. In: *Phys. Rev. Lett.* 101.9 (2008), p. 090502.
- [139] R. Ozeri, C. Langer, J. D. Jost, et al. “Hyperfine Coherence in the Presence of Spontaneous Photon Scattering”. In: *Phys. Rev. Lett.* 95.3 (2005), p. 030403.
- [140] J. M. Pino, J. M. Dreiling, C. Figgatt, et al. “Demonstration of the trapped-ion quantum CCD computer architecture”. In: *Nature* 592.7853 (2021), pp. 209–213.
- [141] J. F. Poyatos, J. I. Cirac, and P. Zoller. “Complete Characterization of a Quantum Process: The Two-Bit Quantum Gate”. In: *Phys. Rev. Lett.* 78.2 (1997), pp. 390–393.

- [142] J. Preskill. “Quantum Computing in the NISQ era and beyond”. In: *Quantum* 2 (2018), p. 79.
- [143] T. J. Proctor, A. Carignan-Dugas, K. Rudinger, et al. “Direct Randomized Benchmarking for Multiqubit Devices”. In: *Phys. Rev. Lett.* 123.3 (2019), p. 030503.
- [144] S. Resch and U. R. Karpuzcu. “Benchmarking Quantum Computers and the Impact of Quantum Noise”. In: *arXiv* 1912.00546 [quant-ph] (2020).
- [145] Rigetti Computing. *Rigetti Aspen-8 on AWS*. 2020. URL: <https://medium.com/rigetti/rigetti-aspen-8-on-aws-236d9dc11613> (visited on Apr. 10, 2021).
- [146] T. Ruster, C. T. Schmiegelow, H. Kaufmann, et al. “A long-lived Zeeman trapped-ion qubit”. In: *Applied Physics B* 122.10 (2016), p. 254.
- [147] C. A. Ryan, J. S. Hodges, and D. G. Cory. “Robust Decoupling Techniques to Extend Quantum Coherence in Diamond”. In: *Phys. Rev. Lett.* 105.20 (2010), p. 200402.
- [148] C. A. Ryan, M. Laforest, and R. Laflamme. “Randomized benchmarking of single- and multi-qubit control in liquid-state NMR quantum information processing”. In: *New Journal of Physics* 11.1 (2009), p. 013034.
- [149] C. A. Sackett, D. Kielpinski, B. E. King, et al. “Experimental entanglement of four particles”. In: *Nature* 404.6775 (2000), pp. 256–259.
- [150] V. M. Schäfer. *Fast gates and mixed-species entanglement with trapped ions*. PhD thesis. University of Oxford, 2018.
- [151] V. M. Schäfer, C. J. Ballance, K. Thirumalai, et al. “Fast quantum logic gates with trapped-ion qubits”. In: *Nature* 555.7694 (2018), pp. 75–78.

- [152] P. O. Schmidt, T. Rosenband, C. Langer, et al. “Spectroscopy Using Quantum Logic”. In: *Science* 309.5735 (2005), pp. 749–752.
- [153] L. Schmöger, O. O. Versolato, M. Schwarz, et al. “Coulomb crystallization of highly charged ions”. In: *Science* 347.6227 (2015), pp. 1233–1236.
- [154] J. A. Sedlacek, A. Greene, J. Stuart, et al. “Distance scaling of electric-field noise in a surface-electrode ion trap”. In: *Phys. Rev. A* 97.2 (2018), p. 020302.
- [155] M. A. Sepiol. *A high-fidelity microwave driven two-qubit quantum logic gate in $^{43}\text{Ca}^+$* . PhD thesis. University of Oxford, 2016.
- [156] M. A. Sepiol, A. C. Hughes, J. E. Tarlton, et al. “Probing Qubit Memory Errors at the Part-per-Million Level”. In: *Phys. Rev. Lett.* 123.11 (2019), p. 110503.
- [157] P. W. Shor. “Algorithms for Quantum Computation: Discrete Logarithms and Factoring”. In: *Proceedings 35th Annual Symposium on Foundations of Computer Science*. Santa Fe, NM, USA, 1994, pp. 124–134.
- [158] P. W. Shor. “Scheme for reducing decoherence in quantum computer memory”. In: *Phys. Rev. A* 52.4 (1995), R2493–R2496.
- [159] J. D. Siverns, L. R. Simkins, S. Weidt, and W. K. Hensinger. “On the application of radio frequency voltages to ion traps via helical resonators”. In: *Applied Physics B* 107.4 (2012), pp. 921–934.
- [160] S. Slussarenko and G. J. Pryde. “Photonic quantum information processing: A concise review”. In: *Applied Physics Reviews* 6.4 (2019), p. 041303.
- [161] C. Smorra, K. Blaum, L. Bojtar, et al. “BASE – The Baryon Antibaryon Symmetry Experiment”. In: *The European Physical Journal Special Topics* 224.16 (2015), pp. 3055–3108.
- [162] A. Somoroff, Q. Ficheux, R. A. Mencia, et al. “Millisecond coherence in a superconducting qubit”. In: *arXiv* 2103.08578 [quant-ph] (2021).

- [163] A. Sørensen and K. Mølmer. “Entanglement and quantum computation with ions in thermal motion”. In: *Phys. Rev. A* 62.2 (2000), p. 022311.
- [164] A. Sørensen and K. Mølmer. “Quantum Computation with Ions in Thermal Motion”. In: *Phys. Rev. Lett.* 82.9 (1999), pp. 1971–1974.
- [165] A. M. Souza, G. A. Álvarez, and D. Suter. “Robust Dynamical Decoupling for Quantum Computing and Quantum Memory”. In: *Phys. Rev. Lett.* 106.24 (2011), p. 240501.
- [166] R. Srinivas, S. C. Burd, H. M. Knaack, et al. “High-fidelity laser-free universal control of two trapped ion qubits”. In: *arXiv* 2102.12533 [quant-ph] (2021).
- [167] A. M. Steane. “Error Correcting Codes in Quantum Theory”. In: *Phys. Rev. Lett.* 77.5 (1996), pp. 793–797.
- [168] A. M. Steane. “How to build a 300 bit, 1 Giga-operation quantum computer”. In: *Quantum Information and Computation* 7.3 (2007), pp. 171–183.
- [169] L. J. Stephenson. *Entanglement between nodes of a quantum network*. PhD thesis. University of Oxford, 2019.
- [170] L. J. Stephenson, D. P. Nadlinger, B. C. Nichol, et al. “High-Rate, High-Fidelity Entanglement of Qubits Across an Elementary Quantum Network”. In: *Phys. Rev. Lett.* 124.11 (2020), p. 110501.
- [171] D. J. Szwer. *High Fidelity Readout and Protection of a $^{43}\text{Ca}^+$ Trapped Ion Qubit*. PhD thesis. University of Oxford, 2009.
- [172] D. J. Szwer, S. C. Webster, A. M. Steane, and D. M. Lucas. “Keeping a single qubit alive by experimental dynamic decoupling”. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 44.2 (2010), p. 025501.
- [173] T. R. Tan, J. P. Gaebler, R. Bowler, et al. “Demonstration of a Dressed-State Phase Gate for Trapped Ions”. In: *Phys. Rev. Lett.* 110.26 (2013), p. 263002.

- [174] T. R. Tan, J. P. Gaebler, Y. Lin, et al. “Multi-element logic gates for trapped-ion qubits”. In: *Nature* 528.7582 (2015), pp. 380–383.
- [175] T. R. Tan, C. Edmunds, A. Milne, et al. “Benchmarking a Trapped-Ion Quantum Memory with a Cryogenic Sapphire Oscillator”. In: *APS Division of Atomic, Molecular and Optical Physics Meeting Abstracts*. APS Meeting Abstracts. 2020, G02.002.
- [176] J. E. Tarlton. *Probing qubit memory errors at the 10^{-5} level*. PhD thesis. Imperial College London, 2018.
- [177] K. Thirumalai. *High-Fidelity Mixed Species Entanglement of Trapped Ions*. PhD thesis. University of Oxford, 2019.
- [178] G. S. Uhrig. “Keeping a Quantum Bit Alive by Pptimized π -Pulse Sequences”. In: *Phys. Rev. Lett.* 98.10 (2007), p. 100504.
- [179] H. Uys, M. J. Biercuk, A. P. VanDevender, et al. “Decoherence due to Elastic Rayleigh Scattering”. In: *Phys. Rev. Lett.* 105.20 (2010), p. 200401.
- [180] L. M. K. Vandersypen and I. L. Chuang. “NMR techniques for quantum control and computation”. In: *Rev. Mod. Phys.* 76.4 (2005), pp. 1037–1069.
- [181] Y. Wan, D. Kienzler, S. D. Erickson, et al. “Quantum gate teleportation between separated qubits in a trapped-ion processor”. In: *Science* 364.6443 (2019), pp. 875–878.
- [182] P. Wang, C.-Y. Luan, M. Qiao, et al. “Single ion qubit with estimated coherence time exceeding one hour”. In: *Nature Communications* 12.1 (2021), p. 233.
- [183] Y. Wang, M. Um, J. Zhang, et al. “Single-qubit quantum memory exceeding ten-minute coherence time”. In: *Nature Photonics* 11.10 (2017), pp. 646–650.

- [184] Z.-H. Wang and V. V. Dobrovitski. “Aperiodic dynamical decoupling sequences in the presence of pulse errors”. In: *Journal of Physics B: Atomic, Molecular and Optical Physics* 44.15 (2011), p. 154004.
- [185] S. Webster. *Raman Sideband Cooling and Coherent Manipulation of Trapped Ions*. PhD thesis. University of Oxford, 2005.
- [186] A. G. White, A. Gilchrist, G. J. Pryde, et al. “Measuring two-qubit gates”. In: *J. Opt. Soc. Am. B* 24.2 (2007), pp. 172–183.
- [187] S. Wimperis. “Broadband, Narrowband, and Passband Composite Pulses for Use in Advanced NMR Experiments”. In: *Journal of Magnetic Resonance, Series A* 109.2 (1994), pp. 221–231.
- [188] D. J. Wineland, C. Monroe, W. M. Itano, et al. “Experimental Issues in Coherent Quantum-State Manipulation of Trapped Atomic Ions”. In: *J. Res. Natl. Inst. Stand. Technol* 103.3 (1998), p. 259.
- [189] S. R. Woodrow. *Linear Paul trap design for high-fidelity, scalable quantum information processing*. Master’s thesis. University of Oxford, 2016.
- [190] K. Wright, K. M. Beck, S. Debnath, et al. “Benchmarking an 11-qubit quantum computer”. In: *Nature Communications* 10.1 (2019), p. 5464.
- [191] J. B. Wübbena, S. Amairi, O. Mandel, and P. O. Schmidt. “Sympathetic cooling of mixed-species two-ion crystals for precision spectroscopy”. In: *Phys. Rev. A* 85.4 (2012), p. 043412.
- [192] Y. Xu, J. Chu, J. Yuan, et al. “High-Fidelity, High-Scalability Two-Qubit Gate Scheme for Superconducting Qubits”. In: *Phys. Rev. Lett.* 125.24 (2020), p. 240503.
- [193] X. Xue, T. F. Watson, J. Helsen, et al. “Benchmarking Gate Fidelities in a Si/SiGe Two-Qubit Device”. In: *Phys. Rev. X* 9.2 (2019), p. 021011.

- [194] G. Zarantonello, H. Hahn, J. Morgner, et al. “Robust and Resource-Efficient Microwave Near-Field Entangling $^9\text{Be}^+$ Gate”. In: *Phys. Rev. Lett.* 123.26 (2019), p. 260503.
- [195] C. Zhang, F. Pokorny, W. Li, et al. “Submicrosecond entangling gate between trapped ions via Rydberg interaction”. In: *Nature* 580.7803 (2020), pp. 345–349.
- [196] J. Zhang, G. Pagano, P. W. Hess, et al. “Observation of a many-body dynamical phase transition with a 53-qubit quantum simulator”. In: *Nature* 551.7682 (2017), pp. 601–604.